

Article

The Price Equation Reveals a Universal Force–Metric–Bias Law of Algorithmic Learning and Natural Selection

Steven A. Frank 

Department of Ecology and Evolutionary Biology, University of California, Irvine, CA 92697-2525, USA;
safrank@uci.edu

Abstract

Diverse learning algorithms, optimization methods, and natural selection share a common mathematical structure despite their apparent differences. Here, I show that a simple notational partitioning of change by the Price equation reveals a universal force–metric–bias (FMB) law : $\Delta\theta = \mathbf{M}\mathbf{f} + \mathbf{b} + \xi$. The force $\mathbf{f}$ drives improvement in parameters, $\Delta\theta$, in proportion to the slope of performance with respect to the parameters. The metric $\mathbf{M}$ rescales movement by inverse curvature. The bias $\mathbf{b}$ adds momentum or changes in the frame of reference. The noise ξ enables exploration. This framework unifies natural selection, Bayesian updating, Newton’s method, stochastic gradient descent, stochastic Langevin dynamics, Adam optimization, and most other algorithms as special cases of the same underlying process. The Price equation also reveals why Fisher information, Kullback–Leibler divergence, and d’Alembert’s principle arise naturally in learning dynamics. By exposing this common structure, the FMB law provides a principled foundation for understanding, comparing, and designing learning algorithms across disciplines.

Keywords: price equation; force–metric–bias; Bayesian updating; momentum-based optimization; information geometry; natural selection; stochastic gradient descent



Academic Editor: Antonio M. Scarfone

Received: 9 September 2025

Revised: 23 October 2025

Accepted: 29 October 2025

Published: 31 October 2025

Citation: Frank, S.A. The Price Equation Reveals a Universal Force–Metric–Bias Law of Algorithmic Learning and Natural Selection. *Entropy* **2025**, *27*, 1129.
<https://doi.org/10.3390/e27111129>

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Learning algorithms pervade modern science. Machine learning improves neural networks through gradient descent. Evolution improves organisms through natural selection. Bayesian inference improves beliefs through probability updates. Despite decades of research in each field, the ultimate relations between these approaches remain unclear.

This article shows that a single force–metric–bias (FMB) law captures the essential structure of algorithmic learning and natural selection. Improvement arises from three components: force, typically expressed by the performance gradient; metric, typically expressed by inverse curvature; and bias, which includes momentum and other changes in the frame of reference. This structure emerges naturally from the Price equation, a simple notational description for the partitioning of change into components [1–3].

Consider how the following two connections arise naturally within this framework. First, the primary equation in evolutionary biology [4], $\Delta\theta = \mathbf{P}\alpha$, and Newton’s method [5] in optimization, $\Delta\theta = -\mathbf{H}^{-1}\nabla U$, are mathematically analogous. Both describe one step of change by multiplying a gradient-like force, α or ∇U , by evolution’s covariance matrix, $\mathbf{P}$, or Newton’s inverse Hessian, $\mathbf{H}^{-1}$, each serving the same metric role of rescaling geometry by inverse curvature.

Second, machine learning algorithms are used to improve the performance of neural networks or other methods of prediction. The progression between a few common algorithms

perfectly illustrates the FMB decomposition. Stochastic gradient descent [6] uses force, $\mathbf{f}$. Polyak [7] adds a momentum bias, $\mathbf{b}$. Adam [8] includes adaptive metric scaling, $\mathbf{M}$. Adam's full structure, $\mathbf{M}\mathbf{f} + \mathbf{b}$, is the same as evolution's primary equation and Newton optimization, with an additional momentum bias term that often improves performance.

These connections reflect a deeper geometric principle. Many learning algorithms face the same fundamental challenge: maximize improvement in performance minus a cost paid for distance moved in the parameter space. Here, we must account for two aspects of geometry. First, there may be a curved relation between parameters and performance. Second, constraints on movement in the parameter space induce a metric that alters how distance is measured. For example, a lack of genetic variation in a particular direction constrains movement in that direction.

The optimal solution is the product of the force and the inverse curvature metric. Different fields have discovered this same result within specific contexts. Here, we see it in its full simplicity and generality, providing a reason for the recurring role of Fisher information as a curvature metric in probability contexts and inverse Hessian metrics in local geometry contexts.

The geometric structure of learning has been partially recognized in prior work. Fisher [9] and Rao [10] established that statistical parameter spaces can have intrinsic curvature. Amari [11] used this insight to develop natural gradient methods.

In evolutionary biology, Shahshahani [12] applied differential geometry to the dynamics of natural selection, introducing metric concepts to evolutionary theory. Newton's method uses curvature to improve stepwise updates. Machine learning algorithms are often designed to estimate local curvature in an efficient way.

These insights about geometry, force, momentum, and bias remained confined to their domains. The full simplicity and universality of algorithmic learning have not been expressed in a clear and formal way. This article demonstrates the underlying unity, revealing the simple mathematical law that governs learning processes.

2. The Force–Metric–Bias Law

2.1. Statement of the FMB Law

I first state the FMB law. The following subsection derives the law and clarifies the notation.

This law is not an empirical hypothesis but rather a universal mathematical structure that underlies learning or selection. The law is

$$\Delta\bar{\theta} = \mathbf{M}\mathbf{f} + \mathbf{b} + \zeta. \quad (1)$$

Here, $\bar{\theta}$ denotes a vector of n mean parameter values that is updated by learning, optimization, or natural selection. The law also applies to updates of a single parameter vector, $\Delta\theta$, instead of mean values. Here, *parameters* are values of any sort. In biology, we call them *traits*.

The $n \times n$ matrix $\mathbf{M}$ describes a metric, which typically expresses the inverse curvature of the parameter space and the rescaling of distances. The nature of the metric varies in different algorithms, as discussed below. Throughout this article, metric matrices that properly rescale distances are positive definite. When a matrix is not positive definite, algorithms typically modify it or use alternative metrics to ensure valid updates.

The force vector $\mathbf{f}$ often includes the gradient $\nabla_{\theta}U$ of the performance function U with respect to the parameters θ . In general, the force vector typically describes processes that push toward increased performance or constrain such an increase.

The bias vector, $\mathbf{b}$, includes processes such as parameter momentum or change in frame of reference. These processes alter parameters in addition to the standard directly acting forces imposed by performance or constraint. The standard form of bias is

$$\mathbf{b} = \mathbf{C}\boldsymbol{\beta} + \boldsymbol{\gamma}, \quad (2)$$

in which $\mathbf{C}$ describes a bias metric, $\boldsymbol{\beta}$ is the slope of performance with respect to biased parameter changes, and $\boldsymbol{\gamma}$ is the bias that is independent of performance. Most algorithms follow this pattern for bias, modifying specific terms according to particular learning goals.

The noise vector, $\boldsymbol{\xi}$, has a mean of zero. Commonly, we partition the noise into a metric term and a simple noise-generating process. For example, many algorithms use some variant of

$$\boldsymbol{\xi} = \sqrt{\mathbf{D}}\boldsymbol{\epsilon},$$

in which $\mathbf{D}$ is a metric that reshapes the noise, and $\boldsymbol{\epsilon}$ is a basic noise process such as a Gaussian with a mean of zero and a standard deviation of one.

The following derivation of the FMB law reveals further key distinctions between directly acting forces, $\mathbf{f}$, and bias, $\mathbf{b}$.

2.2. Derivation from the Price Equation

The generality of the FMB law arises from simple notational descriptions of change. This subsection describes the key steps. Note that, at first glance, the definition of terms in the FMB law may not seem to match many common learning algorithms, such as stochastic gradient descent. Later sections make the connections.

A subsequent section shows that the same simple approach also leads to common methods and measures that frequently arise in learning algorithms, such as Fisher information, Kullback–Leibler divergence, and information geometry. This article ties these pieces together.

(1) We begin with the Price equation, a universal expression for change. A probability vector $\mathbf{q}$ of length m sums to one. Each q_i weights an alternative parameter vector, $\boldsymbol{\theta}_i = \theta_{i1}, \dots, \theta_{in}$, with each parameter vector $\boldsymbol{\theta}_i$ of length n . The symbol $\boldsymbol{\theta}$ without subscript denotes the m -vector of alternative $\boldsymbol{\theta}_i$, each parameter vector associated with a probability q_i .

To begin, assume that we have only one parameter, $n = 1$, with m variant values. Later, I show that the same approach works for $n > 1$, with notation extended for vectors and matrices.

An update to the mean parameter value over the m variants is $\Delta\bar{\theta} = \mathbf{q}' \cdot \boldsymbol{\theta}' - \mathbf{q} \cdot \boldsymbol{\theta}$, in which the dots denote inner products, and Δ is the difference between an updated primed value and the original value. Rearranging yields the Price equation [3]

$$\Delta\bar{\theta} = \Delta\mathbf{q} \cdot \boldsymbol{\theta} + \mathbf{q}' \cdot \Delta\boldsymbol{\theta}. \quad (3)$$

This equation is simply the definition of change in mean value, rearranged into a chain rule analog for finite differences rather than infinitesimal derivatives. The first term is the change in frequencies holding the parameters constant. The second term is the change in parameter values holding frequencies at their fixed updated values.

(2) Define w_i as the relative growth of the i th type, $q'_i = w_i q_i$, such that

$$\Delta q_i = q_i(w_i - 1). \quad (4)$$

In biology, w_i is called the relative fitness of the i th type, describing how survival and reproduction alter the frequencies of the types, with $\bar{w} = \mathbf{q} \cdot \mathbf{w} = 1$.

By the standard definition of covariance, $\Delta \mathbf{q} \cdot \boldsymbol{\theta} = \mathbf{Cov}(w, \boldsymbol{\theta})$, and by the standard definition of expectation

$$\mathbf{q}' \cdot \Delta \boldsymbol{\theta} = \sum_i q'_i \Delta \theta_i = \sum_i q_i w_i \Delta \theta_i = E(w \Delta \boldsymbol{\theta}).$$

With these definitions, the Price equation can be rewritten as [2]

$$\Delta \bar{\boldsymbol{\theta}} = \mathbf{Cov}(w, \boldsymbol{\theta}) + E(w \Delta \boldsymbol{\theta}). \quad (5)$$

These forms of the Price equation are simply notational descriptions for change [3]. We have not assumed anything about the nature of the values or how they change. We have assumed that w_i always describes actual frequency changes.

If the performance function, $\mathbf{U}$, subsumes all of the forces that act on frequency change in a particular time period, then w_i is the performance of the i th type, $\mathbf{U}(\theta_i)$, normalized to $\bar{w} = 1$ for notational simplicity and without loss of generality

$$w_i = \frac{\mathbf{U}(\theta_i)}{\sum_j \mathbf{U}(\theta_j)},$$

in which fitness and relative performance are equivalent descriptions of actual change.

In some cases, the forces acting on frequency change are composed of several distinct processes. One component may arise from a performance function, $\mathbf{U}$. Another component may act as a constraining force that prevents frequency changes from following the forces imposed by $\mathbf{U}$. Then frequency change is no longer aligned with the optimal direction for improving performance, and w_i is not equivalent to the relative value of the performance function, $\mathbf{U}$. Nonetheless, w_i is the actual relative performance in the context of the Price equation's notational conventions.

(3) To derive the first term of the FMB law in Equation (1), write the standard least-squares regression of fitness on parameter values as

$$w_i = f_{w\theta} \theta_i + \zeta_i,$$

in which f is the regression coefficient of w on θ , and ζ is the error uncorrelated with θ . Using this regression in the first Price covariance term yields

$$\Delta_f \bar{\boldsymbol{\theta}} = \mathbf{Cov}(w, \boldsymbol{\theta}) = f_{w\theta} \mathbf{Var}(\boldsymbol{\theta}),$$

in which Δ_f denotes the partial change caused by the force, $\mathbf{f}$, imposed by relative performance, w . In the multivariate case, with $n > 1$ parameters, this same term expands to

$$\Delta_f \bar{\boldsymbol{\theta}} = \mathbf{Cov}(w, \boldsymbol{\theta}) = \mathbf{M} \mathbf{f}, \quad (6)$$

in which $\mathbf{M}$ is the covariance matrix of the parameters, defined by $\mathbf{Cov}(\boldsymbol{\theta}, \boldsymbol{\theta})$, and $\mathbf{f}$ is the vector of partial regression coefficients for fitness, w , with respect to each of the n parameters.

Here, Δ_f changes average parameter values only through changes in frequency.

(4) Bias directly changes a parameter value. For a parameter influenced by bias, $\Delta \theta_i = \theta'_i - \theta_i \neq 0$. To derive the bias term of the FMB law, write the regression of fitness on the changes in parameters

$$w_i = \beta_{w\Delta\theta} \Delta \theta_i + \zeta_i.$$

Then the second Price term yields

$$\mathbf{E}(w\Delta\theta) = \mathbf{Cov}(w, \Delta\theta) + \gamma = \beta_{w\Delta\theta} \mathbf{Var}(\Delta\theta) + \gamma,$$

in which $\gamma = \mathbf{E}(\Delta\theta)$. For $n > 1$, the extended notation is

$$\Delta_b \bar{\theta} = \mathbf{E}(w\Delta\theta) = \mathbf{C} \beta + \gamma = \mathbf{b}, \quad (7)$$

in which $\mathbf{C} = \mathbf{Cov}(\Delta\theta, \Delta\theta)$ is the covariance matrix of $\Delta\theta$, and β is the vector of partial regression coefficients for w on $\Delta\theta$. The symbol Δ_b denotes the partial change caused by bias.

(5) We add a noise term, $\Delta_\xi \bar{\theta} = \xi$, to complete the FMB law. In the infinitesimal limit, the law has the standard form of a stochastic differential equation. The first two components, Δ_f and Δ_b , define the deterministic drift change, and the third Δ_ξ component defines the stochastic diffusion change [13].

(6) As the parameter distributions concentrate near their mean values, the variances and covariances become small. In the limit, we have updates to a single parameter vector, $\Delta\theta$, and the regressions in the Δ_f and Δ_b terms converge to gradients. This limit recovers the common usage of gradients in learning algorithms that update single parameter vectors rather than updating mean parameter vectors over distributions. In this limit, the metrics provided by the covariance matrices are replaced by other aspects of geometric curvature, as discussed in the following subsections.

At this point, the FMB law is simply a notational partition of change into specific parts. The value arises from the insight and unity this notation brings to the diverse and seemingly unconnected applications that arise in different studies of learning and natural selection.

2.3. Metrics, Sufficiency, and Single-Value Updates

The Price equation's force, bias, and noise terms in the FMB law of Equation (1) can be expanded to

$$\Delta \bar{\theta} = \mathbf{M} \mathbf{f} + (\mathbf{C} \beta + \gamma) + \sqrt{\mathbf{D}} \epsilon. \quad (8)$$

For the change in the location of the parameter vector, $\Delta \bar{\theta}$, the metrics $\mathbf{M}$, $\mathbf{C}$, and $\mathbf{D}$, the regression-based gradients, $\mathbf{f}$ and β , and the additional bias component γ are sufficient statistics to reconstruct the update. Additional information about frequencies, $\mathbf{q}$, does not alter the change in the location of the parameter vector.

The sufficiency of the terms in Equation (8) to describe the change in the location of the parameter vector is important. It means that the FMB law, although initially derived from the Price equation's population frequencies, also accurately describes changes to a single parameter vector.

The update depends only on the sufficient statistics, which are the metric matrices, the force vectors, and the intrinsic bias. In other words, we can invoke the common geometry that unifies updates to the mean vector, based on underlying frequencies of different parameter vectors, or updates to a single vector, based on alternative calculations of the sufficient statistics.

The Price equation's metric terms are covariance matrices. However, a covariance matrix is just a metric matrix. In a population interpretation, we call the matrix a covariance. In a geometric interpretation, we call the matrix a metric. Mathematically, they are equivalent.

Similarly, population regressions enter only as slopes that can equivalently be analyzed geometrically. Intrinsic bias can also arise from either a population or a purely geometric interpretation.

Natural selection and some learning algorithms build on population notions of frequency and mean locations. Many other learning algorithms build on single-value updates

of metrics, gradients, and geometry. Both interpretations follow from the Price equation's FMB law. The difference arises in whether we assume that changing population frequencies set the metrics and slopes, or we assume that other attributes of a system set the geometry.

This conceptual shift allows the FMB law to unify disparate fields. As we will see, the metric $\mathbf{M}$ in natural selection is the empirically observed covariance matrix of parameters. In Newton's method for optimization, the metric $\mathbf{M}$ is the analytically calculated inverse Hessian matrix. The FMB law reveals that these are different choices for the metric in different contexts, all within the same underlying mathematical structure.

In the following sections, I first continue to emphasize the Price equation's population-based perspective. Later, I switch emphasis to single-value updates based purely on a geometric perspective. The two perspectives are different views of the same underlying FMB law.

2.4. A Spectrum of Methods: From Local to Population

In practice, the variety of algorithms forms a spectrum of information-gathering strategies. The spectrum runs across the spatial and temporal scope of the information they use to define the curvature metric, $\mathbf{M}$, the force, $\mathbf{f}$, and the bias, $\mathbf{b}$. Here, spatial scope describes a population of parameter vectors considered at a point in time, whereas temporal scope describes a sequence of parameter vectors over time.

Two extremes define the spectral extent. On one side, the purely local methods obtain information for both metric and force from a single parameter vector. For example, Newton's method calculates the force vector as the first derivative of performance and the curvature metric as the inverse Hessian matrix, the second derivative of performance. Both derivatives are calculated with respect to a single parameter vector.

On the other side, purely population-based methods use a full spatial scope to define both a covariance metric and a regression-based force. In this case, curvature and force are averaged over a distribution of alternative parameter vectors. Here, I briefly mention a few classic examples to illustrate how various algorithms fall along this spectrum.

Amari's natural gradient is a hybrid method. It combines a purely local force, the gradient at a point, with a metric of extended spatial scope, the Fisher information metric of a distribution over alternative parameter vectors [11,14].

Stochastic gradient descent samples a batch of local gradients. The average of the several precise local force vectors estimates the force for a population sample. In effect, the statistical sampling transforms the local gradient descent method into a quasi-population method [6,15–17].

Optimization methods like Adam substitute temporal scope for spatial scope. As the optimizer traverses the parameter space over time, it generates a historical sequence of parameter vectors. This sequence provides a population of parameter vectors over which the method combines the local gradients to estimate a momentum-like statistic that augments the local force and to build a diagonal metric that captures aspects of the spatially extended curvature metric [8].

Later sections will develop these analyses in detail, showing how the variety of algorithms arises from particular information-gathering strategies and ways of calculating the components of the FMB law.

2.5. Performance and Cost Functions: Sign Convention

I set $\mathbf{U}$ as a performance function that provides increasing benefits as it rises in magnitude. The choice of a target function to maximize arose from the Price equation's biological convention of fitness as a beneficial attribute.

By contrast, many studies in numerical optimization and other fields take $\mathbf{U}$ as a cost function to be minimized. In this article, I adopt the maximization of $\mathbf{U}$ as the primary goal.

Results for minimizing cost follow by substituting $-\mathbf{U}$ for $\mathbf{U}$. If this substitution is used to minimize cost, then the Hessian calculation for local curvature becomes the curvature of the cost function $-\mathbf{U}$, which inverts the sign of the Hessian used in maximization.

There is no difference except that one has to pay attention to the directions of change and the appropriate signs appended to terms.

3. Natural Selection, Metrics, and Curvature

This section links the metric and force terms to natural selection, a topic that has a well-developed theoretical foundation. The connection illustrates how the familiar concepts in biology associate with the more abstract geometric concepts of the FMB law. In this case, metric and force arise from the spatial extent notion of populations, the basis of the Price equation, and the initial path to the general FMB geometry.

From Equation (6), an update caused solely by the first Price term for frequency change is

$$\Delta_{\mathbf{f}} \bar{\boldsymbol{\theta}} = \text{Cov}(\mathbf{w}, \boldsymbol{\theta}) = \mathbf{M} \mathbf{f}.$$

In biological studies of natural selection, this result is often called the Lande equation [4]. In that case, $\boldsymbol{\theta}$ is a vector of n trait values, $\mathbf{M}$ is the covariance matrix of the trait values, and $\mathbf{f}$ is the vector of partial regression coefficients of fitness, $\mathbf{w}$, on trait values, $\boldsymbol{\theta}$. The Introduction wrote the right side of this equation for biology as $\mathbf{P}\boldsymbol{\alpha}$ to distinguish the biological terms. But now we use our standard notation, $\mathbf{M} \mathbf{f}$.

This classic equation of natural selection matches the primary update process used by most learning algorithms. The slope of performance (fitness) relative to the parameters (traits) creates the primary driving force for updates, $\mathbf{f}$. The covariance matrix, $\mathbf{M}$, defines the update metric. Other algorithms vary in the spatial and temporal extent used to calculate force and metric, the methods for the particular calculations, and supplemental bias and stochasticity components.

A metric changes the length and direction of the update path, $\Delta_{\mathbf{f}} \bar{\boldsymbol{\theta}}$, by modulating the forces acting in each direction of the parameter space. The expected gain in performance is the force, $\mathbf{f}$, multiplied by the displacement, $\Delta_{\mathbf{f}} \bar{\boldsymbol{\theta}}$, yielding

$$\mathbb{E}(\Delta_{\mathbf{f}} \bar{\mathbf{w}}) = \mathbf{f} \cdot \Delta_{\mathbf{f}} \bar{\boldsymbol{\theta}} = \mathbf{f}^{\top} \mathbf{M} \mathbf{f}. \quad (9)$$

A metric alters the sum of squares for a vector, changing its Euclidean length, with the requirement on $\mathbf{M}$ that the resulting length be a nonnegative real value. We can drop the expectation when $\mathbf{f}$ is calculated from an explicit performance function.

A metric has a natural interpretation in terms of inverse curvature. For example, if $\mathbf{M}$ is a covariance matrix for $\boldsymbol{\theta}$, then, along any direction, a small value implies that there is little variation in the values of the parameters in that direction and, therefore, relatively little opportunity to shift the mean of the parameters.

A probability distribution with a small variance is narrow and highly curved, linking large curvature to small variation. Thus, inverse curvature describes variance. The movement in a particular direction becomes the force in that direction multiplied by the variance or inverse curvature in that direction. High variance and a straight surface augment the force. Low variance and a curved surface deter the force. Many publications consider the geometry of evolutionary dynamics [12,18,19].

In biology, natural evolutionary processes set the covariance matrix as the update metric. In learning and optimization algorithms, one chooses the metric by assumption or by particular calculations from the data and the update dynamics. The way in which the metric is chosen defines a primary distinction between algorithms.

4. Geometry, Information, and Work

Learning algorithms provide iterative improvement, a particular type of dynamics. This section reviews general properties of learning, which set a foundation for understanding the variety of algorithms and their unification [20,21].

We will see that the Price equation's notation for frequency change naturally gives rise to Fisher information, Kullback–Leibler divergence, information geometry, and d'Alembert's principle. The simple derivations reveal deep connections between learning dynamics and physical principles.

The Price equation and the consequent classic results follow from the intrinsic geometry of the purely population-based case. The insights from this spatially extended scope of populations provide the foundation to understand how the local geometric analysis of many learning algorithms fits within the broad FMB law.

4.1. Price Equation Foundation

The general expressions for learning updates arise from the Price equation, from which we see that the metric and force terms follow immediately from the basic notational description for the change in frequency.

Many learning algorithms do not have an intrinsic notion of frequency. Earlier, I showed how those algorithms fit into this scheme by considering the FMB terms as sufficient quantities for updates.

In this section, I continue to focus on frequency changes. Frequency here simply means a vector of positive weights with a conserved total value. We normalize the total to be one, which links to notions of probability, frequency, and average values. Several classic measures and methods of learning follow.

Most of the particular results in this section are widely known. Once again, the advantage here is that these aspects emerge simply and naturally as the outcome of our basic Price equation notation, without the need to invoke particular assumptions or interpretations.

4.2. Discrete Fisher–Rao Length

A discrete generalization of the Fisher–Rao length follows immediately from Equation (9), which gave the partial increase in mean fitness, $\bar{w}$, as $\mathbf{f}^\top \mathbf{M} \mathbf{f}$. We can express that same quantity purely in terms of frequencies by the first Price term from Equation (3). To do so, we use fitness as the trait of interest, $\theta_i = w_i$, with $w_i = 1 + \Delta q_i/q_i$ from Equation (4), yielding [20]

$$\Delta_{\mathbf{f}} \bar{w} = \Delta \mathbf{q} \cdot \mathbf{w} = \sum_i \frac{(\Delta q_i)^2}{q_i} = \left\| \frac{\Delta \mathbf{q}}{\sqrt{\mathbf{q}}} \right\|^2 = \mathcal{F}. \quad (10)$$

The notation $\|\cdot\|$ denotes the vector norm, which is the Euclidean length of the vector, and $\mathcal{F}$ denotes the discrete generalization of the squared Fisher–Rao step length that arises from the Fisher information metric [10].

The value of $\mathcal{F}$ measures the divergence between probability distributions for the discrete jump $\Delta \mathbf{q} = \mathbf{q}' - \mathbf{q}$. Equivalently, $\Delta \mathbf{q} \cdot \mathbf{w} = \text{Cov}(w, w) = \text{Var}(w)$, the variance in fitness, describes the same value.

4.3. Kullback–Leibler Divergence

The Fisher–Rao length measures the separation between probability distributions. The Kullback–Leibler (KL) divergence provides another common way to measure that separation [22]. If we write $w_i = e^{m_i}$, so that the discrete update is $q'_i = q_i e^{m_i}$, then we can think of the discrete change as a continuous path arising from the solution of an

infinitesimal process that grows at a nondimensional rate proportional to m_i , which in biology is the Malthusian parameter. Thus

$$m_i = \log \frac{q'_i}{q_i} = \log w_i, \quad (11)$$

and using $\mathbf{m}$ instead of $\mathbf{w}$ in the Price equation, the first Price term for the partial change of mean log fitness caused directly by $\Delta \mathbf{q}$ becomes

$$\Delta_{\mathbf{f}} \bar{m} = \Delta \mathbf{q} \cdot \mathbf{m} = \mathcal{D}(\mathbf{q}' || \mathbf{q}) + \mathcal{D}(\mathbf{q} || \mathbf{q}') = \mathcal{J}, \quad (12)$$

which is known as the Jeffreys divergence [23], a symmetric form of the KL divergence of information theory

$$\mathcal{D}(\mathbf{q}' || \mathbf{q}) = \sum_i q'_i \log \frac{q'_i}{q_i}. \quad (13)$$

For infinitesimal changes $\Delta \mathbf{q} \rightarrow d\mathbf{q}$, we get

$$\mathbf{m} \rightarrow \mathbf{w} - 1 = \frac{d\mathbf{q}}{\mathbf{q}} = d \log \mathbf{q}, \quad (14)$$

so that using $\mathbf{m}$ yields

$$\partial_{\mathbf{f}} \bar{m} = d\mathbf{q} \cdot \mathbf{m} = d\mathbf{q} \cdot \mathbf{w} = \mathcal{F},$$

showing that, for continuous infinitesimal changes $\Delta_{\mathbf{f}} \rightarrow \partial_{\mathbf{f}}$, the Jeffreys divergence for discrete changes based on $\mathbf{m}$ converges to the squared Fisher–Rao step length, $\mathcal{J} \rightarrow \mathcal{F}$.

4.4. Information Geometry

Information geometry analyzes the distance between probability distributions on a manifold typically defined by the Fisher information metric. Simple intuition about information geometry follows if we transform to square-root coordinates for frequencies [14].

Let $\mathbf{r} = \sqrt{\mathbf{q}}$, which leads to $\|\mathbf{r}\| = 1$, creating unitary coordinates such that all changes in $\mathbf{r}$ lie on the surface of a sphere with a radius of one. In the new coordinates, the value of the squared Fisher–Rao step length in Equation (10) for the infinitesimal limit becomes $\mathcal{F} = 4\|\mathbf{dr}\|^2$. This surface manifold for dynamics illustrates the widespread use of information geometry when studying how probability distributions change. Figure 1 shows aspects of the geometry.

4.5. Bayesian Updating

The distinction between Bayesian prior and posterior distributions is another way to describe the separation between probability distributions. Following Bayesian tradition, denote $\tilde{L}(\mathbf{D}|\theta)$ as the likelihood of observing the data, $\mathbf{D}$, given parameter values, θ . To interpret the likelihood as a performance measure equivalent to relative fitness, w , the average value of the force must be one to satisfy the conservation of total probability. Thus, define

$$w_i = L_i = \frac{\tilde{L}(\mathbf{D}|\theta_i)}{\sum_i q_i \tilde{L}(\mathbf{D}|\theta_i)}. \quad (15)$$

We can now write the classic expression for the Bayesian updating of a prior, $\mathbf{q}$, driven by the performance associated with new data, $\mathbf{L}$, to yield the posterior, $\mathbf{q}'$, as $q'_i = L_i q_i$, or [24]

$$\mathbf{L} = \frac{\mathbf{q}'}{\mathbf{q}} = \mathbf{w}. \quad (16)$$

By recognizing $\mathbf{L} = \mathbf{w}$, we can use all of the general results derived from the Price equation. For example, the Malthusian parameter of Equation (11) relates to the log-likelihood as

$$\mathbf{m} = \log \frac{\mathbf{q}'}{\mathbf{q}} = \Delta \log \mathbf{q} = \log \mathbf{L}. \quad (17)$$

We can then relate the changes in probability distributions described by the Jeffreys divergence (Equation (12)) and the squared Fisher–Rao update length

$$\Delta_f \bar{L} = \Delta \mathbf{q} \cdot \mathbf{L} = \left\| \frac{\Delta \mathbf{q}}{\sqrt{\mathbf{q}}} \right\|^2 = \mathcal{F}. \quad (18)$$




Figure 1. Geometry of change by direct forces, Δ_f . (a) Divergence between the initial population with probabilities, $\mathbf{q}$, and the altered population with probabilities, $\mathbf{q}'$. For discrete changes, the probabilities are normalized by the square root of the probabilities in the initial set. The distance can equivalently be described by the various expressions shown, in which V_w is the variance in fitness from population biology, $\mathcal{J}$ is the Jeffreys divergence from information theory, and $\mathcal{F}$ is the squared Fisher–Rao step length. The symbol “ $\rightarrow$ ” denotes the limit for small changes. (b) When changes are small, the same geometry and distances can be described more elegantly in unitary square root coordinates, $\mathbf{r} = \sqrt{\mathbf{q}}$, which sets $\|\mathbf{r}\| = 1$, and $\dot{\mathbf{r}} \equiv d\mathbf{r} = d\sqrt{\mathbf{q}} = (d\mathbf{q}/\sqrt{\mathbf{q}})/2$. From Frank [20].

4.6. d’Alembert’s Principle

We can think of the causes that separate probability distributions during an update as forces. Multiplying force and displacement yields a notion of work. Because we conserve the total weights as normalized probabilities, many learning updates require that virtual work vanishes for allowable displacements, yielding d’Alembert’s principle [25].

From the definition for $\mathbf{m}$ in Equation (11), and for the infinitesimal limit in Equation (14), we have $\bar{\mathbf{m}} = \mathbf{q} \cdot \mathbf{m} = 0$. By the chain rule for differentiation, we can write a Price equation expression

$$d\bar{\mathbf{m}} = d\mathbf{q} \cdot \mathbf{m} + \mathbf{q} \cdot d\mathbf{m} = 0.$$

Using $\mathbf{m} = d\mathbf{q}/\mathbf{q}$ from Equation (14), noting that $\mathbf{q} \cdot d\mathbf{m} = d\mathbf{q} \cdot d\log \mathbf{m}$, and rearranging yields

$$(\mathbf{m} + d\log \mathbf{m}) \cdot \delta \mathbf{q} = 0, \quad (19)$$

in which $\delta \mathbf{q} = d\mathbf{q}$ is a small virtual displacement consistent with all constraints. This expression is a nondimensional form of d’Alembert’s principle, in which the virtual work

of the directly acting force for an update, $\mathbf{F} = \mathbf{m}$, and displacement, $\delta \mathbf{q}$, is balanced by the virtual work of the inertial force, $\mathbf{I} = d \log \mathbf{m}$, and displacement, yielding $(\mathbf{F} + \mathbf{I}) \cdot \delta \mathbf{q} = 0$.

In one dimension, we recover an analog of the familiar Newtonian form, $F = ma$, or $(F - ma)\delta r = 0$, showing that the force, F , has an equal and opposite inertial force, ma , for mass m and acceleration, a . For multiple dimensions, we can rewrite Equation (19) in canonical coordinates and obtain a simple Hamiltonian expression [25].

The conservation of total probability often leads to a balance of direct and inertial components, expressed by the Price equation. For example, when we analyze normalized likelihoods such that the average value is one, $\bar{L} = \mathbf{q} \cdot \mathbf{L} = 1$, we have a conserved form of the Price equation for normalized likelihood

$$\Delta \bar{L} = \Delta \mathbf{q} \cdot \mathbf{L} + \mathbf{q}' \cdot \Delta \mathbf{L} = 0, \quad (20)$$

in which the first term is the gain in performance for the direct force of the data in the likelihood, and the second term is a balancing inertial decay imposed by the rescaling of relative likelihood in each update. Notions of direct and inertial forces and total virtual work provide insight into certain types of learning updates, as shown in later examples.

5. Alternative Perspectives of Dynamics

The preceding sections described the universal geometry of change revealed by the Price equation's notation. Fundamental concepts emerged naturally, including the Fisher–Rao length, information geometry, and Bayesian updating. In this context, the Price equation is a purely descriptive approach that reveals abstract, universal mathematical properties of learning updates.

However, in practice, learning dynamics are more than descriptions of change. Algorithms infer causes or deduce outcomes. This section links the Price equation's description to inductive and deductive perspectives.

By making these alternative perspectives explicit, we see why the same mathematical object, such as the Fisher information metric, arises as a simple notational consequence in one context, an empirical observation in another context, and a chosen design principle for optimal performance in a third context [26,27].

To clarify these perspectives, we begin with the fact that any dynamic process has three key components: an initial state, a rule for change, and a final state. Typically, we know or assume two of the components and infer the third. The following subsections consider the alternative perspectives in detail, connecting each back to the abstract Price equation foundation.

This section's alternative perspectives on dynamics add a complementary axis to the local-to-population spectrum of methods introduced earlier. I develop this complementary axis at the population scale, providing the most general conceptual frame. That population context prepares the ground for later discussion of particular algorithms, many of which blend local geometry with broader spatial or temporal scope.

5.1. Descriptive, Inductive, and Deductive Perspectives

In our Price formulation, the partial change associated with force, $\Delta_{\mathbf{f}}$, describes the initial state as the frequencies, $\mathbf{q}$, the rule for change as the fitnesses, $\mathbf{w}$, and the updated state as $\mathbf{q}'$.

(1) The Price formulation is a purely descriptive and exact expression because it tautologically defines the rule for change from the other two pieces, $\mathbf{w} = \mathbf{q}'/\mathbf{q}$. Consequently, results that follow directly from the Price equation provide the general, abstract basis for understanding intrinsic principles and geometry [3].

(2) In biology, actual frequencies change, $\mathbf{q} \mapsto \mathbf{q}'$. Those changes are driven by an interaction between the current state, $\mathbf{q}$, and unknown natural forces. A biological system has, in effect, direct access to the data for initial and updated states but not for the hidden rules of change.

Natural selection implicitly runs an inductive process [28,29]. It infers aspects of the hidden rules for change, designing systems that use that inferred information to improve future performance. In general, a system may use data on the initial and updated states inductively to infer something about the hidden rules of change.

(3) Most mathematical theories and most learning algorithms run deductively. They start with the initial state and the rule for change and then deduce the updated state. For example, we might have $\mathbf{q}$, and the performance function, $\mathbf{w} = \mathbf{U}(\theta)$, from which we calculate $\mathbf{q}'$.

The updated state $\mathbf{q}'$ is an intrinsic calculation or outcome of the system process. More commonly, in machine learning, the process acts on a single parameter vector, θ , rather than as a population of alternative parameter vectors. Given θ and a performance function, the algorithm calculates an updated vector, θ' .

In summary, dynamics has three components: initial state, rule for change, and final state. Descriptive systems define the rule from the other two, $\mathbf{w} = \mathbf{q}'/\mathbf{q}$. Inductive systems start with the initial and final state and infer the rule, $(\mathbf{q}, \mathbf{q}') \mapsto \mathbf{w}$. Deductive systems start with the initial state and the rule and deduce the final state, $(\mathbf{q}, \mathbf{w}) \mapsto \mathbf{q}'$ for populations or $(\theta, \mathbf{U}(\theta)) \mapsto \theta'$ for single-vector updates.

5.2. Fisher Information in the Three Perspectives

This subsection shows how to interpret the squared Fisher–Rao step, $\mathcal{F}$, in each of the three perspectives. In the pure Price equation, it follows simply and universally from tautological notation. In both the inductive and deductive cases, it is the optimal step in the sense that it maximizes the increase in performance relative to alternative steps of the same length.

5.2.1. Descriptive Perspective

In the abstract mathematical perspective, use Equation (4) to define the focal trait as the average excess in fitness

$$a_i = w_i - 1 = \frac{\Delta q_i}{q_i}.$$

Then the partial change in fitness from the first Price term, from Equation (10), is

$$\Delta_f \Delta \bar{w} = \Delta \mathbf{q} \cdot \mathbf{a} = \left\| \frac{\Delta \mathbf{q}}{\sqrt{\mathbf{q}}} \right\|^2 = \mathcal{F}.$$

By analogy with Equation (9), this quantity can also be written as

$$\Delta_f \Delta \bar{w} = \mathbf{f}^\top \mathbf{M} \mathbf{f} = \mathbf{a}^\top \mathbf{S}^{-1} \mathbf{a},$$

interpreted here purely in frequency space such that the force is $\mathbf{f} = \mathbf{a}$, and the metric is $\mathbf{M} = \mathbf{S}^{-1}$, a matrix with entries q_i along the diagonal.

The matrix $\mathbf{S}$ with entries $1/q_i$ along the diagonal is the Fisher information metric for the probability distribution $\mathbf{q}$, also called the Shahshahani Riemannian metric in certain applications [12].

Here, I use the Fisher metric in its geometric sense as the Shahshahani metric, a full-rank, diagonal matrix that defines the curvature of frequency space [12,14]. In this geometric context, the inverse given here provides the length metric for the space of $\mathbf{a}$

values. In classic statistical estimation theory, the Fisher matrix has a different interpretation that reduces its dimension and leads to a different inverse form [10].

In this pure Price case, the values of $\mathbf{f}$ and $\mathbf{M}$ arise directly from notation, without any additional concepts or assumptions derived from information or particular aspects of geometry.

Thus, the Fisher metric may arise so often in widely different applications because of its fundamental basis in simple definitions rather than in the more complex interpretations commonly discussed. Those extended interpretations are useful in particular contexts, as in the following paragraphs.

The point here concerns the genesis and understanding of fundamental quantities rather than their potential applications. For this pure Price case, $\mathbf{M}$ arises purely from tautological notation and is related to the inverse Fisher metric.

5.2.2. Inductive Perspective

In inductive applications, we begin with or observe $\mathbf{q} \mapsto \mathbf{q}'$. From those data, we may inductively estimate $\mathbf{f}$, the slope of w with respect to trait values, θ . The given frequency changes also inductively improve the system's internal estimate for parameters that perform well by weighting more heavily the high-performance parameter values.

In general, from Equation (6), the update is $\Delta_{\mathbf{f}} \bar{\theta} = \mathbf{M} \mathbf{f} = \mathbf{Cov}(w, \theta)$, and the system's partial improvement in fitness (performance) caused by frequency change is $\mathbf{f}^T \mathbf{M} \mathbf{f}$, which is the squared Fisher–Rao length. Here, $\mathbf{M}$ is θ 's covariance matrix, and $\mathbf{f}$ is the vector of partial regression coefficients of w with respect to θ or, in the infinitesimal limit, the inferred gradient of w with respect to θ .

The covariance matrix $\mathbf{M}$ in parameter space, θ , is related to the Fisher metric $\mathbf{S}$ in probability space, $\mathbf{q}$, by

$$\mathbf{M} = \mathbf{J}^T \mathbf{S}^{-1} \mathbf{J}, \quad (21)$$

in which $\mathbf{J}$ is the matrix of parameter deviations from their mean values, $J_{ij} = \theta_{ij} - \bar{\theta}_j$. This expression reveals the relation of the Fisher metric geometry for probabilities to the covariance metric geometry for parameters.

The Fisher–Rao update is optimal in the sense that, at each step, it maximizes the first-order performance gain minus a penalty proportional to the geometric length of the update. In particular, among all possible frequency changes, $\Delta \mathbf{q}$, that produce the same Fisher–Rao update length, the actual frequency changes for the given trait covariance matrix, $\mathbf{M}$, lead to the greatest improvement in performance.

Equivalently, among all possible frequency changes, $\Delta \mathbf{q}$, that produce the same improvement in performance, $\Delta_{\mathbf{f}} \bar{w}$, the actual frequency changes for the given trait covariance matrix, $\mathbf{M}$, lead to the shortest Fisher–Rao length. In biology, these optimality results are trait-based analogs of Fisher's Fundamental Theorem for genetics [30–33].

Of course, forces other than intrinsic performance can alter frequencies. The more we know about those other forces, such as environmental shifts or directional mutation in biology, the more accurately we can account for the consequences of those forces and improve inductive inference about the causal relation between parameters and performance [34–36].

5.2.3. Deductive Perspective: Frequencies

In deductive studies of frequency, we use $\mathbf{q}$ and $\mathbf{w}$ to calculate $\mathbf{q}'$. Mathematically, there is nothing new here because $q'_i = q_i w_i$. However, this perspective differs because we take the $\mathbf{w}$ as given and deduce $\mathbf{q}'$. In other words, $\mathbf{w}$ is an identified driving force, whereas in the inductive case, $\mathbf{w} = \mathbf{q}'/\mathbf{q}$ is an observation about frequencies that leads to inference about the variety of forces that have acted to change frequency. However, because

the mathematics is the same, we end up with the same covariance and other expressions for change.

5.2.4. Deductive Perspective: Parameters

In deductive studies of parameters, we use $\mathbf{f}$ and $\mathbf{M}$ to deduce system updates to parameter values. This approach becomes interesting when, instead of being constrained by the given frequencies, $\mathbf{q}$, that set the covariance matrix of θ values as the $\mathbf{M}$ metric, we instead choose $\mathbf{M}$ to achieve a better increase in performance.

Suppose, for example, that for $\mathbf{f}$ we use $d\mathbf{U}/d\theta$, the gradient of performance (fitness) with respect to the parameters. If the gradient provides the best opportunity for improvement in a particular direction of the parameter space, but there is no parameter variation in that direction among the given $\mathbf{q}$ values, then the associated covariance of θ and metric $\mathbf{M}$ prevents the potential gain offered by the gradient.

In a deductive application, we may instead choose $\mathbf{M}$ to take advantage of the potential increase provided by the gradient. As before, the parameter update is $\Delta_{\mathbf{f}}\theta = \mathbf{M}\mathbf{f}$, and the gain in performance is $\Delta_{\mathbf{f}}\mathbf{U} = \mathbf{f}^{\top}\mathbf{M}\mathbf{f}$. In this case, we use a local gradient so the steps are accurate to first order, we drop the bar over θ because we no longer have an underlying distribution, $\mathbf{q}$, and we use $\mathbf{U}$ for performance.

An optimal update typically occurs when the metric $\mathbf{M}$ is $\mathbf{G}^{-1}$, in which $\mathbf{G}$ is the Fisher information matrix in θ coordinates. That step is optimal in the sense that it provides the greatest increase in performance among all alternative updates with the same Fisher–Rao path length. For an optimal update, the squared Fisher–Rao length equals the gain in performance.

However, we require an arbitrary assumption to calculate the Fisher matrix. That matrix, and the associated optimality, depend on the positive weights assigned to alternative parameter vectors, which we usually express as probabilities. But, in this case, we are considering an update to a given vector, θ , without any underlying variants associated with probabilities, $\mathbf{q}$. So we must create a notion of alternative parameter values with varying weights.

We are free to choose those probability weights, $\mathbf{q}$. A natural choice is to use Boltzmann probabilities of the performance function,

$$\mathbf{q}(\theta) \propto e^{b\mathbf{U}(\theta)}, \quad (22)$$

in which b is a constant value, the maximum of $\mathbf{U}$ is not infinite, and $\mathbf{U}$ is twice differentiable. Here, b adjusts how quickly the log probabilities rise in proportion to performance, $\mathbf{U}$.

For the parameter vector $\theta = \theta_1, \dots, \theta_n$, the i th row and j th column entries of the Fisher matrix are

$$\mathbf{G}_{ij} = -\mathbf{E}\left(\frac{\partial^2 \log \mathbf{q}(\theta)}{\partial \theta_i \partial \theta_j} \middle| \theta\right).$$

The Boltzmann expression in Equation (22) links the log-probabilities used in the Fisher matrix to the performance function, yielding

$$\mathbf{G}_{ij} = -\mathbf{E}\left(\frac{\partial^2 \mathbf{U}(\theta)}{\partial \theta_i \partial \theta_j} \middle| \theta\right) = -\mathbf{E}[\mathbf{H}(\theta)].$$

Thus, the Boltzmann choice for probability weights means that the Fisher matrix is the negative expected value of $\mathbf{H}$, the Hessian matrix of second derivatives of $\mathbf{U}$ with respect to θ . The expectation is over the Boltzmann probabilities for each Hessian evaluated at a particular θ . Thus, when we choose our metric as $\mathbf{M} = \mathbf{G}^{-1}$, the inverse Fisher matrix, we are choosing a particular metric of inverse curvature.

In the inductive case, $\mathbf{M}$ arises from the given frequencies for alternative parameter vectors. For this deductive case, we allow $\mathbf{M}$ to vary and ask what matrix maximizes the gain in performance minus the cost for the Fisher–Rao path length. The next subsection shows that $\mathbf{M} = \mathbf{G}^{-1}$ is optimal in this context and, in general, that inverse curvature is often a good metric [11].

5.3. Why Inverse Curvature Is a Good Metric

Consider first the case in which the only information we have is the local gradient, $\mathbf{f}$, and the local Hessian curvature, $\mathbf{H}$, near a particular point, θ , the local end of the spatial extent spectrum. Then, a Taylor series expansion of performance at a nearby point up to second order is

$$U(\theta + \Delta\theta) = U(\theta) + \mathbf{f}^\top \Delta\theta + \Delta\theta^\top \mathbf{H} \Delta\theta / 2, \quad (23)$$

in which $\mathbf{H}$ is the Hessian matrix of second derivatives. If we consider a region of the performance surface in which all second derivatives are negative, then $\mathbf{M}^{-1} = -\mathbf{H}$ is a positive definite metric that describes local curvature. We can then write the gain in performance for a step $\Delta\theta$ as

$$\Delta U = \mathbf{f}^\top \Delta\theta - \Delta\theta^\top \mathbf{M}^{-1} \Delta\theta / 2. \quad (24)$$

What step $\Delta\theta$ maximizes the total performance gain up to second order? Equivalently, what step maximizes the first-order gain, $\mathbf{f}^\top \Delta\theta$, for a fixed quadratic cost, $\Delta\theta^\top \mathbf{M}^{-1} \Delta\theta = c$? Lagrangian maximization of the gain subject to the fixed cost yields the optimal direction for an update,

$$\Delta\theta^* \propto \mathbf{M} \mathbf{f}, \quad (25)$$

which in this context is the classic Newton optimization method. For maximization, the positive definite metric is $\mathbf{M} = (-\mathbf{H})^{-1}$, in which $\mathbf{H}$ is the negative-definite Hessian matrix of the performance function U . For minimization of a cost function, the sign of $\mathbf{H}$ reverses, so that the metric becomes $\mathbf{M} = \mathbf{H}^{-1}$, and the force $\mathbf{f}$ also flips its sign. Same idea, different signs [5].

In this case, the metric $\mathbf{M}$ scales each component of the step inversely to the local curvature, pushing far in straight directions and contracting where the surface bends sharply. That inverse-curvature rescaling gains the most in first-order performance change for a given quadratic cost.

Considering that the optimal metric for a local step arises from the inverse Hessian, why use the more complex Fisher metric? One reason is that local optimality requires the Hessian to be negative definite for the maximization of performance or, equivalently, positive definite for the minimization of cost. Another reason is that a local calculation ignores the overall shape of the optimization surface. Therefore, a locally optimal step is not necessarily best with regard to broader goals of optimization.

By contrast, the Fisher metric is essentially an average of the best local curvature metrics over a region of the optimization surface, weighting each location on the surface in proportion to a specified probability. This approach, known as the natural gradient, typically points in a better direction with regard to global optimization [11].

In terms of our local-to-population spectrum of methods, the natural gradient combines the spatial population extent for the calculation of the curvature metric with the local extent for the analysis of the force gradient.

For the Boltzmann distribution, the probability weighting of a location rises with the performance associated with that location, emphasizing strongly those regions of the optimization surface associated with the highest performance. Thus, a Fisher step typically points in a better direction with regard to global optimization.

The Fisher metric also corrects common problems with local Hessians. For example, Hessians are not always proper positive metrics, whereas the Fisher matrix is a proper metric. In addition, local Hessians can change under coordinate transformation, whereas the Fisher metric is coordinate invariant. As the region over which the Fisher metric is defined converges to a local region, the Fisher metric converges to the local Hessian.

The optimality of the Fisher metric follows the same sort of Lagrangian maximization as for the local Hessian [14]. In particular, we maximize the same gain, $\mathbf{f}^\top \Delta \theta$, but this time subject to a fixed KL divergence, $\mathcal{D}(\mathbf{q}' || \mathbf{q})$, between the probability weightings for variant parameter values, taken initially as $\mathbf{q}(\theta)$ and after a parameter update as $\mathbf{q}' = \mathbf{q}(\theta + \Delta \theta)$. Here, $\mathbf{q}(\theta)$ is a general distribution that can take any consistent form. By the Taylor series, the KL divergence to second order is

$$\mathcal{D}(\mathbf{q}' || \mathbf{q}) = \Delta \theta^\top \mathbf{G} \Delta \theta / 2$$

for Fisher matrix $\mathbf{G}$. We can use the right side in place of the quadratic term in Equation (24). Then the same Lagrangian approach yields the optimal update in the context of the Taylor series approximations as

$$\Delta \theta^* \propto \mathbf{G}^{-1} \mathbf{f}, \quad (26)$$

which has the same form as the classic Newton update in Equation (25) but with $\mathbf{M} = \mathbf{G}^{-1}$, the inverse Fisher metric in place of the positive definite form of the inverse Hessian. In the limit of small changes, twice the KL divergence becomes the squared Fisher–Rao path length, $2\mathcal{D} \rightarrow \mathcal{F}$. Thus, the optimality again becomes the maximum performance gain relative to a fixed Fisher–Rao length.

In practice, the Fisher metric does not always provide the best update step. That metric is based on a particular assumption about global probability weightings for alternative parameter vectors, whereas one might be more interested in the local geometry near a particular point in the parameter space. In some cases, a local estimate of the inverse curvature provides a better update or may be cheaper to calculate.

The variety of learning algorithms trades benefits gained for particular geometries against costs paid for specific calculations. However, inverse curvature remains a common theme across many algorithms.

6. The Variety of Algorithms

The following sections step through some common algorithms. The details show how each fits into the FMB scheme, how the various algorithms relate to each other, and why certain quantities, such as Fisher information, recur in seemingly different learning scenarios.

The key distinctions between algorithms arise from how each gathers information about components of the FMB law. The alternatives span the spectrum from local information taken at the current point in parameter space to spatially extended averaging over a population of alternative parameter vectors, and from current values to temporally extended averaging of past or anticipated future locations. This section provides a brief overview.

First, population and Bayesian methods represent the fully extended spatial scope end of the spectrum. These methods focus on changes in frequencies, $\Delta \mathbf{q}$. In biology, frequencies change between ancestor and descendant populations. In Bayesian methods, frequencies change between prior and posterior distributions. The frequencies act as relative weights for alternative parameter vectors.

A particular algorithm can, of course, choose to modify how components of an update are calculated. However, populations set the foundation for analysis. Changes in parameter mean values, $\Delta \bar{\theta}$, summarize updates. The metric $\mathbf{M}$ typically arises from the covariance of

alternative parameter values. Some methods use variational optimization and analogies with free energy, which links learning to various physical principles of dynamics [37,38].

Second, many algorithms update a single parameter vector. These methods fill in the other end of the spectrum and its middle ground. The purely local strategy, such as Newton's method, uses a metric and force gradient analyzed at a single point. Hybrid strategies choose metrics that incorporate a broader spatial scope, such as trust regions [39,40] or the natural gradient [11].

Third, all search methods trade off exploiting the directly available information in their spatial domain against exploring more widely to avoid getting stuck in local optima. Noise provides the simplest exploration method, often expanding the spatial scope of analysis. A broader scope can sometimes push the search beyond a local plateau to find the nearest advantageous gradient to climb.

Finally, modern optimizers often include temporal scope. Extensions to stochastic gradient descent, such as Adam [8] and Nesterov [41,42], explicitly use the history of updates to calculate a bias term, $\mathbf{b}$. In these cases, bias explicitly applies physical notions of momentum, in which past movement of parameter values can push future updates beyond local traps on the performance surface.

Overall, common optimizers combine different spatial and temporal extents of gradient forces, inverse curvature metrics, bias, and algorithmic tricks that compensate for missing information, difficult calculations, and challenging search over complex performance surfaces. We see that the seemingly different algorithms all build on the same underlying principles revealed by the Price equation's FMB law.

I start with spatially extended population methods, which match most closely to the standard interpretation of the Price equation.

7. Population-Based Methods

In machine learning, one often improves performance by directly calculating the gradient of a performance surface. An updated parameter vector follows by moving along the gradient's direction of improved performance.

However, in many applications, one does not have a smooth, differentiable function that accurately maps parameters to performance. Without the ability to calculate the gradient, one has to test each parameter combination in its environment to measure its performance. Such methods are called black-box optimization algorithms.

Covariance matrix algorithms (CMAs) often provide a good black-box optimization method [43]. These evolution strategy (ES) methods proceed by analogy with biological evolution. One starts with a target parameter vector and then samples the performance values for a set of different parameter combinations around the target. That set forms a population from which one can calculate an updated target parameter by an empirically calculated covariance matrix and performance gradient, as in Equation (6).

These methods have three challenges: choosing sample points around the current target, estimating the performance gradient, and calculating the new target from the covariance matrix and performance gradient. I briefly summarize how the popular CMA-ES method handles these three challenges [44].

First, CMA-ES draws sample points from a multivariate Gaussian distribution. The mean is the current target parameter combination. The algorithm updates the covariance matrix to match its estimate of the performance surface curvature. The covariance is reduced along directions with high curvature and enhanced along directions with low curvature.

The size of the parameter changes in any direction increases with the variance in that direction. Thus, the algorithm tends to explore more widely over straighter (low curvature) regions of the performance surface and to move more slowly in curved directions.

In biology, covariance decreases over time in directions of strong selection because better-performing variants increase rapidly and reduce variation. That decline in variation degrades the ability of the system to continue moving in the same beneficial direction because the distance moved in a direction is the selection gradient multiplied by the variance. Eventually, mutation or other processes may restore the variance, but it can take a while to restore variance after a bout of strong selection.

CMA-ES avoids that potential collapse in evolutionary response by algorithmically maintaining sufficient variance to provide the system with good potential to search the performance landscape. In essence, the algorithm attempts to choose the inverse curvature metric, $\mathbf{M}$, to maximize the gain in performance, in which $\mathbf{M}$ is the covariance matrix.

The algorithm chooses $\mathbf{M}$ by modifying the inverse Fisher metric of its Gaussian sampling distribution, steadily blending in the weighted variation of the best-performing samples to estimate its covariance matrix for updates.

The estimated covariance matrix typically converges to an approximation for the local inverse curvature metric. For performance maximization, this metric is given by $(-\mathbf{H})^{-1}$ for the local Hessian, $\mathbf{H}$, in which the negative sign ensures that the metric is positive definite. The Gaussian sampling process introduces stochasticity that corresponds to ζ in the FMB law.

For the second challenge, CMA-ES approximates a modified performance gradient by sampling candidate solutions, ranking them by fitness, and then averaging the parameter differences between the best candidates and the current mean. The average puts greater weight on the higher fitness candidates. By contrast, biological processes implicitly calculate the partial regression of fitness on parameters.

For the third challenge, the direction in which CMA-ES updates the target parameter vector is approximately the product of its internal estimates for the covariance matrix and modified performance gradient. That approximation follows biology's updating by $\Delta_{\mathbf{f}} \bar{\boldsymbol{\theta}} = \mathbf{M} \mathbf{f}$ given in Equation (6), which is the universal combination of force scaled by metric.

The same $\mathbf{M} \mathbf{f}$ structure occurs in other evolution strategy (ES) learning algorithms, which include natural evolution strategies [45], large-scale parallel ES [46], and separable low-rank CMA-ES [47]. The next section moves to the local end of the spectrum, showing that the $\mathbf{M} \mathbf{f}$ structure remains when the population collapses to a single vector.

8. Single-Vector Updates

In population methods, we track a weighted set of alternative parameter vectors over a spatially extended region of the parameter space. We often summarize learning updates by changes in the population mean, $\Delta_{\mathbf{f}} \bar{\boldsymbol{\theta}}$. Here, I use the partial component, $\Delta_{\mathbf{f}}$, of the FMB law to focus on methods that use only the force and metric terms.

Many common algorithms update a single local parameter vector, $\Delta_{\mathbf{f}} \boldsymbol{\theta}$, without tracking any variant parameter values. This section summarizes a few classic single-vector update methods. The next section adds noise to enhance exploration of complex optimization surfaces, as in the commonly used stochastic gradient descent algorithm. The following section adds bias, including a momentum term that often occurs in some common machine learning algorithms, such as Adam.

For the $\mathbf{M} \mathbf{f}$ component of the FMB law, the following algorithms differ mainly in the way that they choose $\mathbf{M}$. In some cases, $\mathbf{M}$ arises from the same local extent as the gradient force calculation. In other cases, an algorithm expands the temporal or spatial extent to obtain a broader calculation of the curvature metric or uses a mirror geometry to design specific curvature attributes into the method.

8.1. Gradient Descent: Constant Euclidean Metric

The update is

$$\Delta_{\mathbf{f}} \boldsymbol{\theta} = \boldsymbol{\eta} \mathbf{f},$$

in which $\boldsymbol{\eta}$ is the step size, and $\mathbf{f}$ is the gradient of the performance function with respect to the parameters, evaluated at the current parameter vector, a purely local method. The implicit metric is $\mathbf{M} = \boldsymbol{\eta} \mathbf{I}$, in which the identity matrix, $\mathbf{I}$, is the Euclidean metric with no curvature. This simple method typically traces a path along the performance surface from the current location to the nearest local optimum.

8.2. Newton's Method: Exact Local Curvature

In Equation (25), we noted that using the positive definite inverse Hessian for the metric, $\mathbf{M} = \tilde{\mathbf{H}}^{-1}$, yields Newton's method

$$\Delta_{\mathbf{f}} \boldsymbol{\theta} = \tilde{\mathbf{H}}^{-1} \mathbf{f},$$

in which $\tilde{\mathbf{H}} = -\mathbf{H}$ if we are maximizing performance, and $\tilde{\mathbf{H}} = \mathbf{H}$ with $\mathbf{f} \mapsto -\mathbf{f}$ if we are minimizing cost.

The Hessian is the matrix of second derivatives of the performance function with respect to the parameters at the current single point in the parameter space, a purely local measure of curvature. We could also include a step size multiplier, $\boldsymbol{\eta}$, for any algorithms. However, we drop that step size term to reduce notational complexity.

The inverse Hessian metric typically improves local optimization by orienting the direction of updates into an improved gain in performance relative to a cost paid for the squared Fisher–Rao path length movement.

On the downside, the second derivatives may not exist everywhere, the Hessian may be computationally expensive to calculate, the information about second derivatives may be lacking, and this method tends to become trapped at the nearest local optimum.

In theory, using the inverse Fisher metric in place of the inverse Hessian often provides a better update [11]. The Fisher metric measures the curvature of the performance function with respect to the parameters when averaged over a spatially extended region of the parameter space.

The section *Why inverse curvature is a good metric* discussed the many theoretical benefits of Fisher information in this context [11,14]. For example, the global Fisher metric often reduces the tendency to be trapped at a local optimum when updating the parameters. However, in practice, one often uses local estimates of curvature to avoid the extra assumptions and complex calculations required to estimate the Fisher matrix.

8.3. Quasi-Newton: Temporal Extension

These methods replace the computationally costly exact Hessian calculation with simpler updates that accumulate curvature information over time. In terms of our FMB analysis, $\mathbf{M}$ is iteratively updated to approximate $\tilde{\mathbf{H}}^{-1}$, so that the temporally extended estimation of the metric becomes part of the algorithm [5].

The Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm and its variants are widely used [48]. After each step, the method compares how the gradient changed to how the parameters moved. That comparison provides information about curvature.

From that curvature information, the algorithm keeps a running update of its estimate for the inverse Hessian matrix. Thus, the method estimates the curvature metric by using only its calculation of first derivatives and parameter updates, without ever directly calculating a second derivative or inverting a matrix.

8.4. Trust Regions: Spatial Extension

Newton's local calculation of the curvature via the inverse Hessian is fragile. The local Hessian may not exist, it may change significantly over space, or it may not be invertible to provide the inverse curvature. Trust region methods may solve some of these problems by defining a local region of analysis [39,40].

One type of trust region method spatially extends the calculation of the curvature metric in a way that matches the FMB law's notion of a population. By combining a local force gradient with a metric that is a positive definite average of the curvature over the region, the hybrid method follows Amari's natural gradient approach [11]. The spatial extension of the metric calculation often improves updates by providing more information about the shape of the local performance landscape.

8.5. Mirror Descent: Transformational Extent

The previous examples in this section match the $\mathbf{M}\mathbf{f}$ force–metric pattern in a simple and direct way. By contrast, the mirror descent algorithm has a more complicated geometry that does not obviously map directly to our $\mathbf{M}\mathbf{f}$ pattern [49,50]. However, the following derivations show that this algorithm does, in fact, use the same basic force–metric approach, but with a more involved method to obtain the curvature metric.

In this case, the challenge is that the curvature either cannot be calculated or is computationally too expensive to calculate. So one calculates curvature in an alternative mirror geometry obtained by transformation of the target geometry, in which the curvature has a simpler or more tractable form. Instead of improving the metric by spatial or temporal extent, one uses a transformational extent to a geometry that can be chosen for its anticipated benefits to algorithmic learning.

For example, in a Newton update, the inverse Hessian is taken with respect to the performance function, $U(\theta)$. In mirror descent, one changes the geometry by choosing a strictly convex potential function $\varphi(\theta)$ that has a positive definite Hessian $\mathbf{H}_\varphi$ over the search domain. The inverse $\mathbf{H}_\varphi^{-1}$ provides a consistent positive definite metric $\mathbf{M}$ for the update, repairing any fragility of the Hessian of U .

Assume throughout this subsection that we are maximizing performance despite the word *descent* in the common name for this approach.

Allowing for variable step size, η , the generalization of the Newton update for maximization becomes

$$\Delta_t \theta = \eta \mathbf{H}_\varphi^{-1} \mathbf{f}. \quad (27)$$

This update is a first-order approximation of the mirror descent method [49,50]. In effect, we are making a Newton-like step in an alternative mirror geometry with metric $\mathbf{H}_\varphi^{-1}$, then mapping that step back to our original geometry for θ .

This approach allows one to control the metric and to compensate for the geometry of the performance surface that does not have a simple or proper curvature. Because the change in mirror space is determined by the metric $\mathbf{H}_\varphi^{-1}$, we end up with our simple force–metric expression from the FMB law.

The optimal update rule in Equation (27) arises as the first-order approximation for the solution of the following optimization problem, which balances the performance gain in the target geometry against a distance penalty in the mirror geometry. The problem is

$$\theta_{t+1} = \arg \max_{\theta} \left\{ \nabla U(\theta_t) \cdot (\theta - \theta_t) - \frac{1}{\eta} B_\varphi(\theta | \theta_t) \right\}. \quad (28)$$

The new update vector θ_{t+1} maximizes the two bracketed terms. The first term is the first-order approximation for the total performance increase. The second term is the Bregman

divergence, B_φ , between new and old parameter vectors measured in the mirror geometry defined by transformation, $\varphi(\theta)$. That divergence, defined below, is an easily calculated distance moved in the mirror geometry. Thus, we are maximizing the performance gain in the target geometry minus the distance moved in the mirror geometry scaled by $1/\eta$.

In some cases, the mirror transformation depends on local information that changes in each time step, denoted φ_t to emphasize the time dependence. Here, for notational simplicity, we drop the t subscript but allow such dependence.

In this update equation, the first right-hand term is the local slope of the performance gain relative to the parameter change, weighted by the amount of parameter change, yielding the total performance increase. In the second term, the Bregman divergence is

$$B_\varphi(\theta||\theta_t) = \varphi(\theta) - \varphi(\theta_t) - \nabla\varphi(\theta_t) \cdot (\theta - \theta_t).$$

We can obtain the first-order approximation for the optimal update given in Equation (27) by the following. Start by differentiating the terms in the brackets on the right-hand side of Equation (28), evaluating at $\theta = \theta_{t+1}$, and setting to zero, which yields

$$\nabla\varphi(\theta_{t+1}) = \nabla\varphi(\theta_t) + \eta\nabla U(\theta_t). \quad (29)$$

Noting that $\theta_{t+1} = \theta_t + \Delta\theta$, a first-order Taylor expansion of $\nabla\varphi(\theta_{t+1})$ around θ_t is

$$\nabla\varphi(\theta_{t+1}) = \nabla\varphi(\theta_t) + \nabla^2\varphi(\theta_t)\Delta\theta + \mathcal{O}(\|\Delta\theta\|^2).$$

Dropping the second-order error and substituting into Equation (29) gives Equation (27) to first order in η , noting that $\mathbf{H}_\varphi = \nabla^2\varphi(\theta_t)$. Thus, even in this relatively complex case, the force–metric law is nearly exact for small updates. This method is purely local in the same sense as Newton’s algorithm, but uses a transformational extent to improve the analysis of the curvature metric.

9. Stochastic Exploration

All gradient algorithms suffer from the tendency to become stuck at a local optimum. This section briefly summarizes two methods that use noise to escape local optima and explore the performance surface more broadly.

These methods match the FMB law, $\mathbf{M}\mathbf{f} + \mathbf{b} + \boldsymbol{\zeta}$, with no bias, $\mathbf{b}$, and noise entering via the stochastic term, $\boldsymbol{\zeta}$. The deterministic component of these methods, $\mathbf{M}\mathbf{f}$, is a single-value update driven by the gradient force, lacking a population.

During a search trajectory, when the magnitude of this deterministic gradient component is greater than the magnitude of the noise, the deterministic single-value update process dominates. When the gradient flattens or the step size weighting for the gradient shrinks, the noise dominates the updates.

In the noise-dominated regime, the temporal trajectory samples a population of parameter locations around the path that the deterministic trajectory would have traced. That temporal wandering changes the search from single-value updates to a quasi-Bayesian population-based method, in which the population distribution is shaped by the covariance of the noise process [16,51].

Technically speaking, as the gradient-to-noise ratio declines, the temporal stochastic sampling effectively becomes an ergodic spatially extended population. Thus, these hybrid methods exploit the simplicity and efficiency of single-value updates when the gradient is strongly informative and exploit the broader exploratory benefits of populations when the step-size weighted gradient is relatively weak.

9.1. Stochastic Langevin Search

This algorithm combines a deterministic gradient step and a noise fluctuation [51]. A simple stochastic differential equation describes the process as follows:

$$d\theta = \nabla U(\theta)dt + \sqrt{2D}dW_t. \quad (30)$$

The parameter vector update, $d\theta$, equals the gradient of the performance surface, $\nabla U = \mathbf{f}$, with respect to the parameters, θ , plus a Brownian motion vector, W_t , weighted by the square root of the diffusion coefficients in the matrix, D , that determines the scale of noise.

In practice, one typically updates by discrete steps, for example, at time t , the update may be

$$\Delta\theta_t = \eta M \nabla U(\theta_t) + \sqrt{2\eta M} \epsilon_t.$$

here, η adjusts the step size.

The metric M scales motion in each direction, typically obtained by estimating the positive definite inverse curvature, such as the inverse of the negative Hessian matrix. The vector ϵ_t is Gaussian noise with mean zero and covariance given by the identity matrix. The overall noise, $\xi_t = \sqrt{D} \epsilon_t$, has a covariance matrix of $D = 2\eta M$. Alternative discrete-step approximations for Equation (30) may be used to improve accuracy or efficiency.

This update matches our FMB standard for learning algorithms, the gradient force multiplied by the inverse curvature metric. The noise term adds exploratory fluctuations in proportion to the inverse curvature metric, M .

When the gradient is relatively large compared to the noise, the deterministic force dominates. When the slope is flat, noise dominates. The weighting of noise by inverse curvature means that fluctuations explore more widely in directions with low gradient and small curvature. Such directions are the most likely to be associated with trend reversals, providing an opportunity to escape a region in which the gradient pushes in the wrong direction.

The relative dominance of the deterministic gradient component compared with the noise component can be adjusted by changing η , the step-size weighting of the gradient term. As the η -weighted deterministic component declines, the method increasingly shifts from single-value updates to population-based updates. Here, dominance by temporal noise creates a similar effect to a spatially extended population [51], as described in the introduction to this section.

In theory, methods such as stochastic Langevin search provide an excellent balance between exploiting gradient information and exploring by noise. However, the algorithm can be very costly computationally for high dimensions and large data sets. Each step uses all of the data to calculate the gradient with respect to the high-dimensional parameter vector. The size of the inverse curvature matrix is the square of the parameter vector length, which can be very large. Stochastic gradient descent provides similar benefits with lower computational cost.

9.2. Stochastic Gradient Descent

In data-based machine learning methods, stochastic sampling of the data creates exploratory noise, which often improves the learning algorithm's search efficacy. To explain, I briefly review the steps used by common gradient algorithms in machine learning.

The data have N input–output observations. The model takes an input and predicts the output. The parameter vector θ influences the model's predictions.

For each observed input, the model makes a prediction that is compared with the observed output. A function transforms the divergence between the model prediction and the observed output into a performance value. One typically averages the performance

value over multiple observations. The gradient vector is the derivative of the average performance with respect to the parameter vector.

In machine learning, the average performance is typically called the loss for that batch of data. In this article, we often use the negative loss as the positive performance value. Climbing the performance scale means descending the loss scale.

The total data set has N observations. If one uses all of the data to calculate the gradient and then updates the parameters using that gradient, the update method is often called gradient descent.

This deterministic gradient descent method will often become trapped in a local minimum, failing to find parameters that produce better performance.

Instead of using all of the data to calculate the gradient, one can instead repeat the process by using many small random samples of the data. Data mini-batches lose gradient precision but gain exploratory noise [6,15–17].

The stochasticity from random sampling can help to escape local optima and find better parameter combinations [52]. Thus, the method is often called stochastic gradient descent. In this process, a parameter update is

$$\Delta\theta = \eta \widehat{\nabla U}(\theta),$$

in which η is a chosen step size weighting, and the hat over the gradient, $\widehat{U}(\theta)$, denotes an estimated value from the sample batch of data.

The estimated gradient is implicitly the true gradient plus sampling noise. Thus, in theory,

$$\Delta\theta = \eta[\nabla U(\theta) + \xi], \quad (31)$$

in which ∇U is the true gradient, and ξ is the sampling noise of the gradient estimate. The variance of the sampling noise scales inversely with the batch sample size. Choosing a good batch size plays an important role in the optimization process [53,54].

The deterministic component of parameter updates dominates when the noise is small relative to the gradient signal. The stochastic component dominates when the noise is large relative to the gradient signal. Larger batch sizes reduce the scale of noise relative to the gradient signal.

As noted in the prior subsection, noisy exploration provides the greatest benefit when both the gradient and the curvature are small. In that case, noise provides the opportunity to jump across a near-zero or mildly disadvantageous gradient to a new base from which the gradient leads to improving performance. The more a region curves in a disadvantageous direction, the more noise one needs to jump across it.

As the batch size declines, the method increasingly shifts from deterministic single-value updates to stochastically sampled population-based updates. Here, dominance by temporal noise creates a similar effect to a spatially extended population [16,55], as described in the introduction to this section.

Many algorithms build on stochastic gradient descent by adding inverse-curvature metric scaling and bias. The next section provides examples.

10. Bias in Modern Optimization

Several widely used machine learning methods include a bias term. The bias alters parameters in addition to the direct force of the performance gradient.

For example, a moving average of past parameter changes describes the update momentum. That momentum includes temporal information about the shape of the performance surface that goes beyond the information in the local gradient and curvature.

This section provides examples of bias in various machine learning algorithms. Before turning to those examples, I briefly review the role of bias within the Price equation and the FMB law.

10.1. Brief Review of Bias

The Price equation's full FMB law from Equation (1) is

$$\Delta \bar{\theta} = \mathbf{M} \mathbf{f} + \mathbf{b} + \zeta.$$

The prior examples focused on the $\mathbf{M} \mathbf{f}$ direct force and ζ noise terms. This section considers the bias term of Equation (2), repeated here

$$\mathbf{b} = \mathbf{C} \boldsymbol{\beta} + \gamma.$$

Bias describes deterministic changes in parameter values, $\Delta \theta_i = \theta'_i - \theta_i$, that are not caused by the directly acting forces, $\mathbf{f}$. Here, we take $\mathbf{f}$ as the regression or gradient of performance with respect to the parameters.

Denote these bias changes as $\Delta_{\mathbf{b}} \boldsymbol{\theta}$. Then $\gamma = \mathbf{E}_{\mathbf{q}}(\Delta_{\mathbf{b}} \boldsymbol{\theta})$ describes the bias that is uncorrelated with the performance. The matrix $\mathbf{C}$ is the covariance of the bias vector. The vector $\boldsymbol{\beta}$ is the regression of performance on the bias vector, which captures correlations between performance and bias.

For single-value updates to the location vector, we use $\bar{\boldsymbol{\theta}} \mapsto \boldsymbol{\theta}$. The components of bias lose their statistical meaning. Instead, $\mathbf{C}$ is a metric for the space of bias vectors. The slope $\boldsymbol{\beta}$ is the gradient of performance with respect to the bias vector. The term γ adds further bias.

To focus on bias in the FMB law, this section drops noise terms. In effect, assume very large batch sizes for single-value updates or very large population sizes for population mean updates. The following examples decompose particular update algorithms into their FMB components.

10.2. Prior Bias: Parameter Regularization

Consider the single-value parameter update

$$\Delta \boldsymbol{\theta} = \eta (\nabla_{\boldsymbol{\theta}} U - \lambda \boldsymbol{\theta}).$$

The metric is $\mathbf{M} = \eta \mathbf{I}$. The force is $\mathbf{f} = \nabla_{\boldsymbol{\theta}} U$. The bias terms are $\mathbf{C} \boldsymbol{\beta} = \mathbf{0}$ and $\gamma = -\eta \lambda \boldsymbol{\theta}$.

The gradient imposes a force that pushes parameters to improve performance. The bias term imposes a static force that pulls all parameters toward a prior value, in this case, the origin.

Those parameters that only weakly improve performance end up close to the prior. In practice, one often prunes the parameter vector by dropping all parameters that end up near the prior, a process called regularization [15].

10.3. Momentum Bias: Polyak

Abbreviate the gradient at time t as $\mathbf{g}_t = \nabla_{\boldsymbol{\theta}} U(\boldsymbol{\theta}_t)$. Then we can calculate the exponential moving average of the gradient as

$$\mathbf{m}_t = (1 - u) \mathbf{g}_t + u \mathbf{m}_{t-1}.$$

A simple parameter update that includes history is [7]

$$\Delta \boldsymbol{\theta}_t = \eta \mathbf{m}_t = \eta (1 - u) \mathbf{g}_t + \eta u \mathbf{m}_{t-1}.$$

On the right side, the first term is a standard gradient term for the update, $\mathbf{f} = \mathbf{g}_t$, with a constant metric $\mathbf{M} = \eta(1 - u)\mathbf{I}$. The second term is the bias caused by the momentum from past updates, $\gamma = \eta u \mathbf{m}_{t-1}$. In this case, the bias is not associated with performance, $\mathbf{C} \beta = \mathbf{0}$.

The idea is that a strong historical tendency to move in a particular direction provides information about the performance surface that supplements the information in the local gradient at the current time. Thus, the algorithm uses the momentum from past updates to push the current update in the direction that has been favored in the past.

Roughly speaking, the method uses temporal extent to gain information about the shape of the performance surface curvature rather than using the spatial extent of populations. However, in this case, the curvature information is used to bias the update rather than to calculate the metric that rescales the local gradient.

10.4. Momentum Bias and Metric: Adam

The widely used Adam algorithm adds metric scaling to the basic momentum update [8]. The metric arises from the exponential moving average of the squared gradient

$$\mathbf{v}_t = (1 - s)\mathbf{g}_t^2 + s\mathbf{v}_{t-1}.$$

The match to FMB follows by

$$\begin{aligned}\mathbf{f} &= \mathbf{g}_t \\ \mathbf{M} &= \eta(\sqrt{\mathbf{v}_t} + c)^{-1} \\ \mathbf{C} &= \mathbf{M} \\ \beta &= u\mathbf{m}_{t-1}/(1 - u) = \mathbf{m}_t/(1 - u) - \mathbf{g}_t \\ \gamma &= \mathbf{0}\end{aligned}$$

for small constant c in $\mathbf{M}$. The update follows as

$$\begin{aligned}\Delta\theta_t &= \mathbf{M}\mathbf{f} + \mathbf{C}\beta \\ &= \frac{\eta\mathbf{g}_t}{\sqrt{\mathbf{v}_t} + c} + \frac{\eta u\mathbf{m}_{t-1}/(1 - u)}{\sqrt{\mathbf{v}_t} + c} \\ &= \frac{1}{(\sqrt{\mathbf{v}_t} + c)(1 - u)}[\eta(1 - u)\mathbf{g}_t + \eta u\mathbf{m}_{t-1}] \\ &= \frac{1}{1 - u} \left(\frac{\eta\mathbf{m}_t}{\sqrt{\mathbf{v}_t} + c} \right) = \frac{1}{1 - u} \mathbf{M}\mathbf{m}_t.\end{aligned}$$

Note in the third line that the bracketed quantity on the right is the Polyak momentum update from the prior subsection. Thus, Adam is proportional to the Polyak update weighted by the metric $\mathbf{M}$.

Here, the metric is based on the exponential moving average of the squared gradient, $\mathbf{v}_t$. Instead of inverse curvature, this metric is an inverse combination of the magnitude of the gradient and the noise in the gradient estimate when using small random batch samples of the data. Thus, the metric reduces step size in directions that have some combination of a large slope or high noise.

In summary, the metric $\mathbf{M}$, derived from the history of squared gradients, creates an adaptive learning rate tuned to the geometry of each direction, and the momentum bias $\mathbf{b}$ provides a temporally smoothed, directed force. Overall, the temporal extent provides an efficient method to estimate spatial aspects of geometry.

11. Gaussian Processes and Kalman Filters

This section returns to population-based algorithms in relation to force and metric. I first introduce Gaussian processes, a population-based Bayesian method that weights alternative functions by how well they describe data rather than weighting alternative parameter vectors. The Bayesian weighting of alternatives provides a natural measure of uncertainty [56].

I then turn to the common Kalman filter method. That method is conceptually similar to Gaussian processes, using a time series sequence of learning updates to track a changing system rather than the typical one-shot learning update by which Gaussian processes describe static systems [57,58].

In the context of our FMB law, the technical details of the calculations do not matter so much. Instead, the point of this section is that a Gaussian process is a spatially extended population algorithm that depends on our typical metric and force terms to learn about a static process.

Similarly, a Kalman filter uses a Bayesian population approach but instead aims to track a dynamically changing system. The Kalman filter update also reduces to our basic metric and force terms, in which the spatial geometry of the metric is estimated by temporal updating, as in Adam.

11.1. Gaussian Processes

Our previous population methods updated the individual probability weightings in $\mathbf{q}$. Each probability weighting is associated with a multivariate combination of parameter values.

In contrast to a parametric model, a Gaussian process looks for the best function rather than the best parameter vector. Suppose, for example, that we are studying air temperature in relation to various predictors, such as cloud cover and humidity. We measure the actual temperature, $\mathbf{y}$, and the vector of predictors, $\mathbf{x}$.

Previously, we had a set of parameters that told us how to link the predictors to the outcome. The relative weighting of each parameter combination, $\mathbf{q}$, described a Bayesian distribution, which we updated from prior to posterior based on the data.

In a Gaussian process, we use a function, $g(\mathbf{x})$, to predict the temperature, with deviations $\mathbf{y} - g(\mathbf{x})$ between observed and predicted values. The goal is to find the best function among the set of candidate functions, $\mathbf{g}$, that describe the data.

Previously, we described the associations between different parameter values by a covariance matrix. In a continuous Gaussian process, we use a kernel function $k(\mathbf{x}, \tilde{\mathbf{x}})$ that tells us, for any two points $\mathbf{x}$ and $\tilde{\mathbf{x}}$ in the domain of inputs, how to calculate $\text{Cov}[g(\mathbf{x}), g(\tilde{\mathbf{x}})]$. For example, a common kernel function is

$$k(\mathbf{x}, \tilde{\mathbf{x}}) = \sigma_g^2 \exp\left(-\frac{\|\mathbf{x} - \tilde{\mathbf{x}}\|^2}{2\ell^2}\right),$$

in which σ_g^2 is the baseline variance when $\mathbf{x} = \tilde{\mathbf{x}}$, and ℓ is the length scale for how quickly a function can change. Different kernel functions encode different assumptions about smoothness or periodicity, and their hyperparameters can be updated as new data arrive.

We do not explicitly specify the functions, $\mathbf{g}$, and their associated probabilities, $\mathbf{q}$. Instead, we assume that, for an input vector $\mathbf{x}$, the average output we expect over all functions is provided by the mean curve $\mu(\mathbf{x}) = \mathbb{E}[g(\mathbf{x})]$. The curve can be thought of as the plot of $\mu(\mathbf{x})$ versus $\mathbf{x}$.

Two outputs covary by the kernel function, $k(\mathbf{x}, \tilde{\mathbf{x}})$, which tells us how similar the functional output values at $\mathbf{x}$ and $\tilde{\mathbf{x}}$ tend to be. For a particular set of inputs, $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, the covariance matrix $\mathbf{K}$ has entries $k(\mathbf{x}_i, \mathbf{x}_j)$.

For inputs $\mathbf{X}$, the associated vector of mean values and the covariance matrix define a particular multivariate Gaussian. That Gaussian is the Bayesian distribution over all candidate function values for these specific inputs.

The notion of stepwise updating arises because, with new inputs, we alter the mean function, $\mu(\mathbf{x})$, and we may also choose to alter the kernel function. The new mean vector and covariance matrix define the updated multivariate Gaussian posterior distribution over functions.

To define the update, note that for the functions $\mathbf{g}$, the prediction, $\mathbf{y}$, for any particular input, $\mathbf{x}$, is given by the mean function, $\mu(\mathbf{x})$. The set of N inputs generates the vector mean, $\bar{\mathbf{g}} = \mu(\mathbf{X})$, and we also have the associated vector of measured values, $\mathbf{y} = [y_1, \dots, y_N]$. The measured values include measurement error, $y_i = h(\mathbf{x}_i) + \zeta$, in which $h(\mathbf{x}_i)$ is the true value associated with input $\mathbf{x}_i$, and ζ is the unbiased Gaussian measurement error with variance σ^2 .

An update is $\Delta \bar{\mathbf{g}} = \mu_1(\mathbf{X}) - \mu_0(\mathbf{X})$, in which μ_0 is the prior mean function, and μ_1 is the posterior mean function. Typically, we set μ_0 by prior knowledge or assumption. In the update, the difference in mean functions defines μ_1 , the posterior mean function by

$$\Delta \bar{\mathbf{g}} = \mathbf{M} \mathbf{f}, \quad (32)$$

our standard FMB form of the metric multiplied by the force, in which metric and force are

$$\mathbf{M} = (\mathbf{K}^{-1} + \sigma^{-2} \mathbf{I})^{-1}$$

$$\mathbf{f} = \sigma^{-2} (\mathbf{y} - \mu_0).$$

In the metric, $\mathbf{K}$ is the covariance matrix of the Gaussian prior, and its inverse, $\mathbf{K}^{-1}$, is the prior precision or, equivalently, the prior Fisher information matrix for the information in an observation about the mean parameter vector. The term $\sigma^{-2} \mathbf{I}$ adds the likelihood Fisher information in the data, $\mathbf{y}$, with Gaussian noise σ^2 .

Adding those two Fisher components yields the total Fisher information in the posterior distribution. For the Gaussian case here, the Fisher matrix is the Hessian of the negative log posterior, $-\log \mathbf{q}(\mathbf{g}|\mathbf{y})$. Thus, $\mathbf{M}$, the inverse of the total Fisher matrix, is the inverse curvature of the Bayesian log posterior probability weights for alternative functions with respect to variations in those functions.

The force is proportional to the deviation between observed values, $\mathbf{y}$, and predicted values by the prior, $\mu_0(\mathbf{x})$. We scale that deviation by the information in the observed data, σ^{-2} , which is the inverse variance of the measurement noise.

11.2. Kalman Filters

One typically uses a Gaussian process model to infer a static functional relation between inputs and outputs. A Kalman filter applies a similar approach to a dynamically changing system [59].

The following description of the Kalman filter takes a bit of notation. However, the point is simple. Suppose there exists a dynamically changing vector of hidden states. We know the basis for the hidden stochastic dynamics, but cannot directly observe the system state.

Instead, we observe a correlated measure of the hidden values. From the observed measurements, we can repeatedly update the estimated mean vector of hidden values, in which those estimates have a Gaussian distribution of error.

The updates of the estimated mean have a simple metric–force expression. The metric is the covariance of the error distribution, which is the local inverse curvature in the space of estimated values. The force is the deviation between observed and predicted values

on the measurement scale, weighted by the information in an observation relative to the hidden values.

Suppose, for example, that the state of some system, $\mathbf{x}_t$, changes with time, t . One cannot directly observe $\mathbf{x}_t$ but can measure a correlate, $\mathbf{y}_t$. We have an explicit stochastic dynamical system

$$\begin{aligned}\mathbf{x}_t &= \mathbf{F}\mathbf{x}_{t-1} + \boldsymbol{\zeta}_t \\ \mathbf{y}_t &= \mathbf{H}\mathbf{x}_t + \boldsymbol{\eta}_t.\end{aligned}$$

Here, $\mathbf{F}$ defines the deterministic dynamics of $\mathbf{x}$, and the noise is $\boldsymbol{\zeta}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$. One can observe $\mathbf{y}$, which is $\mathbf{x}$ measured through the filter $\mathbf{H}$ and subject to $\mathcal{N}(\mathbf{0}, \mathbf{R})$ measurement noise.

How can we use our observation of $\mathbf{y}_t$ to update our prediction for the underlying $\mathbf{x}_t$ values?

A Kalman filter uses the data in each time step to update the Gaussian distribution $\mathcal{N}(\hat{\mathbf{x}}_t, \mathbf{P}_t)$, in which the mean vector $\hat{\mathbf{x}}_t$ is the current estimate for the underlying true values, $\mathbf{x}_t$, and $\mathbf{P}_t$ is the covariance matrix of the estimated values. A Bayesian update maps prior to posterior distribution, $(\hat{\mathbf{x}}_t^-, \mathbf{P}_t^-) \mapsto (\hat{\mathbf{x}}_t^+, \mathbf{P}_t^+)$, given the data, $\mathbf{y}_t$.

The updated prediction, $\Delta\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_t^+ - \hat{\mathbf{x}}_t^-$, follows our standard FMB product of metric and force

$$\Delta\hat{\mathbf{x}}_t = \mathbf{M}_t \mathbf{f}_t.$$

The metric, $\mathbf{M}_t = \mathbf{P}_t^-$, is the prior covariance, which is the inverse of the local curvature of the estimate error. The prior covariance describes a population of plausible state trajectories at each time step.

The recursive temporal updating of this covariance metric, $\mathbf{P}_t^- = \mathbf{F}\mathbf{P}_{t-1}^+ \mathbf{F}^\top + \mathbf{Q}$, is similar to the way in which Adam uses temporal extent to estimate spatial aspects of geometry. However, the Kalman filter implements this principle through a formal dynamic model given by $\mathbf{F}$, rather than by the purely empirical averaging of past gradients used by Adam. Thus, the different algorithms use temporal information in distinct ways to achieve the same type of spatial metric within the FMB law.

The force is

$$\mathbf{f}_t = \mathbf{H}^\top \mathbf{S}_t^{-1} \mathbf{v}_t.$$

Here, $\mathbf{v}_t = \mathbf{y}_t - \mathbf{H}\hat{\mathbf{x}}_t^-$, is the difference between the observed values, $\mathbf{y}_t$, and the predicted values based on the prior mean, $\hat{\mathbf{x}}_t^-$, scaled to the $\mathbf{y}$ coordinates by $\mathbf{H}$. Thus, $\mathbf{v}_t$ is proportional to the force for change expressed on the $\mathbf{y}$ scale. Multiplying by $\mathbf{H}^\top$ rescales the force to the $\mathbf{x}$ coordinates. Finally,

$$\mathbf{S}_t = \mathbf{H}\mathbf{P}_t^- \mathbf{H}^\top + \mathbf{R},$$

which is the covariance matrix of $\mathbf{v}_t$. Thus, the inverse is how much Fisher information there is in an observed $\mathbf{v}_t$ about the hidden $\mathbf{x}_t$.

Overall, the Kalman filter's use of temporal extent to estimate the spatial geometry of the covariance metric provides an interesting contrast with the Gaussian process's typical one-shot purely spatial estimate of geometry.

12. Bayesian Learning

Natural selection and learning are often interpreted in Bayesian terms. Bayesian-inspired methods also provide many important computational learning algorithms. This section summarizes some of those methods, placing them within our broad framework for learning. Once again, I emphasize the simple conceptual unity of seemingly different approaches in the study of learning.

12.1. Brief Review

The first Price equation term links the change in probability distribution, $\Delta \mathbf{q}$, to the change in mean values, $\Delta_f \bar{\theta} = \Delta \mathbf{q} \cdot \bar{\theta}$. If we interpret the change in each frequency value, $q'_i = q_i w_i$, as driven by relative performance, w_i , then we can think of $\Delta \mathbf{q}$ as the change in the probability distribution driven by the improvement caused by learning [24,60–63].

The updating of probability distributions by learning matches the standard Bayesian update process. In Equation (15), we equated relative likelihood with relative fitness, $L_i = w_i$. Thus, $q'_i = q_i L_i$, in which i associates with the θ_i parameter vector, q_i is the prior probability of the i th parameter vector, L_i is the relative likelihood of that parameter vector given some data, and q'_i is the posterior probability of the i th parameter vector. All of that fits exactly into our Price equation expressions for the gain in performance caused by natural selection or learning, leading to Equation (18), repeated here

$$\Delta_f \bar{L} = \Delta \mathbf{q} \cdot \mathbf{L} = \left\| \frac{\Delta \mathbf{q}}{\sqrt{\mathbf{q}}} \right\|^2,$$

which shows that the partial increase in likelihood caused by the direct force of learning is the same enhancement of performance measured by the discrete squared Fisher–Rao path length that we see in many learning models.

We can also write

$$\Delta_f \bar{\theta} = \mathbf{M} \mathbf{f}$$

for the change between the posterior and prior distributions, in which $\mathbf{M}$ is the covariance of θ over the prior, and $\mathbf{f}$ is the slope of the relative likelihood with respect to the parameters. Thus, Bayesian updating is a standard metric–force update for a population model.

The following subsections provide details about the nature of forces, how to partition those forces into components and constraints, and how to link classic physical notions of force, work, variational optimization, and free energy to the FMB framework.

12.2. Variational Bayes

Bayesian updating is, in theory, a simple method. From the prior subsection, $q'_i = q_i L_i$, in which L_i is the relative likelihood given in Equation (15), repeated here

$$w_i = L_i = \frac{\tilde{L}(\mathbf{D}|\theta_i)}{\sum_i q_i \tilde{L}(\mathbf{D}|\theta_i)}.$$

The practical problem arises when it is difficult to calculate the sum in the denominator to obtain the proper normalizing value, which we need to make the total probability of the posterior equal to one. That sum can be difficult to calculate when each likelihood is associated with a large number of parameters, or when we have many alternative parameter vectors to consider.

Suppose we define a performance function that improves as our currently estimated posterior moves toward the true posterior. Then the challenge matches a common learning problem, and we can apply standard learning algorithms.

The variational Bayes method uses this approach [64]. Let our candidate posterior probability distributions, $\hat{\mathbf{q}}(\phi)$, be confined to a particular distributional form that depends on the parameters, ϕ . Then the problem becomes the search for ϕ that minimizes the divergence of the assumed form for the posterior, $\hat{\mathbf{q}}(\phi)$, from the true posterior for θ given the data, $\mathbf{q}'(\theta|\mathbf{D})$.

We measure that difference between estimated and true posterior by the KL divergence, $\mathcal{D}(\hat{\mathbf{q}}(\phi)||\mathbf{q}')$, defined in Equation (13). This minimization problem is a variational

method because we are minimizing the divergence between the function $\hat{\mathbf{q}}(\phi)$ and the true posterior, $\mathbf{q}'$.

Variational Bayes methods [64] provide a common approach to search for $\hat{\mathbf{q}}(\phi)$. The details are simple but require a bit of notation to make the steps clear. In addition, we will use notation that can be plugged easily into the Price equation in the next section. Before starting, I list some notational shortcuts, in which $\mathbf{D}$ refers to data

$q_i = q(\theta_i)$	prior distribution
$\hat{q}_i = \hat{q}(\theta_i; \phi)$	estimated posterior distribution
$q'_i = q(\theta_i \mathbf{D})$	true posterior distribution
$\tilde{L}_i = q(\mathbf{D} \theta_i)$	nonnormalized likelihood, (Equation (15))
$q(\theta_i, \mathbf{D})$	joint distribution θ and data
$p(\mathbf{D})$	probability of the data
$E_{\hat{\mathbf{q}}}$	expectation with respect to $\hat{\mathbf{q}}(\phi)$.

We derive the variational Bayes method by starting with the basic rule of conditional probability, $q(\theta_i | \mathbf{D})p(\mathbf{D}) = q(\theta_i, \mathbf{D})$, which we write in our shorthand notation as

$$p(\mathbf{D}) = \frac{q(\theta_i, \mathbf{D})}{q'_i} = \frac{q(\theta_i, \mathbf{D})}{\hat{q}_i} \frac{\hat{q}_i}{q'_i}.$$

Take the **log** of both sides and then the expectation over $\hat{\mathbf{q}}$, noting that $p(\mathbf{D})$ on the left is a constant given the data, and so the expectation drops out on that side

$$\begin{aligned} \log p(\mathbf{D}) &= E_{\hat{\mathbf{q}}}[\log q(\theta, \mathbf{D})] - E_{\hat{\mathbf{q}}} \log \hat{\mathbf{q}} + \mathcal{D}(\hat{\mathbf{q}} || \mathbf{q}') \\ &= \mathcal{L}(\phi) + \mathcal{D}(\hat{\mathbf{q}} || \mathbf{q}'), \end{aligned} \quad (33)$$

in which

$$\mathcal{L}(\phi) = E_{\hat{\mathbf{q}}}[\log q(\theta, \mathbf{D})] - E_{\hat{\mathbf{q}}} \log \hat{\mathbf{q}} \quad (34)$$

is called the evidence lower bound (ELBO).

The goal of variational Bayes is to minimize the divergence between the estimated posterior, $\hat{\mathbf{q}}$, and the true posterior, $\mathbf{q}'$, measured by the KL divergence term in Equation (33). Because the value of $\log p(\mathbf{D})$ is constant, and the KL divergence term is nonnegative, maximizing the ELBO, $\mathcal{L}(\phi)$, minimizes the divergence. Thus, variational Bayes targets the maximization of $\mathcal{L}(\phi)$.

We can rewrite the ELBO in a more convenient form by starting with the rule for conditional probability and the definition of likelihood

$$q(\theta_i, \mathbf{D}) = q(\mathbf{D} | \theta_i)q(\theta_i) = \tilde{L}(\mathbf{D} | \theta_i)q(\theta_i),$$

which on the right side is the product of the nonnormalized likelihood of the data given the parameters, $\tilde{L}$, and the prior probability of the parameters, $q(\theta_i)$. Using this expansion in the expression for the ELBO in Equation (34) yields an alternative form

$$\mathcal{L}(\phi) = E_{\hat{\mathbf{q}}}(\log \tilde{L}(\mathbf{D} | \theta)) - \mathcal{D}(\hat{\mathbf{q}} || \mathbf{q}), \quad (35)$$

the expected log-likelihood taken over the estimated posterior, $\hat{\mathbf{q}}$, minus the KL divergence of the estimated posterior from the prior.

Intuitively, maximizing ELBO balances the gain from concentrating the updated probability on regions with the greatest log-likelihood versus the cost of departing from

the prior. In other words, the prior sets the default from which one changes only by the weight of new evidence.

12.3. Variational Bayes from the Price Equation

The Price equation elegantly describes how much the ELBO increases as the estimated posterior, $\hat{\mathbf{q}}$, departs from the prior, $\mathbf{q}$.

The Price equation (Equation (3)) can be rewritten with $\mathbf{q}' \mapsto \hat{\mathbf{q}}$ and $\boldsymbol{\theta} \mapsto \mathbf{z}$ as

$$\Delta \bar{z} = \Delta \mathbf{q} \cdot \mathbf{z} + \hat{\mathbf{q}} \cdot \Delta \mathbf{z}.$$

Here, $\Delta q_i = \hat{q}_i - q_i$ is the difference between the estimated posterior and the prior.

We are free to choose the trait values $\mathbf{z}$. For $\Delta \theta_i = \hat{z}_i - z_i$, let

$$z_i = \log \tilde{L}_i - \log \frac{q_i}{\tilde{q}_i} = \log \tilde{L}_i$$

$$\hat{z}_i = \log \tilde{L}_i - \log \frac{\hat{q}_i}{q_i},$$

in which $\tilde{L}_i = \tilde{L}(\mathbf{D}|\theta_i)$ is the likelihood for the i th parameter combination given the data. With these definitions, $\bar{z} = \mathcal{L}(\boldsymbol{\phi})$, the ELBO. Thus, the total change in the ELBO is

$$\begin{aligned} \Delta \bar{z} &= \sum \hat{q}_i \hat{z}_i - \sum q_i z_i \\ &= (\mathbb{E}_{\hat{\mathbf{q}}}(\log \tilde{L}) - \mathcal{D}(\hat{\mathbf{q}}||\mathbf{q})) - (\mathbb{E}_{\mathbf{q}}(\log \tilde{L}) - \mathcal{D}(\mathbf{q}||\mathbf{q})) \\ &= \mathbb{E}_{\Delta \mathbf{q}}(\log \tilde{L}) - \mathcal{D}(\hat{\mathbf{q}}||\mathbf{q}), \end{aligned}$$

which is $\Delta \mathcal{L}(\boldsymbol{\phi})$, the difference between the ELBO at $\hat{\mathbf{q}}$ and the baseline ELBO value when the estimated posterior is equal to the prior, $\hat{\mathbf{q}} = \mathbf{q}$, noting that $\mathcal{D}(\mathbf{q}||\mathbf{q}) = 0$.

To analyze the two Price terms separately, we need

$$\sum \hat{q}_i z_i = \mathbb{E}_{\hat{\mathbf{q}}}(\log \tilde{L}).$$

We can now write the first Price term as

$$\Delta \mathbf{q} \cdot \mathbf{z} = \Delta \mathbf{q} \cdot \log \tilde{L} = \mathbb{E}_{\Delta \mathbf{q}}(\log \tilde{L}),$$

which is consistent with our earlier interpretation in Equation (20) of the likelihood as the direct force driving frequency changes. The second Price term is

$$\hat{\mathbf{q}} \cdot \Delta \mathbf{z} = -\mathcal{D}(\hat{\mathbf{q}}||\mathbf{q}).$$

This term describes how the direct gains made by moving frequencies toward regions of high likelihood alter the frequency context, imposing a cost on performance in proportion to how far the frequencies have moved from the prior.

The idea is that the total change in the ELBO balances the direct gain in moving closer to the likelihood of the data against the changed-context cost of moving away from prior information. We can write that as

$$\Delta \mathcal{L}(\boldsymbol{\phi}) = \Delta \mathbf{q} \cdot \log \tilde{L} - \mathcal{D}(\hat{\mathbf{q}}||\mathbf{q}), \quad (36)$$

in which the first term describes the direct force of the data via the log-likelihood, $\log \tilde{L}$, and the second term is a weighted average of the opposing inertial force imposed by the prior, $\log \hat{\mathbf{q}}/\mathbf{q}$.

For an infinitesimal change in the posterior frequencies, $\delta\hat{\mathbf{q}}$, using variational notation, δ , for small differences that are consistent with the conservation of total probability, the first variational derivative of the ELBO in Equation (35) is

$$\delta\mathcal{L}(\phi) = \left(\log \tilde{\mathbf{L}} - \log \frac{\hat{\mathbf{q}}}{\mathbf{q}} \right) \cdot \delta\hat{\mathbf{q}}, \quad (37)$$

which clearly separates the forces and the displacement. Integrating this variational derivative over the specific frequency changes from $\mathbf{q}$ to $\hat{\mathbf{q}}$ yields the total change in Equation (36).

12.4. Statics, Constraints, and d'Alembert's Principle

In variational Bayes, we choose a family of probability distributions for the estimated posterior, $\hat{\mathbf{q}}(\phi)$, that vary according to ϕ . For example, we might choose a multivariate Gaussian with uncorrelated dimensions and a covariance matrix with diagonal elements given by ϕ .

More complex distributions may improve the potential to increase the ELBO. But those more complex distributions may also make the computational search process for improving the ELBO more difficult.

Whatever the chosen distribution, any learning method that improves the parameter choice for ϕ with respect to the performance measure $\mathcal{L}(\phi)$ can be used. The point in this subsection is not to consider the specific dynamics of learning but rather to place the forces that directly influence improvement by learning into a broader context of direct, inertial, and constraining forces.

In general, setting a distributional family for $\hat{\mathbf{q}}(\phi)$ imposes a constraint on the learning process. The Price equation provides a simple way to connect that particular force of constraint to a broader understanding of forces and dynamics.

Analyzing a static system often provides a simple way to understand various forces. A static system occurs when the forces are in balance, causing the overall system to remain unchanged. Balanced forces provide a strong clue about how the individual components must be changing to conserve the overall system.

We obtain a conserved system in the Price equation by defining

$$\begin{aligned} z_i &= \log \tilde{L}_i \\ \hat{z}_i &= \mathbf{q} \cdot \log \tilde{\mathbf{L}}, \end{aligned}$$

so that $-\Delta z_i$ is the deviation between the force acting on each dimension of $\mathbf{q}$ and the average force acting over all dimensions. Then, the two Price equation terms are

$$\begin{aligned} \Delta \mathbf{q} \cdot \mathbf{z} &= \Delta \mathbf{q} \cdot \log \tilde{\mathbf{L}} \\ \hat{\mathbf{q}} \cdot \Delta \mathbf{z} &= -\Delta \mathbf{q} \cdot \log \tilde{\mathbf{L}} \end{aligned}$$

Recalling that $\tilde{\mathbf{L}} = \mathbf{q}'/\mathbf{q}$, we can expand the second term to include

$$\log \tilde{\mathbf{L}} = \log \frac{\mathbf{q}'}{\mathbf{q}} = \log \frac{\hat{\mathbf{q}}}{\mathbf{q}} + \log \frac{\mathbf{q}'}{\hat{\mathbf{q}}}.$$

If we consider small virtual displacements of the posterior, $\delta\hat{\mathbf{q}}$, a variational form of the Price equation for changes in $\bar{z}$ becomes

$$\delta\bar{z} = \left(\log \tilde{\mathbf{L}} - \log \frac{\hat{\mathbf{q}}}{\mathbf{q}} - \log \frac{\mathbf{q}'}{\hat{\mathbf{q}}} \right) \cdot \delta\hat{\mathbf{q}} = 0, \quad (38)$$

recalling that all allowable deviations $\delta\hat{\mathbf{q}}$ satisfy the conservation of total probability, so that the sum of all deviations is zero. This expression matches d'Alembert's principle from classical mechanics, illustrated by Equation (19).

d'Alembert's principle emphasizes how various forces in a conserved, static system balance, whereas most analyses focus on the dynamics of the moving parts. In other words, d'Alembert changes emphasis to the forces rather than the moving bodies [65]. That perspective helps when the goal is to understand how a system works rather than in the calculation of the actual paths of motion.

In the d'Alembert context, we can parse the balancing forces in Equation (38). Let $\delta\hat{\mathbf{q}}$ be a virtual displacement of the posterior frequencies consistent with all constraints of the system. Then $\delta\hat{\mathbf{q}} \cdot \log \tilde{\mathbf{L}}$ is the virtual work performed by the overall potential force, $\log \tilde{\mathbf{L}}$, acting to change frequencies. That overall potential force is opposed by the reactive inertial force of the prior, $\log \hat{\mathbf{q}}/\mathbf{q}$, within the constrained space of allowable posterior distributions, $\hat{\mathbf{q}}$. The overall potential is also opposed by the residual potential, $\log \mathbf{q}'/\hat{\mathbf{q}}$, a reactive inertial force for the imaginary movement from the constrained posterior, $\hat{\mathbf{q}}$, to the true posterior, $\mathbf{q}'$.

Noting from Equation (37) how the ELBO changes with an infinitesimal update, we can write our static d'Alembert expression in Equation (38) as

$$\delta\mathcal{L}(\phi) - \delta\hat{\mathbf{q}} \cdot \log \frac{\mathbf{q}'}{\hat{\mathbf{q}}} = 0.$$

The virtual work gained by improving the evidence lower bound is balanced by the reactive force from the remaining potential. If $\mathbf{q}^*$ is the optimum within the constrained space, we can split the remaining potential into

$$\log \frac{\mathbf{q}'}{\hat{\mathbf{q}}} = \log \frac{\mathbf{q}^*}{\hat{\mathbf{q}}} + \log \frac{\mathbf{q}'}{\mathbf{q}^*},$$

which is the remaining potential from the current posterior, $\hat{\mathbf{q}}$, to the constrained optimum, $\mathbf{q}^*$, plus the potential from the true posterior, $\mathbf{q}'$, to the constrained optimum.

Overall, the d'Alembert expression for the separation of balancing forces follows naturally from the Price equation description of a conserved system. We gain a clear sense of how the various forces influence the dynamics of learning.

12.5. Friston's Free Energy Models

Friston built a unified brain theory on the principles of variational Bayes analysis. In essence, Friston sought a theory in which the minimization of a single quantity could unify three problems that had previously been treated as separate topics [37,66,67].

First, homeostatic maintenance of biological function requires opposing the inexorable entropic decay of order. Friston suggested that such maintenance arises from minimizing the long-term average surprise from environmental sensory input. In a Bayesian context, surprise is $-\log p(\mathbf{D})$, the negative log probability of the data.

Second, Bayesian analogies have been used for how brains perceive and learn. But Bayesian calculations are often complex and difficult. Perhaps brains use the easier variational Bayes algorithm to approximate Bayesian inference, which maximizes the evidence lower bound (ELBO) by descending the free energy gradient.

Third, theories of behavioral action often derive from optimal control or reinforcement learning. Those theories optimize some measure of reward or value. Friston showed that increased value is typically associated with reduced surprise, providing a direct link to Bayesian methods and the equivalent free energy expressions.

This subsection shows how Friston's methods fit into our broader framework for algorithmic learning. Before starting, it is helpful to emphasize the underlying simplicity of the program.

In essence, all of Friston's analyses reduce to a simple prescription. Keep tuning your model so that it assigns higher probability to the data that you actually observe, while not making the model unnecessarily complicated.

The tuned model is the estimated posterior, $\hat{\mathbf{q}}$. The goal is to drive $\hat{\mathbf{q}}$ close to the true posterior, $\mathbf{q}'$, which means reducing the divergence, $\mathcal{D}(\hat{\mathbf{q}}||\mathbf{q}')$. Friston's wording and technical steps sometimes make it difficult to keep that simplicity in clear focus.

Friston's approach defines free energy, F , in a learning context by starting with the surprise of the data, $-\log p(\mathbf{D})$. Using our previous expansion of that expression in Equation (33) and rearranging the terms yields

$$F(\phi) = -\text{ELBO} + \log p(\mathbf{D}) = \mathcal{D}(\hat{\mathbf{q}}||\mathbf{q}'),$$

in which $-\text{ELBO} = -\mathcal{L}(\phi)$. Because $\log p(\mathbf{D})$ is a constant for a given data input, choosing the parameters, ϕ , to minimize the free energy is equivalent to maximizing the ELBO, as in standard variational Bayes methods.

It is easier to see that F is a plausible analogy for free energy by using the alternative form for $\mathcal{L}(\phi)$ in Equation (35), yielding

$$F(\hat{\mathbf{q}}) = \mathcal{D}(\hat{\mathbf{q}}||\mathbf{q}) - \mathbb{E}_{\hat{\mathbf{q}}}[\log \tilde{L}(\mathbf{D}|\theta)] + \log p(\mathbf{D})$$

in which I have written $F(\hat{\mathbf{q}})$ in this case to emphasize that we can equivalently think of optimizing the estimated posterior, $\hat{\mathbf{q}}(\theta; \phi)$, or optimizing the parameters, ϕ , that determine the estimated posterior.

The difference in free energy, $\Delta F = F(\hat{\mathbf{q}}) - F(\mathbf{q})$ for a change $\Delta \mathbf{q} = \hat{\mathbf{q}} - \mathbf{q}$, can be written from Equation (36) as

$$\Delta F = \mathcal{D}(\hat{\mathbf{q}}||\mathbf{q}) - \Delta \mathbf{q} \cdot \log \tilde{L} = -\Delta \mathcal{L}(\phi), \quad (39)$$

which, in Friston's language, is read as the tradeoff in free energy between the gain from increasing the accuracy, $\Delta \mathbf{q} \cdot \log \tilde{L}$, and the loss from increasing the complexity, $\mathcal{D}(\hat{\mathbf{q}}||\mathbf{q})$. Here, complexity means pulling further away from the maximally disordered state defined by the prior.

Classically, free energy measures the direct force that increases order, opposed by the intrinsic pull toward disorder. Similarly, our Price equation analysis also shows learning as the balanced improvement by direct forces against the opposing decay of performance caused by inertial forces. In my opinion, the Price equation analysis derives the same balance of forces more transparently than the free energy analogy.

With this technical background, I describe Friston's two main conclusions.

First, the brain is designed to behave as if it minimizes surprise. In practice, it selects the variational Bayes posterior that maximizes the ELBO, thereby reducing the free energy, F . With each update, the new posterior becomes the next prior, $\hat{\mathbf{q}} \mapsto \mathbf{q}$, and the corresponding surprise associated with the updated estimate for the probability of the data, $-\log \hat{p}(\mathbf{D})$, should be reduced.

Second, agents choose future actions that minimize expected free energy. To do that, they adopt an active inference policy that balances the tradeoff between exploitation and exploration [68]. For exploitation, choose actions with likely outcomes that match current preferences. For exploration, choose actions that provide information to improve preferences, increasing the value of future outcomes.

Friston embedded an action–inference process within his free energy framework. That approach provides a single quantity to minimize, in which minimization balances the gains from exploitation and exploration. By assuming a common metric within a Bayesian framing, one can develop testable hypotheses about how closely actual behavioral sequences match the predicted learning process.

Actions depend on the current behavioral policy, π . The free energy depends on policy

$$F(\pi) = \text{risk} - \text{information gain}. \quad (40)$$

The exploitation component measures risk, the mismatch between preferred outcomes, and actual outcomes. Friston measures risk by surprise, the mismatch between expected outcomes for a behavioral policy, π , and the actual outcomes. In probabilistic models, we write $-\log q(o|C)$ for the surprise of an outcome, o , given an expected outcome, C . We derive the expected surprise by averaging over the frequency of the actual outcomes given the policy.

The exploration component measures the gain in information by linking the world's outcome state, s , to policy, π . The agent's prior belief is the probability distribution $q(s|\pi)$. After a round of behavior with actual outcomes, o , the agent has an updated posterior belief, $q(s|o, \pi)$. The KL divergence between the posterior and prior measures the gain in information. The negative value of that divergence is the free energy decrease for updating beliefs.

12.6. Matching Friston to Fisher and d'Alembert

As so often happens, we can link Friston's seemingly special results back to our common canonical forms for the Price equation, Fisher information, and d'Alembert's principle. Friston's choice of using logarithms and small changes means that we are considering small path updates, as in Equation (19). That equation shows that the opposing direct and inertial forces typically lead to virtual work components that are equal and opposite Fisher–Rao path lengths.

In Friston's case, with $\hat{q} = q + \delta q$ for small change δq , and thus $\log \hat{q}/q \rightarrow \delta q/q$, the KL divergence becomes

$$\mathcal{D}(\hat{q}||q) \rightarrow \delta q \cdot \log \frac{\hat{q}}{q} = \left\| \frac{\delta q}{\sqrt{q}} \right\|^2.$$

Thus, noting that $\log \tilde{L} = \log q'/q$, Equation (39) becomes

$$\begin{aligned} -\delta F &= \left(\log \frac{q'}{q} - \log \frac{\hat{q}}{q} \right) \cdot \delta q \\ &= \delta q \cdot \log \frac{q'}{\hat{q}}, \end{aligned}$$

in which all displacements, δq , are consistent with the constraint that the posterior remain within the space of allowable alternative probability distributions.

If we split the direct force, $\log q'/q$, into the component between the current prior, q , and the constrained optimum, q^* , plus the component between the constrained optimum and the true posterior, q' , then as the updated prior approaches the local optimum, $q \rightarrow q^*$, we have

$$-\delta F = \left(\log \frac{q'}{q^*} + \log \frac{q^*}{q} - \log \frac{\hat{q}}{q} \right) \cdot \delta q,$$

in which the constraining force, $\log \mathbf{q}'/\mathbf{q}^*$ is orthogonal with respect to allowable displacements $\delta \hat{\mathbf{q}}$ and drops out, yielding

$$-\delta F = \delta \hat{\mathbf{q}} \cdot \log \frac{\mathbf{q}^*}{\hat{\mathbf{q}}} = \left(\log \frac{\mathbf{q}^*}{\mathbf{q}} - \log \frac{\hat{\mathbf{q}}}{\mathbf{q}} \right) \cdot \delta \hat{\mathbf{q}}.$$

Also, for a prior near the constrained optimum and a small displacement that moves to the optimum, $\mathbf{q}^* = \mathbf{q} + \delta \hat{\mathbf{q}}$, we have $\log \mathbf{q}^*/\mathbf{q} \rightarrow \log \hat{\mathbf{q}}/\mathbf{q}$, yielding $\delta F \rightarrow 0$, with a local d'Alembert balance between the direct and inertial forces at the constrained optimum

$$-\delta F = (\log \tilde{\mathbf{L}} - \mathcal{D}(\hat{\mathbf{q}}||\mathbf{q})) \cdot \delta \hat{\mathbf{q}} \rightarrow 0.$$

Overall, Friston's free energy models of learning fit naturally within our Price equation framework.

13. Hierarchical Learning

Biological populations have a natural learning hierarchy. Each individual inherits a set of parameters from its genes. Those parameters guide a learning process over the individual's lifetime.

Within an individual, learning changes an internal nongenetic parameter vector. That learning influences the individual's success in transmitting its genes to the future of the population. In this case, what an individual learns does not get transmitted. The global process influences only the genes that seed the initial state of each individual. Call this nonheritable learning.

In biology, Baldwin [69] was perhaps the first to recognize that such hierarchical separation can greatly accelerate the overall rate at which a system learns. Subsequent studies in biology have extended the idea that nonheritable developmental adjustments or learning by individuals explain many aspects of how populations have actually evolved over the history of life [70].

Alternatively, we may think of each individual or lower-level group as transmitting the parameter vector that it learns over time. Then, the global process aggregates the learned parameter vectors from each lower-level unit, weighting each lower-level contribution by the relative performance associated with its transmitted vector. In biology, many studies of group selection emphasize the enhanced evolutionary potential that arises from hierarchical population structure [71–74]. Call this heritable learning.

In both cases, the success of each lower-level unit in learning determines the fitness or performance of the unit. The distinction between the nonheritable and heritable cases concerns whether the transmitted parameter vector is the same as the initial seeding of the unit or is the updated vector produced by improvement through the unit's learning.

The two cases lead to different types of searches. In nonheritable learning, only the initial seeding vector is transmitted, and the parameters evolve to encode a better learning process within lower-level units. In heritable learning, the final improved vector of each lower-level unit is transmitted, and the parameters evolve to encode a better reaction in direct response to the particular challenge.

Many computational algorithms exploit the potential benefits of hierarchical learning [75]. This section uses the Price equation and the FMB law to show how hierarchical learning fits into our simple and general framework for algorithmic learning and natural selection.

13.1. Nonheritable Learning in Lower-Level Units

In this case, we keep our standard FMB law. However, to account for lower-level learning within individuals, we evaluate the fitness or performance function, $w \equiv U$, at the point

$$\tilde{\theta} = \theta + \Delta_{\tau}\theta,$$

in which θ is the initial parameter vector passed to the lower-level unit, and $\Delta_{\tau}\theta$ is the change in the parameter vector by learning within the lower-level unit over the time period τ .

The global update, $\bar{\theta} = \mathbf{M}\mathbf{f} + \mathbf{b}$, uses $U(\tilde{\theta})$ to calculate the force vector, $\mathbf{f}$. We can also adjust the bias vector in response to lower-level learning.

In computational applications, the parameter change caused by the internal learning process, Δ_{τ} , typically arises from a sequence of discrete learning updates based on a local algorithm. Usually, we can express the local algorithm in FMB form. For example, in the i th individual

$$\Delta_{\tau}\theta_i = \sum_{r=1}^{\tau} \mathbf{M}_i \mathbf{f}_i + \mathbf{b}_i,$$

in which the metric, force, and bias terms may change with context in each time step. Here, if we consider a population or Bayesian process within individuals, we would write $\Delta_{\tau}\bar{\theta}_i$ for the individual change.

As before, we can switch between a population interpretation in which FMB statistics come from the population attributes and a single-value parameter update interpretation in which we set FMB statistics by other methods. A similar equivalence continues at the global level.

This setup provides a good description of many biological scenarios. In those biological cases, genetics fixes the initial values of individuals and the heritably transmitted values. Internal learning affects performance and reproductive fitness but does not influence the transmitted parameters.

13.2. Nonheritable Learning: Algorithmic Examples

Hinton & Nowlan [76] illustrated how individual learning transformed an unsolvable search into an easily solved one. Simplifying the original Hinton & Nowlan model, assume that each inherited genotype is a bit string of length 20. A particular target string is set as success. All other strings have the same low value of performance.

A small population is initialized with random strings. The chance that any single string matches the target is roughly 10^{-6} . No simple combination of mutation, recombination, and reproduction by strings is likely ever to match the target.

Each string was then allowed a learning period. From the initial string value, the bits were mutated randomly over several rounds. If one of the mutated strings matched the target, then that individual had higher fitness and contributed more copies of itself to the next generation. The transmitted copy is the initial seed, not the internally mutated version during learning.

In this scenario, fitness measures the probability that an initial string can be mutated into a target match. That probability increases as a string's divergence from the target declines. Thus, internal learning turned a performance function with all weight on a single point into a graded performance function that increases steadily as the seeded string approaches the target. Learning algorithms are highly effective at following an improving gradient. Thus, this example of a Baldwin learning process shows how simply and powerfully internal learning can improve search.

Newer methods extend the Baldwin approach. For example, Fernando et al. [77] evolved a population of neural networks to solve a particular challenge. Each network

inherits a parameter vector that sets the initial conditions, the way performance is calculated, and the learning rate. Each network then learns through a fixed number of standard stochastic gradient descent rounds of updates.

The final state of the network determines its performance. However, only the initially inherited vector is used to seed the next generation, with the contribution of each vector weighted by its associated performance after local learning. In tests, the nonheritable learning within individuals improved performance.

These population models are split into a nonheritable inner loop of learning within individuals and a heritable outer loop between individuals. The same split can be achieved with single-value parameter updates rather than populations. For example, model-agnostic meta-learning improves a single vector of initial weights for a learning process [78].

Those initial weights seed a task-specific inner loop with multiple steps of stochastic gradient descent. After evaluation of the final state, the updated parameter weights from this inner loop are discarded. Only the gradient of the inner loop's final performance with respect to the initial parameters is fed into the update process for the outer loop.

This update process for a single parameter vector provides the same Baldwin-type separation of the overall learning process as occurred in the population-based methods. Population methods are advantageous when a component of the learning process lacks an easily calculated performance gradient, and one has to use a statistical method to associate particular changes in parameter vectors with changes in performance.

The standard one-level FMB law holds in all of these cases. The force $\mathbf{f}$ in the outer loop is the regression or gradient of fitness on the seed parameters. The inner-loop learning dynamics influence the calculation of fitness for a given seed parameter vector. Otherwise, the parameter updates from the inner loop are discarded.

The seed parameters of the outer loop often include hyperparameters, which are those parameters that control the inner-loop learning process rather than encode the parameters of the particular learned solution. Many other methods optimize hyperparameters and other attributes that control the architecture of the learning process in the outer loop and discard the parameter tuning in each inner-loop round [79–84].

13.3. Heritable Learning: Recursive Price Equation

The prior methods used an inner loop to evaluate fitness, but did not retain updated parameters from the inner loop. Other methods retain the parameter changes from the inner loop, forming a fully realized two-level learning process.

To analyze multilevel learning, we recursively expand the Price equation to describe a population hierarchy [85–87]. This subsection shows the steps.

The expanded Price equation reveals the sufficient statistics for population change as a hierarchical FMB law. The following subsection links the hierarchical FMB sufficient statistics to learning algorithms that update a single parameter vector rather than a population of parameter vectors.

Our initial derivation of the Price equation in Equation (3) set $\bar{w} = 1$ and so dropped that term. For recursive expansion, it is helpful to keep the $\bar{w}$ term, restating the Price equation for the change in the mean parameter vector as

$$\bar{w}\Delta\bar{\theta} = \text{Cov}(w, \theta) + E(w\Delta\theta).$$

This equation describes the change within a single population. Suppose we split the population into distinct groups indexed by g , and individuals within groups indexed by $j|g$.

The total change, which remains the same, can be partitioned into changes between groups and changes within groups. The first step is to write the same one-level Price equation expression with extra notation to emphasize our top-level hierarchy of groups

$$\bar{w}\Delta\bar{\theta} = \text{Cov}(\bar{w}_g, \bar{\theta}_g) + \text{E}(\bar{w}_g\Delta\bar{\theta}_g), \quad (41)$$

in which the covariance and expectation are taken over g . Previously, we had a population of parameter vectors, θ , each parameter vector associated with a fitness value, w . Now we have the same population of parameter vectors, but we call each vector the mean value of a group, $\bar{\theta}_g$. Again, each parameter vector maps to a fitness value, which we now describe as the group mean fitness, $\bar{w}_g$.

We expand recursively by noting that the expectation in the last term on the right includes $\bar{w}_g\Delta\bar{\theta}_g$, which is the same form as the left side of the equation but for the groups, g . Thus, for each group g , we can use the whole equation to expand recursively by writing

$$\bar{w}_g\Delta\bar{\theta}_g = \text{Cov}(w_{j|g}, \theta_{j|g}) + \text{E}(w_{j|g}\Delta\theta_{j|g}), \quad (42)$$

in which the covariance and expectation are taken over j for fixed g . Here, we drop the overbars on the right side because, at the lowest level of our analysis, $j|g$, we can maintain ambiguity about the nature of the values $w_{j|g}$ and $\theta_{j|g}$ with regard to whether or not they are averages over some lower level.

Substituting Equation (42) into Equation (41) yields a two-level recursive expansion of the Price equation. We can expand into any multilevel hierarchy, for example, groups, subgroups, individuals, parts with individuals, and so on. Once again, as long as we maintain consistent notation, the Price equation is just a tautologically true notational expansion.

13.4. Hierarchical FMB Law

To develop the hierarchical FMB law, define

$$\begin{aligned} \mathbf{M}_B &= \text{Cov}(\bar{\theta}_g, \bar{\theta}_g) \\ \mathbf{f}_B &= \text{Reg}(\bar{w}_g, \bar{\theta}_g) \\ \mathbf{M}_g &= \text{Cov}(\theta_{j|g}, \theta_{j|g}) \\ \mathbf{f}_g &= \text{Reg}(w_{j|g}, \theta_{j|g}) \\ \mathbf{b}_g &= \text{E}(w_{j|g}\Delta\theta_{j|g}), \end{aligned} \quad (43)$$

in which the notation $\text{Reg}(w, \theta)$ means the vector of partial regression coefficients of w with respect to θ . Here, subscript B denotes between all groups, and subscript g denotes within group g . In this notation, function arguments with g subscripts denote average over g , and arguments with $j|g$ denote averaging over j .

A potential bias between groups is implicit in the within-group biases, $\mathbf{b}_g$, expressed by a constant component of the bias for each $j|g$, such that $\mathbf{b}_g = \mathbf{b}_B + \tilde{\mathbf{b}}_g$. Thus, we can write the total bias in the population as

$$\mathbf{b} = \text{E}(\mathbf{b}_g) = \mathbf{b}_B + \text{E}(\tilde{\mathbf{b}}_g).$$

The hierarchical FMB law follows the same procedure used in Equations (6) and (7), extended here for hierarchical expansion

$$\Delta\bar{\theta} = \mathbf{M}_B\mathbf{f}_B + \mathbf{b}_B + \text{E}(\mathbf{M}_g\mathbf{f}_g + \tilde{\mathbf{b}}_g). \quad (44)$$

For simplicity, I drop the noise terms in this section. We can interpret this expression as a partition of our standard FMB law by noting that

$$\begin{aligned}\mathbf{M} &= \mathbf{M}_B + \mathbf{E}(\mathbf{M}_g) \\ \mathbf{f} &= \mathbf{M}^{-1}(\mathbf{M}_B \mathbf{f}_B + \mathbf{E}(\mathbf{M}_g \mathbf{f}_g)),\end{aligned}$$

in which the first line is the total covariance, and the second line follows directly from equating $\mathbf{M} \mathbf{f}$ with the parenthetical quantity on the right side of the second line, which is the total change by the product of metric and force taken over all levels. Thus, the hierarchical components add to the total FMB expression

$$\Delta \bar{\theta} = \mathbf{M} \mathbf{f} + \mathbf{b}.$$

In some cases, it may be useful to separate timescales explicitly between the hierarchical levels. For example, if we wish to track K updates within groups for each between-group update, then we can rewrite Equation (42) as

$$\bar{w}_g \Delta \bar{\theta}_g = \mathbf{M}_g \mathbf{f}_g + \mathbf{b}_g = \sum_{k=1}^K \mathbf{M}_g^{(k)} \mathbf{f}_g^{(k)} + \mathbf{b}_g^{(k)},$$

in which each term of the sum is the k th within-group learning update. The expressions in Equation (43) subsume this expansion by defining terms with respect to initial values and final values over the course of a within-group process. However, in practice, additional factors such as noise may enter in each explicit update, making the detailed summation of steps a better guide for matching the recursive Price equation's broad conceptual framing to actual implementations.

13.5. Multilevel Selection: Algorithmic Examples

Group selection has been widely discussed in biology [71–74], often using the Price equation's recursive expansion [85–87]. Many algorithmic learning methods have exploited the potential benefits of splitting populations into a multilevel hierarchy. This subsection mentions a few examples.

Commonly used evolutionary algorithms in machine learning can be extended by dividing populations into groups. These methods explicitly base their approach on the biological analogy of group selection or multilevel selection [88,89].

Other population methods split learning into an outer loop and a heritable inner loop [90]. The multistep inner loops may run within single agents of the population. Occasionally, the current state of one or more of the inner-loop agents is copied to seed some of the agents for a new population. For example, a currently strong agent may be cloned to replace a weaker agent. This approach combines the broad exploratory benefits of populations with the efficient improvement benefits of gradient-based approaches running within individual agents [91,92].

13.6. Single-Vector Updates

The Price equation's population analysis reveals the sufficient statistics for the updates to the mean parameter vector. In practice, instead of calculating those changing statistics for a particular population, we can choose those statistics by other assumptions or calculations. Our choice influences the learning trajectory's rate of gain, cost, and other attributes, allowing us to design learning algorithms to meet different objectives.

Substituting chosen or alternatively calculated values for the sufficient statistics of populations leads to updates of a single location parameter vector rather than a mean

parameter vector. In the earlier subsection *Metrics, sufficiency, and single-value updates*, I discussed how this transformation from population methods to single-value methods unifies natural selection, Bayesian methods, the variety of evolutionary population methods, and the common single-value algorithms.

In this section, we obtain the single-value form of the multilevel Price expansion by dropping the overbars for means, interpreting the previous mean vector as a single-value location descriptor, and choosing how we wish to calculate the sufficient statistics of the FMB expressions.

In this way, we can link the multilevel Price equation's population description to single-vector methods that combine outer and inner learning loops. Often, a hierarchical method of this sort will repeatedly run a fast inner loop by one learning algorithm, then occasionally pass the final result from the inner loop to a slow outer loop that uses a different algorithm. The outer loop then reseeds another round of the inner loop, achieving a separation of time scales. The following paragraphs list three examples.

First, the look-ahead optimizer uses any common algorithm for its fast inner loop, such as stochastic gradient descent or Adam. Multiple inner-loop steps explore the local performance surface. The updated parameters are then passed back to the slow outer loop, which adjusts the previously stored parameters in the direction of the inner-loop update. The next round of the fast inner loop begins with newly adjusted parameters, repeating the cycle [93].

In FMB language, the inner-loop steps are the within-group update. The subsequent blending of the inner-loop result with the prior outer-loop parameters plays the role of between-group selection and transmission, reseeding the next round with an improved state.

Second, we can interpret stochastic weight averaging within our hierarchical framework. The method first trains a model with stochastic gradient descent, forming the outer loop. Then it runs a further sequence of training steps in its inner loop. The outer loop is then updated to a final value by averaging over samples of the inner loop parameters [94].

Roughly speaking, as the trajectory converges near a local optimum, the stochasticity of the inner-loop updates tends to sample more in flatter regions of the performance surface near the optimum rather than in narrow and sharp parts of the performance surface. The flatter regions are associated with parameters that have less sensitivity in their performance and better generalization in their response to previously unseen inputs.

Typically, this algorithm uses one outer loop followed by one inner loop, and so is just a sequence of alternative training methods rather than a hierarchy. However, in challenging cases, the method could be extended to alternate hierarchically between outer-loop updates and tests of performance, followed by further inner-loop sampling and averaging as needed to improve performance.

Third, deep Q networks optimize behavioral sequences. To value an action within the behavioral sequence of a focal network, the algorithm combines the observed reward for that action plus a predicted reward for future actions. To calculate the predicted future reward, the algorithm uses as a target the behavioral sequence of another network with a fixed parameter vector [95].

With that setup, the algorithm runs multiple updates of an inner loop that improves the performance of the focal network measured against the fixed target network. After a round of updates in the inner loop, the inner loop's parameter vector overwrites the target network parameters. In effect, the target network is slowly updated by an outer loop that copies the learned vector of the inner loop.

Put another way, each inner loop is a learning period for the target network. The target network is then updated by inheriting the learned parameters of the inner loop. In biology, this sequence describes cultural or Lamarckian inheritance of acquired traits.

These single-vector procedures separate time scales in a way that matches the hierarchical FMB structure. The fast inner learner refines performance. The slow outer updates select the particular learned refinement that seeds the next round. That pattern is similar to what happens in group selection, in which the fast timescale of within-group selection improves performance, and the slow timescale of between-group selection chooses which within-group improvements seed the next generation.

14. Conclusions

The Price equation reveals a universal mathematical structure of algorithmic learning and natural selection, the FMB law, $\Delta\theta = \mathbf{M}\mathbf{f} + \mathbf{b} + \zeta$. This simple decomposition unifies seemingly disparate approaches, from natural selection's primary equation to machine learning's Adam optimizer and from Bayesian inference to Newton's method.

Each algorithm represents a particular choice of how to calculate or estimate the metric $\mathbf{M}$, the force $\mathbf{f}$, the bias $\mathbf{b}$, and any additional noise ζ . Typically, the metric describes inverse curvature, the force arises from the performance gradient, and the bias includes momentum, regularization of parameters, and changes in the frame of reference. In some cases, the metric encompasses broader notions of rescaling geometry, and the force comes from a different way of pushing toward improvement. But the essence of metric scaling and improving force remains the same.

The framework's power lies in its simplicity. By recognizing that many learning algorithms attempt to maximize the performance gain minus the cost paid for the distance moved in the parameter space, we see why certain mathematical quantities recur across disciplines. For example, Fisher information emerges naturally as a curvature metric in population or probabilistic models, whereas estimates of the Hessian of performance with respect to parameters describe curvature in locally focused methods.

The widespread commonality of the FMB structure suggests that advances in one domain can inform others. Computational methods for estimating curvature may illuminate evolutionary dynamics. The similar structure of Kalman filters and common machine learning optimization methods may suggest new refinements. The simple form of hierarchical learning may extend hierarchy to standard methods.

Ultimately, the FMB law provides a principled foundation for understanding existing algorithms and for designing new ones. Most often, we are choosing among different implementations of the same underlying process.

Funding: The Donald Bren Foundation and National Science Foundation grant DEB-2325755 support my research.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Price, G.R. Selection and covariance. *Nature* **1970**, *227*, 520–521.
2. Price, G.R. Extension of covariance selection mathematics. *Ann. Hum. Genet.* **1972**, *35*, 485–490.
3. Frank, S.A. Natural selection. IV. The Price equation. *J. Evol. Biol.* **2012**, *25*, 1002–1019.
4. Lande, R. Quantitative genetic analysis of multivariate evolution, applied to brain: Body size allometry. *Evolution* **1979**, *33*, 402–416.
5. Nocedal, J.; Wright, S.J. *Numerical Optimization*, 2nd ed.; Springer Series in Operations Research and Financial Engineering; Springer: New York, NY, USA, 2006.
6. Robbins, H.; Monro, S. A stochastic approximation method. *Ann. Math. Stat.* **1951**, *22*, 400–407. <https://doi.org/10.1214/aoms/1177729586>.

7. Polyak, B.T. Some methods of speeding up the convergence of iteration methods. *USSR Comput. Math. Math. Phys.* **1964**, *4*, 1–17. [https://doi.org/10.1016/0041-5553\(64\)90137-5](https://doi.org/10.1016/0041-5553(64)90137-5).
8. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980. <https://doi.org/10.48550/arxiv.1412.6980>.
9. Fisher, R.A. Theory of statistical estimation. *Math. Proc. Cambridge Philos. Soc.* **1925**, *22*, 700–725.
10. Rao, C.R. Information and the accuracy attainable in the estimation of statistical parameters. *Bull. Calcutta Math. Soc.* **1945**, *37*, 81–91.
11. Amari, S. Natural gradient works efficiently in learning. *Neural Comput.* **1998**, *10*, 251–276.
12. Shahshahani, S. *A New Mathematical Framework for the Study of Linkage and Selection*; Memoirs of the American Mathematical Society; American Mathematical Society: Providence, RI, USA, 1979; Volume 17.
13. Rice, S.H. A stochastic version of the Price equation reveals the interplay of deterministic and stochastic processes in evolution. *BMC Evol. Biol.* **2008**, *8*, 262.
14. Amari, S.; Nagaoka, H. *Methods of Information Geometry*; Oxford University Press: New York, NY, USA, 2000.
15. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016.
16. Mandt, S.; Hoffman, M.D.; Blei, D.M. Stochastic gradient descent as approximate Bayesian inference. *J. Mach. Learn. Res.* **2017**, *18*, 1–35.
17. Bottou, L.; Curtis, F.E.; Nocedal, J. Optimization methods for large-scale machine learning. *SIAM Rev.* **2018**, *60*, 223–311.
18. Kirkpatrick, M. Patterns of quantitative genetic variation in multiple dimensions. *Genetica* **2009**, *136*, 271–284.
19. Walsh, B.; Lynch, M. *Evolution and Selection of Quantitative Traits*; Oxford University Press: Oxford, UK, 2018.
20. Frank, S.A. The Price equation program: Simple invariances unify population dynamics, thermodynamics, probability, information and inference. *Entropy* **2018**, *20*, 978.
21. Frank, S.A. Simple unity among the fundamental equations of science. *Philos. Trans. R. Soc. B* **2020**, *375*, 20190351.
22. Kullback, S.; Leibler, R.A. On information and sufficiency. *Ann. Math. Stat.* **1951**, *22*, 79–86.
23. Jeffreys, H. An invariant form for the prior probability in estimation problems. *Proc. Roy. Soc. London Ser. A* **1946**, *186*, 453–461.
24. Harper, M. The replicator equation as an inference dynamic. *arXiv* **2009**, arXiv:0911.1763.
25. Frank, S.A. D'Alembert's direct and inertial forces acting on populations: The Price equation and the fundamental theorem of natural selection. *Entropy* **2015**, *17*, 7087–7100.
26. Jaynes, E.T. *Probability Theory: The Logic of Science*; Cambridge University Press: New York, NY, USA, 2003.
27. Pearl, J. *Causality*; Cambridge University Press: Cambridge, UK, 2009.
28. Frank, S.A.; Fox, G.A. The inductive theory of natural selection. In *The Theory of Evolution*; Scheiner, S.M., Mindell, D.P., Eds.; University of Chicago Press: Chicago, IL, USA, 2020; pp. 171–193.
29. Frank, S.A. The inductive theory of natural selection: Summary and synthesis. *arXiv* **2014**, arXiv:1412.1285.
30. Fisher, R.A. *The Genetical Theory of Natural Selection*, 2nd ed.; Dover Publications: New York, NY, USA, 1958.
31. Ewens, W.J. An interpretation and proof of the fundamental theorem of natural selection. *Theor. Popul. Biol.* **1989**, *36*, 167–180.
32. Ewens, W.J. An optimizing principle of natural selection in evolutionary population genetics. *Theor. Popul. Biol.* **1992**, *42*, 333–346.
33. Frank, S.A. Natural selection maximizes Fisher information. *J. Evol. Biol.* **2009**, *22*, 231–244.
34. Lande, R.; Arnold, S.J. The measurement of selection on correlated characters. *Evolution* **1983**, *37*, 1212–1226.
35. Crespi, B.J.; Bookstein, F.L. A path-analytic model for the measurement of selection on morphology. *Evolution* **1989**, *43*, 18–28.
36. Scheiner, S.M.; Mitchell, R.J.; Callahan, H.S. Using path analysis to measure natural selection. *J. Evol. Biol.* **2000**, *13*, 423–433.
37. Friston, K. The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci.* **2010**, *11*, 127–138. <https://doi.org/10.1038/nrn2787>.
38. Sakthivadivel, D.A.R. Towards a Geometry and Analysis for Bayesian Mechanics. *arXiv* **2022**, arXiv:2204.11900. <http://arxiv.org/abs/2204.11900>.
39. Conn, A.R.; Gould, N.I.; Toint, P.L. *Trust Region Methods*; SIAM: Philadelphia, PA, USA, 2000.
40. Conn, A.R.; Scheinberg, K.; Vicente, L.N. *Introduction to Derivative-Free Optimization*; SIAM: Philadelphia, PA, USA, 2009.
41. Nesterov, Y. A method for solving the convex programming problem with convergence rate $O(\frac{1}{k^2})$. *Dokl. Akad. Nauk. SSSR* **1983**, *269*, 543–547.
42. Sutskever, I.; Martens, J.; Dahl, G.; Hinton, G. On the importance of initialization and momentum in deep learning. In Proceedings of the 30th International Conference on Machine Learning (ICML-13), Atlanta, GA, USA, 16–21 June 2013; pp. 1139–1147.
43. Hansen, N.; Auger, A.; Ros, R.; Finck, S.; Pošík, P. Comparing results of 31 algorithms from the black-box optimization benchmarking BBOB-2009. In Proceedings of the 12th Annual Conference Companion on Genetic and Evolutionary Computation, Portland, OR, USA, 7–11 July 2010; pp. 1689–1696.
44. Hansen, N.; Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *Evol. Comput.* **2001**, *9*, 159–195.
45. Wierstra, D.; Schaul, T.; Glasmachers, T.; Sun, Y.; Peters, J.; Schmidhuber, J. Natural evolution strategies. *J. Mach. Learn. Res.* **2014**, *15*, 949–980.

46. Salimans, T.; Ho, J.; Chen, X.; Sidor, S.; Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv* **2017**, arXiv:1703.03864.
47. Akimoto, Y.; Hansen, N. Diagonal acceleration for covariance matrix adaptation evolution strategies. *Evol. Comput.* **2020**, *28*, 405–435.
48. Fletcher, R. *Practical Methods of Optimization*; John Wiley & Sons: Hoboken, NJ, USA, 2000.
49. Nemirovskij, A.S.; Yudin, D.B. *Problem Complexity and Method Efficiency in Optimization*; Wiley: Hoboken, NJ, USA, 1983.
50. Beck, A.; Teboulle, M. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Oper. Res. Lett.* **2003**, *31*, 167–175.
51. Welling, M.; Teh, Y.W. Bayesian learning via stochastic gradient Langevin dynamics. In Proceedings of the 28th International Conference on Machine Learning (ICML-11), Washington, DC, USA, 28 June–2 July 2011; pp. 681–688.
52. Zhu, Z.; Wu, J.; Yu, B.; Wu, L.; Ma, J. The Anisotropic Noise in Stochastic Gradient Descent: Its Behavior of Escaping from Sharp Minima and Regularization Effects. In Proceedings of the 36th International Conference on Machine Learning, Long Beach, CA, USA, 9–15 Jun 2019; Volume 97, pp. 7654–7663.
53. Keskar, N.S.; Nocedal, J.; Tang, P.T.P.; Mudigere, D.; Smelyanskiy, M. On large-batch training for deep learning: Generalization gap and sharp minima. In Proceedings of the 5th International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2017; pp. 2874–2889.
54. McCandlish, S.; Kaplan, J.; Amodei, D.; OpenAI Dota Team. An empirical model of large-batch training. *arXiv* **2018**, arXiv:1812.06162.
55. Smith, S.L.; Le, Q.V. A Bayesian perspective on generalization and stochastic gradient descent. *arXiv* **2018**, arXiv:1710.06451. <http://arxiv.org/abs/1710.06451>.
56. Rasmussen, C.E.; Williams, C.K. *Gaussian Processes for Machine Learning*; Number 3; MIT Press: Cambridge, MA, USA, 2006.
57. Kalman, R.E. A new approach to linear filtering and prediction problems. *Trans. ASME—J. Basic Eng.* **1960**, *82*, 35–45.
58. Welch, G.; Bishop, G. *An Introduction to the Kalman Filter*; Technical Report TR 95-041; Department of Computer Science, University of North Carolina: Chapel Hill, NC, USA, 1995.
59. Sarkka, S.; Solin, A.; Hartikainen, J. Spatiotemporal learning via infinite-dimensional Bayesian filtering and smoothing: A look at Gaussian process regression through Kalman filtering. *IEEE Signal Process. Mag.* **2013**, *30*, 51–61.
60. Shalizi, C.R. Dynamics of Bayesian updating with dependent data and misspecified models. *Electron. J. Stat.* **2009**, *3*, 1039–1074.
61. Frank, S.A. Natural selection. V. How to read the fundamental equations of evolutionary change in terms of information theory. *J. Evol. Biol.* **2012**, *25*, 2377–2396.
62. Campbell, J.O. Universal Darwinism as a process of Bayesian inference. *Hypothesis Theory* **2016**, *10*, 49.
63. Czégel, D.; Giffar, H.; Tenenbaum, J.B.; Szathmáry, E. Bayes and Darwin: How replicator populations implement Bayesian computations. *Bioessays* **2022**, *44*, 2100255.
64. Blei, D.M.; Kucukelbir, A.; McAuliffe, J.D. Variational inference: A review for statisticians. *J. Am. Stat. Assoc.* **2017**, *112*, 859–877.
65. Lanczos, C. *The Variational Principles of Mechanics*, 4th ed.; Dover Publications: New York, NY, USA, 1986.
66. Friston, K.; Kilner, J.; Harrison, L. A free energy principle for the brain. *J. Physiol.—Paris* **2006**, *100*, 70–87. <https://doi.org/10.1016/j.jphysparis.2006.10.001>.
67. Friston, K.; Mattout, J.; Trujillo-Barreto, N.; Ashburner, J.; Penny, W. Variational free energy and the Laplace approximation. *NeuroImage* **2007**, *34*, 220–234. <https://doi.org/10.1016/j.neuroimage.2006.08.035>.
68. Kaplan, R.; Friston, K.J. Planning and navigation as active inference. *Biol. Cybern.* **2018**, *112*, 323–343.
69. Baldwin, J.M. A new factor in evolution. *Am. Nat.* **1896**, *30*, 441–451.
70. West-Eberhard, M.J. *Developmental Plasticity and Evolution*; Oxford University Press: New York, NY, USA, 2003.
71. Maynard Smith, J. Group selection. *Q. Rev. Biol.* **1976**, *51*, 277–283.
72. Wilson, D.S. The group selection controversy: History and current status. *Annu. Rev. Ecol. Syst.* **1983**, *14*, 159–187.
73. Queller, D.C. Quantitative genetics, inclusive fitness, and group selection. *Am. Nat.* **1992**, *139*, 540–558.
74. West, S.A.; Griffin, A.S.; Gardner, A. Social semantics: How useful has group selection been? *J. Evol. Biol.* **2008**, *21*, 374–385.
75. Wang, J.X. Meta-learning in natural and artificial intelligence. *Curr. Opin. Behav. Sci.* **2021**, *38*, 90–95. <https://doi.org/https://doi.org/10.1016/j.cobeha.2021.01.002>.
76. Hinton, G.E.; Nowlan, S.J. How learning can guide evolution. *Complex Syst.* **1987**, *1*, 495–502.
77. Fernando, C.; Sygnowski, J.; Osindero, S.; Wang, J.; Schaul, T.; Teplyashin, D.; Sprechmann, P.; Pritzel, A.; Rusu, A. Meta-learning by the Baldwin effect. In Proceedings of the Genetic and Evolutionary Computation Conference Companion, Kyoto, Japan, 15–19 July 2018; pp. 109–110. <https://doi.org/10.1145/3205651.3205763>.
78. Finn, C.; Abbeel, P.; Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In Proceedings of the 34th International Conference on Machine Learning, Sydney, NSW, Australia, 6–11 August 2017; pp. 1126–1135.
79. Falkner, S.; Klein, A.; Hutter, F. BOHB: Robust and efficient hyperparameter optimization at scale. In Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, 10–15 July 2018; pp. 1437–1446.

80. Liu, H.; Simonyan, K.; Yang, Y. DARTS: differentiable architecture search. *arXiv* **2019**, arXiv:1806.09055.
81. Nichol, A.; Achiam, J.; Schulman, J. On first-order meta-learning algorithms. *arXiv* **2018**, arXiv:1803.02999.
82. Zoph, B.; Le, Q.V. Neural architecture search with reinforcement learning. *arXiv* **2017**, arXiv:1611.01578.
83. Hospedales, T.; Antoniou, A.; Micaelli, P.; Storkey, A. Meta-learning in neural networks: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *44*, 5149–5169.
84. Li, P.; Hao, J.; Tang, H.; Fu, X.; Zheng, Y.; Tang, K. Bridging evolutionary algorithms and reinforcement learning: A comprehensive survey on hybrid algorithms. *arXiv* **2024**, arXiv:2401.11963.
85. Hamilton, W.D. Innate social aptitudes of man: An approach from evolutionary genetics. In *Biosocial Anthropology*; Fox, R., Ed.; Wiley: New York, NY, USA, 1975; pp. 133–155.
86. Frank, S.A. *Foundations of Social Evolution*; Princeton University Press: Princeton, NJ, USA, 1998.
87. Okasha, S. *Evolution and the Levels of Selection*; Oxford University Press: New York, NY, USA, 2006.
88. Sobey, A.J.; Grudniewski, P. Re-inspiring the genetic algorithm with multi-level selection theory: multi-level selection genetic algorithm. *Bioinspiration Biomim.* **2018**, *13*, 056007.
89. Chand, S.; Howard, D. Multi-level evolution: Automatic discovery of hierarchical robot designs. *Front. Robot. AI* **2021**, *8*, 684304. <https://doi.org/10.3389/frobt.2021.684304>.
90. Jaderberg, M.; Dalibard, V.; Osindero, S.; Czarnecki, W.M.; Donahue, J.; Razavi, A.; Vinyals, O.; Green, T.; Dunning, I.; Simonyan, K.; et al. Population based training of neural networks. *arXiv* **2017**, arXiv: 1711.09846.
91. Khadka, S.; Tumer, K. Evolution-guided policy gradient in reinforcement learning. In Proceedings of the 32nd International Conference on Neural Information Processing Systems, Montreal, QC, Canada, 2–8 December 2018; pp. 1196–1208.
92. Khadka, S.; Majumdar, S.; Nassar, T.; Dwiel, Z.; Tumer, E.; Miret, S.; Liu, Y.; Tumer, K. Collaborative evolutionary reinforcement learning. In Proceedings of the 36th International Conference on Machine Learning, Long Beach, CA, USA, 9–15 June 2019; Volume 97, pp. 3521–3530.
93. Zhang, M.R.; Lucas, J.; Hinton, G.E.; Ba, J. Lookahead optimizer: K steps forward, 1 step back. In Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, BC, Canada, 8–14 December 2019.
94. Izmailov, P.; Podoprikin, D.; Garipov, T.; Vetrov, D.; Wilson, A.G. Averaging weights leads to wider optima and better generalization. In Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence, Monterey, CA, USA, 6–10 August 2018; pp. 876–885.
95. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.; Veness, J.; Bellemare, M.; Graves, A.; Riedmiller, M.; Fidjeland, A.; Ostrovski, G.; et al. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529–533. <https://doi.org/10.1038/nature14236>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.