

Motor memories of objects

Movement is the only way we have of interacting with the world. As such, the primary purpose of the human brain is to use sensory signals to determine future actions. Despite great progress in the field, movement neuroscience is still in its infancy. We have a detailed understanding of a narrow range of constrained tasks, such as single planar arm movements made under simple visual or mechanical perturbations (Figure 1).


There have been clear benefits to this narrowly focused approach. By directing the efforts of many labs toward a few 'model tasks', researchers with different specializations can work together to investigate many relevant aspects (*i.e.*, behavior, modeling, clinical, stimulation, imaging, electrophysiology). Over time, this can lead to an impressive level of understanding. For instance, a recent study by Sun and colleagues (2022) analyzed electrophysiological recordings from thousands of neurons in two monkeys' primary motor and dorsal premotor cortices and successfully pinpointed changes in neural population activity that reflect key behavioral phenomena in this task (*i.e.*, gradual learning, generalization and interference across learning contexts, and faster relearning). Such high-profile discoveries clearly show that we are making technical and theoretical progress toward understanding how the brain controls movement. It is unlikely that this could have been accomplished without a sustained focus on force-field adaptation in motor control research for the past 25+ years.

However, to recognize the pitfalls of a narrowly focused approach, it is instructive to ask: why does a particular task become a model task? One important feature is accessibility: a model task must be simple enough for multiple laboratories to implement, or else it will not become widespread. In practice, this means that real-world complexity must be compressed into a few task variables or simply ignored. The second important feature is replicability: a model task should generate the same results every time or else the risk of unexpected results may become too great to justify the costs of running the study. Given the small sample sizes of motor learning experiments ($N \approx 12$ for human studies and $N \approx 2$ for monkey studies), this effectively means that a model task must reliably produce exactly the same results in each

visuomotor rotation








Learning from a visual error



from Tsay et al., 2022

force field



from Howard et al., 2012

Figure 1. ‘Model tasks’ in motor learning research.

Neural preparation of movements is typically divided into two stages. First, various types of sensory information (e.g., target location, orientation, size; eye position, posture) are transformed into desired movement parameters (e.g., reach destination, shape of the grip). Next, a temporal sequence of muscle activations is prepared that is appropriate to produce the movement specified by the desired movement parameters. To study the first stage, researchers disrupt the normal relationship between visual information and required movement parameters. To study the second stage, researchers use haptic interfaces that produce novel forces on the hand during the movement.

(Left) Visuomotor perturbations. In the early 1900s, Hermann von Helmholtz and George Stratton pioneered the use of prism lenses to perturb the visual input by shifting or inverting the projection of light on the retina. Prism adaptation was popular throughout the late 20th century, with major contributions from Richard Held, Robert Welch, and others. Today, the use of lenses has largely been replaced by the use of computer displays, where the physical movement of the hand on digitizing trackpad is inherently dissociated from the visual movement of the on-screen cursor. With this technique, researchers have focused on two closely related model tasks: *visuomotor rotations*, depicted here, where the cursor moves in a different direction than the hand, and *visuomotor gains*, where the speed of the cursor differs from the speed of the hand.

(Right) Force-field perturbations. To move the hand along a desired trajectory, the motor system must produce an appropriate sequence of muscle activations (a ‘movement plan’). The trajectory that results from a particular movement plan is determined by the laws of physics (the ‘system dynamics’), and depends on various characteristics of the arm (e.g., mass distribution, arm segment lengths, muscle fatigue). The process of generating a movement plan can therefore be understood as computing the *inverse* of this physical transformation, i.e., going from a desired trajectory to an appropriate sequence of muscle activations. One early approach to perturbing the dynamics of reaching movements involved using rotating rooms to induce Coriolis forces (Lackner & Dizio, 1992). Researchers soon turned to programmable haptic interfaces (‘robot handles’; e.g., vBOT) to produce forces on the hand. Although many types of force field can be created with this apparatus, experiments on motor learning have focused on adaptation to curl fields. In a curl field, the instantaneous force exerted on the hand is perpendicular to the direction of motion and proportional to the velocity. The bottom panel shows typical movement trajectory data during adaptation to a counterclockwise (blue) or clockwise (red) force field. At left, normal movements with no force field; center gray region, perturbed movements initially show large deviations but become straighter over time; at right, when the field is turned off, movements show opposing deviations, indicating that the movement plan has adapted to anticipate the altered dynamics.

individual, discouraging the study of phenomena marked by individual differences that cannot (yet) be predicted. Again, from the perspective of accelerating progress at multiple levels (*i.e.*, behavior, algorithm, neural basis), a shared focus on tasks that are highly accessible and consistently replicable certainly makes sense; no one would argue that this approach should be abandoned.

However, it is worth considering that such an intense focus on specific results from specific tasks can lead to an overly restrictive definition of the field itself. That is, an initially broad definition of motor learning has been followed by a great accumulation of data about motor learning in visuomotor rotation and force field tasks, leading to models (often pitched as ‘integrative theories’) that are supposed to be general, but in some cases do little more than redescribe the most reliable observations from the model tasks in formal, mathematical terms. For example, theoretical discussions at motor learning conferences are frequently dominated by two such ‘theories’: associative models of motor memory that implement the principle of gradual, context-specific learning (Figure 2a), and dual-rate models of motor learning that implement the principle of separate implicit and explicit processes (Figure 2b). Without greater task diversity, it is impossible to know whether these principles are truly general, and in what ways they contribute across the various tasks included in the vast human sensorimotor repertoire. This not only limits our ability to translate research findings into the real world, for example in rehabilitation clinics or athletic training, but also increases the likelihood that findings that conflict with these principles will be skeptically regarded as spurious (“the experiment was poorly conducted”), idiosyncratic (“the findings pertain only to a limited context”), or non-motor (“possibly interesting, but not of central importance”).

The long-term goal of my research is to develop a new conceptual framework of motor learning that focuses on skillful interaction with physical objects. Object manipulation is an essential ability in daily life and can be devastatingly disrupted by illness or injury, making it an impactful topic in its own right. Yet, from a broader perspective, a focus on object manipulation may act as a counterbalance to the ‘narrowing’ tendency described above. This is because it naturally recruits many different types of cognitive function while retaining a high degree of ecological validity. Just scratching the surface, it is

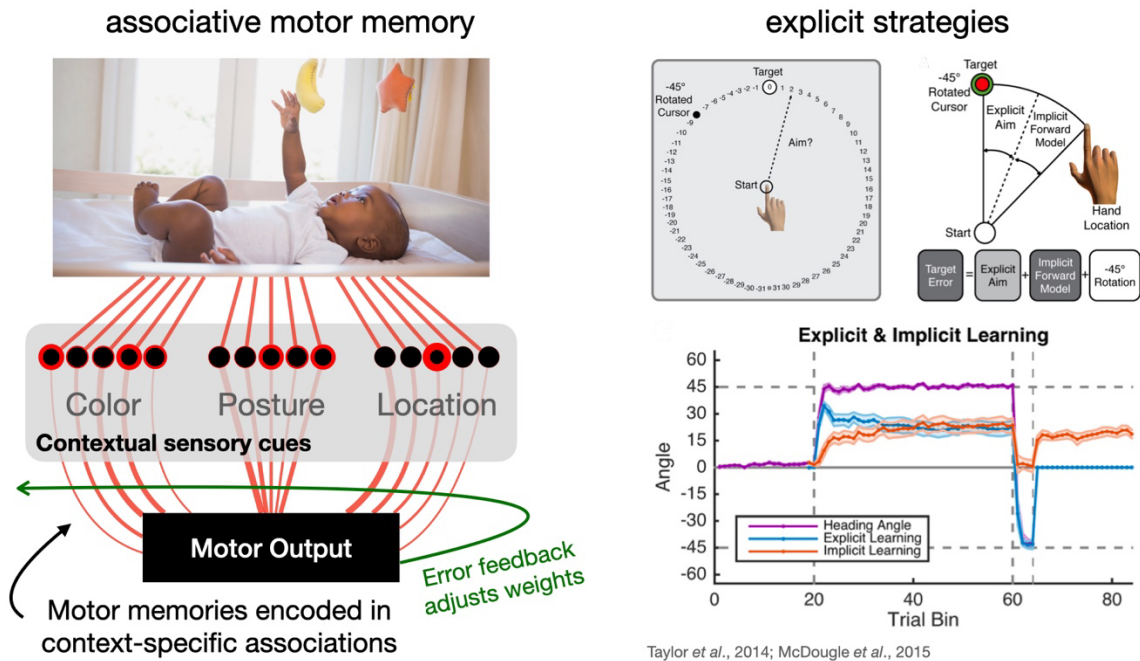


Figure 2. Two dominant ideas in motor learning

(Left) A highly simplified diagram of the ‘associative motor memory’ model. This model posits that when learning a novel motor skill, sensory input information is encoded by populations of sensory ‘neurons’ (black circles) corresponding to different sensory dimensions (e.g., target color, body posture, target location). Each neuron has a *preferred value* along the dimension encoded by its population, meaning it responds most strongly to that value. Here, the activity of each neuron is indicated by the thickness of its outline. For instance, in the neural population that encodes the context dimension *Location*, the center neuron is depicted as responding vigorously (thick red outline) to indicate that its preferred target location is ‘*straight ahead*’, given that the reach target (toy banana) is located straight ahead. This activation specifies one aspect of the sensory context of the movement. The neural activations across all of the input populations thus encode a detailed specification of the movement context. As a result, motor memories can be encoded in the connection weights between specific subsets of these neurons, and the downstream neural population that encodes movement parameters (the black box labeled *Motor Output*). When the movement is unsuccessful, error feedback can be used to update this context-specific motor memory by changing the connection weights.

(Right) Operationalizing the use of explicit strategies in a motor learning task. When participants are learning a visuomotor rotation (see Figure 1a), the experimenter places a ring of numbers around the starting position of the hand. In a visuomotor rotation task, an explicit strategy can be defined as deliberately choosing *not* to move the hand as if you were trying to touch the target, but instead as if you were reaching to a different target. Participants are informed of this strategy and asked to specify where they are aiming prior to each movement by indicating a number along the ring. The observed reaching directions (purple curve) demonstrate an additive combination of the explicitly reported aiming directions (blue curve) and the implicit contribution of motor memory (red curve), which appears to gradually and automatically learn the kinematics of the task based on error feedback.

easy to see the involvement of neural systems for multisensory perception, learning and memory,

reasoning, and in some cases even social cognition and communication. This breadth makes object

manipulation a great starting point for developing new approaches. In addition to potentially changing the

way we think about motor control, it also creates an opportunity for motor control research to influence

theories of object memory, which currently focus on perception and cognition but not action. Though object manipulation has long been an important subfield of motor control research (Westling & Johansson, 1984; Flanagan & Johansson, 2009), research up to this point has primarily investigated the fine mechanical details of grip-force control (*e.g.*, reflex responses to object-slippage events, predicting the timing of liftoff and contact events). As a result, it has served only to complement, not to challenge, the ‘mainstream’ findings from visuomotor rotation and force field adaptation.

Criticisms of the categorical encoding hypothesis of Cesanek *et al.* (2021)

The perspective developed above, stressing the need for greater task diversity, is reflected in my recent work examining the representation of objects’ mechanical properties in motor memory (Cesanek *et al.*, 2021). Briefly, in that study, I tested the novel hypothesis that motor memories of objects are organized categorically, in terms of ‘families’, based on covariation in their visual and mechanical properties. Participants in this study first repeatedly lifted a set of training objects that covaried in size and weight. After they learned the weights of these objects, I introduced an ‘outlier’ object that appeared similar to the training objects, but was much heavier than expected. Remarkably, participants persistently failed to learn the weight of this object, instead using the expected weight based on the density of the other four objects when producing lift forces. This clearly demonstrates that the four training objects were learned as a family, and the outlier was also encoded as a family member. However, when the outlier was sufficiently deviant—and thus exceeded a hypothesized category boundary—it was ‘kicked out’ of the family and learned as an individual. A variety of additional observations reinforced and extended this interpretation. This work provides the first empirical evidence that categorization plays a key role in organizing the memories underlying skilled action.

In the published article, Cesanek *et al.* (2021) argue that the observed categorical effect cannot be accounted for by current theories of motor learning. There is substantial discussion around this point in the article, but systematic evaluation of this claim remained incomplete. Thus, one of the aims of my present research has been to review relevant literature (1) to confirm the extent to which those findings

truly demand revision of current thinking about motor learning, and (2) to motivate new experiments that will help us understand the underlying issues. In general, there have been two aspects to this reappraisal, discussed in sequence below.

The associative encoding assumption shared by current models of motor memory

The first aspect involves pinpointing *why* current models of motor memory are inconsistent with categorical encodings. Ultimately, this rests on a shared, fundamental assumption of all current models of motor memory. This is the assumption that memories are indexed (*i.e.*, separately formed, accessed, and updated) by context-specific sensory cues available during movement planning and/or execution. This assumption is found not only in purely associative models, such as radial basis function networks (*e.g.*, Thoroughman & Shadmehr, 2000), but also in models explicitly designed to simplify the learning of multiple behaviors by clustering similar movements together under one motor memory (*e.g.*, MOSAIC; Wolpert & Kawato, 1997). It should be noted that both of these model classes can, under the right circumstances, produce the same motor output for different sensory inputs. Although this may be considered a primitive form of categorization, it may be helpful to distinguish this mere lack of sensitivity, which we might instead call ‘clustering’, from the more sophisticated function of ‘categorization.’ In particular, the nature of a categorical encoding is that it *avoids* forming associations between specific sensory/motor input/output pairs. Instead, a categorical encoding compactly summarizes the relationship between many potential input-output pairs. For instance, the concept of density is a categorical encoding of a set of objects that can vary in size.

One notable observation from this review is that the empirical support for the assumption of context-specific sensory-motor associations seems to rest on one specific finding. This finding is the bell-shaped profile of the behavioral generalization function that arises in reaching tasks involving multiple, radially arranged targets. The generalization function measures the degree to which learning in one context (*i.e.*, a reach to one target) subsequently affects behavior in other contexts with varying degrees of similarity (reaches to other targets, ranging from only few degrees away to a full 180 degrees away). The bell-

shaped profile of the generalization function in reach adaptation tasks has been one of the most robust results of motor learning research. Yet it has only been observed in one kind of generalization analysis, looking at just one specific contextual cue: the planned movement direction of the hand in center-out reaching tasks. Meanwhile, little attention has been paid to the few conflicting reports demonstrating that generalization over reach direction can show a global, rotational structure (Vetter *et al.*, 1999; Turnham *et al.*, 2011).

In contrast, other types of contextual sensory cues—such as the visual features of manipulable objects (size, surface texture, 3D shape)—have never been experimentally manipulated in a way that could falsify the associative account. This is not to say that other cues have *never* been examined. On the contrary, a wide variety of cues have been employed in numerous experiments investigating the formation of multiple context-dependent motor memories (Cothros *et al.*, 2008; Howard *et al.*, 2013; Ayala *et al.*, 2015; Heald *et al.*, 2018; Schween *et al.*, 2018). However, the design of these experiments always involves presenting two opposing perturbations (*i.e.*, CW and CCW curl fields, or CW and CCW visuomotor rotations), with the perturbation direction cued by one of two values of the context cue (*e.g.*, red cursor vs. green cursor). Such experiments have only two possible outcomes: either they reveal some capacity for context-sensitive learning of both perturbations, or they do not. Importantly, neither outcome sheds any light on the question of whether abstract representational formats, such as categories, are involved in the process of learning.

Defining the role of explicit memory

The second, more complicated aspect of the reappraisal of Cesanek *et al.* (2021) relates to the possibility that the categorical effect could arise from declarative memory (*i.e.*, systems that store factual information for conscious recall) rather than motor memory (*i.e.*, systems that store information relevant to generating a motor plan to perform an action). The basis of this criticism comes from another influential trend in current thinking about motor learning, which focuses on the potential role of explicit strategies (Mazzoni & Krakauer, 2006; Taylor *et al.*, 2014). The idea of adopting an explicit strategy when performing a

motor task is rather intuitive (*e.g.*, “*try to time it such that...*” or “*make sure you grip it like this...*”). Yet it has been rather elusive and difficult to isolate in motor learning experiments. Even in a visuomotor rotation task, where the explicit strategy is quite obvious from the experimenter’s point of view, participants who are not instructed in any way will describe all sorts of idiosyncratic strategies when asked in a post-experiment debriefing. However, Mazzoni and Krakauer (2006) found that if participants were *given* the appropriate strategy at the beginning of the task—“*Since the cursor is rotated with respect to your hand movement, move your hand as if you were trying to hit a target located over _here_*”—all participants would readily adopt this strategy. Surprisingly, in the following trials their performance actually became worse! They kept trying to aim at the instructed location, but actual movements drifted even farther in that direction, overcorrecting for the visuomotor rotation.

This was taken as evidence that the observed learning in a reach adaptation experiment can be decomposed into two additive components. First, an implicit component that results from automatic, error-driven updating of motor memory. This component is typically regarded as *bona fide* motor learning, as it appears to be based on sensory-prediction errors generated by predictive models in the cerebellum that compare actual and expected sensory feedback (Taylor *et al.*, 2010). In parallel, there is an explicit component, which refers to the participant’s conscious decision to apply a particular strategy. This conforms with a longstanding distinction in memory research between implicit and explicit memory systems, and questions a traditional model of motor skill acquisition (Fitts, 1954) which asserted a three-phase process, beginning with an often-clumsy ‘cognitive phase’ in which an explicit understanding of the task procedure is formed, following by an ‘associative phase’ in which repeated practice leads to performance refinements, coalescing in the ‘autonomous phase’ in which explicit processing no longer contributes to the performance of the task. The recent results on the role of explicit processing in visuomotor rotation tasks suggest that, even with extensive practice, explicit processing continues to play a significant role in overtrained participants’ performance.

Figure 2b depicts the labeled-ring technique that was introduced to implicit and explicit processes in visuomotor rotation tasks. While this approach works brilliantly for visuomotor rotations, there are a number of conceptual difficulties that prevent a task-general definition of the explicit component. Notice, for instance, that only in the labeled-ring paradigm is it possible to state the implicit and explicit components in the same units. In other settings, the explicit component is notoriously difficult to define, and questions about how it interacts with implicit learning do not have simple answers. For example, the labeled-ring paradigm has also been used in curl field adaptation. Here, participants are similarly willing to describe an explicit strategy that involves pushing the handle in a different direction, not straight at the target (McDougle *et al.*, 2015). However, the conclusions are not as tidy as in the visuomotor rotation task. In particular, a participant's ability to describe the forces they have learned to generate does not necessarily imply that the explicit memories supporting this description are contributing to the movements they are performing. Thus, unlike in the visuomotor rotation task, the possibility emerges that explicit and implicit memories formed during a motor task can be redundant rather than additive. In recent work, I have aimed to clarify the role of explicit memory in the categorical encoding effects observed in the experiments reported in Cesanek *et al.* (2021). In a variety of web-based experiments conducted during the pandemic, we have recently gained a more detailed understanding of the categorical effect, particularly its modulation by sensory and structural factors. In the hope that in-person laboratory experiments with human subjects can soon resume in a productive way, I have also planned one set of laboratory experiments that aim to resolve outstanding concerns.

References

- Sun, X., O'Shea, D. J., Golub, M. D., Trautmann, E. M., Vyas, S., Ryu, S. I., & Shenoy, K. V. (2022). Cortical preparatory activity indexes learned motor memories. *Nature*, 1-6.
- Westling, G., & Johansson, R. S. (1984). Factors influencing the force control during precision grip. *Experimental brain research*, 53(2), 277-284.

- Johansson, R. S., & Flanagan, J. R. (2009). Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5), 345-359.
- Cesaneck, E., Zhang, Z., Ingram, J. N., Wolpert, D. M., & Flanagan, J. R. (2021). Motor memories of object dynamics are categorically organized. *Elife*, 10, e71627.
- Vetter, P., Goodbody, S. J., & Wolpert, D. M. (1999). Evidence for an eye-centered spherical representation of the visuomotor map. *Journal of neurophysiology*, 81(2), 935-939.
- Turnham, E. J., Braun, D. A., & Wolpert, D. M. (2011). Inferring visuomotor priors for sensorimotor learning. *PLoS computational biology*, 7(3), e1001112.
- Cothros, N., Wong, J., & Gribble, P. L. (2008). Distinct haptic cues do not reduce interference when learning to reach in multiple force fields. *PloS one*, 3(4), e1990.
- Howard, I. S., Wolpert, D. M., & Franklin, D. W. (2013). The effect of contextual cues on the encoding of motor memories. *Journal of neurophysiology*, 109(10), 2632-2644.
- Ayala, M. N., t Hart, B. M., & Henriques, D. Y. (2015). Concurrent adaptation to opposing visuomotor rotations by varying hand and body postures. *Experimental brain research*, 233(12), 3433-3445.
- Heald, J. B., Ingram, J. N., Flanagan, J. R., & Wolpert, D. M. (2018). Multiple motor memories are learned to control different points on a tool. *Nature human behaviour*, 2(4), 300-311.
- Schween, R., Taylor, J. A., & Hegele, M. (2018). Plan-based generalization shapes local implicit adaptation to opposing visuomotor transformations. *Journal of neurophysiology*, 120(6), 2775-2787.
- Mazzoni, P., & Krakauer, J. W. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *Journal of neuroscience*, 26(14), 3642-3645.
- Taylor, J. A., Krakauer, J. W., & Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, 34(8), 3023-3032.
- Taylor, J. A., Klemfuss, N. M., & Ivry, R. B. (2010). An explicit strategy prevails when the cerebellum fails to compute movement errors. *The Cerebellum*, 9(4), 580-586.

- Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology*, 47(6), 381.
- McDougle, S. D., Bond, K. M., & Taylor, J. A. (2015). Explicit and implicit processes constitute the fast and slow processes of sensorimotor learning. *Journal of Neuroscience*, 35(26), 9568-9579.