Second-order Quantile Methods

for Experts and Combinatorial Games

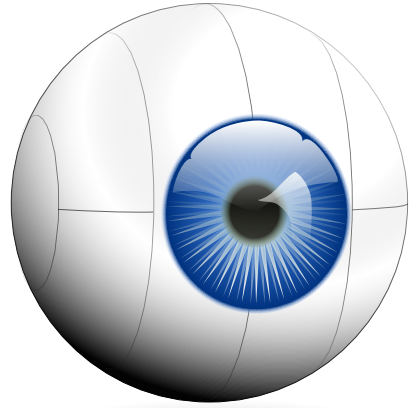

Wouter M. Koolen Tim van Erven

Queensland University of Technology

Berkeley UNIVERSITY OF CALIFORNIA

Universiteit Leiden

Keep it simple: expert setting

Online sequential prediction with expert advice

Core instance of advanced online learning tasks

- Bandits
- Combinatorial & matrix prediction
- Online convex optimization
- Boosting
- ...

Beyond the Worst Case

Two reasons data is often **easier** in practice:

Beyond the Worst Case

Two reasons data is often **easier** in practice:

Data complexity

- Stochastic data (gap)
- Low noise
- Low variance

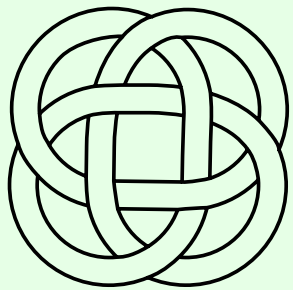

Beyond the Worst Case

Two reasons data is often **easier** in practice:

Data complexity

- Stochastic data (gap)
- Low noise
- Low variance

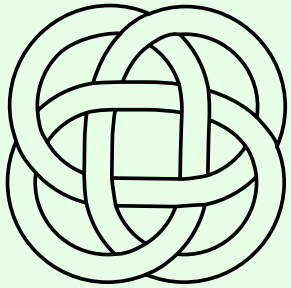

second-order

Beyond the Worst Case

Two reasons data is often **easier** in practice:

Data complexity

- Stochastic data (gap)
- Low noise
- Low variance

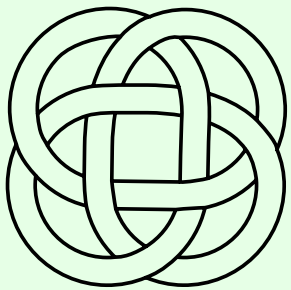

second-order

Model complexity

- Simple model is good
- Multiple good models

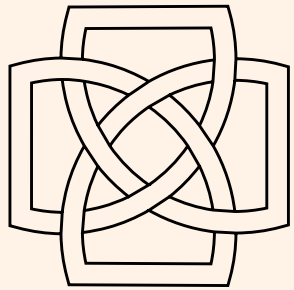

Beyond the Worst Case

Two reasons data is often **easier** in practice:

Data complexity

- Stochastic data (gap)
- Low noise
- Low variance

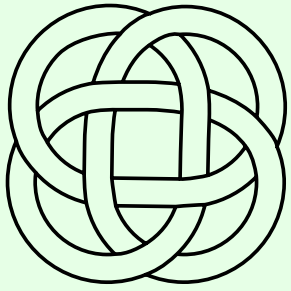

second-order

Model complexity

- Simple model is good
- Multiple good models

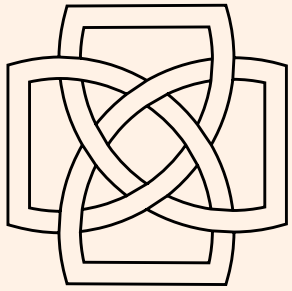

quantiles

Beyond the Worst Case

Two reasons data is often **easier** in practice:

Data complexity

- Stochastic data (gap)
- Low noise
- Low variance

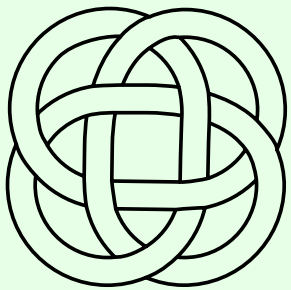

second-order

Model complexity

- Simple model is good
- Multiple good models

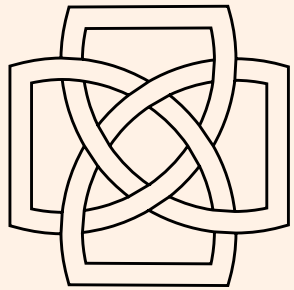

quantiles

Second-order & Quantiles

- Any combination

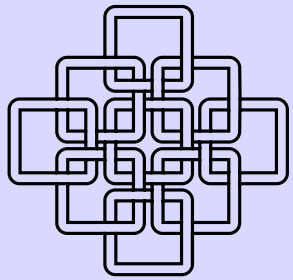

All we need is the right learning rate

Existing algorithms

(Hedge, Prod, . . .)

with

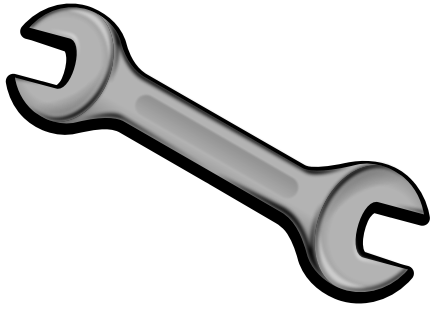

oracle

learning rate η

exploit

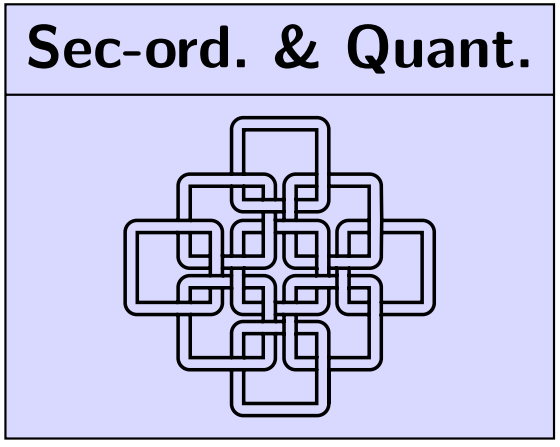

All we need is the right learning rate

Existing algorithms

(Hedge, Prod, . . .)

with

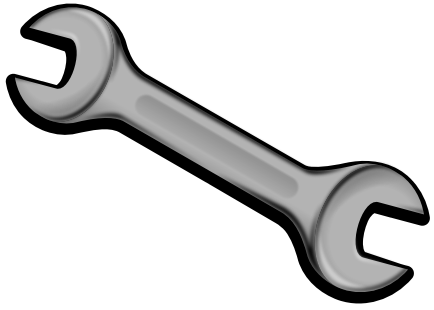

oracle

learning rate η

exploit

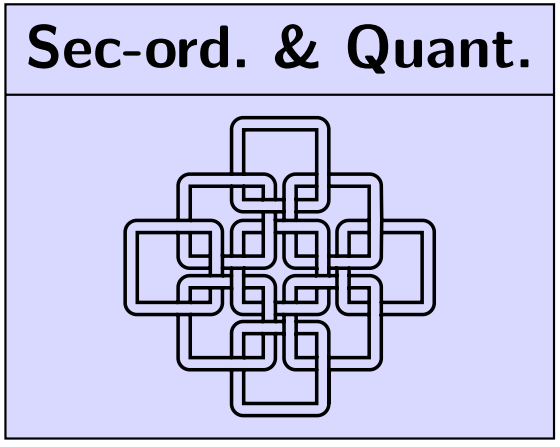

Can we exploit Second-order & Quantiles **on-line**?

But everyone struggles with the learning rate

Oracle η

- **not** monotonic,
- **not** smooth

over time.

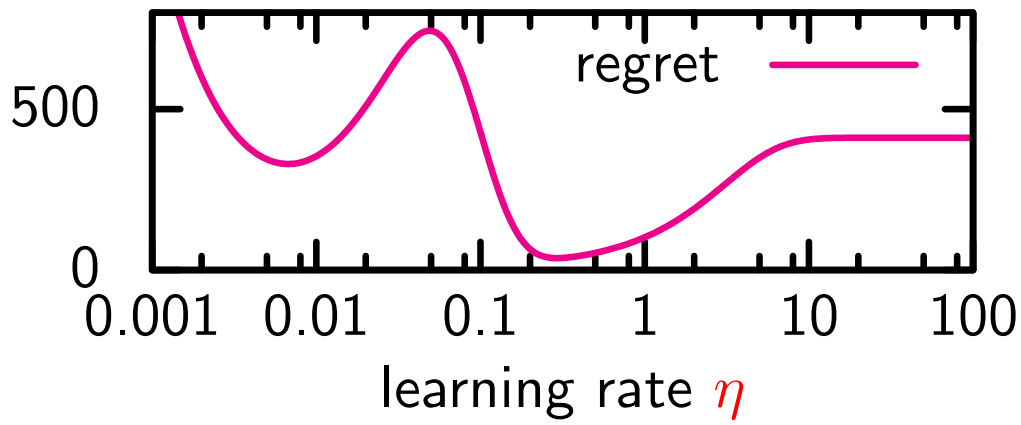

But everyone struggles with the learning rate

Oracle η

- **not** monotonic,
- **not** smooth

over time.

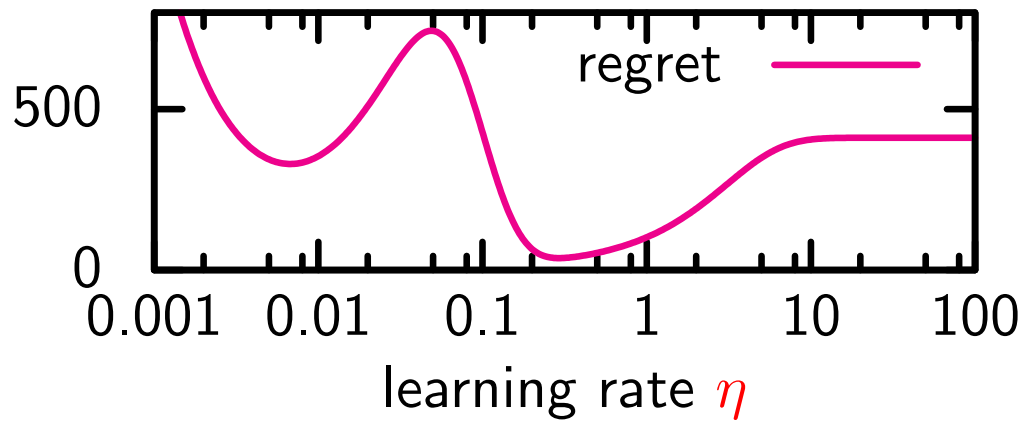

State of the art:

Second-order

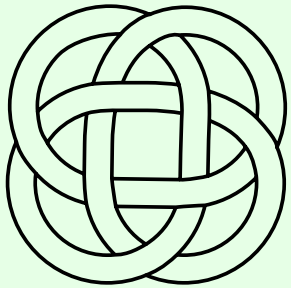

Cesa-Bianchi, Mansour, and Stoltz 2007, Hazan and Kale 2010, Chiang, Yang, Lee, Mahdavi, Lu, Jin, and Zhu 2012, De Rooij, Van Erven, Grünwald, and Koolen 2014, Gaillard, Stoltz, and Van Erven 2014, Steinhardt and Liang 2014

or

Quantiles

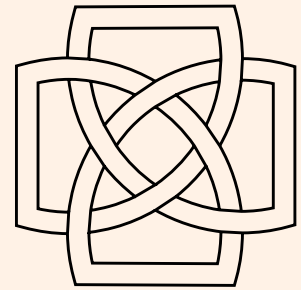

Hutter and Poland 2005, Chaudhuri, Freund, and Hsu 2009, Chernov and Vovk 2010, Luo and Schapire 2014

Learning the learning rate

We found a formulation of the algorithm where simply putting a prior on η works.

Our algorithm **Squint**

$$w_{t+1}^k \propto \pi(k) \mathop{\mathbb{E}}_{\gamma(\eta)} \left[e^{\eta R_t^k - \eta^2 V_t^k} \eta \right]$$

guarantees for each subset $\mathcal{K}$ of experts, at each time $T \geq 0$:

$$R_T^{\mathcal{K}} \prec \sqrt{V_T^{\mathcal{K}} \left(-\ln \pi(\mathcal{K}) + \ln \ln T \right)}$$

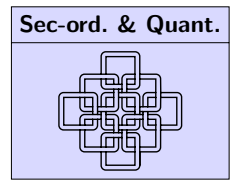

- Run-time of Hedge
- Only $\ln \ln T$ extra over oracle learning rate.
- Extension to Combinatorial Games

More at our poster