Ensemble Methods for Structured Prediction

Corinna Cortes CORINNA@GOOGLE.COM
Google Research, 111 8th Avenue, New York, NY 10011

Vitaly Kuznetsov VITALY@CIMS.NYU.EDU
Courant Institute of Mathematical Sciences, 251 Mercer Street, New York, NY 10012

Mehryar Mohri MOHRI@CIMS.NYU.EDU
Courant Institute and Google Research, 251 Mercer Street, New York, NY 10012

Abstract

We present a series of learning algorithms and theoretical guarantees for designing accurate ensembles of structured prediction tasks. This includes several randomized and deterministic algorithms devised by converting on-line learning algorithms to batch ones, and a boosting-style algorithm applicable in the context of structured prediction with a large number of labels. We give a detailed study of all these algorithms, including the description of new on-line-to-batch conversions and learning guarantees. We also report the results of extensive experiments with these algorithms in several structured prediction tasks.

1. Introduction

Ensemble methods are general techniques in machine learning for combining several hypotheses to create a more accurate predictor (Breiman, 1996; Freund & Schapire, 1997; Smyth & Wolpert, 1999; MacKay, 1991; Freund et al., 2004). These methods in practice often significantly improve the performance and additionally benefit from favorable learning guarantees. However, ensemble methods and their theory have been developed primarily for the binary classification problem or regression tasks. These techniques do not readily apply to structured prediction problems. While it is straightforward to combine scalar outputs for a classification or regression problem, it is less clear how to combine structured predictions such as phonemic pronunciation hypotheses, speech recognition lattices, parse trees, or outputs of several machine translation systems. Areas like natural language processing, optical character recognition and computer vision, in general gives rise to many structured prediction problems with structures or substructures varying with each task.

Proceedings of the 31^{st} International Conference on Machine Learning, Beijing, China, 2014. JMLR: W&CP volume 32.

Consider for example the problem of devising an ensemble method for pronunciation, a critical component of modern speech recognition (Ghoshal et al., 2009). Often, several pronunciation models or experts are available for transcribing words into sequences of phonemes. These models may have been derived using other machine learning algorithms or they may be based on carefully hand-crafted rules. In general, none of these pronunciation experts is fully accurate and each expert may be making mistakes at different positions in the output sequence. One can hope that a model that *patches together* the pronunciation of different experts could achieve a superior performance.

We seek to tackle all mentioned structured prediction problems simultaneously and consider the general setting where the label or output associated to an input $\mathbf{x} \in \mathcal{X}$ is a structure $\mathbf{y} \in \mathcal{Y}$ that can be decomposed and represented by l substructures $y^1, \ldots, y^l$. For the pronunciation example just discussed, $\mathbf{x}$ is a specific word or word sequence and $\mathbf{y}$ its phonemic transcription. A natural choice for the substructures y^k is then the individual phonemes forming $\mathbf{y}$. Other possible choices include n-grams of consecutive phonemes or more general subsequences.

We will assume that the loss function considered admits an additive decomposition over the substructures, as it is common in structured prediction. We also assume access to a set of structured prediction experts $h_1, \ldots, h_p$ that we treat as black boxes. Given an input $\mathbf{x} \in \mathcal{X}$, each of these experts predicts l substructures $h_j(\mathbf{x}) = (h_j^1(\mathbf{x}), \ldots, h_j^l(\mathbf{x}))$. The hypotheses h_j may be the output of other structured prediction algorithms such as Conditional Random Fields (Lafferty et al., 2001), Averaged Perceptron (Collins, 2002), StructSVM (Tsochantaridis et al., 2005), Max Margin Markov Networks (Taskar et al., 2004), the Regression Technique for Learning Transductions (Cortes et al., 2005), or some other algorithmic or human expert. Given a labeled training sample $(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_m, \mathbf{y}_m)$, our objective is to combine the predictions of these experts to form an accurate ensemble.

Variants of the ensemble problem just formulated have

been studied in the past in the natural language processing and machine learning literature. One of the most recent, and possibly most relevant studies for sequence data is that of Nguyen & Guo (2007), which is based on the forward stepwise selection procedure introduced by Caruana et al. (2004). Starting with a possibly empty collection of experts, E_0, that algorithm performs T iterations. To make predictions using a collection of models, E_t, a variant of a majority-vote scheme per position is proposed, and at each iteration t, a new expert h_j from $\{h_1, \ldots, h_p\}$ is added to the collection E_{t-1} in such a way that $E_t = E_{t-1} \cup \{h_j\}$ has the best performance on the training set among all sets $E_{t-1} \cup \{h_j\}$, $j = 1, \ldots, p$. This algorithm always performs at least as well as the best expert among $h_1, \ldots, h_p$ on the training set. If the initial collection, E_0 of experts is empty, then E_1 simply contains the expert with the smallest error on the training set. Further additions to E_t only decrease the error on the training set and hence the performance of this algorithm on the training set cannot be worse than the performance of the best expert.

One disadvantage of this greedy approach is that it may fail to select an optimal ensemble of experts in cases where experts specialize in local predictions. Consider the case where expert h_k is a strong predictor for the kth substructure but does not perform well on other substructures. Assume further that expert h_0 is a jack-of-all-trades and performs better than any of $h_1, \ldots, h_p$ on average, but each h_k beats h_0 at position k. Then, one can show that the stepwise selection routine may end up with an ensemble consisting of only h_0, while an optimal solution would use expert h_k to predict the kth substructure. We provide an explicit construction of such an example in Appendix I and report similar empirical observations in Section 5.

Ensemble methods for structured prediction based on bagging, random forests and random subspaces have been proposed in (Kocev et al., 2013). One of the limitations of this work is that it is applicable only to a very specific class of tree-based experts introduced in that paper. Similarly, a boosting approach was developed in (Wang et al., 2007) but it applies only to local experts. In the context of natural language processing, a variety of different re-ranking techniques have been proposed for somewhat related problems (Collins & Koo, 2005; Zeman & Žabokrtský, 2005; Sagae & Lavie, 2006; Zhang et al., 2009). But, re-ranking methods do not combine predictions at the level of substructures, thus the final prediction of the ensemble coincides with the prediction made by one of the experts, which can be shown to be suboptimal in many cases. Furthermore, these methods typically assume the use of probabilistic models, which is not a requirement in our learning scenario. Other ensembles of probabilistic models have also been considered in text and speech processing by forming a product of probabilistic models via the intersection of lattices (Mohri et al., 2008), or a straightforward combination of the posteriors from probabilistic grammars trained using EM with different starting points (Petrov, 2010), or some other rather intricate techniques in speech recognition (Fiscus, 1997). See Appendix J for a brief discussion of other related work.

Most of the references mentioned do not give a rigorous theoretical justification for the techniques proposed. We are not aware of any prior theoretical analysis for the ensemble structured prediction problem we consider. Here, we aim to bridge this gap and develop ensemble methods that both perform well in practice and enjoy strong theoretical guarantees. Two families of algorithms are introduced. In Section 3 we develop ensemble methods based on on-line algorithms. To do so, we extend existing on-line-to-batch conversions to our more general setting. A boosting-type algorithm is also presented and analyzed in Section 4. Section 5 reports the results of our extensive experiments.

2. Learning scenario

As in standard supervised learning problems, we assume that the learner receives a training sample $S = ((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_m, \mathbf{y}_m)) \in \mathcal{X} \times \mathcal{Y}$ of m labeled points drawn i.i.d. according to some distribution $\mathcal{D}$ used both for training and testing. We also assume that the learner has access to a set of p predictors $h_1, \ldots, h_p$ mapping $\mathcal{X}$ to $\mathcal{Y}$ to devise an accurate ensemble prediction. Thus, for any input $\mathbf{x} \in \mathcal{X}$, he can use the prediction of the p experts $h_1(\mathbf{x}), \ldots, h_p(\mathbf{x})$. No other information is available to the learner about these p experts, in particular the way they have been trained or derived is not known to the learner. But, we will assume that the training sample S available to learn the ensemble is distinct from what may been used for training the algorithms that generated $h_1(\mathbf{x}), \ldots, h_p(\mathbf{x})$.

To simplify our analysis, we assume that the number of substructures $l \geq 1$ is fixed. This does not cause any loss of generality so long as the maximum number of substructures is bounded, which is the case in all the applications we consider. The quality of the predictions is measured by a loss function $L \colon \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$ that can be decomposed as a sum of loss functions $\ell_k \colon \mathcal{Y}_k \to \mathbb{R}_+$ over the substructure sets $\mathcal{Y}_k$, that is, for all $\mathbf{y} = (y^1, \ldots, y^l) \in \mathcal{Y}$ with $y^k \in \mathcal{Y}_k$ and $\mathbf{y}' = (y'^1, \ldots, y'^l) \in \mathcal{Y}$ with $y'^k \in \mathcal{Y}_k$,

$$L(\mathbf{y}, \mathbf{y}') = \sum_{k=1}^{l} \ell_k(y^k, y'^k). \tag{1}$$

We will assume in all that follows that the loss function L is bounded by some $M > 0$: $L(\mathbf{y}, \mathbf{y}') \leq M$ for all $(\mathbf{y}, \mathbf{y}')$. A prototypical example of such loss functions is the nor-

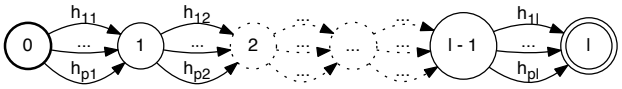

Figure 1. Directed graph G of path experts.

malized Hamming loss, L_{Ham}, which is the fraction of substructures for which two labels $\mathbf{y}$ and $\mathbf{y}'$ disagree.

3. On-line learning approach

In this section, we present an on-line learning solution to the ensemble structured prediction problem just discussed. We first formulate the problem as that of on-line learning with expert advice, where the experts correspond to the paths of a directed graph. The on-line algorithm generates at each iteration a distribution over the path-experts. A critical component of our approach consists of using the distributions to define a prediction algorithm with good generalization guarantees. This requires an extension of the existing on-line-to-batch conversion techniques to the more general case of combining distributions over path-experts (instead of combining intermediate hypotheses).

3.1. Path experts

Each expert h_j induces a set of substructure hypotheses $h_j^1, \ldots, h_j^l$. As already discussed, one particular expert may be better at predicting the kth substructure while some other expert may be more accurate at predicting another substructure. Therefore, it is desirable to combine the substructure predictions of all experts to derive a more accurate prediction. This leads us to considering a directed graph G such as that of Figure 1 which admits $l+1$ vertices $0, 1, \ldots, l$ and an edge from vertex k to vertex $k+1$ labeled with each of the p hypotheses $h_1^k, \ldots, h_p^k$ induced by the experts $h_1, \ldots, h_p$ for the kth substructure. Graph G compactly represents a set of *path experts*: each path from the initial vertex 0 to the final vertex l is labeled with a sequence of substructure hypotheses $h_{j_1}^1, \ldots, h_{j_l}^l$ and defines a hypothesis which associates to input $\mathbf{x}$ the output $h_{j_1}^1(\mathbf{x}) \cdots h_{j_l}^l(\mathbf{x})$. We will denote by H the set of all p^l path experts. We also denote by h each path expert defined by $h_{j_1}^1, \ldots, h_{j_l}^l$, with $j_k \in \{1, \ldots, p\}$, and denote by h^k its kth substructure hypothesis $h_{j_k}^k$. Our ensemble structure prediction problem can then be formulated as that of selecting the best path expert (or collection of path experts) in the graph G. Note that, in general, the path expert selected does not coincide with any of the original experts $h_1, \ldots, h_p$.

More generally, our paths experts can be selected from a directed acyclic graph of experts G' distinct from G, as illustrated by Figure 2. This can be motivated by scenarios where some prior knowledge is available about the expert

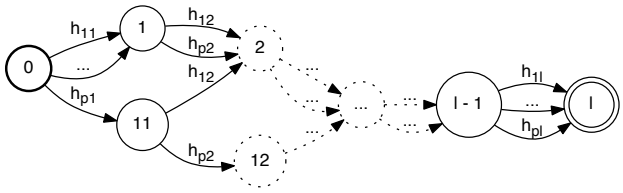

Figure 2. Alternative graph G'.

predictions for different substructures (see Appendix A), which could be related to phonotactic constraints, as in the example of pronunciation sequences, or any other prior constraint on illegal n-grams or other subsequences that would result in ruling out certain paths of graph G.

For convenience, in what follows, we will discuss our algorithms and solutions in the specific case of the graph G. However, the on-line learning algorithms we use apply in the same way to an arbitrary directed acyclic graph G'. The randomized algorithm we describe can also be used in a similar way and our batch learning guarantees for our randomized algorithm can be straightforwardly extended to an arbitrary graph G'. In fact, those guarantees are then somewhat more favorable since the number of path experts in G' will be smaller than in G.

3.2. On-line algorithm

Using G, the size of the pool of experts H we consider is p^l, and thus is exponentially large with respect to p. But, since learning guarantees in on-line learning admit only logarithmic dependence on that size, they remain informative in this context. However, the computational complexity of most on-line algorithms also directly depends on that size, which would make them impractical in this context. But, there exist several on-line solutions precisely designed to address this issue by exploiting the structure of the experts as in the case of our path experts. These include the algorithm of Takimoto & Warmuth (2003) denoted by WMWP, which is an extension of the (randomized) weighted-majority (WM) algorithm of Littlestone & Warmuth (1994) (see also (Vovk, 1990)) to more general bounded loss functions[1] combined with the directed graph Weight Pushing (WP) algorithm of Mohri (1997), and the Follow the Perturbed Leader (FPL) algorithm of Kalai & Vempala (2005).

The basis for the design of our batch algorithms is the WMWP algorithm since it admits a more favorable regret guarantee than the FPL algorithm in our context. However, we have also carried out a full analysis based on the FPL algorithm which can be found in Appendix D.

[1]The extension of the weighted majority algorithm to other losses is also known as the *Hedge algorithm* (Freund & Schapire, 1997) or the *exponentially weighted averaged algorithm* (Cesa-Bianchi & Lugosi, 2006).

Algorithm 1 WMWP algorithm.

Inputs: sample $\{(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T)\}$; set of experts $\{h_1, \ldots, h_p\}$; parameter $\beta \in (0, 1)$.
for $j = 1$ **to** p **and** $k = 1$ **to** l **do**
  $w_{1,kj} \leftarrow \frac{1}{p}$
end for
for $t = 1$ **to** T **and** $j = 1$ **to** p **and** $k = 1$ **to** l **do**
  $w_{t+1,kj} \leftarrow \frac{w_{t,kj}\beta^{\ell_k(h_j^k(\mathbf{x}_t),\mathbf{y}_t)}}{\sum_{j=1}^p w_{t,kj}\beta^{\ell_k(h_j^k(\mathbf{x}_t),\mathbf{y}_t)}}$
end for
Return matrices $\{\mathbf{W}_1, \ldots, \mathbf{W}_T\}$

As in the standard WM algorithm (Littlestone & Warmuth, 1994), WMWP maintains a distribution, p_t, $t \in [1, T]$ over the set of all experts, which in this context are the path experts $\mathsf{h} \in \mathsf{H}$. At each round $t \in [1, T]$, the algorithm receives an input sequence, $\mathbf{x}_t$, incurs the loss $\mathbb{E}_{\mathsf{h}\sim\mathsf{p}_t}[L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t)] = \sum_{\mathsf{h}} \mathsf{p}_t(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t)$ and multiplicatively updates the distribution weight per expert:

$$\forall \mathsf{h} \in \mathsf{H}, \quad \mathsf{p}_{t+1}(\mathsf{h}) = \frac{\mathsf{p}_t(\mathsf{h})\beta^{L(\mathsf{h}(\mathbf{x}_t),\mathbf{y}_t)}}{\sum_{\mathsf{h}'\in\mathsf{H}} \mathsf{p}_t(\mathsf{h}')\beta^{L(\mathsf{h}'(\mathbf{x}_t),\mathbf{y}_t)}}, \tag{2}$$

where $\beta \in (0, 1)$ is some fixed parameter. The number of paths is exponentially large in p and the cost of updating all paths is therefore prohibitive. However, since the loss function is additive in the substructures, the updates are multiplicative, and p_t can be compactly represented and updated by maintaining a potential value stored at each vertex (Takimoto & Warmuth, 2003). The cost of the update is then linear in the size of the graph.

The graph G we consider has a specific structure, thus, our description of the algorithm can be further simplified by maintaining at any round $t \in [1, T]$, an edge weight $w_{t,kj}$ for the jth edge, $j \in [1, p]$, between vertices $k - 1$ and k. This defines a matrix $\mathbf{W}_t = (w_{t,kj})_{kj} \in \mathbb{R}^{l\times p}$ with the following properties:

1. for any path expert h defined by $h_{j_1 1}, \ldots, h_{j_l l}$, $\mathsf{p}_t(\mathsf{h}) = \prod_{k=1}^l w_{t,kj_k}$;
2. the weights of outgoing edges sum to one at any vertex $k \in [0, l - 1]$: $\sum_{j=1}^p w_{t,kj} = 1$.

This clearly ensures that $\sum_{\mathsf{h}\in\mathsf{H}} \mathsf{p}_t(\mathsf{h}) = 1$ with the update rule (2) replaced by the following equivalent and more efficient edge weight update:

$$w_{t+1,kj} = \frac{w_{t,kj}\beta^{\ell_k(h_j^k(\mathbf{x}_t),\mathbf{y}_t)}}{\sum_{j=1}^p w_{t,kj}\beta^{\ell_k(h_j^k(\mathbf{x}_t),\mathbf{y}_t)}}. \tag{3}$$

Algorithm 1 gives the pseudocode of WMWP.

3.3. On-line-to-batch conversion

The WMWP algorithm does not produce a sequence of path experts, rather, it produces a sequence of distributions $\mathsf{p}_1, \ldots, \mathsf{p}_T$ over path experts, or equivalently a sequence of matrices $\mathbf{W}_1, \ldots, \mathbf{W}_T$. Thus, the on-line-to-batch conversion techniques described in (Littlestone, 1989; Cesa-Bianchi et al., 2004; Dekel & Singer, 2005) do not readily apply. Instead, we propose a generalization of the techniques of Dekel & Singer (2005). The conversion consists of two steps: first extract a good collection of distributions $\mathcal{P} \subseteq \{\mathsf{p}_1, \ldots, \mathsf{p}_T\}$; next use $\mathcal{P}$ to define an accurate hypothesis for prediction. For a subset $\mathcal{P} \subseteq \{\mathsf{p}_1, \ldots, \mathsf{p}_T\}$, we define

$$\begin{aligned}\Gamma(\mathcal{P}) &= \frac{1}{|\mathcal{P}|}\sum_{\mathsf{p}_t\in\mathcal{P}}\sum_{\mathsf{h}\in\mathsf{H}} \mathsf{p}_t(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t) + M\sqrt{\frac{\log\frac{1}{\delta}}{|\mathcal{P}|}} \\ &= \frac{1}{|\mathcal{P}|}\sum_{\mathsf{p}_t\in\mathcal{P}}\sum_{k=1}^{l}\sum_{j=1}^{p} w_{t,kj}\ell_k(h_j^k(\mathbf{x}_t), y_t^k) + M\sqrt{\frac{\log\frac{1}{\delta}}{|\mathcal{P}|}},\end{aligned} \tag{4}$$

where $\delta > 0$ is a fixed parameter. The second equality in (4) is a straightforward consequence of the identity $\mathsf{p}_t(\mathsf{h}) = \prod_{k=1}^l w_{t,kj_k}$ and the additive decomposition of L in terms of l_ks (see Lemma 6 in the appendix). With this definition, we choose $\mathcal{P}_\delta$ as a minimizer of $\Gamma(\mathcal{P})$ over some collection $\mathcal{P}$ of subsets of $\{\mathsf{p}_1, \ldots, \mathsf{p}_T\}$: $\mathcal{P}_\delta \in \operatorname{argmin}_{\mathcal{P}\in\mathcal{P}} \Gamma(\mathcal{P})$. The choice of $\mathcal{P}$ is restricted by computational considerations. One natural option is to let $\mathcal{P}$ be the union of the suffix sets $\{\mathsf{p}_t, \ldots, \mathsf{p}_T\}$, $t = 1, \ldots, T$. We will assume in what follows that $\mathcal{P}$ includes the set $\{\mathsf{p}_1, \ldots, \mathsf{p}_T\}$.

Next we define a randomized algorithm based on $\mathcal{P}_\delta$. Given an input $\mathbf{x}$, the algorithm consists of randomly selecting a path h according to

$$\mathsf{p}(\mathsf{h}) = \frac{1}{|\mathcal{P}_\delta|}\sum_{\mathsf{p}_t\in\mathcal{P}_\delta} \mathsf{p}_t(\mathsf{h}). \tag{5}$$

and returning the prediction $\mathsf{h}(\mathbf{x})$. Note that computing and storing p directly is not efficient. To sample from p, we first choose $\mathsf{p}_t \in \mathcal{P}_\delta$ uniformly at random and then sample a path h according to that p_t. Observe that for any fixed $k \in [1, l]$, $\sum_{j=1}^l w_{t,kj} = 1$, thus the non-negative weights $w_{t,kj}$ define a distribution over the edges leaving vertex k that we denote by $w_{t,k\cdot}$. Thus, to sample h from p_t we can simply draw an edge from each of the l distributions $w_{t,k\cdot}$ (the probability mass of a path is the product of the probability masses of its edges). Note that once an input $\mathbf{x}$ is received, the distribution p over the path experts h induces a probability distribution $\mathsf{p}_{\mathbf{x}}$ over the output space $\mathcal{Y}$. It is not hard to see that sampling a prediction $\mathbf{y}$ according to $\mathsf{p}_{\mathbf{x}}$ is statistically equivalent to first sampling h according to p and then predicting $\mathsf{h}(\mathbf{x})$. We will denote by $\mathcal{H}_{\text{Rand}}$ the randomized hypothesis thereby generated.

An inherent drawback of randomized solutions such as the one just described is that for the same input $\mathbf{x}$ the user can receive different predictions over time. Randomized solutions are also typically more costly to store. A collection of distributions $\mathcal{P}$ can, however, also be used to define a deterministic prediction rule based on the scoring function approach. The majority vote scoring function is defined by

$$\widetilde{\mathsf{h}}_{\text{MVote}}(\mathbf{x}, \mathbf{y}) = \prod_{k=1}^{l} \Big(\frac{1}{|\mathcal{P}_\delta|} \sum_{\mathsf{p}_t \in \mathcal{P}_\delta} \sum_{j=1}^{p} w_{t,kj} \mathbf{1}_{h_j^k(\mathbf{x}) = y^k} \Big). \quad (6)$$

The majority vote algorithm denoted by $\mathcal{H}_{\text{MVote}}$ is then defined by $\mathcal{H}_{\text{MVote}}(\mathbf{x}) = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} \widetilde{\mathsf{h}}_{\text{MVote}}(\mathbf{x}, \mathbf{y}), \mathbf{x} \in \mathcal{X}$. In the case of the graph G, the maximizer of $\widetilde{\mathsf{h}}_{\text{MVote}}$ is found efficiently by choosing $\mathbf{y}$ such that y^k has the maximum weight in position k.

In the next section, we present learning guarantees for $\mathcal{H}_{\text{Rand}}$ and $\mathcal{H}_{\text{MVote}}$. We also briefly discuss alternative prediction rules in Appendix E.

3.4. Batch learning guarantees

We first present learning bounds for the randomized prediction rule $\mathcal{H}_{\text{Rand}}$. Next, we upper bound the generalization error of $\mathcal{H}_{\text{MVote}}$ in terms of that of $\mathcal{H}_{\text{Rand}}$.

Proposition 1. *For any $\delta > 0$, with probability at least $1 - \delta$ over the choice of the sample $((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T))$ drawn i.i.d. according to $\mathcal{D}$, the following inequality holds:*

$$\mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})] \leq \frac{1}{T} \sum_{t=1}^{T} L_t + M \sqrt{\frac{\log \frac{T}{\delta}}{T}},$$

where $L_t = \mathbb{E}_{\mathsf{h} \sim \mathsf{p}_t}[L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t)]$.

Proof. Let $\mathcal{P} = \{\mathsf{p}_{t_1}, \ldots, \mathsf{p}_{t_{|\mathcal{P}|}}\}$. Observe that

$$\begin{aligned} &\mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})] - \frac{1}{|\mathcal{P}|} \sum_{s=1}^{|\mathcal{P}|} L_{t_s} \\ &= \sum_{s=1}^{|\mathcal{P}|} \sum_{\mathsf{h} \in \mathsf{H}} \frac{\mathsf{p}_{t_s}(\mathsf{h})}{|\mathcal{P}|} (\mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] - L(\mathsf{h}(\mathbf{x})_{t_s}, \mathbf{y}_{t_s})). \end{aligned}$$

We denote the inner summand by A_s and observe that A_s forms a martingale difference with respect to the filtration $\mathcal{G}_s = \mathcal{F}_{t_s}$ associated with the process $(\mathbf{x}_t, \mathbf{y}_t)$, i.e. $\mathcal{F}_t$ is a σ-algebra generated by $(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_t, \mathbf{y}_t)$. Indeed,

$$\begin{aligned} \mathbb{E}[A_s | \mathcal{G}_{s-1}] = \frac{1}{|\mathcal{P}|} \sum_{\mathsf{h}} &\mathbb{E}[\mathsf{p}_{t_s}(\mathsf{h})\, \mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] | \mathcal{G}_{s-1}] \\ &- \mathbb{E}[\mathsf{p}_{t_s}(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_{t_s}), \mathbf{y}_{t_s}) | \mathcal{G}_{s-1}]. \end{aligned}$$

Since p_t is determined by $\mathcal{F}_{t-1}$ and $(\mathbf{x}_t, \mathbf{y}_t)$ is independent of $\mathcal{F}_{t-1}$, we can write

$$\begin{aligned} &\mathbb{E}[\mathsf{p}_{t_s}(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_{t_s}), \mathbf{y}_{t_s}) | \mathcal{G}_{s-1}] \\ &\quad = \mathop{\mathbb{E}}_{1:t_s-1}[\mathop{\mathbb{E}}_{t_s}[\mathsf{p}_{t_s}(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_{t_s}), \mathbf{y}_{t_s})] | \mathcal{G}_{s-1}] \\ &\quad = \mathop{\mathbb{E}}_{1:t_s-1}[\mathsf{p}_{t_s}(\mathsf{h}) \mathop{\mathbb{E}}_{t_s}[L(\mathsf{h}(\mathbf{x}_{t_s}), \mathbf{y}_{t_s})] | \mathcal{G}_{s-1}] \end{aligned}$$

where $\mathbb{E}_{1:q}$ indicates that the expectation is taken with respect to $(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_q, \mathbf{y}_q)$. This shows that $\mathbb{E}[A_s | \mathcal{G}_{s-1}] = 0$, which implies that A_s is a martingale difference sequence. Since $|A_s| \leq M/|\mathcal{P}|$, it follows from Azuma's inequality that the probability of the event

$$\left\{ \mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})] - \frac{1}{|\mathcal{P}|} \sum_{s=1}^{|\mathcal{P}|} L_{t_s} > M \sqrt{\frac{\log \frac{1}{\delta}}{|\mathcal{P}|}} \right\}$$

is at most δ. Since $\mathcal{P}_\delta$ is a minimizer of $\frac{1}{|\mathcal{P}|} \sum_{s=1}^{|\mathcal{P}|} L_{t_s} + M \sqrt{\frac{\log \frac{1}{\delta}}{|\mathcal{P}|}}$ over $\mathcal{P}$ and $\mathcal{P}$ contains $\{\mathsf{p}_1, \ldots, \mathsf{p}_T\}$, the desired conclusion follows by the union bound. □

The next step consists of relating the expected loss of $\mathcal{H}_{\text{Rand}}$ to the regret R_T of the WMWP algorithm:

$$R_T = \sum_{t=1}^{T} \mathop{\mathbb{E}}_{\mathsf{h} \sim \mathsf{p}_t} [L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t)] - \inf_{\mathsf{h} \in \mathsf{H}} \sum_{t=1}^{T} L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t). \quad (7)$$

Theorem 2. *For any $\delta > 0$, with probability at least $1 - \delta$ over the choice of the sample $((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T))$ drawn i.i.d. according to $\mathcal{D}$, the following inequalities hold:*

$$\mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})] \leq \inf_{\mathsf{h} \in \mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] + \frac{R_T}{T} + 2M \sqrt{\frac{\log \frac{2T}{\delta}}{T}}$$

$$\mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})] \leq \inf_{\mathsf{h} \in \mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] + 2M \sqrt{\frac{l \log p}{T}} + 2M \sqrt{\frac{\log \frac{2T}{\delta}}{T}}.$$

See Appendix B for a proof of this result. We now upper bound the generalization error of the majority-vote algorithm $\mathcal{H}_{\text{MVote}}$ in terms of that of the randomized algorithm $\mathcal{H}_{\text{Rand}}$, which, combined with Theorem 2, immediately yields generalization bounds for the majority-vote algorithm $\mathcal{H}_{\text{MVote}}$. The first proposition, which admits a simple proof, relates the expected loss of the majority vote algorithm to that of a randomized algorithm in the case of the normalized Hamming loss.

Proposition 3. *The following inequality relates the generalization error of the majority-vote algorithm to that of the randomized one:*

$$\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{MVote}}(\mathbf{x}), \mathbf{y})] \leq 2\, \mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})],$$

where the expectations are taken over $(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}$ and $\mathsf{h} \sim \mathsf{p}$.

Proof. By definition of the majority vote, if $\mathcal{H}_{\text{MVote}}$ makes an error at position k on example $(\mathbf{x}, \mathbf{y})$, then, the total weight of incorrect labels at that position must be at least half of the total weight of labels in that position. In other words, the following inequality holds for any k:

$$\mathbf{1}_{\mathcal{H}^k_{\text{MVote}}(\mathbf{x}) \neq y^k} \leq 2 \frac{1}{|\mathcal{P}_\delta|} \sum_{\mathsf{p}_t \in \mathcal{P}_\delta} \sum_{j=1}^{p} w_{t,kj} \mathbf{1}_{h^k_j(\mathbf{x}) \neq y^k}.$$

Summing up these inequalities over k and taking expectations yields the desired bound. □

Proposition 3 suggests that the price to pay for derandomization is a factor of 2. However, this may be too pessimistic. A more refined result presented in the following proposition shows that often this price is lower.

Proposition 4. *The following bound holds for any distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$:*

$$\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{MVote}}(\mathbf{x}), \mathbf{y})] \leq 2\,\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})] - 2\,\mathbb{E}[\gamma(\mathbf{x}, \mathbf{y})],$$

where $\gamma(\mathbf{x}, \mathbf{y}) = \sum_{k=1}^{l} \gamma_k(\mathbf{x}, \mathbf{y})$ *with* $\gamma_k(\mathbf{x}, \mathbf{y}) = \max\left(0, \frac{1}{|\mathcal{P}_\delta|} \sum_{\mathsf{p}_t \in \mathcal{P}_\delta} \sum_{j=1}^{p} w_{t,kj} \mathbf{1}_{h^k_j(\mathbf{x}) \neq y^k} - \frac{1}{2}\right)$.

The proof is a refinement of the proof of Proposition 3 and can be found in Appendix B. Each γ_k in Proposition 4 can be interpreted as the edge of incorrect labels and this result implies that any additional edge of an incorrect hypothesis (beyond $\frac{1}{2}$) should not be included in the bound.

Our methods generalize the results of Dekel & Singer (2005) where $l = 1$ and each p_t is a probability point mass at a hypothesis h_t produced by an on-line algorithm on the tth iteration. It is also possible to extend the cross-validation approach of Cesa-Bianchi et al. (2004) to our setting, but the learning guarantees for this algorithm end up being less favorable than those just given (see Appendix C for a full description and analysis). Our results and algorithms can be extended to the case of other directed acyclic graphs of path experts and other derandomization methods (see Appendix E for a more detailed discussion).

4. Boosting approach

In this section, we devise a boosting-style algorithm for our ensemble structured prediction problem. The variants of AdaBoost for multi-class classification such as AdaBoost.MH or AdaBoost.MR (Freund & Schapire, 1997; Schapire & Singer, 1999; 2000) cannot be readily applied in this context. First, the number of classes to consider here is quite large, as in all structured prediction problems, since it is exponential in the number of substructures l. For example, in the case of the pronunciation problem where the number of phonemes for English is in the order of 50, the number of classes is 50^l. But, the objective function for AdaBoost.MH or AdaBoost.MR as well as the main steps of the algorithms include a sum over all possible labels, whose computational cost in this context would be prohibitive. Second, the loss function we consider is the normalized Hamming loss over the substructures predictions, which does not match the multi-class losses for the variants of AdaBoost.[2] Finally, the natural base hypotheses for this problem admit a structure that can be exploited to devise a more efficient solution, which of course was not part of the original considerations for the design of these variants of AdaBoost.

4.1. Hypothesis sets and loss function

The predictor $\mathcal{H}_{\text{ESPBoost}}$ returned by our boosting algorithm is based on a scoring function $\widetilde{\mathsf{h}}\colon \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, which, as for standard ensemble algorithms such as AdaBoost, is a convex combination of base scoring functions $\widetilde{\mathsf{h}}_t$: $\widetilde{\mathsf{h}} = \sum_{t=1}^{T} \alpha_t \widetilde{\mathsf{h}}_t$, with $\alpha_t \geq 0$. The base scoring functions we consider for our problem are derived from the path experts in H. For each path expert $\mathsf{h}_t \in \mathsf{H}$, we define a scoring function $\widetilde{\mathsf{h}}_t$ as follows:

$$\forall (\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}, \quad \widetilde{\mathsf{h}}_t(\mathbf{x}, \mathbf{y}) = \sum_{k=1}^{l} \mathbf{1}_{\mathsf{h}^k_t(\mathbf{x}) = y^k}. \tag{8}$$

Thus, the score assigned to $\mathbf{y}$ by the base scoring function $\widetilde{\mathsf{h}}_t$ is the number of positions at which $\mathbf{y}$ matches the prediction of path expert h_t given input $\mathbf{x}$. $\mathcal{H}_{\text{ESPBoost}}$ is defined as follows in terms of $\widetilde{\mathsf{h}}$ or h_ts:

$$\begin{aligned} \forall \mathbf{x} \in \mathcal{X},\ \mathcal{H}_{\text{ESPBoost}}(\mathbf{x}) &= \operatorname*{argmax}_{\mathbf{y} \in \mathcal{Y}} \widetilde{\mathsf{h}}(\mathbf{x}, \mathbf{y}) \\ &= \operatorname*{argmax}_{\mathbf{y} \in \mathcal{Y}} \sum_{k=1}^{l} \sum_{t=1}^{T} \alpha_t \mathbf{1}_{\mathsf{h}^k_t(\mathbf{x}) = y^k}. \end{aligned} \tag{9}$$

4.2. ESPBoost algorithm

For any $i \in [1, m]$ and $k \in [1, l]$, we define the *margin of* $\widetilde{\mathsf{h}}^k$ *for point* $(\mathbf{x}_i, \mathbf{y}_i)$ by $\rho(\widetilde{\mathsf{h}}^k, \mathbf{x}_i, \mathbf{y}_i) = \widetilde{\mathsf{h}}^k(\mathbf{x}_i, y^k_i) - \max_{y^k \neq y^k_i} \widetilde{\mathsf{h}}^k(\mathbf{x}_i, y^k)$.

Lemma 5. *The following upper bound holds for the empirical normalized Hamming loss of the hypothesis $\mathcal{H}_{\text{ESPBoost}}$:*

$$\mathop{\mathbb{E}}_{(\mathbf{x}, \mathbf{y}) \sim S}[L_{\text{Ham}}(\mathcal{H}_{\text{ESPBoost}}(\mathbf{x}), \mathbf{y})] \leq \frac{1}{ml} \sum_{i=1}^{m} \sum_{k=1}^{l} \exp\Big(-\sum_{t=1}^{T} \alpha_t \rho(\widetilde{\mathsf{h}}^k_t, \mathbf{x}_i, \mathbf{y}_i)\Big).$$

[2] Schapire & Singer (1999) also present an algorithm using the Hamming loss for multi-class classification, but that is a Hamming loss over the set of classes and differs from the loss function relevant to our problem. Additionally, the main steps of that algorithm are also based on a sum over all classes.

Algorithm 2 ESPBoost Algorithm.

Inputs: $S = ((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_m, \mathbf{y}_m))$; set of experts $\{h_1, \ldots, h_p\}$.
for $i = 1$ **to** m **and** $k = 1$ **to** l **do**
  $\mathcal{D}_1(i,k) \leftarrow \frac{1}{ml}$
end for
for $t = 1$ **to** T **do**
  $\mathsf{h}_t \leftarrow \operatorname{argmin}_{\mathsf{h}\in\mathsf{H}} \mathbb{E}_{(i,k)\sim\mathcal{D}_t}[\mathbf{1}_{\mathsf{h}^k(\mathbf{x}_i)\neq y_i^k}]$
  $\epsilon_t \leftarrow \mathbb{E}_{(i,k)\sim\mathcal{D}_t}[\mathbf{1}_{\mathsf{h}_t^k(\mathbf{x}_i)\neq y_i^k}]$
  $\alpha_t \leftarrow \frac{1}{2}\log\frac{1-\epsilon_t}{\epsilon_t}$
  $Z_t \leftarrow 2\sqrt{\epsilon_t(1-\epsilon_t)}$
  for $i = 1$ **to** m **and** $k = 1$ **to** l **do**
    $\mathcal{D}_{t+1}(i,k) \leftarrow \frac{\exp(-\alpha_t\rho(\widetilde{\mathsf{h}}_t^k, \mathbf{x}_i, \mathbf{y}_i))\mathcal{D}_t(i,k)}{Z_t}$
  end for
end for
Return $\widetilde{\mathsf{h}} = \sum_{t=1}^{T}\alpha_t\widetilde{\mathsf{h}}_t$

In view of this upper bound, we consider the objective function $F\colon \mathbb{R}^N \to \mathbb{R}$ defined for all $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_N) \in \mathbb{R}^N$ by

$$F(\boldsymbol{\alpha}) = \frac{1}{ml}\sum_{i=1}^{m}\sum_{k=1}^{l}\exp\Big(-\sum_{j=1}^{N}\alpha_j\rho(\widetilde{\mathsf{h}}_j^k, \mathbf{x}_i, \mathbf{y}_i)\Big), \quad (10)$$

where $\mathsf{h}_1, \ldots, \mathsf{h}_N$ denote the set of all path experts in H. F is a convex and differentiable function of $\boldsymbol{\alpha}$. Our algorithm, ESPBoost (Ensemble Structured Prediction Boosting), is defined by the application of coordinate descent to the objective F. Algorithm 2 shows the pseudocode of the ESPBoost (see Appendix G.2 for the details of the derivation of the coordinate descent algorithm).

Our *weak learning assumption* in this context is that there exists $\gamma > 0$ such that at each round, ϵ_t verifies $\epsilon_t < \frac{1}{2} - \gamma$. For the graph G, at each round, the path expert h_t with the smallest error ϵ_t can be determined easily and efficiently by first finding for each substructure k, the h_t^k that is the best with respect to the distribution weights $\mathcal{D}_t(i,k)$.

Observe that, while the steps of our algorithm are syntactically close to those of AdaBoost and its multi-class variants, our algorithm is distinct and does not require sums over the exponential number of all possible labelings of the substructures and is quite efficient. We have derived margin-based learning guarantees for ESPBoost which are presented in detail and proven in Appendix G.3.

5. Experiments

We used a number of artificial and real-life data sets for our experiments. For each data set, we performed 10-fold cross-validation with disjoint training sets.[3] We report the average test error for each task. In addition to the $\mathcal{H}_{\text{MVote}}$, $\mathcal{H}_{\text{Rand}}$ and $\mathcal{H}_{\text{ESPBoost}}$ hypotheses, we experimented with two algorithms discussed in more detail in the appendix: a cross-validation on-line-to-batch conversion of the WMWP algorithm, $\mathcal{H}_{\text{CV}}$, and a majority-vote on-line-to-batch conversion with FPL, $\mathcal{H}_{\text{FPL}}$, and a cross-validation on-line-to-batch conversion with FPL, $\mathcal{H}_{\text{FPL-CV}}$. Finally, we compare with the $\mathcal{H}_{\text{SLE}}$ algorithm of Nguyen & Guo (2007).

[3] For the OCR data set, these subsets are predefined.

Table 1. Average Normalized Hamming Loss, ADS1 and ADS2. $\beta_{ADS1} = 0.95$, $\beta_{ADS2} = 0.95$, $T_{SLE} = 100$, $\delta = 0.05$.

	ADS1, $m = 200$	ADS2, $m = 200$
$\mathcal{H}_{\text{MVote}}$	**0.0197 ± 0.00002**	**0.2172 ± 0.00983**
$\mathcal{H}_{\text{FPL}}$	0.0228 ± 0.00947	0.2517 ± 0.05322
$\mathcal{H}_{\text{CV}}$	**0.0197 ± 0.00002**	0.2385 ± 0.00002
$\mathcal{H}_{\text{FPL-CV}}$	0.0741 ± 0.04087	0.4001 ± 0.00028
$\mathcal{H}_{\text{ESPBoost}}$	**0.0197 ± 0.00002**	**0.2267 ± 0.00834**
$\mathcal{H}_{\text{SLE}}$	0.5641 ± 0.00044	0.2500 ± 0.05003
$\mathcal{H}_{\text{Rand}}$	0.1112 ± 0.00540	0.4000 ± 0.00018
Best h_j	0.5635 ± 0.00004	0.4000

Table 2. Average Normalized Hamming Loss, PDS1 and PDS2. $\beta_{PDS1} = 0.85$, $\beta_{PDS2} = 0.97$, $T_{SLE} = 100$, $\delta = 0.05$.

	PDS1, $m = 130$	PDS2, $m = 400$
$\mathcal{H}_{\text{MVote}}$	**0.2225 ± 0.00301**	**0.2323 ± 0.00069**
$\mathcal{H}_{\text{FPL}}$	0.2657 ± 0.07947	0.2337 ± 0.00229
$\mathcal{H}_{\text{CV}}$	**0.2316 ± 0.00189**	**0.2364 ± 0.00080**
$\mathcal{H}_{\text{FPL-CV}}$	0.4451 ± 0.02743	0.4090 ± 0.01388
$\mathcal{H}_{\text{ESPBoost}}$	0.3625 ± 0.01054	0.3499 ± 0.00509
$\mathcal{H}_{\text{SLE}}$	0.3130 ± 0.05137	0.3308 ± 0.03182
$\mathcal{H}_{\text{Rand}}$	0.4713 ± 0.00360	0.4607 ± 0.00131
Best h_j	0.3449 ± 0.00368	0.3413 ± 0.00067

5.1. Artificial data sets

Our artificial data set, ADS1 and ADS2 simulate the scenarios described in Section 1. In ADS1 the kth expert has a high accuracy on the kth position, in ADS2 an expert has low accuracy in a fixed set of positions. More details on the data set and the experimental parameters can be found in Appendix H.1.

Table 1 reports the results of our experiments. In both cases $\mathcal{H}_{\text{MVote}}$, our majority-vote algorithm based on our on-line-to-batch conversion using the WMWP algorithm (together with most of the other on-line based algorithms), yields a significant improvement over the best expert. It also outperforms $\mathcal{H}_{\text{SLE}}$, which in the case of ADS1 even fails to outperform the best h_j. After 100 iterations on ADS1, the ensemble learned by $\mathcal{H}_{\text{SLE}}$ consists of a single expert, which is why it leads to such a poor performance.

It is also worth pointing out that $\mathcal{H}_{\text{FPL-CV}}$ and $\mathcal{H}_{\text{Rand}}$ fail to outperform the best model on ADS2 set. This is in total agreement with our theoretical analysis since, in this case, any path expert has exactly the same performance and the error of the best path expert is an asymptotic upper bound on the errors of these algorithms.

5.2. Pronunciation data sets

We had access to two proprietary pronunciation data sets, PDS1 and PDS2. In both sets each example is an English word, typically a proper name. For each word, 20 possible phonemic sequences are available, ranked by some pronunciation model. Since the true pronunciation was not available, we set the top sequence to be the target label and used the remaining as the predictions made by the experts. The only difference between PDS1 and PDS2 is their size: 1,313 words for PDS1 and 6,354 for PDS2.

In both cases on-line based algorithms, specifically $\mathcal{H}_{\text{MVote}}$, significantly outperformed the best model as well as $\mathcal{H}_{\text{SLE}}$, see Table 2. The poor performance of $\mathcal{H}_{\text{ESPBoost}}$ is due to the fact that the weak learning assumption is violated after 5-8 iterations and hence the algorithm terminates.

5.3. OCR data set

Rob Kassel's OCR data set is available for download from http://ai.stanford.edu/~btaskar/ocr/. It contains 6,877 word instances with a total of 52,152 characters. Each character is represented by $16 \times 8 = 128$ binary pixels. The task is to predict a word given its sequence of pixel vectors. To generate experts we used several software packages: CRFsuite (Okazaki, 2007) and SVM^{struct}, SVM^{multiclass} (Joachims, 2008), and the Stanford Classifier (Rafferty et al., 2014). We trained these algorithms on each of the predefined folds of the data set and used the resulting models to generate expert predictions.

The results reported in Table 7 in Appendix H show that ensemble methods lead only to a small improvement in performance over the best h_j. This is because the best model h_j dominates all other experts and ensemble methods cannot benefit from patching together different outputs.

5.4. Penn Treebank data set

The part-of-speech task (POS) consists of labeling each word of a sentence with its correct part-of-speech tag. The Penn Treebank 2 data set is available through LDC license at http://www.cis.upenn.edu/~treebank/ and contains 251,854 sentences with a total of 6,080,493 tokens and 45 different parts of speech.

For the first experiment (TR1), we used 4 disjoint training sets to produce 4 SVM^{multiclass} models and 4 maximum entropy models using the Stanford Classifier. We also used the union of these training sets to devise one CRFsuite model. For the second experiment (TR2) we trained 5 SVM^{struct} models. The same features were used for both experiments. For the SVM algorithms, we generated 267,214 bag-of-word binary features. The Stanford Classifier and CRFsuite packages use internal routines to generate their features. For more detail, see Appendix H.

Table 3. Average Normalized Hamming Loss, TR1 and TR2. $\beta_{TR1} = 0.95$, $\beta_{TR2} = 0.98$, $T_{SLE} = 100$, $\delta = 0.05$.

	TR1, $m = 800$	TR2, $m = 1000$
$\mathcal{H}_{\text{MVote}}$	**0.0850 ± 0.00096**	**0.0746 ± 0.00014**
$\mathcal{H}_{\text{FPL}}$	**0.0859 ± 0.00110**	**0.0769 ± 0.00218**
$\mathcal{H}_{\text{CV}}$	**0.0843 ± 0.00006**	**0.0741 ± 0.00011**
$\mathcal{H}_{\text{FPL-CV}}$	0.1093 ± 0.00129	0.1550 ± 0.00182
$\mathcal{H}_{\text{ESPBoost}}$	0.1041 ± 0.00056	0.1414 ± 0.00233
$\mathcal{H}_{\text{SLE}}$	**0.0778 ± 0.00934**	**0.0814 ± 0.02558**
$\mathcal{H}_{\text{Rand}}$	0.1128 ± 0.00048	0.1652 ± 0.00077
Best h_j	0.1032 ± 0.00007	0.1415 ± 0.00005

The results of the experiments are summarized in Table 3. For TR1, our on-line ensemble methods improve over the best model. Note that $\mathcal{H}_{\text{SLE}}$ has the best average loss over 10 runs for this experiment. This comes at a price of much higher standard deviation which does not allow us to conclude that the difference in performance between our methods and $\mathcal{H}_{\text{SLE}}$ is statistically significant. In fact, on two runs $\mathcal{H}_{\text{SLE}}$ chooses an ensemble consisting of a single expert and fails to outperform the best model.

6. Conclusion

We presented a broad analysis of the problem of ensemble structured prediction, including a series of algorithms with learning guarantees and extensive experiments. Our results show that our algorithms, most notably $\mathcal{H}_{\text{MVote}}$, can result in significant benefits in several tasks, which can be of a critical practical importance. In Appendix H, we also report very favorable results for $\mathcal{H}_{\text{MVote}}$ when used with the edit-distance, which is the natural measure in many applications. A natural extension of this work consists of devising new algorithms and providing learning guarantees specific to other loss functions such as the edit-distance.

The extension of our algorithms and solutions to other directed graphs, as discussed in Appendix E, can further increase the applicability of our methods and enhance performance. While we aimed for an exhaustive study including multiple on-learning algorithms, different conversions to batch and derandomizations, we are aware that the problem we studied is very rich and admits many more facets and scenarios that we plan to investigate in the future.

Acknowledgments

We warmly thank our colleagues Francoise Beaufays and Fuchun Peng for kindly extracting and making available to us the pronunciation data sets used for our experiments and Richard Sproat and Brian Roark for help with other data sets. This work was partly funded by the NSF award IIS-1117591 and the NSERC PGS D3 award.

References

Breiman, L. Bagging predictors. *Machine Learning*, 24:123–140, 1996.

Caruana, R., Niculescu-Mizil, A., Crew, G., and Ksikes, A. Ensemble selection from libraries of models. In *Proceedings of ICML*, pp. 18–, 2004.

Cesa-Bianchi, N. and Lugosi, G. *Prediction, learning, and games*. Cambridge University Press, 2006.

Cesa-Bianchi, N., Conconi, A., and Gentile, C. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.

Collins, M. Discriminative training methods for hidden Markov models: theory and experiments with perceptron algorithms. In *Proceedings of ACL*, pp. 1–8, 2002.

Collins, M. and Koo, T. Discriminative reranking for natural language parsing. *Computational Linguistics*, 31(1):25–70, 2005.

Cortes, C., Mohri, M., and Weston, J. A general regression technique for learning transductions. In *Proceedings of ICML 2005*, pp. 153–160, 2005.

Dekel, O. and Singer, Y. Data-driven online to batch conversion. In *Advances in NIPS 18*, pp. 1207–1216, 2005.

Fiscus, J. G. Post-processing system to yield reduced word error rates: Recognizer output voting error reduction (ROVER). In *Proceedings of ASRU*, pp. 347–354, 1997.

Freund, Y. and Schapire, R. A decision-theoretic generalization of on-line learning and application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.

Freund, Y., Mansour, Y., and Schapire, R. Generalization bounds for averaged classifiers. *Ann. Stat.*, 32:1698–1722, 2004.

Ghoshal, A., Jansche, M., Khudanpur, S., Riley, M., and Ulinski, M. Web-derived pronunciations. In *Proceedings of ICASSP*, pp. 4289–4292, 2009.

Joachims, T. Support vector machines for complex outputs, 2008. URL http://www.cs.cornell.edu/people/tj/svm_light/svm_struct.html.

Kalai, A. and Vempala, S. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3): 291–307, 2005.

Kocev, D., Vens, C., Struyf, J., and Deroski, S. Tree ensembles for predicting structured outputs. *Pattern Recognition*, 46(3): 817–833, March 2013.

Koltchinskii, V. and Panchenko, D. Empirical margin distributions and bounding the generalization error of combined classifiers. *Ann. Stat.*, 30:1–50, 2002.

Lafferty, J., McCallum, A., and Pereira, F. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML*, pp. 282–289, 2001.

Littlestone, N. From on-line to batch learning. In *Proceedings of COLT 2*, pp. 269–284, 1989.

Littlestone, N. and Warmuth, M. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994.

MacKay, D. J. C. *Bayesian methods for adaptive models*. PhD thesis, California Institute of Technology, 1991.

Mohri, M. Finite-state transducers in language and speech processing. *Computational Linguistics*, 23(2):269–311, 1997.

Mohri, M. and Riley, M. An efficient algorithm for the n-best-strings problem. In *Interspeech*, 2002.

Mohri, M., Pereira, F. C. N., and Riley, M. Speech recognition with weighted finite-state transducers. In *Handbook on Speech Processing and Speech Communication, Part E: Speech recognition*. Springer-Verlag, 2008.

Mohri, M., Rostamizadeh, A., and Talwalkar, A. *Foundations of Machine Learning*. The MIT Press, 2012.

Nguyen, N. and Guo, Y. Comparison of sequence labeling algorithms and extensions. In *Proc. of ICML*, pp. 681–688, 2007.

Okazaki, N. CRFsuite: a fast implementation of conditional random fields (crfs), 2007. URL http://www.chokkan.org/software/crfsuite/.

Petrov, S. Products of random latent variable grammars. In *HLT-NAACL*, pp. 19–27, 2010.

Rafferty, A., Kleeman, A., Finkel, J., and Manning, C. Stanford classifer, 2014. URL http://nlp.stanford.edu/downloads/classifier.shtml.

Sagae, K. and Lavie, A. Parser combination by reparsing. In *Proceedings of HLT/NAACL*, pp. 129–132, 2006.

Schapire, R. and Singer, Y. Improved boosting algorithms using confidence-rated predictions. *Machine Learning*, 37(3):297–336, 1999.

Schapire, R. and Singer, Y. Boostexter: A boosting-based system for text categorization. *Machine Learning*, 39:135–168, 2000.

Schapire, R., Freund, Y., Bartlett, P., and Lee, W. S. Boosting the margin: A new explanation for the effectiveness of voting methods. In *ICML*, pp. 322–330, 1997.

Smyth, P. and Wolpert, D. Linearly combining density estimators via stacking. *Machine Learning*, 36:59–83, 1999.

Takimoto, E. and Warmuth, M. K. Path kernels and multiplicative updates. *JMLR*, 4:773–818, 2003.

Taskar, B., Guestrin, C., and Koller, D. Max-margin Markov networks. In *Advances in NIPS 16*. MIT Press, 2004.

Tsochantaridis, I., Joachims, T., Hofmann, T., and Altun, Y. Large margin methods for structured and interdependent output variables. *JMLR*, 6:1453–1484, December 2005.

Vovk, V. G. Aggregating strategies. In *COLT*, pp. 371–386, 1990.

Wang, Q., Lin, D., and Schuurmans, D. Simple training of dependency parsers via structured boosting. In *Proceedings of IJCAI 20*, pp. 1756–1762, 2007.

Zeman, D. and Žabokrtský, Z. Improving parsing accuracy by combining diverse dependency parsers. In *Proceedings of IWPT 9*, pp. 171–178, 2005.

Zhang, H., Zhang, M., Tan, C., and Li, H. K-best combination of syntactic parsers. In *Proceedings of EMNLP: Volume 3*, pp. 1552–1560, 2009.

A. General path expert graphs

For simplicity, we presented our algorithms and guarantees in the case of the graph G admitting all path experts. In most cases in practice, different acyclic graph of experts such as that of Figure 2 must be considered. This occurs in particular because of the presence of known constraints restricting the set of admissible sequences of substructures.

For example, the learning problem may consist of predicting the pronunciation associated to each sequence of words. In that case, for most languages, there exist phonotactic rules making some phonemic sequences inadmissible. Similarly, in parsing or translation tasks, some word sequences can be ruled out because they do not conform to some clear syntactic or stylistic rule. Let A denote a finite automaton accepting admissible sequences of $\mathcal{Y}_1 \times \cdots \times \mathcal{Y}_l$ and let G_t denote the graph of path experts considered at round $t \in [1, T]$, with $G_1 = G$. At each round $t \in [1, T]$, the learner receives a new input sequence $\mathbf{x}_t$ that is used to derive a finite automaton $\widetilde{G}_t$ from G_t by replacing in G_t the substructure predictor h_j^k, $j \in [1, p]$, $k \in [1, l]$, by its prediction $h_j^k(\mathbf{x}_t)$. Since some sequences of $\widetilde{G}_t$ may not be admissible, we must remove from G_t path experts generating sequences not in $\widetilde{G}_t \cap A$. This can be achieved straightforwardly using the intersection algorithm for finite automata if we keep track, for each substructure predicted, of the original substructure expert generating that prediction. G_{t+1} is the resulting graph of admissible path experts.

The on-line learning algorithm we consider, WMWP, applies to an arbitrary directed acyclic graph and thus can be applied to graph G_t at each round t. The distribution over the path experts maintained by WMWP effectively assigns probability zero to the path experts not present in G_t at round t. Our learning guarantees hold for this more general setting and in fact end up being more favorable since the cardinality of the set of admissible path experts is smaller than that of graph G.

B. On-line-to-batch conversion

Lemma 6. *For any $t \in [1, T]$, the following identity holds:*

$$\sum_{\mathsf{h} \in \mathsf{H}} \mathsf{p}_t(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t)) = \sum_{k=1}^{l} \sum_{j=1}^{p} w_{t,kj} \ell_k(\mathsf{h}^k(\mathbf{x}_t), y_t^k).$$

Proof. Recall that for any $t \in [1, T]$ and $k \in [1, l]$, $\sum_{j=1}^{p} w_{t,kj} = 1$. Thus, let $w_{t,k\cdot}$ denote the distribution defined by the non-negative weights $w_{t,kj}$. Then, the following chain of equalities proves the result:

$$\begin{aligned}
\sum_{\mathsf{h} \in \mathsf{H}} \mathsf{p}_t(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t)) &= \mathop{\mathbb{E}}_{\mathsf{h} \sim \mathsf{p}_t} [L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t)] \\
&= \mathop{\mathbb{E}}_{\mathsf{h} \sim \mathsf{p}_t} [\sum_{k=1}^{l} \ell_k(\mathsf{h}^k(\mathbf{x}_t), y_t^k)] \\
&= \sum_{k=1}^{l} \mathop{\mathbb{E}}_{\mathsf{h} \sim \mathsf{p}_t} [\ell_k(\mathsf{h}^k(\mathbf{x}_t), y_t^k)] \\
&= \sum_{k=1}^{l} \mathop{\mathbb{E}}_{\substack{\mathsf{h}^1 \sim w_{t,1\cdot} \\ \cdots \\ \mathsf{h}^l \sim w_{t,l\cdot}}} [\ell_k(\mathsf{h}^k(\mathbf{x}_t), y_t^k)] \\
&= \sum_{k=1}^{l} \mathop{\mathbb{E}}_{\mathsf{h}^k \sim w_{t,k\cdot}} [\ell_k(\mathsf{h}^k(\mathbf{x}_t), y_t^k)] \\
&= \sum_{k=1}^{l} \sum_{j=1}^{p} w_{t,kj} \ell_k(\mathsf{h}^k(\mathbf{x}_t), y_t^k).
\end{aligned}$$

□

Proposition 4. *The following bound holds for any distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$:*

$$\begin{aligned}
\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{MVote}}(\mathbf{x}), \mathbf{y})] \leq 2\, \mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})] \\
- 2\, \mathbb{E}[\gamma(\mathbf{x}, \mathbf{y})],
\end{aligned}$$

where $\gamma(\mathbf{x}, \mathbf{y}) = \sum_{k=1}^{l} \gamma_k(\mathbf{x}, \mathbf{y})$ with

$$\gamma_k(\mathbf{x}, \mathbf{y}) = \max\left(0, \frac{1}{|\mathcal{P}_\delta|} \sum_{\mathsf{p}_t \in \mathcal{P}_\delta} \sum_{j=1}^{p} w_{t,kj} \mathbf{1}_{h_j^k(\mathbf{x}) \neq y^k} - \frac{1}{2}\right).$$

Proof. The proof is a slight refinement of that of Proposition 3. If $\mathcal{H}_{\text{MVote}}$ makes an error at position k on example $(\mathbf{x}, \mathbf{y})$ then the total weight of incorrect labels at that position must be $\frac{1}{2} + \gamma_k(\mathbf{x}, \mathbf{y})$. In other words, we have the following inequality

$$\mathbf{1}_{\mathcal{H}_{\text{MVote}}^k(\mathbf{x}) \neq y^k} \leq \frac{2}{|\mathcal{P}_\delta|} \sum_{p \in \mathcal{P}_\delta} \sum_{j=1}^{p} w_{t,kj} \mathbf{1}_{h_j^k(\mathbf{x}) \neq y^k} - 2\gamma_k(\mathbf{x}, \mathbf{y})$$

when $\mathbf{1}_{\mathcal{H}_{\text{MVote}}^k(\mathbf{x}) \neq y^k} = 1$. Since the right-hand side of the bound above is always positive, it also holds when $\mathbf{1}_{\mathcal{H}_{\text{MVote}}^k(\mathbf{x}) \neq y^k} = 0$. The rest of the proof is the same as that of Proposition 3. □

Theorem 2. *For any $\delta > 0$, with probability at least $1 - \delta$ over the choice of the sample $((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T))$ drawn i.i.d. according to $\mathcal{D}$, the following inequalities*

hold:

$$\mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}),\mathbf{y})] \leq \inf_{\mathsf{h}\in\mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}),\mathbf{y})] + \frac{R_T}{T} + 2M\sqrt{\frac{\log\frac{2T}{\delta}}{T}}$$

$$\mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}),\mathbf{y})] \leq \inf_{\mathsf{h}\in\mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}),\mathbf{y})] + 2M\sqrt{\frac{l\log p}{T}} + 2M\sqrt{\frac{\log\frac{2T}{\delta}}{T}}.$$

Proof. Since there are only finitely many expert paths h, there is an expert path $\mathsf{h}^* \in \mathsf{H}$ such that $\inf_{\mathsf{h}\in\mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}),\mathbf{y})] = \mathbb{E}[L(\mathsf{h}^*(\mathbf{x}),\mathbf{y})]$. By Hoeffding's inequality, the probability of the event

$$\left\{\frac{1}{T}\sum_{t=1}^{T} L(\mathsf{h}^*(\mathbf{x}_t),\mathbf{y}_t) - \mathbb{E}[L(\mathsf{h}^*(\mathbf{x}),\mathbf{y})] > M\sqrt{\frac{\log\frac{2}{\delta}}{T}}\right\}$$

is at most $\delta/2$. Therefore, by Proposition 1 and the union bound, the following holds with probability at least $1-\delta$:

$$\begin{aligned}
&\mathbb{E}[L(\mathcal{H}_{\text{Rand}}(\mathbf{x}),\mathbf{y})] - \inf_{\mathsf{h}\in\mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}),\mathbf{y})] \\
&\leq \frac{1}{T}\sum_{t=1}^{T} L_t + M\sqrt{\frac{\log\frac{2T}{\delta}}{T}} - \frac{1}{T}\sum_{t=1}^{T} L(\mathsf{h}^*(\mathbf{x}_t),\mathbf{y}_t) + M\sqrt{\frac{\log\frac{2}{\delta}}{T}} \\
&\leq \frac{R_T}{T} + 2M\sqrt{\frac{\log\frac{2T}{\delta}}{T}},
\end{aligned}$$

which proves the first inequality. The regret of the randomized MW algorithm for losses taking values in $[0,1]$ is known to be bounded by $2\sqrt{T\log N}$ where N is the number of experts (Cesa-Bianchi & Lugosi, 2006). In our context, this gives $R_T \leq 2M\sqrt{T\log(p^l)}$. Plugging in this bound in the first inequality of the theorem yields directly the second one. □

C. Cross-validation based on-line-to-batch conversion

Cesa-Bianchi et al. (2004) described an on-line-to-batch conversion technique based on a cross-validation approach. Given a sequence of hypotheses produced by an on-line algorithm, a single hypothesis is selected based on its empirical loss on unseen examples plus a special penalty term. These results can be easily generalized to the case where an on-line algorithm produces distributions over hypotheses rather than just a single hypothesis. More precisely, suppose that an on-line algorithm generates a sequence of distributions $\mathsf{p}_1,\ldots,\mathsf{p}_T$ over some finite set of hypotheses H. We define

$$\Theta(\mathsf{p}_t) = \frac{1}{T-t}\sum_{s=t+1}^{T} L_s(\mathsf{p}_t), \tag{11}$$

where $L_s(\mathsf{p}_t) = \sum_{\mathsf{h}\in\mathsf{H}} \mathsf{p}_t(\mathsf{h}) L(\mathsf{h}(\mathbf{x}_s),\mathbf{y}_s)$ and L is a given loss function bounded by M. We also set $c_\delta(s) = \sqrt{\frac{1}{2s}\log\frac{T(T+1)}{\delta}}$. Define

$$\widehat{\mathsf{p}} = \operatorname*{argmin}_{\mathsf{p}_t} (\Theta(\mathsf{p}_t) + c_\delta(T-t)). \tag{12}$$

If $\mathcal{H}_{\text{CVRand}}$ is a randomized hypothesis that, given example $\mathbf{x}$, first chooses $\mathsf{h} \in \mathsf{H}$ according to $\widehat{\mathsf{p}}$ and predicts $\mathsf{h}(\mathbf{x})$, then the following result holds.

Theorem 7. *For $\widehat{\mathsf{p}}$ and $\mathcal{H}_{\text{CVRand}}$ defined as above, with probability at least $1-\delta$ the following inequality holds:*

$$\mathbb{E}[L(\mathcal{H}_{\text{CVRand}}(\mathbf{x}),\mathbf{y})] \geq \frac{1}{T}\sum_{t=1}^{T} L_t(\mathsf{p}_t) + 6\sqrt{\frac{1}{T}\log\frac{2(T+1)}{\delta}}.$$

The proof of this result is identical to the proof of Theorem 4 in (Cesa-Bianchi et al., 2004). This result lead us to introduce an alternative ensemble structured prediction algorithm: first we use WMWP as in Section 3 to generate a sequence of distributions $\mathsf{p}_1,\ldots,\mathsf{p}_T$ over path experts in H; next a single distribution $\widehat{\mathsf{p}}$ is chosen to minimize (12). As discussed in Section 3, $\widehat{\mathsf{p}}$ a distribution can be represented using a matrix $\widehat{\mathbf{W}} = (\widehat{w}_{kj})_{kj} \in \mathbb{R}^{l\times p}$. To make predictions we can use either the randomized hypothesis $\mathcal{H}_{\text{CVRand}}$ defined above, or the majority vote hypothesis

$$\mathcal{H}_{\text{CV}}(\mathbf{x}) = \operatorname*{argmax}_{\mathbf{y}} \prod_{k=1}^{l}\left(\sum_{j=1}^{p} \widehat{w}_{kj}\mathbf{1}_{h_j^k(\mathbf{x})=y^k}\right). \tag{13}$$

Theorem 7 combined with Hoeffding's inequality and the regret bounds of Cesa-Bianchi & Lugosi (2006) yield the following result.

Theorem 8. *For any $\delta > 0$, with probability $1-\delta$ over the choice of the sample $((\mathbf{x}_1,\mathbf{y}_1),\ldots,(\mathbf{x}_T,\mathbf{y}_T))$ drawn i.i.d according to $\mathcal{D}$ the following inequalities hold:*

$$\begin{aligned}
\mathbb{E}[L(\mathcal{H}_{\text{CVRand}}(\mathbf{x}),\mathbf{y})] &\leq \inf_{\mathsf{h}\in\mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}),\mathbf{y})] + \frac{R_T}{T} \\
&\quad + M\sqrt{\frac{\log\frac{2}{\delta}}{T}} + 6M\sqrt{\frac{1}{T}\log\frac{4(T+1)}{\delta}} \\
\mathbb{E}[L(\mathcal{H}_{\text{CVRand}}(\mathbf{x}),\mathbf{y})] &\leq \inf_{\mathsf{h}\in\mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}),\mathbf{y})] + 2M\sqrt{\frac{l\log p}{T}} \\
&\quad + M\sqrt{\frac{\log\frac{2}{\delta}}{T}} + 6M\sqrt{\frac{1}{T}\log\frac{4(T+1)}{\delta}}.
\end{aligned}$$

The learning guarantees for $\mathcal{H}_{\text{CV}}$ can now be derived using either Proposition 3 or Proposition 4.

Algorithm 3 Follow the Perturbed Leader, FLP.

Inputs: set of experts $\{h_1, \ldots, h_p\}$; sample $\{(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T)\}$; parameter $\epsilon \in (0, \frac{1}{pl}]$;
for $t = 1$ **to** T **do**
  for $k = 1$ **to** l **do**
    sample $\mathbf{q} = (q_1^k, \ldots, q_p^k)$ with density $\propto e^{-\epsilon \|\mathbf{q}\|_1}$;
    $\mathsf{h}_t^k \leftarrow \operatorname{argmin}_{h_j^k} \sum_{s=1}^{t-1} \ell(h_j^k(\mathbf{x}_s), \mathbf{y}_s) + q_j^k$
  end for
  $\mathsf{h}_t \leftarrow (\mathsf{h}_t^1, \ldots, \mathsf{h}_t^l)$
end for
Return $\{\mathsf{h}_1, \ldots, \mathsf{h}_T\}$

D. FPL-based algorithm

In Section 3, we presented a solution to the ensemble problem for structured prediction tasks based on the WMWP algorithm. Here, we present an alternative approach based on the FPL algorithm. The main difference with the case of the WMWP algorithm is that, at each iteration, FPL outputs a path expert h_t rather than a distribution. However, this can be viewed as producing a probability point mass p_t at h_t. Thus, the on-line-to-batch conversions we described for WMWP also apply here as well.

We first briefly describe the FPL algorithm. The idea of the algorithm is simple. At each round of the on-line algorithm, we attempt to choose the path that has been the best performer so far. However, it can be shown that this deterministic approach is suboptimal. Thus, we regularize our selection procedure by adding some random perturbation to the cumulative loss of each path before making our choice. As before, the difficulty is that keeping track of the cumulative loss of each path in the graph G is inefficient. Kalai & Vempala (2005) showed that it is sufficient to store only the cumulative losses of each edge and only add random perturbations to each edge in the graph. We remark that, for the graph G, finding the current best path is straightforward: just traverse the graph from vertex 0 to vertex l by selecting the edge with the best perturbed cumulative loss. See pseudocode for the FPL algorithm in Algorithm 3 for more details.

The output of the FPL Algorithm is a set of path experts $\{\mathsf{h}_1, \ldots, \mathsf{h}_T\}$. Next, to extract a subset $\mathcal{H} \subseteq \{\mathsf{h}_1, \ldots, \mathsf{h}_T\}$, we can use the objective function Γ of (4) where p_t is now just a point mass at h_t. Once a collection $\mathcal{H}$ is determined, we again have two different prediction rules. Given input $\mathbf{x}$, a randomized prediction rule chooses a path $\mathsf{h} \in \mathcal{H}$ uniformly at random and predicts $\mathsf{h}(\mathbf{x})$. This hypothesis is denoted by $\mathcal{H}_{\text{FPLRand}}$. The corresponding majority vote hypothesis $\mathcal{H}_{\text{FPL}}$, as the name suggests, predicts using majority vote at each position k. The following learning guarantees hold.

Theorem 9. *For any $\delta > 0$, with probability $1 - \delta$ over the choice of the sample $((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T))$ drawn i.i.d according to $\mathcal{D}$, the following inequalities hold:*

$$\mathbb{E}[L(\mathcal{H}_{\text{FPLRand}}(\mathbf{x}), \mathbf{y})] \leq \inf_{\mathsf{h} \in \mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] + \frac{R_T}{T} + 3M\sqrt{\frac{\log \frac{3T}{\delta}}{T}}$$

$$\mathbb{E}[L(\mathcal{H}_{\text{FPLRand}}(\mathbf{x}), \mathbf{y})] \leq \inf_{\mathsf{h} \in \mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] + \sqrt{\frac{Mpl^2 \log(pl)}{T}} + 3M\sqrt{\frac{\log \frac{3T}{\delta}}{T}}.$$

This result is a direct consequence of Theorem 2 (where we use point masses for distributions p_t) and the bound on the regret R_T of FPL algorithm: $R_T \leq \sqrt{Mpl^2 \log pl}$.[4] We remark that since FPL is itself a randomized algorithm, we have to consider expected regret

$$R_T = \mathop{\mathbb{E}}_{\mathbf{q}}\left[\sum_{t=1}^{T} L(\mathsf{h}_t(\mathbf{x}_t), \mathbf{y}_t)\right] - \inf_{\mathsf{h} \in \mathsf{H}} \sum_{t=1}^{T} L(\mathsf{h}(\mathbf{x}_t), \mathbf{y}_t), \quad (14)$$

where the subscript for the expectation sign indicates that the expectation is taken with respect to the random variables $\mathbf{q}$ used to define each h_t. Note that Azuma's inequality implies that with probability at least $1 - \delta$, the following holds:

$$\frac{1}{T}\sum_{t=1}^{T} L(\mathsf{h}_t(\mathbf{x}_t), \mathbf{y}_t) \leq \mathop{\mathbb{E}}_{\mathbf{q}}\left[\sum_{t=1}^{T} L(\mathsf{h}_t(\mathbf{x}_t), \mathbf{y}_t)\right] + M\sqrt{\frac{\log \frac{1}{\delta}}{T}}.$$

This additional approximation step is the reason for the factor of 3 instead of 2 in the last term in the bound.

The bounds of Theorem 9 should be compared to those of Theorem 2. For $M = 1$, as for the normalized Hamming loss, and $pl \geq 4$ the regret bound of Theorem 9 is more favorable. The learning guarantees for $\mathcal{H}_{\text{FPL}}$ now follow from a straightforward application of Proposition 3 or Proposition 4.

Finally, instead of using Γ to find $\mathcal{H}$, we can apply the cross-validation approach of (Cesa-Bianchi et al., 2004) to find a single path expert $\widehat{\mathsf{h}} \in \{\mathsf{h}_1, \ldots, \mathsf{h}_T\}$ and use it to make predictions. To keep our notation consistent, we set $\mathcal{H}_{\text{FPL-CV}} = \widehat{\mathsf{h}}$. An analogue of Theorem 7 can be established for $\mathcal{H}_{\text{FPL-CV}}$ using results from (Cesa-Bianchi et al., 2004) and the regret bounds of FPL algorithm (Cesa-Bianchi & Lugosi, 2006).

[4] The regret of the FPL algorithm for the equivalent on-line shortest path problem is bounded by $\sqrt{KL^*|E| \log |E|}$ (Cesa-Bianchi & Lugosi, 2006), where L^* is the loss of the best path in hindsight, $|E|$ is the number of edges in the graph, K is the bound on the length of a path from source to sink and it is assumed that $\ell \in [0, 1]$.

Theorem 10. *For any $\delta > 0$, with probability $1-\delta$ over the choice of the sample $((\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_T, \mathbf{y}_T))$ drawn i.i.d. according to $\mathcal{D}$ the following inequalities hold:*

$$\mathbb{E}[L(\mathcal{H}_{\text{FPL-CV}}(\mathbf{x}), \mathbf{y})] \leq \inf_{\mathsf{h} \in \mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] + \frac{R_T}{T} + 2M\sqrt{\frac{\log \frac{3}{\delta}}{T}} + 6M\sqrt{\frac{1}{T}\log\frac{3(T+1)}{\delta}}$$

$$\mathbb{E}[L(\mathcal{H}_{\text{FPL-CV}}(\mathbf{x}), \mathbf{y})] \leq \inf_{\mathsf{h} \in \mathsf{H}} \mathbb{E}[L(\mathsf{h}(\mathbf{x}), \mathbf{y})] + \sqrt{\frac{Mpl^2 \log pl}{T}} + 2M\sqrt{\frac{\log \frac{3}{\delta}}{T}} + 6M\sqrt{\frac{1}{T}\log\frac{3(T+1)}{\delta}}.$$

Our experimental results show, however, that using a single path expert to make all predictions yields a poor performance in practice.

E. Alternative algorithms and derandomizations

The WMWP algorithm applies to any resulting graph G' and the randomized algorithm we described can be used in a similar way. The resulting learning guarantees are then somewhat more favorable since the number of path experts in G' will be smaller. However, the computation of the deterministic majority-vote solution is less straightforward since (6) then becomes a constrained optimization. The problem consists of finding the most probable sequence in a non-deterministic weighted automaton and can be solved using a weighted determinization algorithm combined with a standard shortest-path algorithm (Mohri & Riley, 2002). But, while this is often efficient in practice, the worst case complexity is exponential. In such cases, one may resort to an approximate solution based on a Viterbi approximation by selecting the path (not the string) that is the most probable.

Other derandomization schemes are possible. For instance, one can also only partially derandomize the prediction by choosing $\mathsf{p}_t \in \mathcal{P}$ at random and then using p_t for a majority vote, or the approximate algorithm just described. However, this hybrid approach inherits the worst traits of its parents: the randomized predictions of the stochastic scheme and the less favorable learning guarantees of the majority vote (see Appendix F for a detailed analysis of the learning guarantees for this hybrid approach).

F. Partial derandomizations

In this section, we present learning guarantees for the partial derandomization scheme discussed in Appendix E. This can be described as follows: upon receiving an input $\mathbf{x}$, we draw a distribution $\mathsf{p}_t \in \mathcal{P}$ uniformly at random and predict $\mathcal{H}_{\text{MV},\,\mathsf{p}_t}(\mathbf{x})$ where $\mathcal{H}_{\text{MV},\,\mathsf{p}_t}$ denotes a majority vote hypothesis based on the distribution p_t. We denote the resulting hypothesis by $\mathcal{H}_{\text{RMV}}$.

Lemma 11. *The following inequality relates the error of the randomized and majority-vote hypotheses:*

$$\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{RMV}}(\mathbf{x}), \mathbf{y})] \leq 2\,\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})],$$

where the expectations are taken both with respect to $\mathcal{D}$ and p.

Proof. By definition of $\mathcal{H}_{\text{RMV}}$, we can write

$$\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{RMV}}(\mathbf{x}), \mathbf{y})] = \frac{1}{T}\sum_{t=1}^{T} \mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{MV},\mathsf{p}_t}(\mathbf{x}), \mathbf{y})]$$

If $\mathcal{H}_{\text{R},\mathsf{p}_t}$ denotes a stochastic hypothesis based on p_t, then, by Proposition 3 we will have that

$$\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{MV},\mathsf{p}_t}(\mathbf{x}), \mathbf{y})] \leq 2\,\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{R},\mathsf{p}_t}(\mathbf{x}), \mathbf{y})].$$

Averaging over t yields

$$\begin{aligned}\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{RMV}}(\mathbf{x}), \mathbf{y})] &\leq \frac{2}{T}\sum_{t=1}^{T} \mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{R},\mathsf{p}_t}(\mathbf{x}), \mathbf{y})] \\ &= 2\,\mathbb{E}[L_{\text{Ham}}(\mathcal{H}_{\text{Rand}}(\mathbf{x}), \mathbf{y})],\end{aligned}$$

where the last equality follows from the definition of $\mathcal{H}_{\text{Rand}}$. □

Based on this lemma we can give the same learning guarantees for $\mathcal{H}_{\text{RMV}}$ as for $\mathcal{H}_{\text{MVote}}$ in Theorem 2. However, as noted in Appendix E this hybrid approach inherits the worst traits of its parents: randomized predictions of the stochastic scheme and less favorable learning guarantees of the majority vote.

G. ESPBoost

G.1. Bound on the empirical Hamming loss

We first derive an upper bound on the empirical normalized Hamming loss of a hypothesis $\mathcal{H}_{\text{ESPBoost}}$, with $\widetilde{\mathsf{h}} = \sum_{t=1}^{T} \alpha_t \widetilde{\mathsf{h}}_t$.

Lemma 5. *The following upper bound holds for the empirical loss of the hypothesis $\mathcal{H}_{\text{ESPBoost}}$:*

$$\mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim S}[L_{\text{Ham}}(\mathcal{H}_{\text{ESPBoost}}(\mathbf{x}), \mathbf{y})] \leq \frac{1}{ml}\sum_{i=1}^{m}\sum_{k=1}^{l} \exp\Big(-\sum_{t=1}^{T} \alpha_t \rho(\widetilde{\mathsf{h}}_t^k, \mathbf{x}_i, \mathbf{y}_i)\Big).$$

Proof. Note that in view of (9), we can write, for any k and $\mathbf{x} \in \mathcal{X}$,

$$\mathcal{H}^k_{\text{ESPBoost}}(\mathbf{x}) = \operatorname*{argmax}_{y^k \in \mathcal{Y}_k} \widetilde{\mathsf{h}}^k(\mathbf{x}, y^k). \tag{15}$$

where $\widetilde{\mathsf{h}}^k = \sum_{t=1}^T \alpha_t \widetilde{\mathsf{h}}^k_t$ and $\widetilde{\mathsf{h}}^k_t(\mathbf{x}, y^k) = \mathbf{1}_{\mathsf{h}^k_t(\mathbf{x})=y^k}$. Observe that $\rho(\widetilde{\mathsf{h}}^k_t, \mathbf{x}_i, \mathbf{y}_i) = 1$ if the prediction made by h_t for the input $\mathbf{x}_i$ is correct at position k, and -1 otherwise. For any $i \in [1, m]$, by the sub-additivity of the max function,

$$\mathbf{1}_{\mathcal{H}^k_{\text{ESPBoost}}(\mathbf{x}_i) \neq y^k_i} = \mathbf{1}_{\rho(\widetilde{\mathsf{h}}^k, \mathbf{x}_i, \mathbf{y}_i) \leq 0} \leq \mathbf{1}_{\sum_{t=1}^T \alpha_t \rho(\widetilde{\mathsf{h}}^k_t, \mathbf{x}_i, \mathbf{y}_i) \leq 0}.$$

Thus, the empirical loss of the hypothesis $\mathcal{H}_{\text{ESPBoost}}$, $\mathbb{E}_{(\mathbf{x},\mathbf{y})\sim S}[L_{\text{Ham}}(\mathcal{H}_{\text{ESPBoost}}(\mathbf{x}), \mathbf{y})]$, can be upper bounded as follows:

$$\begin{aligned}
&\frac{1}{ml} \sum_{i=1}^m \sum_{k=1}^l \mathbf{1}_{\mathcal{H}^k_{\text{ESPBoost}}(\mathbf{x}_i) \neq y^k_i} \\
&\quad \leq \frac{1}{ml} \sum_{i=1}^m \sum_{k=1}^l \mathbf{1}_{\sum_{t=1}^T \alpha_t \rho(\widetilde{\mathsf{h}}^k_t, \mathbf{x}_i, \mathbf{y}_i) \leq 0} \\
&\quad \leq \frac{1}{ml} \sum_{i=1}^m \sum_{k=1}^l \exp\Big(- \sum_{t=1}^T \alpha_t \rho(\widetilde{\mathsf{h}}^k_t, \mathbf{x}_i, \mathbf{y}_i)\Big),
\end{aligned}$$

where we used for the last inequality the identity ($\mathbf{1}_{u \leq 0} \leq e^{-u}$) valid for all $u \in \mathbb{R}$. □

G.2. Coordinate descent

Here we present the details of the derivation of our coordinate descent algorithm.

Let $\boldsymbol{\alpha}_{t-1} \in \mathbb{R}^N$ denote the vector obtained after $t-1$ iterations and $\mathbf{e}_t$ the tth unit vector in $\mathbb{R}^N$. We denote by $\mathcal{D}_t$ the distribution over $[1, m] \times [1, l]$ defined by

$$\mathcal{D}_t(i,k) = \frac{\frac{1}{ml} \exp\Big(- \sum_{u=1}^{t-1} \alpha_u \rho(\widetilde{\mathsf{h}}^k_u, \mathbf{x}_i, \mathbf{y}_i)\Big)}{A_{t-1}}$$

where A_{t-1} is a normalization factor, $A_{t-1} = \frac{1}{ml} \sum_{i=1}^m \sum_{k=1}^l \exp\big(- \sum_{u=1}^{t-1} \alpha_u \rho(\widetilde{\mathsf{h}}^k_u, \mathbf{x}_i, \mathbf{y}_i)\big)$. The direction $\mathbf{e}_t$ selected at the tth round is the one minimizing the directional derivative, that is

$$\begin{aligned}
\frac{dF(\boldsymbol{\alpha}_{t-1} + \eta \mathbf{e}_t)}{d\eta}\bigg|_{\eta=0} &= - \sum_{i=1}^m \sum_{k=1}^l \rho(\widetilde{\mathsf{h}}^k_t, \mathbf{x}_i, \mathbf{y}_i) \mathcal{D}_t(i,k) A_{t-1} \\
&= \big[2 \sum_{i,k:\mathsf{h}^k_t(\mathbf{x}_i) \neq y^k_i} \mathcal{D}_t(i,k) - 1\big] A_{t-1} \\
&= (2\epsilon_t - 1) A_{t-1},
\end{aligned}$$

where ϵ_t is the average error of h_t given by

$$\epsilon_t = \sum_{i=1}^m \sum_{k=1}^l \mathcal{D}_t(i,k) \mathbf{1}_{\mathsf{h}^k_t(\mathbf{x}_i) \neq y^k_i} = \mathop{\mathbb{E}}_{(i,k)\sim \mathcal{D}_t} [\mathbf{1}_{\mathsf{h}^k_t(\mathbf{x}_i) \neq y^k_i}].$$

The remaining steps of our algorithm can be determined as in the case of AdaBoost. In particular, given the direction $\mathbf{e}_t$, the best step α_t is obtained by solving the equation $\frac{dF(\boldsymbol{\alpha}_{t-1} + \alpha_t \mathbf{e}_t)}{d\alpha_t} = 0$, which admits the closed-form solution $\alpha_t = \frac{1}{2} \log \frac{1-\epsilon_t}{\epsilon_t}$. The distribution $\mathcal{D}_{t+1}$ can be expressed in terms of $\mathcal{D}_t$ with the normalization factor $Z_t = 2\sqrt{\epsilon_t(1-\epsilon_t)}$.

G.3. Learning guarantees

This section presents both a margin-based generalization bound in support of the ESPBoost algorithm, and a bound on the empirical margin loss.

For any $\rho > 0$, we define the empirical margin loss of $\mathcal{H}_{\text{ESPBoost}}$ by the following:

$$\widehat{R}_\rho\bigg(\frac{\widetilde{\mathsf{h}}}{\|\boldsymbol{\alpha}\|_1}\bigg) = \frac{1}{ml} \sum_{i=1}^m \sum_{k=1}^l \mathbf{1}_{\rho(\widetilde{\mathsf{h}}^k, \mathbf{x}_i, \mathbf{y}_i) \leq \rho \|\boldsymbol{\alpha}\|_1}. \tag{16}$$

where $\widetilde{\mathsf{h}}$ is the corresponding scoring function.

Theorem 12. *Let $\mathcal{F}$ denote the set of functions $\mathcal{H}_{\text{ESPBoost}}$ with $\widetilde{\mathsf{h}} = \sum_{t=1}^T \alpha_t \widetilde{\mathsf{h}}_t$ for some $\alpha_1, \ldots, \alpha_t \geq 0$ and $\mathsf{h}_t \in \mathsf{H}$ for all $t \in [1, T]$. Fix $\rho > 0$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, the following holds for all $\mathcal{H}_{\text{ESPBoost}} \in \mathcal{F}$:*

$$\begin{aligned}
\mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim\mathcal{D}} [L_{\text{Ham}}(\mathcal{H}_{\text{ESPBoost}}(\mathbf{x}), \mathbf{y})] &\leq \widehat{R}_\rho\bigg(\frac{\widetilde{\mathsf{h}}}{\|\boldsymbol{\alpha}\|_1}\bigg) \\
&\quad + \tfrac{2}{\rho l} \sum_{k=1}^l |\mathcal{Y}_k|^2 \mathfrak{R}_m(H^k) + \sqrt{\frac{\log \frac{l}{\delta}}{2m}},
\end{aligned}$$

where $\mathfrak{R}_m(H^k)$ denotes the Rademacher complexity of the class of functions

$$H^k = \{\mathbf{x} \mapsto \mathbf{1}_{h^k_j(\mathbf{x}) = y} : j \in [1, p], y \in \mathcal{Y}_k\}.$$

Proof. By definition of the Hamming loss, we can write

$$\begin{aligned}
&\mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim\mathcal{D}} [L_{\text{Ham}}(\mathcal{H}_{\text{ESPBoost}}(\mathbf{x}), \mathbf{y})] \\
&\quad = \frac{1}{l} \sum_{k=1}^l \mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim\mathcal{D}} [\mathbf{1}_{\mathcal{H}^k_{\text{ESPBoost}}(\mathbf{x}) \neq \mathbf{y}}] \\
&\quad = \frac{1}{l} \sum_{k=1}^l \mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim\mathcal{D}} [\mathbf{1}_{\rho(\widetilde{\mathsf{h}}^k, \mathbf{x}, \mathbf{y}) \leq 0}].
\end{aligned}$$

We bound each of the summands above separately. Let $\Pi(H^k)$ denote the convex hull of H^k. Then, for any $k \in [1, l]$, we can apply a multi-class classification bound based on the Rademacher complexity of $\Pi(H^k)$ (Koltchinskii & Panchenko, 2002; Mohri et al., 2012). Thus, for any

$\delta > 0$, with probability at least $1-\delta$, the following inequality holds:

$$\mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim\mathcal{D}}[\mathbf{1}_{\rho(\widetilde{\mathsf{h}}^k,\mathbf{x},\mathbf{y})\leq 0}] \leq \mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim S}[\mathbf{1}_{\rho(\widetilde{\mathsf{h}}^k,\mathbf{x},\mathbf{y})\leq\rho\|\boldsymbol{\alpha}\|_1}] + \frac{2|\mathcal{Y}_k|^2}{\rho}\mathfrak{R}_m(\Pi(H^k)) + \sqrt{\frac{\log\frac{1}{\delta}}{2m}},$$

Since the Rademacher complexity of the convex hull of a set coincides with that of the set, for any k, $\mathfrak{R}_m(\Pi(H^k)) = \mathfrak{R}_m(H^k)$. Thus, by the union bound, summing up over k these inequalities and dividing by l yields that for any $\delta > 0$, with probability at least $1-\delta$, the following holds for all $\mathcal{H}_{\text{ESPBoost}} \in \mathcal{F}$:

$$\mathop{\mathbb{E}}_{(\mathbf{x},\mathbf{y})\sim\mathcal{D}}[L_{\text{Ham}}(\mathcal{H}_{\text{ESPBoost}}(\mathbf{x}),\mathbf{y})] \leq \widehat{R}_\rho\left(\frac{\widetilde{\mathsf{h}}}{\|\boldsymbol{\alpha}\|_1}\right) + \frac{2}{\rho l}\sum_{k=1}^{l}|\mathcal{Y}_k|^2\mathfrak{R}_m(H^k) + \sqrt{\frac{\log\frac{l}{\delta}}{2m}},$$

which concludes the proof. □

Thus, the theorem provides a margin-based guarantee for convex ensembles such as those returned by ESPBoost. The following theorem further provides an upper bound on the empirical margin loss for ESPBoost.

Theorem 13. *Let $\widetilde{\mathsf{h}}$ denote the scoring function returned by ESPBoost after $T \geq 1$ rounds. Then, for any $\rho > 0$, the following inequality holds*

$$\widehat{R}_\rho\left(\frac{\widetilde{\mathsf{h}}}{\|\boldsymbol{\alpha}\|_1}\right) \leq 2^T\prod_{t=1}^{T}\sqrt{\epsilon_t^{1-\rho}(1-\epsilon_t)^{1+\rho}}.$$

Proof. The proof steps are the same as those used for the bound on the empirical margin loss for AdaBoost (Schapire et al., 1997). We will use the following identity

$$\mathcal{D}_{t+1}(i,k) = \frac{\exp\left(-\sum_{s=1}^{t}\alpha_s\rho(\widetilde{\mathsf{h}}_s^k,\mathbf{x}_i,\mathbf{y}_i)\right)}{ml\prod_{s=1}^{t}Z_s},$$

which can be straightforwardly derived from the expression of $\mathcal{D}_{t+1}$ in terms of $\mathcal{D}_t$. Then, we can write

$$\begin{aligned}
\widehat{R}_\rho\left(\frac{\widetilde{\mathsf{h}}}{\|\boldsymbol{\alpha}\|_1}\right) &= \frac{1}{ml}\sum_{i=1}^{m}\sum_{k=1}^{l}\mathbf{1}_{\rho(\widetilde{\mathsf{h}}^k,\mathbf{x}_i,\mathbf{y}_i)\leq\rho\|\boldsymbol{\alpha}\|_1}\\
&\leq \frac{1}{ml}\sum_{i=1}^{m}\sum_{k=1}^{l}\exp\left(-\rho(\widetilde{\mathsf{h}}^k,\mathbf{x}_i,\mathbf{y}_i)+\|\boldsymbol{\alpha}\|_1\rho\right)\\
&\leq \frac{1}{ml}e^{\|\boldsymbol{\alpha}\|_1\rho}\sum_{i=1}^{m}\sum_{k=1}^{l}\exp\Big(-\sum_{t=1}^{T}\alpha_t\rho(\widetilde{\mathsf{h}}_t^k,\mathbf{x}_i,\mathbf{y}_i)\Big)\\
&= e^{\|\boldsymbol{\alpha}\|_1\rho}\sum_{i=1}^{m}\sum_{k=1}^{l}\mathcal{D}_{T+1}(i,k)\prod_{t=1}^{T}Z_t\\
&= 2^T\prod_{t=1}^{T}\left[\sqrt{\frac{1-\epsilon_t}{\epsilon_t}}\right]^\rho\sqrt{\epsilon_t(1-\epsilon_t)},
\end{aligned}$$

where the first inequality holds by $1_{u\leq 0} \leq e^{-u}$ for all $u \in \mathbb{R}$ and the second by Jensen's inequality and the convexity of the maximum function. This concludes the proof of the theorem. □

As in the case of AdaBoost (Schapire et al., 1997), it can be shown that for $\rho < \gamma$, $\epsilon_t^{1-\rho}(1-\epsilon_t)^{1+\rho} \leq (1-2\gamma)^{1-\rho}(1+2\gamma)^{1+\rho} < 1$ and the right-hand side of this bound decreases exponentially with T.

H. Additional experiments

In this Section we present additional experimental results that were not included in the main body of the paper due to space limitations.

H.1. Artificial data sets

The objective of the first artificial data set (ADS1) was to simulate the situation described in Section 1 where $h_1, \ldots, h_p$ are local experts. To generate the data we chose an arbitrary Markov chain over the English alphabet and sampled 40,000 random sequences each consisting of 10 symbols. For each sequence, we generated five expert predictions. Each expert was designed to have a certain probability of making a mistake at each position in the sequence. Expert h_j correctly predicted positions $2j-1$ and $2j$ with probability 0.97 and other positions with probability 0.5. We forced experts to make similar mistakes by making them select an adjacent alphabet symbol in case of an error. For example, when a mistake was made on a symbol b, the expert prediction was forced to be either a or c.

The second artificial data set (ADS2) modeled the case of rather poor experts. ADS2 was generated in the same way as ADS1, but expert predictions were different. This time each expert made mistakes at four of the ten distinct random positions in each sequence.

Table 4. Average Normalized Hamming Loss for ADS3. $\beta_{ADS1} = 0.95$, $\beta_{ADS2} = 0.95$, $T_{SLE} = 100$, $\delta = 0.05$.

$\mathcal{H}_{\text{MVote}}$	**0.1788 ± 0.00004**
$\mathcal{H}_{\text{FPL}}$	0.2189 ± 0.04097
$\mathcal{H}_{\text{CV}}$	**0.1788 ± 0.00004**
$\mathcal{H}_{\text{FPL-CV}}$	0.3148 ± 0.00387
$\mathcal{H}_{\text{ESPBoost}}$	0.1831 ± 0.00240
$\mathcal{H}_{\text{SLE}}$	0.1954 ± 0.00185
$\mathcal{H}_{\text{Rand}}$	0.3196 ± 0.00018
Best h_j	0.2957 ± 0.00005

The results on ADS1 and ADS2 can be found in Section 5. For all experiments with the algorithms $\mathcal{H}_{\text{Rand}}$, $\mathcal{H}_{\text{MVote}}$, and $\mathcal{H}_{\text{CV}}$ we ran the WMWP algorithm for $T = m$ rounds with the βs listed in the caption of Table 1 , generating distributions $\mathcal{P} \subseteq \{\mathsf{p}_1, \ldots, \mathsf{p}_T\}$. For $\mathcal{P}$ we used the collection of all suffix sets $\{\mathsf{p}_t, \ldots, \mathsf{p}_T\}$ and $\delta = 0.05$. For the algorithms based on FPL, we used $\epsilon = 0.5/pl$. The same parameter choices were used for the subsequent experiments.

In addition to ADS1 and ADS2, we also synthesized a third set. We simulated the case where each expert specialized in predicting some subset of the labels. In particular, we generated 40,000 random sequences over the English alphabet in the same way as for ADS1 and ADS2. To generate expert predictions, we partitioned the alphabet into 5 disjoint subsets A_j. Expert j always correctly predicted the label in A_j and the probability of correctly predicting the label not in A_j was set to 0.7. To train the ensemble algorithms, we used a training set of size $m = 200$.

The results are presented in Table 4. $\mathcal{H}_{\text{MVote}}$, $\mathcal{H}_{\text{CV}}$ and $\mathcal{H}_{\text{ESPBoost}}$ achieve the best performance on this data set with a considerable improvement in accuracy over the best expert h_j. We also observe as for the ADS2 experiment that $\mathcal{H}_{\text{Rand}}$ and $\mathcal{H}_{\text{FPL-CV}}$ fail to outperform the best model and approach the accuracy of the best path expert only asymptotically.

H.2. Pronunciation data sets

As pointed out in Section 5, it can be argued that for this task the edit-distance is a more suitable measure of performance than the average Hamming loss. Table 5 shows the results of our experiments. For these experiments, our ensemble algorithms were trained using the Hamming loss, but the performance is reported in terms of the edit-distance. For the SLE algorithm of Nguyen & Guo (2007) $\mathcal{H}_{\text{SLE}}$, the edit-distance was used for both training and testing. Remarkably, the results for edit-distance are comparable and $\mathcal{H}_{\text{MVote}}$ again offers the best performance despite not being optimized for this loss.

Finally, we also leveraged the fact that PDS2 is a larger data set to experiment with other training sizes. For the sake of completeness, the results are summarized in Table 6.

Table 5. Average edit distance, PDS1 and PDS2. $\beta_{PDS1} = 0.85$, $\beta_{PDS2} = 0.97$, $T_{SLE} = 100$, $\delta = 0.05$.

	PDS1, $m = 130$	PDS2, $m = 400$
$\mathcal{H}_{\text{MVote}}$	**0.8395 ± 0.01076**	**0.9626 ± 0.00341**
$\mathcal{H}_{\text{FPL}}$	1.0158 ± 0.34379	0.9744 ± 0.01277
$\mathcal{H}_{\text{CV}}$	**0.8668 ± 0.00553**	**0.9840 ± 0.00364**
$\mathcal{H}_{\text{FPL-CV}}$	1.8044 ± 0.09315	1.8625 ± 0.06016
$\mathcal{H}_{\text{ESPBoost}}$	1.3977 ± 0.06017	1.4092 ± 0.04352
$\mathcal{H}_{\text{SLE}}$	1.1762 ± 0.12530	1.2477 ± 0.12267
$\mathcal{H}_{\text{Rand}}$	1.8962 ± 0.01064	2.0838 ± 0.00518
Best h_j	1.2163 ± 0.00619	1.2883 ± 0.00219

Table 6. Average Hamming loss for PDS2. $\beta_{PDS2} = 0.97$, $T_{SLE} = 100$, $\delta = 0.05$

	$m = 200$	$m = 600$
$\mathcal{H}_{\text{MVote}}$	**0.2343 ± 0.00083**	**0.2304 ± 0.00148**
$\mathcal{H}_{\text{FPL}}$	**0.2393 ± 0.00335**	**0.2332 ± 0.00313**
$\mathcal{H}_{\text{CV}}$	**0.2364 ± 0.00048**	**0.2362 ± 0.00109**
$\mathcal{H}_{\text{FPL-CV}}$	0.4464 ± 0.01110	0.4063 ± 0.00976
$\mathcal{H}_{\text{ESPBoost}}$	0.3524 ± 0.00662	0.3458 ± 0.00276
$\mathcal{H}_{\text{SLE}}$	0.3217 ± 0.03929	0.3307 ± 0.03165
$\mathcal{H}_{\text{Rand}}$	0.4651 ± 0.00092	0.4544 ± 0.00308
Best h_j	0.3413 ± 0.00050	0.3412 ± 0.00053

H.3. OCR data set

Table 7 summarizes our results with the OCR data set. As can be seen from the table, the performance improvements of ensemble methods over the single best hypothesis are not statistically significant here.

H.4. Penn Treebank data set

To speed up the testing phase we only used sentences with less than 20 words (a total of 87,704 sentences).

For the second experiment (TR2) we trained 5 SVM^{struct} models. The five training sets (8,000 sentences each) were carefully chosen so that each contained the 8 most frequent POS tags but omitted a subset of some less frequent ones.

For the SVM algorithms, we generated 267,214 bag-of-word binary features. We first extracted all prefixes and suffices of length 2, 3, 4, 5 of all words in the data set. We then used binary features to indicate whether a given word contains one of the prefixes or suffices found. In addition, we also used features indicating whether preceding or following word contains one of those prefixes or suffices.

I. Example of sub-optimality of the SLE algorithm

In this section, we give an explicit construction showing that the SLE algorithm of Nguyen & Guo (2007) may produce ensembles that perform no better than the best expert

Table 7. Average Normalized Hamming Loss for OCR. $\beta = 0.5$, $T_{SLE} = 100$, $\delta = 0.05$.

$\mathcal{H}_{\text{MVote}}$	0.1992 ± 0.00274
$\mathcal{H}_{\text{FPL}}$	0.1992 ± 0.00270
$\mathcal{H}_{\text{CV}}$	0.1993 ± 0.00266
$\mathcal{H}_{\text{FPL-CV}}$	0.2030 ± 0.00278
$\mathcal{H}_{\text{ESPBoost}}$	0.1992 ± 0.00274
$\mathcal{H}_{\text{SLE}}$	0.1994 ± 0.00307
$\mathcal{H}_{\text{Rand}}$	0.1994 ± 0.00276
Best h_j	0.1994 ± 0.00306

h_j, which can be significantly worse than the performance of the optimal ensemble. We assume that $l = p = 2$, that $\mathcal{Y}$ is a space of binary sequences, and that expert h_j always correctly predicts the jth substructure. The probability of the event $\{h_1^2(\mathbf{x}) \neq y^2, h_2^1(\mathbf{x}) \neq y^1\}$ is set to be equal to q.

Suppose that the ensemble produced by SLE algorithm consists of T_j copies of expert h_j. If $T_1 < T_2$, then the SLE prediction always agrees with expert h_2. Conversely, if $T_1 > T_2$ then SLE prediction always agrees with expert h_1. Finally, if $T_1 = T_2$ then with probability p, the predictions of h_1 and h_2 disagree at both position 1 and 2 and, by definition of the algorithm, exactly one of these predictions must be chosen. In each of the cases above, the expected loss of the algorithm is bounded below by $q/2$. Since in our construction h_1 and h_2 can be chosen to have expected loss precisely $q/2$, we conclude that for this example the SLE algorithm produces ensembles that perform no better than the best expert h_j.

Note that in the above we can select $q = 1$, which, will result in an the expected loss of the SLE algorithm being 0.5, while an optimal ensemble for this problem can achieve 100% accuracy.

J. Discussion of other related work

In this section, we briefly discuss several other publications somewhat related to the topic of our work.

In the learning scenario we consider, the learner has access to a set of p predictors $h_1, \ldots, h_p$ mapping $\mathcal{X}$ to $\mathcal{Y}$ to devise an accurate ensemble prediction. No other information is available to the learner about these p experts, which are effectively viewed as black boxes. This scenario covers both the case where $h_1, \ldots, h_p$ are hand-crafted prediction rules and the one where they are the hypotheses returned by some learning algorithms trained on samples typically no longer available.

In contrast, most ensemble methods for structured prediction previously presented in the machine learning literature focus on scenarios where the learner can exploit some specific structure of the given experts $h_1, \ldots, h_p$ or where these experts are trained at the same time as the ensemble learner itself (Grubb et al., 2013; Payet & Todorovic, 2010; Tu & Bai, 2010). For example, weak predictors used in the StructuredSpeedBoost algorithm of Grubb et al. (2013) have a very specific structure based on special selection and update rules. Similarly, the RF^2 algorithm of Payet & Todorovic (2010) uses tree experts to make its predictions. Finally, the auto-context algorithm of Tu & Bai (2010) is based on experts that are assumed to be probabilistic models.

Appendix References

Grubb, A., Munoz, D., Bagnell, J.A., and Hebert, M. Speedmachines: Anytime structured prediction. In *Budgeted Learning Workshop, ICML 2013*, 2013.

Payet, N. and Todorovic, S. (RF)^2 - Random Forest Random Field. In *Advances in NIPS*, pp. 1885–1893, 2010.

Tu, Z. and Bai, X. Auto-Context and Its Application to High-Level Vision Tasks and 3D Brain Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32 (10):1744–1757, Oct 2010.