

# Putting Bayes to sleep



zzzzZZZZZ....

**Wouter M. Koolen**

Dmitry Adamskiy

Manfred Warmuth




UNIVERSITY OF CALIFORNIA  
**SANTA CRUZ**

# Online learning

Real-world tasks:




Many strategies: A, B, ...



Need **adaptive** algorithms that can exploit **repeats**

# Mixing past posteriors





$$\hat{w}_{t+1}(m) = \underbrace{\frac{P(y_t|m)w_t(m)}{\sum_m P(y_t|m)w_t(m)}}_{\text{Bayesian posterior update}}$$

$$w_t(m) = \underbrace{\sum_{s=0}^{t-1} \hat{w}_s(m) \gamma_t(s)}_{\text{bizarre}}$$



## Main result

**Bayesian interpretation** of Mixing Past Posteriors  
using specialist models that **sleep**



## Applications aplenty

We really get it:

- Sharpest bounds
- Optimal efficiency
- Portable trick: great results in multitask learning



Thank you, and see you at poster **Tu57!**