Complex Adaptive Systems Book

A complex adaptive system is a special class of complex system that has the capacity for adaptation. Like all complex systems, they consist of many elements, what are called agents, with these agents interacting in a nonlinear fashion creating a network of connections within which agents are acting and reacting to each other's behavior. Through adaptation, agents have the capacity to synchronize their states or activities with other agents locally.

Out of these local interactions, the system can self-organize with the emergence of globally coherent patterns of organization developing. There develops a complex dynamic between the bottom up motives of the individual agents and the top down macro scale system of organization, both of which are often driven by different agendas but are ultimately interdependent. There are many examples of complex adaptive systems from ant colonies to financial market to the human immune system, to democracies and all types of ecosystems.

Agents

Complex adaptive systems are composed of agents. An agent is an actor that has the capacity to adapt their state, meaning that given some change within their environment they can in response adjust their own state. For example, if the agent is a player within a sports game, then if we throw a ball to the person he or she can catch that ball. They are able to do this because they have what is called a regulatory or control system.

A control system of this kind consists of a sensor, controller, and an actuator. The person is using their optical sense to input information to their brain (the controller), that is then sending out a response to their mussels (the actuator), and through this process, they can adjust to generate the appropriate response to this change in their environment. And it is this same process through which a bird in an ecosystem or a trader within a market is receiving information, processing it and generating a response. Typically, these agents can only intercept and process a limited amount of local information, like a snail following a trail on the ground. It does not have a global vision of the whole terrain around it and it must simply respond to the local information available to it.

Synchronization

With this capacity of adaptation, agents have some degree of autonomy through which they can choose to synchronize or desynchronize their state with that of other agents within their local environment. We might also call this cooperation or competition. They typically do this based upon the costs and payoffs for choosing one of either option, and this cost-benefit ratio varies depending on the scenario, or what we might call the game they are engaged in with other agents. Some scenarios such as playing chess have very low incentives for cooperation while favoring competition. These are called zero-sum games. While other scenarios have a much lower cost and a higher payoff for cooperation; such as driving your car on the correct side of the road.

These different types of games create attractors that result in default positions for agents to cooperate or compete. Added to this are feedback loops, where what one agent does influences what another chooses to do. If you owned a certain stock and upon hearing some negative news about that company all of your fellow traders around you started selling it, this would create a positive feedback loop attracting you to also sell. And if you did, that would again amplify the positive feedback placing a stronger attraction on others to also do likewise. In such a fashion some phenomena can cascade through a population synchronizing their states rapidly.

Emergence

This process previously described is a form of what is called self-organization. From the interaction of the individual agents arises some kind of global pattern which typically could not have been predicted from the behavior of the agents in isolation. For example, in the brain, consciousness is an emergent phenomenon which comes from the interaction between the brain cells. Thus, the global property of consciousness results from the aggregate behavior of individual elements.

Within this macro-scale system that emerges control and regulation is typically distributed out. There is no master neuron or set of neurons that tell the whole brain what to do. No one is in control and no one in the system has complete information of it. This distributed nature of complex adaptive systems may make them very robust, where the system can adapt to some large disturbance.

The internet might be an example of this. Dynamically updated routing tables keep track of how long it takes to send information along any path on the network. If there is a failure in one part of the network, packets are rerouted through another channel. Control over the flow of I.P. packets is distributed out over many different routers and service providers, with a large amount of redundancy making it robust to failure. But equally complex adaptive systems can self-organize into a critical state where feedback loops can work to amplify some small perturbation into a large systemic effect as witnessed during a financial crisis.

Evolution

This emergent macro-scale system of organization then operates within some environment. Whether we are talking about a herd of animals within an ecosystem, the human body, a democracy or a corporation within a market, the whole macro system is periodically subject to perturbations and change within that environment. In order for it to optimize its state there must be some mechanism for performing selection upon the agents within the system. Those creatures within an ecosystem that can best respond to the environment are replicated. Those employees that have proven their value to the company will be promoted while others will be fired. Those products that best fulfill the demand are selected by the consumer, while others go by the wayside, the result being that the whole system evolves to exhibit more of the desired characteristics as they become more prevalent with the system.

Top-Down & Bottom-Up

In this way, this global pattern of organization will feedback to affect the agents on the local level, both enabling them and constraining them. It enables them as it is a mechanism for them to coordinate their activities and thus receive the benefits from forming part of a complex organization in the form of security, shared knowledge, technology, and so on. But it will also constrain them as following regulations and being subject to some form of selection is part of maintaining this global organization.

But of course agents have their own agendas that may or may not be aligned with those of the whole system, and this is where the real complexity comes into the dynamic, as there is now a core tension between the micro and macro levels. The system as a whole, that is, how it appears within its environment, will be primarily defined by how this core tension is resolved. That is to say, is the system driven by the interests of the agents at the expense of the whole? Or by the interest of the whole at the expense of the interests of the individuals? Or has it managed to find some resolution to this conflict? If we take an example of an economy, we can have a free market economy which is driven primarily by the interest of the agents in a bottom-up fashion, or we might have a communist economy driven by a top-down dynamic at the expense of individual motives, or we may have some economic system that manages to integrate the two.

Adaptive Systems

We can define adaptation as the capacity for a system to change its state in response to some change within its environment. An adaptive system then is a system that can change given some external perturbation, and this is done in order to optimize or maintain its condition within an environment by modifying its state. One of the simplest examples of an adaptive system might be a mousetrap. It is designed to respond to some perturbation that triggers a mechanical reaction within the system, and it is simply executing on a logic that was built into its mechanical structure by design.

The growth of a plant or fungus towards a source of light, what is called phototropism, is another example of adaptation. The cells on the plant that are farthest from the light release a chemical causing them to elongate and thus move the plant towards the light source. This adaptability gives the organism some flexibility that improves its performance and chances of survival. Flexibility simply means it can generate an optimal response to a limited set of changes within the state of its environment.

Advanced Adaptation

In both of these previous examples, the capacity for adaptation is simply embedded within the physical mechanics of the system, but entities that are capable of more advanced forms of adaptation have specialized subsystems dedicated to regulating this process of adaption. We call these specialized components regulatory or control systems. For example, an animal like a cat has a nervous system dedicated to sensing, processing, and responding to information it receives from its environment. With an electrical nervous system the creature is able to respond very rapidly and also capable of a much more complex algorithmic processing of information.

With this control system it is able to generate a wide variety of responses to deal with a rapidly changing environment where it may be presented with a large number of different scenarios; as a creature like our cat might encounter during the dynamic activity of hunting. Beyond this type of algorithmic logic that governs basic control systems, adaptive systems can have a much more complex conceptual framework for representing their environment and we call this a schema. Schemata, which is the plural for schema, are mental frameworks or concepts used to organize and structure information. With schemata an adaptive system has a full model for classifying and correlating different information about its environment, which can then be used to interpret new information, learn and generate novel responses to a very wide variety of input stimulus.

Agency

Within complex adaptive systems theory, these adaptive systems are called agents, and they are so called because they have agency. Agency can be defined as an action or intervention designed to produce a particular effect. An agent then is an entity that takes an active role to produce a specific outcome. Thus, agents do not act in a random fashion but actions are performed in order to produce a particular effect. That is to say, all adaptive systems have a goal, whether we are talking about a plant that adapts its state by moving towards the direction of the sunlight or a trader who buys a particular security to diversify her portfolio. These agents are acting according to a set of rules that are specifically formulated or designed to achieve the desired outcome, although these desired outcomes may be very diverse from the plant requiring more sunlight to the trader wanting more capital.

Agent Functioning

Adaptive systems have a particular internal order or structure that enables them to intercept and transform energy and resources of some kind. We may be talking about our plant intercepting photons, combining and transforming them into sugars through the process of photosynthesis, or we may be talking about a business within an economy that takes in some input and transforms it into an output that generates revenue. The aim or goal then of an adaptive system with agency is to maintain and develop this internal order and capacity to process resources. They can only do this by importing energy or resources and exporting entropy to and from their environment. In other words, these systems are dependent upon their environment to ensure their continued functioning and they adapt in order to maintain and optimize their status within this environment, with this whole process being described as homeostasis.

The capacity for adaptation gives rise to autonomy, whereas much of our science is focused on studying deterministic systems – where we search for linear cause and effect interactions that govern them and then encode these in mathematical equations as laws of nature – the capacity for adaptation though gives an element a certain capability to act autonomously from these deterministic linear cause and effect laws. We say to a certain extent because most simple adaptive systems like a thermostat are essentially deterministic they are determined to respond to some cause in their environment with a given effect. It is only when we have advanced adaptive systems with internal agency that we get autonomy and the capacity for a variety of responses given any cause. The more complex the logic governing the adaptive system is the more capable it is of producing a variety of responses to any given input and the more it is able to operate sustainably in a broader complex environment.

Regulatory Systems

Within systems theory, a control system or regulatory system is a specialized subsystem that is designed to monitor and regulate the behavior and operation of the broader system it is a part of in order to maintain its functionality. The primary objective of a control system is to preserve the internal level of order that enables the system to function and develop. Within systems theory this preservation of a stable or equilibrium state to a system's operation is called homeostasis. On its most basic level homeostasis is the maintenance of a system within a given set of parameters or environmental conditions that best enable its internal functioning.
Homeostasis is untimely at the core of what all types of control systems are designed to do. Some examples of control systems include the thermostat that is designed to regulate the temperature of a given system, such as a building, within a defined set of parameters. Another example from the human body is the hypothalamus that regulates the autonomic nervous system. It controls basic body functions such as internal temperature, hunger, sleep and so on. The commander of a military unit is another example. He or she is given the position of supreme control within the organization, responsible for its maintenance and operation. A nation's government is another example of a control system, designed to maintain and develop the socio-economic system within a given jurisdiction.

Control System's Elements

All of these very diverse systems share a basic underlying set of relations and components that are common to all regulatory systems. There are essentially just three components to any given control mechanism. Firstly, we need a sensor for feeding information into the system. Secondly, a controller that contains the logic or set of instructions for processing this information. Lastly, an actuator that executes some action in order to affect the state of the system or its environment.

Sensor

Firstly, a sensor is a component that detects and encodes some stimulus from the system's environment and transfers it to the controller. Any given sensor can of course, only sense a specific stimulus. A sensor has a physical device that is receptive to some change in a parameter that it is measuring, with this change in stimulus then being encoded into information and transferred ultimately to the controller. Examples of this are the visual system within biological organisms where the photons hitting the optical nerves in the retina are encoded into electrical signals to be sent to the brain, or seismometers that sense small movements in the earth’s surface and encode them into graphical representations.

Controller

Secondly, the controller. The controller is the brains of the operation. It contains the critical logic that is governing the whole system and is encoded in some set of instructions. The controller can be modeled as an information-processing unit taking in some input of information, manipulating this information according to its set of instructions with the result being an output of information that is designed to be acted upon. An example of a controller might be a digital circuit board composed of logic gates that physically manipulate an electrical input according to binary operations to produce some output signal.

On this basic level, the set of instructions is what we would call an algorithm. By switching gates on and off, they can respond to a limited set of input signals through an if/then logic, to create some output response. This set of logic gates is an example of a very simple controller. To take an example of a much more complex controller we might think about a democratic nation’s public administration system operating under a set of instructions encoded within a constitution. It is designed to take information about the state of the nation that has been received from a number of different sensors such as the mass media or statistics gathering and process this information according to these sets of instructions to produce the policies and regulations required to maintain and develop the socio-economic system of the nation state.

Actuator

Lastly, the actuator. An actuator is an instrument or set of instruments that act on the instructions produced by the controller. It is designed to physically affect the system that is being regulated in order for it to conform with the instructions produced by the controller. An example of an actuator might be the muscles in the human body. They are controlled by electrical signals sent from the brain. We can actuate them in order to change the state of our environment by simply moving from one location to another. The brakes on a car are another example of an actuator. They execute or act on one’s instructions to regulate the speed of the car.

Control

For a system to be regulated or under control means that for any given change in state presented by its environment the system can generate a response so as to maintain functionality. In order for a system to be able to regulate itself, all of these components to a regulatory system need to be working together. Without a sensor, the controller cannot know the state of its environment, and thus the appropriate response. If you are driving your car with impaired vision, you are not receiving all the required information about your changing environment that is required to generate the appropriate response, with the result being that sooner or later you will fail to receive critical information that will cause an accident, drastically reducing the system's functionality; thus in such a situation one can say you are not fully in control of the vehicle.

Without a controller the system cannot alter an input to a required output, thus cannot adapt, and will be under the control of external influences within its environment. In order for the system to have control over itself, it must have a logical set of instructions that are able to process any input signal to the required output response.

There are two key questions to consider here. Firstly, the basic functioning of the logic unit. Are there errors in the instructions and how they are being processed, such as bugs in computer code or random deformations within DNA? Secondly, does the logic have sufficient variety and complexity to be able to represent all the different states that its environment will present it with? Within cybernetics, this is called requisite variety, which simply means that the system is required to have a set of instructions with sufficient variety and complexity to represent all the diverse states within its environment, or else it will not be able to regulate itself within that environment.

As an example of this, we might think of asking a small child to run a multi-national corporation. The child simply does not have the conceptual capabilities to represent the complexity of the system it is asked to regulate. Thus, it is not in control and the system's functionality will be degraded over time as it moves outside of some homeostatic parameter that the child is not aware of and thus not able to respond to.

Thirdly, without the functioning of the regulatory system's actuator, the instructions created by the logic unit cannot be executed upon and thus the system cannot alter its state to respond to the changes in its environment. If a nation's law enforcement agency refuses to execute on a court's legal decree to disband a popular protest, then the government is essentially out of control, as it no longer has the actuator required to regulate the system.

Schemata

A schema is a conceptual pattern or structure that organizes categories of information and the relationships among them. The term derives from psychology but is used within complex adaptive systems theory to denote the internal logic governing an agent's behavior. The internal logic or schema that governs the behavior of agents within complex adaptive systems spans from the very elementary to the very complex. The most basic type of logic is called an algorithm and more advanced conceptual systems may be called schemata.

Algorithms

The most basic form of logic an agent can have is one that simply responds to a given input signal with an output action that is always the same. For example, if one taps one's knee at the right location it will trigger the nerves to actuate the muscles into generating a sharp reactionary motion. Every time we input the same stimulus to this physiological system we will get the same response. This is the most basic algorithm conceivable, always mapping the same input to the same output.

More advanced algorithms are able to discern between a given set of inputs and use an if-then logic to select an appropriate output. For example, the control system within a chemical processing plant might be able to select from a set of output temperature values based upon a range of input temperature values in order to regulate a chemical process chamber.

An example of a basic algorithm might be the rules that are thought to govern the flocking of birds. They consist of just three simple rules which are; one, separation, meaning always maintain a certain distance from your neighbors; two, alignment, meaning steer towards the average heading of your neighbors; and three, cohesion, meaning to steer towards the average position of neighbors in order to maintain long range attraction. Here the individual bird is continuously inputting a value to these three required parameters, processing this information according to the set of instructions and then selecting from a range of appropriate motion responses in order to maintain its correct positioning.

Schemata

As advanced as these algorithms may become they are essentially designed to just generate a response to a given range of stimuli. As such, they capture much of the logic behind mechanical control systems and those governing many biological systems such as in our bird example above. But the advanced cognitive capability of a modern human being far exceeds a simple set of algorithms. With this cognitive capacity, human agents can create conceptual representations or models of the world and we call these schemata.

The word schema comes from the Greek word meaning to shape, or more generally, plan. A schema is a cognitive framework or concept that helps organize and interpret information. As such, it is a conceptual template that determines how reality is interpreted, and from this what are appropriate responses to a given stimuli. With a schema, an agent can create a model of what it encounters, identify similarities and differences amongst things in order to create categories and relations between categories. This allows an agent to quickly take in new information and classify it with reference to what it already knows. Every time an agent receives new information it references it against the information it already has.

Bayesian Inference

This process of obtaining new information and filtering it to ensure its validity is often modeled using Bayesian inference. Bayesian inference references any new information received by the agent against prior knowledge in order to ascribe a probability value to the likelihood of its validity. If the information is deemed to have a high probability of validity it is incorporated into the agent's schema and used as a reference to infer the validity of any future information it receives. For example, throughout one's life, one have received constant information endorsing the validity to the existence of the force of gravity.

This massive amount of information confirming it gives it a very high probability of being valid, and every day that probability goes up as one receive more confirmation of its existence, the result being that if you are presented with some piece of information that disproves the existence of a gravitational force on planet Earth your immediate reaction will be to ascribe this new piece of information with a very low probability of being valid. In this way, a schema can develop as it receives new information and incorporates this into the framework, both reinforcing pre-existing categories and reducing the overall state of uncertainty as new information confirms or disaffirms the space of unknown possibilities.

Complex Representations

With a schema, we have not only the basic functioning of a control system that is able to respond to an immediate stimulus but by being capable of creating a complex model of a situation we can understand what is generating this stimulus in the first place. A schema allows the agent to identify the causes that create the effects. And not only this, but an agent with an advanced schema is able to also create a model of its own operation, that is, how it responds to any given stimulus, and can then try to alter this basic behavior.

For example, we might be able to identify that every time we get stressed we start smoking and then try to alter this reaction. This somewhat self-referential capacity for a system to model and analyze its own regulatory system is the subject of what is called second-order or new cybernetics.

These advanced schemata, of course, have many benefits to an agent over a simple algorithmic logic. It is ultimately the foundation that has enabled technology, advanced civilization, and human's capacity to dominate its physical environment. But of course it comes at a cost, and not only in terms of the physical energy to maintain the system. But there is now a tension between the basic control system that is designed to react to stimulus, thus ensuring immediate self-preservation, and the schema that creates a broader vision interested in the system's long-term objectives and consequences of its actions, with the possibility of these two levels conflicting and reducing the agent's capacity for action.

Human agents within complex adaptive systems are not only governed by the need for physical self-preservation but being governed by these advanced conceptual frameworks they are required to maintain both conceptual homeostasis as much as physical homeostasis. Through a number of mechanisms, information can be systematically filtered to ensure it does not threaten the basic assumptions that support the schema and that the system is in regular contact with information sources that endorse and preserve this current schema because it is critical to the functioning of the whole system.

Psychology has plenty of examples of this, such as confirmation bias which is a tendency to search for or interpret information in a way that confirms one's pre-existing schema and placing much higher validation standards on information that threatens it. In the same way, agents actively seek out environments that are inductive to their physical requirements, they will often actively seek out information sources that preserve and maintain the status quo of their schema. Thus, we should not expect human agents to be rational or logical. Ultimately, humans are not computers where logic is a precondition to their operation, but there is instead a subjective dimension to humans that is driven by emotions and independent from logical validation.

Culture

This subjective domain to human agents is played out in what we call culture. E.B. Tylor defined culture as “that complex whole which includes knowledge, belief, art, morals, law, custom and any other capabilities and habits acquired by man as a member of society.” Culture often comes in the form of a story or a set of stories about how the world is that endorse what is considered right and wrong, with people then acting out these stories as rituals in order to validate them and feel a part of them.

People buy Nike shoes because advertising agencies have created a story around the brand. People want to be associated with that and they live this story out by wearing the shoes. There is no economic logic as to why people would pay an extra 50 dollars to buy a pair of shoes with a tick on the side of them. Much of human activity only makes sense within the context of the cultural narrative that it is a part of. This may add a whole new level of complexity to our models but we pay a high price when we exclude it in terms of capacity to capture the real-world phenomena exhibited by many complex adaptive systems.

Simple Rules

One of the key premises of complex systems theory is that global coordination and complex behavior can emerge out of very simple rules governing the interaction between agents on the local level without the need for centralized coordination. At the heart of this is the question of how agents synchronize their state or cooperate to create local patterns of organization. We see many examples of self-organization within complex adaptive systems that are composed of elements following simple rules. For example, swarms of fireflies who may start out flashing their light in a random fashion with respect to each other come, through their interaction, to coordinate their behavior into an emergent pattern of the whole swam flashing in synchrony.

Types Of Rules

This type of quite basic self-organization can be modeled using cellular automata where very simple rules are programmed into a computer, and out of the interaction between these simple agents we see emerging surprisingly dynamic patterns that are able to stay evolving over prolonged periods of time to produce novel behavior. Ant colonies are another often-cited example of self-organization through simple rules. Without a centralized coordinator, the colony as a whole exhibits quite sophisticated differentiation and specialization of its functional organs that then work together to maintain the whole system.

Individual ants interact and communicate through exchanging chemical scents that induce other ants to do more or less of a given activity. This type of coordination is the product of what we call feedback loops. Through feedback loops some local pattern or behavior can become amplified to create an attractor state that will draw local elements into a particular synchronized configuration, thus arising some pattern of organization without the need for any form of top-down control system.

Cooperation & Competition

Cooperation and competition are another lens through which we can try to understand this process of synchronization. As an example, we might think about the vast complex adaptive system of our global economy, an organization that is capable of producing things like laptop computers and sports cars that no individual could produce in isolation. They take the coordination of thousands or possibly millions of people in order to complete the full production and distribution process, but no one is in control of this whole operation. No one makes these people coordinate their activities. They have done so according to their own local rules and incentives. The elements in this type of system have agency, that is, some kind of choice over their actions. And thus, we can best understand the coordination of their activities through the concepts of cooperation and competition, where agents choose to synchronize their states in order to maximize their individual payoffs, once again giving rise to local and global patterns of organization.

Cellular Automaton

Cellular automata are algorithmic models that use computation to iterate on very simple rules. In so doing, these very simple rules can create complex emergent phenomena through the interaction between agents as they evolve over time. To illustrate the functioning of a cellular automaton, we will take an example from probably the most famous algorithm called the Game Of Life devised by the mathematician John Conway.
The Game Of Life is played on a grid of square cells. A cell can be alive or dead. A live cell is shown by putting a mark on its square. A dead cell is shown by leaving the square empty, each cell in the grid has a neighborhood consisting of all adjacent cells to it and there are just three rules governing the behavior of an agent. 1. Any live cell with fewer than two live neighbors dies; as if caused by under-population. 2. Any live cell with two or three live neighbors lives on to the next generation. 3. Any live cell with more than three live neighbors dies, as if by overcrowding. 4. Any dead cell with exactly three live neighbors becomes a live cell, as if by reproduction.

Types Of Patterns

Cellular automata rule sets may be classified according to the degree of complexity that they can produce with Stephen Wolfram’s classification being the first attempt to classify the rules themselves in order of complexity. The most simple pattern is called still life. Its product is probably the most simple class of pattern, called class one, where nearly all of these patterns evolve quickly into a stable, homogeneous state and any randomness in the initial pattern disappears. The second class of pattern we may get is where the system evolves into an oscillating structure. The simplest of these being a blinker that has a period two oscillation.

We can also have oscillating structures that cycle over prolonged periods of time. For example, a pulsar has a period three oscillation, but oscillators of many more periods are known to exist. Class three patterns are random, where nearly all initial patterns evolve in a semi-random or chaotic manner. Any stable structures that appear are quickly destroyed by the surrounding noise. Local changes to the initial pattern tend to spread indefinitely.

Here we can get what are called gliders where a group of cells appears to glide across the screen. This is a good example of emergence as we no longer see the simple rules that are producing them but instead this emergent structure of an object gliding. Lastly, automata can also produce patterns that become complex and endure over a prolonged period of time, with stable local structures. With these more complex patterns, cellular automata can simulate a variety of real-world systems, including biological and chemical ones.

Universal Computation

Since the advent of the Game of Life, new similar cellular automata have been developed that can do all sorts of things, such as create fractal patterns, that is, self-similar structures that repeat themselves over various scales of magnitude. Other games create patterns that can reproduce themselves. We might ask, is there anything that these automata can not do? From the perspective of computation, the Game Of Life can do anything that a computer can do. It can count to 100, calculate the volume of a cylinder, or if you wanted to figure out the cube root of 1230 you could encode this into a set of cells on the automaton and have it compute the value. The Game Of Life as simple as it is has been proven by computer scientists to be capable of universal computation.

Computational Models

Cellular automata represent a new approach to mathematical modeling based upon computation. Von Neumann and Ulam originally introduced the concept in the mid-20th century, and then a few decades later the popular Game of Life brought interest to the subject beyond academia. In the 80s, Stephen Wolfram engaged in a systematic study of cellular automata after which he published a book called "A New Kind of Science" claiming that cellular automata could enable a new approach based upon the exploration of these algorithms.

One assumption within modern science is that simple rules can only create simple phenomena, and thus inversely complex phenomena must be the product of complex rules. The advent of chaos theory during the past few decades revealed this to be an invalid assumption as simple systems like a double pendulum proved to be capable of generating complex and chaotic behavior. It is now increasingly accepted that complexity may not be the product of complex rules, but in fact, emerge out of the interaction of simple rules as they evolve over time. Cellular automata are the tools that capture and embody this paradigm within science.

Feedback & Externalities

A feedback loop defines a relationship of interdependency between two or more components where the change in state of one element affects that of another, with this effect then, in turn, feeding back to alter the source element. This dynamic captured by feedback loops plays a fundamental role in the self-organization of elements within complex systems.

When the state of elements within a system is independent of each other, then we can use statistics to model the correlation of states between elements. For example, say we have a hundred people in a town with just two banks, A & B. If all other things are equal then we can model whether two people are customers of the same bank using simple statistics, where approximately fifty percent of the people will be using any one of the banks. But if the usage of each bank is not independent; it is instead interdependent, then it will no longer simply be statistics governing the dynamics. It will now be these feedback loops of interdependence.

Example

To illustrate this, say more people are using bank A, and this leads to overcrowding in the bank. This may then feed back to affect the users as they decide to go to bank B which is now quicker and easier to use. And likewise, if bank B after some time then becomes overcrowded, people may move back to bank A. This is an example of a negative feedback, where the state of one element affects the other in the opposite direction.

We can see how the net result of this would be a stable system. If we had a hundred banks in this town governed by this rule, the result would be a very evenly distributed and stable system where the agents occupy a wide variety of states with respect to the banks that they use. But imagine one day bank A starts a marketing campaign, putting up a big billboard saying for every customer we have we will give you one percent extra interest on your savings. The result of this would be that for every new customer the bank had, it would present itself as a more attractive option for any other prospective customer.

This is an example of a positive feedback, where the more elements that adopt this state the stronger the attraction placed upon any other element is to also synchronize its state with this pattern of organization. Something to take away from this banking example is that in both the first and the second example, that is when we had random correlations or negative feedback between the elements, both of these dynamics led to an overall stable state where the system tended toward an equilibrium. Systems governed by these dynamics are linear and additive. We can create closed formula solutions to model them and they are the focus of most of our scientific framework.

Far-From-Equilibrium

In these first two systems, there is a dynamic that is working to maintain a distribution amongst the states between elements that results in an equilibrium. But this is not always the case. Positive feedback can drive the system far-from-equilibrium. Stock market crashes, outbreaks of war, political movements, growth and decay of ecosystems, traffic jams, and many biological processes are the product of positive feedback that takes place far-from-equilibrium. Take for example a social riot. As the rioting breaks out, your chance of going to jail decreases, and the social benefit of joining increases. This creates an attractor, attracting more elements to align themselves with this new organization. Positive feedback loops are nonlinear and they are often a signal of a system shifting into a new regime.

Externalities

Whereas feedback refers to dependencies between the same actions, externalities refer to dependencies between different actions. An example of an externality might be the relationship between the usage of personal transportation and air quality. The more cars the lower the air quality. This is a negative externality.

A positive externality might be one between the temperature on a given day and the sale of ice-creams. The higher the temperature the higher the ice-cream sales are likely to be. In contrast to a positive feedback loop, positive externalities can reinforce desynchronized states and diversity as two or more different states or classes of things are reinforcing and sustaining each other. This is essentially what we call a synergy. If we have more flowers we can have more bees, if we have more bees we can have more flowers. Thus, they endorse and sustain the diversity of states between them.

Positive Feedback & Externalities

Positive feedback combined with negative externalities can be a powerful force for synchronizing the state of elements within a system, as it both places a strong attraction on elements of the same class to synchronize their states while also depleting a different class. We might think about the rise of the Third Reich in pre-war Germany as an example. Every time a new member adheres and promotes the ideology of a sociopolitical organization like the Nazi party, it has a positive feedback effect amplifying this attractor. But also this social system was having a negative externality on other ethnic minority groups. Thus, it was both reducing the variety within the social group and external to it as all elements became aligned in this sociopolitical regime. The net result of this was totalitarianism as the social system moved far from its equilibrium, ultimately resulting in a phase transition as it collapsed into a post-war economic and social crisis.

In contrast to this, negative feedback combined with positive externalities will create a strong mechanism for maintaining equilibrium through endorsing a diverse set of desynchronized states within the system. This will clearly add to a system’s robustness and long-term sustainability, with mature ecosystems exemplifying this.

Self-Organization & Thermodynamics

Self-organization can be defined as the spontaneous creation of a globally coherent pattern out of local interactions. As such self-organization explores a new approach to understanding the patterns of organization seen in the world around us. Probably the oldest and most fundamental question there is to ask in science is, why and how do we get something instead of nothing, some form of order instead of just randomness? From the formation of galaxies, to the human body, to the structure of snowflakes, or the complex organization within a single biological cell. We live in a world that exhibits an extraordinary order of all kind and on all scales.

The real question is why or how do we get things to work together? How do we get global level coordination within a system? And there are two fundamentally different approaches to trying to answer this question. Firstly, this coordination may be imposed by some external entity, or secondly it may be self-generated internally. For thousands of years, many different societies came to the former conclusion, that this organization we see in the world derives from some external divine entity. Religions and spirituality often depict the world in terms of an interplay between supernatural forces of order and chaos. But of course, modern science has always rejected any form of divine intervention, as core to its foundation is the law of the conservation of energy and matter. The First Law of Thermodynamics is an expression of this fundamental conservation, which states that the total energy of an isolated system remains constant or conserved. Energy and matter can be neither created nor be destroyed, but simply transformed from one form to another. The conservation of energy is a fundamental assumption and keystone of the scientific enterprise.

Thermodynamic Laws

The Second Law of Thermodynamics states that the total entropy, which may be understood as disorder, will always increase over time in an isolated system. To understand where this comes from, we might think about how if we have some object heated that heat will always try to spread out to become evenly distributed within its environment. But the reverse never happens, heat will not spontaneously reverse this process to become concentrated again. Likewise, whenever rooms are cleaned they become messy again in the future. People get older as time passes, and not younger.

All of these are expressions of the Second Law of Thermodynamics, meaning that a system cannot spontaneously increase its order without external intervention that decreases order elsewhere in another system. For many years, the Second Law of Thermodynamics – that systems tend toward disorder – has generally been accepted. Unfortunately, this does not help in answering the question as to why our universe has in fact developed to produce at least some systems with extraordinarily high levels of organization. In fact, the second law of thermodynamics would predict quite the opposite.

Self-Organization

The term “self-organizing” was introduced to contemporary science in 1947 by psychiatrist and engineer W. Ross Ashby. Self-organization as a word and concept was used by those associated with general systems theory in the 1960s, but did not become commonplace in the scientific literature until its adoption by physicists and researchers in the field of complex systems in the 1970s and 1980s. In 1977 the work of Nobel Laureate chemist Ilya Prigogine on dissipative structures was one of the first to show that the Second Law of Thermodynamics may not be true for all systems. Prigogine was studying chemical and physical systems far-from-equilibrium and looking at how small fluctuations could be amplified through feedback loops to create new patterns.

For example, when water is heated evenly from below while cooling down evenly at its surface. Since the warm liquid is lighter than the cold liquid, the heated liquid tries to move upwards towards the surface. However, the cool liquid at the surface similarly tries to sink to the bottom. These two opposite movements cannot take place at the same time without some kind of coordination between the two flows of liquid. The liquid tends to self-organize into a pattern of hexagonal cells called convection cells, with an upward flow on one side of the cell and a downward flow on the other side.

Model Of Self-Organization

The theory of self-organization has come to explore a new approach to this age-old question about the emergence of order. Unlike religion and spirituality that simply ascribes it to exogenous supernatural phenomena or tradition reductionist science that posits that order can only come by transferring it from some other external system.

With self-organization theory, the organization is instead traced back to the interaction between components, where nonlinear interactions between elements can become amplified by positive feedback loops to create attractors that can result in new patterns of order emerging. However, this process requires the system to be far from its equilibrium so as to have sufficient entropy or disorder for new random fluctuations and noise to gain traction and take hold as emergent patterns form.

When the system is far from its equilibrium, it can find a dynamic state between order and chaos that enables it to continue generating novel phenomena and regenerate itself for prolonged periods of time through self-organization. Thus, this new set of theories around self-organization recognizes a complex interplay between order and chaos.

Whether we use the more scientific terminology of a system being far-from-equilibrium or the more catchy term of edge-of-chaos, this new vocabulary has built into it a recognition that self-organization, evolution, and novelty, thrive on a dynamic interplay between order and disorder because it is only when there is a sufficiently high enough level of entropy and disorder within the system that a weak fluctuation can be amplified into a new pattern of order. But when the system settles into an equilibrium or stable configuration, this no longer becomes possible.

Self-Organization Theory Today

Today the study of self-organizing systems is a hot topic that is central to understanding the complex systems that make up our world, with interest in how to model, design and manage complex systems coming from many areas such as the social sciences, computer science, business management, robotics and engineering. The theory and interest in the process of self-organization have arisen in tandem with computing resources. Just as we cannot study galaxies without our telescopes or cells without microscopes, we cannot study complex systems without computers. Whereas before it was very difficult to mathematically model systems with many degrees of freedom, the advent of inexpensive and powerful computers made it possible to construct and explore models composed of many entities. Looking at how out of their local interactions global patterns of organization can emerge and this represents one of the few primary methods which we use to study complex systems.

Self-Organization Far-From-Equilibrium

Far-from-equilibrium self-organization is a hypothesis that describes the process of self-organization as taking place at a critical phase transition space between order and chaos when the system is far from its equilibrium. The essence of the theory of far-from-equilibrium pattern formation is that new forms of organization form when a system is driven far from its stable basin of attraction. Far-from-equilibrium behavior is ubiquitous. The scope of phenomena investigated makes the research of far-from-equilibrium systems an intrinsically interdisciplinary activity that crosses between the physics community and researchers in biology, chemistry, the social sciences, applied mathematics, meteorology and engineering.

Organization

Organization is an ordered structure to the arrangement of elements within a system that enables them to function. As such, we can loosely equate it to the concept of order. Both order and organization are highly abstract concepts, neither of which are well defined within the language of mathematics and science. But probably the most powerful method we have for formalizing them is through the theory of symmetry.

The theory of symmetry within mathematics is an ancient area of interest originally coming from classical geometry, but within modern mathematics and physics, it has been abstracted to the concept of invariance. In this way, symmetry describes how two things are the same under some transformation. For example, if we take two coins, one showing heads and the other tails, by simply flipping one of the coins over it will come to have the same state as the other. Thus we do not need two pieces of information to describe the states within this system. We can describe this system in terms of just one state and a flipping transformation that when we perform it will give us the other state.

If instead of having two coins we had an apple and an orange. Now there is no transformation we know of that can map an apple to an orange. They are different things. There is no trivial symmetry or order between them, and thus we need at least two distinct pieces of information to describe this system. This second system requires more bits of information to describe its state. Thus, we can say it has higher statistical entropy. We can talk about and quantify order and randomness in terms of information theory. Ordered systems can be described in terms of these transformations which we encode in equations. Ordered systems are governed by equations whereas random systems are not. However, because there is no correlation between the element’s states in these random systems, they are governed by probability theory, the branch of mathematics that analyses random phenomena.

Order & Randomness

Complex systems are by any definition nonlinear. Complexity is always a product of an irreducible interaction or interplay between two or more things. If we can just do away with this core dynamic and interplay, then we simply have a linear system. If the system is homogeneous and everything can be reduced to one level, then it might be a complicated system but it is certainly not a complex system. Thus, one of the main ideas or findings of complexity theory is that complexity is found at what is sometimes called the interesting in-between.

If we take some parameter to a system, say its rate of change or its degree of diversity, and turn this parameter fully up, what we often get is randomness or a continuous change or total divinity of states without any pattern. Or if we turn it fully down we get complete stasis and homogeneity with very stable and simple patterns. It is often the case that with too much order the system becomes governed by a simple set of symmetries. Too much disorder results in randomness and the system becomes subject to statistical regularities. It is only between the two that we get complexity. On either side of this, there is a single dominant regime or attractor that will come to govern the system’s behavior.

Edge-Of-Chaos

It is only when a system is far from its equilibrium, away from one of these stable attractor regimes that we get a phase transition area representing the interplay between the two regimes. In this space, the system is much more sensitive to small fluctuations that can take it into either basin of attraction. This phase transition area is also called the edge-of-chaos. The phrase edge-of-chaos was first used to describe a transition phenomenon discovered by computer scientist Christopher Langton. Langton found a small area conducive to producing cellular automata capable of universal computation. At around the same time, physicist James Crutchfield and others used the phrase “onset of chaos” to describe more or less the same concept.

In the sciences in general, the phrase has come to refer to a metaphor that some physical, biological and social systems operate in a region between order and either complete randomness or chaos, where the complexity is maximal. The edge-of-chaos concept remains mainly theoretical and somewhat controversial, but it is often posited that self-organization and evolution can only really happen in this phase transition space.

There may be a number of different interpretations for why this is so, but one way of understanding it is that self-organization requires entropy and evolution requires variety. Unlike external intervention where we can take a well-ordered system and simply reconfigure it by transferring energy to it from some other external source, in this way we go from one ordered regime to another without the need for entropy to enable the process; we simply need some input of energy. But as we know, self-organization does not happen in this fashion. It is internally generated on the local level and this process requires the presence of entropy and randomness for elements to be available for reconfiguration into a new regime through feedback loops that originate as weak signals or fluctuations.

Different Theories

A number of different researchers have posited different theories around this process of self-organization far-from-equilibrium. The principle of “order from noise” was formulated by the cybernetician Heinz von Foerster in 1960. It notes that self-organization is facilitated by random perturbations and noise that let the system explore a variety of states in its state space. A similar principle was presented by Ilya Prigogine as “order through fluctuations” or “order out of chaos.” Researcher Per Bak also looked at this phenomenon in terms of what he calls self-organizing criticality, the mechanism by which complex systems tend to maintain themselves on this critical edge. Many of these theories talk about both the need for entropy and variety in order for the system to stay adapting and evolving over a prolonged period of time.

Robustness & Resilience

Resilience is the capacity of a system to maintain functionality in the face of some alteration within the system’s environment. All systems exist within an environment and are, to a certain extent, dependent upon a specific range of input values from that environment. The system has a set of parameters to these inputs within which it can maintain its structure and functionality, but outside of these critical parameters the system will disintegrate, i.e. become degraded to a lower level of integration or functionality.

Resilience and robustness can then be defined by this set of parameters. The lower the system’s dependency upon its environment and the broader this range of input values that the system can operate within, the more robust it can be said to be. For example, in computer science, robustness is the ability of a computer to cope with errors during execution, that is to say, the ability of an algorithm to continue operating despite abnormalities in input. In this way, robustness defines its independence from a specific range of inputs or inversely its capacity to process a wider range input states.

To illustrate this further, we might think about a tree withstanding the force of wind blowing against it. The tree has a certain tensile strength through its capacity to bend. Within a certain range of input values to the force exerted upon it, it will be able to withstand this perturbation from its environment. The wider the range of these input values the more robust the tree will be.

Types Of Resilience

There are fundamentally just two ways for a system to maintain its integrity given some perturbation. It can resist this change or adapt to it. By resist we mean it creates a boundary or filter condition that prevents the external influence from altering the internal configuration to the system; thus preserving its functionality and structure up to some limit. In our tree example, this might mean the organism developing a sturdy trunk. Inversely, the system can adapt by finding or generating the appropriate response required to counterbalance the perturbation. We might think of this as the tree bending over in response to the force exerted upon it. Robustness and resilience are general characteristics of self-organizing systems, both through their capacity to resist change and their capacity to adapt to it.

Resistance

Firstly, we will talk about their capacity to resist change through distributed control and feedback loops. In centralized systems with top-down control, there are specialized components required for regulating the system. These represent largely irreplaceable hubs that will affect the whole system if removed or degraded. Within complex systems in contrary, control is typically distributed out on the local level, meaning there is much less specialization. Missing or damaged components can often be replaced by others and this gives them a much lower level of criticality.

Secondly, self-organizing systems are held within their current configuration by a set of feedback loops that are also distributed out across the system on the local level. A good example of this might be a magnet which consists of many tiny magnetic spins that are all aligned to produce an overall magnetic force. If some of the spins are knocked out of their alignment, the magnetic field produced by the rest of the spins will quickly pull them back. This force maintaining the system within its current configuration is distributed out, giving it a low level of criticality and thus a higher level of robustness.

Adaptation

Adaptation is another mechanism for resilience. With adaptation, we are talking about the system's capacity to maintain or generate sufficient diversity of states for it to be able to select the appropriate response when required to counterbalance a perturbation from its environment, and thus maintain its internal configuration within the required critical parameters to preserve its structure or function.

To illustrate this, we might think about going hiking on a mountain. In this situation, one needs to be aware of the possible states to the weather that this environmental might present and have sufficient variety of clothing to counter-balance these different possible perturbations in order to maintain one's body within it critical temperature parameters that are required for its continued functioning. If one does not have what is called the requisite variety in order to adapt, then this environment might present me with a blizzard for which I do not have the thermal clothing to maintain my body, and in such a case my body's functionality may be severely or critically degraded.

Another reason for this intrinsic robustness to self-organizing systems is that self-organization thrives on randomness, fluctuations or noise. Without these initial random movements, self-organization cannot happen. A certain amount of random perturbations may facilitate rather than hinder self-organization. If the overall pattern that is generating the system remains intact, the entropy from the perturbation may be used for regeneration and evolution. For example, forest fires are thought to play an important role in the development of ecosystems. Excluding fires from these ecosystems means fire-adapted plants decline in abundance and overstocked forests become more prone to catastrophic fire due to the buildup of woody fuels. Exposing the system to perturbations without destroying it is a core part of the process of evolution and developing resilience.

Self-Organized Criticality

Self-organization does not always lead to robustness. It can also lead to what is called self-organized criticality where the system organizes into a state where some small event can have a large systemic effect. This phenomenon is best described with reference to what is called the sandpile model. This model is simulated by simply dropping grains of sand on a surface. As the pile builds up, grains roll off the side from time to time, typically just one or two at a time, but as we stay adding sand the side of the pile eventually builds up to a critical angle before we get a massive avalanche.

At some critical point, adding just one more grain of sand triggered a massive effect. This sand pile model for self-organization has been used to model everything from the occurrence of earthquakes to neuronal avalanches in the cortex and financial crises. The positive feedback loops that are an inherent part of the process of self-organization can also be a strong force for reducing diversity in the system as they synchronize it into a single regime where all elements become susceptible to the same perturbation. Without diversity to resist the spreading of some phenomenon it can cascade into a systemic shock.

Fitness Landscapes

A fitness landscape – also called an adaptive landscape – is a model that comes from biology where it is used to describe the "fitness" of a creature, or more specifically genotypes within a particular environment. The better suited the creature to that environment the higher its elevation on this fitness landscape will be. As such it visually represents the dynamics of evolution as a search over a set of possible solutions to a given environmental condition in order to find the optimal strategy which will have the highest elevation on this landscape and receive the highest payoff.

As evolution is a fundamental process that plays out across many different types of systems, natural, social and engineered, this model has been abstracted and applied to many different areas in particular within computer science, business management, and economics; but is equally applicable to all complex adaptive systems. Within this more generic model, a location on the landscape is a solution to a given problem. The elevation captures how functional that solution is, and solutions that are similar in nature are typically placed close to each other.

Basic Model

For example, the challenge might be commuting to work in the morning. There are many different strategies we could take from flying to possibly swimming to driving our car or taking the bus. We could then create a fitness landscape to represent this, where each one of these solutions would be given a fitness value based on how well it performs against some measurement of success, such as time or cost. The result being swimming or flying will likely end up at a low elevation relative to taking our car or the bus. We might also note that our car or bus strategy would be located in proximity to each other because they have many similarities while swimming or flying would be placed at very different locations on this landscape.

There are two main things we need to consider. Firstly, the type of landscape we are dealing with and secondly, the types of strategies we might use given these different landscapes. Firstly, to talk about the types of landscapes, what we will call their topology, there are a number of different parameters that will define the overall topology. Starting with how different are the payoffs on the landscape? The lower the range between the height of the peaks the more equal the payoffs between strategies. An example of an even topology might be a scenario where I roll a fair dice and ask you to try and predict the number it will land on. Each number is equally likely to turn up and thus each one of your strategies is an equally viable solution. As we turn up this parameter to the unevenness of the topology, there will come to be a greater disparity to the functionality of the different strategies and their payoffs.

Distribution Of Solutions

A key consideration is how distributed are the optimal solutions on the landscape? Is there just one dominant strategy that will drastically out-perform all others or are there many different viable solutions? For example, in terms of intercontinental passenger transportation, air travel drastically outperforms all other methods with respect to time. If we create a fitness landscape of the different methods, we would see one dominant mountain in the center with lots of other much smaller peaks around it.

Thirdly, how dynamic is the environment? Are we dealing with some ecology where environmental conditions may remain relatively stable for prolonged periods of time, or are we dealing with say some emerging market where the context is changing rapidly, resulting in the peaks and valleys to the landscape moving up and down as the whole landscape dances around? Lastly, how interdependent are events? Does what one agent chooses to do affect the landscape or other agents? A fitness landscape of say a market is created by all the companies, consumers and regulators within that market. Every time one of these players moves it affects the whole landscape, and thus we have a dynamic landscape that will be defined by these sets of interdependencies.

Complexity Of Landscapes

The adaptive landscape represents the different types of environments that agents are operating within and these different environments can span from the very simple to the very complex. On the simple end of the spectrum, we are dealing with a context that is static in nature and with limited interdependencies. On the complex end of the spectrum, we are dealing with environments that are dynamic in nature, consisting of many interdependent interacting parts.

Linear Environments

The most simple environments are static in nature and consist of the least amount of interacting variables, as an example we might think about an absolute monarch or absolute dictatorship where all social, economic and cultural institutions are controlled and held constant through the political hierarchy, within such an environment everything is in relation to one political institution, simply succeeding within that single organization can achieve global success. Or as another example, we might think about some homogeneous cultural system that defines clearly what is considered right and wrong and from this the one correct way to live one's life. These are examples of linear socio-cultural environments that would give the landscape a single dominant peak, one optimal solution that is well-defined, and because of this the agent needs only to follow some linear optimization algorithm.

Interconnected Environments

If we now increase the complexity by turning up the number of equally viable solutions we will get a landscape that has many different peaks and agents now have to invest a certain amount of time searching for the optimal position. As an example of this we might think about a young person having completed high school choosing which university to go to. They will be trying to optimize for a number of different variables, cost of tuition, location, facilities, college ranking etc. and thus there will be a number of different viable solutions, giving the landscape a number of different peaks, a roughed landscape.

But in this situation the variables are not changing over time thus the student could invest quite a bit of time and resources in researching all of the factors involved to find whatever they consider the optimal. Although this environment may represent complicated problems in that there are a number of interacting variables that require a significant amount of computation, it is still a relatively simple environment.

Adaptive Environment

If we now allow for the different interacting variables to adapt and change over time, we then have a complex environment. We now have a landscape where agents are acting and reacting to each other's behavior constantly adapting and it is out of this interdependence and adaptation that we get a landscape where the peaks and valleys are moving up and down over time.

An example of this might be the current international political environment as we move into an increasingly multipolar world, with the rise of China and the other emerging economies we are now no longer in an international environment dominated by the homogeneous Western ideology of the Bretton Woods institutions, but increasingly have many more actors, both public and private, each with their own strategies and interest that are constantly acting and reacting to each other. This means the end target is constantly changing, any solution that may be effective now, may cease to be effective when others adapt to it which once again alters the payoffs on the landscape as it moves up and down over time.

Open Environments

lastly, this whole complex adaptive system of agents acting and reacting is receiving some set of input values from external sources, whether this is the natural environment or the technology infrastructure of a society. A major change in these input values can cause the whole landscape to transform, in such circumstances we are no longer talking about the agents acting and reacting to each other, but instead, we are talking about the whole topology to the landscape transforming.

This is similar to a paradigm shift within science or culture where the whole landscape gets changed. We can think about the paradigm shift in our culture as we moved into the modern era, everything got recontextualized, through a scientific and materialistic context. With this cultural paradigm shift virtually every single social and cultural institution within the entire landscape had to reinvent itself within this new context. Education, governance, work, etc. everything got redefined and those that were not have slowly lost relevance, this is a long-term systemic change where we are no longer talking about adaptation but instead evolution.

Strategies

Knowing the different types of topologies, our attention then turns to thinking about the different types of strategies that agents might use, as the degree of functionality to any solution will alter drastically depending upon the type of landscape it is operating on. Agents within complex adaptive systems can typically only respond to local level information. Whether we are talking about a trader in a financial market or a herd of deer looking for pasture, these agents do not have complete information about their environment. They can only access and thus respond to a limited amount of typically local level information, and they need to have a strategy for processing this information and generating an optimal response.

Evolution

Evolution is a process of adaptation that operates on the macro-level of a system. Adaptation is the capacity to generate a response to some change within the environment. Evolution is this same process but operating on the macro scale, i.e. on the level of a population of agents. Here again, it is the capacity for the system to respond to changes within its environment. Evolution is the adaptive response of a group of entities that occurs over a period of many life cycles. Evolutionary changes reflect the response of the collection of agents to their environment.

The concept of evolution has a strong association with its application in biology. Complexity theory, though, deals with the concept on a slightly more abstract level as it applies to all complex adaptive systems from the development of civilizations to financial markets, cultures, and technologies. As such, we are trying to understand evolution as a continuous and pervasive phenomenon that occurs in all types of natural, social and engineered systems.

Basic Framework

The basic framework of the theory of evolution, as posited by systems theory and complexity theory, consists of just a few key elements. Firstly, a system exists within an environment and that environment changes periodically. For the system to endure, it must be able to generate the appropriate response to these environmental perturbations. Secondly, generating the appropriate response means selecting from a variety of different internal states or strategies, and thus the system needs to maintain and be able to generate a certain degree of variety. Thirdly, that variety is not for free. It costs the system something to maintain, and thus it must select from these variants the most appropriate responses for that environment in order to minimize this cost while maximizing the payoff to the system as a whole. Lastly, these variants that have proven most appropriate for responding to changes within that particular environment will then be selected for replication in order to be more prevalent within the system during its next life cycle, thus achieving the ultimate goal of altering the entire system to make it better suited to that environment.

Response To Change

Our first statement that a system needs to be able to generate an appropriate response to any change from its environment in order for it to endure can be described by the theory of homeostasis, that all systems require a certain state to their environment in order to operate, and they need to somehow maintain that set of input values from their environment or else they will cease to properly function when it changes. If I am driving my car down the road and someone pulls out in front of me, I need to be able to identify this and steer around them. The environment has changed, and if I cannot generate the appropriate response then I am in trouble. So this is not so much a statement of how evolution works, but more a statement of why we need evolution. Put very simply, it is because environments change. If the system cannot change with it, then it will eventually cease to function within that environment; without a centralized control system, evolution is the only way to prevent this.

Requisite Variety

Secondly, that the system needs to maintain a certain level of variety from which it can select the appropriate response given any environmental change. This is the so-called law of requisite variety that a system needs sufficient variety of states to respond to the variety within its environment. For example, the human immune system needs to have the right type of antibodies to neutralize an invader.

The immune system does this by simply producing an astronomical variety of different antibody shapes so that it will have the appropriate response when needed. In general, this requisite variety may be created by randomness – as in the random deformation of DNA – or cross-mixing, such as sexual reproduction, but also may be purposefully generated, as would be the case for an R&D lab.

Selection

Thirdly, selection. The most appropriate responses to the given state of the environment are selected, because diversity typically has some cost associated with it, and thus we can only maintain a limited amount of it. I may want to be appropriately dressed for any given weather condition, but I cannot bring my whole wardrobe with me. Each item I take will have a carrying cost associated with it. Thus, I must perform selection upon the variety of clothing I have based upon an assessment of the environment’s state.

Biological evolution by means of natural selection is another example. There is a limited amount of resources within any ecosystem. Selection takes place as a result of the competition among the members of a population for resources, and this helps by working to ensure that only those that are so-called “fit” for that environment will endure. In this way, the system can adjust its internal configuration to external perturbations, while minimizing the cost of diversity and changes to its overall organization.

Replication

Lastly, replication. Unlike adaptation which is an immediate process operating on the level of an individual, evolution works instead on the level of a population of agents. It does not operate immediately but plays out over the course of several life cycles to the population. Elements that have proven to be functional within that environment during their life cycle are selected for replication, thus increasing the percentage of their representation within the future population in order for the overall system to exhibit more of their desired characteristics.

Environmental Complexity

With this process of evolution, a system can, through its iteration over a prolonged period of time, go from starting simple to becoming more complex through the retention of functional variants. In so doing, expand to become capable of operating within broader more complex environments. Continuing on with the example of the immune system, the immature immune system of a newborn child is dependent upon its mother to produce and provide it with antibodies in order to fend off invaders. As it grows and comes in contact with new antigens, naturally or through vaccines, it develops its own antibodies and retains copies that have proven successful for future application. Thus, building up a catalog of successful antibodies that can provide it with the requisite variety to maintain its physiological homeostasis within more threatening environments.

When a system has requisite variety, then it can be said to have control over itself within that particular environment. However, there is always a broader environment that will present the system with a wider more complex set of eventualities for it to deal with. As the system evolves, it retains the appropriate responses to a given perturbation until it has accumulated the requisite variety for a given environment, and then can expand into a broader one where again it will have to generate more variety in order to deal with a new set of perturbations.

Genetic Algorithms

Now that researchers have come to understand the dynamics of evolution, it is increasingly being used as an optimization algorithm in many areas. For example, computer scientists create programs or formulas that compete against one another to solve a problem, the winners being rewarded with “offspring” in the next generation that then compete again. Over a series of generations, one can use this process to evolve optimal solutions to difficult problems. The resulting method, under the name genetic algorithms, has become a widely used optimization method and a tool for complex systems researchers. Genetic algorithms are good at taking a very large search space and looking for optimal solutions through iteration.