

Online Aggregation of Unbounded Losses Using Shifting Experts with Confidence

Vladimir V'yugin · Vladimir Trunov

Received: date / Accepted: date

Abstract We develop the setting of sequential prediction based on shifting experts and on a “smooth” version of the method of specialized experts. To aggregate experts predictions, we use the AdaHedge algorithm, which is a version of the Hedge algorithm with adaptive learning rate, and extend it by the meta-algorithm Fixed Share. Due to this, we combine the advantages of both algorithms: (1) we use the shifting regret which is a more optimal characteristic of the algorithm; (2) regret bounds are valid in the case of signed unbounded losses of the experts. Also, (3) we incorporate in this scheme a “smooth” version of the method of specialized experts which allows us to make more flexible and accurate predictions. All results are obtained in the adversarial setting – no assumptions are made about the nature of data source. We present results of numerical experiments for short-term forecasting of electricity consumption based on a real data.

1 Introduction

We consider sequential prediction in the general framework of Decision Theoretic Online Learning or the Hedge setting by Freund and Schapire (1997), which is a variant of prediction with expert advice, see e.g. (Littlestone and Warmuth 1994, Freund and Schapire 1997, Vovk 1990, Vovk 1998, Cesa-Bianchi and Lugosi 2006).

The aggregating algorithm updates the experts weights at the end of each trial using losses suffered by the experts in the past. In classical setting (Freund and Schapire 1997, Vovk 1990), the process of an expert i weights updating is based

This paper is an extended version of the conference paper V'yugin (2017). This work was supported by Russian Science Foundation, project 14-50-00150.

Vladimir V'yugin
Institute for Information Transmission Problems, Moscow
E-mail: vyugin@iitp.ru

Vladimir Trunov
Institute for Information Transmission Problems, Moscow
E-mail: trunov@iitp.ru

on exponential weighting with a constant or variable learning rate η :

$$w_{i,t+1} = \frac{w_{i,t} e^{-\eta l_{i,t}}}{\sum_{j=1}^N w_{j,t} e^{-\eta l_{j,t}}}, \quad (1)$$

where $l_{i,t}$ is the loss suffered by the expert i at step t .

The goal of the algorithm is to design weight updates that guarantee that the loss of the aggregating algorithm is never much larger than the loss of the best expert or the best convex combination of the losses of the experts.

So, here the best expert or a convex combination of experts serves as a comparator. By a comparison vector we mean a vector $\mathbf{q} = (q_1, \dots, q_N)$ such that $q_1 + \dots + q_N = 1$ and all its components are nonnegative. We compare the cumulative loss of the aggregating algorithm and a convex combination of the losses $\sum_{t=1}^T (\mathbf{q} \cdot \mathbf{l}_t)$, where $\mathbf{l}_t = (l_{1,t}, \dots, l_{N,t})$ is a vector containing the losses of the experts at time t .

A more challenging goal is to learn well when the comparator \mathbf{q} changes over time, i.e. the algorithm competes with the cumulative sum $\sum_{t=1}^T (\mathbf{q}_t \cdot \mathbf{l}_t)$, where comparison vector \mathbf{q}_t changes over time. An important special case is when \mathbf{q}_t are unit vectors, then the sequence of trials is partitioned into segments. In each segment the loss of the algorithm is compared to the loss of a particular expert and this expert changes at the beginning of a new segment. The goal of the aggregation algorithm is to do almost as well as the sum of losses of experts forming the best partition. Algorithms and bounds for shifting comparators were presented by Herbster and Warmuth (1998). This method called Fixed Share was generalized by Bousquet and Warmuth (2002) to the method of Mixing Past Posteriors (MPP) in which arbitrary mixing schemes are considered. In what follows, MPP mixing schemes will be used in our algorithms.

Most papers in the prediction with expert advice setting either consider uniformly bounded losses or assume the existence of a specific loss function (see Vovk 1990, Cesa-Bianchi and Lugosi 2006). But in some practical applications, this assumption is too restrictive. We allow losses at any step to be unbounded and signed. The notion of a specific loss function is not used.

AdaHedge presented by de Rooij et al. (2014) is among a few algorithms that do not have similar restrictions. This algorithm is a version of the classical Hedge algorithm of Freund and Schapire (1997) and is a refinement of the Cesa-Bianchi and Lugosi (2006) algorithm. AdaHedge is completely parameterless and tunes the learning rate η in terms of a direct measure of past performance.

In de Rooij et al. (2014), an upper bound for regret of this algorithm is presented which is free from boundness assumptions for losses of the experts:

$$R_T \leq 2 \sqrt{S_T \frac{(L_T^* - L_T^-)(L_T^+ - L_T^*)}{L_T^+ - L_T^-} \ln N} + \left(\frac{16}{3} \ln N + 2 \right) S_T, \quad (2)$$

where L_T^* is the loss of the best expert, for other notations see Table 1 below.

In the case where losses of the experts are uniformly bounded the upper bound (2) takes the form $O(\sqrt{T \ln N})$.

We emphasize that the versions of Fixed Share and MPP algorithms presented by Herbster and Warmuth (1998) and Bousquet and Warmuth (2002) use a constant learning rate, while the AdaHedge uses adaptive learning rate which is tuned on-line.

The first contribution of this paper is that we present the ConfHedge-1 algorithm which combines advantages of both these algorithms: (1) we use the shifting regret which is a more optimal characteristic of the algorithm; (2) regret bounds are valid in the case of signed unbounded losses of the experts.

The application we will consider below is the sequential short-term (one-hour-ahead) forecasting of electricity consumption will take place in a variant of the basic problem of prediction with expert advice called prediction with specialized (or sleeping) experts. At each round only some of the experts output a prediction while the other ones are inactive. Each expert is expected to provide accurate forecasts mostly in given external conditions, that can be known beforehand. For instance, in the case of the prediction of electricity consumption, experts can be specialized to a season, temperature, to working days or to public holidays, etc.

The method of specialized experts was first proposed by Freund et al. (1997) and further developed by Adamskiy et al. (2012), Chernov and Vovk (2009), Devaine et al. (2013), Kalnishkan et al. (2015). With this approach, at each step t , a set of specialized experts $E_t \subseteq \{1, \dots, N\}$ is given. A specialized expert i issues its forecasts not at all steps $t = 1, 2, \dots$, but only when $i \in E_t$. At any step, the aggregating algorithm uses forecasts of only “active (non-sleeping)” experts.

The second contribution of this paper is that we have incorporated into ConfHedge-1 a smooth generalization of the method of specialized experts. At each time moment t , we complement the expert i forecast by a confidence level which is a real number $p_{i,t} \in [0, 1]$.

The setting of prediction with experts that report their confidences as a number in the interval $[0, 1]$ was first studied by Blum and Mansour (2007) and further developed by Cesa-Bianchi et al. (2007), Gaillard et al. (2011), Gaillard et al. (2014).

In particular, $p_{i,t} = 1$ means that the expert forecast is used in full, whereas in the case of $p_{i,t} = 0$ it is not taken into account at all (the expert sleeps). In cases where $0 < p_{i,t} < 1$ the expert’s forecast is partially taken into account. For example, with a gradual drop in temperature a corresponded specialized expert gradually loses its ability for accurate predictions of electricity consumption. The dependence of $p_{i,t}$ on values of exogenous parameters can be predetermined by a specialist in the domain or can be constructed using regression analysis on historical data.

In Section 2, we present the ConfHedge-1 algorithm, which is a loss allocation algorithm adapted for the case, where the losses of the experts can be signed and unbounded. Also, this algorithm takes into account the confidence levels of the experts predictions. In Section 3.2, ConfHedge-2 variant of this algorithm is presented for the case when experts make forecasts and calculate their losses using a convex loss function.

In Theorem 1 we present the upper bounds for the shifting regret of these algorithms. The proof of this theorem is given in Section A. Some details of the proof from de Rooij et al. (2014) are presented as a supplementary material in Section B. All results are obtained in the adversarial setting and no assumptions are made about the nature of data source.

N	– number of experts
$\mathbf{l}_t = (l_{1,t}, \dots, l_{N,t})$	– loss vector at step t
$\mathbf{p}_t = (p_{1,t}, \dots, p_{N,t})$	– vector of confidences at step t
$\hat{\mathbf{l}}_t = (\hat{l}_{1,t}, \dots, \hat{l}_{N,t})$	– vector of transformed losses
$l_t^- = \min_{1 \leq i \leq N} l_{i,t}, l_t^+ = \max_{1 \leq i \leq N} l_{i,t}$	– min and max loss
$s_t = l_t^+ - l_t^-$	– loss range
$\mathbf{q}_t = (q_{1,t}, \dots, q_{N,t})$	– comparison vector at step t
$\mathbf{w}_t^\mu = (w_{1,t}^\mu, \dots, w_{N,t}^\mu)$	– experts weights
$\mathbf{w}_t = (w_{1,t}, \dots, w_{N,t})$	– experts posterior weights at step t
$\mathbf{w}_t^* = (w_{1,t}^*, \dots, w_{N,t}^*)$	– the learner prediction, where $w_{i,t}^* = \frac{w_{i,t} p_{i,t}}{\sum_{i=1}^N w_{i,t} p_{i,t}}$ for $1 \leq i \leq N$.
$h_t = (\mathbf{w}_t^* \cdot \mathbf{l}_t)$	– Hedge loss (dot product of two vectors).
$m_t = -\frac{1}{\eta_t} \sum_{i=1}^N w_{i,t} e^{-\eta_t \hat{l}_{i,t}}$	– mixloss
$\delta_t = h_t - m_t$	– mixability gap
α_t	– Fixed Share parameter (we put $\alpha_t = \frac{1}{t}$)
$L_T^- = \sum_{t=1}^T l_t^-, L_T^+ = \sum_{t=1}^T l_t^+$	– cumulative minimal and maximal losses
$S_T = \max_{1 \leq t \leq T} s_t$	– maximum loss range
$H_T = \sum_{t=1}^T h_t$	– algorithm cumulative loss
$M_T = \sum_{t=1}^T m_t$	– cumulative mixloss
$\Delta_T = \sum_{t=1}^T \delta_t$	– cumulative gap
$R_T^{(q)} = \sum_{t=1}^T \sum_{i=1}^N q_{i,t} p_{i,t} (h_t - l_{i,t})$	– confidence shifting regret
$\eta_t = \frac{\ln^* N}{\Delta_{t-1}}$	– variable learning rate, where $\ln^* N = \max\{1, \ln N\}$

Table 1 Basic notations and definitions

In Section 3.3, the techniques of confidence level selection and experts training are presented. We also present the results of numerical experiments of the short-term prediction of electricity consumption with the use of the proposed algorithms.

The approach that sets the confidence levels for expert predictions of electricity consumption is more general than the approach used in the paper Devaine et al. (2013), which uses “sleeping” experts. In our numerical experiments the aggregating algorithm with soft confidence levels outperforms other versions of aggregating algorithms including ones which use sleeping experts.

2 Online loss allocation algorithm

In this section we present an algorithm for the optimal online allocation of unbounded signed losses of the experts. In Section 3.2, a variant of this algorithm will be presented for the case when experts make forecasts and calculate their losses using a convex loss function.

We assume that at each step t , along with the losses $l_{i,t}$ of experts, their confidence levels are given – a vector $\mathbf{p}_t = (p_{1,t}, \dots, p_{N,t})$, where $p_{i,t} \in [0, 1]$ for $1 \leq i \leq N$. We assume that $\|\mathbf{p}_t\|_1 > 0$ for all t .

We can interpret the number $p_{i,t}$ as the algorithm’s internal probability of following the expert i prediction. In this case, we define the auxiliary virtual losses

of the expert as a random variable

$$\tilde{l}_{i,t} = \begin{cases} l_{i,t} & \text{with probability } p_{i,t}, \\ h_t & \text{with probability } 1 - p_{i,t}, \end{cases}$$

where h_t is the aggregating algorithm loss. Denote $\hat{l}_{i,t} = E_{\mathbf{p}_t}[\tilde{l}_{i,t}] = p_{i,t}l_{i,t} + (1 - p_{i,t})h_t$ the mathematical expectation of a virtual loss of an expert i with respect to the probability distribution $\mathbf{p}_{i,t} = (p_{i,t}, 1 - p_{i,t})$.

At any step t we use cumulative weights $w_{i,t}$ of the experts $1 \leq i \leq N$ which were computed at the previous step. The algorithm loss is defined as $h_t = \sum_{i=1}^N w_{i,t} \hat{l}_{i,t}$.

These definitions contain a logical circle – virtual losses are determined through loss of the algorithm, and the latter is determined through virtual losses. Nevertheless, all these quantities can be effectively calculated using the fixed-point method proposed by Chernov and Vovk (2009). We have

$$h_t = \sum_{i=1}^N w_{i,t} \hat{l}_{i,t} = \sum_{i=1}^N w_{i,t} (p_{i,t}l_{i,t} + (1 - p_{i,t})h_t) = \sum_{i=1}^N w_{i,t} p_{i,t} (l_{i,t} - h_t) + h_t.$$

Canceling out the identical terms on the left and on the right sides, we obtain expression for calculating h_t :

$$h_t = \frac{\sum_{i=1}^N w_{i,t} p_{i,t} l_{i,t}}{\sum_{i=1}^N w_{i,t} p_{i,t}} \quad (3)$$

After the value of h_t was calculated by the formula (3), the weights can be calculated as

$$w_{i,t}^\mu = \frac{w_{i,t} e^{-\eta \hat{l}_{i,t}}}{\sum_{s=1}^N w_{s,t} e^{-\eta \hat{l}_{s,t}}} = \frac{w_{i,t} e^{-\eta p_{i,t} (l_{i,t} - h_t)}}{\sum_{s=1}^N w_{s,t} e^{-\eta p_{s,t} (l_{s,t} - h_t)}} \quad (4)$$

using the known value h_t .¹ Also, we compute the mixloss $m_t = -\frac{1}{\eta_t} \sum_{i=1}^N w_{i,t} e^{-\eta_t \hat{l}_{i,t}}$ and the mixability gap $\delta_t = h_t - m_t$, which are used in the construction of the algorithm.

By the method MPP of Bousquet and Warmuth (2002), a mixing scheme is defined by a vector $\beta^{t+1} = (\beta_0^{t+1}, \dots, \beta_t^{t+1})$, where $\sum_{s=0}^t \beta_s^{t+1} = 1$ and $\beta_s^{t+1} \geq 0$ for $0 \leq s \leq t$.

In what follows the vector $\mathbf{w}_t^\mu = (w_{1,t}^\mu, \dots, w_{N,t}^\mu)$ presents the normalized experts weights at step t . The corresponding posterior probability distribution $\mathbf{w}_{t+1} = (w_{1,t+1}, \dots, w_{N,t+1})$ for step $t+1$ is defined as a convex combination $\mathbf{w}_{t+1} = \sum_{s=0}^t \beta_s^{t+1} \mathbf{w}_s^\mu$ with weights β_s^{t+1} , $0 \leq s \leq t$, where $\mathbf{w}_s^\mu = (w_{1,s}^\mu, \dots, w_{N,s}^\mu)$.

The vector β^{t+1} defines the weights by which the past distributions of experts are mixed. It can be re-set at each step t .

¹ In the simple Hedge we put $w_{i,t+1} = w_{i,t}^\mu$. Some other mixing schemes will be given below.

The ConfHedge-1 algorithm for mixing the posteriori distributions of experts is given below. Unlike standard exponential mixing algorithms, this algorithm uses not only the current accumulated weights of experts, but also mixes these weights and all the weights accumulated in past steps.

ConfHedge-1

Put $w_{i,1} = w_{i,0}^\mu = \frac{1}{N}$ for $i = 1, \dots, N$, $\Delta_0 = 0$, $\eta_1 = \infty$.

FOR $t = 1, \dots, T$

Receive confidence levels $\mathbf{p}_t = (p_{1,t}, \dots, p_{N,t})$ of the experts $1 \leq i \leq N$, where $\|\mathbf{p}_t\|_1 > 0$.

Predict with the distribution $\mathbf{w}_t^* = (w_{1,t}^*, \dots, w_{N,t}^*)$, where $w_{i,t}^* = \frac{w_{i,t} p_{i,t}}{\sum_{i=1}^N w_{i,t} p_{i,t}}$ for $1 \leq i \leq N$.

Receive a vector $\mathbf{l}_t = (l_{1,t}, \dots, l_{N,t})$ containing the losses of the experts.

Compute the loss $h_t = (\mathbf{l}_t \cdot \mathbf{w}_t^*)$ of the algorithm.

Update the weights and the learning parameter in three stages:

Loss Update

Define $w_{i,t}^\mu = \frac{w_{i,t} e^{-\eta_t p_{i,t} (l_{i,t}^\mu - h_t)}}{\sum_{s=1}^N w_{s,t} e^{-\eta_t p_{s,t} (l_{s,t}^\mu - h_t)}}$ for $1 \leq i \leq N$.

Mixing Update

Choose a mixing scheme $\beta^{t+1} = (\beta_0^{t+1}, \dots, \beta_t^{t+1})$ and define future weights of the experts

$$w_{i,t+1} = \sum_{s=0}^t \beta_s^{t+1} w_{i,s}^\mu \text{ for } 1 \leq i \leq N.$$

Learning Parameter Update

Define mixloss $m_t = -\frac{1}{\eta_t} \ln \sum_{i=1}^N w_{i,t} e^{-\eta_t (p_{i,t} l_{i,t} + (1-p_{i,t}) h_t)}$. Let $\delta_t = h_t - m_t$ and $\Delta_t = \Delta_{t-1} + \delta_t$. Define the learning rate $\eta_{t+1} = \ln^* N / \Delta_t$ for use at the next step $t+1$.

ENDFOR

We have $m_t \leq h_t$ by convexity of the exponent, then $\delta_t \geq 0$ and $\Delta_t \leq \Delta_{t+1}$ for all t .

We will use the following mixing schemes by Bousquet and Warmuth (2002):

Example 1. A version of Fixed Share by Herbster and Warmuth (1998) (see also Cesa-Bianchi and Lugosi 2006, Vovk 1999) with a variable learning rate is defined by the following mixing scheme. Let a sequence $1 \geq \alpha_1 \geq \alpha_2 \geq \dots > 0$ of parameters be given. Define $\beta_t^{t+1} = 1 - \alpha_{t+1}$ and $\beta_0^{t+1} = \alpha_{t+1}$ ($\beta_s^{t+1} = 0$ for $0 < s < t$). The corresponding prediction for step $t+1$ is defined

$$w_{i,t+1} = \frac{\alpha_{t+1}}{N} + (1 - \alpha_{t+1}) w_{i,t}^\mu$$

for all $1 \leq i \leq N$. In what follows we put $\alpha_t = 1/t$ for all t .

Example 2. Uniform Past by Bousquet and Warmuth (2002) with a variable learning rate. Put $\beta_t^{t+1} = 1 - \alpha_{t+1}$ and $\beta_s^{t+1} = \frac{\alpha_{t+1}}{t}$ for $0 \leq s < t$. The corresponding prediction for step $t+1$ is defined

$$w_{i,t+1} = \alpha_{t+1} \sum_{s=0}^{t-1} \frac{w_{i,s}^\mu}{t} + (1 - \alpha_{t+1}) w_{i,t}^\mu$$

for all i and t .

Bousquet and Warmuth (2002) considered the notion of shifting regret with respect to a sequence $\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_T$ of comparison vectors: $R_T = H_T - \sum_{t=1}^T (\mathbf{q}_t \cdot \mathbf{l}_t)$.²

² The notion of regret with respect to a comparison vector was first defined by Kivinen and Warmuth (1999).

In the presence of confidence values, we consider the corresponding confidence shifting regret $R_T^{(\mathbf{q})} = H_T - L_T^{(\mathbf{q})}$, where

$$L_T^{(\mathbf{q})} = \sum_{t=1}^T (\mathbf{q}_t \cdot \hat{\mathbf{l}}_t) = \sum_{t=1}^T \sum_{i=1}^N q_{i,t} \hat{l}_{i,t} = \sum_{t=1}^T \sum_{i=1}^N q_{i,t} (p_{i,t} l_{i,t} + (1 - p_{i,t}) h_t)$$

and $\hat{\mathbf{l}}_t = (\hat{l}_{1,t}, \dots, \hat{l}_{N,t})$ and $\mathbf{q}_t = (q_{1,t}, \dots, q_{N,t})$ is a comparison vector at step t .

By definition this regret can be represented as

$$R_T^{(\mathbf{q})} = \sum_{t=1}^T \sum_{i=1}^N q_{i,t} p_{i,t} (h_t - l_{i,t}).$$

If $p_{i,t} = 1$ for all i and t then $R_T^{(\mathbf{q})} = R_T$.

The quantity $L_T^{(\mathbf{q})}$ depends on h_t . To avoid this dependence, we will consider its lower and upper bounds: $L_T^{(\mathbf{q}^-)} \leq L_T^{(\mathbf{q})} \leq L_T^{(\mathbf{q}^+)}$, where

$$\begin{aligned} L_T^{(\mathbf{q}^-)} &= \sum_{t=1}^T \sum_{i=1}^N q_{i,t} (p_{i,t} l_{i,t} + (1 - p_{i,t}) l_t^-), \\ L_T^{(\mathbf{q}^+)} &= \sum_{t=1}^T \sum_{i=1}^N q_{i,t} (p_{i,t} l_{i,t} + (1 - p_{i,t}) l_t^+). \end{aligned} \quad (5)$$

Assume that the losses of the experts are bounded, for example, $l_{i,t} \in [0, 1]$ for all i and t . Using the techniques of Section A for $\eta_t \sim \sqrt{\frac{\ln^* N}{t}}$, we can prove that

$$R_T^{(\mathbf{q})} = O\left((k+1) \left(\ln T \sqrt{T} + \sqrt{T \ln^* N}\right)\right). \quad (6)$$

where k is the number of switches of comparison vectors \mathbf{q}_t on the time interval $1 \leq t \leq T$.³

Our goal is to obtain a similar bound in the absence of boundness assumptions for the expert losses.

The following theorem presents the upper bounds for the confidence shifting regret in the case where no assumptions are made about boundness of the losses of the experts.

Theorem 1 *Let the mixing scheme from Example 1 was used. Then for any T and for any sequence $\mathbf{q}_1, \dots, \mathbf{q}_T$ of comparison vectors,*

$$R_T^{(\mathbf{q})} \leq \frac{1}{2} \gamma_{k,T} \sqrt{\sum_{t=1}^T s_t^2 \ln^* N} + \gamma_{k,T} \left(\frac{2}{3} \ln^* N + 1 \right) S_T, \quad (7)$$

$$\begin{aligned} R_T^{(\mathbf{q})} &\leq \gamma_{k,T} \sqrt{S_T \frac{(L_T^+ - L_T^{(\mathbf{q}^-)})(L_T^{(\mathbf{q}^+)} - L_T^-)}{L_T^+ - L_T^-} \ln^* N} + \\ &\quad \gamma_{k,T} \left(\left(\gamma_{k,T} + \frac{2}{3} \right) \ln^* N + 1 \right) S_T, \end{aligned} \quad (8)$$

³ Does this bound is tight is an open question. Some lower bounds for mixloss (for the logarithmic loss function with the learning rate $\eta = 1$) were obtained by Adamkiy et al. (2012). They show an information-theoretic lower bound for mixloss that must hold for any algorithm, and which is tight for Fixed Share.

where $\gamma_{k,T} = (k+2)(\ln T + 1)$ and k is the number of switches of the comparison vectors on the time interval $1 \leq t \leq T$.

The bound (7) is an analogue for the shifting experts of the bound from Cesa-Bianchi et al. (2007) and the bound (8) is an analogue of the bound (16) of Theorem 8 from de Rooij et al. (2014). Proof of Theorem 1 is given in Sections A and B.

A disadvantage of the bounds (8) and (9) below is in the presence of a term that depends quadratically on the number k of switches. Whether such a dependence is necessary is an open question. However, this term does not depend on the loss of the algorithm, it has only a slowly growing multiplicative factor $O(\ln^2 T)$. Corollary 1 below shows that in some special cases this dependence can be eliminated.

The bound (8) of Theorem 1 can be simplified in the different ways:

Corollary 1 *For any T and for any sequence $\mathbf{q}_1, \dots, \mathbf{q}_T$ of comparison vectors,*

$$R_T^{(\mathbf{q})} \leq \gamma_{k,T} \sqrt{S_T(L_T^{(\mathbf{q}^+)} - L_T^-) \ln^* N} + \gamma_{k,T} \left(\left(\gamma_{k,T} + \frac{2}{3} \right) \ln^* N + 1 \right) S_T, \quad (9)$$

$$R_T^{(\mathbf{q})} \leq \gamma_{k,T} \sqrt{S_T(L_T^+ - L_T^{(\mathbf{q}^-)}) \ln^* N} + \gamma_{k,T} \left(\frac{2}{3} \ln^* N + 1 \right) S_T, \quad (10)$$

$$R_T^{(\mathbf{q})} \leq \gamma_{k,T} \sqrt{S_T(L_T^+ - L_T^-) \ln^* N} + \gamma_{k,T} \left(\frac{2}{3} \ln^* N + 1 \right) S_T. \quad (11)$$

The bound (10) linearly depends on the number of switches (for the proof see Section B). The bound (11) follows from (10). If the losses of the experts are uniformly bounded then the bound (11) is of the same order as the bound (6).

An important special case of Theorem 1 is when the comparison vectors $\mathbf{q}_t = \mathbf{e}_{i_t}$ are unit vectors and $p_{i,t} \in \{0, 1\}$, i.e., the specialists case is considered for composite experts i_1, \dots, i_T . Then the confidence shifting regret equals $R_T^{(\mathbf{q})} = \sum_{t: p_{i_t,t}=1}^T (h_t - l_{i_t,t})$ and the corresponding differences in the right-hand side of in-

equality (8) are $L_T^+ - L_T^{(\mathbf{q}^-)} = \sum_{t: p_{i_t,t}=1}^T (l_t^+ - l_{i_t,t}) + \sum_{t: p_{i_t,t}=0}^T s_t$ and $L_T^{(\mathbf{q}^+)} - L_T^- = \sum_{t: p_{i_t,t}=1}^T (l_{i_t,t} - l_t^-) + \sum_{t: p_{i_t,t}=0}^T s_t$.

The bound (10) is important if the algorithm is to be used for a scenario in which we are provided with a sequence of gain vectors \mathbf{g}_t rather than losses: we can transform these gains into losses using $\mathbf{l}_t = -\mathbf{g}_t$, and then run the algorithm. Assume that $p_{i,t} = 1$ for all i and t . The bound then implies that we incur small regret with respect to a composite expert if it has very small cumulative gain relative to the minimum gain (see also de Rooij et al. 2014).

The similar bounds for the mixing scheme of Example 2 also can be obtained, where $\gamma_{k,T} = (2k+3) \ln T + (k+2)$ (see Section A).

3 Numerical experiments

Section 3.1 presents the results of applying ConfHedge-1 to synthetic data. In Section 3.3 the results of the short-term prediction of electricity consumption

are presented. We use in these experiments the ConfHedge-2 algorithm which is a variant the previous algorithm adapted for the case, where experts present the numerical forecasts. The scheme of this algorithm is given in Section 3.2.

3.1 Unbounded signed losses

The first experiment was performed on synthetic data, where one-step losses of experts are signed, unbounded and perturbed by $N(0, 1)$ additive noise. Confidence levels of all experts are always equal to one. Figure 1a shows mean values of these one-step expert losses. Figure 1b shows cumulative losses of three individual experts and cumulative losses of AdaHedge and ConfHedge-1. These experiments show that ConfHedge-1 is non-inferior to AdaHedge, and, after some time, even outperforms it.




Fig. 1 Results of the experiment on the synthetic data. Left subfigure (a) shows the mean values of one-step experts losses (lines 1, 2, and 3). Right subfigure (b) shows cumulative losses of individual experts (thin lines 1, 2 and 3), AdaHedge and ConfHedge-1 cumulative losses (thick lines 4 and 5)

3.2 Aggregation of expert forecasts

In this section we suppose that the losses of the experts are computed using a convex in γ loss function $\lambda(\omega, \gamma)$, where ω is an outcome and γ is a forecast. Outcomes can belong to an arbitrary set, forecasts form a linear space.⁴

Let at any step t the experts forecasts $\mathbf{c}_t = (c_{1,t}, \dots, c_{N,t})$ and their confidence levels $\mathbf{p}_t = (p_{1,t}, \dots, p_{N,t})$ are given. Here $p_{i,t} \in [0, 1]$ for all $1 \leq i \leq N$. Define

⁴ In our experiments, the absolute loss function $\lambda(\omega, \gamma) = |\omega - \gamma|$ was used, where ω and γ are real numbers. In practical applications, we can also use its biased variant $\lambda(\omega, \gamma) = \mu_1|\omega - \gamma|_- + \mu_2|\omega - \gamma|_+$, where $|r|_- = -\min\{0, r\}$ and $|r|_+ = \max\{0, r\}$. The positive numbers μ_1 and μ_2 provide a balance of losses between the deviations of the forecasts γ and outcomes ω in the positive and negative directions.

the auxiliary virtual experts forecasts

$$\tilde{c}_{i,t} = \begin{cases} c_{i,t} & \text{with probability } p_{i,t}, \\ \gamma_t & \text{with probability } 1 - p_{i,t}, \end{cases}$$

where γ_t is a forecast of the aggregating algorithm. Then the mathematical expectation of any expert i forecast is equal to $\hat{c}_{i,t} = E_{\mathbf{p}_{i,t}}[\tilde{c}_{i,t}] = p_{i,t}c_{i,t} + (1 - p_{i,t})\gamma_t$.

Define the aggregating algorithm forecast

$$\gamma_t = \sum_{i=1}^N w_{i,t} \hat{c}_{i,t}. \quad (12)$$

In order to get rid of the logical circle in these definitions, we use the fixed point method by Chernov and Vovk (2009). We have

$$\gamma_t = \sum_{i=1}^N w_{i,t} \hat{c}_{i,t} = \sum_{i=1}^N w_{i,t} (p_{i,t}c_{i,t} + (1 - p_{i,t})\gamma_t) = \sum_{i=1}^N w_{i,t} p_{i,t} (c_{i,t} - \gamma_t) + \gamma_t.$$

Cancel the same terms on the left and on the right sides and obtain

$$\gamma_t = \frac{\sum_{i=1}^N p_{i,t} w_{i,t} c_{i,t}}{\sum_{i=1}^N p_{i,t} w_{i,t}}. \quad (13)$$

The further calculations are given in the scheme of ConfHedge-2 below.

ConfHedge-2

Define $w_{i,1} = w_{i,0}^\mu = \frac{1}{N}$ for $i = 1, \dots, N$, $\Delta_0 = 0$, $\eta_1 = \infty$.

FOR $t = 1, \dots, T$

Receive the expert forecasts $\mathbf{c}_t = (c_{1,t}, \dots, c_{N,t})$ and and their confidence levels $\mathbf{p}_t = (p_{1,t}, \dots, p_{N,t})$.

Compute the aggregating algorithm forecast $\gamma_t = \frac{\sum_{i=1}^N p_{i,t} w_{i,t} c_{i,t}}{\sum_{i=1}^N p_{i,t} w_{i,t}}$.

Receive an outcome ω_t and compute the experts losses $\mathbf{l}_t = (l_{1,t}, \dots, l_{N,t})$, where $l_{i,t} = \lambda(\omega_t, c_{i,t})$, $1 \leq i \leq N$, and the algorithm loss $a_t = \lambda(\omega_t, \gamma_t)$.

Update experts weights and learning parameter in three stages:

Loss Update

Define

$$w_{i,t}^\mu = \frac{w_{i,t} e^{-\eta_t p_{i,t} (l_{i,t} - a_t)}}{\sum_{s=1}^N w_{s,t} e^{-\eta_t p_{s,t} (l_{s,t} - a_t)}} \text{ for } 1 \leq i \leq N.$$

Mixing Update

Define future experts weights $w_{i,t+1} = \frac{\alpha_{t+1}}{N} + (1 - \alpha_{t+1})w_{i,t}^\mu$ for $1 \leq i \leq N$, where $\alpha_t = \frac{1}{t}$.

Learning Parameter Update

Compute the mixloss

$$m_t = -\frac{1}{\eta_t} \ln \sum_{i=1}^N w_{i,t} e^{-\eta_t (p_{i,t} l_{i,t} + (1 - p_{i,t}) a_t)}.$$

Define $\delta_t = h_t - m_t$, where $h_t = \sum_{i=1}^N w_{i,t} (p_{i,t} l_{i,t} + (1 - p_{i,t}) a_t)$, define also $\Delta_t = \Delta_{t-1} + \delta_t$.

After that set $\eta_{t+1} = \ln^* N / \Delta_t$ future value of the learning parameter.

ENDFOR

Let $A_T = \sum_{t=1}^T a_t$ be the loss of ConfHedge-2. We keep the notation $H_T = \sum_{t=1}^T h_t$ and $L_T^{(\mathbf{q})} = \sum_{t=1}^T (\mathbf{q}_t \cdot \hat{\mathbf{l}}_t)$. Theorem 1 also holds for these quantities. Hence, using the same notation as in the Section 2, we obtain a bound (16) and $H_T - L_T^{(\mathbf{q})} \leq \gamma_{k,T} \Delta_T$.

Since by convexity of the loss function $a_t = \lambda(\omega_t, \gamma_t) = \lambda(\omega_t, \sum_{i=1}^N w_{i,t} \hat{c}_{i,t}) \leq \sum_{i=1}^N w_{i,t} \hat{l}_{i,t} = h_t$ for all t , we have $A_T \leq H_T$.

The confidence shifting regret of ConfHedge-2 is equal to

$$R_T^{(\mathbf{q})} = A_T - L_T^{(\mathbf{q})} = \sum_{t=1}^T \sum_{i=1}^N q_{i,t} p_{i,t} (a_t - l_{i,t}).$$

The upper bounds (8) of this regret are given by Theorem 1 and by Corollary 1.

3.3 The electrical loads forecasting

The second group of numerical experiments were performed with the contest data of the GefCom2012 competition conducted on the Kaggle platform (Tao Hong et al. 2014). The main objective of this competition was to predict the daily course of hourly electrical loads (demand values for electricity) in 20 regions according to temperature records at 11 meteorological stations. Databases are available at <http://www.kaggle.com/datasets>. The basic data were provided in the form of the table “temperature-history“ with archive records of temperature monitoring at 11 meteorological stations and the table “load-history“ with hourly electrical load data recorded at 20 power distribution stations of the region for the period from 01.01.2004 to 30.06.2008. The additional calendar information (seasons, days of the week, and working days vs. holidays) could be also used.

As an illustration of how the proposed expert aggregation methods perform, a simplified particular task was designed, namely, electrical load forecasting in one of the power distribution networks (Zone5) one hour ahead based on historical data and current calendar parameters. To account for temperature changes, the temperature measurements of only one meteorological station (Station9) were used; these data provided the best electrical load forecasts in the selected network on the training part of the sample.

Figure 2 shows the averaged curves of the daily electrical loads for each of the four seasons of 2004–2005 in the selected network. We see that the course of the averaged curves clearly depends on the time of day, and also varies from season to season. In addition, the working day and weekend day patterns demonstrate distinct differences in the level of electricity usage. Based on this figure, a simple scheme of forming an ensemble of experts, i.e., specialized algorithms that can only process strictly defined data, was chosen; the scheme includes the following categories: four times of day (night, morning, day, evening); working days and weekend days (two categories); four seasons (winter, spring, summer, fall), all this giving $4 \times 2 \times 4 = 32$ specialized experts (Stepwise Linear Regression). We also use extra four experts, each of which is focused on one of the seasons of the year, and one nonsleeping expert (Random Forest algorithm). Thus, we used a total of 37 experts.

At each moment of time, the confidence function of a given expert is calculated as a product of the confidence functions for each of its specializations. For example, Figure 3 shows the stages of constructing the confidence function for the expert focused on night forecasting (0–6 a.m.) on the working days of January. Thus synthesized confidence functions are used to form individual training samples for

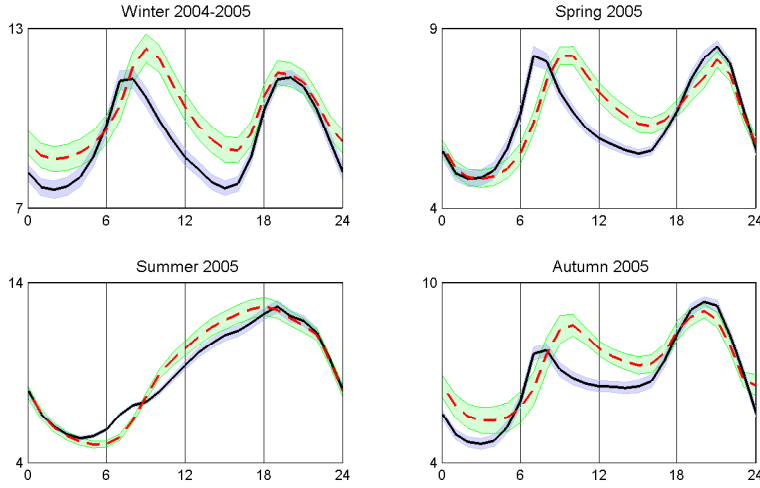


Fig. 2 Averaged curves of the daily electrical loads for each of the four seasons of 2004-2005. The color band around each curve represents standard error of the mean. The solid lines show the average level of electricity usage for working days, and the dashed lines show the same estimates for the weekend days of the same season

each expert at the stage of training and to aggregate expert forecasts at the stage of testing.

To ensure a more smooth switch between experts, the membership functions $p_{i,t}$ were formed as trapezoids, where the function takes the value 1 on the plateau corresponding to the selected calendar interval, and varies linearly from 1 to 0 on the slopes. The slope width depends on the user defined parameter.

At the stage of training, the following steps are taken for each algorithm (expert): (1) For all elements of the training sample, a confidence level is calculated, assuming its values are close to 1, if the sample is to be considered by the expert, or close to 0, if the object is beyond its specialization. Based on the confidence level, an individual training subsample is formed for each algorithm from the full training sample. (2) Based on this individual training subsample, a forecasting model is constructed.

To compare the scheme of “smooth mixing” with the scheme of “sleeping experts”, the experiments on expert decision aggregation were performed in two stages. First, only the scheme of mixing the sleeping and awake experts was used, i.e., the confidence level took only two values (0 or 1), and then the mixing algorithm from Section 2 of this work was used.

The evolution of differences of cumulative losses $L_T^1 - L_T^3$ and $L_T^2 - L_T^3$, where L_T^1 is the cumulative loss of anytime nonsleeping Random Forest algorithm and L_T^2 , L_T^3 are cumulative losses of two schemes of mixing (“sleeping experts” and “smooth mixing”), are shown in Figure 4a.




Fig. 3 Confidence function construction for the expert $\text{January} \times \text{Working days} \times \text{Night}$




Fig. 4 (a) The evolution of differences of cumulative losses $L_T^1 - L_T^3$ and $L_T^2 - L_T^3$: 1 – anytime nonsleeping expert (Random Forest algorithm), 2 – ConfHedge-2 using “sleeping experts” model, 3 – ConfHedge-2 using smooth confidence levels. (b) The mean cumulative losses (MAE) of Random Forest (1) and of two schemes of expert mixing (2 and 3)

The mean cumulative losses (Mean Absolute Error – MAE) $\frac{1}{T}L_T^1$ of Random Forest algorithm and of two schemes of expert mixing: $\frac{1}{T}L_T^2$ and $\frac{1}{T}L_T^3$, are shown in Figure 4b.⁵

⁵ The absolute loss function was used in these experiments.

In this experiment, the “smooth mixing” algorithm outperforms the aggregating algorithm using “sleeping experts” and both these algorithms outperform the anytime Random Forest forecasting algorithm.

4 Conclusion

In this paper we extend the AdaHedge algorithm by de Rooij et al. (2014) for a case of shifting experts and for a smooth version of the method of specialized experts, where at any time moment each expert’s forecast is provided with a confidence level which is a number between 0 and 1.

To aggregate experts predictions, we use methods of shifting experts and the algorithm AdaHedge with an adaptive learning rate. Due to this, we combine the advantages of both algorithms. We use the shifting regret which is a more optimal characteristic of the algorithm, and we do not impose restrictions on the expert losses. Also, we incorporate in this scheme a smooth version of the method of specialized experts by Blum and Mansour (2007), which allows us to make more flexible and accurate predictions.

We obtained the new upper bounds for the regret of our algorithms, which generalize similar upper bounds for the case of specialized experts.

A disadvantage of Theorem 1 and of Corollary 1 is in asymmetry of the bounds (9) and (10) – first of them has a term that depends quadratically on the number k of switches. Whether such a dependence is necessary is an open question.

All results are obtained in the adversarial setting, no assumptions are made about the nature of data source.

We present the results of numerical experiments on short-term forecasting of electricity consumption based on a real data. In these experiments, the “smooth mixing” algorithm outperforms the aggregating algorithm with “sleeping experts” and both these algorithms outperform the anytime Random Forest forecasting algorithm.

A Main lemma

For analysis of the mixing schemes, following Bousquet and Warmuth (2002), we use the notion of relative entropy

$$D(\mathbf{p} \parallel \mathbf{q}) = \sum_{i=1}^n p_i \ln \frac{p_i}{q_i},$$

where n is an arbitrary positive integer number, $\mathbf{p} = (p_1, \dots, p_n)$, $\mathbf{q} = (q_1, \dots, q_n)$ are elements of the n -dimensional simplex of all probability distributions on a set of cardinality n . Put $0 \ln 0 = 0$.

Consider some properties of the relative entropy. The inequalities $\mathbf{p} > \mathbf{q}$, $\mathbf{p} \geq \mathbf{q}$, $\mathbf{p} \geq \mathbf{0}$ for vectors will be understood componentwise; here $\mathbf{0}$ is the vector with zeros components.

Lemma 1 *Bousquet and Warmuth 2002* For each $\mathbf{p}, \mathbf{q}, \mathbf{w}$ such that $\mathbf{q}, \mathbf{w} > \mathbf{0}$,

- $D(\mathbf{p} \parallel \mathbf{q}) \leq D(\mathbf{p} \parallel \mathbf{w}) + \ln \left(\sum_{i=1}^n p_i \frac{w_i}{q_i} \right)$.
- If $\mathbf{q} \geq r\mathbf{w}$ for some real number $r > 0$ then $D(\mathbf{p} \parallel \mathbf{q}) \leq D(\mathbf{p} \parallel \mathbf{w}) + \ln \frac{1}{r}$. In particular, for $\mathbf{p} = \mathbf{w}$ we have $D(\mathbf{w} \parallel \mathbf{q}) \leq \ln \frac{1}{r}$ for each $\mathbf{q} \geq r\mathbf{w}$.

– Let \mathbf{p} be a probability vector, $\mathbf{q} = \sum_{i=0}^t \beta_i \mathbf{w}_i$, where $\mathbf{w}_i > 0$ for $0 \leq i \leq t$, $\beta = (\beta_0, \dots, \beta_t)$, $\sum_{i=0}^t \beta_i = 1$, and $\beta > \mathbf{0}$. Then $D(\mathbf{p} \parallel \mathbf{q}) \leq D(\mathbf{p} \parallel \mathbf{w}_i) + \ln \frac{1}{\beta_i}$ for each i . In particular, if $\mathbf{p} = \mathbf{q}$ then $D\left(\mathbf{w}_i \parallel \sum_{i=0}^t \beta_i \mathbf{w}_i\right) \leq \ln \frac{1}{\beta_i}$ for all i .

Proof. From concavity of the logarithm, we have

$$D(\mathbf{p} \parallel \mathbf{q}) - D(\mathbf{p} \parallel \mathbf{w}) = \sum_{i=1}^n p_i \ln \frac{w_i}{q_i} \leq \ln \left(\sum_{i=1}^n p_i \frac{w_i}{q_i} \right). \quad (14)$$

If $\mathbf{q} \geq r\mathbf{w}$ then $\sum_{i=1}^n p_i \frac{w_i}{q_i} \leq \sum_{i=1}^n p_i \frac{w_i}{r w_i} = \frac{1}{r}$. \square

The notion of mixloss $m_t = -\frac{1}{\eta_t} \sum_{i=1}^N w_{i,t} e^{-\eta_t l_{i,t}}$ and its cumulative variant $M_T = \sum_{t=1}^T m_t$ are used in Hedge analysis.⁶ By definition $m_t \leq h_t$ for all t .

Lemma 2 (Bousquet and Warmuth 2002) For any comparison vector \mathbf{q}_t ,

$$m_t - \sum_{i=1}^N q_{i,t} \hat{l}_{i,t} = \frac{1}{\eta_t} (D(\mathbf{q}_t \parallel \mathbf{w}_t) - D(\mathbf{q}_t \parallel \mathbf{w}_t^\mu)). \quad (15)$$

Proof. By (14),

$$\begin{aligned} m_t - \sum_{i=1}^N q_{i,t} \hat{l}_{i,t} &= \frac{1}{\eta_t} \sum_{i=1}^N q_{i,t} \left(\ln e^{-\eta_t \hat{l}_{i,t}} - \ln \sum_{j=1}^N w_{j,t} e^{-\eta_t \hat{l}_{j,t}} \right) = \\ &= \frac{1}{\eta_t} \sum_{i=1}^N q_{i,t} \ln \frac{e^{-\eta_t \hat{l}_{i,t}}}{\sum_{j=1}^N w_{j,t} e^{-\eta_t \hat{l}_{j,t}}} = \frac{1}{\eta_t} \sum_{i=1}^N q_{i,t} \ln \frac{w_{i,t}^\mu}{w_{i,t}} = \frac{1}{\eta_t} (D(\mathbf{q}_t \parallel \mathbf{w}_t) - D(\mathbf{q}_t \parallel \mathbf{w}_t^\mu)). \quad \square \end{aligned}$$

The following lemma presents a bound for the confidence regret in terms of cumulative mixability gap.

Lemma 3 Let $\alpha_t = \frac{1}{t}$ for all t and the mixing scheme from Example 1 was used. Then for any T , for any sequence of losses of the experts, and for any sequence of comparison vectors \mathbf{q}_t given on-line with no more than k switches on the time interval $1 \leq t \leq T$,

$$H_T - L_T^{(\mathbf{q})} \leq (k+2)(\ln T + 1) \Delta_T. \quad (16)$$

Proof. We apply Lemmas 1 and 2 for mixing schemes of Example 1 (Fixed Share).

Let a sequence $\mathbf{l}_t = (l_{1,t}, \dots, l_{N,t})$ of losses of the experts and a sequence of comparison vectors $\mathbf{q}_t = (q_{1,t}, \dots, q_{N,t})$ be given on-line for $t = 1, 2, \dots$. Assume that T be an arbitrary and the comparison vector \mathbf{q}_t changes k times for $1 \leq t \leq T$.

We let $1 < t_1 < t_2 < \dots < t_k$ be the subsequence of indices in the sequence of comparators $\mathbf{q}_1, \dots, \mathbf{q}_T$, where shifting occurs: $\mathbf{q}_{t_j} \neq \mathbf{q}_{t_j-1}$ and $\mathbf{q}_t = \mathbf{q}_{t-1}$ for all other steps, where $t > 1$. Define also $t_0 = 1$ and $t_{k+1} = T + 1$. We apply Lemma 2 for the distribution β^{t+1} from Example 1. Recall that $w_{i,1} = w_{i,0}^\mu = \frac{1}{N}$ for $i = 1, \dots, N$.

⁶ Mixloss is a very useful intermediate concept, cumulative variant of which is less or equal to the cumulative loss of the best expert (up to a small term) and, on the other hand, the cumulative mixloss is close to the cumulative loss of the aggregating algorithm. For the logarithmic loss function, the mixloss coincides with the loss of the Vovk aggregating algorithm (see Adamski et al. 2012, Cesa-Bianchi and Lugosi 2006, de Rooij et al. 2014).

Summing (15) on time interval where $\mathbf{q}_t = \mathbf{q}_{t-1}$ for $t_j + 1 \leq t \leq t_{j+1} - 1$, we obtain

$$\begin{aligned} \sum_{t=t_j+1}^{t_{j+1}-1} \left(m_t - \sum_{i=1}^N q_{i,t} \hat{l}_{i,t} \right) &= \sum_{t=t_j+1}^{t_{j+1}-1} \frac{1}{\eta_t} (D(\mathbf{q}_t \| \mathbf{w}_t) - D(\mathbf{q}_t \| \mathbf{w}_t^\mu)) = \\ \sum_{t=t_j+1}^{t_{j+1}-1} \left(\frac{1}{\eta_{t-1}} D(\mathbf{q}_t \| \mathbf{w}_t) - \frac{1}{\eta_t} D(\mathbf{q}_t \| \mathbf{w}_t^\mu) \right) &+ \sum_{t=t_j+1}^{t_{j+1}-1} \left(\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) D(\mathbf{q}_t \| \mathbf{w}_t) \leq \end{aligned} \quad (17)$$

$$\sum_{t=t_j+1}^{t_{j+1}-1} \left(\frac{1}{\eta_{t-1}} D(\mathbf{q}_t \| \mathbf{w}_t) - \frac{1}{\eta_t} D(\mathbf{q}_t \| \mathbf{w}_t^\mu) \right) + \sum_{t=t_j+1}^{t_{j+1}-1} \frac{1}{\ln^* N} \delta_{t-1} \left(\ln N + \ln \frac{1}{\alpha_t} \right) \leq \quad (18)$$

$$\begin{aligned} &\sum_{t=t_j+1}^{t_{j+1}-1} \left(\frac{1}{\eta_{t-1}} D(\mathbf{q}_t \| \mathbf{w}_{t-1}^\mu) - \frac{1}{\eta_t} D(\mathbf{q}_t \| \mathbf{w}_t^\mu) \right) + \\ &\sum_{t=t_j+1}^{t_{j+1}-1} \frac{1}{\eta_{t-1}} \ln \frac{1}{1 - \alpha_t} + \sum_{t=t_j+1}^{t_{j+1}-1} \left(\frac{1}{\ln^* N} \ln \frac{1}{\alpha_t} + 1 \right) \delta_{t-1} \leq \end{aligned} \quad (19)$$

$$\begin{aligned} &\frac{1}{\eta_{t_j}} D(\mathbf{q}_{t_j} \| \mathbf{w}_{t_j}^\mu) - \frac{1}{\eta_{t_{j+1}-1}} D(\mathbf{q}_{t_j} \| \mathbf{w}_{t_{j+1}-1}^\mu) + \\ &\sum_{t=t_j+1}^{t_{j+1}-1} \frac{1}{\eta_t} \ln \frac{1}{1 - \alpha_t} + \sum_{t=t_j+1}^{t_{j+1}-1} \delta_{t-1} \ln \frac{1}{\alpha_t} + \sum_{t=t_j+1}^{t_{j+1}-1} \delta_{t-1}. \end{aligned} \quad (20)$$

In transition from (17) to (18), the inequality $w_{i,t} \geq \frac{\alpha_t}{N}$ was used, then

$$D(\mathbf{q}_t \| \mathbf{w}_t) = \sum_{i=1}^N q_{i,t} \ln \frac{q_{i,t}}{w_{i,t}} \leq \sum_{i=1}^N q_{i,t} \ln q_{i,t} - \sum_{i=1}^N q_{i,t} \ln \frac{\alpha_t}{N} \leq \ln N + \ln \frac{1}{\alpha_t}. \quad (21)$$

In transition from (18) to (19), we use the inequality (14), where $s = t - 1$,

$$D(\mathbf{q}_t \| \mathbf{w}_t) \leq D(\mathbf{q}_t \| \mathbf{w}_{t-1}^\mu) + \ln \frac{1}{1 - \alpha_t}.$$

In transition from (19) to (20) the entropy terms within the sections telescope and only for the beginning and the end of each section a positive and a negative entropy term remains, respectively. We also roughen the inequality (19) after division by $\ln^* N$.

For the beginnings of the k sections $t = t_1, \dots, t_k$ define $s = 0$, $\beta_0^{t_j} = \alpha_{t_j}$ in the inequality (14), then

$$m_{t_j} - \sum_{i=1}^N q_{i,t_j} \hat{l}_{i,t_j} \leq \frac{1}{\eta_{t_j}} D(\mathbf{q}_{t_j} \| \mathbf{w}_0^\mu) - \frac{1}{\eta_{t_j}} D(\mathbf{q}_{t_j} \| \mathbf{w}_{t_j}^\mu) + \frac{1}{\eta_{t_j}} \ln \frac{1}{\alpha_{t_j}}. \quad (22)$$

Summing all these inequalities and canceling out the corresponding terms, we obtain

$$\sum_{t=1}^T m_t - \sum_{i=1}^N q_{i,t} \hat{l}_{i,t} \leq \sum_{j=1}^k \left(\frac{1}{\eta_{t_j}} D(\mathbf{q}_{t_j} \| \mathbf{w}_0^\mu) - \frac{1}{\eta_{t_{j+1}-1}} D(\mathbf{q}_{t_j} \| \mathbf{w}_{t_{j+1}-1}^\mu) \right) + \quad (23)$$

$$\sum_{t=2}^T \frac{1}{\eta_t} \ln \frac{1}{1 - \alpha_t} + \sum_{t=2}^T \delta_{t-1} \ln \frac{1}{\alpha_t} + \sum_{t=2}^T \delta_{t-1} + \sum_{j=1}^k \Delta_{t_j-1} \ln \frac{1}{\alpha_{t_j}} \leq \quad (24)$$

$$((k+2) \ln T + k+1) \Delta_T. \quad (25)$$

In transition from (23) to (24) we use the inequality $D(\mathbf{q} \| \mathbf{w}_T^\mu) \geq 0$ for all \mathbf{q} and equality

$D(\mathbf{q} \| \mathbf{w}_0^\mu) = \ln N$. Then $\sum_{j=1}^k \frac{1}{\eta_{t_j}} D(\mathbf{q}_{t_j} \| \mathbf{w}_0^\mu) \leq k \Delta_T$. For $\alpha_t = \frac{1}{t}$ we use the inequality

$$\sum_{t=2}^T \frac{1}{\eta_t} \ln \frac{1}{1 - \alpha_t} \leq \frac{1}{\eta_T} \ln T \leq \Delta_{T-1} \ln T \leq \Delta_T \ln T.$$

Since $H_T = M_T + \Delta_T$, the bound (25) implies (16). \square

We will finish the proof of Theorem 1 at the end of Section B.

The corresponding bounds for mixing scheme of Example 2 can be obtained in a similar way. Since by definition $w_{i,t} \geq \frac{\alpha_t}{Nt}$ for each t , the inequality (21) is changed to $D(\mathbf{q}_t \| \mathbf{w}_t) \leq \ln N + \ln T + \ln \frac{1}{\alpha_t}$. Also, the last term of the inequality (22) is replaced by $\frac{1}{\eta_{t_j}} \ln \frac{1}{t\alpha_{t_j}}$. As a result, we obtain $\gamma_{k,T} = (2k+3) \ln T + (k+2)$.

In the case of bounded losses: $l_{i,t} \in [0, 1]$, set in (17) – (25) $\alpha_t = \frac{1}{t}$ and obtain $M_T - L_T^{(\mathbf{q})} \leq \frac{1}{\eta_T} ((k+2) \ln T + (k+1) \ln N)$. Using the Hoeffding inequality $h_t \leq m_t + \frac{\eta_t}{8}$, where $\eta_t \sim \sqrt{\frac{\ln^* N}{t}}$, we obtain (6).

B Technical bounds

The derivation of the upper bound for Δ_T is similar to that given in de Rooij et al. (2014), except that the losses of experts are replaced by $\hat{l}_{i,t} = E_{\mathbf{p}_{i,t}}[\tilde{l}_{i,t}]$.

$$\text{Let } v_t = E_{j \sim \mathbf{w}_t}[(\hat{l}_{j,t} - E_{j \sim \mathbf{w}_t}[\hat{l}_{j,t}])^2] = \sum_{j=1}^N w_{j,t}(\hat{l}_{j,t} - h_t)^2 \text{ and } V_T = \sum_{t=1}^T v_t.$$

Lemma 4 *The quantity δ_t satisfies the inequality*

$$\delta_t \leq \frac{e^{s_t \eta_t} - 1 - s_t \eta_t}{\eta_t s_t^2} v_t. \quad (26)$$

Proof. The inequality (26) will be proved using Bernstein inequality (see Lemmas 3-5 from Cesa-Bianchi and Lugosi 2006). Let $X \in [0, 1]$ be a random variable and $\text{Var}[X]$ be its variance. Then for any $\eta > 0$, we have $\ln E[e^{-\eta(X - E[X])}] \leq \text{Var}[X](e^\eta - \eta - 1)$.

Consider a random variable which takes the values $\hat{l}_{j,t}$ with probabilities $w_{j,t}$, where $j = 1, \dots, N$. Let us transform it so that its values belong to the segment $[0, 1]$: $X_t^j = \frac{\hat{l}_{j,t} - \hat{l}_t^-}{s_t}$. Then the Bernstein inequality can be written as follows:

$$\ln E_{j \sim \mathbf{w}_t} \left[e^{-\eta(X_t^j - EX_t^j)} \right] \leq \text{Var}_{j \sim \mathbf{w}_t}[X_t^j] (e^\eta - 1 - \eta) \quad (27)$$

for each $\eta > 0$. We rewrite this inequality in more detail for $\eta = s_t \eta_t$.

First, we simplify the left-hand side of the inequality (27)

$$\begin{aligned} \ln E_{j \sim \mathbf{w}_t} \left[e^{-\eta(X_t^j - EX_t^j)} \right] &= \ln \left(\sum_{j=1}^N w_{j,t} e^{-s_t \eta_t \left(\frac{\hat{l}_{j,t} - \hat{l}_t^-}{s_t} - \sum_{i=1}^N w_{i,t} \frac{\hat{l}_{i,t} - \hat{l}_t^-}{s_t} \right)} \right) = \\ &= \ln \left(\frac{\sum_{j=1}^N w_{j,t} e^{-\eta_t \hat{l}_{j,t}}}{e^{-\eta_t \sum_{j=1}^N w_{j,t} \hat{l}_{j,t}}} \right) = \ln \sum_{j=1}^N w_{j,t} e^{-\eta_t \hat{l}_{j,t}} + \eta_t \sum_{j=1}^N w_{j,t} \hat{l}_{j,t} = \eta_t (h_t - m_t) = \eta_t \delta_t. \end{aligned}$$

Then the inequality (27) can be written in the form $\eta_t \delta_t \leq \frac{1}{s_t^2} v_t (e^{s_t \eta_t} - 1 - s_t \eta_t)$, from which we obtain the required inequality (26). \square

The inequality (26) can be presented in the form

$$\delta_t \leq \frac{g(s_t \eta_t)}{s_t} v_t, \text{ where } g(x) = \frac{e^x - x - 1}{x}. \quad (28)$$

Lemma 5 $(\Delta_T)^2 \leq (\ln^* N) V_T + \left(\frac{2}{3} \ln^* N + 1\right) S_T \Delta_T$.

Proof. We have

$$\begin{aligned}
(\Delta_T)^2 &= \sum_{t=1}^T ((\Delta_t)^2 - (\Delta_{t-1})^2) = \sum_{t=1}^T ((\Delta_{t-1} + \delta_t)^2 - (\Delta_{t-1})^2) \leq \\
&\sum_{t=1}^T (2\delta_t \Delta_{t-1} + \delta_t^2) = \sum_{t=1}^T \left(\frac{2\delta_t}{\eta_t} \ln^* N + \delta_t^2 \right) \leq \\
&\sum_{t=1}^T \left(\frac{2\delta_t}{\eta_t} \ln^* N + s_t \delta_t \right) \leq 2 \ln^* N \sum_{t=1}^T \frac{\delta_t}{\eta_t} + S_T \Delta_T. \tag{29}
\end{aligned}$$

The bound for $\frac{\delta_t}{\eta_t}$ is obtained using (28): $\frac{1}{2} v_t \geq \frac{\delta_t s_t}{2g(s_t \eta_t)} = \frac{\delta_t}{\eta_t} - s_t \varphi(s_t \eta_t) \delta_t$, where $\varphi(x) = \frac{e^x - \frac{1}{2}x^2 - x - 1}{xe^x - x^2 - x}$. It is not difficult to prove that $\varphi(x) \leq 1/3$.

Then, summing the inequality $\frac{\delta_t}{\eta_t} \leq \frac{1}{3} s_t \delta_t + \frac{1}{2} v_t$, and combining it with the inequality (29) we obtain the needed inequality. \square

Now, we obtain the bounds for V_T . From definition $v_t \leq (l_t^+ - h_t)(h_t - l_t^-) \leq \frac{s_t^2}{4}$. To obtain (8), we will use the following lemma.

Lemma 6 *If $L_T^{(\mathbf{q})} \leq H_T$ then $V_T \leq S_T \frac{(L_T^+ - L_T^{(\mathbf{q})})(L_T^{(\mathbf{q})} - L_T^-)}{L_T^+ - L_T^-} + \gamma_{k,T} S_T \Delta_T$.*

Proof. The following inequality holds true

$$\begin{aligned}
V_T &= \sum_{t=1}^T v_t \leq \sum_{t=1}^T (l_t^+ - h_t)(h_t - l_t^-) \leq S_T \sum_{t=1}^T \frac{(l_t^+ - h_t)(h_t - l_t^-)}{s_t} = \\
&S_T T \sum_{t=1}^T \frac{1}{T} \frac{(l_t^+ - h_t)(h_t - l_t^-)}{s_t} \leq S_T \frac{(L_T^+ - H_T)(H_T - L_T^-)}{L_T^+ - L_T^-}. \tag{30}
\end{aligned}$$

The inequality (30) is obtained by applying the Jensen inequality to a concave function $B(x, y, z) = (z - y)(y - x)/(z - x)$ on the set $x \leq y \leq z$ (see for detail de Rooij et al. 2014). \square

Recall that $H_T \leq L_T^{(\mathbf{q})} + \gamma_{k,T} \Delta_T$. Then, assuming $L_T^{(\mathbf{q})} \leq H_T$, we have

$$\begin{aligned}
V_T &\leq S_T \frac{(L_T^+ - L_T^{(\mathbf{q})})(L_T^{(\mathbf{q})} + \gamma_{k,T} \Delta_T - L_T^-)}{L_T^+ - L_T^-} \leq \\
&S_T \frac{(L_T^+ - L_T^{(\mathbf{q})})(L_T^{(\mathbf{q})} - L_T^-)}{L_T^+ - L_T^-} + \gamma_{k,T} S_T \Delta_T. \tag{31}
\end{aligned}$$

Denote

$$Q_T = \frac{(L_T^+ - L_T^{(\mathbf{q})})(L_T^{(\mathbf{q})} - L_T^-)}{L_T^+ - L_T^-}.$$

By the inequality (31) and Lemma 5

$$\begin{aligned}
\Delta_T^2 &\leq (S_T Q_T + \gamma_{k,T} S_T \Delta_T) \ln^* N + \left(\frac{2}{3} \ln^* N + 1 \right) S_T \Delta_T = \\
&S_T Q_T \ln^* N + \left(\gamma_{k,T} \ln^* N + \frac{2}{3} \ln^* N + 1 \right) S_T \Delta_T. \tag{32}
\end{aligned}$$

We have the inequality $\Delta_T^2 \leq a + b \Delta_T$, where $a = S_T Q_T \ln^* N$, $b = (\gamma_{k,T} \ln^* N + \frac{2}{3} \ln^* N + 1) S_T$.

Solving this inequality with respect to Δ_T , we obtain: $\Delta_T \leq \frac{1}{2} b + \frac{1}{2} \sqrt{b^2 + 4a} \leq \sqrt{a} + b = \sqrt{S_T Q_T \ln^* N} + ((\gamma_{k,T} + \frac{2}{3}) \ln^* N + 1) S_T$.

If $H_T \leq L_T^{(q)}$ then $R_T^{(q)} \leq 0$ and the inequality (8) is automatically executed. Otherwise, by Lemma 3,

$$R_T^{(q)} \leq \gamma_{k,T} \Delta_T \leq \gamma_{k,T} \sqrt{S_T Q_T \ln^* N} + \gamma_{k,T} \left(\left(\gamma_{k,T} + \frac{2}{3} \right) \ln^* N + 1 \right) S_T.$$

We obtain the inequality (8) using the lower and the upper bounds (5) for $L_T^{(q)}$.

To obtain (7), it is sufficient to use the inequality $V_T \leq \frac{1}{4} \sum_{t=1}^T s_t^2$ and a derivation similar to (32). This completes the proof of Theorem 1.

To prove the inequality (10) of Corollary 1 we simplify the inequality (30) as $V_T \leq S_T(L_T^+ - H_T) \leq S_T(L_T^+ - L_T^{(q)})$ if $L_T^{(q)} \leq H_T$. Finally, $R_T^{(q)} \leq \gamma_{k,T} \Delta_T \leq \gamma_{k,T} \sqrt{S_T(L_T^+ - L_T^{(q-)})} \ln^* N + \gamma_{k,T} \left(\frac{2}{3} \ln^* N + 1 \right) S_T$ for all T .

References

- D. Adamskiy, W.M. Koolen, A. Chernov, V. Vovk. A closer look at adaptive regret. *Lecture Notes in Artificial Intelligence*, 290–304, 2012.
- A. Blum, Y. Mansour. From external to internal regret. *Journal of Machine Learning Research*. 8:1307–1324, 2007.
- O. Bousquet, M. Warmuth. Tracking a small set of experts by mixing past posteriors. *Journal of Machine Learning Research*. 3:363–396, 2002.
- A. Chernov, V. Vovk. Prediction with expert evaluators’ advice. *Lecture Notes in Artificial Intelligence*. 8–22, 2009.
- N. Cesa-Bianchi, G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press. 2006.
- N. Cesa-Bianchi, Y. Mansour, and G. Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*. 66(2/3):321–352, 2007.
- M. Devaine, P. Gaillard, Y. Goude, G. Stoltz. Forecasting electricity consumption by aggregating specialized experts. *Machine Learning*. 90(2): 231–260, 2013.
- Y. Freund, R.E. Schapire. A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting. *Journal of Computer and System Sciences*. 55:119–139, 1997.
- Y. Freund, R.E. Schapire, Y. Singer, M.K. Warmuth. Using and combining predictors that specialize. In: *Proc. 29th Annual ACM Symposium on Theory of Computing*. 334–343, 1997.
- M. Herbster, M. Warmuth. Tracking the best expert. *Machine Learning*, 32(2): 151–178, 1998.
- Tao Hong, P. Pinson, Shu Fan. Global Energy Forecasting Competition 2012, *International Journal of Forecasting* V.30(2): P.357–363, 2014.
- P. Gaillard, Y. Goude, and G. Stoltz. A further look at the forecasting of the electricity consumption by aggregation of specialized experts. Technical report, 2011. pierre.gaillard.me/doc/GaGoSt-report.pdf.
- P. Gaillard, G. Stoltz, and T. van Erven. A second-order bound with excess losses. In *JMLR: Workshop and Conference Proceedings*. 35: 176–196, 2014.
- Y. Kalnishkan, D. Adamskiy, A. Chernov, T. Scarfe. Specialist Experts for Prediction with Side Information. *IEEE International Conference on Data Mining Workshop (ICDMW)*. IEEE, 1470–1477, 2015.
- J. Kivinen, M.K. Warmuth. Averaging expert prediction. In Paul Fisher and Hans Ulrich Simon, editors, *Computational Learning Theory: 4th European Conference (EuroColt ’99)*. 153–167, Springer, 1999.
- N. Littlestone, M. Warmuth. The weighted majority algorithm. *Information and Computation*. 108:212–261, 1994.
- S. de Rooij, T. van Erven, P. Grunwald, W. Koolen. Follow the Leader. If You Can, Hedge If You Must. *Journal of Machine Learning Research*. 15:1281–1316, 2014.
- V. Vovk. Aggregating strategies. In M. Fulk and J. Case, editors, *Proceedings of the 3rd Annual Workshop on Computational Learning Theory*, 371–383, San Mateo, CA, Morgan Kaufmann, 1990.
- V. Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*. 56(2):153–173, 1998.

- V. Vovk. Derandomizing stochastic prediction strategies. *Machine Learning*. 35(3): 247–282, 1999.
- V. V'yugin. Online Aggregation of Unbounded Signed Losses Using Shifting Experts. *Proceedings of Machine Learning Research*. 60: 1–15, 2017. <http://proceedings.mlr.press/v60/>