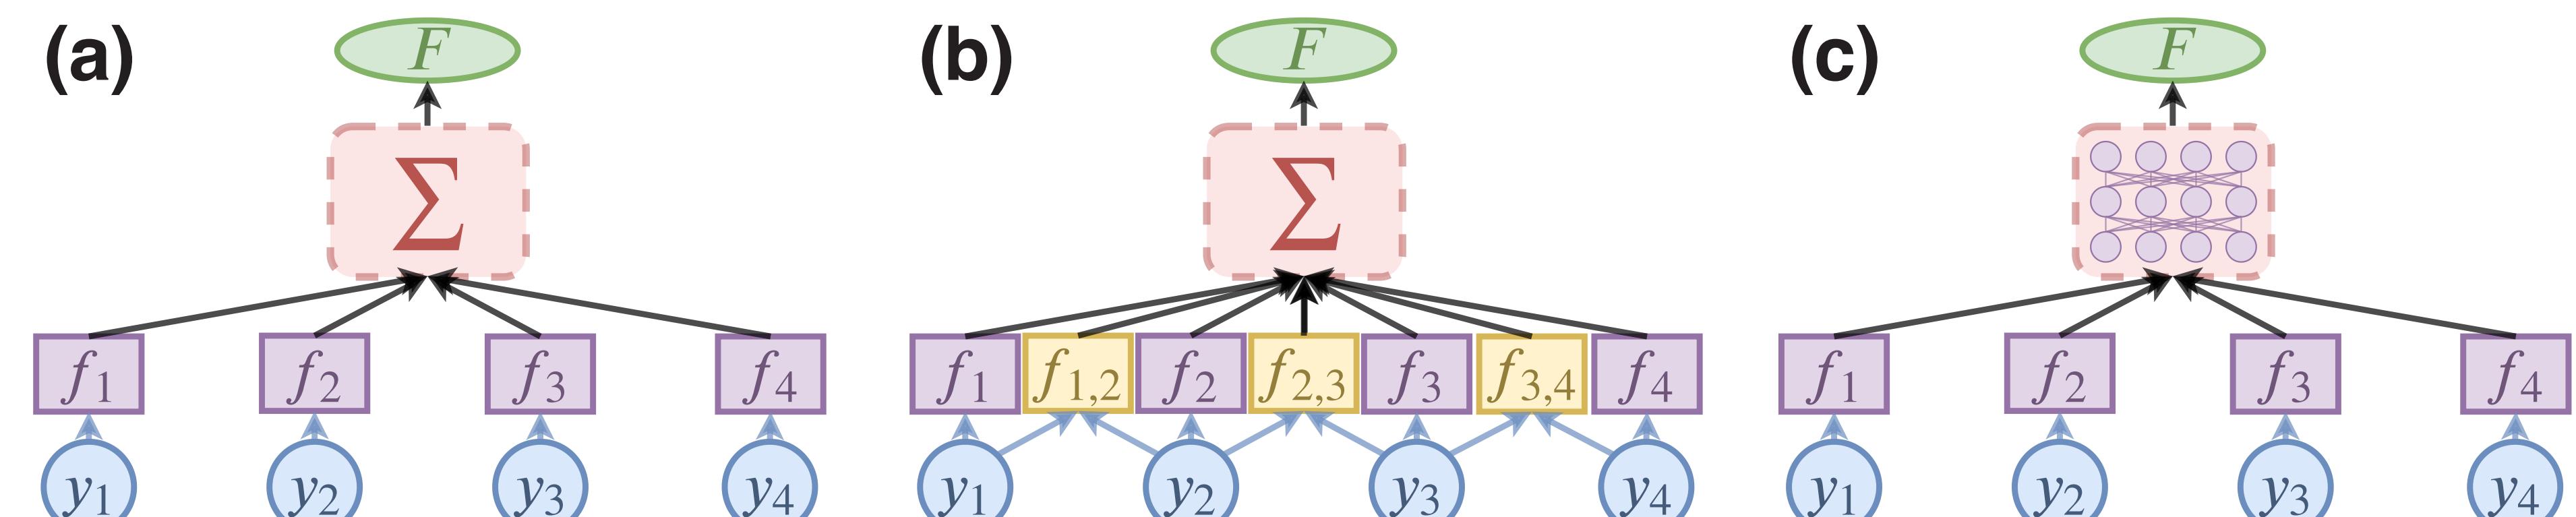


Motivation & Related Work

Objective:

Develop a model that allows known, local structure to be specified while also learning unknown global structure.

Approaches to Structured Prediction:

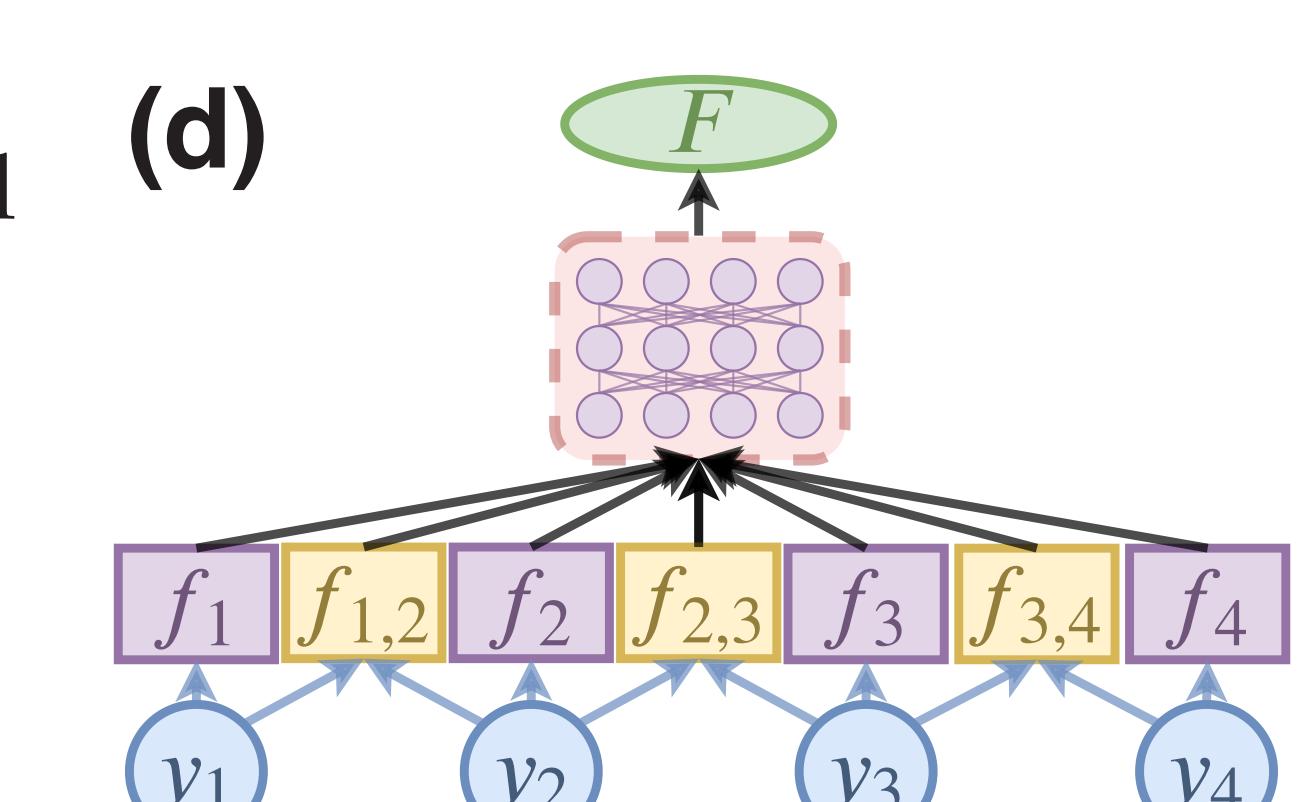


Unstructured Models (a):

Classical Structured Models (b):

Structured Prediction Energy Networks [1] (c):

Ours: Graph Structured Prediction Energy Networks (d):



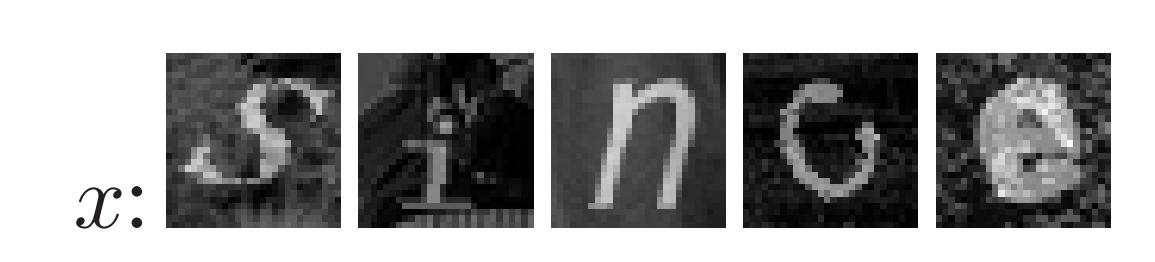
Model Formulation

Notation:

Score Functions: assign score to input-output pair (x, y)

Example:

\mathcal{X} : Set of images with letters
 \mathcal{Y} : ('A', 'B', ..., 'Y', 'Z')⁵



x : S i n G e

y : 'S' 'T' 'N' 'C' 'E'

Graph used by Unary and SPEN:
 1 2 3 4 5

$$\mathcal{R} = \{1, 2, 3, 4, 5\}$$

Graph used by Struct, NLStruct, and GSPEN:
 (1, 2) (2, 3) (3, 4) (4, 5)

$$\mathcal{R} = \{1, 2, 3, 4, 5, (1, 2), (2, 3), (3, 4)\}$$

$$\mathcal{R} = \{(1, 2), (2, 3), (3, 4)\}$$

GSPEN Inference and Learning

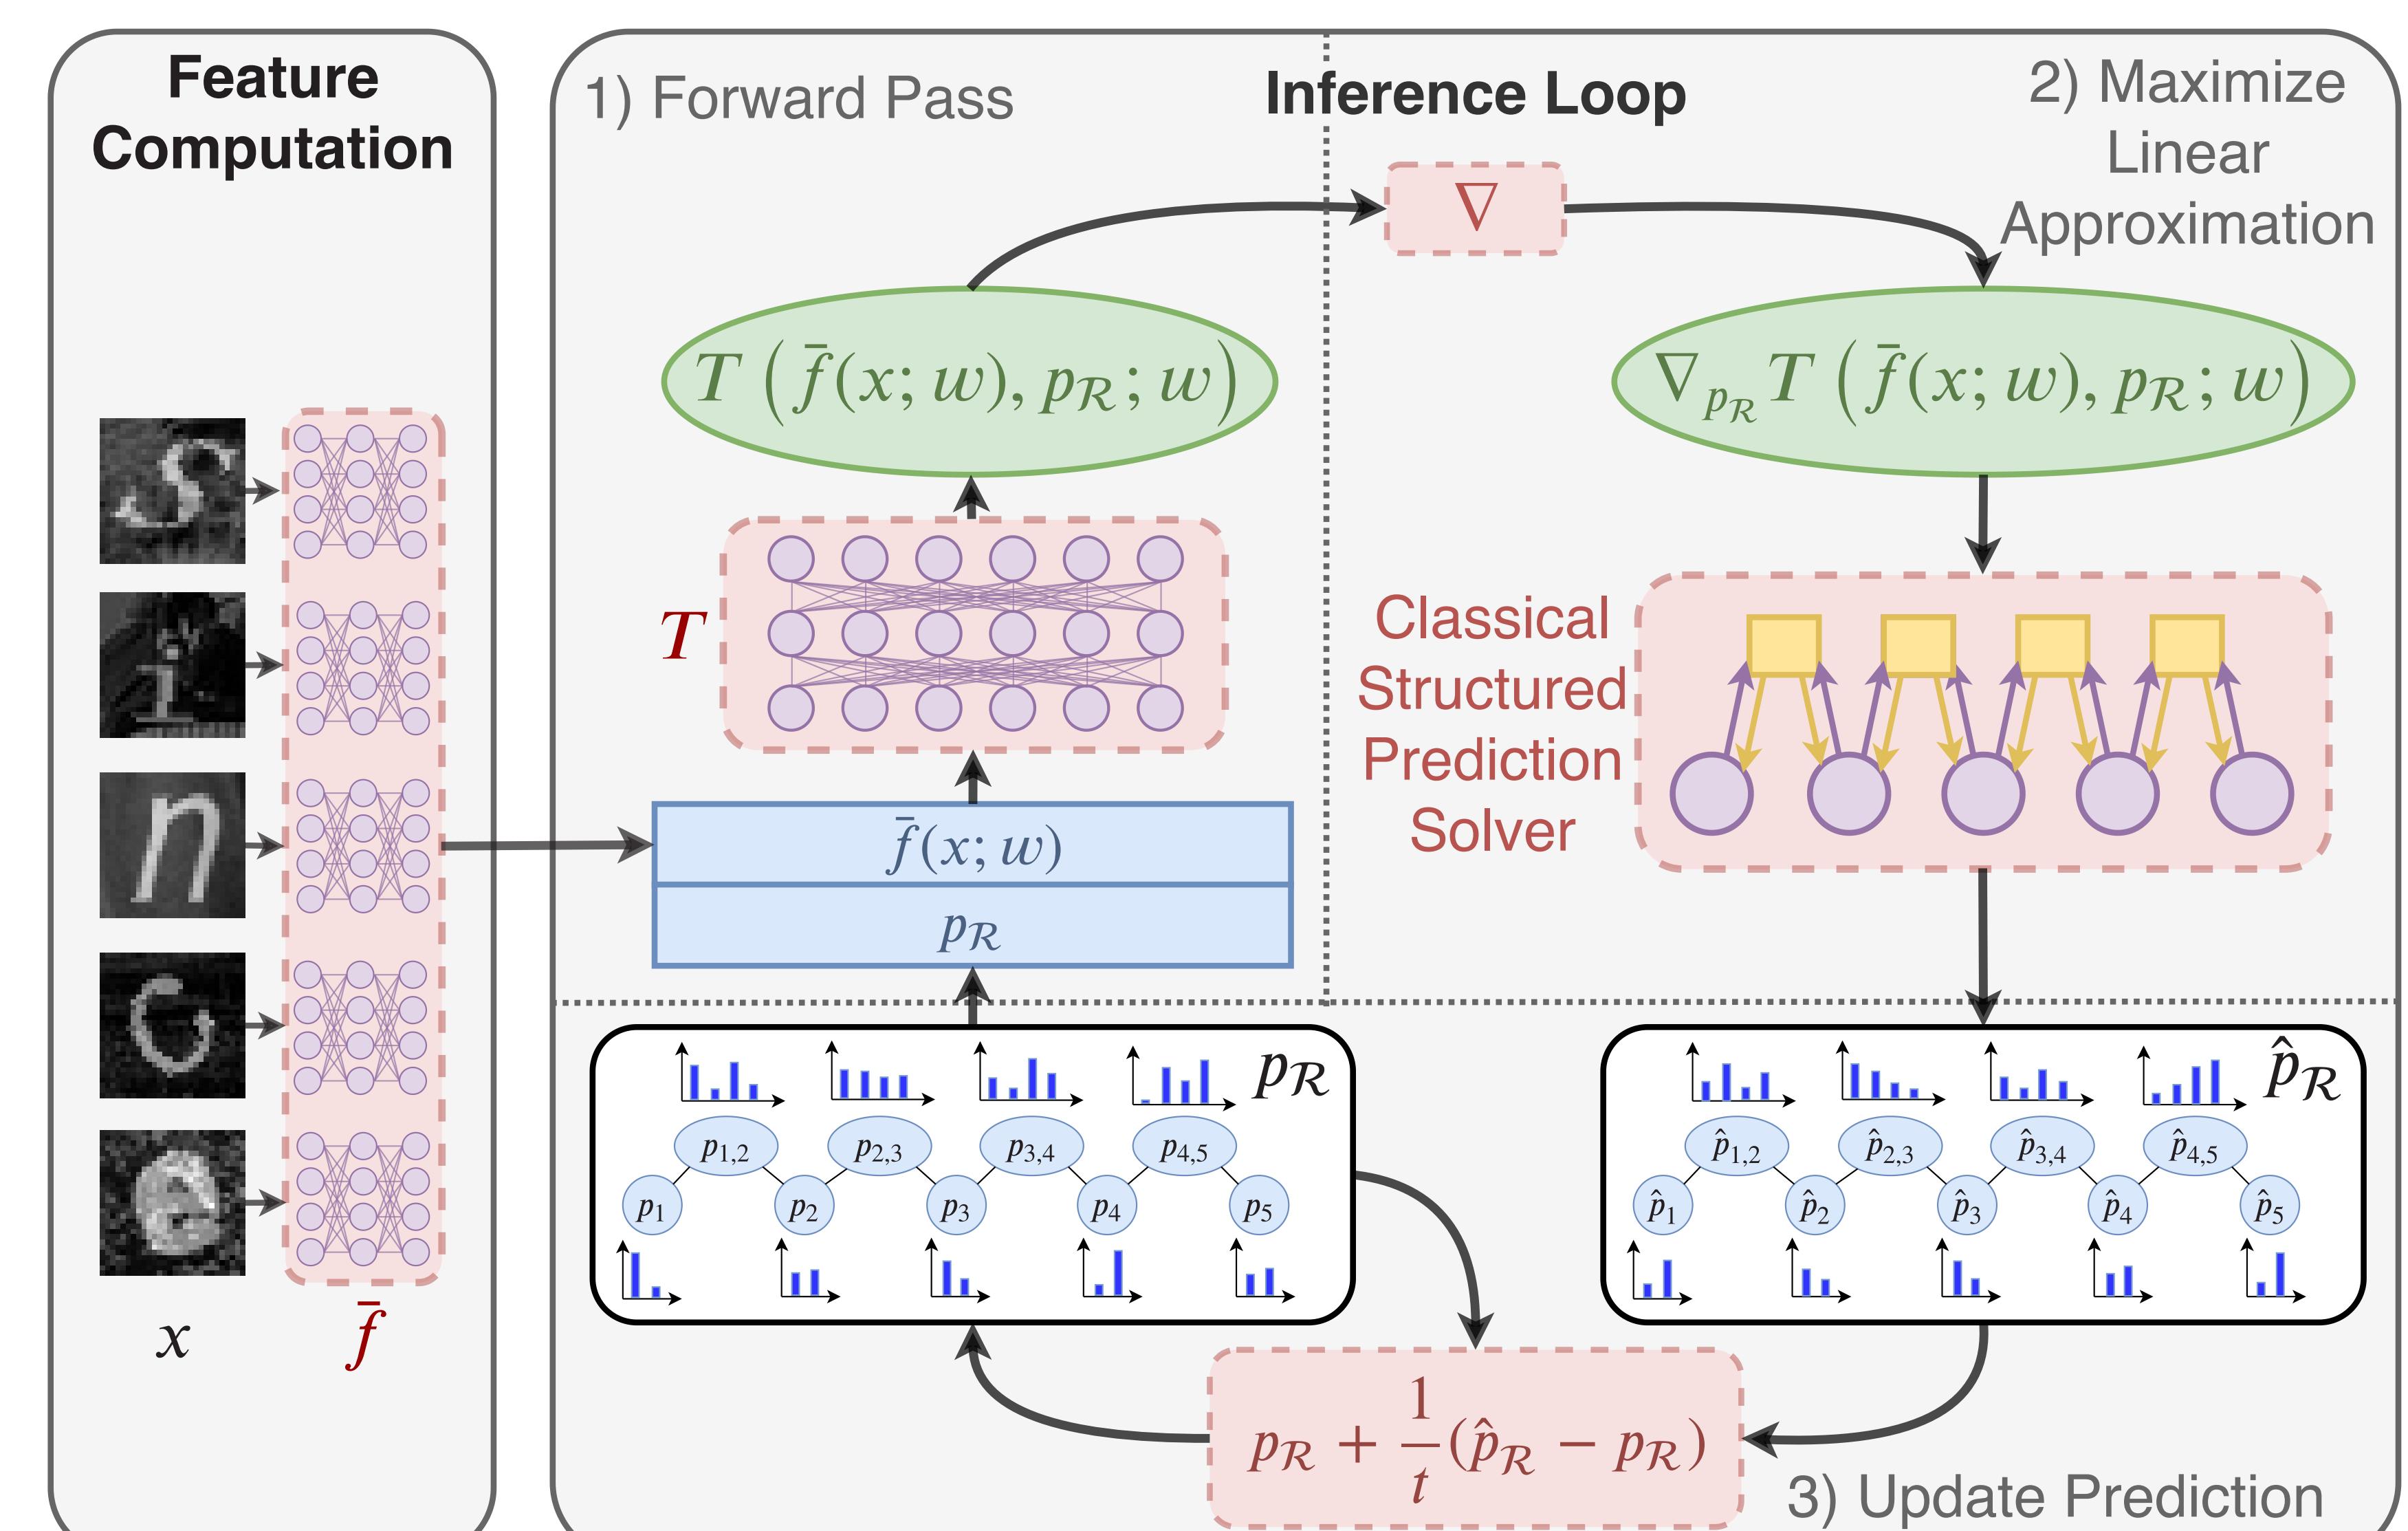
Inference objective: maximize $F(x, p_{\mathcal{R}}; w)$ w.r.t. $p_{\mathcal{R}}$ such that $p_{\mathcal{R}} \in \mathcal{M}$ (marginal polytope)

- Solving inference requires using a constrained inference algorithm
- Because we have a non-linear objective subject to linear constraints, we can use Frank-Wolfe
- Inner maximization problem reduces to classical structured inference
- We also experiment with entropic mirror descent

Algorithm 1 Frank-Wolfe Inference for GSPEN

```

1: Input: Initial set of predictions  $p_{\mathcal{R}}$ ;
   Input  $x$ ; Factor graph  $\mathcal{R}$ 
2: for  $t = 1 \dots T$  do
3:    $g \Leftarrow \nabla_{p_{\mathcal{R}}} F(x, p_{\mathcal{R}}; w)$ 
4:    $\hat{p}_{\mathcal{R}} \Leftarrow \max_{\hat{p}_{\mathcal{R}} \in \mathcal{M}_{\mathcal{R}}, r \in \mathcal{R}, y_r \in \mathcal{Y}_r} \hat{p}_r(y_r) g_r(y_r)$ 
5:    $p_{\mathcal{R}} \Leftarrow p_{\mathcal{R}} + \frac{1}{t} (\hat{p}_{\mathcal{R}} - p_{\mathcal{R}})$ 
6: end for
7: Return:  $p_{\mathcal{R}}$ 
```



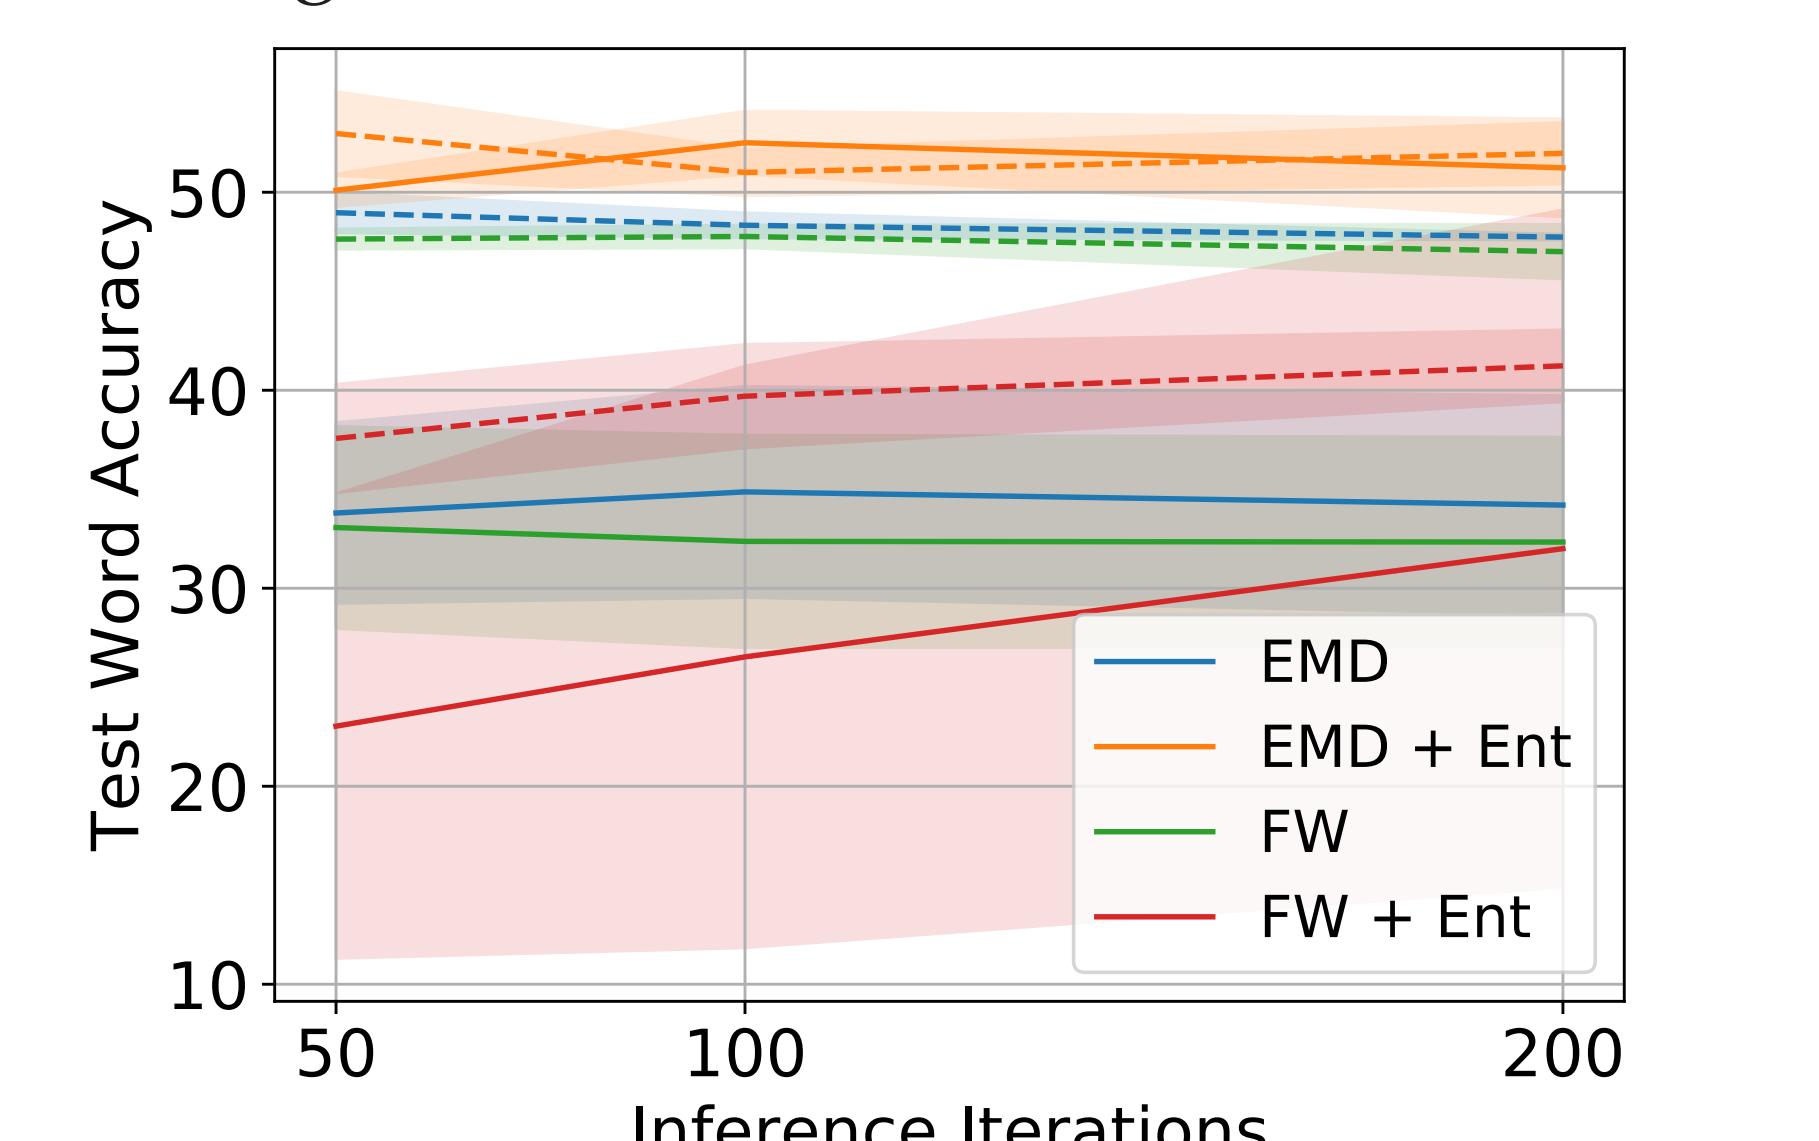
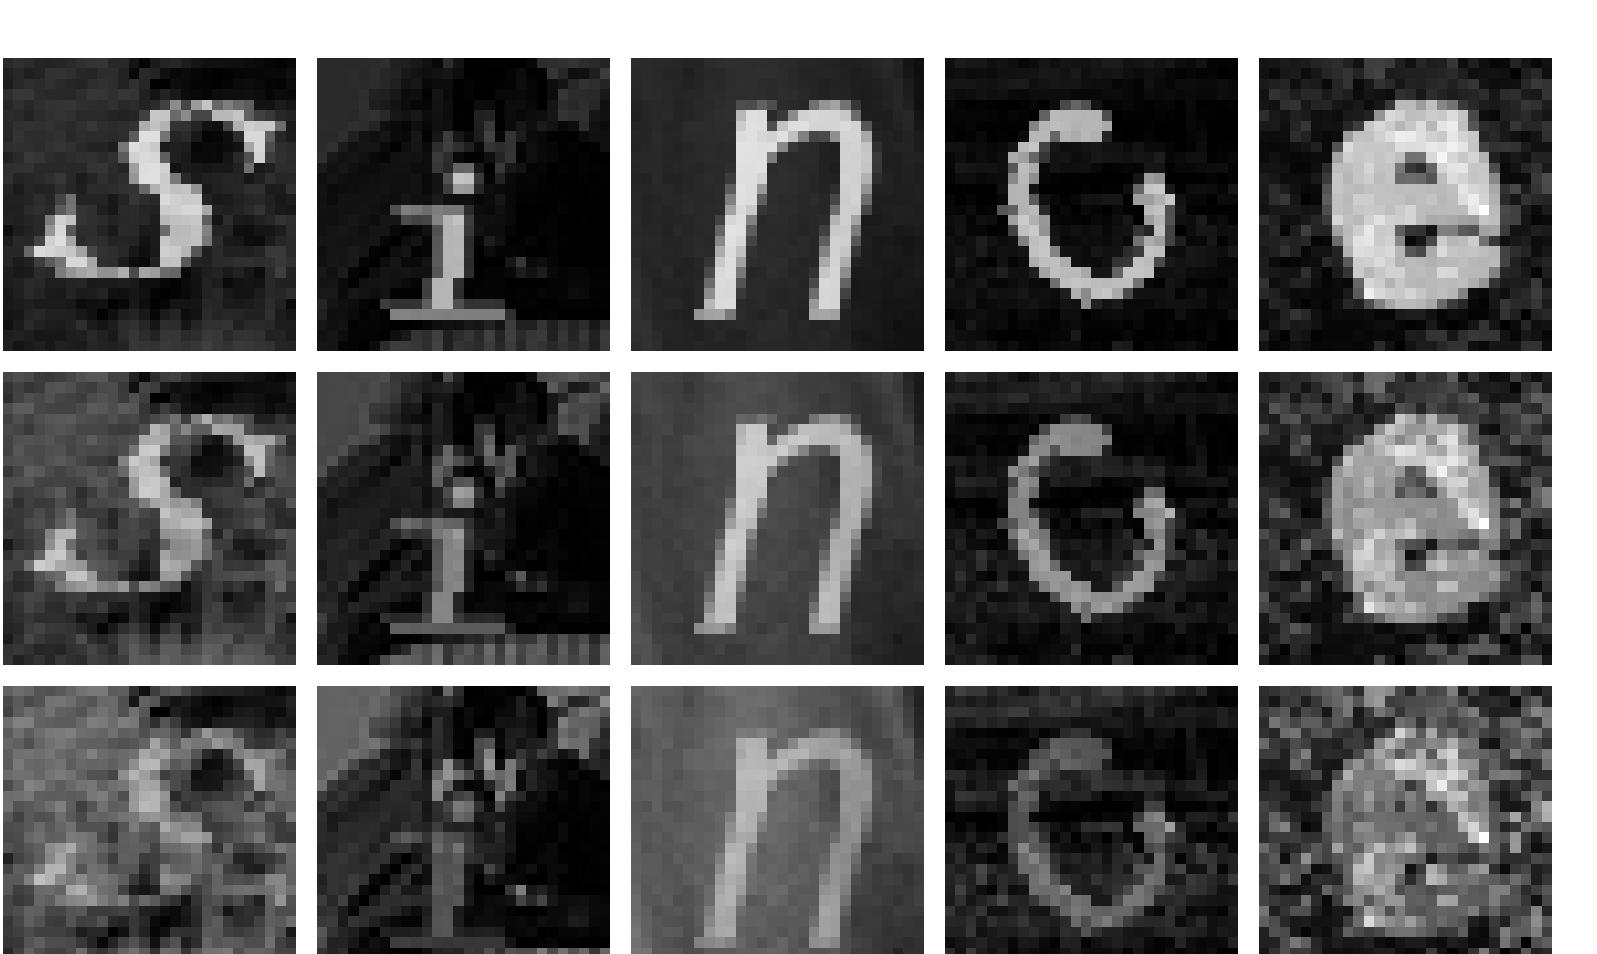
Learning Objective: SSVM objective with loss-augmented inference

$$\min_w \sum_{(x^{(i)}, p_R^{(i)})} \left[\max_{\hat{p}_{\mathcal{R}} \in \mathcal{M}} \left\{ T(\bar{f}(x^{(i)}; w), \hat{p}_{\mathcal{R}}; w) + L(\hat{p}_{\mathcal{R}}, p_R^{(i)}) \right\} - T(\bar{f}(x^{(i)}; w), p_R^{(i)}; w) \right]_+$$

Experiments

Optical Character Recognition

- Characters rendered on high-noise background patches with different interpolation factors
- Examined different inference procedures during learning (below)



Experiments (Continued)

OCR (Continued)

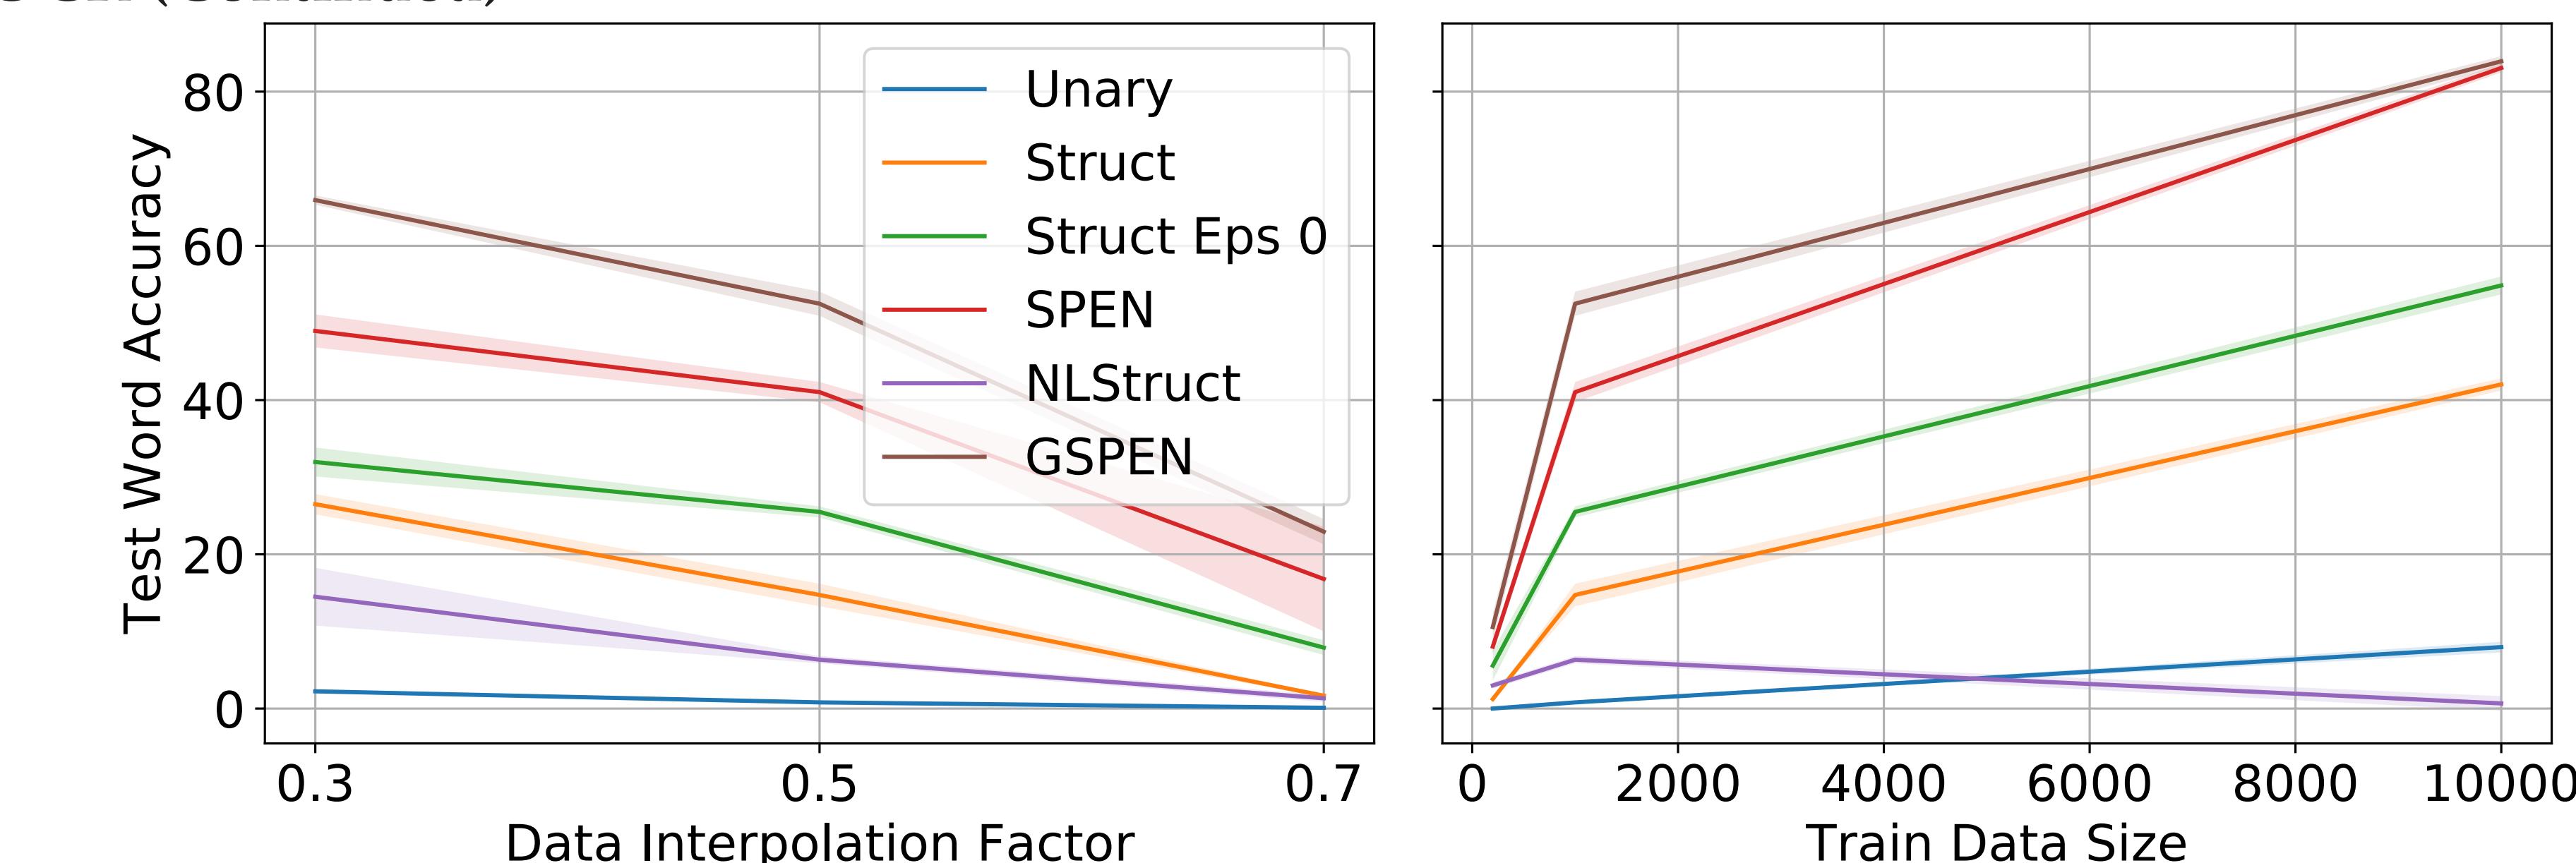
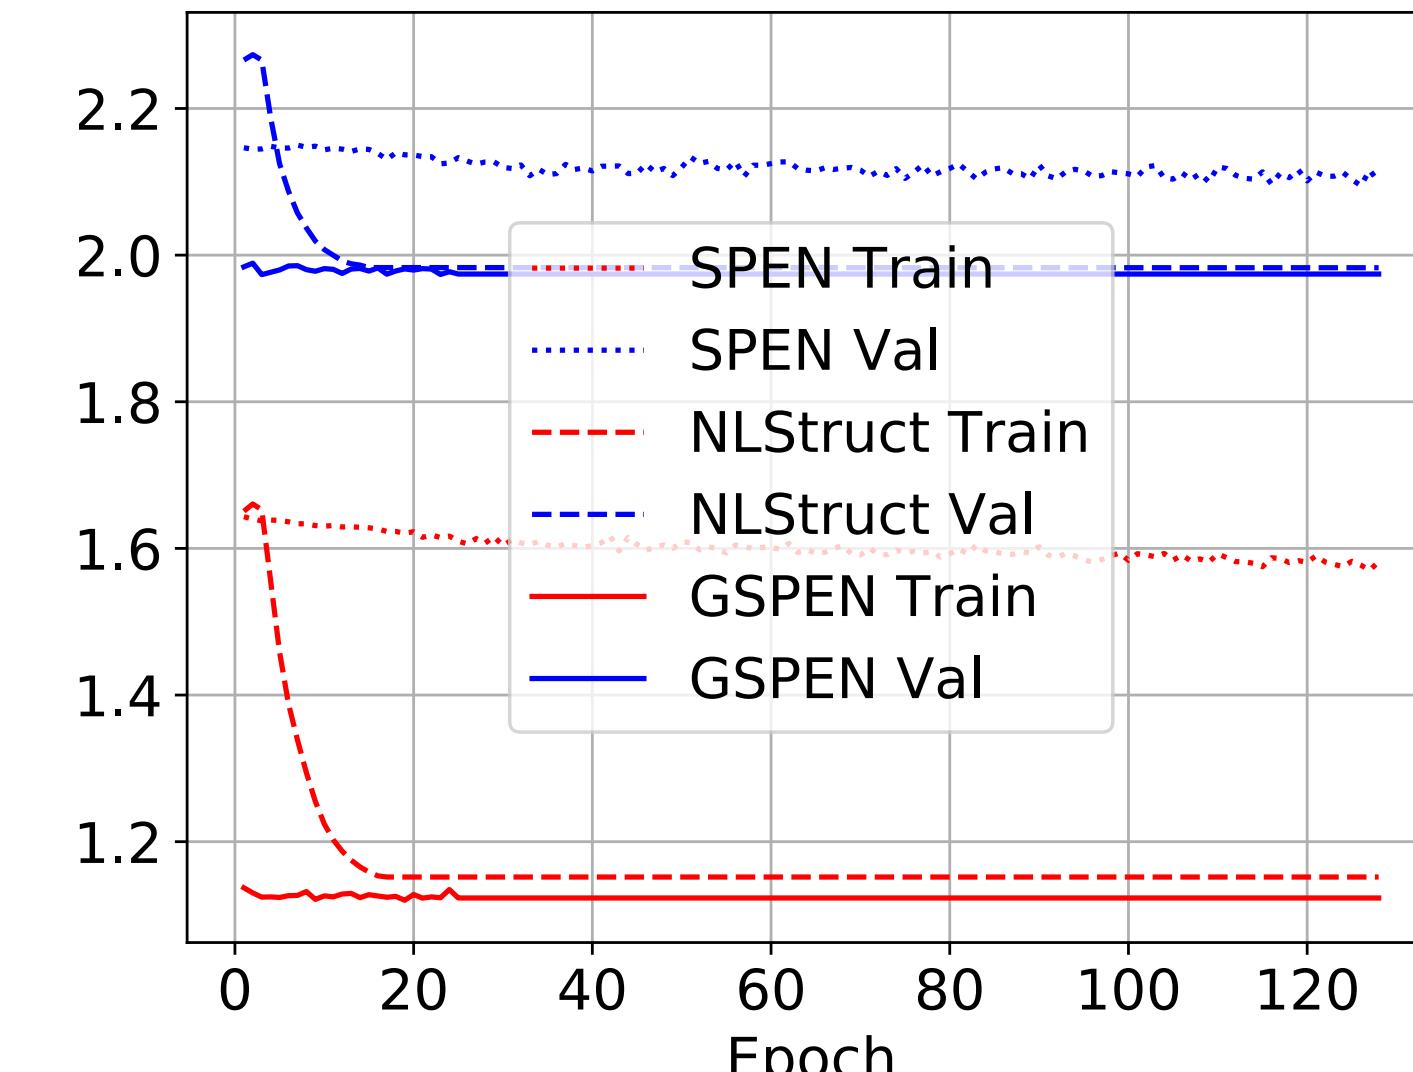
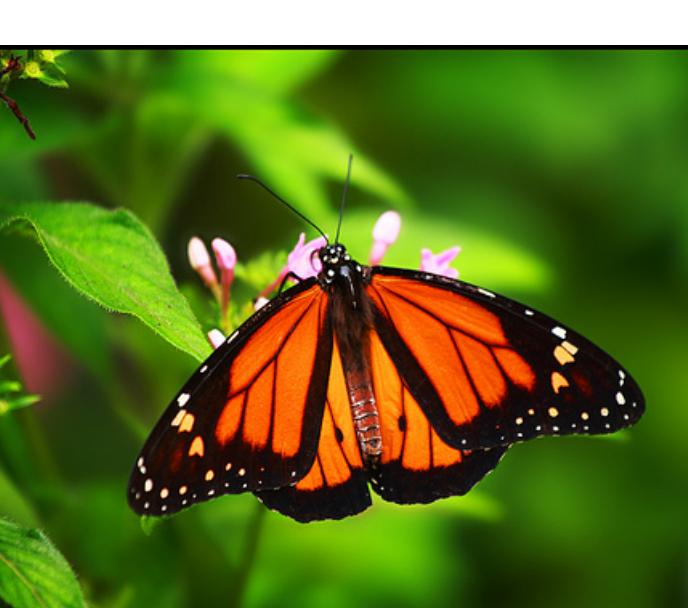


Image Tagging

- Problem: label images with subset of 24 possible tags
- Reported metric: Hamming Loss
- GSPEN trained using pre-trained Struct model from [5]
- Test results:
 - SPEN: 2.158
 - NLStruct: 2.037
 - GSPEN: 2.029



Multilabel Classification

- Problem: given binary input feature vector, assign it to a subset of labels
- Two datasets: Bibtex (159 labels) and Bookmarks (208 labels)
- Reported metric: F1
- Top two rows taken from [3]

	Bibtex	Bookmarks		
	Val.	Test	Val.	Test
SPEN	–	42.2	–	34.4
DVN	–	44.7	–	37.1
Unary	43.3	44.1	38.4	37.4
Struct	45.8	46.1	39.7	38.9
SPEN	46.6	46.5	40.2	39.2
GSPEN	47.5	48.6	41.2	40.7

Named Entity Recognition

B-ORG O B-PER I-PER O O

Celtic 's Jackie McNamara , who

- Problem: label each word in a sentence with tags indicating whether they are part of a named entity
- Metric: F1
- Trained GSPEN models initialized from two different Struct models

	Avg. Val.	Avg. Test
Struct [2]	94.88 ± 0.18	91.37 ± 0.04
+ GSPEN	94.97 ± 0.16	91.51 ± 0.17
Struct [4]	95.88 ± 0.10	92.79 ± 0.08
+ GSPEN	95.96 ± 0.08	92.69 ± 0.17

References

- D. Belanger, A. McCallum. Structured Prediction Energy Networks. In ICML 2016
- X. Ma, E. Hovy. End-to-end Sequence Labeling via Bi-directional LSTM-CNNs-CRF. In ACL 2016
- M. Gygli, M. Norouzi, A. Angelova. Deep Value Networks Learn to Evaluate and Iteratively Refine Structured Outputs. In ICML 2017
- A. Akbik, D. Blythe, R. Vollgraf. Contextual String Embeddings for Sequence Labeling. In COLING 2018
- C. Gruber, O. Meshi, A. Schwing. Deep Structured Prediction with Nonlinear Output Transformations. In NeurIPS 2018