

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY
and
CENTER FOR BIOLOGICAL INFORMATION PROCESSING
WHITAKER COLLEGE

A.I. Memo No. 973
C.B.I.P. Paper 18

June 1987

SYNAPSES THAT COMPUTE MOTION

T. Poggio and C. Koch

ABSTRACT: Biophysics of computation is a new field that attempts to characterize the role in information processing of the several biophysical mechanisms in neurons, synapses and membranes that have been uncovered in recent years. In this article, we review a synaptic mechanism, based on the interaction between excitation and silent inhibition, that implements a veto-like operation. Synapses of this type may underlie direction selectivity to direction of motion in the vertebrate retina.

© Massachusetts Institute of Technology (1987)

This report describes research done at the Massachusetts Institute of Technology within the Artificial Intelligence Laboratory and the Center for Biological Information Processing. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Army contract number DACA76-85-C-0010, and in part by DARPA under Office of Naval Research contract N00014-85-K-0124. Support for the Center is provided in part by the Sloan Foundation, and in part by a grant from the Office of Naval Research, Engineering Psychology Division. This work first appeared as an article in *Scientific American*, vol. 255, 46-52, May, 1987. Some of the figures referred to can be found in the original article; the others are left to the imagination of the reader.

The brain is a computing machine. However, though we know perfectly well how computers work, we know very little of how brains work. The hardware components of our present computers are tiny, interconnected electrical circuits, each one consisting of a few resistors, diodes and transistors, constructed by etching minute grooves in chips of silicon. There are several different types of such circuits, each performing an elementary operation such as logical NOR or NAND. These are some of the logical operations that can be performed among two binary variables. All arithmetic operations, such as multiplication and addition, can be written in terms of NOR and NAND operations. Thus the circuits that perform these NOR and NAND operations are the building blocks of our computers. Adders, registers and memory cells are made out of these building blocks that perform the elementary information processing operations. We understand very well how these elementary components work, on which physical principles they are based, and for which computations they can be used; after all, we have designed and built these computers.

The situation is different when one considers the most powerful and sophisticated computer on earth, the human brain. What are the elementary mechanisms used by this complex piece of “wetware” to process information? What are the operations they perform? Are there biological equivalents of transistors and NOR and NAND gates, and if yes, what are they?

We are far from knowing the answer to those questions. It is quite clear, however, that some old hypotheses are in need of replacement, because they represent a vast oversimplification of the variety of the information processing mechanisms in brains. Consider first what we know about the brain components. Brains are made of specialized cells, called nerve cells or neurons. There are about 10^{12} (a thousand billion) neurons in the human brain. A typical neuron consists of a cell body containing the biochemical machinery supporting the cell, a number of fibrous branches, called dendrites, and the axon. In general terms, the dendrites receive the incoming signals; the axon extends away from the cell body and provides the pathway over which electrical signals are transmitted to connecting neurons. Signals inside a cell are mostly electrical signals, but they are transmitted from cell to cell by special molecules that are secreted across a specialized contact, the synapse. Typically, neurons in the cortex receive thousands of inputs from other cells and may provide hundreds of outputs to other cells. Most neurons can generate nerve impulses—brief changes in the internal potential—and transmit them for long distances in the axons. These impulses are all-or-nothing, that is, either a “standard” impulse of a fixed amplitude, called an action potential or spike, will be generated, or nothing at all. Early on, scientists realized that the mechanism of spike generation could be used to process information. In 1940, McCullough and Pitts at the Massachusetts Institute of Technology described a simplified formal model of the neuron as the basic information processing unit of the brain. In their view, the neuron is a linear threshold device. It summates algebraically in the cell body the electrical signals generated by the input cells at the synapses—both the excitatory (positive) and inhibitory (negative) signals. The model neuron then generates an action potential if this sum exceeds

a certain threshold; otherwise it is silent. It is easy to show—as McCullough and Pitts did—that all logical operations can be synthesized by circuits consisting of such units. In principle, therefore, one can build a universal computer—a Turing machine—using these idealized neurons as building blocks.

Despite the widespread use of this simplified model, it has been clear for many years that neurons are much more complex than simple linear threshold devices. Recent work has shown that many neurons do not have axons and do not generate classical action potentials. Some cells generate local all-or-nothing impulses in their dendrites. On the other hand, the nervous membrane of the cell body, axon and dendrites turns out to contain a large number of different ionic channels that provide numerous sophisticated properties, such as the ability to transduce chemical into electrical signals, to adapt to long-lasting stimuli and to generate and propagate different types of electrical signals. The major aim of our research has therefore been to try to understand what type of neuronal operations can be implemented by this large class of biophysical (and biochemical) mechanisms. It is clear that the generation and propagation of action potentials plays a major role in neuronal information processing, but we do not know which role, and we neither know nor understand what other mechanisms underly the complex transformation from sensory input to motor output occurring in all nervous systems. Thus, we study the **Biophysics of Computation** to better understand the brain.

A simple but nontrivial operation we studied — detecting the direction of a moving stimulus by single neurons — shows that the key biophysical mechanism underlying this computation is likely to rely on a synaptic mechanism and not on the spike threshold of the cell. A large number of neurons in the visual system of all animals are sensitive to the direction of motion of a visual stimulus. Moving, for instance, a bar of light in the *preferred* direction elicits a vigorous response from the cell, while motion in the opposite direction—the *null* direction—yields no response. In 1964, Horace Barlow and Bill Levick, working at the University of California at Berkeley, found that certain ganglion cells in the rabbit’s retina are direction selective and proceeded to characterize their properties. Ganglion cells are only three cell layers removed from the photoreceptors, where light is first transduced into electrical signals. Their axons, making up the optic nerve, are the only output of the retina. Since Barlow and Levick used the so-called extracellular recording technique, that is the electrode recording the electrical activity of the ganglion cell is close to but outside the cell, they could only detect the occurrence of action potentials, but not the effect of the incoming signals on the internal potential of the cell. Even with these limitations, their experiments suggest that sensitivity to the direction of motion is achieved by some kind of *veto* operation. The crucial experiment which led to this conclusion stimulates the direction selective cell with two slits, placed close to each other. Barlow and Levick first measured the response of the cell when each bar was turned on and off in isolation. Subsequently, they recorded the response of the cell to “apparent motion”: the two bars are turned on and off in a sequence corresponding to the imaginary movement of an object

(in the null or preferred direction) across the cell. In each case examined, the combined response to the flashing of the individual bars was at least twice as large as the response to apparent motion in the null direction, indicating that a null direction stimulus is rejected by some kind of veto operation. Their conclusion, from this and several other experiments, was that an inhibitory signal from the photoreceptors is delayed in time before interacting with the excitatory signal from the photoreceptor. If motion is in the preferred direction, the excitatory channel is stimulated first and, due to the delay, the inhibitory signal will reach the gate too late to veto the excitatory signal. If motion is in the opposite direction, the signals from the two channels reach the cell at about the same time and inhibition will cancel the excitatory signal. Moreover, Barlow and Levick's experiments suggested that the veto operation cannot be performed at the cell body of the direction selective ganglion cell since motion within small regions or subunits of the receptive field of the cell is sufficient to elicit an appropriate response. In other words, the veto mechanism must be replicated a number of times throughout the entire receptive field of the direction selective cells. What then is the biophysical mechanism responsible for the veto operation?

In 1978, Vincent Torre of the University of Genoa and Tomaso Poggio, then at the Max Planck Institut für biologische Kybernetik in Tübingen, suggested a synaptic mechanism as the basis for the veto operation. At a synapse, an electrical signal in the presynaptic cell induces the release of a chemical substance, the neurotransmitter. The neurotransmitter diffuses across the small space separating the pre- from the post-synaptic cell and binds to specialized receptors, located in the membrane of the postsynaptic cell, typically in the dendrites where most synapses are located. This binding results in the opening or closing of ionic channels or pores in the membrane. To understand the effects of these channels on the internal potential of a cell, we have to describe the ionic environment of a neuron. Like all cells, the neuron maintains within itself a fluid whose concentration of positive potassium ions is substantially higher than in the external or extracellular medium. Exactly the opposite is true for the positive sodium and calcium ions and for the negatively charged chloride ions, which occur in higher concentration outside the cell. As a consequence of the different composition of the internal and external fluids and of the differential permeability of the membrane to these ions, at rest the inside of a neuron is more negative than the outside by about -60 to -90 millivolts (mV). This potential is usually called resting potential. We will always refer to voltages as relative to the resting potential. Ionic channels selective to specific ions are similar to openings in the membrane through which only certain ions, but not others, can go, following their concentration and electrical gradient. A simple electrical equivalent is a battery — with its voltage depending on the ion concentration on both sides of the membrane — in series with a variable resistance which is controlled by the state of the channel: when the channel is open, the value of the resistance is low, and it is very high when the channel is closed. The difference in the distribution of ions is maintained by another membrane specialization — ionic pumps — that spend metabolic energy to pump sodium and calcium out of the cell and chloride and potassium back in. In a sense they

continuously charge the ionic batteries that drive electrical signaling in neurons.

An excitatory synapse will open a channel specific for an ion such as sodium which can be thought of as being connected to a battery with a positive potential, called synaptic reversal potential, of $+80mV$ relative to the resting potential. Thus, if an “excitatory” neurotransmitter such as Acetylcholine, or ACh, binds to the receptor and increases the membrane conductance to sodium, a current is generated which is proportional to the product of the conductance change and the reversal potential. This is simply Ohm’s law. This outflow of positively charged ions leads to an excitatory postsynaptic potential, or EPSP for short, which is an increase in the intracellular potential. Hyperpolarizing inhibitory synapses behave in a complementary way: they open a channel that lets positive ions, potassium, get inside the cell by increasing the membrane conductance for potassium. Since the synaptic reversal potential for potassium is usually -20 to $-40mV$ with regard to the resting potential, increasing the potassium conductance makes the inside of the cell more negative, or hyperpolarizes the cell. The resulting change in potential is commonly called inhibitory postsynaptic potential, or IPSP.

There is, however, another type of synaptic inhibition called *silent* or *shunting* inhibition, common in the nervous system. The reversal potential of this synapse, increasing the conductance to chloride, is at about the same potential as the resting potential of the cell. Thus, if this synapse is activated alone nothing will happen: there will be no flow of ions across the membrane since there is no voltage difference which could drive any ions. An analogy with water pipes, though not literal, can help understand the difference between depolarizing excitation and hyperpolarizing and silent inhibition. In the case of excitation, the synapse is similar to a valve that gates a pipe coming into the neuron with a higher water pressure. Opening the valve will inject water into the neuron. The water is analogous to electrical charges and the pressure to voltage differences. Hyperpolarizing inhibition corresponds to a gate that connects the neuron’s inside to a pipe with a lower water pressure. Opening the gate will draw water from the inside of the neuron, thereby lowering the pressure inside the cell. Silent inhibition corresponds to a gate in a pipe with the same pressure as the inside of the neuron. Thus, opening or closing the gate will not induce any flow of water.

What then is the function of silent inhibition? Consider a dendritic branch of a neuron that receives an excitatory input next to an inhibitory input of the silent type. Activation of the excitatory gate will inject water into the cell, thereby increasing the internal pressure above the normal resting value. The inhibitory gate will then experience a pressure differential between inside and outside: water will flow outside, reducing the increase in pressure. Thus, activation of silent inhibition is somewhat similar to opening a hole in the membrane. If the hole is large enough and close to the location of the excitatory synapse, it will cancel almost completely the increase in pressure due to the excitatory gate. From this analogy one can see that the effect of a silent inhibitor is somewhat equivalent to a “veto”

or, if one likes the language of digital circuit designers, to an AND-NOT gate: the output of this gate is high only if excitation and no inhibition is present. It can also be regarded as an approximation of a multiplication. Intuitively, this is because the effect of inhibition is present only if there is a non-zero excitatory conductance change, and the effectiveness of inhibition increases with increasing excitation.

Next, we have to address the question of the optimal location of inhibition with respect to excitation, optimal in the sense of maximally reducing the EPSP. For this, we have to take into account the complex morphology of nerve cells with their dendritic trees. Theoretical arguments show that the most effective location for silent inhibition is either at the same location as excitation or some place between the excitation site and the cell body of the cell. The exact best location depends on several parameters, such as the precise morphology of the neuron, the membrane resistance and so on. The important point, however, is that any location on the *direct path* between the excitation site and the cell body yields a quite effective veto effect, whereas locations outside the direct path are relatively ineffective.

To study these properties, we have developed together with Patrick O'Donnell of the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology, computer programs that simulate the electrical properties of a neuron. Anatomical data about the shape of the neuron, the length and diameter of the branches etc. can be fed into a computer, which will then simulate the electrical properties of the dendritic tree. The neuron itself is represented as a large collection of cylindrical cables and synapses. The analysis of current flow in such structures, termed one-dimensional cable theory, was developed by Lord Kelvin for application to transatlantic telephone transmission lines. Its application to the spread of potential in dendritic trees was pioneered by Wilfrid Rall at the Mathematical Research Branch of the National Institute of Health. Thus, our cell is represented in terms of many, many elementary electrical components, like resistors, capacitors and batteries. In order to compute the potential in the cell in response to synaptic input, we use a very popular circuit simulation program, SPICE, developed over 10 years ago at the University of California at Berkeley. This program was first used by Idan Segev, James Fleshman and John Miller, working at National Institutes of Health, to model motoneurons. Our graphic interface to SPICE allows us to plot the timecourse of the transmembrane potential in response to massive synaptic input in color. Simulations of this type show that our proposed veto mechanism can be very specific for relative location of excitation and inhibition. If inhibition is only 20 or 30 μm (one μm is a thousandth of a millimeter) behind excitation — and therefore not on the direct path — its effectiveness in counteracting excitation is reduced many fold. If silent inhibition is on a neighboring branch, removed by 5 or 10 μm from the direct path, it will not be able to reduce the depolarization effectively. Most of these simulations are carried out with the assumption that the neuronal membrane is passive and does not contain any active, amplifying elements. However, these properties also hold true if the dendrites support action potential: in order to block the propagation of an action potential, silent inhibition must be on its direct path to the cell body.

The combination of excitation and silent inhibition seems to provide an important elementary operation underlying information processing. Addition of signals can be easily performed by the integrative properties of the neurons which effectively summate positive and negative inputs at excitatory and hyperpolarizing synapses. The mechanism of silent inhibition and excitation is functionally more similar to a multiplication than to an addition. Next we will consider an example of how this mechanism may be used to compute the direction of motion.

Previously, Torre and Poggio of us had suggested that direction selectivity in retinal ganglion cells relies on the veto mechanism of silent inhibition. The idea has several attractive properties. The veto mechanism requires only a few synapses and a patch of membrane. It can be replicated numerous times in the dendritic tree of the direction selective cell, as required by the physiological evidence for subunits. It is consistent with the experimental data of Barlow and Levick. The model makes several specific predictions. The main prediction concerns the existence of two different types of synaptic inputs to the direction selective cell, one excitatory, the other inhibitory with an equilibrium potential close or equal to the resting potential of the cell. Another prediction concerns the morphology of the cell that should have excitatory synapses on fine and highly branched dendrites, in order to maximize the strength of the veto effect.

How does the theoretical model fare against experimental evidence? Within recent years it has been possible to record from the inside of direction selective cells. Recordings from these cells in the turtle and the bullfrog retina support the existence of the two types of synaptic inputs expected by the model. Silent inhibition, invisible when activated alone, was revealed by first making the cell's potential more positive by injecting a positive current through the recording electrode. Under these conditions, the silent inhibitory input becomes visible as a hyperpolarization, that is, a negative contribution to the intracellular potential. Moreover, it is now possible to identify the anatomy of direction selective ganglion cells. Ralph Jensen and Bob DeVoe, then at John Hopkins University, injected a fluorescent dye into turtle direction selective ganglion cells, and Frank Amthor, Clyde Oyster and Ellen Takahashi at the University of Alabama at Birmingham visualized the rabbit directional selective cells Barlow and Levick had recorded from about twenty years earlier. After appropriate histological processing, the cells can be visualized and drawn. As our model predicted, these cells have highly branched dendrites with very fine processes, maximizing the number of different sites where our synaptic logic can operate.

Furthermore, it has been shown that the amino acid γ -amino-butyric acid (GABA) is the inhibitory neurotransmitter mediating direction selectivity. GABA is a very common inhibitory neurotransmitter in the central nervous system. When the action of GABA in the retina is blocked with the help of various drugs, the formerly direction selective ganglion cell now responds to movement in both directions. Using similar techniques, Richard Masland at the Harvard Medical School, and Mike Ariel and Ariel Adolph at the Retina Foundation

in Boston have provided evidence that Acetylcholine is the excitatory neurotransmitter underlying motion discrimination in the retina of the turtle and rabbit. One prediction that still has to be verified concerns the relative location of excitatory and inhibitory synapses in the dendritic tree of the direction selective cell. Because of the on-the-path property of silent inhibition, inhibitory synapses should be either close to the excitatory ones or between the excitatory synapses and the cell body. Using modern anatomical staining methods, it is now possible to attempt to label the excitatory synapses using ACh, and the GABAergic inhibitory synapses on a stained direction selective ganglion cell. The results of this experiment, begun by S. Zucker and N. Grzywacz in the laboratory of John Dowling at Harvard, may be critical for confirming or disproving our model.

Other models for direction selectivity are possible, and though less likely, are not fully ruled out by available data. The most obvious possibility is a membrane operation such as threshold performed on a linear combination of excitation and inhibition by a cell or a synapse presynaptic to the ganglion cells. This possibility, known for several years, is now being explored in detail by N. Grzywacz. It is also conceivable that excitation and silent inhibition interact on a dendrite of a cell presynaptic to the ganglion cell. One possible candidate for such a cell is the so-called starburst amacrine cell studied by Masland (see a previous issue of *Scientific American*). This cell appears to be the only cell in the retina synthesizing and secreting ACh, thus most likely providing the excitatory input to direction selective ganglion cells. The dendrites of this amacrine cell, which owes its name to its appearance of an exploding star, with its dendrites forming the traces of the explosion, are likely to be electrically isolated from each other. Thus, excitation and inhibition could interact on each dendrite separately. This model predicts that release of transmitter from the starburst amacrine cell would already be direction selective itself. Still other possibilities are opened by the recent findings of DeVoe and collaborators, who found for the first time evidence for direction selective properties in retinal cells different from ganglion cells. This may mean that motion is also computed in other retinal locations as well as within ganglion cells.

A combination of theory, modelling and experiments is thus about to unravel the mechanisms underlying this simple but nontrivial computation performed by neurons. Why is this so interesting? We would, of course, like to understand how direction sensitive neurons work. They have an important function in the first stage of vision, since motion permeates the visual world and gives us essential cues for visual recognition, for guiding our eyes, and for alerting us to potentially dangerous situations. Much more important, however, would be the unravelling of an elementary mechanism used by the nervous system to process information. If the excitation - silent inhibition scheme is used to compute motion in the retina, it is quite likely that it is also used in many other parts of the nervous system. In the meantime, a number of different researchers have proposed that the veto operation underlies the detection of motion discontinuities by the visual system of the housefly, the computation of binocular depth and of expanding and contracting patterns in the visual cortex in the mon-

key, and the computation of the direction of motion across the body surface occurring in the somatosensory monkey cortex. The idea that a veto-like operation may have an important role for computation in the central nervous system is not new. Barlow has stressed before that a veto-like operation may be an important physiological mechanism for the processes that underly perception.

The synaptic mechanism itself is quite attractive as one of the elementary computational elements of the nervous system. It couples the morphology of a cell and its synaptic architecture to the operations that are performed. In other words, different cells with different dendritic trees can perform different operations using the same veto mechanism. The main attraction of the new mechanism is that it only requires a few synapses and a patch of dendritic membrane. It can be replicated many times in a single neuron, in more or less independent subunits. Unlike action potential generation, it is a very local operation; and, since cortical cells receive many thousands of synapses, hundreds of AND-NOT like operations can be carried out in a single cell. To caricaturize the situation, one can think of the old McCullough and Pitts model as equating a neuron with a single transistor, while our model suggests that neurons can be more similar to a silicon chip with many hundreds of analog gates.

However, the generation of action potential and the interaction between synaptic inputs are not the only biophysical mechanisms underlying neuronal operations. A plethora of biophysical phenomena exists to be exploited by the nervous system to process, propagate and store information. In our laboratory, we have studied a variety of different mechanisms implementing operations like temporal differentiation, modification of the functional connectivity between neurons, electrical resonant filtering, and gain control based on the blockage of two specific groups of channels by neurotransmitters. Very specific circuits, termed *microcircuits* by Gordon Shepherd of Yale University, such as dendritic spines, synaptic triades or reciprocal synapses, can subserve neuronal operations within highly localized regions of the dendritic tree. An important task that now confronts both experimentalists and theoreticians is to identify specific biophysical mechanisms underlying neuronal operations and to characterize their role in information processing. In this way, we will eventually learn how to read from the cellular morphology and the synaptic architecture of a given neuronal circuit the operations that are performed in the ultimate dream of the anatomist. Our understanding of the information processing hardware of the brain may even lead, at some point in the future, to the development of new and different types of computing systems, more similar to the delicate and intricate tissues of the brain than to present day computers.

FIGURE LEGENDS

Figure 1: CIRCUITRY COMPUTING MOTION in an artificial and a natural system is demonstrated here. The top half of the figure shows a close up view of a Very-Large-Scale-Integrated (VLSI) circuit developed by Carver Mead, John Tanner and Misha Mahowald at the California Institute of Technology to compute the velocity field from the images of a single rigidly moving object. The chip is build in cMOS technology, using only analog components. Its architecture bears some resemblance to the organization of the vertebrate retina. The distance between the green lines is a tenth of a millimeter. The little red squares are the photoreceptors. About 10 by 20 elementary cells, each cell sensing light and computing one component of the velocity, are shown here. The bottom half shows an electron micrograph taken by Charles Zucker and Norberto Grzywacz at Harvard University of synapses made by amacrine and bipolar cell processes onto horseradish peroxidase filled ganglion cell processes in the rabbit retina. In the left panel, a bipolar cell terminal is making a ribbon synapse (arrow head) onto processes of both an amacrine (A) and ganglion cell (G). This dyadic arrangement is typical of bipolar synapses. The amacrine cell process is making a conventional synaptic contact onto a second filled ganglion cell dendrite at the same time. In the right panel, another amacrine cell process is synapsing onto a ganglion cell dendrite filled with microtubules. The bar represents $0.5\mu m$. Synapses like these may compute the direction of a moving stimulus in retinal ganglion cells.

Figure 2: SYNAPTIC VERSUS SILICON LOGIC is illustrated in this figure. The right figure shows our proposed synaptic logic, based on the nonlinear interaction within a small part of a nerve cell between an excitatory and an inhibitory synapse. One could characterize this operation as an analog version of a digital AND-NOT gate implemented in the CNS (central nervous system). The left figure shows a NAND gate implemented in standard nMOS technology, one of the techniques that etches specific types of transistors in silicon, common on today's digital processors. The spatial dimensions are roughly comparable.

Figure 3: DIRECTION SELECTIVE CELL will respond to movement in one, the preferred, direction but not at all to movement in the opposite, the null, direction. In the null direction, the membrane potential at the cell body — recorded through an electrode indicated in grey — fails to reach threshold and the cell will not generate action potentials. In the preferred direction, the cell discharges vigorously. The bottom figure shows the morphology of a direction selective ganglion cell in the rabbit retina. The cell was recovered by Frank Amthor of the University of Alabama at Birmingham, after the dye horseradish peroxidase was injected into the cell. Its axon (seen leaving the bottom of the cell) leaves the retina. The dendritic tree of the cell is subdivided into two layers, receiving synaptic input from presynaptic cell signaling an increase (On layer) and a decrease (Off layer) in

light amplitude. Both dendritic layers are stacked one upon the other within the plane of the retina. Notice the slender dendrites (less than one thousandth of a millimeter thick).

Figure 4: MODELS OF DIRECTION SELECTIVITY as proposed by different group of researchers. The first scheme (left figure), developed over 30 years ago by Werner Reichardt and Bernhard Hassenstein working on the visual system of the beetle, may also partly underly human motion perception as recent psychophysical experiments indicate. It assumes that inputs from two channels are multiplied after one signal is low-passed (or delayed). Ten years later Horace Barlow and Bill Levick proposed a model for the rabbit retina (middle figure) that can be shown to be a functional approximation of the first one. It assumes that the signal from one channel is delayed before interacting with the signal from the second channel in an AND-NOT like manner. The system signals movement only if the delayed channel carries no signal, that is if the stimulus moves from left to right (in this example). Both schemes assume some sort of delay between the two channels. Thus, in agreement with experimental data, if the stimulus moves very slowly or very fast, both models respond in a similar manner to movement in both the preferred and null direction. Vincent Torre and Tomaso Poggio proposed in 1978 one particular biophysical mechanism, based on the interaction between excitatory and inhibitory synapses, which could represent a good implementation of the scheme of Barlow and Levick (and Hassenstein and Reichardt).

Figure 5: INTERACTION BETWEEN SYNAPSES is illustrated here for two different kinds of synaptic inhibition. The circuit we have used for these simulations mimicks a small patch of dendritic membrane and consists of a resistance in parallel with a capacity and two synapses. The battery of one synapse (the excitatory one) is fixed at $80mV$ while the battery of the second synapse is either set to $0mV$ (silent inhibition) or to $-30mV$ (hyperpolarizing synapse). The top row shows the transient increases in conductance for the excitatory ($g_e(t)$; top left) and the inhibitory ($g_i(t)$; top right) synapse. $g_e(t)$ is delayed by $5msec$ with respect to the onset of g_i . The membrane potential induced by each synapses firing in isolation is plotted in the middle row. The leftmost figure shows V_e (in black) and V_{i1} (in red), that is the EPSP and the IPSP induced by the excitatory and the silent inhibitory synapse. Note that activation of the silent inhibition by itself does not lead to a change in potential, in contrast to the activation of the hyperpolarizing synapse (V_{i2} ; left trace in rightmost figure). Finally, the potential in the presence of both excitation and inhibition is illustrated on the bottom row. If the interaction between synapses were linear, then the potential for excitation and silent inhibition (V_{e+i1} ; bottom left) should be identical to the EPSP caused by the excitatory synapse alone, while activation of excitation and hyperpolarizing inhibition (bottom right) should yield a V_{e+i2} slightly negative. The action of silent inhibition can be thought of as a multiplication while hyperpolarizing acts somewhat more like a simple addition.

Figure 6: ELECTRICAL CABLE MODEL OF A BRANCHING DENDRITE is shown here. A patch of dendritic membrane is modeled by a capacity c_m in parallel with a resistance r_m , while the intracellular cytoplasm is described by a single resistance r_i . An excitatory and an inhibitory synapse are shown to be represented by a fixed battery with a variable conductance. When a synapse is active, the conductances increases temporarily, leading to a change in the potential. If the battery potential is negative, the synaptic induced conductance change will lead to an hyperpolarization; if it is positive, to a depolarization. In the case of a silent inhibition, the battery potential is zero (relative to the resting potential of the cell). Thus, if the cell is at rest, no change in potential will be seen upon activation of silent inhibition.

Figure 7: SIMULATION OF THE POTENTIAL IN A DIRECTION SELECTIVE CELL receiving inputs from 61 excitatory and inhibitory synapses is plotted here in color. The morphology of the computer reconstructed cell is taken from the identified directional selective rabbit retinal ganglion cell stained by Frank Amthor and shown earlier. Part of the axon can be seen leaving the image at the left in each frame. The cell body is the larger segment next to the axon. Note that only the longitudinal dimensions, but not the thicknesses of the cable, are drawn to scale. For our SPICE simulations, we assumed that the cell body contained active sodium and potassium channels like those found in the squid axon, able to generate action potentials. The color code used to plot the change in intracellular potential, relative to the resting state of the cell, is shown at the bottom of each image: high depolarization are represented by “hot” colors while little excitation or even hyperpolarization is represented by “cooler” colors. The silent inhibitory synapses (marked by circles) are always placed, in accordance with our on-the-path rule, within the neighborhood of excitatory synapses (triangles) or on the path between the excitatory synapses and the cell body. The colors within the synaptic symbols code for the relative conductance change (similar to the voltage scale). Thus, a white synapse is inactive, a green one is barely activated and a red one is maximally activated. The top sequence shows the intracellular events occurring upon a simulated movement of a bar from the bottom to the top of the figure. The synapses are activated in the measure as the bar moves across the cell. For this sequence, inhibition is always delayed by $20msec$ with regard to excitation. Thus, the EPSPs can propagate to the cell body, where an action potential is initiated. In the top sequence of six images, mimicking movement in the preferred direction, a small group of excitatory and silent inhibitory synapses would correspond to the subunits of Barlow and Levick. If the delay between excitation and inhibition is reduced to zero (bottom sequence of 4 images), inhibition effectively the veto’s excitation, preventing the cell from firing.

Figure 8: INTRACELLULAR RECORDING FROM A DIRECTIONAL SELECTIVE GANGLION CELL from the turtle retina. The top left (resp. right) image shows the response to movement in the preferred (resp. null) direction. The data is from an experiment

carried out by P. Marchiafava in Pisa. Note that in the null direction, there is no hyperpolarization but only a reduction in the EPSP, which fails to trigger an action potential. The bottom panel shows the intracellular response to movement in the preferred and null direction at the cell body of our computer reconstructed rabbit retinal ganglion cell.

TABLE LEGEND

Table 1: BIOPHYSICS OF COMPUTATION tries to understand the biophysical mechanisms underlying neuronal operations, and to explore the limitations of these operations, as they are implemented in the brain. Understanding the properties and limitations of these mechanism gives us important constraints when trying to understand higher processes like perception and thinking. We have listed at left some biophysical mechanisms which we have studied in recent years. The middle column gives the appropriate neuronal operation implemented by these mechanisms, while the right-hand column cites a particular example where this mechanism is believed to play an important role.

This blank page was inserted to preserve pagination.

**CS-TR Scanning Project
Document Control Form**

Date : 5/26/95

Report # AIM-973

Each of the following should be identified by a checkmark:

Originating Department:

- Artificial Intelligence Laboratory (AI)
- Laboratory for Computer Science (LCS)

Document Type:

- Technical Report (TR)
- Technical Memo (TM)
- Other: _____

Document Information

Number of pages: 13 (18-IMAGES)
Not to include DOD forms, printer instructions, etc... original pages only.

Originals are:

- Single-sided or
- Double-sided

Intended to be printed as :

- Single-sided or
- Double-sided

Print type:

- Typewriter
- Offset Press
- Laser Print
- InkJet Printer
- Unknown
- Other: _____

Check each if included with document:

- DOD Form
- Funding Agent Form
- Cover Page
- Spine
- Printers Notes
- Photo negatives
- Other: _____

Page Data:

Blank Pages (by page number): _____

Photographs/Tonal Material (by page number): _____

Other (note description/page number):

Description :	Page Number:
<u>IMAGE MAP (1) UN # 'ED TITLE PAGE</u>	
<u>(2-13) PAGES # 'ED 1-12</u>	
<u>(14-15) SCAN CONTROL, DOD</u>	
<u>(16-18) TARGETS (3)</u>	

Scanning Agent Signoff:

Date Received: 5/26/95 Date Scanned: 6/7/95 Date Returned: 6/8/95

Scanning Agent Signature: Michael W. Cook

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER 973	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER AD-A190389	
4. TITLE (and Subtitle) Synapses that Compute Motion		5. TYPE OF REPORT & PERIOD COVERED AI Memo/CBIP Paper	
		6. PERFORMING ORG. REPORT NUMBER	
7. AUTHOR(s) Tomaso Poggio & Christof Koch		8. CONTRACT OR GRANT NUMBER(s) N00014-85-K-0124	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Artificial Intelligence Laboratory 545 Technology Square Cambridge, MA 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Blvd. Arlington, VA 22209		12. REPORT DATE June 1987	
		13. NUMBER OF PAGES 13	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Arlington, VA 22217		15. SECURITY CLASS. (of this report) UNCLASSIFIED	
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Distribution is unlimited.			
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)			
18. SUPPLEMENTARY NOTES None			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Parallel processing Synapses Motion Biophysics of computation			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Biophysics of computation is a new field that attempts to characterize the role in information processing of the several biophysical mechanisms in neurons, synapses and membranes that have been uncovered in recent years. In this article, we review a synaptic mechanism, based on the interaction between excitation and silent inhibition, that implements a veto-like operation. Synapses of this type may underly direction selectivity to motion in the vertebrate retina.			

Scanning Agent Identification Target

Scanning of this document was supported in part by the **Corporation for National Research Initiatives**, using funds from the **Advanced Research Projects Agency** of the **United States Government** under Grant: **MDA972-92-J1029**.

The scanning agent for this project was the **Document Services** department of the **M.I.T. Libraries**. Technical support for this project was also provided by the **M.I.T. Laboratory for Computer Sciences**.

