Processes of Life: Towards a Unified Model?

Heikki Hyötyniemi

Cybernetics Group
Helsinki University of Technology, Control Engineering Laboratory
P.O. Box 5500, FIN-02015 HUT, Finland
heikki.hyotyniemi@hut.fi
http://www.control.hut.fi/hyotyniemi

Abstract. What is life? Somehow, it seems that intuitively life-like phenomena emerge from dead materia when enough complexity cumulates. This paper proposes physically/chemically more or less motivated models for capturing the relevant phenomena underlying life processes. It is not surprising that the model structures that pop up are analogous to those ones that have been derived for general cybernetic systems. The studies start in a down-to-earth manner, proceeding towards more complex, and more speculative discussions.

1 Introduction

In *artificial life (AL)* research is typically based on simulations, where evolution of artificial life forms is monitored in artificial environments. However, it seems that after the early enthusiasm, the intuitive feel of the *essence* of life is still missing. The key issue, or the key to "real life" is the capability of adaptation according to unanticipated changes. Something is lost when the environment is preprogrammed. When the basic principles underlying life are still a mystery, it seems that synthetic approaches are of no use; still, the only way to analyze life is to study how it is manifested in nature. Real life has to be studied in real environments, in analytic rather than in synthetic manner, analysing existing biological systems.

After the era of symbolic approaches, data-oriented, computational approaches are flourishing in today's artificial intelligence research; it seems that also in artificial life fresh intuitions can be reached when empirical data is carefully studied — applying modern multivariate statistical tools, and applying new conceptual ideas.

The data-oriented approach can help to avoid deadlocks. For example, folding of proteins should be modeled accurately to reach first-hand knowledge of their activity profiles. However, it has turned out that such studies are extremely elaborate and computer intensive. Whereas such models are necessary when new drugs are being designed and their effects are being predicted, now, when existing metabolics are being modeled, there is no need to study first principles. The actual contributions of chemicals are (hopefully) visible in data, and the system can be studied in a behavior-based manner. Of course, not everything

can be done based only on observation data: The interesting challenge is to find out where is the boundary region between the *possible* and *impossible.*

Methods for measuring different bio-entities have improved considerably during the last decade. DNA microarrays is a very powerful technology for measuring the gene expression responses (in terms of messager RNA's) in organisms. Also protein expression can be measured in a quantitative manner from relatively complex mixtures using mass spectrometry. In parallel, methods such as chromatin immunoprecipitation (ChIP) offer data about protein-protein or protein-DNA interactions. Thus, today there exist very much data but not enough understanding. The wealth of data makes it difficult to see the big picture, and conceptual tools are needed to manage the complexity.

Mathematical modeling methods including statistical approaches have been successfully applied to many biological problems. Such methods are often studied under the name *bioinformatics;* in more general setting, one could speak of *data mining* (for example, see [20]). While the value of old methods has been rediscovered in the biomedical context, there is also a need for deploying new approaches.

To extract real information out from the data, there is need for *systemic approaches* in systems biology [17]. The current models are reductionistic, trying to see the big picture when looking at tiny fragments of reality at a time. One should study the processes in a systemic way, taking the interactions among the system components into account. Life phenomena are processes rather than static structures, and real understanding of these processes cannot be reached without understanding of the underlying interactions and feedback loops. Because of the huge challenges, one should not be too ambitious. Rather than pursuing general system theory (see [3]), this paper concentrates on the engineering-like approaches: System theory is utilized because of the conceptual and practical tools it offers (model structures, and related mathematics).

Finding models for life processes would be very important also from the practical point of view. Understanding the coordination among genes and metabolics is pivotal in biomedical research. It helps, for example, in understanding development and disease states, responses to the environment and adaptation to specific conditions or cell behaviour in normal and triggered situations. Being capable of inferring regulatory or functional networks from present data to help designing new and more precise experiments is a key issue in bioinformatics. Another issue is to simulate biological processes such as cell behaviour, so that savings in drug development process, etc., could be reached. Thus, in 2004 a two-year project "Systemic Models for Gene Expression and Metabolic Dynamics (SyMbolic)" funded by Tekes NeoBio program was launched. Co-operating partners in the project are Helsinki University of Technology (HUT), Center for Scientific Computing (CSC), and MediCel Oy. The goal in the project is to study ways to define simple but powerful models for representing dynamic phenomena within a cell and a nucleus from different, complementary points of view. This paper presents some "background material" perhaps offering some ideas for systemic approaches towards analysis of the life processes.

2 Modeling metabolics

In what follows, analysis of the metabolic system is started from simple static models applying well-known ideas of chemistry. However, it turns out that if additional (more or less plausible) hypotheses are made, more powerful model structures can be derived.

2.1 About models

First, it has to be repeated that a model is just a model, being always simplification of reality. As has been recognized, *all models are false, but some can be useful.* Models are abstractions that ignore details and make it possible to concentrate, hopefully, on the relevant phenomena only. Good models make it possible to make predictions of the system behavior. Sometimes, intuitions and understanding of the system can also be gained; and, further, sometimes a good model helps to see analogues, making it possible to extrapolate beyond the formal validity range of the model. If such a "higher-level" model can be found, a new conceptual level of new explanations is reached.

Higher-level models are based on some more general principles rather than on the immediately visible phenomena. Note that there are many ways to proceed when aiming towards higher-level models. The traditional model structures are not necessarily well suited for modeling biological systems. For example, it is typical that one wants to find causal models, where action/reaction pairs are easily managed. However, in cybernetic systems such studies are doomed: Trying to capture the interactions in explicit dependency structures actually hides the underlying network of interactions. In a biological system, for example, all variables are mutually connected, and pairwise causality models fail to represent the overall picture; indeed, they can promote false intuitions.

The current model structures are typically static, that is, they only capture a single steady state, trying to capture the connections between genes, mRNA's, and metabolites using more or less fancy model structures. There also exist approaches to dynamic modeling applying traditional mass balances and compartment models (for example, see [21]), but the overall picture is still missing. The problem is that in reality the history of a cell dictates its future behavior and its responses to environmental activations, and to develop models for developmental biology purposes, for example, this history should be modeled. For example, whereas the genetic code is identical in each cell, only the *stem cells* are unspecialized, being still capable of developing into any cell type. However, there are problems with dynamic models: As compared to static models, there typically exist more parameters — and as the degrees of freedom are increased, it soon turns out that not enough experiments can be carried out.

This all means that one is facing a huge modeling task. Thinking pragmatically, powerful model structures are needed to compensate for the complexity of the domain area. In practice, this means *linear* models. Linearity makes it possible to efficiently apply multivariate statistical methods and linear algebra. Another key to model compactness is that domain-oriented structures are

employed, so that the expressional power of the model is limited to relevant phenomena only. Luckily enough, as will be shown, at least in some cases it turns out that linearity and domain-orientedness are not necessarily contradictory goals. In more complex cases, linearity assumption is a compromise between physical plausibility and usability: It cannot be determined beforehand whether the model understandability, analyzability, scalability, and easy parameter identifiability outweighs the loss in biological functionalities.

It needs to be emphasized that to explain complex behaviors it is not necessary to always apply complex models. As has been recognized in complexity theory, seemingly complex behaviors can also emerge from simple ones, if they cumulate sufficiently. And, indeed, complex behaviors can emerge also in (almost) linear structures; comprehensibility and analyzability of linear models is then a bonus benefit. Just as in [13], starting from simple starting points, higher levels of abstraction can be reached also when life processes are studied — or, at least, the immense mysteries of Nature can be studied from a fresh point of view. Understandability is more important than absolute accuracy.

Simplicity of the models is also one of the main goals in this research. Counterintuitively, it is through such higher-level models being based on more general modeling principles that one can reach this simplicity: Appropriate models constrain the degrees of freedom automatically, ignoring irrelevant phenomena, making it possible to survive with the wealth of data. Putting it poetically: If one looks in the right direction, one does not need to study shadows only.

2.2 Chemical reactions

First, let us study metabolic processes. It can be assumed that such processes can be characterized as chemical reactions; study a hypothetical example reaction, where the λ reactants on the left hand side are denoted as L_i, $1 \leq i \leq \lambda$, and the ρ products on the right hand side are R_j, $1 \leq j \leq \rho$:

$$l_1\mathrm{L}_1 + \cdots + l_\lambda\mathrm{L}_\lambda \quad \underset{k_2}{\overset{k_1}{\Leftrightarrow}} \quad r_1\mathrm{R}_1 + \cdots + r_\rho\mathrm{R}_\rho, \quad \Delta H. \tag{1}$$

The metabolic processes are typically reversible, so that the reaction can take place in both directions (k_1 being the reaction rate in forward and k_2 in backward direction). Symbol ΔH denotes the change in enthalpy, or inner energy, when the reaction takes place. It needs to be recognized that it is not only chemical reactions that can be expressed using such formulas; also phase transitions, etc., can be expressed in this form.

The practical problem is how to represent such a formula in a useful form for data-oriented approaches. It seems that a practical way to code such reactions in a mathematically compact form is to use vector formulation: Define two vectors

$$\varphi = \begin{pmatrix} -l_1 \\ \vdots \\ -l_\lambda \\ r_1 \\ \vdots \\ r_\rho \\ 1/c_T \end{pmatrix} \quad \text{and} \quad \mu = \begin{pmatrix} \Delta C_{\mathrm{L}_1} \\ \vdots \\ \Delta C_{\mathrm{L}_\lambda} \\ \Delta C_{\mathrm{R}_1} \\ \vdots \\ \Delta C_{\mathrm{R}_\rho} \\ \Delta T \end{pmatrix}, \tag{2}$$

so that φ contains the coefficients, and μ contains the changes in concentrations C of the chemicals (unit mole/liter). Parameter c_T is the (constant) specific heat capacity of the cellular liquids, appropriately scaled, so that one can express the total changes in the system as

$$\mu = \varphi\, x. \tag{3}$$

Here, x is a scalar that reveals "how much" (and in which direction) that reaction has proceeded. It turns out that reactions can be characterized applying linear algebra in the space of chemical concentrations: The above formulation makes it possible to express also more complicated cases in the same framework. When there are many simultaneous reactions taking place, φ is a *matrix* containing individual reactions (as characterized by the vectors of the form (2) as columns), and x is a vector; the weighted sum of reaction vectors μ reveals the total changes in chemical contents.

The row dimension of φ is such that all (active) chemicals in the system are involved. There can be zeros in columns of φ, meaning that respective chemicals are not involved in those reactions — the matrix structure is sparse. The columns of φ determine the degrees of freedom in the high-dimensional "chemical space"; these vectors ("independent components"; see [10]) constitute a (non-orthogonal) subspace basis in the space of chemicals.

Using the above framework, metabolic systems can be modeled: If one knows the rates of reactions, or the vector x, the changes in the chemical contents can be estimated. The derived model is extremely simple; it is easy to understand and operate on, being based on multivariate theories and linear algebra. Using such a model, it is (in principle) possible to determine the unknown quantities applying associative matching techniques when only a subset of chemical contents is known (that is, one first estimates x using the available measurements, and using this, the unknown quantities are determined).

This idea of invariances within a chemical system have been widely applied for metabolic modeling; the key term here is *flux balance analysis (FBA)* (for example, see [7]). However, traditionally the goal has been to construct static networks only; this means that one sets $\mu = 0$ in (3), and utilizes the resulting set of linear equations for determining dependencies among variables. However,

typically the set of equations is not completely determined, and one needs extra assumptions to find additional constraints; there is no common agreement of how this should be done.

In many ways, the model structure (3) is not what one is looking for: Most of all, it is static, and it seems that it cannot be extended into a dynamic model in a natural way. In a sense, the expression (1) only defines *syntax;* the *chemical semantics* is another thing. There is no divine mind mastering the chemical processes; it is local interactions among molecules that is the key to emergent behaviors. This can be also expressed in terms of *arrow of entropy.*

2.3 Balance pursuit

There is a big difference between what is *possible* and what is *probable,* that is, even though something may happen in principle, it will not actually happen. For example, the reaction invariants and their weights in Sec. 2.2 are irrelevant if the reactions simply do not happen. Dynamic modeling gives the tools to understand the *dynamic equilibrium* among chemicals. To understand the dynamic balance, the reaction mechanisms need to be studied closer.

Assume that it takes l_1 molecules of L_1, l_2 molecules of L_2, etc., according to (1), for one unit reaction to take place. This means that all these molecules have to be located sufficiently near to each other at some time instant for the forward reaction to take place. The probability for one molecule to be within the required range is proportional to the number of such molecules in a volume unit; this molecular density is revealed by consentration (when the unit is mole/liter; one mole always contains $6.022 \cdot 10^{23}$ particles). Because the locations of the molecules are independent of each other (this is the first approximation; later, *activities* are employed), the probability for several of them being found within the range is proportional to the product of concentrations. On the other hand, the reverse reaction probability is proportional to the concentrations of the right-hand-side molecules. Collected together, the rate of change for the concentration of the chemical L_1, for example, can be expressed as

$$\frac{d\,C_{\mathrm{L}_1}}{d\,t} = -k_1 C_{\mathrm{L}_1}^{l_1} \cdots C_{\mathrm{L}_\lambda}^{l_\lambda} + k_2 C_{\mathrm{R}_1}^{r_1} \cdots C_{\mathrm{R}_\rho}^{r_\rho}. \tag{4}$$

Note that the expression (1) has to desribe the actual reaction mechanism; stoichiometric net expressions are not good enough, because scalings essentially alter the structure of (4) changing the exponents. What is more, catalytic reactions must also be expressed in terms of actual reaction mechanisms.

In equilibrium state there holds $\frac{d\,C_{\mathrm{L}_1}}{d\,t} = 0$, etc., and one can define the constant

$$K = \frac{k_1}{k_2} = \frac{C_{\mathrm{R}_1}^{r_1} \cdots C_{\mathrm{R}_\rho}^{r_\rho}}{C_{\mathrm{L}_1}^{l_1} \cdots C_{\mathrm{L}_\lambda}^{l_\lambda}}. \tag{5}$$

In practice, the reaction rate factors are functions of temperature T according to the Arrhenius law, so that $k_i = k_i' \mathrm{e}^{c_i T}$. Also, the *activities,* or actual activation probabilities, of the chemicals may vary (for example, if there are complex

enzymes participating in the reaction, it is not only location but also orientation that is of relevance). These activities can still be assumed to be linearly dependent of the concentrations, $\mathcal{A} = aC$, where a is a (typically unknown) constant. This all means that one can write the polished form of (5) as

$$K' = \mathrm{e}^{(c_2-c_1)T}\,\frac{a_{\mathrm{R}_1}^{r_1} C_{\mathrm{R}_1}^{r_1} \cdots a_{\mathrm{R}_\rho}^{r_\rho} C_{\mathrm{R}_\rho}^{r_\rho}}{a_{\mathrm{L}_1}^{l_1} C_{\mathrm{L}_1}^{l_1} \cdots a_{\mathrm{L}_\lambda}^{l_\lambda} C_{\mathrm{L}_\lambda}^{l_\lambda}}. \tag{6}$$

Assume that the nominal temperature is $\bar{T}$, and the nominal concentration value for L_1 is $\bar{C}_{\mathrm{L}_1}$, and so on. Assumedly these values also fulfill (6). When the above expression is differentiated around this nominal state with respect to all variables, there holds for small deviations

$$\begin{aligned} 0 = -l_1\frac{K'}{\bar{C}_{\mathrm{L}_1}}\Delta C_{\mathrm{L}_1} - \cdots - l_\lambda\frac{K'}{\bar{C}_{\mathrm{L}_\lambda}}\Delta C_{\mathrm{L}_\lambda} \\ + r_1\frac{K'}{\bar{C}_{\mathrm{R}_1}}\Delta C_{\mathrm{R}_1} + \cdots + r_\rho\frac{K'}{\bar{C}_{\mathrm{R}_\rho}}\Delta C_{\mathrm{R}_\rho} \\ + (c_2 - c_1)K'\Delta T. \end{aligned} \tag{7}$$

Simplifying, this becomes

$$0 = -l_1\delta C_{\mathrm{L}_1} - \cdots - l_\lambda\delta C_{\mathrm{L}_\lambda} + r_1\delta C_{\mathrm{R}_1} + \cdots + r_\rho\delta C_{\mathrm{R}_\rho} + (c_2 - c_1)\Delta T, \tag{8}$$

where δC denotes proportional deviation from nominal concentration $\bar{C}$, so that $\delta C = \Delta C/\bar{C}$, whereas ΔT denotes absolute deviation from nominal temperature $\bar{T}$. This model is based on local linearization — but it has been observed that in living cells the conditions remain practically constant, no large deviations taking place, and local linearization is truly justified.

When there are several reversible reactions taking place in the system, each of them is governed by an expression of the locally linear form (8). If there are altogether $n-1$ active chemicals (plus temperature) in the system, these linear constraints constitute a *null space* within the n dimensional space. This means that in these directions there is no variability in steady state. Indeed, such a set of constraints gives a very static view of the cell system. There is another option, too: Rather than concentrating on the *constraints,* one can study the *degrees of freedom:* The remaining directions in the data space constitute a linear subspace where all variation among variables is concentrated. Note that the relative concentration quantities along different degrees of freedom are summable, so that these axes truly span a linear subspace in a mathematically solid way. Also, the metabolic processes can again be compactly written in terms of linear algebra:

$$-u = \phi x, \tag{9}$$

where ϕ is the matrix of basis vectors spanning the degrees of freedom, and x contains the corresponding "coordinate values". It seems that this multivariate statistical view of equilibrium systems is a novel one.

As compared to the model in Sec. 2.2, it needs to be recognized that now u contains (proportional) deviations from nominal point, whereas μ contains

changes from any point to another. The minus sign in (9) is added to emphasize that the metabolic system tries to *eliminate* the deviations. Whereas in Sec. 2.2 the errors can cumulate, now it does not matter where the deviations in u originate from. This means that the new model is more practical: For example, the non-isolatedness is no problem. Energy and matter can go through the cell membrane, the dynamic equilibrium model within the cell still holds.

There are dozens of individual underlying reactions controlling the cellular metabolics[1]. It can be assumed that the number of such constraint equations of the form (8) is not much smaller than the variable space dimension n. Or, indeed, it can be even assumed that there are no degrees of freedom left; the visible transients are caused by some of the equilibrium reactions being slower than the others. Then, concentrating on the visible degrees of freedom considerably simplifies the data analysis problem: It does not matter what the actual constraint equations are, one can ignore the null space altogether. When the number of balance reactions grows, the inverse modeling approach involving only the remaining degrees of freedom becomes more and more practical. It can be assumed that all interesting phenomena in the cell are revealed by the "metabolic degrees of freedom".

Now there exist efficient ways of analyzing the metabolic data: When *principal component analysis PCA* (for example, see [2]) is applied to the data set, where the variables are preprocessed as explained above (concentrations expressed in terms of proportional deviations from the nominal values), the relevant degrees of freedom are easily found. This means that matrix ϕ is constructed from those eigenvectors of the data covariance matrix that correspond to the most significant eigenvalues, that is, those directions in the data space where there is most variation.

Phenomena like buffering are typically regarded as a problem when experiments are carried out in a cell system: Huge step inputs may be needed to reach noticeable effects in some specific variable, and these effects cannot be focused, being reflected to the whole set of variables. Now, on the other hand, when applying multivariate analysis methods, such buffering is just a manifestation of the internal null space, and observations of the new balance deliver valuable information concerning the metabolic processes and functions. No one-input/one-output studies are needed. Another traditional problem in metabolic systems is that they seem to be highly redundant (this also applies to gene expression, see later). It seems that there typically is not just a single reaction mechanism explaining the processes, making it difficult to uniquely identify model parameters. Now, these problems are avoided: First, the actual reactions are not searched for, but the "residual" variations; second, PCA is just the right tool to model redundant and noisy phenomena, because it transforms from the

[1] What is more, note that complex reactions can take place in parts, where subprocesses follow each other; each of such intermediate products spans a new dimension in the variable space, and each chemical reaction introduces a new constraint, compensating for the increased dimensionality. The net effect is that there are *practically invisible* dimensions in the variable space

visible variables to new *latent variables,* where noise and redundancies among variables has been ripped off.

2.4 Model structure

When the model structure (9) is to be applied in analysis, the observation data that is collected in u has to be fitted against the vectors in ϕ. The matching can be based on the following criterion:

$$J(x) = \frac{1}{2} \left(u + \phi x\right)^T \left(u + \phi x\right). \tag{10}$$

This means that if (9) can be fulfilled for some $\bar{x}$, this $\bar{x}$ also gives $J(\bar{x}) = 0$, eliminating the deviations. In practice, u has higher dimension than x and exact match cannot be found; Eq. (10) defines the least-squares criterion for finding the best estimate for solution. This $\bar{x}$ can be searched for minimizing the criterion using, for example, the *steepest descent* approach. For this purpose, the gradient of the criterion is needed:

$$\frac{d\,J}{d\,x}(x(t)) = \phi^T \phi\, x(t) + \phi^T\, u. \tag{11}$$

Now the continuous-time version of the gradient algorithm can be written in the state-space form:

$$\frac{d\,x}{d\,t}(t) = \Lambda A\, x(t) + \Lambda B\, u, \tag{12}$$

where A and B are matrices, defined as

$$A = -\phi^T \phi, \quad \text{and} \quad B = -\phi^T. \tag{13}$$

Assuming that u remains constant, x converges to the unique solution of the criterion (note that $-\phi^T\phi$ has all of its eigenvalues in the negative half-plane, so that the process is stable). Above, the adaptation rate Λ is a (diagonal) matrix, meaning that the adaptation rate can be different for each variable; the steady state solution (in this linear case) still remains intact.

When looking at the chemical reactions within the cell, the reactions that restore balance can be rather slow, and the metabolic pathways can be complex. How is this balance restored, then? Again, there is no central control in a cell; for the chemical system, the gradient direction is the most reasonable way to go, it is the only locally visible direction. It can also be motivated in terms of local dynamic balances. This means, that model of the form (12) is not only a computational tool — it can be used also to approximate the dynamic transients within the cell!

Now, we have a simple, dynamic model that can be applied for modeling also transients in a metabolic system, as defined in (12). Parameter identification within the model structure should be simple: Matrices A and B can be determined according to the degrees of freedom in steady-state data, as shown

in (13), and, after that, the adaptation rates on the diagonal of the matrix Λ can be determined from transitory responses.

When a higher-level optimality criterion is being applied, the lower-level mess with the wealth of data becomes better manageable. Seen from the above, one can have an interpretation for dynamic metabolic processes: It can be assumed that the metabolic system implements "chemical pattern matching".

3 Modeling gene expression

There are various levels when studying life processes: The metabolic level is the lowest one. The level of *genetic processes* is the next higher level. Even though these two levels are in close interaction, the natures of the processes, the time scales, etc., are so different that it is reasonable to study them separately. In what follows, modeling of *gene activation* is studied. As compared to modeling of chemical reactions, the underlying processes are very different; still, it may be that similar-looking models can be applied.

3.1 Transcription factors

Perhaps somebody still assumes that being capable of reading the human genome answers the problems of life ... however, as recognized already by the Nobel laureate Joshua Lederberg, this is far from the truth. The code has to be deciphered — and this deciphering is a dynamic process. Again, one needs to start from the bottom, and study how the genes are *expressed.*

Modeling gene expression is extremely challenging: There are various complex subprocesses involving DNA and different kinds of RNA molecules finally producing proteins and enzymes, and, inversely, these enzymes can affect the gene activation. Between the genetic level and the metabolic level, there are complicated transfer and coding processes, etc. However, just as when modeling neurons, it is reasonable to forget about the low-level phenomena, and concentrate on the information processing level. What is the appropriate level of abstraction then, what are the processes to concentrate on? Again, it is *activation levels* that offer a good starting point; now it is gene activation rather than neural activation that is of interest.

Whether or not a gene is active, is a result of a complicated interplay between excitatory and inhibitory factors. There are regions in chromosomes where specific enzymes called *transcription factors* can activate (or inhibit) the gene expression. An activated gene can further produce other transcription factors; complex sequences, or activation pathways, have been identified. Indeed, there are interactions and feedbacks in these genetic regulatory networks.

Transcription factors seems to be a good starting point for modeling of the gene expression process in a compact way. Assume that the concentrations of transcription factors are collected into a vector x (because of redundancies among genes, a practical approach is to first compress the data by applying PCA, for example). The deeply interconnected dependencies among the variables in x can

best be captured in a state-space model; as recognized in [9], linear models seem to be promising when modeling local behaviors among transcription factors:

$$\frac{d\,\xi}{d\,t}(t) = A\,x(t) + B\,u(t), \tag{14}$$

with

$$x(t) = f_{\text{cut}}(\xi(t)). \tag{15}$$

The outside effects are represented by the input vector u. Positive entries in A denote excitatory effects among transcription factors, and negative denote inhibitory ones. To reach enhanced physical plausibility and expressional power, and to extend the validity range beyond a single local region, a minor nonlinearity is included in the above model:

$$f_{\text{cut,i}}\,(x) = \begin{cases} x_i, & \text{if } x_i > 0 \\ 0, & \text{otherwise.} \end{cases} \tag{16}$$

The underlying processes determining the details of activation level, including chromatine packing, etc., are here abstracted away. The main motivation for selecting such a nonlinearity form is based on the chemical fact that concentrations can never become negative, and activations can never be negative (note that in [13] the coupling of the nonlinearity is slightly streamlined). There are also pragmatic benefits what comes to the function form: This nonlinearity means that the model is piecewise linear. Typically, a state vector in such a system becomes sparse and contains zeros. Sparse structures can also be interpreted as *clusters* in the data space; in this sense, the selected nonlinearity combines a wide variety of different modeling approaches in a unified framework.

From the identification point of view, piecewise linearity offers benefits, because transcription factors with zero entries can locally be ignored, and the remaining problem is piecewise strictly linear. This is important, because often there is not enough data to determine the complete covariance structure among variables, and efficient (linear) model structures are necessary to reach anything practical. Yet another aspect is illustrated in the next section: Because of this nonlinearity, the positive feedback structures result in self-organization rather than instability.

Principal component analysis, or the "more professional" version of it, *singular value decomposition SVD,* has successfully been applied for compressing the available gene expression data: The genes seem to be highly redundant, and PCA/SVD efficiently captures the covariation in the data. However, these methods are linear; nonlinear, "sparse coded principal components" supporting the model (14) can be extracted, for example, as presented in [11].

In the previous section, it turned out that powerful models could be derived when some intensional assumptions were made: It was assumed that the metabolic system tries to reach balance. How about gene expression? Can the (sparse) PCA oriented approaches be motivated in a physiologically plausible way also in this case for gaining information and intuition, or is it just a data crunching technique? This question is discussed in more detail later, in Sec. 4.1.

3.2 Example: Morphogenesis

There is one central paradox in developmental biology: How cells and tissues differentiate when their genetic contents are equal? How the abrupt changes are possible when the concentrations within an organism change in a continuous fashion?

To study this issue, the following two-state model "comparator structure" of the form (14) was simulated with two mutually inhibitory (hypothetical) transcription factors:

$$\begin{pmatrix} \dot{\xi}_1(t) \\ \dot{\xi}_2(t) \end{pmatrix} = \begin{pmatrix} -\gamma_1 & -1 \\ -1 & -\gamma_2 \end{pmatrix} \cdot \begin{pmatrix} x_1(t) \\ x_2(t) \end{pmatrix} + \begin{pmatrix} \gamma_1 & 0 \\ 0 & \gamma_2 \end{pmatrix} \cdot \begin{pmatrix} u_1(t) \\ u_2(t) \end{pmatrix}. \tag{17}$$

This is compatible with (14). The negative non-diagonal elements in A matrix implement negative feedback among transcription factors. In simulations, $\gamma_1 = \gamma_2 = 0.75$; this means that without the nonlinearity the system would become unstable, x_1 and x_2 escaping to infinity, either in positive and the other in negative direction. However, as the nonlinearity prevents variables from escaping in negative direction, it simultaneously stabilizes the positive variable as well.

The simulation results (starting from zero initial values) are shown in Figs. 1 and 2. It seems that in this framework inhibition and excitation together define a system where some concentrations stabilize to non-zero values and other to zeroes ("winner-take-all"), depending on the input concentration: Using the above model, x_1 wins and x_2 vanishes altogether if $u_1 > u_2$, and vice versa, the inputs being constant. It turns out that, qualitatively, the behaviour is rather robust regardless of the exact parameter values.

The presented model structure makes it possible to define a genetic functional "state". Minor changes in input concentrations make the resulting environment within the cell completely different: The "flip-flops" take either of the alternative states depending of the ratio between inputs, and once they have ended in some state, it is difficult to change it (the term "attractor" has been used in some other contexts). In this sense, associations to properties of stem cells are easily made: A cell that has specialized cannot any more take some other role. Other bonus intuitions are also available: Today, there is the link missing between strictly biophysical considerations and qualitative ones. The purely numeric, quantitative, continuous approaches and the qualitative and discontinuous approaches are incompatible (this is the same problem as in AI!). The claim here is that the presented model makes it possible to study *emergence of structures* in the form of sparse coding.

The above example can be studied further: Assume that there is a grid of interacting "flip-flop" cells to be analyzed. In such a environment, the operations of *morphogenes* can be simulated (see [18]). It has been demonstrated (using different kinds of model structures) that competition among contradictory factors affecting coloring can explain the dots and stripes in animal furs [19]. To study the grid of cells rather than an individual cell, the model (14) can be extended as follows:

$$\frac{d\,\xi}{d\,t}(t) = N\,A\,x(t) + N\,B\,u. \tag{18}$$

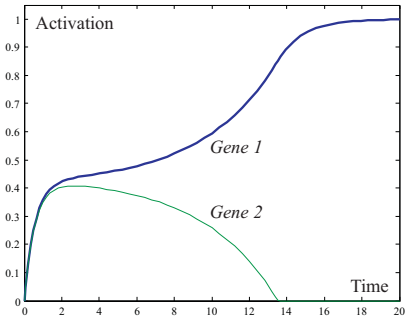

Fig. 1. Incoming concentration ratio $u_1/u_2 = 1.00/0.99$

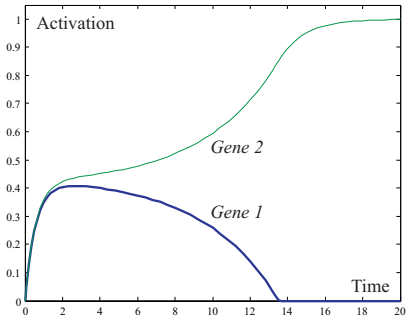

Fig. 2. Incoming concentration ratio $u_1/u_2 = 0.99/1.00$

Simulations of morphogenesis process were carried out applying this model. Vector x contained the states of all cells in the grid; in the simulations the grid had the size 60×60, meaning that the state vector dimension is $2 \cdot 60 \cdot 60 = 7200$. The matrix N represents "spread of activation", that is, diffusion among cells, so that some of the cell activity goes to the neighboring cells. In practice, N implements a set of spatial filters, disretized approximations of the second derivative operators, the diffusion coefficient being d_i and d_2 for morphogenes 1 and 2, respectively.

The simulations revealed (see Fig. 3) that the proposed model structure also can repreduce nontrivial colorings. In the simulations $\gamma_2 = 0.000000001$, u_1 and u_2 having random values in all cells. It is interesting how boundary effects give rise to extra stripes; another theoretically interesting aspect is that the diffusion coefficients are identical in all cases (it has been assumed that the key to emergence of patterns is differing diffusion coefficients; see [19]).

4 Discussions

In the above studies, many ends were left open, and they deserve some additional emphasis. The discussions below are more or less speculative — but intriguing, nonetheless.

4.1 Cybernetic systems

The presented linear and nonlinear model structures above are essentially the same as the model structures in [13], where "simple" cybernetic systems were studied. Indeed, it is intuitively clear that a living system is a cybernetic system; a living cell is perhaps the most characteristic example of cybernetic systems, where local interactions and feedbacks among lowest-level components result in surprisingly expedient behaviors as seen from the level of the complete organism (see [22]).

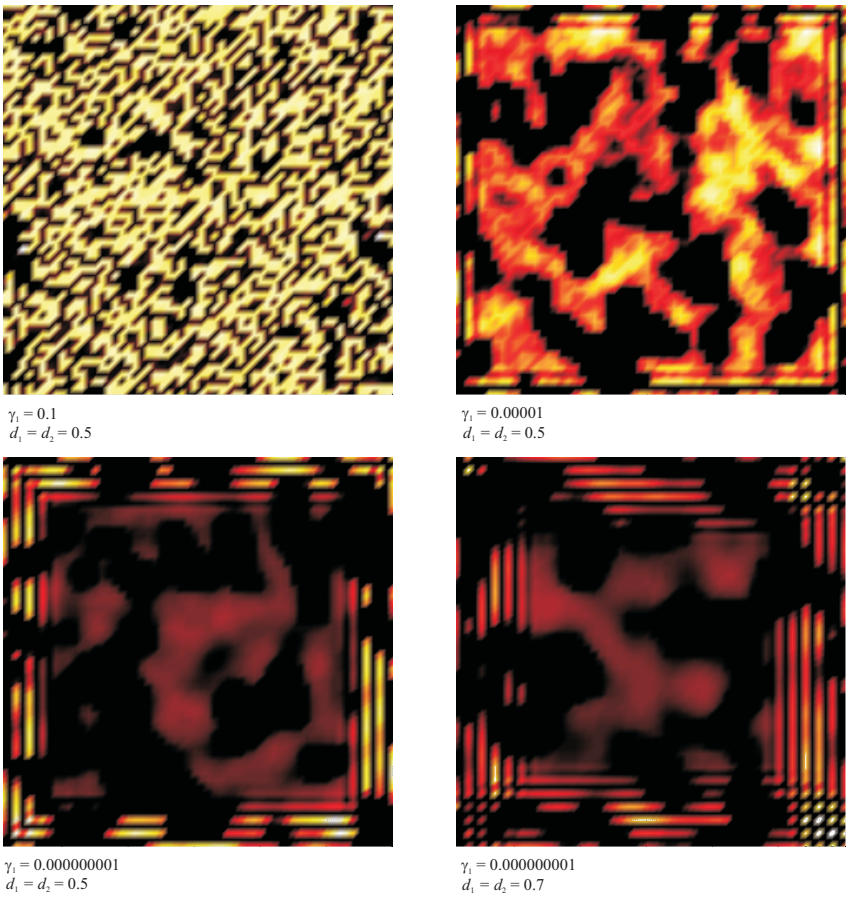

Fig. 3. Colorings of 60×60 cell grids for different parameters (see text)

In a cybernetic system studied in [13], the resulting (nonlinear) structures span the (sparsely coded) principal subspace of the input data. Here, again, PCA was proposed as the mechanism, but its role is to determine the subspace that remains outside the null space, being the first-hand reflection of the underlying chemical processes. Whereas now the system structure is assumed to be "constraints-oriented", in [13] the structures are "freedoms-oriented". The mechanism is also very different — no learning of Hebbian/anti-Hebbian type (see [12]) takes now place as in the "standard" cybernetic system. It is interesting that the same kinds of structures anyway emerge from the local interactions among mindless lowest-level agents even though the physical first principles are completely different. When the principle of local interactions still applies, and as the resulting functionalities seem to be similar, it is not far-fetched to call this metabolic system also cybernetic in the sense of [13].

When studying the genetic system, it needs to be noted that the model (14) was introduced simply because of pragmatic reasons — on the surface level, it seems to capture some of the relevant functionalities in the genetic system. However, this model structure is again familiar from [13] (or, actually, from [12]). Could the genetic system also be described as being a cybernetic system? As noticed above, it would be nice if a higher-level principles governing the system behavior would exist, and this applies also to characterizing the gene behavior.

The gene expression process cannot be characterized in terms of chemical balances in the same way as the metabolic system can: The genetic transcription process is irreversible, just as the mRNA transfer processes are. However, the genetic system is a cybernetic system not in the metabolic chemical sense, but in the original spirit of [13]: The functionalities emerge from competition among agents (even though this "competition" is purely based on statistics). There is scarcity of transcription factors; genes share the available resources. The more there are genes with similar profiles (as measured in terms of transcription factor activation efficiencies), the less activation factors there will be available per gene. Again, if these assumptions hold, ideas of [12] apply, and statistical assumptions about the asymptotic properties of the genetic system can be made. One can start wondering whether it is again some kind of (sparsely coded) principal subspace of the variations in the environment that is spanned by the genes.

It can be assumed that the genetic system is just a more sophisticated implementation of the balance pursuit system with the same goals. Whereas the metabolics system responds to temporal, local, and small changes, the gene expression system responds to spatial, global, and major changes. Spatial here means that the genetic system has to respond to conditions in different locations; of course, there is also temporal gene activity variation within one cell, in the global scale (development from a stem cell into a specialized one), and in the local scale (responding to sporadic changes in the environment). Because of the global range of genetic variables (from "very active" to completely inactive), proportional differences around nominal equilibrium cannot be applied; absolute values need to be used.

The most characteristic feature about a cybernetic system was assumed to be that of *dynamic equilibria.* This balance pursuit assumption applies directly to the chemical level; if it is extended beyond applied beyond that, what are the consequences? This question leads us to yet wider perspectives when starting to study the "total balance" in the system including the cybernetic control.

4.2 About *life*

The explanations of how life once might have emerged on Earth contain so many improbable processes that assuming them all having taken place in succession sounds too marvellous. Applying the balance idea in each phase separately, the problems seem to become solved one by one.

What is interesting in the above studies is that local interactions and feedback on the completely dummy molecular level is enough to implement designed-looking functionalities in a metabolic system, as presented in Sec. 2.3. A key

observation here is that *function* is more relevant than *structure* — and functions are supplied by the chemical equilibrium reactions. No smart agents or "Maxwell Demons", not to speak of centralized control, are needed to implement "ambient intelligence" — everything happens according to the laws of thermodynamics. For example, assume that some chemical level — for example, glucose providing energy — is getting low, or temperature is increasing, or anything like that is happening: The system tries to balance itself, going towards a state where the anomaly becomes compensated according to *Le Chatelier* principle that is familiar from chemistry. On the other hand, if there are excessive resources, metabolic processes are stimulated. Seen from the anthropocentric point of view, when chemicals are given semantical roles ("nutrient", "waste product", etc.), the equilibrium system can look astonishingly clever, in terms of functioning in a seemingly goal-oriented way, and surviving in its environment, and exploiting it.

It seems that something familiar is emerging here. It is just as with intelligence: Also *life* is an emergent phenomenon. You cannot define it, but you can recognize it when you see it. There are some common features that apply to all living systems (see [4]): They can *exploit* the environment and *adapt* according to it; there needs to be some kind of *reproduction* and *modification mechanisms* available in a living system, etc. However the life processes are defined, there always exist counterexamples, and the final "liveness" rate is intuitive. In a living system, some kind of *organization* and *differentiation* is taking place; there is *robustness* against some phenomena, and *sensitivity* with respect to some others. In some sense life forms seem to be locally optimal — but what is the optimality criterion? How everything is orchestrated so that order rather than chaos emerges, and how this orchestration is implemented with no centralized control?

Now, a concrete definition of life can be proposed. However, before such fundamental questions can be addressed, some more intuition is needed. To reach concrete results, the studies are started from the basic principles, from the chemical equilibrium processes, as discussed in Sec. 2.3. Indeed, when studying life in such atomic pieces, one is also facing the questions about the *origin of life.*

To understand the life processes in their simplest form, it is here assumed that chemicals participate in equilibrium reactions, as presented in Sec. 2.3. To facilitate the emergence of something more interesting, three basic hypotheses concerning the reactions are made:

1. *There is a medium available where interactions can take place.* In the simplest case, this means that there is liquid water for chemical solutions. Using the traditional vocabulary, one can speak of *primordial soup,* where there are chemicals and energy available (for example, see [5]).
2. *There are mechanisms available for keeping chemicals together.* It must be assumed that there are partially isolated globules or *micelles* (like the "bags" in [6], or *capsids* of today's viruses) in the soup, with restricted exchange of chemicals and energy, somehow regenerated by the internal reactions within the globule. This isolation can be implemented either by some gel where the reactions take place, or by producing some kind of membrane (based on

lipids?) for isolating the globule contents from the environment. Or perhaps there are originally only water droplets in oily medium, or between layers of clay?

3. *There are autocatalytic chemicals available in the soup.* Autocatalysts are chemicals that act as catalysts, activating reactions, being produced in the same reactions (or through a chain of successive reactions); for example, *peptides,* extremely simple proteins, can be autocatalytic (see [1]).

It is claimed here that the above assumptions are only needed to construct a vision of how life once could have emerged on Earth — or on any other planet.

Assume that within a globule, a set of equilibrium reactions keeps the concentrations of (a subset of) chemicals constant. A globule is a simple functional entity; to start with, the set of reactions can be very reduced. In the globules, no genetic code is needed to control behaviors; no specialized "globule organelles" are needed; not even any complex molecules like nucleic acids or amino acids are needed in the beginning. As shown below, the liquid mixture of globule-specific chemicals suffices to implement the basic "low-level" life functions: Surviving, growing, multiplying.

Such globules, or "archaeocells", are characterized by the autocatalysts that are found in these chemical reaction vessels, determining their chemical processes (or metabolics, if life-like terminology is employed)[2]. Some reactions just do not take place if there are no necessary catalysts; but if one molecule appears in the globule for some reason, more of it will be produced because of the latent reactions that had that far been hibernating. Because autocatalysts are reproduced, the off-spring inherits the properties (or chemical equilibrium concentrations, and also the characteristic reactions) of its single parent — assuming that the globules are somehow split in parts.

Assume that the globule receives chemicals from the environment; the reactions proceed towards new balance. After large amounts of new chemicals are produced, the concentrations increase; water diffuses into the globule because of osmosis, and the globule volume grows. After becoming physically too large, so that the surface tension cannot sustain the sphericity, the globule can constrict into two parts. This way, the first order balance system already can "multiply, and fill the earth", maintaining internal integrity and order.

In the famous *Miller–Urey experiments* only simple amino acids were produced when a chemical soup was tormented — after that, the arrow of entropy is against evolution of more complex peptides or proteins. Mere addition of energy into the chaos only breaks the more complex molecules in parts. Now, on the other hand, the globules in equilibrium are miniature laboratories where dissipated energy can alter the thermodynamic destiny, offering a good testbench for random modifications.

Autocatalysis has long been hypothesized of being the essence of life. However, even though autocatalysis seems to be the key point in the reproduction of elementary life forms, it is intuitively not enough. Through autocatalysis, some

[2] Remember how the *archaeobacteria* illustrate examples of non-convential metabolics in real world

chemical can multiply — but so what? This kind of explosive behavior resembles *cancer,* where some cells start multiplying with no control, resulting in destruction rather than construction of new structures. And it is structures and the corresponding functions that are the essence in life — and structures can only emerge in essentially stable environments. The "instability" of self-organizing processes can start cumulating new order only if the underlying substrate system is stable enough.

The above scenario of mere reproduction is also too trivial to cover more sophisticated life forms. Why one globule type does not exhaust all of the available resources? First, the autocatalytic reactions are also equilibrium reactions, and the chemicals are never completely exhausted however fast the reaction is. But, on the more fundamental level, this question of balance among candidate globule types is the same question that applies to all systems demonstrating (bio)diversity (see [24]), and this question can be answered now: As presented in [13], the *variations* in the environment result in a situation where there are various globule types simultaneously present in optimum state. In stationary state, each of the interacting types has its characteristic "foraging profile" what comes to resources (in this case, available chemicals). The less relevant types are not simply wiped away. Because of this stability of diversity, there is no hurry: There is enough time to wait for advantageous, a bit more complex autocatalysts to emerge by random processes. One molecule suffices to introduce a new line of descendants.

Why is this increased complexity an evolutionary advantage, why the more complex (and typically slower) reactions will survive in this chemical evolution? It is enough that there is some special chemical that can only be utilized by the new globule type; 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.

It is easy to imagine what can happen next: Different globules or globule groups can start exhausting each other's surplus products, and become *symbiotic.* If the different kinds of globules are dependent of each other, they probably grow and divide at the same rate, following the cybernetic balance [13]; this is a rather plausible route to "multiglobular" systems. Globules without partners starve and become outnumbered. As seen from outside, different globules represent different sets of reactions, so that functional differentiation starts taking place. For example, within a single cell the cell organelles can be seen to still constitute such a sub-cellular symbiotic system (mitochondria, etc.).

Symbiotic systems are so common still today that it can be assumed that there really is a huge leap from symbiosis, or chemical cooperation, to more orchestrated strategies, like genetically controlled explicit cooperation. However, assuming that a "multi-purpose" globule with various behavioral patterns once was introduced, its "programmability" between functionalities would have a clear evolutionary advantage outperforming the non-coordinated symbiotic alternatives in complex multiglobular societies with specialization among globules; this is studied in the next section.

Summarizing, it is clear that first-level (chemical) balance (reacting to disturbances) is not enough to explain life. The *cybernetic second-order balance* is needed to maintain *external, inter-globular integrity and order.* Without this principle, the fastest exploding globule type would win the struggle, being the fittest when using too simple criteria, and only structureless "cancer tissue" would remain.

Now, finally, we are ready to express the new definition of life simply as follows:

System striving towards the *higher-order balance* with its environment.

The environmental conditions vary, variations having statistically more or less stationary distribution; a living system (ultimately) constructs a "balanced" (cybernetic) model of the variations in its environment, exhausting the resources in a greedy way. However, the balance is typically never reached. This cybernetic optimality pursuit extends over the scales: Subsystems follow the same cybernetic rules.

There is a continuum from very simple to very complex life forms. At some state, the globules can be called cells, groups of similar globules constitute tissues, cybernetic groups of tissues are organs. Similarly, cybernetic groups of organs are individuals, and cybernetic groups sharing the same genetic control program are called species.

If life is defined in the above way, it can be assumed that some level of life exists on all planets with stationary enough (but not constant!) conditions. Further, neural signals are also based on concentrations of neurotransmitter molecules: It is not difficult to imagine that intelligent life is just the next step in the inevitable development of life forms.

4.3 Natural and artificial evolution

Whereas the null space in the chemical variable space is determined by the equilibria equations, the metabolic degrees of freedom revealing the *cellular functions,* or the intracellular characteristic activity patterns, the genetic degrees of freedom reveal the intercellular activity distribution among the cells, and the longer-term changes in the cellular functions. The role of evolution is to adapt the genetic structures (or the sparse coded principal subspaces) so that, as a whole, the grid of cells balances the environmental variations in a cybernetically motivated way.

Standard presentations about the origin of life are full of mysteries: For example, how could the first genetic code emerge from the primordial soup? DNA molecules — or any molecules that are capable of reproduction in fact — are much too complex to be constructed by chance. And even if they were once constructed, so what? It is in the processes where life resides, not in the structures — otherwise one could not tell the difference between a dead and a living body. Qualitatively new molecules can pop up in chemical systems because of some random events; even an extremely unprobable unique event can happen.

But such events cannot cumulate very long. This is the weak point in today's theories about the origin of life. The key observation is that such a new structure should have immediately evident evolutional advantage, otherwise it will soon be dissolved because of the laws of thermodynamics. Balance pursuit offers such a feedback mechanism with delayless reward.

To understand evolution it is essential to recognize that there are very different levels of hereditary information — complete genes are *not* needed to start with, no DNA molecules, not even RNA's, or ribozymes. There is a continuum from simple to complex mechanisms available. To study these issues, one first needs to distinguish between two separate things: The ability to reproduce, and the ability to modify cellular metabolics (having only the complex DNA and RNA molecules available as examples, these capabilities seem to be mixed). 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, toggling between inert and active state with respect to some specific reaction). The operating modes of the cell being integrated in the chemicals themselves, the cell functionalities are accordingly changed. When some chemical reaction is either active or inactive in different environments (in different globules, for example), a very simple control scheme is implemented, and structures based on such sparse coding can emerge. Of course, it is practical if the two presented capabilities, reproduction and multifunctionality, are combined in a single autocatalytic molecule.

Sooner or later, when more complex life forms are being created, more and more sophisticated control mechanisms need to emerge, though; first this means amino acids, later nucleic acids, these offering the best available combination of flexibility and expressive power. When proteins are involved either as reaction products or as catalysts (enzymes), the environmental conditions (temperature, acidity, chemical conditions) must not change too much, otherwise the proteins denaturate (coagulate) and become inert. To avoid such catastrophes, the conditions within the cell must not change too much. Indeed, this is another motivation for the balance pursuit idea as defining evolutionary fitness: Balance is the natural precondition for non-degeneration.

Another lesson here is that our way of distinguishing between the control machinery and the controlled machinery is incorrect: It is a deeply coupled interaction between the code (information) and the interpreter (formation) that is taking place in a cybernetic living system; the metabolic and genetic systems should not be studied separately. In a way, one must abandon reductionistic thinking — the system itself is simultaneously the computer and the program code. Or, indeed, this holism should be applied to the whole environment.

Even though *structures* are important in living organisms, it is *functions* that are still more important. And, indeed, functions (or properties) can be inherited without structure. The *Lamarckian* theories have been neglected because it has been claimed that there are no necessary mechanisms to implement such views: It must be all in the genes that can be inherited. 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.

Applying the above vision, how can cell differentiation in ontogeny process be explained? Study an embryo in the endometrium: When the conglomerate of cells is still unordered, all cells experience the same temporal variations in the environmental conditions, developing qualitatively in the same direction. However, as the grid of cells becomes denser, there emerge differences in the spatial variation profiles; when matching themselves according to their surroundings, different sets of genes are activated, and different cells start specializing and developing in different directions, as characterized by the active processes taking place within them. When this differentiation has started, further complexification cannot be avoided: Activation of new processes within cells further change the local environments, resulting in continued specialization among the cell groups.

It is still difficult to understand the diversity among cells: Why are there so many different kinds of cells in an organism (or why are there so many species within an ecosystem)? The balance with the environment could perhaps be reached easier? Perhaps a solution to this question is offered by the nature of chemical and genetic realms: When the cell tries to reach a chemical match with its environment, no exact match can typically be found. One cannot simply affect one chemical; typically, various other chemicals are affected as well, determining a further adaptation task. On the other hand, there may exist various reactions *almost* doing the desired task, causing specific side effects, producing different sets of further chemicals. Different approximations (sets of reactions) are not directly summable, and all of the alternatives can flourish side by side. Each of the alternative solutions results in another tree of further candidate reactions, thus resulting in a spectrum of alternatives. Another issue that deserves attention here is that if the models presented in [13] are employed, the system naturally strives towards diversity: Slightly differing populations seem to differentiate further, because of the repulsive effects caused by similarities.

This discrete nature of the chemical realm needs to be studied closer. In the genetic space there exists no continuity; it has been assumed that Darwinian more or less random mechanisms (mutations and crossover processes) is necessary to implement optimization. However, this discontinuity applies only to genetic code; on the other hand, there is piecewise continuity in gene expression, as explained in Sec. 3. Further, there is continuity and differentiability available in output space (in the space of chemical concentrations) as explained in Sec. 3.1. This means that evolution needs not be completely random — there are

clear gradients visible. Even if derivatives cannot explicitly be written, better solutions of chemical match can be searched for in a consistent manner within a single "genetic state".

Everybody who has tried to optimize something applying random search strategies knows that adaptation typically is extremely slow. Still, according to the prevailing paradigm it is assumed that natural evolution could be based on such random optimization strategies. It is lazy thinking to assume that a few billion years would be enough for the immensely complicated life forms to emerge if using such strategies only. Now, not all adaptation needs to based on random search, because local matching (tuning the gene activations) can be carried out applying some gradient-based strategy. What is more, there is a short-term feedback mechanism available here, meaning that optimization can be implemented much more efficiently. 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.

Genes are modified in a Darwinian process of mutation and crossover; however, the genes are not actually optimized. Putting it boldly: When comparing the efficiency of genetic optimization based on selection and the balance-based matching, one can see a remarkable difference. The evolutionary processes produce the genetic mess, whereas the highly streamlined operation of cell functions result from equilibrium processes. The balance pursuit idea can nicely complement the evolution theories. 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 genetic state, as discussed in Sec. 3, can alter the activity of genes within a cell. From this perspective, it is not strange why there are the *introns* left in the genetic code: They are the archive of alternatives.

In evolution theory, one problem has been how to define the selection criterion appropriately. It turns out that the most natural choice — fastest reproduction — results in the simplest versions of the genetic code to prosper, thus, in fact, resulting in degeneration rather than evolution. It may be that the Darwinian idea of "survival of the fittest" is overrated, being utterly optimistic (for example, see [23]); there are more efficient adaptation mechanisms available, at least what comes to lower levels of life. The role of birth and death are very central in Darwinian evolution theory. Now the system is more important than any individual; life is in the system, and in the population of individuals. As long as the system survives, there is no actual death. Another point is that because the

genes only offer the pool of alternatives, the properties of an organism being mainly determined by the environmental conditions, one specific gene combination does not have such a crucial role. The Dawkinsian "selfish gene" view should be accompanied by the system view: Applying the balance principle, the interactions among actors become much more peacefull and the developments are much more gradual — world is not such a draconian place after all?

Is this match (balance) criterion compatible with the generally accepted optimality criteria? Applying the balance principle, adaptation process can be interpreted as follows: If there is plenty of some specific resource available, or if the supply is open-ended, this resource is exploited maximally, the system tries to get into the "resource direction" as far as possible. This means that, as seen from outside, the system behavior looks like standard optimization. The advantage now is that no explicit scalar optimality criterion needs to be defined in terms of individual resources.

Correspondingly, when applying the evolution principles in technical applications, the same considerations could be taken into account. Specially, when studying *genetic algsorithms,* new intuitions are available: Rather than trying the explicitly optimize a single, scalar but complex function, the problem case should be implemented so that the optimum is characterized by balance. The complexity of the function form is substituted with dimensional complexity. The high dimensionality helps to avoid local minima (so that not so large pools of candidate solutions are needed), and the gradient information becomes better applicable. In the same way, perhaps the environments for artificial life research should also be redesigned to implement environments with "integrated fitness"?

4.4 Unified models

The origins in artificial life research are similar to those of artificial intelligence: The goal is to find strategies for surviving in changing environment, and for exploiting the environment. And as it turns out, the cognitive functions can at least to some extent be also interpreted as emergent functionalities in appropriate cybernetic systems (see [25]). Cybernetic studies offer a higher-level framework where the seemingly very different domain fields can be studied in a unified manner. Indeed, this relationship between AI and AL may turn out to be essential when the processes of the mental machinery are explained: It has been recognized (by the Nobel laureate Eric Kandel) that the long-term memory is based on genes becoming locally either active or inactive; this would explain how some memories can last for a very long time. Understanding the mechanisms that implement genetic flip-flops can also be interesting when explaining cognitive phenomena.

When various cell types are being studied, etc., things soon become too complicated. One would need methodology for capturing the substructures, or locally linear models determined by genetic states, within a compact framework. It turns out that the ideas presented in [14] may be applicable again: That is, the hierarchies among structures can be characterized in terms of fractal AND/OR

blocks. In this way, one could perhaps define a *language for life* for describing the cellular diversity.

It is not only the origins of artificial intelligence and artificial life that are the same — if the cybernetic viewpoint is adopted, it may be that research of AI and AL could be united again. More philosophical intuitions have also been presented: According to Humberto Maturana and Francisco Varela *autopoiesis* — essentially duplicating the ideas of cybernetics — is *necessary and sufficient* to characterize a living system. This identification of cybernetic and living systems results in metaphysical conclusions: For example, Maturana concludes that [16]

> *"Living systems are cognitive systems, and living as a process is a process of cognition. This statement is valid for all organisms, with and without a nervous system."*

This is perhaps an overstatement. Principles may be the same, but, however, from the scientific point of view (being based on different kinds of taxonomies), it is reasonable to distinguish between domain fields — substrate makes a difference. For example, in each of the cybernetic domains below, the underlying mechanisms and agents are very different:

- **Life:** Cybernetic system among chemicals, emergent structures being different kinds of tissues corresponding to metabolic functions;
- **Cognition:** Cybernetic system among signals, emergent structures being chunks corresponding to features of patterns;
- **Economy:** Cybernetic system among humans, emergent structures being companies having special activity profiles; and even
- **Science:** Cybernetic system among memes, emergent structures being theories connecting concepts.

According to the definition of life in Sec. 4.2 the proposed artificial life forms (see [8]), being based on explicit, "non-ubiquitous" pieces of code, are *not* forms of life. On the other hand, societies consisting of individuals, or ecosystems consisting of populations, etc., can be regarded as living entities! Perhaps it is reasonable to speak of life only in cases where there is the chemical medium as substrate. It seems that the chemical realm offers an excellent environment for parallel processing.

When looking at today's cybernetic systems, it is the same mystery facing us as when discussing origin of life. When seeing only the optimized result it is tempting to apply divine explanations. How is it possible that such magnificent strategy once emerged? The problem of life alone, being just one example of (extremely) complex systems, is actually *simpler:* It cannot be absolutely denied that an infinitely improbable event once happened; but, looking at the wealth of cybernetic systems around us, it is evident that this has taken place many times in different kinds of complex systems independently. There must exist some underlying principles beneath — why all natural systems seem to inevitably produce life-like complexity.

4.5 Balance in perspective

As presented above, the idea of (higher-order) balance offers new insights in life processes. This can be contrasted to the views of Ilya Prigogine: He has emphasized the role of instability, system being far from equilibrium, as the key to understanding emergent structures. Indeed, it is commonly thought that the key to understanding complex systems is through non-equilibrium state dynamics (for example, see [15]). The Prigoginian approach cannot explain why some dissipative systems get more and more disordered and some go against the arrow of entropy (compare to ideal mixers and "idea mixers", stirred tanks and mental systems, respectively). These paradoxes can be explained in terms of higher-order balance: Either the system tries to reach it, or it is beyond the Edge of Chaos.

It should be recognized that the balance objective is not as restrictive as it sounds. If the environment remains constant so that no changes in chemical levels trigger further processes, the system still does not die, and the processes do not cease in "heat death": There also exist non-equilibrium reactions that are irreversible. Chemicals escaping the system induce steady leakage in the system state, resulting in processes continuing. The chemicals that are produced in these reactions, even though they can be inert in the original environment, can be active in other environments; this way, the system itself produces instability that needs to be compensated by later reactions. This can be interpreted so that a Prigoginian "dissipative system" is constructed, where energy and matter flow through.

Of course, the role of balance has been recognized before. Such studies on *homeostasis,* for example, have been more or less philosophical. Speaking of philosophy — balance is a very old principle in Eastern medicine, and also in the heretic Western traditions. Officially in Western medicine, one symptom and one cure are studied at a time, ignoring the interactions, and forgetting about the wholeness. There have been no concrete tools for approaching such holistic views, and such views have been (aggressively) ignored by the mainstream scientific community. However, it is not only deficiency diseases, but also the autoimmune diseases (like multiple sclerosis, diabetes, allergy, etc.) that could be studied as being caused by some kind of loss of metabolic balance. In the above cases it is chemical balance, or perhaps symbiotic bacterial balance — in the case of mental illnesses, of course, the balance with the environment is also lost. On the level of neuronal/cognitive interactions, the mental images have no match in real life. It is not only jargon that *good life* means living in balance with oneself and in harmony with the environment!

It has been wondered how the geological, climatological, etc., processes on Earth seem to be well suited for maintaining life here rather than wiping it away. Such speculations are concretized in the *Gaia hypothesis:* The Goddess of Earth has purposefully designed these processes to support life; or, *Earth is a living organism itself.* However, after the unstable processes on Earth were exhausted billions of years ago, only equilibrium processes (as seen in the wide scale) are left here; and life is a subset of such processes. It is life that has adapted, not the

Earth's processes. As explained above, stationary variations in environmental conditions is the reason for life, rather than a risk (indeed, in the spirit of Heinz von Foerster's "order-from-noise" principle). Life has adapted to the variation range in climatological and geological processes; changes in environment force the balance systems swallow chemicals — it is a matter of interpretation whether this forced feeding is called active intake of nourishment. In natural history, periods of placid status quo are followed by turmoil transients, where the system overgoes from the previous equilibrium to another one. Just as is the case with the Darwinian theory, one needs to forget about the purpose-oriented or predestinated explanations and unnecessary mystification. Indeed, after adaptation it is "the best of all possible worlds" for life developed here.

As a modeling principle, balance is an extremely powerful concept. This can be compared to modeling of mechanical systems: The today's modeling approach based on *symmetry* (by Emmy Noether) made it possible to apply the same ideas to very different fields (to quantum mechanics, for example). In a way, the principles of balance and symmetry are closely related: A cybernetic system tries to reflect its environment, or to reproduce an (anti)symmetric image of its environment. However, note that such intuitions are far from the realm of the deeply mathematical theory of actual symmetry groups.

5 Conclusion

When trying to understand the metabolic and genetic systems, dynamic models are an important tool, offering the necessary intuitions. Otherwise, applying the static snapshot-style considerations only, the operation of a living system looks completely incomprehensible: There are no conceptual tools for understanding the operation of local distributed control. Indeed, the seemingly goal-oriented nature of the life processes makes the explanations sound like a very improbable fictional story. Teleological assumptions about some *elan vital* seem to be necessary: The wonders taking place there are just too marvellous to be explained in any other way.

However, it may be that understanding of dynamic equilibria can substitute the active creator. To be able to conceptualize the necessary intuitions, mathematics is needed; indeed, mathematics is the *Lingua Franca* of tomorrow's philosophers. As motivated above, it turns out that in a cybernetic system very complicated patterns of self-stabilization and self-organization can emerge. It can be claimed that earlier the scientifically best explanation (simplest, and containing *least* wild hypotheses) was creatonistic, trusting in God; the cybernetic studies show that *the simplest and most plausible explanation for life is non-divine.*

References

1. J. Alander: The origin of life by selection among peptides. *Proceedings of the Finnish Artificial Intelligence Conference (STeP'04),* Vantaa, Finland (September 2004).

2. A. Basilevsky: *Statistical Factor Analysis and Related Methods.* John Wiley & Sons, New York, NY (1994).
3. L. von Bertalanffy: *General System Theory: Foundations, Development, Applications.* George Braziller, New York, NY (1969, revised edition).
4. F. Capra: *The Web of Life.* Anchor Books, New York, 1996.
5. R. Dawkins: *The Blind Watchmaker.* Penguin Books, London (1991).
6. F.D. Dyson: *Origins of Life.* Cambridge University Press, New York (1985).
7. J.S. Edwards, R. Ramakrishna, C.H. Schilling, and B.O. Palsson: Metabolic Flux Balance Analysis. In S.Y. Lee and E.T. Papoutsakis (eds.) *Metabolic Engineering.* Marcel Decker, pp. 13–57 (1999).
8. C. Emmeche: *The Garden in the Machine. The Emerging Science of Artificial Life.* Princeton University Press (1994).
9. N.S. Holter, A. Maritan, M. Cieplak, N.V. Fedoroff, and J.R. Banavar: Dynamic modeling of gene expression data. *PNAS,* Vol. 98, No. 4, pp. 1693–1698 (February 2001).
10. A. Hyvärinen, J. Karhunen, and E. Oja: *Independent Component Analysis.* John Wiley & Sons, New York, NY (2001).
11. H. Hyötyniemi: HUTCH Model for Information Structuring. In P. Ala-Siuru and S. Kaski (eds.): *Proceedings of Finnish Artificial Intelligence Conference (STeP'02),* Oulu, Finland, pp. 241–255 (December 2002).
12. H. Hyötyniemi: Hebbian and Anti-Hebbian Learning: System Theoretic Approach. Submitted to *Neural Networks* (2004).
13. H. Hyötyniemi: Cybernetics — Towards a Unified Theory? Submitted to the Finnish Artificial Intelligence Conference (STeP'04), Vantaa, Finland (September 2004).
14. H. Hyötyniemi: *Modeling Mixtures of Mixtures.* Helsinki University of Technology, Control Engineering Laboratory (2004).
15. S.A. Kauffman: *At Home at the Universe.* Oxford University Press, New York, NY (1995).
16. H. Maturana and F. Varela: *Autopoiesis and Cognition.* D. Reidel, Dordrecht, Holland (1980).
17. H.H. McAdams and L. Shapiro: Circuit simulation of genetic networks. *Science,* Vol. 269, pp. 651–656 (1995).
18. J.D. Murray: *Mathematical Biology. Part 2: Spatial Models and Biomedical Applications.* Springer, New York (2002, third edition).
19. A. Turing: The Chemical Basis of Morphogenesis. *Philosophical Transactions of The Royal Society: Biological Sciences,* Vol. 237, pp. 37–72 (1952).
20. M.S. Waterman: *Introduction to Computational Biology — Maps, Sequences and Genomes.* Chapman & Hall, London (1995).
21. E. Weitzke and P.J. Ortoleva: Simulating cellular dynamics through a coupled transcription, translation, metabolic model. *Computational Biology and Chemistry,* Vol. 27, pp. 469–480 (2003).
22. N. Wiener: *Cybernetics: Or Control and Communication in the Animal and the Machine.* Wiley, New York, NY (1948).
23. C. Wills and J. Bada: *The Spark of Life: Darwin and the Primeval Soup.* Perseus Publishing, Cambridge, MA (2001).
24. E.O. Wilson: *The Diversity of Life.* Harvard University Press (1992).
25. *Additional material on cybernetics will be available in public domain in near future at* `http://www.control.hut.fi/cybernetics`.