# Learning Where to Sample in Structured Prediction

Tianlin Shi Jacob Steinhardt Percy Liang

Tsinghua University 清華大學 Stanford University

AISTATS 2015

# Outline

# Setting

- Have "stolen" a prediction model for structured outputs:

$$p(\underbrace{y_1, y_2, \ldots, y_n}_{\text{Output}} \mid \underbrace{\mathbf{x}}_{\text{Input}})$$

# Setting

- Have "stolen" a prediction model for structured outputs:

$$p(\underbrace{y_1, y_2, \ldots, y_n}_{\text{Output}} | \underbrace{\mathbf{x}}_{\text{Input}})$$

- At run time, use Gibbs sampling to do stochastic search:

$$\arg\max_{y_1,\ldots,y_n} p(y_1, \ldots, y_n | \mathbf{x})$$

# Setting

- Have "stolen" a prediction model for structured outputs:

$$p(\underbrace{y_1, y_2, ..., y_n}_{\text{Output}} | \underbrace{\mathbf{x}}_{\text{Input}})$$

- At run time, use Gibbs sampling to do stochastic search:

$$\arg\max_{y_1,...,y_n} p(y_1, \ldots, y_n|\mathbf{x})$$

- **Goal** Optimize the running time of Gibbs sampler! How?

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

**Example (Part-of-Speech Tagging)**

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | | | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: y | Pronoun | Verb | Adverb | Verb | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: y | Pronoun | Verb | Adverb | Verb | Determiner | Noun | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: y | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | | | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: $\mathbf{y}$ | Pronoun | Verb | Noun | Verb | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: $\mathbf{y}$ | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: $\mathbf{y}$ | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | | | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: y | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: y | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: y | Pronoun | Verb | Adverb | Verb | Determiner | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: $\mathbf{y}$ | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: $\mathbf{y}$ | | | | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| **x** | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | | | Noun | | | | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| x | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: **y** | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: **y** | | | Noun | | | Adjective | |

# Dissection: Gibbs Sampler

**A small cost** for a local move, but **a large number** of moves.

## Example (Part-of-Speech Tagging)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: $\mathbf{y}$ | Pronoun | Verb | Noun | Verb | Determiner | Adjective | Noun |

**Source of inefficiency**. Some parts are harder, while some are easier.

## Example (A Better Strategy)

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| pass 1: $\mathbf{y}$ | Pronoun | Verb | Adverb | Verb | Determiner | Noun | Noun |
| pass 2: $\mathbf{y}$ | | | Noun | | | Adjective | |

A HeteroSampler! ("Heterogeneous Sampler")

– Focus computation to where needed.

# Framework

### Definition

Action $A_j$ updates part $y_j$ based on

$$p(y_j|\mathbf{y}_{-j}, \mathbf{x})$$

# Framework

## Definition

Action $A_j$ updates part $y_j$ based on

$$p(y_j|\mathbf{y}_{-j}, \mathbf{x})$$

## Example

| Example | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| Input $x_j$ | | | $x_3$ | | | | |
| Output $y_j$ | | | $y_3$ | | | | |
| Action $A_j$ | | | $A_3$ | | | | |

$A_3$ samples $y_3$ from $p(y_3|y_{-3}, \mathbf{x})$

# Sampler Template

Our sampler: for total of $T$ rounds, do

# Sampler Template

Our sampler: for total of $T$ rounds, do

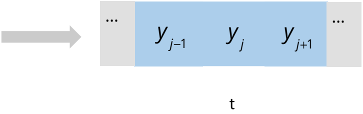

# Sampler Template

Our sampler: for total of $T$ rounds, do

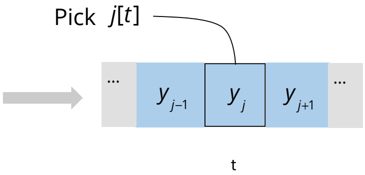


1. **Pick** index $j$ and the action $A_j$.

# Sampler Template

Our sampler: for total of $T$ rounds, do

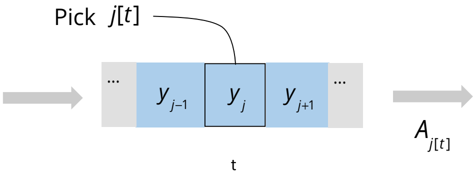


1. **Pick** index $j$ and the action $A_j$.
2. **Sample** $y_j$ using $A_j$

# Sampler Template

Our sampler: for total of $T$ rounds, do

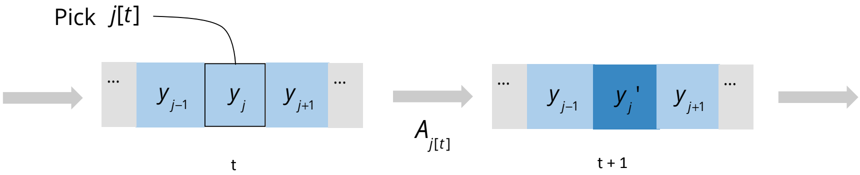


1. **Pick** index $j$ and the action $A_j$.
2. **Sample** $y_j$ using $A_j$

# Sampler Template

Our sampler: for total of $T$ rounds, do

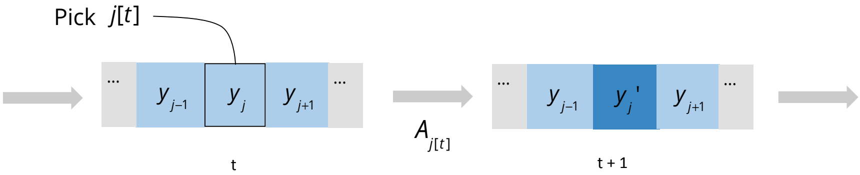


1. **Pick** index $j$ and the action $A_j$.
2. **Sample** $y_j$ using $A_j$

### Example

- **Cyclic Gibbs sampler.** Pick $j$ round-robin.

# Sampler Template

Our sampler: for total of $T$ rounds, do

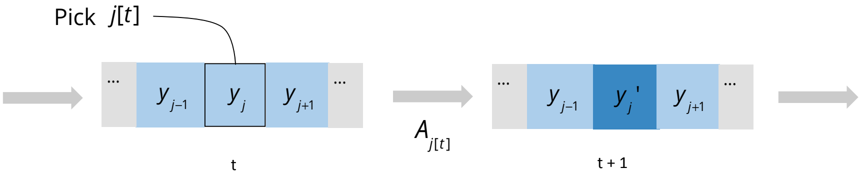


1. **Pick** index $j$ and the action $A_j$.
2. **Sample** $y_j$ using $A_j$

**Example**

- **Cyclic Gibbs sampler.** Pick $j$ round-robin.
- **Random-Scan Gibbs sampler.** Pick $j$ uniformly at random.

# Sampler Template

Our sampler: for total of $T$ rounds, do

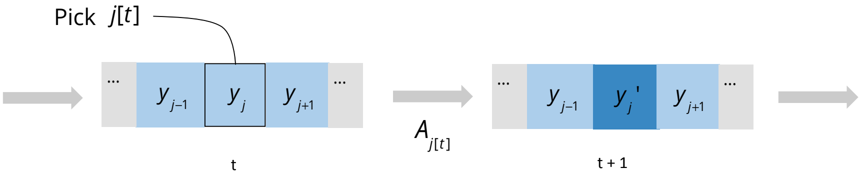


1. **Pick** index $j$ and the action $A_j$.
2. **Sample** $y_j$ using $A_j$

## Example

- **Cyclic Gibbs sampler.** Pick $j$ round-robin.
- **Random-Scan Gibbs sampler.** Pick $j$ uniformly at random.

How to choose $j$?

# Outline

# Reinforcement Learning

- **State** = Entire history with configurations $\mathbf{y}[i]$ and choices $j[i]$

$$s_t = (\mathbf{y}[0] \ldots, \mathbf{y}[t], j[0], \ldots, j[t-1])$$

# Reinforcement Learning

- **State** = Entire history with configurations $\mathbf{y}[i]$ and choices $j[i]$

$$s_t = (\mathbf{y}[0] \ldots, \mathbf{y}[t], j[0], \ldots, j[t-1])$$

- **Action** = Pick a $j[t]$ to sample

$$a_t = A_{j[t]}$$

# Reinforcement Learning

- **State** = Entire history with configurations $\mathbf{y}[i]$ and choices $j[i]$

$$s_t = (\mathbf{y}[0] \ldots, \mathbf{y}[t], j[0], \ldots, j[t-1])$$

- **Action** = Pick a $j[t]$ to sample

$$a_t = A_{j[t]}$$

- **Transition** = Pre-Trained Model $p(y_j | \mathbf{y}_{\neg j}, \mathbf{x})$

# Reinforcement Learning

- **State** = Entire history with configurations $\mathbf{y}[i]$ and choices $j[i]$

$$s_t = (\mathbf{y}[0] \ldots, \mathbf{y}[t], j[0], \ldots, j[t-1])$$

- **Action** = Pick a $j[t]$ to sample

$$a_t = A_{j[t]}$$

- **Transition** = Pre-Trained Model $p(y_j | \mathbf{y}_{\neg j}, \mathbf{x})$
- **Reward** = Improvement in log-probability

$$\mathcal{R}(s_t, a_t, s_{t+1}) = \log p(\mathbf{y}[t+1] \mid \mathbf{x}) - \log p(\mathbf{y}[t] \mid \mathbf{x})$$

- Policy $\pi$: How to pick.

$$\pi : \text{states} \rightarrow \text{actions}$$

- Policy $\pi$: How to pick.

$$\pi : \text{states} \rightarrow \text{actions}$$

- The expected cumulative reward $R_T$

$$\mathbb{E}[R_T] = \mathbb{E}[\sum_{t=0}^{T-1} \mathcal{R}(s_t, a_t, s_{t+1})].$$

- Policy $\pi$: How to pick.

$$\pi : \text{states} \rightarrow \text{actions}$$

- The expected cumulative reward $R_T$

$$\mathbb{E}[R_T] = \mathbb{E}[\sum_{t=0}^{T-1} \mathcal{R}(s_t, a_t, s_{t+1})].$$

**Remark**

$$\textit{Cumulative reward} = \log p(\mathbf{y}[T] \mid \mathbf{x}) - \log p(\mathbf{y}[0] \mid \mathbf{x}) \tag{1}$$

*Maximizing cumulative reward is equivalent to maximizing probability of final sample.*

# Learning Algorithm

Inspired by standard RL (Q-learning [Watkins et al. 1992], SARSA [Rummery et al. 1994]),

$$Q(s, a)$$

# Learning Algorithm

Inspired by standard RL (Q-learning [Watkins et al. 1992], SARSA [Rummery et al. 1994]),

$$Q(s, a) := \textbf{how good it is to take action } a \textbf{ in state } s$$

# Learning Algorithm

Inspired by standard RL (Q-learning [Watkins et al. 1992], SARSA [Rummery et al. 1994]),

$$Q(s, a) := \textbf{how good it is to take action } a \textbf{ in state } s$$

## Example

| $\mathbf{x}$ | I | think | now | is | the | right | time |
|---|---|---|---|---|---|---|---|
| $\mathbf{y}[0]$ | P | V | Adv | V | DT | N | N |
| $Q(s, a)$ | 0.0 | 0.0 | 2.0 | 0.0 | 0.0 | 2.3 | 0.0 |

# Applying RL

## Challenge: Efficiency

# Applying RL

### Challenge: Efficiency

$Q(s, a)$ should be cheap to compute, so it does not become the computational bottleneck.

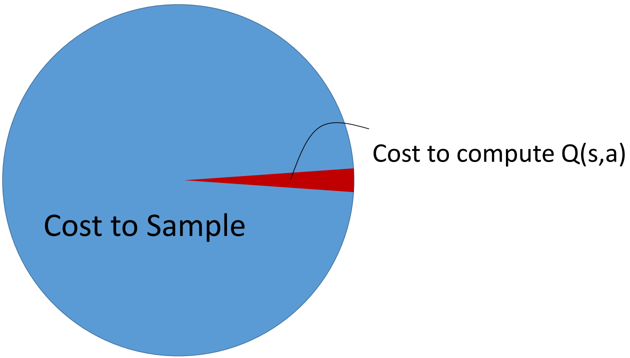

# Applying RL: Efficiency

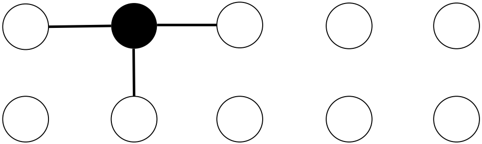

# Applying RL: Efficiency

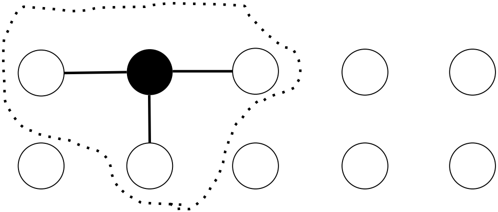

1. **Exploit locality.** Model $Q(s, a)$ using local meta-features only.

# Applying RL: Efficiency

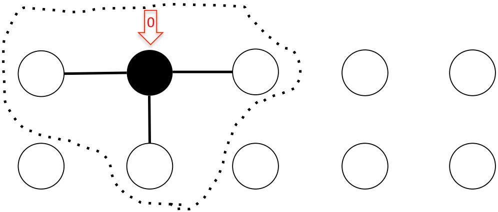


1. **Exploit locality.** Model $Q(s, a)$ using local meta-features only.
2. **Catch.** Local meta-features can't predict cumulative reward.

# Applying RL: Efficiency

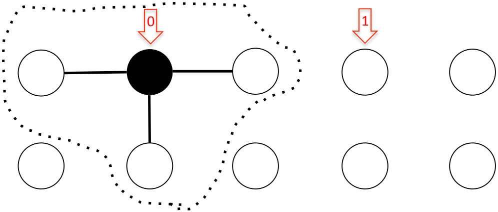


1. **Exploit locality.** Model $Q(s, a)$ using local meta-features only.
2. **Catch.** Local meta-features can't predict cumulative reward.

# Applying RL: Efficiency

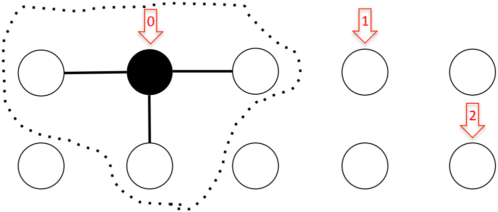


1. **Exploit locality.** Model $Q(s, a)$ using local meta-features only.
2. **Catch.** Local meta-features can't predict cumulative reward.

# Applying RL: Efficiency

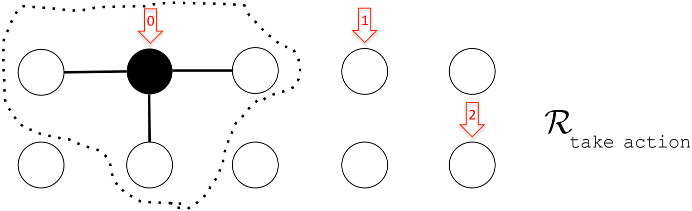


1. **Exploit locality.** Model $Q(s, a)$ using local meta-features only.
2. **Catch.** Local meta-features can't predict cumulative reward.
3. **Credit Assignment.** Isolate the contribution of $a$.

# Applying RL: Efficiency

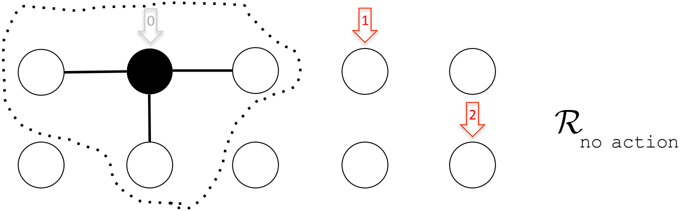


1. **Exploit locality.** Model $Q(s, a)$ using local meta-features only.
2. **Catch.** Local meta-features can't predict cumulative reward.
3. **Credit Assignment.** Isolate the contribution of $a$.

# Applying RL: Efficiency

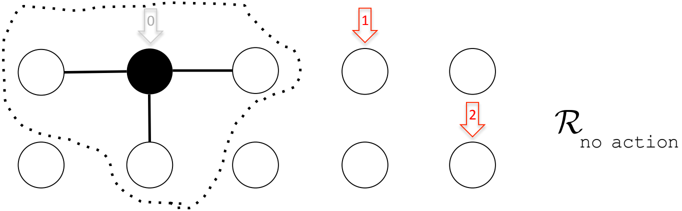


1. **Exploit locality.** Model $Q(s, a)$ using local meta-features only.
2. **Catch.** Local meta-features can't predict cumulative reward.
3. **Credit Assignment.** Isolate the contribution of $a$. Fit $Q(s, a)$ to

$$\mathcal{R}\{\texttt{taking action}\} - \mathcal{R}\{\texttt{no action}\}$$

# Outline

# List of Meta-Features

## Meta-Feature Templates

| Reason about | name | description |
|---|---|---|
| Uncertainty | cond-ent | conditional entropy |
| | unigram-ent | entropy by unigram model |
| | sp | number of times sampled |
| Staleness | nb-vary | number of neighbors changed |
| Discord | nb-discord | discord with neighbors |

# List of Meta-Features

## Meta-Feature Templates

| Reason about | name | description |
|---|---|---|
| | `cond-ent` | conditional entropy |
| Uncertainty | `unigram-ent` | entropy by unigram model |
| | `sp` | number of times sampled |
| Stalness | `nb-vary` | number of neighbors changed |
| Discord | `nb-discord` | discord with neighbors |

# Uncertainty

## Feature I. Entropy

The entropy of $q(y_j|y_{\neg j}, \mathbf{x})$.

# Uncertainty

## Feature I. Entropy

The entropy of $q(y_j | y_{\neg j}, \mathbf{x})$.

**Warning.** Computing entropy is as expensive as sampling.

# Uncertainty

### Feature I. Entropy

The entropy of $q(y_j|y_{\neg j}, \mathbf{x})$.

**Warning.** Computing entropy is as expensive as sampling.

### Principle (Use cheap meta-features)

*Complexity of meta-features* $\ll$ *Complexity of sampling*

# Uncertainty

## Feature I. Entropy

The entropy of $q(y_j|y_{\neg j}, \mathbf{x})$.

**Warning.** Computing entropy is as expensive as sampling.

## Principle (Use cheap meta-features)

*Complexity of meta-features* $\ll$ *Complexity of sampling*

## Very-lazy Evaluation

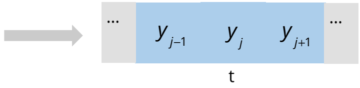

# Uncertainty

## Feature I. Entropy

The entropy of $q(y_j | y_{\neg j}, \mathbf{x})$.

**Warning.** Computing entropy is as expensive as sampling.

## Principle (Use cheap meta-features)

*Complexity of meta-features* $\ll$ *Complexity of sampling*

## Very-lazy Evaluation

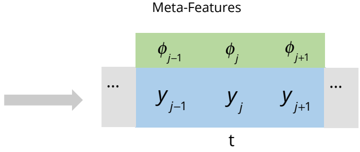

# Uncertainty

## Feature I. Entropy

The entropy of $q(y_j | y_{\neg j}, \mathbf{x})$.

> **Warning.** Computing entropy is as expensive as sampling.

## Principle (Use cheap meta-features)

*Complexity of meta-features* $\ll$ *Complexity of sampling*

## Very-lazy Evaluation

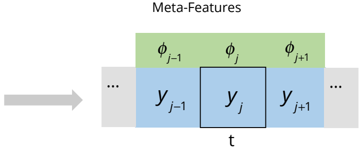

# Uncertainty

## Feature I. Entropy

The entropy of $q(y_j | y_{\neg j}, \mathbf{x})$.

**Warning.** Computing entropy is as expensive as sampling.

## Principle (Use cheap meta-features)

*Complexity of meta-features* $\ll$ *Complexity of sampling*

## Very-lazy Evaluation

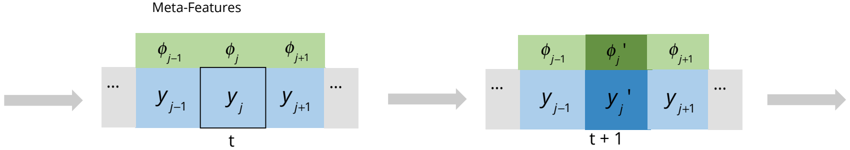

# Uncertainty

**Warning.** Entropy alone can be dangerous.

# Uncertainty

**Warning.** Entropy alone can be dangerous.

## Example

| $\mathbf{x}$ | The | Duchess | was | entertaining | $\text{sp}(\mathbf{y}_3)$ |
|---|---|---|---|---|---|

# Uncertainty

**Warning.** Entropy alone can be dangerous.

## Example

| $\mathbf{x}$ | The | Duchess | was | entertaining | $\mathrm{sp}(\mathbf{y}_3)$ |
|---|---|---|---|---|---|
| $\mathbf{y}[0]$ | Determiner | Noun | Verb | Adjective | 0 |

# Uncertainty

**Warning.** Entropy alone can be dangerous.

## Example

| $\mathbf{x}$ | The | Duchess | was | entertaining | $\text{sp}(\mathbf{y}_3)$ |
|---|---|---|---|---|---|
| $\mathbf{y}[0]$ | Determiner | Noun | Verb | Adjective | 0 |
| $\mathbf{y}[1]$ | Determiner | Noun | Verb | Verb | 1 |

# Uncertainty

**Warning.** Entropy alone can be dangerous.

## Example

| $\mathbf{x}$ | The | Duchess | was | entertaining | $\mathrm{sp}(\mathbf{y}_3)$ |
|---|---|---|---|---|---|
| $\mathbf{y}[0]$ | Determiner | Noun | Verb | Adjective | 0 |
| $\mathbf{y}[1]$ | Determiner | Noun | Verb | Verb | 1 |
| $\mathbf{y}[2]$ | Determiner | Noun | Verb | Adjective | 2 |

# Uncertainty

**Warning.** Entropy alone can be dangerous.

## Example

| $\mathbf{x}$ | The | Duchess | was | entertaining | $\mathrm{sp}(\mathbf{y}_3)$ |
|---|---|---|---|---|---|
| $\mathbf{y}[0]$ | Determiner | Noun | Verb | Adjective | 0 |
| $\mathbf{y}[1]$ | Determiner | Noun | Verb | Verb | 1 |
| $\mathbf{y}[2]$ | Determiner | Noun | Verb | Adjective | 2 |
| $\mathbf{y}[3]$ | Determiner | Noun | Verb | Verb | 3 |

# Uncertainty

**Warning.** Entropy alone can be dangerous.

## Example

| $\mathbf{x}$ | The | Duchess | was | entertaining | $\mathrm{sp}(\mathbf{y}_3)$ |
|---|---|---|---|---|---|
| $\mathbf{y}[0]$ | Determiner | Noun | Verb | Adjective | 0 |
| $\mathbf{y}[1]$ | Determiner | Noun | Verb | Verb | 1 |
| $\mathbf{y}[2]$ | Determiner | Noun | Verb | Adjective | 2 |
| $\mathbf{y}[3]$ | Determiner | Noun | Verb | Verb | 3 |
| $\mathbf{y}[4]$ | Determiner | Noun | Verb | Adjective | 4 |

# Uncertainty

**Warning.** Entropy alone can be dangerous.

## Example

| **x** | The | Duchess | was | entertaining | sp($\mathbf{y}_3$) |
|---|---|---|---|---|---|
| **y**[0] | Determiner | Noun | Verb | Adjective | 0 |
| **y**[1] | Determiner | Noun | Verb | Verb | 1 |
| **y**[2] | Determiner | Noun | Verb | Adjective | 2 |
| **y**[3] | Determiner | Noun | Verb | Verb | 3 |
| **y**[4] | Determiner | Noun | Verb | Adjective | 4 |

## Feature II. Over-exploration

number of times $y_j$ has been sampled thus far.

- simplest measure of the progress in exploration.
- usually has negative weight.

# Staleness

### 3. Change of Markov Blanket

#variables in Markov blanket that have changed.

- Identify outdated variables.

# Staleness

## 3. Change of Markov Blanket

#variables in Markov blanket that have changed.

- Identify outdated variables.

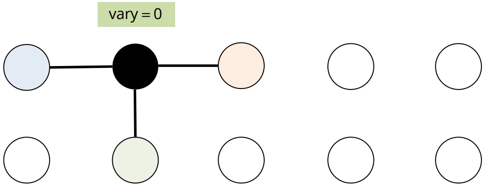

# Staleness

## 3. Change of Markov Blanket

#variables in Markov blanket that have changed.

- Identify outdated variables.

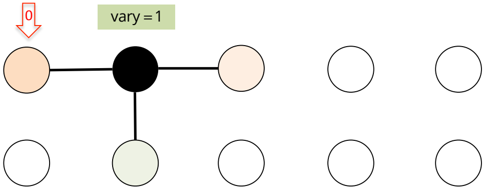

# Staleness

## 3. Change of Markov Blanket

#variables in Markov blanket that have changed.

- Identify outdated variables.

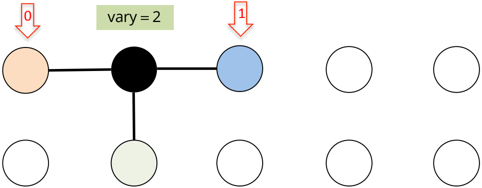

# Staleness

## 3. Change of Markov Blanket

#variables in Markov blanket that have changed.

- Identify outdated variables.

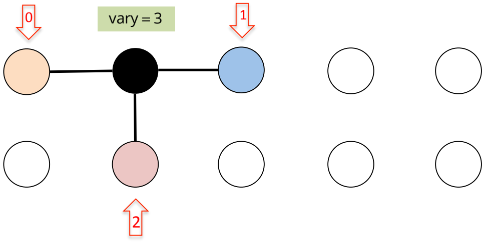

# Staleness

## 3. Change of Markov Blanket

#variables in Markov blanket that have changed.

- Identify outdated variables.

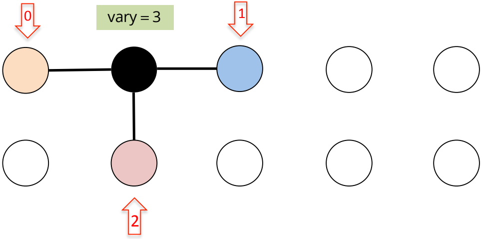


- Reason about very-lazy evaluation.

# Outline

Introduction

Reinforcement Learning

Meta-Features

**Experiments**

# Experiment Outline

## Tasks

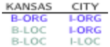


Part-of-speech tagging and name-entity recognition.

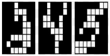

Handwriting recognition.

Color inpainting.

Scene decomposition.

# Experiment Outline

## Tasks

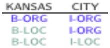


Part-of-speech tagging and name-entity recognition.

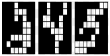

Handwriting recognition.

Color inpainting.

Scene decomposition.

## Brief Setting

1. "Steal" a graphical model with transitions.

# Experiment Outline

## Tasks

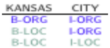


Part-of-speech tagging and name-entity recognition.

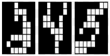

Handwriting recognition.

Color inpainting.

Scene decomposition.

## Brief Setting

1. "Steal" a graphical model with transitions.
2. Train the policy using RL on a training set.

# Experiment Outline

## Tasks

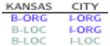


Part-of-speech tagging and name-entity recognition.

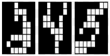

Handwriting recognition.

Color inpainting.

Scene decomposition.

## Brief Setting

1. "Steal" a graphical model with transitions.
2. Train the policy using RL on a training set.
3. Evaluate the policy on a test set.

# Speedup on NER

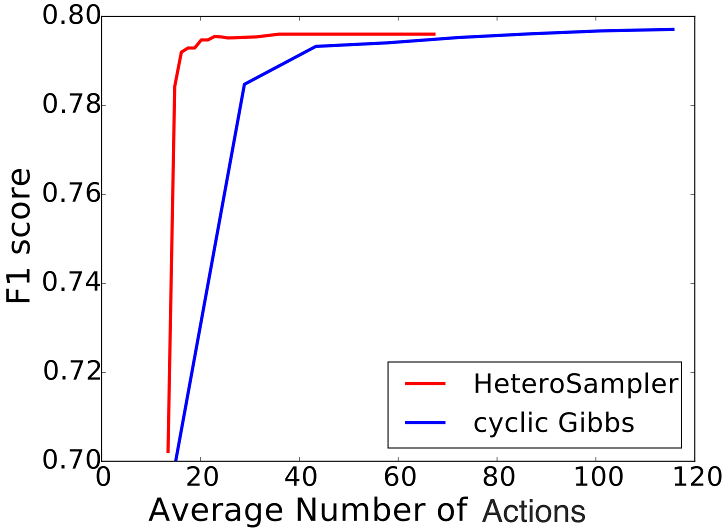

# 2-5X Speedup across tasks

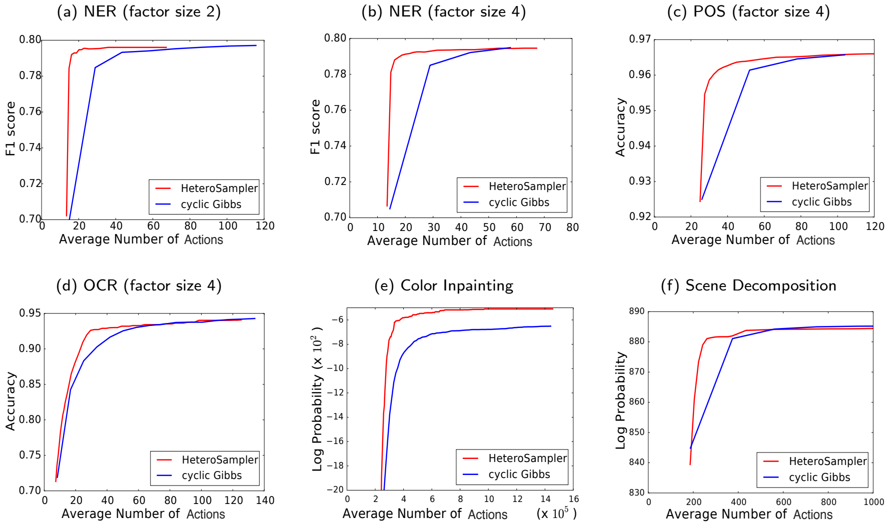

# Overhead of meta-model is minimal

(a) NER (factor size 2)

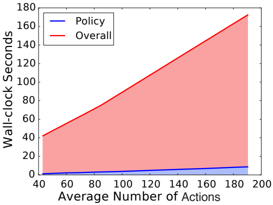


(b) NER (factor size 4)

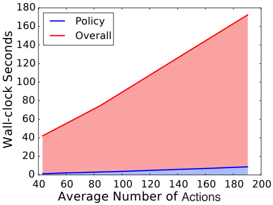


(c) POS (factor size 4)

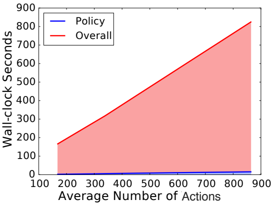


(d) OCR (factor size 4)

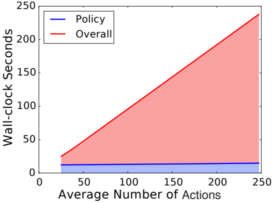


(e) Color Inpainting

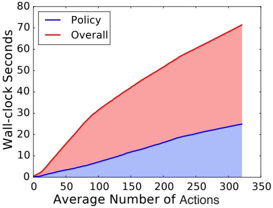


(f) Scene Decomposition

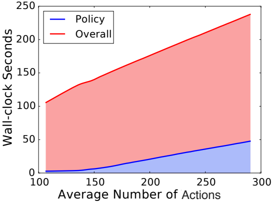

# Qualitative Results

**Visualization of computational resource allocation on NER:**

| Words | Japan | coach | Shu | Kamo | said | : | ' | ' | The | Syrian | own | goal | proved | lucky | for | us |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Truth | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 1 | I-ORG | O | B-PER | I-PER | O | O | O | O | O | B-LOC | O | O | O | O | O | O |
| 2 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 3 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |

(a) Cyclic Gibbs sampler

# Qualitative Results

**Visualization of computational resource allocation on NER:**

| Words | Japan | coach | Shu | Kamo | said | : | ' | ' | The | Syrian | own | goal | proved | lucky | for | us |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Truth | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 1 | I-ORG | O | B-PER | I-PER | O | O | O | O | O | B-LOC | O | O | O | O | O | O |
| 2 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 3 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |

(a) Cyclic Gibbs sampler

| Words | Japan | coach | Shu | Kamo | said | : | ' | ' | The | Syrian | own | goal | proved | lucky | for | us |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Truth | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 1 | I-ORG | O | B-PER | I-PER | O | O | O | O | O | B-LOC | O | O | O | O | O | O |
| 2 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 3 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |

(b) `HeteroSampler`

# Related Work

## Heterogenous Inference

- Schedules asynchronous belief propagation [Elidan et al., 2006]
- Query-aware belief propagation [Chechetka et al., 2010]

# Related Work

## Heterogenous Inference

- Schedules asynchronous belief propagation [Elidan et al., 2006]
- Query-aware belief propagation [Chechetka et al., 2010]

## Use simple models to help computation

- Coarse-to-fine object detection [Viola and Jones, 2001]
- Structured prediction cascades [Weiss and Taskar, 2010]

# Related Work

## Heterogenous Inference

- Schedules asynchronous belief propagation [Elidan et al., 2006]
- Query-aware belief propagation [Chechetka et al., 2010]

## Use simple models to help computation

- Coarse-to-fine object detection [Viola and Jones, 2001]
- Structured prediction cascades [Weiss and Taskar, 2010]

## RL for structured prediction

- SEARN [Daume et al., 2009]
- DAGGER [Ross et al., 2011a]
- RL for dependency parsing [Goldberg and Nivre, 2013]

# Contribution

## Take-way Message: Sample Heterogeneously!

| Words | Japan | coach | Shu | Kamo | said | : | ' | ' | The | Syrian | own | goal | proved | lucky | for | us |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Truth | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 1 | I-ORG | O | B-PER | I-PER | O | O | O | O | O | B-LOC | O | O | O | O | O | O |
| 2 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 3 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |

# Contribution

## Take-way Message: Sample Heterogeneously!

| Words | Japan | coach | Shu | Kamo | said | : | ' | ' | The | Syrian | own | goal | proved | lucky | for | us |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Truth | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 1 | I-ORG | O | B-PER | I-PER | O | O | O | O | O | B-LOC | O | O | O | O | O | O |
| 2 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 3 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |

## Important Ideas

1. It is effective to reason about **uncertainty, staleness**, etc.
2. It is possible to predict where to sample using cheap meta-features (**Very-lazy evaluation**).

# Contribution

## Take-way Message: Sample Heterogeneously!

| Words | Japan | coach | Shu | Kamo | said | : | ' | ' | The | Syrian | own | goal | proved | lucky | for | us |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Truth | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 1 | I-ORG | O | B-PER | I-PER | O | O | O | O | O | B-LOC | O | O | O | O | O | O |
| 2 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |
| 3 | B-LOC | O | B-PER | I-PER | O | O | O | O | O | B-MISC | O | O | O | O | O | O |

## Important Ideas

1. It is effective to reason about **uncertainty, staleness**, etc.
2. It is possible to predict where to sample using cheap meta-features (**Very-lazy evaluation**).

## Thanks!