24
Introduction

(1980)
Stephen Grossberg

How does a brain build a cognitive code?
Psychological Review 87:1–51

This long paper gives access to much of Stephen Grossberg's work. It is effectively two papers. In the body of the paper, Grossberg gives qualitative descriptions of many of his ideas. Then, a series of mathematical appendices provides more precise formal descriptions.

Grossberg develops a series of detailed mechanisms in this paper, but most of them are first used in, and arise from, a consideration of error correction. Widrow and Hoff (paper 10) discuss one error correction technique in some detail, and we have seen others discussed more briefly, for example, some of the perceptron variants. Most of these techniques require the use of teachers who know the right answers. The teacher is not discussed in detail except insofar as it assumed both to exist and to have access to certain kinds of information.

One of Grossberg's key points is that the neural network must generally do error correction by itself, without outside help. This places strong constraints on the way the system *can be* wired up. The initial discussion assumes two systems in series, one communicating with the other. (Lateral geniculate projecting to visual cortex is suggested as an example.) A spatial activity pattern on the first set of cells gives rise to another spatial activity pattern on the second set of cells. Suppose that a different activity pattern on the first set gives rise to the *same* pattern on the second set of neurons as the first pattern did. Let us assume this corresponds to one kind of error. Grossberg asks how the system could know an error was made, if we cannot assume an external teacher with detailed knowledge of the correct responses.

Grossberg makes the reasonable suggestion that there are reciprocal connections between the two sets of cells, so a pattern of activity going upward provokes "learned feedback" going downward. Presumably in the past the descending system learned the output-input association. Then, since both are simultaneously present, the expected input pattern (from the association) and the actual input pattern interact. The details of their interaction are discussed in considerable detail, along with general principles for nervous system organization.

Many supervised error correcting techniques described in this collection (Widrow and Hoff, paper 10; Rumelhart, Hinton, and Williams, paper 41) compute difference signals between the desired and actual outputs. Grossberg computes the sum. Effective use of the sum in an error correcting system requires some assumptions about the dynamics of the network. One mechanism suggested involves network dynamics that predict the existence of what Grossberg calls a "Quenching Threshold" (QT). These nonlinear network mechanisms were suggested in some of Grossberg's earlier papers as good ways for biological information processing systems to work (see appendix C). They have the effect of suppressing small signals and enhancing large signals.

We can then use the nonlinear dynamics to test for match or mismatch of the input and the learned feedback from higher levels. If the learned feedback signal and the input pattern match, then the sum is larger than the input pattern, and has the same pattern; that is, there has been enhancement because expectation matches input pattern. If there is a mismatch between the learned feedback and the input pattern, the mismatch will lead to a more uniform summed signal with lower peak values. The dynamics of the first set of neurons that is forming the sum can be tuned so it will suppress the activity if the Quenching Threshold is set correctly. Grossberg makes the point that a noise suppression system suggested for other reasons, effective information processing in noisy environments, can be used for another purpose, such as here, checking for matches or mismatches in a system that is trying to code information about the input patterns.

The rest of the theoretical discussion considers extensions of these ideas. Grossberg suggests that recurrent feedback can maintain activity in a set of cells if necessary, even if lower level inputs are shut off due to mismatches or the disappearance of the input stimuli. Recurrent feedback allows the maintenance of patterns in a "Short Term Memory" (STM). Recurrent feedback is provided by "shunting on-center, off-surround" networks (nonlinear lateral inhibition). This particular center-surround anatomy has been a consistent theme in Grossberg's work for many years, and he has shown it to have a number of interesting and useful properties, some of which are discussed in this paper in appendix D.

Given a recurrent network capable of forming a persistent short term memory pattern, Grossberg discusses in section 15 what might happen if the feedback and input pattern match. Then, strong signals reinforce each other, giving rise to a dynamical state that Grossberg calls an *adaptive resonance*. These long lasting states become the essential elements of the cognitive apparatus. These resonant states give rise to what Grossberg calls "cognitive codes," which represent states of the system and which can be provoked by appropriate inputs.

It is worth mentioning that long lasting or stable states, generated by recurrent interactions either involving learning or a favorable anatomy, are used in many neural networks in this collection. Systems differ considerably in the details of the feedback mechanisms, but the feedback dynamics often give rise to similar behavior. Identifying stable or long lasting states as key parts of mental information processing is natural in network theory and may be a useful insight into the workings of the brain.

Grossberg spends most of the rest of the paper applying the mechanisms he described to a number of different systems in the brain. These sections contain some intriguing speculations about brain organization.

It is typical of Grossberg's work to bring together in one theory several large scale organizing principles, their effects, and their interactions. His belief is that the same set of techniques for network computation are reused over and over in different contexts. This leads to applications of his models to areas that initially seem far afield from the initial domain of application. However, if the nervous system shows any overall organizing principles and is not just a bundle of ad hoc techniques, then one should be able to see similar principles at work in many places.