Neuron

CellPress

Previews

Plaudits for logits in sensory neuroscience

Joshua J.H. Ryu[1] and Justin L. Gardner[1,*]
[1]Department of Psychology and Wu Tsai Neurosciences Institute, Stanford University, Stanford, CA 94305, USA
*Correspondence: jlg@stanford.edu
https://doi.org/10.1016/j.neuron.2024.08.008

A workhorse tool of economic decision-making has long sought to get inside people's heads through careful examination of their choices. In this issue of *Neuron*, Carandini[1] flips the script, showing how it can model how the brain makes sensory choices.

Before the Bay Area Rapid Transport system was built, economists were able to make a prediction of ridership that turned out seemingly prescient in its accuracy.[2] Making such a prediction required knowing the idiosyncratic, internal preferences people have for different factors, such as cost or travel time, that might guide transit choices. But how should these preferences be estimated, given what people say they prefer might differ from what they reveal through their actual choices? The solution to the problem was to infer preferences from the discrete choices people made for one form of transport or another using multinomial logistic regression, work that helped earn a Nobel for economist Daniel McFadden.

What, if anything, does this have to do with sensory neuroscience? According to a Viewpoint from Matteo Carandini in this issue of *Neuron*,[1] quite a lot.

Logistic regression is now a common way to estimate how people weigh different factors in making decisions (Figure 1, left). Each factor that might influence a discrete decision is multiplied by a weight, then summed together with a bias. The resulting value is turned into a probability of producing a choice by passing it through the logistic function. The goal of logistic regression is to find the unobserved weights and bias that maximize the probability of producing the observed choice data.

By applying it to sensory choices, Carandini shows that logistic regression provides excellent fits to choice data from classic perceptual decision-making tasks such as a monkey trained to choose whether random-dot motion is moving to the left or right.[3] In this case, the percentage of coherently moving dots is a sensory factor that is weighed by the monkey. This weighted value is added to the bias, a value that can account for systematically choosing the more common choice more often. Passed through the logistic function, this results in a sigmoidal (that is, s-shaped) curve relating motion coherence to the probability of making rightward choices in the task. This function is familiar to sensory neuroscientists as a psychometric function, which has traditionally been fit with similar sigmoidal functions such as a Weibull or a cumulative Gaussian.

While an economist uses logistic regression to get inside people's heads, a neuroscientist uses it to get out. That is, the economist uses logistic regression to infer the internal, unobserved preferences a decision-maker has from their externally observed choices. The neuroscientist measures the internal brain signals of a decision-maker and tries to infer how these contribute to the externally observable choice.

Carandini suggests that neuroscientists call the process logistic classification (Figure 1, right). From the internal viewpoint of the decision-maker, the goal is to classify sensory or other evidence into one or more different possibilities; the weights and bias are known and available for them to choose. The result of summing the weighted evidence and bias is a decision variable from which a choice can be made. Stochasticity from noisy sensory evidence, the decision process, or both is implied by the logistic function turning the decision variable into a probability of choice. We will come back to this point, as Carandini's perspective as a sensory neuroscientist on this stochasticity leads to an intriguing solution to a long-standing conundrum of economic choice.

What is the value of logistic classification for sensory neuroscience? Practically, it is easy to fit, and Carandini shows excellent fits to a number of datasets that examine cue combination, choice history bias, and the effects of inactivation experiments on sensory choice. More fundamentally, it is a marked shift away from traditional thinking based on ideal observer models.

Ideal observer models are foundational in the study of perception.[4] Consider a classic cue combination task that Carandini fits: inferring the location of an event from visual and auditory cues. An ideal observer approach determines the optimal solution given constraints such as the reliability of each sensory modality. Bayesian analysis provides the optimal solution by reading out a posterior distribution constructed by multiplying likelihood of location given the auditory and visual cues with the prior. This approach provides a benchmark to compare performance to and answers the question: given sensory constraints, what *should* the nervous system do?

Logistic classification is cut from the same cloth as ideal observer models; it's all based on logits of evidence. For the uninitiated, logit is shorthand for "logistic unit" and represents the log odds. If p is the probability of classifying as one category, the odds are $p/(1 - p)$, which has the familiar meaning of when one says a horse has 3:1 odds of winning. Taking the natural log of the odds makes for logits, which can then be combined using simple summation rather than multiplicative combination of independent probabilities. Logits are what accumulate in drift diffusion decision models, a framework that Carandini shows to be mathematically equivalent.

The decision variable of logistic classification is in units of logits, giving it special

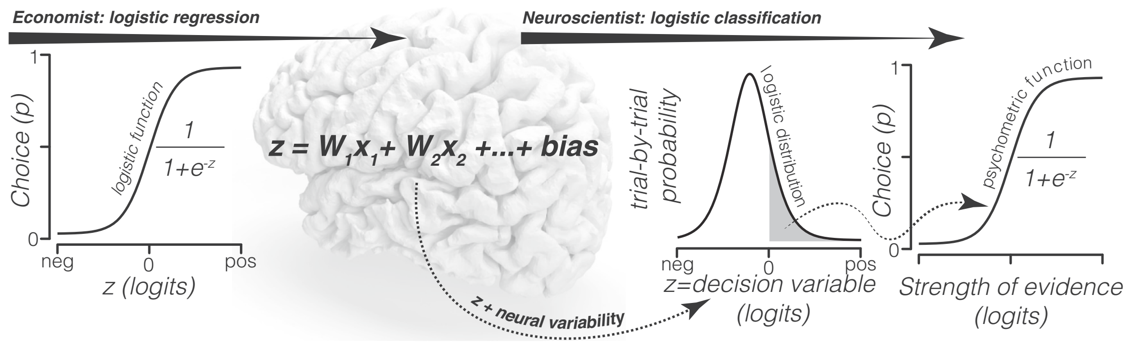

Figure 1. Economists and neuroscientists view the same set of equations from different vantage points
Economists use logistic regression to estimate the weights (W) for different factors (x) and bias that best predict choice data. Neuroscientists use logistic classification to model a decision-maker who weighs sensory and other evidence, combining with bias to form a decision variable. Neural variability results in stochastic choices in the psychometric function.

meaning; it is the decision-maker's estimate, given all the evidence, of the log odds of classification. A logit of 1 corresponds to e (approximately 2.7) odds in favor, −1 to 2.7 against, and 0 to no evidence either way. If you prefer to know the probability of the category rather than the logits, pass the number through the inverse of log odds, which is … the logistic function! Thus, the deeper meaning for using the logistic function: it converts the decision-maker's estimate of logits of evidence into probability of choice.

Logistic classification is compatible with ideal observer models, with a bit of mathematical wrangling. In particular, each weight should convert its corresponding factor into logits of evidence. In cue combination, one might increase the logits of evidence by increasing loudness or contrast of the stimuli. However, the relationship between stimulus strength and logits of evidence cannot be captured by a linear weight. According to Fechner, perceptual strength is logarithmically related to stimulus strength; to Stevens, it is a power law; and for the normalization model, a hyperbolic ratio.[5] To accommodate these relationships and the effect of neural noise,[6] Carandini's formulation allows for a non-linear mapping of stimulus strength before linear weighting. It is worth noting that potentially a lot is packed into that non-linear mapping for it to be optimal, making it non-trivial for the brain to compute exactly.

But how can stochasticity of choice in logistic classification be optimal? A rational decision-maker should always go with the choice with more logits of evidence instead of sporadically choosing the other option. While solely attributing the stochasticity to the decision process might be rationalized as exploration, not exploitation,[7] the insight of signal detection theorists[8] provides a different answer: when sensory evidence is noisy, incorrect choices occur due to random fluctuations of evidence rather than random guessing. For such choices to be fit by a logistic function, the variability of the induced decision variable should be distributed logistically (Figure 1, right). This assumption is not unreasonable given that the logistic distribution is practically indistinguishable from a Gaussian distribution, albeit with slightly heavier tails.

Armed with this insight from sensory neuroscience, Carandini gives back to economic decision-making by providing a sensory solution to a decades-old puzzle concerning the propensity to make choices that match the probability of reward. This probability-matching behavior is puzzling because a decision-maker could maximize their reward if, instead, they always choose the option with higher probability. Carandini proves that an ideal decision-maker using logistic classification will adopt a bias that adds to the evidence for the more probable choice. A decision-maker who is unaware of the absence of any sensory evidence and continues to use logistic classification will perform probability matching by acting on the bias alone.

Pulling back from the brink of optimality, the shift to the framework of logistic classification opens sensory neuroscience to a broader decision-making perspective. Economic decision-making focuses on subjective, idiosyncratic preferences rather than the objective external world of ideal observer approaches. Increasing evidence suggests that perceptual decision-making has both subjective and objective components.[9] Carandini's approach accommodates both, for example, by allowing sensory factors and factors more typically considered by game theorists, such as choice history bias, to commingle. There is liberation from the ideal observer approach, as one need not start with what perceptual systems *should* do but can test hypotheses about what it *does* do.

Most importantly, logistic classification provides a means for sensory neuroscience to discover long-sought-after linking propositions[10] that govern how neural activity translates to perception. It therefore seems appropriate that the logistic function is an instance of what statisticians call a "link function"—one that generalizes the general linear model by linking the domain of linear predictions into the range of probabilities. For sensory neuroscientists, the weights, bias, and factors of logistic classification are signals that should be sought in the brain. Neural gain and summation are all that is needed to compute decision variables that might approximate otherwise complex optimal computations. The logistic function then provides a simple quantitative link to predict effects on behavior.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

1. Carandini, M. (2024). Sensory choices as logistic classification. Neuron *112*, 2854–2868.e1. https://doi.org/10.1016/j.neuron.2024.06.016.
2. McFadden, D. (2001). Economic Choices. Am. Econ. Rev. *91*, 351–378. https://doi.org/10.1257/aer.91.3.351.
3. Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. Nature *341*, 52–54. https://doi.org/10.1038/341052a0.
4. Geisler, W.S. (2011). Contributions of ideal observer theory to vision research. Vision Res. *51*, 771–781. https://doi.org/10.1016/j.visres.2010.09.027.
5. Albrecht, D.G., and Hamilton, D.B. (1982). Striate cortex of monkey and cat: contrast response function. J. Neurophysiol. *48*, 217–237. https://doi.org/10.1152/jn.1982.48.1.217.
6. Ma, W.J., and Jazayeri, M. (2014). Neural Coding of Uncertainty and Probability. Annu. Rev. Neurosci. *37*, 205–220. https://doi.org/10.1146/annurev-neuro-071013-014017.
7. Sutton, R.S., and Barto, A.G. (1998). Reinforcement Learning : An Introduction (MIT Press).
8. Swets, J.A., Tanner, W.P., and Birdsall, T.G. (1961). Decision Processes In Perception. Psychol. Rev. *68*, 301–340. https://doi.org/10.1037/h0040547.
9. Gardner, J.L. (2019). Optimality and heuristics in perceptual neuroscience. Nat. Neurosci. *22*, 514–523. https://doi.org/10.1038/s41593-019-0340-4.
10. Teller, D.Y. (1984). Linking propositions. Vision Res. *24*, 1233–1246. https://doi.org/10.1016/0042-6989(84)90178-0.

Impacts beyond the brain: Unraveling molecular mechanisms linking psychiatric, metabolic, and inflammatory conditions

Nathaniel Napoli[1] and Quenten Schwarz[1,*]
[1]Centre for Cancer Biology, University of South Australia and SA Pathology, Adelaide SA 5005, Australia
*Correspondence: quenten.schwarz@unisa.edu.au
https://doi.org/10.1016/j.neuron.2024.08.001

By establishing semaphorin 6D expression in the amygdala as a central coordinator of brain, metabolic, and immunologic function, the *Neuron* publication by Nakanishi et al.[1] provides new insight to how primary brain deficiency impacts physiological systems beyond the brain.

Psychiatric disorders are complex conditions hypothesized to arise from an interplay of neurotransmitter imbalance, impaired neuroplasticity, chronic inflammation, altered genetics (including epigenetics), altered gut-brain axis, environmental influences, and the dysfunctional hypothalamic-pituitary-adrenal (HPA) axis. Connecting the multitude of hypotheses to the origin of psychiatric conditions has been difficult due to the heterogeneous nature of the disorders. As technology and understanding has progressed, the comorbidity of psychiatric illness with many secondary disorders has been established.[2] Importantly, metabolic and immune disorders observed in psychiatric conditions are connected through the heterogeneous and pleiotropic nature of genes and the interconnectedness of the human body and mind. Fitting with modern psychiatry's holistic approach, which aims to understand the complex interactions between biological, psychological, and social factors, treatment of the "whole person" is being explored to alleviate many of the morbid pathologies associated with psychiatric illness.

Dysregulation of the HPA axis provides an exemplar to the origin of comorbidity between psychiatric illness and other diseases, as this pathway co-regulates the stress response, energy, and immune function.[3] The HPA axis is a complex neuroendocrine system acting as a junction between the central and peripheral nervous systems through its regulation of limbic response (including the amygdala and hippocampus), prefrontal cortex function, and sympathetic-parasympathetic neuronal activity. Dysregulation of the HPA axis either directly or, through regulation of other systems such as the sympathetic nervous system impacts brain function, metabolism, and the immune system.

Beyond the HPA axis, extensive evidence links the immunological and metabolic systems to psychiatric disorders. Notably, most psychiatric patients exhibit elevated levels of pro-inflammatory cytokines, such as IL-6 and C-reactive protein (CRP). This correlation is further supported by the increased risk (45%) of developing autoimmune disorders among individuals living with schizophrenia.[4] These inflammatory markers can serve as predictive indicators for psychiatric diagnosis, illness severity, and treatment response. Genetic studies have also identified shared risk factors between obesity and psychiatric disorders, including variations in genes related to appetite regulation and energy expenditure, such as leptin and ghrelin.[5] The co-occurrence