
BANG: Bridging Autoregressive and Non-autoregressive Generation with Large Scale Pretraining

Weizhen Qi^{1,2} Yeyun Gong^{*3} Jian Jiao^{*4} Yu Yan^{*4} Weizhu Chen^{*4} Dayiheng Liu^{*2,5} Kewen Tang⁴
Houqiang Li¹ Jiusheng Chen⁴ Ruofei Zhang⁴ Ming Zhou³ Nan Duan³

Abstract

In this paper, we propose **BANG**, a new pre-training model to Bridge the gap between Autoregressive (AR) and Non-autoregressive (NAR) Generation. AR and NAR generation can be uniformly regarded as to what extent previous tokens can be attended, and BANG bridges AR and NAR generation by designing a novel model structure for large-scale pretraining. The pretrained BANG model can simultaneously support AR, NAR and semi-NAR generation to meet different requirements. Experiments on question generation (SQuAD 1.1), summarization (XSum) and dialogue generation (PersonaChat) show that BANG improves NAR and semi-NAR performance significantly as well as attaining comparable performance with strong AR pretrained models. Compared with the semi-NAR strong baselines, BANG achieves absolute improvements of 14.01 and 5.24 in the overall scores of SQuAD 1.1 and XSum, respectively. In addition, BANG achieves absolute improvements of 10.73, 6.39 and 5.90 in the overall scores of SQuAD, XSUM and PersonaChat respectively compared with the strong NAR baselines.

1. Introduction

Various pretraining methods (Song et al., 2019; Lewis et al., 2019; Qi et al., 2020; Raffel et al., 2020; Zhang et al., 2019a) have been successfully applied in natural language generation. Most of the pretraining works are based on Transformer and designed with autoregressive (AR) language model. Transformer based pretraining models show consis-

tent improvements with larger model size and larger pretraining corpus. Although the autoregressive generation method achieves high-quality results in many tasks, its latency is a well-known limitation for online real-time usage.

Non-autoregressive (NAR) models (Gu et al., 2017; Lee et al., 2018; Ghazvininejad et al., 2019; Raffel et al., 2020; Zhang et al., 2019a) are proposed to reduce generation latency. Different from AR models which generate tokens sequentially, NAR models generate tokens in parallel. Compared to AR models, NAR models generally come with a much lower inference latency, but a decrease in accuracy. In order to balance latency and accuracy, semi-NAR generation models (Stern et al., 2019; Lee et al., 2018; Gu et al., 2019; Ghazvininejad et al., 2019) are proposed. However, most of the NAR and semi-NAR models focus on translation tasks rather than general natural language generation tasks, which are proved to significantly benefit from pretraining (Qi et al., 2020; Lewis et al., 2019). Some works (Guo et al., 2020b; Su et al., 2021) initialize their NAR models with pretrained natural language understanding model BERT (Devlin et al., 2018) for better performance. To the best of our knowledge, this paper proposes the first large-scale pretraining model designed for NAR and semi-NAR generation.

In this paper, we propose a new model named BANG¹ to bridge the gap between AR and NAR via pretraining a generative model. Specifically, we consider pretraining model using AR, semi-NAR and NAR objectives with different attention mechanisms, which decide what extent previous tokens can be attended to. Precisely, BANG is pretrained to predict each token with arbitrary length of previous golden tokens replaced with special token [MASK]s. For example, with complete previous golden tokens, BANG predicts the next token in the AR manner. With all previous tokens replaced by [MASK], BANG predicts the next token in the NAR manner.

For AR models, the training strategy of teacher-forcing is commonly used, which uses the golden tokens as previous context to predict the next token. For NAR models, [MASK] initialization (Ghazvininejad et al., 2019) or other methods

^{*}Equal contribution ¹University of Science and Technology of China, Hefei, China ²During Internship at MSRA ³Microsoft Research Asia, Beijing, China ⁴Microsoft, Redmond, USA ⁵Sichuan University, Chengdu, China. Correspondence to: Yeyun Gong <yegong@microsoft.com>.

¹<https://github.com/microsoft/BANG>

like encoder copy (Gu et al., 2017) and posterior distribution approximation (Shu et al., 2020) are applied. In BANG pretraining, we consider the previous context of golden and [MASK] tokens, with arbitrary golden tokens’ length and [MASK] tokens’ length. To achieve an efficient implementation for multiple arbitrary alternatives in a same output sequence, we propose a new structure named cross-stream visible n-stream self-attention, which can be used to train BANG with different golden and [MASK] tokens’ combinations. For usage on downstream tasks, the single pretrained BANG model can be directly finetuned for either vanilla AR models or vanilla NAR models. Additionally, BANG can also be finetuned for hybrid semi-NAR models, which support predicting tokens with arbitrary previous golden tokens or [MASK]. Concretely, for semi-NAR generation, BANG predicts the first several tokens one by one as a high-quality sub-sequence hint, then produces all the remaining tokens simultaneously.

Our main contributions are: 1) BANG bridges the gap between AR and NAR by considering arbitrary previous [MASK] length during large-scale pretraining. 2) BANG is pretrained using an efficient cross-stream visible n-stream decoder to realize parallelization. Given multiple arbitrary number of previous tokens replaced with [MASK], every token is trained to predict simultaneously at each time step. 3) BANG supports NAR, semi-NAR and AR finetuning to meet different requirements with the same pretrained model structure. 4) We pretrain BANG with 16GB English language corpora of Wikipedia and BookCorpus, and finetune it on 3 popular natural language generation tasks in AR, semi-NAR and NAR manners, respectively. For NAR and semi-NAR finetuning, BANG achieves significant performance improvements on all the tasks. For AR finetuning with the comparison to strong AR pretrained models, BANG can attain comparable performance.

2. Model Structure

In this section, we first introduce fundamental concepts in § 2.1, and then describe BANG model structure and implementation in § 2.2. Finally the finetuning strategies are given in § 2.3.

2.1. Preliminaries

Natural language generation typically refers to the procedure of generating an output sequence $Y = \{y_1, y_2, \dots, y_{|Y|}\}$ given an input sequence $X = \{x_1, x_2, \dots, x_{|X|}\}$. There are three common generative paradigms: AR, NAR and semi-NAR generation.

AR generation follows the conditional probability:

$$P(Y|X) = \prod_{t=1}^{|Y|} P(y_t|y_{<t}, X) \tag{1}$$

Each token y_t in the target sequence Y is predicted with the dependency of input sequence X and previous tokens $y_{<t}$. Vanilla Transformer realizes this target by shifting decoder inputs one position, each token attending to its previous tokens to predict the next token.

NAR generation follows the conditional probability:

$$P(Y|X) = \prod_{t=1}^{|Y|} P(y_t|X) \tag{2}$$

Each token y_t in target sequence Y is predicted with the dependency of input sequence X but no dependency of previous tokens. Each token in the target sequence should be predicted simultaneously. An intuitive strategy is to feed a list of [MASK] tokens as the decoder initialization (Ghazvininejad et al., 2019).

Semi-NAR generation follows the conditional probability:

$$P(Y|X) = \prod_{t=1}^{|Y|} P(y_t|Y_{c_t}, X) \tag{3}$$

Here, Y_{c_t} means visible context for y_t from target sequence Y . Y_{c_t} is designed differently in different semi-NAR models (Stern et al., 2019; Lee et al., 2018; Ghazvininejad et al., 2019; Gu et al., 2019).

2.2. BANG

In this section, we will introduce the overall model structure of BANG in § 2.2.1 and the proposed cross-stream visible n-stream self-attention in § 2.2.2.

2.2.1. OVERALL

BANG is a sequence to sequence model based on the Transformer Encoder-Decoder structure. It is composed of a multi-layer Transformer encoder with vanilla self-attention mechanism and a multi-layer Transformer decoder with proposed cross-stream visible n-stream self-attention mechanism. We consider the input sequence $X = \{x_1, x_2, \dots, x_{|X|}\}$ and output sequence $Y = \{y_1, y_2, \dots, y_{|Y|}\}$. On the encoder side, BANG firstly encodes X into a list of hidden states by stacked vanilla Transformer encoder layers:

$$H_{enc} = \mathbf{Encoder}(x_1, \dots, x_{|X|}), \tag{4}$$

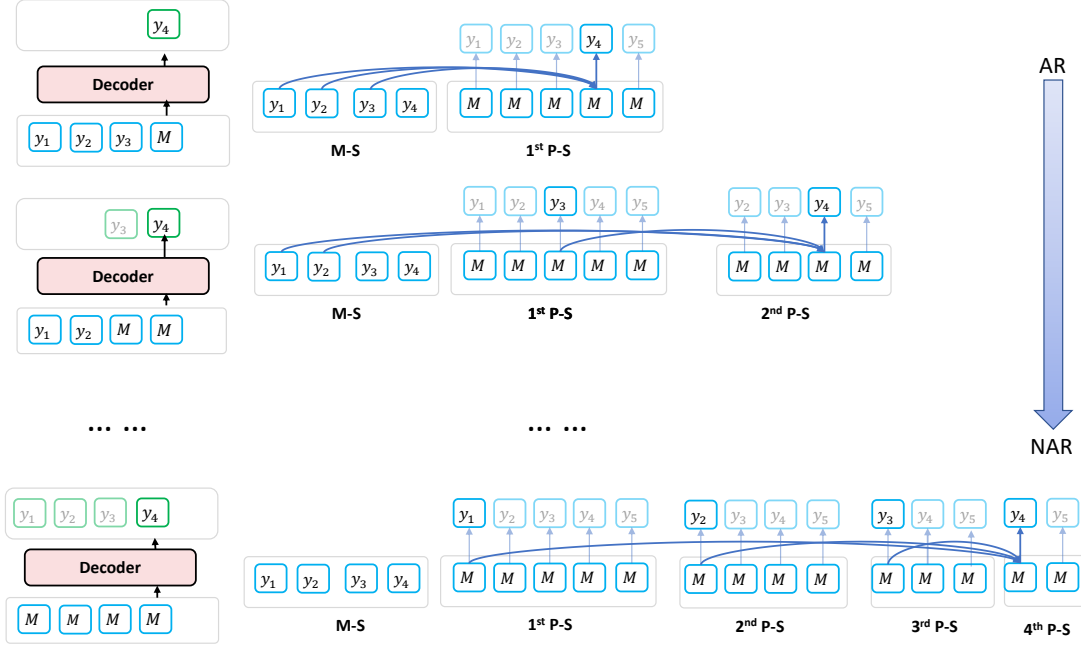


Figure 1. BANG pretraining. The prediction of y_4 is used as an example. “M-S” and “P-S” represent “Main stream” and “Predicting stream”, respectively. The 1st predicting stream predicts tokens in AR manner. First tokens in each predicting stream compose NAR prediction.

On the decoder side, BANG predicts each token y_t in output sequence Y with cross-stream visible n-stream self-attention. Besides predicting each token with all previous golden tokens in the AR manner and all [MASK]s in a NAR manner, BANG targets at predicting them with arbitrary length of previous golden tokens replaced with [MASK]s as the context. Formally, for y_t in Y :

$$P(y_t|y_{<t}, x), P(y_t|y_{<t-1}, x), P(y_t|y_{<t-2}, x), \dots, P(y_t|y_1, x), P(y_t|x) = \mathbf{Decoder}(y_{<t}, H_{\text{enc}}) \quad (5)$$

The conditional probability of target sequence in BANG can be described as:

$$P(\hat{Y}|X) = \prod_{t=1}^{|Y|} \prod_{i=1}^t P(y_t|y_{\leq t-i}, X) \quad (6)$$

BANG targets at optimizing \hat{Y} rather than the original output sequence Y . For each token y_t in Y , \hat{Y} considers replacing its i previous tokens with [MASK] for any $i < t$. For maximum likelihood training with a cross-entropy loss:

$$\begin{aligned} \mathcal{L}_{\text{BANG}} &= \log P(\hat{Y}|X) \\ &= \sum_{t=1}^{|Y|} \sum_{i=1}^t \log P(y_t|y_{\leq t-i}, X) \\ &= \underbrace{\sum_{t=1}^{|Y|} \log P(y_t|y_{<t}, X)}_{\text{AR part}} \\ &\quad + \underbrace{\sum_{t=2}^{|Y|} \sum_{i=2}^{t-1} \log P(y_t|y_{\leq t-i}, X)}_{\text{Bridging part}} \\ &\quad + \underbrace{\sum_{t=1}^{|Y|} \log P(y_t|X)}_{\text{NAR part}} \end{aligned} \quad (7)$$

We can see the BANG optimizing target is composed of three parts: AR part, NAR part and bridging part. AR part and NAR part directly benefit down-stream generation pattern. The bridging part composes a curriculum learning path to benefit NAR from high-accuracy AR learning.

2.2.2. CROSS-STREAM VISIBLE N-STREAM SELF-ATTENTION

To implement the decoder of BANG, we design a cross-stream visible n-stream self-attention based on the Transformer decoder. There are self-attention sub-layers, encoder-decoder attention sub-layers and feed-forward sub-layers in Transformer decoder layers. Encoder-decoder attention sub-layers and feed-forward sub-layers are used to fetch information from encoded hidden states and do linear calculation. We keep these two kinds of sub-layers unchanged. Self-attention sub-layers control the information flow and we propose cross-stream visible n-stream self-attention to replace it.

We duplicate decoder into one main stream and n predicting streams where n equals to target sequence length by sharing parameters. Main stream is fed with golden tokens as golden previous context. Predicting streams are fed with [MASK]s to predict corresponding tokens and these [MASK]s also serve as previous tokens for NAR generation. For the tokens in the i^{th} predicting stream, their $i - 1$ previous tokens are replaced by [MASK] from its previous predicting streams, and further previous tokens are golden tokens from the main stream. Each token in the 1^{st} predicting stream is predicted in AR pattern with all previous golden tokens from main stream visible. The first tokens in each predicting stream compose the NAR generation with all previous visible tokens as previous predicting streams' [MASK]s.

Figure 1 illustrates how BANG bridges the AR-NAR generation with the prediction of y_4 as an example. In the 1^{st} predicting stream, [MASK] for y_4 attends to real y_1, y_2 and y_3 from main stream golden tokens. Each token in the 1^{st} predicting stream works the same as y_4 and is pretrained in AR manner in parallel. In the 2^{nd} predicting stream, [MASK] for y_4 attends to real y_1 and y_2 from main stream golden token and [MASK] for y_3 from 1^{st} predicting stream. y_3 in the 1^{st} predicting stream and y_4 in the 2^{nd} predicting stream are generated with the conditional probability of $P(y_3, y_4|y_0, y_1)$ as shown on the left of the figure. We can see increasing difficulty in the 2^{nd} predicting stream and the transition from AR to NAR as more previous tokens are replaced with [MASK]s. In the 4^{th} predicting stream, y_4 is predicted with no previous golden tokens visible but only [MASK]s from its previous predicting streams are visible. Note that for each token in the target sequence, arbitrary previous golden tokens replaced with [MASK]s are considered in BANG, and all predicting streams are computed simultaneously. At each time step, all tokens in the target sequence from AR pattern to NAR pattern are predicted simultaneously.

To reduce GPU memory and computational cost, we adopt block-wise attention calculation. Since only the previous predicting streams are visible, each predicting stream is only

concatenated with previous streams to avoid extra computation cost. The workflow in decoder layer l can be described in Algorithm 1.

Algorithm 1 Cross-stream Visible N-stream Self-attention

Input: h_l
Output: h_{l+1}
 $K_{cache} \leftarrow \emptyset$
 $V_{cache} \leftarrow \emptyset$
 $h_{l+1} \leftarrow \emptyset$
 $Q, K, V = Linear(h_l)$
for Q_i, K_i, V_i in split(Q, K, V) **do**
 $K_{cache} \leftarrow K_{cache} \oplus K_i$
 $V_{cache} \leftarrow V_{cache} \oplus V_i$
 $h_{l+1} \leftarrow h_{l+1} \oplus Attn(Q_i, K_{cache}, V_{cache})$
end for
 return h_{l+1}

Function *Linear* has three outputs: Q, K, V for self-attention from input hidden states h_l . \oplus denotes the concatenate operator, and function *Attn* is defined in equation 8:

$$Attn(Q, K, V) = softmax\left(\frac{QK^T}{\sqrt{d}} + L\right)V \quad (8)$$

, where d represents the dimension for Q, K and V vectors, L represents the relative positional bias and mask matrix to make sure only visible tokens's V can be attended to.

2.3. BANG Finetuning

In § 2.2.1 and § 2.2.2, we introduce the model structure and its language model for pretraining. In this section, we will introduce BANG finetuning strategy for AR, NAR and semi-NAR generation.

2.3.1. AR GENERATION

BANG AR generation pattern is same as XLNet (Yang et al., 2019) 2-stream strategy. We take the prediction of y_3 as an example in Figure 2 with the encoder part omitted.

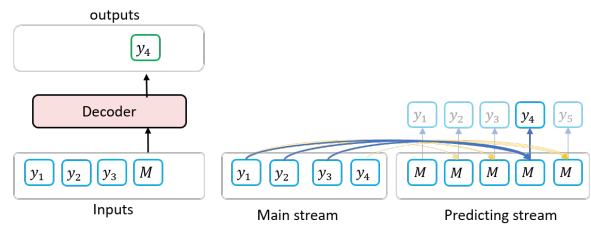


Figure 2. BANG AR generation. The [MASK] can attend to its previous golden tokens to predict y_4 .

On the left, BANG predicts y_4 with y_1, y_2 and y_3 visible as $P(y_4|y_1, y_2, y_3)$. During model training, to predict each

token in parallel, BANG duplicates decoder layers into one main stream and one predicting stream as shown on the right of Figure 2. [MASK] in the predicting stream attends to its previous tokens from main stream to predict corresponding position’s target token.

2.3.2. NAR GENERATION

BANG NAR generation pattern is same as vanilla Transformer. We take the prediction of y_4 as an example in Figure 3 with the encoder part omitted.

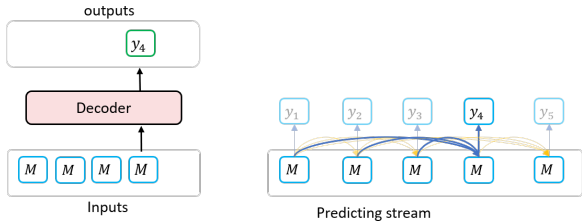


Figure 3. BANG NAR generation. The [MASK] can only attend to its previous [MASK]s to predict y_3 .

On the left, we can see BANG predicts y_4 with no context information from target sequence Y but a list of [MASK]s. On the right, a main difference between BANG NAR and CMLM or NAT is that [MASK] in BANG decoder can only attend to its previous [MASK]s rather than every token in the decoder. We design BANG in this attention manner with two reasons: 1) To be consistent with BANG AR and benefit from AR-NAR bridging pretraining. 2) If [MASK] in the decoder has bi-directional attention, the number of fed [MASK] will influence the result, under which condition a length predictor is needed. If the model predicts the wrong target sequence length, the decoder has to fill all the extra tokens or terminate unexpectedly (early). In BANG, only previous [MASK]s are visible, and the first generated [EOS] token is considered the signal of sentence end token as traditional AR models.

2.3.3. SEMI-NAR GENERATION

For BANG semi-NAR generation, visible context Y_{c_t} is designed as complete golden context visible for first several tokens in AR manner and golden+[MASK] for the rest tokens. BANG semi-NAR training procedure is same as pre-training implementation introduced in § 2.2 and we show the inference procedure in Figure 4.

We can see that BANG predicts first two tokens one by one in AR manner. Then the AR generated sub-sequence is used as a hint for the rest tokens. A list of [MASK]s are appended to be predicted at one time step in NAR manner. The motivation of this semi-NAR generation pattern is that NAR models always mix up target sequence different ex-

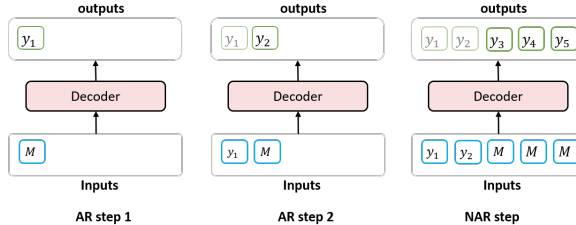


Figure 4. BANG semi-NAR generation. BANG semi-NAR inference supports predicting first several tokens step by step then generating rest tokens at one timestep. Here we present the inference procedure. The training procedure is same to Figure 1.

pressions. High-quality AR generated sub-sequence can help NAR step about choosing the writing style and other dependencies. We can also generate the outputs in vanilla AR or NAR manner with this finetuning strategy.

3. Experiments and Results

In this section, we first introduce the experimental setup and datasets in § 3.1, and report the main experimental results in § 3.2. Next, we conduct a comparison between the NAR pretraining and the BANG pretraining in § 3.3. Last, we put some case studies for each dataset in the Appendix.

3.1. Setup and Dataset

3.1.1. SETUP

We pretrain BANG using the 16GB corpus (Wikipedia and BookCorpus). The BANG model has a 6-layer encoder and a 6-layer decoder, with a dimension of 768 in their hidden states, to be consistent with AR pretrained models (Song et al., 2019; Qi et al., 2020; Lewis et al., 2019) base version model size.

We implement BANG with the span mask and prediction pretraining task, same as MASS (Song et al., 2019). In each block of an input sentence, we mask up to n continuous tokens for prediction, where n is same as the number of predicting streams. We mask 15% tokens for every 64 tokens in the input sequence span, which makes the target sequence length in each span as 9. BANG has 9 predicting streams, same as sequence length. We pretrain BANG from scratch with a learning rate of $3e-4$ for 35 epochs and a batch size of 2048. BANG uses the same dictionary as BERT-uncased (Devlin et al., 2018).

For AR generation, we load the BANG pretrained model and finetune it with teacher forcing strategy. We use the AR model **Transformer** (Vaswani et al., 2017) without pre-training, strong pretrained AR models **MASS** (Song et al., 2019), **BART** (Lewis et al., 2019) and **ProphetNet** (Qi et al., 2020) as our baselines. Most of the AR baseline results are

collected from **GLGE** (Liu et al., 2020). We collect the base version (6-layers encoder, 6-layers decoder, 768 hidden sizes) baseline models results to be compared with same model size BANG for fair comparison. BANG AR finetuning hyper-parameters are: learning rate $1e-4$, warm up steps of 1000, Adam optimizer, the maximum input and output length of 512, and a label smoothness of 0.1. We finetune BANG on each dataset for 10 epochs and select the best model according to the performance on their dev sets. For inference, we set the beam size as 4, length penalty as 1.0 and batch size as 1 to calculate the latency.

For NAR generation, we load the BANG pretrained model and finetune it with all [MASK] inputs. We use **NAT** (Gu et al., 2017), **iNAT** (Lee et al., 2018), **CMLM** (Ghazvininejad et al., 2019) and **LevT** (Gu et al., 2019) as the NAR baseline models. For these baselines, we select the outputs from the first iteration as the NAR outputs if they are semi-NAR models. For NAR finetuning experiments, the hyper-parameters are the same as AR finetuning except the number of finetuning epochs, since NAR finetuning need more epochs to converge. We finetune BANG for 50 epochs and save a checkpoint for every 10 epochs. We select the best checkpoint based on the performance on dev set. For inference, we set the no-repeat-ngram hyper-parameter as 2 to merge consecutive same tokens. Since we consider the first [EOS] token as the end signal rather than predicting the target sentence length, we set the maximal output length as 50, 85, 30 for SQuAD, XSum and PersonaChat, respectively.

For semi-NAR generation, we load the BANG pretrained model and finetune it with multiple predicting streams to simultaneously support AR inference and NAR inference. We set the maximum output length and predicting stream numbers as 30, 30, 40 for SQuAD, PersonaChat, XSum, respectively. In other words, the model produced from multi-stream finetuning can simultaneously support AR inference, NAR inference and semi-NAR inference. For semi-NAR, we predict the first n_{ar} tokens using a token-by-token manner sequentially, followed by predicting the last $n_{nar} = n - n_{ar}$ tokens in parallel via a single step. We set n_{ar} as 5, and n_{nar} as 25, 25, 35 for SQuAD, PersonaChat and XSum, respectively. The semi-NAR strategy in BANG is quite flexible to support different sequential and parallel combinations, and we leave it as future work for further exploration. For semi-NAR baselines, we choose **InsT** (Stern et al., 2019), **iNAT** (Lee et al., 2018), **CMLM** (Ghazvininejad et al., 2019), **LevT** (Gu et al., 2019) and set the maximum iteration steps as 10 (one decoding followed by up to nine iterative refinements).

For NAR and semi-NAR experiments, we do not utilize knowledge distillation (Hinton et al., 2015; Kim & Rush, 2016) from AR models. For NAR generation, knowledge

distillation usually refers to finetuning with the inference outputs from an AR model (Kim & Rush, 2016; Gu et al., 2017), in other words, the finetuning input keeps unchanged, while the target sequence is replaced with the output from an AR model. We do not utilize knowledge distillation for two reasons: We want to use the same fine-tuning corpus as AR models to compare the performance difference; Knowledge distillation with different generated results and different AR models will result in different NAR performance and it’s hard for other researchers to get the same finetuning corpus for reproducing our experiments. NAR and semi-NAR baseline models are integrated into Fairseq library (Ott et al., 2019) and we use the default hyper-parameters² to carry out experiments.

For all downstream tasks, we use 8 NVIDIA Tesla V100 GPUs for finetuning and one single V100 GPU for inference. All the experiments are conducted on the Fairseq (Ott et al., 2019) v0.9.0 codebase and we use the built-in time statistics function to calculate the per-sample inference latency.

3.1.2. DATASETS

We conduct experiments on following three popular generation benchmarks:

XSum (Narayan et al., 2018) contains 227K online article and single sentence summary pairs from the British Broadcasting Corporation (BBC). The average input and output lengths are 358.5 and 21.1, respectively.

SQuAD 1.1 (Rajpurkar et al., 2016) is a dataset created for machine reading comprehension. After preprocessing, the dataset contains 98K ⟨answer, passage, question⟩ data triples. Input is formatted as ⟨answer [SEP] passage⟩ following GLGE. The average input and output lengths are 149.4 and 11.5, respectively.

PersonaChat (Zhang et al., 2018) is a dataset created for multi-turn conversation with personalizing profiles. After preprocessing, the dataset contains 150k ⟨persona profile description text, conversation history, response⟩ data triples. Input is formatted as ⟨profile [SEP] conversation history⟩ following GLGE. The average input and output lengths are 120.8 and 11.8, respectively.

3.2. Main Results

We present the results for question generation task in Table 1, summarization task in Table 2, and dialog task in Table 3. BANG achieves significant performance improvements on all tasks consistently in both the NAR and the semi-NAR settings. Compared with the best semi-NAR baselines, BANG achieves absolute improvements of 14.01 and 5.24 in overall scores of SQuAD and XSum, respec-

²Fairseq NAR Baseline Models

Table 1. Results of semi-NAR, NAR and AR on SQuAD 1.1.

Pattern	Methods	SQuAD 1.1				Latency ms/Sample
		ROUGE-L	BLEU-4	METEOR	OVERALL	
Semi-NAR	InsT (Stern et al., 2019)	29.98	2.34	8.15	13.49 (+0.00)	67.61 (4.3x)
	iNAT (Lee et al., 2018)	32.34	3.16	9.18	14.89 (+1.40)	31.59 (2.0x)
	CMLM (Ghazvininejad et al., 2019)	29.60	3.89	9.70	14.40 (+0.91)	106.84 (6.8x)
	LevT (Gu et al., 2019)	30.81	2.68	9.40	14.30 (+0.81)	116.41 (7.4x)
	BANG	47.39	17.62	21.69	28.90 (+15.41)	111.11 (7.1x)
NAR	NAT (Gu et al., 2017)	31.51	2.46	8.86	14.28 (+0.02)	17.11 (1.1x)
	iNAT (Lee et al., 2018)	32.44	2.33	8.84	14.54 (+0.28)	16.52 (1.1x)
	CMLM (Ghazvininejad et al., 2019)	31.58	2.51	8.85	14.31 (+0.05)	16.41 (1.0x)
	LevT (Gu et al., 2019)	31.38	2.27	9.14	14.26 (+0.00)	27.52 (1.8x)
	BANG	44.07	12.75	18.99	25.27 (+11.01)	15.69 (1.0x)
AR	Transformer (Vaswani et al., 2017)	29.43	4.61	9.86	14.63(+0.00)	159.49(10.2x)
	MASS (Song et al., 2019)	49.48	20.16	24.41	31.35 (+16.72)	N/A
	BART (Lewis et al., 2019)	42.55	17.08	23.19	27.61 (+12.98)	N/A
	ProphetNet (Qi et al., 2020)	48.00	19.58	23.94	30.51 (+15.88)	N/A
	BANG	49.32	21.40	24.25	31.66 (+17.03)	N/A

Table 2. Results of semi-NAR, NAR and AR on XSum.

Pattern	Methods	XSum			OVERALL	Latency ms/Sample
		ROUGE-1	ROUGE-2	ROUGE-L		
Semi-NAR	InsT (Stern et al., 2019)	17.65	5.18	16.05	12.96 (+0.00)	63.37 (4.0x)
	iNAT (Lee et al., 2018)	26.95	6.88	22.43	18.75 (+5.79)	31.27 (2.0x)
	CMLM (Ghazvininejad et al., 2019)	29.12	7.70	23.04	19.95 (+6.99)	113.64 (7.1x)
	LevT (Gu et al., 2019)	25.33	7.40	21.48	18.07 (+5.11)	101.01 (6.3x)
	BANG	34.71	11.71	29.16	25.19 (+12.23)	109.77 (6.9x)
NAR	NAT (Gu et al., 2017)	24.04	3.88	20.32	16.08 (+0.22)	17.47 (1.1x)
	iNAT (Lee et al., 2018)	24.02	3.99	20.36	16.12 (+0.26)	16.94 (1.1x)
	CMLM (Ghazvininejad et al., 2019)	23.82	3.60	20.15	15.86 (+0.00)	16.88 (1.1x)
	LevT (Gu et al., 2019)	24.75	4.18	20.87	16.60 (+0.74)	27.72 (1.7x)
	BANG	32.59	8.98	27.41	22.99 (+7.13)	15.97 (1.0x)
AR	Transformer (Vaswani et al., 2017)	30.66	10.80	24.48	21.98(+0.00)	262.47(16.4x)
	MASS (Song et al., 2019)	39.70	17.24	31.91	29.62 (+7.64)	N/A
	BART (Lewis et al., 2019)	38.79	16.16	30.61	28.52 (+6.54)	N/A
	ProphetNet (Qi et al., 2020)	39.89	17.12	32.07	29.69 (+7.71)	N/A
	BANG	41.09	18.37	33.22	30.89 (+8.91)	N/A

tively. In addition, BANG achieves absolute improvements of 10.73, 6.39 and 5.90 in overall scores of SQuAD, XSUM and PersonaChat compared with the best NAR baselines, respectively. This clearly demonstrates the effectiveness of the BANG pretraining. The NAR or semi-NAR results based on BANG are comparable or even better than the AR generation methods without pretraining.

From Table 1, we can see that via BANG pretraining, the results on semi-NAR and NAR generation are significantly improved. Meanwhile, the results on AR generation via BANG pretraining are comparable with strong pretrained baselines. On the other hand, both BANG semi-NAR and NAR models outperforms Transformer AR baselines with obviously reduced latency. For example, the per sample inference latency of Transformer model is 159.49ms while that of BANG NAR is only 15.69ms, which amounts to a 10.2 times speedup. In terms of the inference speed, we can see that BANG is slightly better than other NAR generation models. This speedup is mainly due to the fact that

BANG NAR uses the special token [SEP] as the ending mark, without the need to predict the length.

Similar to the Question Generation tasks, the improvement on the XSum in Table 2 is also significant. In all the three categories semi-NAR, NAR and AR, BANG can outperform all of the baselines consistently. Meanwhile, BANG’s NAR and semi-NAR results are better than AR generation models without pretraining. Furthermore, we observe an interesting result when studying both the SQuAD and XSum together: BANG semi-NAR consistently outperforms NAR in both tasks, while the baseline models show that semi-NAR via iterative refinements can help XSum summarization but often hurt SQuAD question generation.

For the results on PersonaChat in Table 3, we focus on the dialog outputs diversity Distinct-1 and Distinct-2 metrics. For dialog generation, the outputs are very open such that the outputs diversity is deeply concerned to avoid boring, or useless dialog responses. We can see the D-1 and D-2

Table 3. Results of semi-NAR, NAR, and AR on PersonaChat. D-1(Distinct-1), D-2(Distinct-2) are multiplied by 100 according to GLGE.

Pattern	Methods	PersonaChat				OVERALL	Latency
		BLEU-1	BLEU-2	D-1	D-2		ms/Sample
Semi-NAR	InsT (Stern et al., 2019)	12.63	9.43	0.1	0.3	5.62 (+0.00)	65.27 (4.4x)
	iNAT (Lee et al., 2018)	41.17	32.13	0.1	1.1	18.63 (+13.01)	43.25 (2.9x)
	CMLM (Ghazvininejad et al., 2019)	44.38	35.18	0.1	0.8	20.12 (+14.50)	105.82 (7.1x)
	LevT (Gu et al., 2019)	24.89	18.94	0.1	0.6	11.13 (+5.51)	80.26 (5.4x)
	BANG	39.82	30.72	1.9	14.2	21.66 (+16.04)	109.17 (7.3x)
NAR	NAT (Gu et al., 2017)	31.53	24.17	0.1	0.8	14.15 (+2.20)	17.86 (1.2x)
	iNAT (Lee et al., 2018)	30.56	23.38	0.1	0.7	13.69 (+1.74)	16.40 (1.1x)
	CMLM (Ghazvininejad et al., 2019)	31.44	24.06	0.1	0.6	14.05 (+2.10)	16.26 (1.1x)
	LevT (Gu et al., 2019)	26.92	20.47	0.0	0.4	11.95 (+0.00)	27.56 (1.9x)
	BANG	31.11	23.90	2.5	22.7	20.05 (+8.10)	14.89 (1.0x)
AR	Transformer (Vaswani et al., 2017)	41.56	32.95	0.3	0.8	18.90(+0.00)	138.31(9.3x)
	MASS (Song et al., 2019)	41.06	35.75	1.4	6.9	21.28 (+2.38)	N/A
	BART (Lewis et al., 2019)	47.60	39.36	1.1	6.1	23.54(+4.64)	N/A
	ProphetNet (Qi et al., 2020)	46.00	38.40	1.3	7.3	23.25 (+4.35)	N/A
	BANG	45.77	35.54	1.4	8.4	22.78 (+3.88)	N/A

of models without pretraining are very low, which means even the inputs are different, their outputs are highly overlapped or just repetitive. For NAR baseline models, we observe the outputs are often composed of “i, a, an, the” and punctuation, which may share some repetitive parts with target sequence to achieve a reasonable B-1 or B-2 score, but lacking meaningful responses. For semi-NAR baseline models, we observe the iterative refinement can help produce a fluent response. However, the generated responses are often unrelated to the dialog context with a low accuracy. Baseline AR models have high B-1 and B-2 scores via producing some common phrases like “i have”, “have a”, “a lot”, but meaningless dialog responses. On contrary, from BANG, we observe meaningful response and high diversity in the response, which makes it achieve a good diversity score consistently in all three categories: the high-speed NAR generation, the high-performance AR generation, and the semi-NAR generation. We list more examples in the Appendix for a detailed illustration.

3.3. NAR Results with AR, NAR, BANG Pretraining

Note that NAR baseline models in § 3.2 are not pretrained. Here we provide AR and NAR pretrained models to compare with BANG. We load the parameters from MASS as the AR pretrained model and carry out NAR finetuning. NAR pretraining setting is same as BANG pretraining, except that we replace the cross-stream visible multi-stream decoder with a single decoder filled with [MASK] to make it consistent with the NAR finetuning.

The results for NAR finetuning results with different pretraining settings are shown in Table 4 and Table 5. First, pretraining improves NAR finetuning results significantly and consistently in both tasks. Second, AR pretrained models perform worse than NAR pretrained models because AR pretraining has inconsistent language model between pre-

Table 4. NAR results on SQuAD 1.1 question generation with different pretraining strategy.

Pretrain Strategy	ROUGE-L	BLEU-4	METEOR
No-pretrain	30.69	2.07	8.22
AR(MASS)	40.20	9.50	16.00
NAR	41.15	11.38	17.63
BANG	44.07	12.75	18.99

Table 5. NAR results on XSum summarization with different pretraining strategy.

Pretrain Strategy	ROUGE-1	ROUGE-2	ROUGE-L
No-Pretrain	23.86	4.17	20.12
AR(MASS)	26.10	6.30	22.50
NAR	30.43	7.66	25.50
BANG	32.50	8.92	27.35

training and finetuning. Third, BANG, with the same NAR finetuning procedure and model architecture(the n-stream architecture difference only exists in the pretraining, and in the NAR finetuning procedure, two models have the exactly same architecture), consistently outperforms the NAR pretrained model. This clearly demonstrates that the proposed pretraining strategy via bridging both AR and NAR is critical to achieve a better performance in NAR generation. Lastly, the improvements of BANG are consistent and significant in both tasks across all metrics.

4. Related Work

AR models have been developed for a long time. Recent works show that pretraining on large scale text will guarantee a consistent improvements on downstream generation tasks. GPT series work (Radford et al., 2018; 2019; Brown et al., 2020) pretrains decoders with the task of next token prediction and convert different downstream tasks into language models. MASS (Song et al., 2019) masks continuous

words’ spans from input sentences to predict. BART (Lewis et al., 2019) uses denoising task to pretrain Transformer. ProphetNet (Qi et al., 2020) deploys future tokens’ prediction to enhance AR generation ability. DialogGPT (Zhang et al., 2019b) is pretrained for conversational response generation. XLNet (Yang et al., 2019) utilizes AR pretraining for downstream natural language understanding tasks, and we also borrow 2-stream strategy from XLNet for BANG cross-stream visible n-stream decoder.

NAR and semi-NAR models are proposed to accelerate natural language generation. NAT (Gu et al., 2017) is proposed as a NAR generation model to decode the whole target sequence at one time step. iNAT (Lee et al., 2018) refines outputs with multi-turn post-processing. InsT (Stern et al., 2019) predicts inserting positions and inserting tokens at each iteration. CMLM (Ghazvininejad et al., 2019) firstly predict all target words with NAR generation and maskout-regenerate low confidence words. LevT (Gu et al., 2019) considers two basic operations insertion and deletion at each iteration.

Some works propose to use AR generation facilitate NAR generation and some work propose future tokens’ prediction to facilitate AR generation. Curriculum learning (Bengio et al., 2009) is used to benefit NAR generation from AR generation such as Guo et al. (2020a). Future tokens’ prediction is used to benefit AR generation such as Qi et al. (2020); Goodman et al. (2020). Recently, algorithms (Tian et al., 2020; Mansimov et al., 2019) to support different generation patterns with the same model start to be researched.

BANG benefits from these different NAR and AR generation models to support NAR, AR, semi-NAR generation.

5. Conclusion

We propose a new natural language generation pretraining model named BANG. BANG bridges NAR and AR generation with cross-stream visible n-stream strategy and large scale pretraining. BANG supports NAR, AR and semi-NAR generation with the same pretrained model. Experiments show that BANG can significant improve the NAR and semi-NAR generation performance, and provide comparable results with strong AR pretraining models. NAR and semi-NAR generation can be applied to general natural language generation tasks with acceptable performance with BANG pretraining. BANG is powerful and flexible to support more diverse semi-NAR generation strategies and finetuning strategies, which we leave as future work.

References

- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2018.
- Ghazvininejad, M., Levy, O., Liu, Y., and Zettlemoyer, L. Mask-predict: Parallel decoding of conditional masked language models. *arXiv preprint arXiv:1904.09324*, 2019.
- Goodman, S., Ding, N., and Soricut, R. Teaform: Teacher-forcing with n-grams. *arXiv preprint arXiv:2010.03494*, 2020.
- Gu, J., Bradbury, J., Xiong, C., Li, V. O., and Socher, R. Non-autoregressive neural machine translation. *arXiv preprint arXiv:1711.02281*, 2017.
- Gu, J., Wang, C., and Zhao, J. Levenshtein transformer. In *Advances in Neural Information Processing Systems*, pp. 11181–11191, 2019.
- Guo, J., Tan, X., Xu, L., Qin, T., Chen, E., and Liu, T.-Y. Fine-tuning by curriculum learning for non-autoregressive neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 7839–7846, 2020a.
- Guo, J., Zhang, Z., Xu, L., Wei, H.-R., Chen, B., and Chen, E. Incorporating bert into parallel sequence decoding with adapters. *arXiv preprint arXiv:2010.06138*, 2020b.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Kim, Y. and Rush, A. M. Sequence-level knowledge distillation. *arXiv preprint arXiv:1606.07947*, 2016.
- Lee, J., Mansimov, E., and Cho, K. Deterministic non-autoregressive neural sequence modeling by iterative refinement. *arXiv preprint arXiv:1802.06901*, 2018.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019.

- Liu, D., Yan, Y., Gong, Y., Qi, W., Zhang, H., Jiao, J., Chen, W., Fu, J., Shou, L., Gong, M., et al. Glge: A new general language generation evaluation benchmark. *arXiv preprint arXiv:2011.11928*, 2020.
- Mansimov, E., Wang, A., Welleck, S., and Cho, K. A generalized framework of sequence generation with application to undirected sequence models. *arXiv preprint arXiv:1905.12790*, 2019.
- Narayan, S., Cohen, S. B., and Lapata, M. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. In *EMNLP*, pp. 1797–1807, 2018.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. *arXiv preprint arXiv:1904.01038*, 2019.
- Qi, W., Yan, Y., Gong, Y., Liu, D., Duan, N., Chen, J., Zhang, R., and Zhou, M. Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training. *arXiv preprint arXiv:2001.04063*, 2020.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/languageunsupervised/language_understanding_paper.pdf, 2018.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8), 2019.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., and Liu, P. J. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21 (140):1–67, 2020.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. Squad: 100,000+ questions for machine comprehension of text. In *EMNLP*, pp. 2383–2392, 2016.
- Shu, R., Lee, J., Nakayama, H., and Cho, K. Latent-variable non-autoregressive neural machine translation with deterministic inference using a delta posterior. In *AAAI*, pp. 8846–8853, 2020.
- Song, K., Tan, X., Qin, T., Lu, J., and Liu, T.-Y. Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*, 2019.
- Stern, M., Chan, W., Kiros, J., and Uszkoreit, J. Insertion transformer: Flexible sequence generation via insertion operations. *arXiv preprint arXiv:1902.03249*, 2019.
- Su, Y., Cai, D., Wang, Y., Vandyke, D., Baker, S., Li, P., and Collier, N. Non-autoregressive text generation with pre-trained language models. *arXiv preprint arXiv:2102.08220*, 2021.
- Tian, C., Wang, Y., Cheng, H., Lian, Y., and Zhang, Z. Train once, and decode as you like. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 280–293, 2020.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *NIPS*, 2017.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., and Le, Q. V. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.
- Zhang, J., Zhao, Y., Saleh, M., and Liu, P. J. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. *arXiv preprint arXiv:1912.08777*, 2019a.
- Zhang, S., Dinan, E., Urbanek, J., Szlam, A., Kiela, D., and Weston, J. Personalizing dialogue agents: I have a dog, do you have pets too? In *ACL*, pp. 2204–2213, 2018.
- Zhang, Y., Sun, S., Galley, M., Chen, Y.-C., Brockett, C., Gao, X., Gao, J., Liu, J., and Dolan, B. Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*, 2019b.