Level 5

Role of Information in *Model-Based Control*

The neocybernetic analyses started from simple, reductionistic studies. As the analyses were extended to wider-scale systems, the focus points changed, and new points of view were employed. However, to reach the truly holistic view, yet other interpretations are needed. No new concepts are needed — it turns out that one only has to exploit familiar concepts in new ways. For example, the term "information" has been used routinely, but only intuitively: This is one of the key concepts that open a completely new perspective towards cybernetic worlds.

Many of the cybernetic intuitions become explicitly quantifiable in the neocybernetic perspective. It turns out that when the powerful tools of *control theory* become available, a beautiful new world becomes visible.

5.1 Another view at emergy

The concept of *emergy* was presented in chapter 3, and it turned out that the evolutionary processes could be formulated in that framework. Emergy, the effect that is interpreted as tension, essentially differs from the concepts of energy or power: It is *deviation from the expected* that is crucial — or *information.*

5.1.1 Information vs. noise

Ross Ashby coined the *Law of Requisite Variety* in 1952:

> The amount of appropriate selection that can be performed is limited by the amount of information available.

This is a deep observation — but very "Heraclitus-style", being left obscure. The concept of information is left vague here, and the consequences remain unclear. However, speaking of information seems to offer just the appropriate connotations. To make it possible to efficiently apply mathematical tools for analysis of information flows, the basic concepts necessarily have to be defined in an accurate manner. So, information in the environment is presented by the data, and this data is coded in real-valued signal vectors. How is information manifested?

One is facing a *reverse engineering problem* here: It is known what the cybernetic system (assumedly) does with the data if acquires, and when employing the new terminology, it is assumed that *information is what information processing in natural systems does.* One has to hope that the intuitive notion of information matches with what a cybernetic system is accomplishing. In chapter 3, it turned out that the weighting matrix in the pattern matching is

$$W = \mathrm{E}\{\Delta u \Delta u^T\}. \tag{5.1}$$

This means that data is weighted by the correlation matrix when evaluating matches among patterns: The neocybernetic system must see *information in variation.* The corresponding models are fundamentally based on correlation matrices — principal subspace analysis is just a way of formally rewriting and redistributing this correlation information. The correlation matrices contain atoms of information, entries $\mathrm{E}\{\bar{x}_i \bar{u}_j\}$ revealing cumulated pairwise (co)variations among variables, or *mutual information.*

The correlations and covariances have traditionally been exploited in modeling — what is new in neocybernetic models? Covariances and variances are simple measures for information, being easily expressed and exploited, and they are the basis of modern identification and minimum-variance approaches in systems engineering. The key observation when comparing cybernetic data processing to traditional identification was studied already in chapter 2: Traditionally, when doing parameter fitting applying maximum likelihood criteria for Gaussian data, the approach is opposite — variation is interpreted as something to be avoided — and the weighting matrix is the *inverse* of (5.1). Variation is interpreted as *disinformation,* or noise.

As Gregory Bateson more or less intuitively puts it [7]: "Information consists of differences that make a difference". It is not whatever variation that is thought to be interesting in cybernetic systems: It is *covariation* among data items that is not sensitive to surface-level phenomena like measurement errors, but reveals the underlying common sources or deep patterns. No matter what is the application domain, this covariation is always assumed to be interesting. The role of the cybernetic machinery is to capture the information in compressed form with minimum number of parameters; the correlation matrices that are constructed are essentially storages of the mutual information among data. When the basics are simple and efficiently implementable, accumulation of the information structures makes emergence possible (see chapters 7 and 9).

Such a mechanistic view of information is, however, somehow incomplete. The concept of information also carries something veiled and mysterious that is related to knowledge and *meaning.* One should not lose the power of intuitions; indeed, the concept of information gives tools to attack the problem of *relevance,* too.

When applying Shannons information theory (or Kolmogorov / Chaitin (algorithmic) information theory), the definition of information is strictly syntactical. There is no domain area semantics involved, and thus extreme universality is reached. However, some paradoxes remain: What you expect, contains no information, and it is noise that has the highest information content. When applying the neocybernetic view of information, semantics (in a narrow, formalized sense) is included in manipulations, making the analyses non-universal — but there is *universality among all cybernetic systems.* The approach is intuitively appealing: What is expected, is the most characteristic to the system, and uncorrelated noise has no relevance whatsoever. Capturing the cybernetic semantics and modeling of knowledge is studied in more detail in chapter 7.

5.1.2 State estimation and control

A cybernetic system is a "mirror" of its environment, optimally capturing the information there is available. This is not merely a metaphor — note that the formulas in chapter 3 can be given very concrete interpretations:

- **Model.** It turns out that the neocybernetic strategy constructs the *best possible* (in the quadratic sense) description of the environment by capturing the information (covariation) in the environmental data in the mathematically optimal principal subspace based latent variables:

$$\bar{x} = \left(\mathrm{E}\left\{\bar{x}\bar{x}^T\right\}\right)^{-1} \mathrm{E}\left\{\bar{x}\Delta u^T\right\} \ \Delta u. \tag{5.2}$$

- **Estimate.** It turns out that the neocybernetic strategy constructs the *best possible* (in the quadratic sense) estimate of the environment state by mapping the lower-dimensional latent variable vector back onto the environment applying the mathematically optimal least-squares regression formula (2.22):

$$\hat{u} = \mathrm{E}\left\{\bar{x}\Delta u^T\right\}^T \left(\mathrm{E}\{\bar{x}\bar{x}^T\}\right)^{-1} \ \bar{x}. \tag{5.3}$$

- **Control.** It turns out that the neocybernetic strategy integrates modeling and estimation to maximally eliminate variation in the environment:

$$\tilde{u} = u - \hat{u} \tag{5.4}$$

Even though the operations are represented here in such compact and centralized form, all operations are strictly local, and the represented net effects are only visible as emergent phenomena; for example, the feedback part is implicit. Implicit feedback makes the mappings more conservative: For example, the estimate between $\bar{x}$ and u is indeed implemented applying the regularized least squares formula (2.20), with the role of the regularization parameter q now inverted. The issue of modeling Δu rather than u directly is studied in Sec. 5.2.1; when q increases, u and Δu approach each other what comes to the n most significant eigenvalues.

The above observations mean that a cybernetic system implements *model-based control* of its environment. In terms of information as defined above, this control

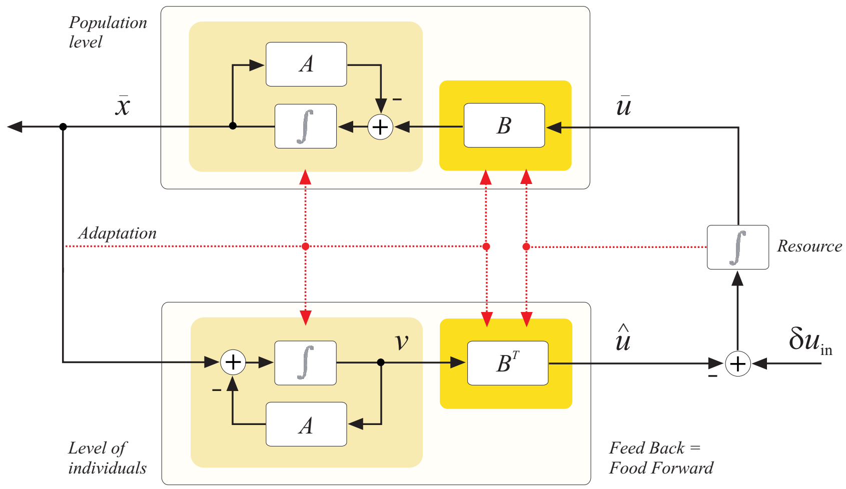

Figure 5.1: Cybernetic system seen through the eyes of a control engineer

is the *best possible.* However, note that the controller is defined as a static structure, control emphasis being shifted from dynamic transients to stationary statistics; the hypothesis here is that however the information acquisition is implemented (for example, as a time-series structure resulting in traditional dynamic control structures; see chapter 7), the cybernetic system maximally compensates that information. The implemented control is far from trivial: It constitutes a multivariate controller where the n most significant variation directions are equalized (or nullified). The symmetric structure of the modeling / estimation loop reminds of Heraclitus' words: "The way up and the way down is the same" (see Fig. 5.1).

In the selected framework, age-old intuitions become concrete. Indeed, the control intuition — cybernetic systems do control — has been clear since Wiener, but the mechanisms have been unclear. Ross Ashby also coined the *Law of Regulatory Models:*

> Regulator must not only have adequate amounts of variety available, but also be or have a *homomorphic representation* of that system.

Since that, the same idea has been known in the field of control engineering as the *internal model control* principle: A controller must contain an (inverse) model of the system to be controlled. Still it needs to be emphasized here that whereas traditional control is always centralized, based on some "master mind", now the control structures are completely distributed: The starting point was local level feedback controls, but the final result is global level feedback control.

Ross Ashby also states that "for appropriate regulation the variety in the regulator must be equal to or greater than the variety in the system" (Ashby's "regulator" being the system, and "system" being the environment). However, here his intuition is *wrong.* The capacity of the cybernetic system must be *less* than that of the environment. If there is no scarcity of resources in the system,

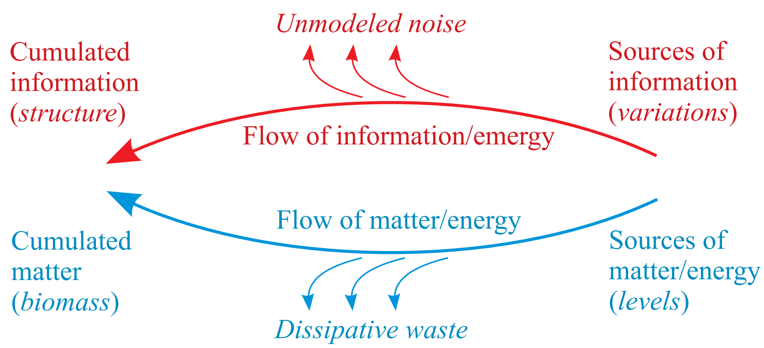

Figure 5.2: Abstract flows in a cybernetic system

no compression — or modeling — needs to take place, and no cybernetic system can emerge. It is the environment that dictates the terms.

5.1.3 Flows of information and matter

Information is also the common denominator capturing the essence in cybernetic systems. Everything that affects the behaviors can be seen as visible (measurable) variation or information; it is information that is being controlled in the environment, and information is being cumulated in the model. Further, information makes different models commeasurable, and information determines the semantics and goals of the system. Yet another viewpoint to the role of information is available here.

The feedback part in the closed-loop structure in Fig. 5.1 is only an abstraction: It does not correspond to a separate real process because it only represents the non-ideality of the information transfer. It is interesting to note that for the closed loop control structure to emerge, two different kinds of processes need to co-operate — first there is the information flow into the model, and then there is the material flow dictated by the model. Without the other flow the other could not exist either. One could say that a cybernetic system constitutes a *marriage mind and matter,* combining these two incompatible dualistic viewpoints (see Fig. 5.2).

In the figure, there are the two flows shown separately: On top, there is the flow of information (or emergy), and on bottom, there is the flow of matter (and energy). Most is wasted — in information flow, the uncorrelated noise becomes filtered, whereas in material flow, it is the dissipative losses that do not get through into the higher-level system. Note that it is often assumed that it is these dissipative material flows that are the manifestation of complex system dynamics [64] — now these are just a side effect. It is the information in the environment (or *variations* in the data) that dictates the structures within the higher-level system, whereas it is the matter (or actual *levels* in the data) that cumulate as some kind of biomass within this predestinated structure of some kind of populations. Whereas the traditional matter and energy oriented views emphasize the level of dissipation, levels of flows being the most essential, in the neocybernetic information oriented perspective constant flows are seen as trivial and not interesting from the point of view of emergent structures.

One could even say that the cybernetic model to some extent captures the Platonian *ideal* beyond the changing world.

5.1.4 Different views at the environment

Here, an example of what are the benefits of applying concrete definitions for concepts is presented. And, again, it is visualized how the fact that real systems are not ideal brings sophistication in the discussions; things do not necessarily become more complex, but new nuances are introduced in the models, and deeper understanding can be reached.

It is assumed that in a long run an *evolutionarily surviving system exploits all information it can see:* Being capable of efficiently exploiting the resources is a prerequisite of surviving in an environment, successful systems are the most active in acquiring for more and more information. This optimality assumption makes behaviors in an environment more or less unique and predictable. When modeling such systems, the optimization task is somewhat trivial, when constraints are given. The interesting challenge is to understand the different mechanisms for information acquisition; why there can still exist different kinds of systems in the same environment, can be studied by assuming that there are different kinds of constraints in the information capture process, and different systems see the environment in different ways. Here, a special aspect is concentrated on: *There can be differences in how systems remember their experiences.* Within the introduced framework these issues have a compact "vocabulary" (distribution of information is further elaborated on in chapter 6).

This far, the expectation operator has been employed in a sloppy way: Indeed, expectation is a mathematical abstraction that cannot be measured, it can only be estimated using the measurement samples. Accurate determination of expectation would necessitate an *infinite* number of samples — this is clearly impossible at least in the changing environments. Instead of employing the mathematically accurate definition, define the "expectation estimate" be an (exponentially) weighted average over the past observations:

$$\frac{d\hat{\mathrm{E}}\{\bar{x}_{\mathrm{s}} u_{\mathrm{s}}^T\}}{dt} = -\gamma_{\mathrm{s}} \hat{\mathrm{E}}\{\bar{x}_{\mathrm{s}} u_{\mathrm{s}}^T\} + \gamma_{\mathrm{s}} \bar{x}_{\mathrm{s}} u_{\mathrm{s}}^T. \tag{5.5}$$

Now, there is an exponential "forgetting horizon" what comes to the covariance estimates: Newest observations are best remembered, whereas old experiences fade away with time. In the similar manner, assume that there is inertia and forgetting taking place in all data processing in the system, so that also the incoming data is seen through such filter:

$$\frac{du_{\mathrm{s}}}{dt} = -\mu_{\mathrm{s}} u_{\mathrm{s}} + \mu_{\mathrm{s}} u_{\mathrm{in}}, \tag{5.6}$$

Here, u_{in} is the original input supplied by the environment, and u_{s} is the filtered input actually seen by the system; the parameters $\mu_{\mathrm{s}} > 0$ and $\lambda_{\mathrm{s}} > 0$ are the filtering coefficients, higher values meaning fast forgetting. This extension makes it possible to take variation structure in time domain into account.

Such linear time domain filtering can most efficiently be represented and analyzed in *frequency domain.* It turns out that information can directly be analyzed in terms of *power spectra.*

To illustrate this, observe that for the Laplace-domain signals $\bar{X}$ and $\bar{U}$, one can express the filtering of signals as $\bar{X} = F\bar{U}$, where the *transfer function* for the first-order filter (5.6) as

$$F(s) = \frac{\mu}{s+\mu} U_{\text{in}}(s), \tag{5.7}$$

and, further, the power spectrum of this becomes

$$H(\omega) = \frac{\mu^2}{\omega^2+\mu^2} H_{\text{in}}(\omega). \tag{5.8}$$

This reveals that the transfer from input power (information) to the power that is actually experienced by the system is a function of angular frequency ω. For low frequencies, $H(\omega) = H_{\text{in}}(\omega)$, but beyond the cut-off frequency μ_{s}, the experienced power decays linearly when studied on the log/log scale.

The filtering effects are visualized in Fig. 5.3 — there it is shown how the information content of a signal can reside in different frequency regions. Frequencies above the cut-off frequency μ_{s} are seen as noise by the system, and gets ignored altogether. Frequencies below that are seen, but assuming that $\mu_{\text{s}} > \gamma_{\text{s}}$, they do not get cumulated in the system's structures — these frequencies are only filtered, or "manipulated" by the cybernetic system. Only variation in the darkest area in the figure becomes cumulated in the model (or in the covariance matrices). Too high frequencies are invisible altogether to the current system, leaving there room for other systems to flourish; but also in the lower frequency range ("environment"), there is competition; even though such signals are visible to the system, there exist probably more customized systems eliminating that variation. The net effect is that the system concentrates on band-limited signals only, signals in other frequency ranges being interpreted either as noise or as constant values — both containing zero information in the cybernetic perspective. The observation from chapter 4 (the behavior of the nominal state, and deviations around it can be modeled by separate systems) can thus be extended and made better quantifiable.

Such differentiation among systems, makes them mutually dependent. Specially, if the lower-range model changes — as it necessarily does in practice when time goes on and the slow phenomena become better visible — the higher-range systems need to adapt to this changing environment; and the needed adaptations can be rather abrupt. Discontinuous changes in the environment are magnified in the subsequent systems.

5.1.5 Cascades of trophic layers

Information is the "nourishment" for systems. It does not matter if the driving force is *loss* of some resource (as when allocating staff labor) or surplus: Posi-

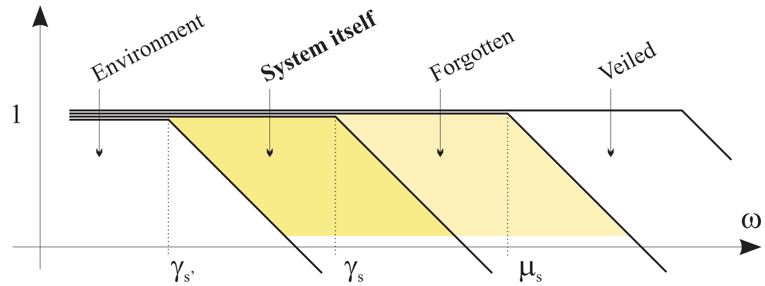

Figure 5.3: Different systems operate on different time scales

tive or negative, the effects are the same. A cybernetic system sees information (emergy) as resources available in the environment, and there is hunger for this information. Again, this sounds teleological — but if some system applies this strategy by accident, it immediately has evolutionary benefit in terms of increasing resources. There is no guiding hand needed — but it is like with Gaia: Even though all behaviors can be reduced to lower levels, simplest models are found if stronger emergent-level assumptions are applied. It turns out that this eternal hunger for information has resulted in very ingenious-looking solutions for reaching more and more information, and, to achieve the necessary sophistication, the systems have typically become ever more complicated. The issues of such information pursuit are studied more in chapter 6.

The systems are hungry, but they are not greedy. Whereas a system exhausts variation in its environment, there is the same variation inherited in the system itself (remember that PCA model maximally relays variation to its latent variables). This gives rise to a *cascade* of trophic layers: Another system can start exploiting the variation that is now visible in the system (being part of the environment as seen by the other systems). When the next trophic layer has been established, there is room for a yet higher trophic layer, etc.

In nature, the basis for all life is the Sun. However, the "non-informative" sunlight alone is *not enough* for cybernetic systems to make them flourish — or, indeed, it is not enough to make them emerge in the first place. Additionally, there are first the physical processes (planets orbiting and rotating) generating more or less cyclic variation in the physical variables, causing temperature gradients. These give rise to second-level chaotic processes: When there are temperature gradients, it is the highly nonlinear Navier-Stokes type equations that produce increasing amounts in randomness in the variables, as being manifested in climatological phenomena, etc. Now, the arena is free for cybernetic systems to start exploiting this non-trivial information; after the information already is there, linear processes are enough to utilize it. The input variables for the lowest-level cybernetic systems (plants) are temperatures, nutrients in the soil, rainfall, etc. On the level of herbivores, it is then the spectrum of plants to forage on, and after that are the carnivores foraging on each other. All loose information seems to give rise to new systems, and, in a way, this can be described as "panspermia". As the number of species increases, the complexity also increases, as the subsystems become more and more interlinked: There emerge pests and diseases to exploit the variety, too. It is only natural that at some stage the lower level species adapt to utilize the higher-level biomass

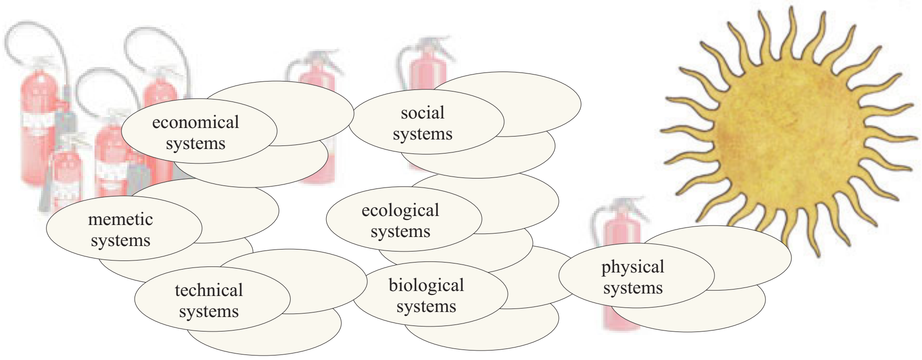

Figure 5.4: Systems in different phenospheres try to extinguish the fire (Heraclitus' *Logos*)

— without recirculation the (dead) biomass would cumulate indefinitely in the resource vector. This makes it a cycle, and finally the natural circulation is established as a consequence of locally controlled information exploitation.

When the succession of systems evolves, the highest-level systems can appear in very different phenospheres. Above the biological systems, there are all the man-made constructivistic systems — but they still live, after all, on the variety resources of the nature: For example, take the scientific systems. Without the simpler cybernetic systems there would be no natural sciences, and without more complex cybernetic systems, there would be no social sciences; without uneven distribution of nature's structures there would be no need for explanations. What science explores, technology exploits — environment being exhausted as a result of such loop. All systems finally try to exploit (or *eliminate* when seen from another point of view) the Sun's fire, either directly or indirectly[1]. Indeed, sun-worship is among the oldest rites. And Heraclitus said that the underlying principle in nature is *fire*. However, in the cybernetic perspective, this is not the key point: It rather seems that the goals of nature could best be explained in terms of a *fire extinguisher*. (see Fig. 5.4).

When the internal inertia in the cybernetic systems is taken into account, one can think of the information transfer between subsystems as some kind of a potential flow from trophic layer to another. There is a "structured leakage" in the information reservoirs; this can also be characterized as "directed diffusion". The subsystems are like (generalized) "ideal mixers" — mixing information (note that the flows are not scalar variables but vectors). As linear systems, the cybernetic mixers can be grouped in different ways; the subsystems seem to be tightly connected and they always define a network, however they are regrouped. When more and more layers are introduced, the ecosystem becomes more and more continuous and smooth from the perspective of information distribution – becoming a lumped parameter approximation of a parabolic partial differential equation (PDE) diffusion model. The evolutionary process of sophistication continues until there are incompletely exploited reservoirs of resources available.

[1] Or, actually, *primus motor* is the fire from the Big Bang: The geological conglomerations and variations in soil properties that also have to be seen as cybernetic resources are not caused by the Sun

Finally the "landscape" should become smooth with no sudden drops, no matter how the intermediate levels are constructed. Changes in resources get filtered when they spread among the systems.

When looking at the wealth of systems that exist to implement the extinction of fire, one cannot help thinking that *the right hand does not know what the left is doing.* It is not about an "intelligent designer"; one could speak of a "hardworking blunderer" instead[2]. The philosophical question is not where the diversity comes from, but why there is *something* instead of *nothing.*

5.2 Control intuitions

Even though truly complex systems cannot be easily quantified, they must share the basic principles: If a system is to remain consistent, there has to exist the balance of tensions deep inside. Qualitatively, identical intuitions apply. When the control notions are employed, it turns out that there are many intuitions directly available for analysis of the behaviors in cybernetic systems — and *vice versa.*

5.2.1 Rise and fall of adaptive control

Adaptation is the key property in truly cybernetic systems, meaning that they are *adaptive control systems,* trying to implement more efficient controls based on simultaneous observations of their environments [3]. If one has control engineering background, one can immediately understand what happens in a truly cybernetic system then: Adaptive controllers are notorious in control engineering, as they can behave in pathological ways. The reason for the "explosions" is *loss of excitation.* Good control eliminates variation in data — and after this there is no information where the model tuning can be based on, and gradually the model becomes corrupted. After that, when the model is no more accurate, the variation cannot all be eliminated, and the control performance can be very poor. But as the control fails, the variation cannot any more be suppressed, and there will exist information in observations once again. The model starts getting better, and after that the control gets better, and the cycle of good and bad closed-loop behavior starts again. This kind of oscillatory behavior is typical in loops of simultaneous model identification and model-based control. This result is paradoxical: Pursuing good balance on the lower level results in high-level instability.

Is it reasonable to compare complex cybernetic systems to simple controllers? This question is motivated as the processes in real life systems are so much more delicate — but still there is some resemblance in the emergent behaviors. Compare to ancient empires: It seems to be so that there is a life-span for all cultures, after which even the strongest civilization collapses. Why is that? For example, during "Pax Romana", there were no enemies, and the once famous Roman army became ruined, morally and otherwise – and then there was a

[2]It would take a truly "intelligent" agent to streamline the natural systems. God forbid that there should be such re-design efforts ...

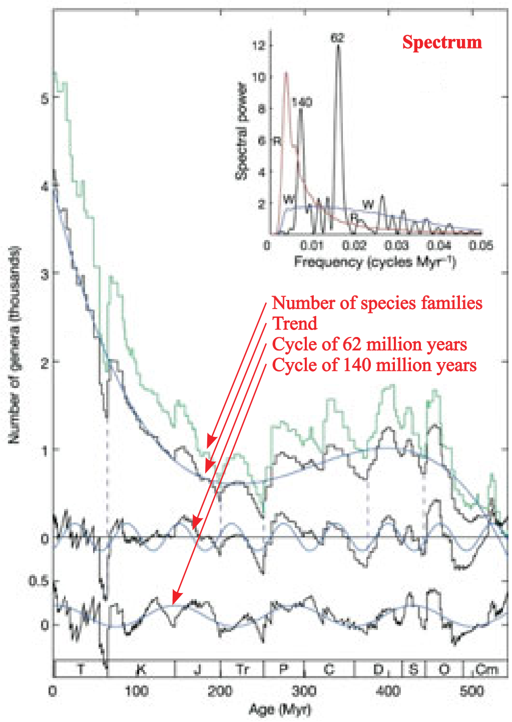

Figure 5.5: For some reason, the ecosystem has periodically become turbulent (diagram adopted from [65]). Note that dinosaurs became extinct about 62 million years ago

collapse after severe disturbances[3]. And this increase of sensitivity does not only apply to human societies (see Fig. 5.5): For some reason, massive extinctions seem to take place in 62 million year cycles [65]. Do you need some meteors to explain extinctions — or is this simply because of evolution dynamics? It seems that current explanations to collapses in general prefer simple solutions (see [22]).

Extreme optimization in some respect results in worsened fitness in changing conditions, and a collapse of the highly specialized subsystem (or the whole ecosystem) is possible. Of course, nature has developed mechanisms to cope with this challenge. For example, in natural systems, there are multiple local minima simultaneously represented. Different species are optimized with respect to their local view of the environment, and as such a pool of structural alternatives is maintained, not the whole system needs to collapse when the environment changes as suitable candidates also exist.

To reach smoother behaviors, there exist other alternatives in addition to the *multiple model approach,* and, again, the technological experience can be exploited here. In control engineering, techniques have been developed to tackle with the adaptive systems: One of the basic techniques is to *add noise* to introduce fresh information in the closed-loop system, preventing the control from becoming too good. A more sophisticated technique can be seen as an extension of this: The controls are designed to artificially make the system roam through

[3]But explicit emphasis on the army results in the Soviet-type collapse: If there is no real need at some time, such investments are cybernetically non-optimal, meaning that the system cannot outperform its competitors in other fields in the evolutionary struggle

the admissible region, thus exciting the modes, and mapping the responses and latent dynamics. For example, in complex industrial plants such control strategies are commonplace, reagents being added until some specific criteria are reached, and after that reagents being reduced until some other criteria are reached. Of course, this results in oscillation (limit cycles) in the closed-loop system, and thus in variability in product properties — but, regardless of its limitations, such cycles are employed also in real natural systems, caused by, for example, the *cell cycle* in cultivations. Formally, a well-behaving system is seemingly permanently on its stability limit.

It seems to be always so that the optimality goal has to be relaxed to reach *good* behavior. The above solutions — messing the control up with more or less stochastic or deterministic noise — add the element of randomness and unpredictability in the system as seen from outside. However, there seems to exist yet another elegant technique that is inherently applied by the natural cybernetic systems. The most important ingredient here is again trivial, caused by the nonideality of nature: It is the *stupidity of agents* that facilitates the emergence of sustainable systems.

5.2.2 Paradox of intelligence

As compared to traditional adaptive controllers, the cybernetic strategy where the feedback is implemented implicitly through the environment, results in "gentle" adaptive control, form of *buffering,* where the variation is not fully eliminated, and the closed loop behavior does not become pathological: There will always remain enough excitation in the signals. One could also speak of *passive control* as only *attenuation* of signals takes place; how near complete elimination of excitation one goes, is determined by the coupling factors q_i. This is because it is Δu rather than the estimate u itself that is being eliminated from the input data, making the overall system evolutionarily stable and sustainable. But such control, leaving some of the input uncompensated, is technically not optimal — and cybernetic systems always pursue better controls ...

Indeed, getting too ambitious, implementing extreme optimization, and full exploiting the information completely wiping out excitation, is also a possible scenario in a cybernetic system — if the system is sophisticated enough. This kind of invasive, fully compensating control can take place if the agents realizing the control are "too smart", implementing the feedbacks explicitly, actively, rather than waiting for the environmental reactions.

To implement such extreme optimization, the different signals have different roles as seen by the agents: The inputs and outputs need to be functionally separated from each other, meaning that the system necessarily has more sophisticated, predetermined structure, as seen from outside. When the competition among agents is explicitly taken into account, one can start the modeling from (3.4) and write

$$\frac{d\,x}{d\,t}(t) = -\Gamma A\, x(t) + \Gamma B\, u(t). \tag{5.9}$$

Here, the gradient expression is extended by taking into account that the diagonal Γ makes it possible for agents to have differing adaptation speeds. Now, when defining

$$A = \Gamma \mathrm{E}\{\bar{x}\bar{x}^T\}, \qquad \text{and} \qquad B = \Gamma \mathrm{E}\{\bar{x}u^T\}, \tag{5.10}$$

one changes the original feedback structure in chapter 3 only minimally. Essentially all signals are handled identically, and weight adaptation is identical for all signals — but there is a twist: If a signal is known to be recirculated, if it belongs to the x variables, its value is additionally multiplied by -1, as shown in (5.9). This is what it takes to actively implement the negative feedback: The agents only need to distinguish between "positive" and "negative" inputs, or information about resources and competitors, respectively. Implementation of the explicit feedback in this way results in combined Hebbian/anti-Hebbian learning (see [92]). The matrix A now defines the communication (or, at least information transfer) among the agents. In large systems, the size of this matrix (having n^2 elements for an n-agent system) can become considerable necessitating structured coordination of signal transfer. In any case, if u varies slowly, the steady state for x is defined through the mapping matrix

$$\phi^T = \mathrm{E}\{\bar{x}\bar{x}^T\}^{-1}\mathrm{E}\{\bar{x}u^T\} \tag{5.11}$$

so that $\bar{x} = \phi^T u$. From discussions in chapter 3, when Δu is now everywhere substituted with u, it is clear that the columns in ϕ span the principal subspace of u, and PSA is implemented explicitly for u. Remember that as the feedback in the "smart" structure is implicit, all signal manipulations taking place within the system, the input data is not disturbed. In this sense, the signal transfer is idealized, information theoretic, assuming that observation can be implemented without exhaustion of the signal source. Also in this sense, the smart agents assumedly operate on a higher abstraction level, not being bound to their immediate surroundings. The disadvantage is that as the input signal is not touched, no control is automatically implemented. In the model-based controller structure in Sec. 5.1.2 two items are also changed:

- The model becomes

$$\bar{x} = \left(\mathrm{E}\left\{\bar{x}\bar{x}^T\right\}\right)^{-1} \mathrm{E}\left\{\bar{x}u^T\right\}\, u. \tag{5.12}$$

- The estimate becomes

$$\hat{u} = \mathrm{E}\left\{\bar{x}u^T\right\}^T \left(\mathrm{E}\{\bar{x}\bar{x}^T\}\right)^{-1}\, \bar{x}. \tag{5.13}$$

However, cybernetic systems are for control purposes — so, if the feedback structured are separately hardwired, applying the "smart" model for explicit control, all available variation in u is exhausted. This results in all the familiar problems of traditional adaptive control. When you can optimize, you typically do it, even though optimal is the enemy of good in the sense of robustness and sustainability: "It is hard to be humble when you are so strong"!

But there are also benefits when feedbacks are optimized — the system can truly be smart, and there is evolutionary advantage. Unnecessary competition can be avoided, resources can be allocated by negotiation (more or less democratically), and the agents can concentrate on more productive issues. As a consequence, a

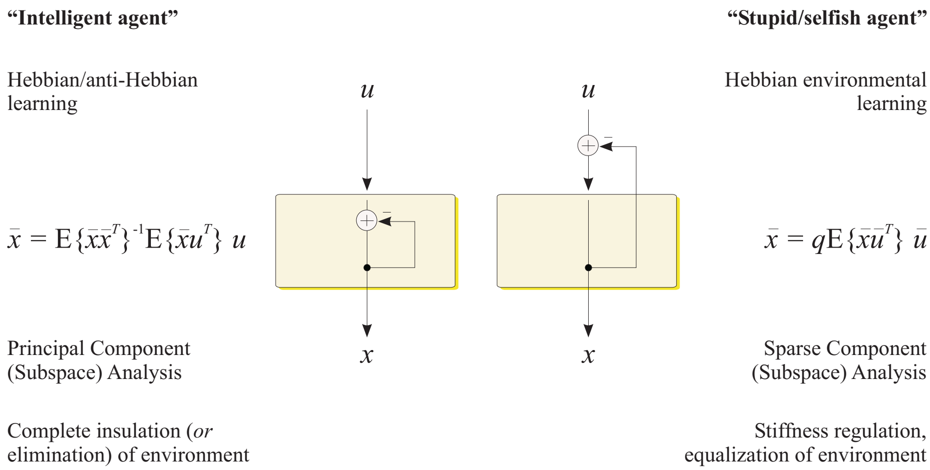

Figure 5.6: Two learning strategies, two ways to see the world and change it. In a system of "intelligent" agents, the interactions among the competing actors are taken explicitly into account, being integrated in the adaptation strategy, whereas in a system of "stupid" agents, adaptation takes place in the direction of visible resources, the interactions becoming evident only implicitly through the exhaustion of the environment (details of differences in input coding are presented in chapter 6)

welfare state need not necessarily be less efficient than a pure capitalist economy — assuming that the model of the (changing) environment (legislation, etc.) remains up-to-date. The two types of feedback implementation strategies are illustrated in Fig. 5.6.

5.2.3 Contribution in inverse direction?

It is not only so that control intuitions would be applicable in analysis of cybernetic systems — there is contribution in the inverse direction, too. It may be that the locally adapting controller schemes could make it possible to implement controls that cannot have been imagined this far. The applications can range from sensor fusion to agent controls and complex networks in general. What is more, the cybernetic systems of humans, the process operators, can perhaps be integrated in the cybernetic models of the processes — issues of "human factors" can perhaps be addressed fluently in the same modeling framework.

Today's main challenge in control engineering is understanding complex automation systems: How emergent properties like *robustness* could be seen from designs, how to find analysis and synthesis methods to address qualitative plant properties?

An industrial plant is "first-level cybernetic" because there are controls implemented so that it can sustain environmental disturbances and it (hopefully) finds a new balance if the conditions change; the industrial process can be seen as an "artificial cell" with its own metabolism, "eating" the raw materials and giving out the products. Applying intuitions concerning natural cells and their robustness, one would like to extend from the first-level to second-level of cybernetics, so that higher-order statistical balance between the system and its environment would be reached, including constant stiffness against disturbances. How to implement "evolutionary adaptation", human acting as the "agent of evolution", then?

Neocybernetic adaptation principles are simple, in principle, and can readily be implemented also in real systems. There are relations to traditional control approaches: Applying the cybernetic view of semantics (together with the "snapshot", its derivative is needed among the measurement data; see chapter 7) it turns out that *multivariate PD controls* can be implemented; there are also connections to *internal model control*. High dimensionality and noise could assumedly be tackled with in unstructured environments ... This sounds like a panacea, and such general solutions probably never exist. Perhaps one should look at the cybernetic models more like methods towards implementing sophisticated data mining and process monitoring, perhaps better matching and supporting the mental views of human experts than what the traditional statistical tools can do (see discussions in chapter 7). The automated "human-like" preprocessing of the huge bodies of the measurement data and historical time series, finding relevant correlation structures among signals, makes it possible for the human expert to explore and perhaps exploit the available information more efficiently.

When extending the idealized cybernetic studies to practical controls, there are many challenges. The key problem is that when trying to impose the cybernetic principles afterwards on top of the existing automation system, where the structures already have been differentiated hierarchically, and when there are predetermined information blockages within those structures (see chapter 6), one somehow has to "bootstrap" the cybernetic machinery. For example, the following issues become acute:

- In industrial plants, there are predetermined goals of the system what comes to the products and their quality. This does not match the self-organization idea, where the system adapts to match its environment; thus, the adaptation process needs to be somehow controlled.
- Related to that, the agents (controllers in the plant) are typically not homogeneous and identical, what has been assumed this far. In real plants, controllers are in different locations, and they are tuned to implement only their specific control tasks — the SISO approach should be

  extended into a MIMO.

In addition to the theoretical aspects, there are also more pragmatic ones: In practical control, there is need of speed. The control quality is measured in terms of real-time reactions, and there is no time to wait until the statistical balance:

- The basic problem in dynamic control is that the time structure cannot be ripped off, it is the signal transients that are to be controlled. The controllers should not only be a simple mirror of the environment; they should be mirrors *between the past and the future.*

- Related to the previous item, there are the causality issues — in a real system, pancausality cannot be assumed: For example, the past measurements cannot be altered by later-time feedbacks. To implement feedback through the environment, external structures are needed (see chapter 7).

One approach is presented in chapter 7, where the ideas of *biomimetic control* are discussed. it turns out that such approaches can be studied in the framework of *model predictive control,* where the model-based estimation of the future is tried to be regulated by applying appropriate actions in current time.

In any case, to implement the cybernetic adaptation, the system must be stable to begin with. The independence of the controls is an advantage, but it is also a disadvantage: *Stability of the overall system cannot be assured* during adaptation. This discourages all practicing engineers — before cybernetic control can become reality, further studies are necessary.

It needs to be recognized that control theory is not in all respects an appropriate framework to understand cybernetics as there are many practices that are in contrast with cybernetic intuitions. Indeed, control is seen as a stereotype of reductionistic engineering-like thinking, systems being localized and divided in separate blocks, and within them control being centralized. One should never underestimate the inertia that is caused by the role of practicing automation engineers and plant operators not willing to alter their practices. The plant-floor level constitutes yet another cybernetic (memetic) system with new sets of tensions. One can expect some of the counterarguments to be rather fierce: For example, a practicing engineer does not want to compromise the plant stability at any cost (there is a big difference here between the engineers and economists who are familiar with risks and complex environments — see next section).

5.3 Towards wider views

The presented ideas of information-oriented control-based perspective are so simple that some comments can be said in general also about truly complex systems without the knowledge of the details of the systems or their numeric parameters. It seems that the most complex of systems, the memetic ones, also share the behaviors that can be motivated more convincingly in better quantifiable environments.

5.3.1 "System cybernetization"

There are two ways to implement enhanced controls in a cybernetic system: Either the controls can be made more accurate, or the controls can be made faster. These objectives can be reached not only through making the model ever better, but specially by implementing tighter coupling. In a cybernetic system, extreme optimization results in "stiffness" of the system, and worsened fitness in changing conditions (see next section).

There are many details in the control structure that can be manipulated to enhance the control — in complex cybernetic systems, the model adaptations can be more complex than in typical adaptive controls, as the system structure also can change; the developments can even take place in separate systems, and in different phenospheres. The structure changes are again related to processing of the critical substance, information: Either enhanced capture, transfer, or usage of this information. When speaking of memetic systems themselves processing information, the critical resource is actually *knowledge,* or "knowhow" about clever usage of the available information. An intelligent agent constructing such a system is always "at the edge of understanding". For example, constructivistic systems (technical, scientific, ...) evolve so that as soon as there is some new understanding about relationships among variables, it is exploited to increase system performance (if there are no compensating drifts, like cost, etc.). This becomes manifested in industrial plants, for example, where new controls are introduced to compensate deviations from the reference values if some new relevant measurements are available, thus making the system remain better in balance — and become more cybernetic. Otherwise there is assumedly evolutionary disadvantage, as the system is "less cybernetic" than it could be. These developments are implemented by humans, but, after all, the system follows its own evolution where individual human signal carriers have little to say.

The cybernetization developments have to be gradual, as the world changes in unpredictable ways as changes in the structures are employed. A clever balance of opposing needs (tensions) cannot easily be determined by a centralized mastermind — if some specific aspect is omitted, all vacuums will be filled somehow through unintended developments. Also the development efforts must be cybernetically balanced. Perhaps the best example is the downfall of the late Soviet Union, where the goal assumedly was to reach a better society — by applying the cybernetic governmental steering following the best theories of centralized control. However, the means and ends were not in balance as they were centrally controlled. Again, the main problems in Soviet can be characterized in terms of information extraction and exploitation: In data input, there were problems as the statistics were forged and not accurate; information was available too seldom in the five-year plan frameworks; information transfer (specially in the low level) was defective because of censorship and scarcity of communication devices; and, finally, when the controls were applied, they could not be enforced because of decline in moral standards — this decline also being caused by ignoring the sophisticated cybernetic balances in social and ethical systems.

So, complex systems seem to develop autonomously towards becoming more and more cybernetic, as being led by a guiding hand (see chapter 9). Regardless of the domain, the limiting factor in this evolutionary process seems to be related to extracting and exploiting information (or knowledge). Typical examples are found in todays working life. First, study the other prerequisite for "cybernetization " — better understanding of the system and gaining more information. This is implemented through supervision, questionnaires, and more paper work in general. And the other prerequisite — applying more efficient controls based on the acquired information — is implemented through increasing administration, organizational changes, etc. This all is introduced in disguise: Who could object to "missions and visions" or "developmental discussions"?

Speaking of terminologies: The system of language use is an interesting example of cybernetization in memetic systems. It seems that as the culture proceeds towards its stagnation, it is a comprehensive decline: For example, when the language becomes more "civilized", certain ways of speaking become obsolete and are substituted with bureaucratic, politically correct ways of speaking. However, small talk with mere cumulating periphrases becomes void, there is loss of dynamics when the variations are eliminated in the well-balanced refined utterances. When concepts lose real content, they are less capable of capturing the "flesh and blood" — and the mental constructs can only receive their meaning through interaction with the brutal reality. As discussed in chapter 7, true understanding goes only through two-way interaction with the environment. — It seems that there exist languages (like Finnish!) where the dynamic range still extends from very fine nuances to extreme bursts, concepts being clear and accurate, but still poetically open-ended. Surprisingly, perhaps it is such "less cultivated", least cybernetized languages that are best suited for expressing oneself — or for doing science, explicating and perceiving the real world outside our standard constructions?

The result of system cybernetization is that diversity becomes eliminated. What happens when finally all degrees of freedom vanish?

5.3.2 Faith of systems

It seems that all development ends in a collapse. If a system of cybernetic systems are let to adapt freely, catastrophes are unavoidable. How to control the adaptive control without paralyzing the system altogether? — at least, Nature has not found the way to do this. One cannot backtrack from a dead-end, after evolution there is a revolution — again see Fig. 5.5 (another perspective to "saltationism" is studied later in chapter 7).

The mathematically oriented *catastrophe theory* flourished together with chaos theory back in 1980's, trying to explain the processes beyond collapses. The goal was to understand continuous mathematical structures that give rise to abrupt behaviors: Why the once stable balances finally become unstable. However, the trivial one-function experiments did not have very much connection to real-life. In the framework of cybernetic systems one can now qualitatively understand such processes with no additional fancy theories: The key point is (again) the nonideal structure of information acquisition, and the resulting hierarchic structure of systems in different time scales.

Above, in Section (5.2.2), it was observed that a cybernetic adaptation strategy does not necessarily collapse — is there not a contradiction? — There is *not,* because now one is studying wider perspectives: In (5.2.2), it was assumed that the environment remains stationary, whereas now *structural changes* in the environment and in the system itself have to be taken into account — after all, true evolution is change in structures, not tuning of the parameters within existing structures. A closer analysis reveals that there are internal and external reasons for catastrophes. The internal reasons can be seen to be caused by the fast-scale structures changing, and the external reasons are caused by the slow-scale ones.

The fastest, catastrophe-like changes in the system balance can be explained in terms of nonlinearities — gradual changes in the system finally push the system onto the watershed boundary, and after that a new attractor is suddenly found. Such behaviors can easily be explained in the framework of sparse-coded nonlinearities, where some degrees of freedom can remain latent and completely inactive until the conditions are favorable (see chapter 6). As the history of memetic developments reveals, new ideas can remain ignored for a long time — after the turning point, developments can be very abrupt. Individuals are, after all, just noise when looking at the cybernetic systems that are based on statistical models, and developments can become relevant only after the whole population is ready to employ them. There is no evolutionary benefit if too smart enhancements are introduced too early — the key point is that the ideas remain available in the systemic memory (genome, or "menome").

The evolutionary changes within a system can often be characterized in terms of increasing coupling, or the parameters q_i increasing, finally the enhancements ending in structural changes. Flourishing systems are living *at the edge of chaos,* trying to capture the most up-to-date information (or knowledge); however, beyond the borderline determined by the information bandwidth, the visible variation is mostly noise, and the once acquired structure will be lost. What is then the appropriate frequency limit? The system guessing right wins it all. Explicit optimization is not easy here. For example, when making controls faster, the continuous processes typically become discontinuous at some stage as the acquisition of information cannot be immediate. And such discrete-time control systems behave in very different ways as the originally assumed continuous ones: As the sampling rate becomes too fast as compared to the system dynamics, increase in the noise sensitivity follows, and robustness is challenged in changing long-term conditions. There are real-life examples of such tendencies: For example, in "quartal capitalism" samples are taken and controls applied every 1/4 of the year, even though the market dynamics has the range of years; also in modern politics, long-term planning becomes impossible as the politicians have to take care of their everyday popularity according to the population polls — and, what is more, the real time constants in a society can be decades! In both cases, too fast adaptation and control actions can lead to loss of informative excitation and problems with stability.

The structural impacts coming from outside, or from the environment, are caused by low-frequency phenomena. Once some dependency structure that a system exploits has been visible for a (too) long time, it is probable that a slower system takes over that resource. The slowest processes are the most dominant in the long run, and the faster ones are left completely empty-handed, becoming unstable, the statistical balance corresponding to their local models being lost. When the universe gets older, ever slower dynamics become visible, and there is room for new systems to be born in the low-frequency end of the spectrum (again, see Fig. 5.3). When the behavior of the nominal state (or when the "fixed" environment, as seen by the faster system) changes, models for variations around that nominal states become outdated. Hierarchy of systems is like a tree, slower ones being nearer to the "root": When the "trunk" is adjusted, the "leaves" can be violently shaken. The overall system structure cannot change without making its subsystems outdated. Remaining fixed to protect its own fine structure would mean system stagnation.

The finer the constructions become, the larger are the catastrophes — this applies also to memetic systems. Indeed, the magnificent span of German philosophies during some 200 hundred years (ending in a complete catastrophe in 1945), starting from Immanuel Kant, continuing with Hegel himself, Arthur Schopenhauer, Karl Marx, and Friedrich Nietzsche, accompanied by the ideologies of Friedrich Schelling and Johann Fichte, and spiced by von Goethe and von Schiller, is itself an example of such ambitious mental endeavors that can only end in a *nemesis.* Indeed, it was Hegel himself who observed that the state of peace is stagnation, and war has positive moral value: One understands the "real values" again, there is *katharsis.* Along the same lines, the larger scale downfall of the entire culture was studied by Oswald Spengler. But the ideas are still there, the latent thoughts someday having an incarnation as some kind of a synthesis.

To avoid deadlocks of development, mechanisms of *regeneration* seem to be programmed deep in the structures of more sophisticated systems: The cycles of death and birth makes it possible to get back to a fresh start.

5.3.3 Coordination of catastrophes

When this dual nature of balances and catastrophes seems to be such a natural part of cybernetic systems, perhaps it cannot be all bad?

The unavoidable fact is that all complex enough environments are changing over time. One reason for this is that the environments are composed of coevolving systems, and these processes never reach the final state — or, if you start waiting for that, you will be hopelessly late. This dynamic nature of the world is general, it can never be escaped by any system, and it applies fractally in all scales; again, according to Heraclitus, "panta rhei". There is a vicious circle here: World evolves as the systems evolve, and as the world evolves, systems need to evolve. What is more, such changes are not only quantitative — when they continue long enough, quantitative becomes qualitative, and the whole system structure becomes outdated. This is typical in evolutionary systems.

To implement up-to-date control of their environments, and to survive in competition, the systems have to constantly update their models of the environments. Only change exists, but, according to the neocybernetic principles, balances are to be modeled. It seems that nature has found a practical way to gather accurate balance information even in changing environments: It seems that in some sense nature "discretizes" the time-variant processes, so that the processes take place in discrete time rather than in continuous time. First the environment is frozen, then a snapshot is taken, and as the internal tensions cumulate, suddenly the tensions are released to burst the old structures to have a fresh start. During the balance periods optimization of parameters within the structural framework takes place, applying the smooth neocybernetic adaptation strategies, but during the collapses, new structures are introduced to escape the local minima. Truly, the catastrophes themselves do not deliver information, they only produce noise and chaos: It is the balance periods between the catastrophes that are the cookers of information. Catastrophes on the lower level are crucial for the well-being on the higher level to reset the information-producing lower-level systems so that fresh information becomes available. The higher-level system is a model over the possible solutions on the lower level.

How can all this be explained — this all sounds very purposeful: It seems that one needs external control to coordinate the actions, to initialize the system, to run the processes, to collect the data, and to exploit the information. Can the above scheme be seen as more than a metaphor? Again, no master mind is needed to orchestrate the alternation of the "sample and hold". It just seems that "perfect control" — the property of the ultimate survivor in evolution — is an internal contradiction, resulting in extreme sensitivity and eventual collapse of the system. This is the nature's mechanism to guarantee the evolution and emergence of ever higher-order systems; at least, when looking back from the higher level, all lower levels have been obeying the this principle. In a way, nature has built this "apoptosis", or programmed death, in all its systems. And it seems there is automatic synchronization: Only after the properties of the environment are mapped, the controls can become complete — and, after that, it is the whole construction becomes unstable at the same time. Overall stagnation can be reached only when all subsystems have found their models, and when a collapse is then launched at some location, because of however small disturbance, the disturbance soon escalates, wiping away all submodels at the same time.

Even though the continuous processes become discretized, there is no one-to-one coupling to the time variable, and the strong tools from discrete-time dynamic system theory are not available. If trying to model the succession of catastrophes and balances, it is the transitions that are relevant, no matter when they happen, and modeling tools for *event-based system* could be applied. Unfortunately, there exist no strong analysis tools for such systems.

Can anything be said about the catastrophes in general? It is evident that individual processes, or unique catastrophes, cannot be individually modeled — but if seems that the catastrophes are by no means unique, they seem to repeat all over again. One can perhaps abstract over individual catastrophes and find a model for them on some slower time scale.

if the lower-level cycles of catastrophes and balances are correlated, it is information to be utilized. The only problem here is that for the most interesting systems, one cannot see the big picture yet, as one is living in the middle of the turmoil and perhaps emerging new order. Whenever the higher-level structure can be seen, it already exists, and our predicting attempts are late. It seems that such behaviors can be analyzed only in retrospect. As a (very) crude approximation, it seems that there is some general constant here: For a system to sufficiently develop, there are about a dozen regeneration stages at the lower level: How many times cells renew during the lifetime of an organism on average; how many generations there are during a life span of a human culture; how many individual species get extinct before the whole ecosystem collapses. However, the variations here are huge: Some species just seem to be less vulnerable to the changing environmental conditions — but, in many cases, such relics seem to be secondary what comes to the main developments in the larger-scale system.

In evolution biology, there are mysteries: The developments in natural evolution seem rather peculiar. One of the questions is *where are the missing links.* It seems that there have been very different kinds of species following each other with no transitional forms; similarly on the ecology level, there was the era of dinosaurs followed by mammals, etc. A species can be there with no changes for millions of years just to be suddenly substituted. This kind of succession of balances and transients is known as *saltationism.* The lack of continuity in evolutionary processes has been used also as an evidence by creationists; however, as discussed above, this kind of behavior of bursts and balances assumedly is characteristic to all evolving cybernetic systems.

5.3.4 Beyond the balances

Balance is needed for healthy functioning of a system, but catastrophes are needed for healthy functioning of a "supersystem". There must exist variation on the lower level, otherwise higher-level developments cease. It would seem that it is the higher-level system that is running experiments on the lower levels, pushing those systems over their limit on purpose — but, again, there is no such master mind. Catastrophes are built in the cybernetic systems themselves, no matter if the generated excitation is ever exploited, or if it remains just noise in the universe. A healthy evolving system follows its *elan vital* until the edge of chaos — and beyond.

In some environments collapses in different scales are commonplace and — as it seems — generally accepted as unavoidable. The *stock market* is a great equalizer of tensions in economy, tensions manifested through sell and purchase prices, being a simple example where the balances should be found according to cybernetic principles. Again, the stock market dynamics is too fast as compared to real market dynamics: Analysts use their mental models reflecting the common beliefs, making the unquantifiable aspirations visible; these beliefs can be very volatile. The agents try to be smart, trying to predict the competitors and market reactions, thus making the stock market a constructivistic system that lives a life of its own, detached from the reality. The money is not necessarily where the needs are: The challenge of a modern society is to match these tensions — needs and means — and it is here where more cybernetic thinking would be needed, more sophisticated models of the interdependencies and their balances, not straightforward centrally-controlled legislation. In any case, it seems that the minor everyday catastrophes are, as seen from outside, only the mechanism of introducing the necessary excitation and information in the market — but inevitably "the big one" also comes some day.

Extending the observations in chapter 4, it can be claimed that a democratic society — if accompanied by transparency — is the most efficient political system in terms of information exploitation. It combines gathering of bottom-up agent-based innovations, and delivers top-down regulatory directives. But to remain "alive", perhaps democracy, too, needs its enemies, or some excitation from outside. The key question is how can the regeneration of the social structures be implemented in the "postmodern" society, where all destructive developments are prevented. Today, it is interesting to see what the alternative is — how long can Europe become older?

Also in natural everyday systems the "catastrophes" are a part of normal behaviors in healthy systems: The limits are being tested all the time. Without pushing the limits, the dynamic range becomes narrower. For example, take the living body: If the machinery is not "calibrated", if there are not the necessary degrees of freedom visible in data, missing compensation capacity against certain excitations is developed. If the body does not get acquainted with microbia, there can be an increase in autoimmune diseases. The genes only determine the gross structure, but the fine-tuning of the system is found as a interaction process with the environment: Diseases are minor catastrophes, extreme cases that determine the dynamic range of the system. And as they say in the United States: You cannot know what the business is all about before you have experienced some bankruptcies.

It is all cybernetic subsystems that are hungry for information: In extreme balance the system starves. This can be extended even to analysis of mental sanity: One needs to have "highs and lows" to experience what life is about. And extreme feelings seem to be the seed for higher-level memetic systems, at least what comes to artistic creativity. Of course, diseases are related to the loss of balance in the biological environment, mental diseases are related to loss of balance in the cognitive environment, and "social diseases" are related to loss of balance in the social environment. If there are no real political issues in a welfare society, the system becomes — concretely — insane. But extreme emphasis on the balance is a fallacy: If there are no real obstacles or problems, these will be imagined — or when life is too easy a healthy mind actively searches for challenges, to find the balance of feelings between danger and security. The "real artists" simply need to experience the highs and lows. Without mental explorations and excitations one has an incomplete model of oneself and of the world. This is where neocybernetics goes even beyond the Eastern wisdom: The goal is not extreme harmony or elimination of variation — as they say it, "in Hell you have merrier company". Such discussions can be extended even to purposeful life and what happiness is about: It is mastery of one's life, or awareness of one's capability of coping with all possible challenges one might face.

It has been observed that evolving morality, etc., are becoming fields of scientific study [83]. This is true, but there is another tendency, too: In the neocybernetic framework all biology is coming back towards more abstract philosophies.

It is tempting to draw some bold conclusions concerning issues that by no means have been seen as subject to scientific study. For example, *why there is evil, why there is poverty in the world,* or, *why there is suffering?* Indeed, suffering seems to be necessary for a cybernetic system to fully develop. There are two ends in the continuum – always somebody is the poorest. If there were no differences, the *heat death* would have been reached. Questions like *why there is death* can also be attacked: Death is dropping out from the dynamic equilibrium to the static balance, it is nature's means to assure regeneration in the system. Whereas death is the final catastrophe from the individual's point of view, it is necessary from the point of view of the wider-scale system. At some stage of the higher-level development, lower-level models are so outdated that it is easiest to start all over again.

Is it perhaps so that engineering disciplines, like understanding of control engineering, can give some mental building blocks for understanding of, for example, what *good life* is? What is more, it is not only ethics, but also other branches of philosophy that can be affected by the cybernetic considerations. These issues are studied in more detail in chapter 10. — However, next it is time to go back into details: It is there where the beauty is.