

Algorithms, Games, and Evolution

Erick Chastain^{*}, Adi Livnat[†], Christos Papadimitriou[‡] and Umesh Vazirani[‡]

^{*}Computer Science Department, Rutgers University, Piscataway, NJ 08854 USA, erick@cs.rutgers.edu., [†]Department of Biological Sciences, Virginia Tech, Blacksburg, VA 24061 USA, adi@vt.edu., and [‡]Computer Science Division, University of California at Berkeley, CA, 94720 USA, christos@cs.berkeley.edu, vazirani@cs.berkeley.edu.

Contributed by Christos Papadimitriou, April 2, 2014 (sent for review November 16, 2013)

Even the most seasoned students of evolution, starting with Darwin himself, have occasionally expressed amazement that the mechanism of natural selection has produced the whole of life as we see it around us. There is a computational way to articulate the same amazement: “What algorithm could possibly achieve all this in a mere three and a half billion years?” In this paper we propose an answer: We demonstrate that in the regime of weak selection, the standard equations of population genetics describing natural selection in the presence of sex become identical to those of a repeated game between genes played according to multiplicative weight updates (MWUA), an algorithm known in computer science to be surprisingly powerful and versatile. MWUA maximizes a trade-off between cumulative performance and entropy, which suggests a new view on the maintenance of diversity in evolution.

Significance statement: Theoretical biology was founded on the mathematical tools of statistics and physics. We believe there are productive connections to be made with the younger field of theoretical computer science, which shares with it an interest in complexity and functionality. In this paper, we find that the mathematical description of evolution in the presence of sexual recombination and weak selection is equivalent to a repeated game between genes played according to the multiplicative weight updates algorithm (MWUA), an algorithm that has surprised computer scientists time and again in its usefulness. This equivalence is informative for the pursuit of two key problems in evolution: the role of sex, and the maintenance of variation.

Weak selection | sex and recombination | coordination games | learning algorithms | multiplicative weights update

Precisely how does selection change the composition of the gene pool from generation to generation? The field of population genetics has developed a comprehensive mathematical framework for answering this and related questions [1]. Our analysis in this paper focuses particularly on the regime of weak selection, now a widely used assumption [2, 3]. Weak selection assumes that the differences in fitness between genotypes are small relative to the recombination rate, and consequently, through a result due to Nagylaki et al. [4], (see also [1] Section II.6.2), evolution proceeds near linkage equilibrium, a regime where the probability of occurrence of a certain genotype involving various alleles is simply the product of the probabilities of each of its alleles. Based on this result, we show that evolution in the regime of weak selection can be formulated as a repeated game, where the recombining loci are the players, the alleles in those loci are the possible actions or strategies available to each player, and the expected payoff at each generation is the expected fitness of an organism across the genotypes that are present in the population. Moreover, and perhaps most importantly, we show that the equations of population genetic dynamics are mathematically equivalent to positing that each locus selects a probability distribution on alleles according to a particular rule which, in the context of the theory of algorithms, game theory and machine learning, is known as multiplicative weight updates (MWUA). MWUA is known in computer science as a simple but surprisingly powerful algorithm (see [5] for a survey). Moreover, there is a dual view of this algorithm: each locus may be seen as selecting its new allele distribution at each generation

so as to maximize a certain convex combination of (a) cumulative expected fitness and (b) the entropy of its distribution on alleles. This connection between evolution, game theory, and algorithms seems to us rife with productive insights; for example, the dual view just mentioned sheds new light on the maintenance of diversity in evolution.

Game theory has been applied to evolutionary theory before, to study the evolution of strategic individual behavior (see, e.g., [6, 7]). The connection between game theory and evolution that we point out here is at a different level, and arises not in the analysis of strategic individual behavior, but rather in the analysis of the basic population genetic dynamics in the presence of sexual reproduction. The main ingredients of evolutionary game theory, namely strategic individual behavior and conflict between individuals, are extraneous to our analysis.

We now state our assumptions and results. We consider an infinite panmictic population of haplotypes involving several unlinked (i.e., fully recombining) loci, where each locus has several alleles. These assumptions are rather standard in the literature. They are made here in order to simplify exposition and algebra, and there is no a priori reason to believe that they are essential for the results, beyond making them easily accessible. For example, Nagylaki’s theorem [4], which is the main analytical ingredient of our results, holds even in the presence of diploidy and partial recombination.

Nagylaki’s theorem states that weak selection in the presence of sex proceeds near the Wright manifold, where the population genetic dynamics become (SI)

$$x_i^{t+1}(j) = \frac{1}{X^t} x_i^t(j) (F_i^t(j)),$$

where $x_i^t(j)$ is the frequency of allele j of locus i in the population at generation t , X is a normalizing constant to keep the frequencies summing to one, and $F_i^t(j)$ is the mean fitness at time t amongst genotypes that contain allele j at locus i (see [4] and the SI). Under the assumption of weak selection, the fitness of all genotypes are close to one another, say within the interval $[1 - \epsilon, 1 + \epsilon]$, and so the fitness of genotype g can be written as $F_g = 1 + \epsilon \Delta_g$, where ϵ is the selection strength, assumed here to be small, and $\Delta_g \in [-1, 1]$ can be called the differential fitness of the genotype. With this in mind, the equation above can be written

$$x_i^{t+1}(j) = \frac{1}{X^t} x_i^t(j) (1 + \epsilon \Delta_i^t(j)), \quad (1)$$

where $\Delta_i^t(j)$ is the expected differential fitness amongst genotypes that contain allele j at locus i (see Figure 1 for an illustration of population genetics at linkage equilibrium).

Reserved for Publication Footnotes

We now introduce the framework of game theory (see Figure 2 for an illustration) and the “multiplicative weight updates” algorithm (MWUA, see SI), studied in computer science and machine learning, and rediscovered many times over the past half-century; as a result of these multiple rediscoveries, the algorithm is known with various names across subfields: “the experts algorithm” in the theory of algorithms, “Hannan consistency” in economics, “regret minimization” in game theory, “boosting” and “winnow” in artificial intelligence, etc. Here we state it in connection to games, which is only a small part of its applicability (see the SI for an introduction to the MWUA in connection to the so-called *experts problem* in Computer Science).

A game has several players, and each player i has a set A_i of possible actions. Each player also has a utility, capturing the way whereby her actions and the actions of the other players affect this player’s well being. Formally the utility of a player is a function that maps each combination of actions by the players to a real number (intuitively denoting the player’s gain, in some monetary unit, if all players choose these particular actions). In general, rather than choosing a single action, a player may instead choose a *mixed* or *randomized* action, that is, a probabilistic distribution over her action set. Here we only need to consider coordination games, in which all players have the same utility function — that is, the interests of the players are perfectly aligned, and their only challenge is to coordinate their choices effectively. Coordination games are among the simplest games; the only challenge in such a game is for the players to “agree” on a mutually beneficial action.

How do the players choose and adjust their choice of randomized (mixed) actions over repeated play? Assume that at time t , player i has mixed action x_i^t , assigning to each action $j \in A_i$ the probability $x_i^t(j)$. The MWUA algorithm [5] adjusts the mixed strategy for player i in the next round of the game according to the following rule:

$$x_i^{t+1}(j) = \frac{1}{Z^t} x_i^t(j) (1 + \epsilon u_i^t(j)), \quad (2)$$

where Z^t is a normalizing constant designed to ensure that $\sum_j x_i^t(j) = 1$, so x_i^{t+1} is a probability distribution; ϵ is a crucial small positive parameter, and $u_i^t(j)$ denotes the expected utility gained by player i choosing action j in the regime of the mixed actions by the other players effective at time t . This algorithm (a) is known to converge to the min-max actions if the game is two-player zero-sum; (b) is also shown here to converge to equilibrium for the coordination games of interest in the present paper (see Corollary 5 in SI); (c) is a general “learning algorithm” that has been shown to be very successful in both theory and practice; and (d) if, instead of games, it is applied to a large variety of optimization problems, including linear programming, convex programming, and network congestion, it provably converges to the optimum quite fast.

It can be now checked that the two processes expressed in equations (1) and (2), evolution under natural selection in the presence of sex and multiplicative weight updates in a coordination game, are mathematically identical (see Theorem 3 in the SI). That is, the interaction of weak selection and sex is equivalent to the MWUA in a coordination game between loci in which the common utility is the differential fitness of the organism. The parameter ϵ in the algorithm, which, when small signifies that the algorithm is taking a “longer-term view” of the process to be solved (see SI), corresponds to the selection strength in evolution, i.e., the magnitude of the differences between the fitness of various genotypes.

The MWUA is known in computer science as an extremely simple and yet unexpectedly successful algorithm, which has surprised us time and again by its prowess in solving sophisticated computational problems such as congestion minimization in networks and convex programming in optimization. The observation that multiplicative weight updates in a coordination game is equivalent to evolution under sex and weak selection makes an informative triple con-

nection between three theoretical fields: evolutionary theory, game theory, and the theory of algorithms/machine learning.

So far we have presented the MWUA by “how it works” (informally, it boosts alleles proportionally to how well they do in the current mix). There is an alternative way of understanding the MWUA in terms of “what it is optimizing.” That is, we imagine that the allele frequencies of each locus in each generation is the result of a deliberate optimization by the locus of some quantity, and we wish to determine that quantity.

Returning to the game formulation, define $U_i^t(j) = \sum_{\tau=0}^t u_i^\tau(j)$ to be the cumulative utility obtained by player i by playing strategy j over all t first repetitions of the game, and consider the quantity

$$\sum_j x_i^t(j) U_i^t(j) - \frac{1}{\epsilon} \sum_j x_i^t(j) \ln x_i^t(j). \quad (3)$$

The first term is the current (at time t) expected cumulative utility. The second term of (3) is the *entropy* (expected negative logarithm) of the probability distribution $\{x_i^t(j), j = 1, \dots, |A_i|\}$, multiplied by a large constant $\frac{1}{\epsilon}$. Suppose now that player i wished to choose the probabilities of actions $x_i^t(j)$ ’s with the sole goal of maximizing the quantity (3). This is a relatively easy optimization problem, because the quantity (3) to be maximized is strictly concave, and therefore it has a unique maximum, obtained through the KKT conditions [8] (see Section 4 of the SI):

$$U_i^t(j) - \frac{1}{\epsilon} (1 + \ln x_i^t(j)) + \mu^t = 0.$$

(Here μ^t is the Lagrange multiplier associated with the constraint $\sum_j x_i^t(j) = 1$ seeking to keep the $x_i^t(j)$ ’s a probability distribution, see the SI.) Subtracting this equation from its homologue with t replaced by $t + 1$, and applying the approximation $\exp(\epsilon u_i^t(j)) \approx (1 + \epsilon u_i^t(j))$, we obtain the precise equation (2) (the normalization Z^t is obtained from μ^t and μ^{t+1} , see the SI for the more detailed derivation).

Thus, since equations (1) and (2) are identical, we conclude that, in the weak selection regime, natural selection is tantamount to each locus choosing at each generation its allele frequencies in the population so as to maximize the sum of the expected cumulative differential fitness over the alleles, plus the distribution’s entropy. Note that quantity (3) is maximized by genes, not by individuals, and that, interestingly, it is maximized with respect to current frequencies while being dependent (through U^t) on all past frequencies; and while there is some precedent to the use of “historical fitness” [9], its importance in this context is unexpected.

This alternative view of selection provides a new insight into an important question in evolutionary biology, namely: How is genetic diversity maintained in the presence of natural selection [10]? That the MWUA process enhances the entropy of the alleles’ distribution (while at the same time optimizes expected cumulative utility) hints at such a mechanism. In fact, entropy is enhanced inversely proportional to s (the quantity corresponding in the population genetics domain to the parameter ϵ), the selection strength: The weaker the selection, the more it favors high entropy. Naturally, entropy will eventually vanish when the process quiesces at equilibrium: One allele per locus *will* eventually be fixed, and in fact this equilibrium may be a local, as opposed to global, fitness maximum. However, we believe that it is interesting and significant that the entropy of the allele distribution is favored by selection in the transient; in any event, mutations, environmental changes, and finite population effects are likely to change the process before equilibrium is reached. This new way of understanding the maintenance of variation in evolution (selection as a trade-off between fitness and entropy maximization) is quite different from previous hypotheses for the maintenance of variation (e.g., [11, 12]). Another rather surprising consequence of this characterization is that, under weak selection, all past generations, no matter how distant, have equal influence on the change in the allele mix of the current generation.

Our discussion has focused on the evolution of a fixed set of alleles; that is, we have not discussed mutations. Mutations are, of course, paramount in evolution, as they are the source of genetic diversity, and we believe that introducing mutations to the present analysis is an important research direction. Here we focus on the selection process, which is rigorously shown to be tantamount to a trade-off, for each locus, between maximizing diversity and maximizing expected cumulative fitness.

We can now note a simple yet important point. Since multiplicative weight updates by the loci operates in the presence of sex, the triple connection uncovered in this paper is informative for the “queen of problems in evolutionary biology,” namely the role of sex in evolution [13, 14]. The notion that the role of sex is the maintenance of diversity has been critiqued [15], since sex does not always increase diversity, and diversity is not always favorable. The MWUA

connection sheds new light on the debate, since sex is shown to lead to a tradeoff between increasing entropy and increasing (cumulative) fitness.

The connection between the three fields, evolution, game theory, and learning algorithms, described here was not accessible to the founders of the modern synthesis, and we hope that it expands the mathematical tracks that can be traveled in evolution theory.

ACKNOWLEDGMENTS. We thank Noam Nisan and Monty Slatkin for helpful comments, and the visitors of the Spring 2014 program on Evolutionary Biology and Computation at the Simons Institute for the Theory of Computing for engaging discussions. Special thanks are due to Nick Barton for his constructive criticism of an earlier version, which led to several important improvements of our presentation, including the correction of a potentially confusing oversight. The research of CHP was supported by NSF Grant CCF-0964033.

1. Bürger R (2000) *The mathematical theory of selection, recombination, and mutation* (Wiley Chichester).
2. Nagylaki T (1993) The evolution of multilocus systems under weak selection. *Genetics* 134:627–647.
3. Nei M (2005) Selectionism and neutralism in molecular evolution. *Mol Biol Evol* 22:2318–2342.
4. Nagylaki T, Hofbauer J, Brunovský P (1999) Convergence of multilocus systems under weak epistasis or weak selection. *J Math Biol* 38:103–133.
5. Arora S, Hazan E, Kale S (2012) The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing* 8:121–164.
6. Maynard-Smith J (1982) *Evolution and the Theory of Games* (Cambridge University Press).
7. Weibull JW (1995) *Evolutionary Game Theory* (MIT Press).
8. Boyd SP, Vandenberghe L (2004) *Convex optimization* (Cambridge university press).
9. Leibler S, Kussell E (2010) Individual histories and selection in heterogeneous populations *PNAS* 107:13183–13188
10. Lewontin R, Hubby J (1966) A molecular approach to the study of genic heterozygosity in natural populations. ii. amount of variation and degree of heterozygosity in natural populations of *Drosophila pseudoobscura*. *Genetics* 54:595–609.
11. Dobzhansky T (1955) *A review of some fundamental concepts and problems of population genetics* (Cold Spring Harbor Laboratory Press), Vol. 20, pp 1–15.
12. Kimura M (1968) Evolutionary rate at the molecular level. *Nature* 217:624–626.
13. Bell G (1982) *The Masterpiece of Nature: the Evolution and Genetics of Sexuality* (University of California Press, Berkeley).
14. Barton NH, Charlesworth B (1998) Why sex and recombination? *Science* 281:1986–1990.
15. Otto SP (2009) The evolutionary enigma of sex. *Am Nat* 174:S1–S14.

Figure 1: The equations of population genetics formulated in the 1930s constitute the standard mathematical way of understanding evolution of a species by tracking the frequencies of various genotypes in a large population. In the simple example shown here, a haploid organism with two genetic loci A and B has three alleles in each of its two loci named A1, A2, A3 and B1, B2, B3 for a total of 9 genotypes. In (a) we show the *fitness* of each genotype, that is, its expected number of offspring. The fitness numbers shown in (a) are all close to each other, reflecting weak selection, where the individual alleles' contributions to fitness are typically miniscule. Initially, each genotype occurs in the population with some frequency; in this particular example these numbers are initially equal (b); naturally, their sum over all nine genotypes is 1 (frequencies are truncated to the fourth decimal digit). Figure (c) shows how the genotype frequencies evolve in the next generation: each individual of a given genotype produces a number of offspring that is proportional to its fitness shown in (a), and the resulting offspring inherits the alleles of its parents with equal probability (since we are assuming, crucially, sexual reproduction). The genotype frequencies in the next generation are shown in (c), calculated through the standard recurrence equations of population genetics. We also show in the margins of the table the allele frequencies, obtained by adding the genotype frequencies along the corresponding row or column. Ten generations later, the frequencies are as shown in Figure (d).

Figure 2: A simple *coordination game* is played by two players: the *row player*, who chooses a row, and the *column player*, who chooses a column. After the two players make a choice, they both receive (or both pay, in case of a negative entry) the same amount of money, equal to the number at the chosen row and column (a). Coordination

games are the simplest possible kind of a game, one in which the strategic interests of all players are completely aligned — that is to say, there is no conflict at all. They are of interest when it is difficult for the players to know these numbers, or to communicate and agree on a mutually beneficial combination (in this example, third row and second column). Notice that this particular coordination game is closely related to the fitness landscape shown in Figure 1(a): If P is a payoff in this game, the corresponding entry of Figure 1(a) is equal to $1 + \epsilon \cdot P$, where ϵ is a small parameter here taken to be 0.01. Suppose that each of the two players chooses each of the three options with some probability, initially $\frac{1}{3}$ for all (b); in game theory such probabilistic play is called a *mixed action*. How do we expect these probabilities to change over repeated play? One famous recipe is the *multiplicative weight updates algorithm*, in which a player “boosts” the probability of each option by multiplying it by $1 + \epsilon G$, where G is the expected amount of money this option is going to win the player in the current round of play, and ϵ is the same small parameter as above. For example, the second action of the row player has G equal to 2 (the average of 3, -1, and 4), and so the probability of playing the second row will be multiplied by 1.02. Then these weights are “renormalized” so they add up to one, yielding the marginal probabilities shown in (c). The probabilities after ten such rounds of play are shown in (d). Comparing now the numbers in the margins of Figures 1(d) and 2(d) we notice that they are essentially the same. This is what we establish mathematically in this paper: the two processes — repeated coordination games played through multiplicative updates, and evolution under weak selection — are essentially identical. This conclusion is of interest because the multiplicative weight updates algorithm is known in computer science to be surprisingly powerful.