Learning with Large Number of Experts: Component Hedge Algorithm

Giulia DeSalvo and Vitaly Kuznetsov

Courant Institute

March 24th, 2015

Learning with Large Number of Experts

- Regret of RWM is $O(\sqrt{T \ln N})$.
  - Informative even for very large number of experts.
- What if there is “overlap” between experts?
  - RWM with path experts
  - FPL with path experts
  - $\longrightarrow$ can we do better?

[Littlestone and Warmuth, 1989; Kalai and Vempala, 2004]

Better Bounds in Structured Case?

- Can "overlap" between experts lead to better regret guarantees?
- What are the lower bounds in the structured setting?
- Computationally efficient solutions that realize these bounds?

Outline

- Learning Scenario
- Component Hedge Algorithm
- Regret Bounds
- Applications & Lower Bounds
- Conclusion & Open Problems

Learning Scenario

Assumptions:

- Structured concept class $\mathcal{C} \subseteq \{0,1\}^d$
  - Composed of components: C^t indicates which components are used for each trial t.
- Additive loss ℓ^t incurred at each trial t.
  - Loss of each concept C is $C \cdot \ell^t \leq M := max_{C \in \mathcal{C}} |C|$

Goal:

- minimize expected regret after T trials

$$R_T = \sum_{t=1}^{T} \mathbb{E}[C^t] \cdot \ell^t - \min_{C \in \mathcal{C}} \sum_{t=1}^{T} C \cdot \ell^t$$

Component Hedge Algorithm

[Koolen, Warmuth, and Kivinen, 2010]

CH **maintains weights** $w^t \in conv(C) \subseteq [0,1]^d$ over the components at each round t.

- **Update**:
  1. weights: $\widehat{w}_i^t = w_i^{t-1} e^{-\eta \ell_i^t}$
  2. relative entropy projection:

$$w^t := argmin_{w \in conv(\mathcal{C})} \Delta(w || \hat{w}^t)$$

where $\Delta(w \parallel v) = \sum_{i=1}^{d} (w_i \ln \frac{w_i}{v_i} + v_i - w_i)$

Component Hedge Algorithm

- **Prediction**:

  1. Decomposition of weights:

  $$w^t = \sum_{C \in \mathcal{C}} \alpha_C C$$

  where $\boldsymbol{\alpha}$ is a distribution over $\mathcal{C}$

  2. Sample C^{t+1} according to $\boldsymbol{\alpha}$

Efficiency

Need efficient implementation of:

- Decomposition (not unique) of weights over the concepts
- Entropy projection step (convex problem)

Sufficient: $conv(\mathcal{C})$ described by polynomial in d constraints

Regret Bounds

Theorem: Regret Bounds for CH

Let $\ell^* = \min_{C \in \mathcal{C}} C \cdot (\ell^1 + \ldots + \ell^T)$ be the loss of the best concept in hindsight, then

$$R_T \leq \sqrt{2\ell^* M \ln(d/M)} + M \ln(d/M)$$

by choosing $\eta = \sqrt{\frac{2M \ln(d/M)}{\ell^*}}$

- Since $\ell^* \leq MT$, regret $R_T \in O(M\sqrt{T \ln d})$.
- Matching lower bounds in applications.

Comparison of CH, RWM and FPL

1. CH has significantly **better regret bounds**:
   - CH: $R_T \in O(M\sqrt{T \ln d})$.
   - RWM: $R_T \in O(M\sqrt{MT \ln d})$
   - FPL: $R_T \in O(M\sqrt{dT \ln d})$
2. CH is **optimal** w.r.t. regret bounds while RWM and FPL are not optimal.
3. Standard expert setting (no structure):
   CH, RWM and FPL reduce to the same algorithm.

Applications

- On-line shortest path problems.
- On-line PCA (k-sets).
- On-line ranking (k-permutations).
- Spanning trees.

On-line Shortest Path Problem (SPP)

- $G = (V, E)$ is a directed graph.
- s is the source and t is the destination.
- Each $s - t$ path is an expert.
- The loss is **additive** over edges.

Unit Flow Polytope

- Convex hull of paths cannot be captured by linear constraints
- **Unit flow** polytope relaxation is used:

$$w_{u,v} \geq 0, \quad \forall (u,v) \in E$$
$$\sum_{v \in V} w_{s,v} = 1$$
$$\sum_{v \in V} w_{v,u} = \sum_{v \in V} w_{u,v}, \quad \forall u \in V$$

- Relaxation does not hurt regret bounds.

Example of Unit Flow Polytope

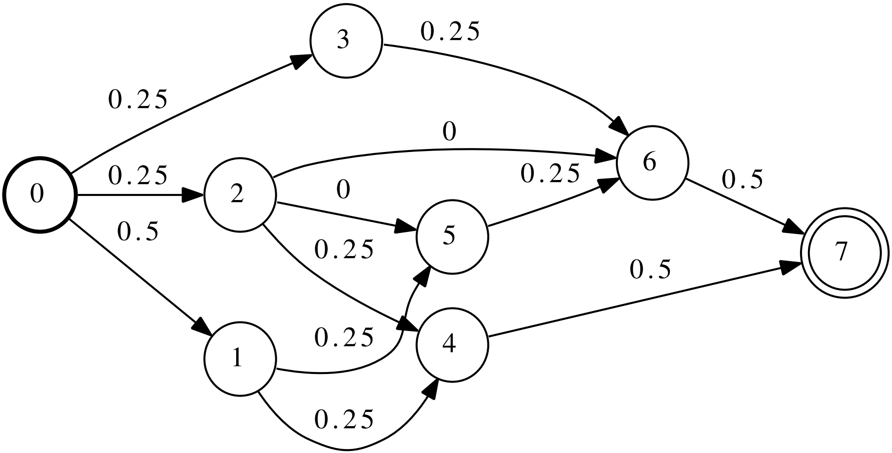

Entropy Projection on Unit Flow Polytope

$$\min_{\mathbf{w}} \sum_{(u,v)\in E} w_{u,v} \ln \frac{w_{u,v}}{\widehat{w}_{u,v}} + \widehat{w}_{u,v} - w_{u,v}$$

subject to:

$$w_{u,v} \geq 0, \quad \forall (u,v) \in E$$

$$\sum_{v \in V} w_{s,v} = 1$$

$$\sum_{v \in V} w_{v,u} = \sum_{v \in V} w_{u,v}, \quad \forall u \in V$$

Dual problem

$$\max_{\boldsymbol{\lambda}} \left\{ \lambda_s - \sum_{(u,v)\in E} \widehat{w}_{u,v} e^{\lambda_u - \lambda_v} \right\}$$

- No constraints.
- Only $|V|$ variables.
- Primal solution: $w_{u,v} = \widehat{w}_{u,v} e^{\lambda_u - \lambda_v}$

Convex Decomposition

1. Find any non-zero path from s to t.
2. Subtract the smallest weight from each edge.
3. Repeat until no path is found.

$\Longrightarrow$ At most $|E|$ iteration is needed.

Example of Convex Decomposition

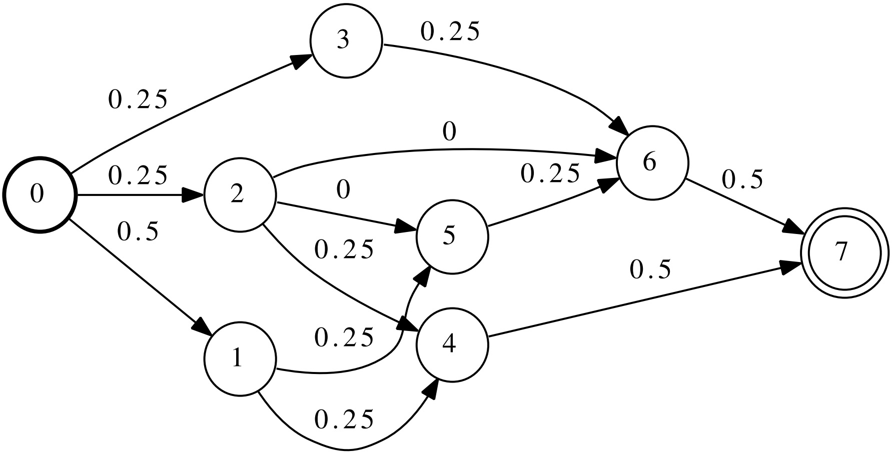

Example of Convex Decomposition

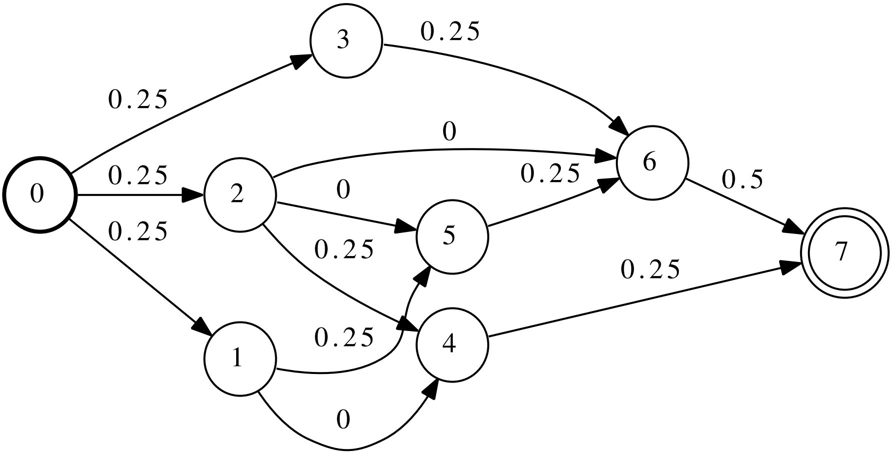

Example of Convex Decomposition

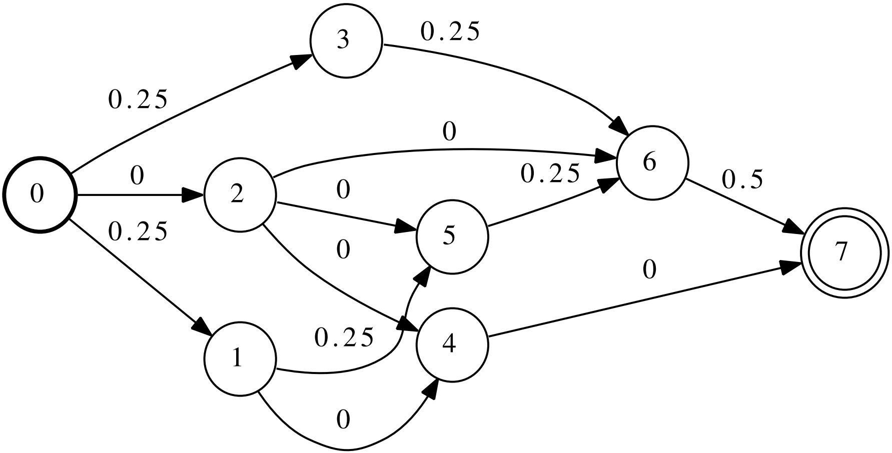

Example of Convex Decomposition

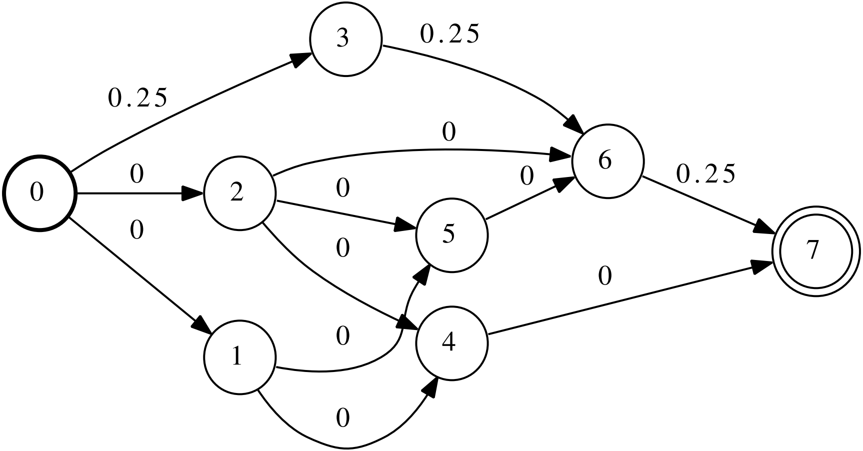

Example of Convex Decomposition

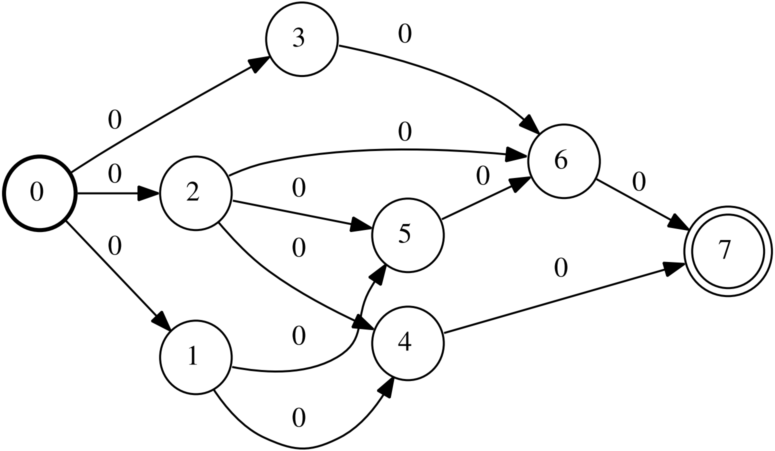

Regret Bounds for SPP

- Expected regret is bounded by

$$2\sqrt{\ell^* k^* \ln|V|} + 2k^* \ln|V| \in O(M\sqrt{T \ln|V|})$$

- Bound holds for arbitrary graphs.

Lower Bounds

Any algorithm can be forced to have expected regret

$$\sqrt{\ell^* k^* \ln \frac{|V|}{k^*}}$$

Idea of the proof:

- Minimize the "overlap".
- Create $|V|/k$ disjoint paths of length k.
- Apply lower bounds for standard expert setting.

Conclusions

- Regret of CH is often **better** than that of RWM or FPL in **structured setting**.
- Regret of CH often matches **lower bounds** in applications.
- Efficient solutions exist for a wide range of applications: on-line shortest path, on-line PCA, on-line ranking, spanning trees.

References

- Wouter Koolen, Manfred K. Warmuth, and Jyrki Kivinen. Hedging Structured Concepts. In *COLT* (2010).
- Nick Littlestone and Manfred K. Warmuth. The Weighted Majority Algorithm. *FOCS* 1989: 256-261.
- Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.* , 2003.
- Eiji Takimoto and Manfred K. Warmuth. Path kernels and multiplicative updates. *JMLR*, 4:773818, 2003.
- T. van Erven, W. Kotlowski, and M. K. Warmuth. Follow the leader with dropout perturbations. In *COLT*, 2014.
- Nicolo Cesa-Bianchi and Gabor Lugosi. Combinatorial bandits. In Proceedings of the 22nd Annual Conference on Learning Theory, 2009.

Regret Bounds

Theorem: Regret Bounds for CH

Let $\ell^* = \min_{C \in \mathcal{C}} C \cdot (\ell^1 + \ldots + \ell^T)$ be the loss of the best concept in hindsight, then

$$R_T \leq \sqrt{2\ell^* M \ln(d/M)} + M \ln(d/M)$$

by choosing $\eta = \sqrt{\frac{2M \ln(d/M)}{\ell^*}}$

Proof of CH Regret Bound

1. Bound:
   $(1-e^{-\eta})w^{t-1}\cdot\ell^t \leq \Delta(C||w^{t-1})-\Delta(C||w^t)+\eta C\cdot\ell^t$.
   - $1 - e^{-\eta x} \geq (1 - e^{-\eta})x$
   - Generalized Pythagorean Theorem
2. Sum over trials t:
   $(1 - e^{-\eta})\sum_{t=1}^{T} w^{t-1} \cdot \ell^t \leq \Delta(C||w^0) - \Delta(C||w^T) + \eta C \cdot \ell^{\leq T}$
   where $\ell^{\leq T} = \ell^1 + \cdots + \ell^T$.
3. Use $w^{t-1} \cdot \ell^t = E[C^t] \cdot \ell^t$:
   $\sum_{t=1}^{T} E[C^t] \cdot \ell^t \leq \frac{\Delta(C||w^0)-\Delta(C||w^T)+\eta C\cdot\ell^{\leq T}}{(1-e^{-\eta})}$

Proof of CH Regret Bound

4. w^0 assumes uniform distribution over concepts $w_i^0 = \frac{M}{d} \Longrightarrow \Delta(C||w^0) = M\ln(\frac{d}{M})$
5. let ℓ^* best concept in hind-sight and choosing $\eta = \sqrt{\frac{2M\ln(\frac{d}{M})}{\ell^*}} \Longrightarrow$ Regret bound R_T.