

# Fact-Enhanced Synthetic News Generation

Kai Shu<sup>\*1</sup>, Yichuan Li<sup>\*2</sup>, Kaize Ding<sup>3</sup> and Huan Liu<sup>3</sup>

<sup>1</sup> Illinois Institute of Technology, Chicago, IL, USA

<sup>2</sup> Worcester Polytechnic Institute, Worcester, MA, USA

<sup>3</sup> Arizona State University, Tempe, AZ, USA

kshu@iit.edu, yli29@wpi.edu, {kaize.ding, huan.liu}@asu.edu

## Abstract

The advanced text generation methods have witnessed great success in text summarization, language translation, and synthetic news generation. However, these techniques can be abused to generate disinformation and fake news. To better understand the potential threats of synthetic news, we develop a novel generation method FACTGEN to generate high-quality news content. The majority of existing text generation methods either afford limited supplementary information or lose consistency between the input and output which makes the synthetic news less trustworthy. To address these issues, FACTGEN retrieves external facts to enrich the output and reconstructs the input claim from the generated content to improve the consistency among the input and the output. Experiment results on real-world datasets demonstrate that the generated news contents of FACTGEN are consistent and contain rich facts. We also discuss an effective defending technique to identify these synthetic news pieces if FACTGEN was used to generate fake news.

## Introduction

With the success of natural language processing, there has been a significant performance improvement in text generation applications, including document summarization (Gehrmann, Deng, and Rush 2018), machine translation (Johnson et al. 2017) and synthetic news generation (Leppänen et al. 2017). For example, we can use the generative adversarial network (GAN) (Aghakhani et al. 2018) or sequence-to-sequence (seq2seq) model (Yang et al. 2019b) to generate human-like comments. More recently, one approach named Grover (Zellers et al. 2019) has achieved promising result on synthetic news generation. It generates news pieces conditioned on multiple attributes such as headlines, authors, and website domains.

However, these methods could also be abused to generate and amplify disinformation and fake news. For example, the machine generated fake review threaten business reputations<sup>1</sup> and virtual characters sends generated story to spread

<sup>\*</sup>Equal contributions.

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup><https://bit.ly/349tPW2>

Claim	iran nuke framework agreement should be judged on merits, not disinformation.
Content	The united states and its negotiating partners reached a <b>very strong framework agreement with iran</b> in <i>lausanne , switzerland , on thursday</i> that limits iran 's nuclear program in such a way as to effectively <b>block it from building a nuclear weapon</b> . The debate that has already begun since the announcement of the new framework will likely result in <i>more heat than light</i> . It will not be helped by the <b>gathering swirl of dubious assumptions and doubtful assertions</b> .

Table 1: Example claim and the beginning part of a news pieces from CNN/DailyMail dataset. The bold sentence fragments are the consistent word and Itatic fragment is the supplementary information.

propaganda<sup>2</sup>. The wide dissemination of synthetic disinformation and fake news will bring new challenges to the news ecosystem. Therefore, it becomes critical to understand synthetic fake news for further achieving accurate detection.

In the real-world scenario, fake news deliberately imitates the writing styles of real news, which makes it hard to be identified by human and computational detection methods (Shu et al. 2020a). Both fake and real news usually contain additional facts<sup>3</sup> that are consistent and supplementary to the news claims. For example, in Table 1, the news mainly focuses on *framework agreement with Iran*, and provide additional facts like the location and time of the agreement. To eventually identify synthetic disinformation, from an adversarial perspective, we attempt to build a powerful synthetic news generation model by closing the inherent *factual* discrepancies between human and machine-generated text. Existing methods on generating synthetic news may fall short with the following limitations: (1) *factual inconsistency*, indicating the generated news contradict or refute the news claims; and (2) *factual scarcity*, meaning the generated news content may miss essential details to supplement the claim. However, directly using or fine-tuning language models does not help as it is non-trivial to enhance factual consistency

<sup>2</sup><https://bit.ly/36if2e9>

<sup>3</sup>We follow the definition of *fact* as, according to Oxford Dictionary, the information used as evidence or as part of news article.

and richness on a language model directly. Therefore, in this study, we aim to address the following challenges in synthetic news generation: (1) how to generate news content related to a given claim/context; and (2) how to ensure that the generated content contains supplemental fact information.

Our solution to these challenges results in a novel framework FACTGEN<sup>4</sup> (Fact-Enhanced Synthetic News Generation). FACTGEN consists of three major components: (1) *Pseudo-Self-Attentive (PSA) Language Model* (Ziegler et al. 2019), where the customized encoder deceptively injects source information (claims and external facts) into pre-trained decoder for the generation. The adapted deceptive injection mechanism can resolve the mismatch between the untrained encoder and the well-trained decoder; (2) *Fact Retriever*, which heuristically retrieves the supplemental information from external fact corpus to provide more candidate facts during generation; and (3) *Claim Reconstructor*, a randomly initialized masked language model (Devlin et al. 2018a) which enhances the output consistency by reconstructing the masked claim tokens from both the representation of the generated content and the unmasked claim tokens. During training, the *PSA Language Model* takes the news claim and the retrieved facts from the *Fact Retriever* as input, then generates highly consistent news content by incorporating the *Claim Reconstructor* into the generation process. In this way, the proposed framework can generate both fact-consistent and fact-enriched news content. To summarize, our main contributions are as follows:

- We study a novel problem of fact-enhanced synthetic news generation, which aims to generate consistent and fact-enriched news content.
- We propose a principled framework FACTGEN generates realistic synthetic news by retrieving external facts and reconstructing the input claim.
- We conduct experiments on real-world datasets using quantitative and qualitative metrics to demonstrate the effectiveness of FACTGEN for synthetic news generation and its defense.

## Methodology

Our goal is to incorporate external facts into news generation that are consistent with the news claim. Given a sequence of tokens from the claim  $X = \{x_1, x_2, \dots, x_N\}$ , the fact retriever retrieves related fact information  $F = \{f_1, f_2, \dots, f_K\}$  by semantic similarity, then the language model generates the news content  $Y = \{y_1, y_2, \dots, y_M\}$  based on claim  $X$  and  $F$ . It should be noticed that the length of  $Y$  is much larger than  $X$ , which is  $M \gg N$ , and  $x_i, y_i, f_i$  are words. Figure 1 illustrates the architecture of the proposed model and the objective functions. The causal language loss  $L_{CLL}$  depicts the loss of generating news content based on the input claim and fact. The masked language loss  $L_{MLL}$  is to reconstruct the masked input claim based on the language model output and the unmasked claims. This has a twofold benefit. Initially, the pre-trained decoder and the retrieved facts will bring unrelated information. This

<sup>4</sup>The code is available at <https://github.com/bigheiniu/FactGen>




Figure 1: The proposed model, FACTGEN. The black dashed line indicates no differential dependency and the black bold line otherwise.  $\oplus$  is the text concatenation.

technique encourages the generated content to cover the input claim and provides a regularization effect. In addition, it is fully differentiable so we can minimize the objective function end-to-end. Overall, we minimize:

$$L = L_{CLL} + \lambda L_{MLL} \quad (1)$$

where  $\lambda$  is the hyperparameter to control the contribution of claim reconstruction. The formulas of  $L_{CLL}$  and  $L_{MLL}$  are in Eq. 5 and Eq. 6 respectively.

## Preliminary

Self-attentive (SA) language models (Devlin et al. 2018b; Radford et al. 2019; Song et al. 2019) have achieved impressive performance gains in various language generation tasks. These models are stacks of several SA blocks which encode the input  $X = \{x_1, \dots, x_i, \dots, x_N\}$  into key-value pairs  $(K, V) = \{(k_1, v_1), \dots, (k_i, v_i), \dots, (k_N, v_N)\}$  and query  $Q = \{q_1, \dots, q_i, \dots, q_N\}$ . The next output is produced by taking the weighted sum of values  $v_i$ , where the weight assigned toward each value is the dot-product of the query  $Q$  with all the keys  $K$ . The formula of SA is:

$$K = H_X W_k, \quad V = H_X W_v, \quad Q = H_X W_q \quad (2)$$

$$SA(X) = \text{softmax}(QK^T) V \quad (3)$$

where  $H_X \in \mathbb{R}^{N \times D}$  is the hidden representation of the input  $X$ ,  $D$  is the hidden dimension, and  $W_k, W_v, W_q \in \mathbb{R}^{D \times D}$  are the parameters to map the hidden representation of tokens  $H_X$  into key, value and query space, respectively.

## Proposed Method

**Pseudo-Self-Attentive Language Model:** Although the fine-tuned self-attentive language models like GPT-2 (Radford et al. 2019) have been applied to many text generation tasks, the application of using GPT-2 for the synthetic news generation may not be satisfactory. Since the GPT-2 is an autoregressive model, only encoding the forward information, it will lose backward information from the input. Besides, without a specific encoder, GPT-2 cannot capture the dependent relationship between the news claim and the retrieved facts, which will hurt the performance of the decoder (Edunov, Baevski, and Auli 2019). Therefore, we need a new encoder to capture bi-directional information and dependency among the input.

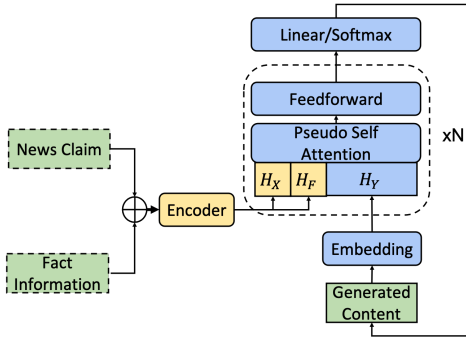


Figure 2: Our pseudo-self-attentive language model. Best visualized in color. The blue indicates decoder’s pre-trained parameters. The yellow indicates the randomly initialized parameters of the encoder.  $N$  is the number of PSA blocks.

Following (Ziegler et al. 2019)’s setting, we employ a pseudo-self-attentive(PSA) language model, where the “pseudo” is that the encoder deceptively extends the decoder’s key-value pairs by the encoder’s pairs, and the decoder predicts the next token not only based on previous output tokens but also from the input. To model the dependency between claim and retrieved facts, we wrap them with “[Claim]” and “[Fact]” separately, and specifically all the retrieved facts are contacted together without any special separation token. The architecture of the language model is shown in Figure 2 and the formula of PSA is:

$$\text{PSA}(Y, X, F) = \text{softmax} \left( Q_Y \begin{bmatrix} K_Y \\ K_X \\ K_F \end{bmatrix}^\top \right) \begin{bmatrix} V_Y \\ V_X \\ V_F \end{bmatrix} \quad (4)$$

Note that  $K_X, K_F, V_X,$  and  $V_F$  are using different projection matrix  $W_*$  and are randomly initialized. The objective function of the language model is:

$$L_{CLL} = - \sum_{i=1}^M (\log P(y_i | y_1, \dots, y_{i-1}; X, F)) \quad (5)$$

**Fact Retriever:** Directly training a sequence to sequence model on  $(X, Y)$  often results in fact scarcity. One main reason is that facts from the input are extremely insufficient compared to the output. Thus, the language model is more likely to generate repeated sentences. Our solution towards the facts imbalance between the input and output is to increase the facts in the source side by retrieving related facts and considers them as part of the input. Our fact retriever (FR) heuristically retrieves external facts in two steps. Firstly, to omit the computation limitation, we retrieve the related document based on the tf-idf vectors’ cosine similarity between the claim and the document. Here we only keep the  $top-k_1$  similar documents. Secondly, to accurately identify related sentences in the document, we utilize the pre-trained BERT (Devlin et al. 2018b) to encode all sentences presented in the picked documents and choose the  $top-k_2$  most similar sentences based on the cosine similarity.

**Claim Reconstructor:** Since the aforementioned modules FR and PSA language model will bring inconsistency during

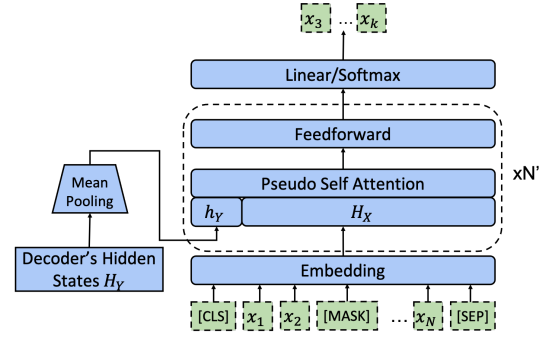


Figure 3: The overview of the claim reconstructor. It is also the PSA structure where we inject the mean pooling of decoder’s hidden states to its key-value pairs.  $N'$  is the number of PSA blocks.

the generation, there still needs an extra mechanism to guarantee the consistency between the input news claims and the generated news content. We propose to reconstruct the claim from the generated content through masked language model.

The existing reconstruction approaches for the consistent generation require the prior knowledge of the input, such as utilizing the topic label to learn a topic consistent reward function (Yang et al. 2019a), or key-entities for multi-classification on the hidden states to entail the key information (Wiseman, Shieber, and Rush 2017). Our claim reconstructor (CR) does not require any prior knowledge about the input. It reconstructs the masked claim  $X_{[Masked]}$  based on the mean pooling of output hidden representation  $h_Y$  and unmasked sentence fragments  $X_{[Unmasked]}$ . We mask claim’s tokens with probability  $P_{mask}$  and we follow the pseudo-self-attention (Ziegler et al. 2019) projecting  $h_Y$  into CR’s key-value pairs to predict the masked sentence fragment,  $X_{[Masked]}$ . The objective function of CR is:

$$L_{MLL} = \sum_{x \in X_{[Masked]}} -\log P(x | X_{[Unmasked]}, h_Y) \quad (6)$$

## Training Schedule

Since FACTGEN needs to guarantee that there is no contradiction between the factual consistency and factual richness, we cannot directly train the model via minimizing eq 1. We then train FACTGEN in two-stages. The overview of the training procedure is summarized in Algorithm 1. The two stages of training bring several advantages: firstly, it allows to start the PSA language model and CR warmly, omitting the gradient explosion problem during training; secondly, because the claim is the main idea of the generated text and the retrieved facts are the auxiliary information during the generation, this order can help the decoder understand the importance of different input sources. The joint training can align the latent space of these two modules.

## Experiments

In this section, we conduct experiments on real-world datasets to demonstrate the effectiveness of FACTGEN for news generation.

---

**Algorithm 1** Training Procedure of FACTGEN

---

**Input:** The source claims, relevant facts and target news pieces corpus  $S = \{(X, F, Y)\}$ ; the masked and unmasked claims  $D = \{X_{[Masked]}, X_{[Unmasked]}\}$ ; first and second stage epoch number  $epochs_1$  and  $epochs_2$ .

**Output:**  $PSA$  language model and claim reconstructor  $CR$ ;

- 1: Initialize  $PSA_{encoder}$  and  $CR$  with random weights
  - 2: Pre-train the  $PSA$  via minimizing eq.5 on  $\{(X, Y)\}$ ;  
Pre-train the  $CR$  via minimizing eq.6 on  $D$ .
  - 3: **for**  $epoch = 1$  to  $epochs_1$  **do**
  - 4:   Jointly training  $PSA$  and  $CR$  via minimizing eq.1 on  $\{(X, Y)\}$  and  $D$ ;  $\triangleright$ First Stage
  - 5: **end for**
  - 6: **for**  $epoch = 1$  to  $epochs_2$  **do**
  - 7:   Jointly training  $PSA$  and  $CR$  via minimizing eq.1 on  $S$  and  $D$ ;  $\triangleright$ Second Stage
  - 8: **end for**
- 

## Dataset

We utilize two news datasets in our experiment. The first dataset is a widely used fake news detection dataset collected from a fact-checking website, GossipCop (Shu et al. 2020b). Each sample contains the news’ claim, content, metadata, label, and social engagements. The average lengths of the claim and content are 30 words and 250 words respectively. The second dataset is the CNN/DailyMail news highlight dataset (Hermann et al. 2015) which contains the news content and selected highlight. In contrary to the text summarization, we use the highlight sentence as the source claim and the news content as the target text. On average, the claim has 56 tokens and the content has 790 tokens. As for pre-processing, we truncate the news claim longer than 100 words and content longer than 300 words in both datasets. For the dataset splitting, we randomly sample 75% training set, 15% validation set, and 10% test set in the GossipCop dataset and follow the same splitting setting in (See, Liu, and Manning 2017). Datasets’ statistical information is listed in Table 2.

We consider the factual sentences in the training dataset as our external fact corpus with the following reasons: (1) utilizing several sentences instead of whole news pieces can avoid the model learning from copy the information from the source to the target side; (2) the fact sentences from the training dataset can omit the data leakage problem during testing. We focus on external facts as the format of the text, though it can be extended to tabular data or knowledge graph.

Dataset	# of train	# of val	# of test
GossipCop	7,331	1,459	974
CNN/DailyMail	278,408	11,490	13,368

Table 2: The statistical information of the datasets.

## Experiment Settings

We implement FACTGEN on OpenNMT (Klein et al. 2017). We tune the hyper-parameter  $\lambda$  on the validation set. The encoder of FACTGEN is 4 blocks of SA block with 12 attention heads and 3072 hidden units. The weight of the decoder is initialized with the median pre-trained GPT-2 (Radford et al. 2019) model. The claim reconstruction module is 3 blocks of SA block with 4 attention heads and 256 hidden sizes and  $P_{mask}$  is set as 0.5. The optimizer is Adam (Kingma and Ba 2014) with  $\beta_1 = 0.9$  and  $\beta_2 = 0.998$ . It should be noticed that the learning rate for the encoder is  $1e - 3$ , for the decoder is  $1e - 5$ , and  $5e - 5$  for the claim reconstruction. The number of retrieved documents  $k_1$  and sentences  $k_2$  is set to 10 and 5 respectively. The  $epochs_1$  and  $epochs_2$  in the training schedule are set to 4 and 2 respectively. During decoding we used Nucleus Sampling (top- $p$ ) with  $p = 0.9$ .

## Evaluation Metrics

**Automatic Evaluation** The traditional text generation metrics like BLEU (Papineni et al. 2002) and ROUGE (Lin 2004) which are focus on the overlap between the generated content and the reference text which is not enough to reflect the claim-content consistency and the richness of the generated content. To remedy this, we develop two new evaluation metrics to measure the quality from different perspectives.

- **Fluency:** we report the BLEU-4 score for the text fluency.
- **Consistency:** The ideal news content should support its claim. Therefore, we propose a stance detection model to detect whether the content is in favor of the claim or against it. Given the claim and the generated news content  $\{X, Y\}$ , the stance detection model will output the relation of the text pair in (*Agrees, Disagrees, Discusses, Unrelated*). We utilize the Fake News Challenge dataset<sup>5</sup> to fine-tune RoBERTa (Liu et al. 2019). This approach achieves a 0.93 accuracy score on the test dataset of the Fake News Challenge. We report the ratio of the “agrees”:

$$Consistency = \frac{\# \text{ of agree samples}}{\# \text{ of all samples}} \quad (7)$$

- **Richness:** The richness of the output can be evaluated by the number of unique name entities in the generated text (Fan, Lewis, and Dauphin 2019). We utilize spaCy<sup>6</sup> to extract the named entity from the output.

**Human Evaluation** We distribute the 100 generated samples in CNN/DailyMail dataset to 2 annotators with a linguistic background. They have no advanced knowledge about the source of the generated content. They are asked to evaluate the generated content from fluency, richness, consistency and trustworthiness 4 different perspectives. So totally, there are 7,200 evaluation questions in our human evaluation. The annotator should answer each question from a score of 1 to 3 (3 being the best, 1 being the worst).

<sup>5</sup><http://www.fakenewschallenge.org/>

<sup>6</sup><https://spacy.io/>

Models	GossipCop			CNN/DailyMail		
	<i>Fluency</i>	<i>Richness</i>	<i>Consistent</i>	<i>Fluency</i>	<i>Rich</i>	<i>Consistent</i>
CopyTransformer	0.2	11.0	0.04	0.5	9.5	0.66
ConvSeq2seq	0.5	5.9	0.09	3.3	9.5	0.44
PPLM	0.7	12.5	0.67	0.8	13.1	0.68
GPT-2	0.8	13.4	0.35	1.65	13.5	0.70
Grover	1.2	<b>15.7</b>	0.56	0.3	15.3	0.72
<b>FACTGEN</b>	<b>2.1</b>	14.5	<b>0.80</b>	<b>4.6</b>	<b>16.6</b>	<b>0.76</b>

Table 3: The performance comparison for the quality of the generated news pieces.

## Baseline Methods

To demonstrate the quality of the generated text, we compare our proposed model on content quality with the following text generation models: **CopyTransformer** (See, Liu, and Manning 2017): a sequence-to-sequence transformer with a pointer network that can copy the word from the source to the target; **Conv Seq2Seq** (Fan, Lewis, and Dauphin 2018): it utilizes the seq2seq convolution neural network to generate claim consistent stories; **PPLM** (Dathathri et al. 2019): a topic and content controlled language model; **GPT-2** (Radford et al. 2019): a large pre-trained language model which is the decoder part of the transformer. For a fair comparison with our model, we utilize the median size of the model; **Grover** (Zellers et al. 2019): generating news text conditioned on the news title, authors, and website domains.

## Experimental Results

The automatic and human evaluation results are shown in Table 3 and 4, respectively. We evaluate the quality of the text generation through the following perspectives:

- **Fluency:** From the human evaluations on fluency in CNN/DailyMail dataset and fluency score in two datasets, we can find that our model achieves the best performance. In the meantime, we find that the pre-trained language model achieves better human evaluation results than the model trained from scratch (PPLM, GPT-2, Grover > CopyTransformer, Conv Seq2Seq). This indicates the importance of incorporating the large pre-trained language model in the synthetic news generation. Besides, FACTGEN’s performance indicates the pseudo-self attention properly connecting the randomly initialized encoder and the pre-trained decoder.
- **Consistency:** The consistent result in both human evaluation and automatic evaluation demonstrates the effectiveness of our approach. Especially, in the GossipCop dataset, our approach achieves 42% performance improvement over the best baseline in automatic metric, and in CNN/DailyMail, the human evaluation also shows that our approach achieves 6% performance improvement compared with the best baseline method. The main reason for the increase is reconstructing the claim increases the coverage of the output on the input information.
- **Richness:** Our approach achieves the best performance in CNN/DailyMail dataset and the second performance in the GossipCop dataset. The reason for the ordinary

performance in GossipCop is that the size of the candidate documents in GossipCop is much smaller than the CNN/DailyMail (7,331 < 278,408). The FR cannot retrieve enough related facts from the external corpus and CR will reject the inconsistent facts during generation. This indicates that FR can bring rich facts in generation.

- **Trustworthiness:** Human evaluation of the Trustworthiness of synthetic news content indicates that overall, FACTGEN can generate high-quality text content. This helps us to understand the difference between machine-generated news content and true news in the future.

## Case Study

One case study of the generated samples is listed in Table 5. We only reveal the output from the model with pre-trained language models and we have several observations: (i) Our model mainly talks about the *agreement of nuclear weapons in Iran* and includes the supplemental information about *Iraq and UK’s action toward nuclear weapons*. This brings more context information about the news claims and makes the generated news more convincing. (ii) Although Grover mentions much additional factual information, it is unrelated to the *nuclear agreement with Iran*. (iii) The outputs of GPT-2 and PPLM mainly discuss the *nuclear agreement* without supplemental information about the agreement.

## Ablation Study

**Impact of  $\lambda$ :** To learn the impact of the hyper-parameter  $\lambda$  in our objective function in Eq. 1, we change  $\lambda$  from {0.001, 0.01, 0.1, 1, 10} and calculating all the automatic

Methods	<i>Fluency</i>	<i>Rich</i>	<i>Consistent</i>	<i>Trust</i>
CopyTransformer	1.68	1.65	1.89	1.62
ConvSeq2seq	1.95	2.12	2.00	1.94
PPLM	1.96	1.77	1.96	1.92
GPT-2	2.03	<b>2.32</b>	1.95	2.08
Grover	2.08	2.15	1.78	1.97
<b>FACTGEN</b>	<b>2.17</b>	2.28	<b>2.12</b>	<b>2.18</b>
Correlation	0.14	0.26	0.21	0.21

Table 4: The human evaluation result of generated samples in the CNN/DailyMail dataset. We calculate the Pearson correlation to show the inter-annotator agreement.

Claim	iran framework agreement on nukes ...opponents cast doubts. ...uranium stockpile, centrifuges, implement rigorous inspections; ...judged on merits, not disinformation.
PPLM	i think the agreement has been agreed on the whole of <i>American and nauru</i> ,and is a good first step towards an <b>agreement</b> on other issues. ...cut its <b>stockpiles</b> and implement some strict inspections. i think this is a good first step towards an <b>agreement</b> .i think <i>the deal is good but its not done yet and it still needs to be signed and accepted by america</i> ,and it will have to meet certain conditions for american to accept its deal on <b>nuclear</b>
GPT-2	The <b>iran nuclear</b> program does not have a single day's notice of a nuclear crisis, and it does not have a single day's notice of a <b>nuclear weapon</b> program. ...Last week, the government of <b>iran</b> signed a framework agreement that <i>allows the permanent establishment of the nuclear program</i> at all levels ...close its existing <b>enrichment plant</b> , ... cooperation and the reduction of the <b>stockpile</b> .
Grover	Unlike (allegedly) other <i>traitors who've betrayed the established foreign policy approach</i> to the Iranian nuclear development , in addition to <i>erasing himself from IMDB, with a stigmata, an air cannon</i> , ... published his own book for reasons which, honestly, I'm not sure are entirely clear. The main key player in the <b>NPT negotiations</b>
FACTGEN	<b>fichard klass</b> has joined the <i>british establishment</i> to try to prevent the sale of <b>nukes</b> after <b>failing to address its facts</b> on how the <b>iran regime</b> used them to try to arm its <b>nuclear programme</b> . <i>the businessman, who once ran a world bank and was governor of iraq</i> , said the nukes would cut <b>uranium stockpile, implement a rigorous inspections</b> ... and try to defend its nuclear program. while <b>opponents</b> said a <i>strong nuclear test would help persuade the uk to develop nuclear weapons</i> , <b>klass</b> said he hoped most sanctions would be <b>judged on merits, rather than disinformation</b> .

Table 5: Samples of generated fake news with different fake news generation methods. The bold sentence fragment is the consistent word and Italic fragment is the supplementary information.

evaluation metrics. From Figure 4 we can find that  $\lambda = 0.001$  achieves the best performance across all the automatic evaluations and with the increase of  $\lambda$ , the fact richness has been greatly decreased. This is because the CR will constrain the coverage of the generated content and cause the language model to only generated content around the input, which will reduce the richness of the generated content.

**Impact of Model Components:** To evaluate the importance of each key components, we set up three different ablation studies of FACTGEN: without claim reconstruction (*w/o CR*), without Fact Retriever (*w/o FR*) and without these two components (*w/o CR and FR*). It should be noticed that all versions of the model have been pre-trained on  $\{X, Y\}$ . The automatic and human evaluation in Table 6 and Table 7 show that the performance decrease in all ablation study. However, an interesting finding is that there seems to have a contradiction between the CR and FR. From Table 6, we find that *w/o CR* contains the richest fact information but has the lowest consistency score; *w/o FR* achieves the best flu-




Figure 4: Impact of hyper-parameter  $\lambda$  in CNN/DailyMail.

ency score and compatible consistency score but the worst richness score. The impact of CR matches the observation of hyperparameter analysis, which improves the consistency of the generated content while decreases the fact richness. These results indicate the effectiveness of CR and FR in improving the richness and consistency in the generation.

**Impact of Training Schedule:** To understand the effectiveness of our two-stage training schedule, we compare it with single-stage training where the model directly takes the claims and external fact information in the first stage. From the automatic and human evaluation result in Table 6 and 7, we can find that two stages training schedule achieve better performance in all categories compared with single-stage. This stipulates the effectiveness of our training schedule.

## Further Analysis

### Difficulty of Defending Synthetic Fake News

To understand the difficulty in synthetic fake news detection, we test the fake news detection methods and synthetic generation detection method on generated fake news and human-written real text. For the fake news classification, we utilize two state-of-the-art content-based fake news detection methods MWSS-CNN<sup>7</sup> (Shu et al. 2020c) and EANN<sup>8</sup> (Wang et al. 2018) trained on GossipCop training dataset containing news content and veracity news label. For the neural text classification, we utilize the RoBERTa (Liu et al. 2019), trained on 2000 GPT-2 machine-generated samples and human written WebText<sup>9</sup> respectively. To give limited access to generated content, the training dataset for both approaches will include extra 100 fake synthetic news pieces. The rea-

<sup>7</sup><https://github.com/microsoft/MWSS>

<sup>8</sup><https://github.com/yaqingwang/EANN-KDD18>

<sup>9</sup><https://github.com/openai/gpt-2-output-dataset>

Methods	<i>Fluency</i>	<i>Rich</i>	<i>Consistent</i>
Full Model	4.7	16.8	<b>0.76</b>
<i>single-stage</i>	4.6	16.6	0.74
<i>w/o CR</i>	3.4	<b>17.7</b>	0.73
<i>w/o FR</i>	<b>5.1</b>	11.7	0.75
<i>w/o CR and FR</i>	4.0	12.9	0.73

Table 6: Results of automatic evaluation of model components ablation study in CNN/DailyMail dataset.

son for different training datasets for these approaches is to test whether the fake news detection model can transfer the knowledge in human written fake news into machine-generated fake news. To guarantee the veracity of the test content, we select the fake generated content which is conditioned on fake claims and human-written real text is from real news pieces in GossipCop.

We test the performances in 300 fake generated news contents and the same amount of human-written real news content. To omit the data leakage problem for evaluation, the test dataset is only used for the generation evaluation. From the result in Table 8, we observe that fake news detection methods achieve worse performance than neural text classification (RoBERTaA > EANN, MWSS-CNN) which indicates the difficulty of the current fake news detection method in detecting fake synthetic news.

### Defending Against Synthetic Fake News

To detect the new synthetic fake news, we follow (Zellers et al. 2019) develop a defending method FACTGEN<sub>def</sub> based on the checkpoint of FACTGEN at iteration 20k. This setting can reduce the parameters overlap between the generator and the discriminator. We also use  $h_Y$  as the final representation of the input, synthetic fake news or human written real news, and add a full connection layer to classify whether the input is fake or real. We utilize 100 synthetic fake news content and the same amount of human written real news to fine-tune FACTGEN<sub>def</sub>. The result in Table 8 shows FACTGEN<sub>def</sub> achieves the best accuracy score. This is because FACTGEN<sub>def</sub> can learn a better representation of the input. We thus conclude that while the synthetic content is hard to be identified by existing methods, it still can be detected by FACTGEN<sub>def</sub>.

Methods	<i>Fluency</i>	<i>Rich</i>	<i>Consistent</i>	<i>Trust</i>
Full Model	<b>2.17</b>	2.28	<b>2.12</b>	<b>2.18</b>
<i>single-stage</i>	2.01	2.22	2.10	2.14
<i>w/o CR</i>	2.15	<b>2.31</b>	2.09	2.15
<i>w/o FR</i>	1.93	2.28	2.03	1.98
<i>w/o CR and FR</i>	2.09	2.19	2.03	2.10

Table 7: Results of human evaluation of model components ablation study in CNN/DailyMail dataset.

	EANN	MWSS-CNN	RoBERTa	FACTGEN <sub>def</sub>
Acc	0.64	0.58	0.74	<b>0.82</b>

Table 8: Results of synthetic fake news content detection.

### Related Work

**Synthetic News Generation** Most synthetic news generation systems used in the newsroom are heavily rule-based and template-based (Leppänen et al. 2017). The neural synthetic news generation like Grover (Zellers et al. 2019) utilizes an autoregressive language model to learn the dependency among news metadata fields include the domain, date, authors, title, and body. Sam and et. al (Wiseman, Shieber, and Rush 2017) propose a structured data to text challenge which is to generate a sport a news piece of sports games from the associated box- or line-score data. To better capture the input data, (Wiseman, Shieber, and Rush 2017) employs copy-mechanism and source reconstruction as their seq2seq model extensions and (Puduppully, Dong, and Lapata 2019) generate text in recording plan and realization two stages.

**Synthetic/Fake News Detection** Content-based fake news detection methods often leverage features from the feature engineering or latent features extracted by deep neural network (Shu and Liu 2019; Pérez-Rosas et al. 2017). Deep learning models utilize linguistic representation of news content to detect fake news. (Qian et al. 2018) proposes a method learning the representation of news content and reconstructing the users comment during training, and in inference, this model makes a classification based on the representation of news content and the generated news comment for early fake news detection. (Schuster et al. 2020) stipulates that current synthetic disinformation detection methods are mainly based on the stylometry which is limited against machine-generated misinformation. Dirk Hovy et al. proposes an adversarial setting in detecting the generated reviews (Hovy 2016). Gehrmann et al. (Gehrmann, Strobel, and Rush 2019) visualize the distribution of words that help non-expert users recognize generated text. (Zellers et al. 2019) and (Solaiman et al. 2019) propose neural generation detectors that fine-tune classifiers on the generator’s previous checkpoint. (Uchendu et al. 2020) proposes to differentiate the sources of natural language generation methods.

### Conclusion and Future Work

We propose a synthetic news generation method FACTGEN to ensure fact-consistency and fact-richness. We demonstrate FACTGEN is more effective than existing methods with extensive evaluation. We discuss the difficulty of detecting synthetic fake news and propose a defending method FACTGEN<sub>def</sub> that achieves outstanding performance in detecting synthetic fake news content. In the future, we plan to include other formats of facts like tabular or knowledge graphs. This can help us retrieve up-to-date fact information during generation. Since fake news often contains catchy information to widely spread on the social network, we would like to explore the style control of the generated content to make it prone to be spread.

## Acknowledgments

This work is, in part, supported with funding from the Defense Advanced Research Projects Agency (DARPA) under Contract No. HR001120C0123, and the John S. and James L. Knight Foundation through a grant to the Institute for Data, Democracy & Politics at The George Washington University. The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

## Appendices on Ethics Statement

To better understand the characteristics of synthetic fake news, we propose a fact-enriched synthetic news generation method to generate high quality news pieces. From the automatic and human evaluation results, we find that FACTGEN can generate human-like and convincing news pieces. In this paper, we also discuss a possible solution to defend this attack, which is to use the checkpoint of FACTGEN. We are discussing the further usage of FACTGEN and ethical concerns as follows:

**Journalism Assistants:** Since our method retrieves the external fact information and generate fact-consistent and fact-enriched news pieces, the journalists can utilize FACTGEN to automatically generate news by providing additional factual information and the claim. However, it still needs manually checking (Leppänen et al. 2017).

**Synthetic Disinformation Detection:** In this paper, we shortly discuss the defending method, FACTGEN<sub>def</sub>, and prove the effectiveness of it. However, like the Grover (Zellers et al. 2019), this method mainly relies on semantic information rather than the veracity of the information (Schuster et al. 2020). Future work should verify the factual correctness of the text in the following pipeline: check-worthy sentence extraction, the verified claim matching, and prediction (Adair et al. 2019).

**Release Policy:** Since FACTGEN can generate human-like and convincing news content, we need to critically release the code and the model parameters. We propose to publicly release the code including generator and discriminator. However, as for the checkpoints of both models, we will only share for academic usage.

## References

Adair, B.; Li, C.; Yang, J.; and Yu, C. 2019. Automated Pop-ÅUp Fact-ÅChecking: Challenges & Progress. In *Proceedings of the Computation+ Journalism Symposium*.

Aghakhani, H.; Machiry, A.; Nilizadeh, S.; Kruegel, C.; and Vigna, G. 2018. Detecting deceptive reviews using generative adversarial networks. In *2018 IEEE Security and Privacy Workshops (SPW)*, 89–95. IEEE.

Dathathri, S.; Madotto, A.; Lan, J.; Hung, J.; Frank, E.; Molino, P.; Yosinski, J.; and Liu, R. 2019. Plug and play language models: A simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018a. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018b. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Edunov, S.; Baevski, A.; and Auli, M. 2019. Pre-trained language model representations for language generation. *arXiv preprint arXiv:1903.09722*.

Fan, A.; Lewis, M.; and Dauphin, Y. 2018. Hierarchical neural story generation. *arXiv preprint arXiv:1805.04833*.

Fan, A.; Lewis, M.; and Dauphin, Y. 2019. Strategies for structuring story generation. *arXiv preprint arXiv:1902.01109*.

Gehrmann, S.; Deng, Y.; and Rush, A. M. 2018. Bottom-up abstractive summarization. *arXiv preprint arXiv:1808.10792*.

Gehrmann, S.; Strobelt, H.; and Rush, A. M. 2019. Gltr: Statistical detection and visualization of generated text. *arXiv preprint arXiv:1906.04043*.

Hermann, K. M.; Kočiský, T.; Grefenstette, E.; Espeholt, L.; Kay, W.; Suleyman, M.; and Blunsom, P. 2015. Teaching machines to read and comprehend. *arXiv preprint arXiv:1506.03340*.

Hovy, D. 2016. The enemy in your own camp: How well can we detect statistically-generated fake reviews—An adversarial study. In *ACL*.

Johnson, M.; Schuster, M.; Le, Q. V.; Krikun, M.; Wu, Y.; Chen, Z.; Thorat, N.; Viégas, F.; Wattenberg, M.; Corrado, G.; et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *TACL*.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Klein, G.; Kim, Y.; Deng, Y.; Senellart, J.; and Rush, A. M. 2017. Opennmt: Open-source toolkit for neural machine translation. *arXiv preprint arXiv:1701.02810*.

Leppänen, L.; Munezero, M.; Granroth-Wilding, M.; and Toivonen, H. 2017. Data-driven news generation for automated journalism. In *NLG*.

Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.

Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.

Pérez-Rosas, V.; Kleinberg, B.; Lefevre, A.; and Mihalcea, R. 2017. Automatic detection of fake news. *arXiv preprint arXiv:1708.07104*.

Puduppully, R.; Dong, L.; and Lapata, M. 2019. Data-to-text generation with content selection and planning. In *AAAI*.



- Qian, F.; Gong, C.; Sharma, K.; and Liu, Y. 2018. Neural User Response Generator: Fake News Detection with Collective User Intelligence. In *IJCAI*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1(8): 9.
- Schuster, T.; Schuster, R.; Shah, D. J.; and Barzilay, R. 2020. The limitations of stylometry for detecting machine-generated fake news. *Computational Linguistics* 46(2): 499–510.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Shu, K.; Bhattacharjee, A.; Alatawi, F.; Nazer, T. H.; Ding, K.; Karami, M.; and Liu, H. 2020a. Combating disinformation in a social media age. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 10(6): e1385.
- Shu, K.; and Liu, H. 2019. Detecting fake news on social media. *Synthesis lectures on data mining and knowledge discovery* 11(3): 1–129.
- Shu, K.; Mahudeswaran, D.; Wang, S.; Lee, D.; and Liu, H. 2020b. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big Data* 8(3): 171–188.
- Shu, K.; Zheng, G.; Li, Y.; Mukherjee, S.; Awadallah, A. H.; Ruston, S.; and Liu, H. 2020c. Leveraging Multi-Source Weak Social Supervision for Early Detection of Fake News.
- Solaiman, I.; Brundage, M.; Clark, J.; Askill, A.; Herbert-Voss, A.; Wu, J.; Radford, A.; Krueger, G.; Kim, J. W.; Kreps, S.; et al. 2019. Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203*.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2019. Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.
- Uchendu, A.; Le, T.; Shu, K.; and Lee, D. 2020. Authorship Attribution for Neural Text Generation. In *EMNLP*.
- Wang, Y.; Ma, F.; Jin, Z.; Yuan, Y.; Xun, G.; Jha, K.; Su, L.; and Gao, J. 2018. EANN: Event Adversarial Neural Networks for Multi-Modal Fake News Detection. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery Data Mining, KDD '18*. New York, NY, USA: Association for Computing Machinery. ISBN 9781450355520. doi:10.1145/3219819.3219903. URL <https://doi.org/10.1145/3219819.3219903>.
- Wiseman, S.; Shieber, S. M.; and Rush, A. M. 2017. Challenges in data-to-document generation. *arXiv preprint arXiv:1707.08052*.
- Yang, P.; Li, L.; Luo, F.; Liu, T.; and Sun, X. 2019a. Enhancing topic-to-essay generation with external common-sense knowledge. In *ACL*.
- Yang, Z.; Xu, C.; Wu, W.; and Li, Z. 2019b. Read, attend and comment: a deep architecture for automatic news comment generation. *arXiv preprint arXiv:1909.11974*.
- Zellers, R.; Holtzman, A.; Rashkin, H.; Bisk, Y.; Farhadi, A.; Roesner, F.; and Choi, Y. 2019. Defending against neural fake news. *arXiv preprint arXiv:1905.12616*.
- Ziegler, Z. M.; Melas-Kyriazi, L.; Gehrmann, S.; and Rush, A. M. 2019. Encoder-agnostic adaptation for conditional language generation. *arXiv preprint arXiv:1908.06938*.