Inference and Prediction in Neocortical Circuits

The American Institute of Mathematics

Table of Contents

Chapter A: How does human vision make good perceptual guesses about objects?

Daniel Kersten

Talk Summary

The speaker provided an explanation of how human vision makes good perceptual guesses about objects using bayesian influence graphs. To this end, he decomposed an object S and its image I into world/object properties and features.The resulting probability distribution was then represented using a graph where nodes represent random variables and links represent the influences. He then illustrated this approach using examples of discounting and cue integration. In the second half of his talk the speaker concentrated on V1 and LOC mechanisms involved in perception. His experiments revealed that perceptual organization correlated with reduced V1 activity and increases LOC activity. He showed that V1 activity can predict percept on the time-scale of behaviour. The decrease in V1 activity could mean two things 1) Predictive Coding 2) Sparsification. In predictive Coding high-level object models project back predictions of the incoming data. In this case a good fit implies a low activity at the lower areas due to subtraction. (The "shut-up" theory). In sparsification, a good high-level fit tells the lower areas to 'stop gossiping". This essentially amplifies the activity for features belong to the object and suppress the rest. Since predictive coding and sparsification have the same empirical fMRI observation, the experiments were inconclusive in deciding which of the two mechanisms (shut-up or stop gossiping) reduces the V1 activity when a higher level (LOC) has a good explanation for an object.

Questions During the Presentation:

- Bill : Could you explain the term Discounting?
- Ans : Discounting is throwing away information that is not required for perception. For eg: lighting
- Jeff: (Regarding shape cues explaining away pinkishness as white paper and pink light)Are these things learnt?
- Ans : Don't know yet. We are in the process of esigning and experiment to find that out.

Discussion after the Presentation:

- Grossberg: In non-stationary environments priors don't exist. Also the bayesian framework breaks apart for some Kanisza illusions.
- Bruno: Nothing in the Bayesian framework is inconsistent with Kanizsa. Bayesian framework is taking out the Kanizsa illusions and explaining those. The correct choice of priors is the trick.
- Dan: For eg: one prior for natural vision is that motion tends to be slow.
- Bill Softky: You showed us data showing suppression of activity in V1. Is there any data for suppression of thalamus?
- Dan: In an fMRI experiment, localizing the thalamus and especially the LGN is not easy. Hence we don't have any data for that.
- This was followed by Steve explaining a Kanizsa illusion.

Chapter B: Anatomical Substrates for Functional Responses of Neurons in the Primate Visual Cerebral Cortex

Alessandra Angelucci

Alessandra's talk addressed specific brain circuits and perception. Specifically, it reported results on the study "How does the spatial scale of classical receptive fields and and surround modulation compare to the spatial scale of horizontal, feedforward and feedback connections?

The study consisted of injecting tracers and correlating the connection spread to the receptive field size at various levels (V1, ...,). The speaker noted that receptive field sizes depend on the measurement technique. Typically encountered measurements of receptive fields are (1) the minimum response field is denoted by mRF, (2) the summation receptive field at high contract denoted by hsRF and (3) the summation receptive field at low contract denoted by hsRF.

Her study showed that horizontal connections are commensurate with the low contrast sRF size of V1 cells and might mediate expansion of the sRF at low contrast and collinear facilitation effects. Moreover, the extent of feedback connections were found to be commensurate with the whole range of V1 cells center surround field sizes.These connections might mediate center-surround interactions and global-to-local signal integration.

She also reported the following; (1)Feedback connections are 10 times faster than horizontal connections, and as fast as feedorward connections. (2)Feedback connections are patchy (3)Feedback connections to V1 are specific (4)Horizontal and feedback connections arise from excitatory neurons (5)There is no such thing as "long-range inhibitory connections" in visual cortex. (6)Horizontal and feedback axons contact excitatory (approx 80%) as well as inhibitory (approx 20%) neurons

Discussion:

Question: Our impression was that feedback is diffused and not orientation specific? Ans : The results presented here say the opposite. It says that feeback is patchy and orientation specific.

Qn: Connections from V1 to LGN? Ans: We don't have any data on that.

Qn: Did you say that 98% of the feedback connections go to excitatory neurons? Ans : No, the division is 80/20

Chapter C: Breakthroughs in Brain Computing

Steve Grossberg

Discussion Following the Talk

Q1: Why do you choose very specific kind of laminar circuits? What is the particular reason for ur circuits? Any predictions for us?

Answer: I would like you to refer to my paper with Razada. There we stimulate in layer 6 and observe what happens in layer 4.

The reason for Laminar architecture and how we came to it: We started working on grouping. What are the units of perception? To get at the grouping properties and analog coherence required certain ordering of mechanisms or else the data would collapse. Needed something which is robust and not parameter sensitive. We could see by analyzing the anatomy that it provided this required ordering of mechanisms in a compact way. We tried knocking down parts of it to see whether this breaks down and we found that the entire thing was necessary. What we basically have is a minimal mechanis for the required ordering property.

Qn:Do you insist that the lateral connections add superlinearly? Ans: Not superlinearly but faster.

Comment: The normalization property with just one inhibitory neron might not be sufficient for all. By variations of the simple circuit you can get a variety of effects from bipole to modulatory.

Qn : What is your prediction for an experiment where there is a stimulus in the center and the flanks are detected?

Qn: Does your model need spikes?

Answer: No you don't need spikes. Spikes are important in some situations (self syncrhonizing nets, order preserving limit cycles) but not required for other.

Qn: Why are spikes there?

Ans: So that we have non attenuated signal transmission over long distance. Also with spikes, we can compensate for axonal delay by increasing the diameter of axons.

Also I talked about a balance between excitation and inhibition. If these balances doesn't occur due to development then the system becomes to unstable. If the system is designed to 'stabilize' using spikes, then spikes become important. These systems show that they love to resynchronize these spikes for stability reasons.

Qn: Does your model do computer vision tasks that are not done yet?

For this the speaker answered with a list of computer vision tasks his algorithms have been employed on.

Qn: The neurophysiological detail is that there are spikes.

A: Why don't u put in all the channels?

Jeff (Comment): Analogy between computers and quantum mechanics. Although the transistors operate based on principles of quantum mechanics, an understanding of quantum mechanics is not required to understand the working of computers.

Speaker(Comment): My equations are mean values of stochastic differential equations.

Speaker(Comment): An example where spikes are not important: Phonemic Restoration. BB noise colored by

eel is on the

wheel wagon

peel orange

heel is on the shoes.

Pentti(Comment): Syncrhonous computing might not need spikes. Need not be the case for asynchronous operations.

Chapter D: A saliency map in primary visual cortex

Li Zhaoping University College London

Talk Summary: The speaker presented a model of V1 that produces a saliency mapl. This model transforms contrast inputs to saliencies using contextual influences. The contextual influences are implemented in a recurrent network with intra-cortical connections.

Saliencies are signalled using firing rates of the neurons.The speaker compared the performance of her model with known psyho-physical results.

Discussion And Questions Following the Talk

Q(Bill): Do all V1 cells signal saliency or a few of them.

A: Yes all cells.

Q: What is the role of other visual areas in this model?

A: No role.

Speakers' Comment: V1 cells firing rates signal saliencies, despite their feature tuning. Strongest response to any visual location signals saliency. Also, this theory is only bottom up. V1's output as saliency map is viewed under the idealziation of the top-down feedback to V1 being disabled. eg. shortly after visual exposure or under anesthesia. (Saliency from bottom up factors only)

Q: Are these results after the system have settled?

A: After one time constant.

Q: Long range excitatory connections are not isotropic. In your figure it seems as if they are not.

Comment: The speaker then explained how they are not isotropic in her model.

Bill: What about color, motion etc. how do u combine in the model?

Ans : These are not included in the current model.

Q: Did you do any of these on natural images?

A: No the limitation is the processing power.

Q: Do you need any endstopping?

A:

Steve (Comment): I am worried about about the use of the word "saliency" map for the visual search tasks you mentioned. For eg we have found that 3D properties (surfaces) can affect pop-out mechanisms and the word 'saliency' is usually used in a more wider context. Have you done any experiments using surface properties?

Answer: I just compared with visual search tasks. This model does not explaining 3D surface properties.

Dana(Comment): The same data can be explained using Signal to Noise Ratio. (Triesmann data) Some of the experimental data could be false because they were done in blocks.

Q:What about multiple scales?

A: Same principle will apply at all scales.

Q: Regarding the disynaptic inhibitory connection, is there a time delay that is required?

A: Yes. To get rid of the symmetry breaking behaviour we need the disynaptic behaviour.

Chapter E: Different Functional Roles of Feedback and Horizontal Connnections

Virginia de Sa

Q: (Regarding the processing of lip movement patterns). Why did you use the motion and not the position.

A: Expediency. That system was already working when I started my work.

Q: Do you have any sense of how this scales for more categories.

A: No we haven't done the experiment on more classes. If we have more modalities the task is clearly easier.

Comment: Feedback and horizontal connections both bring in greater context but feeback conxsn are unique in bringing back information from other sensory modalities. Horizontal conections tend to connect like recpetive fields.

Comment(Regarding the slice experiments on rat): There aren't many lateral connections in rat. So her experiment was on short-range connections.

Q: (Regarding the slice experiment) Where is your recording electrode?

A: In layer 2-3.

Q: Is maximising agreement different from minimizing disagreement?

A: In the way it is defined here, both turn out to be the same.

Comment: Isn't it the independence of the modalities that helped the inter-training of the modalities?

Q: Does the association come about because lip-movement and the associated phoneme have the same physical cause?

A: Yes.

Chapter F: Distributed Syncrhony

Zhohua Zhang, Dana Ballard. (Speaker)

Talk Summary:

The speaker presented a spike-timing based communication mechanism for inter-neuron communication. He suggested reasons why a rate-coding based approach will not work. The proposed model includes a clock of period 20 msec. A signal P to be transmitted is encoded as the multiplication of a probability of firing sqrt(P) and a phase offset of sqrt(P) with respect to the clock. This model can reproduce observed PSTHs and receptive field characteristics. The 'Distributed Synchrony Model' was motivated by three criteria (1) Fast, reliable intercortical communication, (2) the need for a cell to multiplex and (3) the need to reproduce observed cell responses.

Discussion and Questions

Q: Is the cortex trying to reconstruct the image or trying to extract some features from the image?

A: In this case yes..trying to reconstruct.

Q: (Regarding the PSTHs). Your model will reproduce the PSTHs only when the stimulus is presented in synchronization with the internal clock.

A: That is true. But we assume that the cortex picks up the signal from the retina in a syncrhonized fashion.

Q: (Regarding) Coding cost of different strategies. Does the cost come from the learning part or the prediction part?

A: We don't make that distinction.

Q: Does your feedback rule have any relation with EM?

A: The feedback rule is motivated by the EM rule.

Q: (Regarding the synaptic weights used for receptive field simulations) Are the synaptic weights pretty similar or are they different.

A: Depending on how we pick the weights we can get different receptive fields

(Steve): You have a shut up model. But in my talk I gave an example where top-down feedback is excitatory. Could you comment on all these experiments that say that feedback is not shut up.

A: In this model feedback is indeed shutup.

Comment(Alessandra): Anatomical studies show that feedback is excitatory in center and inhibitory in surround.

Comment(Steve): I thinkt that end stopping is caused by feedback from V2 is false.

Q: It seems that phase has done two things with the clock. Multiplexing and Coding.

A: That is right. The width of the 20 ms is related to the number of 'threads' that you have.

Q(Bill): At the beginning of the clock..how does the neuron forget everything?

A: We assume that there is some magical swap space available for storing all the information and then swapping it back in.

Q(Bill): Also how does the cell generate the random number?

A: There are some known mechanisms.

Q: Are spikes are essential?

A: I think spikes are an innovation by biology to send signals reliably over long distance to

Q: If we have a package of spikes ....what is the necessity of clock?

A: So that we can do feedback. If there is no clock, there is no way to align the feedback signal.

Q: How neurons syncrhonize critically depend on the neuron model. For eg the Aertson model

A: He has an integrate and fire in combination with a random number generator.

Chapter G: Helmholtz Inference in Early Vision Areas

Ken Nakayama

Discussion/Questions After the Talk.

Dana: If surfaces are creating the percept of an image space, it seems what you are calling generic views are quite in line with Bayes.

A: When I talk about Bayes, the priors are stochastic and non ergodic. Also the Kanizsa diagrams show that we are not using the priors. Bayesian formulation is far too unconstrained and general.

Comment(Jeff?): I see most of the figures as ambiguous.

Q: Isn't the generic view principle a prior?

(Comment): Bill Freeman has a paper which said the generic view principle in terms of Bayes.

Q: Most of the things can be explained because you see motor experience and not just image ?

Comment: We see the impossible triangle through bottom up process even though we know at the top that its not possible.

Q (Jeff): You are confusing with models and theory. Do you think there is a neurobiological explanation for all this.

Answer : Reductionism is metaphysical.

Comment(Dana): Through development we get the the kinetic depth and stereo vision. We might not be able to get this figured out until we track the developmental trajectory. Also perception and motor systems do interact.

Comment(Jeff): All the speakers yesterday were in the optimistic camp and today we have the speakers in the pessimistic camp

Chapter H: Neural Mechanisms of Perceptual Inference

Rudiger vonder Heydt

Q: What percentage of the recorded cells were in V1 ?

A: 50

Q: Are these unbiased selection of random cells?

A: Yes, as we physiologists do.

Q: How is the response defined?

A: Its the total number of spikes in that window. (800ms).

Q: What is the largest extent of the visual field you tried?

A: 8 degrees.

Q:How do you measure the intitial receptive field of a cell?

A: We use an optimal bar. We move the bar out and figure out the point where this cell does not respond any more. The minimum region outside which you don't get a response.

Comment: Other techniques can give different receptive field sizes.

Q: Figure ground and ownership?

A: Its just different words for the same thing.

Q: Didn't understand the comment about influence of IT.

A: We have the half-strength of the ownership signal at 70ms. IT, upto 100ms is silent.

Q: What about the illlusory contours? ARe they more local?

A:

Q: Does that mean that the illusory contour is withing the summation field of the neuron?

A: What we know as the summation field is quite controversial. Summation os not an explantion of the illusory contours.

Q: What about natural scenes?

A: I suspect that some of this would hold up.

Chapter I: Synaptic Integration in the Early Visual Pathway

Judith Hirsch

Discussion after the presentation.

Q: Isn't there a push-pull mechanism in the retina?

A: Yes I think there is. But it has to be recreated in the thalamus because the connections from retina to LGN are all excitatory.

Q: What is the proportion of cells?

A: Don't know.

Q: How do you select the response window?

A: Centered around the peak of the response.

Q: (About Laminar Distribution of Simple Receptive Fields). Will this hold in monkey as well?

A: No. I expect it to be different. The circuits that form the simple cells in monkeys are very different from than in cats.

Comment : Excitatory and inhibitory cells look differnt. Excitatory cells are spiny. Inhibitory are smooth.

Q: Orientation selectivity is coming from??

A: The pull can contribute to orientation selectivity.

Q: What do the complex cells look like?

A: They don't respond well to the "sparsely dynamic" stimulii. You need just the correct data set to come into a cell to make it fire.(?)

Q: Does anyone try to knock out LGN and try to record from V1?

Q: Do inhibitory cells make contact with other inhibitory cells?

A: Yes I believe they do. It seems that the orientation selective inhibitory cells have the same inhibition pattern as the excitatory cells.

Q: Did you look for evidence of syncrhony now that you have intra-cellular recording.

A: There are lot of blips

Q: What was the relative receptive field sizes of the complex interneurons and the simple excitatory neurons?

A: Roughly the same(?). We don't have an accurate measure.

Chapter J: Resonance Prediction and Priors

Tai Sing Lee

Q: Can you tell us anything about the time course of the response between V1 and V2?

A: This is hard to detect in our experiments.

Q: Is there a connection between predictive ideas and pop out ?

A: Here top down predictions are more like a prior

Qn: What is the attention effect in V2 neurons?

A: The V2 neurons don't get attenuated even when you are attending to other locations.

Q: Did you try turning off the target? If you turn the target off after 100ms. then probably you would see something interesting.

A: We didn't try that. But this looks like an interesting experiment to do.

Q: Did you do experiments where the monkey doesnt know where the pop out ball is?

Q: The latency of the response to the popout seems to be very long to me comapred to the poput of bars.

A: Here the pop out is working on 3D surfaces . Infering this 3D takes longer time.

Chapter K: Notes from Breakout Session I

On the second day of the workshop the participants broke out into four different groups based on their research oreintation. The groups were (1) Pyschology and Psychophysics, (2) Anatomists, and (3) Theorists. Due to the large number of theorists, the theory group was split into two (Theory 1 and Theory 2).

The groups met separately for 45 minutes. Each group was supposed to come up with a list of three things that they want from the other groups. At the end of 45 minutes all the groups met at one place and a discussion ensued on the requirements. What follows is a compilation of suggestions from different groups.

Suggestions From Psychology/Psychophysics group

For theorists

- models that can act on real stimuli real images, moving images
- Are there associative memory models running on the hardware that we have? (spiking models)
- Theory of mid-level vision that is understandable to us.

For anatomists

- more quantitative descriptions organized into ciruits.
  proportions of neurons in one area compared to another area are they myelinated or not
- More oranized quantitative information.
- Timing
- Thalamus
- Anatomists comments on theory slides.

From Theory Group I

For Anatomists

- Naturalize Stimulii
- Make the raw data available.
- Fast-forward development of multi-electrode recordings.
- Specify data on excitatory/inhibitory..which layer they are in.

For Psychologists

- Is the BOLD signal reactive or predictive? Need explanation.
- Request for more mathematical/theoretical training.
- Theoretically driven experiments.
- Theories posited before data.

From Anatomists

For Theorists

- Falsifiable predictions
  crucial tests.
  - test on how bayesian inference occurs in brain.

- Rigorous comparisons between models.
  - surface representation (biologically plausible)
  - Where does the model fail?
- Different feedback systems.
  - predictions from these models.

To Psychologists

- Ken: What approach, what end point?
- More natural stimulii.

From Theory group II

For Both Anatomists and Psychologists

- Better communication
  - theoretical training (academic curriculum)
  - tutorials or invited talks @ conferences.
- Better Data
  A) centralized database, systematic connectivity diagram
  B) Reporting (availability/completeness)
  C) Natural stimulii/ multi-electrode recordings.
- Hire us!!

Chapter L: Notes from Breakout Session II

The four groups met separately to prepare responses to the comments from the previous day.

Reply From Psychologists

The theorists wanted to be hired. The primate experiments would go much better if they did psychophysical experiments.

Giving away raw data looks like giving away ones own work.

Alessandra: Once something is published it is public domain and so the data should me bade public.

Judith: It could be that you have one sort of analysis and more analysis might be waiting. you just don't want to give away the data.

We should be able to publish the data

From Theory group I.

- Are there associative spiking models with spiking neurons? (Yes). Fritz will give a presentation on that.
- Theoreticians will become aware of training in experimentation. This will give us more intuition. But we need more help in doing that. We would also like the data from you!
- Russ: Rigorous model comparisons between different models. One can apply the same inputs to models and get the results. Should document when the models fail.
- Falsifiable predictions? : We alread do.
- The job of theorists is not just to make testable hypothesis.

From Anatomists

- Now more labs are producing less descriptive and more quantitative data. There is some kind of questions that take too much time to address. (For eg: number of synapses). How can computer science help us for this?
- There are computer vision techniques and tools available that could be used.
  Multi-electrode recording? Its a lot of effort. If the cells cannot be activated at the same time, it amounts to recording one after another.
- Natural stimuli.
  We have acquired multichannel recording systems and we are racing to do multi electrode recording.

From Theory II

- Bayesian Framework. : Bayesian framework is a principled way of looking at the computational problems faced by the brian given the ambiguous
  Priors can be implemented as lateral connectivity or synaptic weights or top-down predictions.
  Predictions for testing the bayesian framework: Rajesh Rao's work
- Different forms of feedback : We really don't know how these forms of feedback are. The diffused feedback that is topographic could be doing an spatial priming and the patchy ones are doing feature priming.
- Tony: Natural stimulii. Adaptive models: Natural stimulii. Yes.
  In the case of hardwired models: it will give tighter prediction to the experimentalists.