

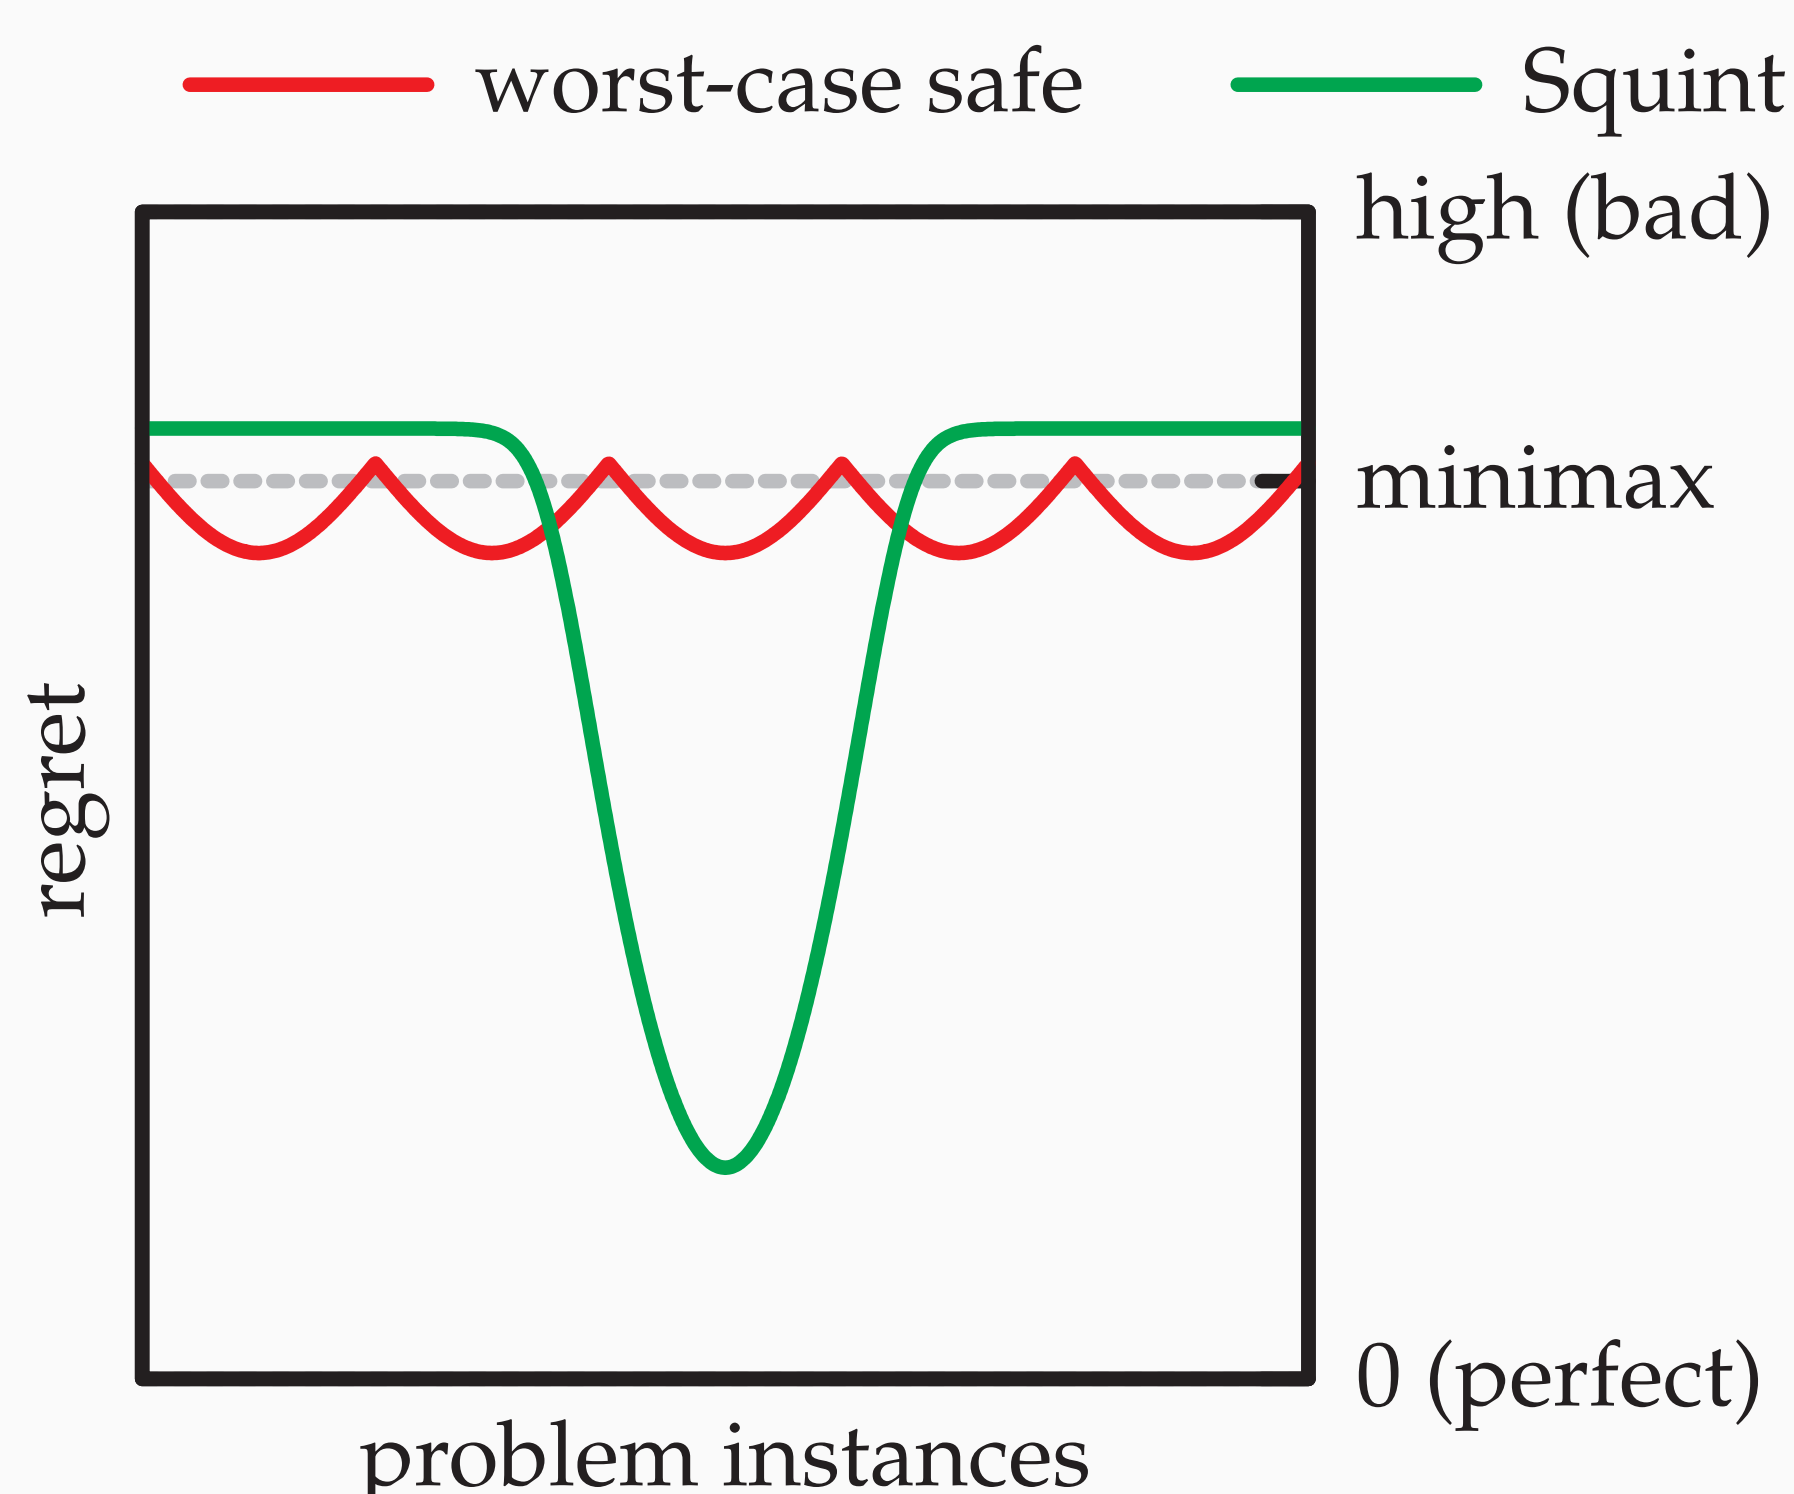
Second-order Quantile Methods for Experts and Combinatorial Games



Wouter M. Koolen Tim van Erven

Our new algorithm Squint

- adapts to the difficulty of the learning problem by **learning the learning rate**,
- thereby integrating both the popular **second-order** and **quantile** adaptivities,
- at the **run time** of standard Hedge.



Second-order adaptivity

Cesa-Bianchi, Mansour, and Stoltz 2007, Hazan and Kale 2010, Chiang, Yang, Lee, Mahdavi, Lu, Jin, and Zhu 2012, De Rooij, Van Erven, Grünwald, and Koolen 2014, Gaillard, Stoltz, and Van Erven 2014, Steinhardt and Liang 2014

$$R_T^k \prec \sqrt{V_T^k \ln K} \quad \text{for each expert } k.$$

for some second-order $V_T^k \leq L_T^k \leq T$

- **stochastic case, learning sub-algorithms**
- **specialized algorithms, hard-coded $\ln K$.**

Regret guarantee

Choose k and fix $\hat{\eta} = \frac{R_T^k}{2V_T^k}$. Now as

$$1 \geq \Phi_T \geq \pi(k)\gamma(\hat{\eta})e^{\hat{\eta}R_T^k - \hat{\eta}^2 V_T^k} = \pi(k)\gamma(\hat{\eta})e^{\frac{(R_T^k)^2}{4V_T^k}}$$

we have

$$R_T^k \leq 2\sqrt{V_T^k(-\ln \pi(k) - \ln \gamma(\hat{\eta}))}.$$

For quantile bound take $\sum_{k \in \mathcal{K}}$

Quantile adaptivity

Hutter and Poland 2005, Chaudhuri, Freund, and Hsu 2009, Chernov and Vovk 2010, Luo and Schapire 2014

Prior π on experts:

$$\min_{k \in \mathcal{K}} R_T^k \prec \sqrt{T(-\ln \pi(\mathcal{K}))} \quad \text{for each subset } \mathcal{K} \text{ of experts}$$

- **over-discretization, company baseline**
- **specialized algorithms, hard-coded T . "impossible tunings". efficiency.**

Three priors

Idea: have prior $\gamma(\eta)$ put sufficient mass around optimal $\hat{\eta}$

1. Uniform prior (generalizes to conjugate)

$$\gamma(\eta) = 2$$

Efficient algorithm, $C_T = \ln V_T^{\mathcal{K}}$.

2. Chernov&Vovk (2010) prior

$$\gamma(\eta) = \frac{\ln 2}{\eta \ln^2(\eta)}$$

Not efficient, $C_T = \ln \ln V_T^{\mathcal{K}}$.


3. Improper(!) log-uniform prior

$$\gamma(\eta) = 1/\eta$$

Efficient algorithm, $C_T = \ln \ln T$.

Hedge setting

K experts



In round $t = 1, 2, \dots$

- Learner plays a probability distribution $w_t = (w_t^1, \dots, w_t^K)$ on experts
- Adversary reveals the expert loss vector $\ell_t = (\ell_t^1, \dots, \ell_t^K) \in [0, 1]^K$



- Learner incurs loss $w_t^\top \ell_t$

The goal is to have small **regret**

$$R_T^k := \underbrace{\sum_{t=1}^T w_t^\top \ell_t}_{\text{Learner}} - \underbrace{\sum_{t=1}^T \ell_t^k}_{\text{Expert } k}$$

with respect to every expert k at every time T .

Squint guarantees both

Squint algorithm with bound

$$R_T^{\mathcal{K}} \prec \sqrt{V_T^{\mathcal{K}}(-\ln \pi(\mathcal{K}) + C_T)} \quad \text{for each subset } \mathcal{K} \text{ of experts}$$

where $R_T^{\mathcal{K}} = \mathbb{E}_{\pi(k|\mathcal{K})} R_T^k$ and $V_T^{\mathcal{K}} = \mathbb{E}_{\pi(k|\mathcal{K})} V_T^k$ denote the average (under the prior π) among the reference experts $k \in \mathcal{K}$ of the cumulative regret $R_T^k = \sum_{t=1}^T r_t^k$ and the (uncentered) variance of the excess losses $V_T^k = \sum_{t=1}^T (r_t^k)^2$ (where $r_t^k = (w_t - e_k)^\top \ell_t$).

pretty two-line proof

Squint potential motivation

Fix prior π on experts $k \in \{1, \dots, K\}$ and prior γ on learning rates $\eta \in [0, 1/2]$.

Potential function (weighted sum of objectives)

$$\Phi_t := \mathbb{E}_{\pi(k)\gamma(\eta)} \left[e^{\eta R_t^k - \eta^2 V_t^k} \right],$$

and associated weights

$$w_{t+1}^k := \frac{\pi(k) \mathbb{E}_{\gamma(\eta)} \left[e^{\eta R_t^k - \eta^2 V_t^k} \right]}{\text{normalisation}}.$$

$$1 = \Phi_0 \geq \Phi_1 \geq \Phi_2 \geq \dots$$

$$\Phi_{t+1} - \Phi_t = \mathbb{E}_{\pi(k)\gamma(\eta)} \left[e^{\eta R_t^k - \eta^2 V_t^k} \left(e^{\eta r_{t+1}^k - \eta^2 (r_{t+1}^k)^2} - 1 \right) \right] \leq \mathbb{E}_{\pi(k)\gamma(\eta)} \left[e^{\eta R_t^k - \eta^2 V_t^k} \eta (w_{t+1} - e_k)^\top \ell_{t+1} \right] = 0$$

by the choice of weights w_{t+1}

Squint with log-uniform prior

Closed-form expression for weights:


$$w_{t+1}^k \propto \pi(k) \int_0^{1/2} e^{\eta R_t^k - \eta^2 V_t^k} \frac{1}{\eta} d\eta \propto \pi(k) e^{\frac{(R_t^k)^2}{4V_t^k}} \frac{\text{erf}\left(\frac{R_t^k}{2\sqrt{V_t^k}}\right) - \text{erf}\left(\frac{R_t^k - V_t^k}{2\sqrt{V_t^k}}\right)}{\sqrt{V_t^k}}$$

constant time per expert per round

Extensions

Combinatorial concept class $\mathcal{C} \subseteq \{0, 1\}^K$:

- Shortest path
- Spanning trees
- Permutations



Component Squint guarantees:

$$R_T^u \prec \sqrt{V_T^u(\text{comp}(u) + KC_T)} \quad \text{for each } u \in \text{conv}(\mathcal{C}).$$

The reference set of experts \mathcal{K} is subsumed by an "average concept" vector $u \in \text{conv}(\mathcal{C})$, for which our bound relates the coordinate-wise average regret $R_T^u = \sum_{t,k} u_k r_t^k$ to the averaged variance $V_T^u = \sum_{t,k} u_k (r_t^k)^2$ and the prior entropy $\text{comp}(u)$.

No range factor

Drop-in replacement for Component Hedge

Koolen, Warmuth, and Kivinen 2010 with same run-time

Future work

Loss range adaptivity, bandits, online convex optimization