

ALGORITHMS WITH PREDICTIONS

WARMING UP AND RENTING SKIS

Adam Polak

SOFSEM PhD School, Kraków, February 9th, 2026

Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances



Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

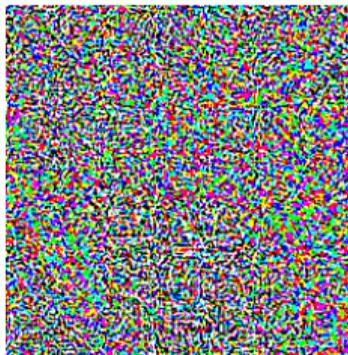


RECALL: ADVERSARIAL EXAMPLES



“panda”
57.7% confidence

+ .007 ×



“nematode”
8.2% confidence

=



“gibbon”
99.3 % confidence

Source: arxiv.org/abs/1412.6572

Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances

Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances

Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

Best of both worlds: learning-augmented algorithms

Input + **black-box predictions** (e.g., coming from ML model)

ALGORITHMS WITH PREDICTIONS = ALPS



Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances

Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

Best of both worlds: learning-augmented algorithms

Input + **black-box predictions** (e.g., coming from ML model)

Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances

Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

Best of both worlds: learning-augmented algorithms

Input + **black-box predictions** (e.g., coming from ML model)

- **Consistency:** close-to-optimal performance when predictions accurate

Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances

Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

Best of both worlds: learning-augmented algorithms

Input + **black-box predictions** (e.g., coming from ML model)

- **Consistency:** close-to-optimal performance when predictions accurate
- **Robustness:** worst-case guarantees, even when predictions adversarial

Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances

Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

Best of both worlds: learning-augmented algorithms

Input + **black-box predictions** (e.g., coming from ML model)

- **Consistency:** close-to-optimal performance when predictions accurate
- **Robustness:** worst-case guarantees, even when predictions adversarial
- **Smoothness:** performance degrades slowly in the prediction error

Classical algorithms

- Worst-case guarantees
- **Overly pessimistic** on easy instances

Machine learning

- Powerful most of the time
- **No guarantees**, can go crazy

Best of both worlds: learning-augmented algorithms

Input + **black-box predictions** (e.g., coming from ML model)

- **Consistency:** close-to-optimal performance when predictions accurate
- **Robustness:** worst-case guarantees, even when predictions adversarial
- **Smoothness:** performance degrades slowly in the prediction error
- **Learnability:** good predictions can be learned efficiently

WARM UP:

SKI RENTAL



Each day of a ski season of unknown length decide whether to:

- rent the skis for one more day paying 1 , or
- buy the skis paying cost B .

Each day of a ski season of unknown length x decide whether to:

- rent the skis for one more day paying 1 , or
- buy the skis paying cost B .

Competitive ratio:

$$\max \frac{\text{cost}(ALG)}{\text{cost}(OPT)}$$

SKI RENTAL

Each day of a ski season of unknown length x decide whether to:

- rent the skis for one more day paying 1 , or
- buy the skis paying cost B .

Competitive ratio:

$$\max \frac{\text{cost}(ALG)}{\text{cost}(OPT)}$$

Without predictions:

(folklore)

- 2 -competitive deterministic algorithm: wait B days and buy
- no deterministic algorithm can do better

Each day of a ski season of unknown length x decide whether to:

- rent the skis for one more day paying 1 , or
- buy the skis paying cost B .

Without predictions:

(folklore)

- 2 -competitive deterministic algorithm: wait B days and buy
- no deterministic algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- predicted length of ski season: y days

Each day of a ski season of unknown length x decide whether to:

- rent the skis for one more day paying 1, or
- buy the skis paying cost B .

Without predictions:

(folklore)

- 2-competitive deterministic algorithm: wait B days and buy
- no deterministic algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- predicted length of ski season: y days
 - blindly following the prediction: if $y > B$, buy on day 1; if $y \leq B$, never buy
- $\text{cost}(\text{ALG}) \leq \text{cost}(\text{OPT}) + |x - y|$ (consistent, smooth, but **not robust**)

Each day of a ski season of unknown length x decide whether to:

- rent the skis for one more day paying 1, or
- buy the skis paying cost B .

Without predictions:

(folklore)

- 2-competitive deterministic algorithm: wait B days and buy
- no deterministic algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- predicted length of ski season: y days
- deterministic algorithm with a trade-off parameter $\lambda \in (0, 1)$
 - if $y > B$, buy on day λB ; if $y \leq B$, buy on day B/λ

Each day of a ski season of unknown length x decide whether to:

- rent the skis for one more day paying 1, or
- buy the skis paying cost B .

Without predictions:

(folklore)

- 2-competitive deterministic algorithm: wait B days and buy
- no deterministic algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- predicted length of ski season: y days
- deterministic algorithm with a trade-off parameter $\lambda \in (0, 1)$
 - if $y > B$, buy on day λB ; if $y \leq B$, buy on day B/λ
 - correct prediction: $(1 + \lambda)$ -competitive (**consistency**)
 - wrong prediction: $(1 + \frac{1}{\lambda})$ -competitive (**robustness**)

Each day of a ski season of unknown length x decide whether to:

- rent the skis for one more day paying 1, or
- buy the skis paying cost B .

Without predictions:

(folklore)

- 2-competitive deterministic algorithm: wait B days and buy
- no deterministic algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- predicted length of ski season: y days
- deterministic algorithm with a trade-off parameter $\lambda \in (0, 1)$
 - correct prediction: $(1 + \lambda)$ -competitive (**consistency**)
 - wrong prediction: $(1 + \frac{1}{\lambda})$ -competitive (**robustness**)
- this is Pareto-optimal

(Angelopoulos, Durr, Jin, Kamali, ITCS 2020)

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better
- simple 1.875-competitive algorithm: wait $\frac{3}{4}B$ days with $p = 0.5$, else wait B

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better
- simple 1.875-competitive algorithm: wait $\frac{3}{4}B$ days with $p = 0.5$, else wait B
- $\frac{e}{e-1}$ -competitive randomized algorithm (Karlin et al. '90)

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better
- simple 1.875-competitive algorithm: wait $\frac{3}{4}B$ days with $p = 0.5$, else wait B
- $\frac{e}{e-1}$ -competitive randomized algorithm (Karlin et al. '90)
- no algorithm can do better

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better
- simple 1.875-competitive algorithm: wait $\frac{3}{4}B$ days with $p = 0.5$, else wait B
- $\frac{e}{e-1}$ -competitive randomized algorithm (Karlin et al. '90)
- no algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- deterministic: $(1 + \lambda)$ -consistent, $(1 + \frac{1}{\lambda})$ -robust

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better
- simple 1.875-competitive algorithm: wait $\frac{3}{4}B$ days with $p = 0.5$, else wait B
- $\frac{e}{e-1}$ -competitive randomized algorithm (Karlin et al. '90)
- no algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- deterministic: $(1 + \lambda)$ -consistent, $(1 + \frac{1}{\lambda})$ -robust
i.e.: $(\frac{w}{w-1})$ -consistent, w -robust

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better
- simple 1.875-competitive algorithm: wait $\frac{3}{4}B$ days with $p = 0.5$, else wait B
- $\frac{e}{e-1}$ -competitive randomized algorithm (Karlin et al. '90)
- no algorithm can do better

Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- deterministic: $(1 + \lambda)$ -consistent, $(1 + \frac{1}{\lambda})$ -robust
i.e.: $(\frac{w}{w-1})$ -consistent, w -robust
- randomized: $(w \ln \frac{w}{w-1})$ -consistent, w -robust

Without predictions:

- 2-competitive deterministic algorithm: wait B days and buy (folklore)
- no deterministic algorithm can do better
- simple 1.875-competitive algorithm: wait $\frac{3}{4}B$ days with $p = 0.5$, else wait B
- $\frac{e}{e-1}$ -competitive randomized algorithm (Karlin et al. '90)
- no algorithm can do better

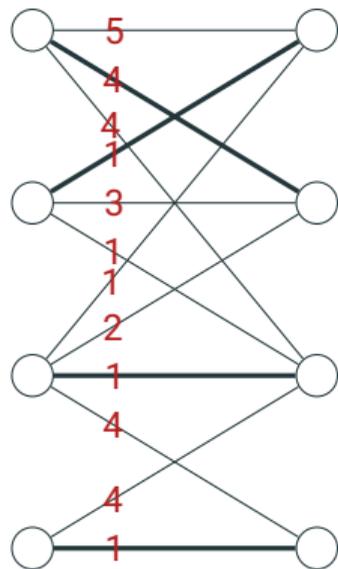
Learning-augmented:

(Purohit, Svitkina, Kumar, NeurIPS 2018)

- deterministic: $(1 + \lambda)$ -consistent, $(1 + \frac{1}{\lambda})$ -robust
i.e.: $(\frac{w}{w-1})$ -consistent, w -robust
- randomized: $(w \ln \frac{w}{w-1})$ -consistent, w -robust
- this is Pareto-optimal (Wei, Zhang, NeurIPS 2020)

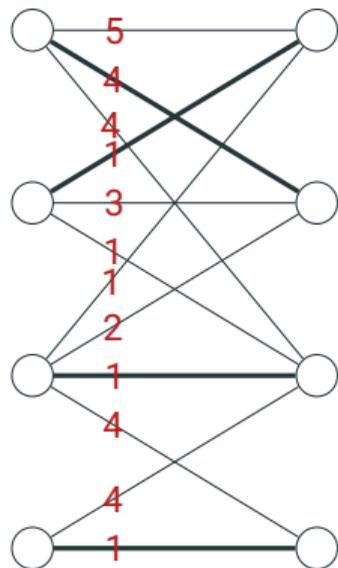
MINIMUM WEIGHT BIPARTITE MATCHING

MINIMUM WEIGHT BIPARTITE MATCHING



Hungarian algorithm: $O(nm)$ time

MINIMUM WEIGHT BIPARTITE MATCHING



Hungarian algorithm: $O(nm)$ time

What if we solve many **similar** instances? E.g.:

- instances sampled from a distribution,
- one instance slowly changing over time.

RECALL: LP FORMULATION OF MATCHING

Primal:

$$\begin{aligned} &\text{minimize} && \sum_{e \in E} c_e x_e \\ &\text{subject to} && \sum_{e \in N(v)} x_e = 1 \quad \forall v \in V \\ &&& x_e \geq 0 \quad \forall e \in E \end{aligned}$$

Dual:

$$\begin{aligned} &\text{maximize} && \sum_{v \in V} y_v \\ &\text{subject to} && y_u + y_v \leq c_{u,v} \quad \forall (u, v) \in E \end{aligned}$$

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input comes with **predicted dual** \hat{y}

There is an $O(m\sqrt{n} \cdot \min\{\|\hat{y} - y\|_1, \sqrt{n}\})$ -time algorithm

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input comes with **predicted dual** \hat{y}

There is an $O(m\sqrt{n} \cdot \min\{\|\hat{y} - y\|_1, \sqrt{n}\})$ -time algorithm

“Can I learn to predict duals?”

(you, maybe)

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input comes with **predicted dual** \hat{y}

There is an $O(m\sqrt{n} \cdot \min\{\|\hat{y} - y\|_1, \sqrt{n}\})$ -time algorithm

“Can I learn to predict duals?”

(you, maybe)

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input (edge costs) comes from distribution \mathcal{D} , maximum cost $\leq C$

Find optimal duals \hat{y} for $\tilde{O}(n^3 C^2)$ samples

W.h.p., \hat{y} approximately (up to $+\epsilon$) minimizes $\mathbb{E}_{c \sim \mathcal{D}} \|\hat{y} - y_{\text{OPT}}(c)\|_1$

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input comes with **predicted dual** \hat{y} There is an $O(m\sqrt{n} \cdot \min\{\|\hat{y} - y\|_1, \sqrt{n}\})$ -time algorithm

“Can I learn to predict duals?”

(you, maybe)

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input (edge costs) comes from distribution \mathcal{D} , maximum cost $\leq C$ Find optimal duals \hat{y} for $\tilde{O}(n^3 C^2)$ samplesW.h.p., \hat{y} approximately (up to $+\epsilon$) minimizes $\mathbb{E}_{c \sim \mathcal{D}} \|\hat{y} - y_{\text{OPT}}(c)\|_1$ **Proof:** bound VC-dim of $\{c \mapsto \|\hat{y} - y_{\text{OPT}}(c)\|_1 \mid c \in \mathbb{R}^n\}$; use PAC-learning toolbox

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input comes with **predicted dual** \hat{y} There is an $O(m\sqrt{n} \cdot \min\{\|\hat{y} - y\|_1, \sqrt{n}\})$ -time algorithm

“Can I learn to predict duals?”

(you, maybe)

Theorem:

(Dinitz et al., NeurIPS 2021)

Suppose input (edge costs) comes from distribution \mathcal{D} , maximum cost $\leq C$ Find optimal duals \hat{y} for $\tilde{O}(n^3 C^2)$ samples (i.e., point-wise **median** of dual solutions)W.h.p., \hat{y} approximately (up to $+\epsilon$) minimizes $\mathbb{E}_{c \sim \mathcal{D}} \|\hat{y} - y_{\text{OPT}}(c)\|_1$ **Proof:** bound VC-dim of $\{c \mapsto \|\hat{y} - y_{\text{OPT}}(c)\|_1 \mid c \in \mathbb{R}^n\}$; use PAC-learning toolbox

THANK YOU!