Neuromodulation Technologies: An Attempt at Classification

by Siegfried Othmer, PhD, BCIAC

INTRODUCTION

This chapter addresses the question of how to classify the neuromodulation effects resulting from widely differing neurofeedback approaches developed over the last four decades. We have seen a proliferation of targets and objectives to which attention is directed in the training. With regard to clinical outcomes, however, one encounters a broad zone of commonality. Why is it that the premises and technological approaches within the neurofeedback network of scholars and clinicians are so disparate, yet they largely achieve common clinical goals? This in-depth analysis may lead us closer to the "essence" of neurofeedback and provide focus for further development efforts.

In its most common applications, EEG feedback typically combines two challenges, one directed to the frequency-based organization of brain communication and one that targets inappropriate state transitions. These two challenges lead to very different rules of engagement. As such rules are unearthed, they must be understood in terms of an appropriate model of brain function. At a more philosophical level, an understanding of this whole process also takes us to the very cusp of the mind-body problem, the neural network relations that provide the nexus where our thoughts are encoded and interact directly and inseparably with network representations of psychophysiological states.

This chapter will attempt to appraise the "state of the field" at this moment. The objective is to discern the commonalities among the various approaches on the one hand, and among the clinical findings, on the other. This will lead to a codification of a "minimal set of claims" that could serve to cover the commonalities among the techniques, and it will lead to a simple classification scheme for the various clinical findings. The evidence in favor of such a minimal set of claims will be adduced largely by reference. Further, the classification of the various clinical findings will serve the objective of a more appropriate or natural language for the field of neurotherapy than is provided in the formalism of the Diagnostic and Statistical Manual of Mental Disorders (DSM-IV, APA, 1994).

Tracing the Historical Threads of Neurofeedback

The Alpha rhythm and "felt states." The field of neurofeedback began in two threads of research that were concerned with one or another of the primary resting rhythms of the EEG. Here the local synchrony of the neuronal assemblies was such that the EEG amplitude would episodically rise above the ambient into dominant spindle-burst activity. In the case of the alpha rhythm, the feature was so obvious in the record that it became the first identified signature of the EEG in the original discovery of Hans Berger (1929). Joe Kamiya then first studied it in relation to our felt states, the question addressed being whether the human subject is able to have any kind of awareness regarding his own alpha activity (Kamiya, 1968). An affirmative finding eventually led to active reinforcement on alpha spindle incidence getting underway (Hardt and Kamiya, 1976). The preoccupation with subjective states of awareness and of feeling, however, was not consonant with the prevailing Zeitgeist, and Kamiya's research found little resonance in the broader field of psychology.

The sensorimotor rhythm and behavioral state. The work of Maurice Barry Sterman very consciously took a very different tack. First of all, the work utilized animal subjects, so there was no question of inquiring into felt states, but that would not have been Sterman's inclination in any event. The thrust was to connect the realm of the EEG with that of overt behavior and of behavioral states. The focus became the sensorimotor rhythm, spindle-burst EEG activity localized to the cat sensorimotor cortex that was observable even in the waking state during periods of motoric stillness. It was observed that training the behavior in order to manifest the SMR spindle was not as efficient as rewarding the animal for the appearance of the SMR spindle and having behavioral stillness as the concomitant. Either way, however, the phenomena were coupled (Sterman, Wyrwicka, and Howe, 1969). (For a review of this early research that ties into later neurofeedback, see Othmer, Othmer, and Kaiser, 1999.)

When attention later turned to the use of this simple reinforcement technique for the suppression of seizure susceptibility in humans, the training had to be done under circumstances in which the EEG was often not well-behaved as it had been in the cats. Sterman was the first to install inhibit functions on this account, but the intent was simply to assure that inappropriate triggers of a reward were suppressed. Significantly, the focus of the work remained entirely on the elicitation of an SMR bursting response. A second issue was that the human waking EEG did not manifest SMR spindles that clearly rose above the background activity, as was the case for cats. But the training philosophy carried over, as only extreme SMR amplitude excursions were rewarded. This was expected to either end up in skewing the SMR amplitude distribution or perhaps in moving the entire distribution to higher amplitudes. The focus on seizure management placed this method within the domain of neurology, but it was unlikely then (and remains unlikely now) that the field of neurology would look favorably upon behavioral interventions. (For a review, see Sterman, 2000)

EEG operant feedback and normalcy patterns. Joel Lubar was the first to employ inhibit functions with the overt objective of training toward more normal distributions. This proscriptive aspect of the training imposed its own rules on the training task, and also made for a non-prescriptive appeal to the brain that differed considerably from what was involved in the reward-based training. It also elevated the issue of EEG normalcy as a guiding principle to EEG reinforcement protocols, with profound implications for the emerging field. (For a review of this early work, see Nash, 2000)

Stimulation-based treatment. Paralleling the above developments were various stimulation-based approaches to brain-state alteration, mainly using audio-visual modes. Indeed this work had its earliest precursors in the work of Adrian and Matthews (1934), who evaluated optical stimulation in their replication of Berger's rhythm. It therefore preceded EEG feedback by some three decades. But audio-visual stimulation had suffered the same fate as Kamiya's work of being taken up by a variety of enthusiasts over the years, which then spoiled it for the attentions of academic researchers. Stimulation-based techniques have since come to be seen as competitive with reward-based feedback in terms of clinical efficacy, and must therefore be included in any comprehensive appraisal of the field. The evidence for this is strongest for ADHD (Siever, 2007). In order to accommodate both neurofeedback and stimulation the more inclusive term of neuromodulation will be used below.

The development of the field subsequent to the early initiatives by Kamiya, Sterman, and Lubar has been modestly evolutionary, but the essential character of the work was laid down during the early days of the field, and threads of continuity carry through to this day. Only the stimulation-based work requires a separate treatment. The subsequent discussion is conducted more at the conceptual level rather than being constructed strictly upon the established empirical basis. Of course empirical data drive the discussion, but it would be premature to make fine distinctions on the basis of the available evidence, or to be too judgmental at this point, for example with respect to the relative efficacy of the various techniques. Most if not all of the approaches remain in a state of both technical and tactical immaturity. Moreover, the clinical world is not restricted to using only one mode but will likely see the emergence of a multiplicity of techniques and combinations of techniques for synergistic effects. The question of which is best therefore does not even merit a response at this time. And by the time the question can be answered well, it will hopefully no longer be relevant.

A Classification of Neuromodulation Technologies

At the top level we may partition the field into volitional and non-volitional approaches, with feedback techniques generally falling into the former category and stimulation-based techniques into the latter. This is in line with the traditional focus on "voluntary controls" in the biofeedback literature, and with the emphasis on recruitment of the individual's efforts and intentions in the service of better self-regulation in traditional biofeedback. Clinical experience with neurofeedback, however, calls even this facile partitioning into question. Neurofeedback training "works," after all, with extremely cognitively compromised infants (e.g., victims of near-drowning) who cannot possibly have any awareness of the objective of the training. Even neurofeedback must therefore be understandable at the brain level, without any assumptions about "higher-level" engagement with the process.

Paraphrasing Robert Bly, one might argue that the "brain is here to seek its own joy," and that it will be attracted to anything that engages it, intrigues it, plays with it, or mirrors it. Given that rather unobjectionable assumption, one might argue that overt instructions to "succeed" may be optional even in the feedback paradigm. Feedback may be sufficiently rewarding intrinsically to mobilize the process even in the compromised brain. One might additionally inquire about the role of volition in Sterman's cat research (Sterman, Wyrwicka, and Howe, 1969). Certainly the cats were strategizing to get food. When one of them happened to trigger food reward while engaged in an elaborate stretch movement, she subsequently repeated the same movement time and again in the hopes of repeating her success. The strategy might well have been counter-productive in the end. But no matter. It was rewarding enough to be sustained. Beyond experiencing food as a reward, for which no special provision needs to be made, nothing more appears to be required. It seems that we cannot give volition an essential or even an exalted role in EEG feedback.

On the other hand, perhaps not much is lost, as the following anecdote illustrates. A person undergoing training with NeuroCarePro software (See www.zengar.com) expressed his satisfaction with the experience at the end of the session, but could not help voicing his irritation with the fact that the CD he was listening to kept interrupting the flow of the music. "I feel much better, sleeping well, but can you ask her to use a new CD that doesn't have skips?"

He had clearly not even been aware that the skips were the bearers of information relevant to his brain---in fact the only ones his brain was receiving---and yet the process clearly influenced him. Not only was he annoyed with the discontinuities in the music, but his brain was also. And the brain had more information to work with than he did. It did not take that exquisite correlation engine long to figure out that it was part of an interactive system in which it was playing an active role. And just as the brain will incorporate a tennis racket as an extension of the arm as soon as it is picked up, the brain will rapidly appropriate the feedback loop as part of its sphere of influence. Clearly we must understand neurofeedback at the brain level. Once that task is accomplished, we can readmit the role of volition to the discussion, as it can strongly enrich the feedback process in most real-world situations.

A less ambiguous distinction may lie in the recognition that feedback is cognitively mediated and stimulation-based approaches are more directly brain-mediated. In nearly all of conventional neurofeedback, the feedback signal is processed as "information," whereas stimulation is experienced by the brain as a direct perturbation of its activity. The feedback signal is processed along with all other news from the environment, and it is appraised similarly. In time, a correlation becomes apparent to the CNS between some feature of the information stream and some aspect of its own internal activity. By contrast, in frequency-based stimulation techniques we are simply creating excess (or in some instances reduced) neuronal synchrony, which then has implications for brain function downstream to which the CNS in turn reacts.

In the feedback case, the work is subject to all of our cognitive and attentional limitations, our susceptibility to distractions, and to the vote of our emotions as to whether we are actually committed to the process. In the stimulation case, the brain has no option but to respond. Repetitively pulsed light sources will ineluctably impose their periodicity on visual cortex, from whence they are propagated throughout cortex. Given the simplicity involved in stimulation-based methods, one might well ask why conventional neurofeedback has not already been largely displaced. A rather general observation appears to hold for neuromodulation technologies, namely that newly emerging techniques don't really displace older ones but rather add to them. It seems that old protocols never die.

A different question therefore needs to be asked, which is why there is so much speciation occurring in EEG feedback at all? What are the evolutionary niches that allow all of the techniques to survive unto the present day? Perhaps we are just in an early proliferation phase, with consolidation to follow later. This issue will hopefully be clarified in what follows. At any rate, our present task of simply classifying the various techniques is challenged by all of the diversity that is already extant.

A simple classification based on the above considerations is that of "active" versus "passive" techniques. We are not writing on a blank page, however, and past usage has tended to regard stimulation techniques as active interventions, and neurofeedback as passive. A more organic view of the matter would hold that the point of reference should be the client, who is the passive recipient of the stimulation and the potentially active participant in feedback.

Focusing then on the 'active' technique of feedback, a major division could be established on the basis of whether the success criterion is defined very narrowly or broadly. The former may be referred as "prescriptive" and the latter as "non-prescriptive." Reward-based training targeting a particular EEG frequency band would be considered prescriptive. Success can only be had if an amplitude objective is met within a narrow band of frequencies. Inhibit-based training is an example of non-prescriptive training. The brain is simply being alerted to an adverse state of affairs, and it is given no particular hint as to how to remedy the problem.

The same partitioning could also be labeled "prescriptive" and "proscriptive" training. In the latter, the brain is simply being notified of some transgression or another with respect to certain established bounds of EEG phenomenology. Deviation (in EEG terms) is deemed to index deviance (in behavioral terms) or dysregulation (in neurophysiological terms). The CNS is left to its own devices to sort things out. At the present state of maturity of our field, prescriptive training tends to be frequency-based, and proscriptive training tends to be event-based. If the dysregulation status of the person is found to have worsened beyond some threshold value, by one or more EEG criteria, then attention is drawn to the event and rewards are withheld for the duration.

Slow Cortical Potential Training

Finally, where does Slow Cortical Potential training fit into this picture? This is the technique, developed by the Tübingen group led by Niels Birbaumer, in which the trainee is asked to alter his or her low-frequency EEG transiently by several microvolts within an 8-second window (Strehl et al., 2006). The rewarded change can be either positive or negative, depending on the assigned goal in a particular trial. Acquisition of control is the principal objective here. The main applications to date have been to locked-in syndrome (for communication purposes), to seizure management, and to ADHD. The impression one gets is that the technique is just as diagnostically non-specific as frequency-based reinforcement. The work has been helpful for the abatement of migraines, for example (Kropp, Siniatchkin, and Gerber, 2002). We have here a case of event-based training that is prescriptive in the prescriptive/proscriptive partitioning while being non-prescriptive in the prescriptive vs. non-prescriptive division. That is to say, the CNS is left to sort out how the training objective is to be met.

Stimulation-based Technologies

How do stimulation-based technologies fit into these categories? Standard audio-visual entrainment has tended toward stimulation of particular EEG frequencies, and as such falls into the category of prescriptive modes of neuromodulation. More recently, however, the technologies have added swept modes of operation, with the general goal of stimulation without a specific target in terms of frequency. These modes of operation come closer to being in the non-prescriptive category.

The derivative technologies of the ROSHI and the LENS, on the other hand, use entrainment techniques effectively for the purpose of disentrainment. The original ROSHI design would pick out the dominant EEG activity in the low-frequency regime and apply a stimulation that was out of phase with the ongoing signal and thus bring about its dephasing and disruption (Ibric and Davis, 2007). The LENS effectively does the same thing by means of electromagnetic stimulation with a carrier frequency in the megaHertz region, modulated by the EEG signal of interest (Ochs 2007). In this case, however, the disruption is achieved with a frequency-shift rather than a phase shift. Over short time intervals, a frequency-shift and a phase shift amount effectively to the same thing.

The particulars probably don't matter nearly as much as ROSHI and LENS practitioners may be inclined to believe. It is the process of disruption that matters. That being the case, these techniques should be lumped into the bin of non-prescriptive modes. The techniques are applied identically regardless of the diagnostic presentation. Once a training site has been selected, these techniques are referenced entirely to the instantaneous behavior of the EEG. Consequently, an understanding of these methods should be possible largely with reference to EEG phenomenology itself. It is true that adjustments in clinical approach (i.e., with respect to site selection and hierarchy of targeting) may be made on the basis of the response to the training. But that is only as it should be, and it does not gainsay the observation that the training is driven entirely by the instantaneous EEG.

The LENS in particular has undergone a developmental pathway that is highly instructive for our more general purposes. Len Ochs observed over the years that the response to the stimulation was far greater than most practitioners were expecting. Already in feedback practitioners were encountering negative effects associated with lingering too long with a protocol that retrospectively can be judged non-optimal. The same was true for the stimulation techniques, only possibly more so. In response, Ochs trimmed back the stimulus duration farther and farther, each time finding that the clinical effectiveness remained robust while the probability of an adverse outcome diminished. The latter never declined to zero, however.

One had the impression that if prescriptive neurofeedback were a rifle that we are increasingly learning how to aim, then LENS would be a cannon whose aim remained ambiguous. The technique in essence remains non-prescriptive, and in application of such a powerful technique to a severely dysregulated nervous system, the outcome must of necessity remain somewhat unpredictable.

The development of the ROSHI was even more overtly in the direction of non-prescriptive training. With the personal ROSHI, the stimulus was provided over a range of EEG frequencies pseudo-randomly selected, and delivered over brief intervals. The principle underlying this approach was that of stochastic resonance. A finite probability exists that the stimulus at any given time would have the right frequency and phase properties to effect the desired disentrainment that in the original ROSHI had been so carefully engineered. At worst one suffers a loss in clinical efficiency, which is not relevant in the personal use applications for which the device was mainly intended.

There are other aspects to the ROSHI design that remain proprietary, and I will honor the wishes of the designer and refrain from bringing these elements into the discussion. Significantly, the personal ROSHI stimulates both hemispheres differentially, thus introducing a phase disparity with which the brain must come to terms. This challenge in the phase domain is intrinsic to its design, and probably accounts for its broad efficacy.

One might reasonably object that a finite probability also exists that the stimulus phase would be such as to induce entrainment rather than disentrainment. It turns out that this does not really matter! When formal attention was finally given to the audio-visual entrainment technique some years ago by Lubar's group, it was found that after the stimulus period was over, the EEG would tend to show disentrainment effects (Frederick, Lubar, and Rasey 1999). The implications are obvious: When the brain is subjected to the interference we call entrainment, it yields to the stimulus but also mounts a defense. The defense is the learned response, and that is what lingers after the stimulus is over. Fortuitously, we are presented with the delightful paradox that the right outcome does not depend strongly on the particulars of the stimulus. It is the disruption itself that matters. So, if standard neurofeedback is the rifle, and the LENS is the cannon, then the ROSHI is the shotgun.

The Evolution of Standard Reward and Inhibit-based Neurofeedback

The exciting developments over the years with the ROSHI and the LENS leave one with the impression that a significant advance over standard neurofeedback practice may have been achieved. Certainly these techniques have added significantly to the clinical arsenal. But in the meantime neurofeedback has seen its own evolution. To complete the picture, these developments also need to be discussed.

The rise of quantitative EEG analysis within the field of neurofeedback that took place in the early nineties had a significant influence on the subsequent development of clinical neurofeedback, as well as determining how the whole process was to be understood. The mandate of quantitative analysis of the EEG is to establish the stationary properties of the system, furnishing measures that will hopefully be valid even the day, the week, and the month after the data are acquired. This focus on the steady-state properties of the EEG already was fraught with implications that were not appreciated generally at the time.

QEEG information is acquired using either Fourier transform or other similar transform techniques (e.g., Gabor). Here a sufficiently long time sample is converted into its constituent frequencies. "Sufficiently long" means that the time window must accommodate the lowest frequencies of interest without compromise. Typically 0.5 Hz is taken as the low end of the range of interest, and at least a half cycle needs to fit comfortably within the window in order to be represented properly. A windowing function is usually installed to minimize aliasing effects, which further narrows the effective length of the time window. A proper representation of 0.5 Hz therefore mandates a time window greater than one second.

Now it will be recalled that Sterman's original interest was in recognizing individual spindle-burst activity in the cat, and the same objective was later translated to human subjects. The steady-state amplitude of the EEG in the SMR-band was never of interest at all. So the question arises: How well do transform-based systems do when the task is to recognize brief transients? The answer is that they don't do well at all. This should be no surprise. The whole intent is to furnish data on the steady-state properties of the EEG, and an analysis scheme oriented to that task cannot be expected also to do well with transient data. In the real EEG one might see several SMR/beta spindles over the course of one sampling window of nominally 1 sec. These are averaged over in the spectral calculation, and individuality is lost. See **Figure 1** in that regard.

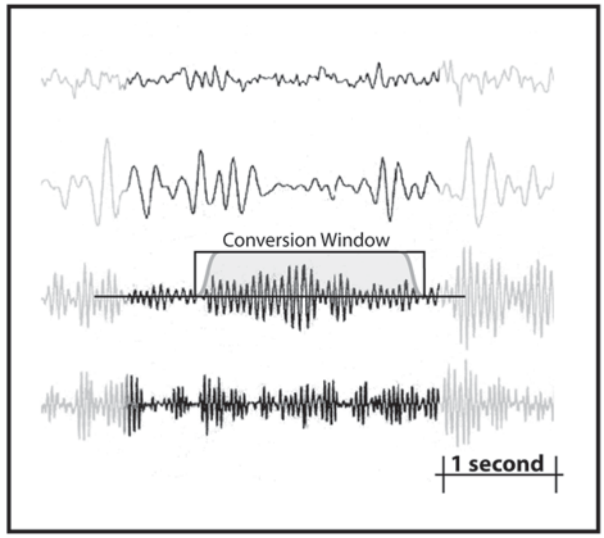

Figure 1. *The data sampling process for transformation from the time to the frequency domain is represented. A finite conversion window, with the windowing function shown as shaded, illustrates the uneasy compromise that needs to be made between the conflicting demands of reporting EEG change promptly and of representing steady-state values. From top to bottom, the graphs reflect the spectral bands of 0-30 Hz; 4-7 Hz; 15-18 Hz; and 22-30 Hz. Gains have been individually adjusted for readability.*

Consider further what happens when a lonely spindle-burst of high amplitude, the very thing Sterman was targeting for success, comes along in the signal stream. First it encounters the windowing function, which means that its full expression in the transform is delayed (hence delaying the issuing of a reward). Some time later it leaves the sampling window (in a subsequent time sample), causing the signal amplitude to decline as it does so. But this decline in signal amplitude does not reflect what is happening at that moment, as one would wish; instead it reflects what happened a second ago. So, we have the disagreeable situation that what enters the window as signal inevitably exits the window as artifact some time later. The simple expedient of moving toward transform-based analysis has cut the signal-to-noise ratio in half for highly dynamic signals such as the EEG (unless an asymmetric windowing function is employed).

A second change that accompanied the transition to QEEG-driven training was the conversion to referential placement from the bipolar montages that Sterman and Lubar had used in their initial research. If the QEEG measures were going to inform neurofeedback then the montages had best be compatible. The localization of brain events that was becoming possible led to a conceptual change in how neurofeedback was to be done, with an increased emphasis on the training of steady-state amplitudes at single target sites referenced to the ear, which was taken to be quasi-neutral.

It was the choice of several instrumentation developers, including the author, to stay with the early systems design in which frequency selectivity was obtained by means of narrow-band filtering. In these designs, the "real-time" incoming signal always carries the greatest weight. But delay in the signal stream was not thereby banished. Some amount of delay is involved with any signal analysis technique. The parameter relevant to filtering is the group delay, the time difference between comparable signatures in the raw signal and in the filter output. This quantity is determined at the center frequency of the filter. The group delay through the filter chain is a parameter that can be managed through suitable choices in the design of the filter to be in a tolerable range of 150-250 msec. This amount of delay still allows the brain to make an identification between the emerging data on the video screen and its own ongoing activity.

A significant change in the way filter-based neurofeedback was actually conducted occurred over a period of years. The change was incremental and cumulative, and was therefore perhaps less than consciously made. At the outset both Sterman and Lubar chose to mete out rewards quite sparingly, with the intent of rewarding only the largest amplitude SMR spindles. This was done straight-forwardly by choice of reward threshold. Done in this fashion, the work was quite tedious for people. In information-theory terms, the brain was not getting a lot of information to work with.

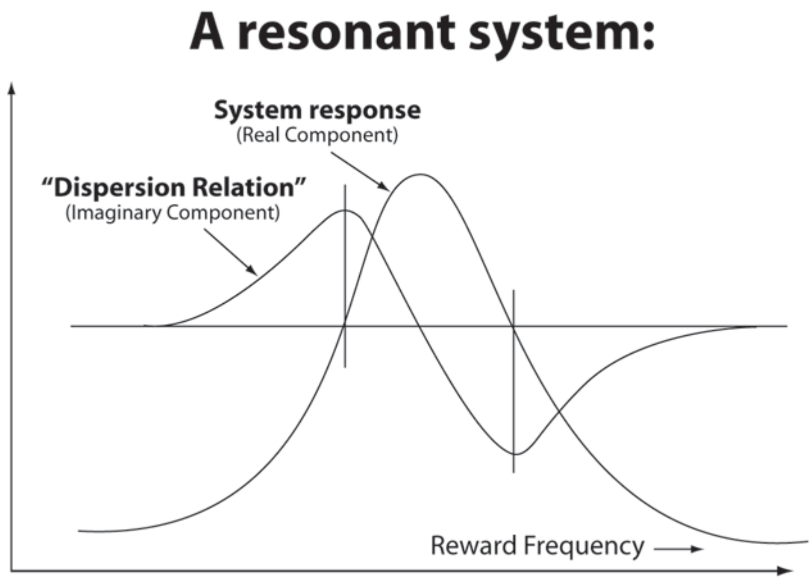

Figure 2. *A response curve is shown for a resonant system. It is assumed that spindle-burst formation in the EEG exhibits the properties of a resonant system. The real and imaginary components of the system response are plotted vertically. The feedback challenge is sufficiently frequency-specific that in sensitive responders an opportunity presents itself to map out the resonance curve for an individual. One EEG frequency clearly dominates the response in such individuals.*

The simple expedient of increasing the reward incidence made the training much more lively, engaging, and rewarding. The payoff in clinical efficiency was dramatic. But over time this success was taken to what appeared to be ridiculous extremes. Thresholds were being set so that the reward incidence was at the 70-85% and even 90% levels. One was reminded of modern American schooling where nearly everybody gets an A. The clinical results were holding up, but what was being discriminated here if 90% of what was happening in the reward band garnered passing marks?

The game had in fact changed underfoot in a manner that was probably not fully appreciated at the time. Typically, the discrete rewards were limited in incidence to a rate of two per second. With the rewards now plentiful, they were arriving in a regular cadence to which the brain rapidly accommodated. With the rewards having become the expectation, the attended event became the occasional dropout of the rewards. (This is reminiscent of the odd-ball design in evoked potential research, in which the occasional odd-ball stimulus evokes the attentional resources as reflected in increased P300 amplitudes.) Effectively, the discrete rewards had come to perform a function we associate with inhibits! In the meantime, the role of the reward had been assumed by the analog signal in the reward band, which was being continuously displayed on the video screen.

The CNS was now in continuous engagement with the analytical representation of its own activity on the screen. The reward here is intrinsic to the process, and is entirely independent of threshold. Give the brain a chance to engage with its own activity, it will quite naturally be inclined to do so. The problem of boredom is resolved by the simple expedient of enlarging the size, the continuity, the promptness, and the salience of the signal stream. The brain will not fail to be interested in its own activity.

In consequence of the above developments, clinical practice then followed the strengths of the respective methods of signal analysis. The relative strength of the filter-based approach was in tracking the dynamics in the reward band, so the preoccupation of filter-based systems has remained with reward-based training. The relative strength of the transform-based systems was in discerning change in ambient band amplitudes with slightly longer effective integration times, and thus the focus of QEEG-based training has been increasingly on the inhibit strategy, with reliance on amplitude-based training.

Resonant-Frequency Training

Although reward-based training has largely been performed under the rubric of SMR/beta training, it has been clear for many years now that people respond quite variedly to the standard protocols. This turns out to be largely a matter of reward frequency, so that the response can be tuned by the mere expedient of adjusting the reward frequency. This diminishes the special role that SMR-band reinforcement has played in our clinical work and in our conceptions. In practice, of course, standard SMR-band training has remained prominent within the field, but that is largely because most practitioners feel obligated to maintain standardization of protocols to the extent possible, and in consequence they have not yet investigated the frequency optimization hypothesis.

It is the responsiveness to optimized-frequency training that makes this training approach practical. The immediate response of the reinforcement is in terms of state shifts in the arousal, attentional, and affective domains. These state shifts are readily perceived within a matter of a minute or two or three by anyone who responds sensitively to this training. Reports on perceived state change are elicited by the therapist, and on this basis the reward frequency is adjusted on the timescale of minutes. As the optimum reward frequency is approached, the trainee achieves a more optimal state in terms of arousal, vigilance, alertness, and euthymia. At the same time, the strength of the training increases perceptibly. For those familiar with the theory of resonant systems, this maps out a conventional resonance curve, and it is our impression that the person's felt states and the responsivity to reinforcement map out essentially the same curve.

This frequency-dependent behavior is shown in terms of a standard resonance curve in **Figure 2**. This curve traces out the frequency response of the "real" component of the resonant system. Both positive feeling states and response to training are thought to be reflected in this single curve, as sketched in **Figure 3**. In any physical resonant system, however, there is also the "imaginary" component to consider, and this is mapped out as well in Figure 2. We have some tantalizing evidence that this quadrature component shows up in terms of an enhanced sensitivity near the resonant frequency, and may be experienced in terms of adverse feeling states. A crude analogy may have to serve us here: the relative calm at the resonant frequency may be like the eye of the hurricane, but turbulence is maximized in the vicinity of that eye.

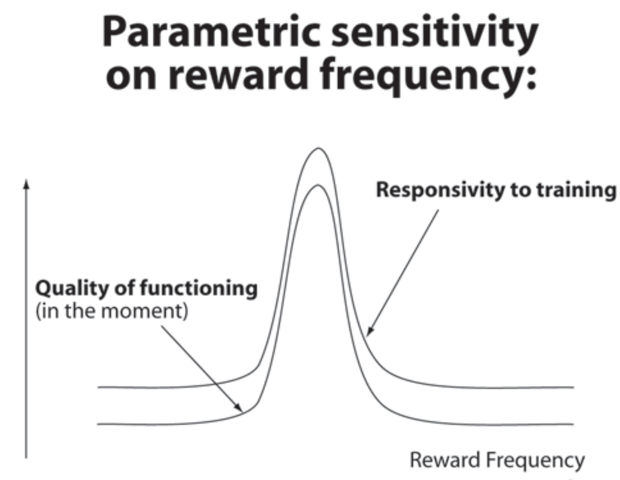

Figure 3. *Both the responsiveness to the reinforcement (responsivity) and the subjective quality of functioning of the individual appear to map out the real component of a resonance curve as the reinforcement frequency is incremented episodically through the band. These dependences are observed on highly sensitive responders, but are likely to hold more generally.*

Since this behavior can be observed in different people across the entire frequency band from 0.01 Hz to 45 Hz, it is likely that the same general organizing principles apply for every part of the band. That is to say, all spindle-burst activity must be organized as resonant systems, even down to the lowest frequency we have characterized. On the other hand, in each person who is sensitive to this training, one frequency band appears to stand out above all others in terms of its relevance to training self-regulation in that individual. The original finding that training one particular band, namely the SMR band, has quite broad (i.e., non-specific) implications for self-regulation status now stands on a more solid foundation, albeit with the proviso that the particular frequency is unique to the individual. The SMR band has just lost its special status in this approach.

An Attempt to Achieve Synthesis

The EEG is organized to a level of detail and precision that is difficult to discern with our conventional measurement tools. Regardless of what elements of the signal we choose to focus on, yet others must remain out of focus or off-screen entirely. If one chooses to view the EEG with high frequency resolution, for example, the segregation into distinct, narrow, rigorously demarcated frequency bands is quite striking. Since this cannot be trivial to arrange, it must be important to brain function. The demarcation line between two frequency regimes is a phase boundary, i.e. a region where the phase can undergo a discontinuity. Within a particular frequency range that defines a spindle, the phase varies smoothly and continuously throughout. This is illustrated schematically in **Figure 4**.

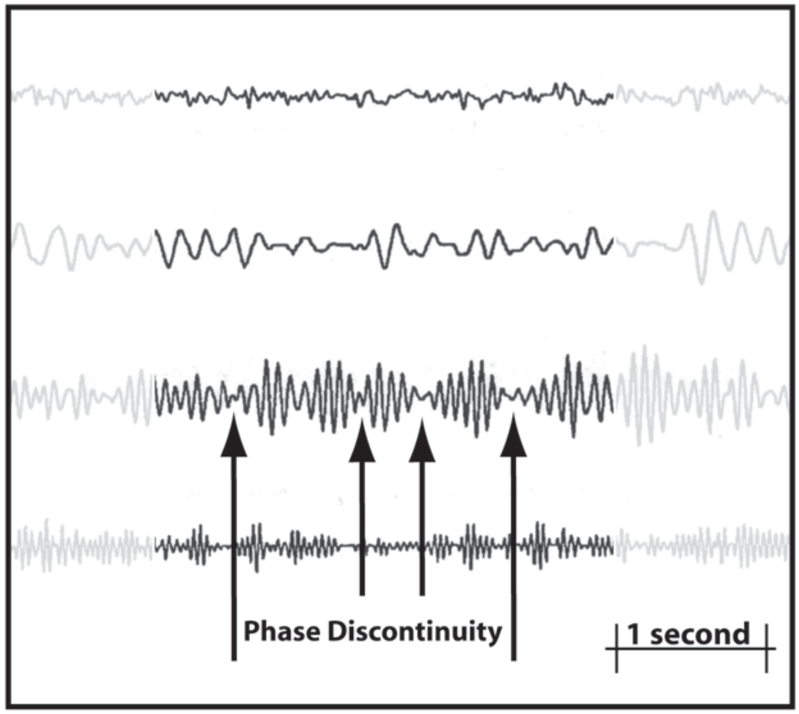

Figure 4. *The phase discontinuities observed between spindle-bursts in the 15-18Hz band make the case that these spindle-bursts represent functionally independent cortical events. Under that assumption, it makes sense to target the individual bursts that are under active management by the brain, as opposed to targeting averages over independent events. This calls for a highly responsive feedback system with the largest possible signal bandwidth.*

Similarly, the spatial distribution of a neuronal assembly must be characterized by a smooth phase variation over the assembly. If a phase boundary exists at the margins of such a neuronal assembly, as one might suspect, it is likely to be obscured in practice by volume conduction. Finally, communication between neuronal assemblies at some remove from one another is contingent on phase alignment. It follows, then, that the CNS must manage phase in exquisite detail in order to regulate the territory that a neuronal ensemble commands in the frequency domain, to delineate the spatial footprint that it occupies on cortex, and to establish and maintain distal communication with neuronal assemblies elsewhere on cortex.

Subtle interference with this process will then provoke the brain's cogent response. The interference is deemed to be subtle if it simply modulates rather than disrupts the ongoing activity. Neurofeedback categorically meets this criterion, and stimulation techniques are capable of meeting it if conducted at sufficiently low drive levels. Both feedback and stimulation techniques are strongest if they impinge upon the aspect of brain function that is under the immediate management of the brain, and that is the relative phase or, equivalently, the instantaneous frequency of the packet.

This model for neuromodulation accounts for the power of the LENS and ROSHI approaches, and for the power of the optimized reward frequency training in feedback. It also accounts for the observation in QEEG-based work that coherence training appears stronger than amplitude-based training. In LENS, the stimulus provides the phase reference. In the case of the ROSHI, the stimulus phase differs between the two hemispheres. In the frequency-optimized feedback, which is typically conducted with a single channel in bipolar montage, the one site is the reference for the other. The same holds true in coherence training with two-channel montage: one site represents the phase reference for the other.

In EEG training with a bipolar montage, the net reward signal is a strong function of the relative phase. This is an essential point, and it is not an obvious one. The reader is referred to a detailed treatment of this topic in Putman and Othmer (2006). By virtue of common-mode rejection in the differential amplifier, activity that is synchronous between the two sites is not seen at the output, and therefore cannot ever be rewarded. And if it cannot be rewarded then with respect to everything else it is effectively being inhibited. The net effect is to reward differentiation of activity between the two sites, which is the real take-away message. The approach was first investigated with inter-hemispheric placements at homotopic sites, which we used rather exclusively for some years (Othmer and Othmer, 2007). Clinical results in terms of continuous performance tests have been published for this method, demonstrating improved outcomes with respect to earlier data (Putman, Othmer, Othmer, and Pollock, 2005).

Looked at in the above way, even this very specifically targeted reward-based training can be seen as having a proscriptive rather than a prescriptive aspect: the state of synchrony, of phase conformity, is proscribed. Conversely, the technique rewards “everything but the condition of synchrony.” And since the only thing precluded from success is the synchrony condition, with respect to the broad remaining phase domain the technique can even be seen as having a non-prescriptive aspect as well: the phase relationship is not being tightly constrained.

The training itself amounts to a subtle, continuous challenge that lies largely in the phase domain. It must be acknowledged at this point that the dominance of phase is not obvious from the mathematics. Indeed, amplitude differences between the two sites play just as strongly into the net reward. Since the relative role of amplitude and phase in real-life situations is not obvious, their respective roles can be clarified with a Monte Carlo calculation in which the experimental situation is simulated in all of its natural variability. This has been done with the assumption of randomness in relative phase and in the amplitude at the two sites (Putman and Othmer, 2006). A nearly complete exclusion of rewards is found for relative phase less than forty degrees in this simulation. In real EEGs there will be some finite correlation in amplitudes, and that only serves to strengthen the posited phase dependence. (If the variance in the amplitude ratio is pinned, then the actual variance must be accounted for in the phase.)

Looking carefully at real EEGs also makes it clear that phase is often the more dynamic variable of the two. The sequential independent spindle-bursts one sees with a fixed narrow-band filter (as shown in Figure 1) may represent neuronal assemblies of slowly varying amplitude that are simply migrating in frequency through the filter band. Others undulate back and forth within the filter pass-band. What really makes the difference is that our experimental situation is so arranged as to highlight phase variations, with the result that these will come to the fore in the reward schema. The argument goes as follows:

A narrow-band filter can be seen for our purposes as a transducer of frequency fluctuations into amplitude fluctuations. Frequency variation and phase fluctuations are obviously directly related. Dynamic, continuous reward-based training using narrow-band filters attempts to shape the EEG frequency distribution toward the middle of the resonance curve, with often immediate and sometimes trenchant consequences for the person’s state. These factors are in play even in ostensibly single-site amplitude training with referential placement (because references are not silent).

Bipolar montage then further augments the role of narrow-band filters as phase discriminants because the amplitudes at the two sites are now more correlated than in referential montage, which shifts the burden of variability more onto the phase. In typical application, the bipolar montage will be deployed either at near-neighbor sites or at homotopic sites. In these cases, the correlation of amplitudes (i.e., comodulation) is typically enhanced with respect to arbitrary site combinations.

Of course we aren’t looking at normal EEGs in the usual clinical situation. In the presence of dysregulation, we typically see enhanced EEG amplitudes, particularly in the low frequency regime. Enhanced amplitudes can be modeled equivalently as excess local synchrony. Bipolar training in the midst of such activity can then be seen as disruptive of that activity, tending us toward better function. The fear about bipolar training under such adverse circumstances is entirely misplaced. In practice, it is all a question of finding the optimum response frequency. Ironically, that criterion enforces an even tighter constraint: All reward frequencies may conceivably be contraindicated, or at least disfavored---except for the narrow band that is favored.

Sometimes, of course, QEEG data reveal deficits in connectivity between sites rather than excesses. Would the standard bipolar training be a mistake under such circumstances, in that promotion of desynchronization is not called for? There is no evidence yet that this presents a problem. Just as the sign of the phase challenge is a secondary issue in LENS, so we believe it to be in feedback as well. Efficacy lies in the subtle challenge to the system, often surprisingly indifferent to the particulars. The intent is to normalize the pathways of communication, and this can be done by challenging them in one way or another. The brain will take it from there. Any kind of response by the brain to the provocation is likely to go in the direction of improved regulation. We don't have to make the pearl. We only have to provide the grain of sand.

The above argument has made the case that clinical efficacy is broadly available in neurofeedback. Clinical efficiency, however, is an entirely different matter. As already implied in the discussion of optimized frequency training, the greatest clinical efficiency may be highly constrained in terms of protocol.

The implication of findings with LENS, with ROSHI, with frequency-optimized bipolar training, and with coherence-based training is that phase-based targeting is more responsive, and more availing, than amplitude-based training. To be effective, the feedback signal must provide information on the relevant timescale, and in that regard we have three timescales to consider. There is first the timescale of the conversion interval in Fast Fourier transform analysis on which QEEG-based training is based. This approach was first featured in the Lexicor unit and is currently used in NeuroCarePro (Gabor transform). Secondly, there is the timescale of individual spindle-burst on which most dynamic training using filtering functions operates. This is on the order of a third to half a second. Thirdly, there is the timescale of as little as a single cycle of the EEG at the relevant frequency, the timescale on which the LENS can operate. At a typical application frequency in the theta range of frequencies, the timescale may be only a fourth to a seventh of a second.

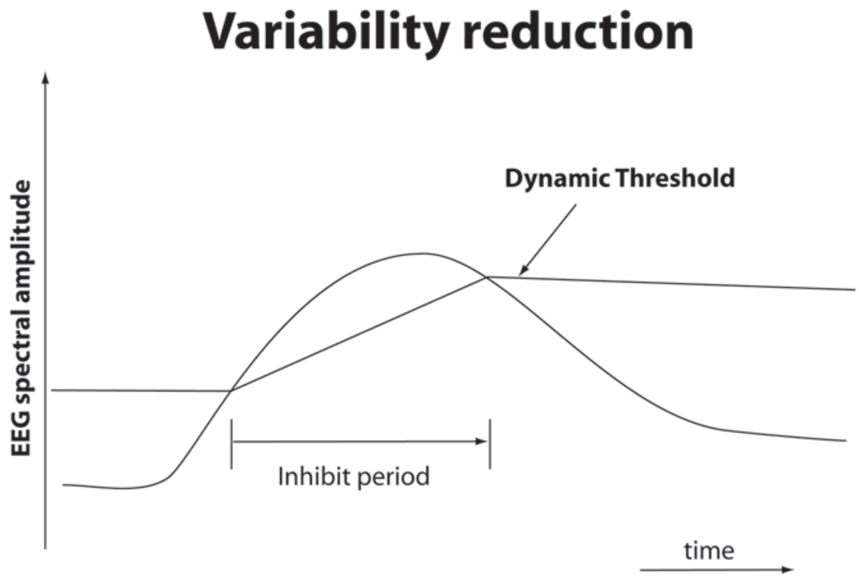

Figure 5. *The process of dynamic thresholding is illustrated for a large excursion in EEG spectral amplitude in the inhibit band. Since the percentage of time over which the inhibit may be engaged is kept fixed, a triggered inhibit will immediately cause an adjustment in the instantaneous threshold. The result is that actual inhibit periods will always be finite. Effectively, we have a transient detector that alerts us to the initiation of an excursion into dysregulation.*

There has been an overall trend toward dynamic training and away from the QEEG-driven focus on reinforcing band amplitudes. This trend has even asserted itself in inhibit-based training, which has been the strength of QEEG-guided training. Over the years there has been a gradual shift from the use of fixed thresholds to dynamic or adaptive thresholding. This was initially driven by a need to keep the level of difficulty within bounds. Software was reworked so as to maintain the level of difficulty fairly constant. The thresholds "breathed" with the ebb and flow of things on longer timescales.

In NeuroCarePro this idea was taken even further. By making the thresholds even more dynamic, inhibit-based training in that system became more like transient detection that zeroes in on inappropriate state shifts. This is shown in **Figure 5**. Effectively we have a slope detector, or derivative detector, which calls attention to unusual excursions in variability. Once again, it is more important to alert the brain immediately to its incipient transgression than to merely inform it as to levels of EEG amplitudes. The brain does not react as well to old news. So in NeuroCarePro the inhibits have become a matter of event detection. One expects that elevated EEG amplitudes are correlated with elevated variability as well. But even if that is not the case, elevated variability makes the better target in any event because it focuses on the present moment.

The nearly universal trend within the field has clearly been from stationary properties of the EEG to EEG dynamics, from discrete rewards to continuous reinforcement, from static to dynamic thresholds, from amplitude-based training to phase-based training. The considerations for effective training have become increasingly referenced to the EEG itself, and in that regard we have moved to feedback "at the speed of thought," operating with the highest response speed of which the signal analysis routines are capable. The emphasis has been on the ongoing dynamics within the reward band (which can be anywhere in frequency) and on how the brain handles state transitions. This sounds rather far removed from where we started, which was with a concern with the resting states of the system, the alpha rhythm and the sensorimotor rhythm. The story is not complete, however. We have in fact omitted from the discussion one basic approach that has been with us since the beginning of the field, and that is Alpha/Theta and alpha synchrony training.

Alpha/Theta and Synchrony Training

Alpha/Theta training, which involves reinforcement in both the alpha and theta bands---typically under eyes-closed conditions---has not been drawn into the discussion up to this point because the implicit focus thus far has been on the application of neuromodulation to the enhancement of brain function. The principal objective in Alpha/Theta training is instead to facilitate certain psychological states that promote healing from trauma reactions, recovery from addiction, etc. (For application to addiction recovery, see Scott, Kaiser, Othmer, Sideroff, 2005.) Inevitably, however, subjecting someone to these reinforcements for hours on end may also have a residual training component. For yet others, the mere exposure to these reinforcements may steep them back into their prior pathology. This is particularly the case for those who exhibit alpha intrusion into their sleep EEG, and those with a history of minor traumatic brain injury.

The promotion of alpha and theta amplitudes effectively enhances local synchrony. Unfortunately, elevated synchrony also characterizes much of the severe pathology that we encounter in clinical practice. For many who could in principle benefit from Alpha/Theta training, the reinforcement of alpha and theta bands is somewhat hazardous terrain. It is also for this reason that the bias of bipolar training toward desynchronization of the EEG makes it preferable as a starting protocol. The bipolar training is always done first, with the result that even vulnerable individuals may then tolerate the Alpha/Theta training later.

Whereas this cautionary tone reflects our own experience over the last twenty years, it is also true that long-term practitioners such as Lester Fehmi have used alpha training routinely in their practice to good effect with a clinical population for many years, more than thirty in the case of Les Fehmi. From the early days, however, Fehmi has used a different approach, which may account for his success on the one hand, and the absence of adverse reports on the other: multi-channel synchrony training (Fehmi, 2007). He was not alone in this. The work of Jim Hardt has also relied on multi-channel synchrony since the early days (Hardt, 1978). And the very first commercially viable computerized instrument to do EEG feedback was developed by Adam Crane on the principle of four-channel synchrony. In fact the first clinical practice in EEG feedback on any significant commercial scale consisted almost entirely of four-channel synchrony training performed on American Biotec's Capscan Prism 5. (Crane, 2007)

Just how is it, then, that enhancing synchrony using straightforward alpha reinforcement with a single channel can be troublesome, whereas promoting even greater synchrony with multiple channels is not similarly fraught with hazards? Most likely it is an issue of control. The precision of control that is required to garner rewards in four- and five-channel synchrony training is simply not available to the injured brain. Facilitating local synchrony with a single-channel setup does not discriminate against the unruly alpha and theta that we observe in the dysregulated brain. Multi-channel synchrony training, on the other hand, promotes global synchrony, which represents a much more specific challenge that is difficult for the dysfunctional brain to organize.

In the final analysis, then, alpha synchrony training makes the same case that has just been made for SMR/beta training and its progeny. Single-channel training of band amplitudes is trumped by multi-site training that focuses on the training of phase relationships. And when such a highly specific phase challenge is mounted, the consequences are not only for enhanced experiences of deep states in the moment but also for improved self-regulation as a learned response.

Another point of similarity is that the training is event-based. The question is one of enhancing phase conformity whenever an alpha spindle is observed in cortex. It is not a matter of training up ambient alpha amplitudes. The analogies go further. Self-regulation is enhanced broadly with reinforcement at only a single frequency, with a technique that uses standard placements and rises above issues of localization of function. In this case, moreover, there isn't much question about the optimum reinforcement frequency. In practice, it is to be found within +-1 Hz of 10.5 Hz in nearly everyone.

In Les Fehmi's mechanization of synchrony training, the reinforcement is delivered with every cycle of the alpha rhythm that meets criterion. It turns out that the timing of the delivery of the reward signal with respect to the underlying alpha signal is crucial. With the phase delay optimized, the reward pulse serves to augment the next cycle of the alpha spindle. This is firstly another demonstration that "phase matters." Secondarily, we have here a stimulation aspect to what is fundamentally a feedback paradigm.

In pursuit of the hypothesis that synchrony training may represent a natural complement to bipolar training, we have begun an investigation into the clinical utility of simple two-channel synchrony training of the resting rhythms of the brain, alpha and SMR. For some years already, we have been doing Alpha/Theta training with two-channel synchrony at P3 and P4, with results that clearly exceed the prior work with single-channel training at Pz. Synchrony training in the alpha band may additionally have its own rationale in the context of awake-state training, as distinct from the induction of deep states in Alpha/Theta.

SMR synchrony training has been evaluated at C3 and C4, with results that have been rewarding in a number of clients. The training tends to be optimized at 14 Hz +- 0.5 Hz, but the sample so far has been small. This may be seen as a derivative of Lubar's and Tansey's reinforcement of SMR at Cz. Driven forward by such success, we have also looked at synchrony training at frontal sites, where frontal midline theta presents an inviting target. Jay Gunkelman has been recommending frontal midline theta rewards for some years now, and that kindled our interest. Indeed, synchrony training at 7 Hz using F3 and F4 placement has had salutary effects in some clients.

It is with synchrony training that the "standard bands" may find their full clinical utility. With bipolar training promoting engagement and activation, synchrony training moves in the direction of disengagement and of resting states. Access to both is required for good brain function. The preliminary findings with SMR synchrony training may (if confirmed) lead to a reinterpretation of Sterman's early results with seizure management. Clinical effectiveness may be attributable more to the desynchronization of the EEG from the use of bipolar placement at C3-T3 than to the focus on the SMR frequency specifically. Moreover Sterman's current emphasis on keeping rewards relatively rare, and mandating a refractory period after the reward, may be seen as an attempt to promote transient local SMR synchrony with a single-channel referential montage. The rewarded event must be clearly distinguishable from the ambient background.

General Self-Regulation and Specific Dysfunctions

Neurofeedback had to assert itself early on in a distinctly inhospitable professional environment. In response, researchers attempted to accommodate by being exceedingly conservative in their claims. Meanwhile, evidence was proliferating even then that the standard protocols were not specific treatments for either ADHD or seizure disorder but instead improved brain function in considerable generality. As feedback researchers were already being hounded by the placebo ghost, it would not serve to mention that neurofeedback was starting to look like a panacea as well. However, the refinement of protocol-based training discussed in this chapter has continued to enlarge the clinical scope of the work. Protocol-based neurofeedback was becoming a generalized approach toward improved self-regulation. The complementary approach of QEEG-based training appeared, for most applications, to be an unnecessary complication.

In many cases, specific functional deficits largely resolved even with the standard approaches. The first priority therefore was to do the best possible job with general self-regulation training. And in the sequencing of training objectives it is prudent to put the deepest and most basic dysregulations first, and to attend to more specific dysfunctions later, if indeed any are left to attend to. In conventional biofeedback, the recent finding that heart rate variability training is more effective in resolving myofascial pain syndrome and asthma susceptibility than more targeted approaches is yet another confirming instance of general self-regulation training trumping the more specific biofeedback approaches.

But sometimes specific deficits remain to be attended to, and QEEG analysis has been shown to be worthwhile in identifying suitable targets for training. In the refinement of QEEG-based approaches over the years there has been a gradual movement from amplitude-based training to the normalization of coherence relationships. This approach has been researched most thoroughly by Kirt Thornton, with a principal focus on specific learning disabilities and on traumatic brain injury. The general thrust of his approach is to work specific linkages that are identified with QEEG measurement under challenge conditions. Remediation is achieved predominantly with coherence up-training in the high-frequency regime (Thornton, 2007).

The Thornton method is nicely complementary to the other techniques so far discussed, which nearly all tend to have their greatest strengths at low EEG frequencies. A partitioning suggests itself that the low EEG frequencies govern the regulation of persistent states, whereas the high-frequency training impinges upon functions that are only episodically engaged in cognitive or other activity. The basic regulatory functions include arousal regulation; affect regulation; autonomic set-point and balance; motor system excitability; interoception; attentional repertoire and executive function; and working memory. The higher frequency training impinges more on the sensorium and on cognitive processes.

The principal hazard at low EEG frequencies is excessive coherence in the injured, the traumatized, the genetically disadvantaged, or the otherwise dysregulated brain. Such excess coherence at low frequency is often associated with brain instability, and hence with gross mental dysfunction, the more intractable psychopathologies, and behavioral volatility. By contrast, the principal hazard at high frequency is lack of task engagement, characterized by dropout of the expected coherence dynamics under challenge. This failure to function is often detrimental only to the affected person, and may be benign as far as the rest of the world is concerned. Hence it may be missed by caregivers, school personnel, mental health practitioners, and even by the person at issue.

Unsurprisingly, then, what populates mental health practices are the problems of general dysregulation, and these should indeed command our primary attentions. But a thorough-going application of all the tools of neuroregulation would include the more specific approaches deployed by Thornton and others.

Finally, coherence-normalization training is also used prominently by clinicians such as Jonathan Walker and Robert Coben (Walker, Kozlowski, and Lawson, 2007) (Coben, Padolsky, 2007). The work was pioneered by the late Joe Horvat (2007). Significantly, however, the focus here is typically on the largest deviations observed in baseline coherence measures, which tend to show up at low EEG frequencies. In contrast to the work of Thornton, the outcomes here are best described as improved self-regulation in considerable generality. Highly specific targeting, just as in resonant frequency training, does not imply narrowly targeted outcomes. The approaches of Horvat, Walker, and Coben therefore belong with nearly everything else into the bin of general self-regulation training, leaving Kirt Thornton's approach as the unique departure into the domain of localized dysfunction.

Summary and Conclusion

Whereas the modes of neuromodulation have undergone considerable proliferation over the years, there are unmistakable trends toward commonalities at a more basic level. From the original model of specific mechanisms-based protocols for limited diagnostic conditions the field has moved toward a more systems-based approach relying ever more heavily on the EEG itself to establish training objectives. From a focus on steady-state properties of the EEG there has been a shift to the brain's organization of transient, episodic states; from a focus on localized phenomena there has been a shift toward multi-site relationships; and from a focus on amplitude and comodulation disparities there has been a shift toward phase and coherence relationships. Effectively there has been a shift from the use of a basic set of static protocols with fixed thresholds to dynamic, open-ended training with adaptive thresholding and multiple targeting. Even more generally, a shift of emphasis from prescriptive to non-prescriptive training has been underway, and the shift in the balance is continuing from prescriptive toward proscriptive training.

Although different technologies and training philosophies each have their particular areas of strength, the overriding impression is one of considerable commonality in outcomes, irrespective of mode and only secondarily dependent on the particulars of targeting. This implies that our appeal is to a highly integrated regulation regime with procedures that mobilize global reorganization fairly generally. It is the integrated nature of neural network functioning and its hierarchical organization that allow such varied challenges to achieve clinical success.(Othmer, 2007)

The story is not complete, however, without reference to the precise targeting of deficits observed under challenge, an approach that has shown itself uniquely efficacious for specific cognitive deficits. Whereas most of our clinical applications involve more general enhancement of the client's self-regulation status for which a variety of methods may be availing, this targeted approach offers relief for more localized, more specifically cerebral deficits. As such, it also offers the most definitive evidence of specific results traceable to a specific intervention, filling a need in this era of evidence-based therapies.

The broad range of applicability of modes of neuromodulation suggests that good brain function depends on tight constraints in the timing of neuronal information transport, and that the failure of such precision in the domain of timing and frequency represents the dominant failure mode of the central nervous system, accounting for much of mental dysfunction. The remedy offered by the neuromodulation technologies ultimately offers the best evidence for the posited causal mechanism.

The minimal set of claims that support neurofeedback efficacy may be articulated as follows: Any sufficiently subtle disturbance of brain function, with respect to a variable that is under active management by the CNS, will evoke a response by the CNS that attempts to restore the desired set-point. The repetition of such a challenge will likely lead in time to improved self-regulatory capacity as a learned response. The evidence at hand suggests that the relative phase prevailing between two sites on the scalp represents a particularly attractive target for neurofeedback intervention because of its criticality to good brain communication, its dynamic features, and its ready accessibility.

The implication of the above is that neurofeedback training can be accomplished in an endless variety of ways, which makes it only too likely that the field will have to contend with continuing proliferation of methods. Nevertheless, the emerging techniques and procedures should still yield to simple ordering and classification. At this point there appears to be no single approach that covers all of the clinical bases. And in order to cover all of the bases, two requirements at least must be met in the prevailing state of the art. There must be some appeal in the frequency domain to the temporal organization of neuronal assemblies, and there must be some means in the time domain to recognize inappropriate changes of state, as these are revealed in the frequency-domain properties of the EEG. The first of these is some kind of narrowly prescriptive reward-based training, and the second is some kind of inhibit-based (i.e., proscriptive) detection scheme for inappropriate state shifts.

The latter of these can be considered a kind of error-correction scheme. The transient, sudden nature of this intervention implies that the brain is restricted here to its already existing resources. That is to say, the brain can only move toward states that are accessible to it, that are already available in state space. This is a necessary constituent of feedback for many conditions, but it is not sufficient by itself to restore optimum brain function for all. There must also be the opportunity for the brain to acquire new patterns of functioning, capacities that are gained incrementally and cumulatively over longer exposure times in training as the state space itself evolves under persistent challenge. This is accomplished by a targeted appeal to specific EEG frequencies. We are only at the beginning of the process of learning how to do this well.

References

Adrian, E. D., Matthews, B. H. C. (1934). The Berger rhythm: potential changes from the occipital lobes in man. Brain, 57, 355-384.

American Psychiatric Association (1994). Diagnostic and Statistical Manual of Mental Disorders, Fourth Edition.

Berger, H. (1929). Über das Elektroenkephalogramm des Menschen, Arch. Psychiatr. Nervenkr., 87, 527-570.

Coben, R., Podolsky, I. (2007). Assessment-Guided Neurofeedback for Autistic Spectrum Disorder, J. Neurotherapy, 11(1), 5-23.

Crane, R. A. (2007). Infinite Potential: A Neurofeedback Pioneer Looks Back and Ahead. In Handbook of Neurofeedback, Evans, J. R. (Ed.), The Haworth Press, 3-21.

Fehmi, L. G., (2007). Multichannel EEG Phase Synchrony Training and Verbally Guided Attention Training for Disorders of Attention. Handbook of Neurofeedback, Evans, J. R. (Ed.), Haworth Press, 301-314.

Frederick, J. A., Lubar, J. F., Rasey, H.W., et al. (1999). Effects of 18.5 Hz auditory and visual stimulation on EEG amplitude at the vertex. J. Neurotherapy, 3(3/4), 23-28.

Hardt, J. V., Kamiya, J. (1976). Anxiety change through electroencephalographic alpha feedback seen only in high anxiety subjects. Science, 201, 79-81.

Hardt, J. V. (1978). A Dedicated Microcomputer for Multi-Channel, Multi-Subject Biofeedback. Proceedings, Biofeedback Society of America, 9th Annual Meeting, 9, 6-8.

Horvat, J. J. (2007). Coherence and the Quirks of Coherence/Phase training: A Clinical Perspective. In Handbook of Neurofeedback, Evans, J. R. (Ed.), Haworth Press, 213-227.

Ibric, V. L., Davis, C.J. (2007). The ROSHI in Neurofeedback. In Handbook of Neurofeedback, The Haworth Press, 185-211.

Kamiya, J., (1968). Conscious control of brain waves, Psychology Today, 1, 56-60 (1968, November).

Kropp, P., Siniatchkin, M., Gerber, W-D (2002). On the Pathophysiology of Migraine---Links for “Empirically Based Treatment” with Neurofeedback, Applied Psychophysiology and Biofeedback, 27(3), 203-213.

Nash, J. K. (2000). Treatment of Attention Deficit Hyperactivity Disorder with Neurotherapy, Clinical Electroencephalography, 31(1), 30-37.

Ochs, L., (2006). The Low Energy Neurofeedback System (LENS): Theory, Background, and Introduction, J. Neurotherapy, 10(2/3), p. 5-40.

Othmer, S., Othmer, S.F., Kaiser, D.A. (1999). EEG Biofeedback: An Emerging Model for Its Global Efficacy, in Introduction to Quantitative EEG and Neurofeedback, James R. Evans and Andrew Abarbanel, editors, Academic Press, San Diego, 243-310.

Othmer, S. (2007). Implications of Network Models for Neurofeedback. In Handbook of Neurofeedback, Evans, J. R. (Ed.), The Haworth Press, 25-60.

Othmer, S.F., Othmer, S. (2007). Interhemispheric EEG training: Clinical Experience and Conceptual Models, in Handbook of Neurofeedback, Evans, J. R. (Ed.), The Haworth Press, 109-136.

Putman, J.A., Othmer, S.F., Othmer, S., Pollock, V.E. (2005). TOVA Results Following Inter-Hemispheric Bipolar Training, J. Neurotherapy, 9(1), 27-36.

Putman, J.A., Othmer, S. (2006). Phase Sensitivity of Bipolar EEG Training Protocols, J. Neurotherapy, 10(1), 73-79.

Scott, W. C., Kaiser, D. A., Othmer, S., Sideroff, S. I. (2005). Effects of an EEG Biofeedback Protocol on a Mixed Substance Abusing Population, American Journal of Drug and Alcohol Abuse, 31(3), 455-469.

Siever, D. (2007). Audio-Visual Entrainment: History, Physiology, and Clinical Studies, in Handbook of Neurofeedback, Evans, J. R. (Ed.), The Haworth Press, 155-183.

Sterman, M. B., Wyrwicka, W., Howe, R. (1969). Behavioral and neurophysiological studies of the sensorimotor rhythm in the cat, Electroencephalography and Clinical Neurophysiology, 27, 678-9.

Sterman, M.B. (2000). Basic Concepts and Clinical Findings in the Treatment of Seizure Disorders with EEG Operant Conditioning, Clinical Electroencephalography, 31(1), 45-55.

Strehl, U., Leins, U., Goth, G., et al. (2006). Self-regulation of slow cortical potentials: a new treatment for children with attention-deficit/hyperactivity disorder, Pediatrics, 118, 1530-40.

Thornton, K. E. (2006). No Child Left Behind Goals (and more) are obtainable with the NeuroCognitive Approach, Vol. 1, BookSurge Publishing Co.

Walker, J. E., Kozlowski, G. P., Lawson, R. (2007). A Modular Activation/Coherence Approach to Evaluating Clinical/QEEG Correlations and for Guiding Neurofeedback Training: Modular Insufficiencies, Modular Excesses, Disconnections, and Hyperconnections, J. Neurotherapy, 11(1), 25-44.