

Evolution is exponentially more powerful with frequency-dependent selection

Artem Kaznatcheev^{1,2,*}

¹Department of Computer Science, University of Oxford

²Department of Translational Hematology & Oncology Research, Cleveland Clinic

*kaznatcheev.artem@gmail.com

Abstract

Valiant [1] proposed to treat Darwinian evolution as a special kind of computational learning from statistical queries. The statistical queries represent a genotype’s fitness over a distribution of challenges. And this distribution of challenges along with the best response to them specify a given abiotic environment or static fitness landscape. Valiant’s model distinguished families of environments that are “adaptable-to” from those that are not. But this model of evolution omits the vital ecological interactions between different evolving agents – it neglects the rich biotic environment that is central to the struggle for existence.

In this article, I extend algorithmic Darwinism to include the ecological dynamics of frequency-dependent selection as a population-dependent bias to the distribution of challenges that specify an environment. This extended algorithmic Darwinism replaces simple invasion of wild-type by a mutant-type of higher scalar fitness with an evolutionary game between wild-type and mutant-type based on their frequency-dependent fitness function. To analyze this model, I develop a game landscape view of evolution, as a generalization of the classic fitness landscape approach that is popular in biology.

I show that this model of eco-evo dynamics on game landscapes can provide an exponential speed-up over the purely evolutionary dynamics of the strict algorithmic Darwinism proposed by Valiant. In particular, I prove that the noisy-Parity environment – which is known to be not adaptable-to under strict algorithmic Darwinism (and conjectured to be not PAC-learnable) – is adaptable-to by eco-evo dynamics. Thus, the ecology of frequency-dependent selection does not just increase the tempo of evolution, but fundamentally transforms its mode.

The eco-evo dynamic for adapting to the noisy-Parity environment proceeds by two stages: (1) a quick stage of point-mutations that moves the population to one of exponentially many local fitness peaks; followed by (2) a slower stage where each ‘step’ follows a double-mutation by a point-mutation. This second stage allows the population to hop between local fitness peaks to reach the unique global fitness peak in polynomial time. The evolutionary game dynamics of finite populations are essential for finding a short adaptive path to the global fitness peak during the second stage of the adaptation process. This highlights the rich interface between computational learning theory, evolutionary games, and long-term evolution.

Evolution is an algorithm. As such, evolution is subject to the same constraints of computational complexity as all algorithms [1–3]. But evolution is not *any* algorithm. It is a kind of local search algorithm acting on the signal of fitness. By better understanding the kinds of algorithms that are evolutionary dynamics, we can better understand the power and limits of evolution in nature.

In this article, I will show that adding ecological interactions – conceived of as frequency-dependent selection – to Valiant [1]’s frequency-independent evolution has a qualitatively transformative effect on evolutionary dynamics. Ecology provides a fundamentally different computational resource for evolution that allows exponential speed-ups in the time it takes to adapt to certain environments. This is not a speed-up in the rate of evolution, but a fundamental shift in how evolution computes and thus which ‘problems’ it can ‘solve’: i.e. which kinds of environments populations can or cannot become well-adapted to. Ecology does not make evolution take steps faster along an exponentially long adaptive paths but instead ecology helps evolution find a short path. Thus, ecology does not just increase the tempo of evolution, but fundamentally transforms its mode. To demonstrate this, I will focus on a family of highly non-linear environments – the noisy-Parity problem defined more formally in Section 4 – that evolution on its own (i.e. without ecology) provably cannot adapt to.

Valiant [1]’s strict algorithmic Darwinism (described in Section 1) formalizes many of our intuitions about evolution and allows for the analysis of a richer set of evolutionary dynamics than most other models. But it lacks an important feedback mechanisms found in nature: eco-evolutionary feedback. I will introduce this feedback as part of the extended algorithmic Darwinism in Section 2.

From first glance, it might seem like ecology would not have a drastic effect on the adaptive power of evolution: especially in settings when ecological dynamics are much faster than evolutionary dynamics. This is probably why this feedback was often ignored or minimized in early models of the modern evolutionary synthesis. Even today, much of the work on eco-evo is motivated by settings where ecology and evolution are on similar timescales [4, 5]. If ecological interactions are on a much faster timescale than evolution, it seems like ecological dynamics can be “absorbed” into a single update of evolutionary dynamics. This might make one “step” of evolution faster or slower, but running faster in no particular direction or along an exponentially long path is not helpful. Thus, it feels intuitive that short-bursts of ecological interaction within a single step of evolution cannot fundamentally change what sort of environments populations can adapt to. This intuition is wrong.

Even if ecological dynamics happen on a much faster timescale than evolution, I will show that the feedback between ecology and evolution due to brief bursts of frequency-dependent selection can help eco-evolutionary dynamics find short adaptive paths that purely evolutionary dynamics cannot find. I will show this by concentrating on if populations can find the peak of the needle-in-a-haystack fitness landscape correspond to the Parity environment. In Section 4, I will define the Parity environment and rehearse a standard hardness argument to show that under strict algorithmic Darwinism, the population cannot adapt to this environment in polynomial time (regardless of which evolutionary dynamic it follows). But in Section 5, I will show that the simple mutation-limited strong-selection dynamics from Section 3 are sufficient for the population to adapt to the parity environment in polynomial time under extended algorithmic Darwinism.

From the perspective of computational learning theory, this exponential speed up from ecology is due to the ability of the ecological dynamics of the evolving population to bias the challenge/example distribution in such a way that creates a new second-order gradient for evolution to follow. This allows for a much richer set of concept classes to be learnable by evolution (or in more biological terminology, a richer set of environments that are adaptable-to). It remains to be seen just how much more computationally powerful eco-evolutionary dynamics can be in comparison to just evolutionary dynamics.

1 Strict algorithmic Darwinism: abiotic/prebiotic interactions

To represent what evolution can achieve just by natural selection on frequency-independent fitness, I will focus on Valiant’s model of strict algorithmic Darwinism [1, 6]. In my presentation, I will change some of the terminology to more closely correspond to biologists. In particular, ‘evolvability’ has a rich meaning in the biology literature that does not always correspond to Valiant [1]’s use of the word. So my primary change will be to use statements like ‘a family of environments is adaptable-to’ instead of ‘a concept class is evolvable’ and correspondingly I will refer to the model as ‘(strict) algorithmic Darwinism’ instead of ‘evolvability’. These are only changes in terminology for talking about the same mathematical model.

Abiotic environment $((D, f))$: Let us consider an abiotic world filled with a primordial soup of complex polymers. For simplicity, let us represent each possible polymer by some string $z \in Z = \{0, 1\}^n$. Let $D \in \Delta_{\{0,1\}^n}$ be the relative distribution of these polymers in the soup. We can think of each polymer z as a ‘challenge’ sampled from D . Let us suppose that for each polymer z there are exactly two ways $\{0, 1\}$ to interact with it and one of these ways is ‘correct’ and the other is ‘incorrect’ – where the correct way liberate useful energy and the incorrect way spoils the polymer. Let the ideal function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ hold the ‘correct’ response to each polymer. Thus, the pair (D, f) can be seen as a representation of a particular abiotic environment.

Genotypes (x) and phenotypes (g_x) : A primordial organism has to encounter these abiotic challenges and respond. We can summarize the organism’s responses as a (behavioral) phenotype $g_x : \{0, 1\}^n \rightarrow \{0, 1\}$ that is encoded by a genotype $x \in X = \{0, 1\}^n$. In general, there will be many particular token organisms that all have the same genotype x and phenotype g_x .

Fitness $(w_x$ and $\hat{w}_x)$: We can imagine a specific token organism of type y that is floating in this primordial soup and bumping into polymers as sampling a challenge $z \in D$. If the response of the organism to this challenge (as given by $g_x(z)$) is the ‘correct’ response (as given by $f(z)$) then the organism gets a small reproductive reward (if $g_x(z) = f(z)$) and otherwise it gets no reward (if $g_x(z) \neq f(z)$).

Of course, evolution itself does not act on tokens but rather on types [7, 8]. Evolution can only respond to statistical properties of many individuals (tokens) of the same type, each encountering many individual challenges sampled from D . Moreover, evolution cannot access arbitrary statistical properties of D – as Kearns [9]’s statistical query model allows – but instead has to use the particular statistical signal of (type) fitness. Formally, we can define $w_z^{(D,f)}$ as the true fitness of z in environment (D, f) :

$$w_x^{(D,f)} = \mathbb{E}_{z \sim D}(f(z) = g_x(z)) \tag{1}$$

which I will also write as w_x when the environment (D, f) is obvious from context. A real population, however, will not have direct access to w_x and instead will have to rely on taking the average fitness over a random sample $S \sim D$ of sM many challenges. I will use $\hat{w}_x^{(D,f)}$ to represent the empirical fitness estimate of the true fitness $w_x^{(D,f)}$. Throughout the analysis, it will be useful to upper bound the error between the empirical and true fitness value by a tolerance τ as $|\hat{w}_x^{(D,f)} - w_x^{(D,f)}| \leq \tau$.

The exact value of τ depends on the population size M and then number of challenges per token s . If we wanted to ensure a tolerance of τ with high probability then we should set $s \in \Theta(\frac{1}{\tau^2})$. Although the above is defined for an f that is error-free, with a reasonable number of extra random

queries, this tolerance can be ensured even if f is sometimes corrupted by noise (i.e. an organism having a small probability of getting a fitness bonus even if it responded to the challenge incorrectly or not getting fitness even if it responded correctly). [9] Thus, we can use the same analysis of adaptability for both standard environments and noisy environments.

Fitness landscape (w or W) view of strict algorithmic Darwinism: It is also possible to express strict algorithmic Darwinism in the long established metaphor of fitness landscapes [10] by setting $w(x) = w_x$. A point $x \in X$ is said to be a local peak if all its point-mutation (single bit flip) neighbours ($N(x)$) are non-improving (i.e. $\forall y \in N(x) w(x) \geq w(y)$) and a global peak if all points are non-improving. For multiplicative fitness measures – like number of offspring in a season – we can instead use $W^{(D,f)}(x) = K^{w^{(D,f)}(x)}$ for some constant $K > 1$ or equivalently $W^{(D,f)}(x) = \exp(\beta w^{(D,f)}(x))$ for some strength of selection β [11].

Adaptable-to families of environments ($(\mathcal{D}, \mathcal{F})$): Since algorithmic biology is after asymptotic results, we are usually concerned not with a particular abiotic environment (D, f) but a family of environments $(\mathcal{D}, \mathcal{F})$ defined over a class of possible distributions \mathcal{D} and set of potential ideal functions $\mathcal{F} = \{f_{n,s}\}$ (i.e. concept class) indexed by some natural number n , index s . Borrowing an idea from his earlier model of probably-approximately correct learning [12], Valiant [1] defines an environment $(\mathcal{D}, \mathcal{F})$ as *adaptable-to* if given an arbitrarily small probability $0 \leq \delta < 1/2$ and arbitrarily small approximation rate $0 \leq \epsilon < 1/2$ there exists an evolutionary dynamic that in time polynomial in n , $|s|$, $\ln 1/\delta$, and $1/\epsilon$ with probability $1 - \delta$ gets the population to a genotype x such that $\Pr_{z \sim D}(g_x(z) = f(z)) \geq 1 - \epsilon$ for any $D \in \mathcal{D}$ and any $f_{n,s} \in \mathcal{F}$.

In the language of fitness landscapes, a family of environments $(\mathcal{D}, \mathcal{F})$ is *adaptable-to* if evolutionary dynamics can find the global fitness peak or get to a type that has fitness ϵ close to it in the corresponding family of fitness landscapes. Or more formally, if for every $(D_n, f_{n,s})$ with probability $1 - \delta$ the population reaches a genotype y with $w^{(D_n, f_{n,s})}(y) \geq 1 - \epsilon$ in time polynomial in n , $|s|$, $\ln 1/\delta$, and $1/\epsilon$. Note that this is different from Kaznatcheev [2] and Kaznatcheev, Cohen, and Jeavons [3]’s focus on local peaks.

Feldman [13] showed that strict algorithmic Darwinism is computationally equivalent to a subset of Kearns [9]’s statistical query model that is restricted to making only correlational queries (CSQ). In other words, a family of environments $(\mathcal{D}, \mathcal{F})$ is *adaptable-to* if and only if the concept class \mathcal{F} is CSQ-learnable over the distributions $\mathcal{D} = \{D_n\}$. As I discuss in more detail in 4, this means that the parity function over the uniform distribution is not *adaptable-to*. Since the parity function is PAC-learnable, Valiant [1] uses this as an argument for why learning is more powerful than evolution.

2 Extended algorithmic Darwinism: biotic/ecological interactions

In this article, I extend the strict Algorithmic Darwinism of Valiant [1] to accommodate frequency-dependent fitness. Let me sketch this in terms of the prebiotic soup of early life.

Abiotic distribution (D_A): For Valiant, the distribution D of environmental challenges is not affected by the resident population: all hypothetical populations in the same world would experience the same distribution of challenges. Consider, for example, challenges as macromolecules in the prebiotic soup: we can imagine the action of the sun on the early chemistry of the Earth producing a consistent supply of random polymers for our organisms to encounter. This would be a sort of constant abiotically generated distribution D_A of challenges.

Full distribution with biotic component ($D(\rho)$): But the abiotic distribution is not the only source of polymers in the environment: there is also other organisms. Thus, we can imagine the distribution D as partitioned into two parts, an abiotic part D_A that is independent of other organisms in the environment and a biotic part $D_B(\rho)$ that is a function of the evolving population ρ . This gives us $D(\rho) = (1 - b)D_A + bD_B(\rho)$, where $b \in [0, 1]$ measures the strength of the biotic component. In this extended algorithmic Darwinism, Valiant’s strict model would correspond to $b = 0$. As such, I focus the rest of the discussion on a family of landscapes where b is a small but polynomial fraction (i.e. $b \approx 1/n^{O(1)}$).

In a general application, the space of possible challenge strings Z and the space of genotypes X might not be the same. In that case, $D_B : \Delta_X \rightarrow \Delta_Z$ needs to map between distributions on these two spaces. We can view this mapping as specifying the ecological part of the evolutionary algorithm, since it would encode what kind of challenges certain organisms tend to present to other organisms. But in this report, I will focus on when $X = Z = \{0, 1\}^n$ and set D_B to be identity thus making $D(\rho) = (1 - b)D_A + b\rho$. In the case of the protocells in a prebiotic soup, we can imagine $D(\rho)$ implemented as all protocells trying to constantly copy their own genotypes but due to their primordial nature and poor lipid membranes, constantly shedding some of these copies without always successfully duplicating. When one cell floats past another, it then encounters this cloud of shed copies as new challenges. Thus, this would bias the abiotic distribution D_A towards the distribution of protocells ρ by an amount b that will depend on various aspects of the environmental and population structure.

Game landscapes (ω): For the sake of this article, I will focus on matrix games as a model of frequency-dependent fitness. For strict algorithmic Darwinism, we could have imagined the w of a fitness landscape as a very long vector in \mathbb{R}^X with entries indexed by $x \in X$ and specified by Equation 1. Similarly, for the extended algorithmic Darwinism, I will build the (matrix) game landscapes by considering a very large matrix $G \in \mathbb{R}^{X \times X}$ with entries given by:

$$G_{xy} = \mathbb{E}_{z \in D_B(e_y)}(g_x(z) = f(z)) \quad (2)$$

where $e_y \in \Delta_X$ is the distribution with all weight at type y (i.e. the unit vector in direction y). When $X = Z$ and D_B is the identity, we can specify fitness functions for the game landscape as:

$$\omega_x^{D,f}(\rho) = (1 - b)w_x^{(D_A,f)} + b \sum_{y \in X} G_{xy}\rho_y \quad (3)$$

Darwinian engine that powers evolution: I visualize this whole model as the Darwinian engine that powers evolution in Figure 1. The Darwinian engine is made of two cycles that together change the distribution of genotypes. On the top, we have the genesis of new variants via the mutation cycle. On the bottom, we have the struggle for existence via the development-ecology-selection cycle. The game landscape acts as a summary of the contribution of development (the mapping from genotype to phenotype) and ecology (from the distribution of phenotypes to fitness) that I described in this section. This is the heart of algorithmic Darwinism and is what specifies the environment that the population is adapting to – the problem instance, for a computer scientist. The other two arrows (selection and mutation) specify the algorithm of evolution. If we want to prove intractability results then we want to reason about arbitrary (polynomial time computable) selection and mutation functions. But since I aim for a surprising tractability result, I will describe in Section 3 the specific evolutionary algorithm of strong-selection weak-mutation dynamics [14, 15].

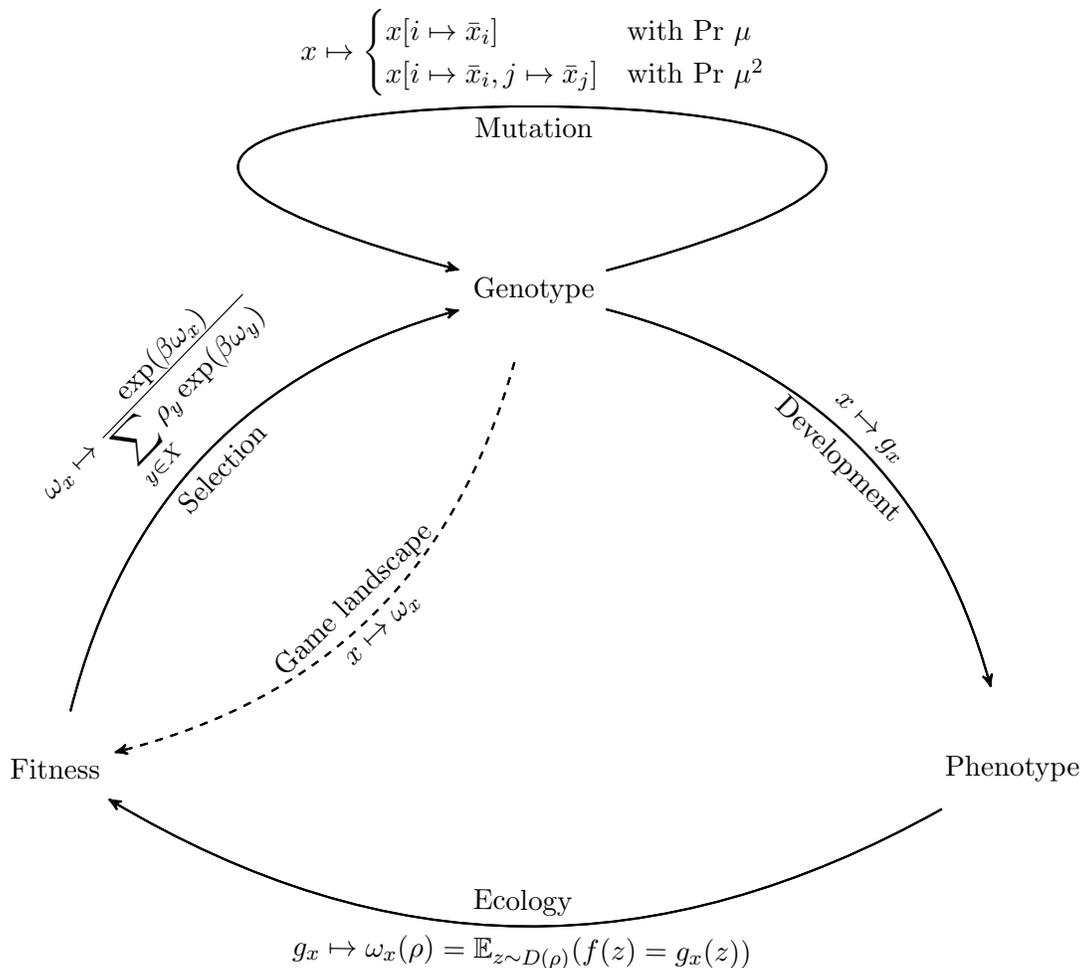


Figure 1: **Darwinian engine powering eco-evolutionary dynamics.** The top cycle captures the genesis of new variants and the bottom cycle captures the struggle for existence. Each edge is labeled by both its conceptual role (in text) and its specific realization in the ecologically extended algorithmic Darwinism (in equations). The two edges (selection and mutation) pointing into the genotype node specify the algorithm of evolution that adapts the population. In general, we could consider arbitrary polynomial time algorithms for these edges, but for the tractability of the Parity environments, I focus on the specific algorithm of strong-selection weak-mutation dynamics in a Moran process with single and double mutations as presented in Section 3. The other edges (development, ecology, and the out-direction of mutation) specify the environment (or ‘problem instance’) that the population is adapting-to. I give a general overview of this in Section 2 and discuss the specific family of Parity environments in Section 4.

3 Mutation-limited dynamics for game landscapes

In general, the strict and extended algorithmic Darwinism allow for a rich set of possible evolutionary dynamics. And when we prove hardness results in these frameworks, we establish that no evolutionary dynamic compatible with the framework is capable of adapting-to some environment. However, when we are establishing a tractability result, it is useful to consider a particular evolutionary dynamic that is simple and biologically plausible.

Monomorphic and briefly polymorphic populations: In the case of traditional fitness landscapes, for the sake of analysis, biologists often consider a mutation-limited population [15]. This population is nearly always monomorphic (i.e. made up of a number of tokens that are all of a single type) except when a new mutant is invading, at which point the population is briefly polymorphic between the resident and invader type. The invading mutant is usually a point mutant of the resident type one on timescales on the order of $1/\mu$ where $\mu \ll 1$ is the small mutation rate. If the resistant cannot be invaded by any point-mutant then it is possible for a double (or more distant) mutant to arise, although the timescales for these events scale as $(1/\mu)^2$ where $\mu \ll 1$ is the small mutation rate. After a successful invasion, the population is again monomorphic. This corresponds to the extreme limit when ecological feedback can act only briefly during a single step of evolutionary dynamics.

In the case of a traditional fitness landscape, strong selection mutation-limited dynamics only allows a genotype y to invade and replace a genotype x if $w(y) \geq w(x)$. In strict algorithmic Darwinism, a mutant invades if it is strictly fitter ($w(y) - w(x) > t$) or nearly-neutral [16] ($|w(y) - w(x)| \leq t$). But the story is more interesting for a game landscape and – instead of a simple threshold – it becomes useful to explicitly study a finite population of size M .

Fitness of resident (ω_x) and mutant type (ω_y) as a function of number of mutants (M_y): Since we are focused on the limit where at most two genotypes $x, y \in X$ co-exist at a time, we can track the state of the population with a single integer $M_y \in [M]$ for the number of invaders. For simplicity, I will assume a well mixed population so that every protocell has the same chance of interaction with every other protocell. This means that in a polymorphic population, a cell of type x has a probability of $M_y/(M-1)$ of interaction with a cell of type y and a probability of $(M-M_y-1)/(M-1)$ of interaction with a cell of type x . Similarly, a cell of type y has a probability of $(M_y-1)/(M-1)$ of interaction with a cell of type y and a probability of $(M-M_y)/(M-1)$ of interaction with a cell of type x . This transforms Equation 3 as a function of M_y into:

$$\omega_x = (1-b)w_x + b \overbrace{\left(G_{xx} \frac{M-M_y-1}{M-1} + G_{xy} \frac{M_y}{M-1} \right)}^{u_x} \quad (4)$$

$$\omega_y = (1-b)w_y + b \underbrace{\left(G_{yx} \frac{M-M_y}{M-1} + G_{yy} \frac{M_y-1}{M-1} \right)}_{u_y} \quad (5)$$

When analyzing evolutionary games, it is useful to look at the (whole) gain function ($\gamma_{yx} = \omega_y - \omega_x$) of switching from x to y [17, 18] – also known as the invasion fitness:

$$\gamma_{yx}(M_y) = (1-b)(w_y - w_x) + b \underbrace{\left((G_{yy} - G_{xy}) \frac{M_y-1}{M-1} + (G_{yx} - G_{xx}) \frac{M-M_y}{M-1} + \frac{G_{xx} - G_{xy}}{M-1} \right)}_{u_y - u_x} \quad (6)$$

and when $w_y = w_x$, it is useful to focus on just the component correspond to the game gain function ($u_y - u_x$). Since f has a boolean range, we have that $G_{xx}, G_{xy}, G_{yx}, G_{yy} \in \{0, 1\}$ and thus only 16 possible game gain functions, as shown in Figure 2. These two strategy matrix games are well studied in evolutionary game theory and can even be measured directly in microscopic experimental systems [8, 19].

Fixation probability of mutant against resident (ϕ_{yx}): Let us use the Moran process [11, 20] based on fitnesses ω_x and ω_y to update the number of invaders M_y . The Moran process is a standard model of invasion of a wild-type by a mutant-type in evolutionary biology, and proceeds as follows: (1) Randomly select an individual m proportional to its multiplicative fitness $e^{\beta\omega}$ [11]; (2) Uniformly select an individual c to be replaced; (3) Replace c by m and repeat the process until the population is monomorphic. This defines a Markov chain on the state space $[M]$ with transition probabilities given by:

$$\Pr\{0 \mapsto 0\} = \Pr\{M \mapsto M\} = 1 \quad (7)$$

$$\Pr\{M_y \mapsto M_y - 1\} = \frac{(M - M_y)e^{\beta\omega_x(M_y)}}{M_y e^{\beta\omega_y(M_y)} + (M - M_y)e^{\beta\omega_x(M_y)}} \frac{M_y}{M} \quad (8)$$

$$\Pr\{M_y \mapsto M_y + 1\} = \frac{M_y e^{\beta\omega_y(M_y)}}{M_y e^{\beta\omega_y(M_y)} + (M - M_y)e^{\beta\omega_x(M_y)}} \frac{M - M_y}{M} \quad (9)$$

$$\Pr\{M_y \mapsto M_y\} = 1 - \Pr\{M_y \mapsto M_y - 1\} - \Pr\{M_y \mapsto M_y + 1\} \quad (10)$$

From this, we can calculate the fixation probability ϕ_{yx} of y (i.e. the probability that the above Markov chain starting at $M_y = 1$ ends at the absorbing state of $M_y = M$) as:

$$\phi_{yx} = \left(-\beta \sum_{k=0}^{M-1} \exp\left(\sum_{M_y=1}^k \gamma_{yx}(M_y) \right) \right)^{-1} \quad (11)$$

Strong selection: I will focus on strong selection ($\beta > M$) and when the game contributes a detectable amount of the fitness ($t \ll \frac{b}{M}$). In that case, it is easy to classify the possible game regimes of the extended algorithmic Darwinism in Figure 2 into 4 classes: (1) In the green region, invasion is very likely: $\phi_{yx} > \phi^*$ for a constant ϕ^* independent of the rest of the parameters; (2) In the yellow region, invasion is possible but not likely (the neutral and nearly-neutral mutation regime): $\phi_{yx} = 1/M$; (3) In the red region, invasion is effectively impossible: $\phi_{yx} < e^{-M}$; (4) In the dotted red region, invasion is likely but the fixation time is exponential in M .

With a fitness landscape, it can be helpful to replace the numeric structure by just the structure of allowed adaptive moves to get the fitness graph [2, 3, 21, 22]. Similarly, with a game landscape, we can imagine it as a set of nodes $\{0, 1\}^n$ connected by the four types of edges outlined above. If type (4) edges are not encountered, my separation of time-scales between ecology and evolution – much like similar separations in the adaptive dynamics literature [5, 23, 24] – allows us to continue to view a population as occupying a single vertex. But unlike the simple continuous landscapes of adaptive dynamics, game landscapes allow us to also preserve the rich combinatorial structure and discrete mutations that in fitness landscapes allows for computationally hard environments [2].

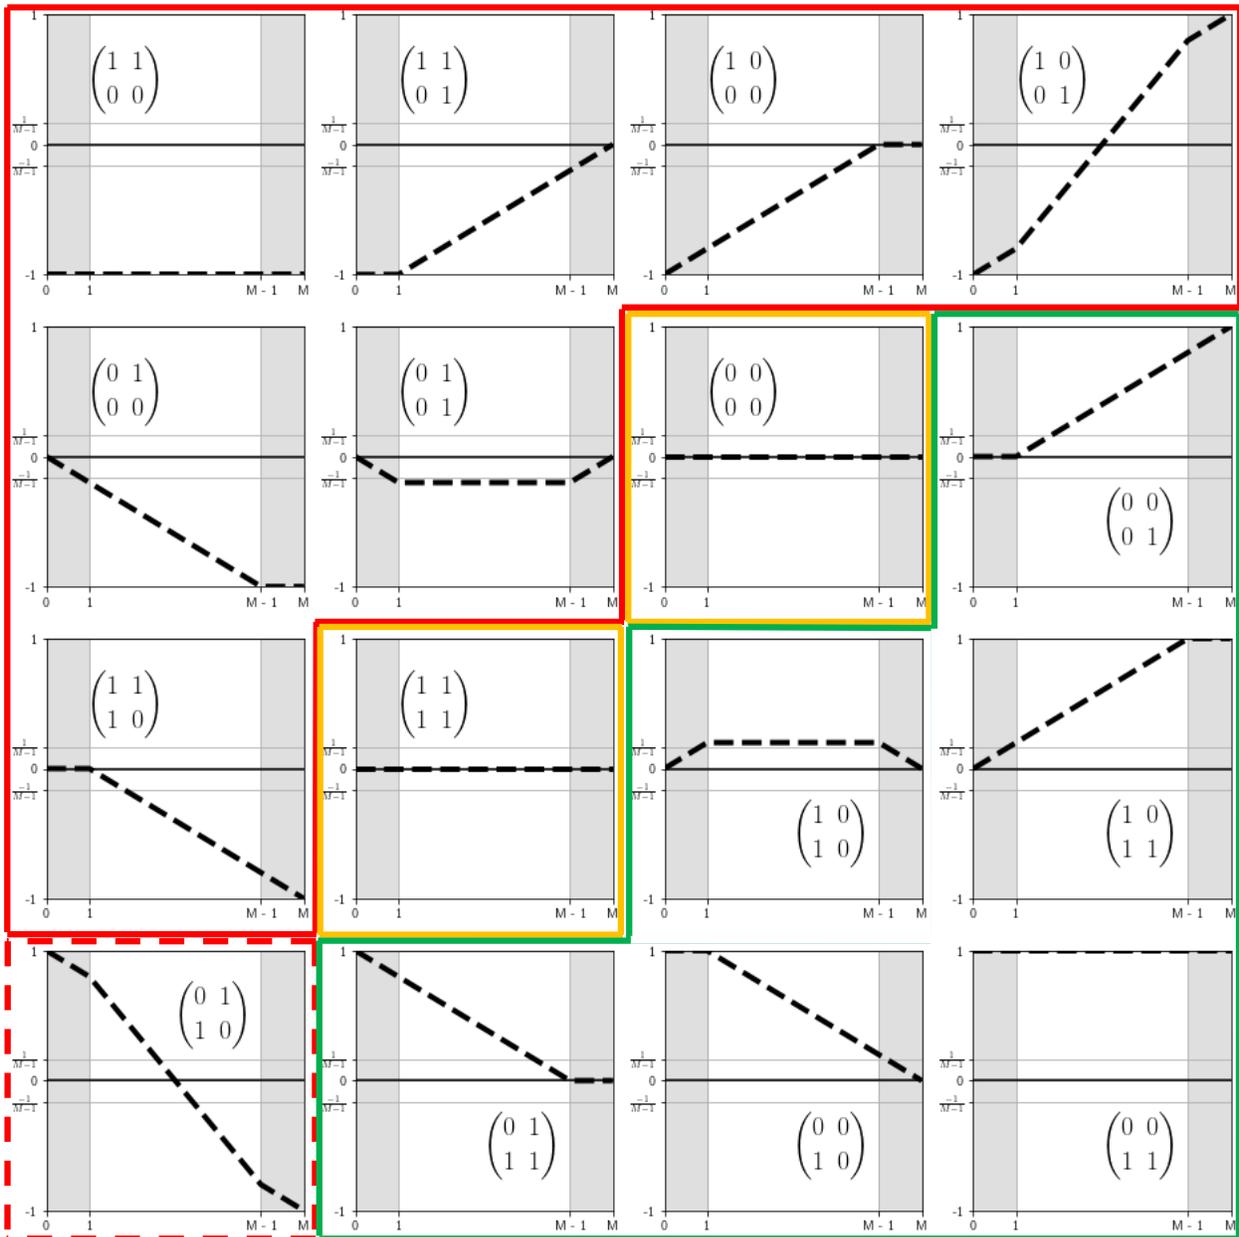


Figure 2: Possible game gain functions for game landscapes of a boolean environment: For each graph, the x-axis is the number (M_y) of invaders of type y and the y-axis is the game gain function ($u_y - u_x$). The game of each graph is inset as the matrix $\begin{pmatrix} G_{xx} & G_{xy} \\ G_{yx} & G_{yy} \end{pmatrix}$. The possible invasion dynamics under strong selection ($\beta > M$) are given by the colour-coding: green means y can invade x with probability greater than some constant p^* ; yellow means y can invade x through random drift with probability $1/M$; red means y cannot invade x (solid red: nearly 0 probability of fixation; dotted red: high probability, but exponentially long time to fixation).

4 Representing parity environments as fitness landscapes

The formal approach of algorithmic Darwinism allows us to view evolution as an algorithm. Since evolution is an algorithm, it is subject to all the same information- and computation-theoretic constraints that limit all algorithms [1–3]. This allows us to identify families of fitness landscapes or environments that cannot be adapted-to. I will call such environments or fitness landscapes ‘hard’ [2]. The canonical example of such a hard environment for strict algorithmic Darwinism is the Parity environment [1].

Parity environments ($f_{n,s}$): For each size n , the challenges come from the uniform distribution U_n : i.e. every string in $\{0,1\}^n$ has an equal probability of $1/2^n$ of being encountered – every string is as likely as every other string to be encountered. Each environment of size n is indexed by a hidden set S of salient bits or – equivalently – a string $s \in \{0,1\}^n$ that has a $s_i = 1$ if $i \in S$ and 0 otherwise. To succeed, organisms have to pay very careful attention to all the positions in S and ignore all other positions. The correct response to a challenge $z \in \{0,1\}^n$ is to return the parity of the bits z_i for $i \in S$ (i.e. $f_{n,s}(z) = \sum_{i \in S} z_i s_i \pmod{2}$). I will call the part of the genotype that overlaps with the hidden set of salient bits ($x[S]$) the *head* of x , and the other part the *tail* ($x[n - |S|]$). A genotype x has an *even-tail* if it has an even number of 1s in the tail (i.e. with index in $[n] - S$) and an *odd-tail* otherwise. Finally, we can set our tolerance parameter τ to account for misclassification noise in f , so without loss of generality, I can treat $f_{n,s}$ as not noisy.

Although Parity might not seem like a very biologically plausible environment, it actually idealizes an important problem that many organisms are likely to face. The Parity environment, divides the sensory/informational world of the organisms into two parts: an important part S and an irrelevant part $[n] - S$. The correct response to an environmental challenge is extreme sensitive to the important part (with any one bit flipping the output) and completely insensitive to the irrelevant part (with the output unaffected by bit flips). But the irrelevant part still gives the organisms opportunities to attend to it (thinking that it matters to what should be done when it doesn’t) and by so doing potentially make mistakes. The difficulty of the Parity environment then is not in the computation of the exclusive OR of the various bits (which is not common in biological systems, but certainly possible for a single cell [25, 26]), but in separating the set of important stimuli from the set of irrelevant stimuli. Thus, in the context of simple biological systems with minimal cognition [27], the Parity environment can be seen as an extreme version of the problem of selective attention.

Parity is not adaptable-to under strict algorithmic Darwinism: After defining SQ-learning, Kearns [9] showed that Parity over the uniform distribution is not efficiently learnable with statistical queries. Since any family of environments that is adaptable-to by strict algorithmic Darwinism has to be an SQ-learnable concept class, [1] gave Parity environments as the prototypical hard family of environments: a family of environments that are not adaptable to by any evolutionary dynamics compatible with strict algorithmic Darwinism. Finally, Parity is PAC-learnable but noisy-Parity is conjectured to not even be PAC-learnable [28, 29].

Representing reaction norms as genotypes: The above hardness result for strict algorithmic Darwinism is true regardless of how the reaction norms g are encoded as genotypes. But for the sake of an intuition about this hardness result, and as preparation for the easiness result in Section 5, I want to consider a specific simple representation of reaction norms as genotypes. Let the genotype space be the same as the example space: $X = Z = \{0,1\}^n$. With a genotype $x \in \{0,1\}^n$ servings

as the indicator set for the parts of the environment that the organism pays attention to, giving the reaction norm phenotype $g_x(z) = \sum_{i=1}^n x_i z_i \pmod 2$.

Parity fitness landscape (w^s): This representation of reaction norms by genotypes specifies a fitness landscape (w^s) corresponding to the environment $(U_n, f_{n,s})$ that shows the intuition for why the Parity environment is hard: fitness simply does not provide much information about s . For a genotype $x = s$ we have $w^s(s) = 1$, but for any other genotype $x \neq s$, we have that $w^s(x) = 1/2$. This is because given any string x , there are exactly 2^{n-1} strings that share an even number of bits in common with x and also 2^{n-1} strings that share an odd number of bits in common with x .

Thus, the fitness landscape is a needle-in-the-haystack landscape: one giant lowland at fitness $1/2$ with a single peak hiding at genotype s out of the 2^n possible genotypes. In this case, where the fitness landscape is not providing any information, mutation limited strong selection dynamics of the kind outlined in Section 3 will with high probability take an exponential number of steps.

5 Parity is adaptable-to under extended algorithmic Darwinism

Note that when $x = s$, $w_x^s = 1$ and for all $y \in \{0, 1\}$, $G_{xy}^s = 1$: the game landscape also has a single unique global fitness peak at s . To see how evolution on the Parity game landscape finds the s peak, we have to focus on how a mutation-limited strong selection dynamics navigates the great low fitness plateau where for all $x, y \in \{0, 1\} - \{s\}$, $w_x^s = w_y^s = 1/2$. Thus, we can focus on just the biotic contribution and the game gain functions from Figure 2 account for every possible invader. This lets us analyze the evolutionary dynamics of adapting to Parity. These dynamics will proceed by two stages, a fast stage where each step happens on the time-scale of $1/\mu$ (Section 5.1) and a slow stage where each step happens on the time-scale of $1/\mu^2$ (Section 5.2).

5.1 Quick convergence to an evolutionary stable set under point-mutations

Even in the monomorphic case, there is a deformation as we sift from a fitness to a game landscape. In particular, we see the great low-fitness plateau at fitness $1/2$ split into two parts. The high fitness even-tail genotypes (that have $G_{xx} = 1$) remain at fitness $1/2$, but the set of odd-tail genotypes (that have $G_{xx} = 0$) form a valley at fitness $(1 - b)/2$. Since any odd-tail genotype is adjacent to $[n] - S$ many even-tail genotypes, this is indeed a valley and not a plateau since any x with an odd-tail can be invaded by any adjacent y with the same head $x[S] = y[S]$ but an even-tail. Thus, the odd-tailed strings are not evolutionary stable under points mutations. But more importantly, this also means that once a string becomes even-tailed, it will not go back to being odd-tailed under point-mutations. So at most one point-mutation will fix in the tail of the string.

What about mutations in the head? Consider genotypes x and y with same parity tail such that there is exactly one $i \in S$ such that $0 = x_i \neq y_i = 1$ (without loss of generality, by relabeling). Since x and y have the same parity tail, we have that $G_{xx} = G_{yy}$ but $G_{xy} \neq G_{yx}$. Here there are two cases to consider: the string is still odd-tailed or it has become even-tailed. For an odd-tailed string, we have $G_{xy} = 1 > 0 = G_{yy} = G_{yx}$, so x can invade y and thus decrease the number of 1s in the head. But if x and y are even-tailed then y can invade x since $G_{xy} = 0$ but $G_{yx} = 1$ (this also means, unsurprisingly, that x cannot invade y). Thus, for even-tailed strings, point mutations in the head will only fix if they increase the number of 1s in the head.

Evolutionary stable set under point-mutations: From this, we can conclude that the only strings that are evolutionary stable under point-mutations have a full-head and even-tail. I will call

this set S_1 . In the space of point-mutations, this evolutionary stable set is a collection of isolated local fitness peaks. More importantly, we will converge to a point in this evolutionary stable set quickly. In the worst case, it will take $2|S| + 1$ many fixations if the population starts with a full-head, odd-tail string and then loses every 1 in the head ($|S|$ fixations), flips to even-tailed, and then regains every 1 in the head. On average, it will take fewer than $\frac{|S|}{2}(1 + \frac{1}{|S|} + \frac{1}{n-|S|})$ fixations.

5.2 Biased random-drift in the evolutionary stable set by double-mutations

Although points in the full-head-even-tail set S_1 are evolutionary stable and isolated under point-mutations (i.e. strict local peaks), they are not isolated under double-mutations. Hence, after the population fixes at a point-mutation peak in S_1 , we need to focus on the dynamics of double-mutations. In terms of timescales, in a mutation limited population with mutation rate μ , this means switching from a timescale of $\Theta(\frac{1}{\mu})$ to $\Theta(\frac{1}{\mu^2})$. There are three kinds of bits of a string in S_1 : 1s in the head, 1s in the tail, and 0s in the tail. Let $2k$ be the even number of 1s in the tail. Call the double mutant genotype y .

Double mutation in head: Suppose that both mutations are in the head. This happens with probability $\frac{|S|(|S|-1)}{n^2}$. Then $G_{yx} = G_{xx} = G_{yy} = G_{xy} = 1$ and thus y can only fix by random drift (i.e. with probability $1/M$). But after it fixes, we have an even-tailed population y outside of S_1 and so it will be quickly returned to S_1 by point-mutations without changing the tail. Thus, this double-mutation does not change our position in S_1 .

Double mutation in tail: Consider that both mutations are in the tail. This happens with probability $\frac{(n-|S|)(n-|S|-1)}{n(n-1)}$. Since this preserves tail parity: $G_{yy} = 1$. If one mutation flips 0 to 1 and the other flips 1 to 0 then $G_{xy} = G_{yx} = 0$ and thus y will be selected against and unable to fix. This happens with conditional probability $\frac{4k(n-|S|-2k)}{(n-|S|)(n-|S|-1)}$. The remaining two cases are if both tail mutations are 0 to 1 or both are 1 to 0. These happen with conditional probabilities $\frac{(n-|S|-2k)(n-|S|-2k-1)}{(n-|S|)(n-|S|-1)}$ and $\frac{2k(2k-1)}{(n-|S|)(n-|S|-1)}$, respectively. In these two cases, $G_{yx} = G_{xx} = G_{yy} = G_{xy} = 1$ and so y can fix only by random drift (with probability of $1/M$). But unlike the double-head case, the resulting y is still in S_1 and so the fixation will not be undone by point-mutations.

One mutation in head, one in tail: Finally, suppose a double mutant arises with one flip in the head and one in the tail. This happens with probability $\frac{2|S|(n-|S|)}{n(n-1)}$ and gives us $G_{xx} = 1 > 0 = G_{yy}$. If the tail mutation flipped a 0 to a 1 then $G_{yx} = 0$, so y is selected against and cannot fix. This happens with conditional probability $\frac{n-|S|-2k}{n-|S|}$. The most interesting case is if the tail mutation flipped a 1 to a 0 then $G_{yx} = 1$ and $G_{xy} = 0$. This happens with conditional probability $\frac{2k}{n-|S|}$. Without self-interactions in finite populations, $u_y - u_x = 1/M$ so for strong-selection, y will fix with probability lower bounded by a constant p^* . After fixing, however, y is outside S_1 and will be quickly returned to S_1 by point-mutations. Upon returning to x' in S_1 , x' will have either one pair fewer or the same number of 1s in the tail than/as x . The former will happen with probability $\frac{2k-1}{n-|S|}$ and the latter with probability $\frac{n-|S|-2k-1}{n-|S|}$.

Biased random walk on S_1 : Putting this all together, we get that the double mutant regime is equivalent to a biased random walk on the integer line $[0, m]$ with $m = \lfloor \frac{n-|S|}{2} \rfloor$ where being at a point $k \in [0, m]$ represents having a genotype with a full-head and an even-tail of length $2k$. Importantly, this walk has only a single absorbing state at $k = 0$ (i.e. when $x = s$). The transition probabilities for the random walk's Markov chain are:

$$\Pr\{k \mapsto k - 1\} \geq \frac{1}{M} \frac{2k(2k - 1)}{n(n - 1)} + p^* \frac{4|S|k(2k - 1)}{n(n - 1)(n - |S|)} \quad (12)$$

$$\Pr\{k \mapsto k + 1\} = \frac{1}{M} \frac{(n - |S| - 2k)(n - |S| - 2k - 1)}{n(n - 1)} \quad (13)$$

$$\Pr\{k \mapsto k\} = 1 - \Pr\{k \mapsto k - 1\} - \Pr\{k \mapsto k + 1\} \quad (14)$$

If $p^*M > \frac{(n-|S|)^3}{|S|}$ then $\Pr\{k \mapsto k - 1\} > \Pr\{k \mapsto k + 1\}$ for all $1 \leq k \leq m - 1$. So in the case that $M = \Theta(n^3)$, the random walk will drift left, converging to $k = 0$ in number of double-mutations that is linear in the initial k and thus also linear in $n - |S|$.

5.3 Total adaptation time

Number of mutations: Putting this together Sections 5.1 and 5.2 this results in an overall $O(|n|)$ number of fixations with evolution proceedings fast then slow. The fast part proceeds via a linear number of point-mutations with each on a time-scale of $O(\frac{1}{\mu})$. The slow part proceeds by a linear number of double-mutations, but since each double-mutation arises on the time-scale of $\Theta(\frac{1}{\mu^2})$, we will expect to see a total of $O(\frac{n}{\mu})$ of point-mutations mutations. Of those numerous point-mutations will not fix with the exception of $O(n)$ point-mutations corresponding to the one-mutation-in-head-one-in-tail double-mutation. Thus, we expect on the order of $O(\frac{n}{\mu})$ total mutations.

Number of birth-death events: To finish the time analysis, we can go past mutations to the number of birth-death events that make up each Moran process. For all the mutations that fail to fix, they will fail with an expected $O(1)$ number of birth-death events. Of the mutations that do fix: the ones that are selected for will fix in an expected $O(M \ln M)$ number of birth-death events and the ones that are neutral will fix in an expected $O(M^2)$ number of birth-deaths [30]. Since invasion by neutral drift happens only for the double mutants, this means that the overall evolutionary time can be bounded by $O(nM^2 + \frac{nM}{\mu})$. But we need to avoid a new mutation arriving while an existing one is fixing, so I need to set $\mu \gg \frac{1}{M^2}$ and from the end of Section 5.2 we have $M = \Theta(n^3)$. Thus, the overall evolutionary time measured in the number of birth-death can be bounded by $O(n^{10})$. For this article, this polynomial bound is sufficient to show that an evolving population can adapt to the Parity environment under extended algorithm Darwinism.

6 Discussion

Theoretical computer science and evolutionary biology provide a rich interface for the development of new theory. Providing a formal model of computation that characterizes which algorithms correspond to evolution is essential to this new interface. Valiant [1] started by defining strict algorithmic Darwinism as a subset of SQ-learning, but this did not account for the rich feedback possible between ecology and evolution [4, 5] through effects like frequency-dependent selection. By extending algorithmic Darwinism to include these ecological interactions, I showed that evolution is computationally more powerful than we previously suspected. Thus, we need to continue to refine our model of evolution to better integrate theoretical computer science into biology.

Protroths vs heterotroths and the origin of life: In the context of the Oparin-Haldane hypothesis for the heterotrophic origin of life, the organisms in both strict and extended algorithmic

Darwinism would, technically, be classified as heterotrophs. This is because these organisms do not generate their own energy but have to harvest it from correctly consuming macromolecules in the environment. But there is also a fundamental difference between these two kinds of organisms. In strict algorithmic Darwinism, all the macromolecules are generated by the environment in some process that is independent of the organisms in that environment. This sort of metabolism is exceptionally rare on Earth, maybe occurring only in deep-sea hydro-thermal vents. As such, I think that the organisms of strict algorithmic Darwinism should be called *prototrophs* to distinguish them from typical heterotrophs. The *heterotrophs* that do occur frequently in nature are much more similar to the organisms of the extended algorithmic Darwinism, in that they harvest energy from correctly eating macromolecules that are generated by a process that depends on the distribution of organisms in the environment. But it is exactly this shift from an abiotic to biotic environment that I have shown to transform the mode of evolution to make it exponentially more powerful. Thus, it is tempting to hypothesize that part of the reason that we have such an abundance of true heterotrophs from a presumably prototrophic origin is due to the computational resource that heterotrophy gave populations in adapting to Earth's early environment.

Many consider the ability to metabolize, replicate and evolve as the central characteristics of living organisms. We should augment this list: the central characteristics of living organisms are the ability to metabolize, replicate, evolve and interconnect in a shared ecology.

Ecology is more transformative than sex: It is also important to highlight just how transformative ecology is to the mode of evolution. For example, by comparing it to the difference in mode between asexual vs sexual reproduction. Sexual reproduction, recombination, horizontal gene-transfer, and fusion are another set of processes that are believed to have the potential to greatly speed-up evolution [31, 32] and even matter in the origins of life [33] and the somatic evolution of cancer [34]. Strict algorithmic Darwinism has previously been expanded to study the transformation from asex to sex by considering recombination [35] and horizontal-gene transfer [36, 37]. Similar to my results for ecology, the authors showed that an exponential speed up is possible in the number of generations required for adapting-to a family of easy environments. But unlike my results for ecology, they also proved that the set of families that are adaptable-to is not changed and is still the set of CSQ-learnable environments. In contrast, in this paper, I presented an exponential speed-up for the Parity environment that is hard for strict algorithmic Darwinism. This moves the ecologically extended Algorithmic Darwinism beyond what is CSQ-learnable. The sexual mode of evolution allowed a speed-ups from polynomial to polylog number of generations. But, unlike the ecological mode of evolution, the sexual mode of evolution did not allow for a speed-up from exponential to polynomial number of generations. In other words, both sex and ecology can dramatically speed up evolution, but it is only ecology that can transform what kind of environments are adaptable-to.

Knowing about this increase in computational power is important not only for a general understanding of evolution, but also so that we do not underestimate the power of eco-evolutionary dynamics in settings where evolution is our adversary. For example, in cancer, the patient and physician are in a battle against somatic evolution [38, 39]. Since we tend to view the change in gene-frequencies as the most important aspect, we tend to treat cancer as a disease of individual aberrant cells. But given that frequency-dependent selection can so drastically transform the mode of evolution, we need to also focus on the aberrant ecology of cells [18, 19, 34, 38, 39]. My hope is that a formal algorithmic theory of eco-evolutionary dynamics like my ecologically extended algorithmic Darwinism can help in this effort.

Acknowledgements

I am grateful to Emily Dolson, Patrick Ellsworth, Nathan Farrokhan, Peter Jeavons, Jeff Maltas, and Julian Xue for feedback on an early draft of this article. The Darwinian engine in Figure 1 was inspired by presentations given by Amitabh Joshi and Joachim Krug.

References

1. Valiant, L. Evolvability. *Journal of the ACM* 1 2009;56:3.
2. Kaznatcheev, A. Computational Complexity as an Ultimate Constraint on Evolution. *Genetics* 2019;212:245–265.
3. Kaznatcheev, A, Cohen, DA, and Jeavons, PG. Representing fitness landscapes by valued constraints to understand the complexity of local search. In: *International Conference on Principles and Practice of Constraint Programming*. Springer. 2019:300–316.
4. Hendry, AP. *Eco-evolutionary dynamics*. Princeton university press, 2017.
5. Govaert, L, Fronhofer, EA, Lion, S, et al. Eco-evolutionary feedbacks—Theoretical models and perspectives. *Functional Ecology* 2019;33:13–30.
6. Kanade, V. *Computational Questions in Evolution*. PhD Thesis at Harvard University 2012.
7. Abrams, M. Measured, modeled, and causal conceptions of fitness. *Frontiers in Genetics* 2012;3:196.
8. Kaznatcheev, A. Two conceptions of evolutionary games: reductive vs effective. *bioRxiv* 2017:231993.
9. Kearns, M. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)* 1998;45:983–1006.
10. Wright, S. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress of Genetics* 1932:356–366.
11. Traulsen, A, Shores, N, and Nowak, MA. Analytical results for individual and group selection of any intensity. *Bulletin of mathematical biology* 2008;70:1410.
12. Valiant, L. A theory of the learnable. *Communications of the ACM* 1984;27.
13. Feldman, V. Evolvability from learning algorithms. In: *Proceedings of the Fortieth Annual ACM symposium on Theory of Computing*. 2008:619–628.
14. Gillespie, J. A simple stochastic gene substitution model. *Theor. Pop. Biol.* 1983;23:202.
15. Orr, HA. The genetic theory of adaptation: a brief history. *Nature Reviews. Genetics* 2005;6:119–127.
16. Ohta, T. The nearly neutral theory of molecular evolution. *Annual Review of Ecology and Systematics* 1992;23:263–286.
17. Peña, J, Lehmann, L, and Nöldeke, G. Gains from switching and evolutionary stability in multi-player matrix games. *Journal of Theoretical Biology* 2014;346:23–33.
18. Kaznatcheev, A, Vander Velde, R, Scott, JG, and Basanta, D. Cancer treatment scheduling and dynamic heterogeneity in social dilemmas of tumour acidity and vasculature. *British Journal of Cancer* 2017.

19. Kaznatcheev, A, Peacock, J, Basanta, D, Marusyk, A, and Scott, JG. Fibroblasts and alectinib switch the evolutionary games played by non-small cell lung cancer. *Nature Ecology & Evolution* 2019;3:450.
20. Moran, PAP. Random processes in genetics. In: *Mathematical proceedings of the cambridge philosophical society*. Vol. 54. 1. Cambridge University Press. 1958:60–71.
21. de Visser, J, Park, S, and Krug, J. Exploring the effect of sex on empirical fitness landscapes. *The American Naturalist* 2009.
22. Crona, K, Greene, D, and Barlow, M. The peaks and geometry of fitness landscapes. *Journal of Theoretical Biology* 2013;317:1–10.
23. Diekmann, O. A beginners guide to adaptive dynamics. *Summer School on Mathematical Biology* 2002:63–100.
24. Waxman, D and Gavrillets, S. 20 questions on adaptive dynamics. *Journal of evolutionary biology* 2005;18:1139–1154.
25. Bonnet, J, Yin, P, Ortiz, ME, Subsoontorn, P, and Endy, D. Amplifying genetic logic gates. *Science* 2013;340:599–603.
26. Gidon, A, Zolnik, TA, Fidzinski, P, et al. Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science* 2020;367:83–87.
27. Van Duijn, M, Keijzer, F, and Franken, D. Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior* 2006;14:157–170.
28. Blum, A, Kalai, A, and Wasserman, H. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM (JACM)* 2003;50:506–519.
29. Regev, O. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)* 2009;56:1–40.
30. Antal, T and Scheuring, I. Fixation of strategies for an evolutionary game in finite populations. *Bulletin of mathematical biology* 2006;68:1923–1944.
31. Becks, L and Agrawal, AF. The evolution of sex is favoured during adaptation to new environments. *PLoS biology* 2012;10.
32. McDonald, MJ, Rice, DP, and Desai, MM. Sex speeds adaptation by altering the dynamics of molecular evolution. *Nature* 2016;531:233–236.
33. Sinai, S, Olejarz, J, Neagu, IA, and Nowak, MA. Primordial sex facilitates the emergence of evolution. *Journal of The Royal Society Interface* 2018;15:20180003.
34. Miroshnychenko, D, Baratchart, E, Ferrall-Fairbanks, M, et al. Spontaneous cell fusions as a mechanism of parasexual recombination in tumor cell populations. *BioRxiv* 2020.
35. Kanade, V. Evolution with recombination. In: *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. IEEE. 2011:837–846.
36. Snir, S and Yohay, B. Prokaryotic evolutionary mechanisms accelerate learning. *Discrete Applied Mathematics* 2019;258:222–234.
37. Yohay, B and Snir, S. Extending the Evolvability Model to the Prokaryotic World: Simulations and Results on Real Data. *Journal of Computational Biology* 2019;26:794–805.
38. Scott, J and Marusyk, A. Somatic clonal evolution: A selection-centric perspective. *Biochimica et Biophysica Acta (BBA)-Reviews on Cancer* 2017;1867:139–150.

39. Rockne, RC, Hawkins-Daarud, A, Swanson, KR, et al. The 2019 Mathematical Oncology Roadmap. *Physical Biology* 2019;16:041005.