



How to read neuron-dropping curves?

Mikhail A. Lebedev*

Department of Neurobiology, Center for Neuroengineering, Duke University, Durham, NC, USA

*Correspondence: lebedev@neuro.duke.edu

Edited by:

Ioan Opris, Wake Forest University, USA

Reviewed by:

Manuel Casanova, University of Louisville, USA

Yoshio Sakurai, Kyoto University, Japan

Ioan Opris, Wake Forest University, USA

Keywords: neuron-dropping curve, brain-machine interface, large-scale recording, neuroprosthetic device, neuronal noise, neuronal tuning, neuronal ensemble recordings

Methods for decoding information from neuronal signals (Kim et al., 2006; Quiroga and Panzeri, 2009) have attracted much interest in recent years, in a large part due to the rapid development of brain-machine interfaces (BMIs) (Lebedev and Nicolelis, 2006; Lebedev, 2014). BMIs strive to restore brain function to disabled patients or even augment function to healthy individuals by establishing direct communications between the brain and external devices, such as computers and robotic limbs. Accordingly, BMI decoders convert neuronal activity into font characters (Birbaumer et al., 1999), position and velocity of a robot (Carmena et al., 2003; Velliste et al., 2008), sound (Guenther et al., 2009), etc.

Large-scale recordings, i.e., simultaneous recordings from as many neurons and as many brain areas as possible, have been suggested as a fundamental approach to improve BMI decoding (Chapin, 2004; Nicolelis and Lebedev, 2009; Schwarz et al., 2014). To this end, the dependency of decoding accuracy on the number of recorded neurons is often quantified as a neuronal dropping curve (NDC) (Wessberg et al., 2000). The term “neuron dropping” refers to the procedure, where neurons are randomly removed from the sample until there is only one neuron left. In this analysis, large neuronal populations usually outperform small ones, the result that accords with the theories of distributed neural processing (Rumelhart et al., 1986).

In addition to BMIs based on large-scale recordings, several studies have adapted an alternative approach, where a BMI is driven by a small neuronal population or even a single neuron that

plastically adapts to improve BMI performance (Ganguly and Carmena, 2009; Moritz and Fetz, 2011). In such BMIs, a small number of neurons serve as a final common path (Sherrington, 1906) to which inputs from a vast brain network converge.

While the utility of large-scale BMIs vs. small-scale BMIs has not been thoroughly investigated, several studies reported that information transfer by BMIs starts to saturate after the population size reaches approximately 50 neurons (Sanchez et al., 2004; Batista et al., 2008; Cunningham et al., 2008; Tehovnik et al., 2013). In one paper, this result was interpreted as mass effect principle, i.e., stabilization of BMI performance after neuronal sample reaches a critical mass (Nicolelis and Lebedev, 2009). However, another recent paper claimed that large-scale recordings cannot improve BMI performance because of the saturation (Tehovnik et al., 2013). This controversy prompted me to clarify here what NDCs show, whether or not they saturate, and how they can be applied to analyze BMI performance.

SIMULATING NDCs IN MATLAB

To produce illustrations of NDC characteristics (Figure 1), I used a simple simulation in MATLAB. The simulation was conducted over a time interval, *Interval*, which consisted of two parts: *Training_interval* used to train the decoder and *Test_interval* used to conduct decoding:

$$Interval = [Training\ Test];$$

To mimic typical BMI conditions, both intervals were set to

10 min, while the sampling rate was 10 Hz.

Neuronal rates were simulated as mixtures of signal and noise:

$$\begin{aligned} Rate(:, i) = & Signal \\ & + Asynchronous_noise(:, i) \\ & + C(i)*Common_noise; \end{aligned}$$

where *Rate* is discharge rate, *i* is index that enumerates neurons, *Asynchronous_noise* is noise unique to each neuron, *Common_noise* is noise common to all neurons, and *C(i)* is the amplitude of common noise for each individual neuron. *Signal* and *Common_noise* were simulated as mixtures of sinusoids. *Asynchronous_noise* was produced by MATLAB random number generator.

Simulated neuronal rates are shown in the left part of Figure 1A. For the purpose of this simulation, it was not important to reproduce single-unit activity in all details (e.g., mimic refractory periods). In fact, many BMIs discard these details and utilize multi-units (Chestek et al., 2011).

I used multiple linear regression as the decoder (Wessberg et al., 2000; Carmena et al., 2003; Lebedev et al., 2005). This decoder was trained using MATLAB *regress* function:

$$\begin{aligned} B = & regress(Signal(Training), \\ & [Rate(Training, :) ones \\ & (size(Training'))]); \end{aligned}$$

where *B* returned regression weights (one weight per neuron).

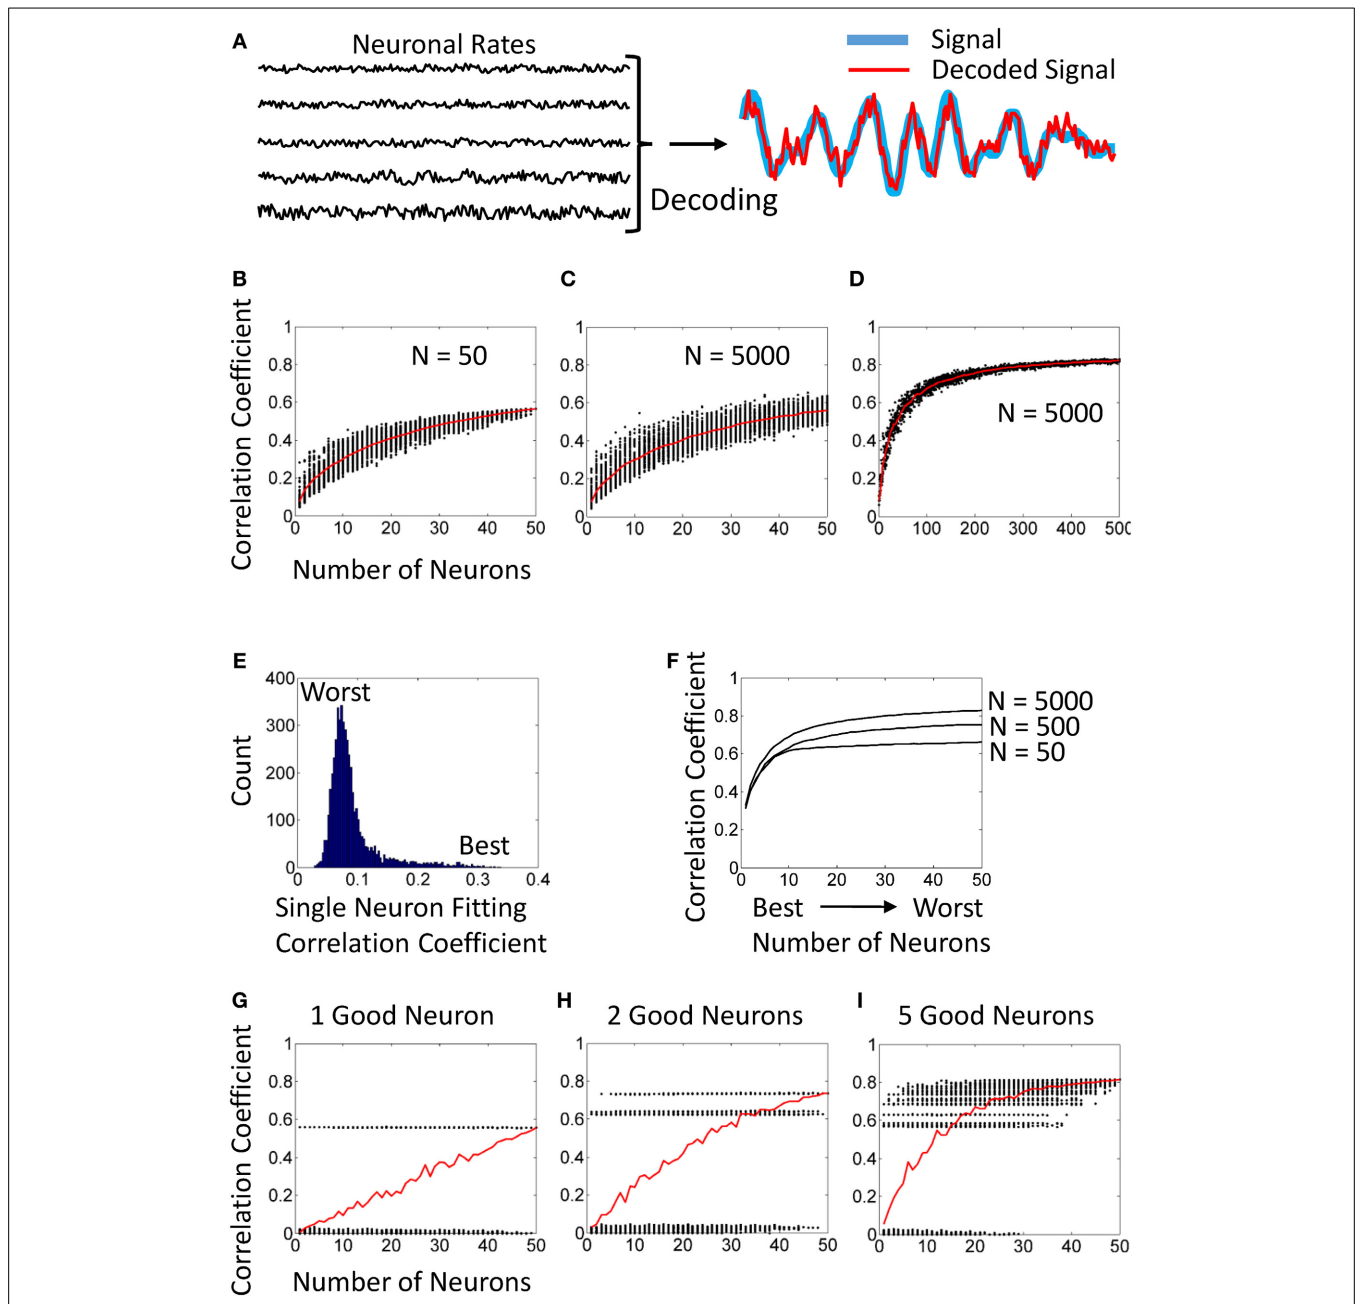


FIGURE 1 | Simulated NDCs. (A) Simulated neuronal rates (left) and their utilization for decoding (right). (B) An NDC generated from a population of 50 neurons. Each dot corresponds to a randomly drawn neuronal subpopulation. (C) An NDC generated from a population of 5000 neurons with the same statistical distributions of tuning characteristics as for 50 neurons in panel (B). (D) The same analysis as in (C) for subpopulation sizes from 1 to 500. (E) Distribution of single-neuron correlation coefficients for fitting. Computed for the same population of 5000 neurons as in panels (C,D). (F) Ranked NDCs for populations of 50, 500, and 5000 neurons. The best tuned neuron was included first, followed by second best, etc. (G) An NDC for a population of 50 neurons, where one neuron is tuned, and 49 produce noise. (H) An NDC for 2 good and 48 noisy neurons. (I) An NDC for 5 good and 45 noisy neurons.

The decoded signal was calculated as:

$$Decoded_signal = [Rate\ ones \\ (size(Rate, 1))] * B;$$

Note that the values of *Decoded_signal* for the training interval represent fitting (i.e.,

decoding is produced from the data that were used to train the decoder), whereas the values for the test interval correspond to decoding *per se* because they are derived from new data. **Figure 1A** shows example traces of the signal (blue curve) and decoded signal (red curve).

Decoding accuracy was evaluated as correlation coefficient, *R*, for the test interval:

$$R = corrcoeff (Decoded_signal (Test), \\ signal (Test); R = R(1, 2);$$

DO NDCs SATURATE?

BMI decoding improves with the size of neuronal sample for a simple reason: noise is reduced and signal is enhanced by ensemble averaging. For the asynchronous noise, overall noise reduction is approximately proportional to the square root of the sample size. Therefore, adding more and more neurons to a population should result in a continuous improvement of decoding. Why, then, do NDCs saturate?

Figure 1B illustrates a typical NDC with a saturation effect. Here, different-size subpopulations were randomly drawn from a fixed neuronal sample ($N = 50$). The scatterplot represents decoding accuracy for each subpopulation as a black dot, and the red curve is the average NDC. Note that the scatter becomes narrow when subpopulation size approaches 50, which seems to indicate that decoding accuracy has saturated.

However, the analysis of **Figure 1B** is flawed because the data points to the right part of the curve are not statistically representative. Indeed, when the subpopulation size is close to 50, repeated draws return mostly the same neurons. For example, in any two subpopulations of size 49, 48 neurons would be identical. Therefore, the NDC of **Figure 1B** converges to the value that depicts the performance of a fixed neuronal sample of $N = 50$ rather than representing random draws of such sample.

A different pattern is revealed when neuronal subpopulations are drawn from a much larger neuronal sample. In **Figure 1C**, subpopulations (1 to 50 neurons) were drawn from a fixed sample of 5000 neurons. This NDC does not saturate. In particular, the scatter of dots does not narrow down as dramatically as in **Figure 1B**, and correctly reflects the standard deviation for different subpopulation sizes.

Furthermore, subpopulation sizes from 1 to 500 are explored in **Figure 1D**. Here the original sample of 5000 neurons is the same as in **Figure 1C**. Clearly, this NDC indicates continuous improvement in decoding beyond 50 neurons.

In conclusion, NDCs lose statistical validity when subpopulation size gets comparable with the fixed population from which it is drawn. This may produce an illusion of saturation.

SOME NEURONS ARE MORE IMPORTANT

NDCs are usually computed as mean values for randomly drawn neuronal subpopulations. This representation masks the fact that individual neurons contribute unevenly to the decoding. There are leaders that represent parameters of interest particularly well, and there are noisy neurons with very small contributions to the decoding. Decoding usually improves if only the good neurons are utilized, and the noisy ones are discarded (Sanchez et al., 2004; Westwick et al., 2006).

Contribution of individual neurons to decoding (also called importance or sensitivity) can be estimated by running the decoder separately for each neuron. Here, depending on the analysis, individual correlation coefficients can be measured for the training interval (i.e., fitting) or for the test interval (i.e., decoding). These two metrics are usually very similar. **Figure 1E** shows a distribution of correlation coefficients for fitting for 5000 individual neurons. Notice a distribution “tail” that corresponds to particularly good neurons.

After individual performance is evaluated, neurons can be ranked by their contributions to decoding. Furthermore, NDC subpopulations can be constructed starting with the best neuron, then adding the second and so on (Sanchez et al., 2004; Lebedev et al., 2008). Examples of such rank-ordered NDCs are shown in **Figure 1D**. Here, the analysis was performed for populations of 50, 500, and 5000 neurons. Notice that in each of these cases, the highest ranked 30–50 neurons performed practically as well, as the entire population. Still, the draws from 5000 neurons always outperformed 500 neurons, and 500 neurons outperformed 50 neurons; simply because better subpopulations could be picked when a large neuronal sample was available.

Notably, a selection of informative neurons for one behavioral parameter (e.g., hand coordinate) can be different from the selection for another parameter (e.g., gripping force or leg coordinate). Therefore, recordings from large neuronal populations become particularly important when several parameters need to be decoded simultaneously (Fitzsimmons et al., 2009).

Currently, little is known about types of neurons in cortical microcircuits

(Casanova, 2013; Opris, 2013) that could be more useful for BMIs. It seems reasonable to assume that output neurons of such microcircuits could provide high quality BMI signals, but this issue needs more investigation. It would be of interest to record from an entire microcircuit, e.g., from a single cortical column, to reconstruct information processing and representation using BMI methods.

AVERAGE NDCs MAY BE MISLEADING

Although average NDCs usually indicate a gradual improvement in decoding accuracy when more and more neurons are added, this representation may be misleading as it conceals the fact that there are only a few informative neurons in the population.

Figure 1G illustrates a population of 50 neurons with only one informative neuron and 49 neurons generating noise. Even in this extreme case, an average NDC indicates a gradual improvement. This is because the probability of the single good neuron to be present in a subpopulation increases with the subpopulation size. Similar gradually rising NDCs can be obtained with 2 (**Figure 1H**) and 5 (**Figure 1I**) good neurons.

These examples show that for the analysis to be complete, an average NDC should be supplemented by a ranked NDC (e.g., **Figure 1F**) and a plot of individual contributions of different neurons (**Figure 1E**).

COMMON NOISE

Unlike asynchronous noise, common noise does not attenuate when firing rates of many neurons are averaged. Correlated variability in neuronal firing has been described in many cortical areas (Abbott and Dayan, 1999; Hansen et al., 2012; Opris et al., 2012). Whereas such correlations have important functions, they may be detrimental to BMI decoding if they are unrelated to the parameter being decoded. Interestingly, transition to online BMI control is accompanied by an increase in correlated variability (Nicolelis and Lebedev, 2009; Ifft et al., 2013).

Decoding in the presence of common noise can be improved by recording from multiple brain areas (Lebedev and Nicolelis, 2006; Nicolelis and Lebedev, 2009) because inter-area correlations are weaker than intra-area correlations

(Ifft et al., 2013). Curiously, adding neurons that are poorly tuned to a parameter of interest, but have a considerable common noise, can improve the decoding. Indeed, the decoder could use such neurons to recognize common noise and to subtract it from the contribution of well-tuned neurons.

Overall, although handling common noise is not as straightforward as handling asynchronous noise, using large populations of neurons is beneficial here, as well.

OFFLINE vs. ONLINE NDCs

NDCs are often used to support the claim that BMIs perform better when they utilize large neuronal ensembles (Nicoletis and Lebedev, 2009). However, in the majority of cases NDCs are calculated offline instead of testing neuronal ensembles of different size in real-time BMIs. To the best of my knowledge, ensembles of different size were tested online in just one study (Ganguly and Carmena, 2009), where removal of neurons from an initial population of 15 neurons deteriorated monkey performance in a reaching task controlled through a BMI. Apparently, more studies of this kind will be needed in the future to clarify this issue in more detail.

An additional caveat of NDC analysis is related to its application to the data obtained during real-time BMI control. Although it might be tempting to apply a decoder or a tuning curve analysis (Ganguly and Carmena, 2009) to BMI control data, these considerations can easily get circular because the signal (e.g., cursor position) in such datasets has been already generated from neuronal activity by a decoding algorithm. To cope with this problem, NDCs could be calculated for a parameter that was not included in the decoding algorithm, for example target position (Ifft et al., 2013).

CONCLUSION

Although NDCs often appear to saturate after neuronal population size reaches a critical value, a more careful consideration indicates that this effect may be an artifact of the overall limited neuronal sample. Large-scale neuronal recordings appear to be a realistic way to attain accuracy and versatility of BMIs.

REFERENCES

- Abbott, L. F., and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Comput.* 11, 91–101. doi: 10.1162/089976699300016827
- Batista, A. P., Yu, B. M., Santhanam, G., Ryu, S. I., Afshar, A., and Shenoy, K. V. (2008). Cortical neural prosthesis performance improves when eye position is monitored. *Neural Syst. Rehabil. Eng. IEEE Trans.* 16, 24–31. doi: 10.1109/TNSRE.2007.906958
- Birbaumer, N., Ghanayim, N., Hinterberger, T., Iversen, I., Kotchoubey, B., Kübler, A., et al. (1999). A spelling device for the paralysed. *Nature* 398, 297–298. doi: 10.1038/18581
- Carmena, J. M., Lebedev, M. A., Crist, R. E., O'Doherty, J. E., Santucci, D. M., Dimitrov, D. F., et al. (2003). Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* 1:e42. doi: 10.1371/journal.pbio.0000042
- Casanova, M. F. (2013). Canonical circuits of the cerebral cortex as enablers of neuroprosthetics. *Front. Syst. Neurosci.* 7:77. doi: 10.3389/fnsys.2013.00077
- Chapin, J. K. (2004). Using multi-neuron population recordings for neural prosthetics. *Nat. Neurosci.* 7, 452–455. doi: 10.1038/nrn1234
- Chestek, C. A., Gilja, V., Nuyujukian, P., Foster, J. D., Fan, J. M., Kaufman, M. T., et al. (2011). Long-term stability of neural prosthetic control signals from silicon cortical arrays in rhesus macaque motor cortex. *J. Neural Eng.* 8:045005. doi: 10.1088/1741-2560/8/4/045005
- Cunningham, J. P., Yu, B. M., Gilja, V., Ryu, S. I., and Shenoy, K. V. (2008). Toward optimal target placement for neural prosthetic devices. *J. Neurophysiol.* 100, 3445–3457. doi: 10.1152/jn.90833.2008
- Fitzsimmons, N. A., Lebedev, M. A., Peikon, I. D., and Nicoletis, M. A. (2009). Extracting kinematic parameters for monkey bipedal walking from cortical neuronal ensemble activity. *Front. Integr. Neurosci.* 3:3. doi: 10.3389/fnint.2009.03.009
- Ganguly, K., and Carmena, J. M. (2009). Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* 7:e1000153. doi: 10.1371/journal.pbio.1000153
- Guenther, F. H., Brumberg, J. S., Wright, E. J., Nieto-Castanon, A., Tourville, J. A., Panko, M., et al. (2009). A wireless brain-machine interface for real-time speech synthesis. *PLoS ONE* 4:e8218. doi: 10.1371/journal.pone.0008218
- Hansen, B. J., Chelaru, M. I., and Dragoi, V. (2012). Correlated variability in laminar cortical circuits. *Neuron* 76, 590–602. doi: 10.1016/j.neuron.2012.08.029
- Ifft, P. J., Shokur, S., Li, Z., Lebedev, M. A., and Nicoletis, M. A. (2013). A Brain-machine interface enables bimanual arm movements in monkeys. *Sci. Transl. Med.* 5, 210ra154. doi: 10.1126/scitranslmed.3006159
- Kim, S. P., Sanchez, J. C., Rao, Y. N., Erdogmus, D., Carmena, J. M., Lebedev, M. A., et al. (2006). A comparison of optimal MIMO linear and nonlinear models for brain-machine interfaces. *J. Neural Eng.* 3, 145. doi: 10.1088/1741-2560/3/2/009
- Lebedev, M. A. (2014). Brain-machine interfaces: an overview. *Transl. Neurosci.* 5, 99–110. doi: 10.2478/s13380-014-0212-z
- Lebedev, M. A., Carmena, J. M., O'Doherty, J. E., Zacksenhouse, M., Henriquez, C. S., Principe, J. C., et al. (2005). Cortical ensemble adaptation to represent velocity of an artificial actuator controlled by a brain-machine interface. *J. Neurosci.* 25, 4681–4693. doi: 10.1523/JNEUROSCI.4088-04.2005
- Lebedev, M. A., and Nicolelis, M. A. L. (2006). Brain-machine interfaces: past, present and future. *Trends Neurosci.* 29, 536–546. doi: 10.1016/j.tins.2006.07.004
- Lebedev, M. A., O'Doherty, J. E., and Nicolelis, M. A. (2008). Decoding of temporal intervals from cortical ensemble activity. *J. Neurophysiol.* 99, 166–186. doi: 10.1152/jn.00734.2007
- Moritz, C. T., and Fetz, E. E. (2011). Volitional control of single cortical neurons in a brain-machine interface. *J. Neural Eng.* 8:025017. doi: 10.1088/1741-2560/8/2/025017
- Nicoletis, M. A. L., and Lebedev, M. A. (2009). Principles of neural ensemble physiology underlying the operation of brain-machine interfaces. *Nat. Rev. Neurosci.* 10, 530–540. doi: 10.1038/nrn2653
- Opris, I. (2013). Inter-laminar microcircuits across neocortex: repair and augmentation. *Front. Syst. Neurosci.* 7:80. doi: 10.3389/fnsys.2013.00080
- Opris, I., Hampson, R. E., Gerhardt, G. A., Berger, T. W., and Deadwyler, S. A. (2012). Columnar processing in primate pFC: evidence for executive control microcircuits. *J. Cogn. Neurosci.* 24, 2334–2347. doi: 10.1162/jocn_a_00307
- Quiroga, R. Q., and Panzeri, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nat. Rev. Neurosci.* 10, 173–185. doi: 10.1038/nrn2578
- Rumelhart, D. E., Hinton, G. E., and McClelland, J. L. (1986). A general framework for parallel distributed processing. *Parallel Distrib. Process.* 1, 45–76.
- Sanchez, J. C., Carmena, J. M., Lebedev, M. A., Nicoletis, M. A., Harris, J. G., and Principe, J. C. (2004). Ascertaining the importance of neurons to develop better brain-machine interfaces. *Biomed. Eng. IEEE Trans.* 51, 943–953. doi: 10.1109/TBME.2004.827061
- Schwarz, D. A., Lebedev, M. A., Hanson, T. L., Dimitrov, D. F., Lehew, G., Meloy, J., et al. (2014). Chronic, wireless recordings of large-scale brain activity in freely moving rhesus monkeys. *Nat. Meth.* doi: 10.1038/nmeth.2936. [Epub ahead of print].
- Sherrington, C. S. (1906). *The Integrative Action of the Nervous System*. New Haven, CT: Yale University Press.
- Tehovnik, E. J., Woods, L. C., and Slocum, W. M. (2013). Transfer of information by BMI. *Neuroscience* 255, 134–146. doi: 10.1016/j.neuroscience.2013.10.003
- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., and Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature* 453, 1098–1101. doi: 10.1038/nature06996

Wessberg, J., Stambaugh, C. R., Kralik, J. D., Beck, P. D., Laubach, M., Chapin, J. K., et al. (2000). Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408, 361–365. doi: 10.1038/35042582

Westwick, D. T., Pohlmeier, E. A., Solla, S. A., Miller, L. E., and Perreault, E. J. (2006). Identification of multiple-input systems with highly coupled inputs: application to EMG prediction from multiple intracortical electrodes. *Neural Comput.* 18, 329–355. doi: 10.1162/089976606775093855

Conflict of Interest Statement: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 06 May 2014; accepted: 12 May 2014; published online: 30 May 2014.

Citation: Lebedev MA (2014) How to read neuron-dropping curves? *Front. Syst. Neurosci.* 8:102. doi: 10.3389/fnsys.2014.00102

This article was submitted to the journal Frontiers in Systems Neuroscience.

Copyright © 2014 Lebedev. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.