Andreas Wagner

Evolutionary Biology and Beyond

Born January 26, 1967 in Linz, Austria. Studied Biology and Mathematics at the University of Vienna. 1990, M.S. in Molecular Genetics, University of Vienna. 1991, visiting student, University of Chicago. 1994, Ph.D., Yale University, Department of Biology. 1995, Postdoctoral Fellow, Yale University. — Address: Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA; E-mail: aw@santafe.edu.

Maybe the most important aspect of my stay at the Wissenschaftskolleg is that it allowed me to venture into two areas, computational neurobiology and the philosophy of science, that are quite different from my main field of research, which is population genetics. Before coming to the Kolleg, I had had only very limited exposure to current research in both areas. Aside from the unlimited freedom to pursue one's academic interests at the Kolleg, two major factors contributed to my success in these excursions. First, the truly exceptional library system, and second, the opportunity of interacting with researchers from these unfamiliar areas. Among the multitude of immensely stimulating conversations I had with many Fellows, I would like to mention especially those with Margie Morrison and Mary Morgan, which familiarized me with an enormous body of literature on causation in the social sciences and in the philosophy of physics. Similarly, many conversations with Holk Cruse and Helge Ritter allowed me to familiarize myself with the current level of knowledge regarding the design principles underlying modularity in neural networks.

Below, I briefly describe the scope of three research projects I carried out at the Kolleg, only the first of which I had originally planned to pursue during my stay.

The Evolution of Redundant Gene Functions

Genes with redundant functions are ubiquitous. Numerous cases of redundant genes have been found in invertebrates, vertebrates, and

microorganisms. Their gene products cover a wide range of biochemical functions, including transcriptional regulators, extracellular matrix proteins, protein kinases, and proteins involved in intracellular force generation. Despite this functional diversity, all genes with redundant functions share one property. The phenotypic effect of eliminating the gene through a loss-of-function mutation is absent, or it is much weaker than expected from independent lines of evidence. Such lines of evidence include strong evolutionary conservation of DNA or amino acid sequence, or the relation of a morphological defect to the size of the expression domain for gene products that presumably act in a cell-autonomous way. Sometimes the expression of such genes might serve no biological purpose at all. However, in general, one or more other genes exist with functions indistinguishable from, but at least overlapping with that of the presumably redundant gene. These two cases are often referred to as full and partial redundancy. The strong phenotype of double (triple, etc.) mutations in the respective genes then serves as proof that they do fulfill a biological function.

Gene duplication events are probably the main source of redundant gene functions. Immediately after a duplication becomes fixed in a population, the two copies (original/duplicate) are completely redundant. Subsequent mutations either cause the silencing of one of the two copies, or a functional diversification between them. In the latter case, redundancy will decrease over time. However, genetic studies on the early development of *Drosophila* have led to the suggestion that previously dissimilar gene functions may have been recruited into redundant control mechanisms of developmental pattern formation events. This would imply that redundancy can somehow be selected for because it "masks" mutations that would otherwise be deleterious. A similar masking of deleterious mutations has been involved in models for the evolution of diploid life cycles and of dominance modifiers.

Because the importance of partial redundancy has only recently been fully appreciated, there is little population genetic insight into the relative contributions of mutation, selection, and genetic drift to the evolution of redundancy. Wherever the subject is discussed in the literature, it is treated on a purely qualitative level. During my stay at the Kolleg, I have established a class of simple mathematical models that allow one to address the relevant issues in a fairly general fashion. I have been able to show, using both analytical and numerical methods, that natural selection can indeed increase redundancy among genes in a population, provided that certain boundary conditions are fulfilled. These include sufficiently large populations and sufficiently large variance of the mutational process that causes functional divergence among gene

products. The models contain a number of simplifying assumptions, such as haploidy and the restriction to pairwise interactions among two genes in a polygenic system. The next step, which I have begun at the Kolleg, is to explore how relaxation of these assumptions affects the predicted evolution of redundancy.

To gain increased understanding of how redundancy evolves is important for at least two reasons. First, redundant gene functions may be an important source of evolutionary novelty on the biochemical level. The more redundancy is sustained by a developmental system, the more pronounced this reservoir function is likely to be. Second, because of its potential "buffering" effect, redundancy may substantially contribute to different degrees of genetic canalization of developmental and metabolic pathways, i.e., to differential robustness of such pathways to mutations. Understanding the evolution of redundancy may therefore aid in solving one of the big unresolved problems in evolutionary biology, namely to what extent genetic variation translates into phenotypic variation, and to what extent the mechanisms responsible for this translation can evolve.

Causation in Complex Biological Systems

The notion of causation is of fundamental importance to the process of scientific inquiry. Curiously, however, the attention paid to the issue varies considerably across fields. In economics, and in the social sciences in general, intense debates on the nature of causality have been going on for decades. In these areas, the main issues regard the possible inference of causal relations from statistical data. Biology, on the other hand, is comparatively untouched by such debates. Clearly, questions regarding the notion of causality lurk behind most contemporary topics in biology, such as the "unit of selection" problem, but rarely are they made explicit. There is ample room for speculation regarding the possible reasons for such differences. A prominent candidate is the relevance of the social sciences and especially of economics to policy decisions. For policy issues the use of causal language, such as "investment causes increased production of capital goods", is central. The debates then result from attempts to clarify the terminology, and from the manifold ambiguities in identifying causes and their effects. Identifying "correct" causal relations among economic variables and eliminating "spurious" correlations from statistical data is thus an activity with potential impact on the lives of many individuals.

Recent developments in biology, specifically in the biomedical field, make it seem likely that similar debates on the nature of causation will become of increasing importance in biology. Genetic factors contributing to diseases are being identified or postulated at an unprecedented rate. And while the importance of genetic factors has been hotly debated for many decades (especially for complex conditions such as schizophrenia), molecular genetic methods have only recently advanced sufficiently far that tests for genetic risk factors for many conditions become feasible. The question of the extent to which complex genetic factors "cause" diseases becomes of immense practical importance if one considers its potential influence on the health care system. Should carriers of susceptibility genes be ineligible for health insurance? How about access to employment for individuals who carry susceptibility markers for certain noxious environmental influences? Such questions will soon be prominent social policy issues, and the necessity to clarify the notion of causation is likely to arise for very practical reasons, similar to the situation in economics.

Interestingly, virtually all of the literature on this issue in the social sciences revolves around systems with linearly interacting variables. This approach is unlikely to be fruitful for the biological problem, because it is well known that epistatic (non-linear) gene interactions are pervasive in many metabolic and developmental pathways whose components may become disease-"causing" genes. However, not much of a framework exists for understanding causation in nonlinear systems. Using several examples from linear and non-linear dynamical systems important in mathematical biology, I developed the argument that a regularity notion of causation (i.e., whenever event A takes place, event B will follow) can only be meaningful if the behavior of a system can be approximated by assuming linear interactions among its state variables. In simple terms, the reason is that in nonlinear systems the effects of changes in individual state variables may depend on the state of all other variables. If this is the case, the context in which an event (cause) occurs may become so important that a regularity notion of causation based on individual variables of a system becomes meaningless. The "background" of all other variables becomes crucially important, so important that one may not be able to identify individual causal factors. Whether this is the case for genetic systems is an open question. However, most of the reasoning in this area is based not on *functional* interactions among state variables (as in the case of dynamical systems), but only on *statistical* associations between genetic factors and phenotypic traits. At the very least, one should be very careful in attributing a disease-causing role to individual genetic factors and in assuming that

effects of the genetic background will somehow "average out" in a population, as long as there is little insight into the functional (and often nonlinear) interactions among genes that may affect phenotypic traits.

Self-Organization of Modular Neural Networks

Modular organization is pervasive in biological neural networks. A neural network is considered modular if it can be subdivided into groups of neurons, such that neurons within a group have many synaptic connections, whereas there are only a few connections between groups. Modularity can be anatomical, i.e., mediated through presence or absence of synapses, or physiological, i.e., mediated through strong connections within modules and weak connections between modules. Probably the best known examples stem from the visual cortex of mammals, where afferent stimuli of different origin or quality are processed by different modules.

My research in this area focused on possible ways in which modular organization of biological and artificial networks might be achieved through principles based only on local information within a network. Whether such modularity exists for functional reasons, or whether it is only a by-product of pattern-formation processes in neural development is not known. Experimental data and mathematical models are available that support either alternative. However, research on artificial neural networks strongly suggests that modular networks are in many ways superior to networks with uniform connections. They often learn faster, generalize better, and solve problems more rapidly. This is because their architecture somehow matches the structure of the task they are designed to solve. In artificial networks, the modular structure is usually given a *priori* by the designer of the network. The design principle is "insight" of the designer into the structure of the task or the data. And while it is likely that there are modular architectures worse than those prescribed by an intelligent designer, it seems also plausible that there are better ones. This is because usually the nature of a task presented to a neural net is sufficiently complex that one can make only an informed guess about an optimal network organization. If it were simpler, one might not need a neural network to solve the task.

One of the achievements of the theory of neural computation was the insight that neural networks can learn based on purely local information, i.e., information that is available to a neuron only via its connection to its immediate neighbors. For several reasons, it would be highly

desirable to find similar self-organizing mechanisms that allow a network to develop a modular architecture reflecting the structure of a task, according to some criterion of optimality (e.g., fast learning). First, it would allow one to design artificial networks that (i) perform better on a given task, and (ii) are easier to implement because of smaller numbers of connections compared to a uniformly dense network. Second, if such mechanisms can be found, they might prove helpful in understanding modularity in biological neural networks. For if the modularity of a biological network indeed reflects optimality with respect to some criterion, there was no intelligent designer available in evolution who would have prescribed such an architecture. In this case, self-organization must have played a role in the evolution of modularity.

To be able to carry out work in three such different areas within only a few months' time does strongly reflect on the unique working conditions at the Kolleg. It is astonishing how smoothly and almost invisibly the staff accompanied me through the academic year, and how it took many time-consuming burdens of "life management" off my shoulders. My warmest thanks to all staff members, in particular to the members of the computer system's administration team and to the librarians.

The results of the above projects will find their way onto printed paper, and have partly already done so. More importantly, however, the Wissenschaftskolleg permits learning experiences that are unlikely to result in publications, and it is nevertheless these experiences that may have the most impact on one's future style of research and self-perception as a scientist. For me, such an experience took place during my stay, although its result is not easy to communicate convincingly. Through the Colloquia series, through the countless discussions over lunch and dinner, and through a number of seminars on the relation of biology and the humanities, I became aware that it is not problems of disagreement on factual issues that cause communication problems across discipline boundaries. It is rather the "facts" themselves that seem to be different. I came to appreciate how profoundly cultural the process of maintaining a common language within a particular field is. The absence of a common language among different fields, the absence of a common system of meanings for common terms, seems to reflect truly different perceptions of "reality", and therefore different "realities", because perceptions are probably all we can have of "reality". Being allowed to have this experience has been a great privilege, albeit somewhat humbling, for someone trained for doing "hard" science.