***Précis: Connectionism, Neural Networks, and Computational Neuroscience ongoing research project.* Morgan Bryant, April 2020. For my website, < taoketao.github.io > .**

In the following, I present a ongoing research project that has formulated over the last several years. I made this document for a few reasons: (1) I wish to present my research to other scholars, streamline resuming this research line in the future (by myself or others), and to write down my research objectives sufficient for a project proposal. To do so, I outline the problem and it's history and current state, provide the necessary preliminary resources and references, and sketch of the reasoning and experiments behind the innovation. (2) I wished to collect my thoughts and debut my ideas. Until now, this has equivocally been a 'pet' project. (3) Cheekily: I can establish priority and accreditation via public post, should that be necessary; and I have a potential Bryant's Last Theorem 'notebook'. Associated with this file is an annotated (and less edited) supplementary version. This précis is part of a series of précis I am composing for the same above reasons, about other projects.

Artificial neural networks are not biologically plausible. The presented theory analyzes the generalizability issues of contemporary neural networks that keep them from being sufficient explanations for human intelligence, and duly presents an alteration to the fundamental neural network model to make them substantially both more effective technologies and more biologically plausible.

To the reader: this document assumes you know: the current state of knowledge about neural networks, deep learning, and general and theoretical machine learning through the 2010s; computational models of cognition; basic neural science and cognitive science; and a undergraduate-math-major's amount of math. Additionally I will discuss the material from the citations below and from my other projects which are on my website also.

Fundamentally, this work upholds the connectionist hypothesis. As such, a broad primary interest of this work is to demonstrate that the brain is primarily a neural network and all our human capacities can be explained (and recreated) using just neural

This work's motivations:

- This work is most immediately initiated by the observation that artificial neural networks underperform at being sufficient biological models for many human intelligence tasks. This work proposes that the shortcomings all maintain a single characteristic, systemic misunderstanding as to why 'sigma' [see citation *] or 'vanilla' neural networks (VNNs) seem to fail at human-level generalization. See notes for the sketch of the argument.
- The primary necessary modification to the VNN concerning the above generalization issue reduces to a need for tensor products in the network (aka, high-order network). Paul Smolensky has already asserted this point convincingly throughout his works in the 1990s. However, basic product networks fail in a way that's 'dual' to sum networks, to which point is the primary contribution of this work.
- There is a demand for theory about deep networks. There is no consensus as to why neural networks have their distinctive shortcomings, nor how to investigate these shortcomings.
- But perhaps even if the theory is unsound, we could still engineer NN systems that suffice at the desired tasks. However, no such system has been built. Such a continued incapacity (see notes) discredits the connectionist hypothesis.

So, this work aims to explain and justify the important, pervasive yet unifying shortcoming of vanilla networks; codify the problem and present the RHO solution; and present either a theoretical basis for RHO and/or demonstrate via experiments the accuracy of the RHO network. I've recently dubbed this project RHO, since ρ is between σ and π.

A sketch of the RHO theory:

VNNs lack generalization but product networks (of any variety) lack the VNN's core advantages of arbitrary pattern learning and model complexity (loosely speaking). Smolensky (1990s) accurately identified this problem, and I concur with his conclusions but not his remedies. A solution he did consider would be to use networks that have units that interpolate between sums and products. However it is not straightforward to design such an interpolative model that does not reduce to either VNNs or product networks and their associated problems. The obvious first-line solutions typical to contemporary neural network work don't work either: For one, manually defining the interpolation suffers from the inflexible, arbitrary (read: biologically implausible, non-connectionist, unintelligent) determination of the hyperparameter, the interpolation function. But also, trying to numerically learn the interpolation with parameters actually reduces to a product network (but even less efficiently) since parameterization is a functional itself (see notes @ 'functionals'). The power of deep VNNs – their deep layered structure – can provide the intuitive basis for an elegant resolution: manage the product-sum duality with function iteration. We seek a special kind of function ϕ for which [redacted]. [Redacted as intellectual property of sorts. Can be recovered from the original documents with ease, which will be committed to private repositories contemporaneously.] Such a neural network unit naturally leads yields the advantage that product networks have concerning generalization while reaping the advantages of the artificial neural network model – relative to orders of magnitude, these characteristics are consistent with biological minds to a novel degree. [redacted] and [redacted] identify this function ϕ as a [redacted]. I'm not sure if they understand the implications as I understand it. [That's homework for me.] This network would be relatively simple to implement, and importantly, would be subject to the usual learning algorithms of contemporary networks.

On the theory side, there are a number of preliminary items that need to be formalized before presenting this work's RHO proposal. Namely, show the unified generalization issue, possibly construct the ResNet-inspired study framework, demonstrate the need for interpolation, demonstrate the inefficiency of the other solutions via development of some expressibility-per-unit metric. Also, see notes for other ones I'm too lazy to transcribe today. Also, it would be helpful to identify the neurological premise/precident as to what they do that VNNs fail to replicate.

Computational experiments would follow the usual sufficiency goals popular in connectionism and deep learning research. This model's efficacy could be evaluated on Andrew Lampinen's Set game setting and dataset (which indeed is the first and only such setting/dataset I've ever seen that precisely controls for the factors concerned by this model). If such a proof-of-concept experiment is successful, then the different possible next experiments & applications grow massively. One subsequent task would be to train such an RHO system on a machine translation task, qualitatively comparing the generalizations that humans and machines demonstrate at different stages of learning (…where 'machines' refers to the RHO model as well as alternate-to-hypothesis species such as general LSTMs, Google's transfer-via-bottleneck machine translation architecture). Should that also be successful, RL game playing and decision making (such as navigational planning: see ***Pathfinder***) would be a reasonable viable venture task that would indeed push the world's research frontier.

Prior works and other citations: P. Smolensky's (1990-1994) computations and philosophy and the derivative problem he also was stopped by; McClelland and Rumelhart's Sigma-Pi networks (1986) [tag for *]; that one paper from ~2016 that simply implements a tensor product; I'm 80% sure someone (Rumelhart Hinton and Williams 1986 "Learning representations…" perhaps?) innovated in the late 80s/early 90s that RNNs have ideal theoretical properties: not sure if high-order networks was the topic; possibly from pre-80s perceptron or Boltzmann machine times. RAM (Mnih+, 2014) as an exemplar attempt to reinvestigate how a network can achieve human abilities; the usage and prevalence of High

Order networks (me, 2017); ModIAN (me, 2017) and the derivative problem; Pathfinder (me, 2017) as a project with convergent goals; Lampinen's 2019 train-to-adapt(?), self-mapping model and Set game task. Distinguish this model from other related models, such as LSTMs/GRUs, hyperparameter-searching/model-initialization networks, metalearning, learn-to-learn models (@ Quoc Le?) and (oh gawd) connectomics, brain mapping, Yamins-style data science. DRAW (arxiv.org/abs/1502.04623). Atari DeepRL ("Playing…" Mnih 2013), DNC ("Hybrid…", Graves+ 2016). Not directly relevant, but I ought highlight the simplistic brilliance of ResNet ("Deep Residual…", He+, 2016) and why its skip-layer connections are a superb way to build more intelligent (ie, more information-trainable-efficient) systems: the theoretical 'skip' paradigm can unifyingly be [iso?]morphed [I need to take algebra…] into many other kinds of modifications such as RNN, RL, Hopfields, etc. LSTMs (Hochreiter), and the other usual deep learning folks (I. Sutskever, F. Pineda, T. Sejnowski, Y. LeCun, …). Gotta imagine I'll need to cite MCMC somewhere. [redacted].

This document:
- version 1-0, .docx file: few-page outline limit on: précis high-order neural networks and the interpolation between elementary operations with iterated Abel functionals. See related files, version 0-0 'two-page' is the preliminary document; version 0-1 'four-page' is the 'notes' referred to in 'see notes' annotations and it is a previously-forked document with technical notes eliminated from this public-perusal-facing document.
- version 1-1, .png file: a screenshot image version so that the blacked-out redacted parts are not recoverable