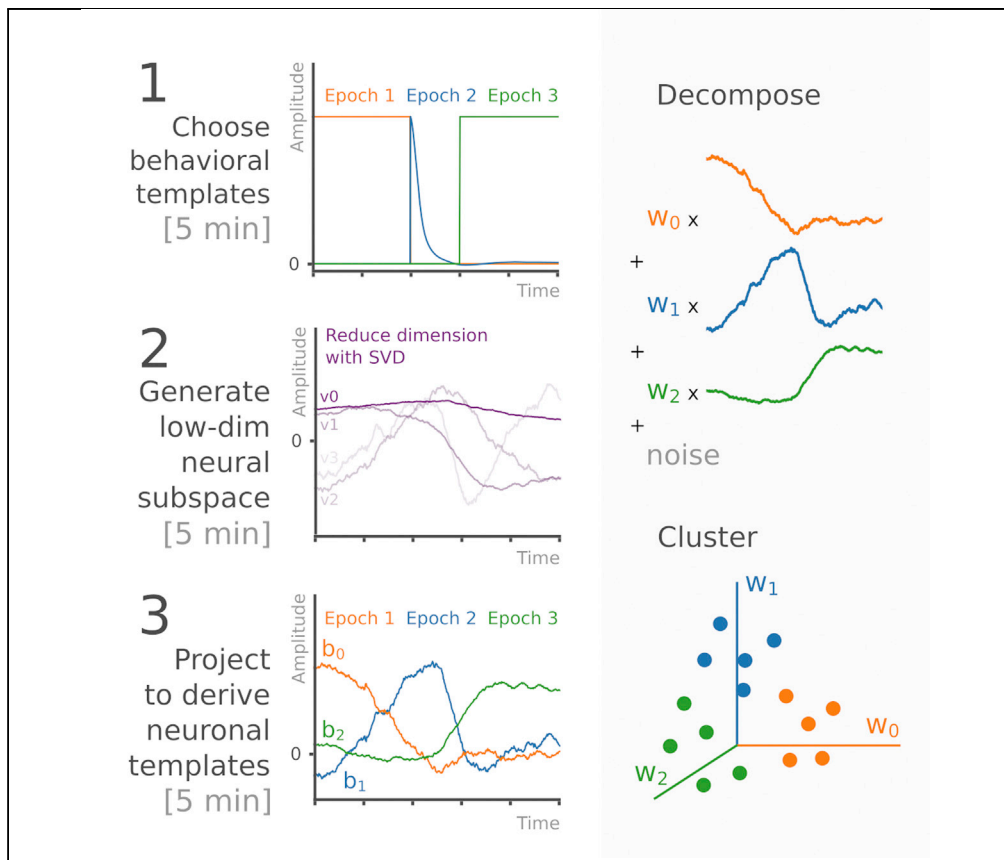


Protocol

Algebraic approach for subspace decomposition and clustering of neural activity



We developed an approach to decompose neuronal signals into disjoint components, corresponding to task- or event-based epochs. This protocol describes how to project behavioral templates onto a low-dimensional subspace of neuronal responses to derive neuronal templates, then how to decompose and cluster neuronal responses using these derived templates. We outline these steps on complementary datasets of calcium imaging and spiking activity. Our approach relies on fundamental, linear algebraic principles and is adaptive to the temporal structure of the neural data.

Publisher's note: Undertaking any experimental protocol requires adherence to local institutional guidelines for laboratory safety and ethics.

Elie M. Adam,
Mriganka Sur

eadam@mit.edu (E.M.A.)
msur@mit.edu (M.S.)

Highlights

Apply singular value decomposition to derive a low-dimensional neural response subspace

Project behavioral templates onto the subspace to derive neuronal templates

Every neuronal response can be decomposed as a weighted sum of neuronal templates

The derived weights can be used for clustering the neurons

Adam & Sur, STAR Protocols 3,
101841
December 16, 2022 © 2022
The Author(s).
<https://doi.org/10.1016/j.xpro.2022.101841>



Protocol

Algebraic approach for subspace decomposition and clustering of neural activity

Elie M. Adam^{1,2,3,*} and Mriganka Sur^{1,*}¹Picower Institute for Learning and Memory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA²Technical contact³Lead contact*Correspondence: eadam@mit.edu (E.M.A.), msur@mit.edu (M.S.)
<https://doi.org/10.1016/j.xpro.2022.101841>

SUMMARY

We developed an approach to decompose neuronal signals into disjoint components, corresponding to task- or event-based epochs. This protocol describes how to project behavioral templates onto a low-dimensional subspace of neuronal responses to derive neuronal templates, then how to decompose and cluster neuronal responses using these derived templates. We outline these steps on complementary datasets of calcium imaging and spiking activity. Our approach relies on fundamental, linear algebraic principles and is adaptive to the temporal structure of the neural data.

For complete details on the use and execution of this protocol, please refer to Adam et al. (2022).¹

BEFORE YOU BEGIN

Advances in systems neuroscience rely heavily on linking neuronal activity to observable outcomes, such as behavior. Technological advances have allowed for simultaneous recording of increasing numbers of neurons. The analysis of these large population data sets has fueled the need for techniques that reduce dimensionality of the neuronal response space and allow clustering of high-dimensional population dynamics while linking them to behavioral correlates.^{2–4} Among current methodologies, principal component analysis (PCA) has been heavily employed to reduce dynamics to dimensions along which the variability in neuronal response is maximized.^{4,5} However, with PCA being unsupervised, the principal components cannot always immediately be tied to behavioral outcomes, as these components are often derived by only considering neuronal activity. Generalized linear models (GLMs) tend to be supervised models and have also provided means of clustering neuronal responses, linking them to features of the task by employing templates^{4,6} (or filters). In GLMs, dimensionality is intrinsically reduced through the templates. However, it is unclear how to best select the training signals from which the templates are fitted to capture a desired neuronal activity feature, especially in the absence of obvious task structure signal. Choosing these ‘design’ training signals can be arbitrary, requiring insight and experience to execute. In this protocol, we developed a method that encapsulates the dimensionality reduction observed in PCA with template-based interpretation often offered through GLMs to decompose and cluster neuronal responses, while linking the decomposition and clustering to behavioral correlates. Our method relies on fundamental linear algebraic principles, and can be easily applied to a wide range of situations.

The protocol below describes the preparatory steps for analyzing the calcium data used in Adam et al.¹ We also apply similar steps for analyzing electrophysiological recordings. In general, the same preparation applies for any response dataset that is temporally structured into epochs.



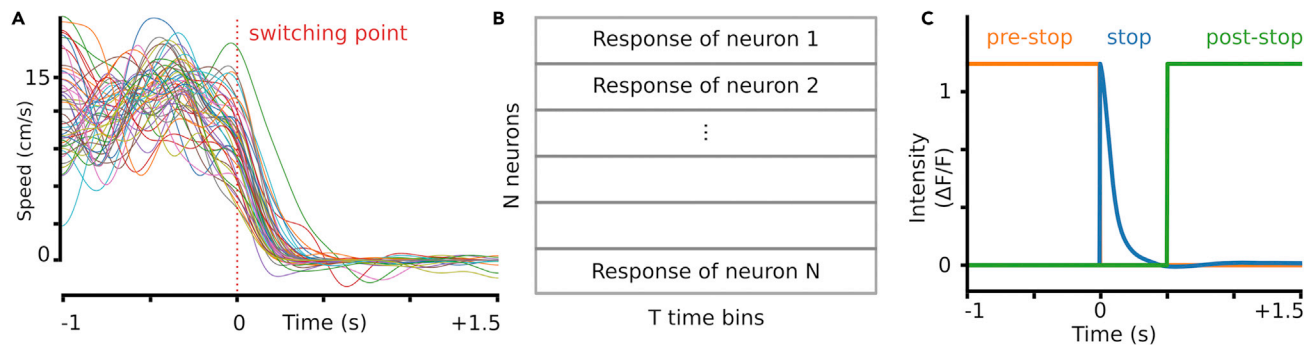


Figure 1. Preparing the data and selecting the behavioral templates

(A) Examples of speed traces in landmark-stop windows aligned to the last peak in velocity before stopping indicating the ‘switching point’ (30 trials drawn from all the sessions across animals. N=10 mice, 3 sessions each). Adapted from Adam et al.¹ Figure 2B.

(B) Schematic of the prepared data matrix M.

(C) Graph showing the three ideal behavioral template signals for a landmark-stop window, **template_pre**, **template_stop** and **template_post**. Adapted from Adam et al.¹ Figure S3A.

The data analyzed here consists of neuronal responses recorded during behavioral sessions where mice underwent a visually guided locomotion task. Mice were trained to run towards a visual landmark, then stop and wait at the landmark for 1.5 s to collect a reward.¹ The preparation phases consist of establishing a matrix of neuronal responses and selecting templates representing task epochs to use for decomposition and clustering.

Prepare the data matrix of neuronal responses to decompose

⌚ Timing: 5 min (for step 1 and 2)

This section describes the steps needed prepare the data matrix of neuronal responses which will be decomposed and clustered.

1. Segment the neural data into task phases windows (Figure 1A).

Note: In our example dataset, a mouse is running towards a landmark, and stops at the landmark and waits to collect reward. A behavioral session consists of time-evolving velocity and position trajectories of the mouse over each trial.

- a. Determine the locomotion stops at the landmark which were rewarded in a behavioral session.
- b. For each such locomotion stop, determine the ‘switching point’ in locomotion corresponding to the time of the last major velocity peak (above 25% of the maximum velocity within 200 ms window before sustained zero velocity) before the rewarded stop.

Note: Other criteria to indicate stopping time are also valid approaches.

- c. From each behavioral session, define landmark-stop windows for each rewarded stop with 2.5 s duration starting 1 s before the switching point and ending 1.5 s after.

Note: The duration of 1.5 s after stop initiation corresponds to the necessary waiting time to collect reward.

Note: In general, the goal is to delimitate temporal windows of similar duration, that are all aligned to a particular reference point (e.g., initiation of action, onset of stimulation, etc).

2. Prepare the data matrix of neuronal response (Figure 1B).
 - a. For each neuron, compute the average change in fluorescence DFF over landmark-stop windows.

Note: $DFF = (F_t(t) - F_0)/F_0$ where F_t is the fluorescence value at a given time and F_0 is the baseline fluorescence value.

- b. Normalize each average response to a maximum DFF of 1.
- c. Form an N-by-T (N: number of neurons and T: number of time bins) matrix M where each row corresponded to the normalized average DFF of a neuron.

Note: Data matrix preparation should be tailored to the specific application but, at the end, should give an N-by-T matrix of responses.

Note: Some form of normalization as performed in 2.b can be essential to ensure that the results are not swayed by only a few extreme/high-amplitude responses. However, it is often desirable for the result to capture high-amplitude responses if these are common as they may indicate strongly task-responsive neurons. The researcher is then faced with a trade-off between undervaluing strong-responses and undervaluing weak-responses, which can be best decided by considering the data under investigation. When in doubt, it is safe to first attempt without normalization, and then assess the obtained solution.

Select behavioral templates over which the signals will be decomposed

⌚ Timing: 5 min (for step 3 and 4)

This section describes the steps needed to create behavioral templates, each representing an epoch in the behavior, over which the neuronal responses will be decomposed.

3. The landmark-stop window consists of three epochs: pre-stop, stop and post-stop. Define three ideal behavioral template signals, denoted by **template_pre**, **template_stop** and **template_post**, where each represents one of the epochs (Figures 1A and 1C).
 - a. Set **template_pre** to be equal to 1 for $-1 \leq t \leq 0$ s and 0 otherwise.
 - b. Set **template_stop** to be equal to the velocity averaged across animals and sessions for $0 \leq t \leq 1.5$ s and 0 otherwise.
 - c. Set **template_post** to be equal to 1 for $0.5 \leq t \leq 1.5$ s and 0 otherwise.

Note: In general, the templates can be chosen using the following rule. Choose one template for each behavioral epoch under investigation. For each epoch in the time window of the neuronal response, create a template which is equal to 1 during the epoch and zero otherwise. If, during that epoch, there exists a signal that can be used to characterize it (e.g., in our setting, the decreasing velocity during the stop epoch), then use a template that is equal to the signal during the epoch, and zero otherwise. Other methods are possible, and the subsequent processing (See “[step-by-step method details](#)” sections 1–6) in projecting the templates to a low-dimensional subspace to obtain neuronal templates ensures that different choices defined over the same epochs ought not have large effects on the results.

⚠ **CRITICAL:** The ideal templates can overlap in time at this stage, and thereby need not be orthogonal. A subsequent processing step will ensure that they are orthogonal to get a basis.

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and algorithms		
Python 3	python.org	Version > 3.7
Numpy	numpy.org	N/A
Other		
Dell OptiPlex	Dell.com	N/A

MATERIALS AND EQUIPMENT

No specific computing power is needed for our methodology other than enough processor speed and memory to load and manipulate the data.

Alternatives: We provide a step-by-step description of our method using the Python 3 language, but any other computing language (e.g., MATLAB) would work as well with applicable modifications to the syntax.

STEP-BY-STEP METHOD DETAILS

Our general method first identifies a low dimensional neural subspace that captures most of the energy in neuronal responses. It then takes the ideal templates that delimitate behavioral epochs and projects them onto that subspace to yield ideal neuronal signals representing each of the ideal templates. It then forms a basis from these ideal neuronal signals, and decomposes the data onto this basis to recover coefficients. Finally, it uses these coefficients to cluster the responses.

Reduce dimensions and identify a neural subspace that captures most of the energy in neuronal responses

⌚ Timing: 5 min (for step 1, 2 and 3)

This section describes the steps needed to derive a low-dimensional subspace of neuronal responses that captures most of the energy in the responses.

1. Perform a singular value decomposition⁷⁻⁹ of the prepared data matrix M as USV^T where the matrix S is a rectangular diagonal matrix of singular values, and the matrices U and V are orthonormal matrices (Figure 2A).

```
> import numpy as np
> U, s, VT = np.linalg.svd(M, full_matrices=False)
> S = np.zeros([U.shape[1], VT.shape[0]])
> for i in range(len(s)):
    S[i, i] = s[i]
```

Note: Check that USV^T is equal to M to verify that all went correctly with the computation.

2. Each singular value in S corresponds to a temporal neuronal response in V^T (Figure 2A). Select the responses \mathbf{v}_0 , \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 corresponding to the highest 4 singular values in S , which together will span a 4-dimensional space (Figure 2B).

```
> v0 = VT[0, :]  
> v1 = VT[1, :]  
> v2 = VT[2, :]  
> v3 = VT[3, :]
```

Note: We reasoned that we had (not necessarily independent) four-degrees of freedom in our analysis: baseline, pre-stop, stop and post-stop activity, and thus sought to begin with a 4-dimensional space.

Note: In general, the number of dimensions can be chosen as follows. If K ideal behavioral templates were selected in the data preparation phase, choose $K+1$ temporal neuronal responses in V^T . The reasoning is that each template would need one distinct dimension, and the remaining dimension would correspond to baseline activity which plays a role as offset, better connecting the behavioral templates to neuronal responses despite the difference in baselines. In the next steps 4–6, the ideal behavioral templates will be projected onto the four (or $K+1$) dimensional subspace, and thus the choice of dimensions only affects the final shape of the templates used for clustering, and will not alter the original neuronal responses.

Note: In some applications, it is also possible that the neural data is inherently high dimensional, and that relevant aspects are captured in dimensions with lower singular values. In such setting, the researcher may not restrict to the highest four (or $K+1$) singular values but other four (or $K+1$) singular values that yield interpretable neuronal responses.

Note: Plot the signals v_0, \dots, v_3 to assess and verify that they represent temporally evolving trajectories.

3. Compute the fraction of energy of the neuronal responses captured by this subspace by squaring all the singular values σ_i on the diagonal in S , and dividing the sum of the greatest 4 by the sum of all the squared singular values.^{8,9}

```
> energy = sum(s[:4]**2)  
> total_energy = sum(s**2)  
> fraction = energy / total_energy
```

Note: The resulting subspace contained more than 80% of the calcium signal energy for the whole population.

Note: In steps 4–6, the ideal behavioral templates will be projected onto the four (or $K+1$) dimensional subspace. Adding additional dimensions will facilitate capturing more of the shape of the ideal behavioral templates following the projection by capturing more energy in the neuronal responses. Generally, fewer dimensions will lead to a more interpretable and robust analysis that avoids over-fitting. However, in the case where either (i) the fraction of the energy captured is low or (ii) the obtained neuronal template derived in a subsequent step appears too coarse to capture the ideal behavioral template, then the researcher ought to increase the number of dimensions. In such a setting, a principled approach consists of cross-validation, where dimensionality is reduced for a random subset of neurons and then

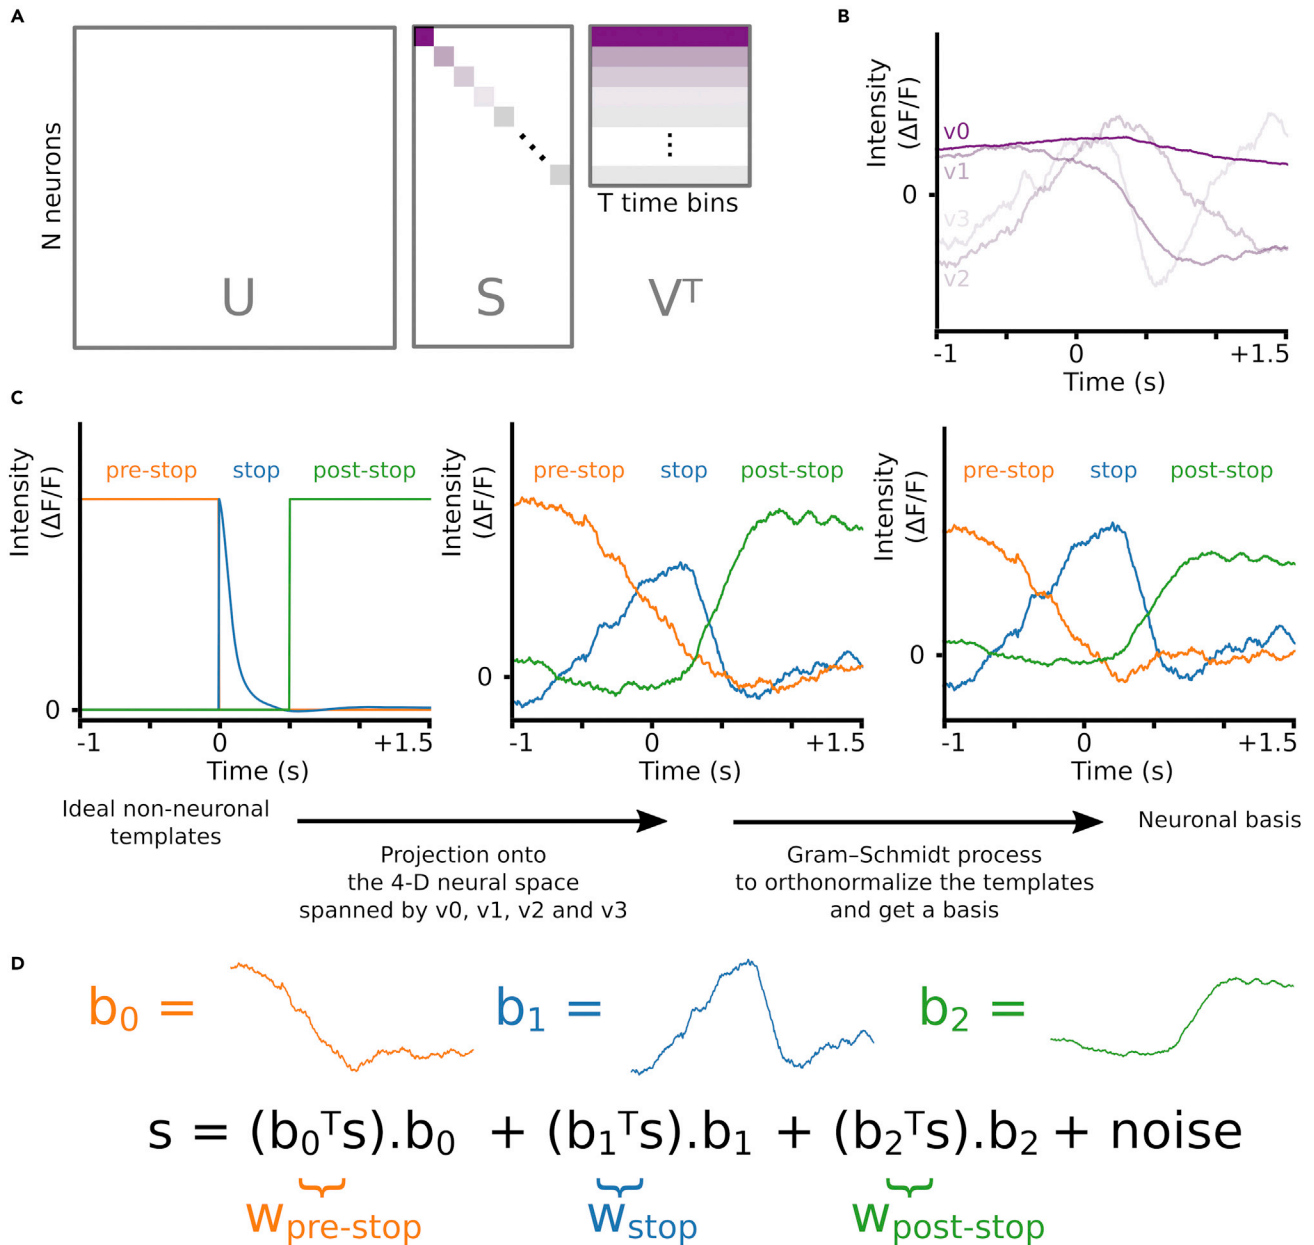


Figure 2. Decomposing neural signals in a low dimensional subspace

(A) Schematic illustrating singular value decomposition of a matrix. Each diagonal element in S corresponds to one row in V^T . These elements can be ordered by decreasing magnitude, with V^T re-arranged accordingly.

(B) Graph showing a basis for a four-dimensional subspace that captures more than 80% of the energy in the responses, all pooled together.

(C) Left. The three initial non-neuronal templates defined for pre-stop, stop and post-stop activity. Middle. The projections of the initial three non-neuronal templates onto a 4-dimensional subspace derived through low-rank approximation. These correspond to the neural responses that represent the non-neuronal templates. Right. The orthonormalized responses that form a basis, generating a 3-dimensional subspace, upon which we can decompose our signals. Adapted from Adam et al.¹ Figure S3A.

(D) Schematic showing the decomposition of a neuronal signal into a linear combination of the basis signals with additional noise.

the captured energy is evaluated on the remaining neurons (See e.g., Bro et al.¹⁰). The reduction is performed over multiple iterations for each number of dimensions to determine at what dimensions noise is more likely being captured by the reduction.

Change the basis of the subspace to recover neuronal template signals

⌚ Timing: 5 min (for step 4, 5 and 6)

This section describes the steps needed to project the behavioral templates selected in the preparatory steps on the low-dimensional subspace, and form a basis upon which to decompose the neuronal signals.

- Take the inner-product (dot-product) of each ideal template signal `template_*` (in Figure 1C) with each temporal neuronal response `v0`, ..., `v3`, to compute a weight coefficient (Figure 2C).

```
> w_stop_0 = np.dot(template_stop, v0)
> w_stop_1 = np.dot(template_stop, v1)
...
> w_pre_0 = np.dot(template_pre, v0)
...
> w_post_0 = np.dot(template_post, v0)
```

- Define three neuronal template signals `v_stop`, `v_pre` and `v_post` as a sum of the 4 neuronal responses `v0`, ..., `v3`, each scaled by its corresponding coefficient in step 4 (Figure 2C).

```
> v_stop = w_stop_0*v0 + w_stop_1*v1 + w_stop_2*v2 + w_stop_3*v3
> v_pre = w_pre_0*v0 + w_pre_1*v1 + w_pre_2*v2 + w_pre_3*v3
> v_post = w_post_0*v0 + w_post_1*v1 + w_post_2*v2 + w_post_3*v3
```

Note: These neuronal template signals represent the best neuronal response in the 4-dimensional subspace that can explain the ideal template signals representing behavioral outcomes.

- Using the gram-schmidt orthonormalization process,⁹ transform `v_stop`, `v_pre` and `v_post` into an orthonormal basis `b_stop`, `b_pre` and `b_post` (Figure 2C).

```
> def norm(v):
    return np.sqrt(np.dot(v, v))
> v = v_stop
> b_stop = v/norm(v)
> v = v_pre - b_stop*np.dot(b_stop, v_pre)
> b_pre = v/norm(v)
> v = v_post - b_stop*np.dot(b_stop, v_post) - b_pre*np.dot(b_pre, v_post)
> b_post = v/norm(v)
```

Note: This step ensures that the templates are orthonormal, and thereby form a basis onto which we can project. Verify that the inner-product of any pair of these two basis vectors is

zero, and that the norm of each is equal to 1. These 3 basis signals will be used to decompose neuronal signals and then cluster them.

Note: We first reduced our space to 4 dimensions, and then further restricted to a 3-dimensional subspace within it. The first reduction relies purely on neural data, while the second relies on the behavioral data. We could have directly reduced to 3 dimensions, but we found that having an additional degree of freedom aids in finding the 3-dimension subspace that is best tailored to the information we want. Indeed, ideal behavioral templates have different baseline offsets compared to the neuronal responses, and the additional dimension aids in offsetting that baseline. As we are only interested in 3 epochs (pre-stop, stop and post-stop) in the task, we do not need to keep the fourth dimension after using it to derive the neuronal template signals.

Decompose neuronal responses onto the neuronal basis signals

⌚ Timing: 5 min (for step 7 and 8)

This section describes the steps needed to decompose the neuronal signals.

7. Take the inner-product (dot-product) of each neuronal response with each basis signal **b_stop**, **b_pre** and **b_post**, to compute a weight coefficient.

```
> coefs_stop = np.dot(M, b_stop)
> coefs_pre = np.dot(M, b_pre)
> coefs_post = np.dot(M, b_post)
```

Note: Each neuronal response can then be written as a weighted sum of the three basis signals plus some additional component orthogonal to the subspace (considered as noise). The weighted sum of the three basis signals gives an approximation of the original neural response, which resides in the 3-dimensional subspace spanned by the basis signals (Figure 2D).

8. Keep the neurons for which the fraction of their calcium signal energy (area of the squared signal) that came from the three-dimensional subspace is above a chosen threshold, and discard the rest from the analysis.

```
> energy_stop = coefs_stop**2
> energy_pre = coefs_pre**2
> energy_post = coefs_post**2
> percentages = (energy_stop+energy_pre+energy_post) / np.sum(s**2)
```

Note: For different approaches, one can either keep the neurons that have their energy ratio above a fixed percentage (e.g., greater than 80% or the fraction of energy captured by the subspace), or above twice the standard deviation from the mean energy, or above a determined percentile in the ratio distribution (e.g., highest 80%).

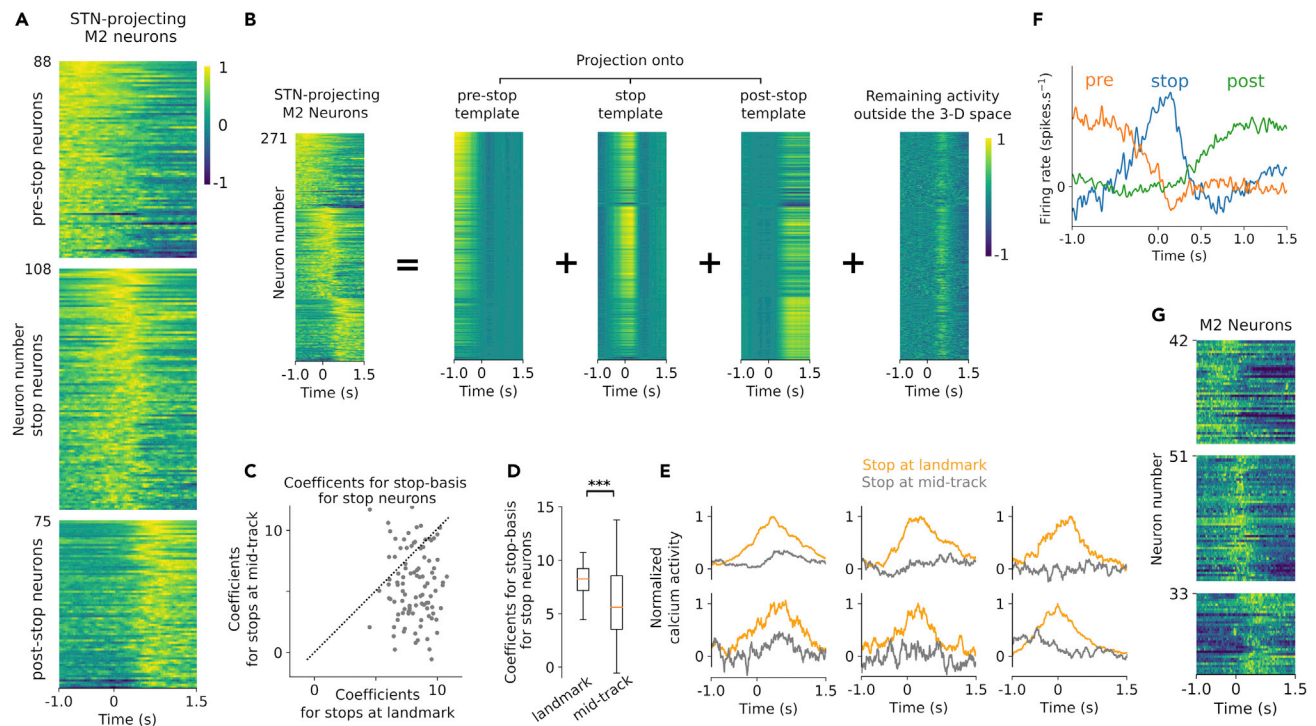


Figure 3. Neural response decomposition reveals three clusters of neurons in M2

(A) Plot showing STN-projecting M2 neurons ($n=271$), whose neural response energy (area under the squared signal) was more than 80% explained by the subspace, clustered into three groups (pre-stop, stop and post-stop) using the templates in Figure 2C. Adapted from Adam et al.¹ Figure 4D.

(B) Each neuronal response can be written as a combination of four components: the first three correspond to weighted versions of the basis functions, and the fourth corresponds to the activity that remains outside the three-dimensional subspace spanned by the basis. The weights on the templates are used to perform clustering. The remaining activity shows some surge after stopping. The 1-dimensional subspace containing the highest energy in the remaining activity will consist of a neural trajectory with activity concentrated after the stop. When added to the three basis functions (pre-stop, stop and post-stop), the new set will span the 4-dimensional space obtained initially. The decomposition is only used for clustering and not to post-process neural signals. Adapted from Adam et al.¹ Figure S3B.

(C) Scatterplot of the coefficient of contribution of the stop template in Figure 2C in spontaneous-stops versus landmark-stops. Each data point represents a stop- neuron ($N = 108$) taken from (A), and the coefficient represents the energy of the response along the dimension of the basis, obtained by taking the dot product of the calcium response with the stop basis vector, b_{stop} . Landmark stops have a higher coefficient than spontaneous stops, indicating that b_{stop} has a greater contribution to responses during landmark stops. Adapted from Adam et al.¹ Figure 2F.

(D) Box plots for the distribution of the coefficients in (C). The orange lines represent the respective medians. The means of the distributions are significantly different (Paired t-test, $***: p=1.09e-6$). Adapted from Adam et al.¹ Figure 2G.

(E) Examples of the calcium activity of six M2 neurons projecting to STN during stops at the landmark and in the middle of the track. The response for spontaneous-stops is normalized to the maximum value of the corresponding response for landmark-stops. Adapted from Adam et al.¹ Figure 2H.

(F) Graph showing 3 templates of ideal neuronal responses related to three epochs in the task: pre-stop, stop and post-stop. The data is derived from extracellular single-unit recordings in M2, and neuronal responses are obtained by deriving firing rates. Each neuronal response can be expressed as a weighted combination of these 3 templates. Adapted from Adam et al.¹ Figure 6B.

(G) Plots of all the M2 neurons whose neural response energy (area under the squared signal) was more than 80% explained by the subspace, clustered into three groups (pre-stop, stop and post-stop, from top to bottom) using the templates in (F). Adapted from Adam et al.¹ Figure 6C.

Note: In step 8, retaining the neurons for which the fraction of their calcium signal energy that came from the three-dimensional subspace is greater than 85% (roughly the fraction of energy captured by the 4-dimensional subspace) kept about 64% of all neurons we started with, and we performed clustering on them as described in steps 9 and 10 below. In these neurons, less than 15% of their energy is explained outside the 3-dimensional subspace, and we expect most of that remainder to be explained by the 4th dimension that complements b_{stop} , b_{pre} and b_{post} to span the 4-dimensional space generated by v_0 , v_1 , v_2 and v_3 . Indeed, we find that the remaining activity outside the subspace is concentrated after the stop (Figures 3A and 3B). As we investigated only 3 epochs (pre-stop, stop and post-stop) in Adam et al.,¹ we considered that 4th dimension to be ‘noise’ for our purposes. In case

we are interested in studying localized surges of activity after stops, that dimension becomes relevant and should be included in the reduced subspace yielding 4 dimensions. More generally, the choice of dimensions reduced to defines everything outside the reduced subspace as irrelevant noise. It is then important to choose the subspace to capture the features essential to the investigation.

Note: The discarded neurons have more than 15% of their calcium signal energy explained outside the three-dimensional subspace. These neurons might have responses related to pre-stop, stop and post-stop, but these responses are considered weak with respect to our criterion. However, these neurons might be tuned to other features of the task, and therefore other exclusion criteria may be considered if they are more adequate to the application at hand.

Cluster neurons based on decomposition weights

⌚ Timing: 5 min (for step 9 and 10)

This section describes the steps needed to cluster the neuronal signals, using the decomposition.

9. Multiply the weight coefficient `coefs_stop`, `coefs_pre` and `coefs_post`, respectively by the maximum value of the basis signals `v_stop`, `v_pre` and `v_post` to account for the template width and get a corrected weight coefficients.

```
> corrected_coefs_stop = coefs_stop*max(v_stop)
```

10. Attribute a neuron to one of three classes (pre-stop, stop and post-stop) whose corresponding corrected weight coefficient is highest.

Note: Most importantly, the calcium responses in each class were not reduced in dimensions; they were the original averages of the raw DFF traces taken over landmark stop windows. The dimensionality reduction was only used to cluster the neuronal responses, leaving the neural responses unaltered.

EXPECTED OUTCOMES

We expect this decomposition to enable clustering of neuronal responses. We applied it in Adam et al.¹ to cluster *in-vivo* calcium signals during behavior imaged using two-photon microscopy, and found three populations of neurons, each preferentially tuned to one epoch of the task (Figure 3A). Figure 3B shows the contribution of each of the basis vector `b_stop`, `b_pre` and `b_post` to every neuronal response, and highlights a residual noise. We further identified stops that occurred spontaneously and were not rewarded, and defined corresponding their stop-windows. By taking inner-product (dot product) of the averaged neural response during such stops with `v_stop`, we recovered coefficients for non-rewarded stops. Our method revealed a significant difference in the coefficients between rewarded and non-rewarded stops (Figures 3C and 3D), which indicated that many M2 neurons projecting to STN are active at visually-guided stops but not spontaneous stops (Figure 3E). Our method can also be applied to electrophysiological data where action potential firing rates, instead of calcium response intensity, are used as the neuronal responses. Single unit recordings in M2 reveal similar basis signals representing the epochs (Figure 3F) and clusters of neurons (Figure 3G).

LIMITATIONS

Our approach is a hybrid between a supervised- and an unsupervised approach. Ideal behavioral templates indicating task epochs need to be chosen at the beginning, which may alter results. From that perspective, our approach is opposite to the direct use of PCA or more generally

canonical correlation analysis¹¹ (CCA) which tends to be unsupervised. However, relying on chosen behavioral templates is crucial to guide the analysis towards the behavioral features or epochs under investigation. While PCA and CCA provide dimensions that maximize information content, using them in an unsupervised manner cannot ensure that these dimensions reflect neural information during specified epochs in the task. Furthermore, our technique relies on projecting these templates onto a low-dimensional neural response subspace, which will ensure that only the prominent features of these templates are used in the analysis, leading to more robust conclusions. Nevertheless, it is possible to incorporate PCA or CCA to our approach. Indeed, we recover a low dimensional subspace using directly SVD onto the data matrix and it is possible to replace this step with PCA or CCA and derive another suitable subspace. The behavioral templates can then be projected onto that newly obtained subspace, and the technique continues thereafter as described in this protocol. Furthermore, CCA can be alternatively used to bypass template projection by deriving, for each behavioral template, the 1st component (dimension) that maximally correlated with the template. A neural basis for decomposition can then be derived using the orthogonalization process described in this protocol.

Our approach also relies on the major patterns in the data that represent most of the energy in the neuronal responses, and can therefore miss more nuanced activity patterns. In particular, the remaining unused dimensions are considered to be task irrelevant. How many and which dimensions to consider becomes an essential issue that can unbiasedly be addressed through cross-validation, as described in step 3. Furthermore, it can be ill-posed to compare the obtained coefficients from Group A (e.g., rewarded stops) to those obtained in Group B (e.g., non-rewarded stops) when the neuronal templates are optimally derived through dimensionality reduction in Group A. The optimality condition may be a confound with any significant difference observed. It is best to perform additional analysis of difference significance relying directly on the raw signals or an alternative visualization and not only rely on the coefficients. On the other end, cross-validation can be employed, e.g., by using a fraction of the rewarded trials to derive the templates, and compute coefficients on the remaining trials, while considering all the neurons in the analysis. Furthermore, our approach can be prone to forcing an identification and interpretation of clusters that is not present in the data. The validity of the findings can again be justified through cross-validation techniques, where templates are derived using a subset of the behavioral windows (e.g., the rewarded stops) and then evaluated on the remaining withheld set of behavioral windows. Alternatively, investigating the distribution of coefficients can provide additional information on how clustered is the data.

Lastly, our approach provides a static decomposition of the data into separate behavioral epochs. It does not explicitly target dynamical evolution of the data. For instance, it is often desirable to decompose the neuronal signals into evolving latent (hidden) variables that reflect different physiological conditions. For such situations, combining our approach with state-space methods would yield the needed perspective. Such a combination is outside the scope of this protocol.

TROUBLESHOOTING

It can be useful to always check and plot the obtained vectors and templates, and assess whether they are consistent with expectations and insight. Throughout the step-by-step description, we added verify steps to minimize troubleshooting needs. Among the problems that can emerge are:

Problem 1

In step 2, the temporal neuronal response obtained do not appear to represent neuronal responses.

Potential solution

Unintentionally transposing the data matrix either in step 1 or step 2 will yield columns in V^T indexed by neuron number instead of time. Furthermore, unintentionally deriving a column from V instead of V^T in step 2 will yield an uninterpretable response. Be sure that no matrix is inadvertently transposed.

Problem 2

In step 2, the temporal neuronal responses are mostly capturing responses with very high amplitude, and are not capture more nuanced responses.

Potential solution

In the situation where the contribution of high amplitude response needs to be discount in the analysis, normalizing the responses is crucial before reducing dimensions. We discuss details and rationales for normalization as a note in step 3.

Problem 3

In step 3, the dimensionality reduction does not capture a large fraction of the energy.

Potential solution

There appears to be a lot of response heterogeneity in the neural responses. Increase the number of dimensions to capture more of this variability and increase the amount of captured energy.

Problem 4

In steps 5 and 6, the derived neuronal response templates do not appear to capture the ideal behavioral templates well.

Potential solution

The sought features through the behavioral templates might be refined, and best captured in higher dimensions. Try increasing the number of dimensions as noted in step 3. The additional first few dimensions that did not fully capture the behavioral templates should also inform you on additional coarser behavioral templates to consider that might be relevant to the investigation.

Problem 5

In step 7, the derived coefficients appear uninterpretable and wrongly scaled.

Potential solution

Before projecting to the low-dimensional neural subspace, it is important to ensure that the template neural responses form an orthonormal basis. Check that the dot product of every pair of template vectors is zero, and that the norm of each is equal to 1. A non-unitary norm would scale the coefficients inappropriately.

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Elie Adam (eadam@mit.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

All data reported in this paper will be shared by the [lead contact](#) upon request.

This paper does not report original code.

Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

ACKNOWLEDGMENTS

This work was supported by a JPB Foundation Fellowship (E.M.A.), NIH grants R01EY028219 and R01MH126351, the Picower Institute Innovation Fund, and the Simons Foundation Autism Research Initiative through the Simons Center for the Social Brain (M.S.). We thank Dr. Grayson Sipe and Karla A. Montejo for their careful reading of the manuscript.

AUTHOR CONTRIBUTIONS

E.M.A. and M.S. conceived the project. E.M.A. analyzed the data and developed the methodology. E.M.A. and M.S. wrote the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

1. Adam, E.M., Johns, T., and Sur, M. (2022). Dynamic control of visually guided locomotion through corticosubthalamic projections. *Cell Rep.* 40, 111139.
2. Saxena, S., and Cunningham, J.P. (2019). Towards the neural population doctrine. *Curr. Opin. Neurobiol.* 55, 103–111.
3. Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* 17, 1500–1509.
4. Pang, R., Lansdell, B.J., and Fairhall, A.L. (2016). Dimensionality reduction in neuroscience. *Curr. Biol.* 26, R656–R660.
5. Jolliffe, I.T., and Cadima, J. (2016). Principal component analysis: a review and recent developments. *Philos. Trans. A Math. Phys. Eng. Sci.* 374, 20150202.
6. Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454, 995–999.
7. Dahleh, M., Dahleh, M.A., and Verghese, G. (2004). Lectures on Dynamic Systems and Control. <https://dahleh.lids.mit.edu/lectures-on-dynamic-systems-and-control/>.
8. Horn, R.A., and Johnson, C.R. (2012). *Matrix Analysis*, 2nd ed. (Cambridge University Press).
9. Strang, G. (2016). *Introduction to Linear Algebra*, 6th edition (Wellesley-Cambridge Press).
10. Bro, R., Kjeldahl, K., Smilde, A.K., and Kiers, H.A.L. (2008). Cross-validation of component models: a critical look at current methods. *Anal. Bioanal. Chem.* 390, 1241–1251.
11. Zhuang, X., Yang, Z., and Cordes, D. (2020). A technical review of canonical correlation analysis for neuroscience applications. *Hum. Brain Mapp.* 41, 3807–3833.