

Fast Image Caption Generation with Position Alignment

Zhengcong Fei^{1,2}

¹Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS),
Institute of Computing Technology, CAS, Beijing 100190, China

²University of Chinese Academy of Sciences, Beijing 100049, China
feizhengcong@ict.ac.cn

Abstract

Recent neural network models for image captioning usually employ an encoder-decoder architecture, where the decoder adopts a recursive sequence decoding way. However, such autoregressive decoding may result in sequential error accumulation and slow generation which limit the applications in practice. Non-autoregressive (NA) decoding has been proposed to cover these issues but suffers from language quality problem due to the indirect modeling of the target distribution. Towards that end, we propose an improved NA prediction framework to accelerate image captioning. Our decoding part consists of a position alignment to order the words that describe the content detected in the given image, and a fine non-autoregressive decoder to generate elegant descriptions. Furthermore, we introduce an inference strategy that regards position information as a latent variable to guide the further sentence generation. The Experimental results on public datasets show that our proposed model achieves better performance compared to general NA captioning models, while achieves comparable performance as autoregressive image captioning models with a significant speedup.

1 Introduction

Image captioning aims to automatically describe the content of a given image through natural language (Kulkarni et al. 2011; Donahue et al. 2015). Inspired by the development of neural machine translation (NMT), recent approaches to image captioning under a general encoder-decoder framework have achieved great success (Vinyals et al. 2015; Xu et al. 2015; 2016; Anderson et al. 2017). In such a framework, an image encoder which is based on a convolutional neural network (CNN) is first used to extract region-level visual feature vectors for a given image, a caption decoder which is based on a recurrent neural network (RNN) is then adopted to generate caption words recurrently. Despite their success, the state-of-the-art image caption models usually suffer from the slow inference speed, which has become a bottleneck to apply deep learning-based models in real-time systems (Ranzato et al. 2016). The slow inference speed of image caption models is due to their autoregressive property, *i.e.*, decoding the target sentence word-by-word according to the given image and generated substance.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

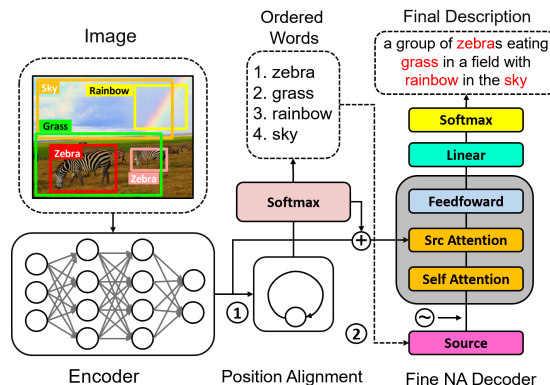


Figure 1: The framework of our proposed NA captioning model. ① Different from general NA decoding models, our model adds a position alignment module between the image encoder module and sentence decoder module to explicitly construct object and position information. ② For a general NA model, the decoder inputs are the copied of source information, but for our model, the decoder inputs are ordered words with image information.

Recently, a non-autoregressive (NA) technology (Gu et al. 2017) has been introduced in NMT which can simultaneously decode all target words to break the bottleneck of the autoregressive generation models. Non-autoregressive sentence generation models usually directly copy the source word representations to the input of the decoder, instead of using previous predicted target word representations (Wei et al. 2019; Guo et al. 2018). Hence, the inference of different target words are independent, which enables parallel computation of the sentence decoder. Non-autoregressive decoding models could achieve 10-15 times speedup compared to autoregressive models while maintaining considerable performance (Gu et al. 2017).

However, existing non-autoregressive systems ignore the dependencies among target words and simultaneously generate all target words, which makes the generated sentences have apparent problems in naturalness and fluency. Specif-

ically, when decoding a target word, the non-autoregressive models will build an indirect distribution of the true target. Consequently, this issue makes non-autoregressive models generate sentences conditioned on less or inaccurate source information, thus leading to missing, repeated and even wrong sentences. Moreover, in an image captioning task, the input of an encoder is a given image, which is a low-level semantic data rather than a high-level semantic sequence in NMT. This brings more difficult for decoder to generation accuracy sentences.

In this paper, to address the above issues, we propose a novel non-autoregressive framework, named *Fast Neural Image Caption* (FNIC), which integrates the position information to guide the generation of descriptions. Specifically, as shown in Figure 1, we first adopt an image encoder, which extracts feature maps from CNNs or object features from object detection models, then maps the dimension of input visual features to the model dimension of the decoder using a multi-layer perceptron (MLP). As for the decoder, we utilize position alignment to generate and order the words that roughly describe the content of the image, and then fine the organized words into the final sentence. We further introduce a decoding inference strategy which adopts position information to guide the sentence generation. Rather than just base on roughly generated words, we utilize the conditional distribution of these ordered words as a latent variable to guide the decoding of the target sentence. Intuitively, with the guide of position information, for each target word, what it is expressed can be nearly limited to the corresponding word of coarse ordered words in the same position.

The experimental results on different widely-used public benchmarks show that our proposed captioning model achieves significant and consistent improvements compared to general non-autoregressive models through explicitly adopting the position information to guide the decoding. Moreover, by introducing a simple but effective autoregressive decoder to construct positive information and generate coarse words, our model immensely narrows the description quality gap between autoregressive and NA sentence generation models, while maintains a considerable speedup (nearly 8 times faster). We will release all source codes and related resources of this work for further research explorations.

Our contributions are summarized as follows:

- We introduced non-autoregressive technology into the captioning task, which effectively accelerates the description generation while maintains a comparable performance.
- To improve the quality of generated sentences and overcome the semantic gap between image and text, we design a light position alignment module based on object detection features to produce ordered words.
- Rather than just base on roughly generated ordered words, we utilize the conditional distribution of these ordered words as a latent variable to guide the decoding of the target sentence.

2 Background

2.1 Image Caption

In an image captioning task, a stand encoder-decoder framework first encodes the given image I to some visual feature vectors I_F and then generates descriptive sentence S according to previously generated words and image features (Vinyals et al. 2015; Fang et al. 2015; Donahue et al. 2015; Lu et al. 2016a). It corresponds to the word-by-word nature of human language generation and effectively learns the conditional distribution of real captions. Formally, autoregressive models consider the distribution over possible output sentences S as a chain of conditional probabilities,

$$P(S|I_F) = \prod_{i=1}^n P(w_i|w_{<i}, I_F) \quad (1)$$

In the training procedure, the model parameters are trained to minimize the conditional cross-entropy loss, where the loss function is defined as:

$$\mathcal{L}_A = - \sum_{i=1}^n \log P(w_i|w_{<i}, I_F) \quad (2)$$

During inference procedure, at each step, a word is produced and then fed into the decoder to predict the next word. Therefore, all autoregressive decoders must remain sequential rather than parallel when generating sentences. Moreover, sequential decoding is tend to copy words from the training data to enhance the grammatical accuracy, which easily causes semantic error and limits the diversity of descriptions (Jiang et al. 2018).

2.2 Non-autoregressive Decoder

Non-autoregressive neural decoder is first proposed by (Gu et al. 2017) to solve the slow generation issue in autoregressive neural machine translation, which could simultaneously generate target words by removing their dependencies. Formally,

$$P(S|I_F) = \prod_{i=1}^n P(w_i|I_F) \quad (3)$$

Instead of utilizing the previous subsentence history, NA models usually directly copy the source feature vectors as the input of the decoder. Hence, when generating a sentence, NA decoder could predict all target words with their maximum likelihood individually by breaking the dependency among the target words, and therefore the decoding procedure of NA models is in parallel and gains very fast generation speed.

However, since NA decoder discards the sequential dependencies among words in the target sentence, it suffers from the potential performance degradation due to the limited information for decision making. To be specific, when decoding a target word, the NA model must be able to figure out not only what target-side information does the word describe but also what is expressed by other target words (Wei et al. 2019; Guo et al. 2018). With the limited information to make predictions, NA models cannot effectively learn the intricate patterns from source information to target

sentences, which leads to inferior language quality. Therefore, the position alignment module is added to increase the location information and reduce the candidate search space.

3 Method

In this section, we detail our proposed FNIC architecture which integrates position alignment structure into decoder, pursuing an improved non-autoregressive sentence generation to accelerate image captioning. We first describe the popular image encoder types, then show how we implement a hierarchical decoder which consists of a position alignment and a fine sentence decoder, to realize a non-autoregressive decoding procedure. Finally, we introduce a non-deterministic probability inference process as well as the model training process.

3.1 Image Encoder

The image encoder in general encoder-decoder framework aims to extract a set of image feature vectors $I_F = \{v_1, v_2, \dots, v_k\}$ for different regions based on given image I , where k is the number of regions. A typical image encoder usually adopts a CNN (e.g. ResNet (He et al. 2016)) to extract features. Moreover, R-CNN based models (e.g. Faster RCNN (Ren et al.)) are employed to improve the captioning performance which utilizes bottom-up attention (Anderson et al. 2017) and provides a better understanding of objects in the image.

To be specific, for a CNN based encoder, a spatially adaptive max-pooling layer is adopted after a classical CNN structure to fix the size of the output vector and streamlining information. As for an R-CNN based encoder, which usually detects objects in two stages. The first stage predicts region proposals, and the second stage predicts class labels as well as bounding box refinements. To obtain a feature vector set, we select all regions in which any class detection probability exceeds a confidence threshold and apply mean-pooling on the convolutional features for each of vectors.

In most cases, we utilize a linear transformation to project all image feature vectors in I_F to a fixed size of d before input to the decoder module.

3.2 Hierarchical Decoder

Here, we will introduce a novel non-autoregressive decoder, which aims to accelerate sentence generation while keeps comparable performance. As shown in Figure 1, our proposed FNIC employs a position alignment module to explicitly model the position information in the decoding. Formally, FNIC first transforms the image feature I_F into a coarse ordered words $X = \{x_1, x_2, \dots, x_m\}$ which the order of coarse words appears is the same as that in the target sentence, and then expands the ordered words to a target sentence S . Our proposed decoder handle the overall image captioning probability as:

$$P(S|I_F) = \sum_X P(X|I_F)P(S|X, I_F) \quad (4)$$

where $P(X|I_F)$ is realized by the position alignment module and $P(S|X, I_F)$ is implemented by the fine decoder module.

Position Alignment The position alignment module determines the order of words corresponding to the target sentence, which represents the objects and scenes in given image I , by learning to transform the image feature I_F into the ordered words X .

Since we conduct an experiment and find that autoregressive models are more suitable for modeling the position information compared to NA models, and even a light autoregressive model with similar decoding speed to a large NA model could achieve better performance in modeling position information. Hence, we integrate a light autoregressive model to build the coarse ordered words probability as:

$$P(X|I_F) = \prod_{i=1}^m P(x_i|x_{<i}, I_F) \quad (5)$$

where $x_{<i} = \{x_1, x_2, \dots, x_{i-1}\}$ represents the ordered word history.

In practical terms, we use a greedy search algorithm to generate ordered words word-by-word. When decoding the i -th ordered word x_i , we obtain its hidden representation as:

$$R_i = \text{GRU}(R_{i-1}||[C_{i-1}; X_{<i}]) \quad (6)$$

where $\text{GRU}(\cdot)$ is a one-layer gated recurrent unit (Cho et al. 2014) and C_{i-1} is calculated by the inter-attention to the extracted image features. After obtaining the hidden representation, the conditional probability can be calculated as,

$$P(x_i|I_F) = \text{softmax}(R_i) \quad (7)$$

Fine Decoder The fine decoder module generates the target sentence with the guiding of coarse ordered words, which regards the generation of each word as non-autoregressive:

$$P(S|X, I_F) = \prod_{i=1}^n P(w_i|X, I_F) \quad (8)$$

Different from the original NA decoder, the inputs of our decoder module is the ordered words plus image features instead of just copied source information, which is used to guide the decoding direction.

In actual work, we employ an NA model with N layers of Transformer decoder which is composed of a multi-head self-attention layer, a multi-head inter attention layer and a feed-forward layer.

$$D(Y, I_F) = \text{FFN}(\text{ATT}_{inter}(I_F, \text{ATT}_{self}(Y))) \quad (9)$$

where $\text{FFN}(\cdot)$ is the feed-ward layer and Y is the input of fine decoder.

3.3 Inference Procedure

Our FNIC model explicitly builds position information of non-autoregressive and aims to utilize it to provide effective limitation information and assist the generation of target sentences. Now the remaining problem is how to perform decoding with the guide of ordered word information. We propose to utilize the coarse ordered words as a bridge

to guide the decoding of the target sentence, which can be formulated as:

$$\begin{aligned} S^* &= \operatorname{argmax}_S P(S|I_F) \\ &= \operatorname{argmax}_S \sum_X P(S|X, I_F)P(X|I_F) \end{aligned} \quad (10)$$

It is impractical to obtain a direct and accurate solution for maximizing Equation 10 due to the high time complexity. One easy way to think about it is that first generates the most probable ordered words based on image features and then decodes the target sentence conditioned on it:

$$X^* = \operatorname{argmax}_X P(X|I_F) \quad (11)$$

$$S^* = \operatorname{argmax}_S P(S|X^*, I_F) \quad (12)$$

This two-stage local optimal approach is simple and effective, but it brings in some noise in the approximation and can not guarantee a global optimum solution.

Inspired from *Hidden Markov Model* (HMM), different from the local deterministic optimal strategy which utilizes a deterministic ordered word set to guide, we introduce a new inference strategy, regards the probability distribution Q of the ordered words as a latent variable, and models the decoder procedure as generating the target sentence according to the latent variable Q , *i.e.*, Equation 10 is re-formulated as:

$$S^* = \operatorname{argmax}_S P(S|Q, I_F) \quad (13)$$

where the probability distribution Q is defined as:

$$Q(X) = P(X|I_F) = \frac{\exp(f(X))}{\sum_{X'} \exp(f(X'))} \quad (14)$$

where $f(\cdot)$ is a score function of ordered words (the input of the softmax layer in the fine decoder). Since the latent variable Q can be viewed as a non-deterministic form of the guiding ordered words, the sentence generation with the non-deterministic strategy is also guided by the hidden ordered words.

To be specific, we use the weighted word embeddings of the conditional probability Q as the input of the decoder, then obtain the final probability:

$$P(S|Q, I_F) = \operatorname{softmax}(D^N(Q^T \operatorname{Emb}(S), I_F)) \quad (15)$$

where $D(\cdot)^N$ is N -layer fine decoder in Equation 9. Please note that the major difference between local deterministic and non-deterministic probability strategy is the inputs of fine decoder module (Location ② in Figure 1), where the local deterministic strategy directly utilizes the word embeddings of generated coarse ordered words and the non-deterministic probability strategy utilizes the weighted word embeddings of the word probability of position information.

3.4 Training Procedure

In the training process, for each training sentence pair (I, S) in data set \mathbb{D} , we first utilize image encoder to extract salient object labels, and then generate its corresponding candidate ordered word set \hat{X} : we adopt a word position alignment tool to select the words that appear in both the target sentence I and the object labels, and then sort the selected

words according to the order in the target sentence. At last, FNIC is optimized by maximizing a joint loss:

$$\mathcal{L} = \mathcal{L}_P + \mathcal{L}_F \quad (16)$$

where \mathcal{L}_P and \mathcal{L}_F represent the position alignment and sentence generation losses respectively. Formally, based on non-deterministic probability approach, the position alignment loss \mathcal{L}_P is defined as:

$$\mathcal{L}_P = \sum_{(I, \hat{X}, S) \in \mathbb{D}} \log P(\hat{X}|I_F) \quad (17)$$

And the fine sentence generation loss \mathcal{L}_F is defined as an overall maximum likelihood of decoding the target sentence from the conditional probability of ordered words:

$$\mathcal{L}_F = \sum_{(I, \hat{X}, S) \in \mathbb{D}} \log P(S|Q, I_F) \quad (18)$$

In particular, we use the trained model for the local deterministic optimal approach to initialize the model before the non-deterministic latent variable approach, since if Q is not well trained, \mathcal{L} will converge very slowly.

4 Experiments

4.1 Dataset

COCO is a standard benchmark for image captioning task which contains 123,287 images (82,783 for training and 40,504 for validation) and each image is annotated with 5 descriptions by humans. Since the annotated descriptions of the official testing set are not provided, we utilize Karpathy split (113,287 for training, 5,000 for validation and 5,000 for testing) as in (Anderson et al. 2017). According to (Karpathy and Fei-Fei 2016), all the training sentences are converted to lower case and we omit rare words which occur less than 5 times. Therefore, the final vocabulary includes 10,201 unique words.

Visual Genome, which contains images with annotated objects, attributes, and relationships, is adopted to train Faster R-CNN for object detection. In this paper, we follow the setting in (Anderson et al. 2017; Yao et al. 2018) and take 98,077 images for training, 5,000 for validation, and 5,000 for testing. As in (Anderson et al. 2017), 1,600 objects and 400 attributes are considered from Visual Genome for training Faster R-CNN with two branches for predicting objects and attribute classes.

4.2 Baselines

Since the non-autoregressive method in the image captioning task has been underinvestigated so far, there are few existing baselines for our comparison. We choose to compare with the following typical autoregressive models for their high relevance with our task as well as their superior performance demonstrated by previous works:

(1) **SCST** (Rennie et al. 2017) employs a self-critical sequence training strategy to train a modified visual attention-based captioning model in (Xu et al. 2015) (2) **ADP-ATT** (Lu et al. 2016b) develops an adaptive attention

Table 1: Performance comparisons with different evaluation metrics on the COCO Karpathy test split. All values except Speed up are reported as percentage (%).

	BLEU@4	METEOR	ROUGE-L	CIDEr-D	SPICE	Speed up
SCST (Rennie et al. 2017)	30.0	25.9	53.4	99.4	-	-
ADP-ATT (Lu et al. 2016b)	33.2	26.6	-	108.5	-	-
LSTM-A (Yao et al. 2016)	35.2	26.9	55.8	108.8	20.0	-
Up-Down (Anderson et al. 2017)	36.2	27.0	56.4	113.5	20.3	-
GCN-LSTM (Yao et al. 2018)	37.0	28.1	57.1	117.1	21.1	-
NAIC	28.5	23.6	52.3	98.2	18.5	12.40 ×
FNIC _{AT}	36.6	27.2	56.2	118.0	20.5	1.00×
FNIC _{NAT}	30.4	25.4	55.1	107.4	19.6	11.21×
FNIC	36.2	27.1	55.3	115.7	20.2	8.15×

based encoder-decoder model for automatically determining whether to attend to the image and which image regions to focus for caption. (3) **LSTM-A** (Yao et al. 2016) integrates semantic attributes into CNN plus RNN captioning model for boosting image captioning. (4) **Up-Down** (Anderson et al. 2017) designs a combined bottom-up and top-down attention mechanism that enables region-level attention to be calculated to boost image caption. (5) **GCN-LSTM** (Yao et al. 2018) further exploits visual relationships between objects through graph convolutional networks.

We also select three models as baselines: (1) **NAIC** develops only non-autoregressive fine decoder (one-layer Transformer decoder) without position alignment. (2) **FNIC_{AT}** adopts autoregressive position alignment (one-layer GRU) and autoregressive fine decoder (one-layer Transformer decoder). (3) **FNIC_{NAT}** utilizes both non-autoregressive position alignment (one-layer Transformer decoder) and non-autoregressive fine decoder (one-layer Transformer decoder). The model architectures of **FNIC_{AT}**, **FNIC_{NAT}** and our proposed **FNIC** are identical. Note that for fair comparison, all the baselines and our model adopt ResNet-101 as the basic architecture of image feature extractor. All models use the same training strategy and are evaluated using greedy decoding.

4.3 Implementation Details

For each image, we apply Faster R-CNN to detect objects within this image and select top $K = 36$ regions with highest detection confidences to represent the image. The dimension of each region is set as 2,048. For fine decoder module, we utilize a one-layer Transformer model ($D_{model}=512$, $D_{hidden}=512$, $n_{head}=2$, $dropout=0.1$). For the GRU position alignment module, we set it to have the same hidden size with the Transformer model. The captioning model is mainly implemented with PyTorch (version 1.1.0), optimized with Adam (Kingma and Ba 2014). We set the initial learning rate as 0.0005 and the mini-batch size as 1,024. The momentum and the weight-decay are 0.8 and 0.999 respectively. The maximum training iteration is set as 35 epochs. In addition, we write a position ranking tool based on python (version 3.6) to rank the object labels appear in the target sentence.

4.4 Automatic Evaluation

Metrics Following previous image captioning work (Yao et al. 2018; Anderson et al. 2017), we select five types of metrics: BLEU@N (Papineni et al. 2002), METEOR (Lavie and Agarwal 2007), ROUGE-L (Flick 2004), CIDEr-D (Vedantam, Zitnick, and Parikh 2015) and SPICE (Anderson et al. 2016). In particular, SPICE focuses on semantic analysis and has a higher correlation with human judgment, and other metrics favor frequent training n-grams and measure the overall sentence fluency, which are more preferred in sequential decoding based methods. For all the metrics, the larger they are, the more relevant or fluent the generated descriptions are. Speed up is computed based on the time to decode a single sentence without minibatching, and the values are averaged over the whole offline test set. The decoding is implemented on a single NVIDIA GeForce GTX 1080 Ti.

Overall Results and Analysis We compare our proposed model Fast Neural Image Caption (FNIC) with all baseline models. All the results are shown in Table 1. From the table, we can find that:

(1) FNIC achieves excellent performance on most of the benchmark metrics, which is even close to the best AT model with smaller than 1 BLEU gap (36.2 vs. 37.0). It is also worth mentioning that although FNIC utilizes a small GRU model to better capture position information, it could still maintain low sentence generation latency (about 8 speedup of **FNIC_{AT}** which only utilizes autoregressive decoder). Compared to **NAIC** which does not keep the position alignment module, our **FNIC** provides effective position and semantic information in the generated coarse ordered words, which limit the candidate word space in the source sentence generation and makes it faster. Another reason is that a major potential performance degradation of non-autoregressive models compared to autoregressive models comes from the difficulty of modeling the sentence structure difference between source information and target sentences, *i.e.*, position and semantic information, which is neglected for most of existing non-autoregressive models but can be well modeled by the small GRU position alignment decoder.

(2) A light GRU position alignment model with close latency to large non-autoregressive models could perform

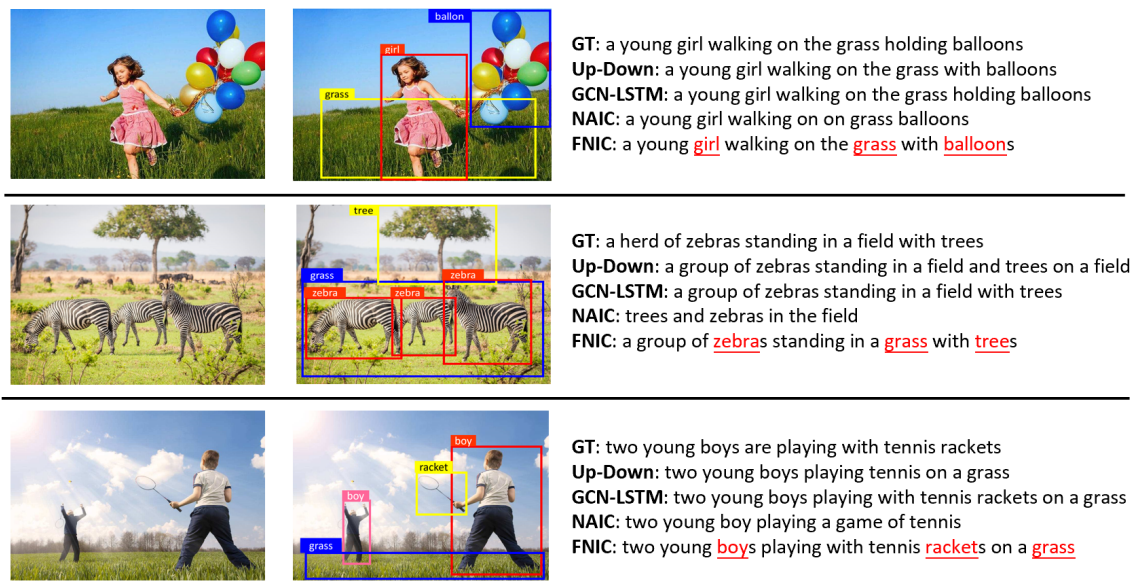


Figure 2: Image examples from (Lin et al. 2014) with object regions and sentence generation results. The output sentences are generated by (1) Ground Truth (GT): One ground truth sentence, (2) Up-Down (3) GCN-LSTM (4) NAIC and (5) our proposed FNIC.

Table 2: Effect of non-deterministic inference strategy on COCO Karpathy test split. DT represents local deterministic inference and NDT represents non-deterministic probability inference. All values are reported as percentage (%).

	BLEU@4	METEOR	ROUGE-L	CIDEr-D	SPICE
FNIC _{NAT} +DT	29.0	24.9	53.4	106.2	19.2
FNIC _{NAT} +NDT	30.4	25.4	55.1	107.4	19.6
FNIC+DT	35.6	26.6	53.2	113.5	19.5
FNIC+NDT	36.2	27.1	55.3	115.7	20.2

much better in modeling position information. On all benchmark metrics, **FNIC** with a small autoregressive GRU position alignment module achieves a much better description quality than that with a large non-autoregressive model (one-layer Transformer decoder has a larger parameter scale than one-layer GRU). Moreover, we find that our proposed model **FNIC** with a one-layer GRU for decoding could even outperform some of the existing autoregressive models such as **SCST**, **ADP-ATT** and **LSTM-A** in most metrics, while maintains acceptable latency. These verify that ordered words modeled by position alignment module could effectively reduce the decoding space and improve the generated sentence quality of the model.

4.5 Effect of Non-deterministic Inference Strategy

We also investigate the effect of two proposed inference strategies including local deterministic and non-deterministic probability on COCO Karpathy test split. In Table 2, we can find that the non-deterministic probability

strategy has better performance compared to the local deterministic inference strategy for both **FNIC** and **FNIC_{NAT}** since the non-deterministic probability strategy could effectively reduce the information loss of the deterministic strategy. On the other hand, we can also find that the non-deterministic probability strategy does not bring many improvements for our best model with the GRU position alignment module. One reason is perhaps that the coarse ordered words generated by the GRU position alignment module are good enough and therefore it does not bring in much information loss in the whole decoding procedure. Therefore, we use non-deterministic inference as the default decoding strategy in the proposed model.

4.6 Case Study

Figure 2 shows example generated descriptions of top-ranking autoregressive model **Up-Down**, **GCN-LSTM**, general non-autoregressive captioning model **NAIC** and our proposed model **FNIC**. We find that the problem of missing objects and repeated descriptions are severe in the general

Table 3: Illustration of the diversity of different methods from different aspects. All values are reported as percentage (%).

	Novel	Unique	Vocabulary Usage
Up-Down	45.05	61.58	7.97
GCN+LSTM	60.23	83.22	8.64
FNIC _{AT}	55.12	78.42	8.33
FNIC	81.25	87.12	12.16

NA model (Such as repeated preposition “on” in the second image and missing object “grass” in third image), while this problem is effectively alleviated in **FNIC** model. Moreover, we find that most of the missing, or wrong word in the generated sentence of general NA model comes from the errors in the coarse ordered words, which demonstrates that the NA model could well transform the coarse ordered words which have the target sentence structure to the final fine decoder.

4.7 Diversity Study

Diversity is an important indicator of image captioning tasks. Since the non-autoregressive decoding approach is a visual-to-linguistic transformation process and can provide more semantic information, it is distinctive to autoregressive decoding that more diverse captions are expected to be generated.

To analyze the diversity of generated captions, according to (Dai, Fidler, and Lin 2019), we compute **novel** caption percentage, **unique** caption percentage and **vocabulary usage**, which respectively account for the percentage of captions that have not been seen in the training data, the percentage of captions that are unique in the whole generated captions and the percentage of words in the vocabulary that are adopted to generate captions. According to Table 3, it is obvious that **FNIC** can achieve better results, which suggest that NA methods can generate more diverse captions than sequential methods, such as **TopDown** and **FNIC_{AT}**.

5 Conclusion

In this paper, we find that the limited information for each word generation is a important factor leading to the poor performance of NA decoding. To cover this problem, we propose a fast and effective image captioning framework named **FNIC** which explicitly integrates the extra information through a light GRU position alignment module in the decoding procedure. Furthermore, we introduce non-deterministic probability inference strategies to utilize the coarse ordered words to narrow the candidate decoding search space and guide the fine sentence generation in the next step. Experimental results on public datasets show that our NA model achieves better performance than general NA captioning models, and keeps comparable sentence quality as popular autoregressive model while keeps a significant speedup. We believe that well model the position and semantic information is a potential way for better NA in the future.

6 Acknowledgments

We would like to thank Haomiao Sun for helpful discussions.

References

- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. Spice: Semantic propositional image caption evaluation. *Adaptive Behavior*.
- Anderson, P.; He, X.; Buehler, C.; Teney, D.; and Lei, Z. 2017. Bottom-up and top-down attention for image captioning and vqa.
- Cho, K.; van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *Computer Science*.
- Dai, B.; Fidler, S.; and Lin, D. 2019. A neural compositional paradigm for image captioning.
- Donahue, J.; Hendricks, L. A.; Guadarrama, S.; Rohrbach, M.; Venugopalan, S.; Darrell, T.; and Saenko, K. 2015. Long-term recurrent convolutional networks for visual recognition and description.
- Fang, H.; Gupta, S.; Iandola, F. N.; Srivastava, R. K.; Deng, L.; Dollár, P.; Gao, J.; He, X.; Mitchell, M.; and Platt, J. C. 2015. From captions to visual concepts and back.
- Flick, C. 2004. Rouge: A package for automatic evaluation of summaries.
- Gu, J.; Bradbury, J.; Xiong, C.; Li, V. O. K.; and Socher, R. 2017. Non-autoregressive neural machine translation.
- Guo, J.; Tan, X.; He, D.; Qin, T.; Xu, L.; and Liu, T. Y. 2018. Non-autoregressive neural machine translation with enhanced decoder input.
- He, K.; Zhang, X.; Ren, S.; and Jian, S. 2016. Deep residual learning for image recognition.
- Jiang, W.; Ma, L.; Jiang, Y. G.; Liu, W.; and Zhang, T. 2018. Recurrent fusion network for image captioning.
- Karpathy, A., and Fei-Fei, L. 2016. Deep visual-semantic alignments for generating image descriptions.
- Kingma, D., and Ba, J. 2014. Adam: A method for stochastic optimization. *Computer Science*.
- Kulkarni, G.; Premraj, V.; Dhar, S.; Li, S.; and Berg, T. L. 2011. Baby talk: Understanding and generating simple image descriptions.
- Lavie, A., and Agarwal, A. 2007. Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments.
- Lin, T. Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context.
- Lu, C.; Krishna, R.; Bernstein, M.; and Li, F. F. 2016a. Visual relationship detection with language priors.
- Lu, J.; Xiong, C.; Parikh, D.; and Socher, R. 2016b. Knowing when to look: Adaptive attention via a visual sentinel for image captioning.

- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W. J. 2002. Bleu: a method for automatic evaluation of machine translation.
- Ranzato, M.; Chopra, S.; Auli, M.; and Zaremba, W. 2016. Sequence level training with recurrent neural networks. In *Iclr*.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. Faster r-cnn: Towards real-time object detection with region proposal networks.
- Rennie, S. J.; Marcheret, E.; Mroueh, Y.; Ross, J.; and Goel, V. 2017. Self-critical sequence training for image captioning.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation.
- Vinyals, O.; Toshev, A.; Bengio, S.; and Erhan, D. 2015. Show and tell: A neural image caption generator.
- Wei, B.; Wang, M.; Zhou, H.; Lin, J.; and Sun, X. 2019. Imitation learning for non-autoregressive neural machine translation.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhutdinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. *Computer Science* 2048–2057.
- Xu, J.; Gavves, E.; Fernando, B.; and Tuytelaars, T. 2016. Guiding the long-short term memory model for image caption generation. In *IEEE International Conference on Computer Vision*.
- Yao, T.; Pan, Y.; Li, Y.; Qiu, Z.; and Tao, M. 2016. Boosting image captioning with attributes.
- Yao, T.; Pan, Y.; Li, Y.; and Mei, T. 2018. Exploring visual relationship for image captioning.