Adv. Complex Systems (1999) **4**, 349–370

A Subjective Measure Of Complexity

Guido Fioretti †
International Institute for Applied Systems Analysis
Laxenburg, Austria

ABSTRACT. This article presents a quantitative measure of complexity, subjectively understood as a property of the relationship between a system and its observer instead of as a property of the system itself. Within this framework, complexity is quantified by assuming to know the mental categories and the mental model by which a system is represented in the mind of its observer.

It is argued that this subjectivist concept of complexity is not in contrast with objective measures of complexity introduced in particular domains, but generalizes complexity to domains where no objective measure is feasible. An extensive numerical example is presented and thoroughly discussed.
KEYWORDS: Mental categories, mental models.

1. Introduction

Complexity is often thought as a property of the system under observation, much in the same way as its mass or its volume. Taking this premise various measures have been proposed, tailored to express complexity in specific settings.

In general, these measures are not in competition with one another, since they only claim validity on a restricted domain. However, their very specificity runs contrary to the development of a general measure of complexity, one that makes sense in any domain of scientific research.

An alternative approach, pioneered by R. Rosen (1985), is that of viewing complexity as a property of the relationship between a system and its observer: the person who is observing a system says that "this system is complex" when she is not satisfied with the mental model she has of it. Note that the shift in perspective implied by a definition of this kind can be easily accommodated with any objectivist idea of complexity, as the following considerations suggest:

† Contact Address: Guido Fioretti, via Angeloni 11, I-05100 Terni

1) *Statistical complexity* is zero when entropy is either zero or maximal, while it takes its maximal value at an intermediate entropy level (Crutchfield and Young 1989; Crutchfield 1994). From a subjectivist point of view we can say that, as long as the observer of a system has a model of it that is based on Newtonian mechanics, the complexity of this system increases with its entropy. But sooner or later a point is reached where complexity is so high that the observer is led to shift to a model based on statistical mechanics, ouch less 'complex' than the previous one: intuitively, subjective complexity rises and falls in much the same way as its objective counterpart.

2) *Computational complexity* (Solomonoff 1964; Kolmogorov 1965; Chaitin 1966) differs from entropy because a given set of characters can be obtained by different rules: for instance, the infinite sequence $\{3.141\ldots\}$ can be either obtained by listing all its characters, or by telling the rule by which π is obtained. In the first case computational complexity is infinite; in the second case computational complexity is finite, and it is also very low. From a subjectivist point of view we would say that, as long as the observer of the sequence $\{3.141\ldots\}$ has no better model to reproduce it other than to list its characters, the computational complexity of this string increases with its length. But if this observer realizes that this sequence is π, he switches to a simpler model, and complexity drops. Again, subjective complexity follows the same pattern of its objective counterpart.

3) *Mean mutual information*, meant as a measure of spatial correlation, is often proposed as a measure of structural complexity (Bennett 1988, 1990). Intuitively, this meaning is not very far from the one employed by evolution theory (Conrad 1990; Kauffman 1993), where 'complexity' denotes the number and the extent of epistatic interactions; in particular, the last of the above authors stresses that evolution tends to select an optimal level of complexity, intermediate between order and chaos. Here the subjectivist could remark that, in practice, any observer of any system must choose to focus on some causal relationships and to neglect the rest of them, until he has constructed a model that is rich enough to yield interesting results and at the same time simple enough to be manageable. Thus, the intermediate level of complexity "at the edge of chaos" can also be seen as a property deriving from the economy of the observer's mind.

The above considerations do not lead me to claim that objective measures of complexity should be abandoned, since they are valuable and useful in the particular contexts they have been created for. Rather, these considerations aimed to suggest that a subjectivist view of complexity is perfectly compatible with any existing objective measure. Later, I shall propose a subjective measure of complexity, and I shall also observe that there are situations where only subjective measures of complexity are feasible.

The rest of this essay is organized as follows: section 2 examines conceptual issues that aim to clarify the subjectivist point of view on complexity, section 3 introduces a characterization of mental models and defines complexity with reference to a given mental model, and section 4 presents a quantitative measure of this complexity. In order to ease understanding, section 5 illustrates this subjective measure of complexity with a numerical example. Finally, section 6 investigates the connection between subjective and objective measures of complexity outlined above.

2. Complexity and Complication

Why is it useful to have a subjective measure of complexity? Can subjective and objective measures of complexity coexist? And if so, which approach should we prefer?

In order to begin to answer these questions, let us first remark that objective measures of complexity are possible if:

i) It is possible to identify some 'elementary particles' the system under investigation is made of. This does not mean that these particles, in their turn, may not be constituted by more basic components; rather, it means that it is not necessary to analyze their inner structure in order to account for the phenomena we are interested in.

ii) The causal relationships that link the events generated by these elementary particles to the actions undertaken upon the system are clearly defined, in the sense that we must not ask whether the egg or the chicken was born first (no self-referentiality).

If these conditions occur, and if it is difficult to describe the behavior of our system, we say that "complex systems are systems that are made by a large number of interacting parts", or similar statements. Furthermore, if the scientific community agrees to some measure of the difficulty to describe this system (e.g. statistical complexity, mean mutual information, etc.), we have an objective measure of complexity for a particular scientific domain.

Conversely, no such objective measure of complexity is possible if the above conditions are not satisfied: if condition (i) is not satisfied we are free to enlarge or contract our possibility space according to (subjective) convenience, while if condition (ii) is not satisfied we must arbitrarily choose one out of two possible causal directions, each of them leading to a different "eigenbehavior" (von Foerster 1976). On the contrary, a subjective measure of complexity is always possible, even if its evaluation might be unpractical: constraint (i) can be overcome if subjective mental categories can be used in place of unknown elementary particles, while constraint (ii) is uneffective because the subject who evaluates complexity can arbitrarily cut any eventual self-referential loop.

Are there scientific domains where at least one of conditions (i) and (ii) above does not hold? Subatomic physics, obviously, as well as the social sciences, since in these disciplines there is no obvious choice of "elementary particles". And at least in the social sciences, also because agents form expectations of one another's behavior, thereby making causal relationships indeterminate (e.g. Do savings cause investments, or is it the other way round?). Biology also knows indeterminacy of causal links (the egg and the chicken), but indeterminacy of elementary particles is limited by our knowledge of DNA code.

On the contrary, computer science and classical physics are the typical domains where conditions (i) and (ii) above hold, and where objective measures of complexity have been found. In computer science, the “elementary particles” are objectively given: they are the symbols a computer works with, and they do not arise out of measurement of a physical magnitude (Pattee 1989, 1993; Cariani 1998a, 1998b). Contrary to computer science, classical physics does involve a translation of continuous and open-ended phenomena into symbols, but since the scientific community agrees on the “elementary particles” it is sensible to focus on, these “elementary particles” can be safely regarded as objectively defined. Finally, in both disciplines self-referentiality cannot arise, since the systems they investigate do not form expectations of their observer’s’ behavior.

However, most situations are not so clear-cut. Objective measures of complexity, with all their advantages in terms of universality, can only arise out of widespread consensus about the “elementary particles” to focus on, as well as on the direction of causal relationships. Thus, to some extent even the possibility of an objective measure of complexity is a subjective matter.

In general, the process of reaching a consensus on which set of particles should be regarded as “elementary” and the process of reaching a consensus on a set of causal directions keep step with one another; thus, it makes little sense to distinguish between cases where only (i) holds, or only (ii), or both. Taking inspiration from Dupuy (1982), let me propose the following preliminary definitions:

- If the observer knows that he did not identify all the relevant elementary particles, or if he knows that he did not identify all the causal relationships that link the events these particles generate to the actions he undertakes upon the system, or both, then he says that the system is *complex.*
- If the observer did identify all the relevant elementary particles, as well as the causal relationships that link the events that these particles generate to the actions he undertakes upon the system, but he is not able to express all these relationships by means of a universal law, then he says that the system is *complicated.*
- If the observer identified all the relevant elementary particles as well as the causal relationships that link the events these particles generate to the actions he undertakes upon the system, and if he is able to express these relationships by means of a universal law, then he says that the system is *simple.*

According to this classification, objective measures of complexity are actually measures of complication. If complexity (in the above sense) can be measured at all, it must be a subjective measure, i.e. one that is relative to the model of the environment the observer is using. A subjective measure of complexity reduces to a measure of complication if conditions (i) and (ii) above hold; furthermore, if different observer’s agree to a single measure of complication, this measure becomes objective. Finally, if the observer is able to subsume the system’s behavior by means of an elegant and universally accepted law, any measure of complexity must yield zero, and the system is said to be “simple”.

At this point, let us turn back to the question this section began with: Should we discard objective measures of complexity in order to use a subjective measure? I claim the contrary. Objective measures of complexity are the outcome of painstaking research, creative theorizing, and time-consuming search for consensus. They are highly valuable, because they yield standard values that subjective measurements, by their very nature, cannot provide. Nonetheless, in many fields the possibility of constructing an objective measure of complexity is not in sight, or it looks like a hopeless enterprise. On the contrary, upon assuming of knowing an individual's mental model a subjective measure of complexity is always possible.

3. The Mental Model

A subjectivist approach to complexity can only produce a quantitative measure if it assumes to know the mental model of the observer of a system. Instead of the elementary particles of the system itself, the object of our measure will be the mental categories that this mental model is grounded upon. In this way, our complexity measure is subjective in the sense that it is grounded on an observer's mental model; at the same time, it is objective with respect to a given mental model.

Mental models are constructed upon mental categories. These are not, in general, sets of objects that share some common features: rather, each member of a mental category has some common feature with some other member, but does not necessarily share it with all of them (Lakoff 1987; Clark 1993). Mental categories are constructed around prototypes suggested by empirical experience; they are flexible, fuzzy aggregations that change with time according to the purposes of the organism they direct (Barsalou 1987; Hampton 1993).

Mental categories are not constructed for their own sake, but are shaped by empirical experience along with a mental model that provides orientation in decision-making. The function of the mental model is that of providing a framework for what is "normal" to expect. It is a net of causal relationships the decision-maker expects to hold - perhaps not always, but at least in the rule (Hebb 1949; Hayek 1952; Johnson-Laird 1983). The idea underlying this essay is that an observer views a system as "complex" if the mental model he has of it does not work, and that if the structure of this mental model is known, complexity can be quantified.

Let us make the following simplifying assumptions:

A1. Let us describe the relationships between system and observer in terms of the actions the observer may undertake, and the results the system returns. One must not suppose that these actions and results are fixed once and for all; on the contrary, in general their number as well as their qualitative features evolve with time in some unpredictable way. Categories of actions and categories of results are the only mental categories we shall consider; consequently, the internal state of the observer's mind can only be represented by the structure of the connections among these categories.

A2. Categories of actions and categories of results generally change with time, in response to the changes that actions and results themselves undergo.

However, the very fact that these categories are there in order to put different items under the same heading to make them comparable over space and time, ensures that action categories and result categories change at a much slower pace than actions and results themselves. Let us express this assumption in a quite extreme form, assuming that the mental categories by which the observer perceives reality do not change during the time interval where we measure complexity. Let $\mathcal{A} = \{\mathrm{A}_1, \mathrm{A}_2, \dots \mathrm{A}_\mathrm{I}\}$ be a finite set of mental categories that we assume to be available to the observer to classify the actions he undertakes, and let us denote them by index $i = 1, 2, \dots I$. In a similar fashion, let $\mathcal{R} = \{\mathrm{R}_1, \mathrm{R}_2, \dots \mathrm{R}_\mathrm{J}\}$ be a finite set of mental categories that we assume to be available to the observer to classify the results he obtains; let us denote them by index $j = 1, 2, \dots J$.

A3. Let us suppose that the mental model consists of a set of causal links from the categories of actions to the categories of results, which represent causal relationships that are deemed to be "normal" by the observer of the system. Fig.(1) illustrates an example of a mental model.

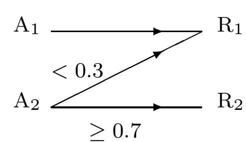

Figure 1: The upper part of this mental model is deterministic: the causal link runs from A_1 to R_1. The lower part is stochastic: the mental model says that the causal link runs from A_2 to R_2 with probability greater or equal than 0.7, from A_2 to R_1 with probability less than 0.3.

A mental model like that of fig.1 says to our observer that if he undertakes an action that belongs to category A_1, a result that belongs to category R_1 is the normal outcome. On the contrary, action category A_2 is connected to two result categories. This means that by undertaking an action of category A_2, either a result in R_2 or in R_1 is the normal outcome. Thus, this mental model is partly deterministic (the link between A_1 and R_1) and partly stochastic (the links between A_2 on one hand, R_2 and R_1 on the other).

With assumptions A1÷A3, the definitions of simplicity, complication and complexity given in the previous section become:

SIMPLICITY: If the observer of a system has been able to identify appropriate "elementary particles" and univocal causal relationships that link the events these particles generate to the actions he undertakes upon the system, if he has been able to build a deterministic model out of these relationships, and if this mental model is never contradicted by experience, then he will say that the system is "simple".

COMPLICATION: If the observer of a system has been able to identify appropriate "elementary particles" and univocal causal relationships that link the events these particles generate to the actions he undertakes upon the system, if he has only been able to construct a macroscopic stochastic model of these relationships, and if this mental model is never contradicted by experience, then he will say that the system is “complicated”.

COMPLEXITY: If the observer of a system has not been able either to identify appropriate “elementary particles”, or to identify univocal causal relationships, or both, and if his tentative mental model of the system is (at least occasionally) contradicted by empirical experience, then he will say that the system is “complex”. Note that it does not matter whether the observer's mental model is deterministic or stochastic.

Let me illustrate these definitions with an example. 1) You throw a die. If you do it just in order to wile away the time, you don't even look at the face that comes up: your mental model has one action category, “throwing the die”, and one result category, “the die rolls, then stops”. It is a deterministic model, it is never contradicted, so you say that the system is “simple”. 2) You throw a die. If you are interested in which face comes up, you may have a mental model with one action category, “throwing the die”, and six result categories, one for each face. Your model is never contradicted, but it is a stochastic one, and it doesn't help you a lot to predict which face will come up: you say that the system is “complicated”. 3) You throw a die. The die explodes, because it was secret agent 007 who left it on your table. Whatever mental model you might have, it is contradicted by empirical experience! So you say that the system is “complex” (if you are still alive).

In general, the whole of all possible actions an observer can undertake towards a system, as well as the whole of all possible results a system can return, can be neither described nor conceived: in principle, they can be anything the universe can produce, including surprising events like exploding dice as in the above example. Actions and results do not even constitute sets, since sets are containers of objects that share some common property, while actions are results that can exhibit absolutely new features.

Mental categories, in their turn, have been supposed to be available in fixed number and types during a certain time interval. Obviously, they cannot recognize any feature beyond those they have been designed to discriminate. However, these mental categories do not stand alone; rather, they are embedded in a net of causal relationships that constitute a mental model. If novel actions and results appear, they can be detected by the fact that they cause the mental model to fail.

Namely, the basic idea underlying this paper is that we can measure complexity by the extent to which a mental model does not work. In other words, given a mental model, complexity must depend on the number and the kind of causal relationships that actually occur, and that were not contemplated by the model.

Let us focus on the case of a deterministic mental model, first. A deterministic mental model looks like the first line of the mental model depicted in fig.(1): it is a set of one-to-one correspondences between action categories and result categories, with $I = J$ (if one action category would be connected to two result categories, it would be a stochastic model; if two action categories would be connected to one result category only, the two action categories would be indistinguishable from one another and could be combined into a single one). Thus, in order to detect complexity we must look whether connections appear other than the one-to-one connections of the mental model.

Complexity is due to qualitative features and causal relationships that surprise the observer as absolutely “new”, since they were not contemplated by his mental model. It follows that only recent information is relevant, and only qualitatively, not quantitatively. Consequently, in order to compute complexity, information must be processed according to the following two rules:

1) When calculating complexity, the observer has a very short memory: connections between action categories and result categories that do not belong to the near past, are forgotten;
2) When calculating complexity, the observer does not care whether a causal connection that does not belong to his mental model occurred one, two, three times or more, but only if it ever occurred, or not.

Let us consider a stochastic mental model now. Here causal relationships are expressed in probabilistic terms, so complexity computation does not only depend on the occurrence of causal connections that are not contemplated by the mental model, but also on how often they occur. For instance, if the mental model thinks the connection from A_1 to R_1 to occur with probability lower than 0.2, and if this connection occurs so often that a probability of 0.9 is measured, this does influence complexity.

However, we can formally reduce the stochastic to the deterministic case by means of the following procedure:

1. Attach the probability of obtaining a result category to the category itself, eventually keeping the original category if it can also be reached by a deterministic causal link.
2. For each probabilistic result category, create a new result category with residual probability; these new result categories are not connected to any action category.
3. Compound all the probabilistic result categories that can be reached by the same action category into one new result category, thereby establishing one-to-one relationships between action categories and result categories.

Fig.(2) illustrates this procedure for the mixed deterministic-stochastic mental model of fig.(1).

The resulting mental model is a set of one-to-one correspondences between action categories and result categories, just like a deterministic model. The only structural difference is the presence of some isolated result categories, that represent residual events.

Conditions (1) and (2) above hold still, but in the following weakened form:

1’) When calculating complexity, the observer has a relatively short memory, barely capable of a rough estimate of connection probabilities between action categories and result categories. Connections that do not appear in this narrow window are forgotten.

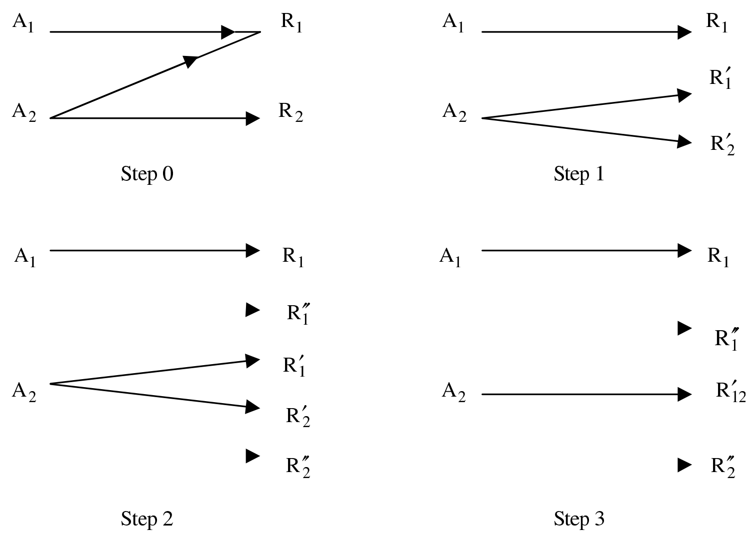

Figure 2: At step 0, the mental model is that of fig.1. At step 1, R'_1 is defined as $R'_1{:=}R_1 \wedge \{p\,(R_1 \mid A_2) < 0.3\}$; R'_2 is defined as $R'_2{:=}R_2 \wedge \{p\,(R_2 \mid A_2) \geq 0.7\}$. At step 2, R''_1 is defined as $R''_1{:=}R_1 \wedge \{p\,(R_1 \mid A_2) \geq 0.3\}$; R''_2 is defined as $R''_2{:=}R_2 \wedge \{p\,(R_2 \mid A_2) < 0.7\}$. Finally, at step 3 category R'_{12} is defined as $R'_{12}{:=}R'_1 \vee R'_2$.

2’) When calculating complexity, the observer does not care at which point of the allowed interval the probability of a causal connection fell, but only whether it is within or outside that interval.

In the end, both in the deterministic as in the stochastic case, our task is to evaluate the extent to which the structure of links between action categories and result categories departs from a set of one-to-one connections. The next section illustrates how this can be accomplished.

4. The Measurement of Complexity

The idea underlying a subjective measure of complexity is that of substituting (supposedly) known mental categories for unknown elementary particles. It assumes that mental categories can be detected even if elementary particles cannot, and also that such a flexible thing as a mental model is stable enough to be described by a set of connections between fixed mental categories. Furthermore, it assumes that the observer’s judgment of a system as being “complex” can be described by means of an algorithm (even if this does not imply to assume that the brain actually carries out algorithms).

As explained in the previous section, both deterministic and stochastic mental models can be described in terms of a set of one-to-one correspondences between action categories and result categories. Thus, we shall assume that the observer of a system has a mental model constituted by a set of one-to-one correspondences between action categories and result categories in his mind, and that he says that the system is "complex" when other correspondences occur. Thus, our measure of complexity is basically a measure of how tightly connected mental categories are.

The subjective measure of complexity presented here owes a lot to R. Atkin (1974, 1981), who used algebraic topology to represent hierarchical relations between the individuals of a society. The same tool will be used here to describe relations between mental categories. However, the "hierarchies" investigated herein actually border triviality since we have just one level of action categories and one level of result categories.

We already denoted by $\mathcal{A}$ and $\mathcal{R}$ the set of action categories and the set of result categories, respectively. Let us denote by $\mathcal{R}_i \subseteq \mathcal{R}$ the subset of result categories that are connected to A_i, and by $\mathcal{A}_j \subseteq \mathcal{A}$ the subset of action categories that are connected to R_j.

A 'simplicial complex' represents geometrically the connections between action categories and result categories. It is composed of as many simplices as the number of action categories, where the vertices of simplex A_i are the elements of $\mathcal{R}_i$. If there are isolated result categories, let us stipulate that each of them belongs to a virtual simplex that coincides with the result category itself.

If empirical experience confirms the mental model, i.e. if the correspondences between categories of actions and results are all one-to-one, simplices are isolated points and no simplicial complex exists: complexity is zero. On the contrary, if at least two simplices have at least one vertex in common, a simplicial complex arises: complexity is greater than zero.

For instance, consider the correspondences between action categories and result categories of fig.(3a): it is $\mathcal{R}_1 = \{\mathrm{R}_1, \mathrm{R}_2, \mathrm{R}_3\}$, $\mathcal{R}_2 = \{\mathrm{R}_2, \mathrm{R}_4\}$, $\mathcal{R}_3 = \{\mathrm{R}_3\}$, $\mathcal{R}_4 = \{\mathrm{R}_2, \mathrm{R}_3, \mathrm{R}_4\}$, $\mathcal{R}_5 = \{\mathrm{R}_4, \mathrm{R}_5\}$. The corresponding simplicial complex is depicted in fig.(3b) (since we are only interested in the structure of connections between simplices, we did not define any metric on them; thus, shape and dimensions of the simplices are unimportant): simplex A_1 is a triangle of vertices $\{\mathrm{R}_1, \mathrm{R}_2, \mathrm{R}_3\}$, simplex A_2 is a segment of extremes $\{\mathrm{R}_2, \mathrm{R}_4\}$, simplex A_3 is point $\{\mathrm{R}_3\}$, simplex A_4 is a triangle of vertices $\{\mathrm{R}_2, \mathrm{R}_3, \mathrm{R}_4\}$, simplex A_5 is a segment of extremes $\{\mathrm{R}_4, \mathrm{R}_5\}$. Note that a simplex may coincide with a part of another simplex, and that two or more simplices may overlap.

Let us represent the structure of the correspondences between mental categories by means of an incidence matrix $\mathbf{\Lambda}$ of dimensions $I \times J$, whose generic element λ_{ij} takes the values:

$$\lambda_{ij} = \begin{cases} 1 & \text{if } \mathrm{R_j} \in \mathcal{R}_\mathrm{i} \\ 0 & \text{otherwise} \end{cases}$$

Element (i, i') of matrix $\mathbf{\Lambda\Lambda}^\mathrm{T}$ is the number of vertices that simplices A_i and $\mathrm{A}_{i'}$ have in common. Thus, element $l_{i\,i'}$ of matrix $\mathbf{L} := \mathbf{\Lambda\Lambda}^\mathrm{T} - \mathbf{11}^\mathrm{T}$ is the dimension of the eventual common face between simplices A_i and $\mathrm{A}_{i'}$. If this number is negative, simplices A_i and $\mathrm{A}_{i'}$ have no common point.

Two simplices that have no point in common may nonetheless be connected by a chain of simplices having common points with one another. Let us say that simplices A_i and $\mathrm{A}_{i'}$ are *q-connected* if there exists a chain of simplices

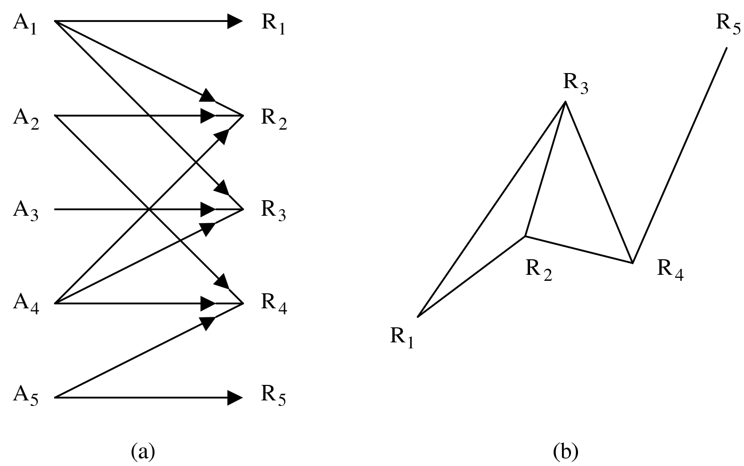

Figure 3: Correspondences between action categories and result categories (a), and the relative mental model (b)

$\{\mathrm{A_u}, \mathrm{A_v}, \ldots \mathrm{A_w}\}$ such that $q{:=} \min\{l_{i\,u}, l_{u\,v}, \ldots l_{w\,i'}\}$ is not less than zero. In particular, two contiguous simplices are connected at level q if they have a common face of dimension q.

Let us consider common faces between simplices, focusing on the face with largest dimension: let Q be the dimension of this face. Note that Q is not necessarily the largest possible dimension of a common face: given J result categories, the largest possible dimension of a common face is $J-1$, and it occurs when two simplices of dimension $J-1$ coincide.

By inspection of matrix $\mathbf{L}$ we can partition the set of simplices that compose the simplicial complex according to connection level q. Let us introduce a *structure vector* $\mathbf{s}$ of dimensions $(Q+1)\times 1$, and let us denote its q-th component by s_q. In general, for any connection level q there is some number of disjoint classes of simplices, such that the simplices belonging to a class are connected at that level: the q-th component of structure vector $\mathbf{s}$ is the number of disjoint classes of simplices that are connected at level q.

In order to avoid repetitions in the calculus of complexity, we do not consider a class of simplices connected at level q to be also a class of simplices connected at levels $q-1$, $q-2$, etc. For example, let simplices A_1 and A_2 be connected at level $q=2$, and let simplex A_3 be connected with A_2 at level $q=1$. Then, $\{\mathrm{A}_1, \mathrm{A}_2\}$ is a class of simplices connected at $q=2$ and $\{\mathrm{A}_1, \mathrm{A}_2, \mathrm{A}_3\}$ is a class of simplices connected at $q=1$; however, $\{\mathrm{A}_1, \mathrm{A}_2\}$ is *not* a class of simplices connected at level $q=0$.

An evaluation of the complexity of Atkin's simplicial complexes can be found in J. Casti (Casti 1989, p.410). The following formula has the same form, but yields different values from Casti's. In fact, Casti considers simplicial complexes connected at level q as also being connected at level $q-1$.

Let the complexity of system $\mathcal{S}$ viewed by observer $\mathcal{O}$ be:

$$C_{\mathcal{O}}\left(\mathcal{S}\right) = \begin{cases} 0 & \text{if} \quad \mathbf{\Lambda} \equiv \mathbf{I} \\ \sum_{q=0}^{Q} \frac{q+1}{s_q} & \text{otherwise} \end{cases}$$

where the sum extends only to the terms such that $s_q \neq 0$. If a mental model does not work in two or more simplicial complexes that are separated from one another by one-to-one connections, let us stipulate that the overall complexity is the sum of the complexities of the isolated simplicial complexes.

In order to visualize the meaning of the above formula, let us for the moment suppose that $s_q = 1$ for $\forall q$. The largest among the $q+1$ terms, namely $Q+1$, represents the nucleus of highest connection of the simplicial complex. The reader can figure it out as an irregular solid in a $Q-$ dimensional space: the more similar to a ball, the more simplices are involved. Lower-order $q+1$ terms represent lower-dimensional bumps that are attached to this solid, filaments that stretch in the surrounding space.

Subsequently, let us consider $s_q \neq 1$ for at least some q. The s_q terms represent the number of the above solids, filaments, etc.: for instance, $s_Q = 2$, $s_1 = 1$, $s_q = 0$ for $\forall q \neq \{1, Q\}$ means that a simplicial complex is composed by two $Q-$ dimensional solids connected by a chain of one-dimensional segments.

Terms $q+1$ measure how intertwined the causal relationships between action categories and result categories are, and each additional level of cross-connection corresponds to an additional $q+1$ term. Terms s_q measure the clustering of cross-connections over parts of the mental model.

Figure (4) is designed to help understand the interplay of these two factors. Fig.(4a) depicts a simplicial complex constituted by two two-dimensional simplices connected at level $q = 1$, where $C_O(S) = 2$. This value does not change if two similar simplices are added to them, as in fig.(4b). On the contrary, fig.(4c) shows that two additional one-dimensional simplices cause complexity to increase to $C_O(S) = 3$, since they introduce a connection level $q = 0$. A third one-dimensional simplex that divides the two-dimensional simplices in two groups, as depicted in fig.(4d), causes complexity to drop at $C_O(S) = 2$ again.

Clearly our complexity measure is definitely non-linear: additional simplices may leave its value unchanged, they may increase it or decrease it, depending on their position. Let me illustrate the meaning of these variations by means of a suggestive, qualitative example. Imagine you are shipwrecked on a desert island.

You may be concerned with two kinds of questions: 1) How can I make a living?, and 2) If there are natives, what a relationship shall I establish with them? The second kind of question is more complex than the first one, since it involves many intertwined causal relationships. Let us suppose that it is represented by simplices that are connected at a higher q than those concerned with mere surviving. The complexity you ascribe to your island is not likely to be affected by the number of natives who eventually inhabit it, or by the number of species you can hunt: adding simplices at the same connection level does not affect complexity, as passing from figure (4a) to (4b). However, it makes a big difference if you consider how to make a living only, or how to deal with natives only, or both. Thus, adding simplices at a different connection level does matter, as a comparison of figures (4b) and (4c) shows. Now, let us suppose that there are natives, who turn out to be very nice people, indeed, and that you wish to marry one of their women: although you view each available woman as very complex, the very fact that your choice is restricted to a finite, possibly small number of women simplifies your picture with respect to the moment you arrived on the island. Thus, complexity decreases due to concentration on a few clusters, as in figure (4d).

Complexity is zero if the simplicial complex reduces to single disconnected simplices, as the isolated result categories that appear in pseudo-deterministic mental models derived from stochastic mental models. Complexity is one if simplices are connected to one another by one vertex only, and it is maximal if there is exactly one class of simplices at each connection level. However, since the number of simplices is limited to $I \leq J$, these classes must contain one simplex only if we want to cover as many connection levels as possible.

The maximum connection level allowed by J result categories is $Q_{\max} = J - 1$. Thus, if two simplices have a common face of dimension $J - 1$, they must necessarily overlap. Consequently, maximum complexity is achieved when there are two simplices at $q = J - 1$, and one simplex at each lower q down to $J - I + 1$. Remembering that $1 + 2 + \ldots N = N\,(N + 1)\,/2$ we find that the maximum value that complexity can attain is $K = (I - 1)\,(J - I/2 + 1)$.

Unless we are measuring complexity along a series of time intervals where the observer changes the number of mental categories he is using, it is useful to normalize complexity to its maximum value:

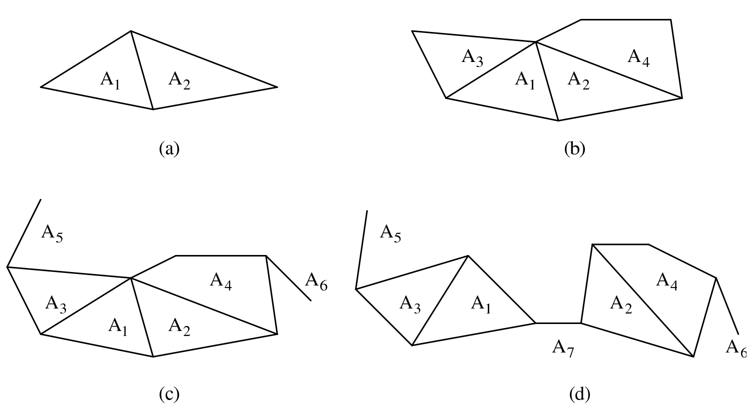

Figure 4: In case (a), there is only one class of simplices at $q = 1$, namely $\{A_1,\ A_2\}$; thus, $s_1 = 1$ and $C_O(S) = 2$. In case (b) there is also one class of simplices only, at $q = 1$; this class is now $\{A_1,\ A_2,\ A_3,\ A_4\}$, but since it is still $s_1 = 1$, complexity evaluation does not change: $C_O(S) = 2$. In case (c) there is one class of simplices at $q = 0$, namely $\{A_1,\ A_2,\ A_3,\ A_4,\ A_5,\ A_6\}$, and one at $q = 1$, namely $\{A_1,\ A_2,\ A_3,\ A_4\}$; thus, $s_0 = 1$ and $s_1 = 1$, which yields $C_O(S) = 3$. In case (d) there is one class of simplices at $q = 0$, namely $\{A_1,\ A_2,\ A_3,\ A_4,\ A_5,\ A_6,\ A_7\}$, and two at $q = 1$, namely $\{A_1,\ A_3\}$ and $\{A_2,\ A_4\}$; thus, $s_0 = 1$ and $s_1 = 2$, which yields $C_O(S) = 2$.

$$C'_{\mathcal{O}}(\mathcal{S}) = \begin{cases} 0 & \text{if } \mathbf{\Lambda} \equiv \mathbf{I} \\ \frac{1}{K} \sum_{q=0}^{Q} \frac{q+1}{s_q} & \text{otherwise} \end{cases}$$

where coefficient $1/K$ ensures that relative complexity $C'_{\mathcal{O}}(\mathcal{S})$ ranges in the $[0,1]$ interval. The following section illustrates the computation of complexity on a practical example.

5. An Example

Tom is a San Francisco based physician who volunteers in poor countries. He knows that in poor countries, due to malnutrition and insufficient hygiene, a number of children suffer from a serious disease called *pink eye*; this disease is absolutely absent in rich countries.

Let us represent Tom's mental model in the following way. Let action categories and result categories have the following meaning: A_1: Tom undertakes examination of 1000 children in a rich country; A_2: Tom undertakes examination of 1000 children in a poor country; R_1: Tom finds no evidence of *pink eye*; R_2: Tom finds that some children suffer from *pink eye*. Tom's mental model is depicted in fig.(5): it is a set of two one-to-one correspondences, one from A_1 to R_1, the other from A_2 to R_2.

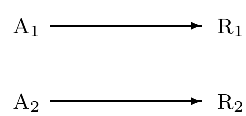

Figure 5: Tom's mental model

Let us now suppose that, surprisingly, Tom finds *pink eye* among children who live in a rich country. Tom can interpret this phenomenon in two possible ways: 1) Some people are hopelessly dirty, even if they can afford to buy soap; or 2) Something strange and new is taking place, e.g. a mutation of the *pink eye* virus that enables it to resist disinfectants.

In case (1) above, Tom's mental model is still confirmed. This mental model is represented by two disconnected simplices (two isolated points), and complexity is zero.

In case (2), the microbiological system has the capability to adapt to the disinfectants man has developed so as to produce new viruses that resist to them. Consequently, man develops a new means to kill the last microorganisms the system developed. This process can continue endlessly without man ever coming to the point of being able to compile a list of all possible microorganisms nature can create (the "elementary particles" of this example), setting this chicken-and-egg problem. However, Tom can construct a category of symptoms that he denotes "*pink eye*" in order to subsume the effects of various poorly known viruses. Furthermore, he can break the above self-referential loop by focusing on the particular viruses that exist at a certain point in time instead of tackling the more general problem of knowing the set of all possible viruses that can eventually cause *pink eye*.

In case (2), correspondences between mental categories are those of fig.(6). They are different from what Tom expected, given his mental model of fig.(5).

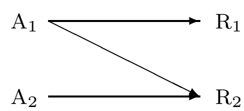

Figure 6: Correspondences between Tom's mental categories if he suspects that a mutation of the *pink eye* virus took place, that enables it to diffuse into rich countries.

Here we have two connected simplices that form a simplicial complex: simplex A_1 is a segment having R_1 and R_2 as extremes while simplex A_2 is just point R_2, so in the end the two simplices have point R_2 in common. A point is a face of dimension 0, and since this is the only common face between the two simplices it is $Q = 0$. Furthermore, since there are only two simplices connected at this level, the number of disjoint classes of simplices connected at $q = 0$ is one, so $s_0 = 1$. Consequently, it is $C_{\mathcal{S}}(\mathcal{S}) = 1$.

Since complexity is greater than zero, Tom prepares to investigate what happened. Imagine that he finds out that a dangerous mutation actually occurred, leading to a new disease that is independent of hygiene and that can be found equally often in poor as well as in rich countries; let us call it *red eye*. Furthermore, the matter is further complicated by the fact that *red eye* eases the occurrence of *pink eye*, so now *pink eye* can be found in rich countries as well. Tom must re-organize his information in order to distinguish this new possibility. Let us say that he introduces a new result category for the outcome of his medical examinations: R_3, for *red eye*. Now, his mental categories and the connections between them are as depicted in fig.(7).

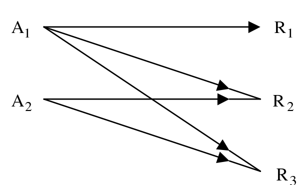

Figure 7: Correspondences between Tom's mental categories after he found out that *pink eye* mutated into *red eye*.

This means that simplex A_1 is made of a triangle having vertices R_1, R_2, R_3 and simplex A_2 is made of the segment having vertices R_2, R_3: so these two simplices have a common face of dimension one, which is the segment R_2-R_3

itself. There are no simplices connected at $q = 0$, so complexity is only calculated at $q = 1$ with $s_1 = 1$, yielding $C_{\mathcal{O}}(\mathcal{S}) = 2$.

Remarkably, if *red eye* had not eased the occurrence of *pink eye*, it could turn out that the *pink eye* cases Tom thought to have found in the rich country were actually *red eye* cases, and complexity would be 1 as before. In fact, the two simplices would be two segments with one point in common. Since complexity evaluation depends on the dimension of the common face between simplices and not on the dimension of the simplices themselves, its value remains unchanged.

Let us now suppose that *red eye* has not spread over the entire globe yet. It makes sense for Tom to distinguish the actions he undertakes for scientific investigations into three categories: A_1: as before; A_2: as before; A_3: examination of children in geographical areas where *red eye* has been found, no matter whether it is a rich or a poor country. Connections between categories are now as follows: A_1 is a simplex constituted by the single point R_1, A_2 is a simplex constituted by the single point R_2, A_3 is a triangle with vertices R_1, R_2, R_3. The three simplices are depicted in fig.(8): they form a chain of simplices connected by faces of dimension 0, and complexity drops to 1.

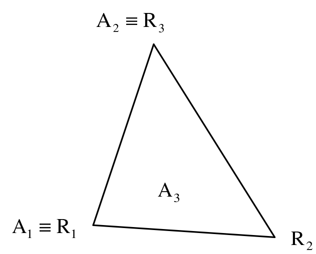

Figure 8: The simplicial complex if *red eye* is confined to a restricted area, no matter wether of a rich or a poor country.

However, if *red eye* spreads out of the geographical area it was initially restricted in, complexity increases again. If *red eye* spreads into rich countries only or into poor countries only, either A_1 or A_2 becomes another triangle of vertices R_1, R_2, R_3. Fig.(9) illustrates the two ensuing simplcial complexes. Both of them have two simplices connected by a face of dimension 0, plus two simplices connected by a face of dimension 2: consequently, complexity takes value 4.

On the contrary, if *red eye* spreads from the above restricted area to the entire globe, all three simplices are connected by a face of dimension 2. Fig.(10) shows a simplicial complex that is made of three overlapping two-dimensional simplices: complexity drops to 3. This might appear counterintuitive at first sight, but it does make sense: Tom understands that after the disease spread it makes little sense to have a complicated mental model that distinguishes three geographical areas, and that a simpler model would do better. It is the same phenomenon discussed in the introduction with respect to diminishing complexity when the observer of a system switches from a Newtonian to a statistical description: too high a complexity suggests to switch to a simpler picture.

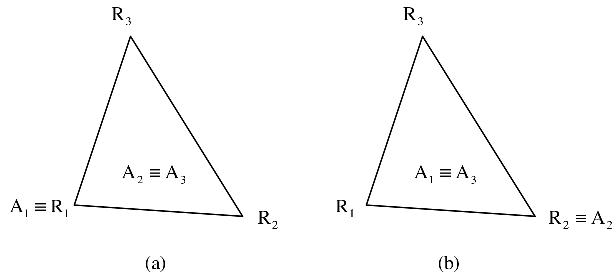

Figure 9: The simplicial complex if *red eye* only spreads into poor countries (a), or only into rich countries (b).

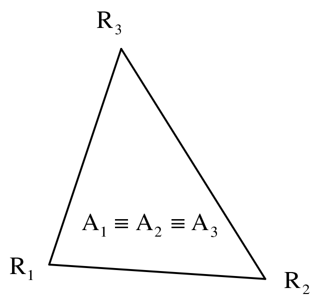

Figure 10: The simplicial complex if *red eye* spreads everywhere

Now, let us reflect on the examples presented. Is it really *complexity* that we are measuring, or is it *complication*?

It depends on Tom. With today's knowledge of microbiology, Tom is neither able to make an exhaustive list of all microorganisms that nature can potentially create, nor to disentangle the directions of all potential causal relationships existing between him and the microbiotic system. But perhaps, one day this will be possible. Perhaps one day man will know everything about all possible genotypes, as well as about their unfolding into phenotypes. There will be an immense, but finite, set containing all microorganisms that can be physically created, with a giant lookup table that will specify any possible causal relationship. Tom is most likely unable to have a reliable model of such an object. His model of the environment will surely fail from time to time. Nevertheless, he should say that the system is "complicated", not "complex".

Let us step again in the time machine: someone has been able to understand the giant lookup table in terms of a few laws, and ultra-speed computers are available: Tom can precisely predict which microorganisms will appear if γ rays increase by 0.0001 %. At that point, he will say that the system is "simple".

6. Subjective Complexity and Objective Measures

In the introduction I hinted at a possible subjective interpretation of objective complexity measures. Later on, I introduced a distinction between "complexity" and "complication", arguing that objective measures of complexity actually evaluate complication, because they can only be defined after elementary particles and causal relationships have been singled out. Subsequently, I introduced a subjective measure of complexity and I illustrated it with an example where, in the end, I suggested some circumstances that would lead the observer to say that the system is no longer "complex", but "complicated". This seems to imply that a subjective measure of complexity can also be used when the observer knows that, in reality, he is measuring complication: does this make sense?

It does, because although elementary particles and clear causal relationships allow to express complication objectively in the sense that this evaluation can be carried out by an algorithm, there must be a machine which "subjectively" computes complication. In this sense, objective measures of complexity do retain a subjectivist flavor. Notably, computational complexity is the length of the program a Turing machine needs to execute in order to produce a certain outcome (Solomonoff 1964; Kolmogorov 1965; Chaitin 1966), and statistical complexity is the amount of memory required by an ε-machine in order to predict the future development of the signal it is observing (Crutchfield and Young 1989; Crutchfield 1994).

By assuming to know a finite set of fixed mental categories, we could displace the complexity algorithm away from a machine that operates on elementary particles, to a brain whose functioning is supposed to be described by (but not actually to be !) an algorithm that operates on mental categories. In the very end, "subjectivity" merely means that the complexity algorithm works on given mental categories instead of on given elementary particles.

But in systems that are only "complicated" the mental categories can be expressed in terms of objective elementary particles. Consequently, subjective and objective complexity measures, in the very end, must have some point in common.

Objective measures of complexity can all be related to the concept of entropy. Computational complexity can be expressed in terms of the entropy of the message to be reproduced by the machine, statistical complexity in terms of the entropy of the past values taken by the signal to be predicted, mean mutual information is by definition the sum of the entropies of two parts of a system, minus the entropy of their combination. Finally, biological complexity, in the sense that a mutation occurring in an organism or organ is likely to affect many others (Conrad 1990, Kauffman 1993), can also be expressed in entropic terms (Conrad 1983). Thus, I shall focus on the relationship between measurement of subjective complexity on one hand, and entropy measurement on the other.

Thermodynamic entropy is the easiest to visualize. Suppose you have two separate cells of equal dimensions A and B, containing different gases a and b. In each cell, entropy $-\sum p_i \, log \; p_i$ is zero, because the probability to find one of the two gases is one. However, if we let the two cells communicate, gases will diffuse and entropy will increase since the p_is tend to equalize. The expression $-\sum p_i \, log \; p_i$ is not the only formula that is able to account for this phenomenon, Boltzmann chose it because he wanted entropy to be additive with respect to the different gases in the system.

Entropy measurements can also be used in biology, but with quite different results. In fact, in the biological realm entropy not only increases because the p_is tend to equalize, but it can also decrease because new p_is can be added to the system (Atlan 1972). Living organisms, by means of mutations, introduce new "elementary particles" into the system. In our example, living organisms are able to add a new cell C containing a different gas c, and the diffusion process starts again. However, it is important to remark that if a new cell C is added to the system, all terms of $-\sum p_i \, log \, p_i$ are affected. Consequently, entropy is not additive with respect to terms eventually introduced by new elementary particles.

Our subjective measure of complexity should behave a lot like its objective counterparts. In fact, mental categories are used to recognize given elementary particles a, b, c in given cells A, B, C, so that mental categories can be expressed in terms of these givens. In particular, let action categories be "I pick up a molecule from cell A", "I pick up a molecule from cell B" and "I pick up a molecule from cell C"; let us denote them by $\mathrm{A_A}$, $\mathrm{A_B}$, and $\mathrm{A_C}$, respectively. In a similar way, let result categories be "I picked up a molecule of gas a", "I picked up a molecule of gas b" and "I picked up a molecule of gas c"; let us denote them by $\mathrm{R_a}$, $\mathrm{R_b}$ and $\mathrm{R_c}$, respectively.

Let us suppose that the system is composed by two cells A and B, and that the observer is picking up gas molecules from cell A. If cells A and B are separated, the observer has no reason to expect to find molecules of gas b in cell A. Thus, his mental model is the one-to-one correspondence illustrated in fig.(11a), which bears zero complexity. But if gas b is allowed to diffuse into cell A, then correspondences between mental categories become those illustrated in fig.(11b): simplex $\mathrm{A_A}$ has now a common vertex with the virtual simplex of isolated point $\mathrm{R_b}$, and the ensuing complexity is $C_{\mathcal{O}}(\mathcal{S}) = 1$.

At this point, our observer is likely to switch from the deterministic mental model of fig.(11a) to the stochastic mental model illustrated in fig.(11c), which in its turn is equivalent to that of fig.(11d). If the probability of finding a molecule of gas a in cell A is greater or equal to 0.9, and if the probability of finding a molecule of gas b in cell A is less than 0.1, then this mental model is satisfied and complexity is zero. But if other molecules of gas b diffuse into A, the probability of finding a molecule of gas a in cell A decreases, while the probability of finding a molecule of gas b in cell A increases. Thus, action category $\mathrm{A_A}$ is connected to result categories $\mathrm{R}''_{\mathrm{a}}$ and $\mathrm{R}''_{\mathrm{b}}$. Simplex $\mathrm{A_A}$ now has two common vertices with the virtual simplices of isolated points R'_{a} and $\mathrm{R}''_{\mathrm{b}}$, and the ensuing complexity is $C_{\mathcal{O}}(\mathcal{S}) = 1$.

While gas b is diffusing into cell A the observer sees a sequence of unit increases of complexity, which is measured with respect to the particular mental model that ensures zero complexity at a given point in time. In other words, complexity is measuring the first difference of the probability of finding a molecule of b in A. This is a measure of the derivative of diffusion, not a measure of the derivative of entropy.

In the above example, the link between entropy and complexity is as follows. Let us denote $p_{\mathrm{A}}^{\mathrm{a}}$ as the probability of finding a molecule of gas a in cell A, and $p_{\mathrm{A}}^{\mathrm{b}}$ as the probability of finding a molecule of gas b in cell A. Clearly, $p_{\mathrm{A}}^{\mathrm{b}} = 1 - p_{\mathrm{A}}^{\mathrm{a}}$ and $dp_{\mathrm{A}}^{\mathrm{b}}/dt = -dp_{\mathrm{A}}^{\mathrm{a}}/dt$. Since we identified complexity as $C_{\mathcal{O}}(\mathcal{S}) = dp_{\mathrm{A}}^{\mathrm{b}}/dt$, from $S = -p_{\mathrm{A}}^{\mathrm{a}} \log p_{\mathrm{A}}^{\mathrm{a}} - p_{\mathrm{A}}^{\mathrm{b}} \log p_{\mathrm{A}}^{\mathrm{b}}$ we obtain:

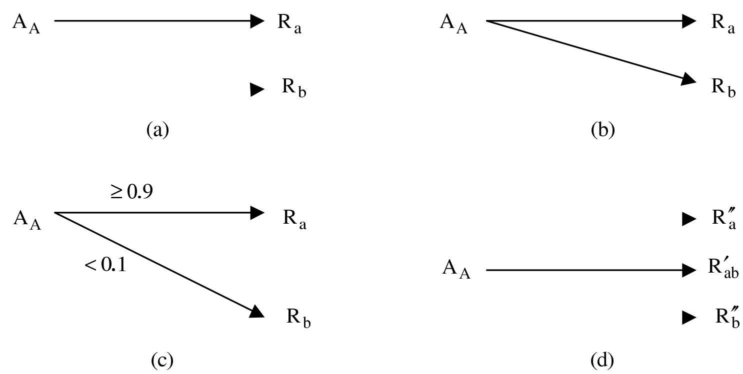

Figure 11: Mental model (a), and the connections that make it fail (b). Mental model (c), and its pseudo-deterministic version (d). Result category $\mathrm{R}'_{\mathrm{ab}}$ is defined as $\mathrm{R}'_{\mathrm{ab}} = \mathrm{R}'_{\mathrm{a}} \vee \mathrm{R}'_{\mathrm{b}}$, where $\mathrm{R}'_{\mathrm{a}} = \mathrm{R}_{\mathrm{a}} \wedge \{p(\mathrm{R}_{\mathrm{a}} \mid \mathrm{A}_{\mathrm{A}}) \geq 0.9\}$ and $\mathrm{R}'_{\mathrm{b}} = \mathrm{R}_{\mathrm{b}} \wedge \{p(\mathrm{R}_{\mathrm{b}} \mid \mathrm{A}_{\mathrm{A}}) < 0.1\}$. Result categories $\mathrm{R}''_{\mathrm{a}}$ and $\mathrm{R}''_{\mathrm{b}}$ are defined as $\mathrm{R}''_{\mathrm{a}} = \mathrm{R}_{\mathrm{a}} \wedge \{p(\mathrm{R}_{\mathrm{a}} \mid \mathrm{A}_{\mathrm{A}}) < 0.9\}$ and $\mathrm{R}'_{\mathrm{b}} = \mathrm{R}_{\mathrm{b}} \wedge \{p(\mathrm{R}_{\mathrm{b}} \mid \mathrm{A}_{\mathrm{A}}) \geq 0.1\}$, respectively.

$$C_{\mathcal{O}}(\mathcal{S}) \;=\; \frac{dS}{dt}\;\log \frac{p_{\mathrm{A}}^{\mathrm{b}}}{1 - p_{\mathrm{A}}^{\mathrm{b}}}$$

According to this expression, complexity decreases with increasing entropy, reaching $C_{\mathcal{O}}(\mathcal{S}) = 0$ when $p_{\mathrm{A}}^{\mathrm{b}} = p_{\mathrm{A}}^{\mathrm{a}} = 0.5$. Unlike statistical complexity, this subjective complexity measure does not increase with entropy when diffusion begins. The reason is that the above expression depends on a mental model based on statistical mechanics. If we would begin with a mental model based on microscopic Newtonian mechanics, we would see complexity increasing along with diffusion of gas b into cell A. Later on, an excessive level of complexity leads the observer to shift to a mental model based on statistical mechanics, and complexity would start decreasing. Note that also statistical complexity has a discontinuity at the peak (Crutchfield 1994).

In order to compare entropy and subjective complexity in biological systems, let us add a cell C to our system while gas b is still diffusing. Let us observe the diffusion of gas b and gas c into cell A. This has two consequences: 1) A new result category R_{c} must be added to the mental model, causing the mental model to fail each time an increase of $dp_{\mathrm{A}}^{\mathrm{c}}/dt$ is detected; 2) Since the speed of diffusion decreases with time, if c begins to diffuse after b there exists an initial phase where $dp_{\mathrm{A}}^{\mathrm{b}}/dt \leq 0$.

While entropy decreases in this initial phase due to the opposite effects above (Atlan 1972), subjective complexity does not. In fact, during this initial phase complexity increases by one unit at each step because a simplex having a single vertex R_{c} is added to the simplicial complex. Subsequently, complexity increases by one unit at each step for this reason *and* because gas concentration increases. However, these two effects do not sum up. This is because complexity is calculated upon classes of simplices connected at levels q, not on single simplices connected at q. Thus, complexity is identified as $C_{\mathcal{O}}(\mathcal{S}) = \max\left\{dp_{\mathrm{A}}^{\mathrm{b}}/dt, \mathrm{dp}_{\mathrm{A}}^{\mathrm{c}}/\mathrm{dt}\right\}$.

The above discussion highlighted that the link between entropy and objective measures of complexity on one hand, and our subjective measure of complexity on the other, cannot be expressed by means of a simple transformation law. In fact, this link depends very strongly on the mental model subjective complexity is calculated upon. This mental model can be such that subjective complexity reproduces the results of objective complexity measures, or not.

Furthermore, even if a mental model is such that subjective and objective measures of complexity keep step with one another, there is still a fundamental difference between subjective complexity and entropy-based measures. In fact, subjective complexity is based on recognizing novelties in the most recent data, while entropy is a statistical measure of the properties of a long series of data.

As shown in the above discussion, a link between these two approaches can only refer to the measurement of the *derivatives* of the probabilities that enter the entropy expression. Thus, relating subjective complexity and entropy-based complexity measures along a whole series of data requires the tremendous effort of specifying the mental model at each time step.

In order to measure subjective complexity over a time span that is long enough for an observer's mental model to change, it is necessary to know how this mental model evolves with time. It is unlikely that this can ever be accomplished by means of an algorithm that specifies any possible reaction of an observer's mind to any possible behavior of any possible system. Rather, the evolution of a mental model probably depends very much on its coupling with other mental models and it should be understood as a co-evolutionary process. This point is high on the agenda of research in artificial intelligence and cognitive sciences, and it is likely to lead us to embed subjective complexity measures in a more comprehensive treatment of the relationships between a system and its observer.

Acknowledgments

A large part of the insights contained in this paper are due to Thomas Lietman, who commented on an earlier draft. Although I bear full responsibility for the ideas exposed herein, I wish to acknowledge his important contribution to their development and do thank him for the intelligence of his remarks. Furthermore, I wish to thank Luis Rocha for pointing to connections with existing literature and and for providing a number of useful and encouraging suggestions. Finally, I thank Song June for pointing to a computational error in a previous draft of this paper.

References

Atkin, R. (1974) *Mathematical Structure in Human Affairs.* New York, Crane, Russak and Company.

Atkin, R. (1981) *Multidimensional Man.* Harmondsworth, Penguin Books.

Atlan, H. (1972) *L'Organisation Biologique et la Théorie de l'Information.* Paris, Herman.

Barsalou, L. (1987) The Instability of Graded Structure: Implications for the Nature of Concepts. *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, ed. by U. Neisser. Cambridge, Cambridge University Press.

Bennett, C.H. (1988) Dissipation, Information, Computational Complexity and the Definition of Organization. *Emerging Synthese in Science*, ed. by D. Pines. Redwood, Addison-Wesley.

Bennett, C.H. (1990) How to Define Complexity in Physics, and Why. *Complexity, Entropy, and the Physics of Information*, ed. by W.H. Zurek. Redwood, Addison-Wesley.

Cariani, P. (1998a) Epistemic Autonomy Through Adaptive Sensing. *Proceedings of the 1998* IEEE ISIC/CIRA/ISAS *Joint Conference,* Gathersburg, MD.

Cariani, P. (1998b) Towards an Evolutionary Semiotics: The Emergence of New Sign-Functions in Organisms and Devices. *Evolutionary Systems*, ed. by G. Van de Vijver *et al.* Dordrecht, Kluwer Academic Publishers.

Casti, J. (1989) *Alternate Realities.* New York, John Wiley & Sons.

Chaitin, G.J. (1966) On the Length of Programs for Computing Finite Binary Sequences. *Journal of the Association for Computing Machinery*, **13**: 547-569.

Clark, A. (1993) *Associative Engines: Connectionism, Concepts, and Representational Change.* Cambridge, The MIT Press.

Conrad, M. (1983) *Adaptability.* New York, Plenum Press.

Conrad, M. (1990) The Geometry of Evolution. *BioSystems*, **24**: 61-81.

Crutchfield, J.P. (1994) The Calculi of Emergence: Computation, Dynamics and Induction. *Physica D*, **75**: 11-54.

Crutchfield, J.P. - Young, K. (1989) Inferring Statistical Complexity. *Physical Review Letters*, **63**: 105-108.

Dupuy, J.P. (1982) La Simplicité de la Complexité. *Ordres et Désordres*, ed. by J.P. Dupuy. Paris, Editions du Seuil.

Foerster, H. von (1976) Objects: Tokens for (Eigen)Behaviors. *Cybernetics Forum,* **8**: 91-96. Reprinted in 1981, *Observing Systems*, ed. by H. von Foerster. Seaside, Intersystems Publications.

Hampton, J. (1993) Prototype Models of Concept Representation. *Categories and Concepts: Theoretical Views and Inductive Data Analysis*, ed. by I. Van Mechelen, J. Hampton, R. Michalski and P. Thenus. New York, Academic Press.

Hayek, F. von (1952) *The Sensory Order*. London, Routledge & Kegan Paul.

Hebb, D.O. (1949) *The Organization of Behavior*. New York, John Wiley and Sons.

Johnson-Laird, P.N. (1983) *Mental Models.* Cambridge, Cambridge University Press.

Kauffman, S. (1993) *The Origins of Order.* Oxford, Oxford University Press.

Kolmogorov, A.N. (1965) Three Approaches to the Quantitative Definition of Information. *Problems of Information Transmission,* **1**: 1-7. Reprinted in 1968: *International Journal of Computer Mathematics,* **2**: 157-168.

Lakoff, G. (1987) *Women, Fire, and Dangerous Things.* Chicago, The University of Chicago Press.

Pattee, H.H. (1989) The Measurement Problem in Artificial World Models. *BioSystems*, **23**: 281-290.

Pattee, H.H. (1993) The Limitations of Formal Models of Measurement, Control, and Cognition. *Applied Mathematics and Computation*, **56**: 111-130.

Rosen, R. (1985) *Anticipatory Systems.* Oxford, Pergamon Press.

Solomonoff, R.J. (1964) A Formal Theory of Inductive Inference. *Information and Control*, **7**: 1-22, 224-254.