

today: • Learning to search  
• SEARNN  
• SPEN

## Learning to search (L2S)

SEARN Hal Daume's PhD thesis

$$h_w: X \rightarrow Y$$

$$h_w(x) = \underset{y \in Y}{\text{argmax}} s(x, y; w) \quad \left. \vphantom{\text{argmax}} \right\} \begin{array}{l} \text{parameterized} \\ \text{search procedure} \end{array}$$

special case of SEARN: learning to do greedy search

split  $Y$  in ordered list of decisions

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

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

classifier "decoding policy"

## L2S framework

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

learn a good classifier/policy  $\pi_w$  s.t.  $\hat{y}_t = \pi_w(\mathcal{F}(\hat{y}_1, \dots, \hat{y}_{t-1}, x))$

$$h_w(x) \triangleq (\hat{y}_1, \dots, \hat{y}_T) \quad \text{"greedy decoder"}$$

s.t.  $l(y^{(i)}, h_w(x^{(i)}))$  is good


⊛ central idea: "reduction" where reduces structured prediction learning to problem of cost sensitive classification learning for  $\pi_w$

method: generate training data for classifier  $\pi_w$

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

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

"roll-out" policy  $\rightarrow$  " " " "  $y_{t+1}, \dots, y_t$  "target completion" to get cost



"reference policy" ideally,  $\pi_{ref}(\hat{y}_1, \dots, \hat{y}_{t-1}, y_{next})$

$$= \arg \min_{y_{completion}} l(y^{(i)}, (\hat{y}_1, \dots, \hat{y}_{t-1}, y_{next}, y_{completion}))$$

intractable (NP hard) in general  
in practice: use heuristic to approximate it

but sometimes can compute exactly

e.g.  $\pi_{ref}$  Hamming loss: is just copy ground truth

LOLS  $\rightarrow$  "locally optimal learning to search"  
ICML 2015 "L2S better than your teacher"

roll-in \ roll-out	(approximate) reference	mixture $\frac{1}{2}\pi_{ref} + \frac{1}{2}\pi_{ref}$	learned $\pi_{ref}$
reference policy <i>"kacher Strategie"</i>	(1) $\leftarrow$		inconsistent $\rightarrow$
learned (use $\pi_{ref}$ )	consistent <u>not locally optimal</u>	consistent and <u>locally optimal</u>	reinforcement learning (not using enough information) e.g. ground truth

*Annotations:*  
 - From the 'inconsistent' cell:  $\rightarrow$  If a dist.  $D$  on  $X \times Y$  s.t.  $\pi_{ref}$  does well on cost sensitive loss but how does poorly for structural pred.  
 - From the 'consistent and locally optimal' cell:  $\rightarrow$  can learn a policy is optimal up "to one step deviation"

(1)

Example of approximate  $\pi_{ref}$ :  $l$  is bleu score (in machine translation)

consider all possible suffixes and pick

15h35

the best BLEU-1 scoring

SEARNN : apply L2S to RNN training

[Leblond & al. ICLR 2018]

i.e.  $\pi_w(y_{1:t+1}, x) \rightsquigarrow$  RNN cell

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

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

and  $\hat{y}_{target} =$  ground truth

then L2S with ref roll-in is standard MLE

\*) cost-sensitive surrogate loss choices:

a) structural SVM:  $\max_{y'} [\Delta(y') + s(y')] - s(y_{target})$   
 $\stackrel{\Delta}{=} c(y') - c(y_{target})$

$y_{target} \triangleq \text{argmin}_{y'} c(y')$

b) "target log-loss":  $-\log p(\hat{y}_{target} | y_{1:t-1}, x)$

↳ differences with MLE:

• address exposure bias using learned roll-in

• make use of structural loss  $l(\dots)$  to predict  $\hat{y}_{target}$  vs. MLE

Kushal Arora: How (not) to Train your Generative Model: Scheduled Sampling, Likelihood, Adversary? by Huszar answers this question why SS might not work well. This is what I observed as well.

Structural prediction energy networks (SPEs)

ICML 2016

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

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

$h_w(x) =$  a few steps of projected G.D. on  $E(x, y; w)$  / approximate prediction  
w.r. to  $y$

$hw(x) =$  a few steps of projected G.D. on  $E(x, y; w)$  (approximate prediction procedure)  
 w.r. to  $\underline{y}$

2016 paper:

SSVM loss  $\rightarrow$  "subgradient" method on  $w$

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

requires an extension

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

$\underline{y}^*$  is approx max.

eg. see Clarke subdifferential

2017 paper: end-to-end learning:

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

"unrolled optimization"

recursively define