

***Hitachi* at MRP 2019: Unified Encoder-to-Biaffine Network for Cross-Framework Meaning Representation Parsing**

Yuta Koreeda*, Gaku Morio*, Terufumi Morishita*, Hiroaki Ozaki*, and Kohsuke Yanai
Hitachi, Ltd.

Research & Development Group
Kokubunji, Tokyo, Japan

{yuta.koreeda.pb, gaku.morio.vn, terufumi.morishita.wp,
hiroaki.ozaki.yu, kohsuke.yanai.cs}@hitachi.com

Abstract

This paper describes the proposed system of the *Hitachi* team for the Cross-Framework Meaning Representation Parsing (MRP 2019) shared task. In this shared task, the participating systems were asked to predict nodes, edges and their attributes for five frameworks, each with different order of “abstraction” from input tokens. We proposed a unified encoder-to-biaffine network for all five frameworks, which effectively incorporates a shared encoder to extract rich input features, decoder networks to generate anchorless nodes in UCCA and AMR, and biaffine networks to predict edges. Our system was ranked fifth with the macro-averaged MRP F1 score of 0.7604, and outperformed the baseline unified transition-based MRP. Furthermore, post-evaluation experiments showed that we can boost the performance of the proposed system by incorporating multi-task learning, whereas the baseline could not. These imply efficacy of incorporating the biaffine network to the shared architecture for MRP and that learning heterogeneous meaning representations at once can boost the system performance.

1 Introduction

This paper describes the proposed system of the *Hitachi* team for the CoNLL 2019 Cross-Framework Meaning Representation Parsing (MRP 2019) shared task. The goal of the task was to design a system that predicts sentence-level graph-based meaning representations in five frameworks, each with its specific linguistic assumptions. The task was formulated as prediction of nodes, edges and their attributes from an input sentence (see [Oepen et al. \(2019\)](#) for details). The target frameworks were (1) DELPH-IN MRS Bi-Lexical Dependencies (DM; [Flickinger, 2000](#);

[Ivanova et al., 2012](#)), (2) Prague Semantic Dependencies (PSD; [Hajič et al., 2012](#); [Miyao et al., 2014](#)), (3) Elementary Dependency Structures (EDS; [Oepen and Lønning, 2006](#)), (4) Universal Conceptual Cognitive Annotation framework (UCCA; [Abend and Rappoport, 2013](#); [Hershcovich et al., 2017](#)), and (5) Abstract Meaning Representation (AMR; [Banarescu et al., 2013](#)).

In this work, we propose to unify graph predictions in all frameworks with a single encoder-to-biaffine network. This objective was derived from our expectation that it would be advantageous if a single neural network can deal with all the frameworks, because it allows all frameworks to benefit from architectural enhancements and it opens up possibility to perform multi-task learning to boost overall system performance. We argue that it is non-trivial to formulate different kinds of graph predictions as a single machine learning problem, since each framework has different order of “abstraction” from input tokens. Moreover, such formulation has hardly been explored, with few exceptions including unified transition-based MRP ([Hershcovich et al., 2018](#)), to which we empirically show the superiority of our system (Section 9). We also present a multi-task variant of such system, which did not make it to the task deadline.

Our *non*-multi-task system obtained the fifth position in the formal evaluation. We also evaluated the multi-task setup after the formal run, showing multi-task learning can yield an improvement in the performance. This result implies learning heterogeneous meaning representations at once can boost the system performance.

2 Overview of the Proposed System

The key challenge in unifying graph predictions with a single encoder-to-biaffine network lays in

* Contributed equally.

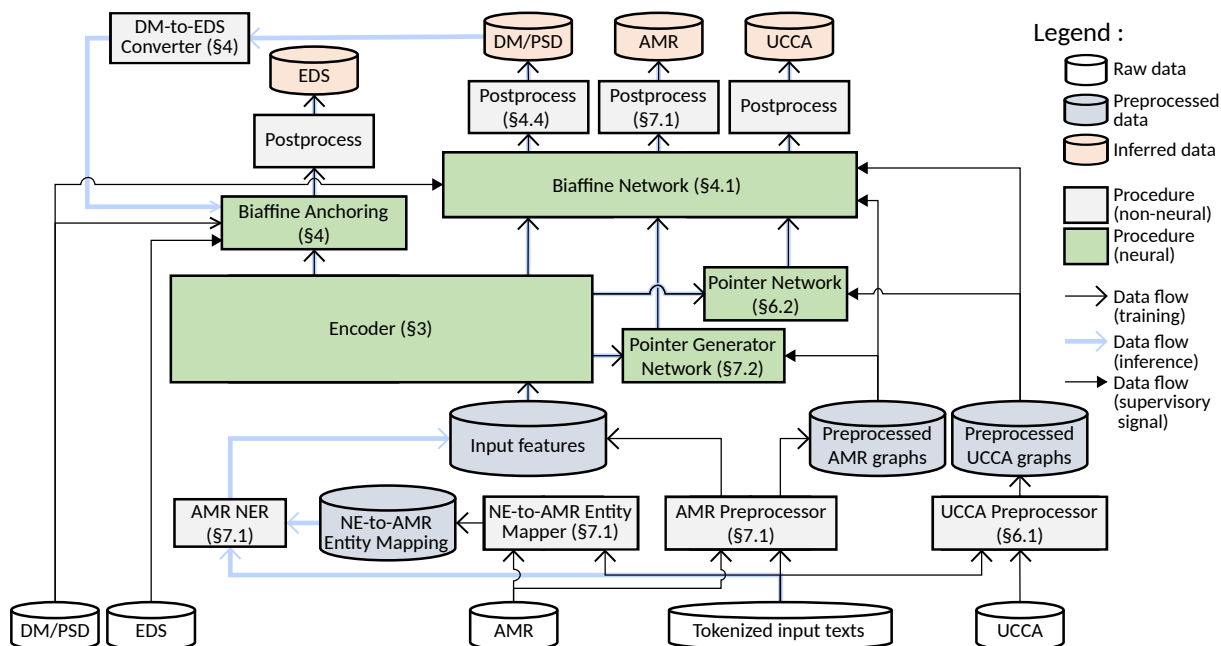


Figure 1: The overview of the proposed unified encoder-to-biaffine network for cross-framework meaning representation parsing.

complementation of nodes, because the biaffine network can narrow down the node candidates but cannot generate new ones. Our strategy is that we start from input tokens, generate missing nodes (nodes that do not have anchors to the input tokens) and finally predict edges with the biaffine network (Figure 1). More concretely, the shared encoder (Section 3.2) fuses together rich input features for each token including features extracted from pretrained language models, which are then fed to bidirectional long short-term memories (biLSTMs; Hochreiter and Schmidhuber, 1997; Schuster and Paliwal, 1997) to obtain task-independent contextualized token representations. The contextualized representations are fed to biaffine networks (Dozat and Manning, 2018) to predict graphs for each framework along with the following framework-specific procedures:

DM and PSD Contextualized representations are fed to biaffine network to predict edges and their labels. They are also used to predict the node property *frame* (Section 4).

EDS The predicted DM graphs are converted to nodes and edges of EDS graphs. Contextualized representations are used to predict node anchors (Section 5).

UCCA Nodes in training data are serialized and aligned with input tokens. Contextualized representations are fed to a pointer network to gen-

erate non-terminal nodes, and to a biaffine network to predict edges and labels (Section 6).

AMR Contextualized representations are fed to pointer-generator network to generate nodes. Hidden states of the network are fed to a biaffine network to predict edges and their labels (Section 7).

All models are trained end-to-end using mini-batch stochastic gradient descent with backpropagation (see Appendix A.1 for the details).

3 Shared Encoder

3.1 Feature Extraction

Following work by Dozat and Manning (2018) and Zhang et al. (2019), we propose to incorporate multiple types of token representations to provide rich input features for each token. Specifically, the proposed system combines surface, lemma, part-of-speech (POS) tags, named entity label, GloVe (Pennington et al., 2014) embedding, ELMo (Peters et al., 2018) embedding and BERT (Devlin et al., 2019) embedding as input features. The following descriptions explain how we acquire each input representations:

Surface and lemma We use the lower-cased node labels and the lemma properties from the companion data, respectively. Surfaces and lemmas that appear less than four times are replaced by a special <UNK> token. We also map

numerical expressions¹ to a special <NUM> token.

POS tags We use Universal POS tags and English specific POS tags from node properties `upos` and `xpos` in the companion data, respectively.

Named entity label Named entity (NE) recognition is applied to the input text (see Section 7.1).

GloVe We use 300-dimensional GloVe (Pennington et al., 2014) pretrained on Common Crawl² which are kept fixed during the training. Surfaces that do not appear in the pretrained GloVe are mapped to a special <UNK> token which is set to a vector whose values are randomly drawn from normal distribution with standard deviation of $1/\sqrt{\text{dimension of a GloVe vector}}$.

ELMo We use the pretrained “original” ELMo³. Following Peters et al. (2018), we “mix” different layers of ELMo for each token;

$$\tilde{s}_j = \text{softmax}(s_j) = \frac{\exp(s_j)}{\sum_k \exp(s_k)},$$

$$\mathbf{h}^{ELMo} = \sum_{j=0}^{N^{ELMo}-1} \tilde{s}_j \mathbf{h}_j^{ELMo},$$

where \mathbf{h}_j^{ELMo} ($0 \leq j < N^{ELMo}$) is the hidden state of the j -th layer of ELMo, \mathbf{h}_0^{ELMo} is the features from character-level CNN of ELMo, and s_j are trainable parameters. \mathbf{h}_j^{ELMo} are fixed in the training by truncating backpropagation to \mathbf{h}_j^{ELMo} .

BERT We use the pretrained BERT-Large, Uncased (Original)⁴. Since BERT takes subword units as input, a BERT embedding for a token is generated as the average of its subword BERT embeddings as in Zhang et al. (2019).

The surface, lemma, POS tags and NE label of a token are each embedded as a vector. The vectors are randomly initialized and updated during training. To allow prediction of the top nodes for DM, PSD and UCCA, a special <ROOT> token

¹Surfaces or lemmas that can successfully be converted to numerics with `float` operation on Python 3.6

²<http://nlp.stanford.edu/data/glove.840B.300d.zip>

³https://s3-us-west-2.amazonaws.com/allennlp/models/elmo/2x4096_512_2048cnn_2xhighway/elmo_2x4096_512_2048cnn_2xhighway_weights.hdf5 and [elmo_2x4096_512_2048cnn_2xhighway_options.json](https://s3-us-west-2.amazonaws.com/allennlp/models/elmo/2x4096_512_2048cnn_2xhighway_options.json).

⁴https://s3.amazonaws.com/models.huggingface.co/bert/bert-large-uncased-pytorch_model.bin, which is converted from the whitelisted BERT model in <https://github.com/google-research/bert>

is prepended to each input sequence. For GloVe, ELMo and BERT, the <ROOT> is also embedded in the similar manner as other tokens with <ROOT> as the surface for the token. A multi-layered perceptron (MLP) is applied to each of GloVe, ELMo and BERT embeddings.

To prevent the model from overrelying only on certain types of features, we randomly drop a group of features, where the groups are (i) lemma, (ii) POS tags and (iii) the rest. All features in the same group are randomly dropped simultaneously but independently from other groups.

All seven features are then concatenated to form input token representation \mathbf{h}_i^0 (where $0 \leq i < L_{in}$ is the index of the token).

3.2 Obtaining Contextualized Token Representation

The input token representations \mathbf{h}_i^0 are fed to the multi-layered biLSTM with N layers to obtain the contextualized token representations.

$$\begin{aligned} \vec{\mathbf{h}}_i^l &= \overrightarrow{\text{LSTM}}(\mathbf{h}_i^{l-1}, \vec{\mathbf{h}}_{i-1}^l, \vec{\mathbf{c}}_{i-1}^l), \\ \overleftarrow{\mathbf{h}}_i^l &= \overleftarrow{\text{LSTM}}(\mathbf{h}_i^{l-1}, \overleftarrow{\mathbf{h}}_{i+1}^l, \overleftarrow{\mathbf{c}}_{i+1}^l), \\ \mathbf{h}_i^l &= \left[\vec{\mathbf{h}}_i^l; \overleftarrow{\mathbf{h}}_i^l \right], \end{aligned}$$

where \mathbf{h}_i^l and \mathbf{c}_i^l ($0 < l \leq N$) are the hidden states and the cell states of the l -th layer LSTM for i -th token.

4 DM and PSD-specific Procedures

4.1 Biaffine Classifier

DM and PSD are Flavor (0) frameworks whose nodes have one-to-one correspondence to tokens. We utilize biaffine networks to filter nodes, and to predict edges, edge labels and node attributes. For each framework $\text{fw} \in \{\text{dm}, \text{psd}\}$, probability that there exists an edge (i, j) from the i -th node to the j -th node $y_{\text{fw},i,j}^{\text{edge}}$ is calculated for all pairs of nodes ($0 \leq i, j < L_{in}$).

$$\begin{aligned} \mathbf{h}_{\text{fw},i}^{\text{edge.from}} &= \text{MLP}^{\text{edge.from}}(\mathbf{h}_i^N), \\ \mathbf{h}_{\text{fw},i}^{\text{edge.to}} &= \text{MLP}^{\text{edge.to}}(\mathbf{h}_i^N), \\ y_{\text{fw},i,j}^{\text{edge}} &= \sigma \left(\text{Biaff}_{\text{fw}}^{\text{edge}} \left(\mathbf{h}_{\text{fw},i}^{\text{edge.from}}, \mathbf{h}_{\text{fw},j}^{\text{edge.to}} \right) \right), \end{aligned} \quad (1)$$

where σ is an element-wise sigmoid function. Biaffine operation $\text{Biaff}_{\text{fw}}^{\text{edge}}$ is defined as:

$$\text{Biaff}_{\text{fw}}^{\text{edge}}(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top \mathbf{U}_{\text{fw}}^{\text{edge}} \mathbf{y} + \mathbf{W}_{\text{fw}}^{\text{edge}}[\mathbf{x}; \mathbf{y}] + b_{\text{fw}}^{\text{edge}},$$

where $\mathbf{U}_{fw}^{\text{edge}}$, $\mathbf{W}_{fw}^{\text{edge}}$ and b_{fw}^{edge} are model parameters. Probability of an edge (i, j) being the c -th edge label $y_{fw,i,j,c}^{\text{label}}$ is calculated for all pairs of nodes.

$$\begin{aligned} \mathbf{h}_{fw,i}^{\text{label_from}} &= \text{MLP}^{\text{label_from}}(\mathbf{h}_i^N), \\ \mathbf{h}_{fw,i}^{\text{label_to}} &= \text{MLP}^{\text{label_to}}(\mathbf{h}_i^N), \\ t_{fw,i,j,c}^{\text{label}} &= \text{Biaff}_{fw,c}^{\text{label}}(\mathbf{h}_{fw,i}^{\text{label_from}}, \mathbf{h}_{fw,j}^{\text{label_to}}), \\ y_{fw,i,j,c}^{\text{label}} &= \text{softmax}_c(t_{fw,i,j,c}^{\text{label}}). \end{aligned} \quad (2)$$

Another form of biaffine operation for the edge label prediction $\text{Biaff}_{fw,c}^{\text{label}}$ is defined as:

$$\text{Biaff}_{fw,c}^{\text{label}}(\mathbf{x}, \mathbf{y}) = \mathbf{x}^\top \mathbf{U}_{fw,c}^{\text{label}} \mathbf{y} + \mathbf{W}_{fw,c}^{\text{label}} \mathbf{y},$$

where $\mathbf{U}_{fw,c}^{\text{label}}$ and $\mathbf{W}_{fw,c}^{\text{label}}$ are model parameters.

A candidate edge (i, j) whose edge probability $y_{fw,i,j}^{\text{edge}}$ ($0 < i, j$) exceeds 0.5 is adopted as a valid edge. Edge label with the highest probability $\arg \max_c y_{fw,i,j,c}$ is selected for each valid edge (i, j) . A candidate top node j whose edge probability $y_{fw,0,j}^{\text{edge}}$ ($0 < j$) exceeds 0.5 is adopted as a top node, allowing multiple tops. Non-top nodes with no incoming or outgoing edge are discarded and remaining nodes are adopted as the predicted nodes.

4.2 DM Frame Classifier

A DM node property `frame` consists of a frame type and frame arguments; e.g. `named:x-c` indicates the frame type is “named entity” with two possible arguments `x` and `c`. The proposed system utilizes the contextualized features to predict the frame types and arguments separately.

Probability of the i -th node being c -th frame type $y_{dm,i,c}^{\text{frame_type}}$ is predicted by applying MLP to the contextualized features:

$$\begin{aligned} t_{dm,i,c}^{\text{frame_type}} &= \text{MLP}_c^{\text{frame_type}}(\mathbf{h}_i^N), \\ y_{dm,i,c}^{\text{frame_type}} &= \text{softmax}_c(t_{dm,i,c}^{\text{frame_type}}). \end{aligned}$$

The number of arguments for a frame is not fixed and the first argument can be trivially inferred from the frame type. Thus, we predict from the second to the fifth arguments for each node. Probability of j -th argument being c -th frame type $y_{dm,i,j,c}^{\text{frame_arg}}$ is also predicted by applying MLP to the contextualized features:

$$\begin{aligned} t_{dm,i,j,c}^{\text{frame_arg}} &= \text{MLP}_{j,c}^{\text{frame_arg}}(\mathbf{h}_i^N), \\ y_{dm,i,j,c}^{\text{frame_arg}} &= \text{softmax}_c(t_{dm,i,j,c}^{\text{frame_arg}}). \end{aligned}$$

4.3 Training Objective

DM and PSD are trained jointly in a multi-task learning setting but independently from other frameworks. The loss for the edge prediction ℓ_{fw}^{edge} is cross entropy between the predicted edge $y_{i,j}^{\text{edge}}$ and the corresponding ground truth label. A top node j is treated as an edge $(0, j)$ and is trained along with the edge prediction. The loss for the edge label prediction ℓ_{fw}^{label} is cross entropy between the predicted edge label $y_{i,j,c}^{\text{label}}$ and ground truth label. The loss for the frame prediction ℓ_{dm}^{frame} is the sum of the frame type prediction loss $\ell_{dm}^{\text{frame_type}}$ and the frame arguments prediction loss $\ell_{dm}^{\text{frame_arg}}$, both of which are cross entropy loss between the prediction and the corresponding ground truth label. Final multi-task loss is defined as:

$$\begin{aligned} \ell_{\text{sdp}} = & \lambda^{\text{label}} (\ell_{dm}^{\text{label}} + \ell_{\text{psd}}^{\text{label}} + \lambda^{\text{frame}} \ell_{dm}^{\text{frame}}) \\ & + (1 - \lambda^{\text{label}}) (\ell_{dm}^{\text{edge}} + \ell_{\text{psd}}^{\text{edge}}). \end{aligned} \quad (3)$$

4.4 Postprocessing

We reconstruct node property `frame` from the predicted frame types and arguments using external resources. For DM, we filter out pairs of predicted frame type and arguments that do not appear in ERG SEM-I⁵ or the training dataset (e.g. a word “parse” has only two possible frames `n:x` and `v:e-i-p`). Then, we select a frame with the highest *empirically scaled likelihood* which is calculated by scaling predicted joint probability $y_{dm,i,c}^{\text{frame_type}} \prod_j y_{dm,i,j,c'}^{\text{frame_arg}}$ proportionally to the frame frequency in the corpus.

For PSD, we use CzEngVallex⁶, which contains frequency and the required arguments of each frame, to reconstruct frames. We identify the frame type of a token from its lemma and POS tag. Then, candidate frames are filtered using the required arguments (extracted by stripping `-suffix` from connected edges) and the most frequent frame is chosen as the node frame.

Token lemma is used for the node label, except for the special node labels in PSD (e.g. `#Bracket` and `#PersPron`) that are looked-up from a hand-crafted dictionary using the surface and POS tag as a key.

⁵<http://svn.delph-in.net/erg/tags/1214/etc>

⁶<http://hdl.handle.net/11234/1-1512>

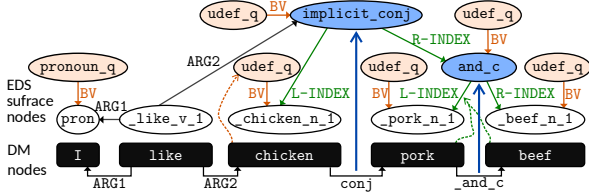


Figure 2: Generation of abstract nodes and their edges from *I like chicken, pork and beef*.

5 EDS-specific Procedure

DM graphs are constructed by lossy conversion from EDS graphs, both of which are derived from English Resource Semantics (ERS; Flickinger et al., 2014). Making use of such relationship, we developed heuristic inverse conversion from DM to EDS graphs by carefully studying EDS-to-DM conversion rules described in the ERG SEM-I corpus. Specifically, our system generates EDS in three steps; the system (i) convert all DM nodes to EDS surface nodes⁷ with simple rules, (ii) generate abstract nodes, and (iii) predict anchors for the abstract nodes.

We explain the generation of abstract nodes (ii) in details using an example in Figure 2:

1. Some abstract nodes (e.g. `_and_c`) and their node labels are generated with rules.
2. Presence of an abstract node on a node or an edge is detected with rules (e.g. `_and_c` implies presence of `_q` node) or with binary logistic regression (e.g. `undef_q` on `_chicken_n_1`).
3. The system predicts labels of the nodes generated in 2 using multi-class logistic regression.
4. The system predicts labels of edges from/to the generated nodes using multi-class logistic regression.

POS tags, predicted DM frames and edge labels of adjacent nodes are used as features for the logistic regression.

We employ another neural network that utilize the contextualized features from the encoder to predict the anchors for the generated abstract nodes (iii). For each abstract node (indexed i), let \mathcal{T}_i be a subset of token indices $\mathcal{S} \equiv \{0, \dots, L_{in} - 1\}$ each of which is selected as a DM node and the corresponding EDS surface node has the abstract node i as an ancestor. First, we create an

⁷For ease of explanation, we adopt a definition that “the EDS surface nodes are the nodes that appear in DM and the abstract nodes are those that do not” which results in slight inconsistency with the original definition.

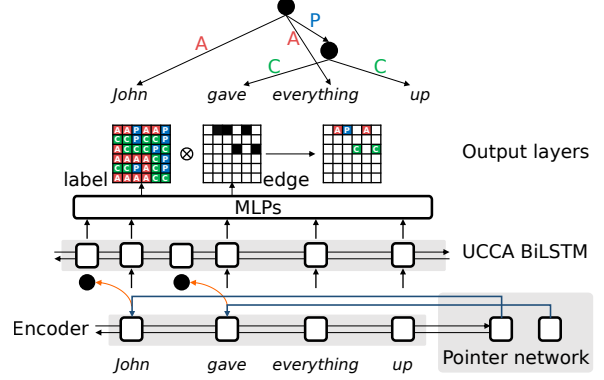


Figure 3: Illustration of UCCA parsing with pointer network and biaffine classifier.

input feature $x_{i,j}^{\text{eds}}$ ($j \in \mathcal{S}$) which is set as the label of node i if $j \in \mathcal{T}_i$ or `<UNK>` otherwise. Then, we embed $x_{i,j}^{\text{eds}}$ to obtain trainable vector $\mathbf{e}_{i,j}^{\text{eds}}$ and feed them to a biLSTM to obtain a contextualized representation $\mathbf{h}_{i,j}^{\text{eds}}$. Finally, we predict a span in input tokens $[\text{argmax}_j y_{i,j}^{\text{eds.from}}, \text{argmax}_j y_{i,j}^{\text{eds.to}}]$ for the i -th abstract node,

$$y_{i,j}^{\text{eds.from}} = \text{softmax}_j \left((\mathbf{h}_{i,j}^{\text{eds}})^\top \cdot \text{MLP}^{\text{eds.from}}(\mathbf{h}_j^N) \right),$$

$$y_{i,j}^{\text{eds.to}} = \text{softmax}_j \left((\mathbf{h}_{i,j}^{\text{eds}})^\top \cdot \text{MLP}^{\text{eds.to}}(\mathbf{h}_j^N) \right).$$

The loss for the anchor prediction ℓ_{eds} is the sum of cross entropy between the predicted span $(y_{i,j}^{\text{eds.from}}, y_{i,j}^{\text{eds.to}})$ and the corresponding ground truth span.

6 UCCA-specific Procedure

A UCCA graph consists of terminal nodes which represent words, non-terminal nodes which represent internal structure, and labeled edges (e.g., participant (A), center (C), linker (L), process (P) and punctuation (U)) which represent connections between the nodes. Motivated by the recent advances in constituency parsing, we predict spans of each terminal nodes at once without using any complicated mechanism as seen in transition-based (Herscovich and Arviv, 2019) and greedy bottom-up (Yu and Sagae, 2019) systems. Our proposed UCCA parser (Figure 3) consists of (i) a pointer network (Vinyals et al., 2015) which generates non-terminal nodes from the contextualized token representations of the encoder, (ii) an additional biLSTM that encodes context of both the terminal and generated non-terminal nodes, and (iii) a biaffine network which predicts edges.

6.1 Preprocessing

We treat the generation of non-terminal nodes as a “pointing” problem. Specifically, the system has to point the starting position of a span which has terminal or non-terminal children. For example, upper part of Figure 3 shows a graph with two non-terminal nodes \bullet . The right non-terminal node has a span of *gave everything up*, and our system points at the starting position of the span *gave*. By taking such strategy, we can serialize the graph in a consistent, straightforward manner; i.e. by inserting the non-terminal nodes to the left of the corresponding span.

The system also has to predict an anchor of a proper noun or a compound expression to merge constituent tokens into a single node. For example, *no feathers in stock!!!!* is tokenized as “(no), (feathers), (in), (stock), (!), (!), (!), (!)” according to the companion data, but the UCCA parser is expected to output “(no), (feathers), (in), (stock), (!!!!)”. To solve the problem, we formulate the merge of tokens as edge prediction; e.g. we assume that there exist virtual edges CT from leftmost constituent token to each subsequent token within a compound expression:



and CT is predicted by the system along with the other edges. There still exists tokenization discrepancy between the companion data and the graphs from EWT and Wiki. The graphs with such discrepancy are simply discarded from the training data.

6.2 Generating Non-terminal Nodes with Pointer Network

Our system generates non-terminal nodes by pointing where to insert non-terminal nodes as described in Section 6.1. To point a terminal node, we employ a pointer network, which is a decoder that uses attention mechanism to produce probability distribution over the input tokens. Given hidden states of the encoder \mathbf{h}_j^N , hidden states of the decoder are initialized by the last states of the shared encoder:

$$\begin{aligned} \mathbf{h}_{-1}^{\text{ucca.dec}} &= \left[\vec{\mathbf{h}}_{L_{in}}^{N-K:N}; \overleftarrow{\mathbf{h}}_0^{N-K:N} \right], \\ \mathbf{c}_{-1}^{\text{ucca.dec}} &= \left[\vec{\mathbf{c}}_{L_{in}}^{N-K:N}; \overleftarrow{\mathbf{c}}_0^{N-K:N} \right], \end{aligned}$$

where K is the stacking number of the biLSTMs in the shared encoder. We then obtain the hidden

states of the decoder $\mathbf{h}_i^{\text{ucca.dec}}$ as:

$$\mathbf{h}_i^{\text{ucca.dec}} = \text{LSTM}_{\text{dec}}(\mathbf{x}_i^{\text{ucca.dec}}, \mathbf{h}_{i-1}^{\text{ucca.dec}}, \mathbf{c}_{i-1}^{\text{ucca.dec}}).$$

Attention distribution $\tilde{a}_{i,j}$ over the input tokens is calculated as:

$$\begin{aligned} a_{i,j} &= \mathbf{v}^\top \tanh(\mathbf{W}^{\text{ucca.dec}}[\mathbf{h}_i^{\text{ucca.dec}}; \mathbf{h}_j^N]), \\ \tilde{a}_{i,j} &= \text{softmax}_j(a_{i,j}), \end{aligned}$$

where $\mathbf{W}^{\text{ucca.dec}}$ and \mathbf{v} are parameters of the pointer network. The successive input to the decoder $\mathbf{x}_{i+1}^{\text{ucca.dec}}$ is the encoder states of the pointed token $\mathbf{h}_{\text{argmax}_j \tilde{a}_{i,j}}^N$. $\mathbf{x}_i^{\text{ucca.dec}}$ is chosen from the gold $\tilde{a}_{i,j}$ when training.

The decoder terminates its generation when it finally points the $\langle \text{ROOT} \rangle$. We obtain new hidden states $\mathbf{h}_i^{\text{ucca.ptr}}$ ($0 \leq i \leq L_{\text{ucca}}$) by inserting pointer representations \mathbf{h}^\bullet before the pointed token. For example, *John gave everything up* (discussed above) will have hidden states

$$(\mathbf{h}_{\langle \text{ROOT} \rangle}^N, \mathbf{h}^\bullet, \mathbf{h}_{\text{John}}^N, \mathbf{h}^\bullet, \mathbf{h}_{\text{gave}}^N, \mathbf{h}_{\text{everything}}^N, \mathbf{h}_{\text{up}}^N).$$

The pointer representation is defined as $\mathbf{h}^\bullet = \text{MLP}_\bullet(\mathbf{r})$, where \mathbf{r} is a randomly initialized constant vector.

We note that the generated non-terminal nodes \mathbf{h}^\bullet lack positional information because all \mathbf{h}^\bullet have the same values. To remedy this problem, a positional encoding Vaswani et al. (2017) is concatenated to each of $\mathbf{h}_i^{\text{ucca.ptr}}$ to obtain position-aware $\mathbf{h}_i^{\text{ucca.ptr}'}$. Furthermore, we feed $\mathbf{h}_i^{\text{ucca.ptr}'}$ to an additional biLSTM and obtain $\mathbf{h}_i^{\text{ucca}}$ in order to further encode the order information.

6.3 Edge Prediction with Biaffine Network

Now that we have contextualized representations for all candidate terminal and non-terminal nodes, the system can simply predict the edges and their labels in the exact same way as Flavor (0) graphs (Section 4.1). Following Equation (1) and Equation (2), we obtain probabilities if there exists an edge (i, j) , $y_{\text{ucca},i,j}^{\text{edge}}$, and its label being c , $y_{\text{ucca},i,j,c}^{\text{label}}$, with the input being $\mathbf{h}_i^{\text{ucca}}$ instead of \mathbf{h}_i^N . We treat the remote edges⁸ independently but in the same way as the primary edges to predict $y_{\text{ucca},i,j}^{\text{remote}}$.

The loss for the edge prediction $\ell_{\text{ucca}}^{\text{edge}}$, the edge label prediction $\ell_{\text{ucca}}^{\text{label}}$, the remote edge prediction

⁸Edges for implicit relations and arguments. They were annotated as unlabeled edges each with an attribute `remote` in MRP.

ℓ_{ucca}^{remote} and the pointer prediction ℓ_{ucca}^{dec} are defined as cross entropy between the prediction $y_{ucca,i,j}^{edge}$, $y_{ucca,i,j,c}^{label}$, $y_{ucca,i,j}^{remote}$ and $\tilde{a}_{i,j}$ with the corresponding ground truth labels, respectively. Thus, we arrive at the multi-task objective defined as:

$$\ell_{ucca} = \lambda_{ucca}^{edge} \ell_{ucca}^{edge} + \lambda_{ucca}^{label} \ell_{ucca}^{label} + \lambda_{ucca}^{remote} \ell_{ucca}^{remote} + \lambda_{ucca}^{dec} \ell_{ucca}^{dec}.$$

7 AMR-specific Procedures

Because AMR graphs do not have clear alignment between input tokens and nodes, the nodes have to be identified in prior to predicting edges. Following Zhang et al. (2019), we incorporate a pointer-generator network (i.e. a decoder with copy mechanisms) for the node generation and a biaffine network for the edge prediction. There are two key preconditions in using a pointer-generator network; i.e. (i) node labels and input tokens share fair amount of vocabulary to allow copying a node from input tokens, and (ii) graphs are serialized in a consistent, straightforward manner for it to be easily predicted by sequence generation. To this end, we apply preprocessing to raw AMR graphs, train model to generate preprocessed graphs, and reconstruct final AMR graphs with postprocessing.

7.1 Preprocessing

We modify the input tokens and the node labels to account for the precondition (i). A node labeled with `.*-entity` or a subgraph connected with `name edge` is replaced with a node whose label is an anonymized entity label such as `PERSON.0` (Konstas et al., 2017). Then, for each entity node, a corresponding span of tokens is identified by rules similar to Flanigan et al. (2014); i.e. a span of tokens with the longest common prefix between the token surfaces and the node attribute (e.g. for `date-entity` whose attribute `month` is `11`, we search for “November” and “Nov” in the token surfaces). Unlike Zhang et al. (2019) which has replaced input token surfaces with anonymized entity labels, we add them as an additional input feature as described in Section 3.1 to avoid hurting the performance of other frameworks. At the prediction, we first identify NE tags in input tokens with Illinois NER tagger (Ratinov and Roth, 2009). Then we map them to anonymized entity labels with frequency-based mapping constructed from the training dataset.

For non-entity nodes, we strip sense indices (e.g. `-01`) from node labels (Lyu and Titov, 2018), which will then share fair amount of vocabulary with the input token lemmas. Nodes with labels that still do not appear as lemmas after preprocessing are subject to normal generation from decoder vocabulary.

Directly serializing an AMR graph, which is a directed acyclic graph (DAG), may result in a complex conversion, which do not fulfill the precondition (ii). Therefore, we convert DAG to a spanning tree by replicating nodes with reentrancies (i.e. nodes with more than one incoming edge) for each incoming edge and serialize the graph with simple pre-order traversal over the tree.

7.2 Extended Pointer-Generator Network

We employ an extended pointer-generator network. It automatically switches between three generation strategies; i.e. (1) *source-side copy*, (2) *decoder-side copy* that copies nodes that have been already generated, and (3) normal generation from decoder vocabulary. More formally, it uses attention mechanism to calculate probability distribution \mathbf{p}_i over input tokens, generated nodes and node vocabulary. Given contextualized token representation of the encoder $H_i^{enc} = \{\mathbf{h}_0^l, \dots, \mathbf{h}_{L_{in}-1}^l\}$, we obtain hidden states of the decoder \mathbf{h}_i^{amr} and \mathbf{p}_i as:

$$\begin{aligned} \mathbf{h}_i^{amr}, \mathbf{p}_i &= \text{Decoder}_{amr}(\mathbf{h}_i^{enc'}, \mathbf{h}_{i-1}^{amr}, \mathbf{p}_{i-1}, H_N^{enc}), \\ \mathbf{h}_i^{enc'} &= \text{Encoder}_{amr}(\mathbf{p}_i, \mathbf{h}_0^{amr} \dots \mathbf{h}_{i-1}^{amr}, H_0^{enc}), \\ \mathbf{h}_0^{enc'}, \mathbf{h}_{-1}^{amr} &= \\ &\text{MLP}_{amr} \left(\left[\begin{array}{c} \overrightarrow{\mathbf{h}}_{L_{in}}^N; \overleftarrow{\mathbf{h}}_0^N; \overrightarrow{\mathbf{c}}_{L_{in}}^N; \overleftarrow{\mathbf{c}}_0^N \end{array} \right] \right). \end{aligned}$$

Encoder_{amr} treats a node as if it is a token, and utilizes the encoder (Section 3) with shared model parameters to obtain representation of $(i-1)$ -th generated nodes $\mathbf{h}_i^{enc'}$. Concretely, Encoder_{amr} combines lemma (corresponds to the node label), POS tags (only when copied from a token) and GloVe (from the node label) of a node, embeds each of them to a feature vector using the encoder and concatenates feature vectors to obtain $\mathbf{h}_i^{enc'}$.

7.3 Edge Prediction with Biaffine Network

Now that we have representations \mathbf{h}_i^{amr} for all nodes, the system can simply predict the edges and their labels in the same way as Flavor (0) graphs (Section 4.1). Following Equation (1) and Equation (2), we obtain probabilities that there exists an

Table 1: MRP F1 scores for the formal run (shown as “score/rank”)

| Team | Mean | DM | PSD | EDS | UCCA | AMR |
|-----------------------|----------------|----------------|----------------|----------------|----------------|----------------|
| HIT-SCIR | .8620/1 | .9508/2 | .9055/4 | .9075/2 | .8167/1 | .7294/2 |
| SJTU-NICT | .8527/2 | .9550/1 | .9119/3 | .8990/3 | .7780/3 | .7197/3 |
| SUDA-Alibaba | .8396/3 | .9226/7 | .8556/8 | .9185/1 | .7843/2 | .7172/5 |
| Saarland | .8187/4 | .9469/4 | .9128/1 | .8910/4 | .6755/6 | .6672/6 |
| Hitachi (ours) | .7604/5 | .9102/8 | .9121/2 | .8374/6 | .7036/5 | .4386/8 |
| ÚFAL MRPipe | .7474/6 | .8495/9 | .7627/9 | .6745/7 | .7322/4 | .7183/4 |
| ShanghaiTech | .6697/7 | .9488/3 | .8949/6 | .8690/5 | - | .6359/7 |
| Amazon | .5132/8 | .9326/6 | .8998/5 | - | - | .7338/1 |
| JBNU | .4652/9 | .9401/5 | .8788/7 | - | .5069/7 | - |
| SJTU | .4303/10 | .4315/11 | .4761/11 | .5321/8 | .3266/9 | .3851/9 |
| ÚFAL-Oslo | .3442/11 | .8051/10 | .6092/10 | .3064/9 | - | - |
| HKUST | .2450/12 | .3699/12 | .3529/12 | - | .5021/8 | - |
| Bocharov | .0655/13 | - | - | - | - | .3273/10 |
| TUPA† single | .5770 | .5554 | .5176 | .8100 | .2756 | .4473 |
| TUPA† multi | .4534 | .4270 | .5265 | .7395 | .2365 | .3375 |

† baseline (Hershcovich and Arviv, 2019)

edge (i, j) , $y_{amr,i,j}^{edge}$, and its label being c , $y_{amr,i,j,c}^{label}$, with the input being \mathbf{h}_i^{amr} instead of \mathbf{h}_i^N . Note that we do not predict the top nodes for AMR, because the first generated node is always the top node in our formalism.

The loss for the edge prediction ℓ_{amr}^{edge} , the edge label prediction ℓ_{amr}^{label} , and the decoder prediction ℓ_{amr}^{dec} are cross entropy between the prediction $y_{amr,i,j}^{edge}$, $y_{amr,i,j,c}^{label}$ and \mathbf{p}_i with the corresponding ground truth labels, respectively. Thus, we arrive at the multi-task loss for AMR defined as:

$$\begin{aligned} \ell_{amr} = & \lambda_{amr}^{biaf} \left(\lambda_{amr}^{label} \ell_{amr}^{label} + (1 - \lambda_{amr}^{label}) \ell_{amr}^{edge} \right) \\ & + \lambda_{amr}^{cov} \ell_{amr}^{cov} + (1 - \lambda_{amr}^{biaf} - \lambda_{amr}^{cov}) \ell_{amr}^{dec}, \end{aligned}$$

where ℓ_{amr}^{cov} is coverage loss (Zhang et al., 2019).

For node prediction, we adopt beam search with search width of five. For edge prediction, we apply Chu-Liu-Edmonds algorithm to find the maximum spanning tree. Postprocessing, which includes inverse transformation of the preprocessing, is applied to reconstruct final AMR graphs.

8 Multi-task Variant

We developed multi-task variant after the formal run. Multi-task variant is trained to minimize following multi-task loss,

$$\begin{aligned} \ell_{mt} = & \lambda^{biaf} \left(\lambda^{label} \left(\sum_{fw} \ell_{fw}^{label} + \lambda^{frame} \ell_{dm}^{frame} \right) \right. \\ & \left. + (1 - \lambda^{label}) \sum_{fw} \ell_{fw}^{edge} \right) + \lambda^{cov} \ell_{amr}^{cov} \\ & + \sum_{fw \in \{ucca, amr\}} \lambda_{fw}^{dec} \ell_{fw}^{dec} + \lambda_{ucca}^{remote} \ell_{ucca}^{remote}. \end{aligned} \quad (4)$$

All training data is simply merged and losses for frameworks that are missing in an input data are

set to zero. For example, if an input sentence has reference graphs for DM, PSD and AMR, losses for UCCA (ℓ_{ucca}^{label} , ℓ_{ucca}^{edge} , ℓ_{ucca}^{dec} and ℓ_{ucca}^{remote}) are set to zero and sum of other losses are used to update the model parameters. Then, the training data (sentences) are shuffled at the start of each epoch and are fed sequentially to update the model parameters as in normal mini-batch training. No under/over-sampling was done to scale the losses of frameworks, each with different number of reference graphs, but we instead applied early stopping for each framework separately (see Appendix A for the details). For EDS, we do not train EDS anchor prediction jointly even in multi-task setting but apply transfer learning; the encoder of the EDS anchor prediction network is initialized from trained multi-task model.

We also experimented with a *fine-tuned* multi-task variant. For each target framework, we take the multi-task variant as a pretrained model (whose training data also includes the target framework) and train the model on the target framework independently to the other frameworks (except for DM and PSD, which are always trained together).

9 Experiments

9.1 Method

Experiments were carried out on the evaluation split of the dataset. We applied hyperparameter tuning and ensembling to our system, which are detailed in Appendix A along with other training details. BERT was excluded for the formal run since it did not make it to the task deadline.

We experimented with enhanced models with BERT after the formal run. For these models, we adopted the best hyperparameters chosen by the submitted model without re-running the hyperparameter tuning.

All models were implemented using Chainer (Tokui et al., 2015; Akiba et al., 2017).

9.2 Results

The official results are shown in Table 1 and Table 2. Our system obtained macro-averaged MRP F1 score of 0.7604 and was ranked fifth amongst all submissions. Our system outperformed conventional unified architecture for MRP (TUPA baselines; Hershcovich and Arviv, 2019) in all frameworks but AMR. This indicates the efficacy

Table 2: MRP and framework specific scores (shown as “score/rank”). Gray background indicates that it is the score on LPPS subset.

| Framework | MRP | | | | | | | Framework specific [†] |
|-----------|----------|----------|------------|----------|----------|------------|----------|---------------------------------|
| | Tops | Labels | Properties | Anchors | Edges | Attributes | All | |
| All | 0.8929/3 | 0.6409/6 | 0.5186/9 | 0.7547/5 | 0.6958/5 | 0.0418/7 | 0.7604/5 | - |
| | 0.9167/3 | 0.6238/6 | 0.3743/9 | 0.7602/6 | 0.7025/5 | 0.0340/7 | 0.7618/5 | - |
| DM | 0.9219/6 | 0.9107/6 | 0.8649/9 | 0.9909/4 | 0.9190/5 | - | 0.9102/9 | 0.9189/5 |
| | 0.9505/5 | 0.8818/8 | 0.8367/10 | 0.9862/6 | 0.9245/5 | - | 0.8939/9 | 0.9272/4 |
| PSD | 0.9538/5 | 0.9494/3 | 0.9118/7 | 0.9896/5 | 0.7948/5 | - | 0.9121/2 | 0.8085/4 |
| | 0.9515/5 | 0.9204/2 | 0.8366/8 | 0.9820/6 | 0.7846/4 | - | 0.8840/2 | 0.8075/4 |
| EDS | 0.7319/9 | 0.8225/7 | 0.5851/7 | 0.8694/6 | 0.8497/7 | - | 0.8374/7 | 0.7826/7 |
| | 0.8515/7 | 0.7763/7 | 0.0670/9 | 0.8737/7 | 0.8427/7 | - | 0.8110/7 | 0.7571/7 |
| UCCA | 0.9965/2 | - | - | 0.9238/6 | 0.5588/6 | 0.2092/7 | 0.7036/6 | 0.4277/6 |
| | 0.9900/2 | - | - | 0.9593/7 | 0.6050/6 | 0.1698/7 | 0.7498/6 | 0.5024/6 |
| AMR | 0.8604/3 | 0.5221/8 | 0.2314/9 | - | 0.3568/8 | - | 0.4386/8 | 0.4254/8 |
| | 0.8400/4 | 0.5404/8 | 0.1311/9 | - | 0.3558/8 | - | 0.4701/8 | 0.4530/8 |

[†] DM/PSD: SDP labeled F1, EDS: EDM all F1, UCCA:UCCA labeled all F1, AMR: SMATCH F1

Table 3: MRP F1 scores for the variants of the proposed system (shown as “score/rank” where the rank is calculated by assuming that it was the submitted model).

| Variant | Average | DM | PSD | EDS | UCCA | AMR |
|-----------------------------|----------|----------|----------|----------|----------|----------|
| SFL | 0.7575/5 | 0.9071/9 | 0.9064/3 | 0.8339/7 | 0.7014/6 | 0.4386/8 |
| SFL (ensemble) [†] | 0.7604/5 | 0.9102/9 | 0.9121/2 | 0.8374/7 | 0.7036/6 | 0.4386/8 |
| BERT+SFL (NT) | 0.7450/6 | 0.9038/9 | 0.9069/3 | 0.8301/7 | 0.6945/6 | 0.3896/8 |
| BERT+MTL (NT) | 0.7144/6 | 0.8726/9 | 0.8791/7 | 0.7987/7 | 0.6422/6 | 0.3794/9 |
| BERT+MTL+FT (NT) | 0.7507/5 | 0.9045/9 | 0.9054/4 | 0.8304/7 | 0.7126/6 | 0.4008/8 |

SFL: single-framework learning, MTL: multi-task learning, FT: fine-tuning, ensemble: with ensembles, NT: random seed is not tuned, [†] formal run

of using the biaffine network as a shared architecture for MRP.

Our system obtained relatively better (second) position in PSD. This was due to relatively good performance on the node label prediction where we carefully constructed postprocessing rule for special nodes’ labels (Section 4.4) instead of just using lemmas.

Our system obtained significantly worse result in AMR (difference of 0.2952 MRP F1 score to the best performing system), even though our system incorporates the state-of-the-art AMR parser (Zhang et al., 2019). One reason is that Zhang et al. (2019) was obtaining a large score boost from the Wikification task, which was not part of the MRP 2019 shared task. Another reason could be that we may have missed out important implementation details for the pointer-generator network, since the implementation of Zhang et al. (2019) was not yet released at the time of our system development.

Table 3 shows the performance of other variants of the proposed system. The single-framework learning variant (SFL) without BERT (SFL) performed better than SFL with BERT (BERT+SFL (NT)), which suggests that impact of hyperparameter tuning was larger than that of incorporating BERT. The multi-task learning variant (MTL) with fine-tuning (BERT+MTL+FT (NT)) outperformed the SFL in the comparable condition (BERT+SFL (NT)). This result implies learning

heterogeneous meaning representations at once can boost the system performance.

10 Conclusions

In this paper, we described our proposed system for the CoNLL 2019 Cross-Framework Meaning Representation Parsing (MRP 2019) shared task. Our system was the unified encoder-to-biaffine network for all five frameworks. The system was ranked fifth in the formal run of the task, and outperformed the baseline unified transition-based MRP. Furthermore, post-evaluation experiments showed that we can boost the performance of the proposed system by incorporating multi-task learning. These imply efficacy of incorporating the biaffine network to the shared architecture for MRP and that learning heterogeneous meaning representations at once can boost the system performance.

While our architecture successfully unified graph predictions in the five frameworks, it is non-trivial to extend the architecture to another framework. It is because there could be a more suitable node generation scheme for a different framework and naively applying the pointer network for partial nodes complementation (or extended pointer-generator network for full nodes generation) may result in a poor performance. Thus, it is our future work to design a more universal method for the node generation.

References

- Omri Abend and Ari Rappoport. 2013. [Universal conceptual cognitive annotation \(UCCA\)](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*.
- Takuya Akiba, Keisuke Fukuda, and Shuji Suzuki. 2017. [ChainerMN: Scalable Distributed Deep Learning Framework](#). In *Proceedings of Workshop on ML Systems in The 31st Annual Conference on Neural Information Processing Systems*.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. [Abstract meaning representation for sembanking](#). In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Timothy Dozat and Christopher D. Manning. 2018. [Simpler but More Accurate Semantic Dependency Parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Jeffrey Flanigan, Sam Thomson, Jaime Carbonell, Chris Dyer, and Noah A. Smith. 2014. [A discriminative graph-based parser for the abstract meaning representation](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*.
- Dan Flickinger. 2000. [On building a more efficient grammar by exploiting types](#). *Natural Language Engineering*, 6(1):15–28.
- Dan Flickinger, Emily M. Bender, and Stephan Oepen. 2014. [Towards an encyclopedia of compositional semantics: Documenting the interface of the English resource grammar](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation*.
- Jan Hajič, Eva Hajičová, Jarmila Panevová, Petr Sgall, Ondřej Bojar, Silvie Cinková, Eva Fučíková, Marie Mikulová, Petr Pajas, Jan Popelka, Jiří Semecký, Jana Šindlerová, Jan Štěpánek, Josef Toman, Zdeňka Urešová, and Zdeněk Žabokrtský. 2012. [Announcing Prague Czech-English dependency treebank 2.0](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation*.
- Daniel Hershcovich, Omri Abend, and Ari Rappoport. 2017. [A transition-based directed acyclic graph parser for UCCA](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Daniel Hershcovich, Omri Abend, and Ari Rappoport. 2018. [Multitask parsing across semantic representations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Daniel Hershcovich and Ofir Arviv. 2019. [TUPA at MRP 2019: A multi-task baseline system](#). In *Proceedings of the Shared Task on Cross-Framework Meaning Representation Parsing at the 2019 Conference on Natural Language Learning*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Computation*, 9(8):1735–1780.
- Angelina Ivanova, Stephan Oepen, Lilja Øvrelid, and Dan Flickinger. 2012. [Who did what to whom? a contrastive study of syntacto-semantic dependencies](#). In *Proceedings of the Sixth Linguistic Annotation Workshop*.
- Diederik Kingma and Jimmy Ba. 2014. [Adam: A Method for Stochastic Optimization](#). In *the Third International Conference on Learning Representations*.
- Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. 2017. [Neural AMR: Sequence-to-sequence models for parsing and generation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Chunchuan Lyu and Ivan Titov. 2018. [AMR parsing as graph prediction with latent alignment](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*.
- Yusuke Miyao, Stephan Oepen, and Daniel Zeman. 2014. [In-house: An ensemble of pre-existing off-the-shelf parsers](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation*.
- Stephan Oepen, Omri Abend, Jan Hajič, Daniel Hershcovich, Marco Kuhlmann, Tim O’Gorman, Nianwen Xue, Jayeol Chun, Milan Straka, and Zdeňka Urešová. 2019. [MRP 2019: Cross-framework Meaning Representation Parsing](#). In *Proceedings of the Shared Task on Cross-Framework Meaning Representation Parsing at the 2019 Conference on Natural Language Learning*.
- Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Dan Flickinger, Jan Hajič, Angelina Ivanova, and Yi Zhang. 2014. [SemEval 2014 task 8: Broad-coverage semantic dependency parsing](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation*.
- Stephan Oepen and Jan Tore Lønning. 2006. [Discriminant-based MRS banking](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global vectors for word](#)

- representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep Contextualized Word Representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Lev Ratinov and Dan Roth. 2009. [Design challenges and misconceptions in named entity recognition](#). In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning*.
- Mike Schuster and Kuldip K Paliwal. 1997. [Bidirectional recurrent neural networks](#). *IEEE Transactions on Signal Processing*, 45(11):2673–2681.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15:1929–1958.
- Seiya Tokui, Kenta Oono, Shohei Hido, and Justin Clayton. 2015. [Chainer: a next-generation open source framework for deep learning](#). In *Proceedings of Workshop on Machine Learning Systems in The 29th Annual Conference on Neural Information Processing Systems*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. [Pointer networks](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2692–2700. Curran Associates, Inc.
- Dian Yu and Kenji Sagae. 2019. [UC Davis at SemEval-2019 task 1: DAG semantic parsing with attention-based decoder](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*.
- Sheng Zhang, Xutai Ma, Kevin Duh, and Benjamin Van Durme. 2019. [AMR parsing as sequence-to-graph transduction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

A Training Details

We split dataset into training dataset which was used to update model parameters, validation dataset (i) which was used for early stopping, and validation dataset (ii) which was used for hyperparameter tuning and construction of ensembles. For AMR and UCCA, we selected sentences that appear in more than one framework to populate the training dataset, and extracted 500 (300) and 1500 (700) data from the rest as validation dataset (i) and (ii) for AMR (UCCA), respectively. For DM, PSD and EDS, we selected data that appear in AMR or UCCA to populate the training dataset, and extracted 500 and 1500 data from the rest as validation dataset (i) and (ii), respectively.

A.1 Model Training

All models are trained using mini-batch stochastic gradient descent with backpropagation. We use Adam optimizer (Kingma and Ba, 2014) with gradient clipping.

For the non-multi-task variant, early stopping is applied for each framework with SDP labeled dependency F1 score (Oepen et al., 2014) (for DM, PSD and UCCA) or validation loss (for EDS and AMR) as the objective. Note that early stopping is applied separately to each framework for the joint training of DM and PSD. Concretely, for the joint training of DM and PSD, we train the model with respect to the joint loss ℓ_{sdp} in Equation (3) but we use a model at a training epoch whose DM-specific (or PSD-specific) SDP labeled dependency F1 score is highest for DM (or PSD) prediction.

For the multi-task variants, we employ a slightly different strategy for early stopping. For the multi-task variant without fine-tuning, we apply early stopping separately to each framework with respect to the framework-specific validation loss. For example, we train the multi-task model with respect to ℓ_{mtl} in Equation (4) but we use a model at a training epoch whose PSD-specific validation loss $\lambda^{label} \ell_{psd}^{label} + (1 - \lambda^{label}) \ell_{psd}^{edge}$ is lowest for PSD prediction. For each framework in the fine-tuned multi-task variant, we adopt the multi-task pretrained model at a training epoch whose framework-specific validation loss is lowest and fine-tune on the model in the same manner as the non-multi-task variant. Note that, for DM and PSD, which are fine-tuned together even in the fine-tuned multi-task variant, we adopt the multi-

task pretrained model at a training epoch whose multi-task validation loss ℓ_{mtl} is lowest.

Dropout (Srivastava et al., 2014) is applied to (a) the input to each layer of the shared encoder, (b) the input to the biaffine networks, and (c) the input to each layer of the UCCA and AMR decoders.

A.2 Hyperparameter Tuning

We random searched subset of hyperparameters for DM, PSD, UCCA and AMR. See Table 4 for hyperparameter search space and the list of hyperparameters chosen by the best performing model in each framework. We tried 20 hyperparameter sets for DM/PSD, 50 for UCCA, and 25 for AMR.

We did not tune the hyperparameters of the multi-task variants. We adopted the best hyperparameters chosen in the non-multi-task variants (Table 4) and hand-tuned the hyperparameters by examining learning curves over few runs. For the fine-tuning, we adopted the best hyperparameters chosen in the non-multi-task variants (Table 4). See Table 5 for the list of hyperparameters used in the multi-task variants.

A.3 Ensembling

We formed ensembles from the models trained in the hyperparameter tuning. Models are added to the ensemble in descending order of MRP F1 score on validation dataset (II) until MRP F1 score of the ensemble no longer improves.

For DM and PSD, we simply averaged edge predictions $y_{fw,i,j}^{edge}$ and label predictions $y_{fw,i,j,c}^{label}$, respectively. On the other hand, the simple average ensembling cannot be applied to UCCA, because number of nodes maybe distinct to each model due to the non-terminal node generation. Hence, we propose to use a two-step *voting ensemble* for UCCA; for each input sentence, (1) the most popular pointer sequence is chosen, and (2) edge and label predictions from the models that outputted the chosen sequence are averaged in the same way as DM and PSD.

For EDS, we do not explicitly use ensemble learning, but utilize DM graphs from ensembled DM models to reconstruct EDS graphs. For AMR, we do not use ensembles.

Table 4: List of hyperparameters. Multiple values indicates that the hyperparameter was tuned within that values. Subscript d (DM), p (PSD), u (UCCA) and a (AMR) denotes the hyperparameter chosen by the best performing model on validation dataset. $\mathcal{U}(a, b)$ is a uniform distribution in $[a, b]$.

| Hyperparameter | Value or search space |
|--|---|
| Common | |
| Word embedding dimension | 100 |
| Lemma embedding dimension | 100 |
| POS embedding dimension | 100 |
| NE embedding dimension | 100 |
| GloVe MLP hidden size | 125 |
| ELMo MLP hidden size | 512 |
| Word drop probability | $0.1_{dpua}, 0.2, 0.4$ |
| POS drop probability | $0.1_{du}, 0.2_a, 0.4_p$ |
| Lemma drop probability | $0.1_p, 0.2_{da}, 0.4_u$ |
| # of layers in encoder | $2_{pu}, 3_{da}$ |
| Encoder LSTM hidden size | $256, 512_{dpua}$ |
| Encoder dropout rate | $0.1_a, 0.25_d, 0.5_{pu}$ |
| Biaffine input dropout | $0.2_{pua}, 0.45_d$ |
| Edge prediction dropout | $0.25_{dpua}, 0.4$ |
| Learning rate | $10^{\mathcal{U}(-3.32, -2.92)}$ → $0.000858_d, 0.000675_p,$ $0.00117_u, 0.00059_a$ |
| Adam $(\beta_1, \beta_2)^\dagger$ | $(0.9, 0.999)_{dp}, (0, 0.95)_{ua}$ |
| DM/PSD | |
| Edge MLP hidden size | 600 |
| Edge label MLP hidden size | 600 |
| Frame prediction MLP hidden size | 600 |
| Frame prediction dropout | $0.2, 0.55_{dp}$ |
| Edge label prediction dropout | $0.33_d, 0.5_p$ |
| Loss coefficient λ_{fw}^{label} | $\mathcal{U}(0.02, 0.03)$ → $0.0210_d, 0.0242_p$ |
| Loss coefficient λ_{fw}^{frame} | 0.5 |
| # of epochs | 50 |
| Batch size | 64 |
| UCCA | |
| Edge MLP hidden size | 400, 500 _u , 600 |
| Edge label MLP hidden size | 400 _u , 500, 600 |
| Edge label prediction dropout | $0.25_u, 0.33$ |
| Decoder dropout | 0.5 |
| Loss coefficient λ_{ucca}^{edge} | 0.3 |
| Loss coefficient λ_{ucca}^{label} | 0.3 |
| Loss coefficient λ_{ucca}^{remote} | 0.2 |
| Loss coefficient λ_{ucca}^{dec} | 0.2 |
| # of epochs | 40 |
| Batch size | 100 |
| AMR | |
| Edge MLP hidden size | 600 |
| Edge label MLP hidden size | 600 |
| Edge label prediction dropout | $0.33_a, 0.5$ |
| Decoder type [‡] | deep small _a , shallow wide |
| Decoder dropout | $0.25, 0.33_a, 0.5$ |
| Loss coefficient λ_{amr}^{label} | $\mathcal{U}(0.1, 0.5) \rightarrow 0.395_a$ |
| Loss coefficient λ_{amr}^{cov} | $\mathcal{U}(0.2, 0.4) \rightarrow 0.339_a$ |
| Loss coefficient λ_{amr}^{gen} | $\mathcal{U}(0.2, 0.4) \rightarrow 0.271_a$ |
| # of epochs | 50 |
| Batch size | 64 |

[†] Commonly used setting and the setting used in Dozat and Manning (2018).

[‡] “deep small” is three-layered LSTM with hidden size of 512 and “shallow wide” is two-layered LSTM with hidden size of 1024.

Table 5: Hyperparameters for the multi-task variants

| Hyperparameter | Value |
|--|----------------------|
| Model architecture | |
| Word embedding dimension | 100 |
| Lemma embedding dimension | 100 |
| POS embedding dimension | 100 |
| NE embedding dimension | 100 |
| GloVe MLP hidden size | 125 |
| ELMo MLP hidden size | 512 |
| # of layers in encoder | 3 |
| Encoder LSTM hidden size | 512 |
| Edge MLP hidden size | 600 |
| Edge label MLP hidden size | 600 |
| Frame prediction MLP hidden size | 600 |
| AMR decoder type [†] | deep small |
| Training conditions | |
| Multi-task (pre)training | |
| Word drop probability | 0.2 |
| POS drop probability | 0.2 |
| Lemma drop probability | 0.2 |
| Encoder dropout rate | 0.5 |
| Biaffine input dropout | 0.45 |
| Edge prediction dropout | 0.25 |
| Edge label prediction dropout | 0.33 |
| Learning rate | 0.00006 |
| Adam $(\beta_1, \beta_2)^\dagger$ | $(0.9, 0.999)$ |
| Loss coefficient λ^{biaf} | 1.0 |
| Loss coefficient λ^{label} | 0.15 |
| Loss coefficient λ^{frame} | 0.5 |
| Loss coefficient λ_{ucca}^{remote} | 0.5 |
| Loss coefficient λ_{ucca}^{dec} | 0.08 |
| Loss coefficient λ_{ucca}^{dec} | 1.2 |
| Loss coefficient λ_{amr}^{cov} | 1.0 |
| # of epochs | 60 |
| Batch size | 128 |
| DM/PSD fine-tuning | |
| Word drop probability | 0.1 |
| POS drop probability | 0.2 |
| Lemma drop probability | 0.2 |
| Encoder dropout rate | 0.25 |
| Biaffine input dropout | 0.45 |
| Edge prediction dropout | 0.25 |
| Learning rate | 0.001^\ddagger |
| Adam $(\beta_1, \beta_2)^\dagger$ | $(0, 0.95)^\ddagger$ |
| Frame prediction dropout | 0.55 |
| Edge label prediction dropout | 0.33 |
| Loss coefficient λ_{fw}^{label} | 0.025 |
| Loss coefficient λ_{fw}^{frame} | 0.5 |
| # of epochs | 50 |
| Batch size | 64 |
| UCCA fine-tuning | |
| Word drop probability | 0.1 |
| POS drop probability | 0.1 |
| Lemma drop probability | 0.4 |
| Encoder dropout rate | 0.5 |
| Biaffine input dropout | 0.2 |
| Edge prediction dropout | 0.25 |
| Learning rate | 0.00117 |
| Adam $(\beta_1, \beta_2)^\dagger$ | $(0, 0.95)$ |
| Edge label prediction dropout | 0.25 |
| Decoder dropout | 0.5 |
| Loss coefficient λ_{ucca}^{edge} | 0.3 |
| Loss coefficient λ_{ucca}^{label} | 0.3 |
| Loss coefficient λ_{ucca}^{remote} | 0.2 |
| Loss coefficient λ_{ucca}^{dec} | 0.2 |
| # of epochs | 40 |
| Batch size | 100 |
| AMR fine-tuning | |
| Word drop probability | 0.1 |
| POS drop probability | 0.2 |
| Lemma drop probability | 0.2 |
| Encoder dropout rate | 0.1 |
| Biaffine input dropout | 0.2 |
| Edge prediction dropout | 0.25 |
| Learning rate | 0.00059 |
| Adam $(\beta_1, \beta_2)^\dagger$ | $(0, 0.95)$ |
| Edge label prediction dropout | 0.33 |
| Decoder dropout | 0.33 |
| Loss coefficient λ_{amr}^{label} | 0.395 |
| Loss coefficient λ_{amr}^{cov} | 0.339 |
| Loss coefficient λ_{amr}^{gen} | 0.271 |
| # of epochs | 50 |
| Batch size | 64 |

[†] See Table 4.

[‡] These are bugs. They should have been different values according to Table 4.