Ensemble Graph Prediction

Anonymous Author(s)
Affiliation
Address
email

Abstract

In many machine learning tasks, models are trained to predict structure data such as graphs. For example, in natural language processing, it is very common to parse texts into dependency trees or abstract meaning representation graphs (AMR). On the other hand, ensemble methods combine predictions from multiple models to create a new one that is more robust and accurate than any individual prediction. In the literature, there are many ensembling techniques proposed for classification or regression problems, however, the problem of ensemble graph prediction has not been studied thoroughly. In this work, we formalize this problem as mining the largest subgraph that is the most supported by a collection of graph predictions. As the problem is NP-Hard, we propose an efficient heuristic algorithm to approximate the optimal solution. To validate our approach, we carried out experiments in AMR parsing problems. The experimental results demonstrate that the proposed approach can combine the strength of the state-of-the-art AMR parsers to create new predictions that are more accurate than any individual model in five standard benchmark datasets. We also contribute the source code for reproducibility.

1 Introduction

Ensemble learning is a popular machine learning practice, in which predictions from multiple models are blended to create a new one that is usually more robust and accurate. Indeed, ensemble methods like XGBOOST are the winning solution in many machine learning and data science competitions [11]. A key reason behind the successes of the ensemble methods is that they can combine the strength of different models to reduce the variance and bias in the final prediction [13, 24].

Research in ensemble methods mostly focuses on regression or classification problems [14]. Recently, in many machine learning tasks prediction outputs are provided in a form of graphs. For instance, in Abstract Meaning Representation (AMR) parsing [5], the input is a fragment of text and the output is a labeled graph. Figure 1 shows an AMR graph for the sentence *You told me to wash the dog*. AMR parsing is an important problem in natural language processing (NLP) research and it has a broad application in downstream tasks such as question answering [15] and common sense reasoning [20]. Recent approaches for AMR parsing leverage the advance from pretrained language models [6] and numerous deep neural network architectures [8, 28].

Unlike methods for ensembling numerical or categorical values for regression or classification problems where the mean value or majority votes are used respectively, the problem of graph ensemble is more complicated. For instance, Figure 2 show three graphs g_1, g_2, g_3 with different structures, having varied number of edges and vertices with different labels. In this work, we formulate the ensemble graph prediction as a graph mining problem where we look for the largest common structure among the graph predictions. In general, finding the largest common subgraph is a well-known computationally intractable problem in graph theory. However, for AMR parsing problems where the AMR graphs have labels and a simple tree-alike structure, we propose an efficient heuristic algorithm (*Graphene*) to approximate the solution of the given problem well.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

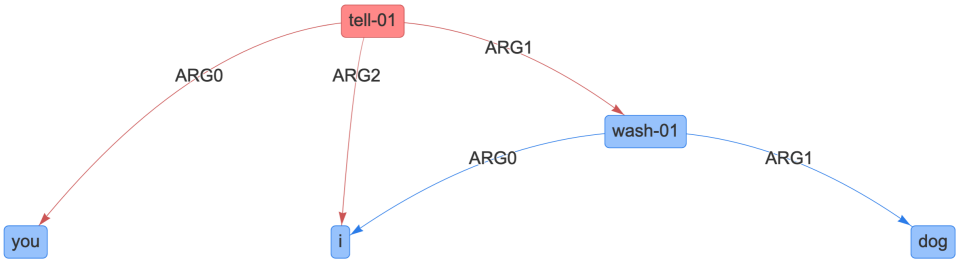

Figure 1: An example AMR graph for the sentence *You told me to wash the dog*.

To validate our approach, we collect the predictions from four state-of-the-art AMR parsers and create new predictions using the proposed graph ensemble algorithm. The chosen AMR parsers are the recent state-of-the-art AMR parsers including a seq2seq-based method using BART [6], transition-based approach proposed in [28] and a graph-based approach proposed in [8]. In addition to those models, we also trained a new seq2seq model based on T5 [21] to leverage the strength of this pretrained language model.

The experimental results show that in five standard benchmark datasets, our proposed ensemble approach outperforms the previous state-of-the-art models and achieves new state-of-the-art results in all datasets. For example, our approach achieves new state-of-the-art results with 1.7, 1.5, and 1.3 points better than prior arts in the BIO (under out-of-distribution evaluation), AMR 2.0, and AMR 3.0 datasets respectively. This result demonstrates the strength of the ensemble methods in leveraging the model diversity to achieve better performance. An interesting property of our solution is that it is model-agnostic, therefore it can be used to make an ensemble of existing model predictions without the requirement to have an access to model training. We also contribute the source code for reproducibility (see the supplemental materials).

Our paper is organized as follows: Section 2 discusses a formal problem definition and a study on the computational intractability of the formulated problem. The graph ensemble algorithm is described in Section 3. Experimental results are reported in Section 4 while Section 5 discusses related works. The conclusion and future work are discussed in Section 6.

2 Problem formulation

Denote $g = (E, V)$ as a graph with the set of edges E and the set of vertices V. Each vertex $v \in V$ and edge $e \in E$ is associated with a label denoted as $l(v)$ and $l(e)$ respectively, where $l(.)$ is a labelling function. Given two graphs $g_1 = (E_1, V_1)$ and $g_2 = (E_2, V_2)$, a *vertex matching* ϕ is a bijective function that maps a vertex $v \in V_1$ to a vertex $\phi(v) \in V_2$.

Example 1. *In Figure 2, between g_1 and g_2 there are many possible vertex matches, where $\phi(g_1, g_2) = [1 \rightarrow 3, 2 \rightarrow 2, 3 \rightarrow 1]$ is one of them (which can be read as the first vertex of g_1 is mapped to the third vertex of g_2 and so forth). Notice that not all vertices $v \in V_1$ has a match in V_2 and vice versa. Indeed, in this example, the fourth vertex in g_2 does not have a matched vertex in g_1.*

Given two graphs g_1, g_2 and a vertex match $\phi(g_1, g_2)$, we define support for every vertex v and edge $e = (v_1, v_2)$ as follows:

- Vertex support with respect to the vertex match ϕ, denoted as $s_\phi(v)$, is equal to 1 if $l(v) = l(\phi(v))$ and 0 otherwise.
- Edge support with respect to the vertex match ϕ, denoted as $s_\phi(e)$, is equal to 1 if $l(e) = l((\phi(v_1), \phi(v_2)))$ and 0 otherwise.

Example 2. *In Figure 2, for the vertex match $\phi(g_1, g_2) = [1 \rightarrow 3, 2 \rightarrow 2, 3 \rightarrow 1]$, the first vertex in g_1 and the third vertex in g_2 shares the same label A, therefore the support of the given vertex is equal to 1. On the other hand, the third vertex in g_1 and the first vertex in g_2 does not have the same label so their support is equal to 0.*

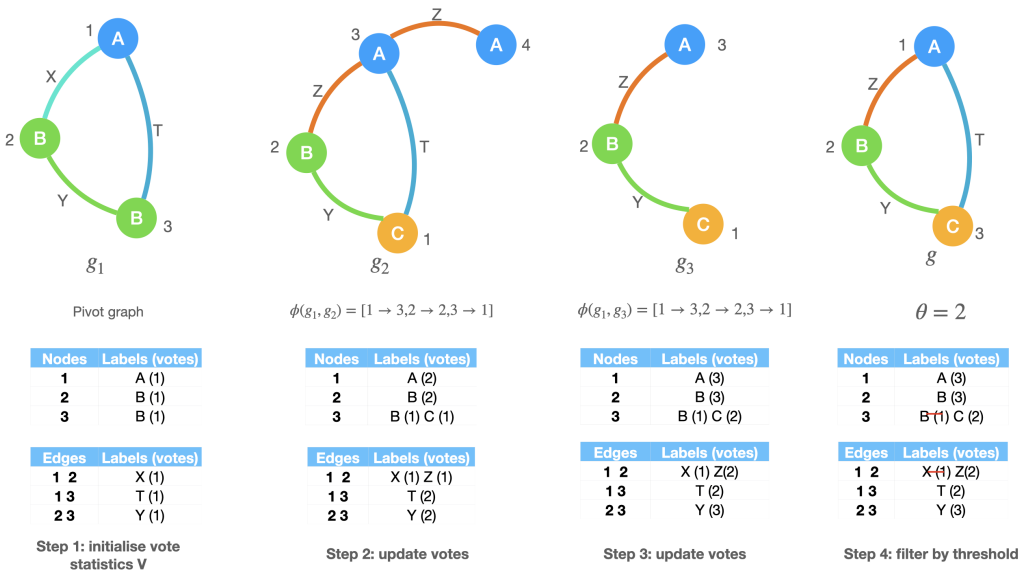

Figure 2: Each node and edge of g occurs in at least two out of three graphs g_1, g_2, g_3. Therefore, g is θ-supported where $\theta = 2$ by the given set of graphs. Graph g is also the graph with the largest sum of node and edge supports among all θ-supported graphs. The tables show the node and edge support (votes) counted in each step of the Graphene algorithm when g_1 is chosen as a pivot graph.

Between two graphs, there are many possible vertex matches, the *best vertex match* is defined as the one that has the maximal total vertex support and edge support. In our discussion, when we mention a vertex match we always refer to the best vertex match.

Denote $G = \{g_1 = (E_1, V_1), g_2 = (E_2, V_2), \cdots, g_m = (E_m, V_m)\}$ as a set of m graphs. Given any graph $g = (E, V)$, for every g_i denote $\phi_i(g, g_i)$ as the best vertex match between g and g_i. The total support of a vertex $v \in V$ or an edge $e \in E$ as is defined as follows:

- $support(e) = \sum_{i=1}^{m} s_{\phi_i}(e)$
- $support(v) = \sum_{i=1}^{m} s_{\phi_i}(v)$

Given a support threshold θ, a graph g is called θ-supported by G if for any node $v \in V$ or any edge $e \in E$, $support(v) \geq \theta$ and $support(e) \geq \theta$.

Example 3. *In Figure 2, graph g is θ-supported by $G = \{g_1, g_2, g_3\}$ where $\theta = 2$.*

Intuitively, an ensemble graph g should have as many common edges and vertices with all the graph predictions as possible. Therefore, we define the graph ensemble problem as follows:

Problem 1 (Graph ensemble)**.** *Given a support threshold θ and a collection of graphs G, find the graph g that is θ-supported by G and has the largest sum of vertex and edge supports.*

Theorem 1. *Finding the optimal θ-supported graph with the largest total of support is an NP-Hard problem.*

Proof. We prove the NP-Hardness by reduction to the *Maximum Common Edge Subgraph* (MCIS) problem, which is known to be an NP-Complete problem [4]. Given two graphs g_1 and g_2, the MCIS problem finds a graph g that is a common subgraph of g_1 and g_2 and the number of edges in g is the largest. Consider the following instance of the Graph Ensemble problem with $\theta = 2$, and $G = \{g_1, g_2\}$ created from the graphs in the MCIS problem. Assume that all vertices and all edges of g_1 and g_2 have the same label A.

Since $\theta = 2$, a θ-supported graph is also a common subgraph between g_1 and g_2 and vice versa. Denote g_s and g_e is the common subgraph between g_1 and g_2 with the largest support and the largest common edge, respectively. We can show that g_s has as many edges as g_e. In fact, since g_s is the

largest supported common subgraph there is no vertex $v \in g_e$ such that $v \notin g_s$ because otherwise we can add v to g_s to create a larger supported graph.

For any edge $e = (v_1, v_2) \in g_e$, since both vertices v_1 and v_2 also appear in g_s, the edge $e = (v_1, v_2)$ must also be part of g_s otherwise we can add this edge to g_s to create a subgraph with a larger support. Therefore, g_s has as many edges as g_e, which is also a solution to the MCIS problem. □

3 Please correct me if I am wrong

In this section, we discuss a heuristic algorithm based on the strategy *"Please correct me if I am wrong!"* to solve Problem 1. The main idea is to improve a pivot graph based on other graphs. Specifically, starting with a pivot graph g_i $(i = 1, 2, \cdots, m)$, we collect the votes from the other graphs for every existing vertex and existing/non-existing edges to correct g_i. We call the proposed algorithm *Graphene* which stands for Graph Ensemble algorithm. The key steps of the algorithm are provided in the pseudo-code in Algorithm 1.

For example, in Figure 2, the algorithm starts with the first graph g_1 and considers it as a pivot graph g_{pivot}. In the first step, it creates a table to keep voting statistics V initialized with the vote counts for every existing vertex and edge in g_{pivot}. To draw additional votes from the other graphs, it performs the following subsequent steps:

- Call the function $\phi(g_1, g_i)$ $(i = 2, 3, \cdots, m)$ to get the best bijective mapping ϕ between the vertices of two graphs g_1 and g_i (with a little bit abuse of notation we drop the index i from ϕ_i when g_i and g_{pivot} are given in the context). For instance, the best vertex match between g_1 and g_2 is $\phi = 1 \to 3, 2 \to 2, 3 \to 1$ because that vertex match has the largest number of common labeled edges and vertices.
- Enumerate the matching vertices and edges to update the voting statistics accordingly. For instance, since the vertex 3 in g_1 with label B is mapped to the vertex 1 in g_2 with label C, a new candidate label C is added to the table for the given vertex. For the same reason, we add a new candidate label Z for the edge $(1, 2)$. For all the other edges and vertices where the labels are matched the votes are updated accordingly.

Once the complete voting statistics V is available, the algorithm filters the candidate labels of edges and vertices using the provided support threshold θ by calling the function $Filter(V, \theta)$ to obtain an ensemble graph g_i^e. For special cases, when disconnected graphs are not considered as a valid output, we keep all edges of the pivot graph even its support is below the threshold. On the other hand, for the graph prediction problem, where a graph is only considered a valid graph if it does not have multiple edges between two vertices and multiple labels for any vertex, we remove all candidate labels for vertices and edges except the one with the highest number of votes.

Assume that the resulting ensemble graph that is created by using g_i as the pivot graph is denoted as g_i^e. The final ensemble graph g^e is chosen among the set of graphs $g_1^e, g_2^e, \cdots, g_m^e$ as the one with the largest total support.

Algorithm 1: Graph ensemble with the Graphene algorithm.

Input: a set of graphs $G = \{g_1, g_2, \cdots, g_m\}$ and the support threshold θ
Output: an ensemble graph g^e
Algorithm: Graphene(G, θ)
for $i \leftarrow 1$ **to** m **do**
  $g_{pivot} \leftarrow g_i$
  $V \leftarrow$ Initialise(g_{pivot})
  for $j \leftarrow 1$ **to** m **do**
    if $j \neq i$ **then**
      $V \leftarrow V \cup$ getVote($\phi(g_{pivot}, g_j)$)
  end
  $g_i^e \leftarrow Filter(V, \theta)$
end
$g^e \leftarrow$ the graph with the largest support among $g_1^e, \cdots, g_m^e$
Return g^e

Recall that $\phi(g_{pivot}, g_i)$ finds the best vertex match between two graphs. In general, the given task is computationally intractable. However, for labeled graphs like AMR a heuristic was proposed [12] to approximate the best match by a hill-climbing algorithm. It first starts with the candidate with labels that are mostly matched. The initial match is modified iteratively to optimize the total number of matches with a predefined number of iterations (default value set to 5). This algorithm is very efficient and effective, it was used to calculate the Smatch score in [12] so we reuse the same implementation to approximate $\phi(g_{pivot}, g_i)$ (report on average running time can be found in the supplementary materials).

4 Experiments

We compare our *Graphene* algorithm against four previous state-of-the-art models on different benchmark datasets. Below we describe our experimental settings.

4.1 Experimental settings

4.1.1 Model settings

SPRING The SPRING model, presented in [6], tackles Text-to-AMR and AMR-to-Text as a symmetric transduction task, showing that with a pretrained encoder-decoder model, it is possible to obtain state-of-the-art performances in both tasks using a simple seq2seq approach. In our experiments, we used the pretrained models provided in [6][1]. In addition, we trained 3 more models using different random seeds following the same setup described in [6]. Blink [19] was used to add wiki tags to the predicted AMR graphs as a post-processing step.

T5 The T5 model, presented in [21], introduces a unified framework that models a wide range of NLP tasks as a text-to-text problem. We follow the same idea proposed in [27] to train a model to transfer a text to a linearized AMR graph based on T5-Large. The data is preprocessed by linearization and removing wiki tags using the script provided in [2]. In addition to the main task, we added a new task that takes as input a sentence and predicts the concatenation of word senses and arguments provided in the English Web Treebank dataset [3]. The model is trained with 30 epochs. We use ADAM optimization with learning rate of 1e-4 and a mini-batch size of 4. Blink [19] was used to add wiki tags to the predicted AMR as a post-processing step.

APT [28] proposed a transition-based AMR parser[2] based on Transformer [26]. It combines hard-attention over sentences with a target side action pointer mechanism to decouple source tokens from node representations. For model training in our experiments, we use the setup described in [28] and added 70K model-annotated silver data sentences to the training data, which was created from the 85K sentence set in [18] with self-learning described in the paper.

Cai&Lam The model proposed in [9] treats AMR parsing as a series of dual decisions (i.e., *which parts of the sequence to abstract*, and *where in the graph to construct*) on the input sequence and constructs the AMR graph incrementally. Following [9], we use Stanford CoreNLP[3] for tokenization, lemmatization, part-of-speech tagging, and named entity recognition. We apply the pretrained model provided by the authors[4] to all testing datasets and follow the same pre-processing and post-processing steps for graph re-categorization.

Graphene (our algorithm) The only hyperparameter of the Graphene algorithm is the threshold θ. A popular practice for ensemble methods via voting strategy [14] is to consider the labels that get at least 50% of the total number of votes, therefore we set the threshold θ such that $\frac{\theta}{m} \geq 0.5$ (where m is the number of models in the ensemble). In all experiments, we used a Tesla GPU V100 for model training and used 8 CPUs for making an ensemble.

[1] Available for download at `https://github.com/SapienzaNLP/spring`

[2] Available under https://github.com/IBM/ transition-amr-parser.

[3] Available at `https://github.com/stanfordnlp/stanza/`

[4] The model "AMR2.0+BERT+GR" can be downloaded from `https://github.com/jcyk/AMR-gs`

4.1.2 Evaluation

We use the script[5] provided in [12] to calculate the Smatch score [10], the most relevant metric for measuring the similarity between the predictions and the gold AMR graphs. The overall Smatch score can be broken down into different dimensions, including the followings sub-metrics:

- Unlabeled (Unl.): Smatch score after removing all edge labels
- No WSD (NWSD): Smatch score while ignoring Propbank senses
- NE: F-score on the named entity recognition (:name roles)
- Wikification (Wiki.): F-score on the wikification (:wiki roles)
- Negations (Neg.): F-score on the negation detection (:polarity roles)
- Concepts (Con.): F-score on the concept identification task
- Reentrancy (Reen.): Smatch computed on reentrant edges only
- SRL: Smatch computed on :ARG-i roles only

4.1.3 Datasets

Similarly to [6], we use five standard benchmark datasets to evaluate our approach [1]. Table 1 shows the statistics of the datasets. AMR 2.0 and AMR 3.0 are divided into train, development and testing sets and we use them for *in-distribution* evaluation in Subsection 4.2. Furthermore, the models trained on AMR 2.0 training data are used to evaluate *out-of-distribution* prediction on the BIO, the LP and the New3 dataset (See Subsection 4.3).

Table 1: Benchmark datasets. All instances of BIO, LP, and New3 are used to test models in out-of-distribution evaluation. For AMR 2.0 and 3.0, the models are trained on the training dataset, validated on the development dataset. We report results on test sets in the in-distribution evaluation.

Datasets	**AMR 2.0**	**AMR 3.0**	**BIO**	**Little Prince (LP)**	**New3**
Training	36,521	55,635	n/a	n/a	n/a
Dev	1,368	1,722	n/a	n/a	n/a
Test	1,371	1,898	6,952	1,562	527

4.2 In-distribution evaluation

In the same spirit of [6], we evaluate the approaches when training and test data belong to the same domain. Table 2 shows the results of the models on the test split of the AMR 2.0 and AMR 3.0 datasets. The metrics reported for SPRING correspond to the model with the highest Smatch score among the 4 models used (the checkpoint plus the 3 models with different random seeds).

For the ensemble approach, we report the result when Graphene is used to make an ensemble of four SPRING checkpoints, denoted as *Graphene 4S*. The ensemble of all the models including four SPRING checkpoints, APT, T5, and Cai&Lam is denoted as *Graphene All*. For the AMR 3.0 dataset, the Cai&Lam model is not available so the reported result corresponds to an ensemble of all six models.

We can see that Graphene successfully leverages the strength of all the models and provides better prediction both in terms of the overall Smatch score and sub-metrics. In both datasets, we achieve the state-of-the-art results with performance gain of 1.5 and 1.3 Smatch points in AMR 2.0 and AMR 3.0 respectively.

Table 2 shows that by combining predictions from four checkpoints of the SPRING model, Graphene 4S provides better results than any individual model. The result is improved further when increasing the number of ensemble models, indeed *Graphene All* improves *Graphene 4S* further and outperforms the individual model in terms of the overall Smatch score.

[5] `https://github.com/mdtux89/amr-evaluation`

Table 2: Results on the test splits of the AMR 2.0 and AMR 3.0 dataset.

Models	Smatch	Unl.	NWSD	Con.	NE	Neg.	Wiki.	Reen.	SRL
AMR 2.0									
SPRING	84.08	87.25	84.57	89.98	90.77	72.65	79.81	74.30	82.90
APT	82.6	86.05	83.10	89.48	90.20	67.27	76.29	73.22	82.01
T5	82.82	86.01	83.31	89.85	90.65	74.43	74.79	72.42	82.01
Cai&Lam	80.06	83.49	80.05	87.39	82.25	**78.09**	**82.95**	66.45	77.37
Graphene 4S	84.68	87.88	85.20	90.65	92.16	75.22	81.06	71.45	83.48
Graphene All	**85.59**	**88.42**	**86.06**	**91.20**	**92.59**	77.01	81.80	**74.35**	**84.15**
AMR 3.0									
SPRING	83.14	86.29	83.60	89.38	87.80	72.94	78.77	**73.34**	81.98
APT	80.45	83.87	80.96	88.38	86.82	68.69	74.44	70.77	80.16
T5	82.07	85.12	82.55	89.03	86.99	72.59	71.07	72.18	81.18
Graphene 4S	83.60	86.75	84.05	90.02	88.33	74.60	79.52	70.11	82.38
Graphene All	**84.43**	**87.28**	**84.77**	**90.54**	**88.71**	**74.89**	**79.73**	71.95	**83.18**

Table 3: Results of out-of-distribution evaluation on the BIO, New3, and Little Prince dataset.

Models	Smatch	Unl.	NWSD	Con.	NE	Neg.	Wiki.	Reen.	SRL
BIO									
SPRING	60.52	65.34	61.40	67.76	**33.92**	65.68	3.80	51.20	62.86
APT	51.23	56.27	51.81	58.22	15.68	52.91	3.62	43.53	54.24
T5	58.89	63.89	59.66	66.63	30.42	65.11	2.46	48.57	61.47
Cai&Lam	42.26	49.84	42.88	47.10	5.19	51.42	**7.32**	39.21	51.00
Graphene 4S	61.5	66.23	62.28	68.52	33.17	68.04	4.46	50.39	63.71
Graphene All	**62.27**	**66.86**	**63.06**	**68.62**	32.30	**69.81**	4.54	**52.07**	**64.15**
New3									
SPRING	74.64	78.96	75.19	82.38	67.52	67.48	67.20	66.45	75.62
APT	71.10	75.79	71.59	80.34	65.65	67.08	57.14	63.09	73.36
T5	73.01	77.32	73.70	82.65	68.24	64.20	56.42	64.69	75.08
Cai&Lam	60.81	66.08	61.33	72.79	45.60	59.57	46.39	57.62	68.85
Graphene 4S	74.85	79.34	75.37	82.58	**69.98**	69.51	**68.34**	63.59	76.26
Graphene All	**75.57**	**79.60**	**76.12**	**83.08**	68.40	**69.62**	67.87	**67.13**	**76.92**
Little Prince									
SPRING	79.60	83.51	80.65	85.77	60.53	78.06	60.53	70.40	80.02
APT	77.00	81.32	77.92	84.37	**65.15**	74.92	**69.70**	64.99	77.52
T5	79.5	83.25	80.38	86.17	58.06	79.01	59.35	68.83	80.54
Cai&Lam	72.76	77.15	73.82	81.24	22.73	63.53	31.50	60.90	74.30
Graphene 4S	79.72	83.60	80.70	85.97	61.54	81.16	60.65	66.50	80.33
Graphene All	**80.48**	**84.01**	**81.43**	**86.65**	63.16	**82.24**	63.45	**69.84**	**81.14**

4.3 Out-of-distribution evaluation

In contrast to in-distribution evaluation, we use the models trained with AMR 2.0 data to collect AMR predictions for test datasets in the domains that differ from the AMR 2.0 dataset. The purpose of the experiment is to evaluate the ensemble approach under out-of-distribution settings.

Table 3 shows the results of our experiments. Similar to the in-distribution experiments, the *Graphene 4S* algorithm achieves better results than other individual models, while the *Graphene All* approach improves the given results further. We achieve the new state-of-the-art results in these benchmark datasets (under out-of-distribution settings). This result has an important practical implication because in practice it is very common not to have labeled AMR data for domain-specific texts. After all, the labeling task is very time-demanding. Using the proposed ensemble methods we can achieve better results with domain-specific data not included in the training sets.

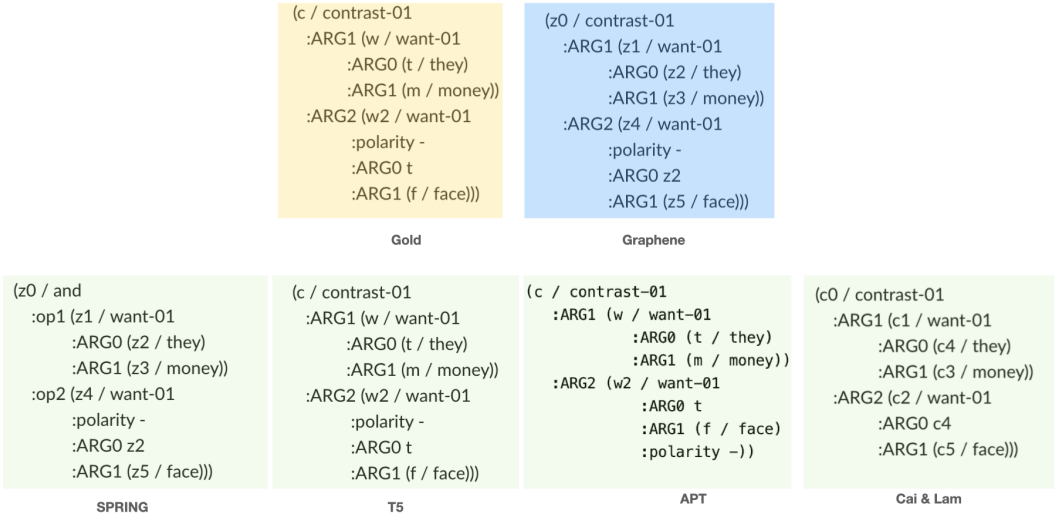

Figure 3: The gold AMR and the ensemble AMR graph of SPRING, T5, APT and Cai&Lam using the Graphene algorithm for the sentence *"They want money, not the face"*.

4.4 How the ensemble algorithm works

We explore a few examples to demonstrate the reason why the ensemble method works. Figure 3 shows a sentence with a gold AMR in Penman format. On the same figure we can see a list of AMRs corresponding to the prediction of SPRING [6], T5 [21], APT [28], Cai and Lam [9] parser and the ensemble graph given by Graphene.

In this particular example, with the sentence *"They want money, not the face"*, the AMR prediction from SPRING is inaccurate. Graphene corrects the prediction thanks to the votes given from the other models. In particular, the label and of the root node z_0 of SPRING prediction was corrected to $contrast-01$ because T5, APT and Cai&Lam parsers all vote for $contrast-01$. On the other hand, the labels $:op1$ and $:op2$ of the edges (z_0, z_1) and (z_0, z_4) were modified to have the correct labels $:ARG1$ and $:ARG2$ accordingly thanks to the votes from the other models.

We can also see that even though the Cai&Lam method misses polarity prediction but since the other models predict polarity correctly the ensemble prediction catches it. Putting everything together, the prediction from Graphene perfectly matches with the gold AMR graph in this example.

Table 4: The average total support and Smatch score of the prediction from SPRING, Graphene with SPRING as a pivot and Graphene considering every prediction as a pivot respectively. Support is highly correlated with Smatch score.

	SPRING		**Graphene SPRING pivot**		**Graphene**	
Data/metrics	Avg. support	Smatch	Avg. support	Smatch	Avg. support	Smatch
AMR 2.0	170.15	84.08	172.7	84.70	**175.73**	**85.59**
AMR 3.0	136.9	83.60	139.42	83.63	**142.07**	**84.43**
BIO	166.86	60.52	169.97	61.57	**179.38**	**62.27**
Little Prince	69.33	79.60	70.97	80.03	**72.64**	**80.48**
New3	118.27	74.64	120.85	74.79	**123.62**	**75.57**

The Graphene algorithm searches for the graph that has the largest support from all individual graphs. One question that arises from this is whether the support is correlated with the accuracy of AMR parsing. Table 4 shows the support and the Smatch score of three models in the standard benchmark datasets. The first model is SPRING, while the second one denoted as *Graphene SPRING pivot* starts with a SPRING prediction as a pivot and corrects the prediction using votes from other models. The last model corresponds to the Graphene algorithm that polls the votes while considering every prediction as a pivot for correction and selecting the best one.

Since the Graphene algorithm looks for the best pivot to have better-supported ensemble graphs, the total supports of the Graphene predictions are larger than the Graphene SPRING pivot predictions. From the table, we can also see that the total support is highly correlated to the Smatch score. Namely, Graphene has higher support in all the benchmark datasets and a higher Smatch score than Graphene SPRING pivot. This experiment suggests that by optimizing the total support we can obtain the ensemble graphs with higher AMR Smatch score.

5 Related work

Ensemble learning. Ensemble learning is a popular machine learning approach that combines predictions from different learners to make a more robust and more accurate prediction. Many ensembling approaches have been proposed, such as bagging [7] or boosting [23], the winning solutions in many machine learning competitions [11]. These methods are proposed mainly for regression or classification problems. Recently, structure output prediction emerges as an important problem thanks to the advance in deep learning research. To apply popular ensembling techniques such as bagging or boosting, it is important to study the ensemble method for combining structure predictions.

Ensemble structure prediction. Previous studies have explored various ensemble learning approaches for dependency and constituent parsing: [22] proposes a reparsing framework that takes the output from different parsers and maximizes the number of votes for a well-formed dependency or constituent structure; [17] uses minimum Bayes risk inference to build a consensus dependency parser from an ensemble of independently trained greedy LSTM transition-based parsers with different random initializations. Note that a syntactic tree is a special graph structure in which nodes for a sentence from different parsers are roughly the same. In contrast, we propose an approach to ensemble graph predictions in which both graph nodes and edges can be different among base predictions.

Ensemble methods for AMR parsing. Parsing text to AMR is an important research problem. State-of-the-art approaches in AMR parsing are divided into three categories. Sequence to sequence models [6, 16, 25, 27] consider the AMR parsing as a machine translation problem that translates texts to AMR graphs. The transition-based method [28] predicts a sequence of actions given the input text, and then the action sequence is turned into an AMR graph using an oracle decoder. Lastly, graph-based methods [9] directly construct the AMR graphs from textual data. All these methods are complementary to each other and thus ensemble methods can leverage the strength of these methods to create a better prediction, as demonstrated in this paper. Ensemble of AMR predictions from a single type of model is studied in [28] where the authors demonstrated that by combining predictions from three different model's checkpoints they gain performance improvement in the final prediction. However, ensemble in sequential decoding requires that all predictions are from the same type of models. It is not applicable for cases when the predictions are from different types of models such as seq2seq, transition-based or graph-based models. In contrast to that approach, our algorithm is model-agnostic, i.e. it can combine predictions from different models. In our experiments, we have demonstrated the benefit of combining predictions from different models, with additional gain in performance compared to the ensemble of predictions from a single model's checkpoints.

6 Conclusions and future work

In this paper, we formulate the graph ensemble problem, study its computational intractability, and provide an algorithm for constructing graph ensemble prediction. We validate our approach with AMR parsing problems. The experimental results show that the proposed approach outperforms the previous state-of-the-art AMR parsers and achieves new state-of-the-art results in five different benchmark datasets. We demonstrate that the ensemble not only works well for in-distribution but also for out-of-distribution evaluation. This result has a significant practical impact, especially when applying the proposed method to domain-specific texts where training data is not available. Moreover, our approach is model-agnostic, which means that it can be used to combine predictions from different models without the requirement of having an access to model training. In general, our approach provides the basis for graph ensemble, studying classical ensemble techniques such as bagging, boosting, or stacking for graph ensemble is a promising future research direction that is worth considering to improve the results further.

References

[1] Amr benchmark datasets. https://amr.isi.edu/download.html.

[2] Amrlib data preprocessing https://github.com/bjascob/amrlib/tree/master/scripts.

[3] English web treebank dataset: https://catalog.ldc.upenn.edu/ldc2012t13.

[4] Laura Bahiense, Gordana Manić, Breno Piva, and Cid C De Souza. The maximum common edge subgraph problem: A polyhedral investigation. *Discrete Applied Mathematics*, 160(18):2523–2541, 2012.

[5] Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. Abstract meaning representation for sembanking. In *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186, 2013.

[6] Michele Bevilacqua, Rexhina Blloshmi, and Roberto Navigli. One SPRING to rule them both: Symmetric AMR semantic parsing and generation without a complex pipeline. In *Proceedings of AAAI*, 2021.

[7] Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.

[8] Deng Cai and Wai Lam. AMR parsing via graph-sequence iterative inference. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1290–1301, Online, July 2020. Association for Computational Linguistics.

[9] Deng Cai and Wai Lam. AMR parsing via graph-sequence iterative inference. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1290–1301, Online, July 2020. Association for Computational Linguistics.

[10] Shu Cai and Kevin Knight. Smatch: an evaluation metric for semantic feature structures. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 748–752, Sofia, Bulgaria, August 2013. Association for Computational Linguistics.

[11] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *KDD*.

[12] Marco Damonte, Shay B. Cohen, and Giorgio Satta. An incremental parser for Abstract Meaning Representation. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 536–546, Valencia, Spain, April 2017. Association for Computational Linguistics.

[13] Pedro Domingos. A unified bias-variance decomposition. In *Proceedings of 17th International Conference on Machine Learning*, pages 231–238, 2000.

[14] Xibin Dong, Zhiwen Yu, Wenming Cao, Yifan Shi, and Qianli Ma. A survey on ensemble learning. *Frontiers of Computer Science*, 14(2):241–258, 2020.

[15] Pavan Kapanipathi, Ibrahim Abdelaziz, Srinivas Ravishankar, Salim Roukos, Alexander Gray, Ramon Astudillo, Maria Chang, Cristina Cornelio, Saswati Dana, Achille Fokoue, et al. Question answering over knowledge bases by leveraging semantic parsing and neuro-symbolic reasoning. *arXiv preprint arXiv:2012.01707*, 2020.

[16] Ioannis Konstas, Srinivasan Iyer, Mark Yatskar, Yejin Choi, and Luke Zettlemoyer. Neural amr: Sequence-to-sequence models for parsing and generation. *arXiv preprint arXiv:1704.08381*, 2017.

[17] Adhiguna Kuncoro, Miguel Ballesteros, Lingpeng Kong, Chris Dyer, and Noah A. Smith. Distilling an ensemble of greedy dependency parsers into one MST parser. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1744–1753, Austin, Texas, November 2016. Association for Computational Linguistics.

[18] Young-Suk Lee, Ramon Fernandez Astudillo, Tahira Naseem, Revanth Gangi Reddy, Radu Florian, and Salim Roukos. Pushing the limits of amr parsing with self-learning. In *Findings of EMNLP 2020*, pages 3208–3214, 2020.

[19] Belinda Z. Li, Sewon Min, Srinivasan Iyer, Yashar Mehdad, and Wen-tau Yih. Efficient one-pass end-to-end entity linking for questions. In *EMNLP*, 2020.

[20] Jungwoo Lim, Dongsuk Oh, Yoonna Jang, Kisu Yang, and Heuiseok Lim. I know what you asked: Graph path learning using AMR for commonsense reasoning. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2459–2471, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics.

[21] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.

[22] Kenji Sagae and Alon Lavie. Parser combination by reparsing. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 129–132, New York City, USA, June 2006. Association for Computational Linguistics.

[23] Robert E Schapire and Yoav Freund. Boosting: Foundations and algorithms. *Kybernetes*, 2013.

[24] Giorgio Valentini and Thomas G Dietterich. Bias-variance analysis of support vector machines for the development of svm-based ensemble methods. *Journal of Machine Learning Research*, 5(Jul):725–775, 2004.

[25] Rik Van Noord and Johan Bos. Neural semantic parsing by character-based translation: Experiments with abstract meaning representations. *arXiv preprint arXiv:1705.09980*, 2017.

[26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, pages 5998–6008, 2017.

[27] Dongqin Xu, Junhui Li, Muhua Zhu, Min Zhang, and Guodong Zhou. Improving amr parsing with sequence-to-sequence pre-training. *arXiv preprint arXiv:2010.01771*, 2020.

[28] Jiawei Zhou, Tahira Naseem, Ramon Fernandez Astudillo, and Radu Florian. Amr parsing with action-pointer transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL 2021)*, pages 5585–5598, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [No]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code will be attached together with the supplementary material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Please see Section 4.1
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Our ensemble algorithm is deterministic.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section 4.1
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We provided links to datasets and github repositories
   (b) Did you mention the license of the assets? [N/A] Licence information is available on github repositories.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We provided source code in the supplemental materials.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A] We used public data.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] We used public data.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]