



The Implications of the No-Free-Lunch Theorems for Meta-induction

David H. Wolpert^{1,2,3,4} 

Accepted: 23 February 2022 / Published online: 13 March 2023
© The Author(s), under exclusive licence to Springer Nature B.V. 2023

Abstract

The important recent book by Schurz (2019) appreciates that the no-free-lunch theorems (NFL) have major implications for the problem of (meta) induction. Here I review the NFL theorems, emphasizing that they do not only concern the case where there is a uniform prior—they prove that there are “as many priors” (loosely speaking) for which any induction algorithm A out-generalizes some induction algorithm B as vice-versa. Importantly though, in addition to the NFL theorems, there are many *free lunch* theorems. In particular, the NFL theorems can only be used to compare the expected performance of an induction algorithm A , considered in isolation, with the expected performance of an induction algorithm B , considered in isolation. There is a rich set of free lunches which instead concern the statistical *correlations* among the generalization errors of induction algorithms. As I describe, the meta-induction algorithms that Schurz advocates as a “solution to Hume’s problem” are simply examples of such a free lunch based on correlations among the generalization errors of induction algorithms. I end by pointing out that the prior that Schurz advocates, which is uniform over bit frequencies rather than bit patterns, is contradicted by thousands of experiments in statistical physics and by the great success of the maximum entropy procedure in inductive inference.

There is only limited value in knowledge derived from experience. The knowledge imposes a pattern, and falsifies, for the pattern is new in every moment.

T.S. Eliot.

✉ David H. Wolpert
dhw@santafe.edu
<http://davidwolpert.weebly.com>

¹ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

² Complexity Science Hub, Vienna, Austria

³ Arizona State University, Tempe, AZ 85287, USA

⁴ Research Staff Associate, International Center for Theoretical Physics, Trieste, Italy

1 Introduction

Inductive inference plays a central role in an extraordinarily wide variety of fields, ranging from traditional statistics to Monte Carlo estimation (Rubinstein and Kroese 2016; Tracey et al. 2013) to community detection and link detection in networks (Peel et al. 2017; Ghasemian et al. 2020; Guimerà 2020). It is also central to essentially all flavors of machine learning (Christopher 2006; Kroese et al. 2019), ranging from supervised learning (in which one is provided samples of an input-output map and wishes to infer the full map) to unsupervised learning (in which one is provided samples of a probability distribution and wishes to infer the full distribution) to active learning / experiment design (in which one is provided samples of an input-output map and wishes to determine what input to sample next, ultimately in order to infer the full map from all the samples) to reinforcement learning (in which one is provided samples of an action-reward map and wishes to infer a sequences of actions that will maximize the expected discounted sum of rewards).

The fact that inductive inference is central to all these fields means that we can sometimes use it as a dictionary, to “translate” techniques developed in one field over to another field, providing novel, powerful tools in that second field. As an example I am personally familiar with, the technique of ‘stacking’ is a powerful meta-inductive algorithm introduced in machine learning and statistics (Breiman 1996; Clarke 2003; Smyth and Wolpert 1999) where it is still being actively investigated and extended (Yao et al. 2018). It can be translated from those fields into the field of Monte Carlo estimation. In that new setting, stacking becomes a technique for post-processing a set of Monte Carlo samples of a distribution to improve the accuracy of the associated estimate of an expectation value. Empirically, this use of stacking seems to improve the accuracy of the estimator *no matter what precise Monte Carlo sampling algorithm is used* to generate the samples (simple sampling, importance sampling, quasi-Monte Carlo, etc.). As another example, stacking has been translated to the domain of link prediction in network science. In that domain it recently outperformed 203 alternative algorithms by optimally combining them, without making any Bayesian assumptions concerning the relative merits of those algorithms (Ghasemian et al. 2020; Guimerà 2020).

Another illustration of how inductive inference techniques are shared among multiple fields is provided by one of the pillars of the scientific method. Science relies deeply on the assumption that one can inductively infer future predictive accuracy from current *out-of-sample* predictive accuracy, even if as Hume emphasized, current in-sample predictive accuracy cannot be used that way (Hume 2003). To be a bit more precise, suppose we have a set of scientific theories, $T = \{T_i\}$, each of which can make predictions about the outcomes $y \in Y$ of any one of a set of experiments $x \in X$. Let d be the set of all experiment-outcome pairs that were observed by the people who made those theories. Let x_{-d}, x'_{-d} be two experiments that are not in d . A central assumption of the scientific method is that whichever of the theories in T is most accurate when predicting the outcome of the experiment x_{-d} would likely be more accurate than the other theories when predicting the outcome of the experiment x'_{-d} . In other words, if we choose among theories based on a set of new experiments S which were not used to create those theories in the first place, then we are likely to find the theory which would also be most accurate in all future experiments we might conduct after the ones in S .

If we translate this assumption underlying the scientific method into the fields of machine learning and statistics, we end up with the technique of cross-validation, which is a core

component of those fields.¹ In turn, the technique of cross-validation can be translated from the field of machine learning to the field of “black box” optimization. In that new field, it provides a way to dynamically set the parameters of an optimization algorithm (e.g., to dynamically set the temperature in simulated annealing). Just as out-of-sample induction has proven so successful in science in general, and was then so successful when translated into machine learning in the form of cross-validation, it has also been found empirically that when translated into the domain of black box optimization it results in faster convergence of the optimizer to the global optimum solution (Wolpert and Rajnarayan 2013; Adam et al. 2019).

Despite sharing inductive inference as a central component, there are many important distinguishers among these fields; they are not simply different expressions of the same underlying phenomenon. In particular, the no-free-lunch (NFL) theorems apply to optimization (Wolpert and Macready 1997) and to certain aspects of both supervised learning (Wolpert 1996a) and community detection (Peel et al. 2017). However, they do not apply to other aspects of supervised learning (Wolpert 1996b) and do not apply to co-evolutionary optimization (Wolpert and Macready 2005); there are “free lunches” in those domains.

Another example of a free lunch was introduced in 1996 by Parrondo and colleagues (Parrondo and Español 1996; Harmer and Abbott 1999). Suppose one has a pair of games that an agent can play, where the agent has higher probability of losing than winning in both games. Parrondo devised a strategy for the agent to play those games in an alternating sequence which results in *higher* probability of winning, despite the fact that in each game, separately, in each round of play, the agent faces a higher probability of losing. This result was called “Parrondo’s paradox”, since it was so surprising when it was discovered. As mentioned above, Parrondo’s result can be viewed as a “free lunch”, related to the meta-induction algorithms considered by Schurz.²

To properly understand the claims made in the recent book by Schurz (2019), it is worth taking a moment to walk through a simplified version of the scenario that Parrondo analyzed. Suppose there are $K + 1$ infinite sequences of bits, $\{v_k(i) \in \{0, 1\} : k = 1, \dots, K + 1, i \in \mathbb{Z}^+\}$. The first K of those sequences are the successive payoffs that a player would have received if they had picked that sequence. So for any counting number n , the accumulated payoff the player would have by iteration n of the sequences if they had always picked sequence k is $\pi_k(n) := \sum_{i=1}^n v_k(i)$. Define the “best” sequence on a given iteration n as $k^+(n) := \arg \max_k \pi_k(n)$, and the “worst” one as $k^-(n) := \arg \min_k \pi_k(n)$.

Next, suppose that on any iteration n , the associated bit $v_{K+1}(n)$ of the $K + 1$ ’th sequence is the value given by choosing the sequence whose accumulated payoff over the *previous* $n - 1$ iterations was highest, and evaluating its payoff bit for iteration n . Formally, $v_{K+1}(n) = v_{k^+(n-1)}(n)$. This rule for how to choose among the sequences on each successive iteration is a simplified version of Parrondo’s strategy.

To see why it can be good to follow this strategy, consider the limiting case where $K = 2$. Suppose that on some iteration n the difference $\pi_{k^+(n)}(n) - \pi_{k^-(n)}(n) = n$. Then it must that in *every* iteration $i \leq n$, $\pi_{k^+(i)}(i) - \pi_{k^-(i)}(i) = i$ (since the maximal difference in any single iteration is 1). So either both $v_1(i) = 1$ and $v_2(i) = 0$ for all $i \in \{1, \dots, n\}$, or vice-versa. As a result, $v_{K+1}(i) = i$ for all those iterations. In other words, following the

¹ To see this relationship, note that cross-validation chooses among a set of learning algorithms (rather than theories), and does so according to which of those performs best at out-of-sample prediction (evaluating that performance by forming “folds” of the single provided data set).

² As an historical aside, it’s interesting to note that Parrondo went on to make some of the seminal contributions to stochastic thermodynamics and non-equilibrium statistical physics (Parrondo 2015).

strategy v_{K+1} results in zero “regret”, in the formal sense of the word, in all iterations before n —there will be zero gap between one’s actual accumulated payoffs and the best possible value of the accumulated payoff, which one could have achieved if one knew the entire sequence of all payoffs for all time before picking among those sequences. Formally, $\pi_{K+1}(n) = \pi_{k^+(n)}(n)$.

If instead $\pi_{k^+(n)}(n) - \pi_{k^-(n)}(n)$ is some large value, but not quite n , then we can still provide guarantees that in *most* of the steps $m < n$, following the strategy $v_{k+1}(m)$ would have resulted in *very little* regret. For example, if $\pi_{k^+(n)}(n) - \pi_{k^-(n)}(n) = n - 2$, then in exactly one of the iterations $i \leq n$, the payoff $v_{k^+(n)}(i) = 0$ and $v_{k^-(n)}(i) = 1$, while in all other iterations, the opposite is true. So for all iterations $j < i$, there will be zero regret for using sequence $K + 1 = 3$, while for all iterations $j \geq i$, there will be regret 2 for using that sequence. So if $\pi_{k^+(n)}(n) - \pi_{k^-(n)}(n) = n - 2$, then we are guaranteed that the regret never exceeds 2 in any iteration $i \leq n$, if we use strategy $K + 1$, and that in half of all sequences, there would have been zero regret for using that strategy for at least half of the n iterations.

As $\pi_{k^+(n)}(n) - \pi_{k^-(n)}(n)$ shrinks, those guarantees on the fraction of iterations with little regret for using strategy $K + 1$ get weaker—but in addition, the maximal regret for using that strategy shrinks. So in general, following the strategy $v_{K+1}(i)$ for all iterations will not frequently result in a large amount of regret. In contrast, choosing some specific strategy $M \in \{1, 2\}$ and using that strategy for all iterations could result in quite bad regret by the iteration n . Conversely, if it *were* the case that using some such specific strategy $M \leq K$ resulted in little regret, then it would also be the case that using strategy $K + 1$ would result in little regret. In this sense, strategy $v_{K+1}(i)$ is superior to the other two strategies, with payoffs $v_1(i)$ and $v_2(i)$ respectively, *no matter what the sequences* $\{v_k(\cdot)\}$ are.

As K grows, this guarantee gets weaker—but it always holds. Moreover, there are strategies v_{K+1} that at each step use the preceding sequences $\{v_k(\cdot)\}$ in a more nuanced way than the all-or-nothing rule v_{K+1} described above. (In particular, the strategy underlying the Parrondo paradox is more nuanced than the all-or-nothing rule.) Importantly, all of these results hold even though no underlying probability distributions have been specified.

In some senses, one might argue that this Parrondo-like strategy for picking among sequences had a precursor in the informal investigations of Reichenbach et al. (1938). Related work was also done in 1997 (Cesa-Bianchi et al. 1997; see also 1996.) That particular analysis was later elaborated and extended in 2006 (Cesa-Bianchi and Lugosi 2006). The general topic of designing and analyzing algorithms for these kinds of scenarios is now known as “online learning under expert advice” (OLEA).

Schurz claims in Schurz (2019), that properly elaborated, these OLEA results “justify (meta-)induction” and “solve Hume’s problem”. In particular, such claims are made on behalf of various attractivity-weighted (AW) algorithms. Of course, since Schurz only considers the (rather limited) version of induction addressed by OLEA, these results cannot be said to “justify induction” in the full sense of inductive inference discussed at the beginning of this chapter. Moreover, since as mentioned above it was already known that there are free-lunch theorems (Wolpert and Macready 2005; Wolpert 1996b), some forms of (meta-) induction already had been “justified”.

It is also important to emphasize that in actual scientific practice, theories are not continually revised with each new experimental datum—in the language of the example above, $v_{K+1}(n)$ is modified to enforce the constraint that it only changes what sequence it chooses quite infrequently, rather than at every iteration, as in its original version described above. Similarly, in actual scientific practice, at any given time scientists are only considering at most a few theories—in the language of the example above, K is quite small. (After all, it would simply be too expensive, in many different aspects, to have many different scientific

theories all continually being updated.) For these and other reasons, the OLEA guarantees are quite weak when applied to actual scientific practice. So they provide little justification for the kind of induction scientists use.

This still leaves open the possibility that the kinds of guarantees given by OLEA could be significant in an idealized version of scientific practice. In this chapter I further analyze Schurz’s claims that this is true, in light of the NFL theorems.

2 The No Free Lunch Theorems

The NFL theorems for supervised learning are the ones most relevant for discussions of “induction” in the sense meant by Schurz. Let X be a finite input space, Y a finite output space. Suppose we have a target distribution $f(y_f \in Y \mid x \in X)$, along with a training set $d = (d_X^m, d_Y^m)$ of m pairs $\{(d_X^m(i) \in X, d_Y^m(i) \in Y)\}$, that is stochastically generated according to a distribution $P(d \mid f)$ (a conditional distribution conventionally called a likelihood). Assume that based on d we have a hypothesis distribution $h(y_h \in Y \mid x \in X)$. The creation of h from d is completely arbitrary. It is specified *in toto* by the distribution $P(h \mid d)$, and is conventionally called the learning algorithm. In addition, let $L(y_h, y_f)$ be a loss function taking $Y \times Y \rightarrow \mathbb{R}$. Next, fix some distribution $P(q \mid d_X)$. Finally, given these distributions, define the associated cost function by

$$C(f, h, d) \propto \sum_{y_f \in Y, y_h \in Y} \sum_{q \in X} P(q \mid d_X) L(y_f, y_h) f(y_f \mid q) h(y_h \mid q) \tag{1}$$

The cost function quantifies how well the algorithm does, averaged over all query points q , when the target is f , the hypothesis generated by the algorithm is h , and the training set is d . Note that the term $P(q \mid d_X)$ in the definition of the cost function governs how a query point q is generated for testing the performance of the algorithm, given the set of points the algorithm has already seen. So for example, under IID sampling to generate both the query point and the training set, $P(q \mid d_X)$ is independent of d_X . In contrast, if we are concerned with the ability of the algorithm to *generalize* from the training set, then we might require that $P(q \mid d_X) = 0$ if $q \in d_X$, since if the query point were the same as an element of the training set, then the cost function would quantify the memorization performance of the algorithm, not the generalization performance.

From now on, for simplicity, I will assume that any f is a single-valued function (i.e., $f(y_f \mid x)$ is a delta function for each x) and similarly for any h . I will also assume that the training-set generation process is “vertical”, in the sense that $P(d_Y \mid d_X, f)$ is independent of the values of $f(x)$ for $x \notin d_X$.

As an example of this framework, in Bayesian statistics analyses of “model mis-specification”, one might investigate the posterior expected cost,

$$\mathbb{E}(C \mid d) = \sum_{f, h} P(h \mid f, d) P(f \mid d) C(f, h, d) \tag{2}$$

$$= \sum_{f, h} P(h \mid d) P(f \mid d) C(f, h, d) \tag{3}$$

(the second line following from the fact that the hypothesis generated by the learning algorithm is conditionally independent of the target, given the training set). In contrast, in sampling theory statistics and computational learning theory, one is typically interested in

$$\mathbb{E}(C | f, m) = \sum_{h,d} P(h | d)P(d | f, m)C(f, h, d) \tag{4}$$

where m is the size of the training set.

This set of definitions is known as the “extended Bayesian framework (EBF)” (Wolpert 1995). The EBF is needed to properly understand the relationship between Bayesian and non-Bayesian statistics. It also allows us to go beyond those two bodies of work. For example, we can use the EBF to analyze the conditional distribution $\mathbb{E}(C | m)$. This allows us to derive a Bayesian correction to the conventional bias-plus-variance decomposition that arises in sampling theory statistics (Wolpert 1997; Kohavi and Wolpert 1996).³

Often when computational learning theory researchers refer to the “generalization error” of a supervised learning algorithm, they have in mind a data-blind cost function, meaning they choose $P(q | d_X)$ in eq. 1 to be independent of d_X . (For example, often one assumes that d_X was formed by IID sampling a distribution $\pi(x)$, and that $P(q | d_X) = \pi(q)$.) However, this choice of $P(q | d_X)$ conflates two very different aspects of induction: being able to recall elements of the training set (i.e., cases where $q \in d_X$), versus truly “generalizing” from the training set, to previously unseen instances (i.e., cases where $q \notin d_X$).

To help disentangle these two aspects of induction, one needs to use a distribution $P(q | d_X)$ that has zero measure on d_X , to focus on the generalization. Any $C(f, h, d)$ with this choice is known as an off-training set (OTS) cost function, and generically written as $C_{OTS}(f, h, d)$. The key feature of an OTS cost function is that it only depends on the partial functions $\{f(x) : x \notin d_X\}$ and $\{h(x) : x \notin d_X\}$ (see Wolpert 1996a). A standard example is any function of the type

$$C_{OTS}(f, h, d) = \frac{\sum_{q \notin d_X} \pi(q)L(f(q), h(q))}{\sum_{q \notin d_X} \pi(q)} \tag{5}$$

In the words of Schurz (2019), “the ultimate goal and evaluation criterion of inductive inferences is success in predictions.” Taken literally, this would imply that we should *only* be interested in OTS error. My personal view is that OTS cost is neither a “right” or “wrong” way to measure performance—rather it is an analytic tool for distinguishing two very different properties of any learning algorithm.

The no-free lunch theorems are also an analytical tool, designed to disentangle what aspects of a given learning algorithm can provide *a priori* guarantees concerning its expected OTS cost. It does by proving that if any given learning algorithm has particularly good OTS cost for one set of target functions, it must have correspondingly *poor* OTS cost on all other target functions. Formally, one of the NFL theorems says that for a broad range of choices of loss function (formally, for any “homogeneous loss function”), for any likelihood function, $\sum_f P(C_{OTS} | f, m)$ is independent of the learning algorithm. Similarly, another NFL theorem says that if $P(f)$ is uniform, then $P(C_{OTS} | d)$ is independent of the learning algorithm. These two NFL theorems imply in particular that whether one uses the the type of expected value of interest in Bayesian statistics (Eq. (3)) or the one of interest

³ The interested reader is directed to Wolpert (1995), Adam et al. (2019) for further-ranging discourse on how to integrate Bayesian and non-Bayesian statistics into an overarching probabilistic model of induction.

in non-Bayesian statistics (Eq. (4)), there are no *a priori*, assumption-free benefits to using one learning algorithm rather than another one.

A secondary implication of the NFL theorems is that *if* it so happens that you make an assumption, but it's that $P(f)$ is uniform, then the average over f 's used in the NFL theorem is the same as $P(f)$. In this case, you must conclude that all learning algorithms perform equally well for your assumed $P(f)$. This second implication is only as legitimate as is the assumption of uniform $P(f)$ it is based on, of course.

However, it must be emphasized that simply allowing $P(f)$ to be non-uniform, *by itself*, does not invalidate the NFL theorems. Arguments that only say that $P(f)$ is non-uniform in the real world, without advocating one particular non-uniformity, do not establish anything whatsoever about what learning algorithm to use in the real world. In fact, allowing $P(f)$'s to vary provides us with a new NFL theorem. In this new theorem, rather than compare the performance of two learning algorithms by uniformly averaging over all f 's, we compare them by uniformly averaging over all $P(f)$'s. The result is what one might expect: if any given algorithm A performs better than algorithm B over a given set of $P(f)$'s, then it must perform corresponding worse on all other $P(f)$'s.

This is the main message of the NFL theorems, not the fact that inference is impossible under the uniform prior $P(f)$. In fact, whether the uniform prior is “induction-hostile” is irrelevant. The NFL theorems do *not* assume that the universe is governed by a uniform prior in some objective sense. Nor do they suppose that the uniform prior somehow best captures our subjective ignorance about the universe—the NFL theorems do not motivate the uniform prior by invoking some variant of the common “maximal ignorance” reasoning underlying various priors found in the Bayesian statistics literature.

To re-emphasize, the NFL theorems are a mathematical *tool*, for analyzing *a priori* relationships between learning algorithms. It is a category error to interpret them as based on any “epistemic assumptions”. Indeed, what they force us to do is try to construct very weak assumptions that are not only reasonable, but also can be exploited to design learning algorithms that perform better than random guessing, (see Wolpert 1990 for earlier work in this vein). Summarizing, Schurz is simply wrong when he states, “Wolpert seems to assume that the state-uniform prior distribution is epistemically privileged.” (Schurz 2019, 240)—the NFL theorems make no assumption whatsoever concerning the epistemic nature of the uniform prior.

I end this section by emphasizing that NFL is completely consistent with many *free lunches*. Crucially, the NFL theorems equate the expected OTS performance of any two learning algorithms *only when they are considered independently*, in isolation from one another. However, in general the OTS performance of any two algorithms can be correlated as one varies over f 's. As an example, depending on the likelihood, loss functions, and other details, it may be that for all f , the expected OTS error of algorithm A , $\mathbb{E}_A(C|f, m)$ equals that of algorithm B , $\mathbb{E}_B(C|f, m)$, without violating NFL. In this case the maximal difference between the expected f -conditioned OTS errors of the algorithms as one varies over f is zero. On the other hand, it may instead be that the two algorithms are *anti*-correlated as one varies over f , again, without violating NFL. In other words, it may be that algorithm A performs better than random guessing on a function f iff algorithm B performs worse than random guessing on that function. In this case, the maximal difference between $\mathbb{E}_A(C|f, m)$ and $\mathbb{E}_B(C|f, m)$ as one varies over f 's can be quite large. As a third possibility, it may be that there are a few f for which algorithm A performs vastly better than algorithm B , but on the large number of other functions f , algorithm B performs just slightly better than algorithm A .

To make this more formal, fix two learning algorithms $\mathcal{A}_1, \mathcal{A}_2$, producing hypotheses $h_{\mathcal{A}_1}$ and $h_{\mathcal{A}_2}$, respectively, and write the associated cost functions as $C_{\mathcal{A}_1} = C(f, h_{\mathcal{A}_1}, d)$ and $C_{\mathcal{A}_2} = C(f, h_{\mathcal{A}_2}, d)$, respectively. Then NFL tells us that $\sum_f P(C_{\mathcal{A}_1} | f, m) = \sum_f P(C_{\mathcal{A}_2} | f, m)$, i.e., the two marginalizations of $\sum_f P(C_{\mathcal{A}_1}, C_{\mathcal{A}_2} | f, m)$ are identical. Nonetheless, in general, $\sum_f P(C_{\mathcal{A}_1}, C_{\mathcal{A}_2} | f, m)$ need not be a symmetric function of its two arguments—it may change if we interchange \mathcal{A}_1 and \mathcal{A}_2 , i.e., if we redefine \mathcal{A}_1 to always produce the hypothesis $h_{\mathcal{A}_2}$ that the original algorithm \mathcal{A}_2 would produce for a given training set d , and redefine \mathcal{A}_2 to always produce the hypothesis $h_{\mathcal{A}_1}$ that the original algorithm \mathcal{A}_1 would produce on d . In this case, we say that there are “head-to-head minimax” distinctions (Wolpert and Macready 1997) between the OTS performances of the two learning algorithms. Such distinctions might have had substantial repercussions for co-evolutionary scenarios like the development of life on Earth under natural selection, as elaborated in Wolpert and Macready (2005). As described below, it is precisely such head-to-head minimax distinctions that underlie the power of OLEA algorithms *in toto*.

3 Counter-Intuitive Implications of the NFL Theorems

In this section I first give a cursory sketch of how the NFL theorems can be consistent with the results of computational learning theory. I then present a scenario that illustrates how to exploit the NFL theorems in specific scenarios to derive counter-intuitive results that do not hold more generally.

1. Suppose that $P(h | d) = \delta(h, h^*)$, where $\delta(., .)$ is the Kronecker delta function. So the learning algorithm always produces the same hypothesis h^* , no matter what d is. Also suppose that the likelihood function is $P(d | f)$ is noise-free. Consider the data-blind cost function $C(f, h) = \sum_x \pi(x)L(f(x), h(x))$, and define the associated f -conditioned expected empirical cost as

$$\hat{C}(f, m) := \sum_{h,d: |d_X|=m, q \in d_X} \frac{P(h | d)L(f(q), h(q))P(d | f)\pi(q)}{\sum_{q \in d_X} \pi(q)} \tag{6}$$

$$= \sum_{d: |d_X|=m, q \in d_X} \frac{L(f(q), h^*(q))P(d_X | f)\pi(q)}{\sum_{q \in d_X} \pi(q)} \tag{7}$$

The law of large numbers assures us that $\mathbb{E}(C | f, m)$ converges to $\hat{C}(f, m)$ as m and $|X|$ both grow with $|X| \gg m$, e.g., for a uniform distribution π . Therefore $\mathbb{E}(C | m)$ as well converges to the m -conditioned empirical cost. (Note this is true for any prior $P(f)$.) One of the primary concerns of the field of computational learning theory is characterizing the precise form of this kind of convergence in different scenarios. Next, note that for $|X| \gg m$, $C(f, h^*) \approx \mathbb{E}(C_{OTS} | f, h, m)$, e.g., using the OTS cost function of Eq. (5) for a distribution π that is uniform over X . One might suppose that this means that if $|X| \gg m$, then $\mathbb{E}(C_{OTS} | m, \hat{C})$ would also become peaked about $C = \hat{C}(f, m)$. In fact though, by NFL, the average over priors $P(f)$ of $\mathbb{E}(C_{OTS} | m, \hat{C})$ is independent of \hat{C} , since \hat{C} has no statistical coupling with C_{OTS} under that average.⁴

⁴ The interested reader is directed to Adam et al. (2019) for further discussion reconciling the NFL theorems and computational learning theory.

2. Given any fixed set of learning algorithms, $\{\mathcal{A}_i\}$, define $\Phi(\{\mathcal{A}_i\})$ to be the learning algorithm that for any data set d determines which of the \mathcal{A}_i has lowest cross-validation error on d and then uses that \mathcal{A}_i to predict the output for all questions $q \notin d_X$, with any convenient tie-breaking mechanism. (For current purposes, there is no need to specify the precise type of cross-validation, e.g., K -fold, leave-one-out, etc.) Similarly define $\hat{\Phi}(\{\mathcal{A}_i\})$ to be the learning algorithm that for any data set d determines which of the \mathcal{A}_i has highest cross-validation error on d and then uses that \mathcal{A}_i to predict the output for all questions $q \notin d_X$, with any convenient tie-breaking mechanism. I will refer to $\Phi(\{\mathcal{A}_i\})$ and $\hat{\Phi}(\{\mathcal{A}_i\})$ as the “method of cross-validation” and the “method of anti-cross-validation”, respectively. Note that for any fixed set of learning algorithms $\{\mathcal{A}_i\}$ that those two methods are applied to, each of them is itself a learning algorithm, i.e., a map from a provided training set d to a hypothesis function h . Next, suppose that $Y = \{0, 1\}$. Again suppose that the likelihood is noise-free. For simplicity consider the case where $\{\mathcal{A}_i\}$ contains exactly two learning algorithms. The “majority” learning algorithm \mathcal{A}_1 predicts $1/0$ for all off-training set queries, depending on whether the output $y = 1/0$ was more common in the data set (with an arbitrary tie-breaking choice). The “anti-majority” learning algorithm \mathcal{A}_2 instead predicts $1/0$ for all off-training set queries, depending on whether the output $y = 0/1$ was more common in the data set. (So \mathcal{A}_1 predicts whatever was the most common output in the training set, independent of the precise question $q \notin d_X$, and \mathcal{A}_2 predicts whatever was the least common output.) Choose the OTS zero-one cost function, for simplicity defined for a uniform sampling distribution over the OTS q ’s:

$$C_{OTS}(f, h, d) = \sum_{x \notin d_X} \frac{1 - \delta(h(x), f(x))}{|X| - m} \tag{8}$$

Consider the prior $P^\dagger(f)$ that allows just the two constant functions: $f(x) = 1 \forall x \in X$, and $f(x) = 0 \forall x \in X$, assigning each of them the probability $1/2$. For either of those two constant functions, for any training set d , $\hat{\Phi}(\{\mathcal{A}_1, \mathcal{A}_2\})$ always makes the wrong prediction for any OTS question. So the expected OTS zero-one loss of anti-cross-validation is 1. This is true whether we condition on a single data set (as in Bayesian statistics) or average over all data sets of a given size (as in sampling theory statistics).⁵ By NFL, this means that there must be some *other* prior, $P^*(f) \neq P^\dagger(f)$, for which the expected OTS zero-one cost of anti-cross-validation is less than $1/2$. (Note that in general such a “compensating” prior may assign nonzero probability to functions f that are not constant over all X .) Next, note that the sum of the expected OTS zero-one cost of cross-validation and anti-cross-validation conditioned on d and one of the two allowed target functions f is independent of f, d :

$$\mathbb{E}_{\Phi(\{\mathcal{A}_i\})}(C_{OTS} \mid d, f) + \mathbb{E}_{\hat{\Phi}(\{\mathcal{A}_i\})}(C_{OTS} \mid d, f) = 1 \tag{9}$$

(This is due to the nature of the majority and anti-majority algorithms and has nothing to do with NFL.) Therefore for *any* prior $P(f)$ over the two constant functions, the sum of the expected errors conditioned on d (as in Bayesian statistics) equals 1:

$$\mathbb{E}_{\Phi(\{\mathcal{A}_i\})}(C_{OTS} \mid d) + \mathbb{E}_{\hat{\Phi}(\{\mathcal{A}_i\})}(C_{OTS} \mid d) = 1 \tag{10}$$

⁵ It is also true if we condition on a particular one of the two allowed f ’s, as in sampling theory statistics, in which case the prior is irrelevant, and NFL does not apply.

Combining, since the expected OTS zero-one loss of anti-cross-validation is less than $1/2$ for the prior $P^*(f)$, the expected OTS zero-one loss of cross-validation must be greater than $1/2$ for that prior. Moreover, no matter what f is, and therefore no matter the prior, expected zero-one OTS loss is $1/2$ for the algorithm that always guesses randomly, with probability $1/2$ of choosing $y = 1$. The NFL theorems already told us that there must be a prior $P'(f)$ for which cross-validation performs worse than random guessing, and that there must be a prior $P''(f)$ for which anti-cross-validation performs *better* than random guessing. However, one might have suspected that in general, those would have to be different priors, i.e., that $P'(f) \neq P''(f)$. For example, one might have supposed that any prior $P'(f)$ for which anti-cross-validation does worse than random guessing is also a prior for which cross-validation does worse than random guessing (and vice-versa). The analysis above shows that this is not the case: For the single prior $P^*(f)$, the method of anti-cross-validation is successful, in the sense of performing better than random guessing. However, for that same prior the method of cross-validation is *not* successful, and performs worse than a random coin-toss. (Note that the experiments recounted in Sect. 9 of Schurz (2019) are consistent with this phenomenon.) This kind of phenomenon holds more generally. For *any* method \mathcal{I} , and associated “anti-” method $\hat{\mathcal{I}}$, there exist priors for which $\hat{\mathcal{I}}$ is successful, but \mathcal{I} is not. (Or as others might put, for any such “meta-induction algorithm” \mathcal{I} and “anti-meta-induction algorithm” $\hat{\mathcal{I}}$.) In this very specific sense, any claim that some such method \mathcal{I} “is guaranteed to be successful, no matter the course of nature, if any method is” Tom (2019) is wrong. This is true even though there *can* be a method \mathcal{I} that performs better than the associated method $\hat{\mathcal{I}}$ in head-to-head minimax distinctions, “no matter the course of nature”. In sum, whether one method can have such guarantees depends on how precisely one is comparing methods.

4 No Free Lunch and OLEA

The central problem in the simplified OLEA scenario introduced in Sect. 1 is how to set the value $v_{K+1}(m)$ based on knowledge of the values of the preceding values of the K other sequences, $\{v_k(i) : i \in \{1, \dots, m-1\}, k \in \{1, \dots, K\}\}$. One can map this problem into a special case of the supervised learning problem of how to generalize from a particular training set d . Take $X = \mathbb{Z}^+$, and identify the successive iterations $i \in \mathbb{Z}^+$ with successive elements of X . Also take $Y = \{0, 1\}$. So each f is a function from $\mathbb{Z}^+ \rightarrow \{0, 1\}$. Identify d_X with the first m counting numbers, and choose d_Y to be any vector of m bits. Also identify each of the first K sequences $\{v_k(x) : k = 1, \dots, K\}$ with the values of K different considered functions, $\{g_k(x) : k = 1, \dots, K\}$, by setting $g_k(x) = d_Y(x)$ iff $v_k(x) = 1$, for all $x \in d_X$. (The values of those candidate functions for values $x > m$ is arbitrary.) So the sequence v_k has a payoff value of $0/1$ for iteration x depending on whether g_k agrees with the training set on $x \in d_X$. Assume a noise-free likelihood function $P(d|f)$, and adopt a OTS cost function $C(f, h, d) = 1 - \delta(h(|d_X| + 1), f(|d_X| + 1))$.

At heart, when the algorithms considered in Schurz (2019) are mapped this way into the realm of supervised learning, they become various learning algorithms for using the m -element training set d to combine the values of the candidate functions g_k evaluated for $x = m + 1$ in order to set the value of the hypothesis function at $x = m + 1$, i.e., in order to set $h(m + 1)$. The associated value $v_{K+1}(m + 1)$ is set to $0/1$ depending on whether that

value $h(m+1)$ produced from the candidate functions equals $f(m+1)$. So translated back into the context of Schurz (2019), the OTS cost function is 0/1 depending on whether $v_{K+1}(x) = 0/1$.

The NFL theorems for supervised learning tell us immediately that averaged over all f , the expected value of this OTS cost is 1/2 for *all* algorithms v_{K+1} (not just those algorithms considered in Schurz (2019)). More generally, if we average uniformly over all priors $P(f)$, then the expected value of $v_{K+1}(x)$ in any iteration x is 1/2. This is true *no matter how we set the sequence* v_{K+1} , no matter what d is, and no matter what the candidate functions g_k are, i.e., no matter what the sequences $\{v_k : k = 1, \dots, K\}$ are.

How can these NFL results for OLEA be reconciled with the regret-reducing results of OLEA in general, and the benefits of AW algorithms in particular? The answer was provided in Wolpert and Macready (1997): OLEA results concern head-to-head distinctions between learning algorithms, and there can be free lunches for head-to-head distinctions. Whether such free lunches are normative, determining how one “should” make predictions is a nuanced topic. (For example, recall the discussion above about natural selection and co-evolutionary free lunches.) Under the most conventional formulations of Bayesian decision theory, the answer is ‘no’, these kinds of distinctions do not provide a reason to prefer one algorithm over another. In this sense, Schurz’s claim to “solve Hume’s problem” results from subtle and rich but ultimately flawed reasoning.

As a final point, while Schurz does not consider the NFL theorems involving an average over priors $P(f)$, he does address the case of a uniform $P(f)$, by contesting its “epistemic validity”. Specifically Schurz argues that one “should” adopt a single, specific prior, in a normative sense (a stance I do not promote). However, Schurz argues that it should be a uniform prior over *frequencies* of the future sequences of bits, rather than (as under uniform $P(f)$) over the patterns of those bits.

In response to this it is important to point out that *all* of statistical physics is based on a uniform distribution over patterns, *not* over frequencies; that uniform distribution over patterns is known as the *microcanonical ensemble*. As an example, under the microcanonical ensemble, the distribution of joint states of all the binary spins in an Ising spin is uniform, and so the distribution of frequencies of the average spin value is highly *non*-uniform. Indeed, the whole validity of standard, macroscopic thermodynamics, relies on the fact that in the thermodynamic limit of an infinite number of spins, the distribution over frequencies becomes a Dirac delta function.

In addition, the highly successful Maximum entropy procedure for inductively inferring (!) a probability distribution from knowledge of its moments relies on a uniform prior over patterns, not over frequencies (Jaynes and Bretthorst 2003; Jaynes 1968). In short, Schurz’s proposal for a uniform prior over frequencies runs afoul of thousands (tens of thousands?) of previous experiments concerning the real, physical world. Again, the central issue is how one is comparing algorithms. In all of those real-world experiments, the key issue is not head-to-head minimax distinctions, which is what allows there to be such strong arguments in favor of a uniform prior.

Acknowledgements I would like to thank the Santa Fe Institute for support.

References

- Adam, S. P., Alexandropoulos, S. A. N., Pardalos, P. M., & Vrahatis, M. N. (2019). No free lunch theorem: A review. In *Approximation and Optimization* (pp. 57–82). Springer.
- Breiman, L. (1996). Stacked regression. *Machine Learning*, 24.
- Cesa-Bianchi, N., Long, P. M., & Warmuth, M. K. (1996). Worst-case quadratic loss bounds for prediction using linear functions and gradient descent. *IEEE Transactions on Neural Networks*, 7(3), 604–619.

- Cesa-Bianchi, N., Freund, Y., Haussler, D., Helmbold, D. P., Schapire, R. E., & Warmuth, M. K. (1997). How to use expert advice. *Journal of the ACM (JACM)*, 44(3), 427–485.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge University Press.
- Christopher, M. (2006). *Bishop pattern recognition and machine learning*. Springer.
- Clarke, B. (2003). Bayes model averaging and stacking when model approximation error cannot be ignored. *Journal of Machine Learning Research*, 683–712.
- Ghasebian, A., Hosseinmardi, H., Galstyan, A., Airolidi, E. M., & Clauset, A. (2020). Stacking models for nearly optimal link prediction in complex networks. *Proceedings of the National Academy of Sciences*, 117(38), 23393–23400.
- Guimerà, R. (2020). One model to rule them all in network science? *Proceedings of the National Academy of Sciences*, 117(41), 25195–25197.
- Hans, R., et al. (1938). Experience and prediction: An analysis of the foundations and the structure of knowledge.
- Harmer, G. P., & Abbott, D. (1999). Losing strategies can win by parrondo's paradox. *Nature*, 402(6764), 864–864.
- Hume, D. (2003). *A treatise of human nature*. Courier Corporation.
- Jaynes, E. T., & Bretthorst, G. L. (2003). *Probability theory: The logic of science*. Cambridge University Press.
- Jaynes, E. T. (1968). Prior probabilities. *IEEE Transactions on systems science and cybernetics*, 4(3), 227–241.
- Kohavi, R., & Wolpert, D. H. (1996). Bias plus variance decomposition for zero-one loss functions. In *ICML* (Vol. 96, pp. 275–283).
- Kroese, D. P., Botev, Z., Taimre, T., & Vaisman, R. (2019). *Data science and machine learning: mathematical and statistical methods*. CRC Press.
- Parrondo, J. M. R., Horowitz, J. M., & Sagawa, T. (2015). Thermodynamics of information. *Nature Physics*, 11(2), 131–139.
- Parrondo, J. M. R. & Español, P. (1996). Criticism of feynman's analysis of the ratchet as an engine. *American Journal of Physics*, 64(9), 1125–1130.
- Peel, L., Larremore, D. B., & Clauset, A. (2017). The ground truth about metadata and community detection in networks. *Science Advances*, 3(5), e1602548.
- Rubinstein, R. Y., & Kroese, D. P. (2016). *Simulation and the Monte Carlo method* (Vol. 10). Wiley.
- Schurz, G. (2019). *Hume's problem solved: The optimality of meta-induction*. MIT Press.
- Smyth, P., & Wolpert, D. (1999). Linearly combining density estimators via stacking. *Machine Learning*, 36(1–2), 59–83.
- Tom, F. (2019). *Sterkenburg. The meta-inductive justification of induction*. Episteme.
- Tracey, B., Wolpert, D., & Alonso, J. J. (2013). Using supervised learning to improve monte carlo integral estimation. *AIAA Journal*, 51(8), 2015–2023.
- Wolpert, D. H. (1995). The relationship between PAC, the statistical physics framework, the Bayesian framework, and the VC framework. In *The mathematics of generalization* (pp. 117–215). Addison-Wesley.
- Wolpert, D. H. (1990). The relationship between occam's razor and convergent guessing. *Complex Systems*, 4, 319–368.
- Wolpert, D. H. (1996a). The lack of a prior distinctions between learning algorithms. *Neural Computation*, 8(1341–1390), 1391–1421.
- Wolpert, D. H. (1996b). The existence of a priori distinctions between learning algorithms. *Neural Computation*, 8, 1391–1420.
- Wolpert, D. H. (1997). On bias plus variance. *Neural Computation*, 9, 1211–1244.
- Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1), 67–82.
- Wolpert, D. H., & Macready, W. (2005). Coevolutionary free lunches. *Transactions on Evolutionary Computation*, 9, 721–735.
- Wolpert, D., & Rajnarayan, D. (2013). Using machine learning to improve stochastic optimization. In *Workshops at the twenty-seventh AAAI conference on artificial intelligence*.
- Yao, Y., Vehtari, A., Simpson, D., & Gelman, A. (2018). Using stacking to average bayesian predictive distributions (with discussion). *Bayesian Analysis*, 13(3), 917–1007.