

- today:
- learning to search
 - SEARN
 - SPENS

Learning to search (L2S)

SEARN Hal Daume's phd thesis

$$h_w: X \rightarrow \mathcal{Y}$$

$$h_w(x) = \arg \max_{y \in \mathcal{Y}} s(x, y; w) \quad] \text{ parameterized search procedure}$$

today: special case of SEARN: learning to do greedy search

split \mathcal{Y} in ordered list of decisions

$$(y_1, y_2, y_3, \dots, y_T)$$

learn a classifier $\pi_w(\text{feature}(\hat{y}_1, \dots, \hat{y}_{t-1}; x)) = \hat{y}_t$
 ↑ classifier "decoding policy"

L2S framework

from $(x^{(i)}, y^{(i)})_{i=1}^n$ and $l(\cdot, \cdot)$

learn a good classifier/policy π_w s.t. $\hat{y}_t = \pi_w^\wedge(\psi(\hat{y}_1, \dots, \hat{y}_{t-1}; x))$
 $h_w(x) \triangleq (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_T)$ "greedy decstr"
 s.t. $l(y^{(i)}, h_w(x^{(i)}))$ is small

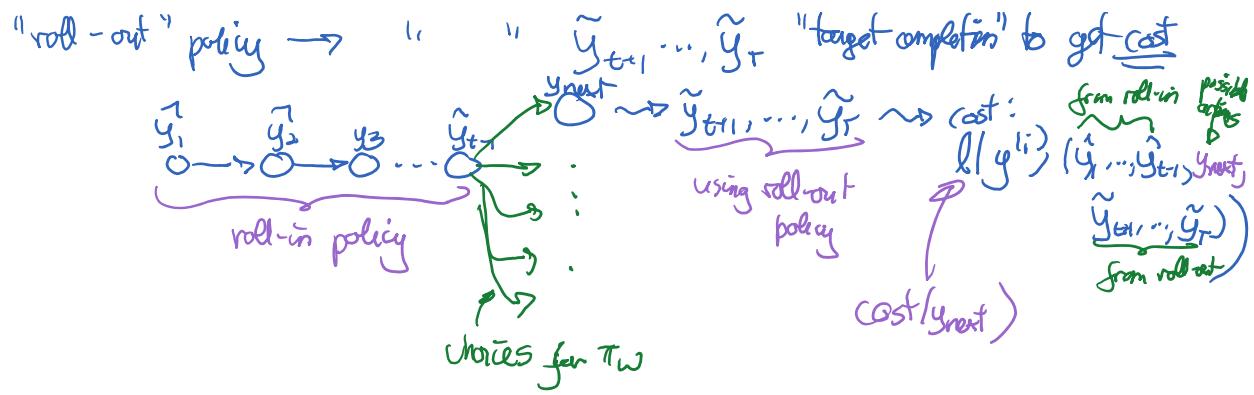
⊗ central idea: "reduction" that reduces the structured prediction problem
 to sequence of cost sensitive classification learning of π_w

method: generate training data for classifier π_w

$$(\underbrace{\hat{y}^{(i)}_{\text{context}}}_{\text{prefix sequence}}, x^{(i)}, \text{cost}(y^{(i)}))$$

prefix sequence to make next prediction

"roll in" policy → determines how $\hat{y}_{1:t-1}$ context



"reference policy", ideally, $\pi_{ref}(y_1, \dots, y_{t-1}, y_{next})$

$$\approx \underset{y_{completion}}{\operatorname{arg\min}} l(y^{(i)}, (y_1, \dots, y_{t-1}, y_{next}, y_{completion}))$$

Intractable. (NP hard) in general

In practice: use heuristics to approximate π

but sometimes, can compute exactly

e.g. π_{ref} for Hamming loss: is just copy ground truth

ICML 2015 "L2S better than your teacher"

L2S \rightarrow "locally optimal learning to search"

| roll-in | roll-out | (approximate) reference | mixture | learned π_W | learned π_T |
|---|----------|---|---|------------------------|---|
| reference policy ↳ use ground truth "teacher forcing" | | (1) | $y_2 \pi_{ref} + \frac{1}{2} \pi_{W^*}$ | inconsistent | ↳ a dist. D on $X \times Y$ s.t. π_W does well on cost-sensitive loss but how does poorly for structured prediction |
| learn (use π_W) | | <ul style="list-style-type: none"> • Consistent • not locally optimal ie. cannot do better than teacher | <ul style="list-style-type: none"> • Consistent • and "locally optimal" | reinforcement learning | ↳ (not using enough information. e.g. ground truth info.) can learn a policy that is "optimal up to one step deviation" |

(1)

example of approximate π_{ref} : l is bleu score
(in machine translation)

consider all possible suffixes from ground truth
and pick the best bleu score one

15h30

SEARNN: apply L2S to RNN training

[Leblond & al. ICLR 2018] ie $\pi_w(y_{1:t-1}, x) \rightarrow$ RNN cell

$p(y_{\text{next}} | y_{1:t-1}, x)$ of a (decoder) RNN

If you use $-\log p_w(\hat{y}_{\text{target}} | y_{1:t-1}, x)$ as a cost surrogate
and $\hat{y}_{\text{target}} = \text{ground truth}$

then L2S with π_{ref} roll-in is standard MLE

④ cost-sensitive surrogate losses (varios)

a) Structured SVM: $\max_{\tilde{y}} [\sum_i \Delta(\tilde{y}) + s(\tilde{y})] - s(y_{\text{target}})$
 $\Delta(\tilde{y}) \triangleq c(\tilde{y}) - c(y_{\text{target}})$

$$y_{\text{target}} \triangleq \arg \min_y c(y)$$

b) "target log-loss": $-\log p_w(y_{\text{target}} | y_{1:t-1}, x)$

↳ differences with MLE:

- address exposure bias using target roll-in
- make use of structured loss $\ell(\cdot)$ to predict y_{target} vs. ground truth for MLE

structured prediction energy networks (SPENs)

ICML 2016

$$E(x, y; w) \quad (-s(x, y; w))$$

• relax $y \in \{0, 1\}^T \rightsquigarrow [0, 1]^T$

$h_w(x) = \text{a few steps of projected G.D. on } E(\cdot, y; w) \quad (\text{approximate prediction procedure})$
 w.r.t. y
 + rounding

in 2016 paper: SSVM loss \rightsquigarrow "subgradient" method on w

$$\text{hinge loss} : \max_{\substack{\bar{y} \in [0, 1]^T \\ \text{augm. } \bar{y} \neq \bar{y}^*}} (l(y^{(i)}, \bar{y}) - E(x^{(i)}, \bar{y}; w)) + E(x^{(i)}, y^{(i)}; w)$$

requires extending \mathbb{Q} to fractional y 's

approximate "subgradient" is

because use \bar{y}^* as approximate max

$$-\nabla_w E(x^{(i)}, \bar{y}^*(w_t); w_t) + \nabla_w E(x^{(i)}, y^{(i)}; w)$$

e.g. see Clarke subdifferential
for generalization on non-convex fct.

2017 paper: end-to-end learning

$$h_w(x) = y_0 - \sum_{t=1}^T \eta_t \frac{d}{dy} E(x, y_t; w) \quad \text{"unrolled optimization"}$$

recursively defined

$$\begin{aligned} y_1 &= y_0 - \eta_1 \frac{d}{dy} E(x, y_0; w) \\ y_2 &= y_1 - \eta_2 \frac{d}{dy} E(x, y_1; w) \end{aligned}$$