

JSMF | SANTA FE INSTITUTE

COMPLEX TIME

AGING, ADAPTATION & THE ARROW OF TIME

Cognitive Regime Shift II

November 12-13, 2019

Post-meeting Reflection by Presenters

Steven Petersen (Washington Univ.-St. Louis)

I found Nikos's presentation enlightening and on point. I will certainly go to his primer article. I think the imbalance between abstraction and empiric work was strong. There was an uncomfortable level of abstraction for me. I am not sure my perspective has changed too much, because of this. I would have liked to spend more time on the questions Russ raised at the end regarding ways to "get things together".

David Krakauer (SFI)

Much of the emphasis was placed on describing the necessary basic principles, models or data, for describing brain functions. These included:

1. Resting state correlations from imaging data
2. Behavioral psychological experiments
3. Local field potentials
4. Deep neural networks
5. Information theoretic formalisms.

Much emphasis was placed on either justifying or discovering appropriate levels for prediction and explanation. On this topic;

1. Is there a preferred level based on fundamental principles?
2. How to reconcile computational models (with strong time separation) with dynamical systems models (with a spectrum of time scales)
3. How to present and justify theoretical frameworks with many free parameters - theory for complex systems (in contrast to mere complication as in physics).
4. How to triangulate among levels of description

My own question dealt with the general problem: does the fact of the brain as a computational organ imply distinct regularities in the way in which it breaks? One approach to this would be to ask about:

1. Robustness and adaptability
2. Critical transitions: order disorder regimes
3. Cascading failure and percolation.

This triplet provides a possible informal coordinate system in which to situate a system to include the brain. The rather unique scale and connectivity and general function of brain might suggest that it sit near a critical point, balanced between robust and adaptive regimes.

JSMF | SANTA FE INSTITUTE

COMPLEX TIME

AGING, ADAPTATION & THE ARROW OF TIME

John Krakauer (Johns Hopkins Univ./SFI)

Very much enjoyed the talks that were more theoretical/philosophical - David Krakauer, Nihat Ay, Niko Kriegeskorte, Artemy Kolchinsky.

Awareness that had a vaguer notion of brain breakage and failure than I previously realized. That embodiment perspective at odds with DNN. We Still have not reconciled networks, dynamical system and representational views. Distributed cognition a but empty.

I think there need to be more meeting like this - very useful for neuroscientists.

Impact on my own work likely through new collaborations.

Nikolaus Kriegeskorte (Columbia Univ.)

Impact of the meeting on my research: The main thing I took away from the meeting was that robustness to damage is property of neural networks that is important from a theoretical as well as an applied perspective. I'm motivated now to think about revisiting the old method of lesioning and perturbing with noise neural network models, so as to better understand the degree to which they are sensitive to local damage and to displacement of their dynamic trajectories.

I would also like to explore how robustness to damage and noise relates to robustness to adversarial attacks and to robust generalization performance in novel situations (i.e. changes to stimuli or more generally the behavior of the environment). Before this meeting, I would have conceptualized these different forms of robustness as largely unrelated. I would have thought that the commonality suggested by the fact that they can all be considered forms of robustness is largely misleading.

Now I think there may be deep links between robustness to damage, internal noise, changes of the body, and changes of the stimuli and environmental behavior. Some half-formed hypotheses:

- Neural noise may help a neural network learn solutions that are robust in a variety of ways.
- Similarly, changes of environmental behavior (including stimulus statistics) during learning may help a neural network learn solutions that are robust in multiple ways.
- Predictive completion of partial neural representations (in symmetric, i.e. energy-based, or asymmetric networks) may provide robustness through redundant representation as well as enabling unsupervised learning through self-supervision.

I'm quite keen to explore these ideas further.

Important open questions:

- How can states close to criticality serve computation in neural networks?
- Might robustness to damage result from the same mechanism that enables robust generalization to new domains? (And what is that mechanism?)

Notes on the meeting

The discussion, though inspiring, was a little more wide-ranging than is optimal for making concrete progress. It might good in future meetings to focus, and to more clearly and specifically define the topics of sessions and particular presentations.

Post-meeting Reflection by Non-presenting Attendees

Jacopo Grilli (ICTP)

I am stuck with a picture of aging and collapse, motivated by catastrophic shifts in ecology, which simply takes the form of a saddle-node bifurcation. A functional and dysfunctional system are separated by some energy barrier. Aging (somewhat by definition) corresponds to decreasing energy barrier height (and therefore increasing probability of transition). This (at this level tautological) view comes with two interesting consequences:

- (critical) slowing down: the typical timescale at which fluctuations relax increases over time
- in multidimensional system there is an effective one dimensional trajectory describing collapse
 1. The latter point, suggests high reproducibility in collapse trajectories. At what scale this framework is useful is unclear to me. At the coarser scale, when only two states exist (functional and not functional) the only thing that matters is transition probability (the when, and there is no how and why). At that scale bridges and brains fail in the same way (as lifetime distributions sort of match). I am very confused about the confusion around the scale(s) at which we want to study aging and breaking of brains.
 2. I found extremely interesting the discussion of machine learning / neural networks as toy models of representation and/or learning in brains.

Caterina Gratton (Northwestern Univ.)

I started my section with the premise that when we discuss the brain "breaking", we often operationalize this in terms of changes in complex behaviors, and that these behaviors may be subserved by large-scale systems of the brain. Much of the work to date has focused on the typical average structure of these systems. I showed some recent work we've done aimed at moving these analyses to the individual level, and discussed some observations we've made based on this.

- (1) I showed that functional network measurements (at rest) can be quite reliable even in single individuals, given enough data
- (2) I showed some data demonstrating that functional network measurements are dominated by stable factors including group commonalities and individual features. Task-state and day-to-day variability is also present, but much smaller in scale.
- (3) I discussed our characterizations of punctate locations of individual differences in functional networks, showing that these locations are present across repeated recordings, relate to altered function, and individuals cluster based on the forms of variants they exhibit. While these individual differences explain some (gross) behavioral differences, the variance they explain is very small. I left off with a

question to the group of why this might be: why do we see relatively stable behavior in the face of some large individual differences in brain organization.

Discussion centered on how we might think about these effects in the context of distributed organization (or not), to what extent these effects can be overcome by functional alignment that does not assume spatial correspondence, and whether manifolds might be a way of modeling variation in brain function that can lead to a similar functional outcome. We also discussed whether behavior has been measured well enough yet, or if we've been too non-specific in our functional assessments.

General meeting reflection: There were some interesting discussions of multiple different scales and ways of thinking about the brain. I would have liked to have seen a little more cross-talk integration, and/or thoughts about practical directions on which to move forward. How can we better unite models with data? What are the right types of data to collect and theories to test?

Dietmar Plenz (NIH)

I enjoyed two talks in particular – David's introduction into 'breaking' which provided a nice meta-overview into brain dysfunction outside the usual context of development and aging. Refreshing and lots of food for thoughts. The triad of 'breaking/perturbation, critical transition, and cascading failure' nicely transitioned into three more concrete directions, which I would have loved to explore more in that workshop:

- 'breaking = scale of anatomy'
- 'critical transition = brain state'
- 'cascading failure = developmental disturbances'?

Also, very much enjoyed Nikolaus's overview talk and insight into convolutional deep networks. Very clear, transparent and a great platform from which discussions emerged.

My favorite open question: What is the computation mechanism/dynamics at the network level? Move away from correlation analyses.

Change in perspective: I would like to move away from the discussion of imaging results and move more towards the nature of computation.

Impact on my own work: Converging ideas on collective decision making and coherence potentials.

Susan Fitzpatrick (JSMF)

I VERY MUCH LIKE THE idea of taking one disorder -- say Parkinson's and see if we can 1) describe what is meant by Parkinson's at multiple levels of analysis 2) accumulate observations (genetic, circuit level, behavioral and environment-social) that contribute to individual variability - again at multiple levels - especially looking at those patients with more rapid or slower disease progression and 3) account for differences among cartoon models of the disease and the actual disease and 4) develop a dynamic understanding of disease progression and what deviations occur and why.

I am not sure why David K disparages the use of the term brain state -- as I could see that there are constellations of factors at multiple levels that lead to a "healthy" functioning state adaptable to the

JSMF | SANTA FE INSTITUTE

COMPLEX TIME

AGING, ADAPTATION & THE ARROW OF TIME

context. ANs this space could be quite large. One could then imagine vulnerabilities or insults that could push the brain in ways that result in a "state" change such that the brain is now dysfunctional in a life context or loses adaptability. And one can further imagine a brain getting itself trapped in a part of brain space where it is hard to see that any perturbation (treatment) or slow recover processes would allow for recovery. It is probably not a coincidence that the numbers of individuals with severe brain disorders are about what one would expect from being in the very tails of a distribution. The individuals who are 2-3 standard deviations out are those for whom treatments could work -- but what moves someone and what keeps them?

We need a framework that gives us some deep multi-level understanding so that we can better access if tweaking X really does impact Z or is the tweak in X actually resulting in an adaptive response in Y that then impacts Z (or maybe stabilizes Z so it does not change or becomes resistant to the X perturbations).

In aging I believe we need a better understanding of what happens to adaptive dynamic systems over time ("as the age" or over life span) so we know whether the changes we see are nothing more than what we should expect and maybe it only seems maladaptive because the environment changes or what we now expect our systems to do at different ages has changed. How do we keep our brains adaptive and responsive -- to continue to explore rather than exploit? This is a different challenge than diseases.

Today I was very struck by our inability to work across levels or to even identify what level is meaningful for what we care about -- and what I care about is using neuroscience and complex systems to advance our understanding of and care for individuals with brain disorders - particularly disorders with no identifiable anatomical lesions. 25 years ago, I initiated a program supporting neurorehab research on the premise that information learned about brain-function relationships should be useful in delineating what is and is not possible for recovery.

We have to understand the dual nature of individual differences 1) the many to one mapping -- there may be lots of ways for us to use of our brains to live adaptably in the world and yet - there seem to be a small number of stereotypical ways that brains break.

Mental health probably offers us the biggest challenges. If we could make a difference there -- even re-framing the way we currently think about these disorders - I think this would be a HUGE contribution.

Could it be that mesoscale dysfunctions -- depression, schizophrenia could benefit from mesoscale interventions -- perhaps all the lower levels changes we come to catalog will then come along for the ride,

For aging -- in pathology -- neurodegen -- treatments might require both a perturbation and a stabilization?

JSMF | SANTA FE INSTITUTE

COMPLEX TIME

AGING, ADAPTATION & THE ARROW OF TIME

Roberto Cabeza (Duke University)

My brief presentation (not a formal talk) focused on the concept of “compensation” in cognitive neuroscience of aging and dementia, and the difficulties of interpreting patterns of changes of brain activity or connectivity as compensatory. I emphasized the need of linking these changes both to a deficit and to enhanced behavior, and the importance that the latter link is established at the intra-individual rather than the inter-individual level.

The meeting was extremely interesting, particularly because it allowed exchanges between researchers with very different perspectives, which doesn't typically interact in standard scientific meetings. I found particularly exciting the idea of generating a theory of how the brain brakes that is not limited to one particular level of neuroscience analysis (e.g., molecular, cellular, systems) or one particular disorder or pathology.

The meeting reminded me of a conference I helped organized in Montreal in 2017, in which the goal was to clarify terminology (such as the term "compensation") rather than just presenting new data. As in this meeting, we also worked with a small group of researchers, without an audience, focusing on thinking rather than on just presenting and seeing new data.

Jack Gallant (UC Berkeley)

I have three brief comments:

- (1) Regarding brain explanations. I agree with others that we should seek to explain behavior, at a fine-grained level, in terms of measurable brain functions. However, it is important to acknowledge that broad analogies involving architectures or cost functions will NOT do this. Those kinds of findings are interesting and possibly necessary, but not sufficient for explaining behavior except in the broadest strokes. What is needed is rich mathematical models that link brain measurements to behavior. When such models are available they can be translated into whatever expressive system is most useful for the purpose.
- (2) Regarding brain measurement. Neuroscience is currently strongly measurement-limited. We have a wide variety of tools, but each tool is limited in spatial resolution, temporal resolution or coverage, and most tools cannot be used in humans. Given this, the best that we can do is to use our measurements as efficiently as we can given our modeling/prediction goals. In the end, all brain measurements are merely different views of the same system, so they will all be correlated with one another to some extent and in the end they should all converge on the same explanation.
- (3) Regarding brain dynamics. The brain is a spatially distributed nonlinear dynamical system. To understand such a system requires that we recover the whole trajectory of the system through space-time. However, as noted above we are measurement-limited. We can recover the spatial marginal alone (e.g., in fMRI) , or the temporal marginal alone (e.g., in EEG), but we can't recover both simultaneously (except in very reduced systems or in very special local cases). The fact that we cannot recover the space-time trajectory of the system inevitably limits the provisional explanations and models that we can construct; it limits how well one can answer different kinds of questions; and it limits the usefulness of dynamical tools for analyzing and modeling our data today.

JSMF | SANTA FE INSTITUTE

COMPLEX TIME

AGING, ADAPTATION & THE ARROW OF TIME

Russ Poldrack (Stanford University)

I thought that the discussions around the nature of distributed versus local function (arising from Caterina's talk) were really interesting and pointed to the way that our field uses these concepts very loosely. This dovetails, I think, with the issue that I raised about the disconnect between computational neuroscience and network neuroscience approaches.

I can't say that my perspective changed, but the way that I think about how to express some of the ideas was definitely changed. In particular it was really useful to talk through the ways that different ontologies might be useful for different purposes.

Paul Garcia (Columbia University)

Brains fail - like bridges, like suspenders, and like relationships. The brain constantly operates at a point of criticality - at least for our optimal cognitive state. Adaptive systems with connectivity often have cascading failures. Can the brain heal itself? Is this robustness? Can that self-organization go wrong, or be improperly applied? It's not a bug - it's a feature. I am reminded of the "swiss-cheese model" familiar in root cause analysis. Multiple failures must be serially associated

Why do we care about a diagnosis? What is meant by a "proper" diagnosis? Does "diagnosis" imply stationarity. Must we have a tight mechanism to have a diagnosis? Or perhaps simply a cluster of symptoms. Or a basis in which to guide therapy I approach my patients based on what is the next thing I am going to do. Sledgehammer solutions. - may be best. Borrowing from vaccines, can we look at treatments as "learning". Pain is an example for a top-down approach to disease. Much like traditional Chinese medicine.

Richard's buzzing that described experimental determination of the boson and how it related to inconsistencies in Alzheimer's Disease was simultaneously the most confusing and the most entertaining part of the conference.

Randy McIntosh (University of Toronto)

The take home message for me was that the principles of how the brain is designed that makes it unique and similar to other complex systems. Some the terminology/features in complex systems can, and should, be applied to the brain, but how these features are realized may be unique to the brain. There are many methods used in empirical neuroscience that can provide a springboard to this, such as graph theory metrics, coherence measures, etc, but these should be conceived in the complex systems framework to how they support features like robustness, criticality, and cascading failures.

The challenge will be to establish the common dialogue to build this bridge and the technological foundation to support it (e.g., modeling platforms for deep learning, dynamical systems that can take the empirical data as direct constraints).

JSMF | SANTA FE INSTITUTE

COMPLEX TIME

AGING, ADAPTATION & THE ARROW OF TIME

Viktor Jirsa (Aix-Marseille Université)

Excellent presentations were given highlighting the descriptive power of convolutional deep networks, also illustrating its partial explanatory power and where it fails. This pointed out some interesting ways forward that have to go beyond their current architecture, in particular taking dynamics into account. Interplay between structural and functional connectivity was highlighted. Limitations of stationary metrics were evident (functional connectivity), but nicely showed how far it can be pushed successfully in applications. Model approaches providing explanatory approaches were often too simplistic, not in terms of realism, but in terms of simplifications of concepts (brain states, behavior, as static entities). In the discussions it was evident that there is a need for a formalisation of the internal state dynamics of the brain, before perturbations can be applied to it (breaking the brain). A formal frame work for provision of and recovery from such perturbations is needed, several good attempts were provided and need to be pursued in the future, and supported by data. Need for individual predictive capacity of these frameworks was highlighted rightfully.

Richard Frackowiak (Ecole Polytechnique Federale de Lausanne)

Presentation highlight was about AI techniques, didactic, informative and comprehensible - thanks Nikolaus Kriegskorte.

There was tension between model and data led approaches.

I had a relatively stable view of the methods by which functional and structural imaging map to anatomy and local function in human brains. Those views were not shared, which meant a rethink is required. I remain unconvinced by what the resting state can inform us about mapping function and structure

Ehren Newman (Indiana Univ.)

Through many of the presentations, I found that there was a tension in identifying the right level or construct by which to think of brain function. In turn, the question was then, what does that mean about how the brain breaks.

The hypotheses that emerged for me was that the healthy / young brain is flexible, with a plurality of ways it can generate an otherwise apparently singular behavior. As such, the loss of a single way is absorbable by the system without apparent change to behavior. To some extent, this seems to be driven by the prior that disfunction or breakage must be defined as a change in behavior. Given this framework, I would argue that 'the way the brain breaks' is that it simply stops being plastic. It ceases to evolve and adapt with the environment, and eventually this creates 'breaks' or apparently inappropriate behavior.

An implication of this is that understanding how the brain breaks must take into consideration change in the of environment, not simply the brain. That is - an environment with a fixed set of statistical structure would be unlikely to ever reveal that a brain is broken.

Artemy Kolchinsky (SFI)

Two scattered thoughts after the meeting:

1. I very much enjoyed Jacopo's brief summary of noise-driven critical transitions & critical slowing down (that is, noise-driven escape across a barrier, from one metastable state to another). Jacopo also described a view of aging as the gradual lowering of the barrier, which results in a gradual increase of the probability of crossing the barrier. To me, it is the only real contender for a universal theory of aging and critical transitions in complex systems that is mathematical and predictive. Unfortunately, it is not at all clear that this theory works well for brain aging or breakdown. In particular, it is not clear that its predictions (about increased variance and/or increase autocorrelation timescales) is what actually occurs in brain aging / decline. Also, it describes aging as an increase in the instantaneous probability of break down (total loss of function). This is quite different from how many people see aging, as a gradual decline of current function. It seems important to make this distinction.
2. I very much enjoyed Niko's talk, showing that correlations uncovered by deep nets are correlated with the representations that are used in the human visual system. However, I am sympathetic to the critique made by Ross, that these correlations do not really provide us with a theory of how e.g. the human visual system functions. Rather, I see the implications of this work (as well as some other deep learning work) for our understanding of cognition to be the following lesson: simple distributed architectures + learning by gradient descent (or some other very simple learning rule, like gradient) can be shockingly effective. In this sense, these theories are similar to previous "emergentist" theories, such as Darwin's or Adam Smith's, in which a simple iterated algorithm produces amazing outcomes... but while the algorithm is easy to understand, the outcomes can be incredibly intricate and complex. This suggests that, similarly to how evolutionary biologists must study environments and niches to understand adaptations uncovered by natural selection, we might have to study the structure of natural environments to understand the representations and mechanisms uncovered by simple learning rules.