

## RESEARCH ARTICLE

# Table-to-Text Generation With Pretrained Diffusion Models

ALEKSEI S. KRYLOV<sup>1,2</sup> AND OLEG D. SOMOV<sup>1,3</sup><sup>1</sup>Moscow Institute of Physics and Technology, 141701 Moscow, Russia<sup>2</sup>Sber, 117312 Moscow, Russia<sup>3</sup>Artificial Intelligence Research Institute (AIRI), 105064 Moscow, Russia

Corresponding author: Aleksei S. Krylov (krylov.as@phystech.edu)


This work was supported in part by the Artificial Intelligence Research Institute.

**ABSTRACT** Diffusion models have demonstrated significant potential in achieving state-of-the-art performance across various text generation tasks. In this systematic study, we investigate their application to the table-to-text problem by adapting the diffusion model to the task and conducting an in-depth analysis. Our experiments cover multiple aspects of diffusion models training. We explore sampling strategy influence by inducing recent diffusion model accelerator DPM-Solver++ into our core model. We have tested different prediction aggregation methods, like ROVER and Minimum Bayes-Risk (MBR). Our studies cover the impact of the pre-training phase in diffusion models and the generation length constraints influence. We also have compared diffusion model generation with auto-regressive text-to-text models with different temperature settings for diversity evaluation. Our key observation is that diffusion models demonstrate the balance between quality and diversity while auto-regressive text-to-text models are not successful at handling both at the same time. Furthermore, we found out that to achieve the highest quality possible, it is preferable to use a regular sampler with the strictest length constraint to create multiple samples, and then use MBR to aggregate the predictions. However, if you are prepared to give up high level of diversity and to accelerate the process, you can also utilize a fast sampler DPM-Solver++. Our findings reveal that diffusion models achieve comparable results in the table-to-text domain, highlighting their viability in the table-to-text challenge as a promising research direction.

**INDEX TERMS** Machine learning, diffusion models, deep learning, data-to-text, table-to-text, natural language processing, artificial intelligence.

## I. INTRODUCTION

Diffusion models, initially prominent in the vision [13], have shown remarkable potential in text generation tasks. These iterative generative models are trained to recover corrupted data through a multi-step denoising process, refining samples from pure noise to produce high-fidelity and diverse outputs. Inspired by their success in vision [13] and audio [19], researchers are exploring diffusion models for text generation [10], [20], [21], where they introduce a novel noising paradigm and a distinct training objective. This approach offers an alternative to traditional token prediction methods, promising enhanced language modeling capabilities and

The associate editor coordinating the review of this manuscript and approving it for publication was Ramoni Adeogun .

paving the way for advancements in generating coherent and contextually appropriate text.

Data-to-text generation is the task of generating a target textual description conditioned on source content in the form of structured data. Table-to-text generation involves creating textual descriptions from tables. Examples include generating sentences from historical data, attractions descriptions, football game summaries from box score statistics, and fun facts from Wikipedia info tables [24].

Current experiments demonstrate that state-of-the-art neural models struggle to generate satisfactory results, despite the high quality of the training data. Therefore it is necessary to explore other generation methods.

In turn, we propose a diffusion model for this task, because it is known for generating highly diverse outputs, robust

generative process, and advanced prediction aggregation techniques, and in natural language processing, in such tasks as summarization, paraphrasing, open domain dialogue, diffusion models showed increased text diversity, while maintaining a comparable accuracy level to modern text-to-text solutions [10], [21].

In our work, we have applied the diffusion model GENIE [21] with a fast sampler [11] to the table-to-text ToTTo challenge [24], because there isn't a pre-trained diffusion architecture with weights other than this one, and pre-training of own architecture is excessively costly and resource-intensive. For the comparison with our solution, we have used models of relevant sizes from the T5 family [17], [26], which are auto-regressive text-to-text models.

As a core part of the paper, we have carried out extensive research on the main parts of diffusion models. We did an in-depth impact analysis of the pre-training phase and generation length of the diffusion model. We replaced the original sampler with efficient DPM-Solver++ [11] and evaluated its performance on our task. We compared the diffusion model diversity with different temperature sampling of the T5 model [17], [26].

Our results show that diffusion models can achieve similar scores in terms of accuracy and diversity in comparison to auto-regressive text-to-text models.

The contribution of this paper is the following:

- We propose an approach for the table-to-text problem that leverages a pretrained diffusion model.
- We compare the proposed method to the auto-regressive baselines and experimentally show that it is effective both in terms of the quality and diversity.
- We adapt fast DPM-Solver++ for GENIE model and compare with regular architecture to test impact to the quality and diversity.
- We discover that trained from scratch diffusion models outperform autoregressive models on the table-to-text issue.
- We assess several features of diffusion models, including the influence of output length limits, different aggregation techniques, and also the impact of sampling temperature to auto-regressive baselines.

## II. RELATED WORK

### A. DATA-TO-TEXT CHALLENGE

Data-to-text is a task of describing structured data adequately and fluently. The three most common datasets for this domain are WebNLG [9], MultiWoz [3], and ToTTo [24]. WebNLG [9] is a dataset with approximately 18K graphs of subject-object-predicate triplets and their textual descriptions. MultiWoz [3] consists of 10K human-human dialogues for developing task-oriented dialogue systems with approximately 56K training samples. The ToTTo dataset [24] is a corpus of over 120K Wikipedia tables paired with natural language descriptions. We conducted our experiments on the

ToTTo dataset [24], because it is the biggest and most modern dataset for the data-to-text task.

Recent research utilized mostly neural solutions for end-to-end generation from a linearized table to its summary. After the triumph of extensively pre-trained sequence-to-sequence Transformer models, modern state-of-the-art systems utilize these models for table-to-text generation, in which the input table is converted into a textual sequence through linearization. Examples of such models include the pre-trained language model Bert-to-Bert [28] and T5 models [17]. Condensing a table's extensive information into one sentence is challenging, but focusing on key details enables selective descriptions, so [32] suggest T5 with structural attention and special position encoding. This helped improve the generation results, but not significantly. Also, some researchers proposed plan-based methods, such as PlanGen [31] and Content Planner [25] to control the structure of the generated output. Eventually, an CoNT [1] developed a new Contrastive Neural Text generation framework and validated it on data-to-text tasks, achieving state-of-the-art quality among base-size models. However, none of existing methods achieved ideal quality, so this still remains an unsolved problem. In addition, there are no works on table-to-text diffusion models and no in-depth exploration of diversity metrics in that challenge, and we decided to do it.

### B. DIFFUSION MODELS IN NLP

Diffusion models have demonstrated significant potential in achieving state-of-the-art performance across various text generation tasks [13], [19]. At first text-to-text diffusion models were discrete, for example, each token can, with some probability, be noisy or replaced with a random token [2], [15], [16]. Then Diffusion-LM [20] appeared as the diffusion model in continuous space for unconditional generation, which maps discrete tokens into continuous latent variables, achieving more complex controllable text generation through continuous diffusion. DiffuSeq [10] proposed a new method for conditional text generation. This represents a BERT-like encoder where it is fed concatenated input and additionally noisy output.

In GENIE [21] was presented a pre-trained encoder-decoder diffusion transformer which improved text generation quality.

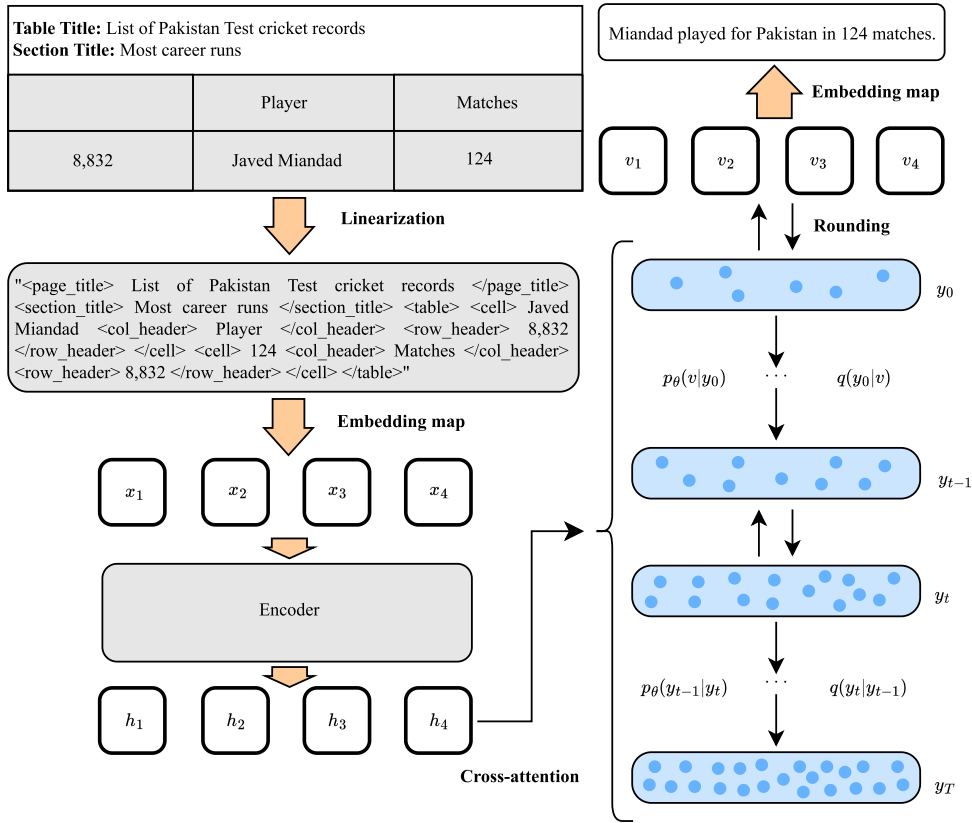
Lately, the most notable advancements in diffusion text generation models are - Difformer [8] with an anchor loss to regularize the embeddings and stabilize the training simultaneously, and DiffuSeq-v2 [11] which proposed a decision for increasing the speed of training in 4 times and speed of inference in 800 times.

## III. METHOD

The proposed method is shown in the Fig. 1 and is described in detail below.

### A. DATA PREPARATION

At first, we translate tables into sequences. We use the standard linearization method, proposed by [24]. Then we



**FIGURE 1.** We take the linearized source table as the input of the Encoder to obtain the hidden information embeddings which interact with the Language Diffusion Model through cross attention. The Language Diffusion Model deletes the random Gaussian noise from the output text embeddings through the iterative denoising and grounding process. Finally, we map received embeddings to tokens and concatenate them to get the output summary of the input table.

represent a transformed table as vector  $x = \{x_1, x_2, x_3, \dots, x_n\}$  with  $n$  tokens  $x_i$ , where  $i$  is the token's position. The target sequence is also represented as vector  $y = \{v_1, v_2, v_3, \dots, v_m\}$  with  $m$  tokens  $v_i$ , where  $i$  is the token's position.

## B. DIFFUSION MODEL

Diffusion models are a type of iter-NAR model [10] where a series of intermediate sequences  $\mathbf{v}^y$  are introduced over  $T$  iterations:

$$p_{\text{iter-NAR}}(\mathbf{v}_{1:n}^y | \mathbf{w}^x) = \sum_{\mathbf{v}_1^y, \dots, \mathbf{v}_{T-1}^y} \left[ \underbrace{\prod_{i=1 \dots n} p(\mathbf{v}_{1,i}^y | \mathbf{w}^x)}_{\text{initial prediction}} \cdot \underbrace{\prod_{t=1 \dots T-1} \prod_{i=1 \dots n} p(\mathbf{v}_{t+1,i}^y | \mathbf{v}_{t,1:n}^y, \mathbf{w}^x)}_{\text{progressive full-context prediction}} \right]. \quad (1)$$

We use the pre-trained encoder-decoder diffusion model GENIE [21] for conditional generation  $p(y|x)$ . This model contains forward and reverse discrete-time Markov processes.

### 1) FORWARD PROCESS

In the forward process, we gradually perturb embeddings of input tokens according to a variance scheduler  $\beta_0, \beta_1, \dots, \beta_T$ , during  $T$  iterations, while we get standard Gaussian noise. The diffusion process starts with the initial state  $x_0$  at time step  $t = 0$ , where  $x_0$  is sampled from the Gaussian distribution of the original data  $q(y_0|v) = \mathcal{N}(\text{Emb}(v), \beta_0 I)$ . At the time step  $t + 1$ , the latent variable  $y_{t+1}$  is only determined by the  $x_t$  at time  $t$ , expressed as:

$$q(y_{t+1}|y_t) = \mathcal{N}(y_{t+1}; \sqrt{1 - \beta_{t+1}}y_t, \beta_{t+1}I) \quad (2)$$

### 2) REVERSE PROCESS

The reverse diffusion process is needed to recover the original  $y_0$  by denoising  $y_t$ :

$$p(y_t|y_{t+1}) = \mathcal{N}(y_t; \mu_\theta^t, \sigma_t^2) \quad (3)$$

$$\mu_\theta^t(y_t, y_\theta) = \frac{\sqrt{\bar{\alpha}_t} * (1 - \bar{\alpha}_{t-1})}{(1 - \bar{\alpha}_t)} * y_t + \frac{\beta_t * \sqrt{\bar{\alpha}_{t-1}}}{1 - \bar{\alpha}_t} * y_\theta \quad (4)$$

$$\sigma_t^2 = \frac{1 - \bar{\alpha}_t}{1 - \bar{\alpha}_{t+1}} * \beta_{t+1} \quad (5)$$

**Algorithm 1** Inference Algorithm

- 1: **Input:** Linearized table as a vector of tokens.
- 2: **Output:** Summary of table as text.
- 3: Transform input tokens to embeddings.
- 4: Pass them through the encoder and get hidden states.
- 5: Initialize input for diffusion decoder  $\mu = 0$  and  $\sigma = I$ .
- 6: **for**  $t = T - 1 \dots 1$  **do**
- 7:  $y_t = \text{Sample from } \mathcal{N}(\mu_\theta^t(y_{t+1}, y_0), \sigma_{t+1})$
- 8: We pass  $y_t$  to the decoder and with the help of hidden states from the encoder non-autoregressive generate  $z_\theta$
- 9:  $y_\theta(y_t, t) = \text{Quantize } z_\theta(y_t, t)$  to nearest embeddings
- 10:  $\mu_\theta^{t-1}(y_t, y_0) = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{(1-\bar{\alpha}_t)} * y_t + \frac{\beta_t * \sqrt{\bar{\alpha}_{t-1}}}{1-\bar{\alpha}_t} * y_\theta(y_t, t)$
- 11:  $\sigma_{t-1}^2 = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t} * \beta_t$
- 12: **end for**

where  $\alpha_{t+1} = 1 - \beta_{t+1}$ ,  $\bar{\alpha}_{t+1} = \prod_{i=1}^t \alpha_i$  and the vector  $y_\theta$  is predicted by a neural network parameterized by  $\theta$ .

## 3) TRAINING STEP

We want to minimize the distance between forward and reverse trajectories to receive a transformer-type model that is capable of removing noise and predicting  $p(y|x)$  in the iter-NAR strategy. We use the variational lower bound to optimize the negative log-likelihood  $\mathbf{E}[-\log p_\theta(y_0)] \leq \mathcal{L}_{VLB}$ . The final objective is a combination of several KL-divergence and entropy terms following [10]:

$$\begin{aligned} \mathcal{L}_{VLB} &= \mathcal{L}_T + \mathcal{L}_{T-1} + \dots + \mathcal{L}_0 \\ &= \mathbf{E}_{q(y_{1:T}|y_0, x)} \left[ \log \frac{q(y_T|y_0, x)}{p_\theta(y_T)} + \sum_{i=2}^T \log \frac{q(y_{i-1}|y_0, y_i, x)}{p_\theta(y_{i-1}|y_i, x)} \right. \\ &\quad \left. + \log \frac{q(y_0|v)}{p_\theta(y_0|y_1)} - \log p_\theta(v|y_0) \right] \end{aligned} \quad (6)$$

$$\begin{aligned} \mathcal{L}_t &= \mathbf{E}_{q(y_{1:T}|y_0, x)} [Const * \|\mu_t(y_t, y_0) - \mu_\theta(y_t, t)\|^2] \\ &= Const * \mathbf{E}_{q(y_{1:T}|y_0, x)} [y_0 - f_\theta(y_t, t)]^2 \end{aligned} \quad (7)$$

We sample  $t$  from the distribution  $r_t \sim \sqrt{\mathbf{E}[\mathcal{L}_t^2]}$ ,  $\sum_{i=0}^{T-1} p_i = 1$ , thus we can sample more often those steps where we have the biggest loss.

## 4) INFERENCE STEP

First of all, we score our input table with the help of an encoder. After that, we generate noise from the Gaussian distribution and consequentially transform it into the target description, as described in Algorithm 1.

## IV. EXPERIMENTS

## A. EXPERIMENT SETUP

## 1) DATASET

We evaluate our method from the quality and diversity aspects on the ToTTo dataset. The validation and testing sets are based on examples from the training samples and other new tables. Test set  $\mathbf{D}_{\text{test}}$  is divided into overlap and non-overlap

sets in the following way:

$$\begin{aligned} \mathbf{D}_{\text{test-overlap}} &:= \{d : h(d) \text{ in } h(\mathbf{D}_{\text{train}})\}, \\ \mathbf{D}_{\text{test-nonoverlap}} &:= \{d : h(d) \text{ not in } h(\mathbf{D}_{\text{train}})\} \end{aligned}$$

where  $h$  is the header value of table  $d$ . Development set  $\mathbf{D}_{\text{dev}}$  is divided into parts similarly [24].

## 2) BASELINES

As baselines, we use T5 models with 60M (small version) and 220M (base version) parameters. We finetune them 10 epochs (about 150000 steps) with AdamW [22] optimizer with learning rate  $2e-4$ , weight decay 0, and batch size 8 on V100, following the paper [32]. All other baseline results on ToTTo can be found in the official leaderboard.

## 3) EVALUATION

For quality evaluation, we use the metrics BLEU [23], [24] to measure intersection with the target sequence by n-grams, PARENT [5], [24] to get the F1-score with table data, and BLEURT [30] to robustly evaluate closeness of the target sequence and the predictions on overall, overlapping and non-overlapping subsets. To assess diversity, we generate 10 samples and use three metrics such as div-4 [4], [10] to measure the percentage of distinct 4-grams in the set of outputs, dist-1 [10] to measure internal diversity as the percentage of distinct unigrams, and self-BLEU [10], [34] as the metric of overlapping between different candidates.

Also, we compute the average length of predictions to understand the reasons for changing BLEU and diversity metrics. In all tables, the length is presented in the number of characters.

## 4) IMPLEMENTATION DETAILS

As a core model we use the official implementation GENIE [21].<sup>1</sup> Our model is a 6-layer transformer as the encoder and a 6-layer cross-attention transformer as the denoising decoder. The latent variable dimension is 768 and the embedding dimension is 128. Our model has 144M parameters. The batch size is 64. The input length is 475 and the output length is 119. The model was trained for 120000 steps with the help of AdamW optimizer [22] with a learning rate  $5e-5$  and weight decay 0 on V100. After the training, we compute the exponential moving average of the model's weights for stability.

## B. MAIN RESULTS

We compare the iter-NAR encoder-decoder diffusion model with modern text-to-text baselines T5-small and T5-base [17]. For the quality of generation, we have followed the existing work [24], reporting BLEU [23] and PARENT [5] on the Overall, Overlap, and Non-Overlap parts of the development and test parts of the dataset ToTTo [24]. For the test part we also public learning robust metric BLEURT [30]

<sup>1</sup><https://github.com/microsoft/ProphetNet/tree/master/GENIE>

**TABLE 1.** Sample outputs in ToTTo [24] development set, conditioned on the same  $x$ .

<i>Input Table:</i>	
<b>Table Title:</b> A. J. Hawk	
<b>Section Title:</b> Career statistics	
row index	TOTAL
	119

---

*Original sentence:* In his rookie season, Hawk led with 119 total tackles.

T5-small	Ours
hawk had 119 tackles.	a. j. hawk played 119 games.
hawk had 119 total tackles	hawk led the team with 119 runs.
hawk had 119 career total tackles.	hawk scored 119 home runs.
hawk scored 119 runs.	a. j. hawk had 119 hits.
hawk scored 119 strikeouts.	hawk earned a total of 119 runs batted in during his career.

**TABLE 2.** Results on the ToTTo development set. We obtained comparable with baselines metrics.

Model	Overall		Overlap		Non-Overlap	
	BLEU $\uparrow$	PARENT $\uparrow$	BLEU $\uparrow$	PARENT $\uparrow$	BLEU $\uparrow$	PARENT $\uparrow$
BERT-to-BERT [28]	44.0	52.6	52.7	58.4	34.8	46.7
T5-Small [17]	45.7	55.9	53.9	60.4	37.7	51.6
T5-Base [17]	47.7	57.1	56.1	61.8	39.6	52.6
T5-Large [17]	48.1	57.3	56.8	62.0	39.8	52.8
T5-3B [17]	48.4	57.8	56.7	62.4	40.4	53.3
T5-base-CONT [1]	49.2	59.4	-	-	41.5	55.0
T5-Small	46.4	56.6	55.0	61.5	38.1	52.0
T5-Base	46.8	56.4	55.0	61.3	38.9	51.7
Ours	<b>52.5</b>	<b>62.5</b>	<b>61.6</b>	<b>67.8</b>	<b>43.7</b>	<b>57.4</b>
Ours (DPM-Solver++)	48.6	57.1	57.3	62.9	40.0	51.5

**TABLE 3.** Results on the ToTTo test set. Our diffusion model outperforms BERT-to-BERT [28] which is twice number of parameters, and achieves comparable with state-of-the-art baselines metrics. BLEURT is original BLEURT score multiplied by 100 for comparable looks to other measures.

Model	Overall			Overlap			Non-Overlap		
	BLEU $\uparrow$	PARENT $\uparrow$	BLEURT $\uparrow$	BLEU $\uparrow$	PARENT $\uparrow$	BLEURT $\uparrow$	BLEU $\uparrow$	PARENT $\uparrow$	BLEURT $\uparrow$
T5-3B [17]	<b>49.5</b>	58.4	23.0	<b>57.5</b>	62.6	35.1	<b>41.4</b>	54.2	10.8
T5-base-CONT [1]	49.1	<b>58.9</b>	<b>23.8</b>	56.7	<b>63.2</b>	<b>35.5</b>	41.3	<b>54.6</b>	<b>12.1</b>
T5-base-Lattice [32]	48.4	58.1	22.2	56.1	62.4	34.5	40.4	53.9	9.9
Pointer-Generator [29]	41.6	51.6	7.6	50.6	58.0	24.4	32.2	45.2	-9.2
Content Planner [25]	19.2	29.2	-57.6	24.5	32.5	-49.1	13.9	25.8	-66.2
BERT-to-BERT [28]	44.0	52.6	12.1	52.7	58.4	25.9	35.1	46.8	-1.7
Ours	47.1	56.9	19.1	55.4	61.9	33.3	38.5	51.9	4.8

**TABLE 4.** Diversity metrics on the ToTTo development set. Our model maintains diversity near the auto-regressive baselines yet.

	Self-BLEU $\downarrow$	div-4 $\uparrow$	dist-1 $\uparrow$	Length
Bert-to-Bert (beam) [27]	75.0	-	-	-
Bert-to-Bert (top-k) [27]	82.2	-	-	-
Bert-to-Bert (top-p) [27]	84.3	-	-	-
T5-3B (beam) [27]	83.7	-	-	-
T5-3B (top-k) [27]	88.9	-	-	-
T5-3B (top-p) [27]	90.1	-	-	-
PlanGen [27]	<b>25.9</b>	-	-	-
T5-small	31.9	69.5	94.9	<b>87.6</b>
T5-base	31.1	<b>70.7</b>	<b>95.5</b>	84.4
Ours	35.4	69.1	93.0	81.4

for completeness of comparison. We evaluated the diversity of these models following the [10] paper by length, dist-1, div-4 [4], and Self-BLEU [34] metrics.

We present the main results of the diffusion model and the baselines on the dataset ToTTo in Fig. 2 and Fig. 4 for the dev set, and in Fig. 3 for the test set. The results are submitted to the ToTTo leaderboard<sup>2</sup> - Diffusion-TT.

#### 1) QUALITY

For the development set, to evaluate the quality of the diffusion model, we sample 10 examples from our model and choose the best prediction. For T5 models decoding is done via the usual greedy search following the paper [17]. For the test set, we apply Minimum Bayes-Risk (MBR) [10], [18] on 30 predictions (10 of each model with different output lengths - 119, 128, 153) to achieve the best results. Our results demonstrate that the diffusion model can achieve comparable quality with T5 models. Moreover, we see that

<sup>2</sup><https://github.com/google-research-datasets/ToTTo>



**TABLE 5.** Impact of different samplers on the ToTto development set. All metrics are estimated on 1 sample.

Model	Overall		Overlap		Non-Overlap	
	BLEU↑	PARENT↑	BLEU↑	PARENT↑	BLEU↑	PARENT↑
T5-Small	36.8	48.3	44.8	53.7	29.1	43.1
T5-base	37.5	47.8	45.2	53.2	30.1	42.5
Ours	<b>38.6</b>	<b>49.0</b>	<b>46.5</b>	<b>54.5</b>	<b>30.7</b>	<b>43.7</b>
Ours (DPM-Solver++)	41.5	50.6	49.7	56.6	33.4	44.9

**TABLE 6.** Impact of different samplers. Results on the ToTto development set. DPM-Solver++ [11] reduces diversity, but improves the quality of each prediction.

	Self-BLEU	div-4	dist-1	Length
Ours	35.4	69.1	93.0	81.4
Ours with Fast Sampler	69.6	30.3	94.7	71.1

for all models, the difference in metrics between the Overlap and Non-Overlap parts of the dataset is essentially the same. From this, we might infer that overfitting affects all models to a roughly equal degree.

## 2) DIVERSITY

From the table 4 can see that the diffusion model achieves comparable diversity metrics, but falls behind the plan-based model [27] and auto-regressive models with temperature sampling. We see that our model confidently outperforms auto-regressive models with other types of diverse generations such as beam search [33], nucleus [14], and top-k generation [6]. Also, we note that auto-regressive models generate longer predictions than our diffusion model, so it may be the cause for the slightly higher diversity metrics. The examples of both models generation are presented in Fig. 1.

## C. ABLATION STUDY

### 1) IMPACT OF DIFFERENT SAMPLERS

DPM-Solver++<sup>3</sup> [11] was adopted for text-to-text diffusion models and increased the speed of training by 4 times and the speed of inference by 800 times without loss of quality [11]. We applied it to our encoder-decoder diffusion model. We see that this solver increases quality metrics, expressed in Fig. 5, but, diversity metrics and the length have deteriorated greatly (Fig. 6). The model starts generating nearly identical, brief, and highly accurate sentences, which when aggregated, do not improve model quality.

### 2) PREDICTION AGGREGATION METHODS

The challenge with diffusion models for text is that we consistently receive various predictions and struggle to obtain the optimal one every time. We can use different strategies to aggregate predictions in text diffusion models. We test ROVER [7], MBR [10], [18], and Best-of-all [21] algorithms

**TABLE 7.** Comparison of different prediction aggregation methods. All algorithms are applied to 10 samples.

	MBR	ROVER	Best-of-all
T5-base	<b>45.5</b>	39.3	<b>54.3</b>
T5-small	45.4	39.7	53.9
Ours	44.6	40.7	52.5
Ours (DPM-Solver++)	42.7	<b>41.3</b>	48.6

on our model and T5. All quality metrics correlate in this experiment, we only use BLEU in Fig. 7. We see that the best results without the target's leak are obtained with the help of the MBR [18].

### 3) IMPACT OF SAMPLING TEMPERATURE

We also study the impact of sampling temperature in T5 models on diversity in comparison to diffusion models. The decrease of the temperature in the sampling of T5 leads to the quality increase but diversity deterioration. In Fig. 8 we explore how diversity and quality depend on sampling temperature in T5 vs our model. The key observation is that our diffusion model demonstrates the balance between quality and diversity whilst T5-models with the temperature near one get worse in quality and near zero get worse in diversity.

### 4) PRE-TRAINING IMPACT

The pre-training is known to increase generalization in text-to-text tasks. We evaluate if it matters on our task. We compared auto-regressive models vs diffusion models with and without pre-training.


As we see in Fig. 2 and in Fig. 9, diffusion models trained from scratch outperform auto-regressive baselines. We believe it is caused by increasing size of the training dataset with help of the adding gaussian noise.

However, the pre-training improves quality for all models evenly, resulting in a 15 to 16 percent increase.


### 5) IMPACT OF OUTPUT LENGTH

We explore how the fixed length of outputs influences quality and diversity. We select three various lengths, 119 as the maximum length of the target sequence in the training subset ToTto [24], 128 as the classical output length of our baselines, and 153 as the length in the pre-training phase of our diffusion model [21]. Our diffusion model generates better results with the small length, as we see in Fig. 10 and Fig. 3. We also observe that the average

<sup>3</sup><https://github.com/Shark-NLP/DiffuSeq/tree/diffuseq-v2>



**FIGURE 2.** Difference between pre-trained and trained from scratch models. The result of our model is the optimal value of 10 Gaussian samples.



**FIGURE 3.** Effect of the maximum output length on the ToTo development set. The result is the optimal value of 10 Gaussian samples.

output prediction length did not change with the increasing maximum allowed output length. The diffusion model uses normal noise during the generation process, so the longer the length, the more noise the model receives and the higher

the likelihood of receiving an inappropriate word. This, we believe, explains why the quality of the model predictions decreases with length. This leads to an increase in inaccuracy, but it appears that these hallucinations only alter words, they

**TABLE 8.** Impact of sampling temperature to BLEU [23] on the overall development set and self-BLEU [34] between 10 samples.

	Self-BLEU↓	BLEU-Overall↑	PARENT-Overall↑
T5-small-1	31.9	36.8	48.3
T5-small-0.97	33.3	37.6	49.0
T5-small-0.85	40.7	41.0	51.9
T5-small-0 (greedy)	100.0	46.4	<b>56.6</b>
T5-base-1	<b>31.1</b>	37.5	47.8
T5-base-0.97	32.3	38.1	48.1
T5-base-0.85	37.5	40.7	50.6
T5-base-0 (greedy)	100.0	<b>46.8</b>	56.4
Ours	35.4	38.6	49.0
Ours (DPM-Solver++)	69.6	41.5	50.6

**TABLE 9.** Results on the ToTo development set for trained from scratch models. Diffusion models outperform auto-regressive baselines in quality metrics.

Model	Overall		Overlap		Non-Overlap	
	BLEU↑	PARENT↑	BLEU↑	PARENT↑	BLEU↑	PARENT↑
T5-Small	38.6	47.3	47.2	53.2	30.1	41.6
T5-Base	40.7	51.3	50.5	57.6	31.5	45.1
Ours	<b>47.2</b>	<b>57.5</b>	<b>56.5</b>	<b>63.6</b>	<b>38.1</b>	<b>51.5</b>
Ours (DPM-Solver++)	40.7	51.0	48.8	56.9	32.6	45.4

**TABLE 10.** Effect of the maximum output length on the ToTo development set. The result is the optimal value of 10 Gaussian samples.

	Maximum Output length	BLEU ↑	Length
Ours	119	<b>52.5</b>	81.4
	128	52.0	81.6
	153	52.2	81.5
Ours (DPM-Solver++)	119	48.6	71.1
	128	48.2	70.7
	153	48.0	70.5

do not add new ones, because the average length remains the same.

## V. CONCLUSION AND FUTURE WORK

We investigated diffusion models application to the table-to-text problem by adapting them to the task and conducting an in-depth analysis. Our experiments covered multiple aspects of diffusion models training. We explored sampling strategy influence by inducing recent diffusion model accelerator DPM-Solver++ [11] into our core model and found out that DPM-Solver++ [11] reduces diversity, but improves the quality of each prediction. In addition, we assessed many prediction aggregation techniques, including Minimum Bayes Risk (MBR) and ROVER, and found that MBR was more effective. We also looked into the effects of generation length limits and the pre-training phase on diffusion models. It was discovered that diffusion models trained from scratch outperformed auto-regressive baselines. In order to assess variety, we also contrasted the performance of diffusion models with that of auto-regressive text-to-text models at various temperature settings. The key observation is that our diffusion model demonstrates the balance between quality and diversity whilst T5-models with the temperature near one get worse in quality and

near zero get worse in diversity. Our findings reveal that diffusion models achieve comparable results in the table-to-text domain, highlighting their diversity in the table-to-text challenge as a promising research direction. We see the future of this topic to explore more modern transformers variations at the core of the model along with the scaling experiments. Another path of investigation is the effect of resource allocation, the impact of the text structure and table complexity.

## VI. LIMITATIONS

We use the pre-trained model GENIE [21], which has the vanilla BERT-to-BERT [28] transformer core. We did not test other contemporary architectures like T5 [26] or DeBERTa-V3 [12].

## REFERENCES

- [1] C. An, J. Feng, K. Lv, L. Kong, X. Qiu, and X. Huang, "CoNT: Contrastive neural text generation," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2023, pp. 2197–2210.
- [2] J. Austin, D. D. Johnson, J. Ho, D. Tarlow, and R. van den Berg, "Structured denoising diffusion models in discrete state-spaces," 2021, *arXiv:2107.03006*.
- [3] P. Budzianowski, T.-H. Wen, B.-H. Tseng, I. Casanueva, S. Ultes, O. Ramadan, and M. Gašić, "MultiWOZ—A large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling," 2018, *arXiv:1810.00278*.
- [4] A. Deshpande, J. Aneja, L. Wang, A. Schwing, and D. A. Forsyth, "Fast, diverse and accurate image captioning guided by part-of-speech," 2018, *arXiv:1805.12589*.
- [5] B. Dhingra, M. Faruqui, A. Parikh, M.-W. Chang, D. Das, and W. Cohen, "Handling divergent reference texts when evaluating table-to-text generation," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics*, Florence, Italy, A. Korhonen, D. Traum, and L. Márquez, Eds. Stroudsburg, PA, USA: Association for Computational Linguistics, 2019, pp. 4884–4895, doi: 10.18653/v1/p19-1483. [Online]. Available: <https://aclanthology.org/P19-1483>
- [6] A. Fan, M. Lewis, and Y. Dauphin, "Hierarchical neural story generation," 2018, *arXiv:1805.04833*.



- [7] J. G. Fiscus, "A post-processing system to yield reduced word error rates: Recognizer output voting error reduction (ROVER)," in *Proc. IEEE Workshop Autom. Speech Recognit. Understand.*, Dec. 1997, pp. 347–354, doi: [10.1109/ASRU.1997.659110](https://doi.org/10.1109/ASRU.1997.659110).
- [8] Z. Gao, J. Guo, X. Tan, Y. Zhu, F. Zhang, J. Bian, and L. Xu, "Empowering diffusion models on the embedding space for text generation," 2022, *arXiv:2212.09412*.
- [9] C. Gardent, A. Shimorina, S. Narayan, and L. Perez-Beltrachini, "The WebNLG challenge: Generating text from RDF data," in *Proc. 10th Int. Conf. Natural Lang. Gener.*, J. M. Alonso, A. Bugarín, and E. Reiter, Eds. Stroudsburg, PA, USA: Association for Computational Linguistics, 2017, pp. 124–133. [Online]. Available: <https://aclanthology.org/W17-3518>
- [10] S. Gong, M. Li, J. Feng, Z. Wu, and L. Kong, "DiffuSeq: Sequence to sequence text generation with diffusion models," 2022, *arXiv:2210.08933*.
- [11] S. Gong, M. Li, J. Feng, Z. Wu, and L. Kong, "DiffuSeq-v2: Bridging discrete and continuous text spaces for accelerated seq2seq diffusion models," 2023, *arXiv:2310.05793*.
- [12] P. He, J. Gao, and W. Chen, "DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing," 2021, *arXiv:2111.09543*.
- [13] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, and T. Salimans, "Imagen video: High definition video generation with diffusion models," 2022, *arXiv:2210.02303*.
- [14] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi, "The curious case of neural text degeneration," 2019, *arXiv:1904.09751*.
- [15] E. Hoogeboom, A. A. Gritsenko, J. Bastings, B. Poole, R. van den Berg, and T. Salimans, "Autoregressive diffusion models," 2021, *arXiv:2110.02037*.
- [16] E. Hoogeboom, D. Nielsen, P. Jaini, P. Forré, and M. Welling, "Argmax flows and multinomial diffusion: Learning categorical distributions," 2021, *arXiv:2102.05379*.
- [17] M. Kale and A. Rastogi, "Text-to-text pre-training for data-to-text tasks," in *Proc. 13th Int. Conf. Natural Lang. Gener.*, Dublin, Ireland. Stroudsburg, PA, USA: Association for Computational Linguistics, Dec. 2020, pp. 97–102, doi: [10.18653/v1/2020.inlg-1.14](https://doi.org/10.18653/v1/2020.inlg-1.14). [Online]. Available: <https://aclanthology.org/2020.inlg-1.14>
- [18] P. Koehn, F. J. Och, and D. Marcu, "Statistical phrase-based translation," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics Hum. Lang. Technol. (NAACL)*, Jan. 2003, pp. 48–54, doi: [10.3115/1073445.1073462](https://doi.org/10.3115/1073445.1073462).
- [19] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "DiffWave: A versatile diffusion model for audio synthesis," 2020, *arXiv:2009.09761*.
- [20] X. Lisa Li, J. Thackstun, I. Gulrajani, P. Liang, and T. B. Hashimoto, "Diffusion-LM improves controllable text generation," 2022, *arXiv:2205.14217*.
- [21] Z. Lin, Y. Gong, Y. Shen, T. Wu, Z. Fan, C. Lin, N. Duan, and W. Chen, "Text generation with diffusion language models: A pre-training approach with continuous paragraph denoise," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 21051–21064.
- [22] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," 2017, *arXiv:1711.05101*.
- [23] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: A method for automatic evaluation of machine translation," in *Proc. 40th Annu. Meeting Assoc. Comput. Linguistics*, Philadelphia, PA, USA, P. Isabelle, E. Charniak, and D. Lin, Eds. Stroudsburg, PA, USA: Association for Computational Linguistics, Jul. 2001, pp. 311–318, doi: [10.3115/1073083.1073135](https://doi.org/10.3115/1073083.1073135).
- [24] A. P. Parikh, X. Wang, S. Gehrmann, M. Faruqui, B. Dhingra, D. Yang, and D. Das, "ToTTo: A controlled table-to-text generation dataset," 2020, *arXiv:2004.14373*.
- [25] R. Puduppully, L. Dong, and M. Lapata, "Data-to-text generation with content selection and planning," 2018, *arXiv:1809.00582*.
- [26] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," 2019, *arXiv:1910.10683*.
- [27] R. Ramamurthy, P. Ammanabrolu, K. Brantley, J. Hessel, R. Sifa, C. Bauckhage, H. Hajishirzi, and Y. Choi, "Is reinforcement learning (not) for natural language processing: Benchmarks, baselines, and building blocks for natural language policy optimization," 2022, *arXiv:2210.01241*.
- [28] S. Rothe, S. Narayan, and A. Severyn, "Leveraging pre-trained checkpoints for sequence generation tasks," *Trans. Assoc. Comput. Linguistics*, vol. 8, pp. 264–280, Dec. 2020.
- [29] A. See, P. J. Liu, and C. D. Manning, "Get to the point: Summarization with pointer-generator networks," in *Proc. 55th Annu. Meeting Assoc. Comput. Linguistics*, Vancouver, BC, Canada, R. Barzilay and M. Y. Kan, Eds. Stroudsburg, PA, USA: Association for Computational Linguistics, Jul. 2017, pp. 1073–1083, doi: [10.18653/v1/p17-1099](https://doi.org/10.18653/v1/p17-1099). [Online]. Available: <https://aclanthology.org/P17-1099>
- [30] T. Sellam, D. Das, and A. P. Parikh, "BLEURT: Learning robust metrics for text generation," 2020, *arXiv:2004.04696*.
- [31] Y. Su, D. Vandyke, S. Wang, Y. Fang, and N. Collier, "Plan-then-generate: Controlled data-to-text generation via planning," 2021, *arXiv:2108.13740*.
- [32] F. Wang, Z. Xu, P. Szekely, and M. Chen, "Robust (controlled) table-to-text generation with structure-aware equivariance learning," 2022, *arXiv:2205.03972*.
- [33] Y. Xiao, L. Wu, J. Guo, J. Li, M. Zhang, T. Qin, and T.-Y. Liu, "A survey on non-autoregressive generation for neural machine translation and beyond," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 10, pp. 11407–11427, Oct. 2023.
- [34] Y. Zhu, S. Lu, L. Zheng, J. Guo, W. Zhang, J. Wang, and Y. Yu, "Texygen: A benchmarking platform for text generation models," in *Proc. 41st Int. ACM SIGIR Conf. Res. Develop. Inf. Retr.*, Jun. 2018, pp. 1097–1100.



**ALEKSEI S. KRYLOV** received the B.S. degree in applied mathematics and computer science and the M.S. degree in machine learning and data analysis from Moscow Institute of Physics and Technology, Moscow, Russia, in 2022 and 2024, respectively. In 2020, he was a Summer Intern at Intel. From 2020 to 2022, he was a Researcher at the Laboratory of Information Technologies and Applied Mathematics, Moscow Institute of Physics and Technology. Since 2022, he has been a Data Scientist with SBER, Russia. His research interests include natural language processing and machine learning in risk management.



**OLEG D. SOMOV** received the B.S. degree in applied mathematics and computer science and the M.S. degree in machine learning and data analysis from Moscow Institute of Physics and Technology, Moscow, Russia, in 2018 and 2020, respectively, where he is currently pursuing the Ph.D. degree. He currently works as a Research Scientist with the Artificial Intelligence Research Institute (AIRI), Moscow.