2019 American Control Conference (ACC)
Philadelphia, PA, USA, July 10-12, 2019

Experimental and educational platforms for studying architecture and tradeoffs in human sensorimotor control

Quanying Liu[1], Yorie Nakahira[1], Ahkeel Mohideen[1], Adam Dai[1], Sunghoon Choi[1], Angelina Pan[1], Dimitar M. Ho[1] and John C. Doyle[1]

Abstract— **This paper describes several surprisingly rich but simple demos and a new experimental platform for human sensorimotor control research and also controls education. The platform safely simulates a canonical sensorimotor task of riding a mountain bike down a steep, twisting, bumpy trail using a standard display and inexpensive off-the-shelf gaming steering wheel with a force feedback motor. We use the platform to verify our theory, presented in a companion paper. The theory tells how component hardware speed-accuracy tradeoffs (SATs) in control loops impose corresponding SATs at the system level and how effective architectures mitigate the deleterious impact of hardware SATs through layering and "diversity sweet spots" (DSSs). Specifically, we measure the impacts on system performance of delays, quantization, and uncertainties in sensorimotor control loops, both within the subject's nervous system and added externally via software in the platform. This provides a remarkably rich test of the theory, which is consistent with all preliminary data. Moreover, as the theory predicted, subjects effectively multiplex specific higher layer planning/tracking of the trail using vision with lower layer rejection of unseen bump disturbances using reflexes. In contrast, humans multitask badly on tasks that do not naturally distribute across layers (e.g. texting and driving). The platform is cheap to build and easy to program for both research and education purposes, yet verifies our theory, which is aimed at closing a crucial gap between neurophysiology and sensorimotor control. The platform can be downloaded at `https://github.com/Doyle-Lab/WheelCon`.**

I. INTRODUCTION

A heavily studied and central topic in neuroscience is speed-accuracy tradeoffs (SATs) [1–4]. At the neuron-component level, the resource limitations (space and metabolic costs) of the brain impose severe speed/accuracy constraints in neural signaling [5], as well as in the muscle actuation [6–8]. At the system level, the tradeoffs associated with the hard limits on SAT are extensively studied with a variety of experiments [9–12] and mathematical theory [2], [3]. However, there is little attention to the theory or experiments that can connect system-level SATs with SATs within the underlying nerve components.

A companion paper [13] develops a theoretical framework to connect the system-level and component-level SATs. The results suggest effectively layering sensorimotor control loops with appropriate diversity in neurons/muscles help achieve systems that are *both* fast and accurate despite being built from components that individually are not.

Motivated by the theory, here we present an experimental platform, some preliminary behavioral results to validate the theory and the relevant neuroscience details for DSSs. In particular, our platform allows us to study the following two questions and potentially more.

1) Combining component and system levels: The system level performance for sensorimotor control is supported by the underlying component hardware, such as lower-layer sensors and muscles and connecting nerves [6–8], and higher-layer nerves in the spinal cord and in cortex [14], [15]. Each component has its own speed and accuracy in the control loops. Inspired by the theory, developing an experimental platform that allows us to test the impact of the speed and accuracy in these components on system performance is essential for connecting neurophysiology and sensorimotor control.

2) Layered architecture in sensorimotor control: The human sensorimotor control system is extremely robust, although the sensing is distributed, sparse, quantized, noisy and delayed; the computing in the central nervous system is slow [5]; and the muscle actuation fatigues and saturates [16]. Effective layered architectures have evolved, to integrate the higher layers of goals, plans, decisions with lower layers of sensing, reflex, and action [17]. Take riding a mountain bike as an example. Two control layers are involved in this task: the plan layer and the reflex layer. For the visible disturbances (*i.e.* the trail), we make a plan before the disturbance arrives. For the invisible disturbances (*i.e.* small bumps), the control heavily relies on reflexes. The layered architecture manages slow or inaccurate hardware but facilitates learning, adaptation, augmentation, and teamwork. However, how to integrate the sensing/communication/computation/actuation component-hardware constraints with the plan / reflex layers in the human sensorimotor system has been unclear.

To answer these questions, a platform with the control elements (e.g. delay, quantization, disturbance, and feedback) is needed. Here we focus on our new, inexpensive and easy-to-use experimental platform that illustrates and tests the new theory that formalizes and explains these connections (Fig. 1). The details for the theory is reviewed in a companion paper [13]. Because it is impossible to noninvasively manipulate the internal delay and data rate in neural signaling, we find an alternative to add the external delay and quantization in actuator which generates an identical effect to simulate the hardware SATs in nerves.

Our sophisticated and versatile platform is a video game that safely simulates riding a mountain bike down a steep, twisted, bumpy trail using a standard display and gaming steering wheel. The virtual trail scrolls down a PC screen which can vary in speed, turns, and visual look ahead (and

[1] California Institute of Technology

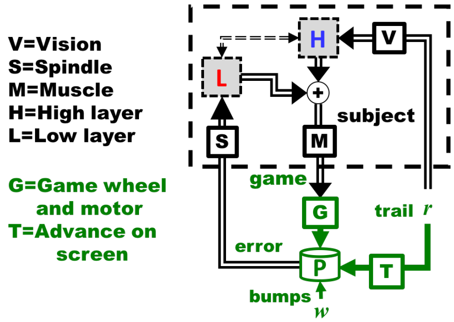

Fig. 1: Basic block diagram for theory and experimental platform with subject and gaming wheel with a motor. Each box is a component that communicates or computes and has potentially both delay and quantization, including within the game in G. The advance warning is also implemented on a computer screen with vision.

thus advanced warning or delay). Subjects can see the trail and turn the wheel to track it with minimum error, while an internal motor can torque the wheel to mimic invisible bumps in the trail. The trail planning and tracking use vision and advanced warning, while the bump disturbances are compensated for with delayed reflex feedback control. Subjects must cope with delays, quantization, disturbances, noise, and component sparsity, all captured in theory, in both their own nervous system and, to more thoroughly verify the theory noninvasively, when added in the experimental hardware.

II. Platform and models

A. Framework of the platform

The video-game platform simulates riding a mountain bike down a steep, twisting, bumpy trail using a standard display and gaming steering wheel. It includes a lower-layer reflex feedback control loop (L in Fig. 1) in charge of the unseen bump disturbance w, and a high-layer advanced plan loop (H in Fig. 1) to see the incoming trail and make a control plan in advance. Notably, the internal delays and quantizations in neural signaling or muscle contraction encoded in the subject are impossible to control noninvasively. However, we find an alternative way to extensively add the external delay and quantizer with a limited data rate in the game. Specifically, the game platform allows us to manipulate the advanced warning or visual delay T_{vis} (T in Fig. 1), along with the action delay T_{act} and action quantizer Q_{act} with data rate R_{act} in the gaming wheel (G in Fig. 1). We haven't yet implemented the quantization Q_{vis} for the visual system. It remains one of our future goals.

While the effects of delay are a combination of internal and external delay, especially when the internal delay is relatively big, such as the delay of flash-evoked responses in the visual cortex (around $120ms$ in healthy subjects [18]), quantization effects are evaluated only considering the external quantizer. The external quantizer we set in the game has a relatively small data rate (up to 10 bits), whereas the data rates for neural signaling and muscle contraction are much higher, such as the 50ms inter-spike intervals in motor neuron [19]. Thus, because the data rate for internal neural signaling is so much higher than that of the external quantizer, we neglect the internal quantization to simplify the analysis.

In addition, the trail disturbance $r(t)$ and bump disturbance $w(t)$ can be designed for the game. The design of these disturbances largely depends on the bounds in the feedback control model. We have to separate the worst-case setting and the average-stochastic setting in the game due to the difference in disturbance (see in Section II-B), although the design of delay and quantization in these two settings are the same.

The output of the game platform is a file including the values of all these manipulated parameters, the control $u(t)$ and the error dynamics $x(t)$. We minimally train subjects to achieve stable performance. To eliminate the learning effects and uncontrollable variables during task switching, we exclude the first 10 seconds from the data analysis.

B. Model and the parameter settings in platform

We consider a simplified feedback control model with the system dynamics given by

$$x(t+1) = x(t) + w(t) + u(t) + r(t) \tag{1}$$

where $x(t) \in \mathbb{R}$ is the error dynamics, $r(t) \in \mathbb{R}$ is the trail disturbance and $w(t) \in \mathbb{R}$ is the bump disturbance, $u(t) \in \mathbb{R}$ is the control action. $r(t)$ and $w(t)$ are independent.

We characterize the impact of delay and data rate on system performance in two settings:

1) Worst-case setting: The uncertainty (noise or changes in the desired trajectory) is bounded. Thus, when we design a game to test the worst-case performance, we set $r(t)$ and $w(t)$ to be periodic to satisfy that the disturbance is infinity-norm bounded, and we analyze the infinity norm of error $\|x\|_\infty$.

2) Average-case setting: The uncertainty is drawn from some distributions. For the average-case experiment design, the time at which there is a turn and/or a bump is decided by Gaussian distribution, and we analyze the mean squared error (MSE) with $\frac{1}{T}\left[\sum_{t=1}^{T} x(t)^2\right]$.

The worst-case performance is more relevant in risk-averse situations, while the average-case performance is more relevant in risk-neutral situations. An example of a risk-averse situation is riding a mountain bike on a cliff because staying on the cliff is a must for survival even given the worst possible uncertainty. This analysis framework is conceptually simpler when compared to the stochastic case. The worst-case control has been reported in the existing literature for risk-sensitive sensorimotor control [20–23]. On the other hand, an example of a risk-neutral situation is riding a mountain bike on a broad field because of no risk of falling out of the field. The control goal, therefore, is to minimize errors from the desired trajectory. To be noted, although the platform supports both the worst-case setting and average-case setting to test different theoretical bounds, we only present the results from the worst-case setting in this paper.

C. Goal of the platform and the experiment design

The experimental platform has two primary goals: 1) to test our new theory which connects the component hardware SATs in control loops with system-level SATs. 2) to serve future sensorimotor control research and educational applications.

1) Validate the new theory: Through manipulating the external delay and quantization (T_{vis}, T_{act} and Q_{act}), and designing the disturbance, $r(t)$ and $w(t)$, the platform allows us to study the impact of component hardware features on system-level performance. In the worst-case setting, we designed the trail changes in the trajectory $r(t)$ and the bump disturbances $w(t)$ to be ∞-norm bounded *i.e.* $\|r\|_\infty \leq 1$, $\|w\|_\infty \leq \epsilon$. Particularly, we will resolve the following three questions with specific game designs.

- What are the effects of vision delay and action delay in a single control loop? (*Section III-A*)
- How do component hardware SATs in layers impose the system-level SATs? (*Section III-B*)
- How does the layering architecture effectively help sensorimotor control? (*Section III-C*)

Due to space constraints, we present a more detailed description of the game settings in the extended version of this paper [24].

2) Beneficial to research and education: The platform is flexible to program, and easy to extend to other tasks and games. Both the source code and the executable file are open-access in our website. It has a user-friendly graphic user interface (GUI) for naive users. It also supports the higher-level individualized task design. Our platform allows users to design specific disturbances and noise in planning/reflex control loops, to add delay and quantization in both visual sensing and action output. It will also provide potential measures to quantify system performance (x), and actuator output (u). These are essential elements in a control system, and its capability of analyzing them suggests its potential use in control education. We hope the game platform will have versatile applications for both research and education.

III. Experiments and Interpretations

A. Delay in sensorimotor control

Delay is an essential topic for both control engineering and neuroscience [18], [25]. Delay in human sensorimotor control is inevitable, existing in each step: sensing, communication, computation, and actuation. In this section, we first overview our theory characterizing the impact of delay [13] in Section III-A.1 and then present our experiments that investigate how visual delay and action delay influence human sensorimotor control performance in Section III-A.2.

1) Simple model for delay in control: In this section, we study the impact of delay in sensing (vision) or actuation (muscle). We model the control action by

$$u(t+T) = \mathcal{K}(x(0:t), r(0:t), u(0:t+T-1)). \quad (2)$$

The game starts with zero initial condition, *i.e.* , $x(0) = 0$. The controller $\mathcal{K}$ generates the control command $u(t)$ using the full information on the histories of state, disturbance, and control input. The control command is executed with delay $T \geq 0$.

Sensorimotor control in the risk-aware setting [23] motivates the use of L_1 optimal control, and as such, our goal is to verify the following robust control problem:

$$\inf_{\mathcal{K}} \sup_{\|r\|_\infty \leq 1} \|x\|_\infty \quad (3)$$

This problem admits a simple and intuitive solution. In particular, the optimal cost is given by

$$\inf_{\mathcal{K}} \sup_{\|r\|_\infty \leq 1} \|x\|_\infty = T. \quad (4)$$

This optimal cost is achieved by the worst-case control policy $u(t+T) = -r(t)$, which yields

$$\inf_{\mathcal{K}} \sup_{\|r\|_\infty \leq 1} \|u\|_\infty = 1. \quad (5)$$

Interestingly, the control effort is not a function of the delay.

2) Trail with delay game and experimental findings: The experimental results for the trail with delay game (see Section II-C for the game design) are illustrated in Fig. 2.

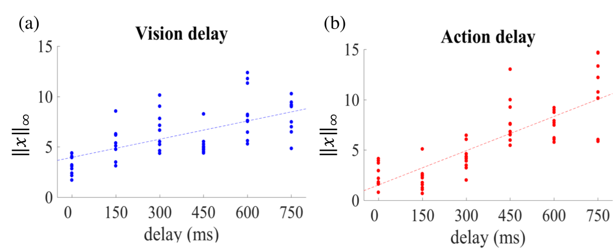

Fig. 2: The experimental results for effects of delay in visual feedback and in the action, respectively. We present the performance $\|x\|_\infty$ for each 2s.

The one-way analysis of variance (ANOVA) for $\|x\|_\infty$ shown in Fig. 2(b) found significant effects of T for both vision delay ($F = 11.49$, $p < 1e-7$) and action delay ($F = 28.66$, $p < 1e{-}14$). It means that the error significantly increases with an increase of delay. The generalized linear model was applied to fit the data. It showed that $\|x\|_\infty = 0.006T + 3.956$ for the vision delay, and $\|x\|_\infty = 0.011T + 1.528$ for the action delay. The results indicate that the error has a linear relation with the delay in a control loop. It is well in line with the prediction from the theoretical model in Eq.(4). Although the theory doesn't distinguish the impacts of vision delay from those of action delay on performance, the experimental data does show a smaller slope with increases in vision delay. This might be because the design of the trail is not random, and the subject can partially predict the future trajectory.

B. Speed-accuracy tradeoffs in a single layer

The impact of the SATs in nerve signaling on the SATs in sensorimotor control performance has been characterized in our previous paper [26] and is reviewed in our companion paper [13]. In this section, we validate the theoretical results using a mountain-bike game in the experimental platform. In this section, we first overview our theory characterizing the impact of delay and quantization [13] in Section III-B.1 and then present our experiments that investigate how they influence human sensorimotor control performance in Section III-B.2.

1) Model connecting the component-level and system-level SATs: Here, we demonstrate the impact of delay and quantization in the sensorimotor control loop. In the game, the control action is generated by the following feedback loop with communication constraints; we therefore model:

$$u(t+T) = Q(K_t(x(0:t), u(0:t-1))). \tag{6}$$

where $K_t : (\mathbb{R}^{t+1}, \mathbb{R}^t) \to \mathbb{R}$ is a controller, and $Q : \mathbb{R} \to S$ is a quantizer with data rate $R \geq 1$, *i.e.* S is a finite set of cardinality 2^R. Here, the net delay is composed of internal delays in the human sensorimotor feedback and the delays externally added. The disturbance is ∞-norm bound and, without loss of generality, $\|r\|_\infty \leq 1$.

The worst-case state deviation is lower-bounded by [26]:

$$\sup_{\|r\|_\infty \leq 1} \|x\|_\infty \geq T + \frac{1}{2^R - 1} \tag{7}$$

and the minimum control effort is given by

$$\sup_{\|r\|_\infty \leq 1} \|u\|_\infty \geq \left(1 + \frac{1}{2^R - 1}\right)\left(1 - \frac{1}{2^R}\right) \tag{8}$$

Measurements (7) and (8) can be recorded by our platform.

2) Experiments showing SATs in a simple control loop: We present the experimental results both for the SATs test in the planning layer and the reflex layer (Fig. 3). Both layers show system-level tradeoffs resulting from the hardware tradeoffs. Although $\|x\|_\infty$ heavily relies on the worst performance (one mistake would largely elevate the error measurement), our results surprisingly show that the effects of the T only test and the Q only test is roughly equivalent to the sum of the infinity error from the T and Q test after systematically subtracting the intrinsic error. It validates the Eq. (7) in theory that the control error in T and R is the sum of T term and R term.

Observed from the data, the tradeoff sweet spot for the planning layer is $R = 4$ bits and $T = -200$ ms when the component tradeoff is $T = 200(R-5)$, whereas the tradeoff sweet spot for the reflex layer is $R = 3$ bits and $T = 400$ ms when the component tradeoff is $T = 200(R-1)$. To be noted, the component level tradeoffs are not the same in the plan layer and the reflex layer in the experiments, because we try to test both the advanced warning and delay in the plan layer. One future objective is to test the plan layer with the same component level tradeoff as the reflex layer.

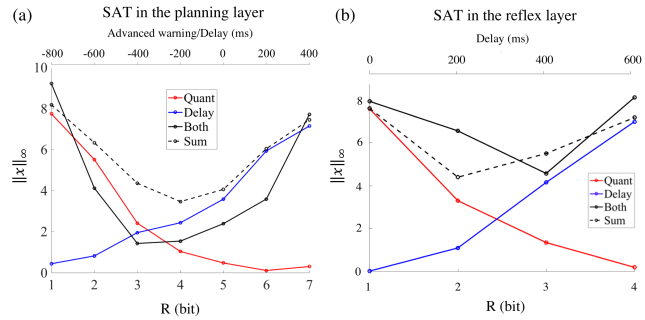

Fig. 3: The experimental results for SATs in the planning layer (a) and reflex layer (b).

Fitting the data with the quantization term in Eq.(7), $\|x\|_\infty = \alpha/(2^R - 1)$, we obtain $\alpha = 8.80$ for the higher-layer plan loop, and $\alpha = 7.83$ for the lower-layer reflex loop in the quantization only test. It indicates that the plan layer is more sensitive to the limited data rate, compared to the reflex layer. Fitting the data with the delay term in Eq.(7), $\|x\|_\infty = \beta T$, we obtain $\beta = 0.12$ for the reflex layer, and $\beta = 0.04$ for the plan layer, suggesting that the performance of the reflex layer is more sensitive to the delay. An increased delay of T has deleterious impacts on the bump rejection.

C. Effective layered architecture in sensorimotor control

Human sensorimotor control is an extremely robust system. It involves multiple control loops in layered architectures, rather than a simple one, as well as a huge diversity of component hardware to manage slow or inaccurate hardware [17]. The effectiveness of layered architectures can be observed in many sensorimotor control tasks. To test the robust control in multiple control loops and the effective layered architecture in human sensorimotor control, we have implemented the 'bump and trail' scenario in our platform, which mimics riding a mountain bike down a steep, twisting, bumpy trail. In this section, we review our theory showing the impact of delay and quantization in a layered system [13] in Section III-C.1 and present our experimental results to verify them in Section III-C.2.

1) Theoretical model for the layered architecture: In the game, the system dynamics in Eq.(1) have the ∞-norm bounded disturbance, with $\|r\|_\infty \leq 1$, $\|w\|_\infty \leq \epsilon$. The actuation action is generated through two controllers, a lower-layer reflex loop mainly in charge of $w(t)$ (e.g. human reflection responding to the bumps), and a high-layer advanced plan loop mainly in charge of $r(t)$ (e.g. our planning for the future road). The theory considers two types of layered architectures: with/without sharing the disturbance information between the two controllers. The layered architecture with shared information is defined by

$$\begin{aligned} &u_L(t+T_L) = L(x(0:t), u(0:t+T_L-1), w(0:t))) \\ &u_H(t+T_H) = H(x(0:t), u(0:t+T_H-1), r(0:t+T_a)) \\ &u(t) = u_L(t) + u_H(t). \end{aligned} \tag{9}$$

Here, L, H are the controllers, and Q_L, Q_H are the quantizers with data rates R_L, R_H, respectively.

The layered architecture without shared information is defined by

$$\begin{aligned} &u_L(t+T_L) = L(u(0:t+T_L-1), w(0:t))) \\ &u_H(t+T_H) = H(u(0:t+T_H-1), r(0:t+T_a)) \\ &u(t) = u_L(t) + u_H(t). \end{aligned} \tag{10}$$

We pose the robust control problem as follows:

$$\inf \sup_{\|r\|_\infty \le \epsilon, \|r\|_\infty \le 1} \|x\|_\infty, \tag{11}$$

where the infimum is taken over the control policies of the form (9) or (10). The worst-case state-deviation of the system with shared information is given by

$$\left\{ T_\ell + \frac{1-2^{-R_\ell(T_h-T_\ell)}}{2^{R_\ell}-1} + \frac{1}{2^{R_\ell+R_h}-1} \right\}(1+\epsilon). \tag{12}$$

The worst-case state-deviation of the system without shared information is given by

$$\left\{ T_\ell + \frac{1}{2^{R_\ell}-1} \right\}\epsilon + T_h + \frac{1}{2^{R_h}-1}. \tag{13}$$

The experiments allow to test whether there is shared disturbance information between two layers or not.

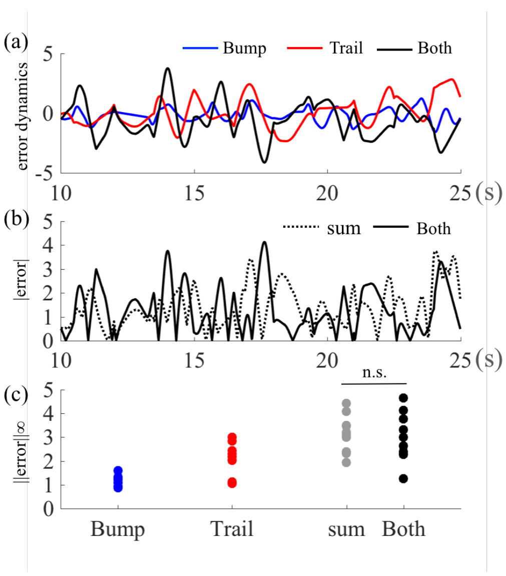

Fig. 4: The experimental results bump and trail task: (a) the error dynamics from bump only task, trail only task, both task; (b) absolute error; (c) infinity norm error. One dot denotes the infinity norm error in 2 seconds.

2) Bump and trail dual task to test the plan/reflex layering: Our experimental results (Fig. 4) demonstrate that the error from the bump and trail session positively correlated with the sum of the error from the bump only session and the trail only session (Pearson correlation, $R = 0.57$, $p < 0.0001$), and they showed no significant difference (paired t-test, $t = 0.21$, $p = 0.83$). The results confirm that the sensorimotor control can be decomposed to a reflex layer and an advanced plan layer. It fits the theory with no shared information case in Eq.(13), rather than Eq.(12). It indicates that the plan layer and the reflex layer can be decomposed and analyzed separately. In this case, the multitasking performance is equal to, not bigger than, the sum of every single task (*i.e.* tracking the trail and rejecting the bump). On the contrary, people usually multitask badly due to dual-task inference [27]. This excellent bump and trail dual-task performance can be explained by the effective layering architecture in human sensorimotor control. The two feedback loops jointly control the error $x(t)$ according to Eq.(10), where the advanced plan layer handle the disturbance $r(t)$ from the curvature of the desired trajectory and the delayed reflex layer due to road bumps $w(t)$.

IV. Discussion

In this paper, we present an experimental platform for studying the layering architecture and tradeoffs in human sensorimotor control. Our experiments focused on the worst-case scenario because of its theoretical simplicity, but it can be easily applied to the average-case scenario.

A. Neuroscience perspectives of experiment and platform

We demonstrate that the system-level tradeoffs in human sensorimotor control result from SATs at component-hardware level. We show clear component-level DSSs to favor the system-level SATs cross control layers (Fig. 3). Furthermore, the results showed that the reflex layer is sensitive to the delay, whereas the advanced plan requires higher accuracy, which is consistent with our theory. Our experiments illustrate the theory that connects the system and component level tradeoffs.

Moreover, the experimental data for the bump and trail dual task shows that optimal controllers with separate layers for trails and bumps can multiplex these tasks well, with errors in the bump with trail task roughly equivalent to, not bigger than, the sum of the errors in the trail only and bump only task, and healthy human subjects achieve this with minimal training (Fig. 4). In contrast, humans multitask badly on the tasks that dont naturally distribute across layers, regardless of whether the tasks are simple (e.g., perceptual discriminations) [27] or complex (e.g., driving and talking on the cell phone) [28]. Our experiments support the theory that the effectively layer sensorimotor control with appropriate diversity in components to achieve systems both fast and accurate despite being built from components that individually are not.

B. Educational value of the platform

Our theory and platform are only the beginning of a new regime. We serve much broader interests of SATs and layered architectures in the brain, from control engineers to neuroscientists to system biologists.

Robust control theory is a powerful tool in the investigation on the effect of noise, disturbances, and other uncertainties in system performance [29]. However, the impact of such theory is limited by its technical accessibility. Given the platform, we can easily study and compare the settings with delay and quantization with those of delay or quantization. This clear separation of constraints in the feedback loop can help explain the basics of control of dynamical systems and allow us to demonstrate how the plant instability, actuator saturation, and unstable zero dynamics impact our sensorimotor system.

C. Future directions

Our theory and platform are only a begining of a new regime. We serve much broader interests of SATs and layered architectures in brain, from control engineers to neuroscientists and system biologists. Along this line, we would like to propose some potential research questions that could be studied with our experimental platform in the future:

- Whether is the optimal control policy obtained the control theory applied in real practice? In this paper, we only examined the system performance. It is valuable to investigate the control policy in human sensorimotor control.
- How do the tradeoffs in speed / accuracy / saturation / energy cost in muscles benefit from the DSSs?
- How does the SATs in the high-layer plan/decision making support human/animals' decision strategies acorss complex environment under uncertainty, limited information, and risks?
- How does the human sensorimotor system tolerate the noise in control loops?
- What are the effects of learning/adaptation in different control layers? Do the tradeoffs exist between fast learning and fast forgetting, between efficiency and plasticity?

In summary, we provide a platform that is cheap, easy to use, and flexible to program. We use this platform is used to validate our theory presented in a companion paper [13]. The code to built the platform can be downloaded at `https://github.com/Doyle-Lab/WheelCon`. The platform can potentially be used in both future research and education.

References

[1] A. J. Nagengast, D. A. Braun, and D. M. Wolpert, "Risk sensitivity in a motor task with speed-accuracy trade-off," *Journal of neurophysiology*, vol. 105, no. 6, pp. 2668–2674, 2011.

[2] R. P. Heitz, "The speed-accuracy tradeoff: history, physiology, methodology, and behavior," *Frontiers in neuroscience*, vol. 8, p. 150, 2014.

[3] R. Bogacz, E.-J. Wagenmakers, B. U. Forstmann, and S. Nieuwenhuis, "The neural basis of the speed-accuracy tradeoff," *Trends in neurosciences*, vol. 33, no. 1, pp. 10–16, 2010.

[4] R. Plamondon and A. M. Alimi, "Speed/accuracy trade-offs in target-directed movements," *Behavioral and brain sciences*, vol. 20, no. 2, pp. 279–303, 1997.

[5] P. Sterling and S. Laughlin, *Principles of neural design*. MIT Press, 2015.

[6] C. J. De Luca and Z. Erim, "Common drive of motor units in regulation of muscle force," *Trends in neurosciences*, vol. 17, no. 7, pp. 299–305, 1994.

[7] V. Brezina, I. V. Orekhova, and K. R. Weiss, "The neuromuscular transform: the dynamic, nonlinear link between motor neuron firing patterns and muscle contraction in rhythmic behaviors," *Journal of neurophysiology*, vol. 83, no. 1, pp. 207–231, 2000.

[8] S. Li, C. Zhuang, M. Hao, X. He, J. C. Marquez Ruiz, C. M. Niu, and N. Lan, "Coordinated alpha and gamma control of muscles and spindles in movement and posture," *Frontiers in computational neuroscience*, vol. 9, p. 122, 2015.

[9] P. M. Fitts and J. R. Peterson, "Information capacity of discrete motor responses." *Journal of experimental psychology*, vol. 67, no. 2, p. 103, 1964.

[10] R. W. Soukoreff and I. S. MacKenzie, "Towards a standard for pointing device evaluation, perspectives on 27 years of fitts law research in hci," *International journal of human-computer studies*, vol. 61, no. 6, pp. 751–789, 2004.

[11] J. Freeston and K. Rooney, "Throwing speed and accuracy in baseball and cricket players," *Perceptual and motor skills*, vol. 118, no. 3, pp. 637–650, 2014.

[12] R. van den Tillaar and A. Ulvik, "Influence of instruction on velocity and accuracy in soccer kicking of experienced soccer players," *Journal of motor behavior*, vol. 46, no. 5, pp. 287–291, 2014.

[13] Y. Nakahira, Q. Liu, N. Bernat, T. Sejnowski, and J. C. Doyle, "Theoretical foundations for layered architectures and speed-accuracy tradeoffs in sensorimotor control (extended version)," http://users.cms.caltech.edu/ ynakahir/.

[14] D. B. Lockhart and L. H. Ting, "Optimal sensorimotor transformations for balance," *Nature neuroscience*, vol. 10, no. 10, p. 1329, 2007.

[15] G. E. Loeb, I. E. Brown, and E. J. Cheng, "A hierarchical foundation for models of sensorimotor control," *Experimental brain research*, vol. 126, no. 1, pp. 1–18, 1999.

[16] J. Blinks, R. Rüdel, and S. Taylor, "Calcium transients in isolated amphibian skeletal muscle fibres: detection with aequorin." *The Journal of Physiology*, vol. 277, no. 1, pp. 291–323, 1978.

[17] J. C. Doyle and M. Csete, "Architecture, constraints, and behavior," *Proceedings of the National Academy of Sciences*, vol. 108, no. Supplement 3, pp. 15 624–15 630, 2011.

[18] A. Halliday, W. McDonald, and J. Mushin, "Delayed visual evoked response in optic neuritis," *The Lancet*, vol. 299, no. 7758, pp. 982–985, 1972.

[19] W. H. Calvin and C. F. Stevens, "Synaptic noise and other sources of randomness in motoneuron interspike intervals." *Journal of neurophysiology*, vol. 31, no. 4, pp. 574–587, 1968.

[20] P. Whittle, "Risk-sensitive optimal control," 1990.

[21] H. Nagai, "Bellman equations of risk-sensitive control," *SIAM journal on control and optimization*, vol. 34, no. 1, pp. 74–101, 1996.

[22] T. D. Sanger, "Neuro-mechanical control using differential stochastic operators," in *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*. IEEE, 2010, pp. 4494–4497.

[23] ——, "Risk-aware control," *Neural computation*, 2014.

[24] Q. Liu, Y. Nakahira, A. Mohideen, S. Choi, and J. C. Doyle, "Experimental and educational platforms for studying architecture and tradeoffs in human sensorimotor control (extended version)," http://users.cms.caltech.edu/ ynakahir/.

[25] W. M. Smith and K. F. Bowen, "The effects of delayed and displaced visual feedback on motor control," *Journal of motor behavior*, vol. 12, no. 2, pp. 91–101, 1980.

[26] Y. Nakahira, N. Matni, and J. C. Doyle, "Hard limits on robust control over delayed and quantized communication channels with applications to sensorimotor control," in *Decision and Control (CDC), 2015 IEEE 54th Annual Conference on*. IEEE, 2015, pp. 7522–7529.

[27] H. Pashler, "Dual-task interference in simple tasks: data and theory." *Psychological bulletin*, vol. 116, no. 2, p. 220, 1994.

[28] P. E. Dux, M. N. Tombu, S. Harrison, B. P. Rogers, F. Tong, and R. Marois, "Training improves multitasking performance by increasing the speed of information processing in human prefrontal cortex," *Neuron*, vol. 63, no. 1, pp. 127–138, 2009.

[29] Y. P. Leong and J. C. Doyle, "Understanding robust control theory via stick balancing," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*. IEEE, 2016, pp. 1508–1514.