AFRL-AFOSR-VA-TR-2020-0065

Characterizing Neural Code from a Minimum-Description-Length Perspective

Ning Qian
TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK

12/28/2020
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
17-06-2020

2. REPORT TYPE
Final Performance

3. DATES COVERED *(From - To)*
30 Sep 2015 to 29 Sep 2019

4. TITLE AND SUBTITLE
Characterizing Neural Code from a Minimum-Description-LengtAFRL-AFOSR-VA-TR-2020-0064h Perspective

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-15-1-0439

5c. PROGRAM ELEMENT NUMBER
61102F

6. AUTHOR(S)
Ning Qian

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
TRUSTEES OF COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK
630 W 168TH ST FL 4
NEW YORK, NY 100323725 US

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AF Office of Scientific Research
875 N. Randolph St. Room 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/AFOSR RTA2

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-AFOSR-VA-TR-2020-0065

12. DISTRIBUTION/AVAILABILITY STATEMENT
A DISTRIBUTION UNLIMITED: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
We have successfully completed the projects in our original proposal and also performed some additional, closely related studies. (1) We have proposed a new theoretical framework for understanding sensory encoding based on the modern minimum-description-length principle. A key finding is our new hypothesis that firing rates of sensory projection neurons are proportional to optimal code length for stimulus features (i.e., negative log estimated probability of stimulus features). This is in sharp contrast to the traditional view that sensory neurons' firing rates represent the probability of the neurons' preferred stimulus in the input. (2) We have proposed a new theoretical framework for understanding visual decoding. Experimental studies have firmly established that encoding is hierarchical, progressing from lower-level representations of simpler and less invariant features to higher-level representations of more complex and invariant features along visual pathways

15. SUBJECT TERMS
Neural Code, Computational Cognition

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			LAWTON, JAMES
Unclassified	Unclassified	Unclassified	UU		19b. TELEPHONE NUMBER *(Include area code)* 703-696-5999

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

DISTRIBUTION A: Distribution approved for public release

Final Report for AFOSR FA9550-15-1-0439: Characterizing Neural Code from a Minimum-Description-Length Perspective

Ning Qian and Jun Zhang

December 25, 2019

Abstract

We have successfully completed the projects in our original proposal and also performed some additional, closely related studies. (1) We have proposed a new theoretical framework for understanding sensory encoding based on the modern minimum-description-length principle. A key finding is our new hypothesis that firing rates of sensory projection neurons are proportional to optimal code length for stimulus features (i.e., negative log estimated probability of stimulus features). This is in sharp contrast to the traditional view that sensory neurons' firing rates represent the probability of the neurons' preferred stimulus in the input. (2) We have proposed a new theoretical framework for understanding visual decoding. Experimental studies have firmly established that encoding is hierarchical, progressing from lower-level representations of simpler and less invariant features to higher-level representations of more complex and invariant features along visual pathways. Much less is known about decoding but most models assume that decoding follows the same low-to-high-level hierarchy of encoding. We found unequivocal evidence against this common assumption, and proposed a new theory that views visual perception as retrospective Bayesian decoding from high-to-low-level features in working memory. (3) We have tested and confirmed a prediction of our decoding theory that repulsion between two orientations in a discrimination task should increase with the duration between the stimulus disappearance and the perceptual report. (4) We have tested and confirmed another prediction of our decoding theory that two orientations should have cross-fixation interactions if they are both task relevant and thus have to be maintained in working memory after their disappearance. (5) We have demonstrated how characteristics of sensory decoding help explain some puzzling results in motor control. For example, by incorporating Weber's law of sensory perception into a stochastic optimal control model, we can explain the observation that initial and terminal visual feedback are, respectively, less and more effective in improving endpoint accuracy. (6) A main purpose of sensory coding is to guide motor response. We investigated visuomotor transformation using the stimulus-response-compatibility (SRC) effect. We found that the sites of both overt and covert visual attention define simultaneous spatial reference centers that influence motor reaction time.

Neuronal Firing Rate as Code Length: A Hypothesis

Many theories assume that a sensory neuron's higher firing rate indicates a greater probability of its preferred stimulus. However, this contradicts 1) the adaptation phenomena where prolonged exposure to, and thus increased probability of, a stimulus reduces the firing rates of cells tuned to the stimulus; and 2) the observation that unexpected (low probability) stimuli capture attention and increase neuronal firing. Other theories posit that the brain builds predictive/efficient codes for reconstructing sensory inputs. However, they cannot explain that the brain preserves some information while discarding other. We propose that in sensory areas, projection neurons' firing rates are proportional to optimal code length (i.e., negative log estimated probability), and their spike patterns are the code, for useful features in inputs. This hypothesis explains adaptation-induced changes of V1 orientation tuning curves, and bottom-up attention. We discuss how the modern minimum-description-length (MDL) principle may help understand neural codes. Because regularity extraction is relative to a model class (defined by cells) via its optimal universal code (OUC), MDL matches the brain's purposeful, hierarchical processing without input reconstruction. Such processing enables input compression/understanding even when model classes do not contain true models. Top-down attention modifies lower-level OUCs via feedback connections to enhance transmission of behaviorally relevant information. Although OUCs concern lossless data compression, we suggest possible extensions to lossy, prefix-free neural codes for prompt, online processing of most important aspects of stimuli while minimizing behaviorally relevant distortion. Finally, we discuss how neural networks might learn MDL's normalized maximum likelihood (NML) distributions from input data.

Visual perception as retrospective Bayesian decoding from high to low-level features in working memory

When a stimulus is presented, its encoding is known to progress from low- to high-level features. How these features are decoded to produce perception is less clear and most models assume that decoding follows the same low-to-high-level hierarchy of encoding. There are also theories arguing for global precedence, reversed hierarchy, or bidirectional processing but they are descriptive without quantitative comparison with human perception. Moreover, observers often inspect different parts of a scene sequentially to form overall perception, suggesting that perceptual decoding requires working memory; yet few models consider how working-memory properties may affect decoding hierarchy. We probed decoding hierarchy by comparing absolute judgments of single orientations and relative/ordinal judgments between two sequentially presented orientations. We found that lower-level, absolute judgments failed to account for higher-level, relative/ordinal judgments. However, when ordinal judgment was used to retrospectively decode memory representations of absolute orientations, striking aspects of absolute judgments, including the correlation and forward/backward aftereffects between two reported orientations in a trial, were explained. We propose that the brain prioritizes decoding of higher-level features because they are more behaviorally relevant, and more invariant and categorical and thus easier to specify and maintain in noisy working memory, and that more-reliable higher-level decoding constrains less-reliable lower-level decoding.

Orientation repulsion increases with noise in working memory

We (PNAS, 2017) proposed that perceptual decoding proceeds from high- to low-level features in working memory, and that more reliable/useful higher-level decoding constrains less reliable/useful lower-level decoding. This retrospective Bayesian theory of perception predicts that in orientation discrimination tasks (where absolute orientations of individual lines and their relative relationship are lower- and higher-level features, respectively), repulsion between two orientations should increase with the noise in their working-memory representations. Here we tested this prediction. Subjects viewed two successively presented lines (50º and 55º from horizontal in random order) while maintaining fixation (enforced with an eye tracker). After a delay of either 0.5 or 5 sec, they rotated marker dots to report the lines' absolute orientations and their ordinal relationship. We also had various control conditions including baseline conditions in which a single line was presented per trial. The variances of the reported orientations under the long-delay condition were greater than those under the short-delay condition, indicating that longer delays led to greater memory noise. Consistent with the theory, the long-delay condition also produced larger repulsion, compared with the short-delay condition. Our results suggest that orientation repulsion indeed increases with working-memory noise, lending further support to the retrospective Bayesian decoding theory for visual perception.

Cross-fixation interactions of orientations suggest that orientation decoding occurs in a high-level area of visual working memory

According to the common, low-to-high-level decoding assumption, orientation decoding must occur in low-level areas such as V1, and consequently, two orientations on opposite sides of the fixation should not interact with each other perceptually. However, we (PNAS, 2017) provided evidence against the assumption and proposed that visual decoding may follow the opposite, high-to-low-level hierarchy in working memory. If two orientations on opposite sides of the fixation are both task relevant and enter a high-level working-memory area in a delay period, then they should interact with each other. We tested this prediction.

Subjects maintained central fixation when two lines with an orientation difference of 5° were flashed on opposite sides of the fixation, with a center-to-center distance of 16° of visual angle. Their eye positions were monitored with an infrared eye tracker and trials with broken fixation were aborted. After a delay, subjects reported the two orientations by drawing and adjusting an indicator line at the fixation. In one condition, the indicator line disappeared after the first report, and was redrawn for the second report, to minimize potential interference. We found that the two lines showed a large and significant repulsion between them, demonstrating the predicted cross-fixation interactions in working memory. The pattern was consistent across 15 subjects. Control conditions and analyses ruled out alternative explanations such as interactions across trials on the same side of the fixation. Moreover, we quantitatively accounted for the repulsion with our retrospective Bayesian decoding model. We conclude that our results support the theory that visual perception may be viewed as high-to-low-level decoding in working memory.

A Single, Continuous Control Policy for Modeling Reaching Movements with and without Target Perturbation

We investigated how characteristics of sensory decoding help explain some puzzling results in motor control. Specifically, there have been debates in motor-control literature on the relative importance of sensory feedbacks at difference stages of movements and on how to understand behavioral data from target-perturbation experiments in which the target jumps to a new location when the hand is reaching for it. We constructed an infinite-horizon optimal feedback control model that unlike previous approaches, does not preset movement durations or use multiple control policies. It contains both control-dependent and independent noises in system dynamics, state-dependent and independent noises in sensory feedbacks, and different delays and noise levels for visual and proprioceptive feedbacks. Importantly, the noise terms and sensory representations in our model implement the well-known Weber's law of perceptual decoding. We analytically derived an optimal solution which can be applied continuously to move an effector toward a target regardless of whether, when, or where the target jumps. In agreement with experimental data, this single policy produces different numbers of peaks and ``sub-movements" in velocity profiles for different trials and conditions. Movements that are slower or perturbed later appear to have more sub-movements. The model is also consistent with the observation that subjects can perform the perturbation task even without seeing the target jump or their hands during reaching. Finally, because the model incorporates Weber's law, it explains that initial and terminal visual feedback are, respectively, less and more effective in improving endpoint accuracy. Our work suggests that the number of peaks or sub-movements in a velocity profile does not necessarily reflect the number of motor impulses, and that the difference between initial and terminal feedback does not necessarily imply a transition between open- and closed-loop strategies.

Sites of overt and covert attention define simultaneous spatial reference centers for visuomotor response

A main purpose of sensory coding is to guide motor response. We investigated visuomotor transformation using the stimulus-response-compatibility (SRC) effect: When subjects press, e.g., a left key to report stimuli, their reaction time is shorter when stimuli appear to the left than to the right of the fixation. Fixation point is the site of overt attention and the SRC effect indicates that overt attention defines a spatial reference center that affects visuomotor response. Covert attention to a peripheral site also appears to define a similar reference center but previous studies did not control for confounding spatiotemporal factors or investigate the relationship between overt- and covert-attention-defined centers. Using an eye tracker to monitor fixation, we found an SRC effect relative to the site of covert attention induced by a flashed cue dot, and a concurrent reduction, but not elimination, of the overt-attention SRC effect. The two SRC effects jointly determined the overall motor reaction time. Since trials with different cue locations were randomly interleaved, the integration of the two reference centers must be updated online. When the cue was invalid and diminished covert attention, the covert-attention SRC effect disappeared and the overt-attention SRC effect retained full strength, excluding non-attention-based interpretations. We conclude that both covert- and overt-attention sites define visual reference centers that simultaneously contribute to motor response.