# Advances in Structured Prediction 



Hal Daumé III
U Maryland me@hal3.name

Examples of structured piediction

## Sequence labeling

$\mathrm{x}=$ the monster ate the sandwich
$\mathrm{y}=\mathrm{Dt}$ Nn Vb Dt Nn
x = Yesterday I traveled to Lille
$y=\quad$ PER - - LOC


## Natural language parsing



## (Bipartite) matching



## Machine translation

## Google

## Translate

This text has been automatically translated from Arabic:

| Moscow stressed tone against Iran on its |
| :--- |
| nuclear program. He called Russian Foreign |
| Minister Tehran to take concrete steps to |
| restore confidence with the international |
| cormunity, to cooperate fully with the IAEA. |
| Conversely Tehran expressed its willingness |

Translate text




 .
from Arabic to English EETA

## Image segmentation



## Protein secondary structure prediction



## Standard solution methods

I.Each prediction is independent
2.Shared parameters via "multitask learning"
3.Assume tractable graphical model; optimize 4.Hand-crafted

Predicting independently

- h : features of nearby voxels $\rightarrow$ class
- Ensure output is coherent at test time
- Very simple to implement, often efficient
* Cannot capture correlations between predictions
* Cannot optimize a joint loss


# Prediction with multitask bias 

- h : features $\rightarrow$ (hidden representation) $\rightarrow$ yes/no
- Share (hidden representation) across all classes
- All advantages of predicting independently
- May implicitly capture correlations
* Learning may be hard (... or not?)
× Still not optimizing a joint loss


# Optimizing graphical models 

- Encode output as a graphical model
- Learn parameters of that model to maximize:
- p(true labels | input)
- cvx u.b. on loss(true labels, predicted labels)
- Guaranteed consistent outputs
- Can capture correlations explicitly
x Assumed independence assumptions may not hold
* Computationally intractable with too many "edges" or non-decomposable loss function


## Back to the original problem...

- How to optimize a discrete, joint loss?
- Input:
- Truth:
- Outputs:
- Predicted: $\hat{\mathrm{y}} \in \mathrm{Y}(\mathrm{x})$
- Loss:
- Data:
loss(y, $\hat{\mathrm{y}}$ )
$x \in X$
$y \in Y(x)$
$Y(x)$
$\hat{y} \in Y(x)$
( $\mathrm{x}, \mathrm{y}$ ) ~ D

| I | can | can | a | can |
| :---: | :---: | :---: | :---: | :---: |
| Pro | Md | Vb | Dt | Nn |
| Pro | Md | Md | Dt | Vb |
| Pro | Md | Md | Dt | Nn |
| Pro | Md | Nn | Dt | Md |
| Pro | Md | Nn | Dt | Vb |
| Pro | Md | Nn | Dt | Nn |
| Pro | Md | Vb | Dt | Md |
| Pro | Md | Vb | Dt | Vb |

# Back to the original problem... 

- How to optimize a discrete, joint loss?
- Input:
$x \in X$
- Truth:
$\mathrm{y} \in \mathrm{Y}(\mathrm{x})$
- Outputs: $\mathrm{Y}(\mathrm{x})$
- Predicted: $\hat{\mathrm{y}} \in \mathrm{Y}(\mathrm{x})$
- Loss:
- Data:


## Goal:

find $h \in H$
such that $h(x) \in Y(x)$
minimizing
$\mathrm{E}_{(\mathrm{x}, \mathrm{y}) \sim \mathrm{D}}[\operatorname{loss}(\mathrm{y}, \mathrm{h}(\mathrm{x}))]$ based on N samples

$$
\left(x_{n}, y_{n}\right) \sim D
$$

## Challenges

- Output space is too big to exhaustively search:
- Typically exponential in size of input
- implies y must decompose in some way
(often: x has many pieces to label)
- Loss function has combinatorial structure:
- Intersection over union
- Edit Distance



## Decomposition of label

- Decomposition of y often implies an ordering

| I | can | can | a | can |
| :---: | :---: | :---: | :---: | :---: |
| Pro | Md | Vb | Dt | Nn |


|  |  | O | C | A | T | O | C | U |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0 | -1 | -2 | -3 | -4 | -5 | -6 | -7 |
| $\mathbf{0}$ | -1 | 1 | -0 | -1 | -2 | -3 | -4 | -5 |
| A | -2 | 0 | 0 | 1 | 0 | -1 | -2 | -3 |
| T | -3 | -1 | -1 | 0 | 2 | 1 | 0 | -1 |
| T | -4 | -2 | -2 | -1 | 1 | 1 | 0 | -1 |
| A | -5 | -3 | -3 | -1 | 0 | 0 | 0 | -1 |
| C | -6 | -4 | -2 | -2 | -1 | -1 | 1 | 0 |
| A | -7 | -5 | -3 | -1 | -2 | -2 | 0 | 0 |

- But sometimes not so obvious....

(we'll come back to this case later...)


## Search spaces

- When y decomposes in an ordered manner, a sequential decision making process emerges


Search spaces

- When y decomposes in an ordered manner, a sequential decision making process emerges


Policies

- A policy maps observations to actions



## Versus reinforcement learning



$$
=\Pi\left(\mathrm{O}_{1}\right)
$$

Classifier

## Goal: $\min _{\pi}$ E [ loss(п) ]

In learning to search (L2S):

- Labeled data at training time
$\Rightarrow$ can construct good/optimal policies
- Can "reset" and try the same example many times


# Labeled data $\rightarrow$ Reference policy 

Given partial traj. $\mathrm{a}_{1}, \mathrm{a}_{2}, \ldots, \mathrm{a}_{\mathrm{t}-1}$ and true label y
The minimum achievable loss is:

$$
\min _{\left(a_{t}, a_{t+1}, \ldots\right)} \operatorname{loss}(y, \hat{y}(\vec{a}))
$$

The optimal action is the corresponding $\mathrm{a}_{\mathrm{t}}$

The optimal policy is the policy that always selects the optimal action

Ingredients for learning to search

- Training data:
$\left(\mathrm{x}_{\mathrm{n}}, \mathrm{y}_{\mathrm{n}}\right) \sim \mathrm{D}$
- Output space:
$\mathrm{Y}(\mathrm{x})$
- Loss function:
loss(y, $\hat{\mathrm{y}}$ )
- Decomposition:
$\{0\},\{a\}, \ldots$
- Reference policy: rref(o, y)


## An analogy from playing Mario

From Mario Al competition 2009

Input:


Output: Jump in $\{0,1\}$ Right in $\{0,1\}$ Left in $\{0,1\}$ Speed in $\{0,1\}$

## High level goal:

Watch an expert play and learn to mimic her behavior

## Training (expert)



Warm-up: Supervised learning
I.Collect trajectories from expert $\boldsymbol{T}^{\text {ref }}$
2.Store as dataset $\mathbf{D}=\left\{\left(\mathrm{o}, \mathrm{m}^{\text {ref }}(\mathrm{o}, \mathrm{y})\right) \mid \mathrm{o} \sim \mathrm{m}^{\text {ref }}\right\}$ 3.Train classifier $\boldsymbol{\pi}$ on $\mathbf{D}$

- Let II play the game!



## Test-time execution (sup. learning)



What's the (biggest) failure mode?
The expert never gets stuck next to pipes
$\Rightarrow$ Classifier doesn't learn to recover!

## Warm-up II: Imitation learning

I. Collect trajectories from expert $\mathrm{T}^{\text {ref }}$
2. Dataset $\mathbf{D}_{0}=\left\{\left(\mathrm{o}, \mathrm{m}^{\mathrm{ref}}(\mathrm{o}, \mathrm{y})\right) \mid \circ \sim \pi^{\mathrm{ref}}\right\}$
3. Train $\Pi_{I}$ on $D_{0}$
4. Collect new trajectories from $\Pi_{1}$

- But let the expert steer!

5. Dataset $D_{I}=\left\{\left(o, \pi^{\text {ree }}(\mathrm{o}, \mathrm{y})\right) \mid \circ \sim \Pi_{\mathrm{I}}\right\}$
6. Train $\Pi_{2}$ on $D_{0} U D_{1}$

- In general:
- $D_{n}=\left\{\left(o, \Pi^{r e f}(0, y)\right) \mid \circ \sim \pi_{n}\right\}$
- Train $\pi_{n+1}$ on $\mathbf{U}_{i \leq n} \mathbf{D}_{\mathrm{i}}$



## Test-time execution (DAgger)


Grentidis



What's the biggest failure mode?
Classifier only sees right versus not-right

- No notion of better or worse
- No partial credit
- Must have a single target answer



# Aside: cost-sensitive classification 

Classifier: $\mathrm{h}: \mathrm{x} \rightarrow[\mathrm{K}]$
Multiclass classification

- Data: $(\mathrm{x}, \mathrm{y}) \in \mathrm{X} \times[\mathrm{K}]$
- Goal: $\min _{\mathrm{h}} \operatorname{Pr}(\mathrm{h}(\mathrm{x}) \neq \mathrm{y})$

Cost-sensitive classification

- Data: $(\mathrm{x}, \mathrm{c}) \in \mathrm{X} \times[0, \infty)^{\mathrm{K}}$
- Goal: $\min _{\mathrm{h}} \mathrm{E}_{(\mathrm{x}, \mathrm{c})}\left[\mathrm{c}_{\mathrm{h}(\mathrm{x})}\right]$


## Learning to search: AggraVaTe

I.Let learned policy $\pi$ drive for t timesteps to obs. o
2.For each possible action a:

- Take action a, and let expert $\Pi^{\text {ref }}$ drive the rest
- Record the overall loss, $\mathrm{c}_{\mathrm{a}}$
3.Update $\pi$ based on example:

$$
\left(\mathrm{o},\left\langle\mathrm{c}_{1}, \mathrm{c}_{2}, \ldots, \mathrm{c}_{\mathrm{K}}\right\rangle\right)
$$

4. Goto (I)


Learning to search:
AggraVaTe
I.Generate an initial trajectory using the current policy

2.Foreach decision on that trajectory with obs. o:
a)Foreach possible action a (one-step deviations)
i. Take that action
ii. Complete this trajectory using reference policy
iii.Obtain a final loss, $\mathrm{C}_{\mathrm{a}}$
b)Generate a cost-sensitive classification example:

$$
(\mathrm{o}, \overrightarrow{\mathrm{c}})
$$

Learning to search:
AggraVaTe
I.Generate an initial trajectory using the current policy

2.Foreach decision on that trajectory with obs. o:
a)Foreach possible action a (one-step deviations)
i. Take that action

Often it's possible to analytically
ii. Complete this trajectory using compute this loss without iii.Obtain a final loss, $\mathrm{C}_{\mathrm{a}}$ having to execute a roll-out!
b)Generate a cost-sensitive classification example:

$$
(\mathrm{o}, \overrightarrow{\mathrm{c}})
$$

Example I: Sequence labeling

- Receive input:

$$
\mathrm{y}=\mathrm{Dt} \quad \mathrm{Nn} \quad \mathrm{Vb} \quad \mathrm{Dt} \quad \mathrm{Nn}
$$

- Make a sequence of predictions:

$$
\begin{aligned}
& x=\text { the monster ate the sandwich } \\
& \hat{y}=D t \quad D t \quad D t \quad D t .
\end{aligned}
$$

- Pick a timestep and try all perturbations there:

$$
\begin{aligned}
& \mathrm{x}=\mathrm{the} \text { monster ate the sandwich } \\
& \hat{\mathrm{y}}_{\mathrm{Dt}}=\mathrm{Dt} \\
& \mathrm{y}_{\mathrm{Nn}}=\mathrm{Dt} \\
& \hat{\mathrm{y}}_{\mathrm{Vb}}=\mathrm{Dt} \\
& \mathrm{Nn} \\
& \mathrm{Vb}
\end{aligned}
$$

- Compute losses and construct example:
( \{ w=monster, p=Dt, ...\},
$[1,0,1])$

Example II: Graph labeling

- Task: label nodes of a graph given node features (and possibly edge features)
- Example:WebKB webpage labeling

- Node features: text on web page
- Edge features: text in hyperlinks

Example II: Graph labeling

- How to linearize? Like belief propagation might!
- Pick a starting node (A), run BFS out
- Alternate outward and inward passes


Linearization: ABCDEFGHI HGFEDCBA BCDEFGHI HGFEDCBA

Example II: Graph labeling
I.Pick a node (= timestep)
2.Construct example based on neighbors' labels 3.Perturb current node's label


## Outline

(1) Empirics
(2) Analysis
(3) Programming
(9) Others and Issues

## What part of speech are the words?



## A demonstration

1 |w Despite
2 |w continuing
3 |w problems
1 |w in
4 |w its
5 |w newsprint
5 |w business
...

## A demonstration

1 |w Despite
2 |w continuing
3 |w problems
1 |w in
4 |w its
5 |w newsprint
5 |w business
vw -b 24 -d wsj.train.vw -c -search_task sequence -search 45 -search_alpha 1e-8 -search_neighbor_features -1:w,1:w
-affix $-\overline{1 w},+1 w-f$ foo.reg
vw -t -i foo.reg wsj.test.vw

## Is this word a name or not?

Named Entity Recognition (tuned hps)


## How fast in evaluation?

Prediction (test-time) Speed


## Entity Relation

Goal: find the Entities and then find their Relations Method $\quad$ Entity F1 $\quad$ Relation F1 $\quad$ Train Time

| Structured SVM | 88.00 | 50.04 | 300 seconds |
| :---: | :---: | :---: | :---: |
| L2S | 92.51 | 52.03 | 13 seconds |

L2S uses ~100 LOC.

## Find dependency structure of sentences.



L2S uses ~300 LOC.

## Outline

(1) Empirics
(2) Analysis
(3) Programming
(9) Others and Issues

## Effect of Roll-in and Roll-out Policies

| roll-out $\rightarrow$ <br> $\downarrow$ roll-in | Reference | Half-n-half | Learned |
| :--- | :--- | :--- | :--- |
| Reference | Inconsistent |  |  |
| Learned | $\cdot$ |  |  |

## Effect of Roll-in and Roll-out Policies

| roll-out $\rightarrow$ <br> $\downarrow$ roll-in | Reference | Half-n-half | Learned |
| :--- | :--- | :--- | :--- |
| Reference | Inconsistent |  |  |
| Learned | $\cdot$ |  |  |

## Theorem

Roll-in with ref:
0 cost-sensitive regret $\Rightarrow$ unbounded joint regret

## Effect of Roll-in and Roll-out Policies

| roll-out $\rightarrow$ <br> $\downarrow$ roll-in | Reference | Half-n-half | Learned |
| :--- | :--- | :--- | :--- |
| Reference | Inconsistent |  |  |
| Learned | Consistent <br> No local opt |  |  |



## Effect of Roll-in and Roll-out Policies

| roll-out $\rightarrow$ <br> $\downarrow$ roll-in | Reference | Half-n-half | Learned |
| :--- | :--- | :--- | :--- |
| Reference | Inconsistent |  |  |
| Learned | Consistent <br> No local opt |  |  |

## Theorem

Roll-out with Ref:
0 cost-sensitive regret $\Rightarrow 0$ joint regret (but not local optimality)

## Effect of Roll-in and Roll-out Policies

| roll-out $\rightarrow$ <br> $\downarrow$ roll-in | Reference | Half-n-half | Learned |
| :--- | :--- | :--- | :--- |
| Reference | Inconsistent |  |  |
| Learned | Consistent <br> No local opt |  | Reinf. L. |

## Theorem

Ignore Ref:
$\Rightarrow$ Equivalent to reinforcement learning.

## Effect of Roll-in and Roll-out Policies

| roll-out $\rightarrow$ <br> $\downarrow$ roll-in | Reference | Half-n-half | Learned |
| :--- | :--- | :--- | :--- |
| Reference | Inconsistent |  |  |
| Learned | Consistent <br> No local opt | Consistent <br> Local Opt | Reinf. L. |

## Theorem

Roll-out with $p=0.5$ Ref and $p=0.5$ Learned: 0 cost-sensitive regret $\Rightarrow 0$ joint regret + locally optimal

See LOLS paper, Wednesday 11:20 Van Gogh

## AggreVaTe Regret Decomposition

$\pi^{\text {ref }}=$ reference policy
$\bar{\pi}=$ stochastic average learned policy $J(\pi)=$ expected loss of $\pi$.

Theorem
$J(\bar{\pi})-J\left(\pi^{r e f}\right) \leq$

## AggraVaTe Regret Decomposition

$\pi^{\text {ref }}=$ reference policy
$\bar{\pi}=$ stochastic average learned policy
$J(\pi)=$ expected loss of $\pi$.
Theorem

$$
\begin{aligned}
& J(\bar{\pi})-J\left(\pi^{r e f}\right) \leq \\
& \quad T \mathbb{E}_{n, t} \mathbb{E}_{x \sim D_{\pi_{n}}^{t}}\left[Q^{\pi^{r e f}}\left(x, \hat{\pi}_{n}\right)-Q^{\pi^{r e f}}\left(x, \pi^{r e f}\right)\right]
\end{aligned}
$$

$T=$ number of steps
$\hat{\pi}_{n}=n$th learned policy
$D_{\hat{\pi}_{n}}^{t}=$ distribution over $x$ at time $t$ induced by $\hat{\pi}_{n}$
$Q^{\pi^{\prime}}\left(x, \pi^{\prime}\right)=$ loss of $\pi^{\prime}$ at $x$ then $\pi$ to finish

## Proof

For all $\pi$ let $\pi^{t}$ play $\pi$ for rounds $1 \ldots t$ then play $\pi^{\text {ref }}$
for rounds $t+1 \ldots T$. So $\pi^{T}=\pi$ and $\pi^{0}=\pi^{\text {ref }}$

## Proof

For all $\pi$ let $\pi^{t}$ play $\pi$ for rounds $1 \ldots t$ then play $\pi^{\text {ref }}$ for rounds $t+1 \ldots T$. So $\pi^{T}=\pi$ and $\pi^{0}=\pi^{\text {ref }}$ $J(\pi)-J\left(\pi^{\text {ref }}\right)$

$$
=\sum_{t=1}^{T} J\left(\pi^{t}\right)-J\left(\pi^{t-1}\right) \text { (Telescoping sum) }
$$

## Proof

For all $\pi$ let $\pi^{t}$ play $\pi$ for rounds $1 \ldots t$ then play $\pi^{\text {ref }}$ for rounds $t+1 \ldots T$. So $\pi^{T}=\pi$ and $\pi^{0}=\pi^{\text {ref }}$ $J(\pi)-J\left(\pi^{\text {ref }}\right)$

$$
=\sum_{t=1}^{T} J\left(\pi^{t}\right)-J\left(\pi^{t-1}\right) \text { (Telescoping sum) }
$$

$$
=\sum_{t=1}^{T} \mathbb{E}_{x \sim D_{\pi}^{t}}\left[Q^{\pi^{r e f}}(x, \pi)-Q^{\pi^{r e f}}\left(x, \pi^{\mathrm{ref}}\right)\right]
$$

since for all $\pi, t, J(\pi)=\mathbb{E}_{x \sim D_{\pi}^{t}} Q^{\pi}(x, \pi)$

## Proof

For all $\pi$ let $\pi^{t}$ play $\pi$ for rounds $1 \ldots t$ then play $\pi^{\text {ref }}$ for rounds $t+1 \ldots T$. So $\pi^{T}=\pi$ and $\pi^{0}=\pi^{\text {ref }}$ $J(\pi)-J\left(\pi^{\text {ref }}\right)$

$$
\begin{aligned}
& =\sum_{t=1}^{T} J\left(\pi^{t}\right)-J\left(\pi^{t-1}\right)(\text { Telescoping sum }) \\
& =\sum_{t=1}^{T} \mathbb{E}_{x \sim D_{\pi}^{t}}\left[Q^{\pi^{\mathrm{ref}}}(x, \pi)-Q^{\pi^{\mathrm{ref}}}\left(x, \pi^{\mathrm{ref}}\right)\right]
\end{aligned}
$$

since for all $\pi, t, J(\pi)=\mathbb{E}_{x \sim D_{\pi}^{t}} Q^{\pi}(x, \pi)$

$$
=T \mathbb{E}_{t} \mathbb{E}_{x \sim D_{\pi}^{t}}\left[Q^{\pi^{\mathrm{ref}}}(x, \pi)-Q^{\pi^{\mathrm{ref}}}\left(x, \pi^{\mathrm{ref}}\right)\right]
$$

## Proof

For all $\pi$ let $\pi^{t}$ play $\pi$ for rounds $1 \ldots t$ then play $\pi^{\text {ref }}$ for rounds $t+1 \ldots T$. So $\pi^{T}=\pi$ and $\pi^{0}=\pi^{\text {ref }}$ $J(\pi)-J\left(\pi^{\text {ref }}\right)$

$$
\begin{aligned}
& =\sum_{t=1}^{T} J\left(\pi^{t}\right)-J\left(\pi^{t-1}\right)(\text { Telescoping sum }) \\
& =\sum_{t=1}^{T} \mathbb{E}_{x \sim D_{\pi}^{t}}\left[Q^{\pi^{\mathrm{ref}}}(x, \pi)-Q^{\pi^{\mathrm{ref}}}\left(x, \pi^{\mathrm{ref}}\right)\right]
\end{aligned}
$$

since for all $\pi, t, J(\pi)=\mathbb{E}_{x \sim D_{\pi}^{t}} Q^{\pi}(x, \pi)$
$=T \mathbb{E}_{t} \mathbb{E}_{x \sim D_{\pi}^{t}}\left[Q^{\pi^{r e f}}(x, \pi)-Q^{\pi^{\text {ref }}}\left(x, \pi^{\text {ref }}\right)\right]$
So $J(\bar{\pi})-J\left(\pi^{\text {ref }}\right)$
$=T \mathbb{E}_{t, n} \mathbb{E}_{x \sim D_{\hat{\pi}_{n}}^{t}}\left[Q^{\pi^{r e f}}\left(x, \hat{\pi}_{n}\right)-Q^{\pi^{\mathrm{ref}}}\left(x, \pi^{\mathrm{ref}}\right)\right]$

## Outline

(1) Empirics
(2) Analysis
(3) Programming
(9) Others and Issues

## Lines of Code



## How？

## Sequential＿RUN（examples）

1：for $i=1$ to len（examples）do
2：$\quad$ prediction $\leftarrow \operatorname{predict}($ examples［ $[i]$ ，examples $[i]$ ．label）
3：$\quad \operatorname{loss}($ prediction $\neq$ examples［ $[$ ］．label）
4：end for

## How?

## Sequential_RUN(examples)

1: for $i=1$ to len(examples) do
2: $\quad$ prediction $\leftarrow \operatorname{predict}($ examples $[i]$, examples $[i]$.label)
3: $\quad$ loss(prediction $\neq$ examples[ $]$. . label)
4: end for

Decoder + loss + reference advice

RunParser(sentence)

1: stack $S \leftarrow\{$ Root $\}$
2: buffer $B \leftarrow$ [words in sentence]
3: $\operatorname{arcs} A \leftarrow \emptyset$
4: while $B \neq \emptyset$ or $|S|>1$ do
5: $\quad$ ValidActs $\leftarrow G e t V a l i d A c t i o n s(S, B)$
6: $\quad$ features $\leftarrow G e t F e a t(S, B, A)$
7: $\quad$ ref $\leftarrow \operatorname{GetGoldAction~}(S, B)$
8: action $\leftarrow \operatorname{predict}($ features, ref, ValidActs)
9: $\quad S, B, A \leftarrow \operatorname{Transition}(S, B, A$, action)
10: end while
11: $\boldsymbol{\operatorname { l o s s }}\left(A[w] \neq A^{*}[w], \forall w \in\right.$ sentence $)$
12: return output

## Program/Search equivalence

Theorem: Every algorithm which:
(1) Always terminates.
(2) Takes as input relevant feature information $X$.
(3) Make 0+ calls to predict.
(a) Reports loss on termination.
defines a search space, and such an algorithm exists for every search space.

## It even works in Python

def _run(self, sentence):
output $=[]$
for n in range(len(sentence)):
pos,word $=$ sentence[ $n$ ]
with self.vw.example('w': [word],
' p ': [prev_word]) as ex:
pred $=$ self.sch.predict(examples=ex,
my_tag $=\mathrm{n}+1$, oracle= pos , condition $=\left[(n, ' p\right.$ ' $\left.\left.),\left(n-1, q^{\prime}\right)\right]\right)$
output.append(pred)
return output

## Bugs you cannot have

(1) Never train/test mismatch.

## Bugs you cannot have

(1) Never train/test mismatch.
(2) Never unexplained slow.

## Bugs you cannot have

（1）Never train／test mismatch．
（2）Never unexplained slow．
（3）Never fail to compensate for cascading failure．

## Outline

(1) Empirics
(2) Analysis
(3) Programming
(9) Others and Issues
(1) Families of algorithms.
(2) What's missing from learning to search?

## Imitation Learning

Use perceptron-like update when learned deviates from gold standard.
Inc. P. Collins \& Roark, ACL 2004.
LaSo Daume III \& Marcu, ICML 2005.
Local Liang et al, ACL 2006.
Beam P. Xu et al., JMLR 2009.
Inexact Huang et al, NAACL 2012.

## Imitation Learning

Use perceptron-like update when learned deviates from gold standard.
Inc. P. Collins \& Roark, ACL 2004.
LaSo Daume III \& Marcu, ICML 2005.
Local Liang et al, ACL 2006.
Beam P. Xu et al., JMLR 2009.
Inexact Huang et al, NAACL 2012.
Train a classifier to mimic an expert's behavior
DAgger Ross et al., AlStats 2011.
Dyna O Goldberg et al., TACL 2014.

## Learning to Search

When the reference policy is optimal
Searn Daume III et al., MLJ 2009.
Aggra Ross \& Bagnell, http://arxiv.org/pdf/1406.5979

## Learning to Search

When the reference policy is optimal
Searn Daume III et al., MLJ 2009.
Aggra Ross \& Bagnell, http://arxiv.org/pdf/1406.5979

When it's not
LOLS Chang et al., ICML 2015.

## Learning to Search

When the reference policy is optimal
Searn Daume III et al., MLJ 2009.
Aggra Ross \& Bagnell, http://arxiv.org/pdf/1406.5979

When it's not
LOLS Chang et al., ICML 2015.

Code in Vowpal Wabbit http://hunch.net/~vw

## Inverse Reinforcement Learning

Given observed expert behavior, infer the underlying reward function the expert seems to be optimizing

## Inverse Reinforcement Learning

Given observed expert behavior, infer the underlying reward function the expert seems to be optimizing propose Kalman, 1968.
1st sol. Boyd, 1994.

## Inverse Reinforcement Learning

Given observed expert behavior, infer the underlying reward function the expert seems to be optimizing propose Kalman, 1968.
1st sol. Boyd, 1994.
from sample trajectories only
Ng \& Russell, ICML 2000

## Inverse Reinforcement Learning

Given observed expert behavior, infer the underlying reward function the expert seems to be optimizing propose Kalman, 1968.
1st sol. Boyd, 1994.
from sample trajectories only
Ng \& Russell, ICML 2000
for apprenticeship learning
Apprent. Abbeel \& Ng, ICML 2004
Maxmar. Ratliff et al., NIPS 2005
MaxEnt Ziebart et al., AAAI 2008

## What＇s missing？Automatic Search order

Learning to search $\simeq$ dependency + search order． Graphical models＂work＂given dependencies only．

## What's missing? The reference policy

A good reference policy is often nonobvious... yet critical to performance.

## What's missing?

## Efficient Cost-Sensitive Learning

When choosing 1-of- $k$ things, $O(k)$ time is not exciting for machine translation.

## What's missing? GPU fun

Vision often requires a GPU. Can that be done?

## How to optimize discrete joint loss？

## How to optimize discrete joint loss?

(1) Programming complexity.

## How to optimize discrete joint loss？

（1）Programming complexity．Most complex problems addressed independently－too complex to do otherwise．

## How to optimize discrete joint loss？

（1）Programming complexity．Most complex problems addressed independently－too complex to do otherwise．
（2）Prediction accuracy．It had better work well．

## How to optimize discrete joint loss?

(1) Programming complexity. Most complex problems addressed independently-too complex to do otherwise.
(2) Prediction accuracy. It had better work well.
(3) Train speed. Debug/development productivity + maximum data input.

## How to optimize discrete joint loss?

(1) Programming complexity. Most complex problems addressed independently-too complex to do otherwise.
(2) Prediction accuracy. It had better work well.
(3) Train speed. Debug/development productivity + maximum data input.
(1) Test speed. Application efficiency

