The effects of uncertain context inference on motor adaptation

Cuevas Rivera, Darío^{1,2} and Kiebel, Stefan J.^{1,2}

¹Chair of Neuroimaging, Faculty of Psychology, Technische Universität Dresden, 01062 Dresden, Germany.

²Centre for Tactile Internet with Human-in-the-Loop (CeTI)

Humans have been shown to adapt their movements when a sudden change to the dynamics of the environment is introduced, a phenomenon called motor adaptation. If the change is reverted, the adaptation is also quickly reverted. Human are also able to adapt to multiple changes in dynamics presented separately, and to be able to switch between adapted movements on the fly. Such switching relies on contextual information which is often noisy or misleading, which affects the switch between adaptations. In this work, we introduce a computational model to explain the behavioral phenomena effected by uncertain contextual information. Specifically, we present a hierarchical model for motor adaptation based on exact Bayesian inference. This model explicitly takes into account contextual information and how the dynamics of context inference affect adaptation and action selection. We show how the proposed model provides a unifying explanation for four different experimentally-established phenomena: (i) effects of sensory cues and proprioceptive information on switching between tasks, (ii) the effects of previously-learned adaptations on switching between tasks, (iii) the effects of training history on behavior in new contexts, in addition to (iv) the well-studied savings, de-adaptation and spontaneous recovery.

Introduction

It has been shown that humans can adapt motor commands to counteract changes in the dynamics of the environment and their own bodies, such as performing reaching movements with a weight attached to the wrist. This is known as motor adaptation. Moreover, human participants have been shown to adapt to different, even opposing, such changes during the course of a single experiment (Gandolfo, Mussa-Ivaldi, & Bizzi, 1996; Shadmehr & Brashers-Krug, 1997). Additionally, humans have been shown to dynamically switch between different learned adaptations (Davidson & Wolpert, 2004; Ethier, Zee, & Shadmehr, 2008; Lee & Schweighofer, 2009).

By introducing blocks of trials in which body dynamics are altered (e.g. a mechanical arm exerts a force on the participant's hand), experimenters are able to observe motor

adaptation through the lens of motor error. Across many different motor adaptation experiments (e.g. Davidson & Wolpert, 2004; Gandolfo et al., 1996; Shadmehr & Mussa-Ivaldi, 1994), well-established phenomena have been observed: (i) the ability to recall previously-learned skills, called savings; (ii) the ability to return to unmodified dynamics, termed de-adaptation; (iii) the interference in motor learning between opposing manipulations in dynamics, called anterograde interference; (iv) spontaneous display of behavior consistent with a previously-learned adaptation, during trials where errors are forced to be zero, called spontaneous recovery.

To explain these phenomena, a number of computational models have been introduced, which adapt their motor commands after observing motor errors. The most studied are linear learners (Forano & Franklin, 2020; Scheidt, Dingwell, & Mussa-Ivaldi, 2001; Smith, Ghazizadeh, & Shadmehr, 2006), but Bayesian accounts have also been presented, providing an alternative explanation for savings and quick de-adaptation in the form of switching between forward models (Körding & Wolpert, 2004; Oh & Schweighofer, 2019).

While these general models of adaptation explain the most common phenomena observed in experiments, other known phenomena remain outside of their scope. For example, it is known that adaptation rate is reduced in situations where the environment is unstable and unpredictable (Herzfeld, Vaswani, Marko, & Shadmehr, 2014), or situations in which errors are small (Marko, Haith, Harran, & Shadmehr, 2012) or adaptations slowly introduced (Huang & Shadmehr, 2009). Action selection has also been found to depend on the history of adaptations learned (Davidson & Wolpert, 2004; Vaswani & Shadmehr, 2013).

In this work, we show that these unexplained phenomena can be explained in terms context inference, an idea that has recently been gaining traction in different fields (e.g. Hunter & Daw, 2021; Sanders, Wilson, & Gershman, 2020). We build on models such as MOSAIC (Wolpert & Kawato, 1998) and Bayesian switching models (Körding & Wolpert, 2004; Oh & Schweighofer, 2019) by expanding the concept of context and introducing an explicit mechanism for inferring the context based on Bayesian inference. In the proposed model, the context encompasses all the information relevant to making decisions, both environmental and internal to the agent making these decisions. To infer the current context, our model integrates several sources of information, including the predictions made by its internal forward models, an estimate of the reliability of these predictions, direct sensory information (such as visual contextual cues), proprioceptive information, and feedback provided by the environment.

We show that the dynamics of context inference provide a principled mechanistic explanation for the experimental phenomena outlined above, providing a single account for motor adaptation under changing contexts, while relying on known mechanisms for adaptation.

Results

We present a motor adaptation model in which an existing model of the environment is updated based on error signals produced by the need of adaptation. We expanded on

existing models (Körding & Wolpert, 2004; Oh & Schweighofer, 2019; Wolpert & Kawato, 1998) by adding an explicit component of context inference. As we will show, context inference guides the selection of the adequate forward model, its updating and the sampling of actions.

We simulated representative experiments from a number of experimental studies on motor adaptation to illustrate how our model explains these experimental findings using the dynamics of context inference. We will present these simulations alongside the experimental results from the representative studies and discuss in detail how context inference explains the experimental phenomena.

Before presenting our simulations, we will briefly go through the fundamentals of the model we present. We leave the mathematical details of the implementation to the Methods section.

Modeling context-dependent adaptation

The model we present has three main components: (1) context inference, (2) motor adaptation and (3) action selection. The processes defined by these component occur in this order, and each component informs the ones that follow.

Central to the model is the concept of context, defined in terms of the task to be performed, the variables of the environment that are relevant to perform the task, the forward models used by the decision-making agent to perform the task, and the update mechanisms necessary to adapt these forward models to the changing environment. Together, these elements allow the agent to make predictions on future observations when this context is active, and these predictions are used to infer the context. For example, when lifting an object of unknown weight, an agent might have learned one context for heavy objects and one for light objects. When observing an object to be lifted, the agent can use its size and texture to estimate the weight of the object, which in turn allows the agent to infer the appropriate context and, with it, decide how to lift the object.

We now briefly describe the three components of the model separately. For clarity, we start with the second component (i.e. motor adaptation), in order to introduce the terminology that we use throughout this work. For an overview, we show in Figure 1 how these three components affect each other. The full mathematical description of the model can be found in the Methods.

Motor adaptation. It is widely accepted that in order to adapt motor commands to a novel environment, we create and update internal forward models of the outcomes of control signals (Wolpert & Kawato, 1998). We used this principle and chose to make use of an exact Bayesian learner. To adapt to altered dynamics (e.g. a mechanical arm exerting a force on a participant’s hand), our model updates the parameters ϕ (see Methods) of a context-specific forward model. These updates are done according to Equation 1, using the difference between predictions made with the current parameters of the forward model, and

Figure 1. Schematic representation of the model. ζ represents the identity of the context as a categorical variable, ϑ are observations, c_t are motor commands at trial t , ϕ are the parameters of the forward model and s are the hidden states of the environment (e.g. the position of the hand). The arrows show flow of information, i.e. the dependencies in the update equations, as well as the flow of a single trial. Trial t starts at the top when an observation ϑ_t is received, from which the current context is inferred ($q(\zeta)$), which informs motor adaptation ($q(\phi)$), which informs action selection ($p(c)$). The update equations can be found in the Methods section. After an action is selected, a new observation ϑ_{t+1} is received and the process is repeated. The posteriors over ζ and ϕ are used as priors at the next trial.

the current observation. With this, a posterior estimate on ϕ , $q(\phi)$ in Figure 1, is calculated after each trial.

Context inference. To infer the adequate forward model to deploy, most previous models (e.g. Imamizu & Kawato, 2008; Oh & Schweighofer, 2019; Wolpert & Kawato, 1998) have focused on the elements of the context which are directly related to the forward models: the model that best predicts the current observation given the past is deemed to be the “correct” model, i.e. the model that best represents the current context. Imamizu et al. (2007) expanded on this idea by showing the effects of visual cues on model switching.

Following these ideas, our model performs context inference that integrates several sources of information, including those typical for experimental studies such as visual cues (Kim, Ogawa, Lv, Schweighofer, & Imamizu, 2015; Lee & Schweighofer, 2009), the place where the task must be carried out (Forano & Franklin, 2020; Shadmehr & Mussa-Ivaldi, 1994) or the start of a new block of trials (Ethier et al., 2008), in addition to proprioceptive information, outcome prediction errors and reward prediction errors. As with motor adaptation, the context is inferred as a probability distribution $q(\zeta)$ with exact Bayesian inference, with the updates given by Equation 4.

Action selection. Action selection is affected by the current context via the context-dependent forward model that is used. In our model, action selection is made by building a distribution over available actions which is a weighted sum of the distributions given by the existing forward models, where the weights are obtained by the context inference component

of the model.

In other words, a distribution $p(c_t)$ over all possible motor commands is created following Equation 9, from which the next action (motor command c_t) can be sampled and carried out.

Experimental results

In this section, we present experimentally-observed phenomena in three sections, and show that the dynamics of context inference provide a unifying explanation for all of them. In the first section, we discuss switches between contexts, and how slow context inference affects these switches. In the second section, we focus on the effects of volatility in the environment on motor adaptation. Finally, in the third section we discuss context inference during so-called error-clamp trials, and its effect on behavior. For each of the three sections, we selected one or two studies which are representative of the phenomenon being discussed.

For clarity, we first introduce necessary terminology that is typically used in experimental studies. As an example, we will use a typical motor adaptation task in which participants have to make reaching movements while holding the handle of a mechanical arm that exerts a curl force on the participant's hand. Depending on the trial, the mechanical arm might exert a curl force in a clockwise or counter-clockwise direction, or no force at all. Let us define the baseline context O as that in which the mechanical arm exerts no force. Contexts A and B can be defined as those with clockwise or counter-clockwise, respectively. Abusing notation, a usual statement is that $B = (-A)$, as the forces have the same magnitude but point in opposite directions. Similarly, one can define context A/2, with the same direction of adaptation as A, but half the magnitude. Finally, many experiments include a block of error-clamp trials at the end of the experiment, in which the mechanical arm forces the participant to make straight-line movements; we represent these with the letter E.

With this terminology, a typical experiment (e.g. Ethier et al., 2008) would have a block structure of O-A-B-E, or O-A-(-A)-E, which means that the participant goes through a block of trials with no external force applied (O), a number of trials with a clockwise curl force (A), a block with counter-clockwise forces (B), and finally a block with error-clamp trials (E). With repeated contexts (e.g. Oh & Schweighofer, 2019), an experiment can be described as $O_1-A_1-O_2-\dots$.

Savings and slow/fast switching. The term 'savings' refer to the ability to remember a previously-learned adaptation and apply it without having to re-learn it. Savings is almost universally observed in humans (Brashers-Krug, Shadmehr, & Bizzi, 1996; Medina, Garcia, & Mauk, 2001; Shadmehr & Brashers-Krug, 1997; Smith et al., 2006; Zarah, Weston, Liang, Mazzoni, & Krakauer, 2008). In an O-A-O-A experiment, for example, savings would express themselves in the second A block in the form of a much higher adaptation rate higher than that observed during the first A block. The related concept of quick de-adaptation occurs in A-O transitions, where participants switch back to baseline without having to re-learn it.

In this section, we discuss savings and de-adaptation in terms of switching between forward models. This has been modeled before (e.g. Imamizu et al., 2007; Oh & Schweighofer, 2019; Wolpert & Kawato, 1998). Here we show in addition that expanding forward model switching into context inference explains facets of savings that these previous studies did not, namely that savings are not immediate, but a relatively fast process that reflects context inference.

To show this, we examined multiple experimental studies in which savings are observed. We categorized these studies based on the amount of contextual information made available to participants: In some experiments (e.g. Kim et al., 2015; Lee & Schweighofer, 2009), the context is clearly revealed to the participant. We call these cued-context experiments. In other experiments, partial information is available to participants (e.g. Davidson & Wolpert, 2004; Zarahm et al., 2008) in the form of large prediction errors, partial contextual information or reward prediction errors; we refer to these as partially-cued experiments.

We selected three representative experiments from two studies (Kim et al., 2015; Oh & Schweighofer, 2019) which differ in the amount of contextual information available to participants. Kim et al. (2015) performed a cued-context visuomotor rotation experiment with three contexts with different rotation: no rotation, 40 degrees and -40 degrees. Participants performed shooting movements in blocks of trials with the same rotation. Importantly, a colored light identified the current context, making this a cued-context experiment. Consistent with our model, the authors found that switching was immediate and accurate. In Figure 2A we show the results of simulations with our model using the parameters of the task, as well as the experimental results from Kim et al. (2015). The correspondence between simulations and experimental results can be seen in the switches between contexts, i.e. when the solid black line (representing the true context) switches between 40, 0 and -40; in both panels (left and right), the blue line, representing the agent's adaptation, quickly follows the switches.

Critically, in Figure 2A it can be seen for the first four context switches (until about trial 250) that participants had not yet completely adapted to the rotation, as evidenced by the blue line not reaching as low or high as the black line during the first blocks. Switching between opposite adaptations was nevertheless immediate in simulations and experimental data. After around 500 trials, participants had adapted completely, as seen by the blue line reaching the heights of the black line, and switches continued to be immediate. These results are similar to those shown by Imamizu et al. (2007), where sensory cues of varying reliability effected immediate or slow contextual switches. To expand on these results, we now turn to performance feedback and its effects on switching behavior.

Oh and Schweighofer (2019) performed two partially-cued experiments with a visual rotation of 20 and 10 degrees, respectively. The results of their experiments can be seen in Figure 2B and 2C, alongside simulations with our model. Participants in the first experiment (Figure 2B) first learned the adaptation in A_1 . From the $A_1 - O_2$ transition onward, participants showed immediate switching (with a one-trial lag) between A and O (both ways), which can be seen in terms of the blue line (adaptation) closely following the switches in the black line (true rotation). In the second experiment (Figure 2C), switching happens

Figure 2. Savings and de-adaptation. Data from our simulations (left column) compared to data adapted from figure 2A by Kim et al. (2015) and figure 4A from Oh and Schweighofer (2019) (right column). Experimental data was, in all three experiments, averaged across all participants; in our simulations, a single simulation is shown. In all panels, the blue line represents the adaptation displayed by the agent as a function of trial number. The black lines represent the optimal adaptation, i.e. the size of the true visuomotor rotation during the task. (A) Experiment by Kim et al. (2015) with two visuomotor rotations (-40 and 40 degrees), in addition to baseline, with all possible transitions between A, (-A) and O. Participants must adapt to the rotations during shooting movements, and the blue line can be understood as the rotation of the participants' shooting movements to compensate for the task's rotation. When black and blue lines match, the participant perfectly counteracted the visuomotor rotation and reached the target. (B) An experiment by Oh and Schweighofer (2019) similar to A, but the contexts were not cued and only one context, with an adaptation of 20 degrees, is introduced. O-A and A-O transitions (i.e. adaptation and de-adaptation) can be observed. Blue and black lines are as in A. Different blocks with the same adaptation are differentiated by the subscript (e.g. A_1 and A_2 have the same 20 degree visuomotor rotation). (C) Same experiment as (B) but with a 10 degree adaptation.

more slowly, with adaptation lagging behind the switches in the real context, and slowly catching up. As can be seen in the left panels in Figure 2B and 2C, the same model parametrization produces fast, accurate switches when the adaptation is large (B), and slow, noisy switches when it is low (C). This difference is explained by our model in terms of the size of the adaptation: as the adaptation is smaller (10 degrees), it becomes more difficult to distinguish errors made by incorrectly inferring the context from the noise due to trial-to-trial variation in motor output. Because of this, the model requires more evidence (i.e. more trials) to infer a switch in contexts.

Note that the results from our simulations, from Oh and Schweighofer (2019) and from Kim et al. (2015) include both savings (O-A transitions) and de-adaptation (A-O transitions), both of which display the same characteristics and are explained by the same mechanism.

Uncertainty over contexts affects action selection. As with learning, our model selects actions (motor commands) based on context inference. If the identity of the current context is known, the forward model for this context is used to select the current action. However, if uncertainty over the context exists, the selected action is influenced by all the possible contexts, with a weight directly related to how likely each one of those contexts is (see Equation 9).

Experimental evidence supporting this view can be found in experiments with context switching. For example, Davidson and Wolpert (2004) reported a curl-force experiment in which participants had to switch from 3A to A in one group, with a block sequence A-3A-A-3A, and from -A to A in another group, with a block sequence A-(-A)-A-(-A). After A and 3A (or -A in the other group) had been learned in the first two blocks, the authors found that the switch from 3A to A was faster than that from -A to A. The authors interpreted this as evidence that switching between adaptations happens more quickly if it is in the same direction as the current adaptation (e.g. both counter-clockwise), and more slowly if they are in the opposite direction (e.g. clockwise to counter-clockwise).

In our context-inference account, the asymmetry is caused by the baseline context, which have a non-zero probability $p(\zeta_O|s_t...)$. When a new block of trials starts (e.g. in the transition from 3A to A), a switch is inferred by the model (given feedback after the first trial) and ζ_O becomes more likely (given that ζ_{3A} has been ruled out). Therefore, in these first trials, action selection has a component guided by the baseline model, in which no extra compensatory force is applied, effectively “pulling” adaptation towards zero (no compensatory force). In the first group, this initial pull towards zero accelerates the transition towards A because $3A > A > 0$, but in the second group, it slows down the switch because $A > 0 > -A$.

We simulated data with our model fitting the experimental structure in Davidson and Wolpert (2004). We show the results in Figure 3, alongside the experimental results from Davidson and Wolpert (2004). It can be seen that the agent exposed to the -A context shifts back to A more slowly than the one exposed to 3A, qualitatively reproducing the data from Davidson and Wolpert (2004). Note that in our simulations, the difference between the

Figure 3. Motor error when switching back to a previously learned adaptation. Following the block structure from Davidson and Wolpert (2004), the first 40 trials of the switch from 3A (yellow) or -A (blue) to A is shown (A) Simulations from our model, with the task parameters from Davidson and Wolpert (2004), running eight simulations per group, each simulation representing one participant. On the y axis, the error is shown in arbitrary units, related to centimeters through a monotonically increasing transformation with the same origin. The shaded area is the SEM across simulated runs (participants). (B) Data adapted from figure 4 by Davidson and Wolpert (2004), showing the error in the same trials as (A).

groups, as well as the variability within each group, is smaller than in the experimental data. However, the effect can be reproduced with any number of simulated runs (each one representing one participant), pointing towards a reliable effect.

Action selection in error-clamp blocks. During error-clamp blocks at the end of block sequences, participants' behavior can be divided in two phases: (1) Spontaneous recovery, during which participants' behavior is consistent with a previously-encountered context (e.g. in an O-A-E experiment, consistent with A); this phase, when present, is seen during the early trials of the E block. (2) A slow return to baseline, which can last as long as hundreds of trials (Brennan & Smith, 2015). However, the direction of the spontaneous recovery, its duration, the delay before it is observed, the speed of the return to baseline and the final asymptote of the adaptation vary greatly depending on the experiment (Brennan & Smith, 2015; Shmuelof et al., 2012; Smith et al., 2006; Vaswani & Shadmehr, 2013).

Because context inference integrates information from different sources, many experiments in which no intentional, overt contextual cues are available indeed contain contextual information that the participant can use to infer the context. For example, proprioceptive signals provide contextual information (Dizio & Lackner, 1995; Shadmehr & Mussa-Ivaldi, 1994). The sudden appearance of motor errors can itself be a cue for contextual change (Herzfeld et al., 2014) and even a pause between two trials could suggest a change in context Ethier et al. (2008).

In this section, we show how context inference can explain these different parameters of behavior by changing the way contextual cues mislead participants' context inference, which in turn influences action selection.

Vaswani and Shadmehr (2013) studied in detail human behavior during an error-clamp block in a shooting movement paradigm with a mechanical arm. The authors found that during an E block at the end of each experiment, there was a lag of a few trials (depending on participant) before their motor behavior changed from that of the previous block. After that, the exerted force slowly dropped towards zero throughout tens of trials, but never reaching values around zero. Participants were divided into four groups, each of which going through a different block sequence: (1) A-E, (2) O-A-E, (3) (-A/2)-A-E, and (4) (-A)-E. No pauses were made during the experiment nor were there any contextual cues, so transitions between blocks were not signaled to participants.

In Figure 4A, we show data simulated with our model, following the parameters of the experiment by Vaswani and Shadmehr (2013), and in Figure 4B we show the experimental plots adapted from Vaswani and Shadmehr (2013). The displayed adaptation is shown during the E trials for the three experimental groups in the experiment. It can be seen that group 1.3 (i.e. the participants who had learned in the -A/2 context in addition to A) more quickly recognized a change in context and lowers the force applied on the mechanical handle, as can be seen in the experimental data.

In Figure 4C, the inference over context is shown for each group separately. Context inference works reliably until the error-clamp trials start, which do not correspond to any of the known contexts. This causes the agent to infer the combination of some of the known contexts that best fits the observations. Depending on the contexts previously learned by the agent: groups 1.1 and 1.4 display the same behavior, where the previous context (A and -A, respectively) slowly dwindles. These agents will slowly lower the force applied. In contrast, group 1.3 has learned the additional -A/2 context, which has a non-zero posterior probability during E trials, pushing the agent's adaptation force more quickly towards zero. Group 1.2 behaves similarly to 1.1, with the exception that the baseline context, which was recently seen, plays a bigger role during E trials, making the agent reduce its force during E trials slightly more quickly than groups 1.1 and 1.4.

Context inference also explains the variability in the lag before adaptation starts to drop to zero after the E block starts. In the model, this lag is due to a period in which context inference has not "caught on" to the change of context; during this period, behavior remains consistent with that of the previously-observed block, as can be seen in the experimental data as well. To show this, we show in Figure 4D the drop in the inferred probability of the previous context for each simulated run (one line per run, 10 runs per group, all groups). The start and speed of the drop depend on the run (i.e. the participant), reproducing the different lags observed in experimental data.

Figure 4. Adaptation during error clamp trials. (A) Simulated adaptation during the error-clamp trials for the three groups of participants in Vaswani and Shadmehr (2013), using the same colors. The groups differ in the sequence of adaptations: 1.1 performed an A-E sequence; 1.2, O-A-E; 1.3, (-A/2)-A-E; and 1.4, (-A)-E. Following Vaswani and Shadmehr (2013), group 1.4 is not shown in A and B, as their behavior is identical to group 1.1. The solid line is the average across 10 runs (i.e. a group of 10 simulated participants) and the shaded area represents the standard deviation. The vertical dashed line is the start of the error-clamp trials. (B) Corresponding experimental data adapted from figure 2C by Vaswani and Shadmehr (2013). (C) Simulations: Inference over the current context, where contexts are color coded: black for baseline, orange for the counter-clockwise force, purple for the clockwise force and brown for counter-clockwise force with half strength. The lines represent the posterior probability of each context in every trial, while the background color represents the true context. An olive-colored background represents error-clamp trials. As in (A), solid lines represent the average across all runs and shaded areas represent the standard deviation. (D) Simulations: Visualization of the lag before a change in context is detected by the agent during the E trials. Each line represents one run (10 runs per group).

Discussion

We presented a hierarchical model for motor adaptation based on Bayesian inference, which not only learns to adapt its motor outputs by observing errors, but also infers the context in which these errors were committed and updates the forward model associated with that context.

We generalized the inference over the best-fitting internal models to the concept of context inference, which includes all sensory information, in addition to feedback provided by the experimenters or the task. With this, we showed how this seemingly small generalization provides a parsimonious explanation for several experimental phenomena across many experiments in motor adaptation that, up to now, had assumed the existence of multiple mechanisms.

We selected representative experimental studies that show the well-established effects of savings, quick de-adaptation, spontaneous recovery and the effects of sensory cues. Using our model, we showed how each of these effects can be explained by the specifics of context inference, which integrates all the available information (e.g. sensory cues, workspace location, reward and endpoint feedback), in some cases throughout many trials.

Our choice of an exact Bayesian learner greatly simplified our simulations, but we expect the results we presented here to hold using alternative models for motor adaptation, such as Kalman filter-based learners (e.g. Baddeley, Ingram, & Miall, 2003; Oh & Schweighofer, 2019) and other Bayesian implementations (e.g. Körding & Wolpert, 2004; Wolpert & Kawato, 1998). An exception is linear learners (e.g. Forano & Franklin, 2020; Lee & Schweighofer, 2009; Smith et al., 2006), as these models do not incorporate uncertainty in the estimate of the parameters of the forward models.

Further experimental evidence

In many cases, the context is not directly observable and context inference takes the form of an evidence-accumulating process that can take any amount of time to be sure of the context. It is in these cases where the effects of context inference are most noticeable. While many experiments exist that give probabilistic contextual information (e.g. Behrens, Woolrich, Walton, & Rushworth, 2007; Nassar, McGuire, Ritz, & Kable, 2019; Scholz & Schöner, 1999), evidence accumulation is not limited to these explicit cases. Indeed, as we noted in the Results section, many experiments inadvertently include partial contextual information used by participants.

The most direct secondary contextual information comes in the form of reward and endpoint feedback. Participants are told whether they obtained the desired reward at the end of a trial and are shown the end point of their movement. When a participants observe an unexpectedly large error, they can infer that the inferred context might be incorrect. This is the case of the experiments by Oh and Schweighofer (2019) shown in Figure 2B-C: if

the adaptation is high, changes in context produce errors much larger than those of motor variability, and a context switch is easily and immediately identified; if adaptation is low, the errors produced by context switching are closer in magnitude to motor variability and evidence accumulation is necessary.

The same rationale explains the results by Davidson and Wolpert (2004): motor learning, which in our model is modulated by context inference, is minimal for errors close to 2 and -2 (see their figure 2E). This is because an error of 2 or -2 signals that the participant incorrectly identified the context (as adaptation has a magnitude of 1).

A subtler source of information can be found in long pauses between blocks of adaptation trials, after which an unprompted partial return to baseline has been observed Ethier et al. (2008). This can be explained by context inference, as a long pause could prompt participants to infer that a switch had occurred, prompting participants to rely on their belief of the underlying probability of observing any of the known contexts, which is dominated by the previously observed context A, but now includes a component of the baseline O, as it is the most common one in everyday life.

Error-clamp (E) trials present another insight. If error is kept at zero, one could assume that participants would continue doing what they were doing before, as there is no reason (no observed error) to infer a change in context. However, this is almost never the case (e.g. Ethier et al., 2008; Forano & Franklin, 2020; Pekny, Criscimagna-Hemminger, & Shadmehr, 2011; Scheidt, Reinkensmeyer, Conditt, Rymer, & Mussa-Ivaldi, 2000; Smith et al., 2006; Vaswani & Shadmehr, 2013). Instead, participants slowly reduce their adaptation, often displaying spontaneous recovery (e.g. Smith et al., 2006). Our model provides a principled account of this behavior: the natural variability in participants' behavior lead them to expect errors, which clashes with the observed zero error. This prompts participants to re-evaluate their inferred context, which can partially activate a previously-observed context, as we showed in Figure 4. Pekny et al. (2011) found similar results, demonstrating that the duration of the previously-observed adaptation block also affects behavior in the E block. Additionally, Criscimagna-Hemminger and Shadmehr (2008) showed that introducing long periods before the E block begins lowers the initial force that participants exerted on the mechanical arm during the E block; longer periods of time make context inference revert to the prior expectation that a new baseline block begins, because participants are free to move their arm about during the pause.

In our account, if all information indicating a change in context is removed from the experiment, participants would continue to behave as they were in the previous block. Evidence for this can be seen in experiments 2 and 3 by Vaswani and Shadmehr (2013), where participants were shown random errors during E trials, with a variance matching that of previously observed motor commands. The authors showed that by matching the errors expected by participants, they eliminated the slow tapering off observed in most E blocks.

Model predictions and future work

The basic principle behind the results we presented is that context inference is a process that develops over time and that carries with it uncertainty. This uncertainty affects learning and behavior during motor adaptation, effecting phenomena that are directly observable during behavioral experiments.

Error-clamp as a known context. The inclusion of reliable sensory contextual cues (e.g. lights whose color uniquely identify a context) makes switching immediate, as in the experiments by Kim et al. (2015). We expect that the same effect would be observed in error-clamp trials. If the E block is learned by participants during training, it might still be difficult for them to infer that an E block has started, which would create delays similar to those in Figure 4. However, the model predicts that if a visual cue is introduced that identifies the E block, participants would immediately switch to their baseline behavior, no longer displaying spontaneous recovery, lag, nor the slow return to baseline. This immediate switch in the presence of contextual cues would persist even if endpoint feedback is manipulated as Vaswani and Shadmehr (2013) did.

Reducing the effect of volatility on learning. Experiments (Herzfeld et al., 2014) have shown an effect of the volatility of the environment (i.e. unpredictable switching between contexts) on measured learning rate. Our model predicts that this effect would be greatly reduced if reliable contextual cues were introduced: if a participant can infer the context of the current trial before any decision or observation has occurred, the learning rate would not be affected by the volatility of the environment.

If this prediction is confirmed, our model would additionally suggest that human participants do not revisit the learned adaptation of the previous trial when a new observation comes in. To see this, consider the following scenario from the experiments by Herzfeld et al. (2014): at trial t , the true context was B but the participant inferred context A, issued a motor command consistent with context A and then observed the outcome at $t + 1$. When the outcome is observed, it becomes clear to the participant that the context was B. Does the participant update the internal model of A or of B? According to our model, participants incorrectly update A and, upon learning of their mistake at trial $t + 1$, do not revert this update. If this were not the case, the volatility of the environment would have no effect even without contextual cues, as the context at trial t can almost always be identified at trial $t + 1$.

Multi-source integration. The model also predicts an effect reminiscent of multi-sensory integration (Ernst & Banks, 2002): in order to integrate contextual information from conflicting sources (e.g. probabilistic visual cues and noisy endpoint feedback), the weight placed on a source increases with its reliability. Such integration would manifest itself in experiments in which observations are noisy, as in the experiments by Körding and Wolpert (2004), in which the position of the finger was obscured and instead participants are shown a blurry cursor which was some times shifted from its real position. If the added observation noise gives evidence for a particular context (the true underlying context or another one) and a visual cue gave partial information for another context, the participants'

behavior would be more consistent with the most reliable source of contextual information.

Methods

Model description

We first describe the general principle behind the model, followed, in the next subsection, by the full mathematical description of each one of its components.

Before proceeding with the model description, let us introduce nomenclature:

- t : Discrete time at the moment an observation is made, before inference is done. Here taken as the trial number.
- ϑ_t : Observation. Note that this is assumed to be a noisy observation of the generalized state (see below)
- s_t : Hidden state, e.g. the position of the hand in reaching experiments. It does not depend on the context.
- c_t : Motor command (action) taken after observing ϑ_t and inferring the context.
- ζ_t : Categorical variable representing the context inferred after observing ϑ_t .
- ω_t : Contextual information. This term includes, e.g. visual cues, proprioceptive feedback, performance feedback.
- z_t : Generalized state, $z_t = \{s_t, \zeta_t\}$. Observed noisily as ϑ_t .
- m_i : Forward model for context ζ_i .
- ϕ_i : Parametrization of the forward model m_i . In adaptation experiments, this can be interpreted as the inferred size of the adaptation.

Context inference can work through context-specific (possibly unreliable) sensory cues, or through the dynamics of the system (e.g. expected outputs of motor commands), using prediction error as the force behind inference. In this case, the generative model can be written as follows:

$$p(z_t, \vartheta_t | z_{t-1}, c_{t-1}) = p(\vartheta_t | z_t) p(z_t | z_{t-1}, c_{t-1})$$

and inferring the current (hidden) generalized state is done with Bayes' theorem:

$$q(z_t | \vartheta_t, \vartheta_{t-1}, c_{t-1}) \propto p(\vartheta_t | z_t) p(z_t | \vartheta_{t-1}, c_{t-1})$$

Additionally, after having inferred the current state, motor adaptation is given by

$$q(\phi | z_t, z_{t-1}, c_{t-1}) \propto p(z_t | z_{t-1}, c_{t-1}, \phi) q(z_{t-1} | \dots) p(\phi) \quad (1)$$

where ϕ are the parameters of the internal model for the dynamics, also called the forward model, i.e. the function $f : (z_t, c_t) \rightarrow z_{t+1}$, which determines the beliefs of the agent regarding how the system evolves after having taken action c_t . $q(z_{t-1}|\dots)$ is the posterior probability calculated in the previous trial.

Having inferred the generalized state (including context), a motor command is sampled from a distribution:

$$p(c_t|s_t, \zeta_t) = \sum_i q(\zeta_i) p(c_t|s_t, \zeta_i) \quad (2)$$

where i runs over all contexts and $p(a_t|\dots m)$ is the probability distribution over actions prescribed by the forward model m_i for context ζ_i .

In the following sections, we present a specific mathematical implementation of the principles outlined above. This is the implementation used throughout the simulations in the Results section.

Mathematical implementation. The generalized state is given by:

$$z_t = \{s_t, \zeta_t\} \quad (3)$$

where s_t is, e.g. in a reaching task, the current position of the hand in Cartesian coordinates, where the origin is the starting point. We assume that motor commands are issued every Δt for simplicity. At the beginning of each time interval, the context is inferred combining the contextual information and the prediction error of the outcome of the previous motor command. This yields:

$$\begin{aligned} q(\zeta_t) &= q(\zeta_t|\omega_t) q(\zeta_t|s_t, s_{t-1}, c_{t-1}) \\ q(\zeta_t|\omega_t) &\propto p(\omega_t|\zeta_t) p(\zeta_t) \\ q(\zeta_t|s_t, s_{t-1}, c_{t-1}) &\propto p(s_t|\zeta_t, s_{t-1}, c_{t-1}) p(\zeta_t) \end{aligned} \quad (4)$$

The terms $p(\zeta_t)$ refer to the prior probability of the context at the beginning of the trial: at the first trial, this refers to prior beliefs over which contexts are more common/likely, which we call $p_0(\zeta_t)$. At each subsequent trial, it incorporates the belief that has so far been accumulated, given previous trials. Note that certain events, such as the start of a new block of trials (see main text) can return an agent to a state in which $p_0(\zeta_t)$ becomes relevant.

We do not directly model hand movements or eye position. Instead, we model systems in the motor error space: the position $s_t = 0$ represents a trial in which no motor error occurred. We make use of the fact that in most experiments, motor adaptation needs to happen in one of two directions (e.g. clockwise vs. counter-clockwise) to keep the error space one-dimensional. With this, $s_t > 0$ represents motor error in one direction (e.g. clockwise) and $s_t < 0$, error in the other direction. We made this choice because in this work we modeled different experiments with different state spaces, but with the commonality of (at most) two-dimensional error spaces. In order to model a specific experiment in its own physical space (e.g. hand positions in a reaching task), only the dynamical system $p(s_t|\zeta_t, s_{t-1}, c_{t-1})$ would need to be adapted.

We assume that the decision-making agent has a baseline internal forward model m_O that maps motor commands to outcomes:

$$p(s_{t+1}|s_t, c_t, m = m_O) = f_O(s_t, c_t) \quad (5)$$

where c_t are motor commands (actions). We further assume that participants have learned f_O throughout their lives and use it to decide which motor command to issue under normal circumstances (i.e. no adaptation). When a perturbation is introduced, to which the participant must adapt, we assume they learn a new forward model f_m^ϕ , where ϕ represents the parameters of the function that participants must learn. For example, in a saccade task, ϕ represents the new gain. Note that ϕ depends on the context and participants may need to use and update multiple forward adapted models f_m^ϕ throughout one experiment.

Given a forward model f_m (or f_m^ϕ), action selection is done by building a probability distribution over possible actions c_t based on how likely they are to produce zero motor error. For example, in saccade experiments, motor commands are issued based on how likely they are to move the fovea to the target. These motor commands are issued to counteract intrinsic sources of error, such as muscle variability and exhaustion, and extrinsic sources of error, such as the experimental manipulations on saccade gain. Thus, the expected observed error can be written as:

$$p(\epsilon|c_t) \propto p(\epsilon_{\text{intrinsic}}|c_t) + p(\epsilon_{\text{extrinsic}}|c_t) \quad (6)$$

and a motor command can be chosen to maximize the chance of observing zero error:

$$p(c_t) = \text{argmin}_{c_t} p(\epsilon|c_t) \quad (7)$$

For simplicity, we assume that motor errors are Gaussian, such that:

$$p(\epsilon|c_t) = N(\mu_{\text{extrinsic}}, \sigma_\epsilon) \quad (8)$$

where σ_ϵ incorporates both intrinsic and extrinsic sources of error and is a free parameter of the model. The mean of the error is assumed to have no intrinsic component.

With this, a forward model will produce a motor command with outcomes centered at $-\mu_{\text{extrinsic}}$, which counteracts the externally-induced errors.

To complete the action-selection picture, we return to context inference. As discussed above, action selection is made via a weighted sum of the motor commands produced by each forward model:

$$p(c_t|\dots) \propto \sum_i q(\zeta_i) p(c_t|\zeta_i, \dots) \quad (9)$$

From this distribution, an action is sampled at each trial.

To update the agent's estimates on the parameters of the forward models, we make use of Bayes' theorem with conjugate priors. This approach has the advantage of not only greatly simplifying computational calculations, but also presenting update equations with intuitive appeal.

In the following equations, we drop the dependence on the context and observations for notational simplicity:

$$q(\phi|s_{1:t}) \propto p(s_{1:t}|\phi)p(\phi) \quad (10)$$

where $s_{1:t}$ are the observed states up until time t . The likelihood is a Gaussian:

$$p(s_t|\phi) = N(\mu_s, \sigma_s) \quad (11)$$

$$\mu_s = \phi + s_t \quad (12)$$

$$\sigma_s = \xi_0 1 \quad (13)$$

This likelihood is the probability of observing s given the previous estimate of the parameters of the forward model, ϕ_{t-1} , and given the dynamics of the system (see below).

To update ϕ , we assumed that for each context, the agent's belief over the magnitude of the adaptation is given by a normal distribution:

$$p(\phi) = N(\mu_\phi, \sigma_\phi) \quad (14)$$

where μ_ϕ and σ_ϕ are parameters to be estimated at each trial. Before the experiment begins, the agent will have a prior distribution over these parameters. A standard Bayesian approach is to choose NormalGamma priors:

$$p(\mu_\phi, \sigma_\phi) = NG(\mu_0, \nu_0, \alpha_0, \beta_0) \quad (15)$$

where $\mu_\phi^0, \nu_\phi^0, \alpha_\phi^0$ and β_ϕ^0 are free hyperparameters of the model.

The update equations for the magnitude parameters are given by:

$$q(\mu_\phi, \sigma_\phi|s_t, s_{t-1}, c_{t-1}) \propto p(s_t|s_{t-1}, c_{t-1}, \zeta_t)p(\mu_\phi, \sigma_\phi) \quad (16)$$

where $p(s_t|s_{t-1}, c_{t-1}, \zeta_t)$ is a Gaussian distribution centered around μ_c with standard deviation σ_c , which is a free parameter of the model. It is assumed that the standard deviation over outcomes (σ_c) is related to that of the parameter μ_ϕ as follows:

$$\sigma_\phi = \sigma_c / \nu_\phi \quad (17)$$

which makes ν_ϕ a parameter to estimate. Each one of the forward models can be updated at every trial, depending on context inference (see main text). The updated parameters for any one forward model i are given by:

$$\begin{aligned} \mu_\phi^{(t)} &= \frac{\nu_\phi^{(t-1)} \mu_\phi^{(t-1)} + q(\zeta_i) s_t}{\nu_\phi^{(t-1)} + q(\zeta_i)} \\ \nu_\phi^{(t)} &= \nu_\phi^{(t-1)} + q(\zeta_i) \\ \alpha_\phi^{(t)} &= \alpha_\phi^{(t-1)} + q(\zeta_i)/2 \\ \beta_\phi^{(t)} &= \beta_\phi^{(t-1)} + \frac{q(\zeta_i) \nu_\phi^{(t-1)}}{\nu_\phi^{(t-1)} + q(\zeta_i)} \frac{(s_t - \mu_\phi^{(t-1)})^2}{2} \end{aligned} \quad (18)$$

where $q(\zeta_i)$ is the posterior probability of context ζ_i from Equation 4. Note that when $q(\zeta_i) = 0$, the model i is not updated.

Interpreting the hyperparameters. μ determines the initial estimate of the adaptation, in the same units as the necessary adaptation. ν encodes how stable this hyperprior is: higher values (e.g. 10,000) all but guarantee that the hypermean will not change its value after observations; In principle, enough evidence should still modify it, but that would not happen during an experiment. Smaller values (i.e. ~ 1) make the hypermean follow evidence more freely. Note that as more observations are accumulated, ν becomes bigger and bigger, stabilizing the value of the hypermean.

The hyperparameters α and β have a more complex effect. Note that the mean of a gamma distribution is $1/(\alpha\nu)$; this mean is being used as the standard deviation of a Gaussian by the rest of the agent, which makes it an important measure of uncertainty. While setting the default hyperparameters, the values used are $\alpha = 0.5/\sigma_0$ and $\beta = 0.5$, where σ_0 is the *a priori* estimate of the standard deviation of the force exerted by the environment, which controls the initial learning rate. This makes the initial standard deviation equal σ_0 , which makes it consistent with the fixed-force model. The 0.5 values ensure that uncertainty is large at the beginning and is greatly reduced during the experiment, but never to a point where it is so small that it makes trial-to-trial variation in the environment surprising. Changing this 0.5 would make the standard deviation change more quickly, making the model more or less precise in its predictions, independently of the volatility of the mean of the adaptation (via the hypermean).

The baseline model defaults to different values that make it a lot more stable. The hyper-standard deviation of the mean is set to 10,000, which makes the mean entirely stable during the duration of the experiment. The values of α and β are fixed regardless of σ_0 such that the standard deviation is 0.001 (compared that to the size of the adaptations in mechanical arm experiments, around 0.0125), and the hyperparameters of the standard deviation are stable during the experiment.

Acknowledgment

Funded by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany’s Excellence Strategy – EXC 2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden.

References

- Addou, T., Krouchev, N., & Kalaska, J. F. (2011, July). Colored context cues can facilitate the ability to learn and to switch between multiple dynamical force fields. *Journal of Neurophysiology*, 106(1), 163–183. doi: 10.1152/jn.00869.2010
- Baddeley, R. J., Ingram, H. A., & Miall, R. C. (2003, April). System Identification Applied to a Visuomotor Task: Near-Optimal Human Performance in a Noisy Changing Task. *Journal of Neuroscience*, 23(7), 3066–3075. doi: 10.1523/JNEUROSCI.23-07-03066.2003

- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007, September). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214–1221. doi: 10.1038/nn1954
- Brashers-Krug, T., Shadmehr, R., & Bizzi, E. (1996, July). Consolidation in human motor memory. *Nature*, 382(6588), 252–255. doi: 10.1038/382252a0
- Brennan, A. E., & Smith, M. A. (2015, June). The Decay of Motor Memories Is Independent of Context Change Detection. *PLOS Computational Biology*, 11(6), e1004278. doi: 10.1371/journal.pcbi.1004278
- Catz, N., Dicke, P. W., & Thier, P. (2008, May). Cerebellar-dependent motor learning is based on pruning a Purkinje cell population response. *Proceedings of the National Academy of Sciences of the United States of America*, 105(20), 7309–7314. doi: 10.1073/pnas.0706032105
- Criscimagna-Hemminger, S. E., & Shadmehr, R. (2008, September). Consolidation Patterns of Human Motor Memory. *Journal of Neuroscience*, 28(39), 9610–9618. doi: 10.1523/JNEUROSCI.3071-08.2008
- Davidson, P. R., & Wolpert, D. M. (2004, November). Scaling down motor memories: De-adaptation after motor learning. *Neuroscience Letters*, 370(2-3), 102–107. doi: 10.1016/j.neulet.2004.08.003
- Dizio, P., & Lackner, J. R. (1995, October). Motor adaptation to Coriolis force perturbations of reaching movements: Endpoint but not trajectory adaptation transfers to the nonexposed arm. *Journal of Neurophysiology*, 74(4), 1787–1792. doi: 10.1152/jn.1995.74.4.1787
- Ernst, M. O., & Banks, M. S. (2002, January). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429–433. doi: 10.1038/415429a
- Ethier, V., Zee, D. S., & Shadmehr, R. (2008, May). Spontaneous Recovery of Motor Memory During Saccade Adaptation. *Journal of Neurophysiology*, 99(5), 2577–2583. doi: 10.1152/jn.00015.2008
- Forano, M., & Franklin, D. W. (2020, October). Timescales of motor memory formation in dual-adaptation. *PLOS Computational Biology*, 16(10), e1008373. doi: 10.1371/journal.pcbi.1008373
- Gandolfo, F., Mussa-Ivaldi, F. A., & Bizzi, E. (1996, April). Motor learning by field approximation. *Proceedings of the National Academy of Sciences*, 93(9), 3843–3846. doi: 10.1073/pnas.93.9.3843
- Herzfeld, D. J., Vaswani, P. A., Marko, M. K., & Shadmehr, R. (2014, September). A memory of errors in sensorimotor learning. *Science*, 345(6202), 1349–1353. doi: 10.1126/science.1253138
- Huang, V. S., & Shadmehr, R. (2009, August). Persistence of Motor Memories Reflects Statistics of the Learning Event. *Journal of Neurophysiology*, 102(2), 931–940. doi: 10.1152/jn.00237.2009
- Hunter, L. E., & Daw, N. D. (2021, October). Context-sensitive valuation and learning. *Current Opinion in Behavioral Sciences*, 41, 122–127. doi: 10.1016/j.cobeha.2021.05.001
- Imamizu, H., & Kawato, M. (2008, October). Neural Correlates of Predictive and Postdictive Switching Mechanisms for Internal Models. *Journal of Neuroscience*, 28(42), 10751–10765. doi: 10.1523/JNEUROSCI.1106-08.2008
- Imamizu, H., Sugimoto, N., Osu, R., Tsutsui, K., Sugiyama, K., Wada, Y., & Kawato, M. (2007, August). Explicit contextual information selectively contributes to predictive switching of internal models. *Experimental Brain Research*, 181(3), 395–408. doi: 10.1007/s00221-007-0940-1
- Kim, S., Ogawa, K., Lv, J., Schweighofer, N., & Imamizu, H. (2015, December). Neural Substrates Related to Motor Memory with Multiple Timescales in Sensorimotor Adaptation. *PLOS Biology*, 13(12), e1002312. doi: 10.1371/journal.pbio.1002312
- Körding, K. P., & Wolpert, D. M. (2004, January). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247. doi: 10.1038/nature02169
- Lee, J.-Y., & Schweighofer, N. (2009, August). Dual Adaptation Supports a Parallel Ar-

- chitecture of Motor Memory. *Journal of Neuroscience*, 29(33), 10396–10404. doi: 10.1523/JNEUROSCI.1294-09.2009
- Marko, M. K., Haith, A. M., Harran, M. D., & Shadmehr, R. (2012, September). Sensitivity to prediction error in reach adaptation. *Journal of Neurophysiology*, 108(6), 1752–1763. doi: 10.1152/jn.00177.2012
- Medina, J. F., Garcia, K. S., & Mauk, M. D. (2001, June). A Mechanism for Savings in the Cerebellum. *Journal of Neuroscience*, 21(11), 4081–4089. doi: 10.1523/JNEUROSCI.21-11-04081.2001
- Nassar, M. R., McGuire, J. T., Ritz, H., & Kable, J. W. (2019, February). Dissociable Forms of Uncertainty-Driven Representational Change Across the Human Brain. *Journal of Neuroscience*, 39(9), 1688–1698. doi: 10.1523/JNEUROSCI.1713-18.2018
- Oh, Y., & Schweighofer, N. (2019, November). Minimizing Precision-Weighted Sensory Prediction Errors via Memory Formation and Switching in Motor Adaptation. *Journal of Neuroscience*, 39(46), 9237–9250. doi: 10.1523/JNEUROSCI.3250-18.2019
- Pekny, S. E., Criscimagna-Hemminger, S. E., & Shadmehr, R. (2011, September). Protection and Expression of Human Motor Memories. *The Journal of Neuroscience*, 31(39), 13829–13839. doi: 10.1523/JNEUROSCI.1704-11.2011
- Sanders, H., Wilson, M. A., & Gershman, S. J. (2020, June). Hippocampal remapping as hidden state inference. *eLife*, 9, e51140. doi: 10.7554/eLife.51140
- Scheidt, R. A., Dingwell, J. B., & Mussa-Ivaldi, F. A. (2001, August). Learning to Move Amid Uncertainty. *Journal of Neurophysiology*, 86(2), 971–985. doi: 10.1152/jn.2001.86.2.971
- Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M. A., Rymer, W. Z., & Mussa-Ivaldi, F. A. (2000, August). Persistence of motor adaptation during constrained, multi-joint, arm movements. *Journal of Neurophysiology*, 84(2), 853–862. doi: 10.1152/jn.2000.84.2.853
- Scholz, J. P., & Schöner, G. (1999, May). The uncontrolled manifold concept: Identifying control variables for a functional task. *Experimental Brain Research*, 126(3), 289–306. doi: 10.1007/s002210050738
- Shadmehr, R., & Brashers-Krug, T. (1997, January). Functional Stages in the Formation of Human Long-Term Motor Memory. *Journal of Neuroscience*, 17(1), 409–419. doi: 10.1523/JNEUROSCI.17-01-00409.1997
- Shadmehr, R., & Mussa-Ivaldi, F. A. (1994, May). Adaptive representation of dynamics during learning of a motor task. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 14(5 Pt 2), 3208–3224.
- Shmuelof, L., Huang, V. S., Haith, A. M., Delnicki, R. J., Mazzoni, P., & Krakauer, J. W. (2012, October). Overcoming Motor “Forgetting” Through Reinforcement Of Learned Actions. *Journal of Neuroscience*, 32(42), 14617–14621a. doi: 10.1523/JNEUROSCI.2184-12.2012
- Sing, G. C., & Smith, M. A. (2010, August). Reduction in Learning Rates Associated with Anterograde Interference Results from Interactions between Different Timescales in Motor Adaptation. *PLOS Computational Biology*, 6(8), e1000893. doi: 10.1371/journal.pcbi.1000893
- Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006, June). Interacting Adaptive Processes with Different Timescales Underlie Short-Term Motor Learning. *PLoS Biology*, 4(6). doi: 10.1371/journal.pbio.0040179
- Vaswani, P. A., & Shadmehr, R. (2013, May). Decay of Motor Memories in the Absence of Error. *Journal of Neuroscience*, 33(18), 7700–7709. doi: 10.1523/JNEUROSCI.0124-13.2013
- Wolpert, D. M., & Kawato, M. (1998, October). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11(7), 1317–1329. doi: 10.1016/S0893-6080(98)00066-5
- Zarahn, E., Weston, G. D., Liang, J., Mazzoni, P., & Krakauer, J. W. (2008, November). Explaining Savings for Visuomotor Adaptation: Linear Time-Invariant State-Space Models Are Not Sufficient. *Journal of Neurophysiology*, 100(5), 2537–2548. doi: 10.1152/jn.90529.2008