System and Hierarchy

by

Peter Cerrato

Submitted in Partial Fulfullment
of the Requirements for the
Degree of Bachelor of Science
at the
Massachusetts Institute of Technology
May, 1983

Signature of Author...
Department of Computer Science
May 17, 1983

Certified by...
Prof. Marvin Minsky
Thesis Supervisor

Accepted by...
Prof. David Adler
Chairman, Departmental Committee on Theses

To Q : Thanks for the reinforcement!

Abstract

This thesis outlines a method for analyzing systems. The organization of the thesis follows the step by step process of the method. The thesis itself can be thought of as a system with its own hierarchical structures. The first section gives the top level abstract definitions of hierarchy and system. The second section studies these concepts in depth by analyzing their hierarchical structure through two examples: the robot control system and the human cerebral cortex. The third section reveals the thesis as a system. Horizontal links are built up between the two hierarchical studies of the robot control system and the human cortex. Feedback is introduced in the final discussion of the method of the thesis itself.

Table of Contents

The Top Level

The first section of my thesis sets out the top level definitions of hierarchy and system.

Hierarchy

A hierarchy is a structure with levels. These levels are ordered by degree of abstraction, the higher levels being those of a greater degree of abstraction.

Levels

Each level in a hierarchy contains elements that follow the same rules as other elements at the same level. An element is linked to a single element in the level above, its parent, and to some number of elements in the level below, its children. Elements at different levels are different in kind.

The notion of closure within a given level is what I am attempting to develop here. And in fact, in a strict hierarchy the mathematical concept of closure does apply: all the elements at a given level are members of the class of objects closed under a certain relation. This relation serves to draw a strong boundary between the given level and the levels above and below. These level defining relations give the rules that elements at the same level must all follow.

In physics these level defining relations would describe the forces of interaction at work among specific particles.

In a data-structure hierarchy these level defining relations would describe the types of the data-structures which all elements at a given level must be. These types would be entities in a strongly typed language.

Structure

A structure is a representation for a given body of knowledge that attempts to get at the skeleton which organizes the information into a coherant whole. To discover a structure is to uncover the major static forces of interaction at work in the body of information you are studying. In a hierarchy the underlying structure is often pictured as the trunk and main branches of a tree with the lowest levels being the leaves. The most important fact about a tree structure is that all its branches and leaves stem from one trunk. All elements in a hierarchy are constituents of the root element, all are focused on one goal or one leader.

In the real world no structure is ever a pure hierarchy in which no closed paths can be traced. Real world structures contain horizontal links which make them networks. When we come to the discussion of systems we will see that networks overlayed on a hierarchy are the structures that best represent systems. The closed paths in a network give rise to feedback which makes systems the dynamic animals they are.

Abstraction

We move from one level in a hierarchy to the next higher level by the process of abstraction. Abstraction is the process by which information is selectively thrown away. The key word here is selectively. The abstraction process is an information filter which removes detail to get at the underlying structure.

The ~~most basic example~~ an example of feedback used as a sensory mechanism is the phase-locked loop. A phase-locked loop performs exactly the function of synching the system's internal rhythms, the frequency of an occilator perhaps, to the frequency of some signal in the environment.

Before proceeding with the next example let me flesh out our model of feedback. To use feedback as a control mechanism the system must have two ingrediants. A model, or representation of the phenomenon to be controled in the environment; and a procedure which gives directions for measuring the difference between expected and actual values. This procedure is a dynamic entity, it must directly address the problems of interacting with the environment in real time.

The thermostat is an example of feedback being used to control a system as it changes its environment. The model of the environment is given by the temperature setting on the thermostat. It is a guess at the state the environment will assume after the system turns the heater on, it is the projected state.

It is important to recognize that the model can never correspond exactly with the environment. The measurement procedure must always allow some slack, a tolerance level, or else the system will occilate without ever reaching a stable state. The heater will constantly be turning on and off as the system tries to get the room to the exact temperature on the thermostat. This is an important fact about all kinds of models and representations of the environment: there is no such thing as a complete model, some information must be lost in the abstraction process to gain the necessary structuring and simplicity for the system to function in real time.

Focusing is an excellent way of thinking about abstraction because it implies a specific point of view from which one is looking, a reason or frame in which to do the abstraction.

In artificial intelligence, vision algorithms provide a good example of the process of abstraction. Vision algorithms take raw data consisting of intensity matrices and produce data structures that identify edges or corners. These new structures are abstractions of the raw data: much information has been thrown away. But from the point of view of finding the solid objects in the visual field of the camera we have created information at a level that did not exist before. We now know about edges instead of just intensities.

It is important to stress that the point of view or frame is necessary when unraveling the abstraction process. In the above example, if instead of solid objects we were hunting for point sources of light in the visual field then the edge information would be fairly useless. Abstraction always takes place with a purpose in mind; it is a motivated process.

<u>System</u>

Hierarchy is involved with the statics of objects and their interactions within a tree structure. System is involved with the dynamics of hierarchy. Systems evolve, systems interact with other systems, and systems interact with open-ended environments. A system changes over time and therefore must have mechanisms for adapting to changes in its environment. A system must be able to match its internal rhythms to the rhythms of change in its environment. Feedback provides the mechanism for adaptation.

In order to evolve a system must be able to act out the process of abstraction. This requires that self-knowledge be embedded in a system. A system must be able to focus on its own hierarchy. Horizontal links which are outside the strictly vertical links in a hierarchy create networks of self-knowledge in a system.

Feedback

This is the nature of feedback: the system makes an initial guess, a hypothesis about how the system expects its environment to behave. This hypothesis is made on the basis of what the system has learned from prior experience. Then the system takes some action that has an effect on its environment and makes a measurement of the change that occurs. The system compares the expected response with the actual response and determines by how much and in what direction its guess was off. This comparison is used to make a better guess and the process is repeated until a stable state is reached or the environment changes independently of the system.

Feedback is used as a sensory mechanism to gather information from the environment or as a control mechanism to guide the system as it performs complex tasks that require interaction with the environment.

Nonsense

Horizontal Links

We will approach the discussion of horizontal links by starting from the definition of a system as that which evolves. Evolution is a self-simplifying process. At each stage in the history of an evolving system old skills and abilities are combined into new higher level skills. New levels appear in the system's hierarchical structure. Because a new skill consists of a grouping of lower level skills the repertoire of skills at the new level will be smaller and simpler.

Self-simplifying evolution requires new information be forged in order to link elements or skills which are at the same level in the hierarchy. A new skill such as throwing a ball is built up out of lower level movement skills. Before the new skill evolved there was no communication among these lower level skills and in the hierarchy no links are provided between elements at the same level. But throwing a ball requires coordination among these lower level skills and therefore horizontal links must be made in the hierarchy, new networks of information flow must be created as part of the evolutionary process.

Two Examples

The first section of my thesis set out the top level definitions of system and hierarchy. The second section can be seen as the hierarchical section in that, through two examples a robot control system, and cortical structures in the human brain, we will explore the concepts of system and hierarchy in depth from two separate points of view.

Hierarchy and System in a Robot Control System

The discussion of the robot cotrol system is divided into four parts. The first part presents the triple hierarchy which provides the underlying structure for the control system. The second part presents the modes of operation of the control system; how the system interacts with its environment. The third part presents a thought experiment in which Sandi hunts the Wabbit. Sandi is a hypothetical robot consisting of a mobile tractor-tread base upon which are mounted robot arms and various sensor devices. The Wabbit is Sandi's nemesis, the human-guided prey for Sandi's hunt. The fourth part is an analysis of how the robot control system fits into the definitions of hierarchy and system as given in the first section of the thesis.

The Triple Hierarchy

The robot control system consists of three cross-coupled hierarchies: the World Picture Hiearchy, the Sensory Processing Hiearchy, and the Task Planning Hierarchy.

The World Picture Hierarchy is the control system's model of the environment it interacts with. The levels in this hierarchy proceed upwards from highly detailed quantitative representations of objects and parts of objects in the robot's world, to more abstract representations that allow the control system to understand the function of the objects in its world, and finally to representations which group together objects in order to capture the functional interactions of the objects.

There are four levels in this hierarchy. Starting at the lowest level and moving to the higher levels they are:

-- 1st level. Quantitative, detailed measurements of the robot and its environment: the room it moves about in, objects in the room, and other entities such as the Wabbit. All these are measured as precisely and completely as the robot's array of sensors will allow.

-- 2nd level. For each sensory modality: sonar, tactile, video, heat, IR, etc. the information from the 1st level is structured according to the requirements of the Sensory Processing Hierarchy which looks to the World Picture Hierarchy for its hypotheses. Visual information, for example, might be structured as edges or intensity gradients or both if both kinds of algorithms are being used.

-- 3rd level. Objects in the robot's environment are structured according to their function. A chair, for instance, would be represented as having legs and a seat and a back, where the system knows that legs support the chair, the seat can support an object of appropriate size, and the back is used when a person is sitting in the chair.

-- 4th level. Here you find information about how the robot interacts will the objects in its environment. This includes a history of past encounters and a list of possible uses for the object. The Task Planning Hierarchy has many links to this level to use when planning high level tasks. For instance the robot may have set a goal for itself to get to the top of a set of stairs in its room. The 4th level representation of a ramp in the room would include an entry stating that the ramp can be used to climb the stairs.

The Sensory Processing Hierarchy identifies objects and environmental conditions in the robot's world. Each level in the Sensory Processing Hierarchy consists of a series of algorithms each of which are driven by a feedback loop which accepts a hypothesis about what the control system can expect to find and compares this guess with data from the robots sensors. At the lowest level input to the Sensory Processing Hierarchy is raw data from the sensors processed by fast and dirty algorithms. Each higher level contains more focused and powerful algorithms which accept better guesses from the World Picture Hierarchy. Each level must filter out extraneous detail from the information in passes up to higher levels in a small scale version of the abstraction and evolutionary processes.

The algorithms in the Sensory Processing Hierarchy fall into levels which correspond to the levels in the World Picture Hierarchy. A vision algorithm that takes intensity levels from the sensors and computes where the edges are would be a 2nd level algorithm because edges would appear in the 2nd level of the World Picture Hierarchy.

The Task Planning Hierarchy sets top level goals for the control system and knows how to break down complex tasks into groups of simpler tasks. While the Task Planning Hierarchy is in control and attempting to accomplish some goal it is responsible for providing hypotheses to the Sensory Processing Hierarchy. These guesses will have been originally embedded in the plans stored in the Task Planning Hierarchy and they take the form of default assumptions about the state of the world after certain actions are taken by the robot. The Sensory Processing Hierarchy returns a measure of the error between the expected and perceived values. The Task Planning Hierarchy can then decide whether to backtrack and try a new strategy or to continue breaking up the plan into simpler pieces until the lowest level of robot motions is reached.

Levels in the Task Planning Hierarchy also correspond to the levels defined in the World Picture Hierarchy. Every plan will either be seeking information at a certain level or will be a heuristic for manipulation of information at a certain level. To grasp the distinction among levels concider a plan to find the distance to the door and a plan to find the shortest path to the door. The first plan would be a 1st level plan requiring only that the correct sensor be properly oriented and a measurement be taken. The second plan would be a 4th level plan which would involve knowledge about obstacles in the room, short cuts the robot had discovered at previous times, etc.. The first plan might even be a sub-plan of the second.

Modes of Operation

The robot control system has three major modes of operation: Action Mode, Sensory Mode, and Simulation Mode. Each Mode dictates which Hierarchy is the nexus of control at any one moment as well as the paths the flow of control may follow.

Action Mode: The Task Planning Hierarchy has selected a goal, move to the door, for example, and must now generate a multi-level plan whose lowest level will consist of basic commands to the robot's motors. The Task Planning Hieararchy must break down the top level goal into sub-goals which are the different parts of the robot's journey. The first layer of sub-goals might look like: move around the desk, look for the door, plan the shortest path to the door. As can be seen these sub-goals can be actions, or sensory scanning, or further planning. During an action such as moving around thedesk the Sensory Processing Hierarchy provides feedback to the Task Planning Hierarchy to tell it whether the goal is being accomplished. In this case the Sensory Processing Hierarchy would be focused on the desk, returning information gathered from its sensors and algorithms as to the desk's position relative to the robot. The World Picture Hierarchy provides the hypotheses, the expected values based on past experiences, which allow the Sensory Processing Hierarchy to focus on the desk.

In Action Mode the Task Planning Hierarchy must be prepared to backtrack and compute a new plan if feedback from the Sensory Processing Hierarchy suggests that a previously generated plan is impossible to achieve. In this example, after moving around the desk the door may be hidden from view and the Task Planning Hierarchy will need to come up with a new plan to find the door.

Sensory Mode: As seen above, the Task Planning Hierarchy can send the control system from Action Mode into Sensory Mode when it called for a plan to look for the door. This is Focused Sensory Mode because the Task

Planning Hierarchy tells the Sensory Processing Hierarchy what it is looking for.

At other times the robot will not be in the midst of any action but will instead be simply observing its environment. This is Unfocused Sensory Mode and it presents an interesting problem for the control system. The heart of this problem is how to generate the expected values or hypotheses needed by the Sensory Processing Hierarchy. In Focused Sensory Mode the control system has a goal, it is performing actions whose effect on the environment can be predicted. In Unfocused Sensory Mode it is the environment that is changing on its own.

To solve this problem the World Picture Hierarchy must contain high level knowledge about "normal" situations. This knowledge takes the form of defaults in the high level frames for "normal" situations. Such defaults might be: Joe comes on Tuesdays, the lights go on at 6AM, no one comes on Sundays, the Wabbit hides behind the desk. Armed with these default situations the Task Planning Hierarchy can go into Simulation Mode and generate possible sequences of events which can be sent as hypotheses to the Sensory Processing Hierarchy to be matched against incoming data. This is a highly dynamic process with defaults constantly being shuffled around until the proper frame is found for the situation at hand.

Simulation Mode: In this mode the Task Planning Hierarchy generates plans with are not to be executed as actual actions. Instead of the Sensory Processing

Hierarchy providing feedback to guide the execution of these actions the World Picture Hierarchy is consulted and asked to supply possible results for these simulated plans. In effect, the control system tries to imagine doing something without actually doing it.

Simulation Mode as well as providing default frames for the control system when it is in Unfocused Sensory Mode can be used as a learning process. By first simulating an action and then actually performing it the validity of parts of the World Picture can be checked. This is an example of feedback at a high level.

Sandi Hunts the Wabbit

Let us peek inside Sandi's control system when he goes Wabbit hunting.

First of all, Sandi's World Picture Hierarchy must contain multi-leveled knowledge of the *Wabbit*, of the *room* he is in, and *Sandi* himself.

Sandi's knowledge of the *Wabbit* fits the four levels outlined in the World Picture Hierarchy. At the 1st level he has precise measurements of the Wabbit to within a thousandth of an inch because he was once able to capture the Wabbit and make these measurements with his precision cameras which only work at close range.

The 2nd level consists of a rough structural representation of the Wabbit which can be fed to the vision algorithms of the Sensory Processing Hierarchy as an expected value for use in identifying the Wabbit at a distance. Such a description would not give the functions the Wabbit's various parts but only the separate parts which could be made out by low level vision algorithms: a blob which rests on the ground for the Wabbit's body, a blob which rests on top of this for the Wabbit's head, and two long skinny blobs above that for the Wabbit's ears.

At the 3rd level Sandi's has functional information about the Wabbit's body parts. The Wabbit's body can propel it at a certain speed Sandi has calculated in a previous encounter, the Wabbit's head can sense Sandi coming if it is in a certain orientation to Sandi's own position. At the moment we are not concerned with how Sandi learned these things, perhaps he had to be told to look for them by his human teachers, perhaps he was taught a high level skill called <u>hunting</u> which contained many defaults which could be tested and altered in various encounters with the Wabbit.

At the 4th level Sandi keeps a record of previous encounters with the Wabbit and the plans he used at those times to hunt the Wabbit. These he can use to help generate plans for a new encounter by looking for old encounters which fit the same patterns. Finally, there are the facts about hunting in general which comprise the above mentioned skill. These are stored as possible plans in conjunction with Sandi's Task Planning Hierarchy. Rules about stalking, optimum sensor scans, and cornering prey are stored in the World Picture Hierarchy.

Sandi's knowledge of the room he is in falls into four levels. At the 1st level Sandi has precise measurements of all the walls and obstacles. At the 2nd level he has blob descriptions, structural representations which are fed as guesses to the Sensory Processing Hierarchy to aid in identifying objects in the room. At the 3rd level he has procedural and functional knowledge about what each object is, and what he can do with it. The big box can be used to trap the wabbit, the door can be shut to keep the wabbit from escaping. At the 4th level he has actual plans for using the objects, plans for optimum movement paths in the room.

Sandi's knowledge of Sandi is highly complex for it is this knowledge that will allow Sandi to learn by changing his own World Picture Hierarchy. Sandi must have complete knowledge of his own physical specifications: the exact size of his body, how fast he can move and turn, the range and freedom of his arms, the range and precision of his sensors, his power consumption, and all the other quantitative information about his own physical capabilities. This is all 1st level knowledge. To get a better idea of what the other levels of self-knowledge look like let us watch Sandi in action hunting the Wabbit.

The top level plan for the Task Planning Hierarchy is capture the Wabbit. It might be possible to have an even higher level of plans which would motivate and initiate the capture the Wabbit plan. Such a level would be characterized by plans involving curiosity and exploration of Sandi's environment.

In any event, this is the situation: the Wabbit is sent into the room by its human controler. The Wabbit darts along behind obstacles and eventually hides behind a big box. Sandi's Task Planning Hierarchy recognizes this as a hunt senario, and the hunt down Wabbit sub-plan of capture the Wabbit is triggered when the Wabbit goes into hiding.

As long as the Wabbit remains in hiding Sandi's Task Planning Hierarchy has time to chose among the various hunting heuristics which appear as alternatives sub-plans to the hunt down Wabbit plan. Choosing the best heuristic involves sending the control system from Action Mode into Simulation Mode. Should Sandi make a frontal assault? The simulation must check if the Wabbit can escape by running to another obstacle, or if the room geometry given by the World Picture Hierarchy shows that the Wabbit is trapped. Should Sandi try to stalk the Wabbit by sneaking up on it? The simulation must check if Sandi's World Picture Hierarchy has enough information about the Wabbit's ability to sense Sandi coming to justify this heuristic. Perhaps waiting is the best choice of heuristic because Sandi has learned from previous encounters that the Wabbit never stays in one place for too long.

As you can see it is very important that the Wabbit be controlled according to a strict set of rules if Sandi is to be expected to discover patterns in its behavior.

The choice among those heuristics: frontal assault, or stalking, or waiting involves extensive simulation and gathering of information from various levels of the World Picture Hierarchy. The simulations proceed by following horizontal links from the Task Planning Hierarchy to the World Picture Hierarchy.

If, during these simulations, the Wabbit appears, then Sandi must make a quick decision. Now it becomes important for the Task Planning Hierarchy to be able to place upper bounds on the time required to finish a simulation. Real time constraints come into play. Decisions must be made based on incomplete or rough estimates. In the Sensory Processing Hierarchy all algorithms must have parameters which can be set by the Task Planning Hierarchy at the same time that it sends its guesses. These parameters control the speed versus accuracy trade-off and are set by the Task Planning Hierarchy because it is aware of the higher level concerns on the system.

Upon seeing the Wabbit move out from behind its box Sandi's control system makes a quick decision to intercept it based on two pieces of information it had gathered from simulations: the Wabbit often moves in a straight line and there is a cul-de-sac nearby. Sandi captures the Wabbit and the senario is over. Now Sandi must assimilate this new experience by comparing his old World Picture of the Wabbit's behavior with these new facts.

The learning process proceeds as a high level feedback loop: old models are measured against new experiences and the models are updated. Learning occurs at all levels:

Perhaps the Wabbit moved faster than it ever had before, or waited longer than expected before emerging. Perhaps certain hunting heuristics have never succeeded while others are seen to have worked given specific orientations of Sandi and the Wabbit. These and many other facets of the Wabbit-hunting experience can be examined in simulation mode by the robot control system as it constantly evolves and explores its environment.

Analysis

The robot control system is of course as of yet only a hypothetical model at this stage in the development of artificial intelligence methods and present-day computational power. It is true that I have exaggerated the intelligence of the system especially in describing the complexities and subtleties of Sandi's encounter with the Wabbit and his ability to learn and plan. But what is important to grasp here is the power of the triple hierarchy model as a framework for examining systems questions in a concrete domain: robotics. The four levels in the hierarchy do capture distinct and seperable levels of knowledge in the factual (World Picture Hierarchy), procedural (Sensory Processing Hierarchy), and planning (Task Planning Hierarchy) domains. It is the clarity of the boundaries between levels that gives a hierarchy its power to structure a system, to focus resourses where they are most needed at any given moment.

Those were hierarchical concidetions, let us move now to systems conciderations. In the robot control system the modes of operation capture the nature of interactions between hierarchies. What we see in the modes is the flow of control from one hierarchy to another, a flow that is always guided by the horizontal links between the hierarchies. It is these links seen as pathways of control that determine how intelligent and flexible the system will be. The mode the system is in at any one time is constantly changing as we have seen in the Sandi example: from Action Mode the Task Planning Hierarchy passes control to the Sensory Processing Hierarchy by going into Focused Sensory Mode, or to the World Picture Hierarchy by going into Simulation Mode in order to choose the best sub-goal to persue. It is the number and variety of horizontal links that facilitate smooth transitions from one Mode to another.

Finally, we saw a brief excursion into learning in the Sandi versus the Wabbit example. Learning is feedback in its most advanced form and it is therefore the most interesting aspect of an evolving system. Learning requires horizontal links between and within hierarchies to provide the necessary level of self-knowledge within the system to allow independant learning. In the Task Planning Hierarchy there must appear plans for learning, plans about plans, heuristics with links to all levels in all three hierarchies. It is here that we begin to get out of the realm of what is possible with today's level of artificial intelligence. The only systems we can honestly claim are independant learning systems are biological. In the next section of my thesis we will look at the human brain as our best example of a self-conscious, learning system and we will see that it is built on a carefully structured hierarchy with overlayed networks of horizontal links just as the robot control system is.

Hierarchy and System in Cortical Structures

To study the hierarchical structure of the brain we will examine the modular construction of the cortex. The structure in the cortex has four levels:

-- 1st level. *mini-columns*

-- 2nd level. *columns*

-- 3rd level. *entities* with inputs mainly from sources outside the cortex.

-- 4th level. *entities* with inputs mainly from other entities within the cortex, including themselves. Two important aspects of hierarchy are exhibited in these structures: modularity and isolation.

The systems aspects of the cortex are numerous and widespread. Horizontal links appear at all levels, columns having links to neighboring columns and entities being intimately connected to other entities in the cortex. These interconnections give rise to *distributed systems*. Feedback in the cortex goes by the name *phasic re-entry*. Outputs from one entity are cycled back as inputs to that entity itself and others in the same network of distibuted systems. Throughout the cortex external sensory inputs always appear alongside these internally generated signals. In this way the brain's internal self-image is continually updated in comparison with sensory information giving rise to the constantly changing focus of attention that is consciousness.

The quotes in the following section are from Mountcastle's keynote paper delivered to the *Neuroscience Research Program's Intensive Study Program* in June of 1977. The paper is titled "*An Organizing Principle* for Cerebral Function: The Unit Module and the Distributed System".

Hierarchy: Entities, Columns, and Mini-Columns

Modules (the columns of the cortex) are grouped into entities by virtue of:

-- a dominant extrinsic connection,

-- the need to replicate a common function over a topographic field, or:

-- certain intermodular interactions.

The large entities of the cortex perform seperate functions and they compose the top level of the brain's structural hierarchy. Let us see how each of the three forces at work in defining the boudaries of entities serve to also define the entities' functions.

A dominant extrinsic connection means that sensory input from a specific part of the thalamus, the switching center of the brain, or from a specific external source: the left eye, or the right hand for example, or from a specific part of the opposite hemisphere of the brain via the corpus callosum form a major portion of the inputs to a single entity.

The need to replicate a common function over a topographic field refers to the way sensory input is mapped onto topographically contiguous fields of columns in an entity. The retina makes a good example: input from adjacent points on the retina are sent to adjacent columns in the same entity.

The boundaries of an entity can also be drawn by examining the density of inter-columnar connections. Phasic re-entry requires that most columns in an entity have outputs which remain within the entity itself. Thus an entity will have more connections to itself that to neighboring entities.

These three critieria: external inputs, the mapping of sensory input, and feedback loops within entities are sufficient evidence that the entities of the cerebral cortex are distinct modules with seperable functions that form the top level in the stnctural hierarchy of the brain.

Entities are composed of modules called <u>columns</u> which form the next level down in the brain's hierarchy. The column is "defined in terms of the static and dynamic properties of its neurons." Statically, in sensory processing entities, the ones defined topographically, each column is responsible for one particular source point on the body and one particular sense modality. (Pain, heat, taste, sound frequency, sight, etc.) Dynamically, adjacent columns exhibit pericolumnar inhibition, "a powerful mechanism for the functional isolation of active columns from their neighbors." When sensory inputs activate one column it sends out signals via specially adapted neurons that inhibit neighboring columns. This mechanism allows the brain to filter out noise which arises from the large number of interconnections in the system and to focus its attention. Inhibition serves two purposes then: one, to strengthen the hierarchy by keeping columns isolated functionally from other columns and two, to provide a filtering and focusing mechanism necessary for unified consciousness.

Each column is composed of a package of several hundred mini-columns which are the "basic modular unit of the neo-cortex." A mini-column "contains about 110 cells and this figure is almost invarient between different neo-cortical areas and different species of mammals." These groups of vertically arranged cells are "heavily interconnected in the vertical axis running across cortical layers and sparsely connected horizontally." Thus at the lowest level the brain is extremely modular being composed of approximately 600 million such highly similar mini-columns each functioning as sperate units.

Now we shall study how this hierarchy of entities, columns, and mini-columns is organized into a system by the massive interconnectivity among areas of the brain and feedback within each entity to create the phenomena we know as consciousness and memory.

System: Horizontal Links, Feedback, and Filtering

Horizontal Links

If the brain consisted only of the structural hierarchy I have outlined in the previous section it would exhibit no signs of intelligence, it would simply be like a great game of Pichinko: sensory inputs appearing like so many silver balls at the appropriate entities and filtering down through columns until they arrived at the bottom, at the mini-columns where they would sit there due to feedback or simply fade away.

It is the interconnectivity of the brain which sets the massive network of parallel distributed systems that are mapped over the hierarchy vibrating. These distibuted systems in turn provide for all the features of the brain that we regard as intelligence: consciousness, memory, the ability to use symbols and metaphors, and the ability to learn. Let's start back at the structural level and build up to these larger concerns.

"The neo-cortex as a whole projects upon almost every other major entity of the nervous system: the basal ganglia, dorsal thalamus, mesencephalon, brain stem and spinal cord." Many of these form massive re-entrant systems with the neo-cortex. These are large scale feedback and feedforward loops, they are the drummers of the human brain. Their response times and cycling rates set the beat for our brains internal rhythms.

Within a single hemisphere there is a step-by-step outward progression of inter-cortical connections from primary sensory areas closest to the spinal cord onto succesively adjacent areas of the cortex. Each of these succesive higher-order convergant regions are reciprocally linked with areas of the frontal lobe. Those areas of the frontal lobe receiving the most highly convergant projections from sensory areas are linked via two-way connections with areas of the limbic lobe. These areas of the limbic lobe are responsible for coordinating a host of convergant connections and are the seats of higher order reasoning power and emotions. Interconnectivity gives rise to intelligence.

Or chaos.

In the past decade, discoveries in the study of the structure of the neo-cortex have turned up a vast amount of information concerning the extrinsic connectivity between large entities in the brain. These horizontal links in the structural hierarchy are now known to be far more numerous, selective, and specific than previously supposed.

Within an entity any one column will contain only a subset of all the entity's extrinsic connections. Thus the total set of columns in an entity is grouped into subsets, each linked bya particular pattern of connections to similarly segragated subsets in other entities. Thus an entity whose function is to coordinate hand and eye will have a subset of its columns with only connections to the eye, and an entity whose function it is to coordinate eye and ear will also have a subset of its columns devoted to the eye. These two subsets of columns form part of a distibuted system devoted to processing sensory input from the eye.

"Information flow through such a system may follow a number of different pathways, and the dominance of one path or another is a dynamic and changing property of the system." In this way the brain can evolve and learn as a system composed of distributed systems by finding the most efficient pathways for information that is to be used repeatedly (such as learning a new motor skill) or which is to be linked in a new way to information circling in feedback loops in other distributed systems.

Certain groups of columns serve as junction boxes that set up these interconnections between entities. Such columns provide for the distribution of certain inputs, sensory or intercortical, to other entities for further processing and they have a high rate of throughput. Very often a single neuron carries the signal to its final destination even if that destination is halfway across the brain.

In the preceeding section on hierarchy in the brain we discussed the mapping of sensory inputs onto topographically contiguous groups of columns. Via throughputs or junction boxes that link entities together into distributed systems two such topographically contiguous regions used in sensory processing can be mapped through a third region thus allowing for the integration of their functions. By learning to use these pathways of interconnectivity a child can make links between senses such as sight and touch, these being important learning steps in the childs development. The integrating regions also have inputs and outputs from other areas of the brain allowing learning and integrating to be watched and controlled.

Feedback

The massive amount interconnectivity in the brain, especially with regard to distributed systens and the integrating regions which inhabit the homotypical cortex (the area of the cortex devoted not to any one sense modality but to the interplay of senses, higher level functions related to intelligence), results in there being numerous sites of feedback in the brain. In neuroscience feedback is called re-entry.

The following quote explains how re-entry in the brain gives rise to consciousness, self-awareness in the complex system of the human brain.

> Phasic cycling of internally generated activity, accessing first primary sensory but then succesively more abstract and general processing units of the homotypical cortex, should allow a continual updating of the perceptual image of self and self-in-the-world as well as a matching function between that perceptual image and impinging external events. This internal readout of internally stored information, and its match with the neural replication of the external continuum, provides a mechanism for conscious awareness.

Consciosness is a dynamic and everchanging process by which high level perceptions arising in integrating regions are compared via feedback to memories and to new sensory information. Memory itself is a dynamic process much like consciousness. Memory consists of local feedback in the homotypical cortex where large scale perceptual patterns are constanly circling and being compared to new input from all over the brain. Memory is no longer to be thought of as static storage in the brain: it is much too complex a process to be explained statically or hierarchically, such a process which is evolving and reacts to external influences must be examined dynamically as a system.

Filtering and Degeneracy

Working our way down the structural hierarchy from the interconnection of large entities, to feedback in distributed systems, we now come to examine system conciderations in the homotypical cortex. Here we find degeneracy: the parallel processing nature of the brain. Degeneracy refers to the fact that any one sensory pattern can trigger activity in widespread areas of the brain.

The homotypical cortex is filled with integrating regions where low level sensory patterns that have already triggered specific mini-columns in the single-mode areas of the cortex (visual, auditory, somatic) are sent via the massive interconnectivity of the brain. Any one moment of sensory experience will trigger many integrating regions in the homotypical cortex. Some integrating regions will be looking for cross-sensory patterns, some will be comparing sensory input to feedback input from memories circling in the cortex, and some will be comparing current sensory input to what remains due to feedback of sensory input from the near past.

Degeneracy provides us with a model of the "curious system", the system that seeks to explore its environment and its own self image. When a sensory pattern resonates in the integrating regions and triggers many columns then we have discovered an "interesting" sensory pattern which will slowly become a memory as feedback from the many areas it has stimulated strengthen its presence in many places in the cortex.

Finally, at the lowest levels, those of columns and mini-columns, we find that pericolumnar inhibition serves as a filter system layed over the network of parallel distributed systems. Because of the massive amount of feedback in the cortex if all patterns were allowed to be constantly feeding back from one area of the brain to another the noise level would pose an insurmountable distraction to consciousness. Therefore the filter system simplifies the mass of signals by focusing attention on the center point of a stimulus. In the visual cortex this means we will notice the moving area against a still visual field because the columns around the point of movement will be inhibited. In the integrating regions of the homotypical cortex and areas responsible for reasoning we can see the limitations on short term memory as arising from the inhibiting influence a new stimulus has on those from the recent past still feeding back in the cortex.

This concludes the discussion of hierarchy and system in the cerebral cortex of the human brain.

This Thesis as a System

In the first two sections I have set out the top level of my thesis, the abstract definitions of hierarchy and system, and the hierarchical structure of my thesis, the two examples. The study of the robot control system and the human cortex gives depth to the abstract definitions from two different points of view, through laying out two independant hierarchies of knowledge. In this final section my thesis becomes a system.

First, I will create a few horizontal links between the two hierarchical examples. I will show that there are structural similarities between the robot control system's four level hierarchies and the four levels of structure in the human cortex. The robot control system and the human cortex also exhibit functional similarities as systems both using feedback, evolving, and learning about their environments and themselves through horizontal links in their hierarchical structures.

Second, I will create a feedback loop by discussing the method I have used in my thesis. The link I make from within the thesis to the thesis itself is a form of self-knowledge: the work contains information within itself about how it too functions as a system.

Horizontal Links: Robot Control System and Human Cortex

The see the structural similarities between the robot control system and the human cortex first note that both have hierarchical structures with four levels. In the robot control system I outlined the four levels in most detail when describing the World Picture Hierarchy. In the human cortex the fur levels were: mini-columns, columns, entities defined by extrinsic connections, and entities defined by connections within the cortex. Let me draw some parallels between each of these levels.

The 1st level of the World Picture Hierarchy contains quantitative, detailed information consisting mostly of raw data from the robot's sensors. Mini-columns serve the same function in the cortex: they are directly linked to sense organs if they are part of entities that do sensory processing.

The 2nd level of the World Picture Hierarchy serves to structure the information in the 1st level in such a way as to make it usable to the Sensory Processing Hierarchy's algorithms. In the cortex, we saw that columns are structured topographically according to the kind of processing their parent entity is doing. In visual processing entities, for example, columns map out the visual field in the cortex.

The 3rd level of the World Picture Hierarchy contains information about each object's function in the robot's environment. This is similar to the way entities with extrinsic links perform integrating functions over two or more sense modalities. Hand to eye coordination, for example, can be thought of as functional knowledge, it is a skill. Integrating regions capture knowledge of how to do things.

The 4th level of the World Picture Hierarchy contains self-knowledge information and historical information. Similarly, entities with inputs from within the cortex serve as sites for consciousness and memory, these being the brain's equivalent to self-knowledge and history.

The functional similarities between the robot control system and the human cortex are harder to lay out because we did not really explore the functioning of the cortex so much as its structure. (This is because of our lack of knowledge about how the brain works.) But some parallels can be drawn, and they are strong ones.

Feedback in the robot control system matches up with phasic re-entry in the human cortex. Both are processes by which a hypothesis is compared to sensory input and the result used to change or control other processes in the system. In both systems feedback and phasic re-entry appear at all levels in the hierarchical structures, most importantly at the highest levels where feedback becomes a learning process , the hypotheses being pieces of the system's self-knowledge. In simulation mode Sandi was able to study his own World Picture Hierarchy and to change it. In the human cortex phasic re-entry at the highest level gives rise to consciousness and memory which are the necessary ingrediants for learning to occur in the human mind.

This comparison of the robot control system and the human brain is not meant to make any claims about the functioning of the human brain. The study of the cortex was strictly structural, the only functional extension of the discussion being that certain structures in the cortex provide the links necessary to give rise to the complex phenomena of consciousness, memory and learning. No claims are made about how these processes work. The main purpose of the comparison is the highlight the power of the method of analyzing systems.

The Method: A Feedback Loop, A Self-Knowledge Link

In the Philosophical Investigations Wittgenstein says: "I am just trying to show you new ways of looking at things." In that sense I have followed in his footsteps. What I have developed in this thesis is a tool for analyzing systems, a method for recognizing systems at work.

The method can be broken down into three parts.

First, take the system you are study and look at it as a static thing, try to seperate it from its environment and pass over the dynamic forces of interaction at work. Here, in the first step, we are looking for the underlying hierarchical structures that form the skeleton of the system. Identify the levels in these structures and most important of all, understand the lines drawn between the levels of abstraction. The stronger these lines are drawn the more powerful the hierarchy will be in structuring the system.

Second, with these hierarchical structures in hand set the system in motion again through your analysis, watch it evolving and learning, study its dynamic behavior. You are looking now for the horizontal links in those hierarchies you found in the first step and the feedback loops which provide for self-knowledge and a balanced interactions with the system's environment. Be aware that the horizontal links take many forms: there are links within hierarchies, links between hierarchies, and links to the system's environment which define the system's self-image in the world. When studying the feedback effects look for the sources of the hypotheses, and discover the processes which compare expected events to actual events, and look for mechanisms which deal with real-time constraints and parallel processing.

Third, and finally, take the knowledge you have gained from the system you are studying and compare it to other systems you or others have studied. This is the most important step in the process. True knowledge requires you to be able to view your subject from many different points of view.

If my "new way of looking at" systems shows anything then it shows that deep down underneath all the superficial details all systems behave in a similar manner and it is by linking one system to others that our studies gain depth. The founder of general systems theory, Ludwig von Bertalanffy, discovered this long before I did. He developed rigorous mathematical models which could be used to lay out isomorphisms between systems which on the surface appear vastly different; systems in physics, biology, sociology and cybernetics. His goal was to train generalists who could see the connections between those fields of study and see how to transfer discoveries in one field to another field creating new insights there. If I have taken even a small step in this direction then I am happy with my work.

Bibliography

Albus, James S. Brains, Behavior, & Robotics
Peterborough, NH : BYTE Books, 1981
Q355.A44

Bertalanffy, Ludwig von General Systems Theory
New York : George Braziller, 1968
Q295.B4

Churchman, C. West The Design of Inquiring Systems
New York : Basic Books, 1971
Q295.C46

Mountcastle, Vernon B. An Organizing Principle for Cerebral Function: The Unit Module and the Distributed System
Cambridge, MA : The MIT Press, 1978
QP395.E36

Whyte, Lancelot Law, ed. Hierarchical Structures
New York : American Elsevier, 1969
Q295.H5