

Win–stay, lose–shift in language learning from peers

Frederick A. Matsen* and Martin A. Nowak

Program for Evolutionary Dynamics, Departments of Mathematics and Organismic and Evolutionary Biology, Harvard University, One Brattle Square, Cambridge, MA 02138

Edited by Kenneth W. Wachter, University of California, Berkeley, CA, and approved November 12, 2004 (received for review September 7, 2004)

Traditional language learning theory explores an idealized interaction between a teacher and a learner. The teacher provides sentences from a language, while the learner has to infer the underlying grammar. Here, we study a new approach by considering a population of individuals that learn from each other. There is no designated teacher. We are inspired by the observation that children grow up to speak the language of their peers, not of their parents. Our goal is to characterize learning strategies that generate “linguistic coherence,” which means that most individuals use the same language. We model the resulting learning dynamics as a random walk of a population on a graph. Each vertex represents a candidate language. We find that a simple strategy using a certain aspiration level with the principle of win–stay, lose–shift does extremely well: stay with your current language, if at least three others use that language; otherwise, shift to an adjacent language on the graph. This strategy guarantees linguistic coherence on all nearly regular graphs, in the relevant limit where the number of candidate languages is much greater than the population size. Moreover, for many graphs, it is sufficient to have an aspiration level demanding only two other individuals to use the same language.

graph | population | random walk

Consider a population of individuals dispersed over a graph (Fig. 1). The task of these individuals is to find each other in an arbitrary node of the graph. In one time step, each individual can remain where it is or move to an adjacent node. The individuals have no global information about the graph or the configuration of the population. They cannot see each other at a distance. They can only determine whether other individuals are in the same node as themselves or in a different node. Which strategy could lead to success?

This question about the random walk of a population on a graph is motivated by the problem of language acquisition. Learning theory is a fascinating field at the interface of mathematics, computer science, and formal linguistics (1–7). The original studies (8) were motivated by Chomsky’s claim (9) that any given learning algorithm can only succeed on a limited set of grammars. Consider a teacher and a learner. The teacher provides sample sentences from a language. The learner tries to infer the underlying grammatical rules. It can be shown formally that no learner can solve this task if he entertains an unlimited set of hypotheses. In other words, each particular learning strategy can only learn a limited set of languages. This impossibility of *tabula rasa* learning has been shown in the Gold framework (1, 8), in statistical learning theory (2), in learning with queries (10), and in the framework of “probably almost correct” (4).

As outlined above, the traditional setting of formal language learning is an idealized teacher–student pair. In contrast, this paper considers a population of identical agents that are trying to find a common language. This approach is motivated by the fact that in reality children learn many aspects of language neither from their parents nor from other members of the previous generation, but from each other (11). A remarkable example of this is the development of Nicaraguan Sign Language, where a group of students developed a full language in a




Fig. 1. The basic operation of the model. The figure shows a possible sequence of events for a hypercube language space of 16 nodes (languages) and six aspiration learners using the $k = 2$ strategy. In the first frame, all individuals are using different languages. The black dots indicate that the node is occupied, and the number gives how many agents are occupying that node. The next frame shows the next time step: Because there are no languages with strictly more than two individuals at that language, each individual takes a random step along an edge of the graph. In the third frame, enough time has passed for three individuals to arrive at the same node. They will not move for the remainder of the simulation. The last frame shows the last three individuals collecting at this node. It is possible that these individuals could form a second cluster, but our analysis shows that under certain conditions a second cluster is unlikely when the graph is large.

short span of time (12, 13). We would like to model this process and find conditions for the population to develop linguistic coherence, which means that all of the individuals end up with the same language. Convergence to linguistic coherence in a population has been studied before in evolutionary and non-evolutionary settings (14–20).

Typically, in the language learning literature, one fixes a type of “learner” with certain capabilities and then establishes the class of languages this agent can learn under certain conditions. The learner in this paper is a type of “memoryless learner.” The memoryless learner has a finite set of possible grammars. Initially, the learner chooses one of these grammars as a hypothesis. The learner receives sentences. As long as the sentences are sufficiently compatible, the learner will remain with the current hypothesis. If a sentence is not compatible, the learner will jump to another hypothesis. The algorithm is memoryless, because the learner keeps no record and may return to a hypothesis that was previously rejected.

For such a learner, we need to decide on some sort of structure for the set of grammars: Which “jumps” are legal? A main concern of linguistic or cognitive research is to identify the learning algorithm that is used by human children and the

This paper was submitted directly (Track II) to the PNAS office.

See Commentary on page 17885.

*To whom correspondence should be addressed. E-mail: matsen@math.harvard.edu.

© 2004 by The National Academy of Sciences of the USA

underlying structure of candidate languages (universal grammar). The approach here is simple but general: we think of the set of possible grammars as a graph, where the nodes represent languages and the edges represent possible “jumps” from one candidate language to another. For example, if the set of languages is given by the complete graph, where all vertices are connected, then the learner can move from any one hypothesis to any other. If, on the other hand, the graph of languages is represented by a cycle, then an individual can move from its present language only to one of the two adjacent languages. The hypercube is a graph that lies between these two extreme cases: It has more edges than the cycle and fewer edges than the complete graph.

There is precedent for thinking of the set of languages as a graph in the “trigger learning algorithm” (21). In this framework, the set of grammatical hypotheses is generated by k binary parameters (22). At any one time, the learner holds a particular hypothesis given by a particular setting of these parameters. The learner does not change the hypothesis as long as the sentences are compatible. If a sentence arrives that is not compatible, then the learner will change one of the k parameters. In this model, the set of languages can be thought of being represented by the vertices of a k -dimensional hypercube. The process of language acquisition is described by the learner moving along the edges of this hypercube graph.

Let us consider N individuals initially randomly distributed on a language graph with n languages (nodes). Naturally, we assume $n \gg N$: There are many more possible candidate languages than people. At each time step, each individual is allowed to update its language. To do so, the individual can attempt to communicate with other individuals. The only information that can be inferred, however, is whether or not other individuals use the same language as the given individual. If an individual decides to change its current language, it can move to any adjacent language. In one version of the model, all individuals are updated synchronously. In another version, one individual is updated at each time step.

The goal is to find a single strategy that ensures (with high probability) that all individuals use the same language in the end. This state is called linguistic coherence. Thus, we want to find a mechanism that governs the random walk of a population on a graph with the aim that all individuals settle in the same place with high probability after a reasonably short time.

A natural strategy is the following: An individual stays with a certain language if some fixed number of other individuals use that same language. We call such agents “aspiration learners”: The learner aspires to use a language that is shared by a minimum number of other people. Because communication can be seen as an evolutionary game (14), the individual aspires to at least a certain level of payoff. In this case, payoff is coming from communicative success. The individual is happy when this goal is reached and stays with the current language. The individual shifts to another language when the goal is not reached. Thus, the strategy is one of win–stay, lose–shift (23).

Win–stay, lose–shift is a successful strategy in many different contexts. For example, it can outperform tit-for-tat in the repeated Prisoner’s Dilemma (23, 24). Strategies using this principle do well in noisy situations and even when a 2×2 game is randomly assigned to the players (25). It is a fundamental and simple rule of animal behavior in numerous situations (26, 27). The aspiration level determines what the agent considers a win versus a loss. The crucial question is what is the most successful aspiration level in a given situation (28). This is also the main question that we want to address here.

In our context, the aspiration level, k , is given by the minimum number of people with the same language such that the individual will stay with this language. Once this threshold is met or exceeded, the agents form a “cluster” that does not move




Fig. 2. The convergence rates of different k values with a fixed N on the complete graph. The line represents the value given by Eq. 1, and the marks represent data from simulations. One can see that the time is approximately proportional to n^k .

anymore. For example, an aspiration learner with $k = 1$ will keep its language if at least one other person uses this language. For a learner with $k = 2$, at least two other individuals are required. An obvious guarantee for success is if every individual uses an aspiration level, k , given by any integer $\geq N/2$. In this case, all individuals will continue to move until more than half of them use the same language. Then all others will converge to the same cluster. This strategy, however, will take a very long time to form the first cluster. Strategies with a lower aspiration level are faster to form the first cluster, but they admit the possibility that several independent clusters might be formed, which prevents convergence to complete coherence.

For this paper, we consider only memoryless learners whose set of possible grammars forms a “nearly regular” graph. If we define the “local degree” of a node to be the number of edges touching that node, we call a graph nearly regular when the maximum local degree is close to the minimum local degree. A regular graph is one in which these two numbers are equal, so all regular graphs are nearly regular. The hypercube, the complete graph, and the cycle are examples of regular graphs. Binary trees are examples of what we call nearly regular graphs. These ideas are formalized in *Appendix: Proofs*.

We found the following surprising result: If there are many more candidate languages than people ($n \gg N$), then $k = 3$ will almost always converge to complete coherence on any nearly regular graph. Moreover, for many nearly regular graphs, such as the hypercube, the complete graph, the balanced d -tree, or the torus of at least two dimensions, even $k = 2$ leads to complete coherence. Thus, even though each agent only demands to be able to communicate with a few others, in most cases the dynamics will lead to everyone speaking the same language. The proof of these statements are given in *Appendix: Proofs*.

Note that some structural assumption is needed on the class of graphs under consideration. For example, if we consider two star graphs connected by an edge, we will need a very large k to gain coherence. Our “nearly regular” assumption excludes these sorts of graphs.

We can estimate the time it takes to reach coherence. This time is the sum of two terms. The first term approximates how long it takes to form the first cluster. It is approximately the inverse of n times the probability that a cluster of size $k + 1$ forms at a given node, assuming that we begin with a uniform distribution. The second term describes the amount of time it takes

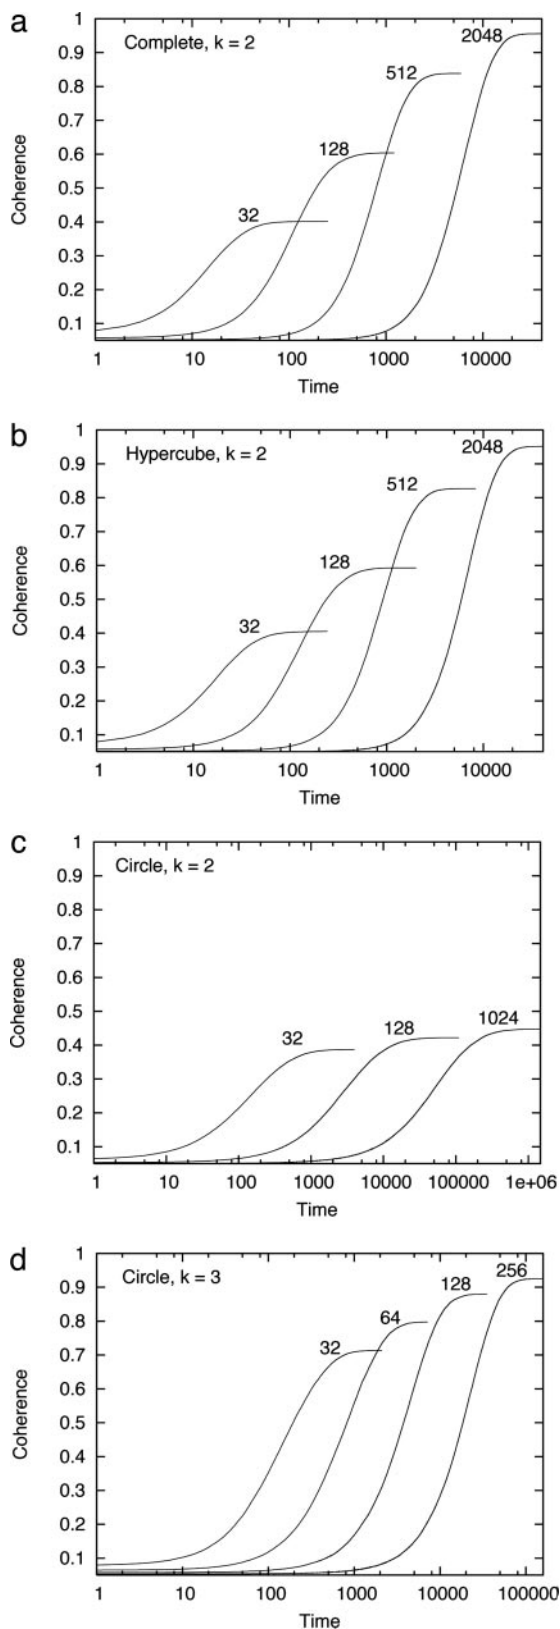


Fig. 3. Results from simulations of the model with a population size of $n = 20$. In *a* and *b*, we clearly see that $k = 2$ leads to coherence with high probability when n is large for both the complete graph and the hypercube. The curves follow a sigmoidal shape, which comes from the two time scales involved in the analysis. *c* shows that on the cycle, $k = 2$ does not lead to complete coherence with high probability. *d* shows that one gets the coherence result when $k = 3$ as predicted by the theory.

for all individuals to be absorbed into this cluster. This time depends explicitly on the graph structure. For the complete graph we estimate the total time to coherence as

$$T = \frac{1}{\binom{N}{k+1}} \frac{n^k}{(1-n^{-1})^{N-k-1}} + n \log(1 - 2^{-1/(N-k-1)}). \quad [1]$$

This is an excellent estimate, as shown in Fig. 2. Note that in the case $k \geq 2$ the first term dominates, and that for large n this term is $\approx n^k / \binom{N}{k+1}$. Therefore, although we may get a higher probability of complete coherence with a larger k , the individuals will take a much longer time to converge when n is large.

In Fig. 3, we show simulation results. We define the level of coherence as the probability that two randomly (with replacement) chosen individuals speak the same language (16, 29). The population size is $n = 20$. The expected coherence (averaged over 2,000 or more runs) is shown as a function of time. Fig. 3*a* shows the performance of the aspiration level $k = 2$ for complete graphs of sizes $n = 32, 128, 512$, and $2,048$. Similarly, Fig. 3*b* shows the performance of $k = 2$ on hypercubes of four different sizes. Fig. 3*c* and *d* shows the performance of $k = 2$ and $k = 3$ on cycle graphs. For large complete graphs and hypercubes, $k = 2$ is sufficient for the average coherence to converge to one as n increases. For cycles, $k = 2$ is not sufficient, but $k = 3$ is required.

The sigmoidal shape of the curves suggests that, on average, one has to wait a long time for the first cluster to form, but once this cluster has appeared coherence is reached relatively quickly. Eq. 1 describes these two time scales in the case of the complete graph.

In summary, we have modeled language learning from peers as the random walk of a population on a graph. The set of learnable languages is given by a graph, where each node represents a language (grammar) and each edge indicates a possible change of hypothesis during the learning process. We consider a memoryless learner with an aspiration level. The learner is satisfied if k others use the same language. In this case, the learner will stay with its current hypothesis. If less than k others use this language, then the learner will jump to a new hypothesis. We show that $k = 3$ guarantees success with high probability on all nearly regular graphs. Moreover, on many of those graphs, $k = 2$ is sufficient and leads to coherence more quickly. Thus, a simple strategy of win–stay, lose–shift is highly successful in language learning from peers. Our approach has applications beyond language acquisition to learning in other social or game theoretic settings (30).

Appendix: Proofs

In this section, we formalize the ideas of the paper and prove the propositions. We note that although the analysis is performed here for the case of synchronous updating (each agent may move each time step), the case of asynchronous updating (each time step a single randomly chosen agent is allowed to move) is similar and the corresponding statements are proven with minor modifications.

Let us assume G_1, G_2, \dots is a infinite sequence of graphs with at least n nodes, such that the local degree (the number of edges touching a given node) is bounded below by L_k and above by U_k on graph G_k . We call the sequence of graphs “nearly regular” if the ratio U_k/L_k is bounded above by a fixed β .

We assume such a sequence, and that the N agents start distributed on each graph according to a uniform probability distribution. We also assume that the agents all use strategy k . If the probability of converging to complete coherence on G_n goes to 1 as n goes to infinity, we will say that strategy k leads

to complete coherence in the limit for the given sequence. Note that as long as N is fixed, this does not depend on the value of N . We have the following result:

Proposition 1. Any k strategy with $k \geq 3$ leads to complete coherence in the limit for any nearly regular sequence of graphs.

Moreover, for many nearly regular sequences of graphs $k = 2$ leads to complete coherence in the limit. Whether $k = 2$ or $k = 3$ is needed depends on the mean hitting time weighted by the stationary distribution, which we write as

$$\tau_0 = \sum_{i,j} \pi_i \pi_j H_{ij}.$$

Here, H_{ij} is the hitting time, i.e., the expected time it takes to diffuse from node i to node j , and π_i is the stationary distribution for the graph diffusion. The quantity τ_0 is well known; it can be computed, for example, via the eigenvalues of a graph. For more details, see ref. 31.

Knowledge of this quantity allows us to make a stronger statement about certain types of graphs:

Proposition 2. Assume that G_1, G_2, \dots is a nearly regular sequence of graphs such that G_n has at least n nodes and such that we have a little- o bound on τ_0 :

$$\tau_0(G_n) \leq o(n^2). \quad [2]$$

Then the $k = 2$ strategy leads to complete coherence in the limit.

Note that such bounds exist for sequences of complete graphs, hypercubes, balanced d -trees, and tori of at least two dimensions (see <http://stat-www.berkeley.edu/users/aldous/RWG/book.html>).

Now, we begin the analysis proving these statements. From the uniform distribution of agents eventually one cluster will develop, say of size m . It is possible that more than one cluster will develop in a single time step, but that has strictly lower order of probability, and we ignore this case. From the single cluster either we eventually reach complete coherence or a second cluster will develop. We will show that the probability of this latter event is small for large n .

We fix a graph G and use the following notation. Let $1, \dots, n$ be the nodes of the graph, i.e., the languages under consideration. Let d_i be the local degree of node i . Recall that by the assumption of near regularity, d_i is bounded above by a U and below by an L . Let node one be the location of the first cluster and time $t = 0$ be the time of its formation. Let R_t be the probability that at time $t - 1$ there was not yet a second cluster. Let $s_{i,t}$ be the probability that an agent is at node i at time t assuming that there is not yet a second cluster.

The probability of a second cluster developing is equal to the probability that there is no second cluster already times the probability that a new cluster develops under this condition. Thus, we can write the probability of a second cluster of size q developing at time t in node i as

$$R_t \binom{N-m}{q} s_{i,t}^q (1 - s_{i,t})^{N-m-q}.$$

Therefore, the probability of getting a second cluster of size q at time t in an arbitrary node is bounded above by

$$K_{q,t} = R_t \sum_{i=2}^n \binom{N-m}{q} s_{i,t}^q (1 - s_{i,t})^{N-m-q}.$$

To get the desired results, one can simply compare the probability distribution $s_{i,t}$ with the probability distribution of a pure random walk (no absorption) with the same initial conditions. We call this corresponding probability $p_{i,t}$. We have then $s_{i,t} \leq p_{i,t}$

for all $i \geq 2$, which can be checked by looking at the corresponding matrices of the Markov chains.

The probability distributions are updated by

$$p_{i,t+1} = \sum_{j \in \text{adj}(i)} \frac{p_{j,t}}{d_j},$$

where $\text{adj}(i)$ means the nodes that are adjacent to i .

We claim that $p_{i,t}$ never gets above $d_i/(nL)$ or below $d_i/(nU)$. This is certainly true at time 0, and other times can be checked by induction using the above equation. Then, by the definition of β , we have

$$\frac{1}{\beta n} \leq p_{i,t} \leq \frac{\beta}{n}. \quad [3]$$

This gives an upper bound for $K_{q,t}$:

$$K_{q,t} \leq \sum_{i=2}^n \binom{N-m}{q} s_{i,t}^q \leq \sum_{i=2}^n \binom{N-m}{q} p_{i,t}^q \leq \binom{N-m}{q} \frac{\beta^q}{n^{q-1}}.$$

Because we are taking the limit as n goes to infinity, we ignore terms of higher order in $n - 1$. Therefore, we only consider $q = k + 1$ in what follows.

The number of steps until all agents are absorbed into one cluster or another is a random variable that is bounded above by the amount of time it would take for the same $N - m$ agents, all diffusing on the graph independently, to be absorbed by node one. We call the latter random variable T . Now, we can split up cases according to how long absorption takes:

$$\begin{aligned} P(\text{incoherence}) &\leq \sum_{r \geq 0} \left[P(T = r) \sum_{t=1}^r K_{k+1,t} \right] \\ &\leq \sum_{r \geq 0} \left[P(T = r) r \binom{N-m}{k+1} \frac{\beta^{k+1}}{n^k} \right] \\ &= E(T) \binom{N-m}{k+1} \frac{\beta^{k+1}}{n^k}. \end{aligned}$$

Now, say that S is the random variable representing time it takes for just one of these noninteracting agents to hit node number one. The bound

$$E(T) \leq N E(S)$$

can be gained by taking the derivative of the N th power of the cumulative distribution function of T . Furthermore, by an argument similar to others in this paper using Eq. 3, $E(S)$ can be bounded above:

$$E(S) \leq \frac{\beta^{2m+1}}{n^2} \sum_{i,j} H_{i,j}.$$

Also, since

$$\pi_i = \frac{d_i}{\sum_j d_j} \geq \frac{1}{\beta n},$$

we have

$$\frac{1}{n^2} \sum_{i,j} H_{i,j} = \beta^2 \sum_{i,j} (\beta n)^{-2} H_{i,j} \leq \beta^2 \sum_{i,j} \pi_i \pi_j H_{i,j} = \beta^2 \tau_0.$$

The final result is

$$P(\text{incoherence}) \leq \binom{N-m}{k+1} \frac{N\tau_0\beta^{2m+k+4}}{n^k}.$$

From this equation it is clear that upper bounds on the mean hitting time τ_0 are going to imply upper bounds on the probability of incoherence. In fact, this equation proves *Proposition 2*.

General bounds on τ_0 exist for nearly regular graphs. Kahn *et al.* (32) proved that the “cover time,” and therefore τ_0 , is at most of order $O(n^2)$. Substituting this bound into the above inequality proves *Proposition 1*, that $k \geq 3$ is sufficient to ensure that we get a single cluster in the limit of large n .

Now, we address the question of why on the cycle $k = 2$ is not sufficient to guarantee complete coherence in the limit. We present a plausibility argument that indicates why this value of k is important.

Assume that the first cluster is of size m . Modulo terms of higher order in the $s_{j,t}$, the probability of getting a second cluster is

$$\sum_t \left[R_t \sum_{j=2}^n \binom{N-m}{k+1} s_{j,t}^{k+1} (1-s_{j,t})^{N-m-k-1} \right].$$

For the purposes of this analysis, we only care whether this quantity goes to zero or not as the number of languages n goes to infinity with a fixed N . Through a series of simplifying steps, one can show that this is equivalent to the question of whether the sum

$$\sum_t \sum_{j=2}^n s_{j,t}^{k+1} \quad [4]$$

goes to zero or not.

It is well known that our Markov process on the cycle becomes one-dimensional Brownian motion in the limit (see, for example, ref. 33). Therefore, for large n we can approximate

$$s_{j,t} = F((j-1)/n, t),$$

where $F(x, t)$ is a solution to the diffusion equation

$$\frac{\partial F}{\partial t} = \frac{\partial^2 F}{\partial x^2}$$

with fixed boundary conditions $F(0, t) = 0$ and $F(1, t) = 0$. We take initial conditions $F(x, 0) = 1$ corresponding with the uniform distribution.

The analog of increasing n is increasing the size of the line, say by a multiplicative factor $q > 1$. Now, we wish to find another solution of the diffusion equation with the initial conditions $F(x, 0) = q^{-1}$, as we start with a uniform distribution on the interval $(0, q)$. The solution in this case is $q^{-1}F(x/q, t/q^2)$. The fact that time is scaled by the square of the space scaling factor is the crucial fact of this argument.

The continuous analog of Eq. 4 specialized to the case $k = 2$ is

$$\int_0^\infty \int_I (F(x, t))^3 dx dt. \quad [5]$$

On the larger interval, it will look like

$$\int_0^\infty \int_{qI} (q^{-1}F(x/q, t/q^2))^3 dx dt.$$

After making the usual variable substitutions, this becomes

$$q^{-3} \int_0^\infty \int_I (F(x, t))^3 q dx q^2 dt,$$

which is, of course, Eq. 5 again. Thus, increasing the size of space does not change the probability of multiple clusters in the case $k = 2$, and so we conclude that the probability does not go to zero as the line becomes long. Note that this argument depends very much on the value of k , for if $k = 3$, then increasing q decreases the value of the integral. This decrease corresponds to the fact that $k = 3$ will lead to coherence for large values of n .

We thank Igor Pak and Lorens Imhof. The Program for Evolutionary Dynamics is supported by Jeffrey Epstein.

- Osherson, D., Stob, M. & Weinstein, S. (1986) *Systems That Learn* (MIT Press, Cambridge, MA).
- Vapnik, V. N. (1998) *Statistical Learning Theory* (Wiley, New York).
- Vapnik, V. N. & Chervonenkis, A. Y. (1971) *Th. Prob. Appl.* **17**, 264–280.
- Valiant, L. G. (1984) *Commun. ACM* **27**, 436–445.
- Pinker, S. (1979) *Cognition* **7**, 217–283.
- Niyogi, P. (1998) *The Informational Complexity of Learning* (Kluwer, Boston).
- Poggio, T., Rifkin, R., Mukherjee, S. & Niyogi, P. (2004) *Nature* **428**, 419–422.
- Gold, E. M. (1967) *Inf. Control* **10**, 447–474.
- Chomsky, N. (1959) *Inf. Control* **2**, 137–167.
- Angluin, D. (1987) *Inf. Comput.* **75**, 87–106.
- Pinker, S. (1995) *The Language Instinct* (HarperPerennial, New York).
- Senghas, A. & Coppola, M. (2001) *Psychol. Sci.* **12**, 323–328.
- Kegl, J., Senghas, A. & Coppola, M. (1999) in *Language Creation and Language Change: Creolization, Diachrony, and Development*, ed. Degraff, M. (MIT Press, Cambridge, MA), pp. 179–237.
- Nowak, M. A. & Krakauer, D. (1999) *Proc. Natl. Acad. Sci. USA* **96**, 8028–8033.
- Nowak, M. A., Plotkin, J. B. & Jansen, V. A. A. (2000) *Nature* **404**, 495–498.
- Nowak, M. A., Komarova, N. L. & Niyogi, P. (2001) *Science* **291**, 114–118.
- Nowak, M. A., Komarova, N. L. & Niyogi, P. (2002) *Nature* **417**, 611–617.
- Kirby, S. (2000) in *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, eds. Knight, C., Hurford, J. & Studdert-Kennedy, M. (Cambridge Univ. Press, Cambridge, U.K.).
- Steels, L. (2003) *Trends Cognit. Sci.* **7**, 308–312.
- Hurford, J. & Kirby, S. (2001) in *Simulating the Evolution of Language*, eds. Parisi, D. & Cangelosi, A. (Springer, Berlin).
- Gibson, E. & Wexler, K. (1994) *Linguistic Inquiries* **25**, 407–454.
- Chomsky, N. (1981) *Lectures on Government and Binding: The Pisa Lectures* (de Gruyter, Dordrecht, The Netherlands).
- Nowak, M. A. & Sigmund, K. (1993) *Nature* **364**, 56–58.
- Kraines, D. & Kraines, V. (1993) *Theor. Decis.* **35**, 107–150.
- Posch, M. (1997) *Win Stay-Lose Shift: An Elementary Learning Rule for Normal Form Games* (Santa Fe Institute, Santa Fe, NM), <http://ideas.repec.org/p/wop/safire/97-06-056e.html>.
- Millinski, M. (1993) *Nature* **364**, 12–13.
- Domjan, M. & Burkhard, B. (1986) *The Principles of Learning and Behavior* (Brooks/Cole, Monterey, CA).
- Posch, M., Pichler, A. & Sigmund, K. (1999) *Proc. R. Soc. London Ser. B* **266**, 1427–1436.
- Komarova, N. L. & Nowak, M. A. (2003) *J. Theor. Biol.* **3**, 445–457.
- Fudenberg, D. & Levine, D. K. (1998) *The Theory of Learning in Games* (MIT Press, Cambridge, MA).
- Lovasz, L. (1993) in *Paul Erdős Is Eighty* (János Bolyai Mathematical Soc., Budapest), Vol. 2, pp. 1–46.
- Kahn, J., Linial, N., Nisan, N. & Saks, M. (1989) *J. Theor. Prob.* **16**, 121–128.
- Karlin, S. & Taylor, H. (1975) *A First Course in Stochastic Processes* (Academic, San Diego).