

- today : - learning to search
- SeaRNN
 - SPENS

Learning to search (L2S)

SEARN Hal Daume's PhD thesis

$h_w: X \rightarrow Y$

$h_w(x) = \underset{y \in Y}{\text{argmax}} s(x, y; w)$ } parameterized search procedure

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

split y in ordered list of decisions

(y_1, y_2, \dots, y_T)

learn a classifier $\pi_w(\text{feature}(\hat{y}_1, \hat{y}_2, \dots, \hat{y}_{t-1}; x)) = \hat{y}_t$
 ↑ 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. $\hat{y}_t = \pi_w(\mathcal{L}(\hat{y}_{1:t-1}; x))$
 $h_w(x) \triangleq (\hat{y}_1, \dots, \hat{y}_T)$ "greedy decoder"
 s.t. $l(y^{(i)}, h_w(x^{(i)}))$ is small


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

method: generate training data for classifier π_w

$(\hat{y}_{\text{context}}^{(j)}, x^{(j)}, \text{cost}(y_{\text{next}}))$
 ↑ cost function of y_{next}
 prefer sequence to make next prediction

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

"roll out" policy → " " " " y_{t+1}, \dots, y_T "forget completion" to get





using roll-out policy

loss $(y_{next}) \dots (y_{t+1}, y_{next}, y_{t+1:t})$
 from roll-in from roll-out

choices of y_{next} for π_w

then
 idea is to define surrogate loss for π_w as function of $\text{cost}(y_{next})$

ex of roll-out:

"reference policy", ideally, $\pi_{ref}(\hat{y}_{1:t-1}, y_{next})$

$$= \underset{y_{completion}}{\text{argmin}} \ell(y_{li}, (\hat{y}_{1:t-1}, y_{next}, y_{completion}))$$

intractable (NP hard) in general

in practice: use heuristic to approximate π_{ref}

but sometimes, can compute exactly

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

LOLS \rightarrow "locally optimal learning to search"

ICML 2015 "LDS better than your teacher"

roll-in use ground truth reference policy "teacher forcing"	(approximate) π_{ref}	mixture $\frac{1}{2}\pi_{ref} + \frac{1}{2}\pi_w$	learned π_w
	(i)		inconsistent
learned (use π_w)	consistent not locally optimal	consistent and locally optimal	reinforcement learning

\rightarrow \exists a dist D on $X \times Y$ s.t. π_w does well on cost sensitive loss but π_w does poorly for structural prediction

\rightarrow (not using enough information e.g. ground truth predictions)

\rightarrow can learn a policy is optimal up to "one step deviation"

(i)

 cannot do better than teacher

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

consider all possible suffixes of ground truth and pick best bleu score

10/83

PEARL: apply LDS to RNN training

[Leland et al., ICML 2018] is $\pi_w(y_{1:t-1}, x) \rightarrow$ RNN cell

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

if you use $-\log p(\hat{y}_{target} | y_{1:t-1}, x)$ as a cost surrogate

if you use $-\log p(\hat{y}_{\text{target}} | y_{1:t-1}, x)$ as a cost surrogate
 and $\hat{y}_{\text{target}} = \text{ground truth}$
 then LAS with learned roll-in is standard MLE

⊗ cost-sensitive surrogate loss choices

a) structured SVM : $\max_y [\underbrace{\Delta(y')}_{\substack{\triangleq c(y') - c(y_{\text{target}})} \\ \substack{y_{\text{target}} \triangleq \text{argmin}_y c(y')}}] - S(y_{\text{target}})$

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

↳ differences with MLE:

- addresses exposure bias using learned roll-in
- make use of structured loss $\Delta(\cdot)$ to predict y_{target} v.s. ground truth in MLE

Structured prediction energy networks (SPENs)

ICML 2016

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

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

$hw(x) =$ a few steps of projected G.D. on $E(x, y; w)$ (approximate prediction procedure)
 w.r.t. to y
 + rounding

2016 paper:

SSVM loss \rightarrow "subgradient" method on w

large loss $\max_{y \in [0, 1]^T} (l(y^{(i)}, \bar{y}) - E(x^{(i)}, \bar{y}; w) + E(x^{(i)}, y^{(i)}; w))$
 requires estimation through PGD

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

because \bar{y}^* is approximate max

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

2017 paper: end-to-end training

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

$$y_0 \leq \int_{t=1}^t \frac{dy}{dy} E(x, y, \omega)$$

numerical optimization -

$$y_1 = y_0 - \eta_1 \frac{d}{dy} E(x, y_0, \omega)$$

$$y_2 = y_1 - \eta_2 \frac{d}{dy} E(x, y_1, \omega) \text{ etc...}$$