Movement chunking as locally optimal control

Pavan Ramkumar, Daniel E. Acuna, Max Berniker, Scott Grafton, Robert S. Turner, Konrad P. Kording

Chunking is a fundamental and widely recognized aspect of working memory [1]. To illustrate, a string of 10 digits such as in an American phone number, is typically learned, stored and recalled as 3 chunks of 3, 3, and 4 digits each. Similarly, when one learns a sequence of movements, some of its components become temporally integrated into chunks — contiguous groups of movements executed with short and relatively invariable response times. Why are composite motor behaviors chunked rather than executed in one fell swoop? Although the dominant hypothesis from the working-memory literature suggests that cognitive storage limits inform chunk structure [1], here, we hypothesize that the computational costs of generating lengthy sequences of movements may be just as important. In optimal control, the search space grows exponentially with the length of the temporal horizon. Therefore, we propose that the formation of motor chunks represents a trade-off between the desire to optimize movement efficiency and the computational effort involved therein. According to this hypothesis: (a) through learning movements become more efficient, and (b) sequences are partitioned in chunks, each of which can be optimized for efficiency independently. We show evidence to support both these predictions.

First, we analyzed the kinematics of a monkey learning to execute a sequence of center → out → center reaches comprising 5 elements (Fig. 1) over 41,865 trials and 63 days. Using the minimum-jerk cost as a proxy for efficiency, we observed that learning results in increased movement efficiency, where inefficiency is characterized as squared jerk corrected for differences in trial duration (Fig. 2). Next, to model the kinematics of complex motor sequences, we allowed each peripheral position (out) as well as center position (center), to be either a via point (arm does not come to a stop) or a halt point (arm comes to a stop) [2]. Thus parameterized, the search space included two extreme canonical models (Fig. 3). At one extreme, the *halt model* assumes a halt after every center → out or out → center segment; each sequence is a series of N isolated minimum-jerk movements. This model is least efficient and best approximates arm velocities during the execution of sequences with completely unpredictable target locations. At the other extreme, the *via model* assumes only via points between all elements; each sequence is one continuous minimum jerk reach. This model is maximally efficient (minimizes jerk) over the entire sequence and represents a hypothetical overlearned sequence where all elements within the sequence are grouped into a single chunk. The minimum-jerk modeling framework allows us to ask where on this continuum between the halt and via models actual behavior lies. To this end, we exhaustively searched through all resulting 2^{2N-1} predictions ($N = 5$ is the number of elements) for the locally optimal model that best explained kinematic data. We found that a locally optimal model (Fig. 4A) — a series of minimum-jerk trajectories — outperformed the canonical models (Fig. 4B), with the number of inferred halt points decreasing with practice (compare Figs. 3 and 4A). Finally, using only response times (independent of efficiency considerations) we identified transitions between motor chunks using a hidden Markov model [3] and estimated the expected number of chunks for each trial (Fig. 5A). We observed that over the course of learning, the expected number of chunks was highly correlated with the number of halt points in the locally optimal minimum-jerk model (Fig. 5B; $r = 0.84$, $p < 0.0001$). Thus, each trial was closely approximated as a series of reaches that minimized jerk.

The optimality of within-chunk movements rather than of the entire movement sequence, even after extensive practice, suggests that the existence of discrete chunks may be a fundamental feature of the motor system. Future studies may focus on identifying what the nervous system optimizes, by seeking regularities with chunks specifically, rather than across the movement as a whole, and how difficult these optimizations are, by examining the number of chunks.

References

[1] Cowan N (2000) Behav Brain Sci 24, 87–185.

[2] Todorov E, Jordan M (1998) J Neurophys 80, 696–714.
[3] Acuna D et al. (2014) J Neurophys [Epub ahead of print].

Figures

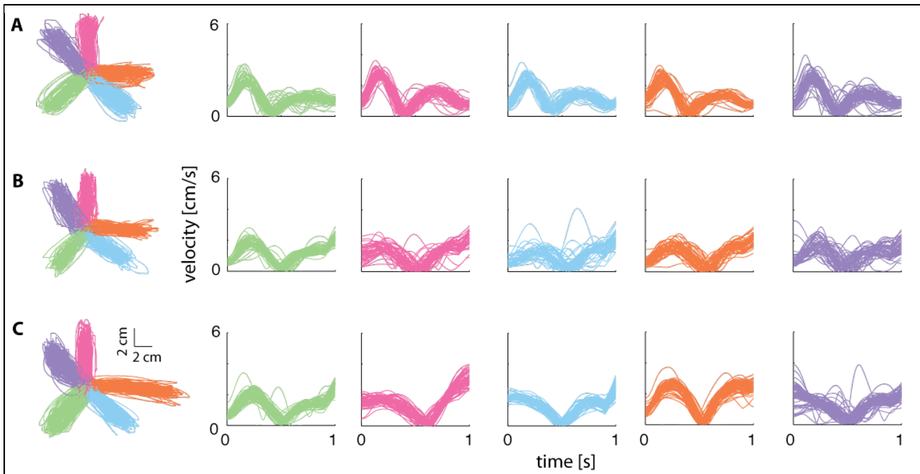

Figure 1. Arm kinematics of the reaching sequence task. Hand positions and velocity traces for each center->out->center movement in the sequence; each trial is warped to a 1 s duration. (A) Trials 1–1000, (B) Trials 20001–21000, (C) Trials 40001–41000.

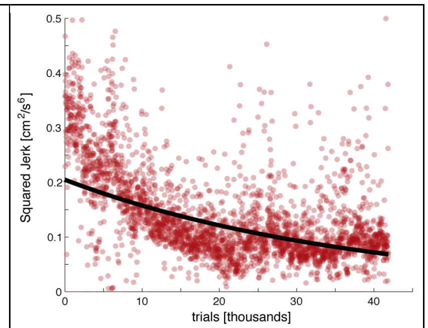

Figure 2. Squared jerk of the reaching sequence warped to have a 5 second duration over 41865 successive trials (63 days of learning), suggests a clear increase in efficiency with learning. The black curve shows an exponential fit.

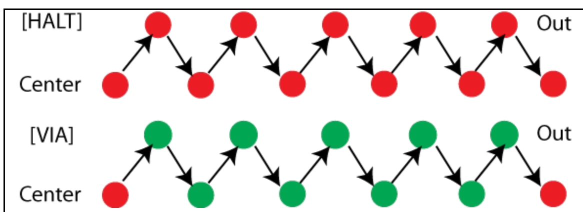

Figure 3. Glyph illustrating the two canonical halt and via models for the 5-element center–out reaches. Green: via points. Red: halt points. A continuum of models exists in between these canonical models for which the intermediate center and out points can either be halt or via points.

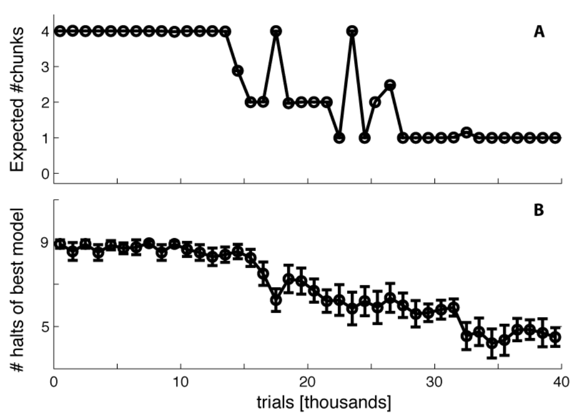

Figure 5. (A) Expected number of chunks from response times (independent of efficiency considerations) (B) Estimated number of halts of the best-fitting locally optimal model. The two quantities are highly correlated (r = 0.84; p < 0.0001)

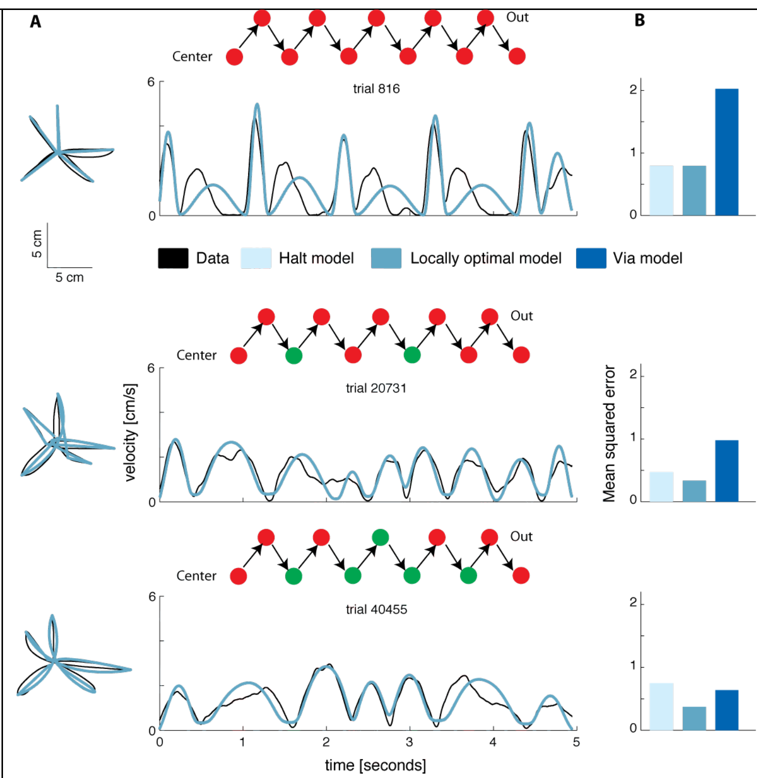

Figure 4. (A) Generative model of arm kinematics based on the minimum-jerk principle for three exemplar trials over the course of learning. The data are shown (black) along with the best fitting locally optimal model (blue). The halt and via points of this model are illustrated according to the glyph in Figure 3. (B) MSE of the halt model, the locally optimal model and the via model.