Brain-like variability in convolutional neural networks reveals evidence-, uncertainty- and bias-driven decision-making

Johan Nakuci[1*]

[1]U.S. Army DEVCOM Army Research Laboratory, Aberdeen Proving Ground, Aberdeen, MD, 21005, USA

*Corresponding author. Email: jnakuci@gmail.com

Keywords: convolutional neural networks, trial-to-trial variability, perceptual decision-making, sensory evidence, choice-bias

Acknowledgements: We would like to thank Dobromir Rahnev for his thoughtful comments and suggestions. This research was supported by the U.S. Army DEVCOM Army Research Laboratory through mission funding (JOG) and army educational outreach program (W911SR-15-2-0001). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army DEVCOM Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Competing interests: Authors declare that they have no competing interests.

Author contributions:
Conceptualization: JN
Methodology: JN
Data Curation: JN
Visualization: JN
Funding acquisition: JN
Writing – original draft: JN
Writing – review & editing: JN

Abstract

Even when stimuli and tasks are held constant, brain activity fluctuates markedly across trials, yet it is not well understood how these fluctuations affect decision-making and behavior. Here we address this gap in knowledge using a convolutional neural network (CNN) trained on a perceptual decision-making task. Applying an established analytical framework, we show that CNN activity exhibits trial-level variability reflecting findings from human neuroimaging. Examination of the network's internal state revealed three distinct activation patterns reflecting: (i) decisions dominated by strong sensory evidence, (ii) decisions under low discriminability with weak or ambiguous evidence, and (iii) decisions in which sensory evidence was opposed by bias. Notably, choice bias shifted the decision boundary, improving performance despite low or conflicting evidence. Together, these findings show that variability in CNN activity can serve as a model for understanding how the brain transforms changes in internal states into adaptive behavior, providing a bridge between fluctuations in activity, decision-making, and behavior.

Introduction

Brain activity exhibits pervasive trial-to-trial variability despite consistent sensory inputs and task demands (1, 2). Although often dismissed as noise in decision-making models, yet growing evidence indicates that fluctuations in brain activity play a critical role in adapting behavior to the demands of a changing environment (3). However, it is not well understood how fluctuations in brain activity impact decision-making and behavior.

In fact, brain regions can exhibit stark variation in activity yet these fluctuations can have minimal impact on behavioral performance (4). For instance, during a perceptual decision-making task, on a subset of trials the Default Mode Network (DMN) was activated — brain regions commonly associated with mind-wandering (5) and being off task (6) — but accuracy, reaction time, and confidence were comparable to when the DMN was inactive, and regions linked to focused attention were engaged. This dissociation underscores the need to determine how trial-level variation in activity impacts decision-making and behavior.

Unraveling the relationship between neural activity and decision-making requires a system that provides complete access to the internal states. Convolutional neural networks (CNNs) are such a system, enabling direct investigation of internal states. Trained on objectives comparable to those of the brain, CNNs have been used to gain insight into object recognition (7, 8), representation (9) and processing in the visual cortex regions (10–12). Building on this foundation, we use a CNN to determine how decision-making changes with fluctuation in activity.

To anticipate our findings, first we show that a during a perceptual decision-making task, a CNN exhibits trial-level variability in neural activity reflecting observations in human neuroimaging data. Inspecting the internal state, we find that the trial-level activation patterns corresponded to decisions made when sensory evidence strongly favored one option. A second state reflected decisions under low discriminability, where sensory evidence provided little or no basis for distinguishing between options. A third state reflected bias-driven decisions when sensory evidence provided conflicting information. Collectively, these results suggest that trial-to-trial variation in activity reflects evidence-, uncertainty-, and bias-driven decision-making.

Results

Implementation of a perceptual decision-making task in a CNN

The brain exhibits marked fluctuations across trials, despite constant stimuli and task demands (**Fig. 1A**). However, it is not well understood how these fluctuations affect decision-making processes and behavior. To address this gap in knowledge, we trained a CNN to perform a perceptual decision-making task and examined trial-level neural activity. The CNN performed a modified version of a perceptual decision-making task which we have previously used to show that there are mutliple activation patterns for the same task (4). Specifically, a CNN was tasked to determine the dominant color in an image containing red and blue dots (**Fig. 1B**). The CNN was trained on 100,000 unique images balanced for red and blue dominance. To reflect previous findings that the visual system exhibits consistent activation during the task, the weights of the convolutional layers were fixed and only the weights in the fully connected layers were modifiable during training (4). The CNN was then tested on 1,000 new images balanced for red and blue dominance. In total, we trained and tested five different CNN models.

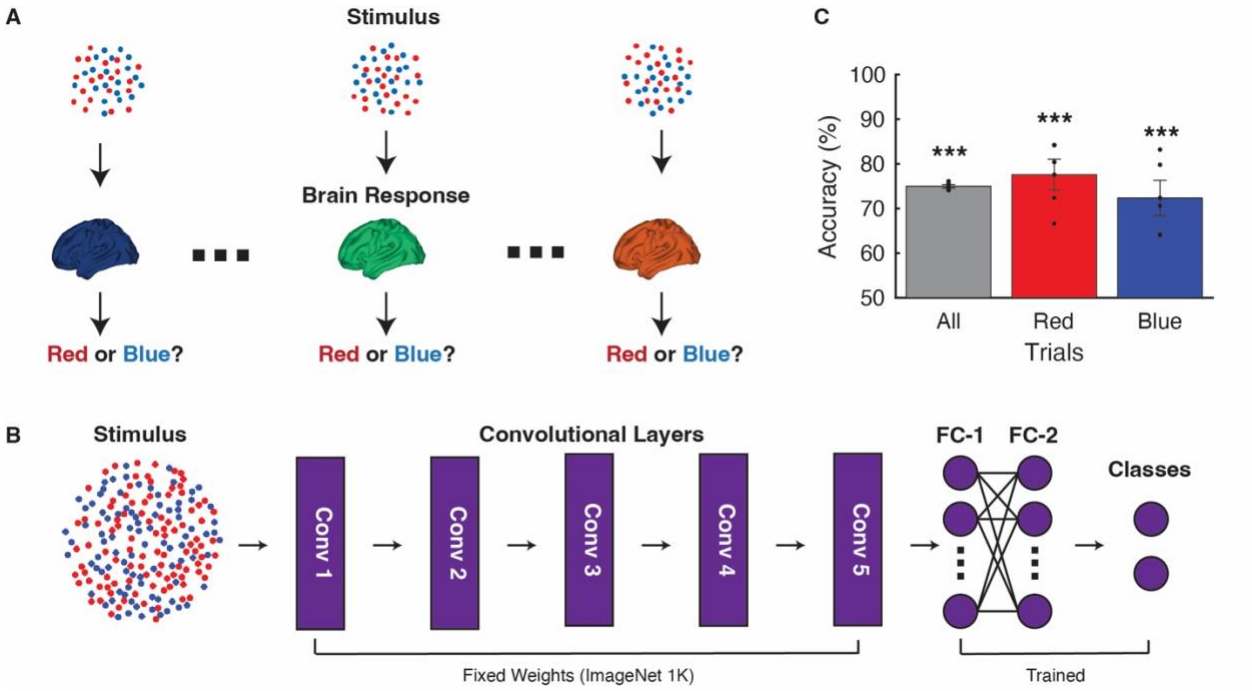

Figure 1. CNN implementation of a perceptual decision-making task. (A) Brain activity can vary in response to the same stimulus over time. (B) Schematic of CNN was trained to identify the dominant color in an image consisting of red and blue dots. During training only the weights of the fully connected layers were modifiable. Whereas the weights in the convolution component were fixed based on ImageNet. (C) Across five different models, CNN performed

better than chance (All = 74.96 ± 0.38 %; Red = 77.57 ± 3.48 %; Blue = 72.36 ± 3.95 %; $P < 0.001$) and no significant differences ($t(4) = 0.07$; $P = 0.52$) were observed in accuracy between red and blue dominant trials. ***$P < 0.001$.

Across models, CNNs performed significantly above chance in classifying the dominant color (All = 74.96 ± 0.38 %; Red = 77.57 ± 3.48 %; Blue = 72.36 ± 3.95 %; $P < 0.001$; **Fig. 1C**). Additionally, CNNs exhibited comparable accuracy in classifying red and blue dominant trials ($t(4) = 0.70$, $P = 0.52$), confirming that the CNN was able to perform the task.

<u>CNN contains multiple distinct but stable subtypes of trials during perceptual decision-making</u>

A characteristic of brain activity during a perceptual decision-making task are consistent trial-to-trial fluctuations (4). These fluctuations are driven by both top-down reflecting changes in attention (13) and by bottom-up processes driven by the stimulus (14). In CNNs, trial-level variability originates solely from bottom-up processing, providing a unique opportunity to isolate and determine how stimulus-driven fluctuations reflect decision-making processes.

For the analysis, we focus on the activity of neurons in the last fully connected layer (FC-2; **Fig. 1B**). The activity pattern in FC-2 encodes a transformed representation of the input image that is used by the classifier to determine the dominant color (i.e. red or blue). Therefore, trial-to-trial variation in FC-2 directly influences the decision output.

For the analysis, we used an established framework based on modularity-maximization to determine if a CNN exhibits trial-level variation similar to what is observed in humans (15). Briefly, the approach groups individual trials based on the similarity in neuronal activation estimated using Pearson correlation. To minimize noise, Pearson correlation was computed using only neurons that were activate on at least one trial.

Focusing on a representative CNN model (Model-4), modularity-maximization based clustering identified three groups of trials, or subtypes (**Fig. 2A**). Subtype 1 accounted for 29.6% of trials, Subtype 2 for 4.9%, and Subtype 3 for 65.5%. All three subtypes contained both red- and blue-dominant trials, but their composition differed (**Fig. 2B**). Subtypes 1 and 2 were composed

primarily of red-dominant trials (86.8% and 61.2%, respectively), whereas Subtype 3 contained primarily blue-dominant trials (67.5%). Trials associated with Subtype 2 were initially assigned to their own individual cluster because in these trials all neurons were inactive, therefore we grouped them to form Subtype 2.

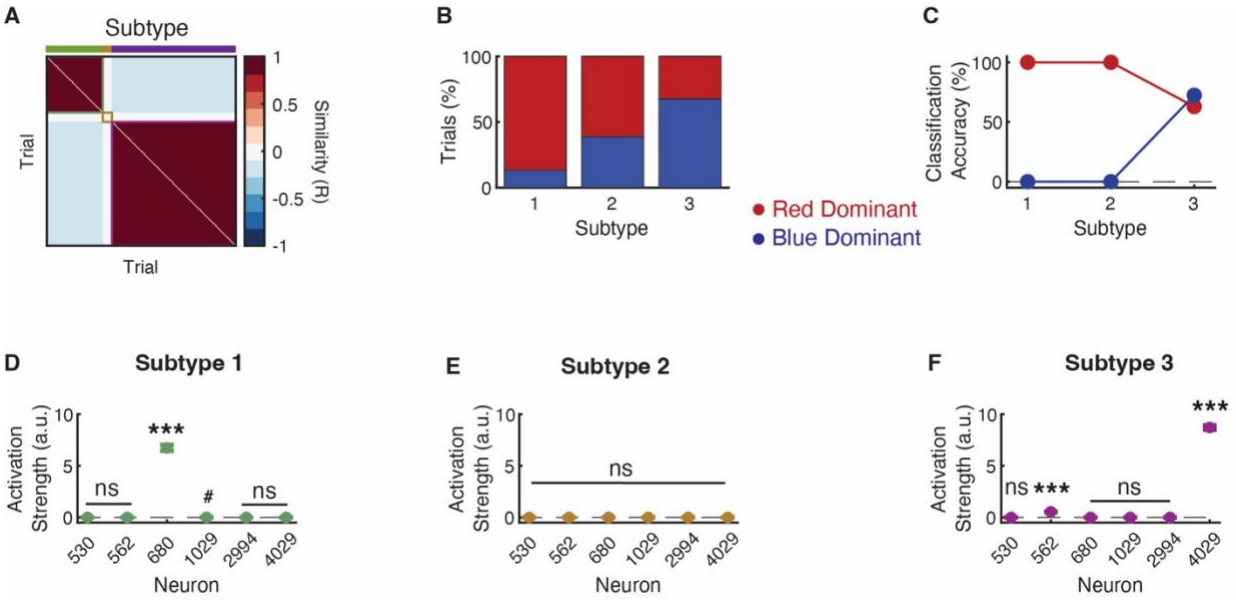

Figure 2. CNN exhibits structured variability during perceptual decision-making task. (A) Modularity-maximization-based clustering identified three subtypes of trials based on the activation pattern in layer FC-2. The colored squares correspond to the trials composing each subtype. Pearson correlation was used to calculate the similarity in activity among individual trials. (B) Composition of trials in each subtype. (C) Classification accuracy per subtype. (D) Activation strength of neurons in Subtype 1. Dot plot shows the average activation value of neurons that were active in at least one trials. Plotted activation value was estimated based only on trials within the subtype (mean ± sem). (E, F) Same as Panel C but for Subtype 2 and Subtype 3, respectively. Statistical significance was determined using one-samples *t*-test and Bonferroni corrected for multiple comparison. ***$P_{Bonferroni} < 0.001$; $^{\#}P < 0.05$; ns = not significant.

In Subtype 1 and Subtype 2, all red-dominant trials were correctly classified (**Fig. 2C**). Even in Subtype 3, 62.9% of red-dominant trials were correctly classified. On the other hand, blue-dominant trials were exclusively correctly classified in Subtype 3.

Turning to the activity patterns of the neurons driving the clustering, qualitatively distinct patterns of activity can be observed. In Subtype 1, activation was observed in Neuron-680 (Activation = 6.73 ± 0.28; t(295) = 23.84, $P_{Bonferroni} < 0.001$) and Neuron-1029 (Activation = 2.37

x 10^{-5} ± 2.37 x 10^{-5}; t(295) = 2.05, P = 0.04; **Fig. 2D**). By contrast, in Subtype 2, no measurable activation was observed in any neuron (**Fig. 2E**). In Subtype 3, two neurons exhibited an average activation that was significantly greater than zero – Neuron-562 (Activation = 0.55 ± 0.03; t(654) = 19.10, $P_{Bonferroni} < 0.001$) and Neuron-4029 (Activation = 8.7 ± 0.21; t(654) = 42.37, $P_{Bonferroni} < 0.001$; **Fig. 2F**). Trial-level activation are plotted in Figure S1. These findings confirm that a CNN can exhibit trial-level variation in activity similar to what is observed in humans.

<u>Trial-to-trial variability reflects decision-making under varying sensory evidence</u>

Importantly, the key challenge is to determine how fluctuations in activity shape decision-making. To address this unknown, we leverage having full access to the network's internal state to directly probe the signal underlying classification. This approach allows us to isolate the components driving decision-making in each subtype.

In a CNN, the decision-making process reflects the contribution of *sensory evidence* and *choice-bias.* Formally, classification is based on two signals, Z_{red} and Z_{blue}, estimated as:

$$Z_{red} = (A \cdot W_{red}) + C_{red}$$
$$Z_{blue} = (A \cdot W_{blue}) + C_{blue}$$

where A is the activation vector in FC-2, W are the corresponding weights and C is the constant term. Z_{red} and Z_{blue} are then converted to probabilities, $P(red)$ and $P(blue)$. We can estimate the decision preference (DP) for one option over another as:

$$DP = P(red) - P(blue)$$

with positive DP values indicating that the classifier outputs red and a negative DP favors blue. The advantage of decision preference is that it provides information on how strongly one option is favored over another. The strength of *sensory evidence* for red- or blue-dominance is represented in their respective $A \cdot W$ component and *choice-bias* by the constant term with the highest value.

Variation in neural activation between subtypes suggests that the subtype reflect decision-making under varying levels of *sensory evidence*. To test this, for each subtype, we estimated the decision preference using only the strength of the *sensory evidence*. This can be accomplished by omitting the constant terms C_{red} and C_{blue} when estimating Z_{red} and Z_{blue}.

First, we focus on red-dominant trials since these could be correctly classified in all three subtypes. As can be observed in Figure 3A, for Subtype 1, the constant terms, C_{red} and C_{blue}, were not critical for classification because the decision preference favored a classification of red dominance with (DP = 0.42 ± 0.006; t(256) = 66.00; $P_{Bonferroni} < 0.001$) and without the constant terms (DP = 0.18 ± 0.007; t(256) = 24.2, $P_{Bonferroni} < 0.001$; **Fig. 3A**). These results indicate that in Subtype 1 the sensory evidence provided sufficient information favoring one option – red dominant.

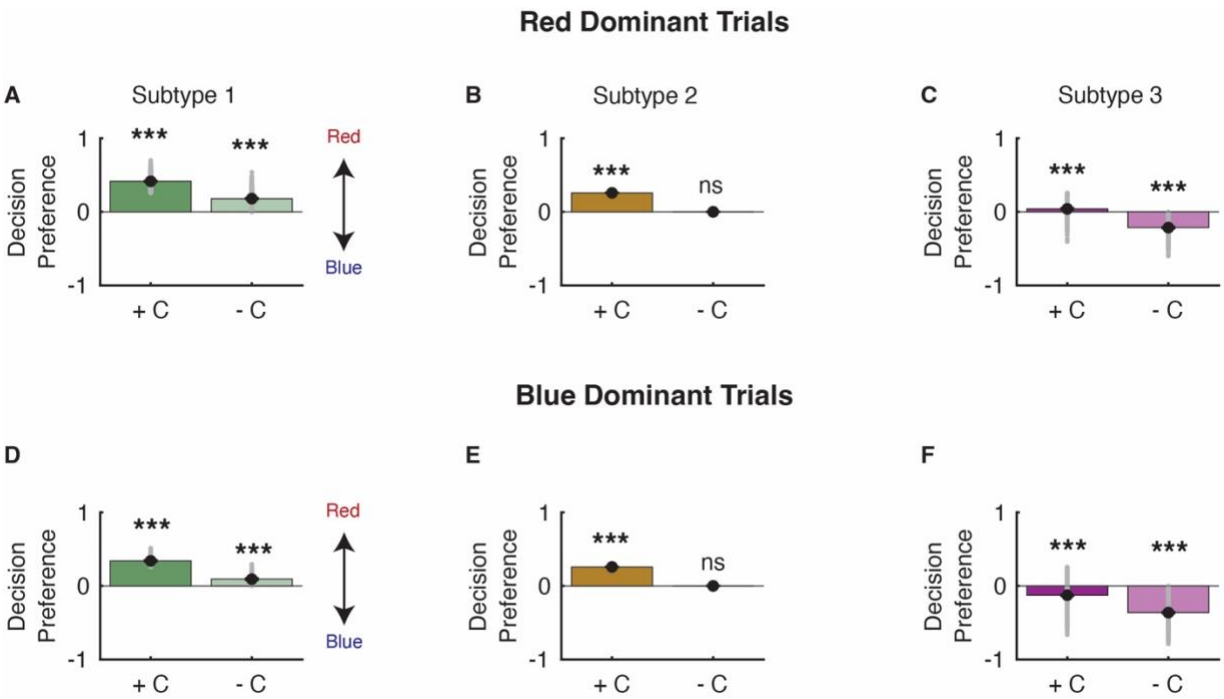

Figure 3. **Trial-level variability reflects decision-making under varying sensory evidence.** (A-C) Decision preference estimated with (+) and without (-) constant terms, C, for red dominant trials in Subtype 1, Subtype 2, and Subtype 3, respectively. Bar plot shows the average decision preference value across trials (Mean ± SEM). Gray dots show individual trials. (D–F) Same as panels A–C, but for blue-dominant trials. Statistical significance was determined using one-samples *t*-test and Bonferroni corrected for multiple comparison. ***$P_{Bonferroni} < 0.001$; ns = not significant.

For Subtype 2, the constant term was critical in classifying trials as red dominant (DP = 0.25 ± 0.0; $P_{Bonferroni} < 0.001$) because without the constant term the decision preference dropped to zero as expected given no neurons were active (DP = 0.0; t(29) = 0.00, P = 1.00; **Fig. 3B**). Similarly, in Subtype 3, inclusion of the constant terms, C, resulted in a subset of trials favoring a red classification (DP = 0.04 ± 0.01; t(212) = 3.9, $P_{Bonferroni} < 0.001$), but without the constant terms then the decision preference favored classifying the trials as blue-dominant (DP = -0.21 ± 0.01; t(212) = -22.43, $P_{Bonferroni} < 0.001$), indicating classification sensory evidence provided conflicting information (**Fig. 3C**).

Applying the same analysis to blue-dominant trials, in Subtype 1, the strength of the sensory evidence alone was not sufficient to correctly classify the trials (DP = 34.02 ± 0.01; t(38) = 32.05, $P_{Bonferroni} < 0.001$), and this was compounded when including the constant terms (DP = 0.09 ± 0.01; t(38) = 7.58, $P_{Bonferroni} < 0.001$; **Fig. 3D**). Similar findings were observed in Subtype 2 (**Fig. 3E**). Moreover, in Subtype 3 inclusion of the constant term resulted in a subset of trials favoring a red classification (DP = -0.13 ± 0.01; t(441) = -13.66, $P_{Bonferroni} < 0.001$) compared to when the classification decision is based only on the strength of *sensory evidence* (DP = -0.36 ± 0.01; t(441) = 45.00, $P_{Bonferroni} < 0.001$; **Fig. 3F**).

Additionally, we combined all the trials together to examine where the trials in each subtype reside along the evidence boundary. As can be observed in **Figure 4A**, subtypes were separated by the evidence boundary, linking trial-level variability to classification accuracy (**Fig. 4B**). Together, the results suggest that choice-bias is a critical factor in decision-making particularly when sensory evidence is weak or provides conflicting information.

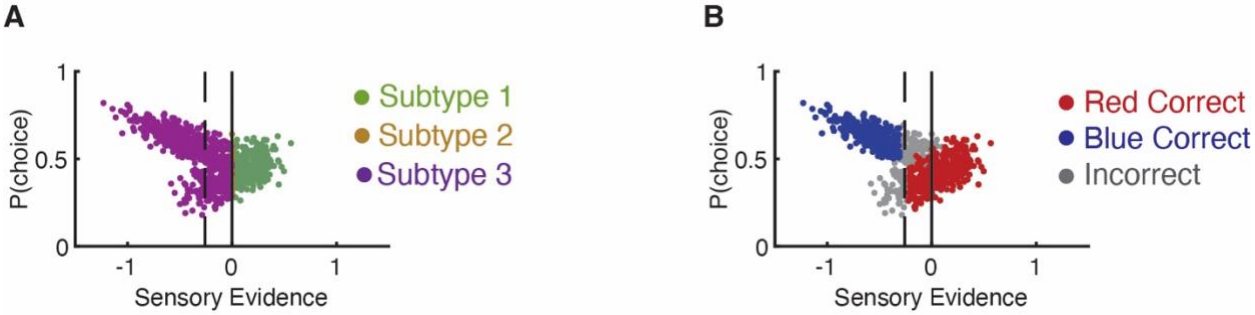

Figure 4. Evidence boundary separates subtypes. (A) Trials in each subtype residing along the evidence boundary. Individual dots correspond to the trials. Dark black line corresponds to the evidence boundary. Dashed line corresponds the shift reflecting the constant term, C_{blue}. (B) Same as panel A but colored based on the accuracy of the classification. Red and blue dots reflect trials that were correctly classified as red or blue-dominant, respectively. Gray dots reflect trials that were incorrectly classified.

Corroborating the findings in additional CNN models

Lastly, we corroborated these findings in four additional CNN models. As can be observed in **Fig. S2-S5**, in all models, trial-level variation in activity could be clustered into distinct subtypes. Moreover, *choice-bias* was a critical component in decision-making. Collectively, these results reveal that trial-level variation in activity reflects evidence-, uncertainty-, and bias-driven decision-making.

Discussion

Fluctuations in brain activity are ubiquitous during both simple and cognitively demanding tasks and are often indicative of variation in cognitive processing (16, 17). However, how such fluctuations relate to decision-making and behavior is not well understood. Here, we address this gap in knowledge using a convolutional neural network (CNN) trained to perform a perceptual decision-making task. Leveraging an established analytical framework, we show that trial-to-trial variability in CNN activity organizes into consistent patterns of activation, mirroring variability observed in human brain activity during the same task. Examination of the network's internal state revealed that trial-level variability reflects making decisions under varying level sensory evidence: (i) decisions dominated by strong sensory evidence, (ii) decisions under low discriminability, and (iii) decisions in which bias overrode sensory evidence.

Trial-to-trial variability in brain activity is driven by both top-down (13) and bottom-up processes (14). However, it is challenging to determine how they individually shape decision-making (18). In CNNs, trial-level variability is driven only by bottom-up processing, providing an opportunity to assess the contribution to trial-level variation in activity. We observed that bottom-up processing was sufficient to result in discrete patterns of activation suggesting that

bottom-up processes may be the primary driver of multiple activation patterns in a perceptual decision-making task (4).

Interestingly, the CNN could correctly classify trials even when sensory evidence was weak or provided conflicting information. The analysis revealed that the CNN was able to accommodate varying level of sensory evidence and uncertainty by shifting the decision boundary using choice bias (i.e. constant terms). How choice bias arises in artificial neural networks and the brain is open question, but prior work suggests that bias may reflect competition between different components of the network (19).

In humans, choice-bias has been shown to play a role in decision-making (20). Modeling using signal detection theory suggests that choice-bias reflects a shift in the criterion (21). In perceptual decision-making tasks, choice bias is thought to arise from experimental factors and measures are taken to minimize the effect (22). However, choice bias can be beneficial in cases where one alternative is preferred over another (23). In humans, choice-bias may play a similar role as in the CNN and may not simply be a confounding factor (24, 25).

Take together, our findings outline a unifying framework on the relationship between fluctuation in activity and decision-making: both biological and artificial systems organize trial-to-trial variability into a small set of discrete states. Understanding how systems make decisions as evidence waxes and wanes could provide a foundation for studying disrupted decision-making where abnormal fluctuations may reflect impairment. Overall, variation in activity over time is not a byproduct of noise but a feature of adaptive decision-making, enabling flexible responses under varying levels of evidence and uncertainty.

Methods

<u>Task description</u>

The CNN was trained to perform a perceptual decision-making task, identifying whether red or blue dots were more frequent in a cloud of randomly distributed dots. Each stimulus consisted of either 90 red- or 100 blue-colored dots, presented within an imaginary circular aperture. Dot positions were randomized for every trial to minimize low-level visual regularities. In total,

100,000 unique stimuli were generated for training, and an additional 1,000 novel stimuli were created for testing with equal number of red- or blue-dominant trials. All images were generated using Python.

CNN model, training and testing

We implemented a convolutional neural network (CNN) based on AlexNet architecture, modified for binary classification (red- versus blue-dominant images). CNN architecture is shown in Table S1. The weights associated with the convolutional layers were fixed corresponding to pretrained on ImageNet (1K classes), as provided by the PyTorch library (https://pytorch.org/). Only the fully connected layers were trainable, allowing the network to adapt the decision stage while maintaining a stable feature-extraction stage.

The CNN was trained on 100,000 randomly generated trials of the perceptual decision-making task. Training was conducted in five batches of randomly chosen trials. Model optimization was performed using stochastic gradient descent as implemented in PyTorch. After training, the CNN was evaluated on 1,000 unique trials not included in the training set. Accuracy was assessed separately for red-dominant and blue-dominant images, as well as across all trials. In total five different CNN models were trained and tests. All training and testing were conducted using Google Colab on an NVIDIA A100 GPU.

Activation of fully connected layer and modularity-maximization based clustering

For each trial in the test set, we extracted the activation patterns from the second fully connected layer (FC-2). These activation patterns formed the basis for subsequent analyses of trial-to-trial variability. To identify structure in variability, we applied clustering to FC-2 activations, grouping trials with similar patterns of activity into discrete subtypes. A trial-by-trial similarity matrix was created using the Pearson Product Correlation based on the activation. To minimize noise, the Pearson correlation was computed using only neurons that were activate on at least one trial. Clustering of the similarity matrix was conducted using modularity-maximization (26). Modularity-maximization does not require the number of clusters to be specified, and the resolution of the clusters was controlled with resolution parameter, $\gamma = 1$. Modularity-

maximization was implemented with the Generalized Louvain algorithm part of the Brain Connectivity Toolbox (27).

The community detection is not deterministic, and the results can depend on the specific random seeds. This limitation can be overcome using consensus clustering to identify stable clusters out of a set of partitions (28). Moreover, in our previous work of clustering single trials, we had found that 100 iterations were sufficient to identify stable clusters (29). Specifically, consensus clustering identifies a single representative partition from the set of 100 iterations. This process involves the creation of a thresholded nodal association matrix which describes how often nodes were placed in the same cluster. The representative partition is then obtained by using a Generalized Louvain algorithm to identify clusters within the thresholded nodal association matrix. We have previously used this framework to identify trial subtypes in brain activity measured with fMRI (4) and electroencephalography in perceptual decision-making tasks (30) and working memory task (31).

Code and Data Availability

The analysis was based on a combination of publicly available toolboxes, analysis specific scripts. Specifically, the CNNs were based on AlexNet architecture from the PyTorch library (https://pytorch.org/). Formal analysis was conducted in MATLAB 2021b. Trained CNN models, scripts for training and formal analysis are available at https://osf.io/efnq6/. Clustering analysis was based on the Community Louvain part of the Brain Connectivity Toolbox (https://sites.google.com/site/bctnet/). Consensus clustering was determined using *consensus_iterative.m* (http://commdetect.weebly.com/).

References

1. R. L. T. Goris, J. A. Movshon, E. P. Simoncelli, Partitioning neuronal variability. *Nature Neuroscience* **17**, 858–865 (2014).

2. A. Arieli, A. Sterkin, A. Grinvald, A. Aertsen, Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science* **273**, 1868–1871 (1996).

3. L. Waschke, N. A. Kloosterman, J. Obleser, D. D. Garrett, Behavior needs neural variability. *Neuron* **109**, 751–766 (2021).

4. J. Nakuci, *et al.*, Multiple brain activation patterns for the same perceptual decision-making task. *Nature Communications* **16**, 1785 (2025).

5. K. Christoff, Z. C. Irving, K. C. R. Fox, R. N. Spreng, J. R. Andrews-Hanna, Mind-wandering as spontaneous thought: a dynamic framework. *Nature Reviews Neuroscience* **17**, 718–731 (2016).

6. M. F. Mason, *et al.*, Wandering Minds: The Default Network and Stimulus-Independent Thought. *Science* **315**, 393–395 (2007).

7. R. Rajalingham, *et al.*, Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *Journal of Neuroscience* **38**, 7255–7269 (2018).

8. D. L. Yamins, J. J. DiCarlo, Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience* **19**, 356–365 (2016).

9. N. Kriegeskorte, R. A. Kievit, Representational geometry: integrating cognition, computation, and the brain. *Trends in Cognitive Sciences* **17**, 401–412 (2013).

10. R. M. Cichy, A. Khosla, D. Pantazis, A. Torralba, A. Oliva, Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific reports* **6**, 27755 (2016).

11. M. Eickenberg, A. Gramfort, G. Varoquaux, B. Thirion, Seeing it all: Convolutional network layers map the function of the human visual system. *NeuroImage* **152**, 184–194 (2017).

12. S.-M. Khaligh-Razavi, N. Kriegeskorte, Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS computational biology* **10**, e1003915 (2014).

13. N. Al-Aidroos, C. P. Said, N. B. Turk-Browne, Top-down attention switches coupling between low-level and high-level areas of human visual cortex. *Proceedings of the National Academy of Sciences* **109**, 14675–14680 (2012).

14. M. Corbetta, G. L. Shulman, Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience* **3**, 201–215 (2002).

15. J. Nakuci, D. Rahnev, A Practical Guide to Identifying Robust Clusters in Neuroimaging Data. *Human Brain Mapping* **46**, e70330 (2025).

16. A. A. Faisal, L. P. J. Selen, D. M. Wolpert, Noise in the nervous system. *Nature Reviews Neuroscience* **9**, 292–303 (2008).

17. D. D. Garrett, *et al.*, Moment-to-moment brain signal variability: A next frontier in human brain mapping? *Neuroscience & Biobehavioral Reviews* **37**, 610–624 (2013).

18. E. Awh, A. V. Belopolsky, J. Theeuwes, Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in cognitive sciences* **16**, 437–443 (2012).

19. L. Lebovich, R. Darshan, Y. Lavi, D. Hansel, Y. Loewenstein, Idiosyncratic choice bias naturally emerges from intrinsic stochasticity in neuronal dynamics. *Nature human behaviour* **3**, 1190–1202 (2019).

20. J. I. Gold, L. Ding, How mechanisms of perceptual decision-making affect the psychometric function. *Progress in Neurobiology* **103**, 98–114 (2013).

21. D. Linares, D. Aguilar-Lleyda, J. López-Moliner, Decoupling sensory from decisional choice biases in perceptual decision making. *eLife* **8**, e43994 (2019).

22. S. A. Klein, Measuring, estimating, and understanding the psychometric function: A commentary. *Perception & psychophysics* **63**, 1421–1455 (2001).

23. M. J. Mulder, E.-J. Wagenmakers, R. Ratcliff, W. Boekel, B. U. Forstmann, Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience* **32**, 2335–2343 (2012).

24. R. Akaishi, K. Umeda, A. Nagase, K. Sakai, Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron* **81**, 195–206 (2014).

25. A. E. Urai, A. Braun, T. H. Donner, Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nature communications* **8**, 14637 (2017).

26. P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, J. P. Onnela, Community structure in time-dependent, multiscale, and multiplex networks. *Science* **328**, 876–878 (2010).

27. M. Rubinov, O. Sporns, Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage* **52**, 1059–1069 (2010).

28. A. Lancichinetti, S. Fortunato, Consensus clustering in complex networks. *Sci. Rep.* **2** (2012).

29. D. S. Bassett, *et al.*, Robust detection of dynamic community structure in networks. *Chaos* **23** (2013).

30. J. Nakuci, J. Samaha, D. Rahnev, Brain signatures indexing variation in internal processing during perceptual decision-making. *iScience* **26**, 107750 (2023).

31. J. Nakuci, T. J. Covey, J. L. Shucard, D. W. Shucard, S. F. Muldoon, Single trial variability in neural activity during a working memory task reveals multiple distinct information processing sequences. *NeuroImage* **269**, 119895 (2023).