Introduction

Abstract

The problem of our age is how to go about understanding, deeply, the kinds of complex systems that we need to understand today. Our world is in trouble. Between suffering the unforeseen consequences of technologies that we thought were brilliant (when they were invented) and those of political decisions that seemed right (under the then circumstances) we seem to be witnessing complications in every aspect of societies, the environment, public health, and just about anything you can name.

The problem is complexity. Or, rather, the problem is one of managing complexity so that we benefit and chaos does not overwhelm us. And in order to manage anything you need to know it deeply. That is, you need to know how it works (and how it doesn't work at times) down to a level where interventions (if needed) can be used to keep the system stable.

We suggest that the pathway to deep understanding of any system, no matter how complex, is in front of us. Our guide to the pathway is to apply systems thinking and systems science to how we go about understanding systems. Many will argue that that is exactly what they do. And many people do so to one extent or another. Systems science has been developing for more than half a century and many thinkers have advanced concepts about systemness and many have used those concepts to pursue avenues in the sciences and engineering. There are a plethora of tools and methodologies developed around those concepts. But heretofore there has been no central focus on a set of comprehensive methodologies that could unify the various approaches to using systems science to gain deep understanding. There has not been a systemic systems approach. Many examples of disparate and sometimes contradictory views of systems science, or rather some part of systems science, will be provided throughout the book while still developing approaches to unification.

I.1. What this Book is About

Fundamentally this book addresses the problem of how to deeply understand the systems we encounter in the Universe or create from our own desires. There is a difference between deep understanding and more shallow understanding. This introduction will explore these differences and explain why deep understanding is preferable in every domain of knowledge. Briefly, though, superficial understanding is characterized by an ability to build an abstract model of a phenomenon and use that model to predict future behaviors under various conditions. Take a very simple system for example, a standard regression model can be constructed on the basis of correlations between two variables, one an independent variable like investment dollars and the other a dependent variable like returns. We can build this model by analyzing data from past experience. The regression curve can then be used to predict the value of the dependent variable given the value of the independent variable with some probability based on the correlation in the original samples. Correlation, however as the famous admonition goes, is not causation. Such models can only tell us what might happen (within certain error bounds) given the value of an independent variable. What is not known in this kind of model is what the causal relation is between the two variables. Shallow understanding is certainly better than no understanding (or guessing). But we know that shallow understanding can lead to surprises when systems do not behave as predicted. All too often the problem is that the original model failed to capture some internal causal mechanics that, perhaps only rarely, were operating to generate not understood aspects of the system.

At a slightly deeper level of understanding we can build causal models based on the relations between macrostates of a dynamic system at a course-grained level of resolution. There needs to be an assumed mapping from microstates, which are not known in detail, to those macrostate variables of interest. Wolpert, et al (2017), describe this as a state space compression, i.e., the macrostate variables are a compressed (in the sense of data compression) version of the aggregate of microstates. One example they give is of weather prediction based on measuring discrete points on the geography (weather stations) and using algorithms that map such measurements to larger phenomena such as pressure fronts or hurricanes. The latter are characterized as macrostates of the larger system whereas the former data points represent microstates. The point of having such mappings or models is to be able to predict future macrostates by also modeling the temporal evolution of the system at the macro level. The success of this procedure depends on how well the mapping from microstates to macrostates approximates the reality of the system.

Models based on observed relations between macrostate variables, based on the compression mapping from microstates to macrostates, are well developed, particularly with modern computational power. But they still are based on an inferred mapping process. The origins of these models can be traced back to work in the statistical mechanics of properties of, for example, gasses where the macrostate variables of temperature and pressure could be related to the behavior of the gas particles in containment, possible by knowing the average behavior of particles in collisions at average kinetic energies and in particular densities. Whenever a system, at the micro-level, has a large number of degrees of freedom, it becomes impossible to work from the microstates to describe the behavior of the system. Even the notion of having knowledge of the initial conditions (at first observation) is fundamentally impossible. But because these, what we will call “simple,” systems can be understood from their physics and because their overall behavior at the macro-level is generally what interests us, this form of understanding and modeling is completely sufficient[1].

[1] This should not be interpreted to mean we have “perfect” knowledge/understanding of these systems. There are probably more details of such systems that physics may yet uncover as they continue to be studied. For example we might do a better job of connecting the microstates with quantum-level laws.

In this book we are interested in phenomena that transpire at higher levels of organization and complexity, all of which are composed from these simpler systems[2]. This is the domain of complex systems, which includes living systems and human-derived artifacts. The methods of modeling that have been so successful for simple systems start to break down in terms of providing veridical predictions (or anticipations) of future behaviors of such systems under varying environmental conditions. The mapping between microstates and macrostates needs to be based on a new way to obtain understanding. The accuracy and precision of computing the future behavior of extremely complex systems requires a deeper understanding of the causal links between those microstates and the macrostates we observe.

Deep understanding involves obtaining knowledge of what goes on inside the phenomenon that produces the behaviors themselves. This is the meaning of systems. A system is composed of components that individually behave but in concert with the other components so as to produce the behavior of the whole. Knowing what those subsystems are and how they behave, essentially having models of them, leads to much better models of the system as a whole. Deep understanding provides a causal model of the system working. Henceforth, we explore the nature of obtaining deep understanding of complex systems and the construction of deep causal models as opposed to either statistical modeling (e.g. regression) or presumptive micro-to-macro mappings (e.g. as is often done in both dynamical systems and systems dynamics[3]).

Why this deep-understanding-based kind of model is superior may not be immediately obvious. The reason that deep understanding should be preferred over shallow understanding, even when the latter is adequate for many practical purposes, is that system behaviors can surprise us over longer time scales; the Universe is evolving. Things change and on all scales of space and time. This fact shows up most clearly in very complex dynamic systems such as living and supra-living systems (e.g. human societies or ecosystems). That means components of a system may change what they do by mechanisms that are just becoming clear to us. And if and when they do change it has an effect on the whole system that could not have been predicted by the shallow understanding models. A general rule is that the more complex a system is, the more potential there is for component subsystems to change[4].

With deep understanding, coming from deeper analysis of components as subsystems, it is feasible in principle to construct models that can also demonstrate what we call *evolvability*. Such models might generate an ensemble of predictions or scenarios of anticipated behaviors of

[2] An excellent review of the nature of composition of more complex systems from simpler systems as the deep history of the Universe is Tyler Volk (2017). His process of combogenesis is the story of how simple systems (like atoms) combine to form more complex composites.

[3] Dynamical systems are generally modeled as sets of differential equations integrated over time. Systems dynamics models are based on computer simulations of causal relations between macrostate variables.

[4] The capacity for complex systems, particularly ones with nonlinear internals, to undergo fundamental behavioral changes is called ‘evolvability.’ This topic will be taken up below in this Introduction and in several sections of the book.

systems that would provide advanced warning about possible changes that we should know about. That is to say, deep understanding provides us with a path toward anticipation of different future states of affairs that might help us avoid threats or exploit opportunities. Armed with advanced knowledge of how systems might behave differently due to changes reduces the surprise factor.

Or deep understanding might simply provide us with the satisfaction of knowing better how the world works. For example, cosmologists seek deep understanding of the Universe not so much to exploit it, you can't very well control a super cluster of galaxies, but to better grasp the meaning of human existence within it.

Deep understanding in biology has led to the genetic revolution, the ability to decode genes and, now, even modify genes to achieve arguably practical purposes. All of the sciences achieve deep understanding through the approach known as reductionism, the process by which phenomena in one level of organization can be understood by knowing about phenomena at a lower level. For example, living systems (one level of organization) can be deeply understood if we take into account the chemistry that operates at the molecular level of organization. What happens in the smaller scales of time and space (molecules) materially explains how living systems achieve their behavior at the larger scales of time and space (cells).

In this book, we will attempt to generalize the concept of reductionism from a systems science perspective to produce a methodology for gaining deep understanding of every kind of system. Our main concern will be directed toward what we call *concrete*, or real physical systems. This is distinguished from *abstract* systems (see below). The reason for doing this is that the greatest needs for systems analysis, modeling, and design are in relation to concrete systems, particularly human designed or modified systems. We need to understand our world, our ecosystems, our organizations, etc. because we are constantly doing things that impact them and too often do not appreciate the consequences before doing so. Having a method to gain deep understanding before taking such actions might help reduce the numbers and severity of unintended consequences.

This is not to say that abstract systems are not, themselves, subject to the same sort of methodology. Toward the end of the book we will provide a few examples of abstract systems that are in need of 'engineering' and thus need to be deeply understood. But for the bulk of this book we will focus on concrete systems.

I.1.1. Talking About Systems

It will be important to understand a significant difference between systems science and the other sciences with respect to how the subjects are described. In the sciences the subject of description is the specific substrate of the science. Biologists describe biological phenomena; chemists describe chemical phenomena. Their subjects are specific and their methods of description are particular to their domains. This is true even when the subject is a general pattern of phenomenon within that domain, such as when chemists derive laws of reaction rates or biologists expound on ‘laws[5]’ of behavior in animals. When we read a paper reporting on a biological phenomenon we expect to be told about the specific elements of the phenomenon, including names of the elements, the relations between them, and the dynamics of their interactions that, collectively produce the phenomenon. We expect to be provided with specific measurements and analysis of the data relevant to the phenomenon itself. This is basically the same for all of the sciences and is just as much the case for various comparative studies between different representative elements, e.g. between different species in biology, as for direct studies of single elements.

In systems science the situation is quite different and we have a different way of talking about systems as phenomena[6]. The concepts of systems science apply to all domains of the sciences. When we describe a system phenomenon, we are not describing a specific phenomenon in a specific domain. Rather we are describing a general or meta-pattern of phenomena that is applicable across all domains. Going beyond mere comparison studies, or even analogous phenomena descriptions, we enter a domain of description that finds isomorphic relations that are true regardless of any specific subject domain. This general mode of description is different from how we approach descriptions in the sciences as a rule[7]. In Chapter 2, on the ontology of systems, we will be providing a guide to the generality of structural and functional patterns in systems science. In Chapter 3, we will develop these ontological elements into a language of systems that will provide us with a way to talk about systemness applied to specific phenomena in the sciences and engineering fields. We will assert (with evidence) that *systemese is a universal language that allows us to describe any concrete phenomenon regardless of the substrate’s domain*.

Hence, this book will provide descriptions of general patterns, those that systems analysts will be able to use as guides to discovery of specifics within particular systems of interest. As much as possible we will also provide several examples of how the patterns pertain to several different phenomena in different knowledge domains. For example, we rely on examples from different areas of biology, physics, social sciences, and engineering to demonstrate the applicability of the patterns.

---

[5] The quotes here, and not on the chemists’ laws, is a reminder that law-like phenomena seem to get rarer as we go from physics up the ladder of complexity to biological and psychological phenomena. There has been a long active debate within the sciences and the philosophy of science as to whether there really are laws in a strict sense, or emergent ‘rules’ that govern the phenomena. See Unger & Smolin (2015), Unger’s Chapter 5 in particular.

[6] George Klir referred to Thinghood vs. Systemhood (2001). The former is the content of a science – what the thing is – while the latter are the set of properties that define systemness and are applicable across all sciences.

[7] Somewhat similar but still restricted descriptions occur in more general fields such as evolution theory, ecological theory, and reaction theory (to name a few) where the emphasis is on categories of phenomena. Ironically, however, these are the same areas that are currently morphing into systems sciences applied, e.g. systems biology!

This way of talking about systems in general terms as the major mode of description may tend to put off domain scientists who are used to reading about the specifics of a phenomenon. To some it may even seem like a kind of 'arm-waving' as a distraction from substance. But it is not. Systems science operates at a meta-science level, but we feel strongly, a necessary level in order to better talk about concrete phenomena within any of the sciences. And as will become evident as we work through specific examples using the language of systems, how having a basic knowledge of systemness will facilitate transdisciplinary communications between the domain experts. Having systems language as a kind of Rosetta Stone translator from domain to domain should significantly improve those communications, particularly with regard to the transfer of meaning (sematics).

I.1.2. The Nature of Systems Science

The study of systemness has a long history when viewed in retrospect[8]. Just as the way we talk descriptively about patterns that are isomorphic across different domains, the study of those patterns is the domain of systems science (c.f. Friendshuh & Troncale, 2012). Consider a dominant pattern of organization that is found ubiquitously in nature and human designed systems, a hierarchy of organization. Any system can be shown to be comprised of interconnected subsystems. And each of those will in turn be comprised of interconnected sub-subsystems. Each level in this hierarchy contains simpler elements or components (in Chapter 2 we will clarify the meaning of components). Eventually there is a level of the 'simplest' components (also clarified in Chapter 2). So the regress is not infinite.

Systems science is a transdisciplinary approach to discovery and characterization of the general nature and aspects of patterns such as hierarchy of organization. It is transdisciplinary in that the systems scientist must examine a wide array of systems from many different subject disciplines in order to derive the isomorphic nature of the patterns of interest. Systems science seeks understanding of the nature of reality, which we believe is reflected in the fact that these patterns can be found so pervasive.

Systems science is concerned with the general theory of systemness[9]. It seeks to investigate the general nature of patterns that apply to all systems and derive a general theory that can, in turn, inform the disciplinary sciences as they explore their own subject domains. It should also be informative to engineers when designing complex systems. This is why we think of systems science as a meta-science.

That having been said, it is important to establish a basic understanding that what is going to be covered in this book is emphatically *NOT* presented as being a general systems theory as

[8] See for example the excellent summary in Rousseau, et al (2018), Chapter 1.

[9] Ludwig von Bertalanffy (1901 – 1972), a mathematical biologist, proposed the idea of a general systems theory (GST). See the Wikipedia article: https://en.wikipedia.org/wiki/Ludwig_von_Bertalanffy for background (accessed 7/8/2019) and also the review presented in the reference in the previous footnote.

envisioned by von Bertalanffy (1969) or general systemology as described by Rousseau, et. al (2018). Rather, our claim is much less grand in scope. In Chapter 2 we present *AN* ontology that we believe covers the notion of systemness and the nature of an ontogenesis process that has produced the complex Universe we observe today. In Chapter 3 we present what we term a "*formal framework for understanding complex systems*" based on the ontological commitments made in Chapter 2. The aim of this framework is not to be taken as a general systems theory per se, but to guide the exploration of system properties in analysis and assist in designs of systems artifacts. It is possible that future research may arrive at a true General Systems Theory (labeled GST*, where the apostrophe, pronounced 'star' signifies completeness) that will contain, at its core, some of what we are arguing in this book. Time will tell.

Chapter 1 will review a set of principles regarding what is known today about the patterns that recur in all areas of knowledge. This review comes from the author's previous book, with co-author Michael Kalton, *Principles of Systems Science*, published by Springer in 2015.

I.1.3. Concrete and Abstract Systems

It will be important to differentiate between concrete or real physical systems and abstract systems (Miller, 1978; Ackoff, 1971). Examples of concrete systems include animals, organizations, nations, and the whole Earth. Abstract systems are conceptual, i.e., represented in a medium other than being embodied in a real physical, working structure and come in two flavors. Pure abstract systems are those, like the natural or real numbers, exist as mental constructs that can be used generally for purposes such as measurement (in the other form of abstract system) or mathematics. Such abstract systems are applicable across a broad array of abstract systems of the second kind. The second form of abstract system is what we generally call a 'model,' of which there are several forms (see below). We will be concerned, in this book, with the relation between the concrete and second form of abstract systems, linked by the first form of abstract system. That is, models of concrete systems using abstract representations provided by the first form of abstract system. For example, an equation of motion that describes the progression of a pendulum's motion back and forth through an arc is an abstraction of the second form of a real pendulum. It is a model of a pendulum that captures its dynamics represented in the real number domain.

The distinction between abstract and concrete systems is a little less clear than suggested by this introduction. In Chapter 1 those subtleties will be explored more thoroughly. Anticipating that exploration, however, consider that even abstract systems take on a hint of concreteness when instantiated in the human mind. Indeed, all of our mental concepts could be considered abstract systems that, because they are based on real functioning neural firing patterns, are approachable as a form of concrete system!

As one example of an abstract system of the first kind consider the theory of deterministic chaos. It can be invoked to explain a category of nondeterministic behavior in some physical systems such as the activities of the atmosphere we call weather. The models of chaos are important to our enterprise of understanding concrete systems, for certain. However, our focus will be on the internal aspects of concrete systems that give rise to chaotic behavior rather than the models that explore the concept of chaos (or any of the various forms of complexity generation). There are a large number of volumes that deal with such models (e.g. the Lorenz attractor) and we will be referencing some of them as might apply to understanding a specific example concrete system.

The relation between concrete and abstract systems follows a straightforward pattern. Our objective is to show how to construct an abstract system, a model, from a deep understanding of a concrete system, using abstract systems of the first kind (mathematical constructs). Our approach is, however, not in line with the traditional use of modeling in the sciences and engineering. That is, the use of models has been seen as the way to come to understand systems outside of deep analysis. Based on some preliminary observations of a system's behavior (as a black box) a mathematical model is constructed and solved (run) to produce a result that is then tested against the behavior of the concrete system under experimentally controlled (or naturally occurring) conditions. The tradition is based on the notion that the model need only be constructed to answer questions of interest to the modeler and need not involve details considered to be outside the scope of those questions. As will be argued, this approach becomes increasingly problematic in the face of higher orders of complexity.

Complexity is the enemy of deep understanding in a timely fashion, unless the methods employed are systemic. We advocate that the program of science be turned around; that it is possible to use principled systems science to first obtain deep understanding and then construct models that will provide veridical predictions that can be used for design of complex artifactual systems and policies for governance of human affairs.

Our objective is to demonstrate how to come to deeply understand concrete complex systems and their environments first and then construct models from this understanding. What this book will show is a methodology based solidly on systems science theory that will lead to that kind of understanding. What will be demonstrated is that a deep understanding of a system leads to much better outcomes in terms of prediction or, in the case of engineered systems, better products.

Science, to date, has been the societal process of gaining deep understanding. But, it has pursued a program of 'discovery' that involves generating hypotheses regarding the behavior of interesting phenomena (in various domains such as chemistry or biology), constructing abstract models of these phenomena, e.g. laws in physics and law-like relations in biology, and then making observations on the actual behavior of the phenomena in comparison to the mathematical solutions generated by the models. Upon discovering deviations of the models from the actual behavior (e.g. relativistic deviations from Newton's laws), science proceeds to dissect the phenomena to seek internal mechanisms, deconstruction for the purpose of discovering deeper understanding. In other words, science quite ordinarily starts with building models and then proceeds to open up the phenomenon to deeper inspection when the model fails to adequately predict the real system.

This formula – construct a model, test it, and dig deeper when the model fails – has been adequate in the history of the sciences. It works well enough when there is no hurry to gain deep understanding. However, there is reason to consider that we might not have the luxury of time anymore. In the next chapter we will consider a new approach to gaining deep understanding that turns the program on its head. The main thrust of this book is to show how systems analysis, performed in advance of constructing models, can lead to deep understanding more rapidly. As will be argued, this approach is made possible by significant developments in our ability to decompose systems without destroying the inherent interconnections between component parts.

I.1.4. Systems and Complexity

For the last half of the last century and the first decades of this century people in many different disciplines, from the sciences and engineering to politics and governance to history studies, have been confronting increasingly complex issues[10]. The nature of the objects of their primary disciplines is now seen to be more complex than was historically true before World War II. The various disciplines, in fact, were established along the lines of silos such as they are organized in the universities because traditional reductionist methodologies operating on relatively simple objects of study worked amazingly well. Figure I.1 shows a traditional hierarchy of major scientific disciplines and the root problem that runs through all of the sciences.

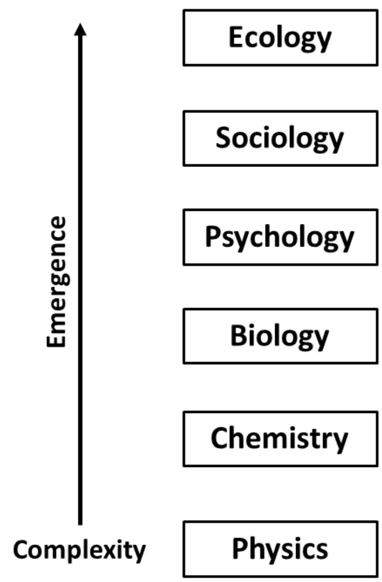

[10] Chapter 5, "Complexity," in Mobus & Kalton (2015) provides a general overview of various kinds of complexity but settles on Herbert Simon's hierarchy-based definition (see section 5.2.2, starting on page 173). The full development of the concept of complexity would require a similar chapter in this book, so we refer the reader back to this reference. Otherwise, the nature of complex systems will become evident in the descriptions that follow.

Fig. I.1. There is a natural hierarchy of major science disciplines based on several aspects. The lower disciplines are fundamentally simpler to explore using the scientific process. Their objects of study are inherently simpler in terms of numbers of components. The emergence of greater complexity is explained in the text.

What the sciences are exploring and attempting to explain are phenomena that play out at each of these levels. Physics, for example, is concerned (among many other sub-subjects) with atoms, their structures and properties. Chemistry is concerned with their interactions. So chemical bonds (explained by physics) generate chemical reactions and molecular structures. Chemistry has a direct connection to physics in this regard but works with a new level of complexity that emerged from the atomic structures studied by physics[11]. The field of study called chemistry depends on the field of study called (atomic) physics just as the existence of molecules and chemical reactions depend on the properties and dynamics of atoms.

Biological phenomena depend on chemical phenomena. The study of biology now very much depends on the study of (organic and bio-) chemistry. And that same kind of relation proceeds at each level up the hierarchy. The gaps between the major subject boxes are filling in with multi-disciplinary sub-fields. As the sciences have matured we find that knowledge from lower levels is reaching higher into the upper levels. For example, modern psychology now rests on the biology of the brain, but that in turn rests on specific chemical phenomena (e.g. neurotransmitter-receptor chemistry). The higher we go in these subjects the more we find we must understand the specific roles of phenomena that are studied at lower levels. This is a natural outcome of the fact of systemness – that systems are hierarchies of subsystems that, at some point, are based on the lower level phenomena. But it gets more complicated than that. As you go up the hierarchy of subjects (and phenomenal types) the levels below are themselves impacted by the kinds of actual structures and functions that do emerge and set patterns of phenomena at the higher level. One example to make this clear is the fact that all amino acids that form proteins in living systems have a left-handed chirality. Somewhere in the origin of life a presumed accident of nature broke the symmetry between left and right handed amino acid molecules being preferred (a form of symmetry breaking?) and it stuck in all subsequent evolution. Thus, causal arrows of influence point both upward and downward in the phenomenal processes contained in each of the boxes in Figure I.1.

At first, in the late 19th and early 20th centuries each of the sciences managed to stay within the bounds of its surface level phenomena. But the methods of reduction were increasingly successful as instrumentation and models of phenomena improved.

Then, early in the 20th century people began to run into three aspects that started complicating the grasp of knowledge. One was the increase in structural and functional complexity of the objects under study as well as the complexities of how those objects operated

[11] Physical chemistry is the interface between these disciplines so is, essentially transdisciplinary. Atoms emerged from the physical forces (strong, weak, and electromagnetic) during the emergence of the Universe (Big Bang).

in their environments (thanks to the aforementioned success of reductionist approaches but also to the fact that the objects of study were including more components at a given level). Another aspect was the realization that the behaviors of the objects being studied could not be described quite as simply as had been the case with simpler objects. Non-linear dynamics and feedback structures started complicating the approaches to analyzing objects of study. The third aspect was the recognition that the objects of study often times spanned different fields or disciplines (those causal influence arrows going both up and down).

To really understand an object of study required more interaction between experts from different fields. But this is problematic because people in different disciplines often speak very different languages and have different approaches to study. Researchers in many fields began using what we call “systems thinking” to analyze the objects. The recognition that these more complex objects could not be understood in the standard disciplinary ways led to the explosion of new disciplines that arose to address these issues of complexity and interactions between disparate kinds of objects.

Fortunately for everyone interested in developing methods for dealing with complex objects and their environments the advent of the digital computer coincided with the recognition of the problems associated with complexity and heterogeneity of component parts. Boosted with the development of the transistor and later the integrated circuit, computing became widely available and the capacity to build computational models of systems, as the complex objects of study came to be called, grew rapidly. Purely mathematical modeling, for example using systems of differential equations, was augmented rapidly by discrete recursive modeling which allowed researchers to explore the dynamics of systems with feedback. The main difference between these approaches, however, was the requirement to know how a system worked from the inside in order to specify the internal behaviors and structures that gave rise to the whole system function. It was no longer possible to simply model a system based on data collected from observations of past behavior. That approach is still used, but still suffers the same problems described above. One needed a causal explanation of how the parts fit together and worked together.

There are three basic ways to approach modeling a system with complex structures and behaviors. The method applied to many kinds of systems, especially those where exploring the internals through dissection leads to loss of function, was to use a process called system identification, which looks strictly at the mathematical relation between inputs and outputs and build a (usually) statistical function that can be used to “predict” behavior of the whole system under variations in the inputs. Such a model is often referred to as a “black box” since one cannot see inside the system to understand how it produces the results it does. As long as systems were seen to be relatively linear (or piece-wise linear) in their functions, this method would suffice to make predictions.

A second approach, used for example in biological work, is to take what one hopes is a representative sample of a population of similar systems, dissect them to find out something about what is going on inside, and use that knowledge to infer the input-output relations needed to, again, make predictions about behaviors under variations of the inputs. The latter can be determined by experiments on the still intact systems that can verify the understanding derived by the system identification approach. This kind of model is often referred to as a "grey box" since the inferences made between inner parts and outward behavior are still largely uncertain and require multiple observations and experiments to reduce that uncertainty.

A third approach that has been largely restricted to the engineering fields uses what we call "reverse engineering" or taking a complex machine apart to see what its parts are, and from a previous deep understanding of what the various kinds of parts do, work out the functional relations with fairly high certainty. The model is then called, as you have already guessed, a "white box.[12]"

This third approach leads to a system model with the highest level of capability to predict future behavior because it includes detailed understanding of what the components are, what relations they have with one another, and what independent behaviors can be expected from them (highly deterministic). Additionally, and perhaps most important, this kind of model allows the inclusion of non-linearities resulting from internal amplifications and feedback phenomena. These are the aspects of systems that give rise to really interesting (and important) complexities that need to be understood for real anticipation (prediction).

Increasingly white box-style analysis and modeling is being applied in the other sciences where the accumulation of lower levels of organization, thanks to reductionist methods, provides the same kind of knowledge of component behaviors as was the case for engineered systems. For example, the approach can be used in modeling and predicting the behavior of whole ecosystems since it is relatively easy to obtain information on the component plants, animals, fungi, bacteria, etc. and much is now known about species-determined behaviors as well as food webs and other interaction relations. It still takes a lot of work to obtain this information, but the approach of studying a system from the inside is feasible. Another example of this is when IT specialists analyze the internal workings of an organization that is in need of information systems able to support management decisions.

However, the trend toward white box analysis of living systems in biology has made a tremendous amount of progress over the last several decades as well with the development of non-intrusive instrumentation such as functional magnetic resonance imaging (fMRI), ultra-sound imaging, and now fairly spectacular abilities to see inside the brain and other tissues using

[12] Alternatively, we call these "transparent boxes" and black boxes are referred to as "opaque boxes." The author prefers the latter terms since the former are too strongly associated with machine descriptions and the latter terms are more neutral with respect to substrate media (a living system for example). Additionally the latter terms are not associated with specific disciplines.

"optogenetics."[13] These methods for examining the insides of opaque systems while the latter are actually behaving have given biologists dramatic abilities to pursue the same kinds of models that have been helpful in engineering and physics/chemistry since the last century.

Why is the ability to produce a white box model superior to grey box ones? It is the reduction in uncertainty that comes with what we have called "deep understanding" of the system. If we know more about how a system works on the inside we can more reliably predict how it will behave given new environmental conditions.

I.1.5. Deep Understanding

There are many different kinds of *knowledge*. We can know facts about things and relations between different things. But the epitome of knowledge is *understanding*[14]. This is the kind of knowledge that gives us the ability to anticipate future outcomes of processes and dynamic relations. Deep understanding comes from knowing not only what a thing is made of and how it is connected to other things in its environment, but how it works internally as well, which includes how its internal processes respond to changes in its environment. With understanding comes an ability to predict with some confidence that if the environment does something the system will respond in a particular way.

To be clear, all understanding comes from the knowledge we put into our models, be they purely mathematical, computer-based, or mental. The model is not the thing being modeled, but the intention is to obtain enough knowledge about how a system works to bring our models as close to representing the thing-in-the-world as realistically as possible. We will have much more to say about models, model construction and the construction of a knowledgebase from which those models can be generated. But for now, recognize that we are pursuing the development of a method, a process, by which we can ultimately build more reliable models by

Deep understanding is more than just being able to anticipate behaviors. A predator has to anticipate the movements of its prey in order to be ready to pounce. This it accomplishes by having learned a set of causal relations between what is going on in the environment and what the prey normally does in response. If this is the dry season, the prey is likely to be gathered at a few remaining water holes so that is where the predator goes for dinner. The predator doesn't think about the prey being thirsty, its motivation for going to the water hole. It only has a model of what the prey does, its behavior, not what causes its behavior. This is effectively the same as the correlation-based regression model discussed above.

---

[13] A very good article in *Scientific American* reports on an important new imaging technology developed by one of the developers of optogenetics. This technology allows the imaging of the living connectome (Deisseroth, 2016). See also this interesting Wikipedia article for details: https://en.wikipedia.org/wiki/Optogenetics

[14] Several authors have considered higher levels of knowledge such as wisdom – the knowledge of how to use knowledge (Potter, 1971), and higher still is conscience or moral sentiment (Damasio, 1999, page 230). These higher forms of knowledge are relevant to high-order system patterns such as agency and governance, which will be discussed in Chapter 11.

Rather, deep understanding goes beyond mere associations based on historical observations. It involves a deeper knowledge of causal mechanisms and motivations.

There are a considerable number of books covering knowledge discovery and construction methodologies. The traditional sciences, physics, chemistry, biology, psychology, etc. have long employed what is called the scientific method for this purpose. Other fields of inquiry have a variety of methods that work along the same lines as the scientific method but may not employ mathematics in the same way that the sciences do. History, for example, relies on the discovery of meanings in texts and historical artifacts. It is an interpretive method (i.e. involves some subjective thought) rather than a rigorous objective process. But it seeks the same sort of understanding about the human world that the sciences work toward.

The nature of understanding comes under the philosophical framework of epistemology (see that discussion in Chapter 2). Some of the major questions posed in this branch of philosophy are: What is knowledge? How do we humans gain knowledge? Why do we do so? There are many more, but this is the basic framework. This book will basically be about an approach to the gaining of deep understanding of complex systems.

I.1.6. Understanding Systems

How do we human beings succeed in living in the world? What does it mean to be successful? The evolutionary definition of success is ‘fitness,’ an organism’s match of physical and behavioral characteristics meeting the demands of the environment in which it lives. A species is considered fit if its member individuals are fairly successful at staying alive and reproducing at least a replacement number of their kind. This species-level fitness is compared with any possible competing species that would usurp what is called the ecological niche. For example, if two different species of predators relied on the same prey for food, the one that was best able to catch the prey would tend to out-compete the other and, all other things being equal, eventually dominate the ecological niche.

Within the species individuals tend to have a range of variations regarding any traits or behaviors such that the population average defines the species fitness as above. Some individuals may have trait variations that reduce their average fitness relative to the species norm. Others might have variations that give them advantages above the population norm. In this case if there is a selection pressure, say from a competing species, and the individual is generally more fit it will also tend to produce a larger proportion of offspring and help strengthen the species position in the niche.

Fitness, both species and individual, is a well understood concept in biological evolution and is a core explanation, along with selection, for speciation or the long-term development of organisms. But, what about human beings? We are biological organisms. In what way are we fit?

There are a number of dimensions to this question. The dominant one is the fact that human beings have evolved to be able to adapt their environments to their biological needs or invent some technology that protects them from the vagaries of different environments. In other words, human beings have the ability to construct their own niches rather than simply fit into the existing conditions. We have succeeded in the evolutionary sense. And we have done so spectacularly. Almost too spectacularly, as it turns out.

Since there are myriad books on the subject of how clever humans are, how the various revolutions in technology, the Agricultural, the Metallurgical, the Industrial (I, II, and III), and the second agricultural or Green revolutions, have allowed us to literally cover the planet, and how doing so, along with the expansion of consumption, have turned around to put us in jeopardy (changed the fitness equation), that ground will not be covered here. The question we are really asking is: What characteristic of our thinking capacity has endowed us with the ability to do clever things? The answer to this, we think, is *our ability to understand how the world, or parts of it, works*. We have the basic capacity to understand the systems we encounter and the system in which we live. Our brains have the ability to observe phenomena carefully and postulate theories about causes and effects. We can build mental models of how things work and we have conscious access to those models to use them to anticipate the future (to some extent). We have, moreover, the capacity to test our models by carefully intervening in phenomena to see if our understanding is correct. We test our understanding through ‘experiments.’ The human brain is pre-equipped to do what amounts to science, at least an intuitive form of it.

This ability is what made it possible for humans to first exploit the effects of fire, to test the way in which stones could be chipped to produce a sharp edge, to discover the heat preserving benefits of wrapping themselves in animal skins, and every cultural innovation that has emerged since.

It is the understanding of how things work, understanding the regularities in nature, which is at the root of human success as a species and as individuals. And it is the understanding of systemness, how things are interrelated and causally interact with consequences (e.g. feedback) that has made us so successful.

Unfortunately, we are not sufficiently equipped to understand broadly. We can, through our modern rigorous practice of science, understand very deeply but the vast majority of people are capable of specializing in only one or a few subjects. Polymaths are rare. Our sciences have operated for several centuries, quite successfully, by ignoring interactions between phenomena of interest with other phenomena that seem remote. We have been very successful at drilling down deeply in subjects but too often at the cost of understanding how the low-level phenomena relate to distant other kinds of phenomena. Sometimes we find out later that there actually are causal connections that were hidden from view, or at least not in our focus. Sometimes we find out because the interactions can have negative consequences. The connection between burning fossil fuels and climate change is a dramatic case in point.

We have also been extremely successful in exploiting what we have learned within various domains of knowledge. Engineering, in many different forms, involves taking the knowledge derived about an area and figuring out how to build technologies that serve our purposes. Historically this has been realized within domains of phenomena such as mechanical or chemical or electrical. Engineers, trained in the formal methods (mathematics) of the domains have been able to invent and refine technologies that have improved the standards of living for large portions of the population[15].

Once again there have been thousands of books (probably many more) on the subject of how engineering exploits understanding of phenomena so we will not spend any time rehashing the obvious. But the point we wish to emphasize is, again, it is *understanding* that is at the root of this successful strategy of living.

I.1.7. Understanding Causal Relations

In all of this what is being understood is the causal relations that exist between events that occur and consequent events, or, equivalently, states that lead to other states. These relations are most often probabilistic, that is they are related statistically, not absolutely deterministically. The great successes of science and engineering have been in refining our ability to understand these stochastic relations and still do a very good job of predicting outcomes from varying initial conditions.

I.1.8. New Kinds of Causal Relations

For the entire history of science and engineering the kinds of causal relations that had been explored and understood have been what we call linear; that is, event *X* causes transition *Y*. And event *X* was caused by transition *W* that happened in the immediate past. We have learned how to trace linear causations backward (abduction) to discover what caused the observed event, and forward (deduction and induction), discovering what an event will cause in the future. We have learned how these causal relations are moderated by stochastic (probabilistic) factors like noise or nonlinearities that give rise to chaotic dynamics.

But now we are realizing this is not sufficient for complete understanding[16]. As we have gained greater understanding of how things work in our various disciplinary silos we have also begun to realize how connected these things are across silos. Even within silos we have begun to realize that multiple phenomena at lower levels of organization contribute to the phenomena we observe at higher levels. For example, we have known for a long time that organic chemistry

[15] Note that it is the political-social framework that prevents these advantages from being shared globally. There is no reason that people all over the world should not benefit from the standard of living provided by engineering, except for political machinations and the fact that the human population growth is out of the normal controls.

[16] Complete understanding is a relative term! Complete suggests absolute, and science does not deal in absolutes (except possibly in temperature). The progression in science and engineering is based on “more complete” understanding, which characterizes a process of getting closer to the “truth” but not claiming to have arrived as the ultimate truth.

must be understood in order to grasp biochemistry, which, in turn, needs to be understood in order to grasp metabolism, which, in turn, needs to be understood in order to grasp cells… and so on. We've also known for a long while that organic chemistry is predicated on not just the peculiarities of carbon, but the quantum properties of the electron shells of carbon atoms. Chemistry depends on physics.

The sciences have been able to move ahead because it is possible to "abstract" the aggregate features of the lower levels of reality as we move up the hierarchy of physical organization. They have been able to act as if the lower levels are just "details" so long as those lower levels of phenomena acted in consistent ways. Fortunately, they mostly do.

What has started to become obvious is that as we move up the hierarchy of organization we also move up in terms of complexity. Phenomena that could not have been predicted strictly on the mechanics of the substrates start to emerge as those mechanics give rise to more complex structures. The origin of life problem is a case in point. In theory, we should be able to explain how metabolism in ancient cells originated and entered into biological evolution. In practice, we have not yet achieved that goal.

I.1.9. Complexity and Interactions

What we humans are coming to understand, though we don't understand it well enough yet, is that our whole world is a system, an extremely complex one. It is composed of myriad heterogeneous subsystems (see principle 1 in chapter 1). They are coupled through various cyclical flows (e.g. carbon cycle, hydro cycle, etc.) that are driven by massive energy flows through the planet. The subsystems are themselves composed of sub-subsystems that interact within their parent system with one another and some of them interact with other subsystems in the larger system. These interactions can be characterized also as flows of material, energy, or influence (information) that act to shape and regulate the structures and functions within.

In a very real sense, though difficult to perceive at times, everything is connected to everything else. Pull on this string and it is likely that many things will result in different parts of the system changing in some way. Any change to one part of the system is likely to affect every other part of the system even if the effect is infinitesimal at the outer reaches. Some changes may be amplified by nonlinear processes and have significant impact even though rather minute in origin.

And this is where our problem arises. We humans are part of the larger Earth system. Whatever we do there will be some effect generated that affects the other parts of the system. We have grown in size and power as a collective force affecting a complex network of components. And, as we are now aware, it is too often in negative ways at the scale of our global civilization.

Even at the other end of the scale, at the local level, our technological wizardry is showing signs of fraying. Each new release of computing and communications technologies seems to have more bugs than the last release. Our attempts at producing artifacts like self-driving cars are beginning to demonstrate the limits of our ability to design super complex devices that are supposed to operate in a super complex environment. As this is being written there have been two deaths reported associated with such cars along with property damage. Those are unwanted interactions between one kind of complex system and others (human beings, the legal system, the economic system, etc.)

The simple fact is that we have reached a likely point of diminishing returns on the ununderstood increase in the complexity of our inventions. And the negative consequences are starting to show. That complexity arises from the desire to create new kinds of machines that combine diverse technologies in seeking additional functionality (the current buzz words are ‘cyber-physical systems’ and ‘Internet of Things’, IoT). For the most part, today, that means integrating computational and communications capabilities and the programs that exercise the mechanical parts of the machines. In turn, this means bringing together engineers from different backgrounds with different languages and different methodologies who must nevertheless work together to produce an advanced design. On top of that newer complexities have arisen. Environmental hazards, energy efficiency, material resource efficiency, recyclability, and many other factors have to be considered in designs.

Complexity is the bane of achieving understanding and designing new artificial systems. Systems science provides a direct way to grapple with complexity. The principles discussed in Chapter 1 provide a pathway to overcoming complexity in the pursuit of understanding. As discussed below, intuitions about systemness, what we call ‘systems thinking,’ is a necessary condition but insufficient when striving to overcome the vagaries of complexity. What is required is the discipline of systems science and the methods of systems analysis that are designed to work directly with complexity so as to manage it. The first step is to be clear on what we mean by a system and the principles of systemness. This will be the main task of the chapters in Part 1.

I.1.10. Simple to Complex Adaptive and Evolvable Systems

Throughout the book we will be referring to four categories of systems based on degrees and forms of complexity. The degree of complexity is characterized by factors such as the number of elements and kinds of elements, and the number of levels in the structural/functional hierarchy as described in Mobus & Kalton (2015, Chapter 5). The forms of complexity are related but have to do with a system’s capacity to interact with changes in its environment that deviate from operating norms. In general, the more complex systems have more capacity to cope with changing environments. Figure I.2 shows these categories and provides some examples.

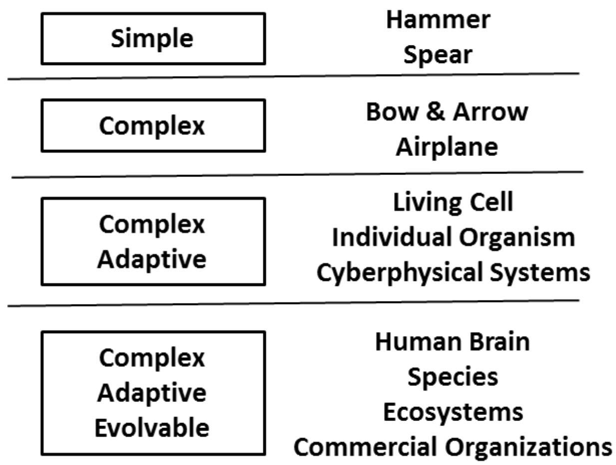

Fig. I.2. Systems may be classified into these four categories according to measures of complexity based on the heterogeneity and number of parts as well as the numbers of levels of organization and by behavioral characteristics.

The four system categories are: Simple, Complex, Complex Adaptive, and Complex Adaptive and Evolvable systems (SS, CS, CAS, and CAES). We will have very little to say about simple systems (e.g. clocks or automobiles). We have some to say about complex systems (e.g. jet liners of space shuttles), which may be complicated but still not capable of the facility which seems to differentiate non-living from living systems. That capability is an ability to adapt to changes in the environment that go outside of a nominal operating range. Adaptive change generally requires the system to alter its internal distributions of material and energy to compensate for changes. It is a key characteristic of living systems that can alter physiology or behavior to address external stresses brought on by environmental changes. Homeostasis is the paradigm example of an adaptive response.

Much has been written about CASs in a variety of arenas, such as complexity science. Most writers seemed not to distinguish between mere adaptivity in individual systems versus evolutionary reconfiguration in more complex systems such as populations (species). In (Mobus, 2015) that distinction and a new class of complex systems was introduced and explained. Complex adaptive and *evolvable* systems (CAES) go beyond mere alterations in internal distributions and shifts in work processes for compensation of stresses. Evolvable systems are those that are able to adopt or eliminate functions (and the structures that perform them) de novo. That is, a system that has the capacity to incorporate a completely new mechanism for obtaining and processing a new material resource (e.g. a company starting a new product line), performing a new function and perhaps finding a new purpose, has evolved. In the living world, pre-humans, individuals within a species were adaptive, but not evolvable. The species, or more generally the genus, as the system, was evolvable. Throughout the book we will refer to these categories of systems as they are pertinent to the theories or examples. It won't be until Chapter 9, however, that we will examine the full meaning of the categories. There we will elaborate on what we call model archetypes, starting with the CAS and CAES generic models. We will also explain the nature of several subsystem model archetypes, an agent, an economy, and a governance system that constitute the inner workings of the CAS/CAES archetypes.

I.2. A Systems Understanding Framework

An early approach to a knowledge framework can be found in Ackoff (1971). Ackoff calls for a "system of system concepts," noting that many terms and concepts used by systems thinkers were not sufficiently organized. The objective of this book is to develop a holistic system of systems understanding methodology based on the principles of systems science but for the purpose of doing system science and engineering. The same methodology, acting as a framework, applies to all of the sciences when they are directed at exploring the systemness within their individual domains.

I.2.1. Conventional Systems Understanding

Before proceeding to the outline of the framework proposed in this book, we should make comment on the current widely practiced framework for understanding systems. The concept of modeling (especially mathematical and computer simulation) and systems understanding, have gone hand in hand since the outset of the systems approach. In this framework the researcher starts with a hypothesis-like[17] conjecture about how a system works, what its relevant parts are, how they interact with one another, and what the relevant inputs and outputs are. They then construct a model using formal tools like a system of ordinary differential equations to describe the system as they conceive it. Using computing tools to 'solve' the system output at some future time, given the inputs at time t_0[18] they project the behavior of the system. Then it is necessary to observe the real concrete system in action to see if the predicted behavior is borne out in real life. Under the gold standard of scientific investigation the researcher may be able to set up controlled experimental conditions to test the hypothesis. Otherwise they need to hope for conditions that closely approximate the model inputs and record the behavior of the system as it reacts.

In any case the traditional approach is:

1. Make educated guesses about the nature of the concrete system and what are its important features
2. Construct an abstract model of the system based on the suppositions made
3. Run the model (do the computations) capturing the outputs as data
4. Compare the model output with the concrete system's behavior to see if there is correspondence

[17] It is hypothesis-like, in that the conjecture says, in effect, the model I have created constitutes the important features of the real concrete system. If the two systems, the concrete and the abstract, both have the same behavior, within some arbitrary degree of accuracy and precision, then the hypothesis is NOT disproved. Otherwise, it's back to the drawing board.

[18] Most models require continuous input data as time passes.

5. If yes – conclude that the hypothesized mechanisms are correct, otherwise tweak what you guess (educated, of course) is off in the model and run it again.
6. Repeat this process as often as necessary until the model behavior sufficiently replicates that of the concrete system.

This approach is really a reflection of the conventional scientific method and empiricism dating back to, for example, Galileo's attempts to grapple with how things fall to earth. Observe, hypothesize, test, and revise as needed. The history of science has been a gradual teasing out of details when the testing indicated more knowledge was needed – further dissection was required to gain a better understanding of the phenomenon of interest.

The framework being proposed in this book turns this conventional approach on its head. Note that the decomposition of a phenomenon or system ends up taking place in any event, driven by the need to make the models better. The model, in this conventional framework, has been used as a spur to the action of deeper analysis more than a confirmation that our original guesswork was spot on. In what follows we suggest that the process of science, and systems science in particular, has reached a level of maturity in which we can effectively do the deep systems analysis before necessarily producing models. In other words, the models we end up with are generated after we know what the system is and how it works at deeper levels. We are aided in this new approach by the fact that our knowledge of, for example, measurement theory, along with the state-of-art in advanced and inexpensive sensor technology, makes it possible to 'instrument' systems in ways not achievable even a few decades ago. We have non- or minimally-intrusive sensing methods (e.g. functional magnetic resonance imaging, fMRI) that allow deconstruction of concrete complex adaptive systems in ways never imagined previously. We've had advanced digital computing for many decades and so the ability to create models has been the main recourse for science and systems investigation. But now, for many systems of interest, we can analyze first and model second (and verify third is still a good idea).

One might be tempted to ask, then, what are models for if not to help us gain better knowledge? Well, they can still work in this way as needed. The process of functional/structural decomposition that will be described can still be spurred by models used as a tool when it is absolutely necessary to 'guess' about mechanisms that we cannot easily directly analyze. However, much of what we really want from models is the ability to predict the future or, at least, project scenarios to aid us anticipate possible futures and make some decisions on actions that would be preemptive to prevent harm or take advantage of novel opportunities. With the proposed methodology, we will have models that are readily useful in this context with much less of the 'development' phase shaping the model from the test-it-and-see approach.

A particularly important use of models that we will explore in later chapters is that they are used in the cybernetic problem of control, management, and governance (see Chapter 11). What we call "decision models" are used to process real-time data and make control decisions to correct errors. Such models cannot be based on guesses, no matter how educated they are. We must be able to generate veridical models early if they are to be used for this purpose.

I.2.2. Outline of the Framework

This framework is comprised of seven basic subsystems, four of them activities directly involved in obtaining system knowledge, one a governing process, one a monitoring process that provides information to the governing process, and one a core process, the knowledgebase, the structure of which links the activities together. Chapter 4 will provide a comprehensive introduction to these activities and structure. Here we will only introduce the ideas and a preliminary justification for why they are necessary for the entire enterprise of system understanding.

The first part is the activity called systems analysis (SA). This is the activity in which a system is first identified as an opaque-box[19] entity and then, using a structured analytic approach, converted into a transparent-box entity. Or rather a set of entities organized in a hierarchy of increasing details. The core structure (the second part) is a system knowledgebase, implemented in a hybrid of relational-object database technology along with Web technologies seen in examples like Google™ (for indexing and search) and Wikipedia (for editable content) to keep track of the system knowledge obtained through analysis. The third part involves the methods for generating various kinds of models from the knowledgebase for various uses. Most important among these are aids to mental modeling (i.e. diagrams) and building computer simulations for testing hypotheses. The fourth activity concerns the development of human engineered systems, artifacts, procedures, and policies. These address the issues of systems engineering that apply to all kinds of complex designed systems from products to organizational systems to governance. Finally, all of these parts require that the concrete systems themselves, natural or engineered, be monitored on an on-going basis with the predictions provided by the models (hypotheses) being verified by actual behaviors. These five will be expanded below.

However, the framework requires something more than a set of activities. Essential to the whole enterprise is the ability to use language to communicate the knowledge gained between multiple stakeholders or participants. In complex heterogeneous systems there will be many of these with many different (yet related) perspectives, and often from different disciplines using different languages. So before any sort of mechanical framework can be developed and used it is essential to consider what language will allow these multiple stakeholders to work together and gain the full value of understanding the system of interest. Chapters 2 and 3 will provide the explanation of the language that will make the understanding of systems feasible for all parties. Below is a brief expansion of the ideas of these five parts, but starting with the nature of a system language.

[19] We have adopted the term 'opaque-box' to replace the conventional 'black-box' terminology from reverse engineering. Opaque implies the possibility of transforming the opaque boundary to a transparent one.

I.3. Sharing Understanding

Understanding is a community process. No one human can possibly grasp every aspect of every phenomenon of interest. But, as a social enterprise, we should, at least in principle, be capable of deriving arbitrarily deep understanding of any phenomenon we encounter. Just as importantly we should be able to understand the connections between that phenomenon and all other phenomena as we expand our vision from a local to a universal scale.

Assuming a language that permits effective communications between members of the social system, the gradual buildup of multi-perspective understanding, along with methods for synthesizing seemingly different perspectives, through intersubjective integration provides a route to shared understanding.

I.3.1. Communications

Humans talk to each other. They use language to communicate ideas, facts, concepts and meaning. Natural languages, which evolved as different cultural traditions spread out from Africa, resemble the speciation represented in phylogenetic trees in biology. And thus people from different language groups have had difficulty communicating complex ideas, etc. Throughout history peoples from different backgrounds have had to work hard to learn translations of terms and interpretations of syntactical structures in order to share thoughts.

The same situation exists today for the various tribes of sciences. Each science has evolved different languages tuned to their individual disciplines. Even though English has become the dominant natural language of the sciences the vast differences in terminology, meaning, and technical details of the objects of study in each discipline remains a drag on cross-disciplinary communication. Why this is problematic is that in today's world of complex objects of interest multiple kinds of disciplines have to be called upon in order to study (analyze) them. The Human Genome Project is a case study in this problem. Calling upon scientists from computer science, genetics, chemistry and several other support sciences, the project required that all of the participants basically learn a common language that borrowed terms and concepts from all of the disciplines. Geneticists had to become familiar with the ideas of computational complexity while the computer scientists were called upon to learn the chemistry and structural aspects of DNA and genes. Initially this was challenging. Fortunately, DNA is organized as a string of "letters" that encode a "message". Computer scientists rapidly recognized the similarity between this and data structures with which they were intimately familiar and were able to develop efficient algorithms for deconstructing long strings of DNA to decipher genes. Thus a sufficient commonality between genetics and computer science concepts allowed for the rapid deciphering of the genetic code.

It is possible that there is a similar Rosetta stone underlying the commonality among all science languages. In fact, this is the basis for why systems science can be considered as a meta-science. The claim, explained in Mobus and Kalton (2015) and in Chapter 1, is that systemness is universal, meaning that it is at the base of all objects and structures in the world. It follows that the language of systems should be a universal language that covers all phenomena (ideas and concepts) of interest. It could, therefore, provide a means for translating specific terms and meanings in one scientific language into any other language. An example will be provided below.

Chapters 2 and 3 are devoted to the development of a language of system. This is more than a language used to model systems. There are many of those that have been developed specifically to build abstract models of systems. Most have been designed to aid software and systems engineers capture their thoughts about what a to-be-designed system should be. A few, such as system dynamics (SD) were designed to capture abstract representations of existing systems (like the world – Meadows, et. al, 1976). All of these languages exist in order to produce abstractions of the real systems. Some of these abstractions can be turned into computer code and run in simulation to produce an abstract representation of the behavior of the system model (which is presumed to represent the behavior of the real system, either an existing one or one yet-to-be built).

These languages serve an important purpose in allowing researchers or engineers to visualize the large-scale aspects of systems and their behaviors but do not produce the kind of deep understanding we described above. The models produced using these languages are simply great elaborations on the kinds of models discussed regarding shallow understanding. They are only as good at producing predictions or anticipations as the languages used permit and the modelers skills allow. While this has proven good enough in many historical instances, it is not given that they will continue to be sufficient as we move inexorably into levels of complexity of systems that seem certain in the future. What is needed, in our opinion, is a language of system that is meant to actually describe in sufficient detail all aspects of a system of interest (and its environment). An abstract model might always be obtained from such a description if the language is sufficiently formal in its structure. At the same time, such a language permits many different disciplinarians to share their understanding as we try to do with natural language. Rather than reduce a system description to a set of state variables and functions alone, and those at a very abstract level for computational tractability, the language of system should be ‘speakable’ as well as ‘viewable’ as well as ‘computable’. People should be able to talk to each other without over-abstraction causing the loss of details and meanings that support real understanding.

I.3.2. The Language of Systems

From another angle consider an interesting idea in linguistics. A number of linguists, psychologists, and philosophers of science are proposing that the human brain has an internal, subconscious language that is essentially preprogrammed into the circuitry. They call it the Language of Thought (LoT) or “mentalese” (Fodor, 1975; Pinker, 2008). The reasons for this proposal are too complex to go into here, but suffice it to say that there is a growing body of evidence that this is a valid notion. Let’s go one step further to ask what are the symbols and syntax of this language? Assuming it is a universally possessed internal language – that is, all humans possess the same language – then that means there is already a basis for a common language. One more step. Suppose this mentalese evolved specifically to represent the world as humans found it[20]. What would be the most successful way to think about the world? If the world is a system of systems, then the language of thought should reflect this. It should be a language of systems, or *systemese*!

As it turns out modern functional neuroimaging is providing some preliminary support for this conjecture (see Mobus & Kalton, 2015, chapters 8, section 8.2.5 and 9, and Appendix B in this book). All human beings recognize and talk about system concepts but using different terms and syntax and often having different surficial meanings. For example, the system concept of a “stock,” a reservoir for a substance (matter or energy) or data is universally used in many different contexts with seemingly different meanings. A retail store manager talks about her inventory of goods to be sold. A private citizen talks about the money he has in a bank account. An ecologist talks about the amount of water in a lake. An electronics engineer talks about the amount of charge contained by a capacitor in a circuit. All four use different words to describe the components of a system that temporarily ‘holds’ a substance, along with associated notions of how the substance comes into the stock and how it goes out. All four have intuitions about the phenomenon of how much stuff is in the stock and that inflow and outflow rates work to determine this and it doesn’t matter what spoken language they speak. A stock element of a system is a general and universal concept as is the dynamical behavior of its ‘level’ due to difference in inflows and outflows.

The general claim is that the human brain of every individual on this planet is designed to grasp this general idea and all of the ideas of systemese. Chapters 2 and 3 will show how we derive a general systemese and give explicit (and English) names to them. Below we will introduce just a few systems concepts to make this notion more concrete.

The human brain doesn’t just deal with names of things (nouns), relations (e.g. prepositions), and actions (verbs). It is also capable of quantification (and giving names to generalized forms) as well as constructing abstract relations between quantities (e.g. measurement, comparison, and calibration). That is, it attaches to concepts like inflow, stock, and outflow, magnitudes of those flows as per units of time. These may be qualitative, as in “that stream is bringing a lot of water into the lake.” Or they can be more explicit, “the stream is deeper and wider after the rain so much more water is coming into the lake.” Or it can be more explicit still, “the stream depth is twice as deep and twice as wide after the storm…” The brain has the ability to quantify measures, usually by comparison, and attach those to the phenomena

[20] One implication of this notion is that some form of mentalese (systemese) is present in all brains throughout evolution. In other words, worms’ very primitive brain-like structures process a very primitive version of systems, namely they are ‘aware’ of limited environmental sources (e.g. food in the soil) and have sensory and processing ability to determine the internal states of their bodies and make response decisions that result in overall behavior.

of interest. Thus, communications consist not only of terms and syntax, but measures of magnitude (starting with sensory intensity measures) that add meaning to the terms. Handling quantification requires an addendum to ordinary language. We call it math.

What ties our quantification (concrete and abstract) to linguistic (concrete and abstract) representations is a kind of middle ground. As soon as humans evolved the language capability they needed to also construct arguments to present to one another when making social decisions. Those arguments had to make sense in order for the species to achieve and maintain fitness in the late Pleistocene environment (relatively frequent climate shifts). Argumentation gives rise to logic that is valid reasoning forms such as deduction and useful heuristic forms such as induction, and abduction.

A full description of a system requires all three modes of thinking. Linguistic descriptions may be thought of as the skeleton onto which the flesh of mathematical and logical relations is attached for clarity and precision. Linguistic elements of language are most closely attached to our perceptual experiences (see Appendix B for a treatise on representation of concepts of systems in the brain). Thus, it forms the core of description. Mathematics and logic refine these descriptions, quantitatively and relationally[21].

I.3.3. The Story of S – System Narratives

Communications is not a simple description of something for human beings. We do not share detailed descriptions. We share stories that tell of things that happen to the subject of the story, that tell of things the subject did, that tell of the subject's responses to situations, and so on. Human beings are story tellers. Indeed, every sentence that can be uttered is a micro-story. Every paragraph tells a scenario. A chapter tells a history. We don't just blandly describe the subject, we tell its story and that includes the story of what went on around the subject.

Systems thinking is effectively a form of third-person perspective in a narrative. The story is told from the vantage point of a quasi-omniscient narrator who can look down upon the subject and its environment and describe the unfolding story of its interactions with the other entities in its world.

One way to see this is to consider the story told by a system dynamics model using one of the popular SD languages, like STELLA. What the model simulation does is generate a story about what the values of key variables are over the course of the time of the simulation. The graphs that an SD language produces are graphical representations of what happened as the story unfolded.

[21] In addition, our vision processing module is interoperable with the linguistic module so that visual images play a role in describing models, e.g. flow diagrams, visual maps, etc. We will describe this more fully in Chapter 3.

A really good language of systems must provide adequate descriptive power, but also adequate narrative power. In being able to tell us what the subject has done it can tell us something about what we can expect the subject to do in the future.

I.3.4. The Mathematics of Systems

The formal language of science and engineering is often said to be mathematics. This has been demonstrated abundantly in all of the sciences and since what all sciences investigate is the systems in their domains of study, and all of the sciences use various levels of mathematics, it follows that all of mathematics is relevant to systems science. But the notion that mathematics is *THE* language of science, engineering, and systems is too simplistic. For human understanding it requires linguistic, mathematical, and logical reasoning elements to describe things and situations. The linguistic components provide the semantic grounding and the logical and mathematical provides two basic augmentations, the quantification of state variables (e.g. the level of substance in a stock of fixed volume) and a way to produce abstract models of the system. All three of these aspects are needed in order to talk about systems.

However, this book seeks to apply what is learned in systems science to a generalized methodology for gaining deep understanding of systems regardless of the domain of study. We employ the framework described above, the description of a universal structure that captures all of the aspects of systemness, to act as a form of skeleton upon which special forms of mathematics can be attached as needed for specific systems. For example, we will not be using differential equations to attempt to describe systems. However, specific elements of specific systems contained in our framework may have places for the use of differential equations during the building of specific models. Moreover, the logic of systemness is built into the framework. Where it might improve understanding of why a specific structure is the way it is, this logic will be explored more formally (e.g. temporal logic and its role in causal reasoning).

While a main objective of the author's is to make systems science and these methodologies as widely accessible as possible, to help any educated person elevate their systems thinking and awareness, it will take more than just spoken/written language. We will be explicating all of the concepts and methods in this book in both language and mathematics. And we will see that the two go very much together toward gaining understanding. The level of math, however, is not so high as to exclude the average educated reader. Most high school graduates, these days, have had to take basic algebra and much of our mathematical formulations will be such. However, the mathematics that will be used to build our framework is called discrete math or sometimes discrete structures.

I.3.5. Discrete Mathematics

We will be using four interrelated topics within discrete math (there are many more topics typically covered in a course in discrete math). The term refers to the fact that the math itself is based primarily on the idea of working with integers as opposed to real numbers. We can't avoid the real numbers entirely of course. The discrete math we will encounter is based on discrete "operations" which may or may not involve real numbers. But those operations are generally easier to grasp for most people, say as compared with the calculus! If the use of continuous math would improve an understanding in the pages that follow, we will isolate the math into a boxed-off section and work extra hard to improve the descriptive material so that readers will not lose anything if they do not understand the continuous math.

We can get by with avoiding or even ignoring the continuum because of an interesting property of systems. Systems are composed of subsystems in a hierarchy of smaller scales of time and space. And what appears to be a continuous substance at one scale is actually found to be composed of discrete objects at a much smaller scale. For example, the water flowing in a stream appears to be a continuous substance at the scale of human perception. But if we could put a really powerful microscope on the water we would discover that water is actually composed of independent molecules made of two atoms of hydrogen with one atom of oxygen. Water only appears to be a continuum from the perspective of ordinary human perception. The whole world that we humans deal with is composed of atoms and molecules, so technically, if we choose our scale of measurement properly, we can work with purely integer values. Of course, this is totally impractical for many situations. Counting molecules of water or electrons passing a particular point during a particular time is really infeasible. So we resort to mass averages (e.g. statistical mechanics) and real number representation (and calculus!) Most of the concrete systems we will be using as examples will require some real number representations but will not require calculus. For example, rates of flows will be couched in discrete time intervals (called a time series) even if the mass flowing during a discrete interval has to be represented by a real number of finite precision.

Readers unfamiliar with the subjects: set theory (including fuzzy set theory), graph theory (including flow networks), probability theory, and combinatorics, may want to spend some time studying the main results in these maths. We will not be doing math, in the sense of proving theorems or the like. We will, however use these mathematical concepts in the construction of a basic and (we hope) universal definition of system that will provide a foundation for the language of system as well as the design of a knowledgebase in which to hold what we discover in a structured way, recoverable for a number of different purposes including building models.

Even if the reader is not certain of their level of understanding of the mathematical statements made, we are dedicated to providing enough explanatory verbiage that they will not feel they are missing anything essential to at least form a 'better' understanding, if not a deep one.

I.4. Deep Systems Analysis (DSA)

Before anything about a system can be truly understood it is necessary to do a thorough analysis. By thorough we mean deconstruction of the system and all of its subsystems down to the level where the components are so fundamental (and well understood) that no further deconstruction is needed. For very complex systems this can seem like a monumental, if not impossible, task. What this book is going to show is that, while the task is monumental in some cases, it is far from impossible. In fact, we will argue that it is absolutely necessary in order to assure success of the overall objective of deep understanding. Furthermore, in the case of human engineered systems it is absolutely mandatory for a successful deployment of the system upon implementation.

The term ‘systems analysis’ has been used in a number of different contexts (Checkland, 1999, page 134). It conveys roughly the process of inquiry into what a system of interest is, or what a system of ‘desire’ should be. The methods of inquiry have varied significantly from one domain (e.g. information systems) to another (e.g. policy design). Yet they all have two things in common. They are all based on a relatively loose kind of systems thinking, or more appropriately, systems awareness. And they all attempt to go about the inquiry systematically, even if not guided by a principle-based approach. The systems analysis methodology introduced in this book attempts to base the process on the principles of systems science directly (Chapter 1). What we will consider is a procedure that works in two different contexts because it operates from the basis of a single set of principles. Namely it provides a way of finding out *what exists* (science) and it provides an aid to thinking about *what needs to exist* (design). The latter context can be split into two major arenas: the engineering of physical systems and the design of human activity systems (Checkland, 1999). Figure I.3 shows the general relations between these contexts. Checkland makes a significant distinction between system types, in particular between engineered (hard) systems and organizational (soft) systems and argued that the systems approaches appropriate to each of those are quite different. The argument advanced in this book, however, is that even though there are differences in these systems, they are still systems in a general sense and so many of the methods given here, especially systems analysis, provide a unified approach – there need not be different methods applied so long as the methods are based on system principles and not just ad hoc systems thinking[22].

DSA is a top-down discovery process with bottom-up refinement that seeks deep understanding. More traditional engineering and organizational approaches have been based on a variety of methods, including a bottom up, integration of modules, to a reliance on ‘requirements’ of the users/customers that presume the latter know what they need. The typical engineering analysis (including software engineering) targets the technical details (and costs) of

[22] Checkland’s concern for differentiating were largely due to the attempts that were made in the 60’s to apply engineering approaches to the design of ‘soft’ systems like companies or government agencies, where human factors could not be treated as parameters in an equation. At the time systems engineering was still in a nascent state and a number of systems engineers and mathematicians were attempting to define systems from their perspective alone. This led to what this author considers some premature and limited notions of a definition of system (c.f. Wymore, 1967, Chapter 2).

how to meet the requirements without ever necessarily asking whether those requirements are the right ones[23]!

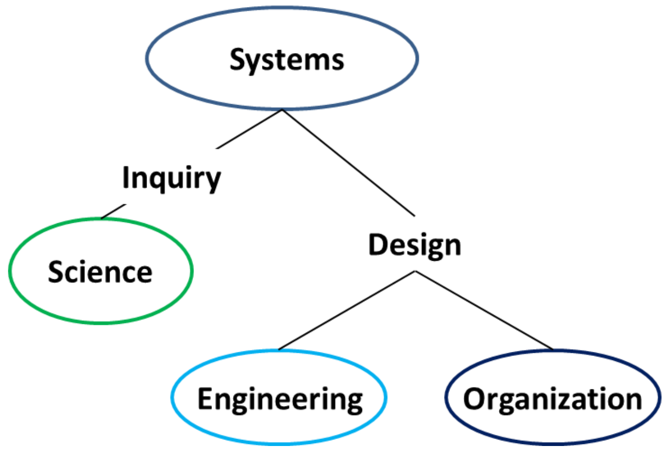

Fig. I.3. Science (broadly defined) uses systems analysis to discover the nature of what already exists. Design seeks to produce something that needs to exist. Engineering is directed toward the design of machines, whereas organization is directed toward those systems that include human agents, such as the economy, businesses, schools, etc. This classification roughly follows Checkland's (1999) classification of 'natural', 'hard', and 'soft' systems.

As will be explained in Chapter 4, and then thoroughly developed in Part 2 of the book, DSA involves several phases.

I.4.1. System Identification – Boundary Determination

The first phase involves actually identifying the system of interest (SOI) and establishing a boundary that demarks the system's insides from the rest of the environment. Sometimes this is a simple as identifying a living cell as the SOI and seeing the cell membrane as the boundary. More often this is a complicated task. As explained in Mobus & Kalton (2015) there are many kinds of boundaries, including conceptual boundaries that are often hard to recognize. In Chapter 5 we will cover some of the various techniques used to identify boundaries and "boundary conditions." Then in Chapter 6 we will provide a few examples of kinds of boundaries in real systems and how to identify them.

Along with the determination of the boundary itself is the identification of all of the inputs and outputs that pass through the boundary. This means finding and quantifying (as best possible in this phase) the substances that enter and exit the system as well as the messages that it receives

[23] As discussed in Chapter 5, this 'starting point' of analysis, wherein some questionable user requirements are included in the initial specification of a system has caused innumerable failures of projects in the past and has led modern systems and software development methodologies collectively known as 'agile.' These methods essentially abandon getting an upfront right specification of a system and instead adopt an iterative refinement method that includes having the user on the development team so that as the errors of design become apparent, the user member can change their specification on the fly.

and sends, and the disturbances that, strictly speaking, are not inputs but may be found to affect the inner workings of the system when deconstruction is being done.

Outputs include a system's overt behavior. Some systems move or generate forces all of which have to be accounted for in the analysis.

A major objective of this phase is to characterize the relation between inputs and outputs. To the degree possible given the complexity of the system and environment, the objective is to establish a quantitative function describing the output given the input. There are a number of techniques that will be discussed in Chapter 5 for accomplishing this. The importance of doing so is that this knowledge will provide guidance to the deconstruction phase.

I.4.2. Environment Analysis – Sources, Sinks, and Disturbances

Every real system, with the possible exception of the Universe as a whole, exists embedded within a larger supra-system. The Earth is embedded in the solar system and obtains its energy input from Sol. It receives a smattering of material inputs in the form of meteorites and comet dust (an occasional meteor of some mass). The biosphere is a subsystem within the whole Earth system (which we refer to as the 'Ecos' – Greek for 'home'). The human species and its cultures are a subsystem of the Ecos as well. A for-profit enterprise is embedded in a market/economic system. A child is embedded (we hope) in a family system. And thus it goes.

No system exists as an isolated thing. If any such system did so, it would soon succumb to the second law of thermodynamics and cease to be a system at all – just a collection of bits and pieces in thermodynamic equilibrium. Thus, the understanding of any system depends on some understanding of its environment as well (what we will call the 'embedding' environment).

The environment is comprised of other entities that provide resource and disturbance inputs (sources) and product/waste outputs (sinks) to/from the system of interest. Each and all of these need to be identified and their output/input flows quantified. Multiple sources could contribute the same kind of substance to an input to the SOI. Multiple sinks might absorb the outputs of a system. The relative contributions of all of them need to be identified and measured.

This is of fundamental importance. Every concrete, real, system is subject to the influences of inputs and outputs. Therefore, it is essential to identify all such source and sink entities and measure their rates of flow of material, energy, and messages in order to even begin understanding the nature of the SOI. Inputs other than disturbances (e.g. sudden changes in flows or thermal noise) are considered resources for the system, or otherwise they would not get the status of inputs. Similarly, outputs are either useful to other entities in the environment, or they are wastes that need to be absorbed and neutralized. For example, consider that every system involves a work process in which input energy at a high potential (and appropriate form – called exergy) is used to drive the work. The work might involve the transformation of matter, other forms of energy, or messages (e.g. computation) and will always result in the loss of some portion of that energy in the form of waste heat (the rest presumably used up in the work). Again, this is the second law of thermodynamics and absolutely must be taken into consideration in any truly deep understanding of systems. The heat must be transmitted to the environment and the total energy flow must account for it. Almost none of the existing modeling languages make provisions for this important aspect of concrete systems.

The first step in systems analysis is an analysis of the SOI's environment. We do not attempt to model the sources and sinks other than their outputs/inputs (usually via a flow function). As we will discuss in Chapter 5, there are times when accounting discrepancies are noted during subsequent systems analysis between the environmental flows and the internal system flows. At such times it becomes necessary to reverse the direction of SA and do what we will call a 'supra-system analysis,' that is including the source or sink in the boundary in order to resolve the discrepancy. An example of this can be found in the analysis of a supply chain for a manufacturing company in which the input of a specific part comes from a single source and thus the entity that supplies that part is more strongly coupled to the original SOI (the manufacturer) than initially understood. The single source supplier should probably be included within the boundary of the original SOI because it is acting as a resource obtainer for the manufacturer, i.e. is actually a subsystem of the SOI.

I.4.3. Recursive System Deconstruction

The general method of SA will be a top-down recursive search of the system level tree. After the environmental analysis has established the main inputs and outputs of the SOI the procedure opens up the opaque box to look at what's inside. It looks at the subsystems that comprise the whole system and determines the mapping of inputs to internal subsystems (those that work to obtain resources or deal with disturbances) and outputs from subsystems (those that export the products and wastes from within the SOI).

As will be shown in Chapter 2 subsystems of an original SOI may be treated (once discovered) as systems in their own right. That is, by focusing in on a subsystem as if it were an SOI, we can deconstruct it in the same way we do the original. The trick, here, is to treat the other subsystems as if they are the environment of the newly chosen SOI. We ignore their internal workings (for the time being) and simply concern ourselves with the mapping of inputs and outputs in exactly the same fashion as we did with the original SOI.

This constitutes a recursive method for analyzing the system from the top down to the level of components. The procedure used assures us that we account for all material/energy/message flows in accordance with the conservation laws and the second law of thermodynamics. By capturing these in the structured knowledgebase, we insure that discrepancies will be flagged and provide a way to re-analyze system to ensure we get it right.

Systems analysis as describe in this work is an expensive proposition. This will not sit well with managers or scientists concerned with short-term profit or budgets. Our justification for this intensive procedure, to be explained in Chapter 5, is summed up in an epithet often heard among working engineers, "Why is it there isn't time to do it right the first time, but there is always time to fix it?" Profit-motive driven (or cost averse-driven) managers are often motivated to accelerate the designs of systems. Even budget-motivated scientists are driven to short-circuit analysis in order to produce a result within (or under) budget approved by a national funding agency, and on time. Yet with the methodologies in current practice they invariably find themselves having to repeat or redo their work. We will argue that the problems being faced by large-scale and complex projects are the result of a failure to deeply understand a problem or requirement and thus lead to a failure of production of a final result.

I.5. The System Knowledgebase

With the definition of system to be worked out in Chapter 3 we will provide a structure of knowledge that emulates that of the human brain, and so, is compatible with human communications of ideas. We will, in Chapter 7, explore the construction of a knowledgebase – a database with a specific structure for encoding knowledge[24]. This knowledgebase will be shown to emulate the knowledge of concepts held in the human brain and thus be able to provide a basis for communications between humans and between humans and machine (computational models). This is a major leap forward in the conceptualization of systems and models.

Fundamentally, the argument is that the human brain is designed to capture, and build, concepts that reflect the real world. We learn about the world from experience in it. The method by which this is accomplished is to encode percepts, concepts, and thoughts in neural networks. But those networks are direct representations of the real world. This will be explored in Chapter 3 in the context of the language of system, then further developed in Chapter 7 in describing the knowledgebase.

A major feature of the knowledgebase design is that the knowledge may be checked for internal completeness and consistency automatically. These checks will report any lack of adequate knowledge, thereby providing the systems analysts with guidance about what they need to go back and find. This is similar to the syntax checking that computer language compilers do to flag errors before trying to generate machine code.

I.5.1. Constructing Models from the Knowledgebase

Part 3 of this book is devoted to the methods of obtaining models from the knowledgebase. The knowledgebase is, in fact, a very large, detailed model, just as the concepts encoded in neural tissues are mental models. In principle it would be possible to generate a simulation model of the whole system from the knowledgebase once it is filled from analysis. In practice this might be infeasible even with today's massively parallel computing systems. What will be shown in Chapter 8 is how to generate abstract models and sub-models (models of the subsystems down

[24] The idea of a knowledgebase, as opposed to a mere database comes from Principle 9, discussed in Chapter 1. Knowledge is organized and effective data – models, rather than just aggregations of data. Knowledgebases are accessed differently than databases – by association rather than by indexical search.

the hierarchy). Once the knowledgebase is sufficiently complete and consistent internally various kinds of models, including current model types such as SysML or System Dynamics can be produced for the various purposes such models serve.

I.5.2. Engineered Systems

The process of systems analysis done on to-be-designed systems provides design-for-free. That is, at the point of a completed knowledgebase everything about an engineered system to be specified and constructed has been captured through DSA. Chapters 12 and 13 will be devoted to showing how specifications for some number of designed systems can be derived directly from the knowledgebase. These chapters provide a principled approach to systems engineering. Part of the current engineering process involves modeling systems after the design has been specified in order to check the functionality and veracity of the system. Since in this methodology the models will have already been simulated and tested as part of the systems analysis process, all that is left is to extract the specifications for the new system directly from the knowledgebase. This is possible because the process of functional/structural deconstruction led to complete descriptions of the systems from the top level down to the component level. The behavior knowledge captured at each level is the basis for generating specifications. Assuming that another knowledgebase – e.g. one of standard components available from vendors – is available the generator need only search through that knowledgebase for component specifications meeting the functional/structural description is the system knowledgebase. The process of generation runs in a bottom-up manner, that is, specifying the lowest level subsystems first in a breadth-first manner. Since the knowledgebase already includes the information needed to integrate modules at a level of deconstruction, that task can be automated as well. Finally, the generator would be able to generate test specifications – the modeling should have included determining correct outputs given sets of inputs. Those can be used to create test results to be used in constructing test specs.

The level of automated specification generation under this process is extensive. But unless some form of exceptional AI is employed it is not completely without need for human intervention. As with current engineering methodologies, it will rely on human observations and recognition of problems when they arise. Humans will be needed to conduct the analysis, to control the generation of models, to check the results at each stage, and to intervene when there is misinformation in the knowledgebase (e.g. from faulty analysis). These difficulties are natural parts of any large-scale, complex system design process. They will not be avoided by using this process. Rather, what we claim, is the process will tend to minimize errors in the early stages and thus produce more reliable designs at the specification stage. It can be no more ‘perfect’ than the people who use and manage it. As with any tool set, this one depends on the craftspeople who use it.

I.5.3. Monitoring and Feedback

Recognizing that there is no perfect way to design a system, or understanding an existing system, the process ultimately depends on learning from experience. We learn from mistakes and from when a system fails to behave as expected. But in order to learn we have to continue to collect data on the system's actual behavior over time. The last stage of our process requires the instrumentation of systems to continue monitoring their behavior. For natural or existing systems this will tell us if our initial analysis was correct (or good enough). Or it might alert us to changes in the environment of the system that need to be factored into our understanding. The situation isn't much different for engineered systems. Detecting errors in comparing our behavior expectations against actual behavior will tell us that either we got the initial design wrong (or a bit off) or that something else has changed, either in the environment or in one or more components of the system.

Either way we are alerted to the need to iterate the process to find out what has changed (or where lay the error). No matter how good our tools for analysis and design, these conditions will always prevail. The world changes. Complex systems change (at very least they degrade over time). Fortunately, the formal description of systemness presented in Chapter 3 and developed into the schema for the knowledgebase (and hence both the generation of models and of specifications) includes the ability to capture life-cycle information where it is known. The fact that we include an ability to designate both environmental and component parts as "evolvable" means that we can model life history aspects of systems. Currently there does not appear to be any such facility in the stock of modelling languages. If changes occur the system has to be modified manually or sent to decommissioning. Yet even with a capacity to anticipate future states of the system as it undergoes changes we will still need to collect information about the actual life of the system to better inform our knowledgebase. The need for monitoring and feedback never goes away. It is part of the cost of doing business in the real world.

I.6. Synopsis and Motivation

Part 1 – Foundations of Systems Understanding will provide, first, the theoretical underpinnings of systemness and why it can be understood. We start, in Chapter 1, with a review of the general principles or attributes of being a system that were given in Mobus & Kalton (2015). Then chapters 2 and 3 provide the developments of an ontology of systems, that is the question of 'what' exists in a systems framework, and a language of systems that will provide the basis for doing DSA and capturing the results in a knowledgebase. Chapter 4 then tries to situate what had been developed in the overall process of systems understanding. That is we provide more details about how of the different stages of the process, introduced above, take advantage of the theory of systems developed in chapters 2 and 3 (with the principles from Chapter 1 forming the backdrop).

In Part 2 – The Approach to Gaining Knowledge, we unpack the first two stages of the process, the analysis of a system and the knowledge that we acquire therefrom. Chapter 5 goes into the details of analysis, demonstrating the nature of systems analysis as generally a top-down recursive procedure using the system language (SL) that we developed in Chapter 3. We show how that language guides the process of analysis in that its syntax tells us what we should be looking for next as we construct an abstract model of the system. In Chapter 6 we will pause to let what we showed in Chapter 5 sink in. We provide several examples of the application of analysis to a wide variety of things we would call systems, each from different domains of knowledge, e.g. biological or organizational. We will show how the process of analysis is carried out exactly the same in all cases, hopefully demonstrating its universality.

Chapter 7 then returns to the main thrust of the process. We take a much closer look at how the knowledgebase is designed in accordance with the mathematical definition given in Chapter 3 and is thus in conformance with the language structure developed in that chapter. It will be made clearer how the information captured in analysis, by being stored in a specific database architecture becomes knowledge in the sense it will be defined in Chapter 2. In anticipation of that development, we will just say that knowledge and information are not the same thing even though they are related deeply from a cosmological ontology point of view.

The conclusion of Part 2, Chapter 8, will be a single example of an incredibly complex system being analyzed. We will show how something as complex and "fuzzy" as the economy might be tackled using the methods thus far presented. This chapter holds some surprises for both seasoned systems scientist and traditional economists. We explore new ways to look at the economy as a system rather than as a free-form phenomenon.

In the next two chapters, Part 3 – Archetype Models, we consider the nature of complex, adaptive, and evolvable systems (CAES), the most complex kinds we know of. Chapter 9 is devoted to an overview of the CAES concept. This archetype model is the integration and amalgamation of a number of historically important systems models that other researchers/thinkers have proposed in the past based on their observations of systems that were of interest to them, such as the Viable Systems Model (VSM) of Stafford Beer and Living Systems of James Grier Miller (and many other contributions) along with updated theories applicable to the sub-models of a CAES.

Chapters 10 through 12 present these sub-models. Chapter 10 is devoted to an updated understanding of Agents and Agency as the computational and decision processing elements in the next two sub-models. These are the Governance and Economy sub-models respectively. These are treated in more detail than is described in Chapter 9.

Finally, in Part 4, we address the practical uses of the process for deep systems understanding for producing complex systems artifacts. These include everything that humans invent and produce for use, including not just machines and cities, but procedures and policies; essentially anything that is the product of the human mind and did not arise from "natural" evolution. In Chapter 2 we describe this so-called natural evolution as the result of non-intentional auto-organization and selection by environmental factors. In Chapter 13 we introduce the notion that human's in their ability to imagine new organizations that are 'intended' to solve problems or provide new capabilities transcend auto-organization and natural selection with intentional-organization and intentional-selection. What we typically call human culture, the aggregate of all artifacts along with beliefs, norms of behavior, and other mental properties of humans, has become a major part of the human environment, even more important, for most people today, than the so-called natural environment.

In Chapter 14 the process by which CAESs can be brought into existence, through intentional-organization bolstered by systems engineering, resulting in artifacts that serve their intended purposes while minimizing unintended consequences. The latter aspect of human culture has been the source of many collapses of societies throughout history. Those societies were small and reasonably local as compared with the situation today in which cultures have global reach and impact.

There are growing signs that this cultural environment has developed in such a way that threatens human well-being. We will address this issue in the final chapter of the book in which we ask, can the use of deep understanding of the human social system (HSS) as an exercise in design and engineering produce a culture that is more conducive to both the human psyche and the rest of the natural world, and if so, what would it look like in broad outline?

The last chapter, 15, is devoted to exploring the larger implications of designing systems that involve both human subsystems and the larger supra-system of the Ecos. Human beings and their social activities have already impacted the whole Ecos in ways we never would have imagined even a half century ago. The question now is how might systems science and the deep systems understanding process described in this book actually be put to the purposes of designing management subsystems to restore balance to the Ecos and make the human subsystem provide a purpose to that larger system.

Why should we consider the contents of this book? There are practical reasons in terms of doing better systems engineering. Some of this motivation will be covered in Chapter 1. There are purely intellectual reasons, in terms of the satisfaction of knowing how things work to a deep level. But as much as anything a more compelling motivation is that humanity is finally coming to realize that our world is in danger and the problems are systemic. We are just at the beginning of a process of understanding the systemness of the world. So far problems such as global warming and climate change and peak oil/energy and the role of neoliberal capitalism in driving these are just coming into focus in their own lights. But more and more, people are coming to realize that they are all interconnected. Thus, we have need of a scientific way to make those connections and grasp the various dynamics involved if we hope to begin solving these existential threats. Perhaps the process of gaining deep understanding described in these pages might help.

I.7. References and Further Reading

Ackoff, RL (1971) Toward a system of systems concepts, *Management Science*, Vol 17, No 11: 661-671

Damasio, A (1999) *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*, Harcourt Brace & Co., New York

Deisseroth, K (2016) *A Look Inside the Brain*, Scientific American, Vol. 315, No.4, Oct. 2016.

Fodor, JA (1975) *The Language of Thought*, Harvard University Press, Cambridge MA

Friendshuh, L, & Troncale, LR (2012) "Identifying fundamental systems processes for a general theory of systems," In *Proceedings of the 56th Annual conference, International Society for the Systems Sciences (ISSS)*, July 15–20, San Jose State University, 23 pp

Klir, GJ (2001) *Facets of Systems Science*, Second Edition, Kluwer Academic, New York

Mobus, GE & Kalton MC (2015) *Principles of Systems Science*, Springer, New York

Mobus, GE (2015) "A Systems Science Framework for Understanding the Nature of Governance", Proceedings of the 59th Annual Meeting of the International Society for Systems Science - 2015 Berlin, Germany, Vol 1, No 1 (2015). Available on-line: http://journals.isss.org/index.php/proceedings59th/article/view/2497, accessed 2/19/2019

Pinker, S (2008) *The Stuff of Thought: Language as a Window into Human Nature*, Penguin Books, New York

Potter, VR (1971) *Bioethics: Bridge to the Future*, Prentice-Hall, New York

Rousseau, D, Billingham, J, Wilby, J, & Blachfellner, S (2018) *General Systemology: Transdisciplinarity for Discovery, Insight and Innovation*, Springer, Singapore

Unger, RM & Smolin, L (2015) *The Singular Universe and the Reality of Time*, Cambridge University Press, Cambridge UK

Bertalanffy, L. von, (1969) *General System Theory*, George Braziller, New York

Volk, T (2017) *Quarks to Culture: How We Came to Be*, Columbia University Press, New York

Wymore, AW (1967) *A Mathematical Theory of Systems Engineering – The Elements*, John Wiley and Sons, Inc. New York

[more to come]