

One motor cortex, two different views

TO THE EDITOR—In a recent paper in *Nature Neuroscience*, Todorov¹ referred to our finding that a motor cortical representation of hand trajectory during spiral drawing precedes the hand's movement by an interval that varies with path curvature^{2,3}. Although there are several possible explanations for this finding, Todorov, using a simplistic model, argued that because cortical cells share common properties with muscles, this relationship could be due to a combination of inertia, viscosity and stiffness acting on the acceleration, speed and position of the arm, respectively. Although simple, his model is flawed and cannot support this conclusion.

The author models a multijoint arm as a simple cantilever that is converted to single point-mass equation using a Jacobian transformation (web supplement A, http://www.nature.com/neuro/web_specials/). The arm's properties were derived from a simplified version of muscle whose activity is a linear combination of motor cortical activity. This model was used to reinterpret our results^{2,3}. In our study, monkeys drew spirals on a vertically oriented computer touchscreen. The center of the spiral was located in front of the monkey, between its shoulders. According to Todorov's model, this location corresponded to the equilibrium point of the arm—the location where the parameters in his model would force the arm to rest. Todorov assumed that cortical activity reflects the inertia, viscosity and stiffness of the arm and showed that his model produces the same variable lags as our cortical population vectors. However, any acceleration representation in the cortical activity would actually decrease lags as a function of curvature, which is exactly opposite to our finding (web supplement B, http://www.nature.com/neuro/web_specials/).

The increased lag with increasing curvature shown in Todorov's article is due to his positional term. The idea that extrinsic position may be a factor in motor cortical activity is not new^{4–6}. However, Todorov's method of equating extrinsic position representation to muscle stiffness is incorrect. This model assumes that muscle viscoelastic properties are independent of muscle activation. Thus, even an inactivated muscle will act as a large spring pulling the arm back to some equilibrium position. In real muscle, the force–length and force–velocity relationships are modulated by muscle

activation such that at zero activation, the muscle is essentially a non-force producer. In the real world, the combination of gravity and inactive muscles will force the arm to fall to the side. In Todorov's model, the combined effect of gravity and muscle stiffness on inactive muscles would make the hand float at mid-chest level; muscle activity would be required to force the arm down below chest level. This, of course, is unrealistic. Viscoelastic models like the ones used by Todorov are only valid for perturbation studies where both posture and neural activity are assumed to be constant. Using such equations to solve for time-varying muscle activations violates the basic assumptions of perturbation models. Simple dynamic models can be useful to explain arm mechanics. However, when the models are not consistent with basic physiology, exclude important phenomena, and violate inherent assumptions, they cannot be compared to empirical data.

Daniel W. Moran and Andrew B. Schwartz
The Neurosciences Institute, 10640 John Jay Hopkins Drive, San Diego, California 92121, USA
email: d Moran@nsi.edu or aschwartz@nsi.edu

TO THE EDITOR—Here we refute claims by Todorov¹ and Scott⁷ that the importance of target direction as an explanatory factor for cortical activity in a regression analysis we performed⁵ is an 'artifact' of a square-root transformation of neural discharge rates. Specifically, it was touted by Scott⁷ that "squaring [*sic*] the discharge rate of neurons in order to stabilize the variance ... causes a dramatic increase in the percentage of neurons that appear to represent movement direction (from 17% [*sic*] to 43% in Todorov's model)." The data to which Todorov¹ referred concerned the percentages of cells for which a particular variable yielded the highest R² when used alone in the regression. We re-analyzed these data using the regression analysis we used previously⁵ but without any transformation of the discharge rate. The results of the two analyses were practically identical, the average absolute difference being only 1.9% (http://www.nature.com/neuro/web_specials/). However, there was a statistically significant improvement of the regression fit when the square-root transformation was used. The median R² for the square-root transformed data was 0.5811, as compared

to 0.544 for the non-transformed data ($P < 0.0001$, Wilcoxon's signed-rank test). This was anticipated, because the square-root transformation is expected to make the distribution of counts more symmetrical. This transformation is routinely used when analyzing counts^{8–10}, given the commonly highly skewed distribution of such data. Finally, we analyzed the data without any transformation or smoothing. In this case, the agreement with the original analysis was even closer, the average absolute difference being only 0.86%. We conclude that the relationship between neural activity and movement parameters found earlier⁵ holds irrespective of the specific transformation and/or smoothing used. Finally, while we dealt above with the issue of square-root transformation because of the more general importance of this transformation for analyzing neuronal spike counts, there are also numerous other points raised by Todorov¹ which we also dispute, including the force direction/magnitude issue, which we cannot critically discuss due to space limitations.

Apostolos P. Georgopoulos and James Ashe
Brain Sciences Center, Veterans Affairs Medical Center and Department of Neuroscience, University of Minnesota Medical School, Minneapolis, Minnesota 55417, USA
email (A.P.G.): omega@tc.umn.edu

REPLY TO MORAN AND SCHWARTZ—Assuming that M1 cells control the activation of muscle groups, I have previously derived an equation¹ relating the M1 population vector (PV) to hand kinematics and kinetics. In addition to force and acceleration terms, this equation includes velocity and positional terms needed to compensate for muscle visco-elasticity. The interplay among these terms offers a simple explanation to several puzzling phenomena¹ including the curvature-dependent time-lag between PV direction and tangential velocity^{2,3}. The strength of the model is that multiple phenomena are explained simultaneously, using the most basic properties of the musculoskeletal apparatus and thus avoiding the danger of curve fitting.

Moran and Schwartz claim that the fit to their data^{2,3} is somehow an artifact of the approximation I used, and that the results will change if additional details are considered. It is not explained how a first-order approximation could produce such

artifacts. The absence of gravity and activation-dependent stiffness in my model are discussed at length, without any explanation as to why adding them should change the results. Unlike hand acceleration, which is time-varying, the gravitational force is roughly constant for small variations in limb configuration. Therefore, its effect is absorbed in the baseline (defined as postural activity in the center of the workspace). The constant stiffness approximation is reasonable once a certain activation level is reached¹¹. It is true that setting stiffness to zero abolishes the time lag–curvature effect; however, that point is irrelevant—whether or not muscle stiffness is constant, it certainly exists and there is no justification for setting it to zero. To compensate for muscle stiffness, M1 cells have to exhibit well-documented positional gradients—which, in combination with the acceleration term, produce the negative time lag–curvature relationship¹.

Moran and Schwartz have only shown that my model is approximate—which is very different from being “flawed.” Still, is it possible that the results are an artifact of the approximation, for reasons that these authors did not identify? To assess the sensitivity to previously unmodeled details, I repeated the analysis using a state-of-the-art muscle model¹² (http://www.nature.com/neuro/web_specials/). Muscle force was expressed as a complex function of muscle length, velocity and stimulation frequency; this function depends on 19 experimentally derived parameters and incorporates numerous results from muscle physiology¹². For a wide range of parameters, the predicted relationship between PV time lag and path curvature was quantitatively similar to the original result¹ as well as to the experimental data^{2,3}. Thus, when Moran and Schwartz wrote that my model “is not consistent with basic physiology, excludes important phenomena, and violates inherent assumptions,” their concerns were misplaced.

REPLY TO GEORGOPOULOS AND ASHE—In their letter, Georgopoulos and Ashe address the issue of whether data preprocessing affects M1 cell classification. Their results do not refute my main point¹, which is that previous classification procedures^{5,13} can be seriously biased—with or without data preprocessing. Here I focus on the procedure⁵ for classifying cells as direction- (D), position- (P), velocity- (V) or acceleration-related (A) according to the largest R².

By definition, the bias of a statistical estimator is the expected difference between the correct and estimated values

of a given parameter. The identical percentages found by Georgopoulos and Ashe with and without data preprocessing only prove that, for this particular dataset, the bias is equal in both cases. The value of this bias cannot be inferred from their results (or anything else computed on real data), because the correct answer is unknown. In the absence of analytical insight, the only way to identify the bias of an estimator is to apply it to synthetic datasets where the correct answer is known. When applied to synthetic data¹ with no separate directional component, the above classification procedure finds D 43%, V 39%, P 16%, A 2% on smoothed square-root-transformed spike trains, and D 26%, V 56%, P 16%, A 2% on raw binned spike trains (different from continuous mean firing rates which were labelled MFR previously¹). The effects of the data transformation (8% on average) are to be expected in general, and could exist in other datasets. With or without the transformation, the above percentages are very far from the correct answer: D 0%, V 49%, P 46%, A 5%—that is, the classification procedure itself is biased. Thus the burden of proof lies on Georgopoulos and Ashe. Unless they identify the exact conditions under which their procedure is unbiased, and ascertain by independent means that these conditions hold for the M1 population, their results remain hard to interpret.

To gain more insight into why the R² classification procedure fails, I analyzed the family of synthetic responses misclassified as directional (http://www.nature.com/neuro/web_specials/). These responses do not look directional: the temporal fluctuations of the underlying position, velocity and acceleration terms do not cancel out. The artificially created ‘directional’ region of parameter space is centered at the point where the correct classification boundaries meet. Therefore responses are misclassified as directional just because they do not fit well in the other categories. It would be interesting to re-analyze the data of Georgopoulos and Ashe for that possibility. Do responses labeled as directional vary only with movement direction and contain no systematic temporal fluctuations (which is how a truly directional cell should behave), or do they fluctuate over time in ways that do not happen to fit in any of the alternative categories? The latter type of response is more properly labeled ‘unknown’ rather than ‘directional.’

Finally, this debate obscures a more fundamental problem^{1,14} with M1 cell classification, a problem that remains even if unbiased procedures are developed. The

different components of the cell response are not fixed, but instead increase monotonically with the magnitude of the corresponding kinematic and kinetic terms. Thus a cell classified in one task as ‘velocity-related’ could become ‘position-related’ in another task if the movement is slow enough, ‘acceleration-related’ if the movement is fast enough, and ‘load-related’ if a large enough external load is imposed (http://www.nature.com/neuro/web_specials/). Given this sensitivity to task parameters, classifying M1 cells according to the largest component of their response should perhaps be avoided altogether.

Emanuel Todorov

Gatsby Computational Neuroscience Unit,
University College London, 17 Queen Square
London WC1N 3 AR, UK
email: emo@gatsby.ucl.ac.uk

REPLY—The article by Todorov¹ and associated letters illustrate clear opinion differences regarding the function of motor cortex during goal-directed arm movements. This controversy is partially generated by the different experimental protocols used to examine motor cortex function in non-human primates. The first, introduced by Evarts, examines single-joint movements and relates neural activity to muscle-based or joint-based variables¹⁵. The second, introduced by Georgopoulos, examines whole-arm movements and relates neural activity to hand-based variables¹⁶. Practitioners of the former find correlates of muscle-based or joint-based variables; practitioners of the latter find correlates of hand-based variables.

T. S. Kuhn captures the present situation: “proponents of competing paradigms practice their trades in different worlds ... the two groups of scientists see different things when they look from the same point in the same direction.”¹⁷ With regard to motor cortex function, neural correlates of hand direction are seen as evidence by one group that hand direction is an important and potentially dominant signal, whereas the other group views this observation as an obvious by-product of neural activity that controls muscles to move the limb.

These differences of opinion are important for understanding not only the function of motor cortex, but also the function of other cortical and subcortical regions of the CNS, such as the spinal cord. At the extremes, the spinal cord could be viewed as the central location where all decisions on the details of motor selection are generated from a simple descending command specifying the global goal of the task. Alter-