# Level 9

# Arrow of Entropy and *Origin of Life*

As was observed before, in some cases there are credibility problems when trying to model complex systems applying simple methods. However the credibility problems are fixed, at some stage the models become not only incredible but truly impossible. When very improbable phenomena are assumed to cumulate *ad infinitum,* no time in the universe is enough to make emergence to happen.

There exist plenty of examples of such very improbable processes in biology: For example, the gene transcription from DNA to RNA consists of a huge number of marvelous coordinated-looking steps that are needed, and so does the translation process of RNA further to proteins. How do the locally controlled atoms know when to adhere and when to let loose when the sequential reading of the codes is being carried out? And these are only subprocesses — above them, there are the developmental processes in an individual and evolutionary processes in a population that are equally astonishing. How do the systems climb the endless steps of increasing complexity?

Admitting that there are still challenges is the first step towards more plausible models. Unprejudiced analyses make it possible to see things in a perspective — and, suddenly, it turns out that all is clear. When correct interpretations are applied, it turns out that actually the systems are not going in the direction of increasing improbability; they go down towards *maximum* probability. The systems struggling against the flow of entropy is just an illusion (see Fig. 9.1).

In this chapter, this viewpoint is applied to the analysis of how life could have emerged from the non-living. Indeed, it does not matter how long the ladders are; when you are going in the right direction it does not matter how long it takes. There is enough time — as long as the processes go in the right direction.

## 9.1 Thermodynamic view of cybernetics

The most universal framework that governs all physical systems is *thermodynamics.* The thermodynamic concept of *entropy* is among the most fundamental

Figure 9.1: The paradox of entropy flow is just an illusion (graphics by Maurits Escher)

ones in nature, and when searching for universal laws governing cybernetic systems, among others, these issues need to be addressed.

### 9.1.1 Entropy and order

Applying the thermodynamic interpretation (as defined by Rudolf Clausius), entropy reveals the extent to which the energy in a closed system is available to do work (as defined in a somewhat sloppy manner). The lower the entropy level is, the more there is *free energy.* In a closed system, entropy level cannot decrease; it remains constant only if all processes within the system are reversible. However, because the natural processes typically are irreversible, entropy in the system increases, so that energy becomes "inert". Even though the total amount of energy remains constant, according to the *first law of thermodynamics,* it becomes less useful, according to the *second law of thermodynamics.* Ultimately, the system ends in a thermodynamic balance, or "heat death", where there is no more free energy available.

This direction of increasing entropy seems to be opposite to what takes place in cybernetic systems. The accumulation of complexity in the evolving structures seems to fight against the second law of thermodynamics. The easy answer here, of course, is that cybernetic systems are *open systems,* where there is energy transfer between the system and its environment — the total entropy level in the whole universe increases despite some "countercurrents" in the flow. However, there are some rough edges in this explanation: This assumption means that the strongest of theories, thermodynamics that should govern everything, is not applicable in cybernetic systems, becoming void and useless *if the system just decides to develop.* There seems to exist a gap between "normal" and "abnormal", evolving systems (see Fig. 9.2). Why do not the cybernetic systems not choose the "easy way", following the flow of entropy?

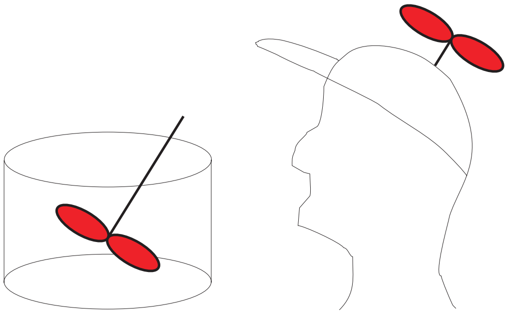

Figure 9.2: Two vessels — an ideal mixer and an "idea mixer". Two systems where there seemingly is something very different from the thermodynamic point of view: In the former case, in the perfectly stirred tank, addition of energy decreases order and structure, whereas in the latter case, in the cognitive system, activity increases order, new structures being constructed

There also exist different, more or less closely related definitions to entropy. In statistical mechanics (by Ludwig Boltzmann and Willard Gibbs), and analogously in information theory (by Claude Shannon), entropy is related to probability: More probable states (observations) reflect higher entropy than less probable ones. In a sense, entropy is the opposite of information — less probable observations contain more information about the system state. In such discussions, the second law of thermodynamics, or the increase in entropy, is reflected so that systems tend to become less ordered, and information becomes wasted.

It seems that intuitions concerning entropy are to some extent contradictory, or at least obscure. One hypothesis assumes that entropy, being among the only one-directional quantities in physics — defines the direction of time. Perhaps the most marvelous conclusion is that the universe cannot shrink because that would make particles be closer to each other — thus the system being more ordered, total entropy in the universe going down. This would also mean that time would start going backwards! Perhaps there is room for yet other interpretations.

The probability-bound interpretation of entropy is appealing, but it also seems to result in paradoxes: For example, a symmetrical pattern is intuitively more ordered, containing more information, and consequently having lower entropy than a completely random pattern — on the other hand, symmetric pattern can be seen to contain *less information* than a random pattern, because the redundancies caused by the symmetricity can be utilized to represent the patterns more efficiently, so that the entropy level should be *higher* now. Indeed, it can be claimed that the "algorithmic entropy" is higher in a symmetric pattern than in a non-symmetric one. To confuse concepts concerning order and symmetry even more, or, rather, to reveal the inconsistencies in our intuitions, think of

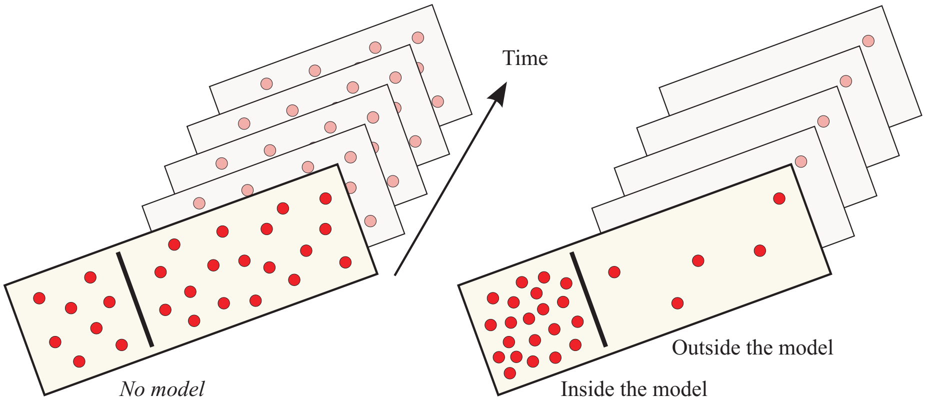

Figure 9.3: Illustrating the effects of cybernetic control (dots denote "information units" and their conglomeration): The case *without* model (on the left) has high probability at start time, and thus high *momentary* entropy; the case *with* model (on the right) has low total information over time, and thus high *sustainable flow* of entropy. Note that the information stored in the model is constant and thus negligible, being defined once for all

the following claim: A totally unordered system can be said to be extremely symmetric as the components cannot be distinguished from each other.

The intuition has been seen as the basic tool in neocybernetics to reach good models — such thinking seems to collapse here; but can the power of intuitions still be preserved? The answer is yes — here it is assumed, according to the original intuition, that orderliness, or loss of disorder, is a manifestation of low entropy. The key point here is that the simplicity of symmetric patterns, or ordered patterns in general (loss of information in them), is just an illusion: The missing information of the pattern is buried in our mental pattern recognition capability. If the same data is to be presented without the supporting underlying cognitive machinery, or specialized interpretation and analysis tools, there is no handicap — the redundancy cannot be exploited, and no compression of data can be reached. In general, a higher-level representation makes it possible to abstract the domain area data; in other words, as has been observed, a model is the key to a compressed representation. And this idea can be extended to cybernetic systems in general: It need not be our personal cognition machinery that constructs the model storing the excess information; any cybernetic system can do that in its own more or less narrow environment.

In the cybernetic perspective, the two views of entropy can be combined in a natural way: On the one hand, it is about balances and pursuit towards heat death, and, on the other, it is pursuit towards least information, as measured in terms of variation. What is more, it turns out that *evolution of structures increases entropy.*

### 9.1.2 Control changes it all

The cybernetic systems, as studied before, are characterized by balances: First, the determination of $\bar{x}$ is based on finding the dynamic equilibrium as determined by the system model. Second, the system structure as determined by the matrix $\mathrm{E}\{\bar{x}\bar{u}^T\}$, is also a dynamic equilibrium as determined by the statistical properties of the environment. Indeed, in a cybernetic system there are balances at each level — and, in this sense, the convergence towards a steady-state model is completely in line with the second law of thermodynamics: Variation (information) around the balance is maximally being eliminated, the "local heat death" almost being reached.

Where does this balancing property come from? It is the structure and order on the higher level — or the model — that makes it possible to control the lower level, or to reach the information elimination there. Evolution is the process of introducing ever more complicated structures that facilitate ever better control of the environment, either implicitly, as in lower-level biological systems, or explicitly, as in man-made systems. In any case, the cybernetic controls boost entropy — and the more sophisticated the control is, the higher is the rate of entropy production. In this way, rather than opposing entropy, the cybernetic system tries to maximize entropy — quite in accordance with normal physical systems. It is all about correct viewpoint, and selecting the system boundaries appropriately. It is the control system intuition that is needed to solve the "arrow of entropy" paradox.

Because of the simple definition of information (information being manifested as variation), it is possible to distinguish between information being captured in the structure (the model) and information being left in the signals (unmodeled noise). The cybernetic system acts like a Maxwell Daemon, distinguishing between two "containers" of information and noise, compressing information and pumping "negative entropy" into the emerging structures, thus causing positive entropy be left outside the structures (see Fig. 9.3). The key point here is that the single container of negative entropy (the model) is outweighed by the large number of samples with increased entropy level (data variation in the environment being suppressed), thus being a thermodynamically sustainable scheme. Whereas the momentary entropy increases, the "emergent entropy", the average of entropies over the whole environment and over all future decreases when the cybernetic strategy is employed. The same thinking applies to entropy as to other quantities when the neocybernetic perspective is applied: The time axis is abstracted away, only the average over the long-term loss of information is considered.

Perhaps the most important consequence of the new interpretation of the cybernetic systems is that reductionistic approaches become possible: Traditionally, the only systemically consistent level of studying such complex systems, with the whole environment being involved, was the holistic level, the whole universe being seen as one entity. Now each subsystem can be studied independently, as an independent thermodynamically consistent entity. The traditional view of seeing the relationship between the system and the environment is turned upside-down, or actually inside-out, environment now being the innermost part, being controlled by the system (see Fig. 9.4). It is the environment that is seen as the object, and the system is the subject, manipulating the environment, and a

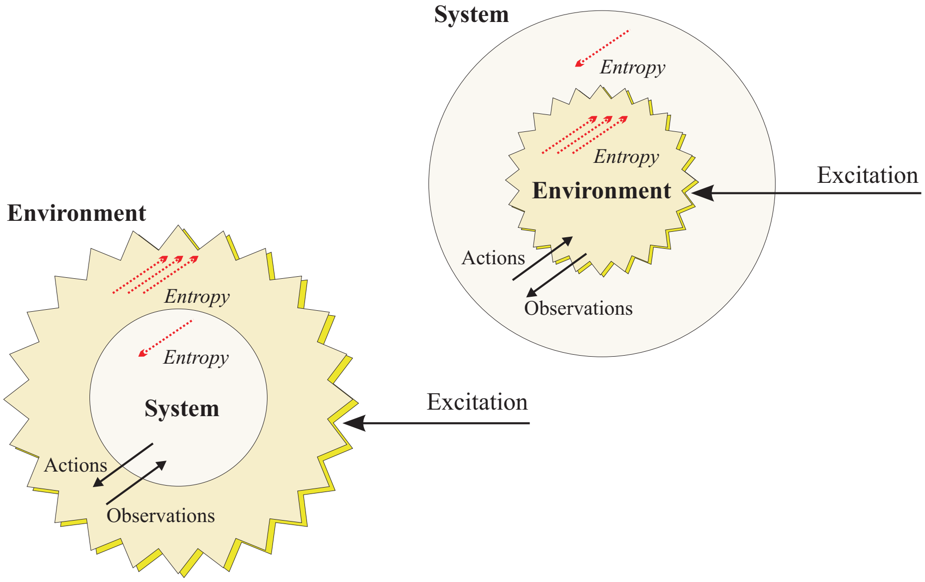

Figure 9.4: Reaching entropic consistency within systems. Left — traditional view, right — cybernetic view

more complicated system always sees the lower levels as through a looking-class. The highest-level model where the negative entropy is concentrated remains outside the boundaries. The original input into the environment is white noise; as seen by the highest-level system, the lower-level systems distort this noise, and the systems tries to capture this distortion, or the redundancy there is in the observations. It does not matter how many levels of systems there are, the same principle of modeling always applies. Note that the low-level systems only see a narrow view of the complex environment, and it is only this limited information that is relevant to that system — the whole complexity of the world needs not be captured.

### 9.1.3 Another view at model hierarchies

When the view of "information units" is employed, it is perhaps motivated to take a closer look at models themselves. The model is a container of information that characterizes the patterns that distinguish the system in question. To have some perspective, note that the cybernetists Norbert Wiener and Arturo Rosenblatt have argued that

> The best material model of a *cat* is another, or preferably the same, cat.

However, this view only applies to a trivial structure of models, when just a isolated single cat is being analyzed. In the beginning, modeling truly starts from representing all available information, the "world model" consisting of the data directly, but when abstracting over individuals, the model becomes more compact, the set of common patterns becoming smaller. When consistent vari-

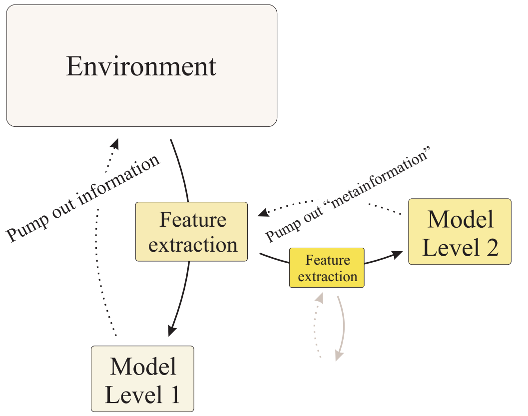

Figure 9.5: The challenges in "universal neocybernetics" are compressed into the question: How are the domain-specific features (pieces of information) extracted from environment

ations from the nominal are detected, the best model for the whole category of cats summarizes the similarities or invariances characterizing all of the "cat" samples. Neocybernetic models are the collections of *invariances over variances.* This model compression, or separation of information between the containers is implemented automatically by the cybernetic adaptation and control mechanisms. At this level, the model contains the detected *similarities.*

But this kind of simplification is not all that happens. As has been observed, neocybernetic systems are not alone — and, similarly, also neocybernetic models form a hierarchic interlinked structure. If there is a hierarchy of models, further compression of the sub-level model takes place: Redundancy in information gives rise to a higher-level "model of models", or, as seen in the control perspective, "control of controllers". A generalized view of control can be based on the view of *eliminating information invariants,* or transferring them onto the higher level, being shared by different domains. As there are common patterns among features, the higher-level model captures this redundancy. This means that the same information is represented in the lower-level model *only once.* At the higher level, when differentiating submodels from each other, the model contains the detected *differences.* The hierarchy of cybernetic models optimizes among the representations of similarities and differences, assumedly minimizing the size of the overall model.

Seeing information as bits — in the spirit of information theory — makes different levels of controllers commeasurable. No matter how a feature is defined, directly in terms of a measurement or through a complicated algorithm, there is no qualitative leap in their algorithmic complexities; information can be collapsed onto the same format, and structures within the models also become a matter of analysis and control. At the lower level, there is identity among information sources (features) that deliver the same information (distinguishing between categories in an identical way), whereas when seen from above, the information (algorithmic complexity) in the model is minimized, so that the simplest representation remains. The cybernetic adaptation and compression of information always follows the same principles (as studied in chapter 3) — the key issue when escaping in a phenosphere to a higher level is that of determining the features (see Fig. 9.5).

This analysis that is based on the formalized view of information applies also to memetic systems — for example, when looking at *science,* one can even find new perspectives into the groundings of model thinking. Simplification of a model is a manifestation of existence of a higher-level control, and when studying science, these models are *theories,* also becoming more compact as being "controlled" by the higher-level controls; these controls are defined by the paradigm of doing sciences in general, governing the principles of all scientific work. This means that as any science is a subsystem in a controlled hierarchy of cybernetic sciences, it is bound to become more and more simplified as the hierarchy matures: Only the most powerful explanations survive. Indeed, such simplicity pursuit (compare to "Ockham's razor") is traditionally taken as the philosophical foundation of all modeling without any attempt to justify it. As the scientific discipline is a cybernetic system as is the subject of its study, this simplification is perhaps not only an engineering-like shortcut: The same kind of simplification takes place also in nature.

The qualitatively separate levels in the models are also visible in practical real-life systems — study the flows and information hierarchies in an industrial process plant:

1. **Physical flows** are the real flows of matter and energy in the process.
2. **Information flows** typically consist of the feedback controls governing the physical flows.
3. **"Knowhowflows"** consist of supervision and optimization of the underlying control structures.

The goal of traditional control is balancing of the time-domain dynamics by exploiting the causalities; this process-specific layer supplies for the cybernetic medium to be exploited by the domain-independent cybernetic structures. Minimization of variances in product quality, and robustness against environmental disturbances, is implemented finally applying the concrete controllers. The ideas of feedback control are the same in all kinds of physical systems. In the similar way, the ideas of cybernetics are still more general, covering all kinds of control systems, abstract or concrete. One has to proceed from "bulk information" to *metainformation* or knowledge (information on information). The level of "cybernetic controls" on the metainformation level is somewhat ill-defined — indeed, as soon as all information flows in a cybernetic loop are unambiguously fixed, it becomes a traditional control loop. Even though (as observed above) one always operates on the same kind of information units, it is reasonable to distinguish between levels: The cybernetic framework combines systems from different phenospheres. The plant-level idea of an industrial system is functional, combining subsystems, both memetic and physical in appropriate, ingenious ways. Information is transferred between phenospheres, or "parallel universes of information"; one could distinguish between *eksoinformation,* or "inter-domain information", and *endoinformation,* or "intra-domain information". The lower-level controls exploit the available endoinformation that can be included in the formal loops, whereas the eksoinformation can be called "knowhowflow". In knowhowflow one can exploit expert understanding and common sense reasoning, and humans are integrated in such closed loops. This flow is typically very stochastic, and there are only the tensions visible; it is impossible to formulate the actual processes explicitly, but, again, the final state is well-defined. The key point in expertise exploitation is selection and preprocessing of the appropriate variables and weighting of them. The drifts are manifested in the engineering-like pursuit towards better solutions — cheaper, faster and more accurate measurements, actuators, and algorithms. Whenever the features are formulated, or the "probes" are defined (see chapter 4), model adaptation takes place in the familiar way. There is a balance among technological possibilities and economical constraints.

### 9.1.4 Principle of maximum entropy production

Traditionally, the second law of thermodynamics is thought of as being a universal, more or less metaphorical principle. The existence of systems with inverted, entropy-decaying nature has made it difficult to motivate explicit utilization of this principle in practice: It seems that the entropy principle cannot be applied in a reductionistic way for analysis of concrete large-scale systems.

Now, according to the above discussions, the entropy in a subsystem always increases when seen from the higher-level system. In a cybernetic system, entropy increases in a consistent manner, there is "balance pursuit" at all levels, completely in line with the second law where thermodynamical balance is the ultimate goal. Because of this consistency, any subsystem at any level — when its boundaries are appropriately determined — can be studied separately, and also holistic systems can be analyzed in a reductionistic manner. In this sense there is no more difference between different kinds of complex systems: Living systems and non-living ones, for example, can be modeled in the same framework. Whereas the first law of thermodynamics (energy principle) offers powerful tools for deriving static models, it seems that the second law (entropy principle), being a fundamentally flux-based concept, offers generic tools for deriving *dynamic models* — also for complex adaptation processes. As long as there are in a thermodynamic system heat resources, or heat differences, there is capacity to do work; similrly, as long as there is emergy in a cybernetic system, there is capacity to adapt and "live". There is directed (generalized) diffusion, or "leakage" of emergy from the environment, evolution making this leakage from the reservoirs become faster.

If entropy production can be seen as a consistent process, the next step is to assume that it happens *as fast as possible.* It can be assumed that it is the most efficient strategies that only remain visible, characterizing the whole system when seen from outside. In the spirit of principles of *least action* or *minimum energy,* as originally proposed by Maupertuis, and later extended by Euler, Lagrange, and Hamilton, one can propose the *principle of maximum entropy production* for characterizing the processes of information decay. Such somewhat teleological modeling principle give strong tools for looking systems and their adaptations in a perspective.

When information is seen as a concrete measurable quantity, formally incompatible systems can be put in the same framework, and the intuitive visions concerning behaviors in cybernetic systems can be functionalized. The entropy levels, or, rather, changes in the levels, determine the *free energy,* and they can be applied as a measure for tensions in a cybernetic system; this measure can be expressed as bits of information. Let us study what this means in practice when doing entropy pursuit — what is the maximum speed of information container separation, or what is the rate of the "emergent dynamics" in cybernetic adaptation processes?

Assume there are $k$ samples of data. Information is extracted from this data in terms of more or less computational features, defining different ways of looking at the system. In the spirit of information theory, the features are reduced to one bit: They contain elementary characterizations of the form "yes" or "no". Further, for simplicity assume that the probability that a random independent feature gives the correct classification is $p$; then the probability that the feature remains indistinguishable from the others all the time during sampling is

$$p^k. \tag{9.1}$$

If the acquired information is optimally exploited, the probability that a superfluous feature does not become ripped off decays exponentially. In practice, the optimum speed of adaptation becomes dependent of the signal-to-noise ratio in data — but at last in this (extremely) simplified case, the rate of elimination of bits in both environment and in the model takes place exponentially. Indeed, this is what one would expect.

The above discussions are not only a theoretical exercise: They offer powerful conceptual tools to attack complex evolutionary systems. The new view turns the direction of thermodynamic tensions, changing the destructive-looking tendencies into constructive ones, making the originally improbable developments probable after all. The power of the new intuition is illustrated here by discussing one of the biggest mysteries there is — the *origin of life.* The question about the origin of life is not only a philosophical problem: When trying to extend biology from the analysis of distinct examples, individual animals or species, to the scientific analysis of general principles, it is necessary to understand the common principles all of them share: Emergence of a cybernetic system is a birth of an individual. Fundamentally, it is this origin of life that is faced by each living system — after all, the individuals repeat this process in their development starting from non-living chemicals.

## 9.2 Ladders towards life

The processes of DNA transcription and translation into proteins are much too amazing to be credible — and still it happens all the time. It seems so cleverly orchestrated that it is easy to assume that a guiding hand is necessary in these processes. Thinking of this all having emerged all by itself, by some "blind watchmaker" — this is the basic dilemma where creationists have all the memetic weapons at their hands. However, in the cybernetic setting the things turn upside down, and it can be claimed that *the most plausible explanation for origin of life is non-divine.*

### 9.2.1 Paradoxes of living systems

There exist various seemingly reasonable presentations about the origin of life in the literature — for example, see [13] and [23]. There seem to exist no real problems whatsoever, and even questioning the trivializations is seen as next to insane [21]. However, it seems that it is these explanations that are missing the common sense.

When the origin of life is discussed, it is often claimed that the key problem is to explain where the first DNA molecule came from — after that, reproduction etc. should be no more problem. And when a reproduction machinery is available, and mutations in the code take place, it is the Darwinian principles that only need to be followed to explain today's diversity in nature. Even though such emergence of DNA is highly improbable, there was billions of years time. Unfortunately, as revealed by the famous Miller-Urey experiments, just adding mindless energy in the potion of chemicals, only simple amino acids are produced — more complex molecules just are not energetically stable enough. But it is still *possible,* is it not?

However, the question of whether a single molecule ever came on stage is still rather irrelevant. To understand the true nature of the problem, study the following scenarios:

- *Assume* that the DNA sequence once were produced in the primordial sea. However, there are no mechanisms to read to code, and in a matter of days in the harsh conditions of the early earth the molecule breaks apart. Alone the single molecule is completely void: There is information for the structure but the structures never become actually constructed.

- *Assume* that there is a complete body of an animal being washed on the shore of the primordial sea. However, if it is dead it is dead and that is it, it never comes alive again. Now there are all the structures ready, but the functions have forever ceased; the dynamic attractors cannot be instantiated by any amount of energy.

- Finally, *assume* that some living Robinson lands on the shore of that primordial sea. However, there are no *other* live forms already there, and there is no food to eat. Even though there are all the necessary structures and processes appropriately running in the body, there will be death in a matter of weeks.

Indeed, it is evident that there cannot exist life alone, isolated from other life, or from its natural inhabitat. Life is manifested in interaction with the environment — or, when putting this in more pointed way, life (or ability to host life) is the *property of the environment.* The problem of life is not about explaining a single molecule; the whole ecosystem should be explained simultaneously. It is evident that this observation makes the problem even more difficult. True, the one fixed molecule would be simpler to explain than a dynamic system with active interactions — but this is what life is. The role of the single molecule in explaining life is like the role of a single logic formula when explaining intelligence.

When seen in the correct perspective, there are no paradoxes here — it is only the traditional ways of thinking that are paradoxical. Concerning the evolutionary processes there also exist strange intellectual dead ends: For example, it is often assumed that the evolutionary fitness determined by the ability to produce offspring most efficiently, or by the ability to adapt to new environments most efficiently. However, the simpler an organism is, the faster it is to reproduce and to modify itself — leading to degeneracy of structures, and world power of bacteria! What is more, the assumed power of the mechanisms of natural selection also seem to be a myth: Those who have done random search in a high-dimensional space, know how notoriously inefficient it can be.

Alongside with artificial intelligence, there are efforts to construct *artificial life* (for example, see [28]). However, these efforts are plagued by the same problems as the behavior-based AI research: It seems that the emphasis is on superficial patterns. The criterion of relevance is based on how interesting a simulation looks. And it is all computer programs and algorithmic procedures — what you program there is the only what you get out; only what you can think of, you can implement. But it seems that life is an emergent phenomenon so that its essence cannot be captured in definitions. Trivialization of the complex questions only results in what one could call (following the AI terminology) "shallow life".

The new view of cybernetic systems as pursuing balance is fundamentally different from traditional intuitions. It has been assumed that "interesting" complex systems are "at the edge of chaos". When studying the processes of life, the mainstream view is expressed by Ilya Prigogine: Life is "as far as possible" from balance, whereas death means final balance. Erwin Schrödinger phrased this as "What an organism feeds upon is negative entropy; it continues to suck orderliness from its environment". Also in cybernetic systems, static balance means death — but a living system is characterized by *(thermo)dynamic,* non-static balances. The ways of thinking have to be inverted: Whereas a living thing is traditionally assumed to play an active role, now it just has to adapt to its environment; it is the environment that pumps disorder into the system, and life processes try to restore balance, or homeostasis. The system controls the environment, yes, but it is the environment that dictates how it is to be controlled.

It seems that the neocybernetic starting point is useful, capturing the correct intuitions here: The goal of a system is the ability to find the best balance with the environment. To implement this, there are not only structural adaptations available, but continuous matching processes take place to fine-tune the structures. And in evolution, becoming structurally more complex is the method to reach better match with the complex environment. These issues are studied closer in what follows.

### 9.2.2 Balanced autocatalysis

The definitions of what life is are very intuitive, and no matter which set of characterizations is selected, there always exist counterexamples. In the neocybernetic perspective, the following definition is employed here: *Life is higher-order balance with the environment.* What kind of assumptions are needed to make this definition plausible?

Looking from the point of view of the end result, seeing the living organism as being at the mercy of the environment makes it seem very volatile. However, analyses must be started in the bottom-up direction: Starting from the simplest of environments and proceeding towards more sophisticated ones, always making the outer system supervise and control the inner one towards local heat death, puts the system into an active role. Applying the vision of inversion of the arrow of entropy, the problems seem to become solved one by one, balances being restored in each phase separately. Getting to the higher levels, bigger picture is seen, the entity becoming better and better controlled, keeping the emergence of sophistication in the developing system thermodynamically consistent. The assumed balance with

So, start from the bottom and select a narrow view of the environment, let the system adapt there, and only after that widen the view. The transition from intuitively non-living structures to living ones becomes smoother; on the bottom of hierarchies one has *chemical evolution.*

One of the central prerequisites for life is the capability of reproduction. The simplest example of chemical reproduction is demonstrated in *autocatalysis,* where a chemical catalyses a reaction where this same chemical is produced; the autocatalyst thus can make copies of itself. Assume that for some chemicals A, B, and C there holds

$$\begin{array}{lcl} \mathrm{A} & \Leftrightarrow & 2\mathrm{A} \\ \mathrm{A}+\mathrm{B} & \Leftrightarrow & \mathrm{C}. \end{array} \tag{9.2}$$

The autocatalyst A acts like switch, activating the reaction producing chemical C from B. If there is chemical A present in the system to begin with, it will forever continue to be there no matter how much the solution is diluted. The chemical A thus characterizes the functioning of the system, selecting functions by making the corresponding reactions possible that otherwise would never take place. In practice, the autocatalytic reaction chains must be more complicated than the above one; it has been shown that the probability that a random set of chemicals is autocatalytic becomes high under certain assumptions.

Autocatalysis makes it possible to explain inheritance of functions between chemical systems. Indeed, autocatalytic sets are seen as the explanation for origin of life by Stuart Kauffman and others [44]. However, there are theoretical problems: Looking the chemical reactions syntactically, as a cookbook, there seems to be an explosion of chemicals. There is no self-regulation in the system, and there seems to be no emergence of structure. Indeed, it seems that autocatalytic systems typically only produce sticky tars, ending in deadlocks. In the neocybernetic framework this problem is solved: The system consists of balance reactions that proceed only in favorable conditions. It is the environment — or the reaction set itself — that takes care of self-regulation, the autocatalysts determining the spectrum of possible degrees of freedom in the chemical system. And the *function* is more relevant than the *structure.* In a system of autocatalysts function is manifested without solid form; structure is of secondary importance. But as illustrated in the next section, the physical properties of the world can make structures automatically emerge without explicit maintenance.

### 9.2.3 Chemical evolution

To understand the life processes in their simplest form, it is here assumed that chemicals participate in equilibrium reactions, as presented in chapter 1. To facilitate the emergence of something more interesting, three basic hypotheses concerning the reactions are made:

1. There are autocatalytic chemicals present among chemicals.
2. There is a medium available where interactions can take place.
3. There are mechanisms available for keeping chemicals together.

In the simplest case, this means that there is liquid water for chemical solutions to react in. Using the traditional vocabulary, one can speak of *primordial soup,* where there are chemicals and energy available (for example, see [21] and [88]). To keep chemicals together, it can be assumed that there are some kind of partially isolated "droplets" in the medium (see Fig. 9.6). The physical environment makes the droplets behave like "proto-cells". The growth of such droplets can be explained in terms of good match between the environment and the autocatalytic set characterizing the contents of the droplet: The reactions are active, keeping up the "metabolics". Chemical properties determine the internal balance in the droplet, but it is physical phenomena like osmosis and surface tension that together determine the size of the droplet, and whether it splits up. Because of the geometric constraints, the chemical reactions in the droplet are also affected: As there is less surface, there is smaller total intake of chemicals; and if some of the droplets is surrounded by other droplets, it experiences a very different environment, thus perhaps exhibiting different reactions, and different chemical functions.

The droplet has to maintain its integrity, so that it does not dissolve in the surrounding water. This can be assured if the contents of the droplet are, for example, based on fatty acids or some gels. It is the chemical reactions within the cell that have to provide also for the supply of this substrate. In more sophisticated cases the proto-cell can have some membranes that are based on phospholipids or other compounds with *amphipathic* character, having hydrophilic and hydrophobic parts. The more complex scenarios can employ the ideas of *vesicles, globules,* or *micelles* to host the reactions, having restricted exchange of chemicals and energy with the environment. Similar scenarios have been proposed a lot in the literature — but even though structures that resemble "cells" can emerge in a rather autonomous manner, it is clearly not such physical structures only that characterize living systems.

The key difference here as compared to the standard autocatalysis models is that the proto-cells are not whatever droplets, but they can host complicated sets of balance reactions. Such an equilibrium system is a *local mill of entropy.*

The entropy considerations in the beginning of this chapter were rather abstract, and can be applied only as evolutionary processes are seen from outside. But what are the mechanisms — how is the increase in entropy manifested in the lowest level of a developing proto-cell? A well-balanced proto-cell offers a good platform for biologically relevant functionalities to emerge. The neocybernetic structure has evolutionary advantage as seen in the thermodynamic perspective.

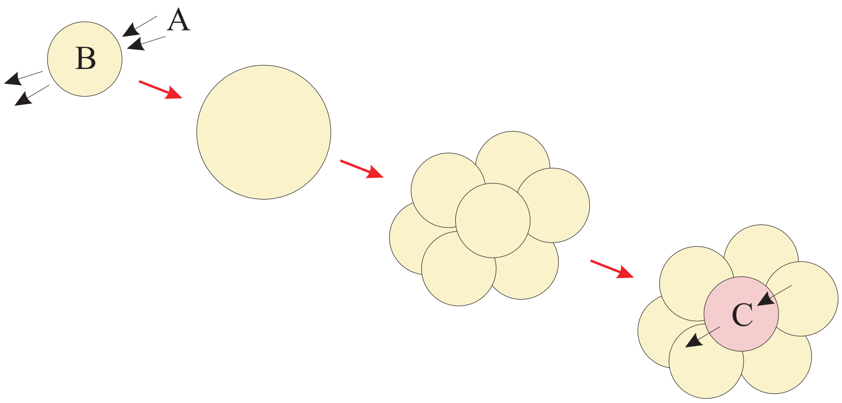

Figure 9.6: Illustration of how there can exist “cells” also without actual hereditary material: There can be nutrient intake, cell growth, reproduction, and even differentiation in the low level with no explicit developmental control. **a.** Droplet containing a set of autocatalysts is characterized by a reaction, say, A→B. **b.** There is plenty of chemical A available, so that concentration of B increases, and osmosis makes the droplet absorb water and grow in size. **c.** The droplet splits up in smaller ones because of the weakening of the surface tension; adhesion keeps the droplets still together. **d.** The middle droplets experiencing a new environment, different chemical reactions become appropriate, reaction now being B→C; such developments continue depending of the mixture of autocatalysts

First, the “heat death” within the system makes it possible for the very fragile molecules, like proteins, to remain whole, balance thus promoting survival; second, the balance makes it possible for very vague phenomena to become magnified — the “signal-to-noise ratio” in the system becomes high, revealing the remaining information in the signals, balance thus promoting more consistent evolution by making the appropriate adaptation direction better visible. In a well-maintained balance, it is enough that there is some special chemical that can only be utilized by the new cell type with a special set of autocatalysts; this gives it the adequate competitive advantage. Only if there exist alternative solutions (in terms of chemical solutions), secondary aspects like reproduction speeds become relevant in competition. Thus, around the balance, increased complexity is an evolutionary advantage, and proto-cells hosting more complex (and typically slower) reactions will survive in the chemical evolution.

No specialized “cell organelles” are needed to implement the basic cell-like functionalities; it is just assumed that the proto-cells are not completely isolated but search for the balance of reactions in the prevailing environment. It is easy to imagine what can happen next: Different proto-cells or cell groups can start exhausting each other’s surplus products, and become *symbiotic.* From the point of view of a single functional unit, other ones start to do the “sanitation”, exploiting the excess “waste products” — otherwise the reactions cease, the proto-cell suffocates, or, at least, its well-being becomes jeopardized. But, as seen from the point of view of the neighbors, such excess products are available resources, deviations from the nominal to be exploited. If the different kinds of proto-cells are dependent of each other, they probably grow and divide at the same rate, following the cybernetic balance; this is a rather plausible route to "multiglobular" systems (again, see 9.6). Cells without partners starve and become outnumbered. As seen from outside, different cell groups represent different sets of reactions, so that functional differentiation starts taking place. Control of the flourishing diversity is local, following the cybernetic principles; because of the physical realm, the experienced environments differ, and differentiation among cells start taking place.

When looking at the early development of an animal embryo, this kind of scenarios of differentiation is exactly what seems to take place even today in a fertilized egg: First the prototype cells form a *morula* or *blastula,* where the totipotent cells start differentiating depending on their environments. In a way, the Haeckel's intuition (see 7.3) should apply also to the earliest phases of life, giving motivation to the above studies: Individuals repeat the whole sequence of becoming alive. Perhaps there is something to learn in today's developmental processes when trying to reconstruct the origin of life. Of course, there still is an essential difference: Even in the simplest cell today there are the instructions, or the DNA code readily available, and the development is in this sense preprogrammed. But in the stem cell phase of morulas, no genetic imprinting has yet taken place. The simplest processes need no instructions; genetic code is only needed when there are alternative routes to select in the development. Perhaps the genes only started orchestrating the natural processes.

## 9.3 Codes and beyond

Can the above-like non-genetic, strictly chemical behaviors be called life? Today's life forms in biosphere are all characterized by genetic code, and it seems that there is a huge leap from non-genetic to genetically controlled. However, it seems that evolution towards such more sophisticated control of structures can still be explained in a rather consistent way, and no giant leaps are needed.

### 9.3.1 Towards programmed structures

The functionalities in the proto-cells need not be something clever or preplanned, as long as they exploit the chemicals available in their environments and produce something else — in short, being successful is capability of being active, exploiting the available chemical resources. The environment is not predestinated, as it is the surrounding set of successful proto-cells that *create* this environment. When there is appropriate accommodation, the system as a whole starts looking "clever" — but only as seen in retrospect.

When studying the possibilities of more complicated functionalities to emerge, one needs to distinguish between two separate things: First, there is the ability to reproduce, and, second, there is the ability to modify cellular metabolics. Traditionally, it is assumed that it is the same solution (genetic code) that is responsible for both of these capabilities — but this need not be the case. It is the autocatalysts that have the reproduction capability; some other chemicals can be multifunctional ones. In the lowest level, it is enough that some chemical operates in different ways in different chemical environments. For example, inert and active states can be toggled depending on the environmental conditions like pH or temperature. This means that the reactions are nonlinear. The operating modes of the cell being integrated in the chemicals themselves, the cell functionalities are accordingly changed when the environment changes.

In the proto-cells, genetic code is also not necessarily needed to control behaviors; not even any complex molecules like nucleic acids or amino acids are needed in the beginning. No code reading capability is necessary to begin with. Of course, it is practical if the two presented capabilities, reproduction and multifunctionality, are combined in a single autocatalytic molecule. And — as seen in retrospect — it seems that DNA has outperformed all other mechanisms. The combination of DNA as the code and proteins as the tools for implementing functionalities is a very versatile combination, offering almost as flexible platform for different kinds of chemical structures as the neural machinery offers for cognitive structures.

Still, all information that is inherited needs not be transferred in the form of DNA. The *Lamarckian* theories have been neglected because it has been claimed that there are no necessary mechanisms to implement such views — however, also in the highly developed forms of life, there *are* other mechanisms available. It need not be assumed that the initial state of the stem cells is completely null; there can be some chemicals that follow the genetic material into the gametes, being manifested in the tsygote. This kind of inheritance can be called *epigenetic,* being also related to *genetic imprinting.* However, it is not any acquired properties that can be inherited this way; it is the commands of which of the available genes are activated in the beginning. Another issue is that it has been recognized that the microbial symbiotic fauna seems to be also inherited from the mother. As has been recognized, this symbiotic inheritance can essentially affect the metabolic processes that are activated in off-spring.

Such symbiotic systems illustrate that everything needs not be coded in the same genome in a centralized manner. Coordinated operation and reproduction is possible without sharing any genetic material. For example, *mitochondria* in cells have their own genetic codes; *lichens* are associations of a fungus with a photosynthetic partner that can produce food for the lichen from sunlight. All the subsystems are still based on DNA of their own; these codes need to have coevolved.

Even though there were only a single set of codes, there is need for coevolution. For example, the trinity of DNA, RNA, and proteins necessarily had to be there from the very beginning in some simple form, even though the roles of the components need not have been so clear-cut — the theory of the “RNA world” as the immediate predecessor of the modern life forms probably cannot hold. As discussed in the following section, the developments from the beginning of life to the present day have to be smooth — the basic structures cannot be changed abruptly. Even though there are tensions towards more sophisticated structures, it would be difficult to understand huge sudden steps in developments. The interesting question that remains is what the actual autocatalytic set of simplest proto-DNA, proto-mRNA, and “protoin” is.

Continuity (and differentiability) of functions is the key to efficient optimization in mathematics; otherwise, one only can do random search, and in a high-dimensional space this is extremely inefficient. How can continuity and consistent adaptation be reached when the functions are based on discrete genes? And, further — how can the very discrete nature of structures in the phenotype be explained if the functions are continuous?

The key point is that it is not a static one-to-one mapping from the genotype to the phenotype, but it is dynamic processes that implement the mapping, the static-looking patterns being the final dynamic equilibria. As the genetic system is in contact with its environment, it searches the balance; redundancy of the genes, and the quantitative nature of gene expression makes it possible to reach continuity. And because of the sparse-coded nonlinear nature of the genes, there can exist various equilibria: Minor changes in environmental conditions can result in very different outcomes, giving raise to emergent structures.

Genes are modified in a Darwinian process of mutation and crossover; however, the genes are not actually optimized. The main role of evolutionary processes is to generate variation: The goal is to supply material, a pool of alternatives, whereas the local balances within a cell finally select the appropriate genes, revealing the actual potential and limits of the new genetic combination. The genetic process determines the (sparse coded) subspace in the metabolic space, and other processes are utilized for final optimization within those subspaces. The genes only determine the *potential* in terms of available degrees of freedom, whereas the environment determines the *actual,* the location of the equilibrium in the search space. In each cell the same functionalities in latent form wait to be activated. Optimality of solutions is defined in a very local and immediate fashion, there is no need to wait feedback from explicit "goodness" evaluator, with the delay being of the order of one generation — level of match with the surroundings suffices. Yet another fact needs to be recognized: There is no global single fitness criterion. Each variable is being matched more or less independently, so that, in a sense, "parallel processing" for fitting the data is implemented, further enhancing the adaptation speed.

The genes are hierarchic, and there is often accumulation of various individual genes that is needed to implement some more sophisticated functionalities. The benefits of the genes are visible only after the whole structure is completed — how can the sudden emergence of such complete functionalities be explained? However, the local minima are not necessarily very far apart, and the chain of gene activations can still be reasonably cut in subparts, as studied below.

### 9.3.2 Case: Development of an *eye*

It has been claimed that evolution theory cannot hold — there exist a plenty of highly complex structures that are functional only when they are complete. As long as the structure is not yet fully developed, the infrastructure for it is only a burden, and evolutionarily disadvantageous; this should mean that the barrier between the local fitness maxima is too wide to be crossed. The complex organ should have emerged immediately, without intermediate steps, and this is simply too improbable. A typical example is the *eye* — an example that has been studied widely in literature.

However, it turns out that there is a path from no-eye to a complete eye consisting of simple gradual steps where each stage is evolutionarily beneficial.

Indeed, it has been observed that the eye has developed various times in different branches of the "tree of life", and the solutions are not unique. Below, one simple scenario is presented.

Still, it seems that the population-level feedback loop between better properties and consequently more probable survival is too inefficient to support the consistent development of structures — and, especially, the simultaneous development of separate functionalities seems like a too lucky coincidence. For example, enhancements in the eye cannot be exploited if the processing of the neural signals is neglected; the eye and the brain have to develop in a somehow orchestrated fashion. — Indeed, it seems that neocybernetics may offer some tools to understand such dilemmas, as in that framework the *fitness criterion can be distributed.* All cells simply try to maximize the emergy they receive from their environments, whatever form this variation takes; emergy is not only physical nourishment but information in general. There is no need for external evaluation; variation or information can be regarded as beneficial no matter whether that information can be exploited locally or not. As the system itself is quite well-balanced, increasing variation reflects enhanced coupling with the environment. Increasing excitation in the eye means increase in nerve cells, and increasing excitation in the nerve cells means motivation for brains to develop. It may be (it is) so that at the higher (proto)animal level better processing of the visual signals means better possibilities of responding to the threats and opportunities in the environment, thus improving the survival of the organism as a whole, but the underlying organs and tissues can see the developmental gradients more instantly.

In Fig. 9.7, a simple scheme is depicted where the development of an eye can be understood. Each step in the series of stages means more accurate detection of behaviors in the environment — first, there is the capability of seeing whether it is dark or not; second, the direction of the light source can be detected; after that, it becomes possible to tell patterns from each other, with ever increasing accuracy and sensitivity. The presented development process is by no means unambiguous — for example, compare it to the compound eye structure of a house fly.

As compared to discussions in chapter 8, it is again interesting to study the relationship between the genotype and the resulting phenotype — what the "interpreter" is like, and how "semantics" in the physical and physiological environment can be defined in the above hypothetical case. Evidently, the results of interpretation are this time determined by the limited information delivery in the physical domain. The physical dimensions are crucial — it can be assumed that the signals are transferred in terms of chemicals; as there is no way to control the spreading of the chemicals, the environment of cells is determined in terms of *concentration gradients.* This means that the codes have to be essentially location-based, the fixed points in the configuration of cells being determined by some activated cells expressing some special genes and producing some signaling chemicals: It can be assumed that the concentrations of the diffusing chemicals decay monotonically (exponentially?) as the distance to the imprinted cells increases.

Chemicals, or, indeed proteins, are the only immediate outcome of gene expression — how are the chemical concentrations changed to physical properties of

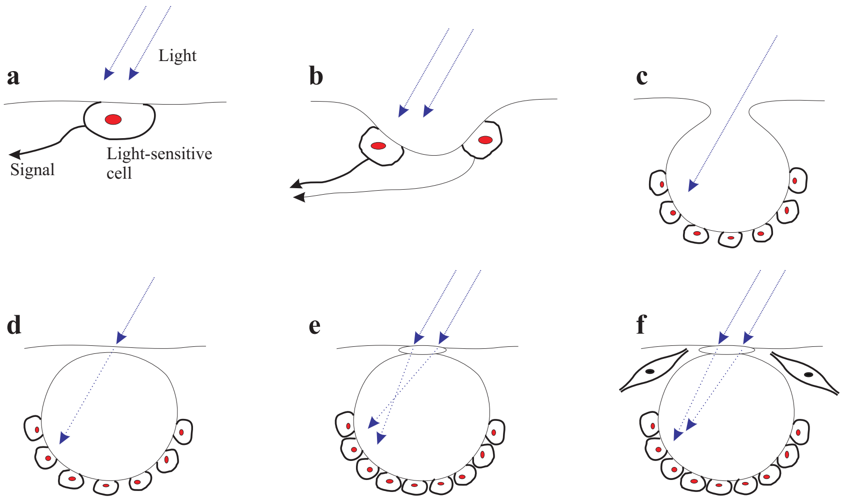

Figure 9.7: The development of an eye can be explained as a continuous process of enhanced information retrieval. In **a**, there is just a single light-sensitive cell — but knowledge of whether it is dark or light is already valuable information. In **b**, physical reasons can make the cells differentiate, as the amount of absorbed light depends on the direction where the light is coming from — and knowing *where* there is light can be crucial information. In **c** this differentiation has proceeded, so that what one has is a simple "needle's eye" camera — it is already possible to distinguish between different light sources. Later, in **d**, the process of increasing the eye resolution ends in the opposite walls of the cavity merging together — this, too, is beneficial as the robustness of the proto-eye increases, the sensor cells becoming isolated from the environment. It is clear that there is evolutionary advantage if this filter layer becomes more transparent — and as it does, in **e**, one has a lens, making it possible to reach much higher light sensitivity and resolution at the same time. Finally, in **f**, the ready-to-use muscle functionality is employed to deform the lens, thus increasing the adaptation capability of the eye. It is clear that the brain has to co-evolve to make use of the available new information — if the information is not cleverly exploited, the evolutionary pressures vanish and the consistent developments cease

tissues? Enzymes and transcription factors can promote cell metabolism and reproduction, and possibility of increasing activity means increasing biomass. The thickness of tissues is related to numbers of cells; further, the transparency of cell layers is related to thickness, etc. — when the physical properties of cells become manifested, all physical functionalities are available that are prerequisites for imaging and image processing. Again, the physical constraints result in emergence of smart-looking structures — as seen in retrospect: Development of cavities facilitates differentiation among the light-sensitive cells, etc.

According to the above lines of thought, below is a simplified example of what the "eye program" could look like, when a "transcription" from DNA into an explicit pseudocode is carried out:

```
1.  IF ``location'' $\approx p_1$
    THEN imprint ``EYE'':
              utilize light sensitivity, emit ``eye'' and NGF

2.  IF ``SKIN'' AND ``eye'' $\approx p_2$ AND ``lens'' < eps
    THEN split up

3.  IF ``location'' $\approx p_1$ AND ``EYE'' $\approx p_3$ AND ``lens'' < $p_4$
    THEN imprint ``LENS'': develop transparency, emit MGF

4.  IF ``lens'' $\approx p_5$ AND ``muscle'' < $p_6$
    THEN imprint ``MUSCLE'': grow towards MGF, emit ``muscle''
              etc.
```

In the above code, each of the four rules represents a function of its own, or a gene (or set of redundant genes), as listed in order of assumed activation. The first row in the rules describes the control part, the rest determines the “actions”. It needs to be recognized that the “interpreter” for the code is distributed, running separately in each cell; the code is identical in the cells, so that differing reactions are caused by the local environments. The communication and coordination among the cells is implemented trough special signaling chemicals that are in the code denoted by quoted lowercase names. The quoted uppercase names denote imprinting — that is, the cell is assumed to have reached some specific state and its role has been determined.

The first gene becomes activated if one (or more) chemical levels match the preset value(s) $p_1$; such “location” signals truly exist in a real embryo where, for example, they implement the anterior–posterior an dorsal–ventral asymmetries. The comparison operation is streamlined here — this genetic control perhaps has to be composed of several elementary toggles as presented in 6.2.4. When this gene is activated, the corresponding cell will forever have the role of an eye, its special property being light sensitivity. Simultaneously, it produces chemicals: when the signal “eye” diffuses outside the cell, the neighbors can detect its existence, and nerve growth factor NGF starts persuading nerve cells to connect to the cell. The second gene can only be activated in a skin cell: If there is an appropriate distance to the eye cells (chemical “eye” in the surrounding tissue having decayed to the level $p_2$) and there is no “lens” signal present, the cell starts reproducing, thus making the number of cells grow, causing the skin get wrinkled, leaving the eye cells on the bottom of a cavity. It is assumed that neighboring cells automatically attach to each other when they are in contact. The reproduction ceases as the lens develops — this happens in the cells that are in the center of the eye area but are not eye cells. In evolution, such lens cells develop towards better transparency. The lens cells also secrete nerve growth factor, trying to persuade muscle cells to attach to the lens. The last rule represents the muscle as being exploited by the eye — however, muscles are of course general purpose structures, and they can be activated also through other sequences.

The "muscle" example above is characteristic to genetic systems — once the muscle functionality has been "invented", it is readily available and can be exploited in different organs. It can pop up in different structures — and if it is beneficial (as it is in the eye), this new functionality is supported by later evolution. The genetic substructures are ready to pop up as soon as such functionality is needed; in this sense "genetic design" is like *functional programming.*

The functional structure of the genes as proposed above is very simple offering just the coarse framework for physiological structures; within this discrete representation, there is continuity in the structures. Quantitative fine-tuning is possible in terms of the parameters $p_1$, $p_2$, $p_3$, $p_4$, $p_5$, and $p_6$ affecting the eye dimensions. As the two halves of the genome are inherited from differing individuals, the threshold values in the parents typically become averaged in the offspring. This is an efficient optimization scheme when there is only a narrow region of acceptable parameter values available. Different-looking species are possible when different parameter values are selected — and, indeed, as most of the genes are common to all life forms, some kind of fine-tuning of their effects is necessary.

### 9.3.3 Optimality in mechanical structures

Not everything can be coded in genes: After all, genes represent central control, and if employing only them, behaviors would not be tuned maximally — or, at least, adaptation towards the optimum would be very slow. The genetic machinery needs to be accompanied with better adjustable mechanisms to reach the fine tuning. In a way, it is as it is with cognitive systems: The lowest level (chemical concentrations or synaptic weights) is continuous, the "intermediate level" (organ structures or conscious thinking) is discontinuous, and the highest level (complete optimized organism or automated behaviors) is again more or less continuous and optimizable; to facilitate real-life survival, automation of slow cognitive processes has to take place, and, similarly, the final polishing of the structures in living bodies takes place after the actual implementation of the codes. Especially, this means that evolution of fitness cannot be based on so delayed mechanisms as it is assumed when speaking of natural selection; more immediate feedback mechanisms are necessary.

In principle, the neocybernetic optimization principles can be applied in any environment and at any level. However, the intuitions about uniformity among signals collapse when the system is a sophisticated functional entity; the vision of a cybernetic system as reflecting its environment also becomes far-fetched. The relationship between the system and its environment becomes blurred as it is the other organs that deliver the input signals to an organ — this environment also changes, and one should implement optimization for all systems simultaneously. Can the basic neocybernetic model of separate system and environment be applied any more when all signal are internal ones? This issue can be studied when looking at Fig. 9.5 again: When the system and the environment become one, the only thing that remains outside is the feature extraction. Now this generation of features, or manipulation of measurements, is not artificially constructed by some designer, but it reflects the effect of how the real world distorts the signal transmission process among the organs.

To make the above discussion more concrete, let us concentrate on the question in what sense the *outlook* of an organism can be captured in compact formalisms and optimized therewith. To proceed, one needs to imagine how the organs act as probes deforming the "steel plate" around them (see chapter 3). The outside world is not known, but iterative adaptation within the individual organs (or cells) still optimizes the system.

To have some more background, it is necessary to get acquainted with the techniques of modeling mechanical systems. In *Lagrangian mechanics* it is observed that the Newtonian laws of motion can be reformulated as optimization problems: Along the motion trajectory, the time integral of the quantity $L = W_{\mathrm{kin}} W_{\mathrm{pot}}$ reaches its minimum value, where $W_{\mathrm{kin}}$ is the kinetic energy of the system, and $W_{\mathrm{pot}}$ is the potential energy. Applying the vector of generalized coordinates $q$, the kinetic total energy can be expressed as

$$W_{\mathrm{kin}} = \frac{1}{2}\dot{q}^T I \, \dot{q}, \tag{9.3}$$

where $I$ is the *inertia matrix,* and the vector$\dot{q}$ stands for the generalized velocities (translational or rotational). Does this not look familiar? Indeed, when defining $x = \dot{q}$, the basic neocybernetic cost criterion can be interpreted in this framework:

$$J = \frac{1}{2}x^T \mathrm{E}\left\{\bar{x}\bar{x}^T\right\} x - x^T \mathrm{E}\left\{\bar{x}F^T\right\} F. \tag{9.4}$$

Now, $\mathrm{E}\left\{\bar{x}\bar{x}^T\right\}$ can be interpreted as the inertia matrix. If $F$ is the vector of forces and torques that can sustain the corresponding velocities, $\mathrm{E}\left\{\bar{x}F^T\right\}$ becomes some kind of a *viscosity matrix* and the latter term in (9.4) is the viscous work (or power lost in movement). This is an extension of the Lagrangian thinking: Forces in the assumed system are *non-conservative,* as the "potential" is not free of the velocity variable.

It is an open question whether the above cost criterion truly has relevance in real life. Yet, if it does, this cost criterion offers a framework for analysis of natural life forms; what is more, it makes it possible to create "cybernetic designs" in life-like (biomimetic) structures. It is clear that the balanced designs are optimal in such a sense that *maximum amount of correlated forces are transferred into movement.*

Assume that muscle cells (forces) and sensory neural cells (velocity measurements) are modeled together. The relationship between these variables is determined mainly by the limb configuration — this relationship implements the "feature generation". When the integrated system becomes optimized in the neocybernetic sense, constructing a statistically balanced model for the relationships between the variables, the system speed and agility become optimized automatically: There will be maximum possible velocity in structured (sparse coded) directions with the minimum effort. The iterative optimization process can end in outlooks that differ very much from the initial. Local adaptations in the structure (as accumulated in the inertia and viscosity matrices) are reflected in the increasing overall "fitness" of the global structure. Even though the goals of adaptation are not fixed beforehand, the direction of "better performance" is known by the local actors, and the post-genetic developments are not random — there are gradient directions visible. For example, rehearsing of muscles makes it possible to learn the model between the forces and corresponding limb velocities; adaptation of this model results in ever more optimal and economical (and thus more "beautiful") trajectories — the final outlook of the body need not be coded in the genes.

## 9.4 Are we alone?

As discussed above, it can be assumed that life unavoidably emerges if the conditions are favorable, and if there is enough variation in the conditions so that the modeling task is non-trivial. And as the arrow of entropy is inverted, it is not difficult to imagine that intelligent life is just the next step in the inevitable development of life forms. But if the origin of life and intelligence can be explained in such a straightforward manner, one is facing yet another paradox. When there assumedly exist millions of planets that can host life, and as the evolution sooner or later results in intelligence emerging among the life forms, there is a question that was originally coined by Enrico Fermi: "Where are they?". Why cannot we see the activity of the other civilizations? There simply must be other (more) intelligent civilizations in the universe in addition to us.

When the radio frequency spectrum has been scanned, nothing "intelligent-looking" has been found, only noise has been detected in the signals coming from the stars. But there is a simple (partial) explanation available here. One can only search for redundancies in the signals — but, from the point of view of transmission efficiency, redundancy necessarily means unoptimality. A message with all redundancy ripped off looks like noise if the decoding scheme is not known. It is a very short period in the history of a civilization that transmissions are not optimized and packed; in our case it is something like 100 years only. — But it even seems that the other civilizations actively try to keep the distance, why is that? One can make some hypotheses here.

Why are they not trying to contact us — as we do, sending easily decodable signals to us on purpose? But, on the other hand, why should a civilization make a big number of itself? Only civilizations being in the early stages of their intellectual development make a big fuzz of themselves — the older and more mature ones observing this blustering sympathetically. If a civilization is to survive the turmoil periods there are in the development, the periods of chauvinistic arrogance need to be overcome. Indeed, knowing that there are civilizations millions of years ahead of us we should perhaps be a bit ashamed. It is plausible that we could not even recognize the systems far ahead of us: After the chemospheres and biospheres, our frontier systems today reside in infosphere. But after the principles of intelligence (or infosphere cybernetics, really) are fully implemented in the computer, developments in infosystems become very fast — so fast that when the time axes in information modeling processes collapse into singularities, qualitatively yet another level is perhaps reached. It is impossible for us to understand the higher-level systems: It is like biological systems facing cognitive systems — trying to "understand" the structures on the emergent level is an intellectual contradiction.

But why is this passivity so categorical — why do none of the higher-level civilizations even play with ourselves? Indeed, this consistency promises that we are on the track of something big here.

It seems that this complete silence is purposeful: They do not want to disturb us, they want to see how we manage on our own. Preservation of life and natural diversity is seen important by all intelligent civilizations — and there is a reason for that. The analyst does not want to disturb the processes he/she/it is observing. The problem of life seems to remain an eternal challenge for intelligent minds: Understanding the mystery of life is understanding survival, and it seems that it is cybernetic-like modeling over the spectrum of possible forms of life that will continue as long as the civilization lives. More material, more fresh data is needed to map different local solutions in different environments.

But it is not only curiosity that drives such galactic research — such research is necessary to maintain sustainable development. Also extraterrestrial life is facing the limits of its home planet, and the only way forward is available in the infinite space. All evolving civilizations have to be based on science and information pursuit, searching for new frontiers. The link between science and the society is the question of life and death to a civilization, even more than what we can understand today. The more intelligent a species is the more it is dependent of scientific research, as acquiring new information seems to be the survival strategy. The intelligent species necessarily have found the principles of cybernetics, and they must understand that the only way to avoid cybernetic stagnation and catastrophes is to receive ever new information, and new sources of information. It might be so that it is us as a peculiar example of living systems that provide a piece of this crucial information. What is the nature of this information, then? Of course, such knowledge is beyond our capabilities of understanding. We cannot yet see "higher-order life" in such a wider perspective where our world would just be a single sample case.

From the point of view of the higher-level intelligence, we are running just another experiment. As Douglas Adams observed in his "Hitchhiker's Guide to the Galaxy": The Earth is a giant simulator. Perhaps some day — if we pass the test of survival — we receive an invitation to the "Galactic Board of Intelligent Species" where the universal experiment designs are coordinated. And perhaps it is our maturity test — being faced by any developing civilization — to understand the cybernetic processes, tame them, and avoid the downfalls, making the succession of ever more deadly catastrophes a steady process of sustainable development.

The risk of mankind committing an explicit suicide is today a well-understood risk — but it seems that there also exist more latent threats to a developing civilization. As studied in [22], the developments in Tasmania may give us a hint of this risk:

> Some 10000 years ago Tasmania was cut off from the Australian mainland, and about 4000 Aboriginal Australians remained totally isolated. When Europeans discovered Tasmania in the 17th century, it was technologically the simplest, most primitive human society. Native Tasmanians could not light a fire from scratch, they did not have bone tools, they did not have multi-piece stone tools, they

> did not have axes with handles, they did not have spear-throwers, they did not have boomerangs, and they did not even know how to fish. Incredibly, archeological investigations have shown that during those 10000 years of isolation, the Tasmanians actually *lost* some technologies that they had carried from the mainland to Tasmania. What caused this decay in civilization?

There were no catastrophes in the Tasmanian culture — before Europeans, there were no rivals and no external disturbances to shake the system that had reached the stagnation. The smaller the population is, the faster it seems to reach the stasis — but, similarly, a planet-wide monoculture can decay; what is the difference between evolution and devolution?

Living systems seem to share the property of vitality — there is some kind of *arrogance,* tendency to grow and conquer. Growth has to be eternal, but this growth need not be physical, it can also be mental. The internal spirit is a matter of life and death — and even though the cybernetic principles are universal, the evolutionary processes cease if this spirit is missing. It seems that vitality must be explicitly maintained. What does such loss of memetic vital force look like in concrete terms? It can be claimed that it is *pessimism* — as seen in the scale of the whole civilization — that characterizes the end of culture: When everybody starts looking back into some lost paradise, trying to oppose the change, there will be decay. The system as a whole must keep up optimism and curiosity, looking forward even if facing the “cosmic angst” in front of the unknown future.

It has been said that if an individual human being wants to be happy throughout one’s life, keeping up optimism and good humor, one should do *gardening,* letting living systems grow and seeing them prosper. Perhaps the same ideas are the key to sustainable, non-explosive developments in the wider scale, too — if an individual civilization wants to live “happily” ever after, it should do “gardening” of lower-level civilizations. And just as a good gardener protects its plants, nourishing and eliminating hazards, perhaps the “universal gardener”, the “cosmic philanthropist”, also protects its planets, looking after us ... perhaps the belief in personal gods (or UFO’s), our higher-level protectors, is not completely unjustified?