

Computer Simulated Evolution of a Network of Cell-Signaling Molecules

Dennis Bray* and Steven Lay†

*Department of Zoology and †Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, United Kingdom

ABSTRACT We have trained a computer model of a simple cell-signaling pathway to give specified responses to a pulse of an extracellular ligand. The pathway consists of two initially identical membrane receptors, each of which relays the concentration of the ligand to the level of phosphorylation of an intracellular molecule. Application of random "mutational" changes to the rate constants of the pathway, followed by selection in favor of certain outputs, generates a variety of wave forms and dose-response curves. The phenotypic effect of mutations and the frequency of selection both affect the efficiency with which the pathway achieves its target. When the pathway is trained to give a maximal response at a specific concentration of the stimulating ligand, it gives a consistent pattern of changes in which the two receptors diverge, producing a high-affinity form with excitatory output and a low-affinity form with inhibitory output. We suggest that some high- and low-affinity forms of receptors found in present-day cells might have originated by a similar process.

INTRODUCTION

Living cells respond to their environment through elaborate networks of biochemical reactions that have evolved to perform certain input/output transformations. The input signals may be soluble or bound to the extracellular matrix or bound to the surface of a neighboring cell. Although an organism produces receptors for only a limited number (perhaps hundreds) of such signals, the signals act in combinations to generate a code with many millions of possibilities. Integration between different extracellular signals depends, in large part, on interactions between various protein phosphorylation cascades. Some individual proteins in the cascades function as integrating devices; in response to multiple signal inputs, they produce an output that is calibrated to cause the desired biological effect (Nishizuka, 1992; Stoddard et al., 1993; Pawson and Schlessinger, 1993).




It is apparent that any system as complex as a living cell will be capable of performing an enormous number of transformations of its input signals. In 1961, Monod and Jacob (1961) pointed out that even with the limited number of biochemical switches and other regulatory devices known at that time, one could in principle perform a vast number of computations. More recently, the enormous proliferation of signal mechanisms together with the extensive cross talk that exists between signal cascades led to the suggestion that they might function as neural networks and serve as pattern recognition devices (Bray, 1990). A formal demonstration that linked chains of chemical reactions (not, in this case, biochemical reactions) can act as Turing machines or as neural networks has been obtained (Hjelmfelt et al., 1991; Hjelmfelt, et al., 1993).

Although the potential computational capacity of cell-signaling pathways is enormous, we have little information at present on the actual computations they perform. Because even the most important signaling pathways have been known only for a few years and new signaling components are still being discovered on a regular basis, very little kinetic data is available. Nevertheless, we do have information on the general features of signaling components: on the molecular structure and mode of operations of cell-surface receptors, for example, and the way in which they relay information about their cognate ligands to the interior of the cell. Therefore, it is possible to explore the computational possibilities inherent in small "circuits" composed of membrane receptors and other cell-signaling proteins and to try to learn how these are shaped by biological evolution.

MODEL PATHWAY

To explore the properties of networks of cell-signaling reactions, we have constructed a computer simulation of a simple idealized pathway (Fig. 1). This comprises two cell-surface receptors, R_1 and R_2 , each of which has a binding site for an extracellular ligand, L , and an intracellular target molecule, T . The concentration of L serves as an input to the pathway that produces a corresponding concentration of phosphorylated target molecule, T_p , as output.

Seven reactions or binding steps are performed by the pathway shown in Fig. 1




Received for publication 25 October 1993 and in final form 16 December 1993.

Address reprint requests to Dennis Bray, Department of Zoology, University of Cambridge, Cambridge CB2 3EJ, UK.

© 1994 by the Biophysical Society

0006-3495/94/04/972/06 \$2.00

FIGURE 1 Signaling pathway used in the simulation. Two transmembrane receptors (R_1 and R_2) each have an extracellular domain with a binding site for a low molecular weight ligand (L) plus one cytoplasmic domain with catalytic activity. Binding of the extracellular ligand causes, through an unspecified mechanism, changes to the cytoplasmic domain that enable it to bind the target molecule (T). This receptor-ligand complex then catalyzes formation of Tp, which dissociates immediately to enter the cytosol. An autocatalytic phosphatase activity in the cytosol restores the target protein to its initial, unphosphorylated state. The kinetic performance of the pathway is specified by 7 variable rates (indicated by *thick arrows*); association rates (*thin arrows*) are assumed to be diffusion-limited constants.



Steps 1, 2, 4, and 5 are treated as simple binding equilibria governed by the Law of Mass Action. Steps 1 and 3 are further simplified by the assumption that the extracellular ligand is present in large excess. Each binding equilibrium is characterized by a constant rate of association (on rate) here assumed to be a diffusion-limited rate of $2 \times 10^7 \text{ M}^{-1} \text{ s}^{-1}$ (Fersht, 1985; Northrup and Erickson, 1992)-and a dissociation rate that is variable and subject to mutation. Phosphorylation and dephosphorylation steps (3, 6, and 7) are treated as first-order, irreversible, reactions, each characterized by a single rate, which is also variable.

The total concentrations of the three cell-signaling molecules, R_1 , R_2 , and T, summed over their various complexes, are each held constant at $5 \mu\text{M}$ throughout the simulation. The performance of the pathway is consequently defined by 7 rates, each of which is potentially variable and subject to independent mutation. The total number of possible networks that can be derived from the pathway shown in Fig. 1 is very large; if each rate is allowed to take 10 different values, then there will be 10^7 networks, each with a distinct performance.

The steps of the model cell-signaling pathway are intended to be simplified, generic representations of signaling reactions and do not correspond in detail to any known system. The ligand could be any diffusive molecule present in the extracellular fluid, such as a hormone, growth factor, neurotransmitter, or chemattractant. The cell-surface receptors of the model resemble, in broad features, the many transmembrane receptors that possess an intrinsic kinase associated with their catalytic domain (Hunter et al., 1993). They resemble most closely the family of histidine kinase receptors found in prokaryotic cells that propagate signals across the plasma membrane, it is believed, by a propagated conformational change (Parkinson and Kofoid, 1992). Receptors for epidermal growth factor (EGF), platelet-derived growth factor (PDGF), and insulin all have an associated tyrosine

kinase activity, although signal transduction by such receptors depends on their ligand-induced dimerization (Fantl et al., 1993; Hunter et al., 1993), a step not included in the present model. Alternatively, the receptors in the model might be viewed as receptors whose cytosolic domain activates a cytoplasmic protein kinase, such as the Src protein (Cooper and Howell, 1993); or they could be G protein-linked receptors, in which case the target molecule will be




FIGURE 2 Conversion from random numbers to phenotypic factors. Random numbers between 0 and 1.0, generated by the program, were converted to phenotypic factors using one of the above curves. Factors were then used to change rate constants of the network, individually selected on a random basis, according to the relationship, $\text{new_rate} = \text{phenotypic_factor} \times \text{old_rate}$. Note that curves A, B, C, and D are symmetrical about 0.5, whereas curves E and F give a preponderance of factors less than 1.0.

a G protein carrying GDP (guanosine-5'-diphosphate), and its phosphorylated output, the alpha subunit carrying GTP (guanosine-5'-triphosphate) (Gilman, 1987).

PATHWAY EVALUATION

The program used to evaluate the pathway shown in Fig. 1, was modified from a simulation of signaling reactions in bacterial chemotaxis (Bray et al., 1993). Many of the same data-handling algorithms and test routines were used, whereas the differential equations and the method of numerical integration were changed. Two forms of stimulus were used to evaluate each network of reactions. In one, a 40-s square-wave pulse of 5×10^{-8} M ligand (50 nM) was supplied to the system, and the output concentration of Tp was monitored with time. The other stimulus was a series of ligand concentrations ranging from 10^{-16} M to 10^{-4} M, with the steady-state concentration of Tp being measured at each ligand concentration.

The response of networks to either type of stimulus was evaluated by computing the output values of Tp by numerical integration of reactions 1 to 7. The following differential equations were used for the pulse stimulus

$$\Delta_i = (R_i \times L \times \text{on_rate} - R_i L \times k_i) \times \delta t \quad (i = 1, 4)$$

$$\Delta_i = (R_i L \times T \times \text{on_rate} - R_i L T \times k_i) \times \delta t \quad (i = 2, 5)$$

$$\Delta_i = ([\text{ATP}] / ([\text{ATP}] + K_{\text{atp}})) \times R_i L T \times k_i \times \delta t \quad (i = 3, 6)$$

$$\Delta_i = T_p \times k_i \times \delta t \quad (i = 7);$$

Where Δ_i is the rate of change of the i^{th} reaction, k_i its variable rate, and δt the time increment used in the integration. [ATP] is the concentration of ATP in the cytoplasm (usually taken as 3 mM) and K_{atp} the K_M of the phosphorylation reaction (taken as 0.1 mM); on_rate is the diffusion-limited rate of association, taken as $2 \times 10^7 \text{ M}^{-1} \text{ s}^{-1}$. Reactions 1 and 4 were modified in the case of the dose-response curves to:

$$\Delta_i = \frac{(R_i + L) \times L}{L + k_i / \text{on_rate}} - R_i L \quad (i = 1, 4)$$

implying that the receptors were already equilibrated with ligand at the start of the simulation. In this way very low concentrations of ligand could be used without requiring an excessively long period of time to attain equilibrium.

In most cases a fourth-order Runge-Kutta integration (Press et al., 1992) was used to evaluate networks at each selection step, although a simpler sequential Euler method was found to be adequate for some dose response curves and had the advantage of reducing CPU time. Final values of rate constants were in every case evaluated by fourth-order Runge-Kutta integration.

EVOLUTION ALGORITHM

Networks of reactions were trained to give desired output characteristics by a process analogous to the evolution or, more accurately, selective breeding of a novel species. This

entailed making a set of random "mutational" changes in the 7 variable rates, each set being followed by a selection step. Mutations were simulated by a series of random numbers between 0 and 1, generated by the routine *ran1* described by Press et al. (1992). These numbers were used (i), to select one of the 7 variable rates of the network, and then, (ii), to specify a multiplicative factor (phenotypic factor) by which this rate was to be changed. Various relationships were employed in this second step to convert random numbers to phenotypic factors (Fig. 2).

The starting network was allowed to produce a number of "offspring" networks by an alternating sequence of groups of mutations followed by a selection step. Offspring networks were then used as founder networks for further sequences of mutation and selection. This strategy was employed so as to





FIGURE 3 Schematic diagram showing the sequence of steps used to train networks. Each line represents a lineage of networks related by sequential, random changes in their rate constants. Within each lineage, each group, or cluster, of "mutations" (typically 5 in number) was followed by a selection step (indicated by intermittent cross bars); there were usually many more selection steps in each lineage than indicated here (typically 300). The basis for selection was the performance of the network in response to either a pulse of ligand or a series of pulses of different concentration, as described in the text. Only those altered networks that gave an improved performance (closer to the target) were adopted for further modification. After the desired number of offspring networks (○) had been produced (usually between 5 and 15, depending on the time taken by the program) the network that had the closest output to the target (here called an "isolate" and indicated by a ●) was used as a new founder network. The entire process was then repeated a specified number of times (usually 5–10). Kinetic rates of each of the "isolates" were recorded, and their performances in the appropriate test were assessed graphically.

FIGURE 4 Results of two simulations in which networks were trained to give a specified time course. Each individual curve shows the concentration of Tp produced in response to a square-wave pulse of 5×10^{-8} M ligand, applied for 40 s. The target performance used as a basis for selection is shown shaded gray in the final pattern of each series. In (A) a series of curves shows an initially positive square wave pulse that becomes inverted after the period of training. In (B), the response of the network changes from a square wave pulse that closely follows the stimulus, to one that remains elevated for a long period after the stimulus (thereby behaving as an "ON" switch).



reduce the incidence of networks that became trapped in false minima, a common occurrence in most optimization routines. A schematic outline of the selection procedure is shown in Fig. 3.

RESULTS OF TRAINING

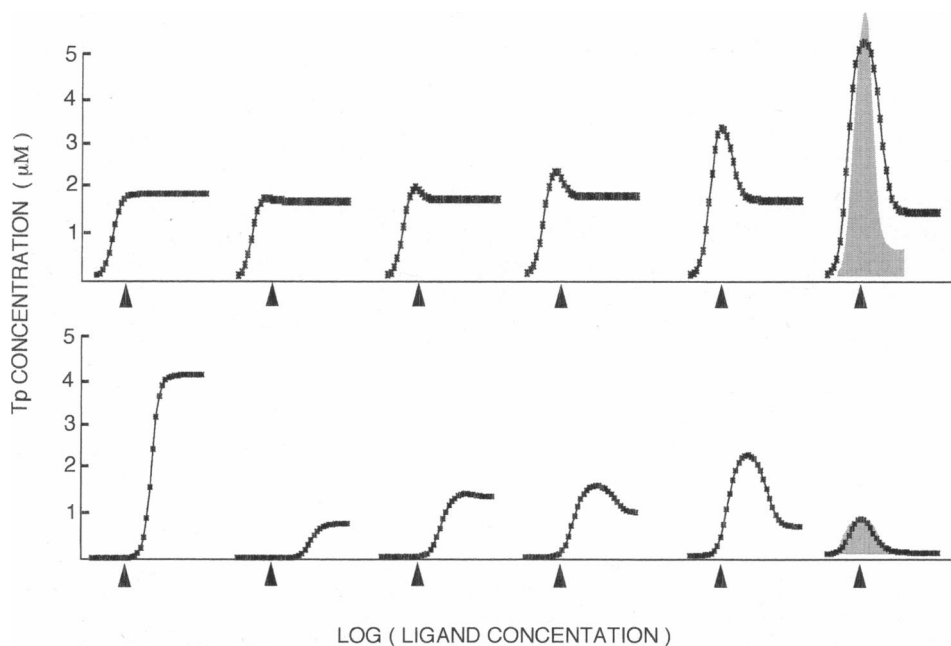
The results of four typical training sessions are shown in Figs. 4 and 5. Each individual curve in these figures represents the output of a different network based on the pathway shown in Fig. 1; each series of curves illustrates the progression of one training session, with the target output shown as a shaded curve superimposed on the final curve of the series.

In Fig. 4, the signal applied was a square wave pulse of ligand L. Rate constants of the initial network were chosen

so that its output was a positive pulse of Tp with a similar time course to the applied stimulus. In the first training session illustrated (Fig. 4 A), the target output was a negative-going square wave of similar duration to the applied pulse, so that the network achieved an inversion of the original wave form. In the second series (Fig. 4 B), the target pulse had a similar magnitude to that of the starting output, but its duration was longer; in this case the network behaved like an irreversible "on switch." Other simulations (not shown) were used to produce changes in the magnitude of the output square wave by a specified amount (amplification) and to generate a rising ramp of Tp concentration of the same duration as the input square wave.

In Fig. 5, the test applied was a series of fixed concentrations of the stimulating ligand from 10^{-16} M to 10^{-4} M.

FIGURE 5 Results of two simulations in which networks were trained to give a specified dose-response curve. Each curve shows concentrations of the output molecule, Tp, produced by a range of ligand concentrations, from 10^{-16} M to 10^{-4} M. The target performance used as a basis for selection is shown shaded gray in the final pattern of each series. The two series, A and B, begin with different networks, but both evolve to an optimal value of about 5×10^{-12} M ligand (indicated by the position of the arrowheads).



At each concentration, the network was allowed to come to steady state before the concentration of Tp was recorded. The starting rate constants had identical values for the two receptors (that is, $k_1 = k_4$; $k_2 = k_5$; $k_3 = k_6$), and this resulted in a simple hyperbolic dose-response curve (see the first curves in Fig. 5, A and B). In the two training sessions illustrated, the target was in each case a dose-response curve that showed a maximum Tp concentration at 5×10^{-12} M ligand, the target maximum in the case of Fig. 5A was chosen to show an amplification of the response, and in Fig. 5B a diminution.

In every case in which a network was trained to produce a maximum concentration of Tp at a certain concentration of ligand, the rate constants of the initial pathway changed in a characteristic way. As illustrated in the representative network of Fig. 6, the dissociation constants of the two receptors, which were initially identical, diverged to give one receptor with a relatively high-affinity binding for ligand and a second lower-affinity receptor. At the same time, the catalytic efficiency of the two receptors also changed, with the high-affinity receptor giving a rapid phosphorylation of the target protein and the low-affinity receptor a slow, or non-existent, phosphorylation. Because the low-affinity receptor still bound strongly to the output protein, its low catalytic activity made it act as a nonproductive, "inhibitory" complex. At low concentrations of ligand, therefore, only the binding site of the high-affinity receptor was occupied to any appreciable extent leading to an increase in Tp levels. At high concentrations of ligand, the low-affinity receptors were also occupied; this reduced Tp, producing the desired bell-shaped dose-response curve.

MUTATIONAL PARAMETERS

Each training run was distinguished by a unique combination of starting network and mutational parameters. The




FIGURE 6 Rate constants obtained in a network trained to respond optimally to a concentration of 5×10^{-12} M ligand. These are the rate constants that generated the final curve in Fig. 5A.

conversion between each random number and its effect on a rate constant (the phenotypic factor) had a major influence on the success of the simulation. For most simulations, the conversion factor lay on a linear range from 0.8 to 1.2, corresponding to random numbers ranging from 0.0 to 1.0 (curve B, Fig. 2). Wider ranges (larger changes per mutation) allowed the network to reach the target faster, as shown in Table 1. The simulation was also more effective when factors on either side of 1.0 were equally possible; that is, when an increase in rate was as probable as a decrease. Generation of the desired dose-response curve thus depended on the ability of the pathway to improve catalytic efficiencies and binding affinities as well as reduce them; the program would be rendered ineffective if any of the rates had reached a maximum possible value.

DISCUSSION

The network used in this study is far simpler than any naturally occurring pathway of signal reactions. Nevertheless, as we show in this paper, even this minimal array of three signalling proteins can be trained to generate a variety of transformations of its input. Amplification, diminution, and inversion of input wave forms were readily achieved. Similarly, the dose response of the pathway to a range of ligand concentrations could be shifted at will over the entire range studied and could be modified to show a maximum response at a specified ligand concentration. Other target performances were more difficult to achieve. The pathway could produce a rising ramp of Tp in response to a square wave input, but a sinusoidal output seemed unobtainable. Nor was it found possible to generate a curve that followed the first differential of an input pulse. It is likely that for these and other more complicated responses additional cell signalling elements are necessary. Feedback from Tp that modifies the catalytic ability of the receptors, for example, should considerably enhance the range of responses produced.

TABLE 1 Effect of changing the phenotypic factor on the rate of training of a network

Isolate	Cost Value (mM)					
	A	B	C	D	E	F
0	14.0	14.0	14.0	14.0	14.0	14.0
1	13.1	13.1	12.9	6.10	13.2	13.2
2	12.9	12.6	12.2	1.10	13.2	13.2
3	12.5	3.6	3.2	0.10	13.2	13.2
4	11.1	3.0	1.5	0.10	13.1	13.1
5	6.0	2.7	1.5	0.10	13.0	13.1
6	2.1	2.6	1.5	0.10	13.0	13.1

A network was trained to give a maximum response at 5×10^{-12} M ligand. Rate constants of the starting network and the target network were the same as in Fig. 5A. Cost values represent the cumulative absolute difference between the concentration of Tp generated by the network and the corresponding concentration of Tp of the target response, summed over a range of ligand concentrations. These cost values were generated for a series of isolates, as described in Fig. 3. The transformations A-F, used to convert random numbers to phenotypic factors, are displayed graphically in Fig. 2.

The algorithms used in this study are not intended to mimic biological evolution with any precision. Changes are "mutation driven" and large, "genetically diverse" populations of networks are not produced. The procedure described is therefore closer to the artificial selection used to breed new strains of domestic animals, and we do not expect it to teach us anything about evolution in the wild.

Nevertheless, this exercise does raise a number of questions regarding the constraints on the evolution of systems of interacting reactions. One question arises from our observation that the most efficient way to train a network is to allow "mutational" changes both to increase and decrease rate constants; when there was a preponderance in one direction, then the efficiency fell off dramatically (Table 1). Could it be that a similar feature is needed to remain responsive to the needs of the organism and to allow evolution to produce signaling networks of the desired range of performance? Are kinetic rates and binding constants in living cells, for a similar reason, also poised at a level that allows them to move in either direction? Note that it has been argued that, for the majority of metabolic enzymes, there is little if any selection pressure to maximize their reaction rates (Kacser and Burns, 1973).

Perhaps the most interesting response of the simulated pathway occurred when it was trained to give a maximum output at a specified concentration of ligand (Fig. 6). The same pattern of changes in rates was seen with a wide range of conditions, raising the possibility that a similar path may have been taken by biological evolution. Many cell-surface receptors exist in low- and high-affinity forms, and it is possible that some of these serve, as in our simulated system, to give an optimal response at a certain concentration of ligand. Such a response could arise by a sequence of events similar to that observed in this model system, namely: (i) duplication of the gene for the receptor, (ii) reduction of its affinity for the ligand, and (iii) a change in its cytoplasmic domain so that instead of sending a positive signal it sends a negative one.

This scenario fails to account adequately for the two forms of nerve growth factor receptor, which are structurally distinct and bear no obvious homology to each other (Meakin and Shooter, 1992). Neither does it fit with EGF receptors with different affinities (Berkers et al., 1991), inasmuch as these seem to be the same protein, which is modulated by an association with other cell components. On the other hand, the two forms of PDGF receptor have different affinities for various isoforms of PDGF and have closely similar protein sequences (Heldin and Westermark, 1990), so that they remain possible candidates.

One pathway, in particular, that seems to fit our model is that mediating the chemotactic response to pH in coliform bacteria. Two membrane receptors in *Escherichia coli*, Tar and Tsr, which recognize different amino acid attractants, also show different responses to changes in pH (Krikos et al., 1985). Tsr mediates an avoidance response in acid media, whereas Tar mediates an avoidance response in alkaline media. Chemotactic pH sensors, believed to be specific histidine residues, are located in both the cytoplasmic and the external

domains of both receptors. The histidines present in Tar evidently generate a positive, attractant response when protonated; those in Tsr generate a negative, or repellent, response. Operation of these two opposing receptors in concert ensures that wild type bacteria migrate toward neutral pH. Interestingly, the two receptors, Tar and Tsr, are homologous and show strong evidence of having evolved from a common ancestral gene, as they should according to our model (Krikos et al., 1983).

REFERENCES

- Berkers, J. A. M., P. M. P. van Bergen en Henegouwen, and J. Boonstra. 1991. Three classes of epidermal growth factor receptors on Hela cells. *J. Biol. Chem.* 266:922-927.
- Bray, D. 1990. Intracellular signalling as a parallel distributed process. *J. Theor. Biol.* 143:215-231.
- Bray, D., R. B. Bourret, and M. I. Simon. 1993. Computer simulation of the phosphorylation cascade controlling bacterial chemotaxis. *Mol Biol. Cell* 4:469-482.
- Cooper, J. A., and B. Howell. 1993. The when and how of Src regulation. *Cell.* 73:1051-1054.
- Fantl, W. J., D. E. Johnson, and L. T. Williams. 1993. Signalling by receptor tyrosine kinases. *Annu. Rev. Biochem.* 62:453-481.
- Fersht, A. 1985. *Enzyme Structure and Mechanism*. W. H. Freeman, New York. 475 pp.
- Gilman, A. G. 1987. G proteins: transducers of receptor-generated signals. *Annu. Rev. Biochem.* 56:615-649.
- Heldin, C.-H., and B. Westermark. 1990. Platelet-derived growth factor: mechanisms of action and possible in vivo function. *Cell Regul.* 1:555-566.
- Hjelmfelt, A., J. D. Weinberger, and J. Ross. 1991. Chemical implementation of neural networks and Turing machines. *Proc. Natl. Acad. Sci. USA.* 88:10983-10987.
- Hjelmfelt, A., F. W. Schneider, and J. Ross. 1993. Pattern recognition in coupled chemical kinetic systems. *Science (Wash. DC).* 260:335-337.
- Hunter, T., R. A. Lindberg, D. S. Middlemas, S. Tracy, and P. van der Geer. 1993. Receptor protein tyrosine kinases and phosphatases. *Cold Spring Harbor Symp. Quant. Biol.* 57:25-30.
- Kacser, H., and J. A. Burns. 1973. The control of flux. *Symp. Soc. Exp. Biol.* 32:65-104.
- Krikos, A., N. Mutoh, A. Boyd, and M. I. Simon. 1983. Sensory transducers of *E. coli* are composed of discrete structural and functional domains. *Cell.* 33:615-622.
- Krikos, A., M. P. Conley, A. Boyd, H. C. Berg, and M. I. Simon. 1985. Chimeric chemosensory transducers of *Escherichia coli*. *Proc. Natl. Acad. Sci. USA.* 82:1326-1330.
- Meakin, S. O., and E. M. Shooter. 1992. The nerve growth factor family of receptors. *Trends Neurosci.* 15:321-331.
- Monod, J., and F. Jacob. 1961. General conclusions: teleonomic mechanisms in cellular metabolism, growth and differentiation. *Cold Spring Harbor Symp. Quant. Biol.* 389:401.
- Nishizuka, Y. 1992. Signal transduction crosstalk. *Trends Biochem. Sci.* 27:367.
- Northrup, S. H., and H. P. Erickson. 1992. Kinetics of protein-protein association explained by Brownian dynamics computer simulation. *Proc. Natl. Acad. Sci. USA.* 89:3338-3342.
- Pawson, T., and J. Schlessinger. 1993. SH2 and SH3 domains. *Current Biology* 3:434-442.
- Parkinson, J. S., and E. C. Kofoid. 1992. Communication modules in bacterial signaling proteins. *Annu. Rev. Genet.* 26:71-112.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. 1992. *Numerical Recipes in C*. Cambridge University Press, Cambridge, UK. 994 pp.
- Stoddard, B. L., H.-P. Biemann, and D. E. Koshland. 1993. Receptors and Transmembrane Signaling. *Cold Spring Harbor Symp. Quant. Biol.* 57:1-5.