

HACovGNN: Hierarchical Attention Based Convolutional Graph Neural Network for Code Documentation Generation in Jupyter Notebooks

Xuye Liu ‡
University of Waterloo

Dakuo Wang ‡
IBM Research

April Yi Wang
University of Michigan

Yufang Hou
IBM Research Europe

Lingfei Wu *
JD.COM Silicon Valley Research Center

Abstract

Jupyter notebook allows data scientists to write machine learning code together with its documentation in cells. In this paper, we propose a new task of code documentation generation (CDG) for computational notebooks. In contrast to the previous CDG tasks which focus on generating documentation for single code snippets, in a computational notebook, one documentation in a markdown cell often corresponds to multiple code cells, and these code cells have an inherent structure. We proposed a new model (*HACovGNN*) that uses a hierarchical attention mechanism to consider the relevant code cells and the relevant code tokens information when generating the documentation. Tested on a new corpus constructed from well-documented Kaggle notebooks, we show that our model outperforms other baseline models.

1 Introduction

In recent years, computational notebooks such as Jupyter have become popular programming platforms for data scientists and machine learning researchers to document ideas, write code, and visualize results, all in a single document (Wang et al., 2021a). Documentation in a notebook provides a rich medium for users to record not only what the code does, but also why they code it. This richness of content is one distinctive nature of code documentation in a notebook versus in traditional software source code.

Code documentation is found critical for data scientists to share or reuse code (Zhang et al., 2020; Chattopadhyay et al., 2020). However, research has shown that many data scientists still neglect to write appropriate documentation for their code in notebooks, as they feel writing documentation will slow down their coding process. Rule et al. (2018) report that among one million computational notebooks on Github, 25% of them have no comment.

* Equal contributions from the first authors: x827liu@uwaterloo.ca, dakuo.wang@ibm.com.

Documentation	
ground truth	Implementing Neural Network
our Model	Implementing Neural Network
code2seq	The following function of the model
graph2seq	After perturbations
T5-small	Model

Code Cells	
<pre>import keras from keras.utils import plot_model from keras.models import ↳ Model, Sequential, load_model ...</pre>	
<pre>def nn_model(X, y, optimizer, kernels): input_shape = X.shape[1] if len(np.unique(y)) == 2): op_neurons = 1 op_activation = 'sigmoid' loss = 'binary_crossentropy' else: op_neurons = len(np.unique(y)) op_activation = 'softmax' loss = 'categorical_crossentropy' classifier = Sequential() ... classifier.summary() return classifier</pre>	
<pre>model = ↳ nn_model(X_train, y_train, 'adam', 'he_uniform') history = model.fit(X_train, y_train, ↳ batch_size = 64, epochs = 1000, ↳ validation_data=(X_test, y_test))</pre>	
<pre>pd.DataFrame(abs(train.corr())['Survived']) .sort_values(ascending = False)</pre>	

Table 1: An example of multiple code cells after one documentation block

As a first step towards building an automated documentation generation system for notebooks, in this paper we focus on the code documentation generation (CDG) task for Jupyter notebooks. Since there is no publicly available CDG dataset for notebooks, we construct a new dataset (**notebookCDG**) which contains around 28k processed code-documentation pairs extracted from 2,476 highly-ranked notebooks from Kaggle competitions (details in Section 3)

Part of work was done when Xuye, April, and Lingfei were at IBM.

A few previous literature have explored techniques to generate documentation for software code snippet one at a time (LeClair et al., 2020; Haque et al., 2020, 2021; Xu et al., 2018). However, in computational notebooks, one documentation (in a *markdown cell*) can cover more-than-one code cells after it. For instance, the ground truth text in Table 1 is a single documentation covering four code cells. Existing work on CDG (Kery and Myers, 2017; Iyer et al., 2016; Hu et al., 2018; Alon et al., 2019; LeClair et al., 2020) does not consider such structure information since they only focus on documentation generation for single code snippet (i.e., one function, or one expression).

To account for the above mentioned properties of documentation in computational notebooks, in this paper, we propose a graph-augmented encoder-decoder model to generate documentation for notebooks (Section 4). In particular, our model consists of three parts: a code sequence encoder, an auxiliary documentation text encoder based on the already predicted documentation tokens, and a Hierarchical Attention-based Convolutional Graph Neural Network (HACnvGNN) component.

The first two sequence encoders encode the semantic information in code and documentation text, respectively. The graph encoder encodes the contextual abstract syntactic trees (i.e., AST extracted from the code sequence). In order to capture the relations between code sequences and the corresponding text documentations, we further employ a hierarchical attention mechanism consisting of a low-level attention module and a high-level attention module. The former attends to the token in a code sequence and the latter attends to the corresponding code cells in the AST tree.

Experiments show that our model achieves better performance on the *notebookCDG* dataset compared to baseline models on ROUGE scores, and in a multi-dimensional human evaluation study.

Based on this result, we integrated our approach into a user-facing downstream application (Wang et al., 2021c) to further explore the Human-AI collaboration opportunity in the code documentation scenario. In the follow-up user study (reported separately (Wang et al., 2021b)), users found that the automatically generated documentation reminded them to document code they would have ignored, and improved their satisfaction with their computational notebooks.

In summary, the main contributions of our work

are: (1) a large-scale high quality dataset for the CDG task in the computational notebook context; (2) a graph-based neural network architecture with hierarchical attention for the notebook CDG task which considers the structure information between multiple code cells and the relations between code tokens and text tokens; and (3) human evaluations to validate our model for real world application. The experiment code and data are shared¹.

2 Related Work

In order to automate the machine learning and AI workflow, researchers have applied automation techniques on various code-related tasks (Wang et al., 2020), including code summarization (Iyer et al., 2016; LeClair et al., 2020; Haque et al., 2020, 2021), source code generation from natural language (Agashe et al., 2019), and source code transformation (Roziere et al., 2020).

In this work, we focus on the code documentation generation (CDG) task. Our work is closely related to code summarization. Most existing datasets for code summarization contain one summary per one code snippet. For instance, CodeSearchNet (Husain et al., 2019) contains two million function-documentation pairs across six programming languages (e.g., java, php, python). In contrast, our new dataset (*notebookCDG*) is designed for computational notebooks. The difference from previous CDG datasets is that in our dataset, a documentation text can correspond to several code snippets.

Previous work on code summarization focuses on summary generation for a single standalone code snippet. Iyer et al. (2016) collected Stack Overflow question titles as code summaries and paired them with top-rated code snippets. They then used an attention seq2seq model to generate a summary for each code snippet. Several studies explored the abstract syntactic tree (AST) information of source code to better capture the relation between different elements (Hu et al., 2018; Alon et al., 2019). Recently, Xu et al. (2018) and Chen et al. (2020) have proposed a general graph to sequence model to learn node embeddings and then reassemble them into the graph embeddings.

Unlike the aforementioned works that only focus on summary generation for a single standalone code snippet, in our new CDG task for computational notebooks, multiple adjacent code cells can

¹<https://github.com/dakuo/HACnvGNN>

	Overall	Train	Dev	Test
Notebooks number	2,476	2,426	1,390	1,394
Code-documentation pairs	28,625	22,851	2,856	2,856
Code vocabulary size	20,522			
Code AST vocabulary size	67,211			
Documentation vocabulary size	13,053			
Avg. # token in documentation	9.15	9.13	9.37	9.18
Max. # token in documentation	202	202	130	104
Std. # token in documentation	8.40	8.44	8.27	8.25
Avg. # token in code cell(s)	65.38	65.50	65.41	64.39
Max. # token in code cell(s)	400	400	400	395
Std. # token in code cell(s)	68.93	69.16	68.23	67.71
Avg. # token in code AST	181.08	181.47	180.77	178.24
Max. # token in code AST	1732	1548	1732	1167
Std. # token in code AST	192.19	193.00	190.43	187.40

Table 2: *notebookCDG* dataset statistics. The overall code-to-markdown ratio is 2.2195, which suggests one markdown corresponds to more than one code cells.

correspond to one documentation and these code cells may have a hierarchical structure, and use a graph to represent it (Kipf and Welling, 2016). We thus propose Hierarchical Attention-based Convolutional Graph Neural Network (HACovGNN) to handle the hierarchical AST graph structure of multiple code cells.

3 *notebookCDG* Dataset

CDG for notebooks is a relatively new task. To our best knowledge, we could not find an appropriate dataset for this task. Thus, we decided to construct a new dataset and share it with the community. Publicly shared notebooks on Github are often ill-documented (Rule et al., 2018), thus are not suitable for constructing the training dataset for CDG task. A recent work (Wang et al., 2021a) manually analyzed 80 publicly available notebooks on two Kaggle challenges (i.e. out of 12,000 notebooks submitted to Titanic and HousePrice). Kaggle allows community members to vote up and down on those notebooks, and Wang et al. (2021a)’s findings show that the highly-voted notebooks are of good quality and quantity in code documentation. Inspired by their work, we decided to utilize the top-voted and well-documented Kaggle notebooks to construct the *notebookCDG* dataset².

We collected the top 10% highly-voted notebooks from the top 20 popular competitions on

Kaggle (e.g. Titanic). We checked the data policy of each of the 20 competitions, none of them has copyright issues. We also contacted the Kaggle administrators to make sure our data collection complies with the platform’s policy. In total, we collected 3,944 notebooks as raw data.

3.1 Data Preprocessing

We performed various preprocessing steps to prepare the dataset, following LeClair and McMillan (2019). For example, we removed notebooks in non-English language. One major difference between our dataset and previous datasets is that in previous datasets, each documentation unit is corresponding to one code snippet, whereas in our dataset, one documentation unit may correspond to upto four code snippets (code cells). We first located the markdown cells that have code cells beneath them. According to Wang et al. (2021a), there are nine categories of documentations in a notebook, some are related to code, some are not related to code. For those types closely related to code (Process and Headline), which take up 80% of the cases, we can directly use the markdown cell as documentation. For some other types, such as the Result type, which interprets the rendered result table or plot thus are often long and irrelevant to the code, we used a list of keywords (e.g., shows) to filter out the key sentences from the markdown cell as the documentation. Another special types of documentation are Reason and Education, which also uses long word sequence

²We share the *notebookCDG* dataset with processed 28k code-document pairs at <https://ibm.biz/Bdfpk6>

to explain why the author did something. In these cases, based on our observation, we used the first sentence as the documentation, as the first sentence is often related to the code cells.

Our analysis shows that for one markdown cell, there could have maximum four code cells following it. We construct our dataset to have a structure with one documentation unit and four code sequence units, and fill with empty sequence if there is less than four code sequences. As part of the data preparation, we also parse each of code sequence to an AST graph structure through a Python AST library³. While doing so, we removed all the non-Python notebook magic (e.g. `%matplotlib`).

3.2 Dataset Core Statistics

After data preprocessing, the final dataset contains 2,476 notebooks out of the 3,944 notebooks from the raw data. It has 28,625 code-documentation pairs. The overall code-to-markdown ratio is 2.2195, which suggests one markdown corresponds to more than one code cells. Then, the code-documentation pairs are randomly split into train, dev, and test subsets, following a 8:1:1 ratio (Table 2).

Our *notebookCDG* dataset has a vocabulary size of 13,053 for the documentation sequence, a vocabulary size of 20,522 for the code sequence, and 67,211 for the parsed code AST node. On average, each pair of code-documentation has 65.38 code tokens, and 9.15 documentation tokens. When code is translated to AST structure, on average it has 181.08 tokens.

4 Approach

Our model is built upon the standard encoder-decoder structure. To handle multiple code cells in computational notebooks, we propose a hierarchical attention mechanism based on convolutional graph neural network (HACnvGNN) for capturing the relevant code cells during the decoding stage.

The system architecture is illustrated in Figure 1. Below, we describe each module in detail.

4.1 Model Input

As mentioned in Section 3, we found that there are up to four adjacent code cells under a markdown cell, thus we constructed the *notebookCDG* dataset to have one documentation mapping to four code

cells, and used empty code cell as padding. Therefore, when generating the abstract syntactic tree (AST) for a code cell, we can assemble up to four AST trees into a higher level graph structure.

In summary, each training data point has four parts: the tokenized code sequence, the tokenized documentation sequence, the nodes of the AST graph generated from the code sequence, and the edges (topology) of the AST graph generated from the code sequence. We denoted code sequence input as $S = \{s_1, s_2, \dots, s_n\} \in \mathcal{S}$ where s_i is sequence consisting of a sequence of code token embeddings $s_i = \{w_1, w_2, \dots, w_k\} \in \mathcal{W}$ in which \mathcal{W} is the token embedding space and k is the length of s_i . Next we construct the AST graph input $A = (V, E)$ where V are the nodes containing the original code, E are the edges which denote whether two nodes are connected or not in the AST graph.

4.2 Embeddings

We use three embedding layers to generate embeddings for the tokenized code sequence, the nodes in an AST graph, and the documentation decoder, respectively.

4.3 Encoder

We use one encoder to encode the source code sequence, and additional four encoders to encode up to four code cells' AST graphs. In addition, we have a high-level GRU encoder layer for all the four AST graphs to generate one high-level output. More specifically, the encoder for the tokenized code sequence is a GRU with an output length of 256. An AST graph encoder is a collection of Convolutional Graph Neural Networks layers followed by a GRU layer of output length 256. We use four AST graph encoders for up to four code cells. Following LeClair et al. (2020), the number of hops in our GNN layers is set to 2.

4.4 HACnvGNN

The key design of our HACnvGNN model is the hierarchical attention. When handling AST graphs input, instead of blending these 4 code cells as a whole sequence, we propose to use a hierarchical attention mechanism (low-level attention and high-level attention in HACnvGNN in Figure 1) on these AST graphs to better preserve the graph structure.

Firstly, the four code cells' AST graph can be represented as $G = \{G_1, G_2, G_3, G_4\}$. We denote

³<https://docs.python.org/library/ast.html>

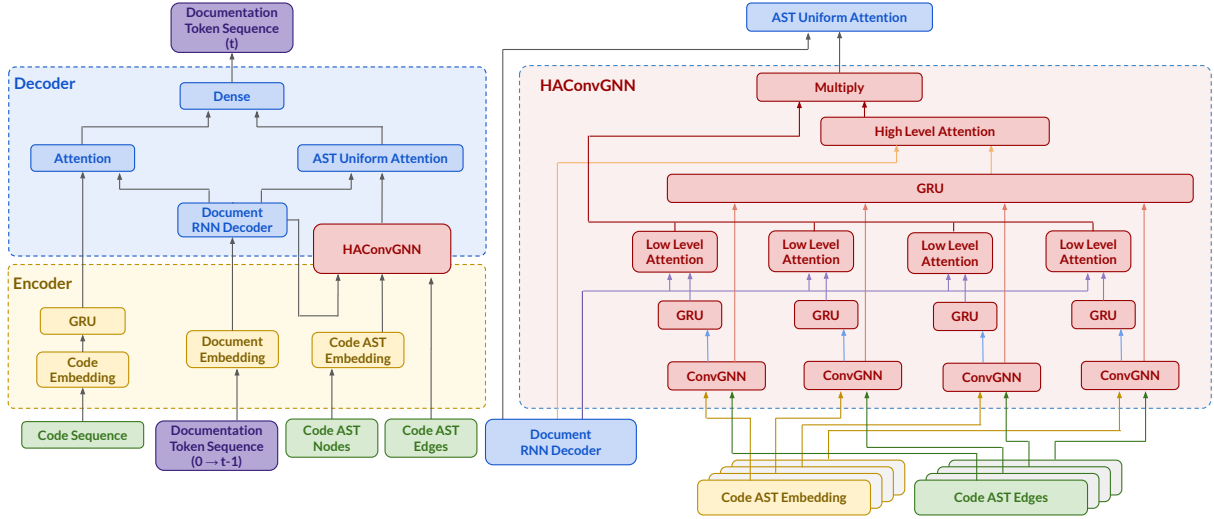


Figure 1: HACovGNN model architecture

the decoder output (i.e., the predicted documentation tokens up till $t - 1$) as $D \in \mathbb{R}^{n \times d}$ where d is the dimension. We further denote each code cell’s AST graph as $G_i \in \mathbb{R}^{m \times d}$ where m is the number of nodes. After using a high-level encoder to encode the AST graph input, we execute a graph-level attention to get high-level attention score:

$$\alpha(G_i, D) = DG_i^\top / \sqrt{d} \quad (1)$$

Then we apply softmax on α , given by:

$$b^i = \frac{\exp(\alpha(G_i, D))}{\sum_j \exp(\alpha(G_j, D))} \quad (2)$$

In this way, we get the results denoted as $\alpha = \{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}$. This is our high-level attention weights indicating the relations between each code cell and the already predicted documentation sequence D .

Secondly, we apply an attention mechanism on each code cell to find the relations between nodes in a code cell’s AST and the predicted documentation sequence D . For each code cell’s AST tree $G = \{G_1, G_2, G_3, G_4\}$, we apply the same operation as in EQ.1 and EQ.2. As a result, for each code cell G_i , we are able to get a new low-level attention weight β_i . For all code cells, we can denote these attention scores as $\beta = \{\beta_1, \beta_2, \dots, \beta_m\}$.

Eventually, we fuse these attention weights (α and β) with code cells:

$$O = \sum_{i=1}^4 \alpha_i \sum_{j=1}^m \beta_{i,j} G_{i,j} \quad (3)$$

Now we get the AST matrices from HACovGNN. It is then concatenated with code matrices

into a single context matrix. Note that code matrices are based on the code sequence input with a separate uniform attention (see the left “Code Sequence” in Figure 1). Next, we apply a linear projection to project the merged context matrix into a 256 dimension space. This is an effective way to avoid overfitting during the training process. Finally, we flatten the new context matrix and apply another linear layer to project it into an output. The output layer size is the vocabulary size. By applying the Argmax function to the output layer, we can obtain the predicted next token (i.e., documentation token at time step t) in the output sequence.

5 Experimental Setup

5.1 Implementation Details

We split our dataset into training, development, and test datasets at a 8:1:1 ratio. We use the Adam optimizer (Kingma and Ba, 2014) with a batch size of 20. The learning rate is 0.001 and the code sequence embedding size is 100. In the encoder, we use GRU (Cho et al., 2014) with the hidden size of 256. The hop size of our GNN is 2. The dropout rate of our attention layer is 0.5.

5.2 Baselines

We compare our model against two baseline models for single code snippets.

code2seq. Alon et al. (2019) proposed a *code2seq* model to generate a summary for a C# function. The model creates a vector representation for each AST path separately through an encoder. During decoding, the model uses attention to select

Models	ROUGE-1			ROUGE-2			ROUGE-L		
	P	R	F1	P	R	F1	P	R	F1
Baselines									
code2seq	11.45	8.46	8.23	1.67	1.11	1.11	13.13	10.28	10.24
graph2seq	13.21	9.87	9.51	2.86	1.99	2.03	14.46	11.40	11.18
Our Model & Ablation Study									
HACConvGNN (Our Model)	22.87	16.92	16.58	6.72	4.86	4.97	24.03	18.60	18.54
HACConvGNN with low-level attention without high-level attention with uniform attention	20.66	15.65	14.91	4.74	3.92	3.80	21.84	17.27	16.81
HACConvGNN with low-level attention without high-level attention without uniform attention	19.57	14.59	14.23	4.87	3.56	3.63	20.83	16.24	16.12
HACConvGNN without low-level attention without high-level attention with uniform attention	11.39	7.73	7.82	1.58	1.06	1.08	13.13	9.47	9.82

Table 3: ROUGE scores for the baselines, our model, and the ablation models. Results show that our model has higher scores for all three metrics, demonstrating a robust advantage over the code2seq and graph2seq models.

the relevant paths. We re-implement this model and apply it on our dataset.

graph2seq. Xu et al. (2018) proposed a graph-to-sequence learning framework that maps an input graph to a sequence of vectors and uses an attention-based LSTM method to decode the target sequence from these vectors. The authors tested the model on natural language question generation from the SQL query task. We re-implement this model using all recommended parameters from the original paper.

5.3 Experimental Details

The training time of code2seq model is around 2.5 hours per epoch; the training time of graph2seq is around 2.75 hours per epoch; the training time of T5-small is around 3.25 hours per epoch; the training time of our HACConvGNN model is around 2.65 hours per epoch.

The training environment of code2seq, graph2seq, and HACConvGNN is three GPUs using Parallelism. The training environment of T5-small is two GPUs.

code2seq and graph2seq are implemented in Keras framework⁴. T5-small model is implemented based on Huggingface repo⁵.

⁴<https://github.com/Attn-to-FC/Attn-to-FC>

⁵<https://github.com/huggingface/>

6 Automated Evaluation

We use ROUGE scores (Lin, 2004) to evaluate our model’s performance with regard to the ground-truth documentation content. We report ROUGE-1, ROUGE-2, and ROUGE-LCS (longest common sub-sequence). As shown in Table 3, our HACConvGNN model outperforms the other two baselines in all ROUGE metrics.

Ablation study. In order to better understand the impact of the attention components in our model, we also perform an ablation study (Table 3). Our ablation study evaluates how low-level attention, high-level attention, and AST uniform attention contribute to the model. More concretely, we generate ablation models as the following:

(1) without high-level attention in the hierarchical attention: we remove high level attention component in Figure 1 in our HACConvGNN structure. That means we do not compute attention weights for separated code cells.

(2) without AST uniform attention: we do not apply uniform attention mechanism (i.e., the attention component above *HACConvGNN* in Figure 1 for our HACConvGNN output with the decoder. (3) without low-level or high-level attentions: we

[transformers](https://github.com/huggingface/transformers)




Figure 2: Attention visualization for the data point illustrated in Table 1. Each row represents a code cell, and each column is a code token. In this example, it shows the second and third token in the second code cell (“ nn_model”, “X”) contribute the most to the predicted documentation in Table 1.

remove separated low-level attention components in Figure 1) in our HAConvGNN structure. Note that when we remove these separated attentions, we also remove the high-level attention (thus the entire hierarchical attention structure). We treat multiple code cells as a standalone code snippet in this situation and process graph data with the original GNN layer (see the last row in Table 3).

In general, we found that the hierarchical structure in our HAConvGNN is proven to enhance our final performance. It is worth noting that the separated attention mechanism is essential in our model. Remember that we use the attention mechanism for our four code cells separately. Treating them as a single big code snippet leads to a considerable performance drop (see the last row in Table 3). This demonstrates that the hierarchical structure in our model can better handle the code documentation generation task for multiple code cells.

Attention Visualization. Our high-level attention mechanism can indicate the most relevant code cell when generating the documentation for several code cells. Figure 2 illustrates the attention heatmap for the code example in Table 1. Each row represents a code cell, and each column corresponds to a code token. It seems that the model pays more attention to the second code cell (especially the first few tokens) when generating the documentation “*Implementing Neural Network*”.

7 Human Evaluation

We also conduct a human evaluation to further evaluate our model against the two baselines and the ground truth.

Participants. Our human evaluation task involves reading code snippets and rating the generated documentation of the codes. We recruited participants with data science and machine learning backgrounds ($N = 15$).

Task. We randomly selected 30 pairs of documentation and code(s) from our dataset. Note that each pair has only one summary, but may have multiple code snippets. Each participant is randomly assigned 10 trials, and the order of these 10 trials is also randomized. Each pair is evaluated by 5 individuals. In each trial, a participant reads 4 candidate documentation for the same code snippet(s): three generated by the three models, and the other one is the groundtruth. Participants do not know which documentation text is from which model. The participant is asked to rate the 4 documentation texts along three dimensions using a five-point Likert-scale from -2 to 2.

- *Correctness*: The generated documentation matches with the code content.
- *Informativeness*: The generated documentation covers more information units.
- *Readability*: The generated documentation is in readable English grammar and words.

Evaluation Results. We conducted *pairwise t-tests* to compare each model’s performance. The result (Table 4) shows that for the *Correctness* dimension, our model (avg=0.21) is significantly better than the other two baselines (avg=-0.59 for code2seq, avg=-0.30 for graph2seq, both $p < .01$). Our model is also the only model that has a positive rating. For the *Informativeness* dimension, groundtruth also has the best rating. Our model (avg=0.17) comes in second and outperforms code2seq (avg=-0.72, $p < .01$) and graph2seq (avg=-0.21, $p < .01$).

For the *Readability* dimension, in which we consider whether generated documentation is a valid English sentence or not, groundtruth outperforms all ML models again, but our model (avg=0.67) also significantly outperforms baseline models code2seq (avg=0.03 $p < .01$) and graph2seq (avg=0.32 $p < .01$). Our model can generate more readable documentation than baselines.



Figure 3: Average rated scores given by human evaluators to each method across three dimensions.

Model	Correctness	Informativeness	Readability
Groundtruth	$\bar{x} = 1.09, \sigma=0.95$	$\bar{x} = 0.85, \sigma=0.97$	$\bar{x} = 1.03, \sigma=1.01$
Our model	$\bar{x} = 0.21, \sigma=1.33$	$\bar{x} = 0.17, \sigma=1.18$	$\bar{x} = 0.67, \sigma=1.20$
Code2seq	$\bar{x} = -0.59, \sigma=1.29$	$\bar{x} = -0.72, \sigma=1.17$	$\bar{x} = 0.03, \sigma=1.35$
Graph2seq	$\bar{x} = -0.30, \sigma=1.40$	$\bar{x} = -0.21, \sigma=1.25$	$\bar{x} = 0.32, \sigma=1.35$

Table 4: Human Evaluation Result

All the results suggest that our model has above-zero ratings, which suggests it reaches an acceptable user satisfaction along all three dimensions.

8 Comparison With Transformers

We also carried out an additional experiment to compare our model with T5 (Raffel et al., 2020), which is a state-of-the-art transformer encoder-decoder model. In order to fairly compare our model against T5, we do not use any pre-trained embeddings for the T5 model. Also, T5 input has limitation for the input token length thus we did not feed AST hierarchy into it. More specifically, we initialize a T5-small model⁶ with random weights and train this model using our training data. Our code adapts the transformer models from HuggingFace (Wolf et al., 2020). We use the dev dataset to choose the hyperparameters and evaluate the trained model on our test dataset. The ROUGE F1 scores for the trained T5-small model are as follows: ROUGE-1 = 17.55, ROUGE-2 = 4.57, ROUGE-L = 19.53.

We found that the trained T5-small model achieves slightly better results than our model in ROUGE-1 and ROUGE-L. In practice, we found that the T5-small model relies on a much more hyperparameters and tends to generate less informative content compared to other models (see the documentation generated from different models in Table 1 for an example).

But in our dataset, as reported in Table 2, the max

⁶In a pilot study, training a T5-base model (with random initialization) on our dataset leads to worse results.

AST token sequence is 1,732, which is too long as T5 input (512) or BART input (1,024). That is why T5 in Sec 8 can only take the raw code sequence as input, instead of the AST hierarchy. It is known that programming code has a tree-based hierarchy and leveraging such AST hierarchy can enhance the baseline model (e.g., (Alon et al., 2019)). Our contribution is that we provide a hierarchical attention architecture that is well suited for the programming code nature and can generalize to a much longer length of code inputs. Imagine in a scenario where we can feed a whole code repo as training input by treating each code file as a lower layer, and connecting them through function/variable referencing – our architecture can also handle that. In general, we think our model is orthogonal to the standard transformer models. One interesting future work is to integrate our hierarchical attention mechanism into the transformer-based structure instead of a GRU-based structure.

9 Downstream User Application

To demonstrate the application of the HACnvGNN model, we designed a Jupyter Notebook plugin to assist document writing in data science programming (as shown in Figure 4).

The plugin is triggered when detecting users focusing on a code cell (Figure 4.A). The plugin then reads the contents from the focused cell and its adjacent cells, and sends the content to the backend. The backend server first generates a code summarization using the HACnvGNN model (Figure 4.B). In addition, we implemented two other

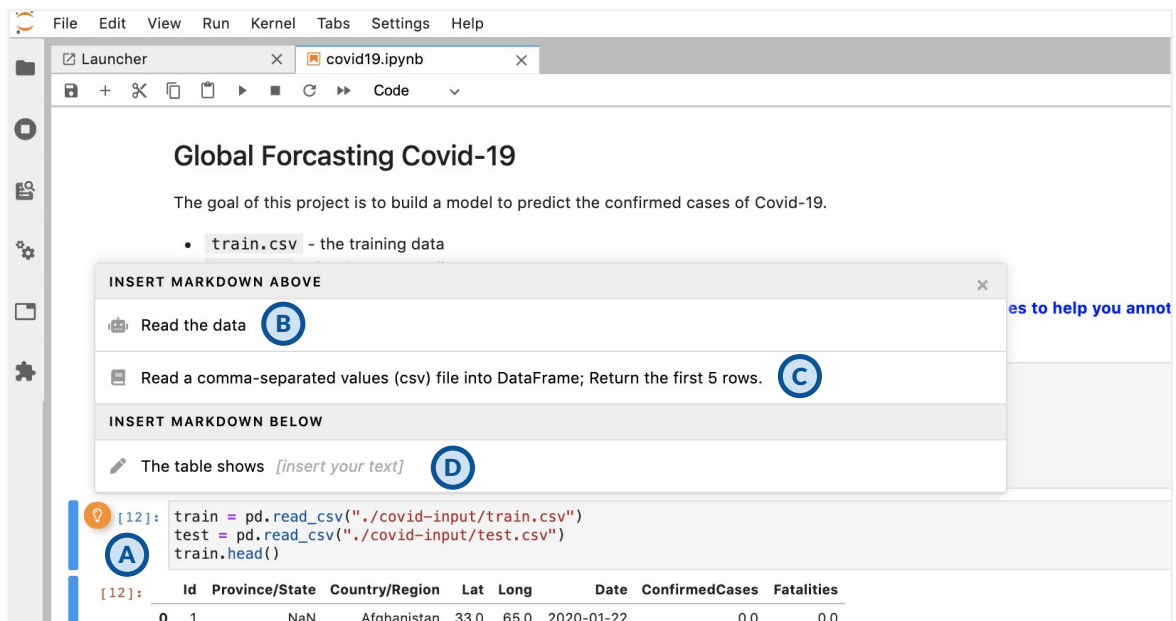


Figure 4: We implement a downstream application as a Jupyter Notebook plugin (A) to assist users documentation writing, incorporating the HAConvGNN-predicted results (B) next to an IR-based approach (C), and a user-prompt approach (D).

approaches to generate documentation that was intended for explaining a design decision or explaining a technical concept for educational purposes. We retrieved the relevant documentation from the API webpage for educational purposes (Figure 4.C) and we used prompts to nudge users to explain an output (Figure 4.D). If the user likes one of these three candidates, they can simply click on one of them, and the selected documentation candidate will be inserted into above the code cell (if it describes what and why for the code), or below it (if it interprets the result of the code).

Our plugin went through several rounds of pilot testing and iterative design. Participants found it reminds them to document code they would have ignored, reduce the time for developing documentation while they were actively exploring the data science task. The implementation details and a formal evaluation of understanding the benefits of the human-AI collaborative effort for automatic documentation generation are reported separately in (Wang et al., 2021b).

10 Conclusion and Future Work

This work targets a new application that aims to automatically generate code documentation (CDG) for a computational notebook. This project is part of our longterm research initiative of designing AI to automated the various tasks in an AI project's

lifecycle (Wang et al., 2021d). The notebookCDG context imposes unique challenges to the current code documentation generation approaches which only consider a single code snippet. We construct a dataset from Kaggle challenge notebooks, and present a novel HAConvGNN model to encode the multiple adjacent code cells as a hierarchical AST graph to enhance a sequence model architecture. Both automated evaluation and human evaluation show that our model outperforms the baseline models. We also incorporate our algorithm into a Jupyter Notebook plugin to assist code documentation creation.

11 Ethical Concern

Our task is an instance of natural language generation task, thus it may have potential risk and ethical issues similar to any other NLG tasks, such as the generated content may have offensive language. However, we believe our task and our approach has minimum risk of such ethical issues, due to two reasons: firstly, the language used in the context of machine learning code documentation is more strict to technical terms, offensive language is less likely to appear in the dictionary thus in our model; secondly, the dataset construction method is to use highly-voted notebooks from a publicly available Kaggle community, there is unlikely to have offensive languages in these highly-voted notebooks.

References

- Rajas Agashe, Srinivasan Iyer, and Luke Zettlemoyer. 2019. [JulCe: A large scale distantly supervised dataset for open domain context-based code generation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5436–5446, Hong Kong, China. Association for Computational Linguistics.
- Uri Alon, Omer Levy, and Eran Yahav. 2019. [code2seq: Generating sequences from structured representations of code](#). In *International Conference on Learning Representations*.
- Souti Chattopadhyay, Ishita Prasad, Austin Z Henley, Anita Sarma, and Titus Barik. 2020. What’s wrong with computational notebooks? pain points, needs, and design opportunities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12.
- Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2020. Reinforcement learning based graph-to-sequence model for natural question generation. In *The Eighth International Conference on Learning Representations (ICLR 2020)*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Fethi Bougares, Holger Schwenk, and Y. Bengio. 2014. [Learning phrase representations using rnn encoder-decoder for statistical machine translation](#).
- Sakib Haque, Aakash Bansal, Lingfei Wu, and Collin McMillan. 2021. Action word prediction for neural source code summarization. In *2021 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 330–341. IEEE.
- Sakib Haque, Alexander LeClair, Lingfei Wu, and Collin McMillan. 2020. Improved automatic summarization of subroutines via attention to file context. In *Proceedings of the 17th International Conference on Mining Software Repositories*, pages 300–310.
- Xing Hu, Ge Li, Xin Xia, David Lo, Shuai Lu, and Zhi Jin. 2018. Summarizing source code with transferred api knowledge. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI’18*, page 2269–2275. AAAI Press.
- Hamel Husain, Ho-Hsiang Wu, Tiferet Gazit, Miltiadis Allamanis, and Marc Brockschmidt. 2019. [Code-searchnet challenge: Evaluating the state of semantic code search](#). *CoRR*, abs/1909.09436.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2016. [Summarizing source code using a neural attention model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2073–2083, Berlin, Germany. Association for Computational Linguistics.
- Mary Beth Kery and Brad A. Myers. 2017. [Exploring exploratory programming](#). In *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, pages 25–29.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- A. LeClair and C. McMillan. 2019. Recommendations for datasets for source code summarization. In *2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Alexander LeClair, Sakib Haque, Lingfei Wu, and Collin McMillan. 2020. Improved code summarization via a graph neural network. In *2020 IEEE International Conference on Program Comprehension (ICPC)*.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#).
- Baptiste Roziere, Marie-Anne Lachaux, Lowik Chanussot, and Guillaume Lample. 2020. [Unsupervised translation of programming languages](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 20601–20611. Curran Associates, Inc.
- Adam Rule, Aurélien Tabard, and James D Hollan. 2018. Exploration and explanation in computational notebooks. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–12.
- April Yi Wang, Dakuo Wang, Jaimie Drozdal, Xuye Liu, Soya Park, Steve Oney, and Christopher Brooks. 2021a. What makes a well-documented notebook? a case study of data scientists’ documentation practices in kaggle. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–7.
- April Yi Wang, Dakuo Wang, Jaimie Drozdal, Michael Muller, Soya Park, Justin D Weisz, Xuye Liu, Lingfei Wu, and Casey Dugan. 2021b. Themisto: Towards automated documentation generation in computational notebooks. *arXiv preprint arXiv:2102.12592*.
- April Yi Wang, Dakuo Wang, Xuye Liu, Lingfei Wu, et al. 2021c. Graph-augmented code summarization in computational notebooks. *IJCAI’21 Demo*.

- Dakuo Wang, Q Vera Liao, Yunfeng Zhang, Udayan Khurana, Horst Samulowitz, Soya Park, Michael Muller, and Lisa Amini. 2021d. How much automation does a data scientist want? *arXiv preprint arXiv:2101.03970*.
- Dakuo Wang, Parikshit Ram, Daniel Karl I Weidele, Sijia Liu, Michael Muller, Justin D Weisz, Abel Valente, Arunima Chaudhary, Dustin Torres, Horst Samulowitz, et al. 2020. Autoai: Automating the end-to-end ai lifecycle with humans-in-the-loop. In *Proceedings of the 25th International Conference on Intelligent User Interfaces Companion*, pages 77–78.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Kun Xu, Lingfei Wu, Zhiguo Wang, Yansong Feng, Michael Witbrock, and Vadim Sheinin. 2018. Graph2seq: Graph to sequence learning with attention-based neural networks. *arXiv preprint arXiv:1804.00823*.
- Amy X Zhang, Michael Muller, and Dakuo Wang. 2020. How do data science workers collaborate? roles, workflows, and tools. *arXiv preprint arXiv:2001.06684*.

A Appendix: Code snippets-documentation Pair Examples

Documentation	
Ground Truth	Feature scaling
Our Model	Feature scaling
Code2seq	We can have the model
Graph2seq	The next step is a lot of the training set
T5-small	Scaling
Code Cells	
<pre>from sklearn.preprocessing import StandardScaler scaler = StandardScaler().fit(train_df) train_scale = ↳ pd.DataFrame(scaler.transform(train_df))</pre>	

Table 5: Example: Feature Scaling

Documentation	
Ground Truth	handle missing values in X test
Our Model	we can deal with missing values
Code2seq	We can have the categorical data
Graph2seq	We can also make any numeric variable in the model
T5-small	Filling the missing values in the test set
Code Cells	
<pre>cols_with_missing_val = [col for col in ↳ X_test.columns if X_test[col].isnull().any()] print(cols_with_missing_val)</pre>	
<pre>from sklearn.impute import SimpleImputer my_imputer = ↳ SimpleImputer(strategy='most_frequent') my_imputer.fit(X_train) imputed_X_test = ↳ pd.DataFrame(my_imputer.transform(X_test)) imputed_X_test.columns = X_test.columns</pre>	

Table 6: Example: Handle Missing Values

Documentation	
Ground Truth	Data Augmentation
Our Model	Data Builder
Code2seq	We can have the model
Graph2seq	LSTM
T5-small	Visualize the images
Code Cells	
<pre>import warnings from imgaug import augmenters as iaa warnings.filterwarnings("ignore") augmentation = iaa.Sequential([iaa.OneOf([iaa.Affine(rotate=0), iaa.Affine(rotate=90), iaa.Affine(rotate=180), iaa.Affine(rotate=270),]), iaa.Fliplr(0.5), iaa.Flipud(0.2), iaa.OneOf([iaa.Cutout(fill_mode="constant", ↳ cval=255), iaa.CoarseDropout((0.0, 0.05), ↳ size_percent=(0.02, 0.25)),]), iaa.OneOf([iaa.Snowflakes(flake_size=(0.2, 0.4), ↳ speed=(0.01, 0.07)), iaa.Rain(speed=(0.3, 0.5)),]), iaa.OneOf([iaa.Multiply((0.8, 1.0)), iaa.contrast.LinearContrast((0.9, ↳ 1.1)),]), iaa.OneOf([iaa.GaussianBlur(sigma=(0.0, 0.1)), iaa.Sharpen(alpha=(0.0, 0.1)),])]), random_order=True) def get_ax(rows=1, cols=1, size=7): _, ax = plt.subplots(rows, cols, ↳ figsize=(size*cols, size*rows)) return ax limit = 4 ax = get_ax(rows=2, cols=limit//2) for i in range(limit): image, image_meta, class_ids, \ bbox, mask = modellib.load_image_gt(dataset_train, config, image_id, ↳ use_mini_mask=False, augment=False, augmentation=augmentation) visualize.display_instances(image, bbox, mask, ↳ class_ids, dataset_train.class_names, ↳ ax=ax[i//2, i % 2, show_mask=False, ↳ show_bbox=False)</pre>	

Table 7: Example: Data Augmentation

Documentation	
Ground Truth	Plot the model s performance
Our Model	Plot the model s performance
Code2seq	We can have the model
Graph2seq	The next step is a lot of the training and test set
T5-small	Plot model performance
Code Cells	
<pre>plt.plot(history_size_val_1) plt.plot(history_size_val_2) plt.plot(history_size_val_3) plt.plot(history_size_val_4) plt.plot(history_size_val_5) plt.plot(history_size_val_6) plt.title('Model accuracy for different Conv ↪ sizes') plt.ylabel('Accuracy') plt.xlabel('Epoch') plt.ylim(0.98,1) plt.xlim(0,n_epochs) plt.legend(['8-16', '16-32', '32-32', '24-48', ↪ '32-64', '48-96', '64,128'], loc='upper left') plt.savefig('convolution_size.png') plt.show()</pre>	

Table 8: Example: Plot Model Performance

Documentation	
Ground Truth	Count Monthly Mean
Our Model	Monthly Count
Code2seq	We can have a look at the training set
Graph2seq	Feature Engineering
T5-small	Creating a new column
Code Cells	
<pre>for year in year_list: for month in range(num_months_per_year): start_date = datetime.datetime(year, ↪ month+1, 1, 0, 0, 0) end_date = datetime.datetime(year, month+1, ↪ 19, 23, 0, 0) count_mean = ↪ train_data[start_date:end_date] ['count'].mean() train_data.loc[start_date:end_date, ↪ 'count_mean'] = count_mean start_date = datetime.datetime(year, ↪ month+1, 20, 0, 0, 0) last_day_of_month = ↪ calendar.monthrange(year,month+1)[1] end_date = datetime.datetime(year, month+1, ↪ last_day_of_month, 23, 0, 0) test_data.loc[start_date:end_date, ↪ 'count_mean'] = count_mean test_data.head()</pre>	

Table 9: Example: Count Monthly Mean