

Z-Score Training, Combinatorics, and Phase Transitions

Lincoln Stoller, Ph.D.

LS@tengerresearch.com

— DRAFT v1.3 —

Abstract

Live Z-score neurofeedback training introduces operant conditioning based on simultaneously meeting multiple quantitative EEG-based conditions related to power, dispersion (amplitude asymmetry), coherence, and phase¹. While it's possible to train only one or two criteria the greater strength of the approach lies in being able to train tens, hundreds, or even thousands of criteria simultaneously. In taking multi-threshold training to a quantitatively new level Z-score training suggests we examine the space of the brain's electrical activity. This article develops a picture of this space with the object of coming to a better understanding of how we learn, and how we might help the brain adjust itself.

Introduction

Z-score training asks a trainee to configure their brain's activity in a way that can be measured with reference to a normalized database containing averages of the quantitative EEGs of a large population. This does not mean that one is necessarily training "to the average," although that is an option.

In this article we focus on the 4-channel "Percent Z-Score OK" protocol, developed by Tom Collura and others² and abbreviated as PZOK, which compares a trainee's EEG real-time absolute power, relative power, amplitude asymmetry, coherence, and phase at four sites in each of eight frequency bands (delta, theta, alpha, low beta, medium beta, high beta, total beta, and gamma) to population averages. These measurements generate 248 simultaneous criteria.

The clinician chooses the four sites at which the EEG is measured, the degree to which the trainee is asked to conform their EEG signal to that of the norm — which is to say the range about each population value to which the trainee is to configure their EEG signal — and the number of requested conditions (out of the total 248) that need to be met in order for the trainee to receive positive reinforcement.

Other Z-score training protocols can be employed. You can ask the trainee to place themselves in any relation to any particular values in the normed database, either conforming to them or deviating from them. The PZOK protocol displays interesting features and allows insight into, and improvements of this and other protocols.

Assumptions

There are a number of assumptions inherent in using population norms as benchmarks for individual performance, and we will not review them here.³ We

introduce new assumptions that support a simple analysis and provide insights into the general nature of multi-variable conditioning.

The first assumption regards the number of independent Z-score parameters. Some of the 248 Z-score parameters are inter-dependent. Some are manifestly dependent, such as the Total Beta absolute and relative amplitudes that include in their computation the Low, Medium, and High Beta absolute and relative amplitudes. Similarly the Relative Amplitude Ratios, formed from the real-time Relative Amplitude and the average population relative amplitudes, provide another measure of the relative amplitudes contained in other Z-score parameters. The use of dependent values as separate variables distorts the presumption of equal reward weightings for the different parameters of the PZOK protocol.

We will consider the fully independent variables to be the power (relative or absolute, but not both), at 7 bands at each of 4 sites (Delta, Theta, Alpha, Low Beta, Middle Beta, High Beta, and Gamma, but not Total Beta), and the asymmetry, coherence, and phase at 7 bands at each of 6 pairs of sites. This results in 154 independent variables. We are still describing the same Z-score training, we're just assuming it to be represented by 154 and not 248 independent variables.

The second assumption I make is that the trainee is equally able to meet the reward criteria of the various parameters. That is, that the trainee can just as easily raise their alpha amplitude by one standard deviation as their theta amplitude by one standard deviation, or their alpha amplitude at one site by a certain percent of a standard deviation versus their alpha amplitude at another site by the same percent of the standard deviation measured at the other site. This is most likely a false assumption, but it will not impact our general conclusions.

The third assumption we make is that Z-score feedback training works purely through operant conditioning, an assumption common to other forms of neurofeedback training. This assumption has not been well tested and is not true in both the area of LENS protocols, which employs exactly the same hardware and software but different software algorithms, and in protocols that involve active entrainment, which are numerous.

These assumptions enable us to discuss important issues that would otherwise escape generalization and go unexplored.

Multi-dimensional Training Spaces

A one-dimensional training space employs a single criterion. The trainee is asked to modify a single variable and feedback informs him or her of their progress. There are three basic types of feedback: positive reward, negative inhibit, and no feedback. Feedback can also reflect related conditions, such as rewarding rate of progress.

To succeed in a situation of this kind the trainee needs to control the parameter in question, correlate the feedback with their internal changes, and remember the results of a series of actions. We take these powers for granted, but in the case of multi-dimensional feedback these requirements become complex and the trainees' abilities should be measured rather than assumed.

In multi-dimensional conditioning the trainee is asked to satisfy multiple criteria simultaneously or in sequence. With many dimensions involved there are more directions in which the trainee can configure their brain state. This leads to new questions concerning how the trainee can distinguish in which direction they are changing, how they identify those aspects of change that are being rewarded, and how they assemble, or move toward, a multi-parameter goal.

The number of different ways the trainee can configure their EEG grows geometrically with the number of conditions they're asked to meet. In a one-parameter space the trainee responds to the three conditions regarding their proximity to the desired action or condition: getting closer, getting further away, or not changing. For brevity we'll describe these relative to the goal as positive, negative and neutral.

Based on these three criteria (positive, negative, and neutral) a two-parameter space (a description in which there are two parameters that can each be positive, negative, or neutral) supports to $3^2=9$ descriptions of the trainee's change of state. These nine descriptions include all the combinations of positive, negative, or neutral change in one parameter combined with the three possible changes in the second parameter. For example, one's change of state could be positive in the first and second parameters, positive in the first and neutral in the second, neutral in the first and positive in the second, and so on.

A three-parameter space responds to $3^3=27$ conditions, and a space with 154 independent parameters responds to 3^{154} conditions, which is a number on the order of 10 to the 73rd power. That is to say there are as many ways to mix and match 3 possible responses with 154 criteria as there are particles in the known universe (estimated at between 10^{72} and 10^{81}). These counting exercises are not very useful and some approximations are needed.

The PZOK protocol asks the trainee to configure their EEG so as to accomplish a certain degree of simultaneous success in all of the 248 criteria, where each "success" entails putting one of the measured parameters within a certain number of standard deviations of the mean population value.

Each successfully met condition in the parameter space of 248 conditions adds one point, as it were, to the trainee's instantaneous score. Whenever that score exceeds a certain limit, or a certain percent of the possible 248, a feedback reward is given. The trainee is allowed some time to interact with the feedback, something on the order of 20 to 40 minutes, and during that time it is assumed

that their brain moves toward a general conformity with all 248 criteria, in some as-yet unquantified way.

The maps in Figure 1 show a measure of relative amplitudes in five frequency regions. The maps color code the deviation of the amplitude from the population norm in units of standard deviations. This measure can be different for each frequency region as it's derived from the distribution of amplitudes of individuals at this frequency in the sample population.

In Z-score training the deviations from the norm are recomputed several times each second to give a series of “instantaneous deviations from the norm.” It is not really instantaneous, but is rather a value computed over a relatively short time period and based on an EEG sample that, in the protocol considered here, are taken over an “epoch,” or truncated time series of 1 second or longer⁴.

If the maps in Figure 1 give the results over this short time period, at an instant as it were, and the PZOK protocol is set to accept relative amplitudes within +1 or -1 standard deviations of the norm, then each of the 36 sites circled in the maps in Figure 2 are sites where the amplitudes meet this condition. These are the sites located within the green regions in the maps in Figure 1.

If the PZOK protocol is run based only on these figures 95 sites and relative amplitude scores — which it is not done in real training but which we might assume for the sake of this example — then the 36 sites with acceptable amplitudes would yield a score of 35/95, or 37%. If the acceptance threshold is set below 37%, then the instantaneous presentation in Figure 1 results in a single reward feedback event.

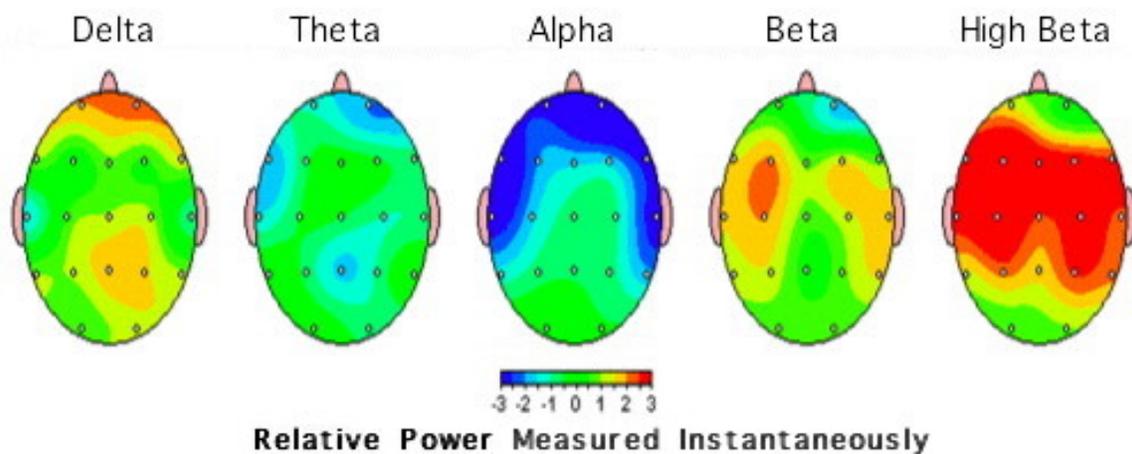


Figure 1

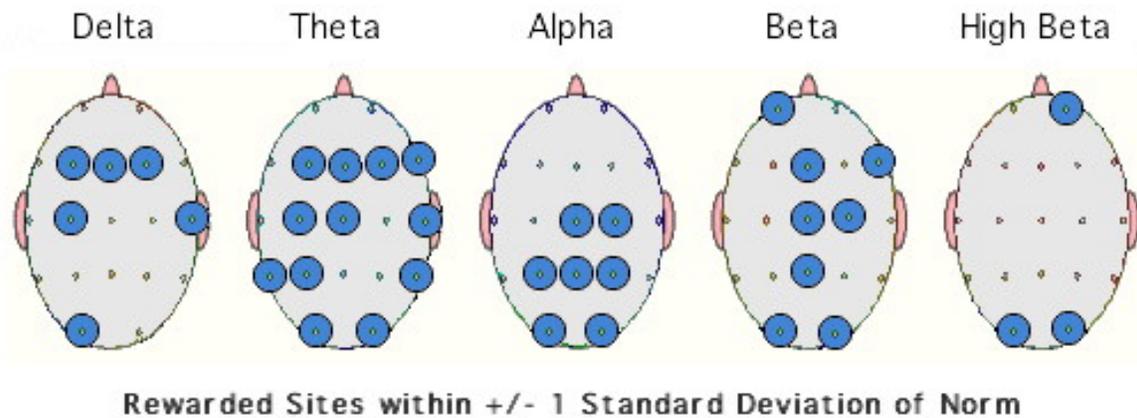


Figure 2

In the PZOK protocol the clinician adjusts the range and threshold parameters until the reward is given at a rate that the clinician judges to be sufficiently frequent. This usually obtains from requiring compliance to within 1 to 2 standard deviations of the norm for any 50% to 70% of the 248 parameters. The protocol does not use inhibitory feedback.

A 1-second epoch is the shortest epoch that can discriminate frequencies down to 1 Hz. Feedback criteria can be generated more frequently than this, up to the hardware's sampling rate normally around 256 Hz, but EEG frequency components computed more frequently than the inverse of their frequency are interdependent, which means that feedback based on their measure is not immediately responsive to changes in the trainee's EEG.

Another constraint on the feedback pertains to the rapidity of feedback events. For tone-reward based sound feedback to be effective the trainee must be able to discriminate between sounds, so feedback events are limited to something equal to or less than 4 Hz, or 4 per second. Some people may be able to discriminate the presence of absence of a reward at twice this rate, and presumably visual feedback could be delivered more frequently, but for the purposes of this analysis I'll assume a feedback rate of 4 Hz. That means the trainee is either getting and comprehending a reward, or getting and comprehending the absence of a reward, 4 times each second.

We will refer to the number of standard deviations from the norm within which the trainee must constrain each parameter as the range, and we'll refer to the percent of all parameters that must simultaneously fall within this range of the norm in order to generate a reward as the threshold. Feedback rewards can be made more frequent by widening the range, lowering the threshold, or doing both at the same time.

The Probability of Success and the Exploration of States

Imagine you are learning to dance ballet by standing in front of 248 judges each of whom will rate your performance on separate criteria. You know nothing about

ballet, have no knowledge of the criteria, and have no idea of how the criteria are assigned to the judges. To make matters more confusing, all that each judge can tell you is whether you're performing better, performing worse, or not exhibiting any change in your performance. More over, you don't know which judge is evaluating what criteria!

You make some movements and you're confronted with 248 evaluations. How could you ever learn ballet in this manner? Yet this is exactly what the brain confronts in PZOK training, and somehow the brain learns to "dance."

Learning in this context is not as hard as it might seem. The game Mastermind is a code breaking game in which one player, the defender, thinks of a key composed of 4 separate colored pegs, each peg being one of six colors. The challenger tries to figure out the key with as few guesses as possible.

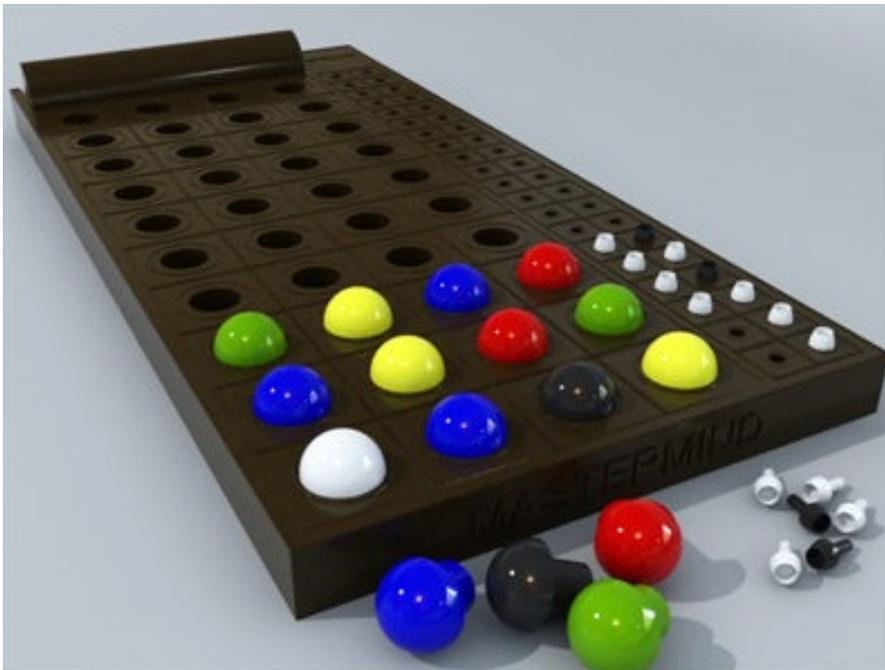


Figure 3. Mastermind, a code-breaking game published in 1970.

On each turn the challenger proposes some color combination and the defender gives feedback telling the challenger whether they've guessed any of the colors in the key, and whether they've guessed the correct location of that color in the key sequence. There are $6^4 = 1,296$ different key sequences, yet it usually takes a challenger less than 10 guesses to decrypt the key, and the game can be solved in 5 guesses or less⁵. It is helpful to consider Z-score training as a variation of Mastermind.

Considered as a game of Mastermind, the Z-Score database presents a "key" composed of 248 ranges about the norm of the EEG that the trainee is asked to meet. The trainee plays as the challenger and makes a new "guess" every 1/4 of a second. The feedback sound tells them if they've gotten sufficiently many parts

of this key correct.

A 20-minute training session admits 4,800 1/4 second feedback events, which means that 40 sessions provides the trainee with 200,000 “tests” of their EEG configuration. In this time many trainees manage to make significant progress in moving their EEG to a small range of “normal” configurations out of the total 10^{73} configurations that are possible.

Parameter Space and the Fitness Landscape

Much of physics is based on understanding how systems configure themselves in the state of all possible configurations, and we can use some of those concepts here. The notions of learning discussed here are implicit in computer simulations that have provided revelations into the interaction of quarks, the evolution of ecosystems, the mutation of viruses, and the percolation of oil through rock, to name just a few.

In dissipative systems, of which biological systems are an example, the manner in which a system changes over time is determined by way in which they explore their environment. This is true regardless of whether that environment is deterministic, chaotic, or sentient. All such systems admit a notion of “reward,” be it greater stability, greater reproductive success, or positive subjective experience.

“Fitness Landscape” is a term used to describe the full range of configurations available to a system at any point in time, or over a period of time. The landscape metaphor refers to the effect of gravity to reward downhill movement, inhibit uphill movement, and encourage movement down and along valleys. This is a picture of the landscape as a two-dimensional surface, but in most systems the landscape is multi-dimensional and the notion of “movement” is abstract. What actually takes place is that systems change their configuration to improve their survival and reduce the energy they must expend in order to persist.

The fitness landscape is easily understood in terms of balls rolling on surfaces or species evolving over time, but it is just as effective in describing a subsystem whose behavior is modified within the context of a circumscribed environment. In particular, the notion of a fitness landscape well describes the brain’s regulation of its frequency components within its own sensory, cognitive, and control environment.

Two Strategies

Two strategies for exploring parameter space appear in one form or another in all systems. The two strategies are stochastic and deterministic⁶, although they can be called mutative and adaptive, or discontinuous and incremental. Here we’ll use the terms inductive and deductive because they have direct meaning in the context of learning. These two aspects of change are distinct, universal, and usually distinguishable.

“Evolutionary histories should generally display two distinct regimes: Periods of stabilizing selection, where the population resides near a local fitness maximum, and peak shifts in which the population moves quickly from one fitness peak to another of greater height.”⁷

Dissipative physical systems, systems whose components interact with each other and exchange energy, exhibit both stochastic and deterministic properties. In species evolution we observe both random mutation and incremental change. In the game Mastermind we experience these two approaches when we begin the game by exploring sequence configurations as dissimilar as possible until we find one that we judge sufficiently promising. From this point we proceed in incremental steps. I believe the brain’s behavior during Z-score training may embody these principles.

We do not know how the brain reconfigures its EEG signature. We observe a limited range of voltage fluctuations, and we infer that due to mechanical origins, such as the slow cortical potentials, certain frequencies are coupled and their changes constrained. We also operate under the assumption that over time robust brains can reconfigure their EEG in a manner that reflects healthy neuro-regulation.

We make no reference to underlying EEG generators in this analysis. We rely instead on a combinatoric understanding of the number of states available to the brain, and a general knowledge of how systems navigate environments of this kind. We’ll look at the scale of the fitness landscape, rather than at its shape, or at the mechanisms of the brain’s passage across it.

Distinguishing Random from Incremental Changes

Inductive and deductive development strategies differ in the kind of changes they utilize. The utility of one strategy over the other depends on the character of the fitness landscape that’s being explored.

During the short 10 million year period of the Cambrian Explosion, 530 million years ago, most of the basic marine animal body forms developed as we observe them today. How these divergences originated so quickly remains in dispute, but there is little doubt that they involved the alternative paths of stochastic versus deterministic change, what we’re here referring to as inductive versus deductive strategies.

Evolution?

The reader may be wondering how an article on clinical biofeedback has wandered into a discussion of species evolution. You may be satisfied with the following short answer, but you should also appreciate its deep implications that we cannot fully explore here.

The short answer is that species evolve and individuals learn in much the same manner despite vast differences in the time spans and mechanisms involved. This revelation has resulted in breakthroughs in a variety of fields where the systems being studied can be seen as trying to optimize their interactions within complex environments. This approach directs our focus away from the often insolvable problem of deterministic prediction, and toward an exploration of the adaptive process itself.

At its deepest level this describes a feedback process that applies wherever the systems being studied share the following properties:

- 1 – They change within a well-behaved space of constrained possibilities, which we call the fitness landscape.
- 2 – The changes that occur do not qualitatively change the systems, the landscape, or the systems' ability to interact with their environment.
- 3 – The systems can discriminate the degree to which their change of state changes their fitness in the environment.
- 4 – Changes in fitness constitute a reward that impacts each system in a direct way.
- 5 – The combination of past changes and related rewards create a feedback loop that impacts systems' future changes.

Neurofeedback research meets these criteria when the frequencies we train play a consistent, constructive role in the brain's larger regulatory systems. In addition, changes in the EEG appear to correlate with changes in the quality of the brain's performance.

The drawback of this approach is that it tends to obscure the key dynamics. The inquisitive reader will be interested to learn that the key to how systems' evolve lies in the topography of the fitness landscape, and in what defines a configuration's fitness. In EEG neurofeedback these translate as the range of the brain's electrical patterns — that is the plasticity of the EEG — and an understanding of how the brain is rewarded for changes in its EEG.⁸

Inductive Strategies

Inductive strategies aim to widely explore the fitness landscape in an effort to appreciate its topography. Inductive strategies aim to develop a complete picture of the landscape's hills and valleys in order to find optimal locations or configurations. In various systems this leads to water collecting in puddles, species populating niches, and the brain's facility to move between optimal regulating states that, in the case of Z-score, might mean optimally rewarded EEG signatures.

The rate at which a system accomplishes inductive development depends on its speed and the landscape's topography. Here "speed" is not so much how quickly a system changes, but how rapidly it obtains the feedback that determines the success of each new location.

A flat landscape would be one in which each state, or location, was equally accessible. The use of the word “landscape” is just a heuristic device. In nearly all systems the landscape is composed of various dimensions, some of which are qualitatively different. In these cases the notion of distance, while problematic, is important because it allows us to compare different parts of the landscape.

A space’s topology refers to its most basic features such as holes, islands, and bridges between separated areas. If the space consists of entirely disconnected regions, then the system exploring it will be stuck on the “island” on which it first finds itself.

In the evolutionary example ocean and land environments were sufficiently disconnected to inhibit marine species from exploring dry land. In the realm of brain dynamics we might conjecture that a trainee suffering TBI will have difficulty in reaching functional states when starting from dysfunctional ones no matter what form of feedback or entrainment we provide.

Consider the inductive strategy in playing a variation of the Mastermind game. Imagine a Mastermind game in which you need to guess a series of 10 separate and unrelated 4-color sequences. Since the sequences are unrelated, determining any one 4-color sequence will not help in finding others. An inductive approach would be to test a selection of demonstrably different sequences in order to find the best initial configuration. This configuration would then provide the starting point of a path to the correct solution.

Reward is Crucial

The hallmark of inductive change is the coverage of large areas of the fitness landscape, the exploration of divergent configurations, and the development of a map of successful configurations. Here “success” is entirely defined in terms of reward although, truth be told, we really don’t understand the brain’s reward mechanism.

This notion of reward is an essential ingredient in how the brain changes. We would like to know what rewards really are, including what’s being rewarded, and how rewards correlate with changes. We say that we reward the brain for conforming to the Z-score norms using computer generated feedback. This greatly simplifies the description of the brain’s process of self-organization, and it isn’t true.

When, on occasion, our feedback encourages a dysfunctional state, the tones and colors that we present are not a full description of the incentive, or disincentive, experienced by the trainee. The true reward, if it even exists as an independent and quantifiable entity, is something experienced internally. Our computer-generated feedback is just our approximation or reflection of it. Our computer generated “rewards” are more a means for communication than for incentive.

A noted benefit of multi-channel Z-score protocols is that they allow the brain to decide for itself those aspects of the normalized state to which it will conform. Presumably this means that the brain will use its discretion in conforming to the suggested norm⁹. In making its own decisions regarding how to configure itself, the brain is using the computer's feedback more as a landmark rather than a goal. And in this case who is generating the reward, and what is being rewarded?

This question is important because an understanding of the reward is central to an understanding of how a system navigates its fitness landscape. In the case of inductive learning we assume that the system's large changes in state result in a reward that can be compared to the reward it has received from its previous state. Without this feedback the system cannot understand the fitness landscape, and if it cannot understand the landscape it will wander aimlessly.

Regardless of the true nature of the feedback it is necessary that the brain be able to discern the reward. This is the central assumption that enables us to claim that the brain learns. Without this assumption the graphs in Figures 4, 5, and 6 would display no improvement over time.

Deductive Strategies

A deductive strategy makes incremental changes to a known configuration in search of measurable improvement. A deductive strategy only works when small changes result in small improvements; it requires continuity. Lacking continuity there is no way to generalize the nature of the changes that result in more reward from those that result in less. Without this information the brain would not be able to "deduce" how to make improvements and the deductive strategy would fail.

In an evolutionary example the development of a webbed foot provides an aquatic animal with better propulsion, and this webbing can develop incrementally as connective tissue between the toes. In the Mastermind game, changing one color will either get a higher, lower, or unchanged score. In a brain training example changing the amplitude of oscillations in a particular band, while keeping others relatively constant, will either result in a feedback reward or not.

Taking a purely deductive approach the brain would consider an EEG pattern for which it received a reward and it would proceed to independently modify each of the 248 parameters of the PZOK protocol and move to configurations that receive a greater reward.

One might object that this anthropomorphizes the brain by treating the unconscious as something with rational intent, but I know of no neuro-anatomical theory or observation to suggest that this could not be the case.

Measurable Differences

The difference between the inductive and deductive approaches lies in the speed at which they explore the fitness landscape. Inductively made changes affect all

aspects of the state — the changes are global — and the sequence of states tested are disparate and weakly correlated.

In order to determine a state's fitness its full Z-score measure must be recorded. This means that a sufficient amount of time must elapse to allow feedback at all frequencies. This assumes the brain is trying to explore both high and low frequencies, or the fast and the slow parts of the EEG spectrum. It could be that some brains ignore some portions of the fitness space. If this is the case, then as therapists we would like to know it.

When exploring its configuration landscape the rate of feedback can be no faster than what's required for the slowest frequencies in that part of the landscape. In the case of the PZOK protocol the slowest part of the landscape is the Delta band that includes frequencies of 1 to 4 Hz. This means that feedback pertaining to changes in the Delta states is provided only once per second, or thereabouts. And this means that while the inductive approach explores all EEG configurations, the rate at which it receives the feedback needed to compare one full-spectrum state to another is constrained by the slowest components of the landscape.

To the extent that the brain takes an inductive approach and the fitness landscape is uniform across frequencies, we expect to see widely dispersed EEG configurations converging to all frequency norms at the same average rate. Both this wide dispersion across the fitness landscape and the rate of convergence should be measurable using the metric described below.

In contrast, a deductive approach is manifestly local in the changes it makes to an EEG state. Deviations are made about a configuration of "known" fitness and these deviations are made sequentially in the different variables. In the Delta portion of the spectrum rewards obtain roughly once per second, but in the 25 – 30 Hz Gamma region feedback is provided more than 25 times per second.

To the extent that the brain takes a deductive approach we expect to see narrowly dispersed EEG configurations gradually converging to their norms. We expect to see Gamma frequency bands converge 20 times faster than Delta bands. This localization in the fitness landscape and the relative rates of convergence should be measurable with the EEG hardware and software that we currently employ.

Picturing the Inductive Learning Process

We can graphically represent how these learning processes might appear based on parameters that we can easily measure. The three axes in these graphs are Training Time, Frequency Band, and Distance from the Norm.

Clinical experience has shown us that deviations from the norm (i.e. success in meeting Z-score criteria) are reduced by half or more over a period of 20 to 40 training sessions. We can use these scales even though we don't know what

mode of learning lies behind the process. This gives us the graph in Figure 4 where deviations from the norm lie within a couple of standard deviations and show substantial reversion to the norm over the course of 40 training sessions.

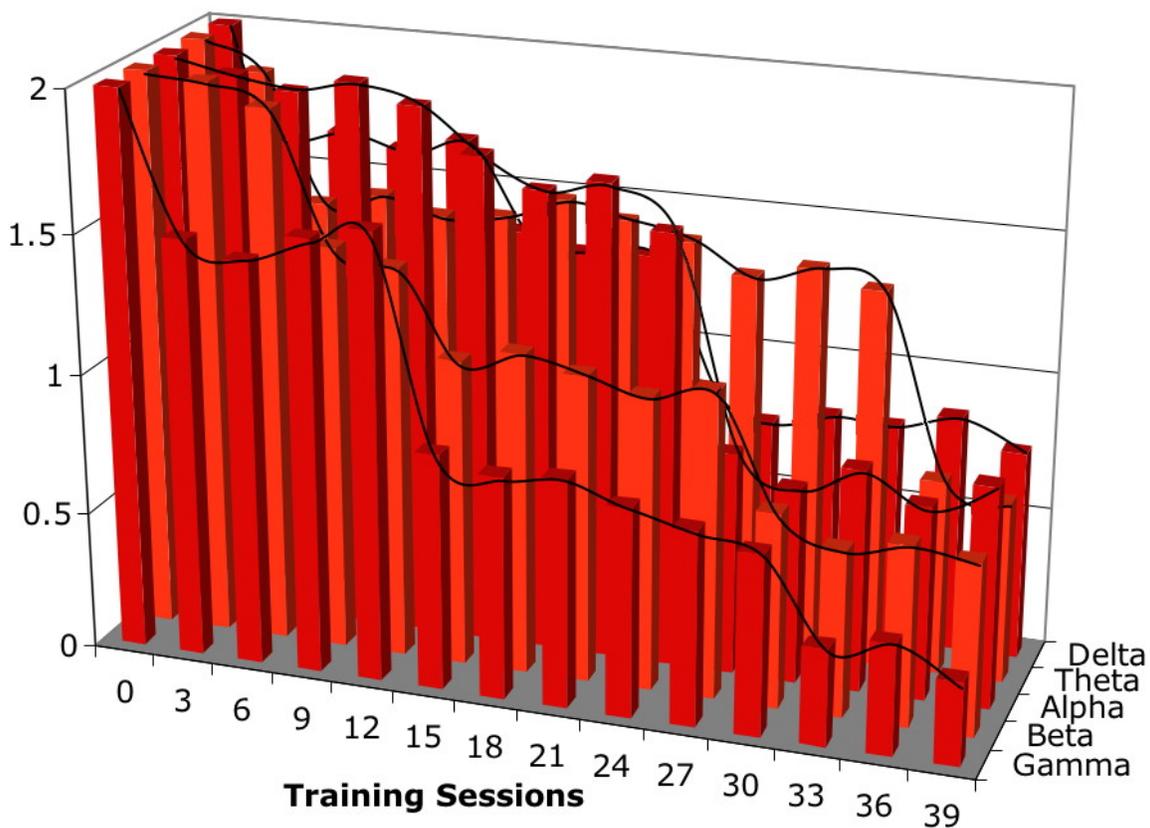


Figure 4

Figure 4 shows what we expect from an inductive learning strategy. It pictures what might obtain if only inductive learning took place, which we expect to be true occasionally if at all. This graph is only a heuristic and its details should be understood.

Convergence to the norm occurs stochastically. Large deviations in the brain's state are made in an erratic attempt to find a more parsimonious state. The changes occur independently across the frequencies. When the brain finds an improved state it explores this neighborhood of the fitness landscape. It does this because it's receiving a reward that makes this state more "comfortable."

The lines connecting the tops of each column help the reader follow the progression in each frequency band. The columns are colored red to indicate that this type of learning involves the system traversing relatively large distances of its fitness landscape. Two shades of red are employed to better distinguish columns of the different frequency bands.

The column heights all start at 2 standard deviations from the norm and this is arbitrary. One or two standard deviations from the norm are normal, but each trainee will be different. Equally arbitrary is the rate and extent of convergence to

the norm over the course of many trainings. The progression of column heights, their rates of changes, and the relative changes across frequency bands is hypothetical.

Column heights grow and shrink because some new configurations are better and others worse. When new configurations are rewarded the brain remains in that neighborhood; when they fail to be rewarded, or are otherwise inhibited, the brain moves to a different location.

The brain wanders around like a drunk on a pogo stick bouncing over a rough and hilly landscape, falling down steps and slopes, and bumping into or jumping over banks and walls.

This discontinuous improvement over time and across frequencies embodies two assumptions. First, it assumes that the brain is efficiently exploring all accessible states. This is known as ergodic hypothesis that is believed to be true for most physical systems that are not otherwise constrained.

Second, it assumes that as the brain explores this space it remembers where it was rewarded and retains these configurations while it continues to search for additional rewards. Unlike the drunk on the hillside the adaptive brain can remember where the valleys are located, and return to them.

One's ability to locate more optimal positions in the fitness landscape decreases with time in accordance with a power law: the probability of discovering an improved state decreases as some inverse power of the time. This reflects the topology of the fitness landscape and the fact that as time goes on there is increasingly less unexplored space. This power law form is a general property of inductive learning and does not depend on the reward criteria or the mechanism by which the system traverses the landscape.

Picturing the Deductive Learning Process

The graph in Figure 5 reflects our expectations of a purely deductive strategy, which is also something we do not expect to prevail at all times.

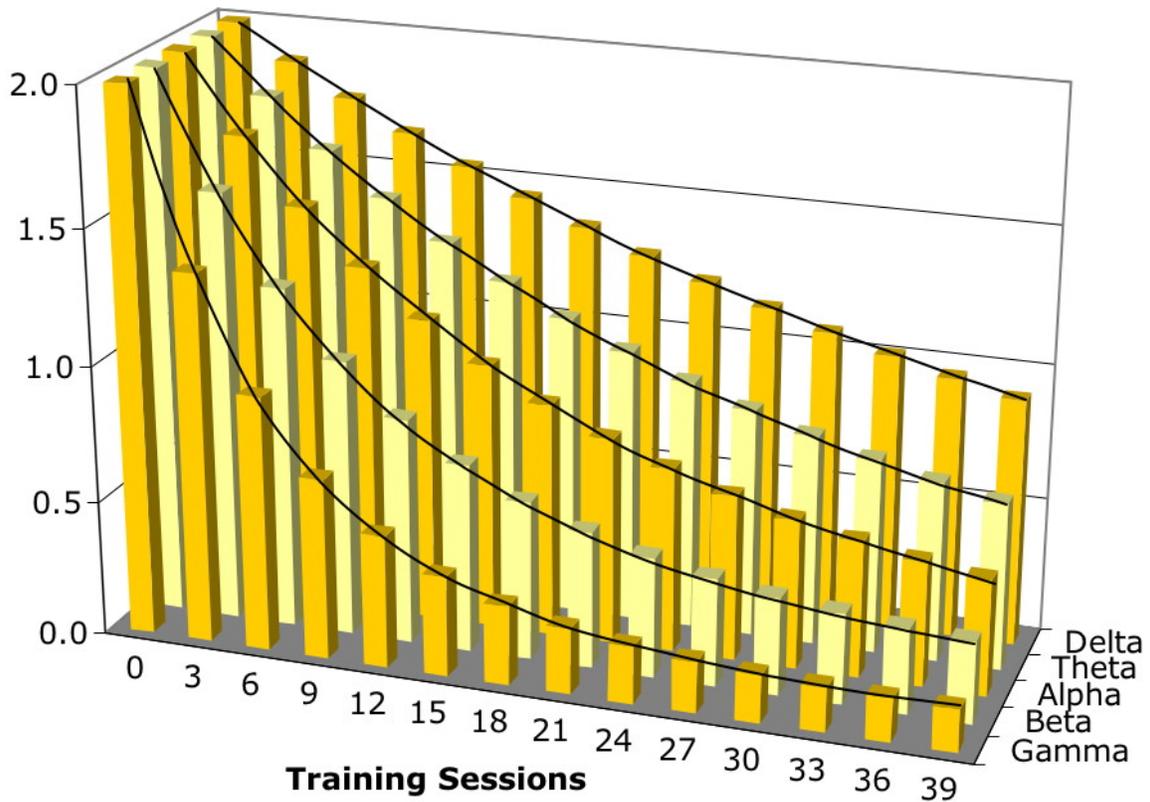


Figure 5

As in Figure 4 we've chosen to start our hypothetical trainee from a state that departs from the norm by 2 standard deviations in each frequency band. The columns are yellow to indicate that the paths of deductive learning states traverse small neighborhoods of the fitness landscape. Two shades of yellow are employed to better distinguish the columns in different frequency bands.

Learning occurs more rapidly in the Gamma frequency bands because feedback in these bands is delivered 25 times more rapidly than in the Delta band. This allows for faster movement through the higher frequency bands of the fitness landscape, resulting in faster high frequency convergence to the norm.

The most unrealistic aspect of the learning curves in Figure 5 is that they all progress monotonically toward zero. In any system that's described by a complex fitness landscape this progression, while monotonic, can get stuck and cease to show further improvement.

This stasis happens when the deductive strategy finds itself in a kind of sinkhole. When this occurs all the incremental changes are for the worse. The system remains in this state, which is the best it can find in its neighborhood. In light of this property Figure 5, which shows nothing getting stuck, only applies to a system with a smoothly curved fitness landscape, one that doesn't have any pockets (local minima) away from the target state. We would not expect Figure 5 to depict the fitness landscape of a real brain, and certainly not a dysfunctional one.

The Measure of Learning

When engaging in inductive learning the brain attempts to explore the fitness landscape widely, and it generates states that cover the landscape to a maximal degree. When engaging in deductive learning the brain makes small, incremental changes to its state in order to optimize the state and maximize the reward.

An essential property distinguishing these two strategies is the distance between EEG configurations that the brain generates as it explores the fitness landscape. This property requires a new kind of measurement.

Measuring the “length” of the path that the brain takes through the fitness landscape is problematic because the fitness landscape is parametrized by qualitatively different variables. Defining the “distance” that the brain “travels,” as it changes amplitude, coherence, and phase presents ambiguities. It is similar to defining the distance that an object travels through space-time: how does one combine extension and duration? In this case the question is how does one create a single parameter that measures a change in relative amplitude at one frequency in combination with a change in coherence between two sites measured at another frequency?

We can create a natural measure based on the scale set by the standard deviation in each parameter. This is a dimensionless scale that allows us to combine different measurements.

Using this measure we can compute the distance between two EEG states that differ from one another by some number of standard deviations in amplitude at one site and frequency, call this value x_1 , some other number of standard deviations in phase at a pair of sites at another frequency, call this x_2 , and some third number of standard deviations in coherence at a pair of sites a some frequency, call this x_3 .

Using the simple analogue of the length of a line in three dimensions, the distance D between these three states would be $(x_1^2 + x_2^2 + x_3^2)^{1/2}$. This generalizes to a distance D in any number of dimensions, 248 in the case of 4-site Z-score training, as:

$$D(248) = (x_1^2 + x_2^2 + \dots + x_{248}^2)^{1/2}$$

We don't have experimental record of these distances, but we already know all the x_1, x_2, \dots, x_{248} component values that are involved. These are the same values that are color coded in Figure 1, and whose measurements are instantaneously taken in the course of Z-score training. As a result, adding this measurement to Z-score training requires a trivial additional computation after each epoch.

We can now define $D(j,k)$ as the distance in standard deviations of the EEG state

at time j from the EEG state at time k . If we take k to be one epoch after j and add up all of these distances over a sequence of epochs, then we obtain D_{path} , which is the length of the path over which the EEG state has traveled over that period.

$$D_{\text{path}}(j,k) = D(j,j+1) + D(j+1,j+2) + \dots + D(k-1,k)$$

Inductive learning presupposes that the learner widely explores the space of possibilities. In the EEG model this translates to changes in the EEG signature that traverse a long path. Conversely, deductive learning involves exploring closely spaced states, so the EEG signature will traverse a short path. We have used a color code of yellow Figure 4 to indicate short paths, and a color code of red in Figure 5 to indicate long paths.

Combined Learning Strategies

Experience with computer simulations and the evolutionary record suggests that either one strategy or the other is pursued depending on the environment. Figure 6 shows rates of extinctions taken from the fossil record. Recognition of these periods of mass extinctions contributed to the evolutionary theory of punctuated equilibrium. We might suspect a similar alteration between stability and instability to occur in the realm of neuro-regulation.

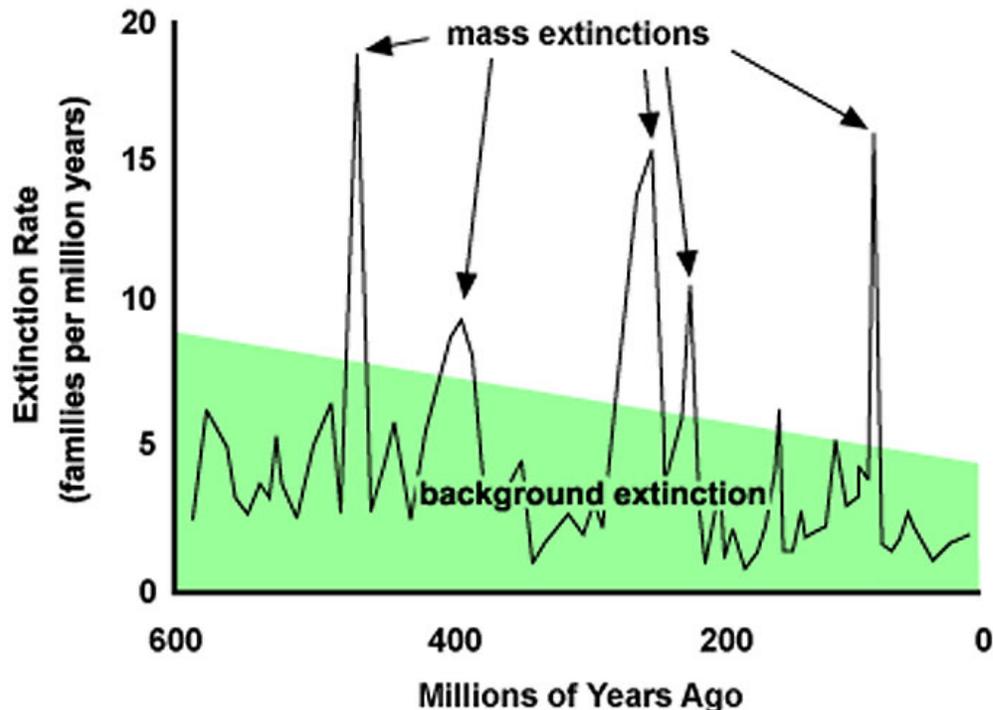


Figure 6. Rates of mass species extinctions taken from the fossil record.¹⁰

Figure 7 superimposes learning rates in five frequency bands over a full course of training. The relative length of the paths through the fitness landscape, what

we've defined as D_{path} , are shown by the color of the lines. The lines are predominately yellow as is characteristic of deductive learning. The occasional columns in which the lines are red correspond to episodes where the brain is frustrated, induced, bored, or otherwise directed to engage in inductive learning.

We expect the brain employs a combined learning strategy since a combination of strategies is employed by many systems. Deductive strategies are best when the fitness landscape is smooth. An inductive strategy always obtains when relatively large amounts of noise push a system up the hillsides that surround the landscape's stable locations. Whether this impetus is sufficient to reveal new valleys depends on how deeply stuck the system is (how high the hills are that surround the local minima), how powerful the random forces are that tend to kick it out of local equilibrium, and on what other valley exist in the landscape.

Figure 7 portrays the behavior of any dissipative system. The only thing about this picture specific to EEG training is that the frequency bands are separated, and the higher frequencies learn faster.

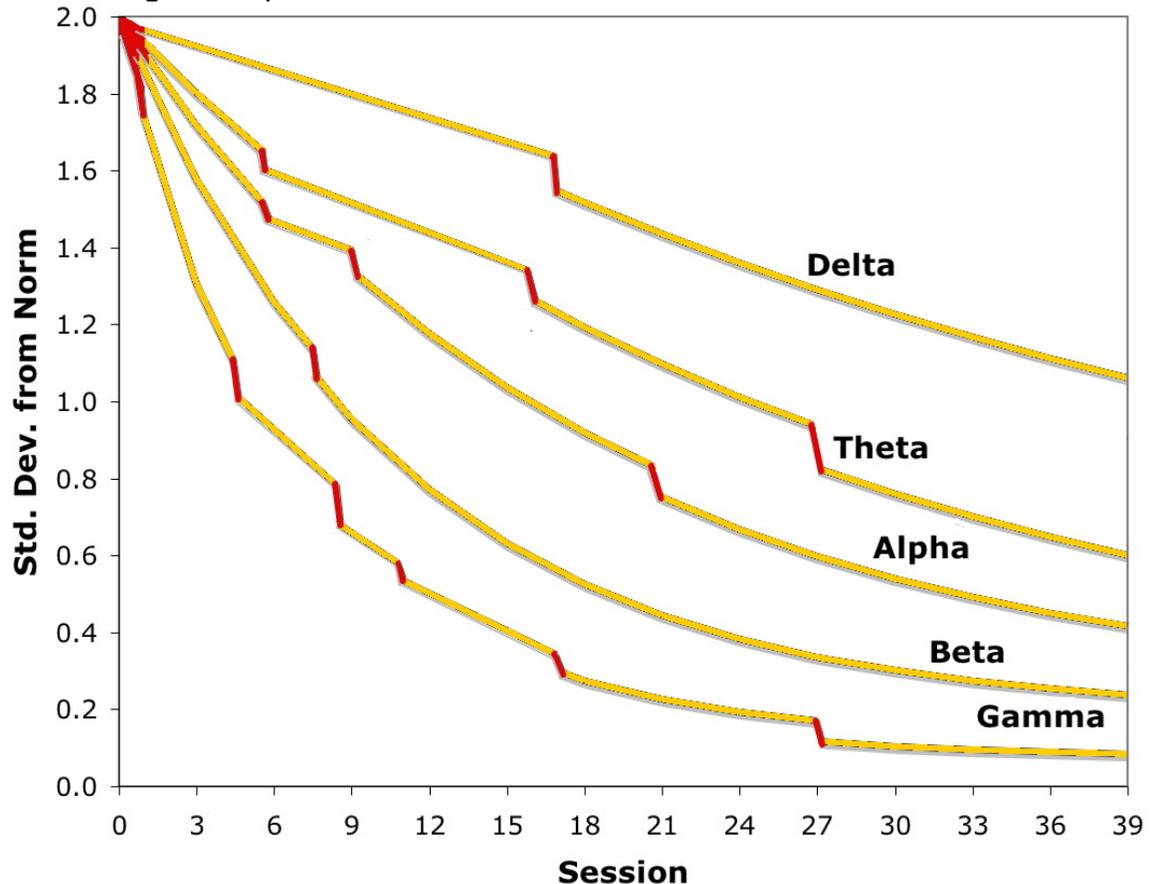


Figure 7

In Figure 7 lines have replaced the columns in the previous figures, and the column colors by line colors. The colors remind us that inductive learning explores the fitness landscape widely with resultant rapid improvements, while deductive learning tends to stay in a small neighborhood and slowly seek improvements. Charles Stark reports that patterns of this sort are observed in Z-

score training, namely gradual and abrupt changes in deviations from the norms desynchronized across the frequency bands.¹¹

Figure 7 is idealistic in its portrayal of monotonic improvement at all frequencies because, while we expect learning to alternate between sudden and incremental improvement, we expect trainees to hit thresholds that stall their progress. We are describing a process, not an outcome.

Phase Transitions

The length scale D_{path} plays the role of an “order parameter,” and the red line-segments in Figure 7, showing the sudden drops in the standard deviation from the norm, can be thought of as phase transitions in neural state.

Phase transitions provide a powerful paradigm for understanding complex systems. They play a central role in the Haken-Kelso-Bunz (HKB) neurological model formulated in 1985 to account for experimental observations on human bimanual coordination¹². Phase transitions and related concepts, including inaccessible regions, divergences, and hysteresis, are explored in developmental psychology in a 2001 paper by Jansen and Van der Maas.¹³

Phase transitions are defined with respect to the fitness landscape. The term refers to a change between two distinct states that are separated by an inhibitory barrier. In our context this means that there exist at least two stable, different, and functionally distinct EEG patterns and that do not gradually “morph” from one into the other through a continuous series of incrementally small changes. In other words, going from one phase to the other necessarily passes through inferior or “less fit” regions of the fitness landscape.

This is logically the case because if one could move from one phase to the other along a path of constant rewards, then everyone would eventually pass to the preferred phase. The inferior phase would be unstable and, while potentially long-lasting, it would not persist as a separate phase.

It’s worth exploring the meaning of a phase transition at greater depth because there are ambiguities in the phase transitions of finite systems that we will encounter when applying these notions to brain states. Some of these ambiguities arise due to the brain’s finite size, and not because properties specific to the brain.

There are no infinite systems in nature, but in physics and chemistry the mechanisms that govern a system’s state operate on a microscopic scale while the phase changes are measured on a macroscopic scale. For this reason observed phase changes for systems with astronomically many atoms are well described by the equations for infinite systems.

The following three properties describe phases in infinite systems. These properties do not hold in finite systems.

Absolute Stability: phases in noise-free systems are stable when external conditions are held constant.

Flawless finite systems obey the condition of absolute stability but the existence of a flaw that allows “leakage” from one state to another undermines stability. Such a leak provides a way around the barrier that separates one state from another.

You can think of this metaphorically as either a crack in the dike that holds the reservoir from flooding the valley, or as a stream that exits the back of the reservoir and meanders its way around the mountain to irrigate the valley. When this passage between the states is large enough we need to think of the less advantageous state as meta-stable, and not absolutely stable. In this case we apply the term “neural phase transition” with some latitude.

Kelso suggests that certain meta-stable states of neuro-motor coordination (referring to his studies of human hand movements) persist through a form of mode locking rooted somewhere in the neuro-mechanical complex. This system’s exploration of its neuromuscular fitness landscape is described by a process of phase slippage which, when it reaches a critical point, loses resonance. At this point the system behaves chaotically and widely explores its fitness landscape before locking in to another meta-stable mode¹⁴.

Relative Stability: phases in systems subject to fluctuations are stable to the extent that fluctuations are small relative to the thresholds that separate different phases.

The brain is certainly a fluctuating system and we know little about the extent, duration, or specificity of its fluctuations. That being the case, states that may be relatively stable in one person, at one time, or under certain conditions may be unstable in others. When the conditions that maintain phase separation are lost, the brain will change state. The point is that phases may or may not obtain depending upon conditions, and the nature of the phase transition may depend on the manner in which conditions change.

Some conditions will support a barrier between the states and allow us to view these states as phases. Consequently we’ll be able to speak of protocols that enhance phase transitions. Other conditions may allow the emergence of a gradual that leads from one neural state to the other such that the notion of the two states as being separate phases has no support. In this case we would be better off thinking of protocols that facilitate stalled change, rather than protocols that breakdown barriers between distinct patterns.

Recall in our previous description of evolving systems we stated that the changes that occur should not qualitatively change the system, or the landscape. In applying this description to the brain it should be recognized that the brain might

do either or both of these things. That does not vitiate our attempt to apply this model, but it does imply that the model might be limited.

Comprehensive: a pure system undergoes a complete and pervasive change when a new phase becomes advantageous.

The flakes in any given snowstorm are all about the same size, yet the transition to ice crystals admits a infinity of structures. Phase changes in finite systems are limited in extent but, as the snowflake example illustrates, the resulting structure need not be unique. The phase transition describes a process, not an outcome.

Most physical systems undergoing change manifest the phenomenon of “pinning,” which is the hindering of motion due to impurities and something similar occurs when real, finite systems undergo phase changes. These impurities, exceptions, or injuries impose something like friction, and this leads to the cessation of change when the impetus becomes weak. Impurities limit the size of the region over which the phase change occurs.

Looking at a sheet of galvanized metal, for example, you see a hotch-potch of thumb-sized regions — whose appearance is known as “the spangle” — within which the metallic crystallites have lined up. In the case of a two-dimensional surface the size of these regions is the order parameter that describes the crystallization transition. When undergoing its transition from disordered to ordered crystals this parameter grows by a factor of 100,000,000, starting on the atomic scale of a few angstroms (10^{-8} cm) and growing to a macroscopic scale on the order of 1 cm.

Order parameters diverge in the phase transitions of certain infinite systems, such systems are said to undergo “second order” phase transitions. In a crystallization transition the size of crystal would become infinite. The path length we’re using as an order parameter would also diverge if a finite change had to be made everywhere in a code consisting of an infinity of parameters. But the brain we’re describing is finite, and in a finite-sized system the order parameter gets large without diverging.

The term phase transition should be taken generally. There are nuanced distinctions between phase transitions, bifurcations, and meta-stable states that depend on a better knowledge of the fitness landscape than that given by the Z-score norms, which is all that we currently have. Without better knowledge of this landscape we don’t know exactly which description might underlie the behavior shown in Figure 7.

The positions of the red line segments in Figure 7 are only suggestive. We suspect the inductive approach is more suitable at early stages when the least is known about the fitness landscape, and we only guess that additional excursions away from equilibrium may occur, or may need to occur, in order to reach new plateaus in regulation. What might actually be observed remains an open question.

The Process of Neuronal Adaptation

“...between the hard and the easy, the regular and the irregular where the creative pulse beats.”¹⁵ — J.A. Scott Kelso

We propose that inductive and deductive strategies are employed to explore the fitness landscape. These strategies may operate in a person at different times, or in different people to differing degrees.

Some neuronal problems may stem from an inability to execute one or the other of these strategies. This suggests new feedback methods targeted at strengthening the way in which the brain learns, rather than the EEG signature that it manifests.

Viewing brain states as paths through a fitness landscape naturally leads us to wonder if there are healthy and appropriate path-finding skills, which is an extension of the EEG neurofeedback’s traditional goal of healthy and appropriate states.

Some Neurofeedback protocols may already be enhancing a trainee’s facility to move between states. In the case of Alpha-Theta training a trainee learns a certain manner of entering a Theta-dominated state over the course of 10 to 20 sessions. Is it the Alpha-Theta state that’s being trained, or the neural skills to find it? And does the benefit that accrues arise from manifesting the Alpha-Theta state itself, or from the facility in navigating that area of the fitness landscape? This may seem to be a fine distinction, but the two views have quite different implications: one is goal oriented, the other process oriented.

A path description may also provide a framework in which to better understand the LENS protocol named “Rocking the Spectrum,” developed by Nick Dogris. In this protocol a series of short bursts of low amplitude stimulation are delivered at 5 to 15 Hz offsets from that frequency at which the maximum power is recorded over a specific site on the cranium. This protocol is known to facilitate new progress in trainees who have reached an achievement plateau.¹⁶

Enhancing the Learning Process

Given that these strategies are fundamental, how might we facilitate the training process so that the brain can fully and quickly explore its fitness landscape?

Consider the inductive approach first. From playing the Mastermind game we know that more specific feedback allows for a faster determination of the quality of the guess. In the neurofeedback context this translates as providing feedback that differentiates both the aspects of the EEG’s success in meeting the norm, and the degree to which it meets the norm.

For example, different feedback elements could simultaneously provide indications of success in different areas of the EEG configuration. The sound of a piccolo could indicate achieving Gamma normalcy, an oboe for Beta normalcy, all the way down the scale to a tuba or timpani for Delta normalcy. These instruments could be voiced together as the instrumental parts, fading in and out, of a familiar symphony.

Visuals could simultaneously decompose a picture into its different color components, where each color component corresponds to achieving normalcy in parameters such as amplitude asymmetry, coherence, and phase. It would be difficult to present 248 separate feedback components in such a way that they remained consonant and distinct, but it is not impossible. We do not know the limits of the brain's discrimination, but it appears to be more discriminating than we had suspected.

Next consider the needs of the deductive approach to locating the EEG reward state. According to this strategy we begin with a state of known fitness and attempt incremental improvements.

In the inductive case, where changes are made globally, it is important to differentiate the reward according to the variable that was changed, but that is not needed for the deductive strategy. The reason it is not needed is that changes are made one at a time in an incremental strategy. The system already knows which parameter it is varying. It only needs to know if whether it's varying this parameter in the right direction.

A proportional reward is optimal for the deductive approach. A proportional reward provides one kind of feedback to a degree that varies in accordance with the improvement in the state, and another kind of feedback to indicate greater degrees of discouragement.

A simple rising or falling tone might be sufficient, but people have imperfect pitch so a better approach might be to offer tones of distinct character. For example one might offer consonant tones of rising pitch, or a sequence of consonant tones in a rising scale, to indicate greater improvements, and a dissonant tone of falling pitch, or a sequence of dissonant tones in a falling scale, to indicate a growing deviation from the norm.

This begs the question of whether some feedback could support both inductive and deductive learning strategies. I suspect not because the strategies are so different. But that does not mean they shouldn't be offered consecutively, or that the trainee couldn't switch from one to the other during a single session.

The existence of these learning modalities is conjecture, and while they are universally relevant in theory, they may not be observable. Abrupt transitions as shown in Figure 4 are apparent, but we know little about them and they may not manifest in a regular or uniform way. In some people, or after certain training, the gradual transition processes of deductive learning, shown in the schematic of

Figure 5, might prevail.

Z-score maps provide a picture of recurring patterns¹⁷ that are different from what we identify in EEG traces. Z-score's D_{path} measure and fitness landscape enables us to speak of phase transitions in neural states and other chaotic phenomena¹⁸. The utility of this model is supported by anecdotal reports¹⁹ that the feedback refinements we describe do improve PZOK training results.

Notes

¹ Robert W. Thatcher, "Z-Score EEG Biofeedback: Conceptual Foundations," Neuroconnections Newsletter, April 2008, p.9.

² Thomas F. Collura, "Whole-Head Normalization Using Live Z-Scores for Connectivity Training," Neuroconnections Newsletter, April 2008, p.12; and Thomas Collura, Joseph Guan, Jeffrey Tarrant, John Bailey, Fred Starr, "EEG Biofeedback Case Studies Using Live Z-Score Training and a Normative Database," Journal of Neurotherapy, 14:22-46, 2010.

³ Robert W. Thatcher, Carl J. Biver, Duane M. North, "Z Tunes: Technical and Conceptual Foundations," published by Applied Neuroscience, Inc. 2002, at www.appliedneuroscience.com

⁴ The duration of the epoch determines the slowest wavelength component that can be extracted from the signal, as discussed in: Lincoln Stoller, "Making Sense of Infra-Low Frequency Neurofeedback," Neuroconnections Newsletter, Summer 2010, p. 19.

⁵ Donald Knuth, "The Computer as a Master Mind", Journal of Recreational Mathematics, 1976-77, p.1-6, Vol. 9.

⁶ I. M. Rouzine, A. Rodrigo, and J. M. Coffin, "Transition between Stochastic Evolution and Deterministic Evolution in the Presence of Selection: General Theory and Application to Virology," Microbiology and Molecular Biology Reviews, March 2001, p. 151-185, Vol. 65, No. 1 at mmbbr.asm.org/cgi/content/full/65/1/151

⁷ Kavita Jain and Joachim Krug, "Adaptation in Simple and Complex Fitness Landscapes," in Structural Approaches to Sequence Evolution, Eds. U. Bastolla, M. Porto, H.E. Roman, M. Vendruscolo, Springer Berlin, 2007, pp. 299-340, at xxx.lanl.gov/abs/q-bio.PE/0508008

⁸ An important phenomena widely noted in the study of systems of this kind is that the mechanisms of change, while important for setting the time and length scales over which change occur, generally do not determine the way a system changes. For example, the patterns of social change are similar regardless of

whether wars are between tribes battering each other with clubs, or nations annihilating each other with bombs.

⁹ This is discussed in the context of individual phenotypes by Kerson, Gunkelman and Collura, "Neurofeedback Using the Phenotype and Z-Score Modalities," Neuroconnections Newsletter, July 2008, p.24.

¹⁰ Graph taken from "Understanding Evolution for Teachers," an online course created by University of California Museum of Paleontology, at evolution.berkeley.edu/evosite/evo101/VIIB1dMassExtinctions.shtml

¹¹ Charles Stark, "Consistent Dynamic Z-Score Patterns Observed During Z-Score Training Sessions," Neuroconnections Newsletter, April 2008, p. 37.

¹² J. A. Scott Kelso, "Haken-Kelso-Bunz Model." Scholarpedia (2008), 3(10), p.1612; at www.scholarpedia.org/article/Haken-Kelso-Bunz_model

¹³ Brenda R. J. Jansen and Han L. J. Van der Maas, "Evidence for the Phase Transition from Rule I to Rule II on the Balance Scale Task," *Developmental Review* (2001) **21**, p.450-494, at psychology.stanford.edu/~jlm/pdfs/JansenvanderMaas01Hysteresis.pdf

¹⁴ J.A. Scott Kelso, "Dynamic Patterns, the Self-Organization of Brain and Behavior," MIT Press (1995), p.127.

¹⁵ J.A. Scott Kelso, *Ibid*, p. 135.

¹⁶ Mary St. Clair, "LENS Case Study: Traumatic Brain Injury from AVM (Arterial Venous Malformation)", Neuroconnections Newsletter, July 2008, p. 29.

¹⁷ Charles Stark, *Ibid*.

¹⁸ György Buzsáki considers similarities between cortical and critical phenomena in his section titled, "Cycle 5: A System of Rhythms, from Simple to Complex Dynamics," in *Rhythms of the Brain*, Oxford (2006), pp.111-135.

¹⁹ Penijean Rutter, personal communication pertaining to her techniques for optimizing Z-score training.