Expectation and Perception: A New Dimension Unfolds

(Bayesian Computation through Cortical Latent Dynamics)

Hansem Sohn, Devika Narain, Nicolas Meirhaeghe, Mehrdad Jazayeri

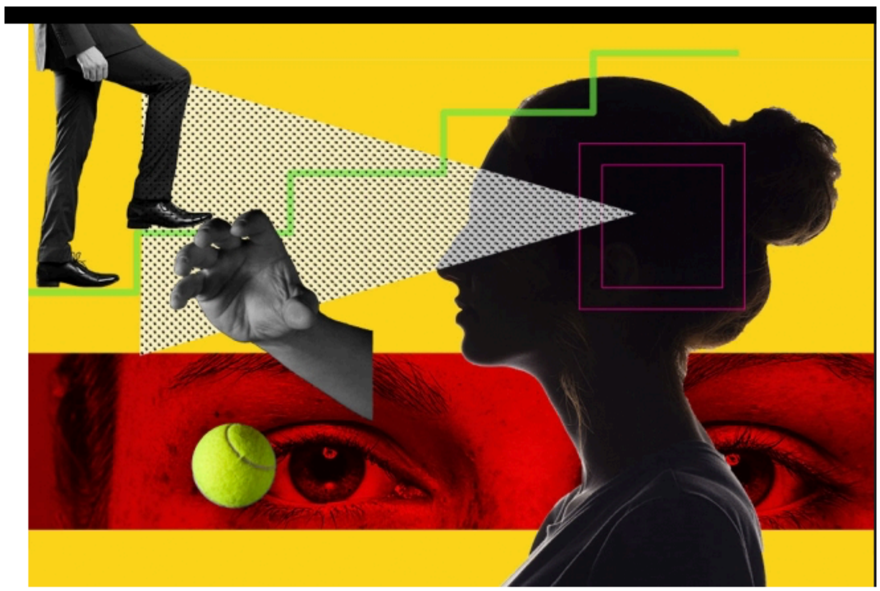

We take some interest in this recent research because it exemplifies a progressive shift of modeling of scientific concepts which can enable escape from the perpetual muddle of trying to relate language usage based narratives of 'psychological events" to various mathematical equations and expression. The authors do so by recourse to a 'geometric" expression which transcends the gap between the language based narrative and the quantitative measurement based equations is being seen throughout science.

When they speak of their 'observations" in terms of various geometric conceptions, they have more latitude for expression and for articulation concepts and "predictions" that are not as likely to reveal 'inconsistency" with the mathematical representations such as those that come with the use of Bayesian concept.

We first came across this shift in reliance on an attempted correspondence between language based narratives of our world and mathematical expressions…to an underlying more permissive reliance on the use of geometric models, in the Nobel Prize winning work of Richard Feynman, and his amazing 'diagrams" which managed to provide the basis for intuitive understandings of quantum phenomena that then led to vast developments in the articulation of the mathematical expression of whatever it was that the diagrams more directly and powerfully 'represented. Today we see the role in quantum physics of concepts such as the amplituhedron which take Feynman's virtuosity in two dimensional rendering into many higher dimensions.

For decades, the authors note, "research has shown that our perception of the world is influenced by our expectations" Precisely what that means is anybody's guess since the notions of 'perception' and "expectation" and even of 'world" are more or less a sloppy pudding and there has never been nor will be any "proof in the pudding that such words present'.

Statistical regularities in the environment create prior beliefs that we rely on to optimize our behavior when sensory information is uncertain. Bayesian theory formalizes how prior beliefs can be leveraged, and has had a major impact on models of perception , sensorimotor function and cognition However, it is not known how recurrent interactions among neurons mediate Bayesian integration

The classic formulation of Bayesian models assumes that the observer integrates a sensory likelihood function with the prior probability distribution to compute a posterior distribution, and applies a cost function to the posterior to extract an optimal estimate. We came across an essay published this week which put it simply, ‘The way the brain identifies and interprets external stimuli and then executes appropriate behaviors remains largely a black box. It is known, for example, that following certain visual stimuli, distinct ensembles of neurons in the visual cortex are activated, but whether this activity is an initial neurological process driving behavior has been difficult to test.”

Statisticians have known for centuries that Bayesian integration is the optimal strategy for handling uncertain information. When we are uncertain about something, we automatically rely on our prior experiences to optimize behavior. “If you can’t quite tell what something is, but from your prior experience you have some expectation of what it ought to be, then you will use that information to guide your judgment,” Jazayeri says. “We do this all the time.”

The process of combining prior knowledge with uncertain evidence is known as Bayesian integration and is believed to widely impact our perceptions, thoughts, and actions. Inspired by this formulation, theoretical and experimental studies have sought to find representations of various components of Bayesian inference at the level of single neuron Unsurprisingly this has and will lead nowhere fast.

Of course, here again, what they believe. when it is expressed in the vocabulary of the shopping mall, has been shown in the past is basically a cartoon like story of some ‘observer” residing within the sphere of action of an organismic being and being the beneficiary of various sensory in puts. Then the ‘observer” is tasked with ‘integrating” some ‘likelihood” function which will allow the arrival at a decision to then engage in one or another action.

As we believe this is simply fishing for yet another, among dozens. of red herrings as long as the underlying mythology of the inside and the outside being expressed in disparate grammars of usage need to somehow connecting by a homuncular decision-making present. We cringe when we hear that “This challenge was clearly stated by Nobel laureate Richard Axel “we do not know the language by which […] patterns of neural activity are […] translated into appropriate behavioral or cognitive output.”

As Axel put it "in the case of Bayesian integration, we need a language that has the potential to explain how synaptic coupling between neurons could mediate prior-dependent biases in behavior." For us that means an expression in terms of geometry which will hopefully put an end to the dependence on the endless chatter of language based narratives..

Thus we were thus rather pleased to note that this current research has been aimed to bypass some of the verbal clutter that strews the path to any reasonable narrative of the goings-on within the brains of organisms that enable them to act intelligently and, by that we mean to "learn from experience".

The authors here state "The central challenge in understanding Bayesian computations is the need for a framework that could bridge explanations at multiple scales. As they say "At one end are cellular-level explanations of how past experiences alter synaptic coupling between neurons, and on the other, are explanations of behavior based on the abstract notion of prior knowledge{

The finding that prior experiences warp the patterns of neural activity provides a window onto how experience alters synaptic connections.   When Jazayeri writes that he seeks to understand what it is that "enables the brain to coordinate internally generated dynamic patterns of activity with external events." He is tacking an inevitable and conceptual problem, not just a problem of the complex world, because the problem itself is already shaped as a difficult   if not " iconically hard", one. Once we have a situation presented to our intuitions with "internally generated patterns of activity" and "external event" , that situation, by virtue of the very way it is represented to us in all the feebleness of ordinary language is going to make any concept of 'coordination" virtually impossible.

We can see here that we have an assumption of 'internal generation" and with as assumption of such a thing as an "event" which is external The project of understanding
"the capacity of the brain to flexibly stretch or compress neural activity patterns to anticipate

external events and plan for upcoming action" then becomes a project which needs to free the scientists intuitions from the constraints of ordinary language use and its limited logical potential for expression.

There is nowhere to be found a more glaring instance of the failure of neuroscience to free itself from the quicksand of ancient models of perception and expectation than in the silence in most areas of neuroscience research in regard to what have been called 'efference copies" , and indeed that silence extends to this publication as well, as if by some magic they can rely only upon cortical events in their narrative.

For those in neuroscience, the event which accompanies all motor action and which has lamentably been called an 'efference copy" is the alleged transmission from any motoric event of a neuronal signal somehow mirroring the expected external response to the movement. This ridiculously designated 'copy" is then somehow to serve as a comparison by means of which the action which is chosen or decided upon fits with the 'expected" sensory/perceptual information from the world.

Presumably this immediate and direct transmission of the efference signal can enable the organism to adjust on the fly, as it were, any discrepancies between the motor 'command" (OMG> .yes this is how they talk) and the actual movements before the actual movements result in adverse consequences to the organism. So the best that neuroscience can do up till now is to give us picture of a captain of a ship, issuing motor commands, and each motor command somehow generating a mini-copy of itself to directly reach his sensory area (the dashboard???) before the senory area is clobbered by an adverse external event so the captain can see before allowing the execution of the action whether it matches up to his expectations?

However rather than relying on the always available homuncular 'observer" the authors here look for other modes of expressing the events of the brain's minding' Theoretical studies and recent artificial neural network models have established a framework that could potentially address this challenge. According to this framework, the key to a deeper understanding of how neural circuits perform computations is an analysis of the geometry and dynamics of activity across the population

"The brain seems to embed prior experiences into synaptic connections so that patterns of brain activity are appropriately biased," Jazayeri says Researchers believe that prior experiences change the strength of connections between neurons. The strength of these connections, also known as synapses, determines how neurons act upon one another and constrains the patterns of activity that a network of interconnected neurons can generate

To make sense of these signals, the researchers analyzed the evolution of neural activity across the entire population over time and found that prior beliefs bias behavioral responses by warping the neural representation of time toward the middle of the expected range.

Recent studies have focused on an analysis of the geometry and structure of *in-vivo* cortical activity in trained animals and *in-silico* activity in trained recurrent neural networks (RNNs) to gain a deeper understanding of how neural dynamics might give rise to behaviorally-relevant computations

The MIT neuroscientists have discovered distinctive brain signals that encode these prior beliefs. They have also found how the brain uses these signals to make judicious decisions in the face of uncertainty. "How these beliefs come to influence brain activity and bias our perceptions was the question we wanted to answer,"says Mehrdad Jazayeri, the Robert A. Swanson Career Development Professor of Life Sciences, a member of MIT's McGovern Institute for Brain Research, and the senior author of the study.

They wanted to understand how the brain encodes prior beliefs, and put those beliefs to use in the control of behavior. Once they had established that the animals relied on their prior beliefs, the researchers set out to find how the brain encodes prior beliefs to guide behavior.

They recorded activity from about 1,400 neurons in a region of the frontal cortex, which they have previously shown is involved in timing.

When Jazayeri tells us, “The brain has a remarkable ability to generate complex behaviors by combining sensory evidence, prior experience, and cost-benefit considerations…and that his research probes the neural mechanisms that allow the brain to integrate this plethora of cues, resulting in flexible, goal-directed behavior”. he cannot afford to allow his own intuitions to be guided by the primitive sequence of causes and interactions which are presented in the verbal narrative, and his recourse to geometry to facilitate understanding both allows him to find some ‘fit”….if not a completely perfect one to the data of their measurements and their equations of Bayesian sort, but then such ‘understanding” also enhances the articulation of further questions and hypotheses that derive from the adequate geometric model….that would never arise without the shift to geometric representation and conceptualization.…

We found that prior statistics that were presumably embedded in the coupling between neurons, established low-dimensional curved manifolds across the population. Recording from neurons in the frontal cortex revealed a simple mechanism for Bayesian integration: Prior experience warped the representation of time in the brain so that patterns of neural activity associated with different intervals were biased toward those that were within the expected range.

This curvature, in turn, warped the underlying neural representations and afforded biases in accordance with Bayes-optimal behavior. This mechanism was evident across multiple behavioral conditions including different prior distributions and different effectors suggesting that it may be a general computational strategy for Bayesian integration.

Based on their understanding of the geometry of the problem they reasoned that a good candidate for the encoding axis was the vector pointing from the states associated with the shortest to the longest *"times"* for each prior condition. They projected neural states to generate a one-dimensional representation

We all have a view of science as producing 'theories" of one kind or another which are verified by experimentation and the observation/ measurement that is part of the classic scientific method. However a close look at any theory, even the most exemplary ones, is that they cannot be said to somehow capture the 'truth", and, indeed if we we wish to play games with our own language usage, we can claim that every theory, ultimately is shown to "not be true", as we sometimes hear in the discussions of the archetypal 'paradigm shift' between Newtonian and Einsteinian theories.

But the fact of the matter is that theories are not meant to be, nor can they ever be 'true' in some idealist Platonic…and truly nonsensical use of the term. They are assessed by their 'adequacy', that is, insofar as we guide ourselves by their conceptions, can we then go on to generate actual technological events that make the world a better place…or, in the limited sense can be conduct experiments testing the particulars of these theorist, that may be shown to verify them or not. However, as we can see from the Newtonian example, what verification entails is not being 'inconsistent with result to be expected" and those expected results are directly related to the details and nuances of the circumstances

So the aim of science is actually give us ways of expressing ourselves that are 'falsifiable" as Popper would say, but are such that they continue to not be 'inconsistent" with measures and events expressed mathematically by our paradigms. The advantage of a geometric model is that it allows a great latitude and range that the clumsy domain of language use and its pedestrian logic. That is presumably why geometric models are now so prevalent in theoretical physics.

They now believe that they have found a novel computational principle for how neural circuits perform Bayesian integration. It turns out that structured connectivity creates low-dimensional activity patterns across the population with powerful computational capacities for integration   categorization   gating timing , learning movement control and forming addressable memories

These experiments allowed us to reveal the role of bias in compensating uncertainty (i.e.. larger biases for noisier measurements), and validated the role of the curved manifold for integrating prior knowledge into behavioral responses.

Numerous studies have found an important role for natural scene statistics in vision, and have further shown that the organization of tuning in neurons of the primary visual cortex follow those statistic. This observation is often explained in terms of efficient coding   which is a statement about the nature of the representation, and not about the underlying computations

However, what distinguishes this geometric approach is that it does not stop at the representational description. Instead, our results show how biased tuning across single neurons leads to warped representations in population dynamics whose geometry can explain the underlying Bayesian computations.

Instead here, to investigate the neurobiological instantiation of Bayesian integration at the level of cells and synapses, future experiment they examine how functional and causal measures of coupling between neurons may change while such prior-dependent curved manifolds are formed across the population

Remarkably, this computational strategy ,, when represented in these low dimensional geometric terms that allowed further examinations, also emerged spontaneously in an artificial neural network trained on the same sensorimotor task. Analysis of recurrent neural network models performing the task revealed that this warping was also enabled by a low-dimensional curved manifold and allowed a further probe of the potential causal underpinnings of this computational strategy

This thinking of their data in terms of the geometric model and then the 'thinking ' of their potential implications in terms of that model that it generates in the mind of the scientist is either right or 'wrong" because it offers more possibilities, but that as a tool to guide our thinking, when it is used in relation with the measurements and their equations and formulations it remains 'adequate " in a wider range of circumstances…and much of that enhanced adequacy is because a geometric representation can side step the paradoxes and pitfalls of logic that plague language based narrative of the world. It is simply more adequate to the task of science.

Reverse-engineering the model thus revealed that it solved the task the same way the monkeys' brain did. That is, the model also had a warped representation of time according to prior experience. Although they focused on Bayesian integration in the domain of time, the key insights gleaned from our results are likely to apply more broadly to the general problem of Bayesian integration in perception, sensorimotor function and cognition.

Moreover, the network allowed them to probe the causal role of the underlying mechanisms *in-silico* using a set of targeted-dimensionality perturbation experiments that are currently not possible *in-vivo* They speculate that the same framework may provide valuable insights into the link between efficient coding and Bayesian perception as well as numerous other sensorimotor and cognitive functions…but without the switch in representation by means of the which the world of the brain is studied to a geometric one rather than language base ordinary account, there would no such optimism possible.

Footnote: There really is no sense to the distinction between 'expectation" and "perception". They are not different…and form another perspective they may be 'the same thing". The actions we said to "take" and the world we are said to "see" are truly not two aspects of life, or of biology which have to be reconciled. Of course, once we fail to see them except as ‘separate”, he have dug a huge hole for ourselves to cross…and there will no bridge that can carry our intuitions across that gap.

As we look more closely at what the representational mode which we know as ‘geometry” provides, it is far more than a picture of some world beyond us….as a map is not at all in any picturing relation to a territory.   Just as a map is not but a limited of possible ways to cross a terrain that are coordinated as actions we may or may not take….so does the brain not have anything like the simple ‘cognitive’ map as a representation of a world, but any such ‘map”is not merely a metric…or geometric rendering but is an awareness of countless possible movements or actions in a ‘space” all of which might get us to a destination…and each of which constitutes a component of a’decision” on how to achieve that result ….