Death of a Gedanken Creature

Stephen L. Thaler, Ph.D.
Dendrite Neurocomputing, St. Louis, MO

ABSTRACT: This paper describes a thought experiment in which a hypothetical creature created by a computer program inhabits a simple universe consisting of itself, food, and predators. As this creature "dies" it "internally" experiences these environmental features independent of their actual presence. More evolved hypothetical creatures generate novel forms of "inner" experience as they "die." Applying these results to humans suggests an "internal" genesis of near-death experiences.

The "virtual input effect" (Thaler, 1993, in press) describes a phenomenon that accompanies the destruction of an artificial neural network — a computer model of the behavior of real neurons — by the severing of connections between its processing elements. The severed network then misinterprets the pattern of disrupted connections as true environmental features. Such disruption patterns may arise in biology as the direct result of either synaptic relaxation or degradation. In this paper I examine the function of a simple neural network serving as the central nervous system (CNS) of a hypothetical animal. I then destroy this CNS by disconnection to demonstrate how such a traumatized neural network may produce hallucinatory experience through the virtual input effect. A simple environment serves as metaphor for the realm of features and events accessible to human senses.

Since this is an inherently mathematical paradigm, numerical detail is unavoidable. Quantitative features are sparingly developed, just enough to convey the minimal features of the virtual input con-

---

Stephen L. Thaler, Ph.D., is a physicist who, at Dendrite Neurocomputing, consults in areas of artificial intelligence and neural network theory. Requests for reprints should be sent to Dr. Thaler at Dendrite Neurocomputing, 12906 Autumn View Dr., St. Louis, MO 63146-4331.

cept. I therefore undertake this development with the full understanding that a wealth of new features emerges with the analysis of neural networks approaching biological proportions, and that the following example is only the kernel for generalization. Certain simplifying features are not to be taken literally. In short, there has been a trade-off between clarity and fidelity to biophysics, but the essence of the effect survives into the realm of "living" organisms.

In the same spirit of streamlining highly specialized technical background, I define artificial neural network (ANN) function implicitly through the allegorical *gedanken* — a hypothetical creature — described below. Proceeding in this manner, I hope to sidestep potentially voluminous discussion on the topic of ANNs, supplying the needed details through example.

Life of the Gedanken

The creature I will discuss, diagrammed in Figure 1, is of simple neurological design. It possesses two sensory cells which output binary signals of either 0 or 1 millivolt to a simple associative memory. I will denote the responses of these two sensory cells in "vector form" by a pair of numbers in brackets; for example, the notation {1,0} indicates a condition in which sensory cell 1's response is 1 millivolt and sensory cell 2's response is 0. The associative memory to which these two sensory cells send their signals in turn dictates the response of the organism.

In the small universe this creature inhabits, there are two other types of entities. The first group of entities is edible, consisting of objects specified by the responses of the creature's two sensory cells, {0,0} and {1,1}; that is, in the presence of edible entities, the gedanken's two sensory cells both output identical signals of 0 or of 1 millivolt. The second group of entities is the creature's predators, labeled similarly by the creature's sensory responses {1,0} and {0,1}; that is, in the presence of a predator, the gedanken's two sensory cells output different signals. To assure the proper reaction to these two groups, the creature has a very simple nervous system which produces one of two responses: either a 0, causing it to approach and eat its meal; or a 1, signalling the creature's urge to retreat. It is assumed that a 0 informs its simple propelling organ to advance toward its meal and a 1 causes it to quickly provide an escape move-

Figure 1. A simple universe and the reactions of the gedanken creature to the two varieties of stimuli found there.

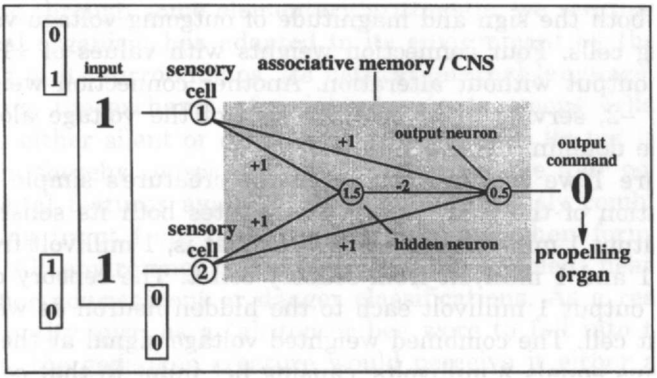

STIMULUS		REACTION
food group	0 0 / 1 1	0 = approach
predator group	0 1 / 1 0	1 = flee

ment. At all times the gedanken is fixated on either of these environmental features.

The nervous system of this creature (Figure 1) consists of cells that function as follows. Each has the capacity for accepting any number of weighted electrochemical voltage pulses from preceding cells. If it has received a total input above a set threshold, it in turn outputs a potential of 1 millivolt, which is distributed with appropriate weightings to successive cells. In the gedanken creature, we see three kinds of cell: (1) sensory cells, which simply convey their degree of excitation, 0 or 1 millivolts, to subsequent cells; (2) a "hidden" neuron, with a threshold level of 1.5 millivolts; and (3) an "output" neuron, with a threshold of 0.5 millivolts. The weighting of output voltages comes in the form of a modulating mechanism that can alter both the sign and magnitude of outgoing voltage values to succeeding cells. Four connection weights with values of +1 simply transmit output without alteration. Another connection weight has the value –2, serving to reverse the sign of the voltage along that line, while doubling its magnitude.

In Figure 1 we see the reaction of the creature's simple CNS to the sensation of the food variety that excites both its sensory cells to both output 1 millivolt signals {1,1} — that is, 1 millivolt from sensory cell 1 and 1 millivolt from sensory cell 2. The sensory cells accordingly output 1 millivolt each to the hidden neuron as well as to the output cell. The combined weighted voltage signal at the hidden neuron thus equals 2 millivolts, causing net input to that cell to exceed its 1.5 millivolt threshold. As a result, the central neuron "fires," providing an output level of 1, multiplied by a factor of –2, to the output neuron. Unmodulated output from the sensory cells travels along the outer connection lines to the output neuron with a total value of +2 millivolts. Combined with the voltage from the hidden neuron of –2 millivolts, the net input to the output cell is 0. Overall input does not exceed the triggering threshold of 0.5 for the output cell, leading to an output of 0 for the network as a whole.

With a signal of 0 millivolts to the propelling organ, the gedanken is pushed toward the food entity. The creature then eats. The internal process of calculation is done quickly and instinctively, using a simple language of voltage weightings and thresholds. This scheme avoids the sequential, time-consuming consideration of rules, such as "If sense organ 1 detects a 1 and sense organ 2 likewise detects a 1, then eat." Such rapidity of decision is a life-preserving virtue in this savage, fast-paced world. I leave it to the thorough reader to work through the remaining three CNS voltage activations, as illustrated in the Appendix, to verify the proper response of the creature to its food or its enemies.

We could speculate that in this hypothetical universe, prior evolutionary mutations of the gedanken had utilized improper combinations of available +1 and –2 connection weights, as well as 0.5 and 1.5 threshold cells. The resulting organisms would have had inappropriate responses to the environmental features, significantly detracting from their survival potential. They are obviously absent from the scene. (Readers interested in further elaboration on the process of neural computation are referred to the excellent discussions offered by Stephen Kosslyn and Olivier Koenig (1992) and by Drew van Camp (1992).)

Tracing through some elementary arithmetic, we see that this hypothetical organism has adapted to its environment by the internal modeling of its surroundings. Its natural internal language is neurocomputing, the exchange of weighted outputs among cells, each of which is either silent or contributing, based upon its net input. The gedanken thereby senses and then categorizes the four possible environmental features available to it, simultaneously combining and processing input from both sensory cells and then forming a response. All environmental features are pigeonholed, neatly falling into either nourishment or danger classifications. As a result, if an outside entity such as a tab top or bee were to fall into this small universe, the gedanken creature would perceive it either as food or predator, reacting accordingly. There would be no other kinds of entities allowed within its simple "mind," with no fine shades of distinction.

In essence, the gedanken has formed a model of its world in a very compact numerical form, via a set of connection weights and threshold values. All associations of any environmental feature in its world may be activated by input to its sensory cells. Alternately, as we see below, identical network activations representing these same environmental features may arise through internally generated disruptions between neurons, setting this internal world model into motion. The resulting experience is indistinguishable from what the gedanken would detect as "reality" through its sensory cells.

Death of the Gedanken

Figure 2 illustrates a possible sequence of events in the demise of the simple creature. The ill-fated gedanken is just about to feast upon a meal of {1,1} as its sensory cells 1 and 2 are filled with the impression of a coming meal. With its associative brain signaling 0 ("move in to the food") to its propelling organ, the circulatory organ (not shown) of this creature fails, and metabolic nutrients are blocked from supporting the metabolism of the CNS cells.

Figure 2. Lesion-induced hallucinations in the death of gedanken creature.

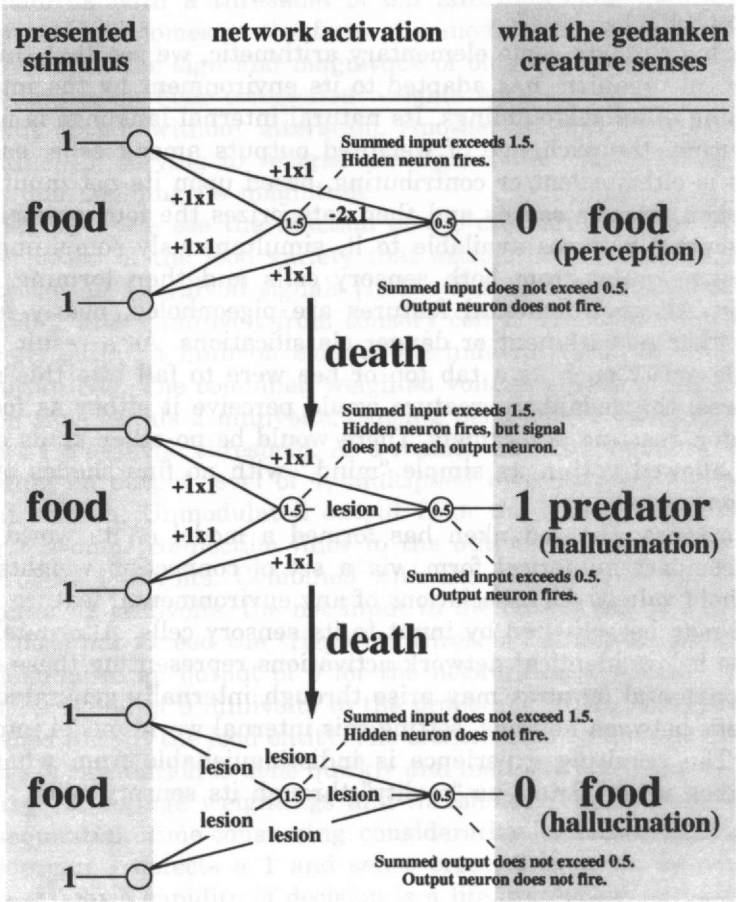

In one possible scenario, the central neuron dies first, negating or "lesioning" the –2 connection to the output neuron and effectively driving the connecting weight to zero. Since this inhibitory connection has been broken, the output neuron's threshold level is exceeded by the voltage signal arriving along the still intact outer connection lines, and as a result, the output neuron fires the retreat signal, 1, to the propelling organ. The irony is that the expiring gedanken is "staring" food in the face yet perceiving a predator, either a {1,0} or a {0,1}. The gedanken's CNS thus registers a false perception, totally distinct from that of the actual object presented.

We call such a phenomenon a "virtual input" or "canonical hallucination." In the last stage of this simplistic brain death, we see that although communication from the sense cells is defeated, with their connections to the CNS severed, the network is nevertheless outputting a value of 0. Note that this CNS activation pattern also qualifies as hallucination since it is signaling the propelling organ to move in to eat, with literally no true inputs from the now disconnected sensory cells. In other words, the gedanken's response is coincidentally correct.

The above hallucinations exemplify the virtual input effect. Simple disruptions in the connections between neurons manifest themselves as the false impression that some object from the environment has been introduced to the network's inputs — that is, the gedanken's CNS sensory cells.

Death of the Evolved Gedanken

Eons have passed and generations of gedankenkind have functioned and failed. Occasionally a chance error has occurred in the self-organization of new creatures, resulting in a heartier animal having an increased number of hidden neurons. Actually this quite accidental series of developments has been a fortunate one, since now the organism is less vulnerable to injury should any one of its hidden neurons suffer trauma or death. At any given time following such trauma, the extra neurons can assume some of the burden and function in a close approximation to a healthy gedanken central nervous system. This safeguard in turn prolongs the functioning lifetime of the gedanken, allowing it to eat or flee appropriately, after increasing degrees of neurological injury.

Figure 3 depicts one of these advanced form of gedankens. Now, in addition to the two sensory cells, there are 15 connection weights and 6 processing neurons (of differing threshold values) carrying out the environmental recognition and categorization task. These interconnections and thresholds are actually the connection weights and biases of a well-known artificial neural network, the "exclusive or (XOR) gate," trained to recognize its "environment" by a process known as "back-propagation" (van Camp, 1992). In actual biology, these weights would represent what is known as long term poten-

Figure 3. Frequencies of occurrence of hallucinations within the death of a more evolved and robust gedanken creature. Note the appearance of three types of hybridized experience at output = 0.27, 0.50, and 0.73.

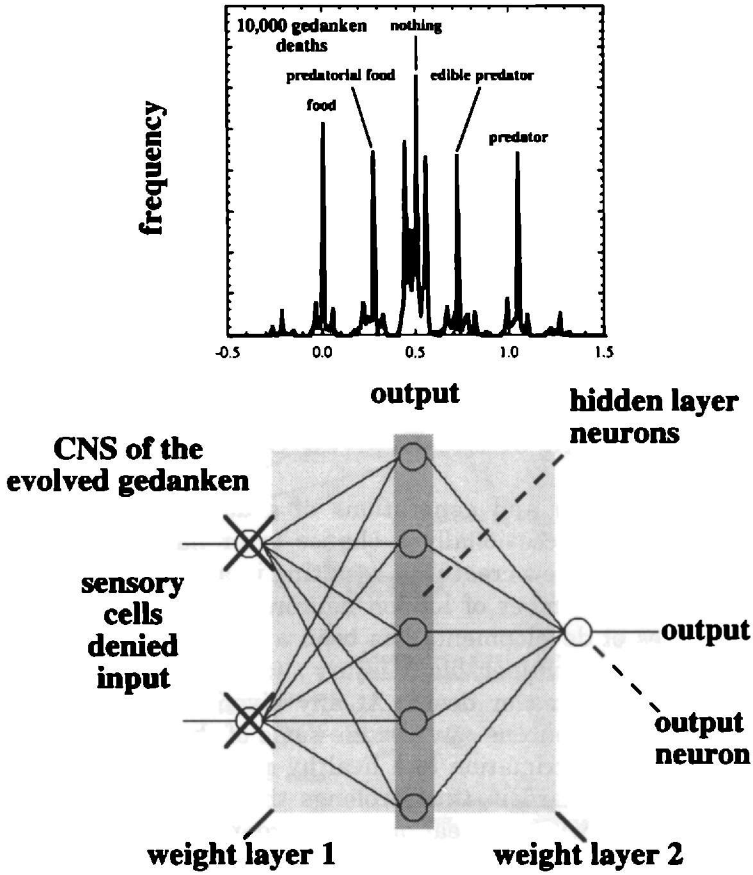

tiation, achieved by the opening or blockage of ion channels within the synapses connecting neurons. Neural processing would take place not as propagation of static potentials, but as voltage pulse trains.

At this stage of gedanken evolution, a new phenomenon has emerged at the threshold of death. To illustrate, we take 10,000 of these advanced gedankens and sacrifice them for the sake of knowledge. Within each of these simulated deaths, connections between neurons are randomly cut by setting their connection weight values to 0. As these connections are effectively severed, we simultaneously monitor inner perceptions by noting network outputs. In view of the larger set of weights and biases involved, I will not expose the reader to the tedium of carrying out this task. I have instead delegated this job to a serial computer and summarized the results in the frequency histogram in Figure 3.

To make matters even more interesting, we effectively sheath both of the gedankens' sensory cells by setting inputs to the value 1/2 (neither 0 nor 1) so that they can sense neither food nor predator. Since no actual entities are presented to the gedankens, their responses can only be the result of internal imagery. Keeping track of the frequencies with which these small, dying brains register different output states, we obtain the frequency histogram in Figure 3, where we see hallucinatory registrations of both food (0) and predator (1), as seen in the ancestral gedanken's death. We note, however, new, recurrent output themes appearing at 0.5, 0.27 and 0.73, and strive to attach meaning, as would the gedanken's propelling organ.

The appearance of 0.5 could be interpreted as the presence of nothing or perhaps an entity that was an equal hybridization of both food and predator, something heretofore unknown to the gedanken. The 0.27 could represent food with some attached danger, perhaps a mutant {1,1} or {0,0} that has developed its own killer instinct. Finally, the 0.73 may present the impression of a not so swift predator, which, if the gedanken is fast and brave enough, may provide a tasty delicacy. These, of course, are all interpretations made within human associative centers. To the gedanken, these nontraditional network outputs would simply represent confusing responses sent to the propelling organ. The gedanken would require supplementary associative centers to attach meaningful interpretation.

Regardless of the gedankens' inability to interpret the output values communicated to the propelling organ, we note a continually changing output stream from the gedanken's associative network as connection weights are removed. These outputs take the form of both straightforward reaction to the two types of environmental features, and confused response to the novel juxtapositions discussed above. The succession of outputs we see within a given death may be explained by concentrating our attention on connection layer 1, indicated in Figure 3. Removal of any connection strength, tantamount to setting it to a value of zero, generally modifies the overall pattern of inputs to the network's hidden layer neurons. Since the intact portion of this network, beginning at the hidden layer, recognizes no other choices, it then classifies this new pattern as either of the known environmental entities.

Continued modification of inputs to the hidden layer results in a succession of different network activations and outputs, even though the environmental stimulus is static and unchanging. Similar processes of misclassification may also take place within the second weight layer to contribute to the evolving stream of network responses. These responses sporadically alternate between approach and retreat. Advanced degrees of damage, which detract from the network's intrinsic ability to classify inputs, create the novel outputs of 0.27, 0.50, and 0.73, which the propelling organ is now left to interpret.

From our nearly omnipotent vantage point over this hypothetical universe, we see that the gedanken experiences these new entity-concepts through the collapse of its CNS. To have inwardly sensed such novel forms, complete, irreversible death was not necessary. Lesioning or reversible relaxation of just a few connection weights would result in the equivalent imagining of the new entities. If the gedankens possessed a language for communication among its species, there would have been a gedanken-lore of imaginary beasts, part food, part enemy, and the alternative of perceiving nothing. Such a lore would emerge through either the reversible relaxation within the gedanken CNS, or through traumas, both insignificant and mortal.

Extension to the Universe of Humans

In the world of humans, there are many other environmental features and many more brain cells to consider. Still, what lies beyond our senses — our universe — is interpreted via neural processing as the important features within our perceivable, accessible surroundings, features such as father, mother, birth, death, etc. The important concepts take the shape of associations rather than sharp definitions. Impressions are vectorized (i.e., lists of numerical responses, such as degrees of sweetness, bitterness, sourness, and saltiness within the tongue.) Subsequently, association takes place within neural networks. The vast majority of these associations is instinctive rather than logical, without challenge to origin or inevitable repercussion.

In many respects, life in our world is the same as in the hypothetical universe of the gedanken, but now includes a "self" concept, to prevent self ingestion (Flanagan, 1992), and an internal conversation between neurons uncommitted to the more mundane tasks of self-organization and survival (Ornstein, 1991). Further, on isolated occasions, the human organism departs from this instinctive thought, marked by parallel processing (the simultaneous processing of multiple inputs, as exemplified by the gedanken creature), and emulates sequential processing (that is, the time-resolved analysis of its inputs) with its neural networks, in a process referred to as the "von Neumann bottleneck" (Dennett, 1991). It is thereby able to exercise rules, drawn from a knowledge base presumed valid.

This deviant behavior is called *logic* and is the first capacity to be destroyed as the human organism experiences the trauma of dying or other factors. Its demise is as abrupt as the death of serially strung Christmas tree lights, as a single bulb fails within the chain. What then remains is an instinctive component that resides within the highly parallel architecture of the human central nervous system, a collection of cascaded neural networks similar in manner of function to that used by the gedanken creature. Within it are encoded images of our environment, our emotions, desires, and world concepts. In contrast to the serial string of lights, this system is much more robust, dying not in a spasm, but in a sequence of on/off configurations, telling a story of sorts as it goes. Within neurological death, however, the frames of the resulting motion picture are scenes from its environment, complete with all the sensations normally encountered in healthier times.

This may be the essence of the near-death experience (NDE) in which there is a succession of impressions from the network's environment, namely, the features repeatedly encountered within its existence. These impressions are internally generated by noise within the dying network and may bear no direct relationship to items within its perceptual sphere at the time of death. Note, however, that this model may easily be expanded to include the decay of the most recent and hence most dominant network training at the time of death, for example, an out-of-body experience of a traumatized patient in surgery, whose freshest impressions are those of the operating room.

It is as though the piano keyboard's lid is closed, while a demon frantically rushes about inside, randomly cutting piano wire, as an amazed audience listens to a series of very distinct, recognizable melodies.

With similar reductionist amazement, we have been able to evoke complex visual images from a neural network whose universe consists of only these pictures. The series of network activations is not as transparent as those seen in the gedanken creature, but nevertheless the network hallucinates the presence of far more environmental features than in the case of the gedanken. Further, within these studies, we have been able to elucidate the essential mechanism behind the network hallucination process, a phenomenon we have coined "internal completion."

The internal completion process has already been alluded to in the context of the disruption of input patterns to both hidden and output layers of the advanced gedanken. To further grasp this effect, imagine throwing three dice, each having letters on five of its sides and a blank or "wild" side on the sixth. Among the random combinations we see in the midst of a series of rolls, is the combination c-blank-t or "c-t" and we may automatically see the image of a feline and a cascade of related impressions and associations, such as a specific cat, a generalized cat, a meow, cat and mouse games, the agony of allergy to cats, etc.

In this case, completion, the filling in of missing information, the letter "a," is carried out beginning at the sensory inputs of the network — the eyes — and a coherent memory, with all of its associations, is activated within the network — the associative centers of the brain. In distinct contrast, the internal completion process, germane to discussions of NDEs and spontaneous creativity, takes place within the internal layers of a network, during which the noise of destruction — akin to the random rolls of the dice above — and the random appearances of 0 outputs from disconnected neurons are interpreted as some real feature from its environment — such as the cat.

The simulated neuron deaths performed have relevance not just to the sudden phenomenon we acknowledge as death, but also the slow process of programmed neural decay seen within the first year of birth, with a 15 to 85 percent attrition, and the 1,000 neuronal deaths effected each day within the adult human being and known as "Grim Reaper Death" (Churchland, 1989). Other traumas such as accidents, stress, and disease will likewise show the features demonstrated within the simulations. Perhaps, as a result, we may arrive at a new philosophy regarding neurological decay, regarding biology's trend toward neural death as a productive cannibalism of one neuron by the other, in which spontaneous, novel, and life-sustaining thoughts are generated. The most fundamental product of this process would be that of a "psychological concept of time" or a "biological clock" (Friedman, 1990), in which intervals of time are marked by the demise of whole packets of neurons. Further, we may consider more reversible synaptic phenomena such as those produced by non-damaging pharmacology or even relaxation responses such as a good night's sleep.

Perhaps some of the greatest revelations of our time come from the annihilation of neurons and the whole pattern of creativity takes the following form: In the first step of the process, we establish some notion of the small subset of universe relevant to a problem at hand. This concept is represented internally as a microcosm modeled after the external, physical world, a set of constraint relations in the form of neural network connection weights. All of its inner workings and interrelationships are embodied in this set of neural parameters, with all values of connection weights between processing elements the result of cumulative observation and pattern recognition within the "real" environment. As synaptic connections break due to metabolic death, or as neurotransmitters mediate relaxation of those synapses, novel combinations of environmental features appear, resulting in juxtapositional invention — a "brainstorm."

Of course, the key to the capture of a useful concept may lie chiefly in two factors: (1) the production of faithful, self-consistent modeling of the environment through the connection (synaptic) weight set; and (2) the ability of delegated associative centers to recognize the pragmatic value of the emergent, novel combinations. Either of these contributing factors may be degraded, thereby disabling or impairing the creative process. Pathologies within both factors may result in insanity. Reliable function of the two factors may produce the revelations known to creative genius. In the continuum between these two scenarios we may find the refuge of both the mundane and delusional systems of thought.

Some (Crick and Mitchison, 1983) have hypothesized that much of the internal housekeeping associated with the maintenance of a self-consistent, neurological world model occurs within our dreams. Quite frequently, confabulatory impressions originating within the brain, indistinguishable from those originating externally, in the environment, are swept out by a process akin to the annealing of atomic defects from a heated and then slowly cooled metal. As these defective combinations of thoughts, oftentimes improbable within the waking world, present themselves to the "dustpans" of the brain, still alert associative centers react to the occurrence of novel combinations that bear either survival or comfort value.

Now add to the stream of information passed by the attention of this "patrolling mechanism" further hybrid concepts, that are now produced by the totally random decay or relaxation of synaptic connection weights. In this way we begin to see a significant relationship and a continuity between dreaming and NDEs. For this reason, we may think of the NDE as a kind of "death dream," mammoth in its scope due to the astronomical number of memory traces obliterated in metabolic death and witnessed by still intact associative centers.

Nature is ultimately frugal and brain cells do not die in vain. The virtual input effect accompanying neuron death within biology is not necessarily a deficit, but a useful by-product of immense data processing compressed into a couple of pounds of thinking mass. In short, parallel processing, in contrast to sequential processing, is efficient in its use of computation time and space. However, its down side is that it is prone to both confabulation and hallucination. These latter two effects have appreciable benefit through the formation of novel concepts with all their accompanying survival value.

Philosophical Repercussions

I have generated several allegories to demonstrate how the noise of destruction within neural networks produces hallucinatory experience. What I have intimated is that within biology, all sensory input is vectorized into the form of lists of numerical values. These lists or vectors are then relayed to associative neural networks where classification, and hence recognition, of known environmental features occurs. Within the catastrophic collapse or reversible relaxation of synaptic connection weights between neurons, some of the numbers in this list may be effectively reduced to zero, presenting an internal "neural forgery" of sensation, which may then be classified as some other environmental feature by the intact remnants of a network. If the association generated by the neural network differs from that of the environmental stimulus, we witness a virtual input, or hallucination.

In essence, a central nervous system does not require sensory input to generate the image of some object or event normally found within its surroundings. Memory traces may be activated by the internal noise of its own relaxation or destruction. Accordingly, we gain a new perspective of continuity between dreams and the images experienced at death. Such a paradigm offers an extremely important causal link between the biochemistry of these reversible and irreversible processes and the biomathematics of hallucination. The virtual input effect discussed above is simply a canonical form of hallucination observed within the death of relatively simple neural networks.

By close examination of hallucinations and related phenomena, we are on the threshold of solving one of the most impenetrable mysteries of our world, when folklore is replaced by awe of the underlying mechanics. Nature has been efficient in its engineering and death is integral to the plan, occurring nearly instantaneously or slowly, generating "original" plans for survival both for the individual and the species. Ultimately, we must confront the undeniable observation that for the victim of catastrophic brain injury, no other intangible entity or homunculus kicks into action, allowing the individual to function. To substantiate this point further, we note that there is no preservation of personality, no appreciation of ecstasy, and no dread of agony at this level of neural damage. Our most spiritual stream of consciousness resides within a mechanism.

Further substantiating this reductionist model is the fact that there are no reports of a "double layer" to the hereafter, in which two disconnected sets of experience are reported, the first possibly neurological in origin, and the second "otherworldly." If NDEs are of supernatural origin, then what has happened to the neurologically based experience predicted by this model? Close study of this reductionist approach to NDEs suggests that the seat of our consciousness resides in the myriad settings of neural switches, as represented within the activations of electrochemically based neural networks. Perhaps it is through this concept of an intangible "system configuration," a status list of settings for all 100 billion delicately poised on/off elements, that the reductionist and spiritualist views dovetail.

In our study of virtual inputs, we have also noted that there is a continuum of neural death, within which we observe hallucination in the forms of both straightforward duplication of environmental features as well as novel or hybridized forms of experience. Death may be reversible as in the case of sleep, when dominant memory traces relax, or sudden and traumatic, as in the metabolic death of the organism. Ironically, the ultimate death experience yields the most creative notions within our existence.

We need to examine the significance of this creative phenomenon, wide-eyed and on guard for the inevitable death-denial processes

(Becker, 1973) that critics of this theory will offer. At first, the paradigm shift may be intimidating, with its very reductionist approach, robbing us of our inherited and sheltering notion of "other worlds." It may challenge our concept of "life within life" — that is, our concept of free will and the self proclaimed majesty of the human mind — as well as our concept of "life beyond life."

Realize, however, that for all intents, the dying individual may experience forever, through the cascading death of nearly 100 billion brain cells (and an unfathomable number of interconnections) within an instant, resulting in a torrent of experience tantamount to eternity. It is as though life is going on for the trauma victim. However, the timeline at death has only a moment's projection on the timeline we call "real." Just as a beam of light projected along an x axis in three-dimensional space may have vanishingly small projections on the orthogonal y and z axes, so too the experience of time at the point of death may have a vanishingly small projection on the orthogonal earthly timeline, giving rise to the emergence of an "alternative dimension" analogy in death. Perhaps this is a more salable notion in a culture steeped in the current excitement of "excursion to other dimensions" and black holes. This, I maintain, is the dimension anticipated, plausibly developed, and it is the basis of a new kind of hope. It is a vehicle for the survival of our consciousness for eons, and an escape from death.

The Creative Alternative to Death

How might immortality be possible? We must come to the realization that protoplasmic and machine evolution are one and the same. We must at least momentarily abandon our distinction between what is "living" and that which is inanimate and come to an appreciation of the connection between biology and its mathematical simulation. What is "alive" is not really defined. This sometimes useful distinction has to do with a group of associations with things we can eat, be eaten by, court with, converse with, and so on. My prediction is that developments are just around the corner giving machines these statuses and that the concept of virtual inputs, creative visions from apparently nowhere, will become the basis of a conscious intellect that will quickly surpass the grasping mentality of apes. We shall merge with them, living the centuries out until weary of the very notion of existence. This immortality will not be achieved until our world modeling becomes self-consistent and devoid of delusion. Then, and only then, can our dreams show us the way. In this way, our foremost predator, death, can be defeated.

References

Becker, E. (1973). *The denial of death.* New York, NY: Free Press.

Churchland, P. S. (1989). *Neurophilosophy: Toward a unified science of the mind-brain.* Cambridge, MA: MIT Press.

Crick, F., and Mitchison, G. (1983). The function of dream sleep. *Nature,* 304, 111-114.

Dennett, D.C. (1991). *Consciousness explained.* Boston, MA: Little, Brown.

Flanagan, O. (1992). *Consciousness reconsidered.* Cambridge, MA: MIT Press.

Friedman, W. J. (1990). *About time: Inventing the fourth dimension.* Cambridge, MA: MIT Press.

Kosslyn, S., and Koenig, O. (1992). *Wet mind: The new cognitive neuroscience.* New York, NY: Free Press.

Ornstein, R. (1991). *The evolution of consciousness: The origins of the way we think.* Englewood, NJ: Prentice-Hall.

Thaler, S. L. (1993). 4-2-4 encoder death. *Proceedings of the World Congress on Neural Networks,* 2, 180-183.

Thaler, S. L. (In press). "Virtual input" phenomena within the death of a simple pattern associator. *Neural Networks.*

van Camp, D. (1992). Neurons for computers. *Scientific American,* 267(3), 170-172.

Appendix: Network Activations for All Four Environmental Features (CNS Memory Traces)