

Logician and Orator: Learning from the Duality between Language and Knowledge in Open Domain

Mingming Sun^{1,2}, Xu Li^{1,2}, Ping Li¹

¹Cognitive Computing Lab (CCL), Baidu Research

²National Engineering Laboratory of Deep Learning Technology and Application, China
{sunmingming01, lixu13, liping11}@baidu.com

Abstract

We propose the task of Open-Domain Information Narration (OIN) as the reverse task of Open Information Extraction (OIE), to implement the dual structure between language and knowledge in the open domain. We then develop an agent, called Orator, to accomplish the OIN task, and assemble the Orator and the recently proposed OIE agent – Logician (Sun et al., 2018) into a dual system to utilize the duality structure with a reinforcement learning paradigm. Experimental results reveal the dual structure between OIE and OIN tasks helps to build better both OIE agents and OIN agents.

1 Introduction

The duality between language and knowledge is natural for human intelligence. The human can extract knowledge from natural language to learn or remember, and then narrate the knowledge back to natural language to communicate. Information extraction (IE) is a task to simulate the first part of the duality, which is a long-term hot spot for NLP research. Recently, the task that fulfills the last part of the duality, that is, assembling a set of relation instances/facts or database records into natural language sentences/documents, has also attracted many interests (Wiseman et al., 2017; Chisholm et al., 2017; Agarwal and Dymetman, 2017; Vougiouklis et al., 2017; Yin et al., 2016). In the literature, this task has been referred to as “data to document generation” (Wiseman et al., 2017) or “knowledge-to-text” (Chisholm et al., 2017). In this paper, we name the task as *information narration (IN)*, to emphasize the reverse relationship to the information extraction (IE) task.

The duality between language and knowledge (and thus between the IE and IN tasks) can be examined in closed-domain or open-domain. For the closed-domain problem, the closed-domain IE (CIE) task is often referred to as “relation ex-

traction” or “relation classification”, which identifies instances of a fixed and finite set of relations from natural language corpus, using supervised methods (Kambhatla, 2004; Zelenko et al., 2003; Miwa and Bansal, 2016; Zheng et al., 2017) or weakly supervised methods (Mintz et al., 2009; Lin et al., 2016). In the meantime, the close-domain IN (CIN) task (Wiseman et al., 2017; Chisholm et al., 2017; Agarwal and Dymetman, 2017; Vougiouklis et al., 2017; Yin et al., 2016) transforms a set of facts with a pre-defined schema or relation types (such as facts from Freebase (Bollacker et al., 2008), DBpedia (Auer et al., 2007), or database tables), into natural language sentences/documents. Furthermore, the dual structure between CIE and CIN tasks has been noticed and utilized in (Chisholm et al., 2017).

For the open-domain problem, the open-domain IE (OIE) task is to investigate how the natural language sentences express the facts, and then use the learned knowledge to extract entity and relation level intermediate structures from open-domain sentences (Christensen et al., 2011; Etzioni et al., 2011; Schmitz et al., 2012; Pal and Mausam, 2016). Although the OIE task has attracted much interests and obtained many applications (Christensen et al., 2013, 2014; Mausam, 2016; Stanovsky et al., 2015; Khot et al., 2017; Fader et al., 2014), the OIN task has not been stated, neither the duality between the language and knowledge in the open domain.

	Open-Domain	Closed-Domain
Extraction	OIE	CIE
Narration	OIN	CIN

Table 1: Taxonomy: Tasks between knowledge and natural language.

The tasks involved in the duality between language and knowledge is shown in Table 1, where

the OIN task has not been stated. In this paper, we focus on the OIN task and the duality between the OIE and OIN tasks, for following reasons: 1) the OIN task is an essential component for open-domain information processing pipeline. For example, it is helpful for building natural and informative response for open-domain KBQA systems (Khot et al., 2017; Fader et al., 2014). 2) (as the results in this paper will illustrate) the duality between tasks can be valuable for building better agents for both tasks (Xia et al., 2017, 2016).

A major historical obstacle for investigating the duality between OIE and OIN tasks is the absence of parallel corpus between natural language sentence and open-domain facts. Recently, the SAOKE dataset (Sun et al., 2018) was released, which contains more than forty thousand of human-labeled open-domain sentence-facts pairs, and thus essentially eliminates the obstacle for our investigation.

The contribution of this paper lies in following aspects:

- We propose the concept of OIN task, which is potentially an important component for open-domain information pipeline. We develop the Orator agent to fulfill the task;
- We build a multi-agent system with Logician and Orator to exploit the dual structure between language and knowledge in open domain. Experimental results reveal that the dual information is beneficial for improving the performance of both agents.

The paper is organized as follows: Section 2 discusses the related work. Section 3 explains the Orator agent for OIN task. Section 4 describes the multi-agent system with Logician and Orator and its algorithm to learn from the duality between language and sentence. The experimental results of the fine tuned agents are shown and discussed in Section 5. We conclude our work and discuss the future direction in Section 6.

2 Related Work

2.1 Information Narration

The closed-domain information narration (CIN) task has been studied in (Wiseman et al., 2017; Chisholm et al., 2017; Agarwal and Dymetman, 2017; Vougiouklis et al., 2017). These

CIN agents face problems from different problem domains, from people biographies to basketball game records, but most of them follow the same sequence-to-sequence pattern. First, the algorithm encodes a sequence of facts into a set of annotations and then decodes the annotations into a natural language text. Mechanisms such as attention (Bahdanau et al., 2014) and copying (Gu et al., 2016) are employed into the decoder to improve the performance. Then, the models are trained on a supervised dataset with back-propagation.

In this work, we adopt a similar sequence-to-sequence architecture to build our baseline Orator agent, but with following differences: 1) the Orator is proposed to narrate open domain facts, where the encoder must encode words rather than the entities and relations in the closed domain; 2) the baseline Orator will be fine tuned using the dual learning algorithm proposed in this paper.

2.2 Dual Learning Systems

For many natural language processing tasks, there exist corresponding reverse/dual tasks. One example of a pair of dual problems is the question answering (QA) and question generation (QG). In (Tang et al., 2017), the duality between QA problem and QG problem was considered as a constraint that both problems must share the same joint probability. Then, a loss function that implemented the constraint was involved in the supervised learning procedures for both agents. Furthermore, researchers (Tang et al., 2017; Duan et al., 2017; Sachan and Xing, 2018) use both the question-answering agent and the question-generation agent to identify extra high-confident question-answering pairs, which are further used to fine tune the pre-trained agents.

Back-and-forth translation (or round-trip translation)¹ is another example of duality, in the field of machine translation. It has been employed to evaluate the quality of machine translation systems (Van Zaanen and Zwarts, 2006), or to test the suitability of text for machine translation (Gaspary, 2006; Shigenobu, 2007). Recently, (Xia et al., 2016) implemented the duality in a neural-based dual learning system, in which the quality of each translation agent was improved on the unlabeled dataset using the rewards provided by its

¹https://en.wikipedia.org/wiki/Round-trip_translation

	Chinese	English Translation
Sentence	他发明了有两个固定点（水的沸点和冰点）的100摄氏度计温尺，是世界大多数地区通用的摄氏温度计的前身。	He invented a 100 degree centimeter temperature scale with two fixed points (the boiling point of water and the freezing point), which is the precursor of the Celsius thermometer in most parts of the world.
Facts	(他,发明,100摄氏度计温尺)(100摄氏度计温尺,有两个固定点)(水的[沸点 冰点],ISA,两个固定点)(100摄氏度计温尺,是X的前身,世界大多数地区通用的摄氏温度计)	(He, invented, the 100 degree Celsius temperature scale) (100 Celsius temperature scale, has, two fixed points) (water [boiling point freezing point], ISA, two fixed points) (100 degrees Celsius temperature scale, is the predecessor of X, most Celsius thermometers in most parts of the world)

Table 2: An example sentence and the corresponding facts in the SAOKE dataset, where “ISA” is a symbol denoting the “is-a” relationship in the SAOKE format.

corresponding dual agent, using the reinforcement learning technique.

Parsing-reconstruction is also a pattern of duality. (Konstas et al., 2017) considered the AMR (Abstract Meaning Representation)(Banarescu et al., 2013) parsing problem (text to AMR) and AMR generation problem (AMR to text) in one system, in which the AMR parser generated extra text-AMR pair data to fine tune the AMR generator. The AMR generator, however, does not contribute to the performance improvement of the AMR parser. The CIE agent and CIN agent in (Chisholm et al., 2017) also follow this pattern, where both agents help each other to improve by sharing weights. Nevertheless, the sharing weight strategy cannot be applied to agents with different architecture, which is a typical situation in practice.

From these practices, it can be seen that the duality can be implemented with two major approaches: 1) by providing additional labeled samples via bootstrapping, and 2) by adding losses or rewards to the training procedure of the agents. In this paper, we follow the second approach. We design a set of rewards, among which some are related to OIE and OIN tasks respectively, and some are related to the duality of the problems. Then we optimize both agents using the reinforcement learning technique. The learning algorithm is similar to the dual-NMT algorithm described in (Xia et al., 2016), but with adaption for the OIE and OIN tasks, especially on the task related rewards. Compared to the approach of applying the regularization about sharing the same joint probability (Tang et al., 2017), our approach directly optimizes the task objective by introducing task

related rewards. Furthermore, our approach is more adaptable than the weight sharing approach adopted in (Chisholm et al., 2017).

3 Orator

3.1 SAOKE Dataset

Symbolic Aided Open Knowledge Expression (SAOKE) is proposed in (Sun et al., 2018) as the form to honestly record the facts that humans can extract from sentences when humans read them. SAOKE uses a unified form - an n-ary tuple:

$$(subject, predicate, object_1, \dots, object_N),$$

to express four categories of facts: 1) Relation: Verb/preposition-based n-ary relations between entity mentions; 2) Attribute: Nominal attributes for entity mentions; 3) Description: Descriptive phrases of entity mentions; 4) Concept: Hyponymy and synonymy relations among concepts and instances.

Using this SAOKE format, Sun et al. (2018) manually labeled the SAOKE dataset \mathbb{D}_{SAOKE} by crowdsourcing, which includes more than forty thousand sentence-facts pairs $\langle \mathcal{S}, \mathcal{F} \rangle$.² The labeling procedure is under the supervision of the “Completeness” criterion (Sun et al., 2018), so the facts recorded information in the sentence as much as possible (only auxiliary information and relation between facts are omitted (Sun et al., 2018)). As a result, the SAOKE dataset is a valid open-domain sentence-facts parallel dataset for both OIE and OIN tasks. Table 2 is an example from the SAOKE dataset for an easy understanding of the dual relationship between sentence and facts.

²<http://ai.baidu.com/broad>

3.2 Model

The Orator is an agent \mathcal{D} that assembles a set of open-domain facts \mathcal{F} into a sentence \mathcal{S} with probability $P_{\mathcal{D}}(\mathcal{S}|\mathcal{F}, \Theta_{\mathcal{D}})$, where $\Theta_{\mathcal{D}}$ is the set of parameters of \mathcal{D} :

$$\text{Orator } \mathcal{D}: \mathcal{F} \rightarrow P_{\mathcal{D}}(\mathcal{S}|\mathcal{F}, \Theta_{\mathcal{D}}).$$

For each pair $\langle \mathcal{S}, \mathcal{F} \rangle \in \mathbb{D}_{\text{SAOKE}}$, the set of facts \mathcal{F} is actually expressed as a sequence of facts, in the order of the labeler wrote them. So, the deep sequence to sequence paradigm is suitable to model the Orator. In this work, we build the base Orator model with the attention-based sequence to sequence model, together with copy and coverage mechanism, in a similar way of the implementation of the Logician in (Sun et al., 2018).

3.2.1 Attention based Sequence-to-sequence Learning

The attention-based sequence-to-sequence learning (Bahdanau et al., 2014) first encodes the input fact sequence \mathcal{F} (actually the sequence of N_e -dimensional word embedding vectors) into a N_h -dimensional hidden states $H^{\mathcal{F}} = [h_1^{\mathcal{F}}, \dots, h_{N_S}^{\mathcal{F}}]$ using bi-directional GRU (Gated Recurrent Units) network (Cho et al., 2014). Then, when generating word w_t of the target sentence, the decoder computes the probability of generating w_t by $p(w_t|\{w_1, \dots, w_{t-1}\}, c_t) = g(h_{t-1}, s_t, c_t)$, where s_t is the hidden state of the GRU decoder, g is the word generation model, and c_t is the dynamic context vector which focuses attention on specific location l in the input hidden states $H^{\mathcal{F}}$.

For the Orator, we use the copy mechanism to implement the word generation model g and use the coverage mechanism to compute the dynamic context vector c_t .

3.2.2 Copy Mechanism

In the SAOKE dataset, the words in the set of facts (excluding the external symbols) must be in the corresponding sentence, so the problem is suitable to be modeled via the copy mechanism (Gu et al., 2016). In the copy mechanism, when the decoder is considering generating a word w_t , it can either be copied from the source fact sequence \mathcal{F} or select from a vocabulary V :

$$p(w_t|w_{t-1}, s_t, c_t) = p_{\mathcal{F}}(w_t|w_{t-1}, s_t, c_t) + p_V(w_t|w_{t-1}, s_t, c_t),$$

where $p_{\mathcal{F}}$ is the probability of copying from \mathcal{F} and p_V is the probability of selecting from V . The details can be found in (Gu et al., 2016).

3.2.3 Coverage Mechanism

To cope with the problem of information lost or redundancy in the generated sentence, the copied histories of previous generated words should be remembered to guide future generation. This could be done through the coverage mechanism (Tu et al., 2016), in which a coverage vector m_j^t is introduced for each word $w_j^{\mathcal{F}}$ in \mathcal{F} and updated at each step t as a gated function of $h_j^{\mathcal{F}}, \alpha_{tj}, s_{t-1}, m_j^{t-1}$. By this means, the coverage vectors remember the historical attentions over source sequence and can be incorporated in the alignment model to generate complete and non-redundant sentences. Detailed formulations can be found in (Tu et al., 2016) and (Sun et al., 2018).

4 Learning the Dual Structure between Knowledge and Natural Language

4.1 Dual Structure between Orator and Logician

In (Sun et al., 2018), an agent \mathcal{L} , called Logician, was trained to convert a sentence \mathcal{S} into a set of facts \mathcal{F} with probability $P_{\mathcal{L}}(\mathcal{F}|\mathcal{S}, \Theta_{\mathcal{L}})$, where $\Theta_{\mathcal{L}}$ is the set of parameters of \mathcal{L} :

$$\text{Logician } \mathcal{L}: \mathcal{S} \rightarrow P_{\mathcal{L}}(\mathcal{F}|\mathcal{S}, \Theta_{\mathcal{L}}).$$

Obviously, the Logician and Orator can cooperate to supervise each other. Given $\langle \mathcal{S}, \mathcal{F} \rangle \in \mathbb{D}_{\text{SAOKE}}$, the Logician produces a predicted set of facts \mathcal{F}^* for the sentence \mathcal{S} , and the Orator can calculate the probability $P_{\mathcal{D}}(\mathcal{S}|\mathcal{F}^*, \Theta_{\mathcal{D}})$ of reconstruction \mathcal{S} from \mathcal{F}^* . Intuitively, if \mathcal{F}^* loses major information of \mathcal{S} , honestly reconstructing \mathcal{S} from \mathcal{F}^* would be impossible, and thus the probability $P_{\mathcal{D}}(\mathcal{S}|\mathcal{F}^*, \Theta_{\mathcal{D}})$ would be small. Thus, it is a strong signal to evaluate the quality of \mathcal{F}^* . Similarly, when the Orator produces a sentence \mathcal{S}^* for the set of facts \mathcal{F} , the probability $P_{\mathcal{L}}(\mathcal{F}|\mathcal{S}^*, \Theta_{\mathcal{L}})$ provided by the Logician is a strong signal for evaluating the quality of \mathcal{S}^* . These signals are helpful to conquer several problems of the original agents, including information lost, information redundancy, and non-fluency.

Note that the supervision signals $P_{\mathcal{D}}(\mathcal{S}|\mathcal{F}^*, \Theta_{\mathcal{D}})$ and $P_{\mathcal{L}}(\mathcal{F}|\mathcal{S}^*, \Theta_{\mathcal{L}})$ do not rely on any supervised parallel corpus. Thus, similar to the application of dual learning paradigm on NMT

task (Xia et al., 2016), it is theoretically possible to use unparallelled sentences and sets of facts to compute these signals. However, unsupervised collections of fact-groups that can be reasonably narrated in a sentence are not naturally available. Currently, the only available collection is the sets of facts provided by the SAOKE dataset, where the supervised information is available. As a result, we implement our dual learning system in a supervised approach, which uses the reinforcement learning algorithm to optimize the Orator and the Logician. The involved rewards are described in the next subsection, and then the algorithm is detailed in the last subsection.

4.2 Rewards

Given $\langle \mathcal{S}, \mathcal{F} \rangle \in \mathbb{D}_{\text{SAOKE}}$, we sample a set of facts \mathcal{F}^* from distribution $P_{\mathcal{L}}(\cdot|\mathcal{S}, \Theta_{\mathcal{L}})$ and a sentence \mathcal{S}^* from distribution $P_{\mathcal{D}}(\cdot|\mathcal{F}, \Theta_{\mathcal{D}})$. Following rewards are introduced into the proposed dual learning system, and the relationships between them are shown in Figure 1.




Figure 1: Illustration of the dual learning system of Logician and Orator.

4.2.1 Reconstruction Rewards

Following the idea described in above subsection, we design the reconstruction reward for the Orator as:

$$R_{\mathcal{D}}(\mathcal{S}^*, \mathcal{F}) = \log P_{\mathcal{L}}(\mathcal{F}|\mathcal{S}^*, \Theta_{\mathcal{L}}),$$

and that for the Logician as:

$$R_{\mathcal{L}}(\mathcal{F}^*, \mathcal{S}) = \log P_{\mathcal{D}}(\mathcal{S}|\mathcal{F}^*, \Theta_{\mathcal{D}}).$$

4.2.2 Similarity Rewards

Since the SAOKE dataset has label information, the similarities between the predicted results and the ground truths can be used as rewards.

For the Orator, since the \mathcal{S}^* can be viewed as the summarization of \mathcal{S} , we use the widely used ROUGE-L (Lin, 2004) measure in the text summarization field to evaluate the quality of \mathcal{S}^* :

$$S_{\mathcal{D}} = \text{ROUGE}_L(\mathcal{S}, \mathcal{S}^*).$$

For the Logician, we use following procedure to calculate the similarity between \mathcal{F} and \mathcal{F}^* . First, we compute the similarity between each predicted fact $f^* \in \mathcal{F}^*$ and each ground truth fact $f \in \mathcal{F}$ with following measure:

$$\text{SimFact}(f^*, f) = \frac{\sum_{i=1}^{\min(|f^*|, |f|)} \text{SimStr}(f_i^*, f_i)}{\max(|f^*|, |f|)},$$

where f_i^* and f_i denote the i -th element of tuples of fact f^* and f , $\text{SimStr}(\cdot, \cdot)$ denotes the gestalt pattern matching (Ratcliff and Metzner, 1988) measure for two strings, and $|\cdot|$ is the cardinality function. Then, each predicted fact in \mathcal{F}^* is aligned to its corresponding ground-truth fact in \mathcal{F} by solving a linear assignment problem (Wikipedia, 2017) to maximize the sum of similarities between the aligned facts. Finally, the similarity reward for the Logician is calculated by:

$$S_{\mathcal{L}}(\mathcal{F}^*, \mathcal{F}) = \frac{\sum \text{SimFact}(f^*, f)}{\max(|\mathcal{F}^*|, |\mathcal{F}|)},$$

where $f^* \in \mathcal{F}^*$, $f \in \mathcal{F}$ are aligned pair of facts.

4.2.3 Validity Rewards

For the Orator, the output is expected as a valid natural language sentence, so the validity reward can be defined as:

$$V_{\mathcal{D}}(\mathcal{S}^*) = LM(\mathcal{S}^*),$$

where the $LM(\cdot)$ is a language model.

For the Logician, the output should represent a valid collection of facts, which means: 1) the output can be parsed into a collection of facts; 2) there is no duplicated fact (identified by the SimFact value larger than 0.85) in the parsed collection. The validity reward for Logician is defined as:

$$V_{\mathcal{L}}(\mathcal{F}^*) = \begin{cases} 0 & \text{if } \mathcal{F}^* \text{ is valid;} \\ -1 & \text{otherwise.} \end{cases}$$

Algorithm 1 A simple dual-learning algorithm for facts extraction and expression

Require:

- A set of sentence-facts pairs $\{ \langle \mathcal{S}, \mathcal{F} \rangle \}$;
- An initial Logician \mathcal{L} and an initial Orator \mathcal{D} ;
- Beam size K ;

repeat

- 1: Sample a sentence-facts pair $\langle \mathcal{S}, \mathcal{F} \rangle$;
- 2: Logician produces K sets of facts $\mathcal{F}_1, \dots, \mathcal{F}_K$ from \mathcal{S} via beam search;
- 3: **for** each set of facts \mathcal{F}_i **do**
- 4: Compute the reward for \mathcal{F}_i as:

$$r_i^{\mathcal{F}} = \alpha_1 R_{\mathcal{D}}(\mathcal{S}, \mathcal{F}_i) + \alpha_2 V_{\mathcal{F}}(\mathcal{F}) + \alpha_3 S_{\mathcal{F}}(\mathcal{F}, \mathcal{F}_i).$$

- 5: **end for**
- 6: Compute the total reward $r = \frac{1}{K} \sum_{i=1}^K r_i^{\mathcal{F}}$;
- 7: Compute the stochastic gradient of $\Theta_{\mathcal{L}}$:

$$\nabla_{\Theta_{\mathcal{L}}} \hat{E}[r] = \frac{1}{K} \sum_{i=1}^K r_i^{\mathcal{F}} \mathcal{D}_{\Theta_{\mathcal{L}}}(\mathcal{S}, \mathcal{F}_i)$$

- 8: Compute the stochastic gradient of $\Theta_{\mathcal{D}}$:

$$\nabla_{\Theta_{\mathcal{D}}} \hat{E}[r] = \frac{\alpha_1}{K} \sum_{i=1}^K \mathcal{D}_{\Theta_{\mathcal{D}}}(\mathcal{F}_i, \mathcal{S})$$

- 9: Model updates:

$$\begin{aligned} \Theta_{\mathcal{L}} &\leftarrow \Theta_{\mathcal{L}} + \eta_{\mathcal{L}} \cdot \nabla_{\Theta_{\mathcal{L}}} \hat{E}[r], \\ \Theta_{\mathcal{D}} &\leftarrow \Theta_{\mathcal{D}} + \eta_{\mathcal{D}} \cdot \nabla_{\Theta_{\mathcal{D}}} \hat{E}[r]. \end{aligned}$$

- 1: Sample a sentence-facts pair $\langle \mathcal{S}, \mathcal{F} \rangle$;
- 2: Orator produces K sentences $\mathcal{S}_1, \dots, \mathcal{S}_K$ from \mathcal{F} via beam search;
- 3: **for** each sentence \mathcal{S}_i **do**
- 4: Compute the reward for \mathcal{S}_i as:

$$r_i^{\mathcal{S}} = \beta_1 R_{\mathcal{L}}(\mathcal{F}, \mathcal{S}_i) + \beta_2 V_{\mathcal{S}}(\mathcal{S}_i) + \beta_3 S_{\mathcal{S}}(\mathcal{S}, \mathcal{S}_i).$$

- 5: **end for**
- 6: Compute the total reward $r = \frac{1}{K} \sum_{i=1}^K r_i^{\mathcal{S}}$;
- 7: Compute the stochastic gradient of $\Theta_{\mathcal{D}}$,

$$\nabla_{\Theta_{\mathcal{D}}} \hat{E}[r] = \frac{1}{K} \sum_{i=1}^K r_i^{\mathcal{S}} \mathcal{D}_{\Theta_{\mathcal{D}}}(\mathcal{S}_i, \mathcal{F})$$

- 8: Compute the stochastic gradient of $\Theta_{\mathcal{L}}$:

$$\nabla_{\Theta_{\mathcal{L}}} \hat{E}[r] = \frac{\beta_1}{K} \sum_{i=1}^K \mathcal{D}_{\Theta_{\mathcal{L}}}(\mathcal{F}_i, \mathcal{S})$$

- 9: Model updates:

$$\begin{aligned} \Theta_{\mathcal{D}} &\leftarrow \Theta_{\mathcal{D}} + \eta_{\mathcal{D}} \cdot \nabla_{\Theta_{\mathcal{D}}} \hat{E}[r], \\ \Theta_{\mathcal{L}} &\leftarrow \Theta_{\mathcal{L}} + \eta_{\mathcal{L}} \cdot \nabla_{\Theta_{\mathcal{L}}} \hat{E}[r]. \end{aligned}$$

until convergence

4.3 Algorithm

For each pair $\langle \mathcal{S}, \mathcal{F} \rangle \in \mathbb{D}_{\text{SAOKE}}$, the following procedures are performed respectively (details are shown in Algorithm 1):

4.3.1 Learning from Sentence to Facts

We sample \mathcal{F}^* from the Logician $P_{\mathcal{L}}(\cdot | \mathcal{S}, \Theta_{\mathcal{L}})$ and calculate the total reward for \mathcal{F}^* by

$$r_{\mathcal{L}} = \alpha_1 \cdot R_{\mathcal{L}}(\mathcal{F}^*, \mathcal{S}) + \alpha_2 \cdot V_{\mathcal{L}}(\mathcal{F}^*) + \alpha_3 \cdot S_{\mathcal{L}}(\mathcal{F}^*, \mathcal{F}),$$

where $\sum \alpha_i = 1$. The gradients of the expected reward $E[r_{\mathcal{L}}]$ to the parameters of agents can be computed as follows, according to the policy gradient theorem (Sutton et al., 1999):

$$\begin{aligned} \nabla_{\Theta_{\mathcal{L}}} E[r_{\mathcal{L}}] &= E[r_{\mathcal{L}} \mathcal{D}_{\Theta_{\mathcal{L}}}(\mathcal{F}^*, \mathcal{S})], \\ \nabla_{\Theta_{\mathcal{D}}} E[r_{\mathcal{L}}] &= E[\alpha_1 \mathcal{D}_{\Theta_{\mathcal{D}}}(\mathcal{S}, \mathcal{F}^*)]. \end{aligned}$$

where $\mathcal{D}_{\Theta_{\mathcal{L}}}(\mathcal{F}, \mathcal{S}) = \nabla_{\Theta_{\mathcal{L}}} \log P_{\mathcal{L}}(\mathcal{F} | \mathcal{S}, \Theta_{\mathcal{L}})$ and $\mathcal{D}_{\Theta_{\mathcal{D}}}(\mathcal{S}, \mathcal{F}) = \nabla_{\Theta_{\mathcal{D}}} \log P_{\mathcal{D}}(\mathcal{S} | \mathcal{F}, \Theta_{\mathcal{D}})$.

4.3.2 Learning from Facts to Sentence

We sample \mathcal{S}^* from the Logician $P_{\mathcal{D}}(\cdot | \mathcal{F}, \Theta_{\mathcal{D}})$ and define the total reward for \mathcal{S}^* by:

$$r_{\mathcal{D}} = \beta_1 R_{\mathcal{D}}(\mathcal{S}^*, \mathcal{F}) + \beta_2 V_{\mathcal{D}}(\mathcal{S}^*) + \beta_3 S_{\mathcal{D}}(\mathcal{S}^*, \mathcal{S}),$$

where $\sum \beta_i = 1$. The gradients can be computed as follows:

$$\begin{aligned} \nabla_{\Theta_{\mathcal{L}}} E[r_{\mathcal{D}}] &= E[\beta_1 \mathcal{D}_{\Theta_{\mathcal{L}}}(\mathcal{F}, \mathcal{S}^*)], \\ \nabla_{\Theta_{\mathcal{D}}} E[r_{\mathcal{D}}] &= E[r_{\mathcal{D}} \mathcal{D}_{\Theta_{\mathcal{D}}}(\mathcal{S}^*, \mathcal{F})]. \end{aligned}$$

In practice, we use beam search (Sutskever et al., 2014) to obtain high-quality samples as \mathcal{F}^* and \mathcal{S}^* , and estimate the true gradient with the empirical average of gradients over these samples.

5 Experimental Results

5.1 Experimental Design

First, we evaluate the performance of each agent fine-tuned by the dual learning procedure on the

SAOKE dataset. Then we evaluate the Orator on noisy facts, which accords with real OIN application scenarios. Last, we investigate the behavior of agents in the dual system.

In the experiments, the SAOKE dataset is split into the training set, validating set and testing set with ratios of 80%, 10%, 10%, respectively. For each algorithm involved in the experiments, we perform grid search to find the optimal super-parameters, and the model with the best performance on the validating set is chosen as the learnt model to be evaluated on the testing set.

5.1.1 Evaluation Metric

For the Orator, BLEU-4 and ROUGE-L are used to measure how well the output matches the ground truth sentence.

For the Logician, based on the fact-equivalence judgment proposed in (Sun et al., 2018), we compute the Precision(P), Recall (R) and F1-score over the testing set of the SAOKE dataset as the evaluation metric.

5.1.2 Agent Implementation

For the Orator, we make a vocabulary V with size 72,591 by collecting all web pages from Baidu Baike website³ (a Chinese alternative to Wikipedia) and identifying the words occurred in more than 100 web pages. For the Orator, the dimension of embedding vectors is set to $N_e = 256$, and the dimension of hidden states is set to $N_h = 256$. We use a three-layer bi-directional GRU with dimension 128 as the encoder. All dimensions of hidden states in the decoder are set to 256.

For the Logician, we implement the model described in (Sun et al., 2018), including the shallow tag information and the gated dependency attention mechanism.

Furthermore, to provide an intuitive comprehension of the OIN task, we implement a rule-based method for OIN task. For each sequence of facts in the SAOKE dataset, the method first identifies the subsequences in which the facts share the same subject. Then it preserves the subject of the first fact in each subsequence and removes the subjects of following facts (by replacing it with an empty string). It is necessary since the SAOKE dataset requires the shared subject to be repeated for completeness of the related facts. At last, each fact is formatted into a string by filling the objects into the placeholders of the predicate and these

strings are concatenated with commas to form the final sentence.

5.1.3 Reward Implementation

For the validity reward of the Orator, the language model is trained using an RNN based method (Mikolov et al., 2010) with the same vocabulary V and the web pages from Baidu Baike website.

For the reconstruction reward of the Orator, since the Logician needs the shallow tag and dependency information of \mathcal{S}^* as inputs, the information is extracted using the LTP tool-set (Che et al., 2010) and then fed to the Logician.

5.1.4 Training

When training the base model for each agent, the batch size is set to 20. When training two agents in dual learning, the batch size is set to 12, and the beam size is set to 3. Both agents are trained using the stochastic gradient descent (SGD) with RM-SPROP strategy (Hinton et al., 2012) and early-stop strategy on the validating set. In dual learning, the super-parameters, including α_i , β_i , is determined by grid-search.

5.2 Evaluation of Agents on the SAOKE dataset

First, we evaluate the performance of the agents optimized by the dual learning method. To identify the contribution of the dual structure, we train another pair of agents with $\alpha_1 = 0$ and $\beta_1 = 0$ in Algorithm 1 to exclude the dual information. Without the dual information, these two agents are trained independently to each other with reinforcement learning on their own supervised information. We name these two agents as R-Logician and R-Orator, where ‘‘R’’ means ‘‘Reinforced’’. In the experimental results of this paper, the symbol at the top mark means that the marked result is significantly different (with $p = 0.05$) with the corresponding result of the agent with the specific mark.

Methods	Precision	Recall	F1
Logician*	0.449	0.400	0.423
R-Logician [‡]	0.462*	0.432*	0.446*
Logician@Dual	0.494* [‡]	0.426*	0.457* [‡]

Table 3: Performance of the Logicians.

The experimental results for the Logician agents are shown in Table 3, from which we can observe a significant performance improvement

³<http://baike.baidu.com>

from Logician to R-Logician and also from R-Logician to Logician@Dual. The experimental results for the Orator agents are shown in Table 4. The neural based Orator agents significantly outperform the rule-based agent. For both evaluation metric, the R-Orator and Orator@Dual are both significantly outperform the original Orator. The Orator@Dual significantly outperforms the R-Orator on the BLEU-4 score, but is not significantly different on the ROUGE-L score.

Methods	BLEU-4	ROUGE-L
Rule*	0.257	0.434
Orator*	0.401*	0.556*
R-Orator [‡]	0.405**	0.559**
Orator@Dual	0.419** [‡]	0.559**

Table 4: Performance of the Orators.

By comparing the performance of R-agents and the agents@Dual, we can observe that agents@Dual generally achieve better performance on precision, but may recall less information, resulting in smaller advances in the balanced evaluation metric (F1 and ROUGE-L). This may imply that the agents tend to provide easy input for each other for higher accuracy, by neglecting some difficult part of the problem which they currently cannot handle properly. This interesting phenomenon is the subject of our future research.

5.3 Evaluation of Orator on Noisy Facts

Experiments in the previous subsection show the performance of Orators to narrate a set of human-labeled facts. In practice, however, the input to the Orator might not be the human-labeled perfect facts, but some noisy facts automatically extracted by OIE algorithms. In this subsection, we make a collection of sets of noisy facts by feeding the sentences in the testing set of the SAOKE dataset to the base Logician model and collecting the outputs. Then we evaluate the series of Orator models on these noisy facts, and report their performance at Table 5, from which we can see the performance improvement from the Orator to Orator@Dual.

Methods	BLEU-4	ROUGE-L
Orator*	0.428	0.565
R-Orator [‡]	0.431*	0.567*
Orator@Dual	0.458* [‡]	0.572* [‡]

Table 5: Performance of the Orators on noisy facts.

5.4 Evaluation of the Dual System

In this section, we investigate the behavior of agents in the dual system. We first examine the procedure $\mathcal{F} \xrightarrow{\text{Orator}} \mathcal{S}^* \xrightarrow{\text{Logician}} \mathcal{F}^{**}$, that is, for each \mathcal{F} in the testing set of the SAOKE dataset, let the Orator narrate it into a sentence \mathcal{S}^* , and then let the Logician to extract facts \mathcal{F}^{**} from \mathcal{S}^* . Then the quality of \mathcal{F}^{**} is measured by comparing it with \mathcal{F} . Then we examine $\mathcal{S} \xrightarrow{\text{Logician}} \mathcal{F}^* \xrightarrow{\text{Orator}} \mathcal{S}^{**}$, which is the reverse procedure. The comparison is made between the family of base agents and that of the dual-trained agents. The results are shown in Table 6, and two instance of these two experiments are shown in Table 7 and 8 respectively. From these results, we can observe large improvements of reconstruction quality on both directions.

Methods	$\mathcal{F} \xrightarrow{\text{Orator}} \mathcal{S}^* \xrightarrow{\text{Logician}} \mathcal{F}^{**}$		
	Precision	Recall	F1
Base*	0.574	0.488	0.527
Dual	0.657*	0.565*	0.608*
Methods	$\mathcal{S} \xrightarrow{\text{Logician}} \mathcal{F}^* \xrightarrow{\text{Orator}} \mathcal{S}^{**}$		
	BLEU-4	ROUGE-L	
Base*	0.428	0.565	
Dual	0.635*	0.635*	

Table 6: Reconstruction performance for the Logicians and the Orators.

6 Conclusion

In this paper, we investigate the OIN task and its duality to the OIE task. The proposed Orator has shown its ability to fulfill the OIN task, that is, assembling open-domain facts into high quality sentences. Furthermore, our attempt to utilize the duality between OIN and OIE tasks for improving the performances for both OIN and OIE agents accomplishes a preliminary success.

Our work suggests at least three future research topics: Firstly, one can enrich the theoretical study of the duality between the OIE and OIN tasks. Secondly, one can investigate how to conquer the barrier of the absence of an extensive collection of reasonable sets of open-domain facts and incorporate unsupervised information into this Logician-Orator dual learning structures for further improvement. Lastly, one can also interested in developing task-oriented rewards for adapting the agent to a specific task, for example, the answer generation task for open-domain KBQA system.

S	大宗货物吞吐量均保持两位数增长，其中铁矿石吞吐量突破4500万吨，木材吞吐量突破600万立方，均创历史新高，成为全国进口木材第一大港。	
S in English	The throughput of all integrated cargoes kept a double-digit growth. Among them, the throughput of iron ore exceeded 45 million tons and the throughput of timber exceeded 6 million cubic meters, all of which hit record highs, and became the country's largest port of timber imports.	
	Base Models	Dual Models
$S \xrightarrow{\text{Logician}} \mathcal{F}^*$	(大宗货物吞吐量, 保持, 两位数增长) (铁矿石吞吐量, 突破, 4500万吨) (木材吞吐量, 突破, 600万立方) (创历史, DESC, 新高) (__, 成为, 全国进口木材第一大港)	(大宗货物吞吐量, 均保持, 两位数增长) (铁矿石吞吐量, 突破, 4500万吨) (木材吞吐量, 突破, 600万立方) (__, 均创, 历史新高) (__, 成为, 全国进口木材第一大港)
$S \xrightarrow{\text{Logician}} \mathcal{F}^*$ in English	(The throughput of all integrated cargoes, kept, double-digits growth) (Throughput of iron ore, exceeded, 45 million tons) (Throughput of wood, breakthrough, 6 million cubic meters) (Hit historical, DESC, new high) (__, become, the country's largest port of timber imports)	(The throughput of integrated cargoes, all kept, double-digit growth) (Throughput of iron ore, exceeded, 45 million tons) (Throughput of timber, exceeded 6, million cubic meters) (__, all hit, record highs) (__, become, the country's largest port of timber imports)
$S \xrightarrow{\text{Logician}} \mathcal{F}^*$ $\mathcal{F}^* \xrightarrow{\text{Orator}} \mathcal{S}^{**}$	大宗货物吞吐量保持两位数增长，突破4500万吨，突破600万立方，新高，成为全国进口木材第一大港。	大宗货物吞吐量均保持两位数增长，铁矿石吞吐量突破4500万吨，木材吞吐量突破600万立方，均创历史新高，成为全国进口木材第一大港。
$S \xrightarrow{\text{Logician}} \mathcal{F}^*$ $\mathcal{F}^* \xrightarrow{\text{Orator}} \mathcal{S}^{**}$ in English	The throughput of all integrated cargoes kept double-digit growth, breaking 45 million tons, breaking 6 million cubic meters, a new high, became the country's largest port of timber imports.	The throughput of all integrated cargoes kept double-digit growth. The throughput of iron ore exceeded 45 million tons and the throughput of timber exceeded 6 million cubic meters, all hit a record high and became the country's largest port of timber imports.

Table 7: Examples for sentence reconstruction.

\mathcal{F}	(营业法人, 具备, 中专 (高中) 以上学历) (营业法人, 有, 一定的[管理经营]能力) (营业法人, 具备, [较强的服务理念 团队合作精神])	
\mathcal{F} in English	(Business legal person, is with, technical secondary school (high school) or above) (Business legal person, has, certain [managerial operational] ability) (Business legal person, is with, [strong philosophy of service teamwork spirit])	
	Base Models	Dual Models
$\mathcal{F} \xrightarrow{\text{Orator}} \mathcal{S}^*$	营业法人具备中专 (高中) 以上学历, 有一定的管理和经营能力, 具备较强的服务理念和团队合作精神。	营业法人具备中专 (高中) 以上学历, 营业法人有一定的管理和经营能力, 具备较强的服务理念和团队合作精神。
$\mathcal{F} \xrightarrow{\text{Orator}} \mathcal{S}^*$ in English	Business legal person is with secondary (high school) or above, has certain managerial and operational capabilities, and is with strong philosophy of service and teamwork spirit.	Business legal person is with secondary (high school) or above, business legal person has certain managerial and operational capabilities, and is with strong philosophy of service and teamwork spirit.
$\mathcal{F} \xrightarrow{\text{Orator}} \mathcal{S}^*$ $\mathcal{S}^* \xrightarrow{\text{Logician}} \mathcal{F}^{**}$	(营业法人具备中专 (高中) 以上学历, 有, 一定的[管理经营]能力) (营业法人具备中专的[管理经营]能力, 具备, 较强的[服务理念 团队合作])	(营业法人, 具备, 中专 (高中) 以上学历) (营业法人, 有, [一定的管理经营能力]) (营业法人, 具备, 较强的[服务理念 团队合作])
$\mathcal{F} \xrightarrow{\text{Orator}} \mathcal{S}^*$ $\mathcal{S}^* \xrightarrow{\text{Logician}} \mathcal{F}^{**}$ in English	(Business legal person with technical secondary school (high school) or above, has, certain [managerial operational] ability) (Business legal person with technical secondary school (high school) or above, is with, [strong philosophy of service teamwork spirit])	(Business legal person, is with, technical secondary school (high school) or above) (Business legal person, has, [certain managerial operational ability]) (Business legal person, is with, [strong philosophy of service teamwork spirit])

Table 8: Examples for fact reconstruction.

References

- Shubham Agarwal and Marc Dymetman. 2017. A Surprisingly Effective Out-of-the-Box Char2char Model on the E2E NLG Challenge Dataset. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, August, pages 158–163.
- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. DBpedia: A Nucleus for a Web of Open Data. *The Semantic Web*, 4825 LNCS:722–735.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. In *International Conference on Learning Representations*.
- 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. *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: a Collaboratively Created Graph database for Structuring Human Knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, pages 1247–1250. ACM.
- Wanxiang Che, Zhenghua Li, and Ting Liu. 2010. LTP: A Chinese Language Technology Platform. In *Proceedings of the Coling*, pages 13–16.
- Andrew Chisholm, Will Radford, and Ben Hachey. 2017. Learning to Generate One-sentence Biographies from Wikidata. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, volume 1, pages 633–642.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734.
- Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 2011. An Analysis of Open Information Extraction based on Semantic Role Labeling. In *Proceedings of the Sixth International Conference on Knowledge Capture*, pages 113–120. ACM Press.
- Janara Christensen, Mausam, Stephen Soderland, Oren Etzioni, Mausam, Stephen Soderland, and Oren Etzioni. 2013. Towards Coherent Multi-Document Summarization. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Section 3, pages 1163–1173.
- Janara Christensen, Stephen Soderland, and Gagan Bansal. 2014. Hierarchical Summarization: Scaling Up Multi-Document Summarization. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 902–912.
- Nan Duan, Duyu Tang, Peng Chen, and Ming Zhou. 2017. Question Generation for Question Answering. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 866–874.
- Oren Etzioni, Anthony Fader, Janara Christensen, Stephen Soderland, and Mausam. 2011. Open Information Extraction: The Second Generation. In *Proceeding of International Joint Conference on Artificial Intelligence*, pages 3–10.
- Anthony Fader, Luke S Zettlemoyer, and Oren Etzioni. 2014. Open Question Answering Over Curated and Extracted Knowledge Bases. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1156–1165.
- Federico Gaspari. 2006. Look Who’s Translating: Impersonations, Chinese Whispers and Fun with Machine Translation on the Internet. In *EAMT-2006: 11th Annual Conference of the European Association for Machine Translation*, pages 19–20.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. Incorporating Copying Mechanism in Sequence-to-Sequence Learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1631–1640.
- Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. 2012. Overview of Mini-batch Gradient Descent. Technical report.
- Nanda Kambhatla. 2004. Combining Lexical, Syntactic, and Semantic Features with Maximum Entropy Models for Extracting Relations. In *Proceedings of the ACL 2004 on Interactive Poster and Demonstration Sessions*.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2017. Answering Complex Questions Using Open Information Extraction. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 311–316.
- 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*, pages 146–157. Association for Computational Linguistics.
- C Y Lin. 2004. Rouge: A Package for Automatic Evaluation of Summaries. In *Proceedings of the 2004 ACL Workshop on Text Summarization Branches Out*, 1, pages 25–26.

- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural Relation Extraction with Selective Attention over Instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 2124–2133.
- Mausam. 2016. Open Information Extraction Systems and Downstream Applications. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 4074–4077.
- T Mikolov, M Karafiat, L Burget, J Cernocky, and S Khudanpur. 2010. Recurrent Neural Network based Language Model. In *Proceedings of Interspeech*, September, pages 1045–1048.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant Supervision for Relation Extraction without Labeled Data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language*, volume 2, page 1003, Morristown, NJ, USA. Association for Computational Linguistics.
- Makoto Miwa and Mohit Bansal. 2016. End-to-End Relation Extraction using LSTMs on Sequences and Tree Structures. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1105–1116, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Harinder Pal and Mausam. 2016. Demonyms and Compound Relational Nouns in Nominal Open IE. In *Proceedings of the 5th Workshop on Automated Knowledge Base Construction*, pages 35–39.
- John W Ratcliff and David E Metzener. 1988. Pattern Matching: The Gestalt Approach. *Dr Dobbs's*, 13(7).
- Mrinmaya Sachan and Eric Xing. 2018. Self-Training for Jointly Learning to Ask and Answer Questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pages 629–640.
- Michael Schmitz, Robert Bart, Stephen Soderland, and Oren Etzioni. 2012. Open language learning for information extraction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 523–534.
- Tomohiro Shigenobu. 2007. Evaluation and Usability of Back Translation for Intercultural Communication. In *International Conference on Usability and Internationalization*, pages 259–265. Springer.
- Gabriel Stanovsky, Ido Dagan, and Mausam. 2015. Open IE as an Intermediate Structure for Semantic Tasks. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 303–308.
- Mingming Sun, Xu Li, Xin Wang, Miao Fan, Yue Feng, and Ping Li. 2018. Logician: A Unified End-to-End Neural Approach for Open-Domain Information Extraction. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, February, pages 556–564, New York, New York, USA. ACM Press.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to Sequence Learning with Neural Networks. In *Advances in Neural Information Processing Systems 27*, volume 155, page 9.
- Richard S. Sutton, David Mcallester, Satinder Singh, and Yishay Mansour. 1999. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems 12*, pages 1057–1063.
- Duyu Tang, Nan Duan, Tao Qin, Zhao Yan, and Ming Zhou. 2017. Question answering and question generation as dual tasks. *arXiv preprint arXiv:1706.02027*.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling Coverage for Neural Machine Translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 76–85.
- Menno Van Zaanen and Simon Zwarts. 2006. Unsupervised Measurement of Translation Quality using Multi-engine, Bi-directional Translation. In *Australasian Joint Conference on Artificial Intelligence*, pages 1208–1214. Springer.
- Pavlos Vougiouklis, Hady Elsahar, Lucie-Aimée Kaffee, Christoph Gravier, Frederique Laforest, Jonathon Hare, and Elena Simperl. 2017. Neural Wikipedian: Generating Textual Summaries from Knowledge Base Triples. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*.
- Wikipedia. 2017. Assignment problem— Wikipedia, The Free Encyclopedia.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. Challenges in Data-to-Document Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2243–2253.
- Yingce Xia, Di He, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. Dual Learning for Machine Translation. In *Advances in Neural Information Processing Systems 29*, pages 1–9.
- Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, and Tie-Yan Liu. 2017. Dual Supervised Learning. In *Proceedings of 34th International Conference on Machine Learning*.
- Jun Yin, Xin Jiang, Zhengdong Lu, Lifeng Shang, Hang Li, and Xiaoming Li. 2016. Neural Generative Question Answering. In *Proceedings of*

2016 NAACL Human-Computer Question Answering Workshop, pages 36–42.

Dmitry Zelenko, Chinatsu Aone, Anthony Richardella, Jaz Kandola, Thomas Hofmann, Tomaso Poggio, and John Shawe-Taylor. 2003. Kernel Methods for Relation Extraction. *Journal of Machine Learning Research*, 3:1083–1106.

Suncong Zheng, Feng Wang, Hongyun Bao, Yuexing Hao, Peng Zhou, and Bo Xu. 2017. Joint Extraction of Entities and Relations Based on a Novel Tagging Scheme. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1227–1236.