

Structured Attention for Unsupervised Dialogue Structure Induction

Liang Qiu¹, Yizhou Zhao¹, Weiyan Shi², Yuan Liang³, Feng Shi¹,
Tao Yuan¹, Zhou Yu², Song-Chun Zhu¹

¹UCLA Center for Vision, Cognition, Learning, and Autonomy

²University of California, Davis

³University of California, Los Angeles

{liangqiu, yizhouzhao, liangyuandg, shifeng, taoyuan}@ucla.edu

{wyshi, joyu}@ucdavis.edu

sczhu@stat.ucla.edu

Abstract

Inducing a meaningful structural representation from one or a set of dialogues is a crucial but challenging task in computational linguistics. Advancement made in this area is critical for dialogue system design and discourse analysis. It can also be extended to solve grammatical inference. In this work, we propose to incorporate structured attention layers into a Variational Recurrent Neural Network (VRNN) model with discrete latent states to learn dialogue structure in an unsupervised fashion. Compared to a vanilla VRNN, structured attention enables a model to focus on different parts of the source sentence embeddings while enforcing a structural inductive bias. Experiments show that on two-party dialogue datasets, VRNN with structured attention learns semantic structures that are similar to templates used to generate this dialogue corpus. While on multi-party dialogue datasets, our model learns an interactive structure demonstrating its capability of distinguishing speakers or addressees, automatically disentangling dialogues without explicit human annotation.¹

1 Introduction

Grammatical induction for capturing a structural representation of knowledge has been studied for some time (De la Higuera, 2010). Given the achievement in related areas like learning *Hidden Markov* acoustic models in speech recognition (Bahl et al., 1986) and sentence dependency parsing in language understanding (Covington, 2001), our work aims to explore a more sophisticated topic: learning structures in dialogues. Figure 1 shows the underlying semantic structure of conversations about bus information request from SimDial dataset (Zhao and Eskenazi, 2018), with one

¹The code is released at <https://github.com/Liang-Qiu/SVRNN-dialogues>.

example dialogue as shown in Table 1. Another interesting type of dialogue structure is the interactive structure in multi-party dialogues. Figure 2 illustrates the interactive structure we learned from a dialogue sample in Ubuntu Chat Corpus (Lowe et al., 2015). Each node represents an utterance from different speakers in the dialogue with darker linkages represent stronger dependency relations between utterances. When speaker/addressee information is unavailable in the corpus, learning such a structure allows disentangling the conversation (Serban and Pineau, 2015) and estimating the speaker labels. Discovering dialogue structures is crucial for various areas in computational linguistics, such as dialogue system building (Young, 2006), discourse analysis (Grosz and Sidner, 1986), and dialogue summarization (Murray et al., 2005; Liu et al., 2010). Through looking into this topic, we can further improve the capability of machines to learn more generalized, interpretable knowledge representation from data.

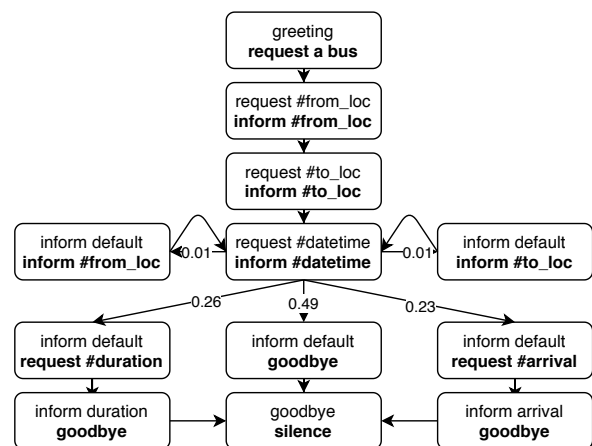


Figure 1: Original dialogue structure of the bus information request domain in SimDial (Zhao and Eskenazi, 2018). User intents are marked in bold.

However, capturing structure from the conversation is still much under-explored. The complexity of dialogues could range from several-round task-oriented dialogues to tens-round multi-party chitchat. It is unclear that for these different categories of dialogues, what types of inductive biases or constraints we could add to reduce the search space. It also remains an unsolved question for formally evaluating the performance of dialogue structure induction algorithms. In this paper, we propose to use a combination of structured attention and unsupervised generative model to infer the latent structure in a dialogue.




Figure 2: Learned interactive structure from a multi-party dialogue sample in Ubuntu Chat Corpus (Uthus and Aha, 2013).

Specifically, instead of simply applying a softmax function on potentials between a decoder query and encoder hidden states, dynamic programming algorithms like *Forward-Backward* (Devijver, 1985) and *Inside-Outside* (Lari and Young, 1990) could be used to efficiently calculate marginal probabilities from pairwise potentials with a structural constraint. Through embedding such structured attention layers in a *Variational Recurrent Neural Network* (VRNN) model, we can learn latent structures in dialogues by jointly re-generating training dialogues. Such a process requires no human annotation and is useful for dialogue analysis. In addition, by selecting appropriate structural biases or constraints, we can learn not only semantic structures but also interactive structures. A linear *Conditional Random Field* (CRF) attention layer is used in two-party dialogues to discover semantic structures. A non-projective dependency tree attention layer is embedded to learn an interactive structure that could help identify speaker/addressee information in multi-party dialogues that have tangled conversation threads, such as forum discussions.

This paper makes the following contributions. We propose to incorporate a structured attention layer in VRNN to learn latent structures in dialogues. To our knowledge, no work connecting

structured attention with unsupervised dialogue structure learning has been done. We prove our proposed VRNN-LinearCRF learns better structures than the baseline VRNN on the SimDial dataset for semantic structure learning in two-party dialogues. For interactive structure learning in multi-party dialogues, we combine VRNN with a non-projective dependency tree attention layer. It achieves similar generation performance as the baseline GSN model (Hu et al., 2019) on Ubuntu Chat Corpus (Uthus and Aha, 2013; Lowe et al., 2015), while our model can identify the speaker/addressee information without trained on explicit labels. We release our code as well as the processed datasets to help stimulate related researches.

2 Related Work

Attention mechanism (Vaswani et al., 2017) has been widely adopted as a way for embedding categorical inference in neural networks for performance gain and interpretability (Jain and Wallace, 2019; Wiegrefe and Pinter, 2019). However, for many tasks, we want to model richer structural dependencies without abandoning end-to-end training. *Structured Attention Networks* (Kim et al., 2017) can extend attention beyond the standard soft-selection approach by attending to partial segments or subtrees. People have proven its effectiveness on a variety of synthetic and real tasks: tree transduction, neural machine translation, question answering, and natural language inference (Rush, 2020). In this paper, we propose to utilize structured attention to explore dialogue structures. Specifically, we work on two types of dialogue structures, semantic structures (dialogue intent transitions), and interactive structures (addressee/speaker changes).

Semantic structures have been studied extensively. Some previous works, such as (Jurafsky, 1997), learned semantic structures relying on human annotations, while such annotations are costly and can vary in quality. Other unsupervised studies used *Hidden Markov Model* (HMM) (Chotimongkol, 2008; Ritter et al., 2010; Zhai and Williams, 2014). Recently, *Variational Autoencoders* (VAEs) (Kingma and Welling, 2013) and their recurrent version, *Variational Recurrent Neural Networks* (VRNNs) (Chung et al., 2015), connects neural networks and traditional Bayes methods. Because VRNNs apply a point-wise non-linearity to the output at every timestamp, they are

also more suitable to model highly non-linear dynamics over the simpler dynamic Bayesian network models. Serban et al. (2017) proposed the VHRED model by combining the idea of VRNNs and *Hierarchical Recurrent Encoder-Decoder* (HRED) (Sordani et al., 2015) for dialogue generation. Similarly, Zhao et al. (2018) proposed to use VAEs to learn discrete sentence representations. Shi et al. (2019) used two variants of VRNNs to learn the dialogue semantic structures and discussed how to use learned structure to improve reinforcement learning-based dialogue systems. But none of the previous work has tried to incorporate structured attention in VRNNs to learn dialogue structure.

Compared to semantic structures, the interactive structure of dialogues is not clearly defined. Elsner and Charniak (2008) initiated some work about dialogue disentanglement, which is defined as dividing a transcript into a set of distinct conversations. Serban and Pineau (2015) tested standard RNN and its conditional variant for turn taking and speaker identification. Both of the tasks are highly related to understanding the interactive structure but not identical. Our task, different from both of them, aims to construct an utterance dependency tree to represent a multi-party dialogue’s turn taking. The tree can not only be used to disentangle the conversations but also label each utterance’s speakers and addressees. We compare our model with *Graph Structured Network* (GSN), recently proposed by Hu et al. (2019). GSN builds a conversation graph utilizing explicit speaker/addressee information in Ubuntu Chat Corpus (Uthus and Aha, 2013) to improve the dialogue generation performance. Our model shows similar generation performance as them while demonstrating its capability of learning the utterance dependency tree.

3 Problem Formulations

We discuss the semantic and interactive dialogue structure learning separately. In task-oriented two-party dialogues (between system and user), we want to discover a semantic probabilistic grammar shared by dialogues in the same domain. While for multi-party dialogues, e.g., conversations in a chatroom, which may have multiple conversations occur simultaneously, we are more interested in finding an interactive structure that could help disentangle the conversation and identify the speakers/addressees. Our method of structure learning is flexible to handle both problems with the formula-

tions as shown below.

For semantic dialogue structure learning, we formulate the problem as labeling the dialogue with a sequence of latent states. Each conversational exchange x_i (a pair of system and user utterances at time step i) belongs to a latent state z_i , which has an effect on the future latent states and the words the interlocutors produce. The latent dialogue state is defined to be discrete, i.e., $z_i \in \{1, 2, \dots, N\}$, where N is the number of states predefined from experience. Our goal is to generate the current sentence pair x_i that maximizes the conditional likelihood of x_i given the dialogue history while jointly learning a latent state sequence $\mathbf{z} = [z_1, z_2, \dots, z_n]$:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmax}} \sum_{i=1}^{|\mathbf{x}|} \log(P(\mathbf{z}_{<i} | \mathbf{x}_{<i}) P(x_i | \mathbf{z}_{<i})). \quad (1)$$

Then, we can induce a probabilistic dialogue grammar by estimating the state transition probabilities through maximizing the likelihood of the parsed latent state sequences.

A multi-party dialogue session can be formulated as an utterance-level dependency tree $\mathbf{T}(V, E)$, where V is the set of nodes encoding the utterances, $E = \{e_{i,j}\}_{i < j}^m \in \{0, 1\}$ indicates whether utterance i is the parent of utterance j , and m is the maximum number of possible edges.

$$\begin{aligned} \hat{\mathbf{x}} &= \underset{\mathbf{x}}{\operatorname{argmax}} \sum_{i=1}^{|\mathbf{x}|} \log(P(\mathbf{T} | \mathbf{x}_{<i}) P(x_i | \mathbf{T})) \\ &= \underset{\mathbf{x}}{\operatorname{argmax}} \sum_{i=1}^{|\mathbf{x}|} \log\left(\prod_{j < k}^{i-1} P(e_{j,k} = 1 | \mathbf{x}_{<i}) \cdot P(x_i | \mathbf{T})\right) \\ &= \underset{\mathbf{x}}{\operatorname{argmax}} \left[\sum_{i=1}^{|\mathbf{x}|} \sum_{j < k}^{i-1} \log(P(e_{j,k} = 1 | \mathbf{x}_{<i})) + \sum_{i=1}^{|\mathbf{x}|} \log(P(x_i | \mathbf{T})) \right] \end{aligned} \quad (2)$$

Each path of the dependency tree represents a thread in the multi-party conversation in chronological order. Our goal is to generate the response $\hat{\mathbf{x}}$ that maximizes the conditional likelihood of the response given the dialogue history while jointly learning a latent utterance dependency tree as shown in Equation 2. The conditional likelihood is factorized into two parts, representing the encod-

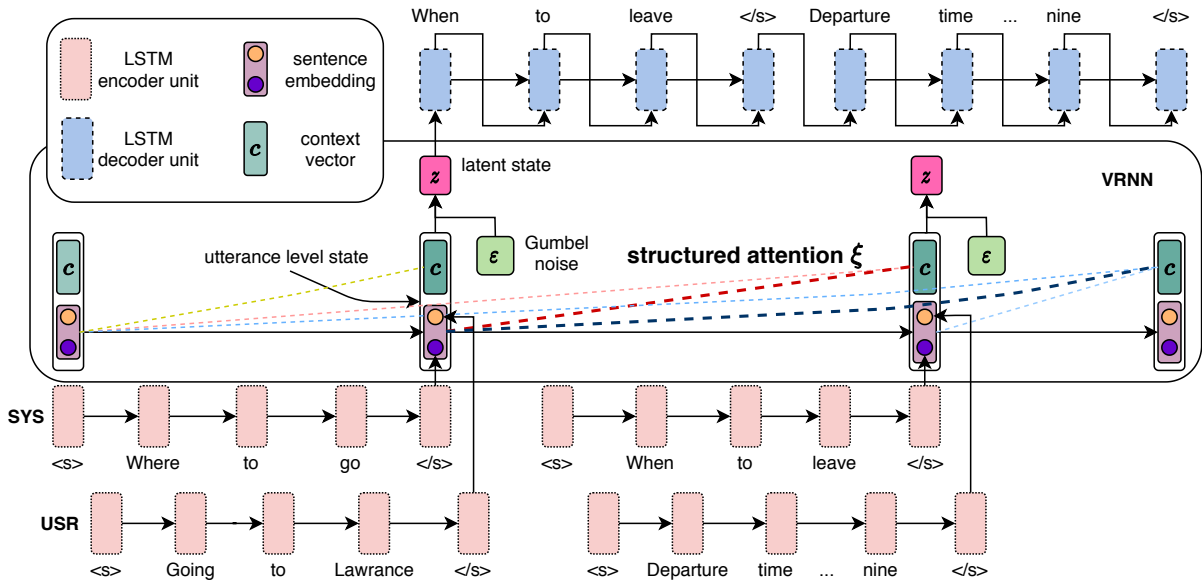


Figure 3: Structured-Attention Variational Recurrent Neural Network (SVRNN)

ing and decoding processes respectively. We can further reason about the speaker/addressee labels or disentangle the conversation by clustering the utterances from the learned tree.

4 Variational Recurrent Neural Network with Structured Attention

The overall architecture of Structured-Attention Variational Recurrent Neural Network (SVRNN) is illustrated in Figure 3. The LSTM (Hochreiter and Schmidhuber, 2001) word-level encoder marked in pink encodes each utterance into a sentence embedding. Then an utterance-level encoder VRNN with different structured attention layers encodes the dialogue history into a latent state z . A decoder marked in blue will decode the next utterances from the latent state. We describe more details about the key components of our model in the following subsections.

4.1 Variational Recurrent Neural Network

The pursuit of using an autoencoder like *Variational Recurrent Neural Network* (VRNN) is to compress the essential information of the dialogue history into a lower-dimensional latent code. The latent code z is a random vector sampled from a prior $p(z)$ and the data generation model is described by $p(x|z)$. The VRNN contains a *Variational Autoencoder* (VAE) at each time step. The VAE consists of an encoder $q_\lambda(z|x)$ for approximating the posterior $p(z|x)$, and a decoder $p_\theta(x|z)$

for representing the distribution $p(x|z)$. The variational inference attains its maximum likelihood by maximizing *evidence lower bound* (ELBO):

$$\mathbb{E} [\log p_\theta(x|z)] - \text{KL} (q_\lambda(z|x) \| p(z)) \leq \log p(x). \quad (3)$$

For sequential data, the parameterization of the generative model is factorized by the posterior $p(z_t|x_{<t}, z_{<t})$ and the generative model $p(x_t|z_{\leq t}, x_{<t})$, i.e.,

$$p(x \leq T, z \leq T) = \prod_{t=1}^T [p(x_t|z_{\leq t}, x_{<t}) \cdot p(z_t|x_{<t}, z_{<t})]. \quad (4)$$

The learning objective function becomes maximizing the ELBO for all time steps

$$\mathbb{E} \left[\sum_{t=1}^T (-\text{KL} (q(z_t|x_{\leq t}, z_{<t}) \| p(z_t|x_{<t}, z_{<t})) + \log p(x_t|z_{\leq t}, x_{<t})) \right]. \quad (5)$$

In addition, to mitigate the *vanishing latent variable problem* in VAE, we incorporate Bag-of-Words (BOW) loss and Batch Prior Regularization (BPR) (Zhao et al., 2017) with a tunable weight λ . By adjusting the λ , the VRNN based models can achieve a balance between clustering the utterance surface formats and attention on the context.

4.2 Linear CRF Attention

As we formulate the semantic structure learning in two-party dialogues as a state tagging problem, we find it suitable to use a linear-chain *Conditional Random Field* (CRF) attention layer with VRNN. Define ξ to be a random vector $\xi = [\xi_1, \dots, \xi_n]$ with $\xi_i \in \{0, 1\}$. n is the number of utterances in a dialogue. The context vector \mathbf{c}_j given the current sentence hidden state \mathbf{h}_j and hidden state history \mathbf{h} can thus be written as:

$$\mathbf{c}_j = \sum_{i=1}^{j-1} p(\xi_i = 1 | \mathbf{h}, \mathbf{h}_j) \mathbf{h}_i. \quad (6)$$

We model the distribution over the latent variable ξ with a linear-chain CRF with pairwise edges,

$$p(\xi_1, \dots, \xi_n | \mathbf{h}, \mathbf{h}_j) = \text{softmax} \left(\sum_{i=1}^{j-2} \theta_{i,i+1}(\xi_i, \xi_{i+1}) \right), \quad (7)$$

where $\theta_{i,i+1}(k, l)$ is the pairwise potential for $\xi_i = k$ and $\xi_{i+1} = l$. The attention layer is a two-state CRF where the unary potentials at the j -th dialogue turn are:

$$\theta_i(k) = \begin{cases} \mathbf{h}_i \mathbf{W}_1 \mathbf{h}_j, & k = 0 \\ \mathbf{h}_i \mathbf{W}_2 \mathbf{h}_j, & k = 1 \end{cases}, \quad (8)$$

where $[\mathbf{h}_1, \dots, \mathbf{h}_n]$ are utterance level hidden states and $\mathbf{W}_1, \mathbf{W}_2$ are parameters. The pairwise potentials can be parameterized as

$$\theta_{i,i+1}(\xi_i, \xi_{i+1}) = \theta_i(\xi_i) + \theta_{i+1}(\xi_{i+1}) + \mathbf{h}_i^\top \mathbf{h}_{i+1}. \quad (9)$$

The marginal distribution $p(\xi_i = 1 | x)$ can be calculated efficiently in linear-time for all i using message-passing, *i.e.*, the *forward-backward* shown in Algorithm 1.

\mathcal{C} denotes the state space and $\langle t \rangle$ is the special start/stop state. Typically the forward-backward with marginals is performed in the log-space semi-field $\mathbb{R} \cup \{\pm\infty\}$ with binary operations $\oplus = \text{logadd}$ and $\otimes = +$ for numerical precision. These marginals allow us to calculate the context vector. Crucially, the process from vector softmax to *forward-backward* algorithm is a series of differentiable steps, and we can compute the gradient of the marginals with respect to the potentials (Kim et al., 2017). This allows the linear CRF attention layer to be trained end-to-end as a part of the VRNN.

Algorithm 1: Forward-Backward for LinearCRF Attention

Input: potential θ
 $\alpha[0, \langle t \rangle] \leftarrow 0$
 $\beta[n+1, \langle t \rangle] \leftarrow 0$
for $i = 1, \dots, n; c \in \mathcal{C}$ **do**
 $\alpha[i, c] \leftarrow \bigoplus_y \alpha[i-1, y] \otimes \theta_{i-1,i}[y, c]$
end for
for $i = n, \dots, 1; c \in \mathcal{C}$ **do**
 $\beta[i, c] \leftarrow \bigoplus_y \beta[i+1, y] \otimes \theta_{i,i+1}[c, y]$
end for
 $A \leftarrow \alpha[n+1, \langle t \rangle]$
for $i = 1, \dots, n; c \in \mathcal{C}$ **do**
 $p(\xi_i = c | x) \leftarrow \exp(\alpha[i, c] \otimes \beta[i, c] \otimes -A)$
end for
return p

4.3 Non-projective Dependency Tree Attention

For interactive structure learning in multi-party dialogues, we want to learn an utterance dependency tree from each dialogue. Therefore, we propose to use a non-projective dependency tree attention layer with VRNN for this purpose. The potentials $\theta_{i,j}$, which reflect the score of selecting the i -th sentence being the parent of the j -th sentence (*i.e.*, $x_i \rightarrow x_j$), can be calculated by

$$\theta_{i,j} = \tanh(\mathbf{s}^\top \tanh(\mathbf{W}_1 \mathbf{h}_i + \mathbf{W}_2 \mathbf{h}_j + \mathbf{b})), \quad (10)$$

where $\mathbf{s}, \mathbf{b}, \mathbf{W}_1, \mathbf{W}_2$ are parameters, $\mathbf{h}_i, \mathbf{h}_j$ are sentence hidden states.

The probability of a parse tree ξ given the dialogue $x = [x_1, \dots, x_n]$ is,

$$p(\xi | x) = \text{softmax}(\mathbb{1}\{\xi \text{ is valid}\}) \cdot \sum_{i \neq j} \mathbb{1}\{\xi_{i,j} = 1\} \theta_{i,j}, \quad (11)$$

where the latent variable $\xi_{i,j} \in \{0, 1\}$ for all $i \neq j$ indicates that the i -th sentence is the parent of the j -th sentence; and $\mathbb{1}\{\xi \text{ is valid}\}$ is a special global constraint that rules out configurations of $\xi_{i,j}$'s that violate parsing constraints. In our case, we specify each sentence has one parent and that must precede the child sentence, *i.e.*,

$$\sum_{i=1}^n \xi_{i,j} = 1 \quad \xi_{i,j} = 0 (i \geq j). \quad (12)$$

It is possible to calculate the marginal probability of each edge $p(\xi_{i,j} = 1 | x)$ for all i, j in $O(n^3)$

time using the *inside-outside* algorithm with details explained in Appendix, which is a generalization of the *forward-backward* algorithm.

Then the soft-parent or the context vector of the j -th sentence is calculated using parsing marginals, *i.e.*,

$$\mathbf{c}_j = \sum_{i=1}^n p(\xi_{i,j} = 1 | \mathbf{h}, \mathbf{h}_j) \mathbf{h}_i. \quad (13)$$

The original embedding is concatenated with its context vector to form the new representation

$$\hat{\mathbf{h}}_j = [\mathbf{h}_j; \mathbf{c}_j]. \quad (14)$$

4.4 Decoder

In order to generate a response to an utterance i , the decoder calculates a distribution over the vocabulary then sequentially predicts word w_k using a *softmax* function:

$$\begin{aligned} p(\mathbf{w} | \hat{\mathbf{h}}) &= \prod_{k=1}^{|\mathbf{w}|} P(w_k | \hat{\mathbf{h}}, \mathbf{w}_{<k}) \\ &= \prod_{k=1}^{|\mathbf{w}|} \text{softmax}(MLP(\mathbf{h}_k^{dec}, \mathbf{c}_k^{dec})) \\ \mathbf{h}_0^{dec} &= \hat{\mathbf{h}}_i \\ \mathbf{h}_k^{dec} &= LSTM(\mathbf{h}_{k-1}^{dec}, MLP(\mathbf{e}_{w_{k-1}}, \mathbf{c}_{k-1}^{dec})) \\ \mathbf{c}_k^{dec} &= \sum_{j=1}^i \text{softmax}(\mathbf{h}_k^{dec} \mathbf{W}_a \hat{\mathbf{h}}_j) \hat{\mathbf{h}}_j, \end{aligned} \quad (15)$$

where $\hat{\mathbf{h}}_i$ is the hidden state for utterance i with structured attention, \mathbf{h}_k^{dec} is the hidden state of the decoder LSTM, $\mathbf{e}_{w_{k-1}}$ is the embedding of the predicted word at decoding time stamp $(k-1)$, and \mathbf{c}_k^{dec} is the attention-based context vector at decoding time stamp k . Note that the context vector here is calculated with the simple attention different from the structured attention we described before. \mathbf{W}_a is a matrix to learn the match degree of \mathbf{h}_k^{dec} and $\hat{\mathbf{h}}_j$.

5 Experiments

We incorporate structured attention in VRNNs to explore two types of dialogue structure, semantic structure, and interactive structure.

5.1 Semantic Structure Learning in Two-party Dialogues

5.1.1 Datasets

We test the VRNN with Linear CRF Attention on the SimDial dataset (Zhao and Eskenazi, 2018)

From	Utterance
SYS:	Ask me about bus information. How can I help?
USR:	Hi. I need a bus.
SYS:	Where do you want to take off?
USR:	Going to Lawrance.
SYS:	What time do you need the bus?
USR:	Departure time is 9.
SYS:	Bus 137 can take you there. What else can I do?
USR:	Not done yet. How long will it take?
SYS:	The ride is 45 minutes long. What else can I do?
USR:	No more questions. Thank you.
SYS:	Goodbye.

Table 1: An example two-party bus information request dialogue in SimDial (Zhao and Eskenazi, 2018).

of simulated conversations. Dialogues are generated for information requests in four domains: bus, restaurant, weather, and movie. Table 1 shows an example dialogue in bus schedule request domain. Despite significant variations exist between dialogues of the same domain, we aim to explore a shared semantic structure among each dialogue domain. We validate our algorithm on this simulated dataset because these dialogues are generated using pre-defined templates that make recovering ground truth structures much easier. One recovered ground truth structure with transition probabilities is shown in Figure 1. We have 800 dialogue samples for training, 100 for validation, and 100 for testing in each dialog domain. The length of the dialogues ranges from 6 to 13 utterances. The maximum length of an utterance is 33 words.

5.1.2 Evaluation Metrics

Since the number of states is unknown during unsupervised training, we set the state number empirically to 10. Then the learned structure is essentially a state transition matrix of size 10×10 . However, the original structure could be another state transition matrix of any size depending on the domain complexity. This makes the model evaluation on the ground truth a problem because it requires us to measure the difference between two state transition matrices of different sizes. To alleviate this problem, we define two metrics: Structure Euclidean Distance (SED) and Structure Cross-Entropy (SCE). We first estimate a probabilistic mapping P_{s_i, s'_i} between the learned states $\{s'_i, i = 1, 2, \dots, M\}$ and the true states $\{s_i, i = 1, 2, \dots, N\}$, through dividing the number of utterances that have the ground truth state s_i and learned state s'_i by number of utterances with the ground truth state s_i . And we let the reversed mapping probability $P_{s'_i, s_i}$ be the normalized transpose of

P_{s_i, s'_i} . Then SED and SCE are defined as:

$$T'_{s_a, s_b} = \sum_{i, j \in \{1, 2, \dots, M\}} P_{s_a, s'_i} \cdot T'_{s'_i, s'_j} \cdot P_{s'_j, s_b}$$

$$SED = \frac{1}{N} \sqrt{\sum_{a, b \in \{1, 2, \dots, N\}} (T'_{s_a, s_b} - T_{s_a, s_b})^2}$$

$$SCE = \frac{1}{N} \sum_{a, b \in \{1, 2, \dots, N\}} -\log(T'_{s_a, s_b}) T_{s_a, s_b},$$
(16)

where T'_{s_a, s_b} is the learned transition probability from state s_a to state s_b and T_{s_a, s_b} is the true transition probability.

5.1.3 Results and Analysis

We compare the proposed VRNN-LinearCRF against other unsupervised methods: K-means clustering, Hidden Markov Model, D-VRNN (Shi et al., 2019) and VRNN with vanilla attention. D-VRNN is similar to our work but without structured attention. We use a bidirectional LSTM with 300 hidden units as the sentence encoder and a forward LSTM for decoding. 300-dimensional word embeddings are initialized with GloVe word embedding (Pennington et al., 2014). A dropout rate of 0.5 is adopted during training. We set the BOW-loss weight λ to be 0.5. The whole network is trained with the Adam optimizer with a learning rate of 0.001 on GTX Titan X GPUs for 60 epochs. The training takes on average 11.2 hours to finish.

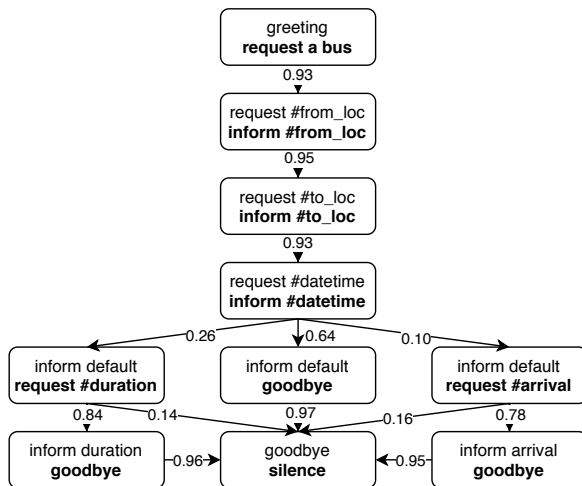


Figure 4: Learned semantic structure of SimDial bus domain (Zhao and Eskenazi, 2018). User intents are marked in bold. Transitions with $P < 0.1$ are omitted.

To evaluate the learned structure, we compare VRNN-LinearCRF’s output in Figure 4 with the

ground truth dialogue structure in Figure 1. A dialogue structure learned by VRNN without structured attention is also shown in the Appendix. We find our method generates similar structure compared to ground truth in the bus domain. Figure 5 shows all models’ quantitative results. Having a lower value in SED and SCE indicates the learned structure is closer to the ground truth and better. Our method with BERT, VRNN-LinearCRF-BERT performs the best. K-means clustering performs worse than VRNN-based models because it only considers utterances’ surface format and ignores the context information. Hidden Markov Model is similar to VRNN but lacks a continuous propagating hidden state layer. VRNN-LinearCRF observes the entire history of latent states but ignores the redundant transitions due to the structure attention. The model’s performance further improves when replacing the vanilla LSTM encoder with a large scale pre-trained encoder like BERT (Devlin et al., 2019), as BERT provides better representations.




Figure 5: All models’ performance in (a) Structure Euclidean Distance (SED) and (b) Structure Cross-Entropy (SCE) in four dialogue domains.

5.2 Interactive Structure Learning in Multi-party Dialogues

We extend our method to learn interactive structure in multi-party dialogues. Specifically, we detect each utterance’s speaker and addressee by constructing an utterance dependency tree.

5.2.1 Datasets

We use Ubuntu Chat Corpus (Uthus and Aha, 2013) as the dataset to study interactive structure since

Model	BLEU1	BLEU2	BLEU3	BLEU4	METEOR	ROUGE _L
HRED	10.54	4.63	2.67	1.53	4.22	10.14
GSN No-speaker (1-iter)	9.23	3.32	1.89	1.24	3.57	8.12
GSN No-speaker (2-iter)	11.32	4.89	2.94	1.54	4.12	10.15
GSN No-speaker (3-iter)	11.42	4.81	3.11	1.87	4.51	10.29
GSN W-speaker (1-iter)	10.11	3.75	1.93	1.31	3.56	9.89
GSN W-speaker (2-iter)	11.43	4.90	2.99	1.63	4.32	10.34
GSN W-speaker (3-iter)	11.52	4.93	3.23	1.91	4.77	11.21
VRNN-Dependency-Tree	11.23	4.92	3.24	1.92	4.69	10.88

Table 2: Different methods’ experiment results on Ubuntu dataset.

From	To	Utterance
p_1	p_2	I know upgrading always got hardon settings to new system..
p_3	—	And the description of the settings is even wrong
p_1	p_2	So these days i always clean install
p_2	p_1	Yeah, i think i will end up doing it
p_2	p_1	Do you happen to know if 12.10 install will let me install grub2 to partition instead of mbr without any extra tweaks?
p_1	p_2	I think default clean install will install grub2 on first section of your hd
p_4	p_2	No

Table 3: Multi-party dialogue example in Ubuntu Chat Corpus (Uthus and Aha, 2013).

it provides the ground-truth of speaker/addressee information for evaluation. Though every record of Ubuntu Chat Corpus contains clear speaker ID, only part of the data has implicit addressee ID, coming as the first word in the utterance. We select addressee ID that appeared in a limited context and extract dialogue sessions with all utterances having verified speaker ID and addressee ID. We extract 20k dialogues with length ranging from 7 to 8 turns. Table 3 shows an example dialogue.

5.2.2 Results and Analysis

Considering Ubuntu Chat Corpus have a large number of technical terminologies, we use a relatively larger vocabulary size of 30k. We use LSTMs and BERT as the sentence embedding encoder and two GRU (Chung et al., 2014) layers with 300 hidden units each as the decoder. The model converges after 100 epochs on GTX Titan X GPUs. The training procedure takes about 54 hours.

To evaluate the learned utterance dependency tree, we compare it with the annotated speaker-addressee relation and find 68.5% utterances are assigned the correct parents. This is a reasonable number because the dependency relationship does not fully rely on the speaker/addressee informa-

tion in a chatroom. A different interlocutor could answer others’ questions even when the questions were not addressed to him/her. Figure 2 visualizes the learned interactive structure from the example in Table 3. Specifically, utterance 4 largely depends on utterance 3, while utterance 6 and 7 are answering the question from utterance 5.

We also compare the model’s generation performance with *Hierarchical Recurrent Encoder-Decoder* (HRED) and *Graph-Structured Network* (GSN) (Hu et al., 2019). The GSN model uses the annotated speaker/addressee information to construct a dialogue graph for utterance encoding iteration. However, this is not required by our VRNN-Dependency-Tree since we generate the original dialogues while learning a dependency structure. For consistent comparison with previous work, we evaluate all models with BLEU 1 to 4, METEOR, and ROUGE_L with the package in (Chen et al., 2015). All results are shown in Table 2. We observe that the proposed VRNN-Dependency-Tree model without using any speaker annotation achieves similar generation performance compared to the state-of-the-art method, GSN with speaker annotation.

6 Conclusion

This paper proposed to inject structured attention into variational recurrent neural network models for unsupervised dialogue structure learning. We explored two different structure inductive biases: linear CRF for utterance-level semantic structure induction in two-party dialogues; and non-projective dependency tree for interactive structure learning in multi-party dialogues. Both models are proved to have a better structure learning performance over the state of the art algorithms. In the future, we will further explore how to explicitly incorporate linguistics information, such as named entities into the latent states.

Acknowledgments

We would like to thank Ziheng Xu for providing scripts to process dialogue datasets. We also thank the outstanding reviewers for their helpful comments to improve our manuscript.

References

- Lalit Bahl, Peter Brown, Peter De Souza, and Robert Mercer. 1986. Maximum mutual information estimation of hidden markov model parameters for speech recognition. In *ICASSP'86. IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 11, pages 49–52. IEEE.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Ananlada Chotimongkol. 2008. *Learning the structure of task-oriented conversations from the corpus of in-domain dialogs*. Ph.D. thesis, Carnegie Mellon University, Language Technologies Institute, School of
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*.
- Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, pages 2980–2988.
- Michael A Covington. 2001. A fundamental algorithm for dependency parsing. In *Proceedings of the 39th annual ACM southeast conference*, pages 95–102. Citeseer.
- Pierre A De Vijver. 1985. Baum’s forward-backward algorithm revisited. *Pattern Recognition Letters*, 3(6):369–373.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Micha Elsner and Eugene Charniak. 2008. **You talking to me? a corpus and algorithm for conversation disentanglement**. In *Proceedings of ACL-08: HLT*, pages 834–842, Columbus, Ohio. Association for Computational Linguistics.
- Barbara J. Grosz and Candace L. Sidner. 1986. **Attention, intentions, and the structure of discourse**. *Computational Linguistics*, 12(3):175–204.
- Colin De la Higuera. 2010. *Grammatical inference: learning automata and grammars*. Cambridge University Press.
- Sepp Hochreiter and Jürgen Schmidhuber. 2001. Long short-term memory.
- Wenpeng Hu, Zhangming Chan, Bing Liu, Dongyan Zhao, Jinwen Ma, and Rui Yan. 2019. Gsn: A graph-structured network for multi-party dialogues. *arXiv preprint arXiv:1905.13637*.
- Sarthak Jain and Byron C. Wallace. 2019. **Attention is not Explanation**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dan Jurafsky. 1997. Switchboard swbd-damsl shallow-discourse-function annotation coders manual. *Institute of Cognitive Science Technical Report*.
- Yoon Kim, Carl Denton, Luong Hoang, and Alexander M Rush. 2017. Structured attention networks. *arXiv preprint arXiv:1702.00887*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Karim Lari and Steve J Young. 1990. The estimation of stochastic context-free grammars using the inside-outside algorithm. *Computer speech & language*, 4(1):35–56.
- Jingjing Liu, Stephanie Seneff, and Victor Zue. 2010. **Dialogue-oriented review summary generation for spoken dialogue recommendation systems**. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 64–72, Los Angeles, California. Association for Computational Linguistics.
- Ryan Lowe, Nissan Pow, Iulian Serban, and Joelle Pineau. 2015. **The Ubuntu dialogue corpus: A large dataset for research in unstructured multi-turn dialogue systems**. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 285–294, Prague, Czech Republic. Association for Computational Linguistics.
- Gabriel Murray, Steve Renals, and Jean Carletta. 2005. Extractive summarization of meeting recordings.
- 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 (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.

Alan Ritter, Colin Cherry, and Bill Dolan. 2010. **Unsupervised modeling of twitter conversations**. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 172–180, Los Angeles, California. Association for Computational Linguistics.

Alexander M. Rush. 2020. **Torch-struct: Deep structured prediction library**.

Iulian V Serban and Joelle Pineau. 2015. Text-based speaker identification for multi-participant open-domain dialogue systems. In *NIPS Workshop on Machine Learning for Spoken Language Understanding*. Montreal, Canada.

Iulian Vlad Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A hierarchical latent variable encoder-decoder model for generating dialogues. In *Thirty-First AAAI Conference on Artificial Intelligence*.

Weiyang Shi, Tiancheng Zhao, and Zhou Yu. 2019. **Unsupervised dialog structure learning**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1797–1807, Minneapolis, Minnesota. Association for Computational Linguistics.

Alessandro Sordoni, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob Grue Simonsen, and Jian-Yun Nie. 2015. A hierarchical recurrent encoder-decoder for generative context-aware query suggestion. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*, pages 553–562.

David C Uthus and David W Aha. 2013. The ubuntu chat corpus for multiparticipant chat analysis. In *2013 AAAI Spring Symposium Series*.

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 *Advances in neural information processing systems*, pages 5998–6008.

Sarah Wiegrefe and Yuval Pinter. 2019. **Attention is not not explanation**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics.

Steve Young. 2006. Using pomdps for dialog management. In *2006 IEEE Spoken Language Technology Workshop*, pages 8–13. IEEE.

Ke Zhai and Jason D. Williams. 2014. **Discovering latent structure in task-oriented dialogues**. In *Proceedings of the 52nd Annual Meeting of the Association*

for Computational Linguistics (Volume 1: Long Papers), pages 36–46, Baltimore, Maryland. Association for Computational Linguistics.

Tiancheng Zhao and Maxine Eskenazi. 2018. **Zero-shot dialog generation with cross-domain latent actions**. In *Proceedings of the 19th Annual SIG-Dial Meeting on Discourse and Dialogue*, pages 1–10, Melbourne, Australia. Association for Computational Linguistics.

Tiancheng Zhao, Kyusong Lee, and Maxine Eskenazi. 2018. **Unsupervised discrete sentence representation learning for interpretable neural dialog generation**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1098–1107, Melbourne, Australia. Association for Computational Linguistics.

Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. **Learning discourse-level diversity for neural dialog models using conditional variational autoencoders**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 654–664, Vancouver, Canada. Association for Computational Linguistics.

A Appendices

A.1 Learned Structures of SimDial

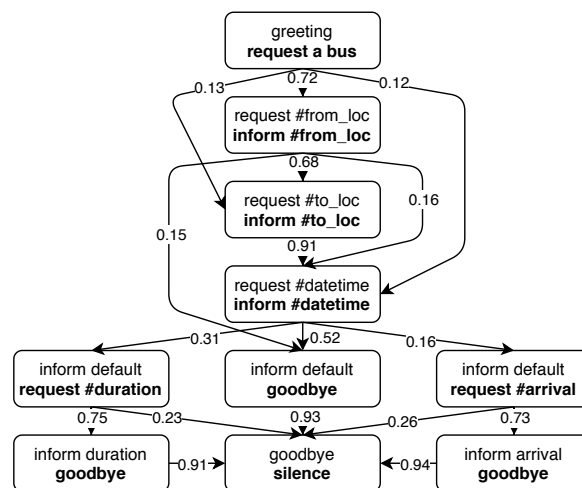


Figure 6: Learned dialogue structure from VRNN without structured attention in SimDial bus domain.

A.2 Inside-Outside Algorithm

Algorithm 2: Inside-Outside for Non-projective Dependency Tree Attention

Input: potential θ_{ij}
 $\alpha, \beta \leftarrow -\infty$
for $i = 1, \dots, n$ **do**
 $\alpha[i, i, L, 1] \leftarrow 0$
 $\alpha[i, i, R, 1] \leftarrow 0$
end for
 $\beta[1, n, R, 1] \leftarrow 0$
for $k = 1, \dots, n$ **do**
 for $s = 1, \dots, n - k$ **do**
 $t \leftarrow s + k$
 $\alpha[s, t, R, 0] \leftarrow \bigoplus_{u \in [s, t-1]} \alpha[s, u, R, 1] \otimes \alpha[u + 1, t, L, 1] \otimes \theta_{st}$
 $\alpha[s, t, L, 0] \leftarrow \bigoplus_{u \in [s, t-1]} \alpha[s, u, R, 1] \otimes \alpha[u + 1, t, L, 1] \otimes \theta_{ts}$
 $\alpha[s, t, R, 1] \leftarrow \bigoplus_{u \in [s+1, t]} \alpha[s, u, R, 0] \otimes \alpha[u, t, R, 1]$
 $\alpha[s, t, L, 1] \leftarrow \bigoplus_{u \in [s, t-1]} \alpha[s, u, L, 1] \otimes \alpha[u, t, L, 0]$
 end for
end for
for $k = n, \dots, 1$ **do**
 for $s = 1, \dots, n - k$ **do**
 $t \leftarrow s + k$
 for $u = s + 1, \dots, t$ **do**
 $\beta[s, u, R, 0] \leftarrow \bigoplus_{\oplus} \beta[s, t, R, 1] \otimes \alpha[u, t, R, 1]$
 $\beta[u, t, R, 1] \leftarrow \bigoplus_{\oplus} \beta[s, t, R, 1] \otimes \alpha[s, u, R, 0]$
 end for
 if $s > 1$ **then**
 for $u = s, \dots, t - 1$ **do**
 $\beta[s, u, L, 1] \leftarrow \bigoplus_{\oplus} \beta[s, t, L, 1] \otimes \alpha[u, t, L, 0]$
 $\beta[u, t, L, 0] \leftarrow \bigoplus_{\oplus} \beta[s, t, L, 1] \otimes \alpha[s, u, L, 1]$
 end for
 end if
 for $u = s, \dots, t - 1$ **do**
 $\beta[s, u, R, 1] \leftarrow \bigoplus_{\oplus} \beta[s, t, R, 0] \otimes \alpha[u + 1, t, L, 1] \otimes \theta_{st}$
 $\beta[u + 1, t, L, 1] \leftarrow \bigoplus_{\oplus} \beta[s, t, R, 0] \otimes \alpha[s, u, R, 1] \otimes \theta_{ts}$
 end for
 if $s > 1$ **then**
 for $u = s, \dots, t - 1$ **do**
 $\beta[s, u, R, 1] \leftarrow \bigoplus_{\oplus} \beta[s, t, L, 0] \otimes \alpha[u + 1, t, L, 1] \otimes \theta_{ts}$
 $\beta[u + 1, t, L, 1] \leftarrow \bigoplus_{\oplus} \beta[s, t, L, 0] \otimes \alpha[s, u, R, 1] \otimes \theta_{ts}$
 end for
 end if
 end for
 $A \leftarrow \alpha[1, n, R, 1]$
for $s = 1, \dots, n - 1$ **do**
 for $t = s + 1, \dots, n$ **do**
 $p[s, t] \leftarrow \exp(\alpha[s, t, R, 0] \otimes \beta[s, t, R, 0] \otimes -A)$
 if $s > 1$ **then**
 $p[t, s] \leftarrow \exp(\alpha[s, t, L, 0] \otimes \beta[s, t, L, 0] \otimes -A)$
 end if
 end for
end for
