

Can We Trust the Performance Evaluation of Uncertainty Estimation Methods in Text Summarization?

Jianfeng He^{†*}, Runing Yang[†], Linlin Yu[‡], Changbin Li[‡],
Ruoxi Jia[†], Feng Chen[‡], Ming Jin[†], Chang-Tien Lu^{††}

[†]Department of Computer Science, Virginia Tech, Falls Church, VA, USA

[‡]Department of Computer Science, The University of Texas at Dallas, Richardson, TX, USA

[†]{jianfenghe, yruning, ruoxijia, jinming, ctlu}@vt.edu,

[‡]{Linlin.Yu, Changbin.Li, Feng.Chen}@utdallas.edu

Abstract

Text summarization, a key natural language generation (NLG) task, is vital in various domains. However, the high cost of inaccurate summaries in risk-critical applications, particularly those involving human-in-the-loop decision-making, raises concerns about the reliability of uncertainty estimation on text summarization (UE-TS) evaluation methods. This concern stems from the dependency of uncertainty model metrics on diverse and potentially conflicting NLG metrics. To address this issue, we introduce a comprehensive UE-TS benchmark incorporating 31 NLG metrics across four dimensions. The benchmark evaluates the uncertainty estimation capabilities of two large language models and one pre-trained language model on three datasets, with human-annotation analysis incorporated where applicable. We also assess the performance of 14 common uncertainty estimation methods within this benchmark. Our findings emphasize the importance of considering multiple uncorrelated NLG metrics and diverse uncertainty estimation methods to ensure reliable and efficient evaluation of UE-TS techniques. Our code and data are available [here](#).

1 Introduction

Text summarization (Tas and Kiyani, 2007) is a representative NLG task that generates summaries for given texts. This study researches abstractive summarization (Nallapati et al., 2016), which is more flexible than extractive summarization (Gupta and Lehal, 2010). In many scenarios (e.g., finance and health), there are serious consequences if relying on false predicted summaries. For instance, an inaccurate financial report summary could lead to incorrect financial decisions, resulting in financial losses (Gómez-Guillamón, 2003). Consequently,

^{*}This work was done before joining Amazon.

[†]Ming Jin and Chang-Tien Lu are co-corresponding authors.

the task of UE-TS, which measures the likelihood that a generated summary is low-quality, has garnered significant interest (Gidiotis and Tsoumakas, 2021b; Fadeeva et al., 2023; He et al., 2024). However, an overlooked question persists regarding the *validity* (the ability to accurately measure what it intends) and *robustness* (consistency across various scenarios and datasets) of the UE-TS evaluation framework.

This concern arises due to two primary factors. First, common UE-TS evaluation metrics, such as forced-choice evaluation (He et al., 2024) and Prediction Rejection Ratio (PRR) (Malinin et al., 2017a; Fadeeva et al., 2023), rely on the alignment between two types of sample rankings: those based on uncertainty scores from an estimation method (e.g., entropy of generation semantics (Kuhn et al., 2023)) and those derived from a specific NLG metric (e.g., ROUGE (Lin, 2004)). While higher alignment suggests better uncertainty estimation, the reliance on a single NLG metric may not fully capture the nuances of summary quality.

Second, the inherent label diversity in NLG tasks, particularly text summarization, necessitates the use of various NLG metrics. For instance, in text summarization, summaries like “The cat chased the mouse” and “The mouse was pursued by the kitty” convey identical messages using different expressions. Due to the multiple valid responses possible in such tasks (Vijayakumar et al., 2016), it is unreasonable to expect NLG predictions to perfectly match the labels, unlike in NLU tasks (Hu and Khan, 2021; He et al., 2020). Consequently, each NLG metric emphasizes different aspects of the diverse labels. For instance, ROUGE (Lin, 2004) and BLEU (Papineni et al., 2002) assess the similarity between the prediction and label based on n-gram overlap, while SummaC (Laban et al., 2022) examines the level of hallucination between the generation and reference text.

Collectively, previous evaluations of UE-TS

models (Gidiotis and Tsoumakas, 2021b; Fadeeva et al., 2023; He et al., 2024) typically relied on one or two uncertainty estimation metrics, each dependent on a single type of NLG metric. This limited evaluation raises a critical question:

How reliable and valid are current methods for evaluating uncertainty estimation in text summarization?

Answering this question is essential for developing reliable UE-TS methods that can mitigate the risks associated with inaccurate summaries in critical applications, such as financial decision-making or healthcare. Moreover, our research aims to explore the broader relationship between uncertainty estimation and NLG metric scores, the impact of NLG metric diversity on experimental design and analysis, and the effective selection of diverse uncertainty estimation methods. This comprehensive exploration will inform and enhance future UE-TS research. To address these questions, we introduce a UE-TS benchmark incorporating a diverse set of text generation metrics. Our contributions are summarized below.

- To the best of our knowledge, we are the first to highlight a critical issue regarding the reliability of performance evaluation for UE-TS methods. We propose addressing this concern by evaluating UE-TS methods using a diverse set of NLG metrics. These NLG metrics span four dimensions crucial for NLG evaluation, including coherence, consistency, fluency, and relevance, following Zhong et al. (2022).
- We present a UE-TS benchmark for evaluating uncertainty estimation through various NLG metric perspectives, marking the first endeavor to our knowledge. This benchmark assesses two Large Language Models (LLMs) and one Pre-trained Language Model (PLM) across three datasets focusing on UE-TS. Within this benchmark, we incorporate 31 NLG metrics and fourteen uncertainty estimation methods. To facilitate future research, the intermediate results encompass each sample’s NLG metric and uncertainty scores. An intermediate result for an original text includes the generated summaries obtained from respective summarization models, the uncertainty scores for the sample via different uncertainty estimation methods, and the NLG

metric scores via different NLG metrics for the sample. The statistics reported in our tables and figures can be reproduced with these intermediate results.

- We also conduct human experimental analysis via a publicly available human-annotation dataset. It finds that the NLG metrics do not always correlate closely with the human annotation. However, the NLG metrics can still be used to find the best uncertainty estimation method, which saves the cost and shows some similar uncertainty estimation method rank as that ranked via human annotation.
- We have uncovered intriguing findings outlined in Sec. 6. Because text summarization is a representative NLG task, our findings in text summarization tasks could also motivate design and analysis on uncertainty estimation in other NLG tasks.

2 Related Work

Text summarization. There are two types of text summarization: extractive text summarization (Wong et al., 2008; Liu, 2019), which extracts the original sentences from the text for summarization, and abstractive text summarization (Liu and Lapata, 2019; Nallapati et al., 2016), which directly generates summaries from the text. Due to the flexibility of abstractive text summarization, we focus on abstractive text summarization. Recent models for abstractive text summarization are typically divided into two categories: PLMs (e.g., BART (Lewis et al., 2019)) and LLMs (e.g., Llama 2 (Touvron et al., 2023)). To conduct comprehensive research on our question, we test abstractive summarization models from both PLM and LLMs.

Uncertainty estimation on text summarization. The UE-TS methods primarily fall into four categories (Fadeeva et al., 2023): information-based, density-based, ensemble-based, and prompt-based methods. Information-based methods use middle output (e.g., token probability) to obtain uncertainty scores (Beigi et al., 2024). For example, Tsvigun et al. (2023) calculate uncertainty scores based on token logits in two ways: mean and Monte Carlo Dropout (Wang et al., 2019). Similarly, Simpson et al. (2020); Zhang et al. (2022) also calculate uncertainty based on token logits. Density-based methods leverage latent representations of instances, which are further used to construct a probability density. For example, Ren et al. (2022) detect

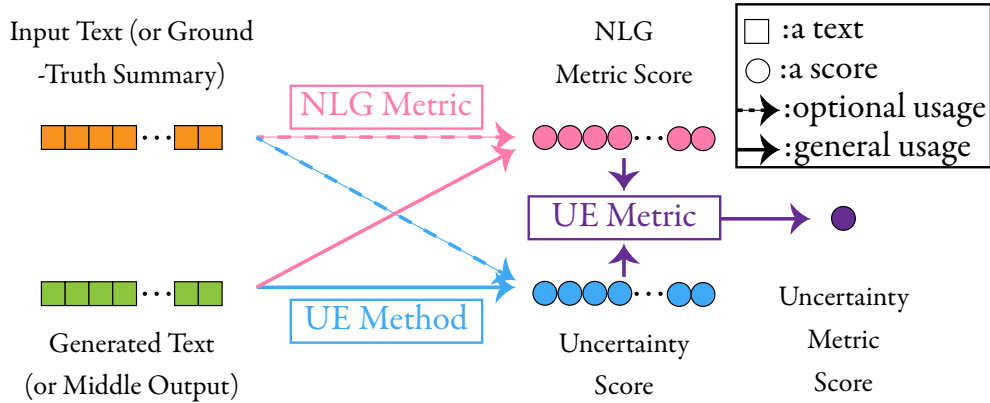


Figure 1: Diagram of the relationship between the Uncertainty Estimation (UE) metric, NLG metrics, and UE methods in the evaluation process. Specifically, the evaluation process for UE-TS methods involves using the generated texts (or intermediate outputs, such as token probabilities) and the optional input text (or ground-truth summary) to obtain NLG metric scores and uncertainty scores for all test samples, through an NLG metric and a UE method, respectively. Finally, the NLG metric scores and uncertainty scores for all testing samples are both inputted into a UE metric to obtain an uncertainty metric score of the UE method.

low-quality generations by a density learned on the given sample embeddings. Ensemble-based methods use ensembles to approximate Bayesian Neural Networks (BNN) (Mukhoti et al., 2023) or use variance of ensembled generation to obtain uncertainty scores. For example, dropout (Gal and Ghahramani, 2016) is used in Gidiotis and Tsoumakas (2021b) to approximate BNN. Also, Chuang et al. (2024) obtain the uncertainty score by the variance of ensembling predictions. As for prompt-based methods, they refer to the methods that ask the generation model via a prompt to obtain the uncertainty score (Kadavath et al., 2022). Some methods, like the SiCF score (He et al., 2024), integrate aspects of multiple uncertainty estimation methods, such as information and ensemble methods. To ensure comprehensive research on UE-TS methods, our benchmark includes representative methods from each category, following Fadeeva et al. (2023).

Performance evaluation of uncertainty estimation in text summarization. As for the uncertainty evaluation methods, Gidiotis and Tsoumakas (2021b,a) obtain the uncertainty score by measuring the variance among ensembled generations based on an NLG metric, BLEU (Papineni et al., 2002). BLEU solely assesses uncertainty estimation method performance based on n-gram similarity. Kolagar and Zarcone (2024) compare the annotated expression of uncertainty between human annotations and annotations made by LLM based on an NLG metric, semantic similarity via SentenceBERT (Reimers and Gurevych, 2019). Lei

et al. (2024) evaluates UE-TS models based on the polarity score between prediction and ground truth summaries. Zhao et al. (2022); Zablotskaia et al. (2023) evaluate uncertainty estimation methods solely using one NLG metric, ROUGE (Lin, 2004). Only two NLG metrics, ROUGE and position accuracy, have been used separately for evaluation in Zhang et al. (2022). Additionally, ROUGE and BERTScore (Zhang et al., 2019), measuring semantic similarity, are both considered in He et al. (2024); Fadeeva et al. (2023).

Thus, we observed that current evaluations of uncertainty estimation in text summarization rely on a limited set of NLG metric scores, which may lead to inconsistent performance rankings across other NLG metrics. Hence, our research aims to comprehensively explore the relationship between different uncertainty methods and various NLG metrics.

3 Metrics & Methods in Benchmark

3.1 Our Benchmark: A Global Overview

To answer “How does the choice of NLG metric affect the evaluation of uncertainty estimation methods in text summarization?”, our benchmark aims to explore the uncertainty estimation metric score across different NLG metrics for a given uncertainty estimation method. The relationship of these three items in the model evaluation process is shown in Figure 1. Below, we formalize the definition of the three scores we used in our work.

- **NLG metric score:** reflects the quality of a

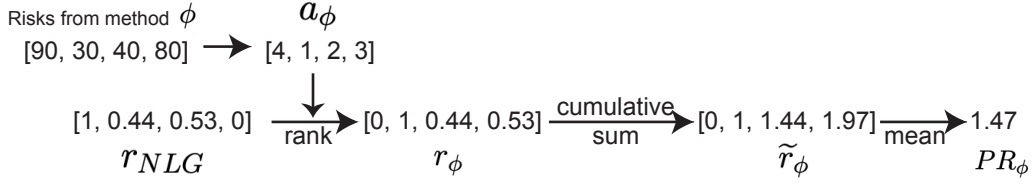


Figure 2: Diagram of the PR_ϕ calculation example with testing sample size $N = 4$. In this example, we have min-max normalized $\hat{s}_{NLG} = [0, 0.56, 0.47, 1]$, which is not drawn in the figure. Once we have obtained the sample rank a_ϕ based on a score list from method ϕ . We rerank r_{NLG} via a_ϕ to get r_ϕ . Then, we use Eq. 3 to cumulatively sum the elements and obtain \tilde{r}_ϕ . Finally, the PR_ϕ is the mean of \tilde{r}_ϕ .

generation at the sample level via an NLG metric (e.g., ROUGE).

- **Uncertainty score:** reflects the likelihood that a model generation is low-quality at the sample level. The likelihood comes from an uncertainty estimation method (e.g., BNN).
- **Uncertainty estimation metric score:** reflects the performance of an uncertainty estimation method at the method level.

Regarding the *NLG metrics*, we evaluate 31 different ones across four NLG evaluation dimensions, as proposed by Zhong et al. (2022). For the *uncertainty estimation methods* to obtain uncertainty scores, we examine fourteen common approaches outlined in Fadeeva et al. (2023). For the *uncertainty estimation metric*, we adopt a widely used one, the Prediction Rejection Ratio (PRR) (Malinin et al., 2017b; Fadeeva et al., 2023). Further details about these three components are provided below.

3.2 Uncertainty Estimation Metric: PRR

Although uncertainty estimation metrics, such as force-truth evaluation (He et al., 2024) and PRR (Malinin et al., 2017b; Fadeeva et al., 2023), rely on NLG metrics, we opt for PRR in our benchmark due to its efficiency. Unlike force-truth evaluation, which necessitates repeating NLG metric calculations ten times, PRR offers a more streamlined approach. It is formatted as follows:

$$PRR = \frac{PR_{uncertainty} - \frac{1}{\alpha} \sum_{i=1}^{\alpha} PR_{random}}{PR_{oracle} - \frac{1}{\alpha} \sum_{i=1}^{\alpha} PR_{random}} \quad (1)$$

where PR_ϕ is a scalar representing the cumulative risk, calculated based on a predicted sample rank a_ϕ from a method $\phi \in \{uncertainty, random, oracle\}$ ¹ and a list of

¹The usage of i and α for PR_{random} will be introduced in Sec. 3.3.

NLG metric scores s_{NLG} . Specifically, the $s_{NLG} \in \mathbb{R}^N$ for N testing samples, is first min-max normalized to $\hat{s}_{NLG} \in \mathbb{R}^N$. Then, the risk for N testing samples is as follows,

$$r_{NLG} = 1 - \hat{s}_{NLG} \quad (2)$$

here, we assume that our chosen NLG metric score is positively correlated with generation performance, and common NLG metrics (e.g., ROUGE) meet this assumption. Thus, the NLG metric score (an element in \hat{s}_{NLG}) is negatively correlated with the risk (an element in r_{NLG}) of inaccurate generation. A higher risk indicates a higher chance of inaccurate generation. We then obtain r_ϕ by ranking r_{NLG} based on sample rank in a_ϕ . This process is shown in Fig. 2. Further, we obtain a cumulative risk vector $\tilde{r}_\phi = [\tilde{r}_{\phi,1}, \tilde{r}_{\phi,2}, \dots, \tilde{r}_{\phi,N}]$. Its k -th element is calculated via cumulative summing the first k elements,

$$\tilde{r}_{\phi,k} = \sum_{j=1}^k r_{\phi,j} \quad (3)$$

finally, PR_ϕ is the mean of cumulative risk \tilde{r}_ϕ .

Because the risk is defined based on the normalized NLG metric score list \hat{s}_{NLG} in Eq. 2, the PR_ϕ is smaller if the predicted sample rank a_ϕ is more aligned with the sample rank based on the NLG metric score. A simple intuition is that if the predicted sample rank a_ϕ is identical to the sample rank based on r_{NLG} , PR_ϕ is minimized.

3.3 Details about Three PR Calculations

Then, we introduce PR_{oracle} , PR_{random} , and $PR_{uncertainty}$ as follows.

For PR_{oracle} , the predicted sample rank a_{oracle} is obtained via an oracle score vector $s_{oracle} = -1 \times s_{NLG}$. The PR_{oracle} is the lowest cumulative risk we can obtain because a_{oracle} is aligned with the sample rank via the NLG metric score.

For PR_{random} , we generate a random permutation of N numbers each time. $\frac{1}{\alpha} \sum_{i=1}^{\alpha} PR_{random}$ represents the average of α different PR_{random} values, each calculated using a random permutation. In our experiments, we set α to 1000.

For $PR_{uncertainty}$, the predicted sample rank is obtained from the uncertainty score, which is calculated via an uncertainty estimation method. We list our uncertainty estimation methods in Sec. 3.4.

As a result, PRR in Eq. 1 calculates the relative risk between uncertainty scores from an uncertainty method and NLG metric scores from an NLG metric. The relative risk is normalized to random risk expectation $\frac{1}{\alpha} \sum_{i=1}^{\alpha} PR_{random}$. A higher PRR means a more accurate uncertain estimation model. This is because a smaller PR_{ϕ} leads to better alignment for method ϕ and the denominator in Eq. 1 is negative. Thus, a larger PRR means smaller $PR_{uncertainty}$ and higher generation quality.

3.4 Uncertainty Estimation Methods

The uncertainty estimation methods utilized in our benchmark adhere to the framework established by Fadeeva et al. (2023), as the framework provides a publicly available uncertainty estimation tool, enhancing our results’ reproducibility.

The uncertainty estimation methods can be divided into two kinds: white-box methods and black-box methods. White-box methods refer to uncertainty estimation methods using the intermediate output, model structure, or model parameters of the text summarization model (e.g., BART). Black-box methods refer to uncertainty estimation methods that require only the final text output from a text summarization model (e.g., GPT-3.5 (OpenAI, 2023)) for the uncertainty estimation. Below, we introduce our uncertainty estimation methods one by one in Sec. A., with a summary provided in Table 1.

3.5 NLG Metrics

We chose thirty-one commonly used NLG metrics. To better understand the relationship between these NLG metrics, we categorize them into four dimensions (Zhong et al., 2022). We first introduce these four dimensions, then present all thirty-one NLG metrics with a summary in Table 2. In this table, except for the LLM-based metrics, most traditional NLG metrics are tested in Fadeeva et al. (2023).

3.5.1 Dimensions for NLG Metrics

We choose four commonly used dimensions for NLG metrics, proposed in Zhong et al. (2022) and listed below.

Relevance refers to whether the generated text contains only the important information from the input text.

Consistency is the factual alignment between the generated text and the input text.

Coherence refers to whether all the sentences in the given generated text form a coherent body.

Fluency represents the quality of individual sentences in the generated text.

Additionally, we also consider the **overall** dimension that rates the generated texts based on all the above four dimensions.

The detailed explanation for each NLG metric in Table 2 is in Sec. B.

4 Experiments

4.1 Experimental Settings

Dataset. We employ two widely used text summarization datasets in our experiments. Firstly, we utilize the AESLC dataset (Zhang and Tetreault, 2019), comprising 1,906 testing texts from the email domain. Secondly, we incorporate the XSUM dataset (Narayan et al., 2018), featuring articles collected from the British Broadcasting Corporation (BBC) and encompassing 11,334 testing samples.

Besides these two, we are also interested in human-related experiments. To meet our experimental requirements, we used the TofuEval (Tang et al., 2024) dataset. This dataset contains 100 dialogue-summary pairs with human annotations in terms of quality, evaluated from seven dimensions, which will be introduced in Sec. C.2.

Related methods and metrics. We provide detailed descriptions of the uncertainty estimation methods in Sec.3.4. Additionally, an introduction to the NLG metrics is presented in Sec.3.5. By combining the uncertainty scores obtained from the uncertainty estimation methods with the NLG metric scores, we calculate the uncertainty metric score using the PRR method introduced in Sec. 3.2.

Implementation details. To limit diverse generation, the temperature was set to 0, and a random seed of 42 was used. The experiments were conducted on a server equipped with a single A100 GPU. For the AESLC dataset, it took approximately 17 hours to run on GPT-3.5 generation and

White-box	Information-based methods	Maximum Sequence Probability (MSP) Mean Token Entropy (MTE) Monte Carlo Sequence Entropy (MCSE)
	Density-based methods	Mahalanobis Distance (MD) Robust Density Estimation (RDE)
	Ensemble-based methods	Token-level Total Uncertainty (T-TU) Token-level Reverse Mutual Information (T-RMI) Sequence-level Total Uncertainty (S-TU) Sequence-level Reverse Mutual Information RMI (S-RMI)
	Prompt-based	P(True)
Black-box	Mixture types	Number of Semantic Sets (NumSets) Eccentricity (ECC) Lexical Similarity (LexSim) Sum of Eigenvalues of the Graph Laplacian (EigV)

Table 1: A summary of the fourteen uncertainty methods that are used in our benchmark.

Relevance	ROUGE-L Spearman Kendall-Tau UniEval (Relevance) wo-GPT-3.5 (Relevance) wi-gt-GPT-3.5 (Relevance) wi-in-GPT-3.5 (Relevance) wi-ingt-GPT-3.5 (Relevance)
Consistency	BARTSCORE SummaC CTC UniEval (Consistency) wo-GPT-3.5 (Consistency) wi-gt-GPT-3.5 (Consistency) wi-in-GPT-3.5 (Consistency) wi-ingt-GPT-3.5 (Consistency)
Coherence	UniEval (Coherence) wo-GPT-3.5 (Coherence) wi-gt-GPT-3.5 (Coherence) wi-in-GPT-3.5 (Coherence) wi-ingt-GPT-3.5 (Coherence)
Fluency	UniEval (Fluency) wo-GPT-3.5 (Fluency) wi-gt-GPT-3.5 (Fluency) wi-in-GPT-3.5 (Fluency) wi-ingt-GPT-3.5 (Fluency)
Overall	UniEval (Overall) wo-GPT-3.5 (Overall) wi-gt-GPT-3.5 (Overall) wi-ingt-GPT-3.5 (Overall) wi-in-GPT-3.5(Overall)

Table 2: A summary of the thirty-one NLG metrics that are used in our benchmark.

around 22 hours to run on Llama 2. For the XSUM dataset, the experiments took approximately six times longer to run compared to those on AESLC. The total cost is about 1000 USD for the GPT-3.5-based generation and GPT-3.5-based evaluation. For the density-based method, we randomly sampled 1000 training samples with a fixed random seed from the respective dataset to represent the training data distribution, following the settings outlined in [Fadeeva et al. \(2023\)](#).

4.2 Experimental Results about Uncertainty Estimation Methods

Analysis on ensemble-based uncertainty estimation methods. In the case of ensemble-based uncertainty estimation methods, T-TU, S-TU, and S-RMI generally exhibit positive correlations with all other white-box uncertainty estimation methods, except for T-RMI. This trend is evident in Figures 6 and 12. Additionally, T-RMI demonstrates negative correlations with T-TU and S-TU, as well as a weakly positive correlation with S-RMI. Consequently, we propose that for future applications of uncertainty estimation methods, focusing on one of T-TU, S-TU, and S-RMI could be advantageous, with T-RMI serving as a supplementary baseline.

Analysis on information-based uncertainty estimation methods. In the realm of information-based uncertainty estimation methods, MSP and MCSE typically exhibit strongly positive correlations with each other, as evidenced in Figures 6, 8, 12, and 14.

Conversely, MTE exhibits inconsistently strong positive correlations with MSP and MCSE. For instance, while Figure 8 demonstrates a weakly positive correlation between them, Figure 12 depicts a strongly positive correlation. Therefore, we recommend that in future comparison of uncertainty estimation methods, when considering MSP and MCSE, opting for one of them and optionally utilizing MTE as a supplementary baseline could be beneficial.

Analysis on density-based uncertainty estimation methods. In the domain of density-based uncertainty estimation methods, MD and RDE typically demonstrate strongly positive correlations with each other. This trend is observed in Figures 6, 12, and 14, with the exception of a weakly positive correlation in Figure 8.

NLG Metrics	MSP	MTE	MCSE	MD	RDE	P(True)
ROUGE-L	0.2107	0.1668	0.2082	0.2650	0.1608	-0.0233
BARTSCORE	0.0372	0.2015	0.0418	0.2451	0.1573	0.0615
SummaC	-0.1301	-0.0440	-0.1189	0.0739	-0.0182	0.0998
CTC	0.0736	-0.1515	0.0685	0.0457	0.0347	-0.1074
Spearman	0.0656	0.1430	0.0640	0.1429	0.0929	-0.0080
Kendall-Tau	0.0649	0.1412	0.0630	0.1404	0.0904	-0.0084
UniEval (Relevance)	-0.0572	0.0769	-0.0309	0.1995	-0.0334	-0.2365
UniEval (Consistency)	0.1408	0.1007	0.1143	0.0224	0.1392	-0.3789
UniEval (Coherence)	0.1804	0.1802	0.1932	0.2163	0.2332	-0.7554
UniEval (Fluency)	-0.0524	-0.3809	0.0079	0.0568	-0.1954	-0.2343
UniEval (Overall)	0.0300	0.0000	0.0563	0.1702	0.0148	-0.4410
wo-GPT-3.5 (Relevance)	-0.0863	0.2315	-0.0451	0.1784	-0.0533	0.1029
wo-GPT-3.5 (Consistency)	-0.0280	0.1825	0.0025	0.2586	0.0157	0.0849
wo-GP-T3.5 (Coherence)	-0.0796	0.0771	-0.0479	0.1672	-0.0712	0.0969
wo-GPT-3.5 (Fluency)	-0.0148	0.1498	0.0023	0.1992	-0.0347	0.0615
wo-GPT-3.5 (Overall)	-0.0634	0.0976	-0.0794	0.1807	0.0187	0.0399
wi-gt-GPT-3.5 (Relevance)	-0.2205	0.0841	-0.2692	0.0835	0.0833	-0.1177
wi-gt-GPT-3.5 (Consistency)	-0.2533	-0.0078	-0.2726	-0.0076	-0.0268	-0.0773
wi-gt-GPT-3.5 (Coherence)	-0.2991	-0.0493	-0.2979	0.0333	-0.0597	-0.1096
wi-gt-GPT-3.5 (Fluency)	-0.2344	-0.0392	-0.2084	0.1485	0.0882	-0.3175
wi-gt-GPT-3.5 (Overall)	-0.1829	0.0045	-0.2133	-0.0836	-0.0842	-0.1202
wi-in-GPT-3.5 (Relevance)	0.0185	0.0798	0.0735	0.0608	-0.0820	-0.0239
wi-in-GPT-3.5 (Consistency)	0.1015	0.1054	0.1041	0.1972	0.0545	-0.1937
wi-in-GPT-3.5 (Coherence)	-0.0170	0.0653	-0.0206	0.0283	-0.1453	-0.2302
wi-in-GPT-3.5 (Fluency)	0.0808	-0.1020	0.0955	0.1685	0.0308	-0.4904
wi-in-GPT-3.5 (Overall)	-0.1816	0.1285	-0.1818	0.0457	-0.0564	-0.0854
wi-ingt-GPT-3.5 (Relevance)	-0.1114	0.1486	-0.1015	0.1182	0.0500	-0.0379
wi-ingt-GPT-3.5 (Consistency)	-0.1391	0.1558	-0.1462	0.0934	0.0114	-0.0660
wi-ingt-GPT-3.5 (Coherence)	-0.1079	0.0909	-0.0959	0.1531	0.0058	-0.1189
wi-ingt-GPT-3.5 (Fluency)	-0.0283	0.0799	0.0065	0.1402	-0.0077	-0.2216
wi-ingt-GPT-3.5 (Overall)	-0.1149	0.0485	-0.1011	-0.0094	-0.1268	-0.0594
Col Mean	-0.0451	0.0634	-0.0364	0.1204	0.0092	-0.1263

Table 3: Main results of the relationship between the uncertainty estimation methods and NLG metrics on AESLC dataset using generation from Llama 2.

Hence, we propose that in future applications of uncertainty estimation methods, utilizing either MD or RDE alone would suffice rather than using both of them as the baselines.

Analysis on prompt-based uncertainty estimation methods. In the case of prompt-based uncertainty estimation methods, P(True) typically exhibits a weak or negative correlation with other uncertainty estimation methods. This trend is evident from Figures 8 and 14.

Analysis on black-box uncertainty estimation methods. In terms of black-box uncertainty estimation methods, ECC, LexSim, and EigV typically exhibit positive correlations with each other, with ECC and EigV showing particularly strong correlations. Conversely, NumSets demonstrates a weak or negative correlation with ECC, LexSim, and EigV. These trends are shown in Figures 7 and 13.

Hence, we recommend that in future applications of uncertainty estimation methods involving ECC, LexSim, and EigV, opting for one of them and employing NumSets as a complementary measure could be beneficial.

Analysis on specific uncertainty estimation methods. Besides, we observed that each UE method consistently achieves positive UniEval (Consistency) scores across all eight methods listed in Table 8. This indicates that all of the aforementioned methods outperform random ranking in terms of the NLG metric, UniEval (Consistency).

Each uncertainty estimation method exhibits dissatisfaction performance with negative ROUGE-L scores across all eight methods in Table 9. This suggests that all the methods listed above perform worse than random ranking in terms of the NLG metric, ROUGE-L (Consistency).

Each uncertainty estimation method demonstrates strong performance, achieving positive scores for both UniEval (Consistency) and UniEval (Coherence) across all four methods presented in Table 6. This indicates that all aforementioned methods outperform random ranking in terms of both consistency and coherence according to UniEval assessment. We summarized our findings in Sec. C.2.

NLG Metrics	EI	MR	SOAF	RE	TME	CO	NMS	Row Mean
ROUGE-L	-0.5357	0.2070	-0.3931	-0.7112	-0.4778	0.0084	0.0713	-0.2616
SummaC	-0.1109	0.0525	0.0981	0.0838	-0.2313	-0.5984	0.1242	-0.0831
CTC	-0.4205	0.2899	0.2378	-0.3425	-0.4008	0.0238	0.1431	-0.0670
Spearman	-0.4254	0.2311	-0.0991	-0.2111	-0.1136	-0.2723	0.0096	-0.1258
Kendall-Tau	-0.4216	0.2387	-0.1086	-0.2151	-0.1111	-0.2905	0.0112	-0.1281
UniEval (Relevance)	0.0388	-0.1822	0.3394	-0.2541	-0.1908	-0.1578	-0.1172	-0.0748
UniEval (Consistency)	-0.6647	-0.1144	0.3829	-0.1344	0.2585	-0.0797	-0.1606	-0.0732
UniEval (Coherence)	-0.5794	0.2330	0.6794	-0.4737	-0.2264	0.1767	-0.1442	-0.0478
UniEval (Fluency)	-0.0665	0.2300	-0.6612	-0.3210	-0.3194	0.2168	-0.3231	-0.1778
UniEval (Overall)	-0.5925	-0.0105	0.5262	-0.3591	-0.0376	0.0002	-0.2080	-0.0973
wi-ingt-GPT-3.5 (Relevance)	-0.0991	0.1672	-0.2758	-0.2624	-0.1556	-0.1358	-0.2762	-0.1482
wi-ingt-GPT-3.5 (Consistency)	-0.3081	0.1279	0.0396	-0.3236	-0.4197	-0.0768	-0.0286	-0.1413
wi-ingt-GPT-3.5 (Coherence)	-0.2569	0.0856	-0.0278	-0.3163	-0.4721	0.0666	0.1028	-0.1169
wi-ingt-GPT-3.5 (Fluency)	-0.1679	0.0791	0.0177	-0.3292	-0.2071	0.0501	0.0548	-0.0718
wi-ingt-GPT-3.5 (Overall)	-0.2945	0.0100	0.1369	-0.2878	-0.3577	0.0233	0.2251	-0.0778

Table 4: Main results of the relationship between the NLG metrics and human annotations on TofuEval dataset using generation from GPT-3.5.

4.3 Results about NLG Metrics

The experimental results and experimental analysis from the view NLG metrics are in Sec. C.1. Based on the information presented in these tables and figures in Sec. C.1, we can address our key question: “How does the choice of NLG metric affect the evaluation of uncertainty estimation methods in text summarization?” The answer is that **using different NLG metrics could lead to different ranks for uncertainty estimation methods**. Therefore, it is important to design uncertainty estimation metrics that are robust across various NLG metrics. Besides Sec. C.1, more findings are summarized in Sec. 6.

4.4 Experimental Results Involving Human Annotations

The experimental results and experimental analysis involving human annotations are in Sec. C.2. Based on the human annotation from seven dimensions, we found that the rankings of certain uncertainty estimation methods differ between NLG metrics and human annotation. Also, current uncertainty estimation methods struggle to attain positive PRR scores across various perspectives of human annotations. This could be a future direction for the research. Besides Sec. C.2, more findings are summarized in Sec. 6.

5 Conclusion

Recognizing the dependency of uncertainty model metrics on diverse NLG metrics, we introduce a comprehensive UE-TS benchmark to investigate the impact of NLG metric choice on the evaluation of uncertainty estimation methods in text sum-

marization. This benchmark encompasses **three** datasets, **fourteen** uncertainty estimation methods (both white-box and black-box), and **thirty-one** NLG metrics spanning **four** evaluation dimensions. Additionally, we assess one black-box LLM, one white-box LLM, and one PLM for text summarization.

Our findings not only highlight the current limitations in UE-TS model evaluation but also contribute to a deeper understanding of the complex relationship between NLG metrics, uncertainty estimation methods, and human annotation. This knowledge can guide future research, inform the development of more effective uncertainty estimation methods, and ultimately enhance evaluation protocols for UE-TS and other NLG tasks.

6 Summary of Our Findings

Below, we summarize our findings, which are take-away knowledge. The evidence to obtain these findings is detailed in Sec. 4.2, C.1 and C.2.

From the view of **NLG metrics** (see Sec. C.1),

- It is evident that evaluating uncertainty estimation models using different NLG metrics leads to variations in the performance ranking of these models.
- Some evaluations of uncertainty estimation models using different NLG metrics could result in different performance ranks. However, some evaluations of uncertainty estimation models using different NLG metrics may result in the same performance rankings.
- Generation models of the same type across different datasets usually result in similar cor-

relations among various methods. However, this does not apply to the fluency dimension.

- When utilizing LLMs as a type of relevance NLG metric, the choice of target text source can greatly impact the final conclusion. Specifically, using ground-truth summaries versus using input text as the target text source can result in different performance rankings.
- When using LLMs as an NLG metric, if both ground-truth summaries and input text are employed together as the target text source, ground-truth summaries typically impact metric results more than using only the input text in most cases.
- GPT-3.5 knows the concept of relevance, overall. But it might not know the concept of consistency, coherence. As for the concept of fluency, it is hard to tell whether the GPT-3.5 itself understand it or not. Therefore, using GPT-3.5 as an evaluation tool, it would be better to provide the concept definition in the prompt.

Specific recommendations for **NLG Metrics**:

- **Relevance:** Spearman and Kendall-Tau metrics show positive correlations, so either can be used. When using LLMs as generation models, UniEval (Relevance) and wo-GPT-3.5 (Relevance) typically show positive correlations with most other NLG metrics. Therefore, either of them could serve as a representative NLG metric in the relevance dimension (see Sec. C.1.2).
- **Consistency:** CTC may be preferable to SummaC or UniEval due to positive correlations. For GPT-3.5-based metrics, it is recommended to use wi-in-GPT-3.5 (Consistency) alongside wo-gt-GPT-3.5, wi-gt-GPT-3.5 (Consistency), or wi-ingt-GPT-3.5 (Consistency) (see Sec. C.1.3).
- **Coherence:** Determining the superior GPT-3.5-based metric in the coherence dimension is challenging. But, due to their strongly positive correlation, either wi-gt-GPT-3.5 (Coherence) or wi-in-GPT-3.5 (Coherence) is a suitable choice. In the coherence dimension, UniEval (Coherence) could complement either wi-gt-GPT-3.5 (Coherence) or wi-in-GPT-3.5 (Coherence) (see Sec. C.1.4).

- **Fluency:** We recommend wi-gt-GPT-3.5 (Fluency) and wi-ingt-GPT-3.5 (Fluency) due to their correlation patterns. UniEval (Fluency) can supplement a GPT-3.5-based metric (see Sec. C.1.5).
- The UniEval (Overall) can serve as an optional supplementary tool for GPT-3.5-based NLG metrics in the overall dimension (see Sec. C.1.6).

Key insights for **uncertainty estimation methods** (see Sec. 4.2): For future applications of uncertainty estimation methods:

- Prioritize one of T-TU, S-TU, and S-RMI, with T-RMI serving as a supplementary baseline.
- When considering MSP and MCSE, opting for one of them and utilizing MTE as an optional supplementary measure could be beneficial.
- Using either MD or RDE alone would suffice rather than using both of them as the baselines.
- When considering ECC, LexSim, and EigV, opting for one of them and employing NumSets as a complementary measure could be beneficial.

From the view of **human annotations** (see Sec. C.2.2, C.2.3, C.2.4),

- Rankings of certain uncertainty estimation methods differ between NLG metrics and human annotation.
- Utilizing NLG metrics as an oracle remains meaningful, as it saves the cost but leads to an uncertainty estimation method rank that is inconsistent with human annotations.
- Current uncertainty estimation methods struggle to attain positive PRR scores across various perspectives of human annotations. This could be a future direction for the research.

7 Acknowledgment

This research is supported in part by the National Science Foundation (NSF) grants IIS-2107449, FAI-2147375, and IIS-1954376.

8 Ethical Consideration

This study explores the overlooked relationship between uncertainty metrics and NLG metrics. During our study, we answer the question “Can We Trust the Performance Evaluation of Uncertainty Estimation Methods in Text Summarization?”.

Our research employs datasets that are publicly available, ensuring transparency and accessibility. The datasets integral to our work are utilized in adherence to their respective licenses (Tang et al., 2024; Fadeeva et al., 2023).

We recommend that any future expansion of this research into areas involving personal or sensitive data should be approached with stringent ethical guidelines in place.

9 Limitations

Due to the resources (GPU and budget for calling GPT APIs), we only conduct one-time experiments. In contrast, the experiments on BART are set to use greedy generation for each sample rather than diverse generations for the robustness of the experiments.

References

- Mohammad Beigi, Ying Shen, Runing Yang, Zihao Lin, Qifan Wang, Ankith Mohan, Jianfeng He, Ming Jin, Chang-Tien Lu, and Lifu Huang. 2024. Internalinspector i^2 : Robust confidence estimation in llms through internal states. *arXiv preprint arXiv:2406.12053*.
- Yu-Neng Chuang, Ruixiang Tang, Xiaoqian Jiang, and Xia Hu. 2024. Spec: A soft prompt-based calibration on performance variability of large language model in clinical notes summarization. *Journal of Biomedical Informatics*, page 104606.
- Mingkai Deng, Bowen Tan, Zhengzhong Liu, Eric P Xing, and Zhiting Hu. 2021. Compression, transduction, and creation: A unified framework for evaluating natural language generation. *arXiv preprint arXiv:2109.06379*.
- Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, Kirill Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, et al. 2023. Lmpolygraph: Uncertainty estimation for language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 446–461.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. Unsupervised quality estimation for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:539–555.
- Yarin Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR.
- Alexios Gidiotis and Grigorios Tsoumakas. 2021a. Bayesian active summarization. *arXiv preprint arXiv:2110.04480*.
- Alexios Gidiotis and Grigorios Tsoumakas. 2021b. Should we trust this summary? bayesian abstractive summarization to the rescue. *arXiv preprint arXiv:2105.10155*.
- Antonio Duréndez Gómez-Guillamón. 2003. The usefulness of the audit report in investment and financing decisions. *Managerial auditing journal*, 18(6/7):549–559.
- Vishal Gupta and Gurpreet Singh Lehal. 2010. A survey of text summarization extractive techniques. *Journal of emerging technologies in web intelligence*, 2(3):258–268.
- Jianfeng He, Hang Su, Jason Cai, Igor Shalyminov, Hwanjun Song, and Saab Mansour. 2024. Semi-supervised dialogue abstractive summarization via high-quality pseudolabel selection. *arXiv preprint arXiv:2403.04073*.
- Jianfeng He, Xuchao Zhang, Shuo Lei, Zhiqian Chen, Fanglan Chen, Abdulaziz Alhamadani, Bei Xiao, and ChangTien Lu. 2020. Towards more accurate uncertainty estimation in text classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8362–8372.
- Yibo Hu and Latifur Khan. 2021. Uncertainty-aware reliable text classification. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 628–636.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Zahra Kolagar and Alessandra Zarccone. 2024. Aligning uncertainty: Leveraging llms to analyze uncertainty transfer in text summarization. In *Proceedings of the 1st Workshop on Uncertainty-Aware NLP (UncertainNLP 2024)*, pages 41–61.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*.

- Philippe Laban, Tobias Schnabel, Paul N Bennett, and Marti A Hearst. 2022. Summac: Re-visiting nli-based models for inconsistency detection in summarization. *Transactions of the Association for Computational Linguistics*, 10:163–177.
- Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31.
- Yuanyuan Lei, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Ruihong Huang, and Dong Yu. 2024. Polarity calibration for opinion summarization. *arXiv preprint arXiv:2404.01706*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2023. Generating with confidence: Uncertainty quantification for black-box large language models. *arXiv preprint arXiv:2305.19187*.
- Zihao Lin, Mohammad Beigi, Hongxuan Li, Yufan Zhou, Yuxiang Zhang, Qifan Wang, Wenpeng Yin, and Lifu Huang. 2024. Navigating the dual facets: A comprehensive evaluation of sequential memory editing in large language models. *arXiv preprint arXiv:2402.11122*.
- Pengfei Liu, Jinlan Fu, Yang Xiao, Weizhe Yuan, Shuaichen Chang, Junqi Dai, Yixin Liu, Zihuiwen Ye, and Graham Neubig. 2021. **ExplainsBoard: An explainable leaderboard for NLP**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 280–289, Online. Association for Computational Linguistics.
- Yang Liu. 2019. Fine-tune bert for extractive summarization. *arXiv preprint arXiv:1903.10318*.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*.
- Andrey Malinin and Mark Gales. 2020. Uncertainty estimation in autoregressive structured prediction. *arXiv preprint arXiv:2002.07650*.
- Andrey Malinin, Anton Ragni, Kate Knill, and Mark Gales. 2017a. **Incorporating uncertainty into deep learning for spoken language assessment**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 45–50, Vancouver, Canada. Association for Computational Linguistics.
- Andrey Malinin, Anton Ragni, Kate Knill, and Mark Gales. 2017b. Incorporating uncertainty into deep learning for spoken language assessment. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 45–50.
- Jishnu Mukhoti, Andreas Kirsch, Joost van Amersfoort, Philip HS Torr, and Yarin Gal. 2023. Deep deterministic uncertainty: A new simple baseline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24384–24394.
- Leann Myers and Maria J Sirois. 2004. Spearman correlation coefficients, differences between. *Encyclopedia of statistical sciences*, 12.
- Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- OpenAI. 2023. Gpt-3.5. <https://platform.openai.com/docs/models/gpt-3-5-turbo>.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, and Peter J Liu. 2022. Out-of-distribution detection and selective generation for conditional language models. In *The Eleventh International Conference on Learning Representations*.
- Philip Sedgwick. 2012. Pearson’s correlation coefficient. *Bmj*, 345.
- Edwin Simpson, Yang Gao, and Iryna Gurevych. 2020. Interactive text ranking with bayesian optimization: A case study on community qa and summarization. *Transactions of the Association for Computational Linguistics*, 8:759–775.
- Liyan Tang, Igor Shalyminov, Amy Wing-mei Wong, Jon Burnsky, Jake W Vincent, Yu’an Yang, Siffi Singh, Song Feng, Hwanjun Song, Hang Su, et al. 2024. Tofueval: Evaluating hallucinations of llms on topic-focused dialogue summarization. *arXiv preprint arXiv:2402.13249*.
- Oguzhan Tas and Farzad Kiyani. 2007. A survey automatic text summarization. *PressAcademia Procedia*, 5(1):205–213.

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Akim Tsvigun, Ivan Lysenko, Danila Sedashov, Ivan Lazichny, Eldar Damirov, Vladimir Karlov, Artemy Belousov, Leonid Sanochkin, Maxim Panov, Alexander Panchenko, et al. 2023. Active learning for abstractive text summarization. *arXiv preprint arXiv:2301.03252*.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2016. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*.
- Shuo Wang, Yang Liu, Chao Wang, Huanbo Luan, and Maosong Sun. 2019. Improving back-translation with uncertainty-based confidence estimation. *arXiv preprint arXiv:1909.00157*.
- Kam-Fai Wong, Mingli Wu, and Wenjie Li. 2008. Extractive summarization using supervised and semi-supervised learning. In *Proceedings of the 22nd international conference on computational linguistics (Coling 2008)*, pages 985–992.
- Yuan Xue. 2022. *Relationships between US Universities and International-Student Recruitment Agents*. Ph.D. thesis, Edgewood College.
- KiYoon Yoo, Jangho Kim, Jiho Jang, and Nojun Kwak. 2022. Detection of word adversarial examples in text classification: Benchmark and baseline via robust density estimation. *arXiv preprint arXiv:2203.01677*.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277.
- Polina Zablotskaia, Du Phan, Joshua Maynez, Shashi Narayan, Jie Ren, and Jeremiah Liu. 2023. On uncertainty calibration and selective generation in probabilistic neural summarization: A benchmark study. *arXiv preprint arXiv:2304.08653*.
- Rui Zhang and Joel Tetreault. 2019. This email could save your life: Introducing the task of email subject line generation. *arXiv preprint arXiv:1906.03497*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Xingxing Zhang, Yiran Liu, Xun Wang, Pengcheng He, Yang Yu, Si-Qing Chen, Wayne Xiong, and Furu Wei. 2022. Momentum calibration for text generation. *arXiv preprint arXiv:2212.04257*.
- Yao Zhao, Mikhail Khalman, Rishabh Joshi, Shashi Narayan, Mohammad Saleh, and Peter J Liu. 2022. Calibrating sequence likelihood improves conditional language generation. In *The Eleventh International Conference on Learning Representations*.
- Xin Zheng, Yusi Feng, and Hongkai Chen. 2022. Analysis of each components of glass samples based on the spearman correlation coefficient model. *Highlights in Science, Engineering and Technology*, 22:241–245.
- Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. Towards a unified multi-dimensional evaluator for text generation. *arXiv preprint arXiv:2210.07197*.

A More Explanation of Uncertainty Estimation Methods

Below we give a detailed explanation of the uncertainty estimation methods in Table 1.

A.1 White-Box Methods

The UE-TS methods primarily fall into four categories (Fadeeva et al., 2023): information-based, density-based, ensemble-based, and prompt-based methods.

Information-based methods utilize token probability to obtain uncertainty scores. Within this category, we employ the following methods:

- (1) **Maximum Sequence Probability (MSP)**: estimates uncertainty score as the log-probability of the generation in a greedy search way. It is calculated as the sum of log probabilities in each token.
- (2) **Mean Token Entropy (MTE)**: estimates the uncertainty score as the mean entropy for all tokens in a generation.
- (3) **Monte Carlo Sequence Entropy (MCSE)**: calculates the generation entropy estimations using Monte-Carlo estimation. It is the “predictive entropy” in Kuhn et al. (2023).

Density-based methods utilize latent representations of instances to construct a probability density. Below, we list the related methods compared in our benchmark.

- (4) **Mahalanobis Distance (MD)**: calculates a Gaussian distribution for training samples. Then, the MD calculates the distance between the testing sample and the Gaussian distribution as the uncertainty score (Lee et al., 2018).
- (5) **Robust Density Estimation (RDE)**: improves over MD by reducing the dimensionality via principal component analysis decomposition (Yoo et al., 2022).

Ensemble-based methods utilize ensembles to approximate Bayesian neural networks (BNN) or the variance of ensemble generations to obtain uncertainty scores. Below are the related methods we used:

- (6) **Token-level Total Uncertainty (T-TU)**: calculates the entropy of the predictive posterior at token level as the uncertainty score (Malinin and Gales, 2020; He et al., 2024).
- (7) **Token-level Reverse Mutual Information (T-RMI)**: uses reverse-KL divergence counterpart to the mutual information at token level as the uncertainty score (Malinin and Gales, 2020).

- (8) **Sequence-level Total Uncertainty (S-TU)**: calculates the entropy of the predictive posterior at sequence level as the uncertainty score (Malinin and Gales, 2020; He et al., 2024).

- (9) **Sequence-level Reverse Mutual Information (S-RMI)**: uses reverse-KL divergence counterpart to the mutual information at the sequence level as the uncertainty score (Malinin and Gales, 2020).

Prompt-based methods refer to methods that prompt the generation model to obtain the uncertainty score (Kadavath et al., 2022). We employ (10) **P(True)** (Kadavath et al., 2022), which takes the uncertainty score as the probability of asking whether the proposed answer is true or false. This method assumes that the generation model possesses superior zero-shot prompt abilities.

A.2 Black-Box Methods

For scenarios where only the final textual output is available, we utilize the following black-box uncertainty estimation methods, which have demonstrated effectiveness in previous studies (Fadeeva et al., 2023).

- (11) **Number of Semantic Sets (NumSets)**: takes the number of diverse semantic interpretations for the generation as uncertainty score (Lin et al., 2023).
- (12) **Eccentricity (ECC)**: gets the uncertainty via calculating a distance between all eigenvectors that are informative embeddings of graph Laplacian (Lin et al., 2023).
- (13) **Lexical Similarity (LexSim)**: obtains the uncertainty scores via calculating mean similarity between all pairs of sampled generations (Fomicheva et al., 2020).
- (14) **Sum of Eigenvalues of the Graph Laplacian (EigV)**: extends the NumSets from integer case into a continuous case, where the uncertainty score is calculated based on a matrix trace (Lin et al., 2023).

B More Explanation of NLG metrics

Below we give a detailed explanation of the NLG metrics in Table 2.

B.1 Specific NLG Metrics

- (1) **ROUGE-L**: measures the longest common subsequence between the generated text and ground-truth text (Lin, 2004). It emphasizes relevance.
- (2) **BARTSCORE**: uses an encoder-decoder pre-trained model to compute a similarity score for

each token in the generation with each token in the reference text Yuan et al. (2021). It emphasizes consistency.

(3) **CTC**: uses a pretrained model to measure the alignment between the generation text and ground-truth text (Deng et al., 2021). It emphasizes consistency.

(4) **SummaC**: uses pretrained models to segment both generated and input texts into sentence units and aggregate scores between pairs of sentences (Laban et al., 2022). It emphasizes consistency.

(5) Spearman Correlation (**Spearman**) and (6) Kendall-Tau Correlation (**Kendall-Tau**) are designed to measure the semantic overlap between the model output and the reference text via text embeddings. They have relatively high correlations in the relevance dimension (Liu et al., 2021; Zhong et al., 2022).

(7-11) **UniEVAL**: takes NLG evaluation as a boolean question answering (QA) task and guides the model with different questions. UniEVAL (Zhong et al., 2022) can use one evaluator to evaluate one of the dimensions in Sec. 3.5.1. (12-16) GPT-3.5 without dimension concepts **w-gt-GPT-3.5**: uses GPT-3.5 (OpenAI, 2023) to prompt about the generation quality. Among the prompts, we do not provide any prior knowledge about each dimension concept described in Sec. 3.5.1. This method compares the difference between the generated summaries and *ground-truth summaries*.

(17-21) GPT-3.5 with dimension concepts **wi-gt-GPT-3.5**: uses the GPT-3.5 (OpenAI, 2023) to prompt about the generation quality. However, the prior knowledge about each dimension concept is given in the prompt. This method compares the difference between the generated summaries and *ground-truth summaries*.

(22-26) **wi-in-GPT-3.5**: is very similar to wi-gt-GPT-3.5. The only difference is that this method compares the difference between the generated summaries and *input text*.

(27-31) **wi-ingt-GPT-3.5**: is very similar to wi-gt-GPT-3.5. The only difference is that this method compares generated summaries to the *input text and ground-truth summaries*.

Although different NLP tasks may use various evaluation metrics, such as those in Lin et al. (2024), we adopt the evaluation metric settings from Zhong et al. (2022) due to their generalizability to other NLG tasks. Consequently, our research conclusions may be applicable to a broader range

of NLG tasks.

C More Experimental Results

In this subsection, we list more experimental results below. We first give the experimental results from the view of NLG metrics. Then, we show the experimental results involving human annotations.

C.1 Experimental Results about NLG Metrics

Tables 8, 6, 3, 9, 7, and 5 present various uncertainty metric scores obtained from different uncertainty methods alongside different NLG metrics. Figures 6, 7, 8, 12, 13, and 14 illustrate the correlation of uncertainty methods in terms of uncertainty metric scores. Additionally, Figures 3, 4, 5, 9, 10, and 11 depict the correlation of NLG metrics with uncertainty metric scores.

Based on the information presented in these tables and figures, we can address our key question: “How does the choice of NLG metric affect the evaluation of uncertainty estimation methods in text summarization?” The answer is that **using different NLG metrics could lead to different ranks for uncertainty estimation methods**. Therefore, it is important to design uncertainty estimation metrics that are robust across various NLG metrics.

Next, we provide a detailed analysis from the perspectives of NLG metrics and uncertainty estimation methods, respectively.

C.1.1 Analysis Based on All NLG Metrics

Analysis from all five dimensions. Figures 3, 4, and 5 show the Spearman correlation between NLG metrics in a comprehensive view of the AESLC dataset. Figures 9, 10, and 11 show the Spearman correlation between NLG metrics in a comprehensive view of the XSUM dataset. We use Spearman correlation (Myers and Sirois, 2004; Zheng et al., 2022) rather than Pearson correlation (Sedgwick, 2012; Xue, 2022). This is because the Spearman correlation is calculated based on the ranks of the data while the Pearson correlation is calculated via the raw data values. Since our work wants to know the impact on the performance rank of the different uncertainty estimation methods. We use the Spearman correlation. From these figures, we can draw the following conclusions.

It is evident that evaluating uncertainty estimation models using different NLG metrics leads to variations in the performance ranking of these models. This discrepancy arises because there are no rows or columns in these figures where all elements

NLG Metrics	MSP	MTE	MCSE	MD	RDE	P(True)
ROUGE-L	0.1300	0.1171	0.1257	0.0096	-0.0232	-0.1698
BARTSCORE	0.1657	0.2000	0.1789	0.2132	0.1891	-0.1828
SummaC	-0.0659	-0.0580	-0.0642	-0.1556	-0.2557	0.0392
CTC	0.0448	-0.0112	0.0699	-0.0021	-0.0759	-0.3161
Spearman	0.0189	0.0516	0.0094	0.0569	0.0190	-0.2855
Kendall-Tau	0.0138	0.0437	0.0034	0.0495	0.0151	-0.2735
UniEval (Relevance)	0.4496	0.4654	0.4586	0.3709	0.1655	-0.8063
UniEval (Consistency)	0.5711	0.5589	0.5746	0.4069	0.2830	-0.7399
UniEval (Coherence)	0.7198	0.7069	0.7219	0.5884	0.2795	-1.0973
UniEval (Fluency)	0.3714	0.3008	0.3969	0.3526	0.2316	-0.5880
UniEval (Overall)	0.6020	0.5904	0.6113	0.4893	0.2567	-0.9452
wo-GPT-3.5 (Relevance)	0.3447	0.3852	0.3756	0.2945	-0.0065	-0.3841
wo-GPT-3.5 (Consistency)	0.0789	0.1104	0.0990	0.0993	0.0142	-0.0140
wo-GP-T3.5 (Coherence)	-0.0197	-0.0238	-0.0012	-0.0316	0.0102	0.0612
wo-GPT-3.5 (Fluency)	0.1701	0.1997	0.1935	0.1850	-0.0118	-0.0961
wo-GPT-3.5 (Overall)	0.1249	0.1626	0.1442	0.1562	0.0138	-0.1296
wi-gt-GPT-3.5 (Relevance)	-0.1070	-0.1089	-0.0933	-0.1349	-0.1222	-0.1304
wi-gt-GPT-3.5 (Consistency)	-0.0985	-0.1148	-0.0832	-0.0625	-0.0538	-0.0531
wi-gt-GPT-3.5 (Coherence)	-0.0546	-0.0504	-0.0270	-0.0813	-0.0471	-0.0909
wi-gt-GPT-3.5 (Fluency)	-0.1158	-0.1192	-0.1139	-0.1233	-0.1168	-0.1173
wi-gt-GPT-3.5 (Overall)	-0.0327	-0.0285	-0.0130	-0.0718	-0.0911	-0.0383
wi-in-GPT-3.5 (Relevance)	-0.0105	-0.0084	-0.0035	-0.0576	-0.0280	-0.0266
wi-in-GPT-3.5 (Consistency)	-0.0276	-0.0289	-0.0069	-0.0458	-0.0211	0.0216
wi-in-GPT-3.5 (Coherence)	-0.0454	-0.0456	-0.0234	-0.0526	-0.0171	0.0429
wi-in-GPT-3.5 (Fluency)	-0.1349	-0.1423	-0.1161	-0.1589	-0.1039	0.0210
wi-in-GPT-3.5 (Overall)	-0.1067	-0.1069	-0.0861	-0.1170	-0.0915	-0.0361
wi-ingt-GPT-3.5 (Relevance)	-0.0269	-0.0038	-0.0295	-0.0894	-0.1017	-0.0178
wi-ingt-GPT-3.5 (Consistency)	-0.0940	-0.0760	-0.0917	-0.1211	-0.0636	-0.0263
wi-ingt-GPT-3.5 (Coherence)	-0.0587	-0.0155	-0.0530	-0.1187	-0.1089	-0.0441
wi-ingt-GPT-3.5 (Fluency)	-0.1545	-0.1254	-0.1790	-0.1492	-0.1588	-0.0200
wi-ingt-GPT-3.5 (Overall)	-0.0362	-0.0156	-0.0098	-0.0193	-0.0316	-0.0162
Col Mean	0.0844	0.0906	0.0957	0.0542	-0.0017	-0.2084

Table 5: Main results of the relationship between the uncertainty estimation methods and NLG metrics on XSUM dataset using generation from Llama 2.

NLG Metrics	NumSets	ECC	LexSim	EigV
ROUGE-L	0.0415	0.1071	0.2121	0.0964
SummaC	0.0919	-0.1259	-0.1362	-0.0994
CTC	0.0387	0.0999	0.0506	0.1132
Spearman	0.0326	0.0101	0.1344	0.0112
Kendall-Tau	0.0326	0.0102	0.1324	0.0114
UniEval (Relevance)	0.1131	-0.0728	-0.0165	-0.0747
UniEval (Consistency)	0.0559	0.0898	0.1260	0.1211
UniEval (Coherence)	0.0992	0.1668	0.2259	0.1754
UniEval (Fluency)	-0.0252	-0.1101	-0.2024	-0.0921
UniEval (Overall)	0.0946	-0.0379	-0.0055	-0.0284
wo-GPT-3.5 (Relevance)	0.0714	-0.0908	0.1165	-0.1115
wo-GPT-3.5 (Consistency)	0.0977	-0.0382	0.0806	-0.0488
wo-GP-T3.5 (Coherence)	0.1004	-0.0981	0.0624	-0.1043
wo-GPT-3.5 (Fluency)	0.1217	-0.1114	0.0466	-0.1265
wo-GPT-3.5 (Overall)	0.1008	-0.0548	0.0785	-0.0734
wi-gt-GPT-3.5 (Relevance)	0.1465	-0.0595	0.1097	-0.0893
wi-gt-GPT-3.5 (Consistency)	0.1002	-0.0101	0.0344	-0.0330
wi-gt-GPT-3.5 (Coherence)	0.1418	-0.0267	0.0512	-0.0519
wi-gt-GPT-3.5 (Fluency)	0.1501	0.0117	0.1928	0.0091
wi-gt-GPT-3.5 (Overall)	0.1226	-0.0892	0.0206	-0.0953
wi-in-GPT-3.5 (Relevance)	0.0303	-0.0291	0.0783	-0.0334
wi-in-GPT-3.5 (Consistency)	0.0792	0.1707	0.2616	0.1626
wi-in-GPT-3.5 (Coherence)	0.1428	0.0940	0.2092	0.0840
wi-in-GPT-3.5 (Fluency)	0.0167	0.1683	0.0820	0.1765
wi-in-GPT-3.5 (Overall)	0.0774	-0.1700	0.0117	-0.1711
wi-ingt-GPT-3.5 (Relevance)	0.1462	-0.1226	0.0659	-0.1473
wi-ingt-GPT-3.5 (Consistency)	0.0677	-0.1344	0.0135	-0.1543
wi-ingt-GPT-3.5 (Coherence)	0.0395	-0.1203	0.0317	-0.1438
wi-ingt-GPT-3.5 (Fluency)	0.0971	-0.0075	0.1304	-0.0011
wi-ingt-GPT-3.5 (Overall)	0.1281	-0.0642	0.0760	-0.0918
Col Mean	0.0851	-0.0215	0.0758	-0.0270

Table 6: Main results of the relationship between the uncertainty estimation methods and NLG metrics on AESLC dataset using generation from GPT-3.5.

NLG Metrics	NumSets	ECC	LexSim	EigV
ROUGE-L	-0.0207	0.0867	0.0759	0.0565
SummaC	0.0445	-0.0485	-0.1196	-0.0297
CTC	0.0058	0.1267	0.0278	0.0956
Spearman	-0.0045	-0.0721	-0.1082	-0.0781
Kendall-Tau	-0.0046	-0.0664	-0.1084	-0.0716
UniEval (Relevance)	0.0468	-0.0740	-0.1099	-0.0728
UniEval (Consistency)	-0.0325	0.3751	0.4828	0.3161
UniEval (Coherence)	-0.0033	0.3327	0.4128	0.2822
UniEval (Fluency)	0.0573	0.2495	0.1347	0.2297
UniEval (Overall)	0.0259	0.1686	0.1814	0.1385
wo-GPT-3.5 (Relevance)	0.0714	-0.1213	0.0464	-0.1533
wo-GPT-3.5 (Consistency)	0.0266	-0.0478	-0.0526	-0.0624
wo-GP-T3.5 (Coherence)	0.0325	-0.0464	-0.0550	-0.0558
wo-GPT-3.5 (Fluency)	0.0105	-0.0094	-0.0039	-0.0272
wo-GPT-3.5 (Overall)	-0.0007	-0.0300	-0.0094	-0.0379
wi-gt-GPT-3.5 (Relevance)	-0.0598	-0.1062	-0.0832	-0.0932
wi-gt-GPT-3.5 (Consistency)	-0.0066	-0.0486	-0.0582	-0.0516
wi-gt-GPT-3.5 (Coherence)	-0.0918	-0.0589	-0.0727	-0.0541
wi-gt-GPT-3.5 (Fluency)	-0.0402	-0.1248	-0.0733	-0.1125
wi-gt-GPT-3.5 (Overall)	-0.0563	-0.0768	-0.0341	-0.0771
wi-in-GPT-3.5 (Relevance)	0.0122	-0.0074	-0.0050	-0.0112
wi-in-GPT-3.5 (Consistency)	-0.0130	-0.0686	-0.0892	-0.0739
wi-in-GPT-3.5 (Coherence)	-0.0295	-0.0117	-0.0832	-0.0214
wi-in-GPT-3.5 (Fluency)	0.0059	-0.1526	-0.1198	-0.1627
wi-in-GPT-3.5 (Overall)	-0.0506	-0.1875	-0.1962	-0.1827
wi-ingt-GPT-3.5 (Relevance)	-0.0152	-0.0612	-0.0431	-0.0502
wi-ingt-GPT-3.5 (Consistency)	-0.0176	-0.0290	-0.0326	-0.0167
wi-ingt-GPT-3.5 (Coherence)	-0.0870	-0.1265	-0.0924	-0.1268
wi-ingt-GPT-3.5 (Fluency)	-0.1197	-0.0844	-0.0136	-0.0609
wi-ingt-GPT-3.5 (Overall)	-0.0136	-0.1040	-0.1022	-0.0979
Col Mean	-0.0109	-0.0142	-0.0101	-0.0221

Table 7: Main results of the relationship between the uncertainty estimation methods and NLG metrics on XSUM dataset using generation from GPT-3.5.

NLG Metrics	MSP	MTE	MCSE	MD	RDE	T-TU	T-RMI	S-TU	S-RMI
ROUGE-L	0.1763	-0.0385	0.1539	-0.0333	-0.0844	-0.0571	-0.0752	-0.0849	0.0577
BARTSCORE	-0.0541	0.0314	0.0252	0.1228	0.0930	0.0525	-0.0867	0.0026	0.0590
SummaC	0.0106	-0.0024	0.0304	0.0215	-0.0719	-0.0247	-0.0422	0.0229	0.0267
CTC	0.1101	-0.2225	0.0569	-0.2217	-0.2708	-0.1256	-0.0538	-0.1675	0.0765
Spearman	0.0581	0.0108	0.1008	0.0358	-0.0096	0.0393	-0.0444	-0.0038	0.0472
Kendall-Tau	0.0659	0.0080	0.1043	0.0316	-0.0141	0.0345	-0.0443	-0.0077	0.0458
UniEval (Relevance)	-0.0262	0.1157	0.0480	0.0851	0.0233	0.1165	0.1175	0.1388	0.1046
UniEval (Consistency)	0.3493	0.4330	0.2544	0.3288	0.2752	0.2014	0.0470	0.2281	0.1353
UniEval (Coherence)	0.4194	0.5342	0.2205	0.4139	0.3656	0.2693	0.0839	0.3005	0.0952
UniEval (Fluency)	0.1141	0.0538	0.0782	-0.0489	-0.0556	0.0228	0.1079	0.0917	0.0844
UniEval (Overall)	0.1508	0.2423	0.1274	0.1586	0.1071	0.1490	0.1154	0.1902	0.1172
wo-GPT-3.5 (Relevance)	-0.0183	0.1433	0.0632	0.1613	0.0862	0.0948	-0.1168	0.0700	0.0527
wo-GPT-3.5 (Consistency)	0.1177	0.1945	0.0912	0.1519	0.1255	0.0422	-0.0873	0.0435	0.0546
wo-GP-T3.5 (Coherence)	0.2006	0.2308	0.0798	0.1344	0.1176	0.0535	-0.0777	0.0568	0.0143
wo-GPT-3.5 (Fluency)	0.1277	0.2753	0.0995	0.2308	0.1792	0.0913	-0.0589	0.1319	0.0489
wo-GPT-3.5 (Overall)	0.0287	0.2327	0.0661	0.1784	0.1434	0.1107	-0.1695	0.0788	0.0688
Col Mean	0.1144	0.1402	0.1000	0.1094	0.0631	0.0669	-0.0241	0.0682	0.0681

Table 8: Main results of the relationship between the uncertainty estimation methods and NLG metrics on AESLC dataset using generation from BART.

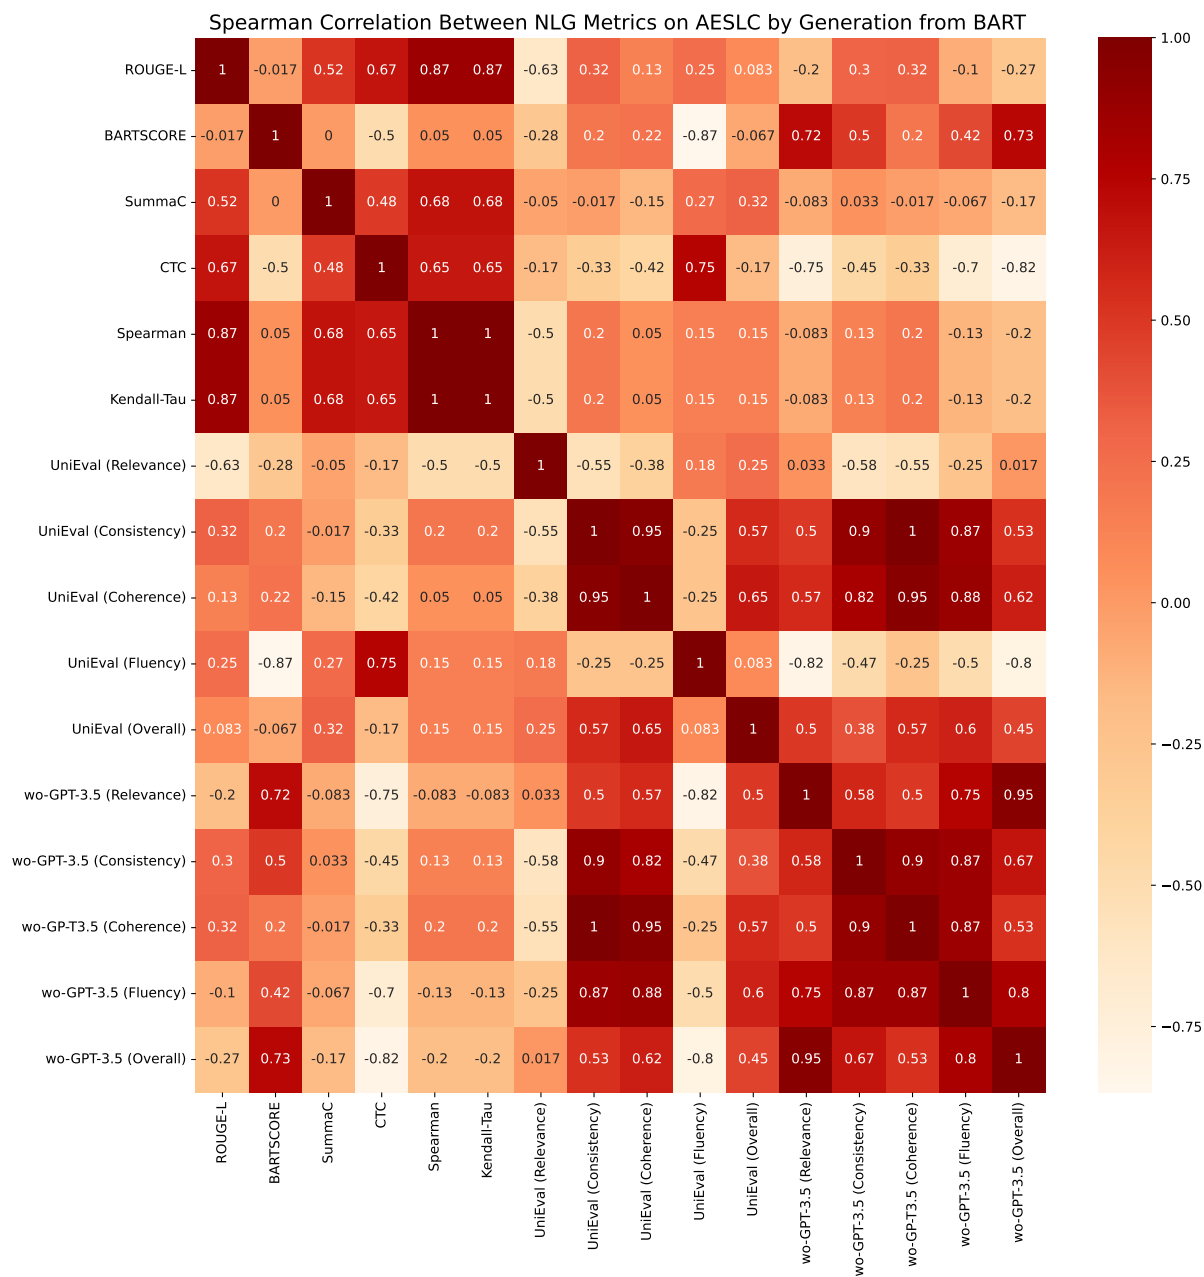


Figure 3: Diagram of Spearman correlation between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 6. The generated summaries are from BART. For the GPT-3.5-based NLG metrics, we only conduct wo-GPT-3.5 on the BART generation model setting.

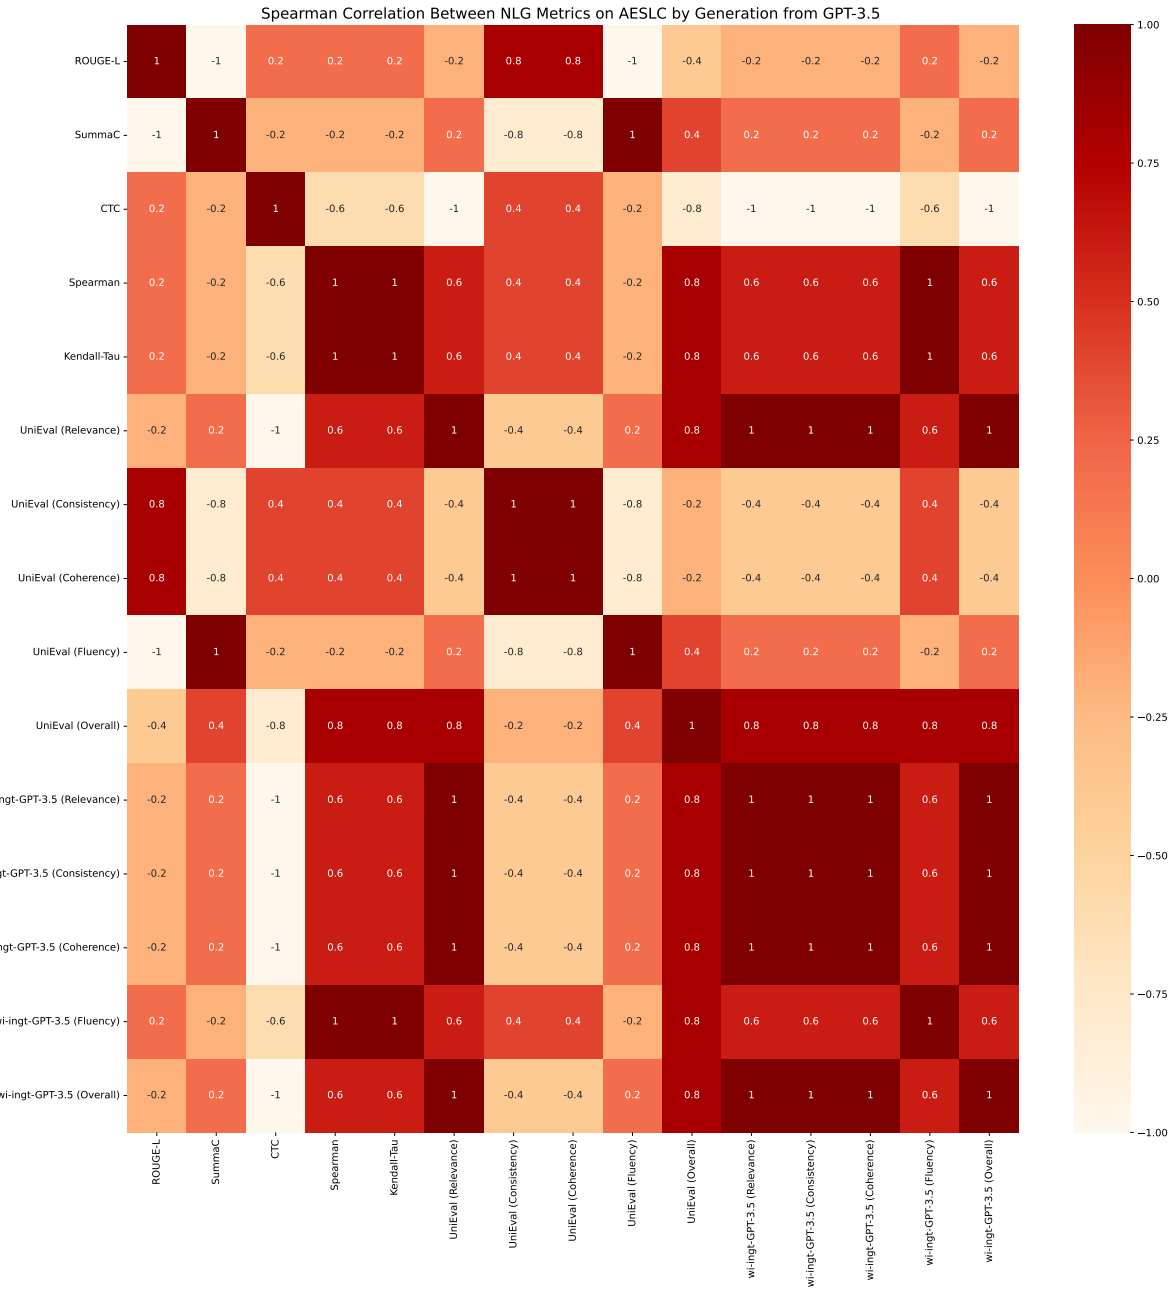


Figure 4: Diagram of Spearman correlation between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 7. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

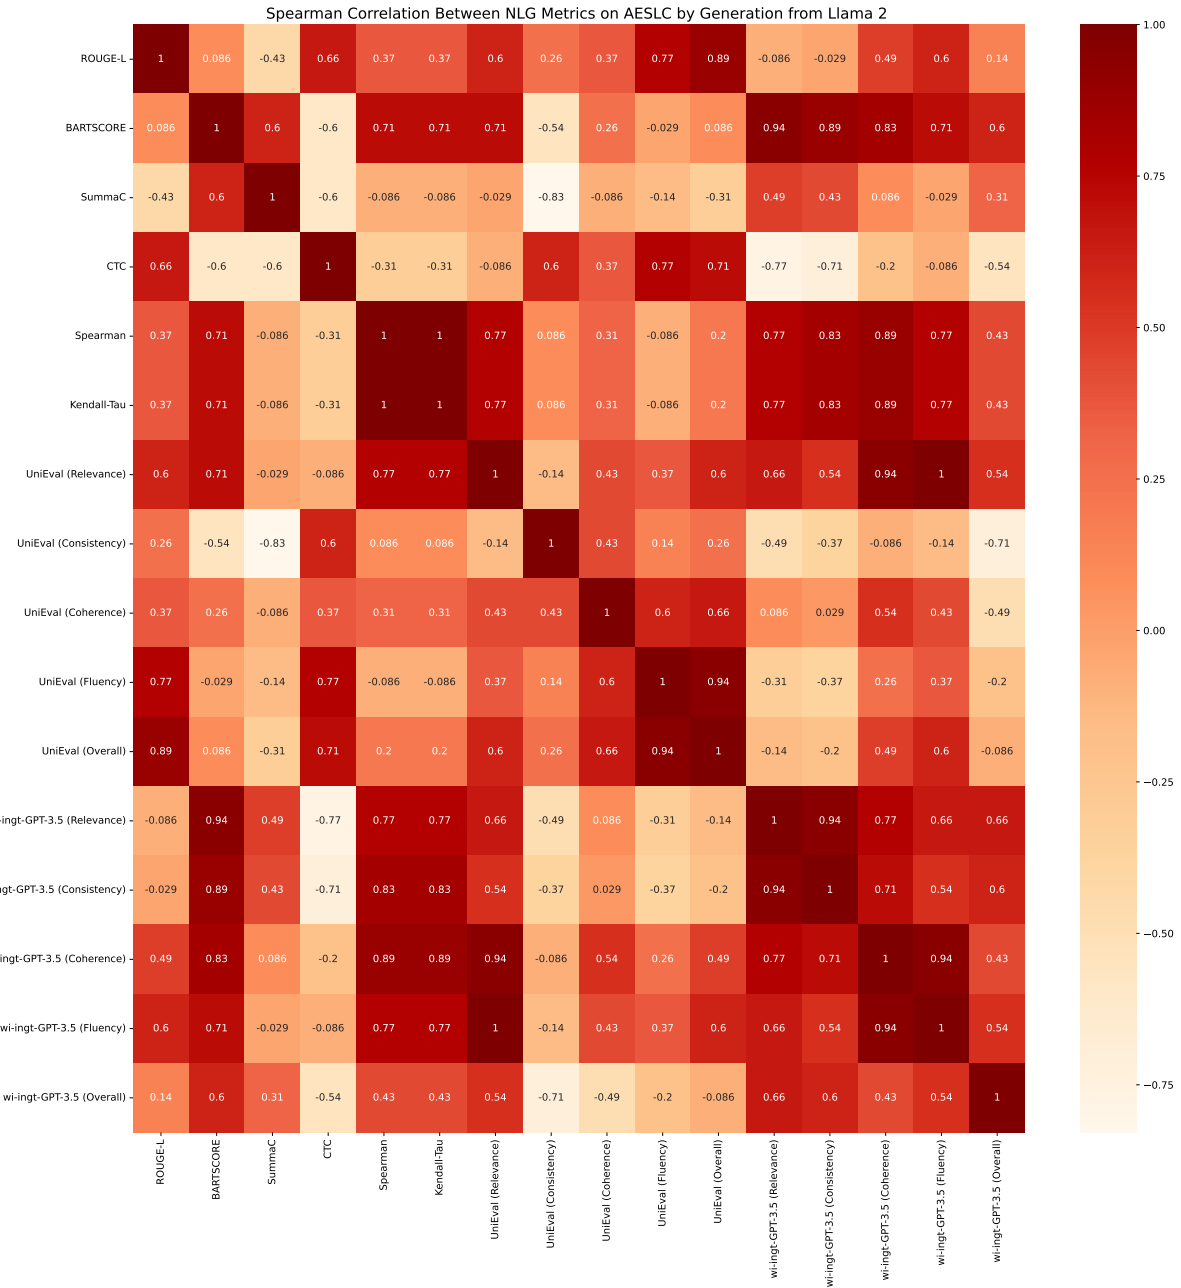


Figure 5: Diagram of Spearman correlation between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 8. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

Spearman Correlation Between Uncertainty Estimation Methods on AESLC by Generation from BART

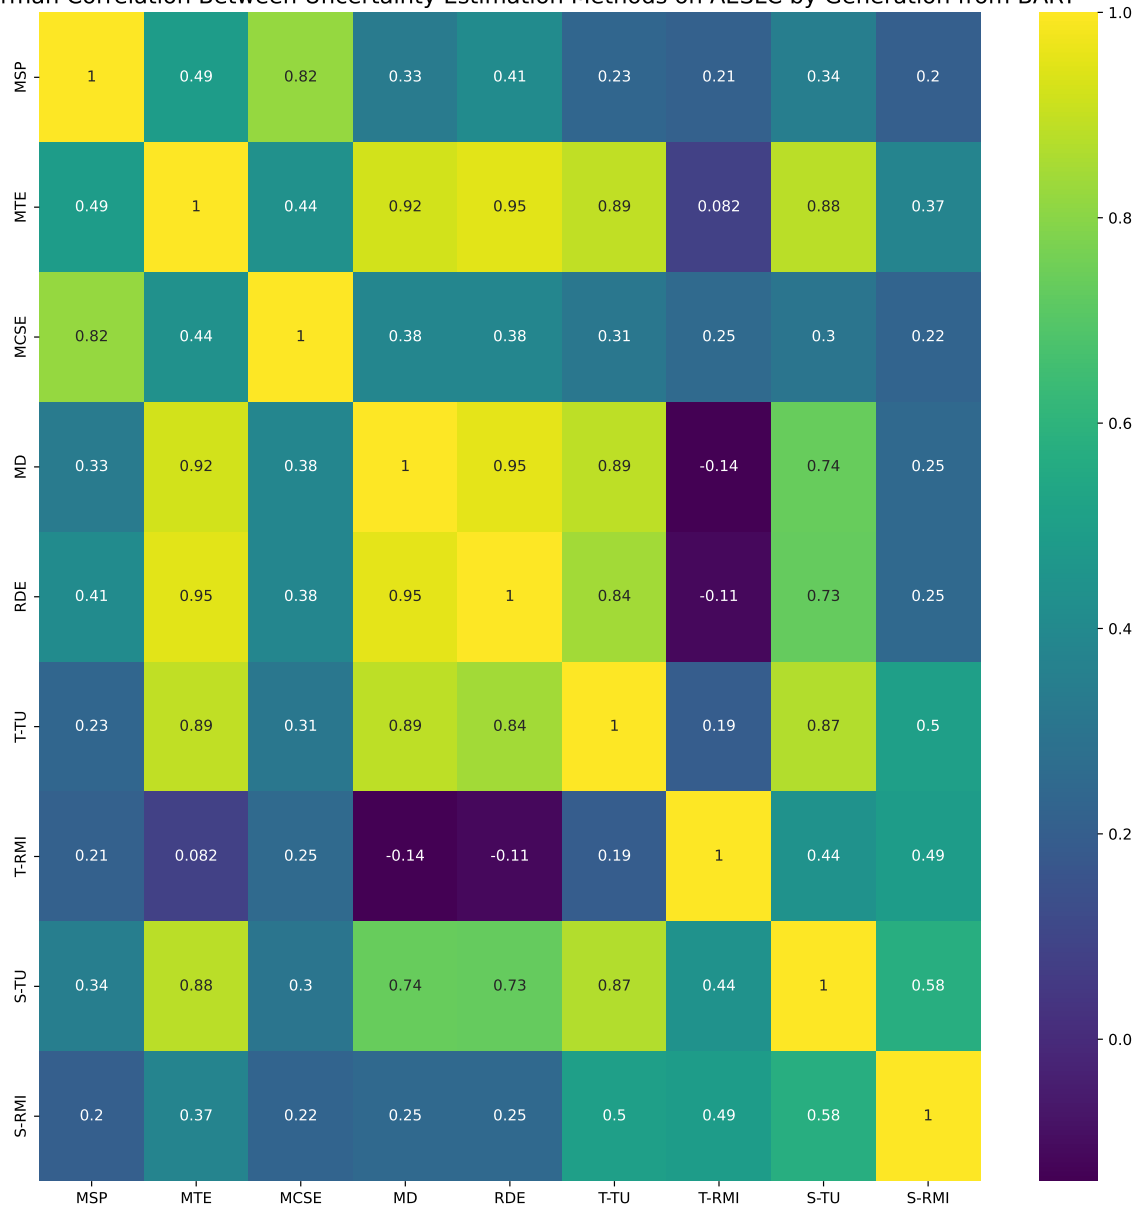


Figure 6: Diagram of Spearman correlation between uncertainty estimation methods on AESLC dataset from the view of NLG metrics used in Fig. 3. The generated summaries are from BART.

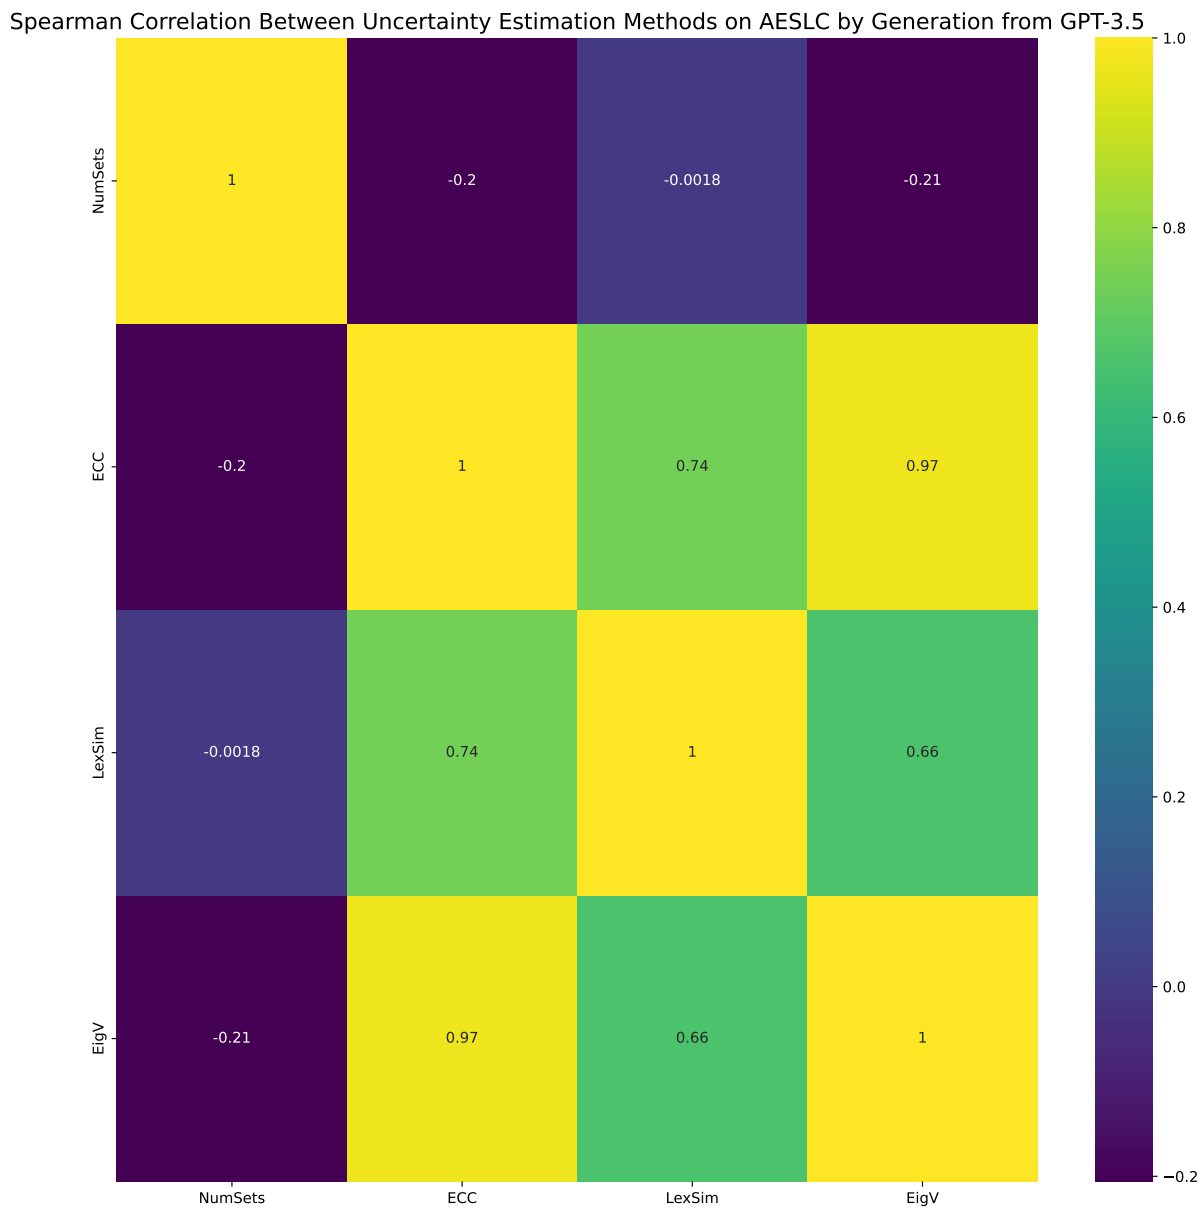


Figure 7: Diagram of Spearman correlation between uncertainty estimation methods on AESLC dataset from the view of NLG metrics used in Fig. 4. The generated summaries are from GPT-3.5.

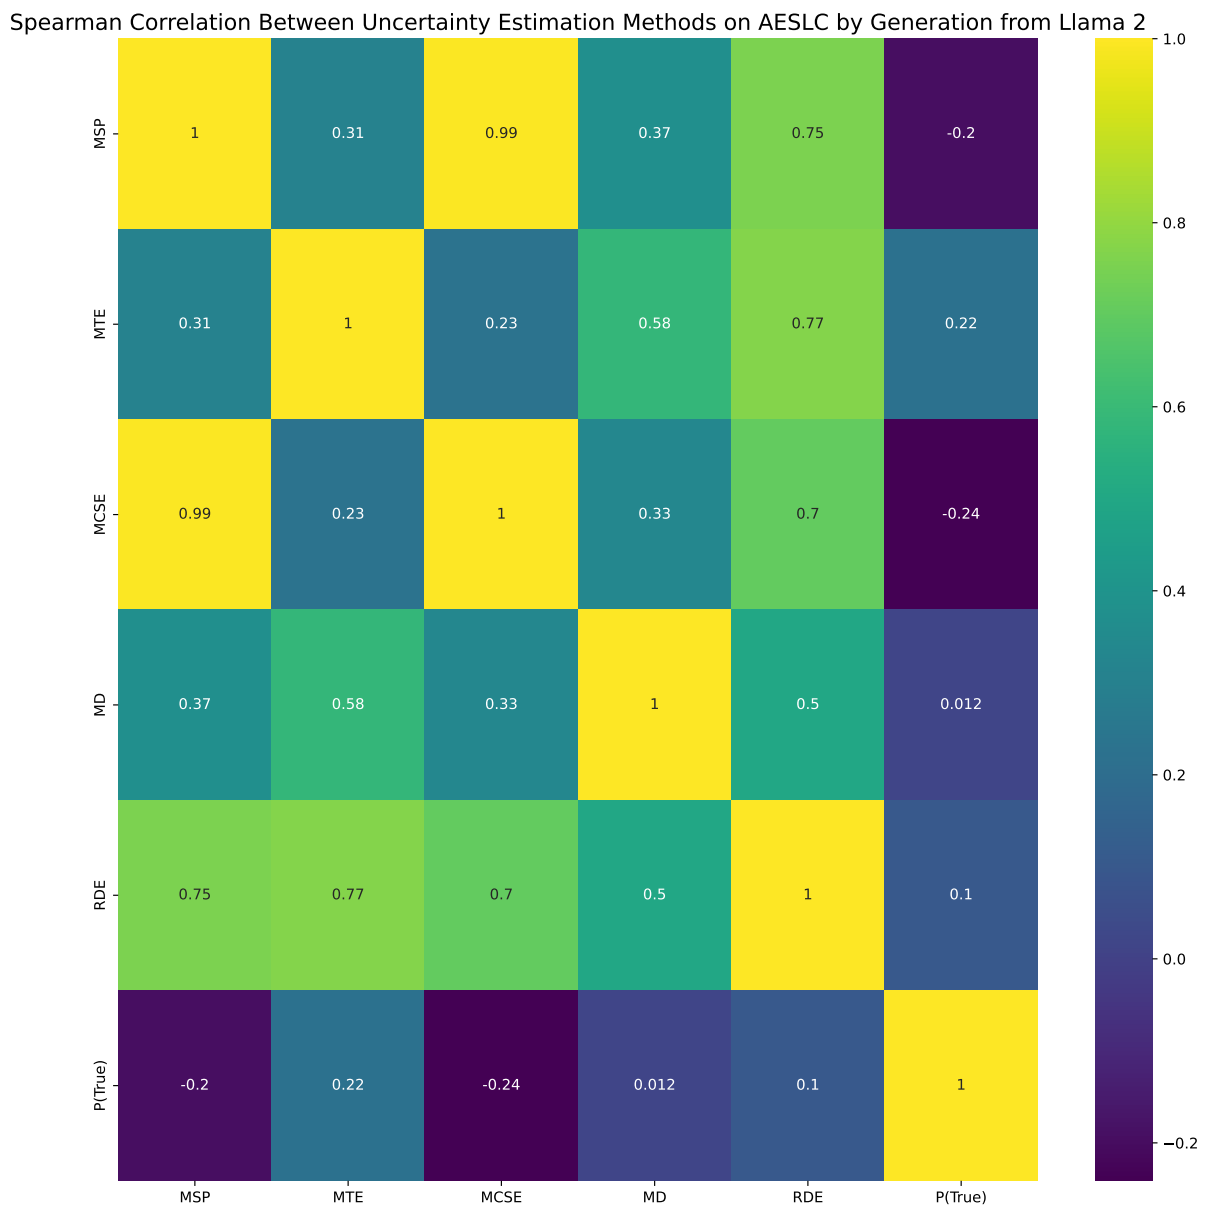


Figure 8: Diagram of Spearman correlation between uncertainty estimation methods on AESLC dataset from the view of NLG metrics used in Fig. 5. The generated summaries are from Llama 2.

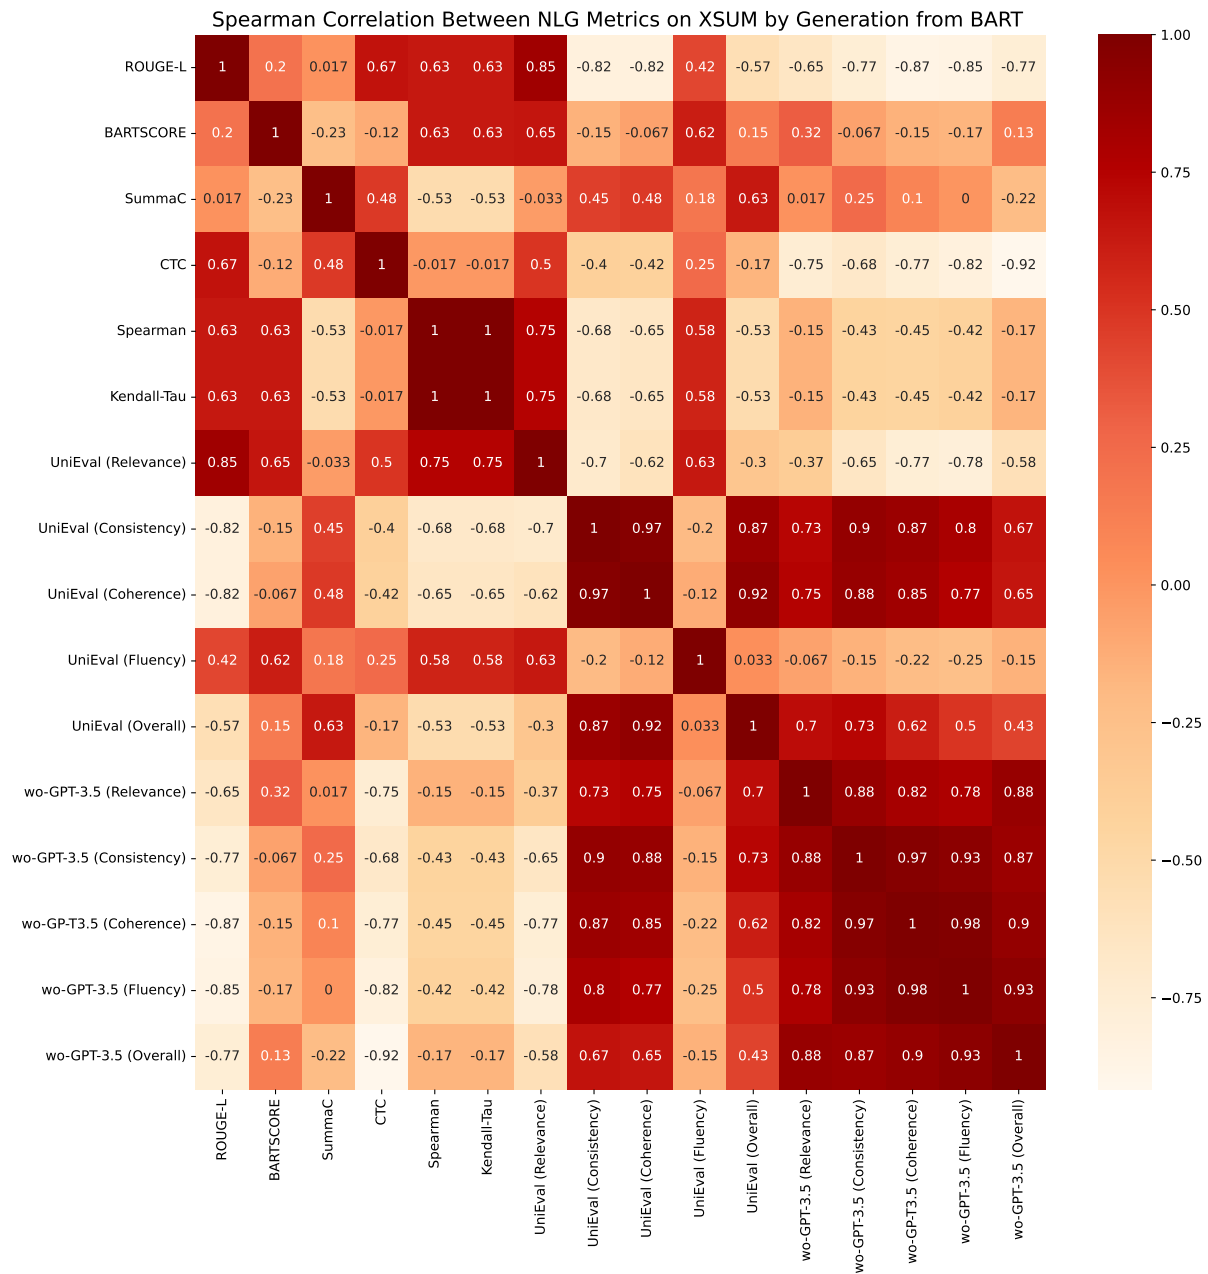


Figure 9: Diagram of Spearman correlation between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 12. The generated summaries are from BART. For the GPT-3.5-based NLG metrics, we only conduct wo-GPT-3.5 on the BART generation model setting.

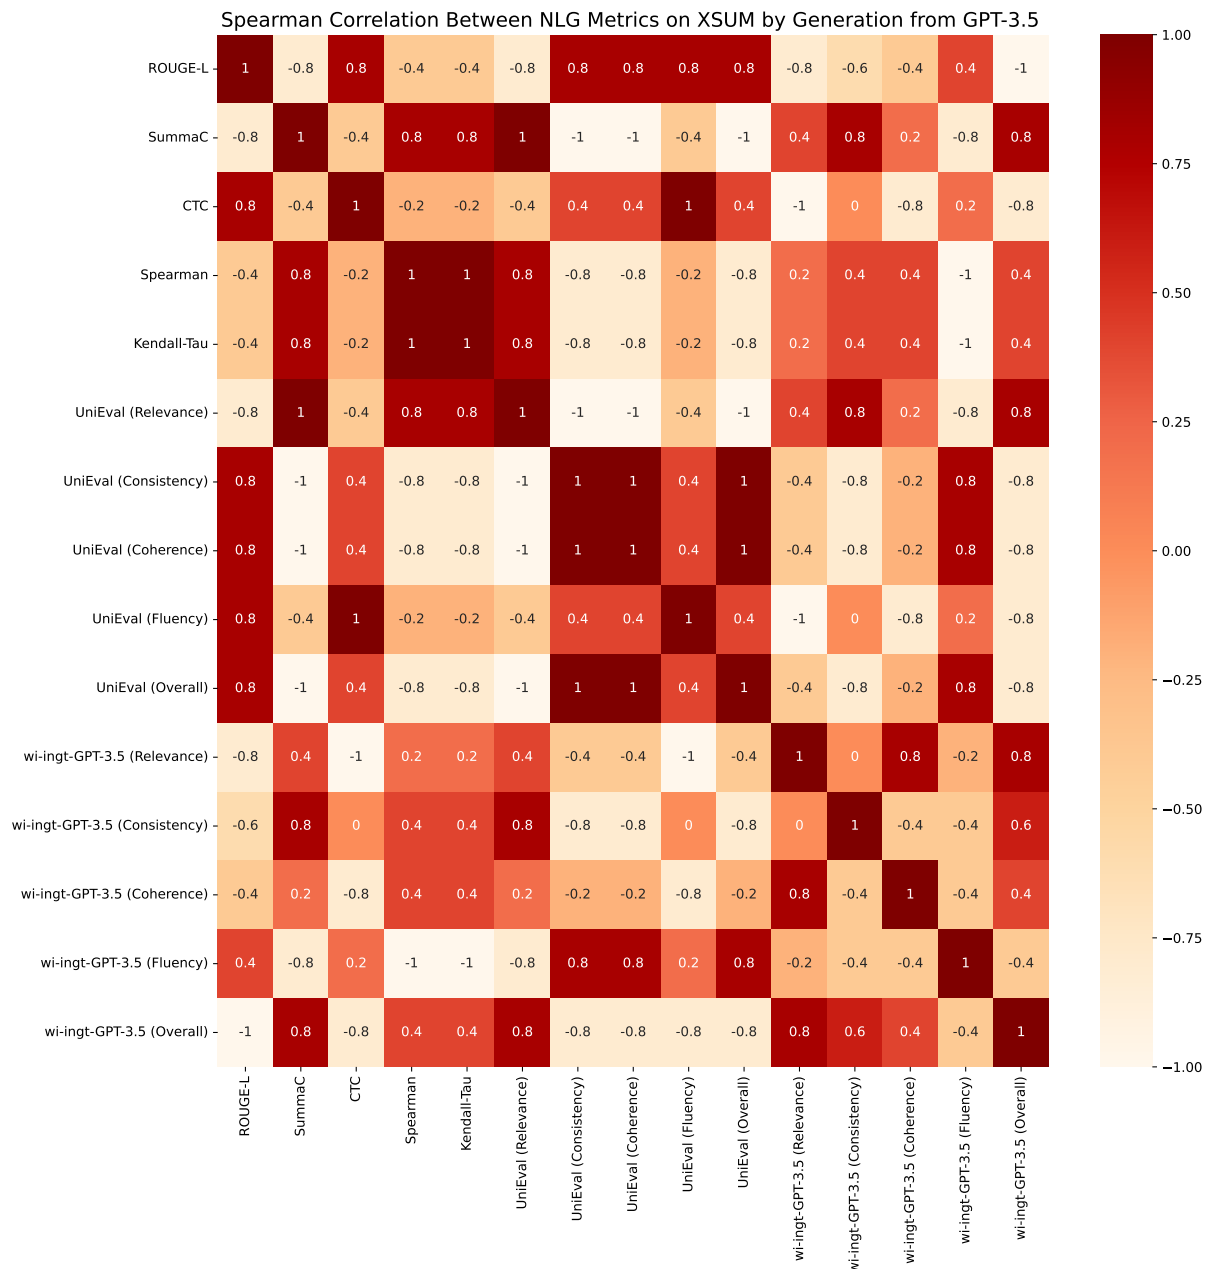


Figure 10: Diagram of Spearman correlation between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 13. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

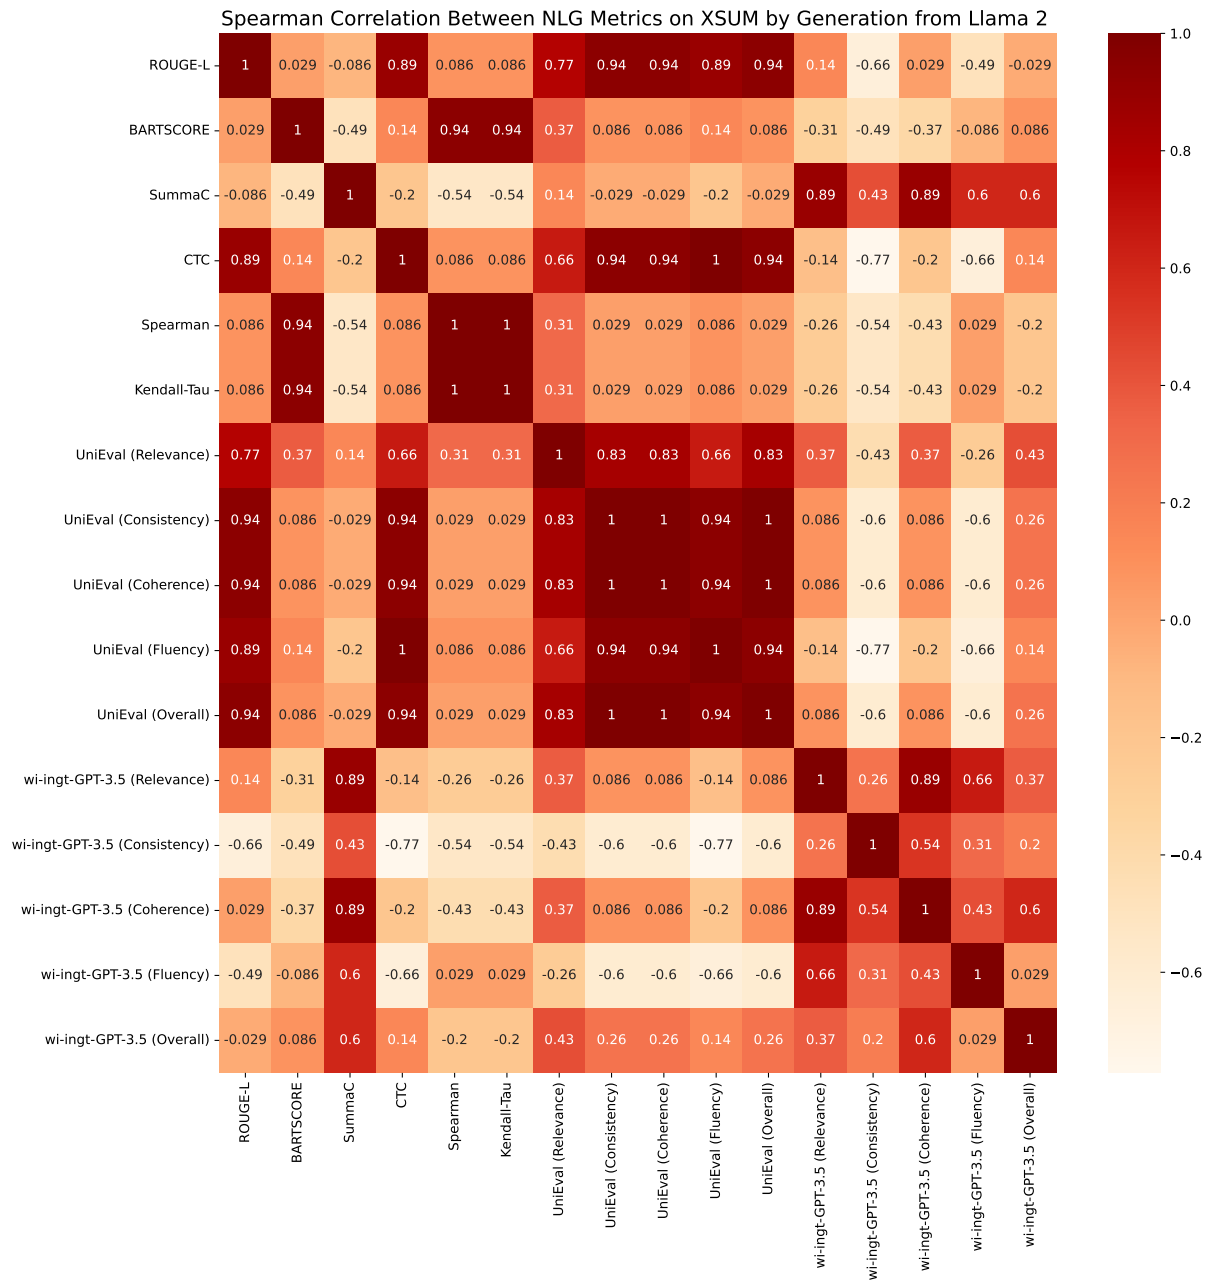


Figure 11: Diagram of Spearman correlation between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 14. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

Spearman Correlation Between Uncertainty Estimation Methods on XSUM by Generation from BART



Figure 12: Diagram of Spearman correlation between uncertainty estimation methods on XSUM dataset from the view of NLG metrics used in Fig. 9. The generated summaries are from BART.

Spearman Correlation Between Uncertainty Estimation Methods on XSUM by Generation from GPT-3.5

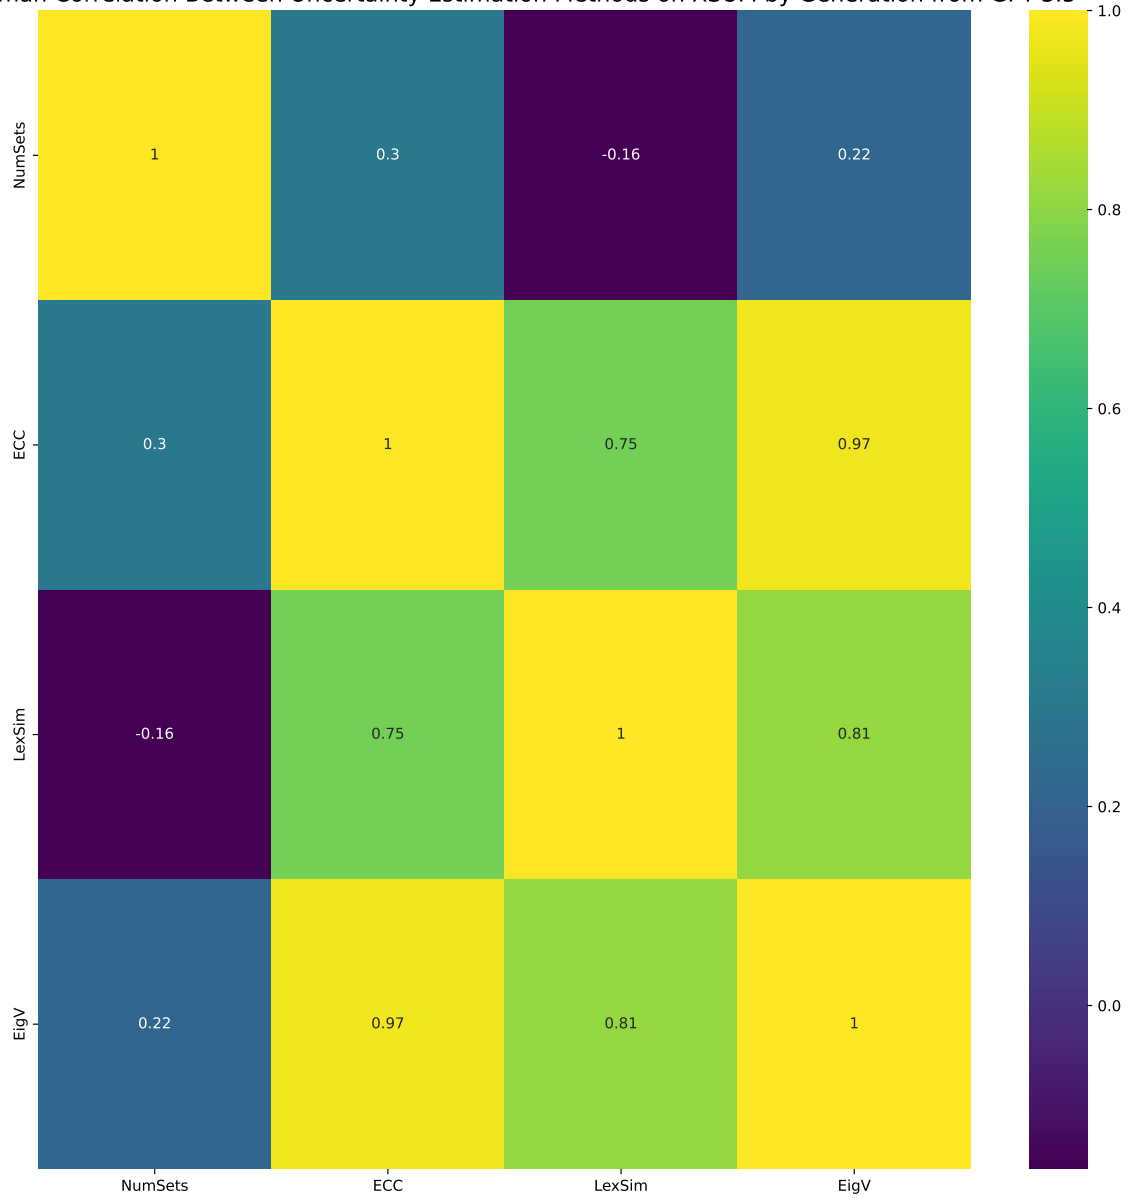


Figure 13: Diagram of Spearman correlation between uncertainty estimation methods on XSUM dataset from the view of NLG metrics used in Fig. 10. The generated summaries are from GPT-3.5.

Spearman Correlation Between Uncertainty Estimation Methods on XSUM by Generation from Llama 2

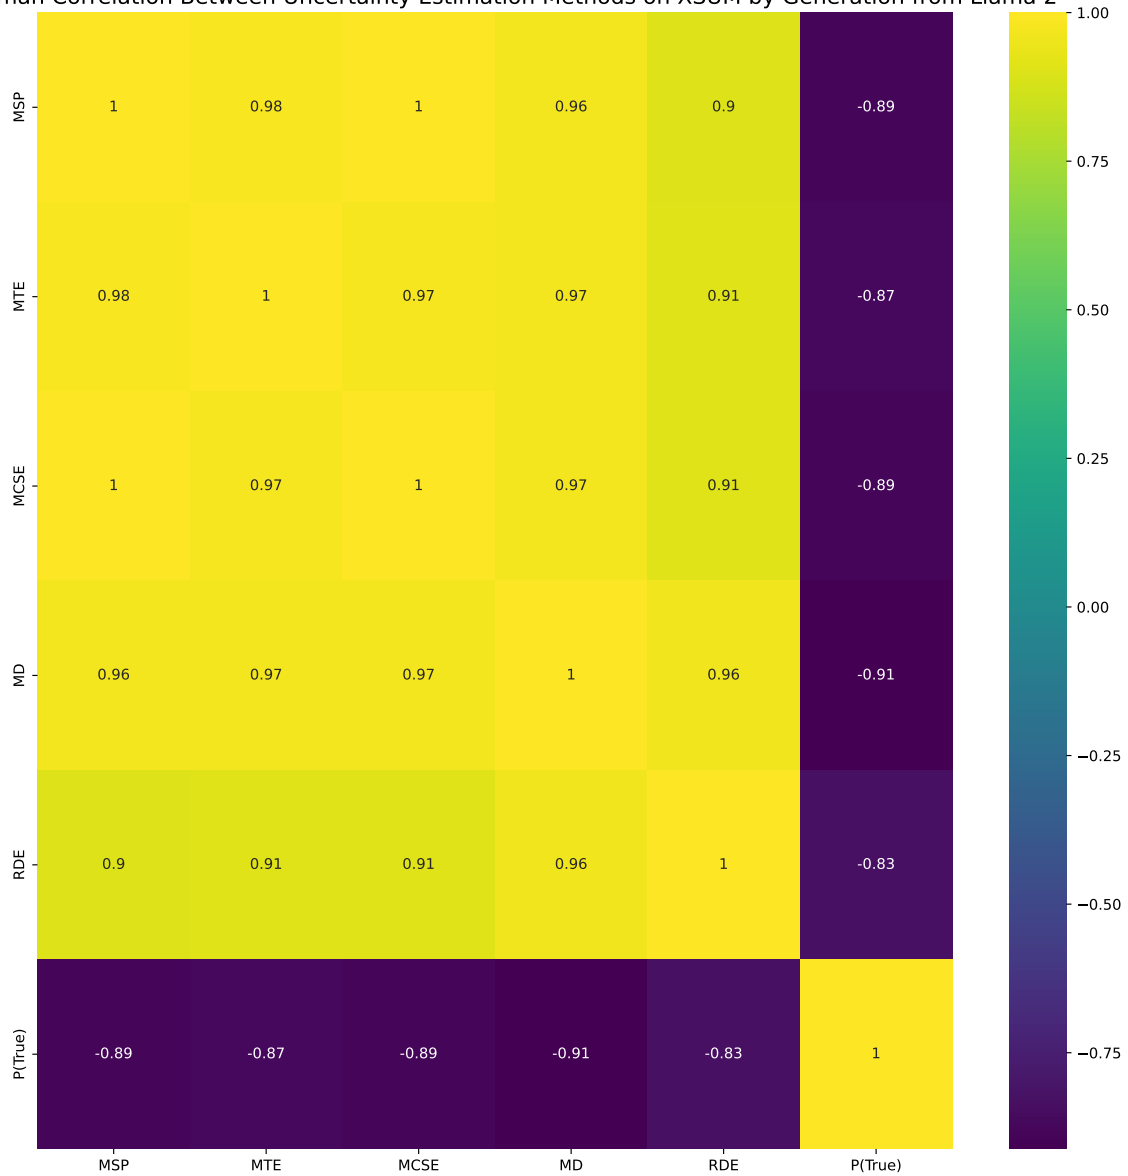


Figure 14: Diagram of Spearman correlation between uncertainty estimation methods on XSUM dataset from the view of NLG metrics used in Fig. 11. The generated summaries are from Llama 2.

NLG Metrics	MSP	MTE	MCSE	MD	RDE	T-TU	T-RMI	S-TU	S-RMI
ROUGE-L	-0.0978	-0.1359	-0.0587	-0.0375	-0.0996	-0.0511	-0.0126	-0.0733	-0.0041
BARTSCORE	-0.0558	0.0009	0.0752	0.1515	0.1324	0.0037	0.1136	-0.0457	-0.0040
SummaC	0.0760	0.0148	0.0696	-0.0004	-0.0542	-0.0432	0.0611	-0.0227	-0.0144
CTC	-0.0486	-0.3754	-0.0001	-0.2301	-0.3298	-0.1879	0.0176	-0.2027	-0.0095
Spearman	-0.2040	-0.1310	-0.1209	0.0028	-0.0424	-0.0579	-0.0500	-0.0895	-0.0048
Kendall-Tau	-0.2013	-0.1321	-0.1206	-0.0005	-0.0452	-0.0586	-0.0511	-0.0906	-0.0048
UniEval (Relevance)	-0.1475	-0.1092	-0.0562	-0.0126	-0.0966	-0.0858	0.0156	-0.1038	-0.0235
UniEval (Consistency)	0.7162	0.7064	0.6192	0.5025	0.5037	0.1936	0.1523	0.2284	0.0223
UniEval (Coherence)	0.6467	0.6474	0.5416	0.4556	0.4740	0.1358	0.1423	0.1790	0.0095
UniEval (Fluency)	0.0345	-0.1334	0.0338	0.0493	0.0469	-0.1619	0.1130	-0.2010	0.0375
UniEval (Overall)	0.2380	0.2421	0.2456	0.2264	0.1792	0.0076	0.0920	0.0113	-0.0033
wo-GPT-3.5 (Relevance)	0.0667	0.1933	0.0761	0.1921	0.1629	0.0292	-0.0060	0.0261	0.0081
wo-GPT-3.5 (Consistency)	0.1220	0.1545	0.0544	0.0828	0.0777	0.0235	-0.0382	0.0299	0.0051
wo-GP-T3.5 (Coherence)	0.1436	0.2104	0.0479	0.1176	0.1313	0.0454	-0.0740	0.0527	0.0046
wo-GPT-3.5 (Fluency)	0.1924	0.2679	0.0675	0.1547	0.1815	0.0769	-0.0528	0.0816	0.0102
wo-GPT-3.5 (Overall)	0.1137	0.2039	0.0448	0.1446	0.1546	0.0848	-0.1038	0.0681	0.0115
Col Mean	0.0997	0.1015	0.0950	0.1124	0.0860	-0.0029	0.0199	-0.0095	0.0025

Table 9: Main results of the relationship between the uncertainty estimation methods and NLG metrics on XSUM dataset using generation from BART.

are greater than 0.5. Thus, the reliability of evaluating uncertainty estimation models becomes a concern.

Additionally, some evaluations of uncertainty estimation models using different NLG metrics may result in different performance rankings. For instance, in Figure 10, it is observed that the correlation between ROUGE-L and wi-ingt-GPT-3.5 (overall) is -1, indicating a completely different ranking.

However, some evaluations of uncertainty estimation models using different NLG metrics could result in the same performance ranks. For instance, Figure 4 shows that the correlation between wi-ingt-GPT-3.5 (Fluency) and UniEval (Relevance) is 1, indicating identical ranks.

C.1.2 Analysis on NLG Relevance Dimension

Figures 15, 16, 17, 27, 28, and 29 show the evaluation of UE-TS models from the relevance-dimension NLG metrics.

Based on the BART generation model, Figure 15 shows strongly positive correlations among ROUGE-L, Spearman, and Kendall-Tau. However, the UniEval (Relevance) and wo-GPT-3.5 (Relevance) show weak correlations with all other NLG metrics in the relevance dimension.

Also, based on the BART generation model, Figure 27 shows strongly positive correlations among ROUGE-L, Spearman, Kendall-Tau, and UniEval (Relevance). However, the wo-GPT-3.5 (Relevance) shows weak correlation with all other NLG metrics in the relevance dimension.

Based on the GPT-3.5 generation model, Figure 16 shows strongly positive correlations among

Spearman, Kendall-Tau, UniEval (Relevance), and all GPT-3.5-based NLG metrics in the relevance dimension. However, Rouge-L shows weak correlation with all other NLG metrics in the relevance dimension. As for GPT-3.5-based NLG metrics, except for wi-in-GPT-3.5, we see that all other GPT-3.5-based models show identical performance ranks with each other. This implies that the relevance metric using GPT-3.5 is strongly related to the target text source.

Also, based on the GPT-3.5 generation model, Figure 28 shows a positive correlation among Spearman, Kendall-Tau, UniEval (Relevance), and all GPT-3.5-based NLG metrics in the relevance dimension, except for ROUGE-L.

Based on the Llama-2 generation model, in Figure 17, UniEval (Relevance) shows a strongly positive correlation with all other NLG metrics in the relevance dimension. This is because each element in the row of UniEval (Relevance) is greater than 0.5. Similarly, wo-GPT-3.5 (Relevance) also shows a strongly positive correlation with all other NLG metrics in the relevance dimension except for ROUGE-L. Additionally, we found that wi-in-GPT-3.5 shows weak correlation with other GPT-3.5-based NLG metrics. This weak correlation indicates that if we only provide the ground-truth summaries for the relevance evaluation, the relevance metrics may ignore many details.

Also, based on the Llama-2 generation model, Figure 29 shows that both UniEval (Relevance) and wo-GPT-3.5 (Relevance) show a positive correlation with other methods. The other metrics do not exhibit consistent positive or negative correlations

with each other.

Besides, by comparing wo-GPT-3.5 (relevance) and wi-GPT-3.5 (relevance), where the only difference lies in whether GPT-3.5 is given the concept of relevance or not, their high positive correlation indicates that GPT-3.5 incorporates the concept of relevance. Consequently, the difference in providing the relevance concept to GPT-3.5 is not apparent.

For the evaluation of UE-TS models using relevance-dimension NLG metrics, we have the following conclusions.

1. Generation models of the same type across different datasets could result in similar correlations among various methods.
2. Spearman and Kendall-Tau usually exhibit positive correlations. Therefore, in future experiments, choosing one of them is sufficient.
3. When employing LLMs as generation models, UniEval (Relevance) and wo-GPT-3.5 (Relevance) tend to exhibit positive correlations with most other NLG metrics. Therefore, one of them could serve as a representative NLG metric.
4. When utilizing LLMs as a type of relevance NLG metric, the choice of target text source can greatly impact the final conclusion. Specifically, using ground-truth summaries versus using input text as the target text source can result in different performance rankings.
5. When using LLMs as a type of relevance NLG metric, if both ground-truth summaries and input text are employed together as the target text source, the ground-truth summaries will dominate the metric results.
6. GPT-3.5 knows the concept of relevance. Consequently, the difference in providing the relevance concept to GPT-3.5 is not apparent.

C.1.3 Analysis on NLG Consistency Dimension

Figures 18, 19, 20, 30, 31, and 32 display the evaluation results of UE-TS models based on the consistency-dimension NLG metrics.

Specifically, based on BART generation, Figure 18 illustrates that UniEval (Consistency) and wo-GPT-3.5 (Consistency) exhibit a strongly positive correlation. SummaC and CTC show a positive correlation as well. Additionally, wo-GPT-3.5 (Consistency) displays a strongly positive correlation with BARTScore. However, there is no or even a negative correlation between the group of

UniEval (Consistency) and wo-GPT-3.5 (Consistency) and the group of SummaC and CTC.

Also, based on BART generation, Figure 30 illustrates that UniEval (Consistency) and wo-GPT-3.5 (Consistency) exhibit a strongly positive correlation. SummaC and CTC show a positive correlation as well. However, SummaC displays a positive correlation with UniEval (Consistency) and wo-GPT-3.5 (Consistency), whereas CTC does not.

Based on GPT-3.5 generation, in Figure 19, positive correlations are found in a group consisting of CTC, UniEval (Consistency), and wi-in-GPT-3.5 (Consistency). Additionally, positive correlations are observed in another group comprising SummaC, wo-GPT-3.5 (Consistency), wi-gt-GPT-3.5 (Consistency), and wi-ingt-GPT-3.5 (Consistency). However, these two groups display negative correlations. Thus, different metrics in the NLG consistency dimension can lead to different evaluation performance rankings of UE-TS models.

Also, based on GPT-3.5 generation, Figure 31 depicts positive correlations between CTC and UniEval (Consistency). Additional positive correlations are observed in a group comprising SummaC and all GPT-3.5-based NLG metrics. The only exception is that wi-ingt-GPT-3.5 (Consistency) and wo-ingt-GPT-3.5 (Consistency) exhibit negative correlations.

Based on Llama 2 generation, Figure 20 illustrates positive correlations in a group comprising BARTScore and all GPT-3.5-based NLG metrics. Additionally, positive correlations are found in a group consisting of CTC and UniEval (Consistency). However, negative correlations are observed for both of these groups. Regarding SummaC, it shows positive correlations with most of the first group except for wi-in-GPT-3.5 (Consistency) and negative correlations with the second group.

Also, based on Llama 2 generation, Figure 31 shows strong positive correlations in a group comprising CTC, UniEval (Consistency), and wo-GPT-3.5 (Consistency). Another set of strong positive correlations is found in a group including wi-gt-GPT-3.5, wi-in-GPT-3.5, and wi-ingt-GPT-3.5 (Consistency). However, these two groups exhibit negative correlations with each other. BARTSCORE and SummaC display negative correlations.

According to the above analysis, we can draw the following conclusions regarding the evaluation of UE-TS models using consistency-dimension NLG

metrics.

1. Different metrics within the NLG consistency dimension can result in varying evaluation performance rankings of UE-TS models.
2. CTC tends to exhibit a positive correlation with SummaC or UniEval (Consistency). Thus, when faced with scenarios where we must choose between CTC, SummaC, or UniEval, we can opt for CTC due to its positive correlations with the other two in most cases.
3. For GPT-3.5-based NLG metrics, differences in the target text source could lead to discrepancies in some cases. Using ground-truth summaries as the target text source will have a more dominated impact compared to using input text as the target text source.
4. For GPT-3.5-based NLG metrics, GPT-3.5 might not fully comprehend the concept of consistency, as indicated by the negative correlation between wo-GPT-3.5 and wi-GPT-3.5 in Figure 34. However, the impact of understanding the concept of consistency is not as pronounced as the impact of using different target text sources.
5. For GPT-3.5-based NLG metrics, it is recommended to use wi-in-GPT-3.5 (Consistency) along with one of wo-gt-GPT-3.5, wi-gt-GPT-3.5 (Consistency), or wi-ingt-GPT-3.5.

C.1.4 Analysis on NLG Coherence Dimension

Figures 21, 22, 33 and 34 display the evaluation results of UE-TS models based on coherence-dimension NLG metrics. We do not illustrate correlations for BART generation models in the coherence dimension. This is because we only utilize UniEval (coherence) and wo-GPT-3.5 (coherence) for BART generations, and these two metrics may not adequately represent the correlation.

Based on GPT-3.5 generation, Figure 21 indicates that all GPT-3.5-based NLG metrics exhibit a strongly positive correlation with each other in terms of evaluating UE-TS models. However, UniEval (Coherence) demonstrates weak or even negative correlation with these GPT-3.5-based NLG metrics.

Also, based on GPT-3.5, Figure 33 demonstrates that a group comprising wo-GPT-3.5 (Coherence) and wi-ingt-GPT-3.5 (Coherence) exhibits a strongly positive correlation. Another strong positive correlation is observed in a group consisting of wi-gt-GPT-3.5 (Coherence) and wi-in-GPT-3.5 (Coherence). However, these two groups show no or even negative correlation with each other. Addi-

tionally, UniEval (Coherence) demonstrates weak or even negative correlation with these GPT-3.5-based NLG metrics.

Based on Llama 2 generation, Figure 22 illustrates that all GPT-3.5-based NLG metrics display positive correlations. Among them, wi-gt-GPT-3.5 shows a strongly positive correlation with other GPT-3.5-based NLG metrics except for wi-in-GPT-3.5 (Coherence). However, UniEval (Coherence) demonstrates weak or even negative correlation with these GPT-3.5-based NLG metrics.

Also, based on Llama 2 generation, Figure 34 indicates that all GPT-3.5-based NLG metrics exhibit positive correlations. wo-GPT-3.5 (Coherence) and wi-in-GPT-3.5 (Coherence) demonstrate a strongly positive correlation of 1. However, wi-gt-GPT-3.5 shows a relatively weak positive correlation with the other metrics. Regarding UniEval (Coherence), it displays a strongly positive correlation with wi-gt-GPT-3.5 (Coherence) but shows no or even negative correlation with the other three GPT-3.5-based NLG metrics.

According to the above analysis, we can obtain the following conclusions regarding the evaluation of UE-TS models using coherence-dimension NLG metrics:

1. It is difficult to determine which GPT-3.5-based metric is better in the coherence dimension. However, based on the strongly positive correlation, either wi-gt-GPT-3.5 or wi-in-GPT-3.5 could be a good choice.
2. UniEval (Coherence) exhibits weak or negative correlation with most of the GPT-3.5-based metrics. Therefore, UniEval (Coherence) could serve as a supplement to either wi-gt-GPT-3.5 or wi-in-GPT-3.5.
3. Based on Figures 33 and 34 on XSUM datasets, GPT-3.5 exhibits significant divergence between wo-GPT-3.5 and wi-gt-GPT-3.5. This suggests that GPT-3.5 itself might not fully grasp the concept of coherence, and providing this concept could improve evaluation.

C.1.5 Analysis on NLG Fluency Dimension

Figures 23, 24, 35, and 36 display the evaluation results of UE-TS models based on fluency-dimension NLG metrics. We do not illustrate correlations for BART generation models in the fluency dimension. This is because we only utilize UniEval (fluency) and wo-GPT-3.5 (fluency) for BART generations, and these two metrics may not adequately represent correlation.

Using GPT-3.5 as a generation model, Figure 23 indicates that wi-in-GPT-3.5 (Fluency) has a negative correlation with wo-GPT-3.5, wi-gt-GPT-3.5, and wi-ingt-GPT-3.5 (Fluency). UniEval (Fluency) demonstrates weak or negative correlation with these GPT-3.5-based NLG metrics.

Also, using GPT-3.5 for generation, Figure 35 illustrates that all the GPT-3.5-based NLG metrics except wi-ingt-GPT-3.5 (Fluency) exhibit strongly positive correlations. In contrast, these three GPT-3.5-based NLG metrics show negative correlation to UniEval (Fluency) and wi-ingt-GPT-3.5 (Fluency).

Using Llama 2 as a generation model, Figure 24 demonstrates that all five metrics have positive correlations. Among them, wi-gt-GPT-3.5, wi-in-GPT-3.5, and wi-ingt-GPT-3.5 (Fluency) exhibit strongly positive correlations.

Figure 36 illustrates that UniEval (Fluency) has strongly positive correlations with wo-GPT-3.5 (Fluency) and wi-gt-GPT-3.5 (Fluency). However, it exhibits strongly negative correlation with wi-in-GPT-3.5 (Fluency) and wi-ingt-GPT-3.5 (Fluency). Among the GPT-3.5-based NLG metrics, only wi-gt-GPT-3.5 (Fluency) and wi-in-GPT-3.5 (Fluency) show positive correlations, while the other correlations are weak or negative.

Based on the above findings, we can conclude the following observations regarding the fluency dimension.

1. It is challenging to determine which NLG metrics consistently exhibit positive correlations with others. However, based on the positive correlations, we recommend wi-gt-GPT-3.5 (Fluency) and wi-ingt-GPT-3.5 (Fluency) due to their correlation patterns.
2. Because of the negative correlations between UniEval (Fluency) and GPT-3.5-based NLG metrics, we also recommend using UniEval (Fluency) as a supplement to a GPT-3.5-based NLG metric.
3. In the context of GPT-3.5-based NLG metrics, the choice of target text source influences the performance ranking of UE-TS models. Furthermore, employing ground-truth summaries as the target text source exerts a more significant impact compared to using input text.
4. The same generation methods cannot achieve similar correlations for fluency NLG dimension. Consequently, correlations within fluency dimensions are more closely tied to the dataset.
5. Due to the positive or negative correlations between wo-GPT-3.5 (Fluency) and wi-GPT-3.5 (Flu-

ency), it is difficult to determine whether GPT-3.5 itself understands the concept of fluency.

C.1.6 Analysis on NLG Overall Dimension

Figures 25, 26, 37, and 38 display the evaluation results of UE-TS models based on overall-dimension NLG metrics. Correlations for BART generation models in the overall dimension are not illustrated. This is because we only utilize UniEval (overall) and wo-GPT-3.5 (overall) for BART generations, and these two metrics may not adequately represent correlation.

Utilizing GPT-3.5 as the generation model, Figure 25 shows that all GPT-3.5-based NLG metrics have identical performance rankings for UE-TS models. Additionally, these GPT-3.5-based NLG metrics exhibit a strongly positive correlation with UniEval (Overall).

Also, utilizing GPT-3.5 as the generation model, Figure 25 illustrates that wo-GPT-3.5 (Overall) has positive correlations with the other three GPT-3.5-based NLG metrics. Furthermore, strongly positive correlations are found between wo-GPT-3.5 (Overall) and wi-gt-GPT-3.5 (Overall). Another strongly positive correlation is observed between wi-in-GPT-3.5 (Overall) and wi-ingt-GPT-3.5 (Overall).

Utilizing Llama 2 as the generation model, Figure 26 indicates that all GPT-3.5-based NLG metrics have strongly positive correlations. However, UniEval (Overall) exhibits negative correlations with these GPT-3.5-based NLG metrics.

Also, utilizing Llama 2 as the generation model, Figure 38 shows that almost all five metrics have positive correlations, with the exceptions being that wi-in-GPT-3.5 (Overall) has a negative correlation with UniEval (Overall) and wo-GPT-3.5 (Overall).

Based on the above findings, we can conclude the following regarding the overall dimension.

1. When it comes to the GPT-3.5-based NLG metrics, given their positive correlations with other metrics, it is recommended to use either wi-in-GPT-3.5 (Overall) or wi-gt-GPT-3.5 (Overall).
2. The UniEval (Overall) exhibits high correlations with other GPT-3.5-based methods in most cases and negative correlations in a few instances. It could serve as an optional supplementary tool for GPT-3.5-based NLG metrics.
3. Due to the positive correlations between wo-GPT-3.5 and wi-GPT-3.5 NLG metrics in most cases, it suggests that GPT-3.5 itself grasps the concept of overall.
4. Unlike previous NLG metric dimensions, it is

difficult to determine which one dominates the evaluated performance ranks between using ground-truth summaries and using input text.

C.2 Experimental Results Involving Human Annotations

Besides the aforementioned experiments on AESLC and XSUM, we are also interested in conducting human-related experiments. We utilize the TofuEval dataset (Tang et al., 2024), which provides human annotations across seven dimensions as will be introduced later.

These experiments consist primarily of three types:

- UE-HUM experiments: These experiments involve using the human annotations as the PR_{oracle} and employing uncertainty estimation metric scores as the $PR_{uncertainty}$ in Eq. 1.
- NLG-HUM experiments: In these experiments, the human annotations serve as the PR_{oracle} while the negative NLG metric scores are utilized as the $PR_{uncertainty}$ in Eq. 1. The negative NLG metric score is obtained by multiplying the original NLG metric scores by -1. We use a negative NLG metric score because $PR_{uncertainty}$ expects a larger value for lower quality, and the NLG metric scores used in our work are positively correlated with quality.
- UE-NLG experiments: These experiments utilize NLG metric scores as the PR_{oracle} and also employ uncertainty estimation metric scores as the $PR_{uncertainty}$ in Eq. 1.

The first two types are specifically tailored to the human annotations provided in the TofuEval dataset. The third type follows the same setup as the experiments conducted on the AESLC and XSUM datasets.

C.2.1 Seven Dimensions of Human Annotations.

The seven dimensions of human annotation used in TofuEval correspond to seven error types.

- Extrinsic Information (EI): The summary sentence contains new information not grounded in the source document.

- Mis-Referencing (MR): A property or an event in the summary sentence can be found in the source material, but are associated with the wrong entity.
- Stating Opinion As Fact (SOAF): The summary sentence entails a proposition that’s mentioned in the source material not as a fact, but as someone’s opinion.
- Reasoning Error (RE): The summary sentence makes one or more wrong inferences from information in the source document.
- Tense/Modality Error (TME): The tense or modal (e.g. can, may, must) used in the summary sentence does not match the tense/modality of the source document.
- Contradiction (CO): The summary sentence contradicts the source material.
- Nuanced Meaning Shift (NMS): The summary sentence twists information from the source material in a subtle way.

For the human annotation score for an error type T , we calculate the PR_{oracle} as $1 - Ratio(T)$, where $Ratio(T)$ represents the ratio of the incorrect word count for error type T to the total word count in a summary.

C.2.2 UE-HUM Experimental Results

Table 11 and Figure 40 show the UE-HUM experimental results. From Table 11, we can draw the following conclusions.

1. All uncertainty estimation methods used in our work have positive PRRs for EI.
2. There is no uncertainty estimation method that has positive PRRs in all dimensions.
3. The LexSim dimension achieves the highest performance among the seven human annotation dimensions, exhibiting the largest mean value across rows.

From Figure 40, we observe that ECC and EigV perform similarly from the perspective of human annotation.

C.2.3 NLG-HUM Experimental Results

Table 4 and Figure 41 show the NLG-HUM experimental results.

Based on Table 4, the following conclusions can be drawn.

1. It is challenging for any NLG metric to achieve a positive PRR for all human annotation dimensions

Spearman Correlation Between NLG Metrics on AESLC by Generation from BART (Relevance)

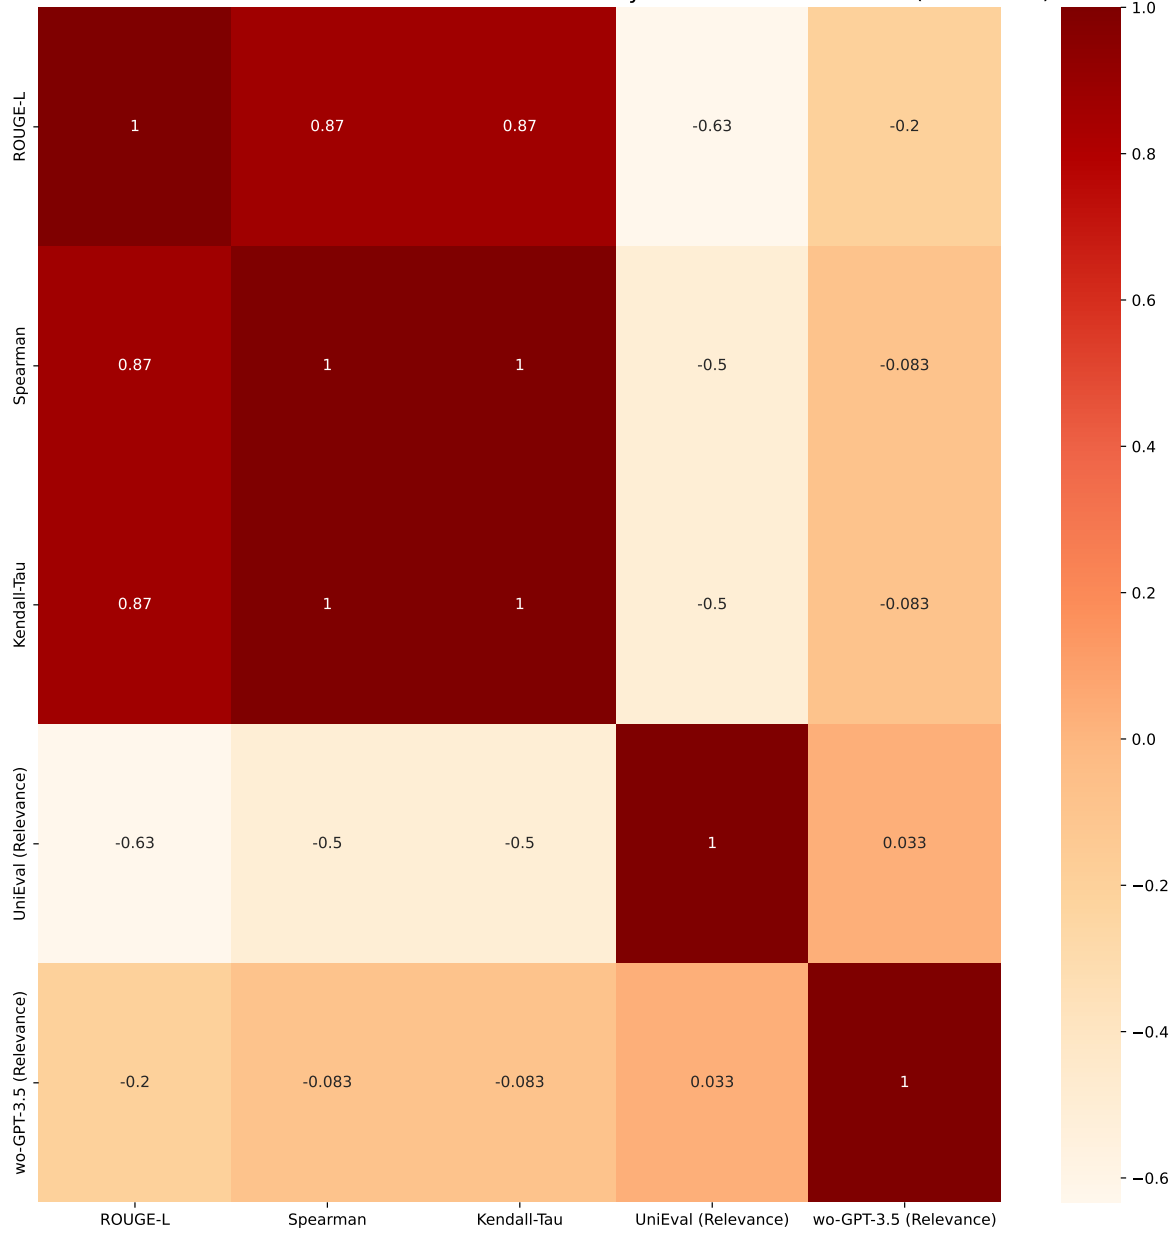


Figure 15: Diagram of Spearman correlation in terms of relevance between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 6. The generated summaries are from BART. For the GPT-3.5-based NLG metrics, we only conduct wo-GPT-3.5 on the BART generation model setting.

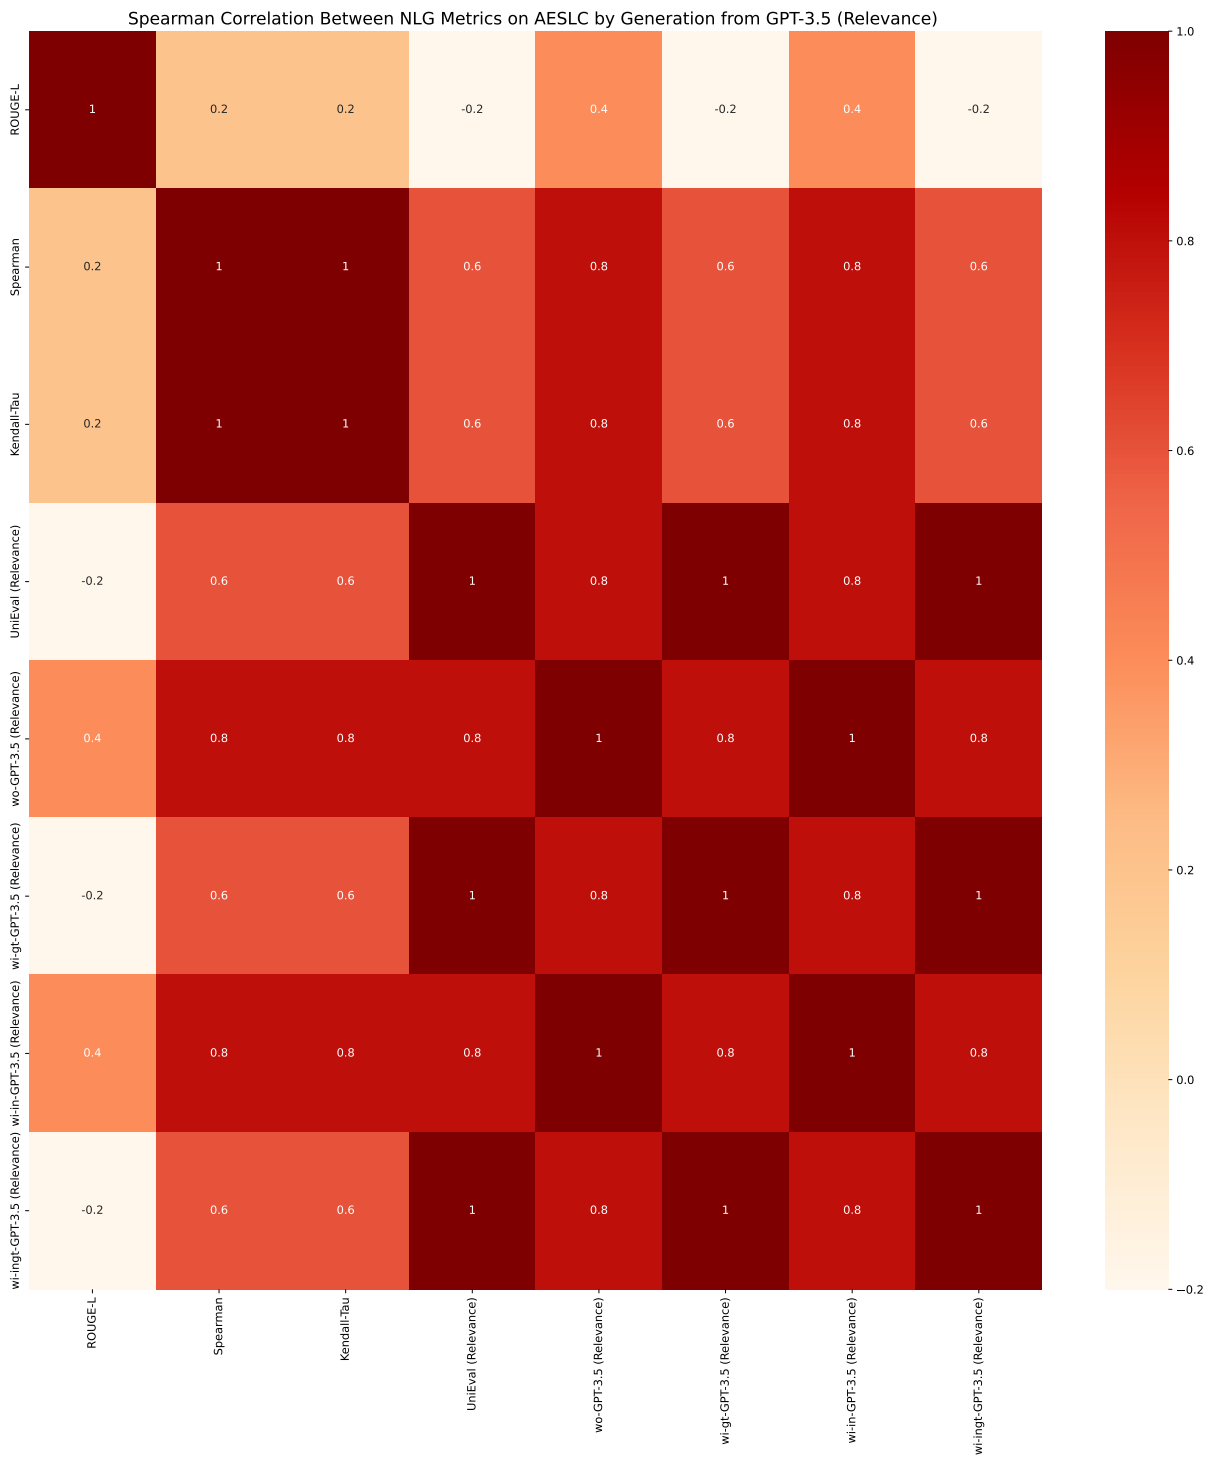


Figure 16: Diagram of Spearman correlation in terms of relevance between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 7. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

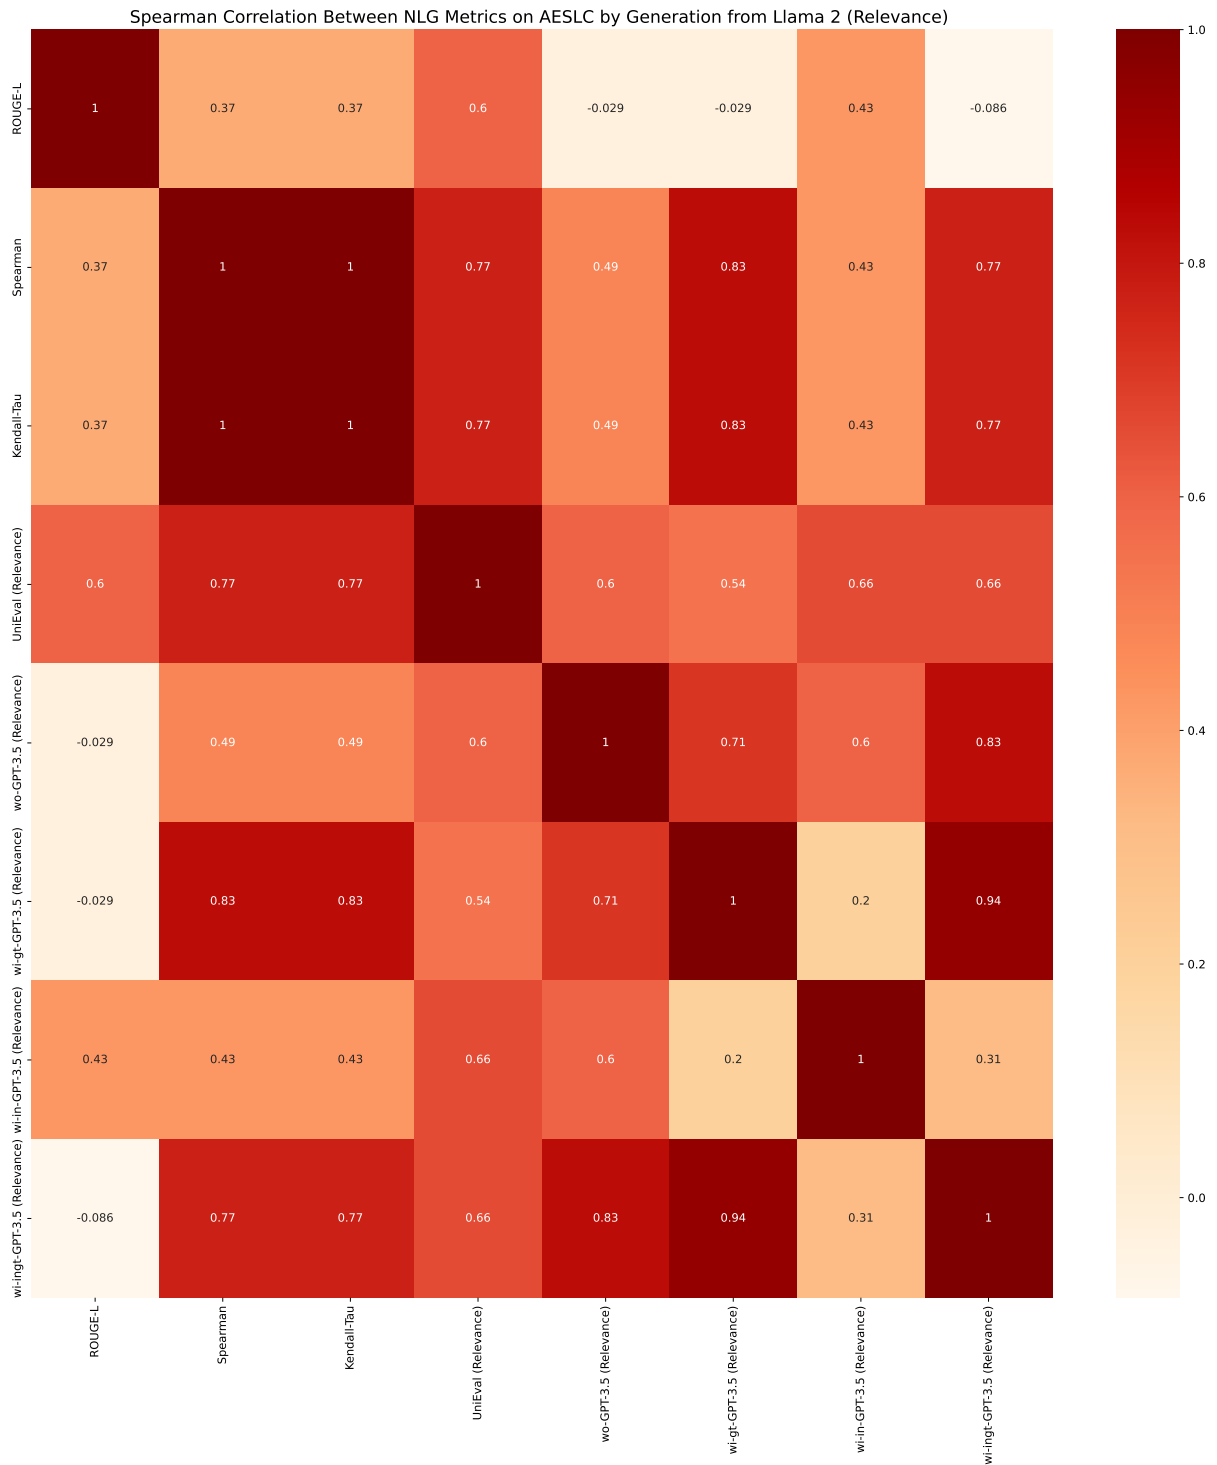


Figure 17: Diagram of Spearman correlation in terms of relevance between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 8. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

Spearman Correlation Between NLG Metrics on AESLC by Generation from BART (Consistency)

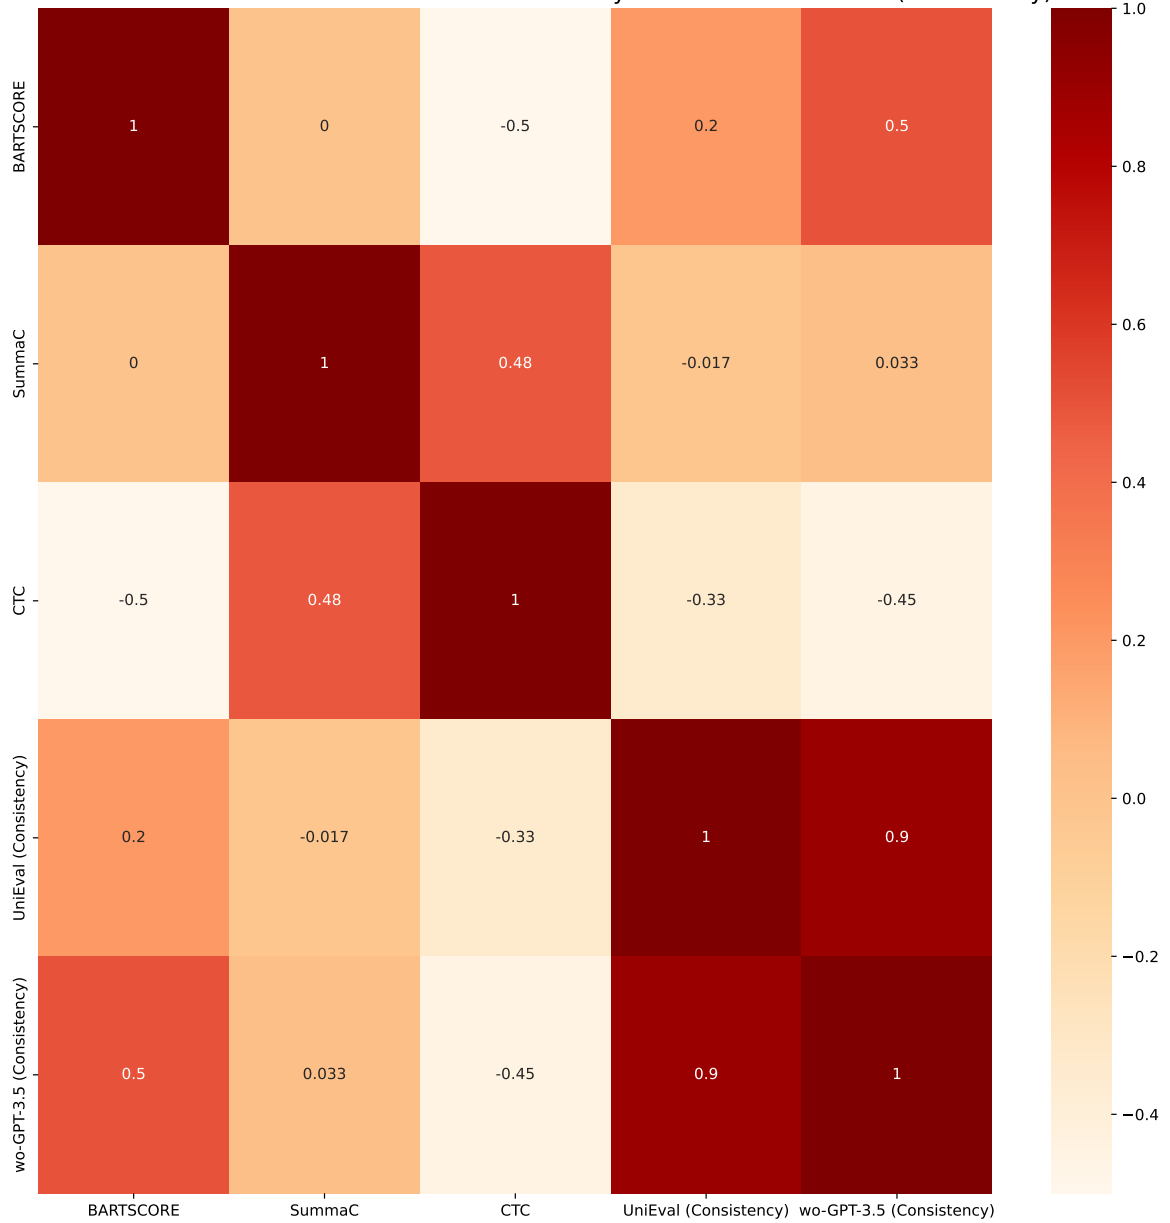


Figure 18: Diagram of Spearman correlation in terms of consistency between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 6. The generated summaries are from BART. For the GPT-3.5-based NLG metrics, we only conduct wo-GPT-3.5 on the BART generation model setting.

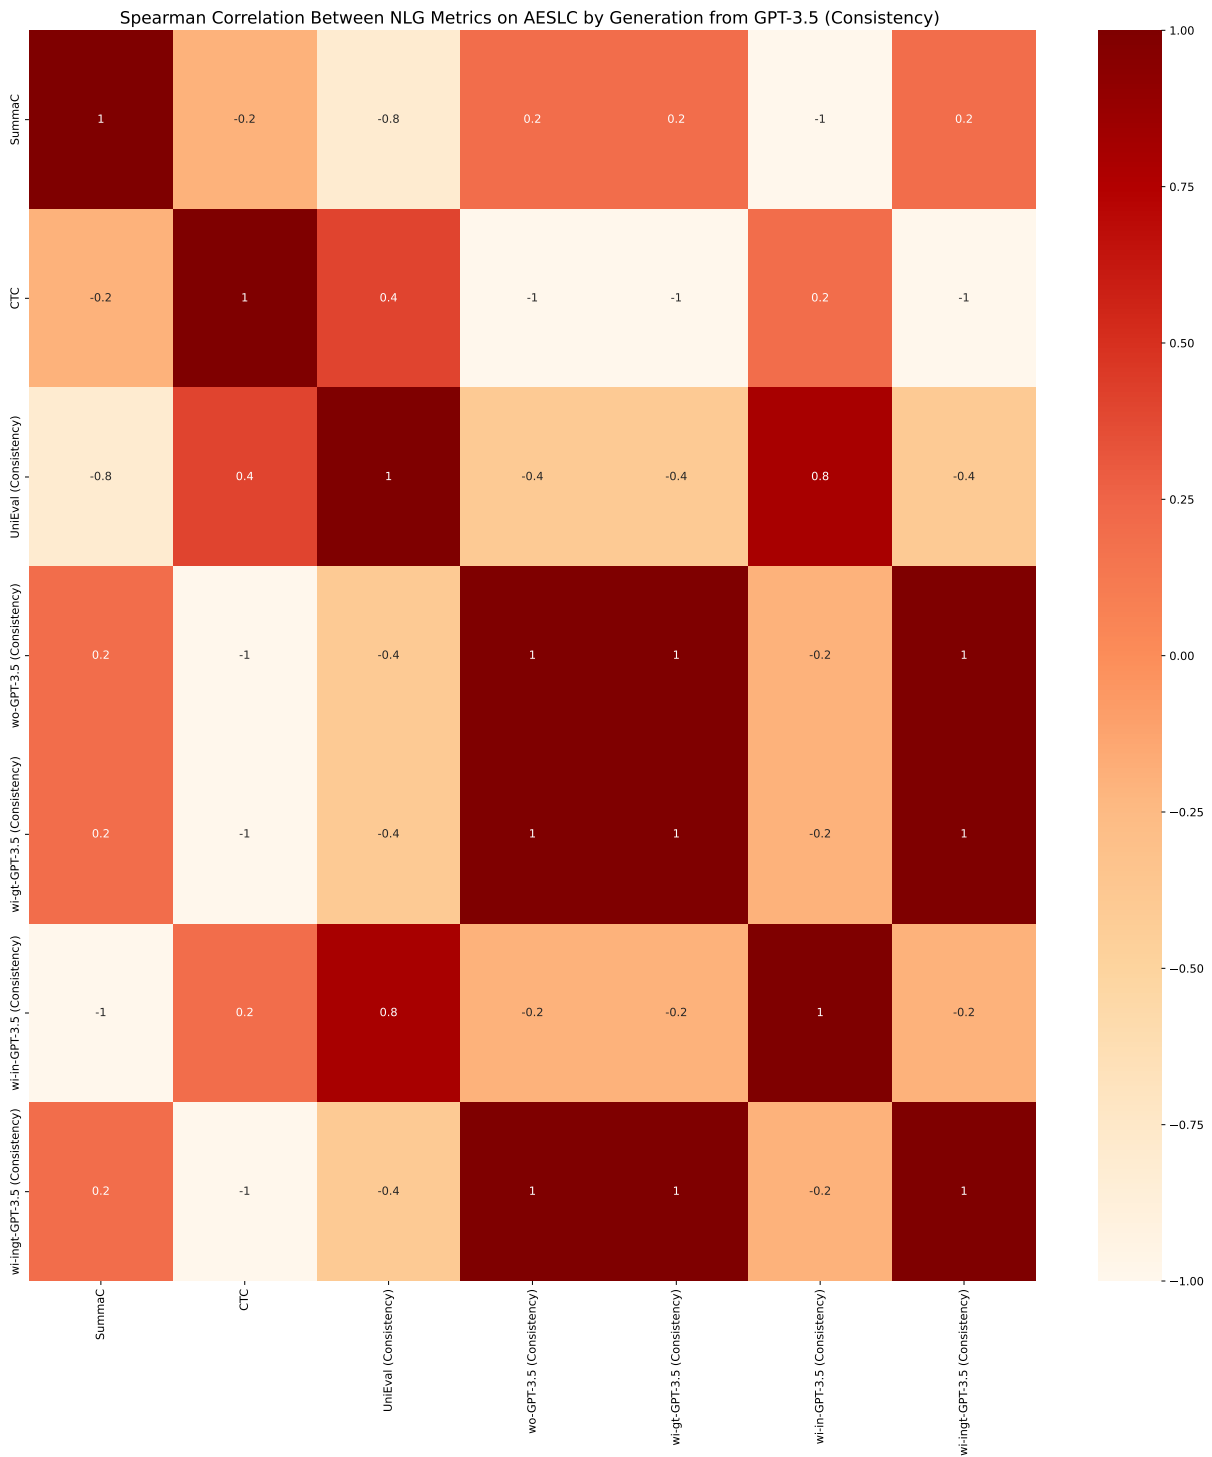


Figure 19: Diagram of Spearman correlation in terms of consistency between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 7. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

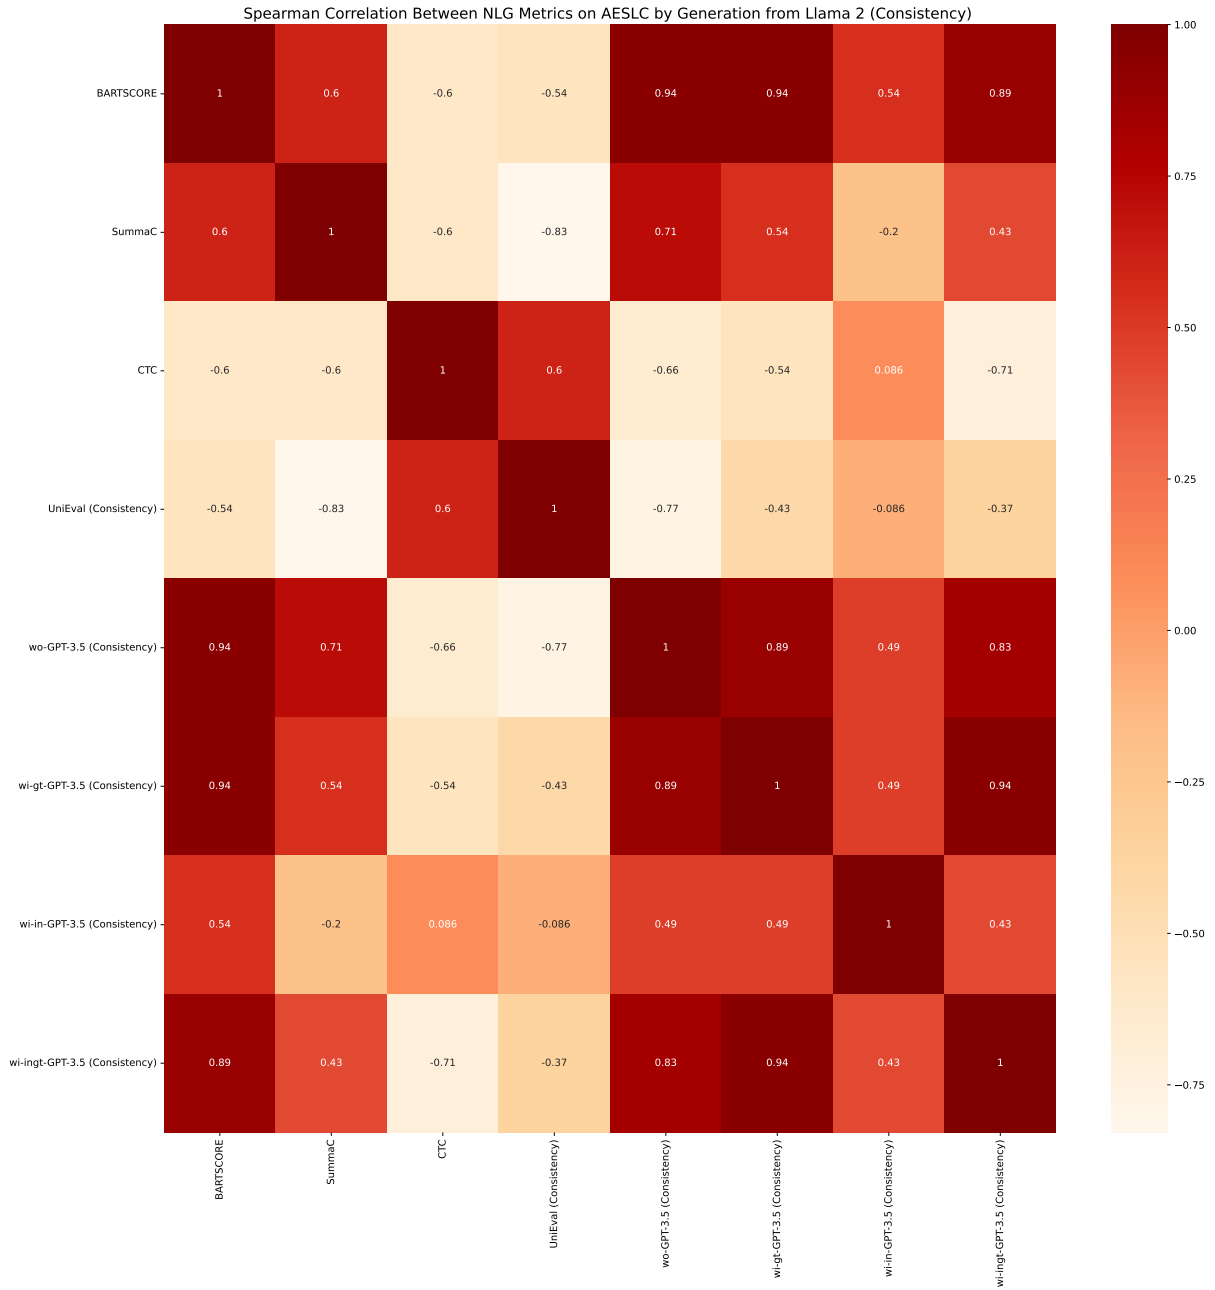


Figure 20: Diagram of Spearman correlation in terms of consistency between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 8. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

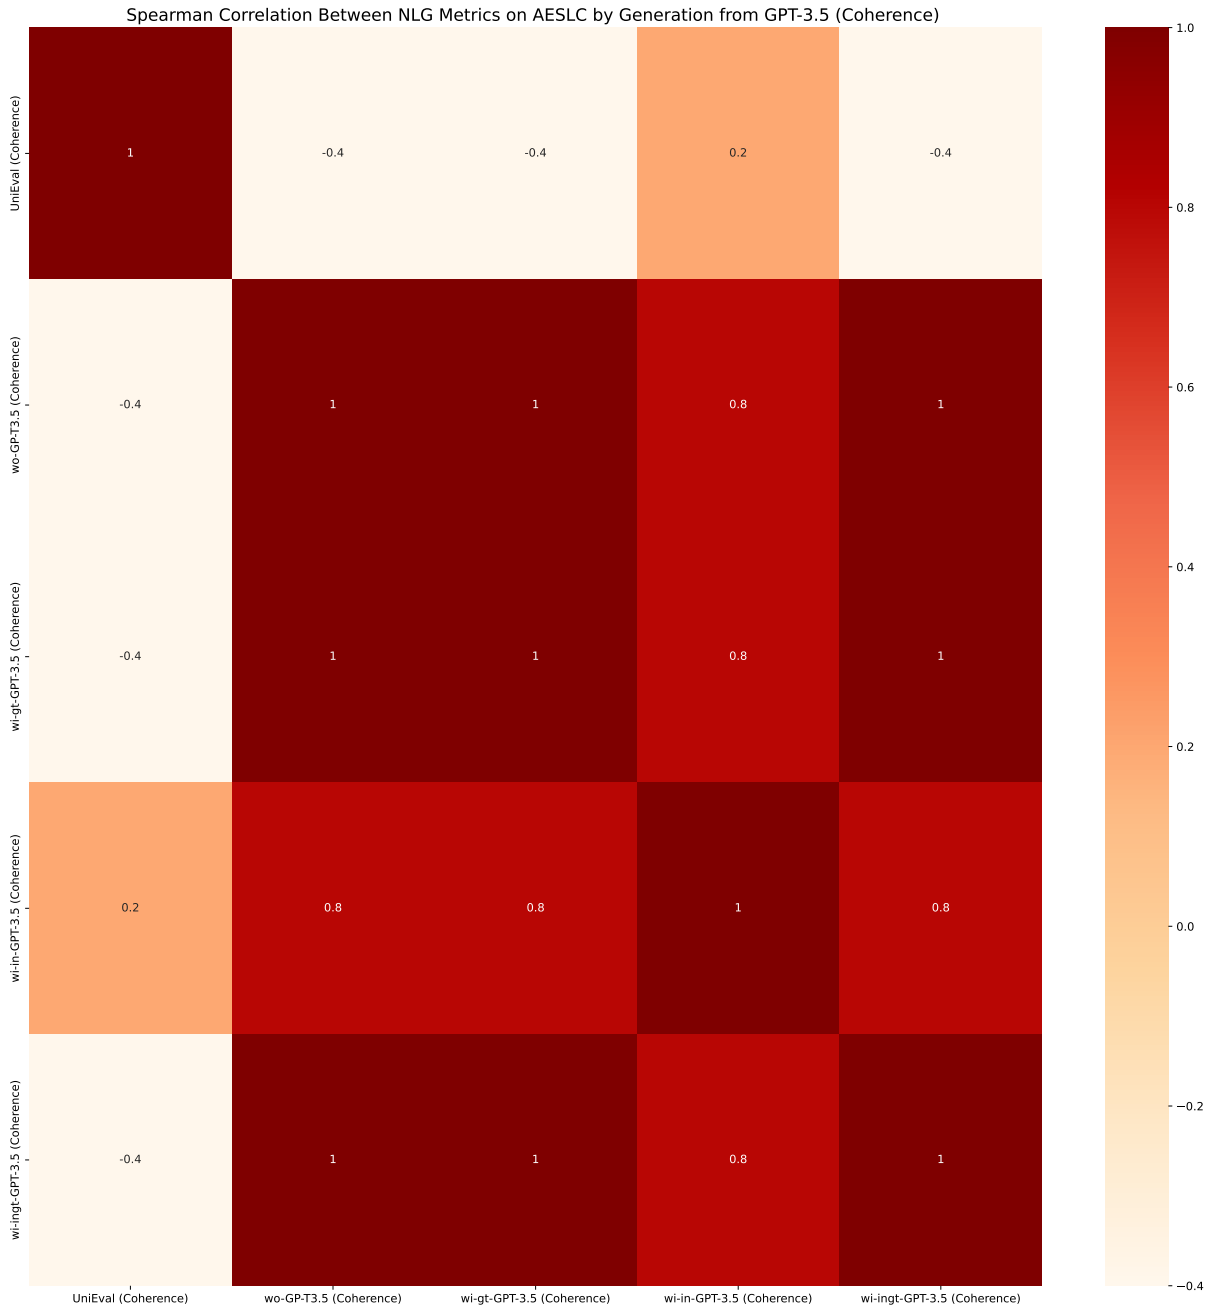


Figure 21: Diagram of Spearman correlation in terms of coherence between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 7. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

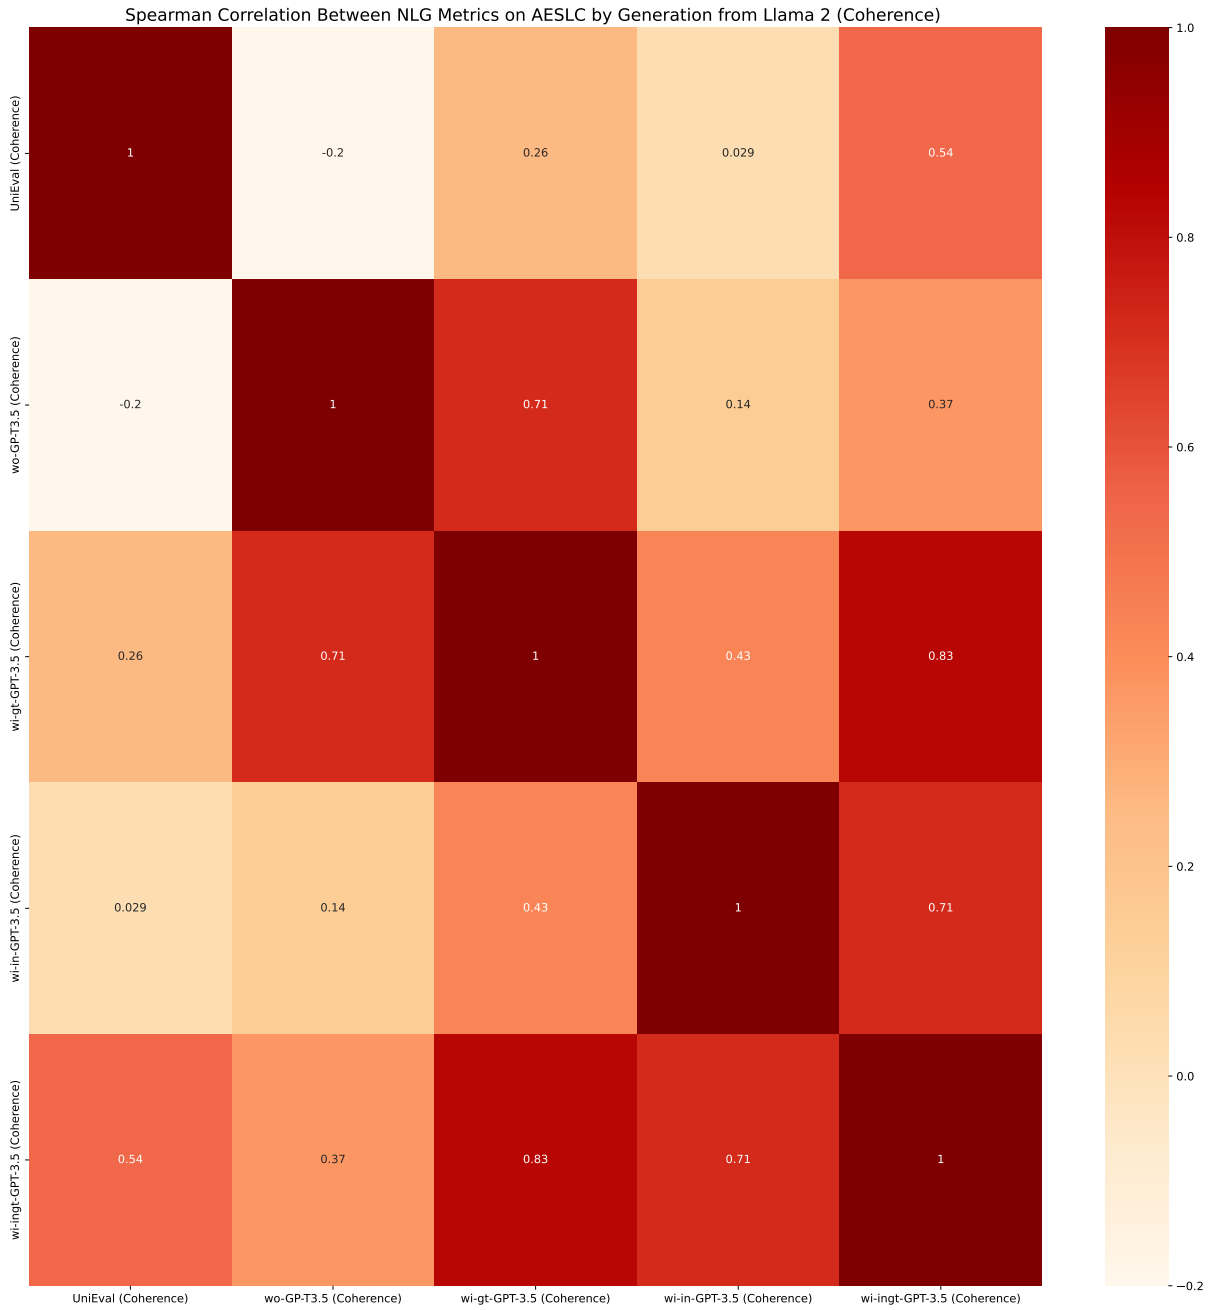


Figure 22: Diagram of Spearman correlation in terms of coherence between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 8. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

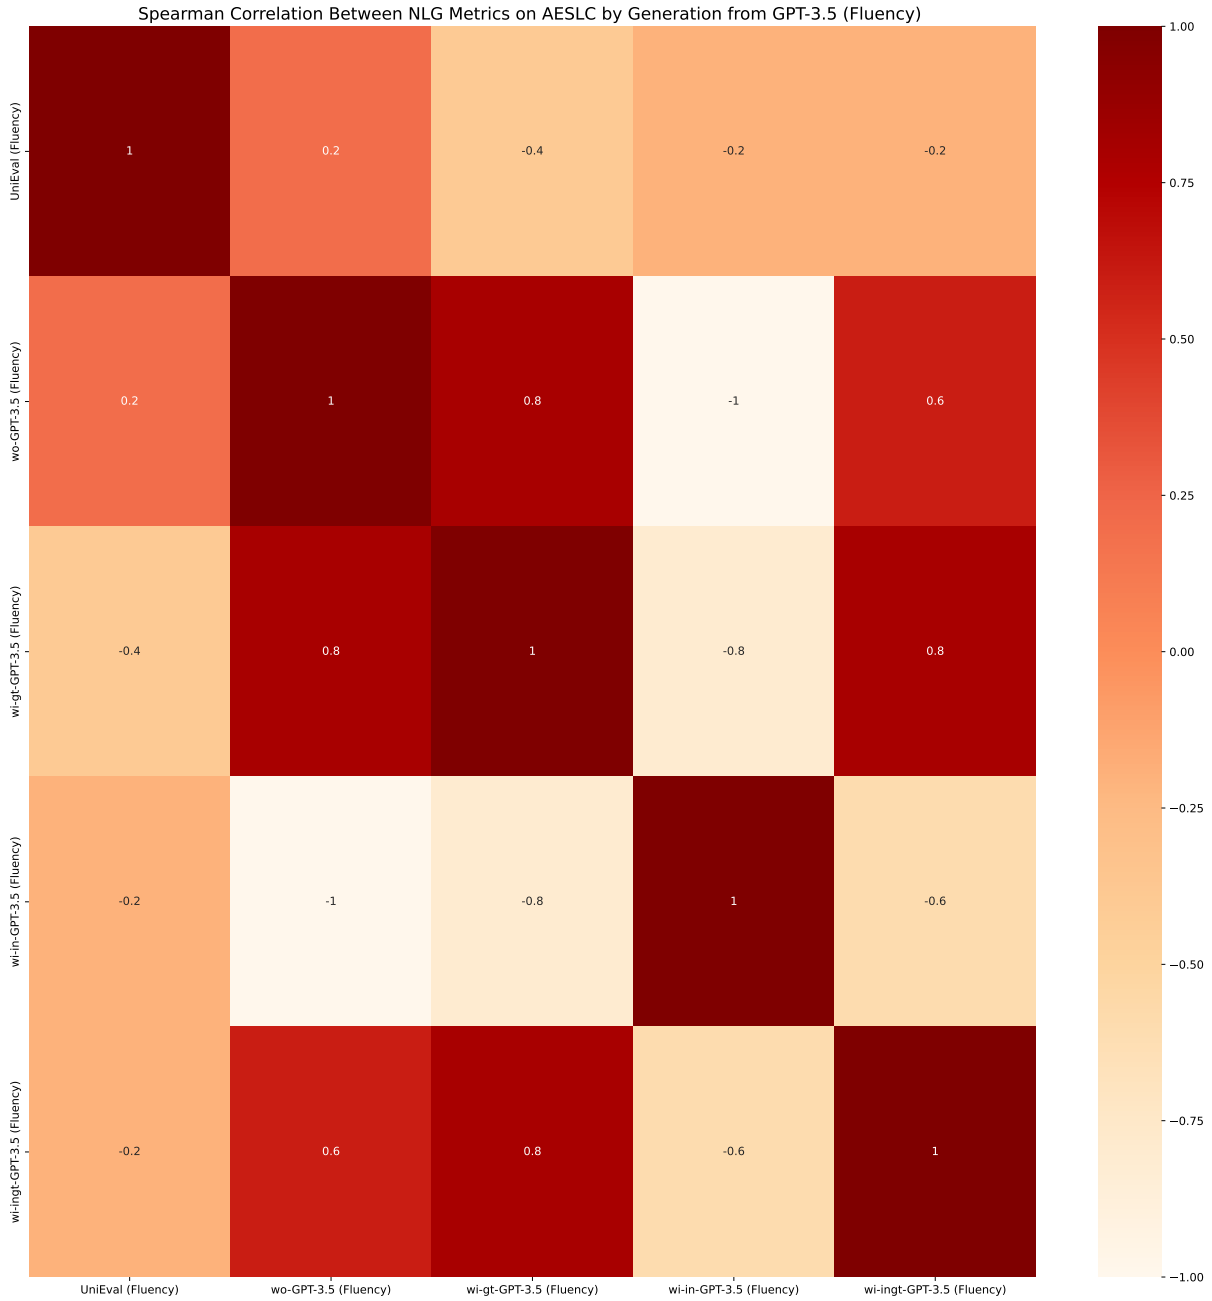


Figure 23: Diagram of Spearman correlation in terms of fluency between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 7. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

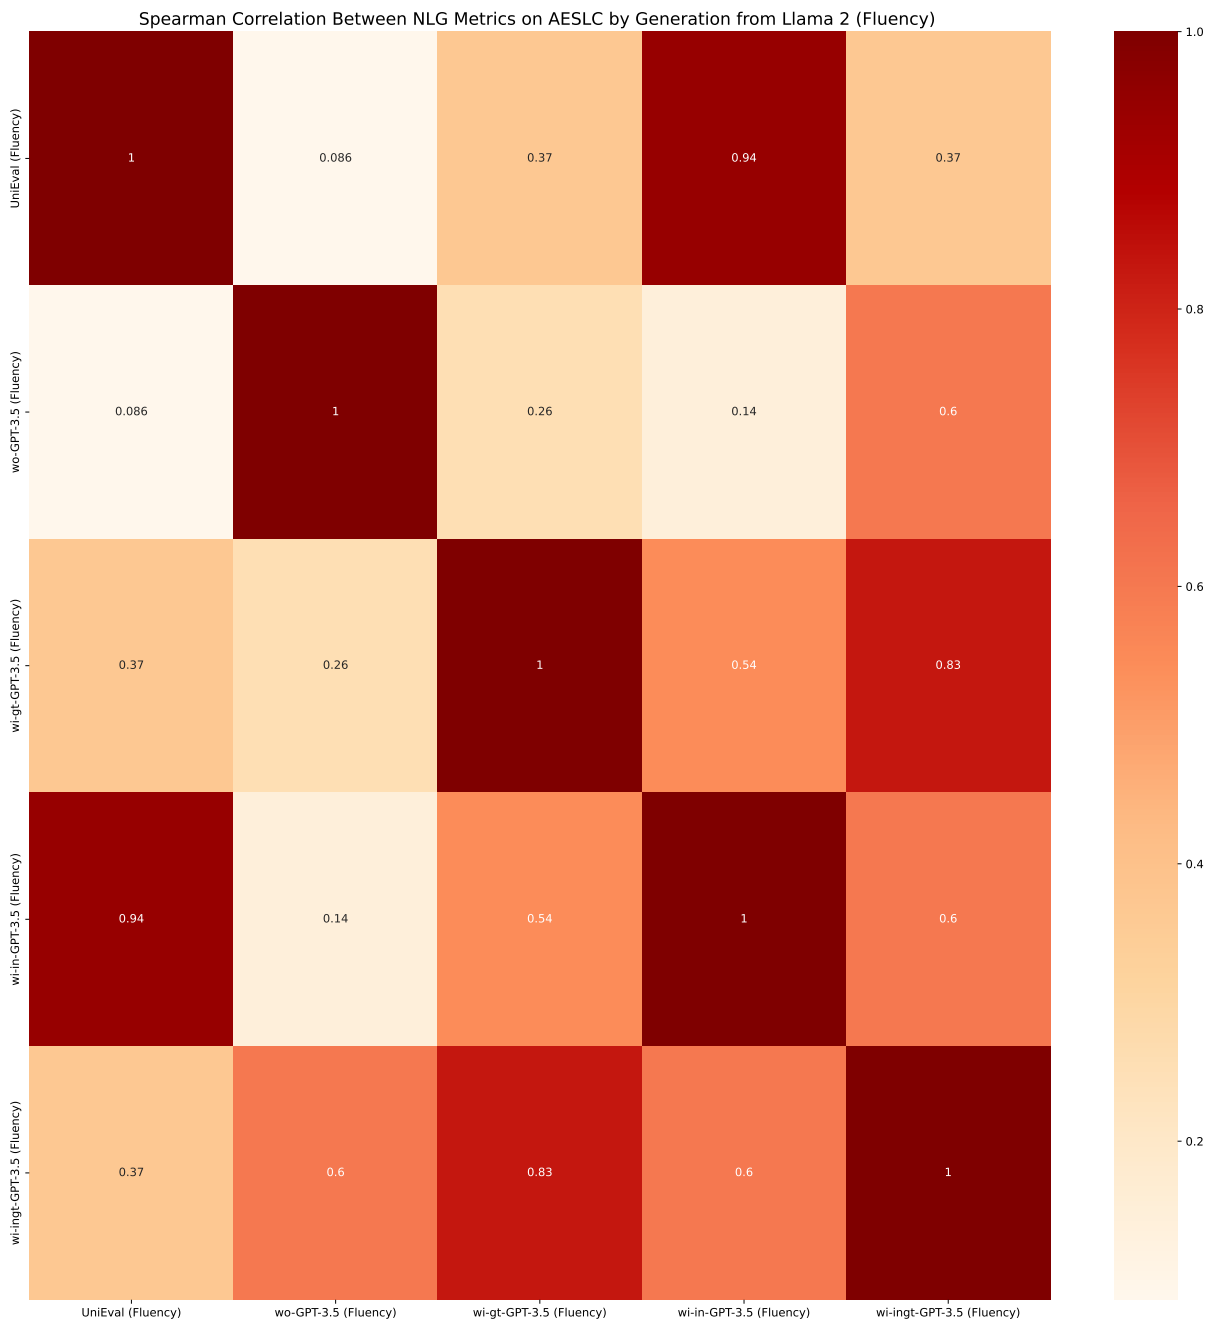


Figure 24: Diagram of Spearman correlation in terms of fluency between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 8. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

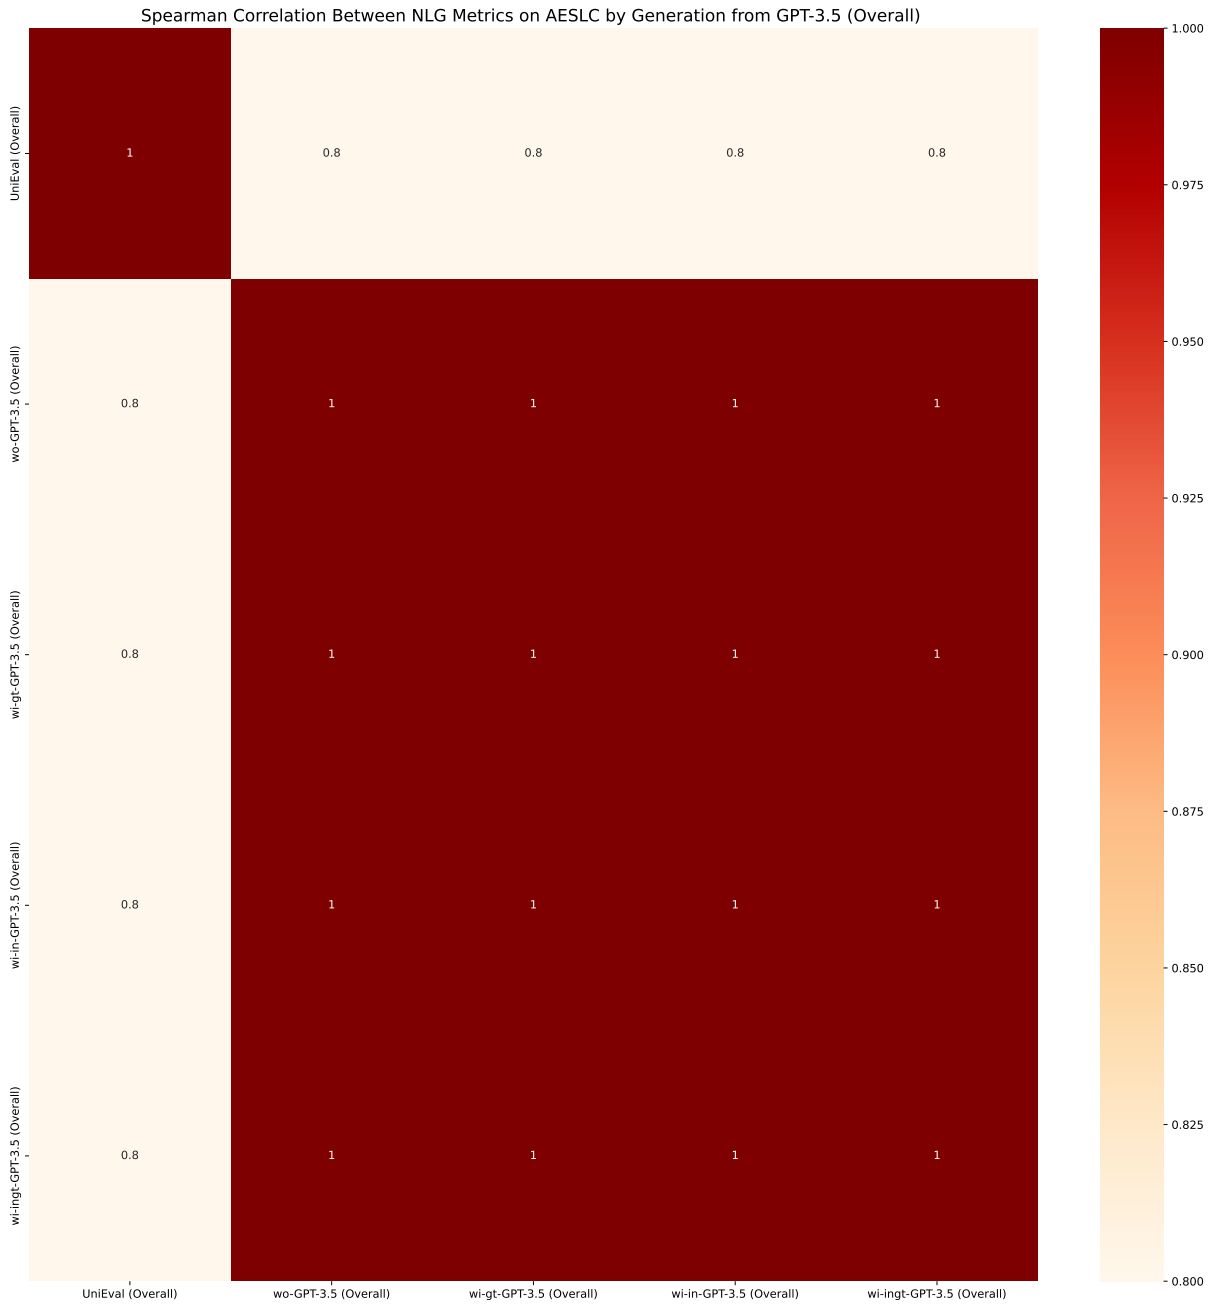


Figure 25: Diagram of Spearman correlation in terms of overall between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 7. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

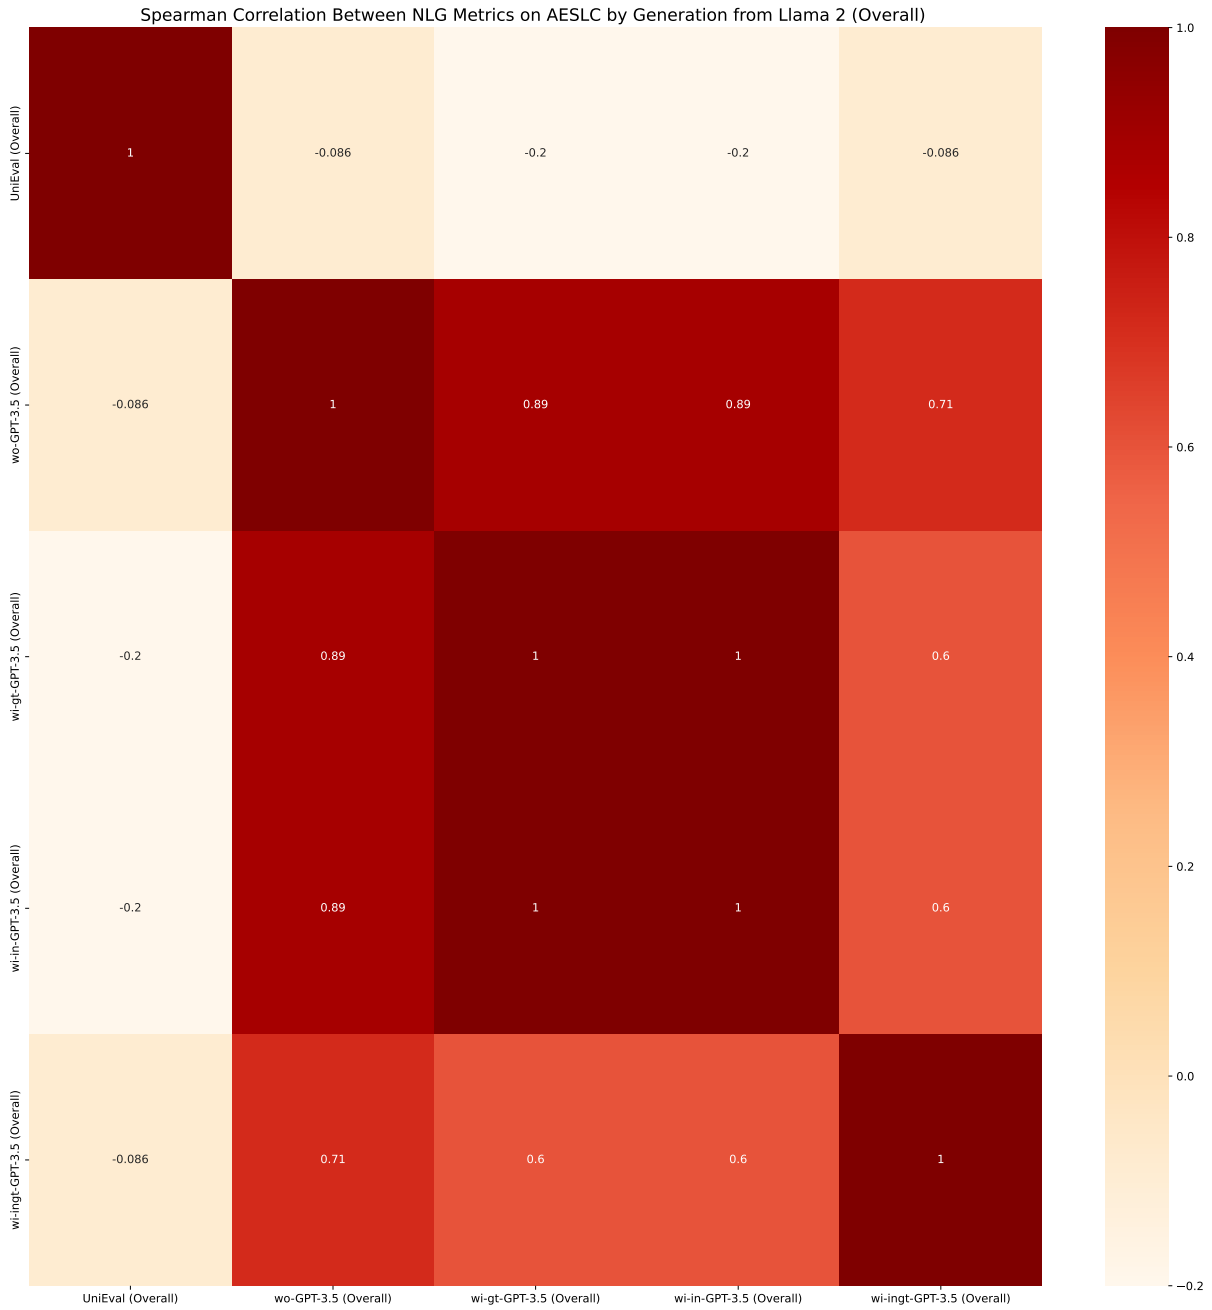


Figure 26: Diagram of Spearman correlation in terms of overall between NLG metrics on AESLC dataset from the view of uncertainty estimation methods used in Fig. 8. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

Spearman Correlation Between NLG Metrics on XSUM by Generation from BART (Relevance)

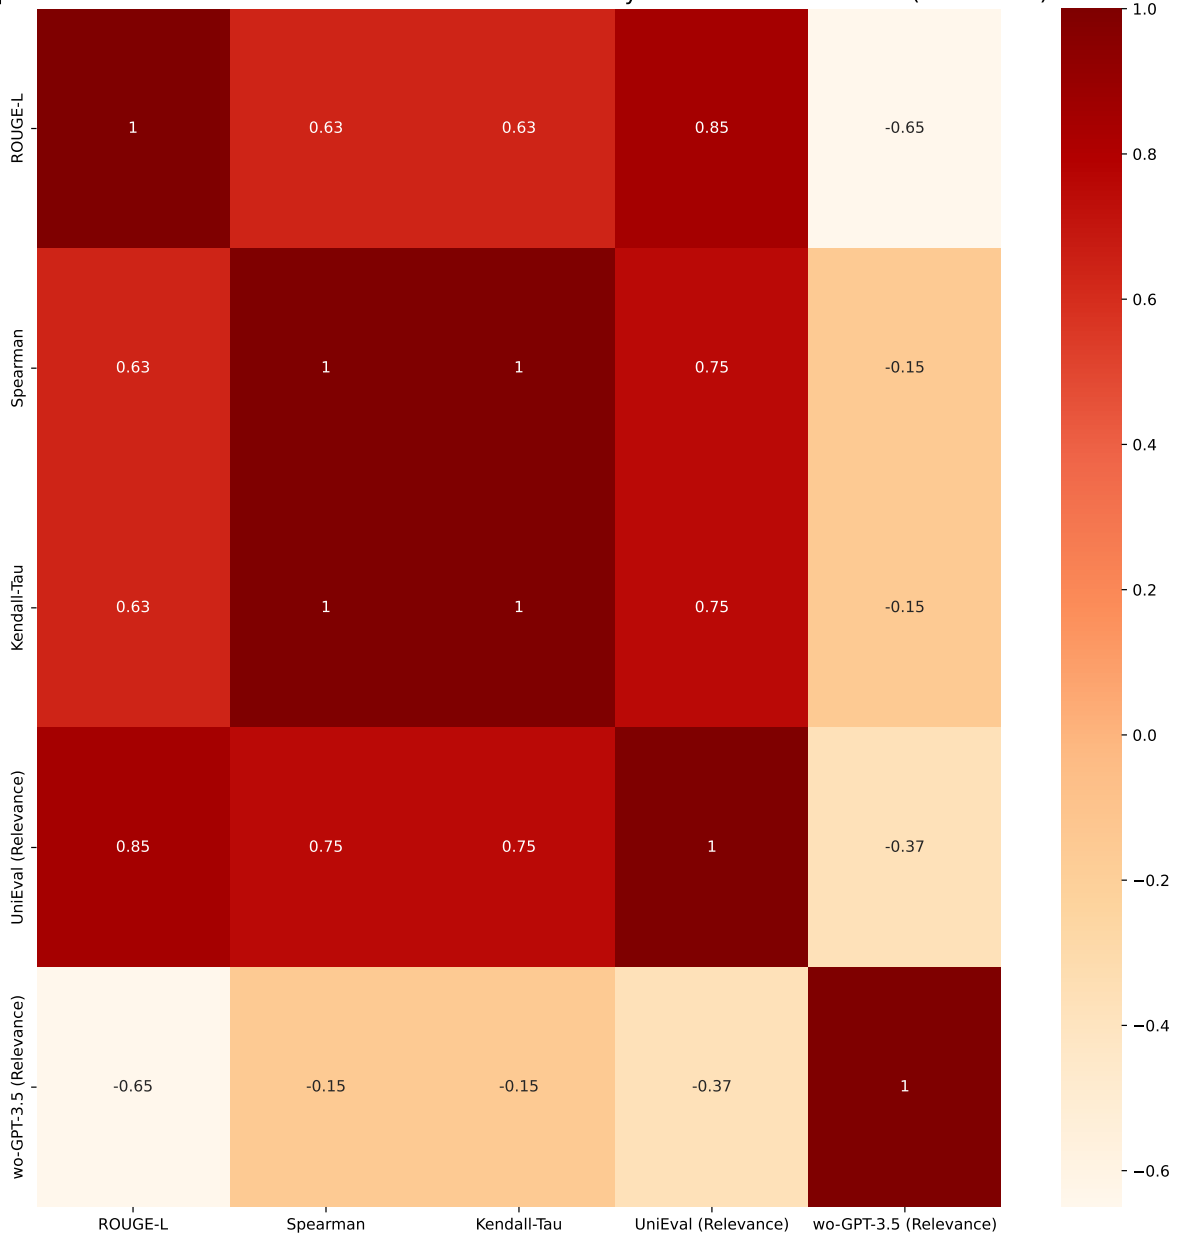


Figure 27: Diagram of Spearman correlation in terms of relevance between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 12. The generated summaries are from BART. For the GPT-3.5-based NLG metrics, we only conduct wo-GPT-3.5 on the BART generation model setting.

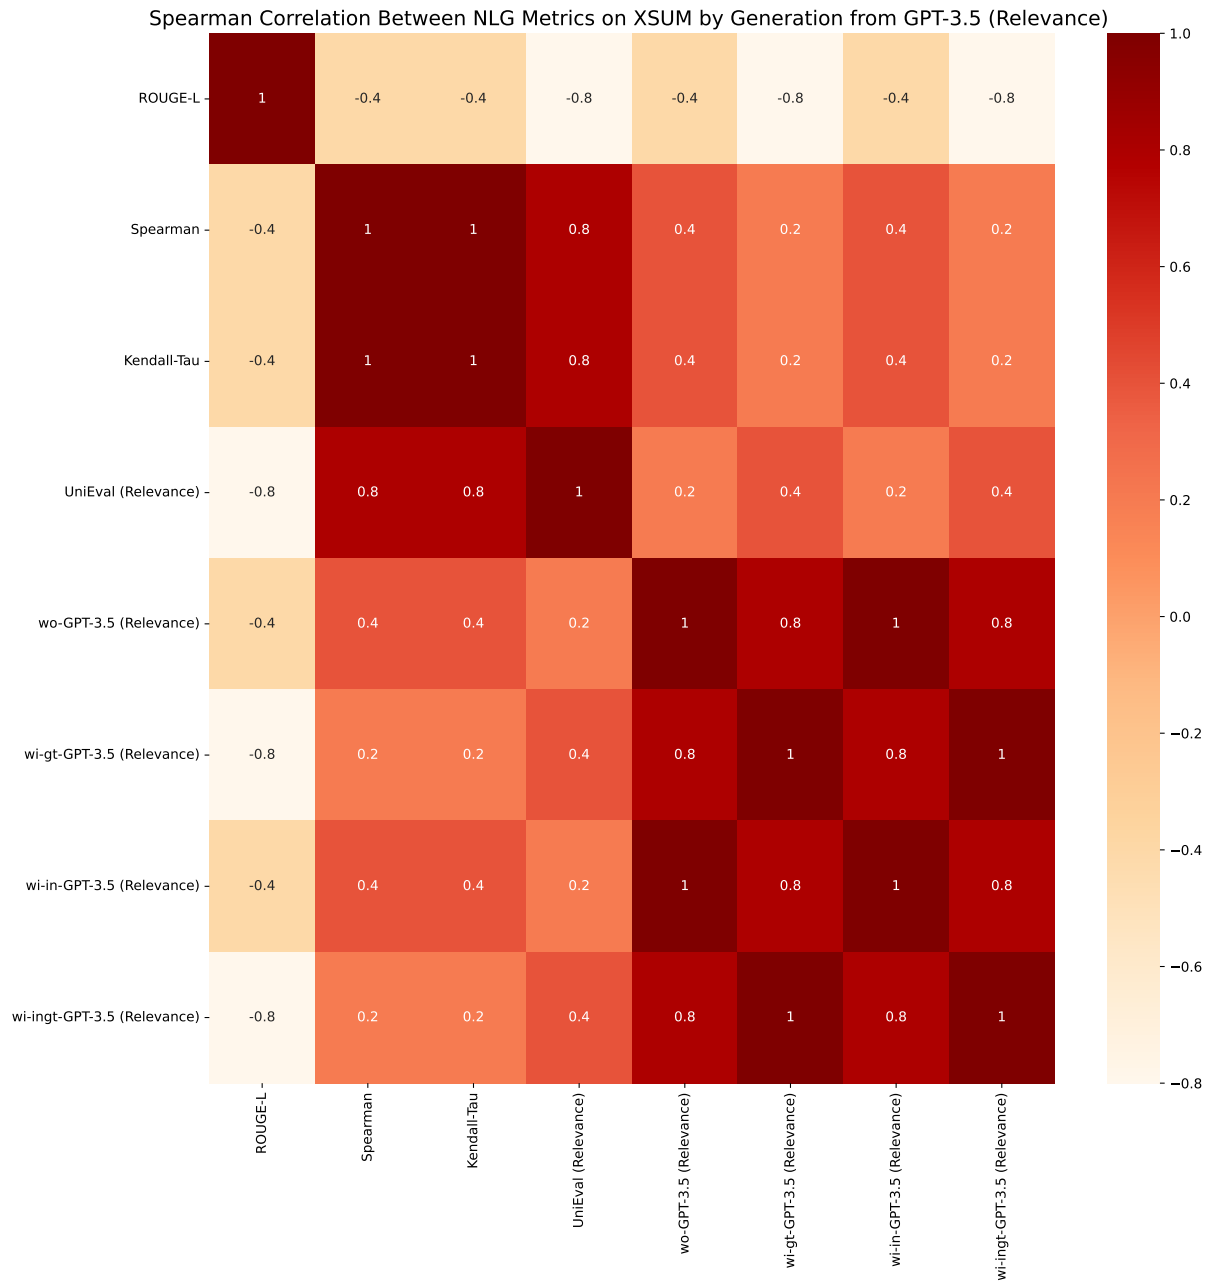


Figure 28: Diagram of Spearman correlation in terms of relevance between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 13. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

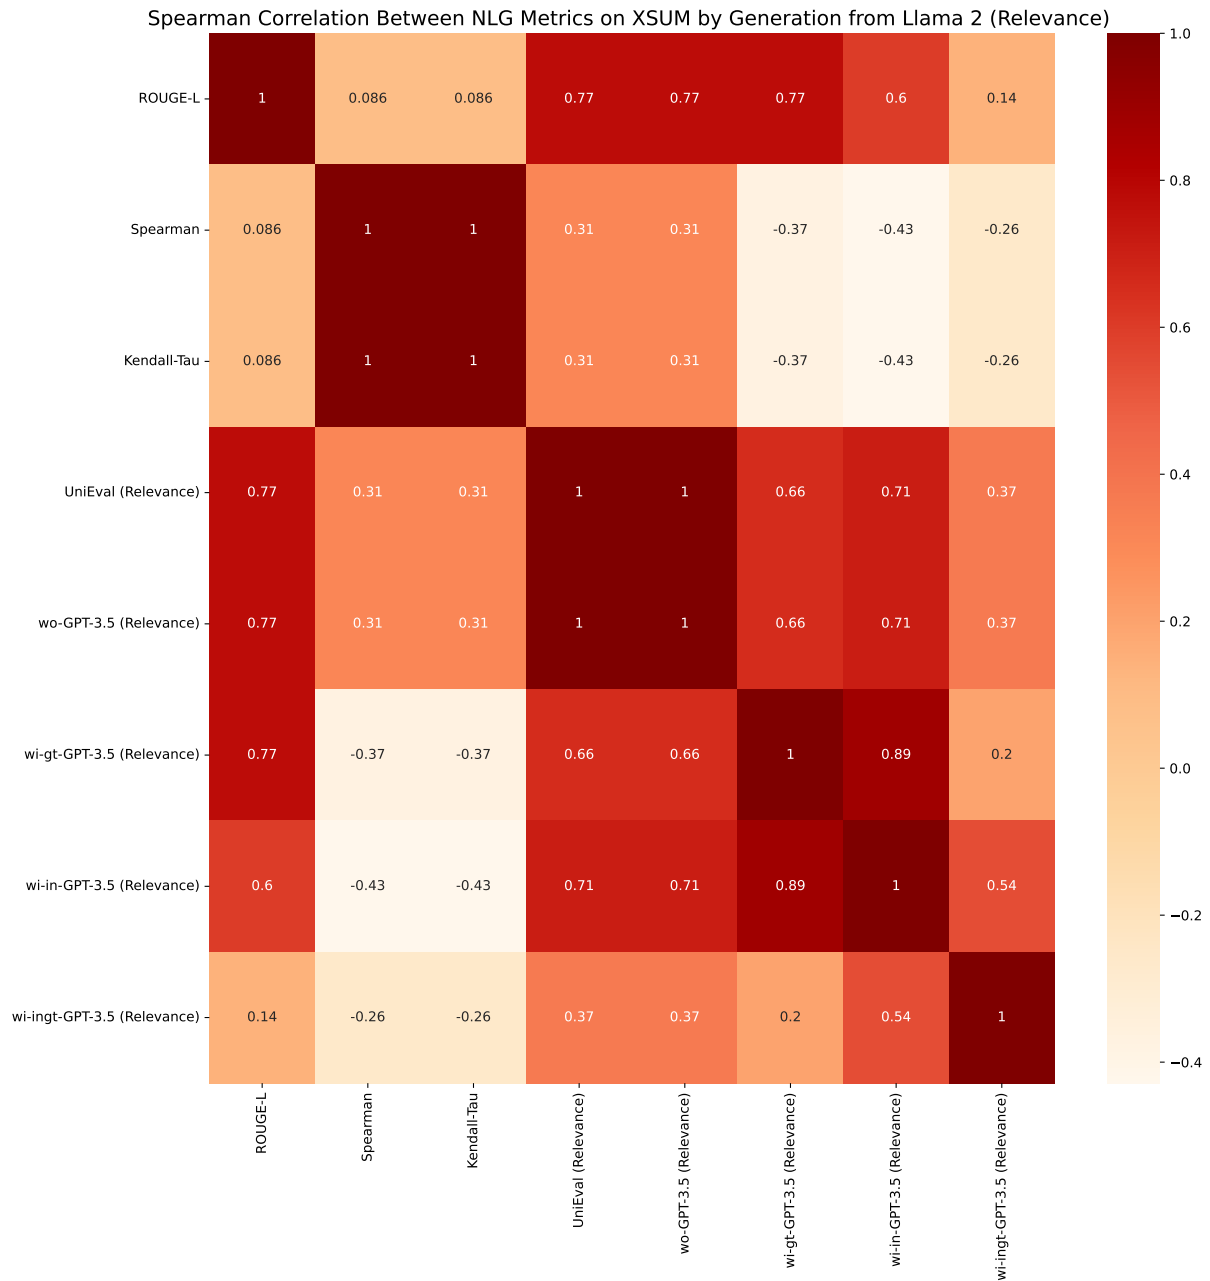


Figure 29: Diagram of Spearman correlation in terms of relevance between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 14. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

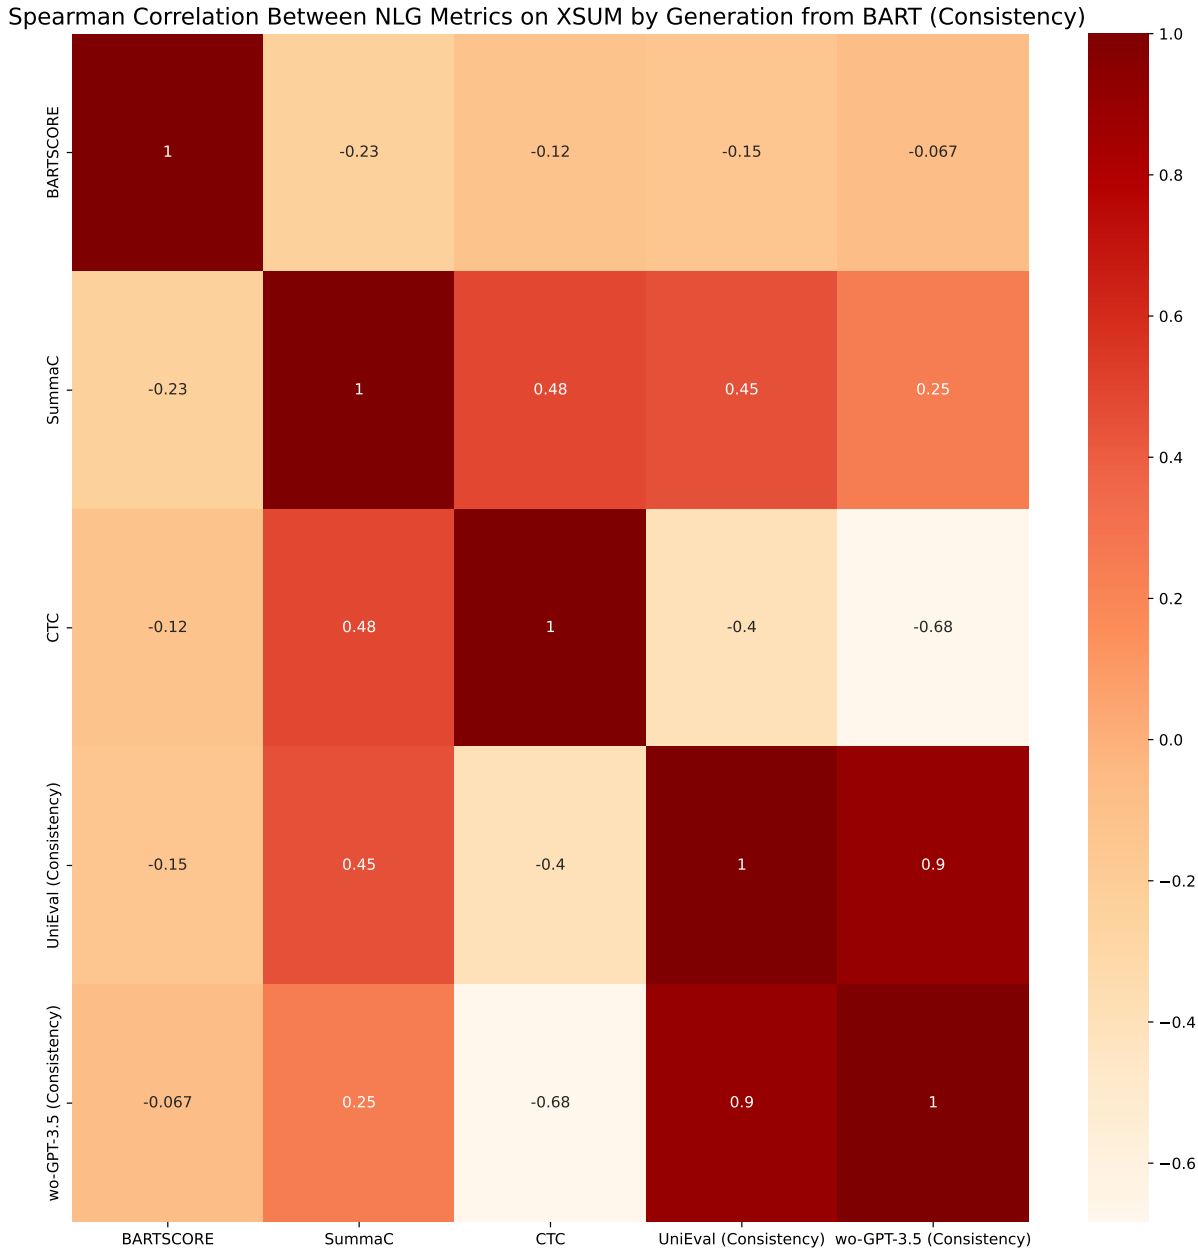


Figure 30: Diagram of Spearman correlation in terms of consistency between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 12. The generated summaries are from BART. For the GPT-3.5-based NLG metrics, we only conduct wo-GPT-3.5 on the BART generation model setting.

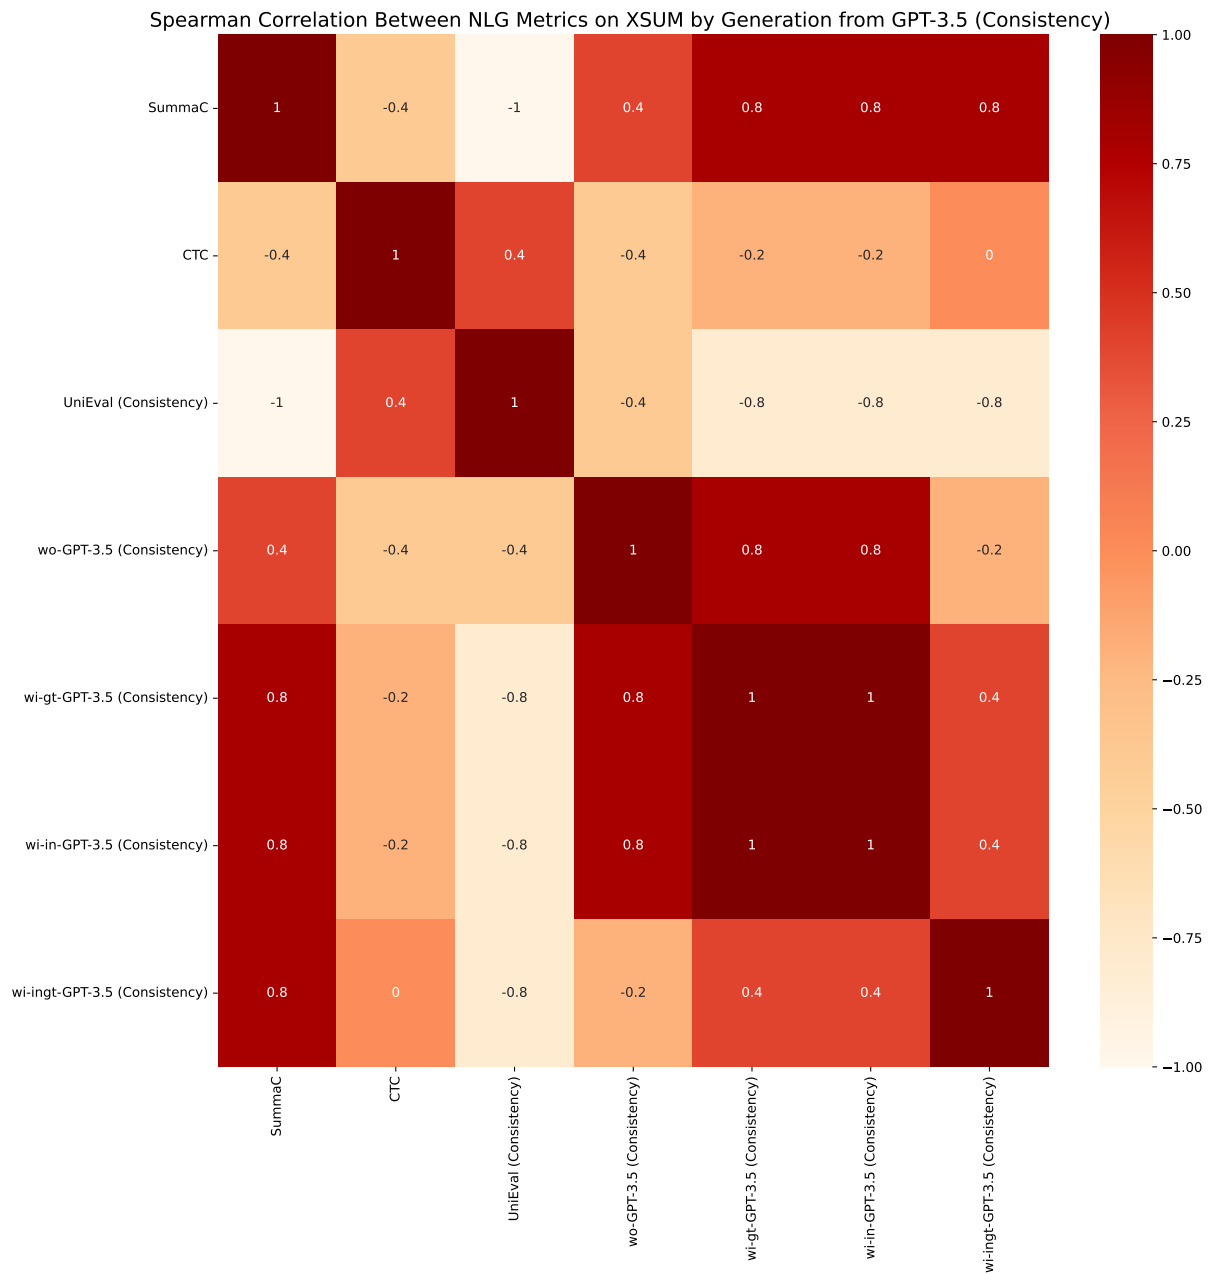


Figure 31: Diagram of Spearman correlation in terms of consistency between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 13. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

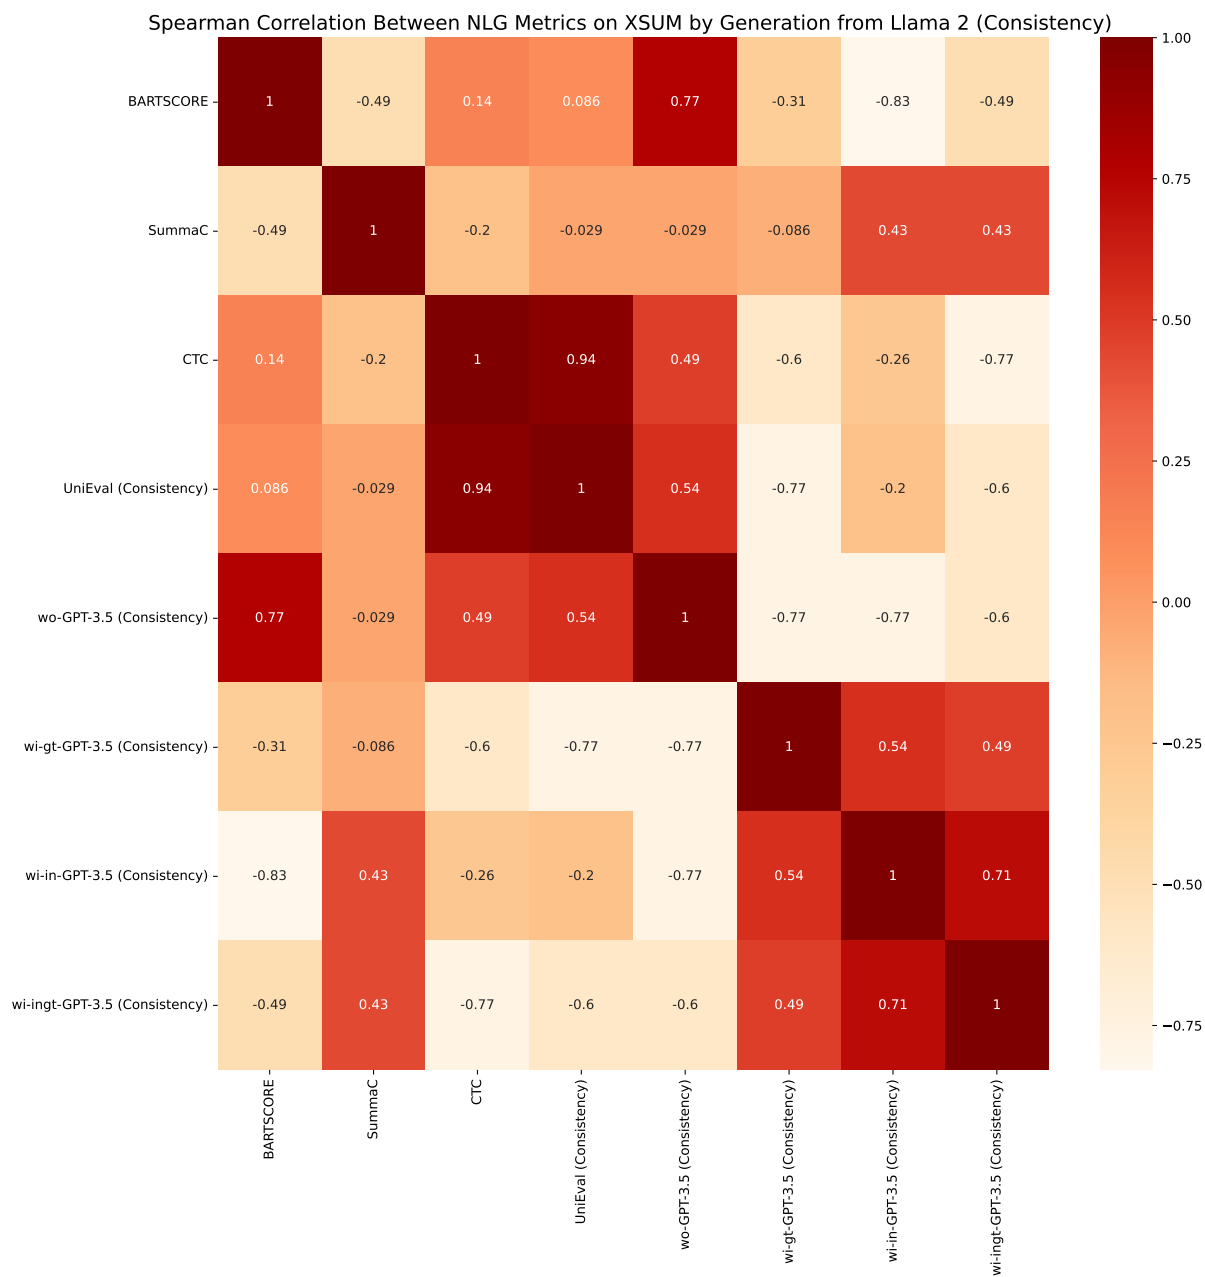


Figure 32: Diagram of Spearman correlation in terms of consistency between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 14. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

Spearman Correlation Between NLG Metrics on XSUM by Generation from GPT-3.5 (Coherence)

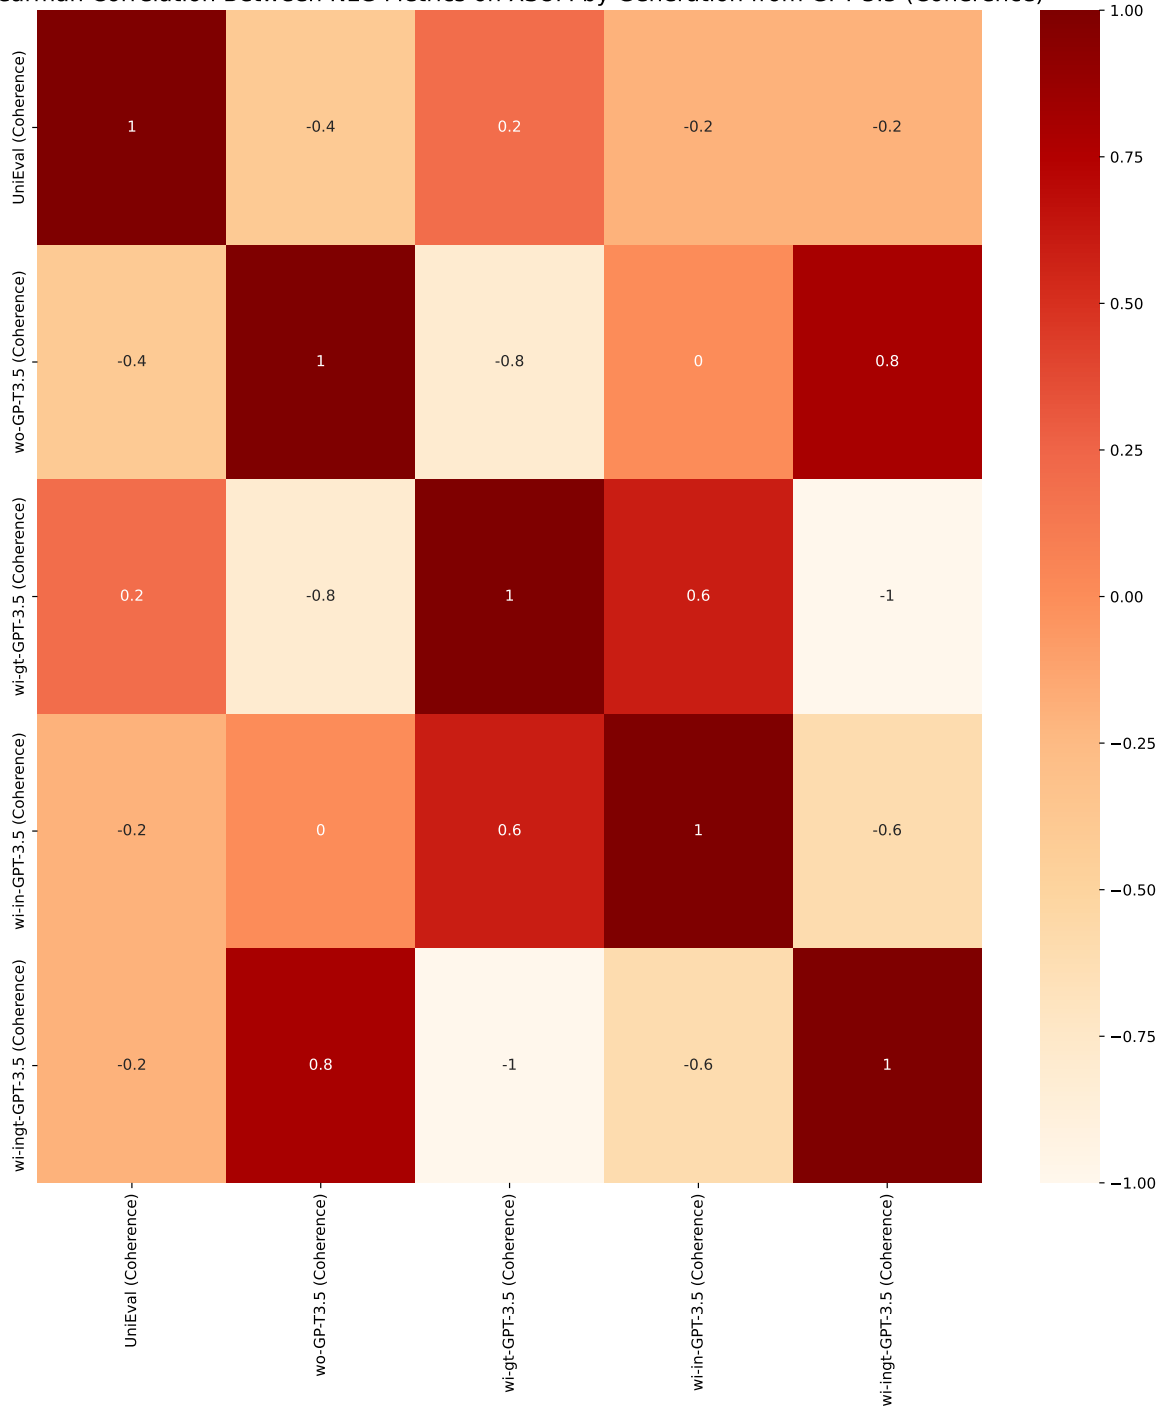


Figure 33: Diagram of Spearman correlation in terms of coherence between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 13. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

Spearman Correlation Between NLG Metrics on XSUM by Generation from Llama 2 (Coherence)

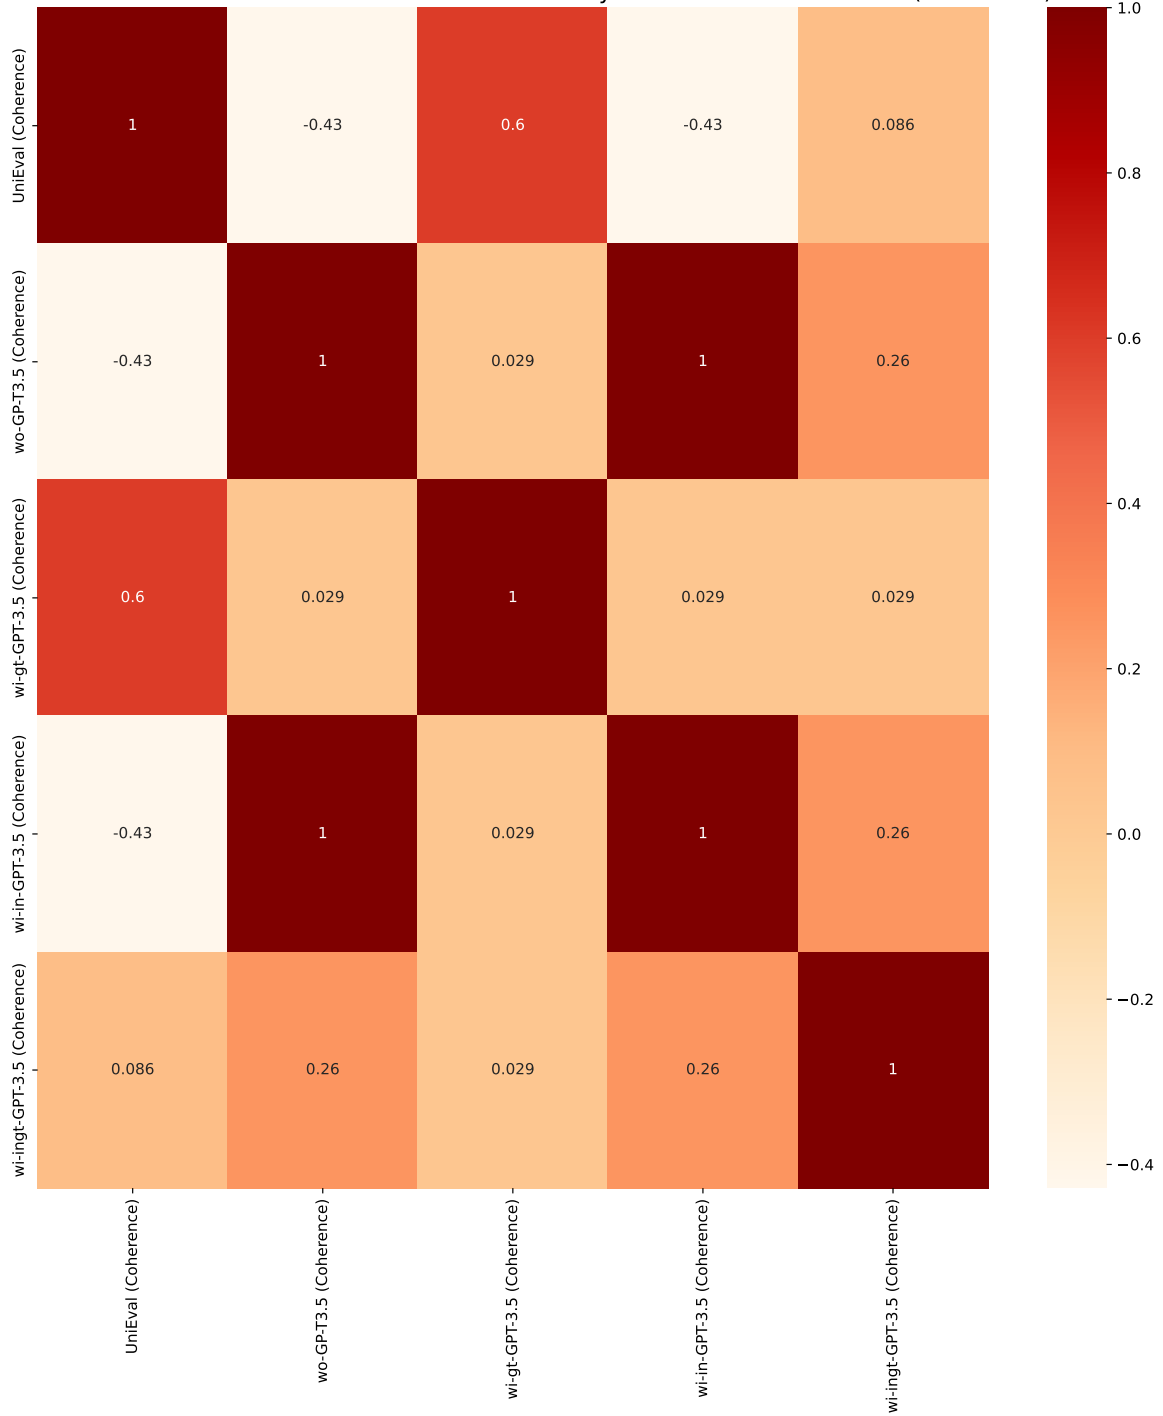


Figure 34: Diagram of Spearman correlation in terms of coherence between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 14. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

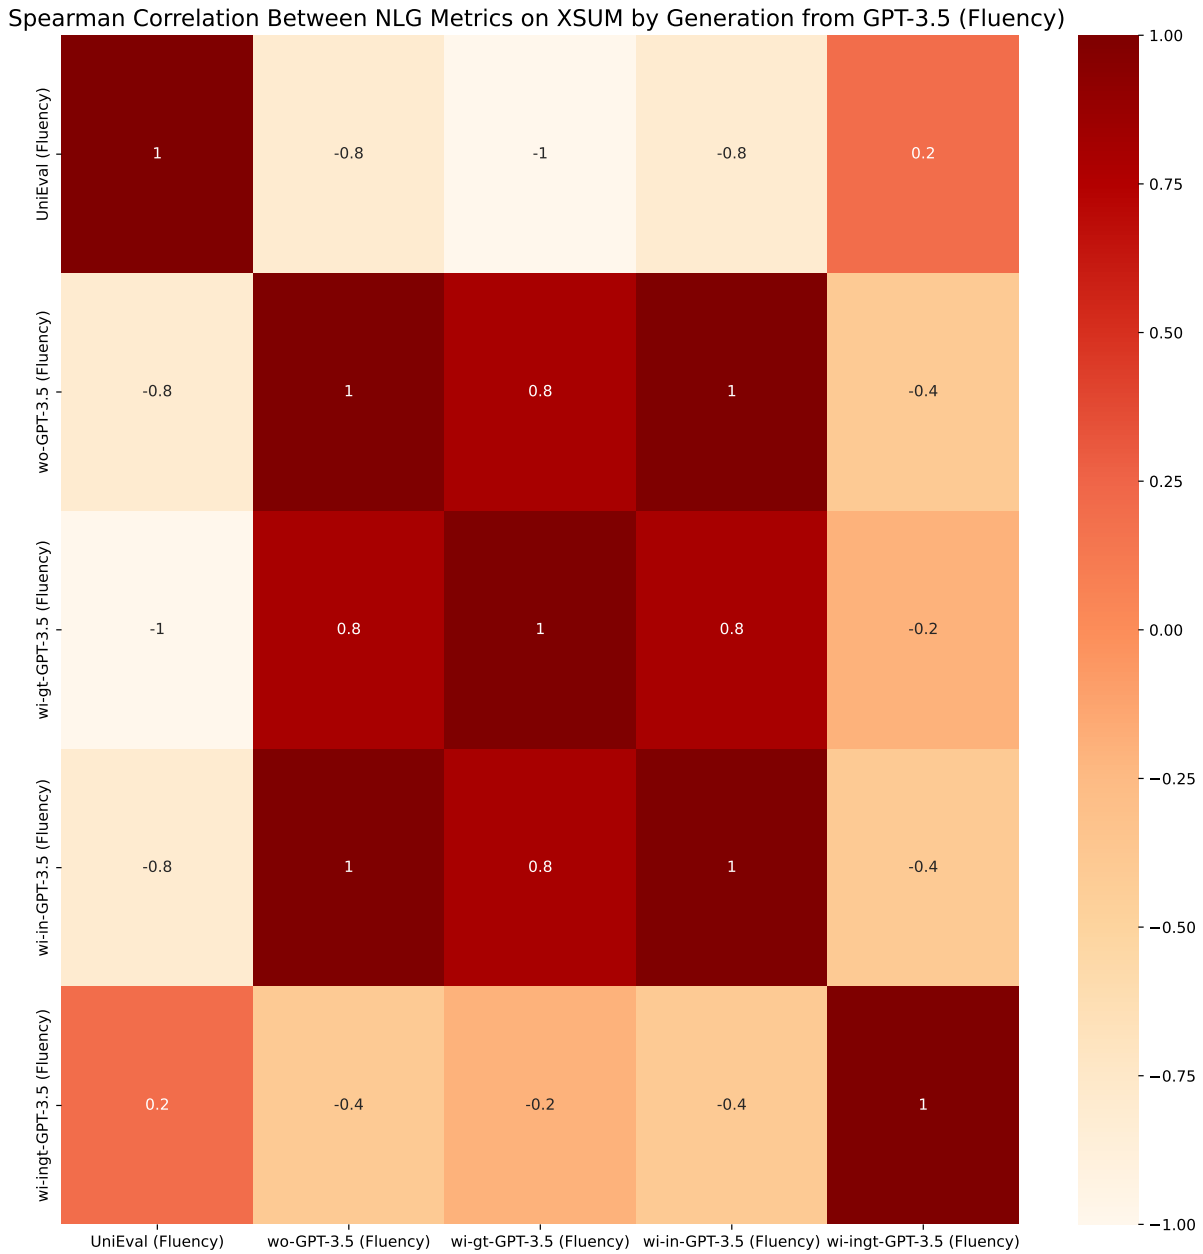


Figure 35: Diagram of Spearman correlation in terms of fluency between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 13. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

Spearman Correlation Between NLG Metrics on XSUM by Generation from Llama 2 (Fluency)

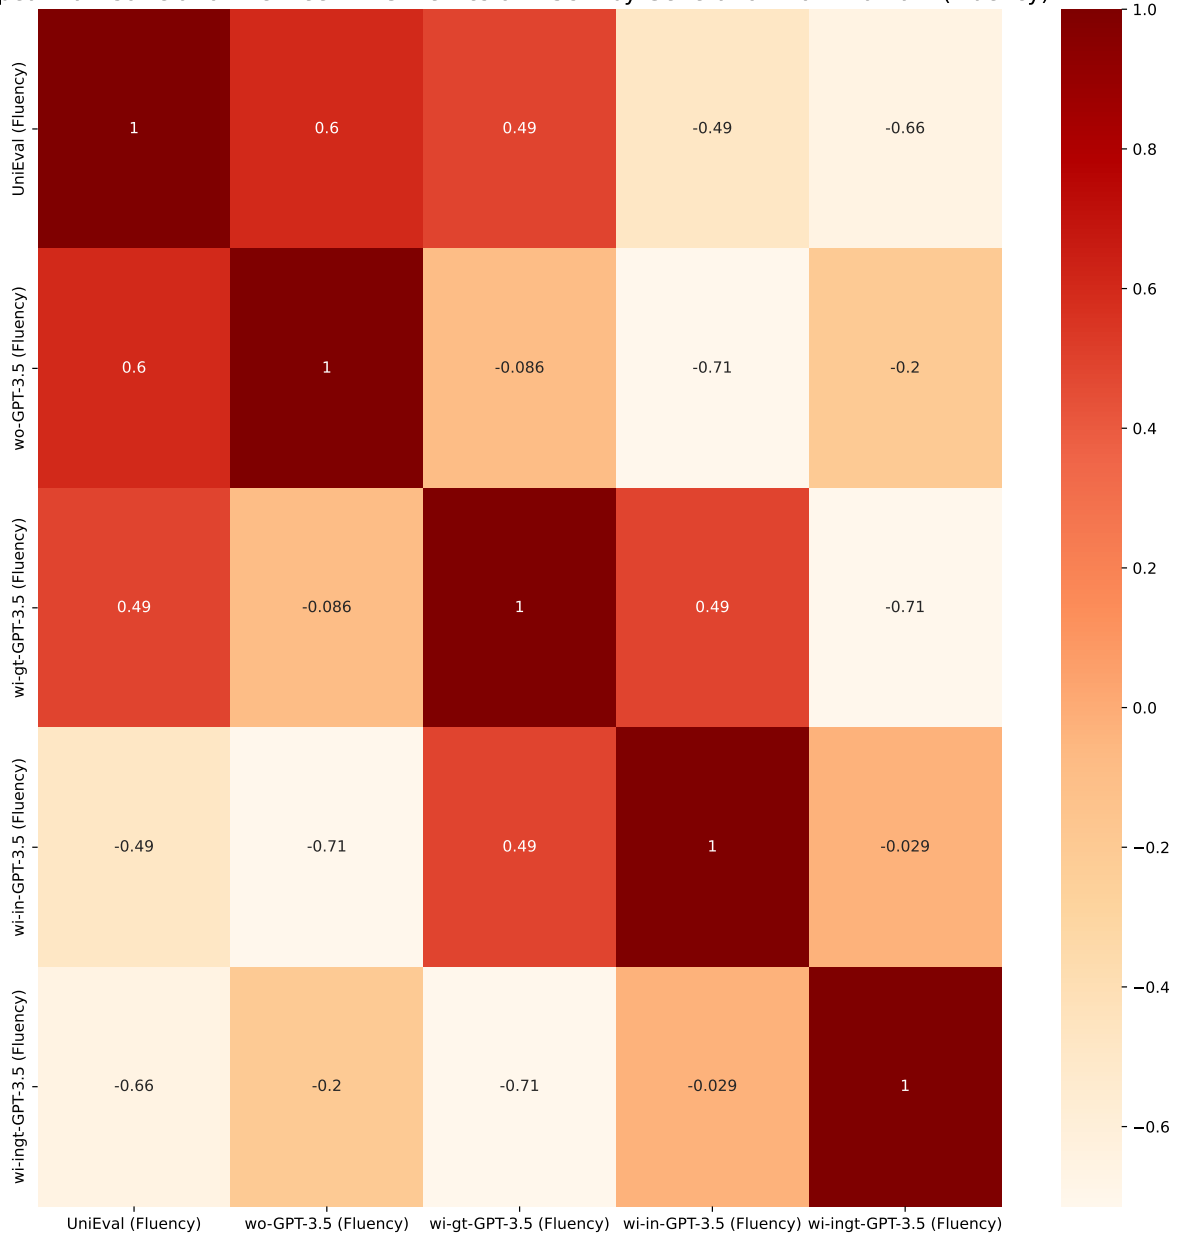


Figure 36: Diagram of Spearman correlation in terms of fluency between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 14. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

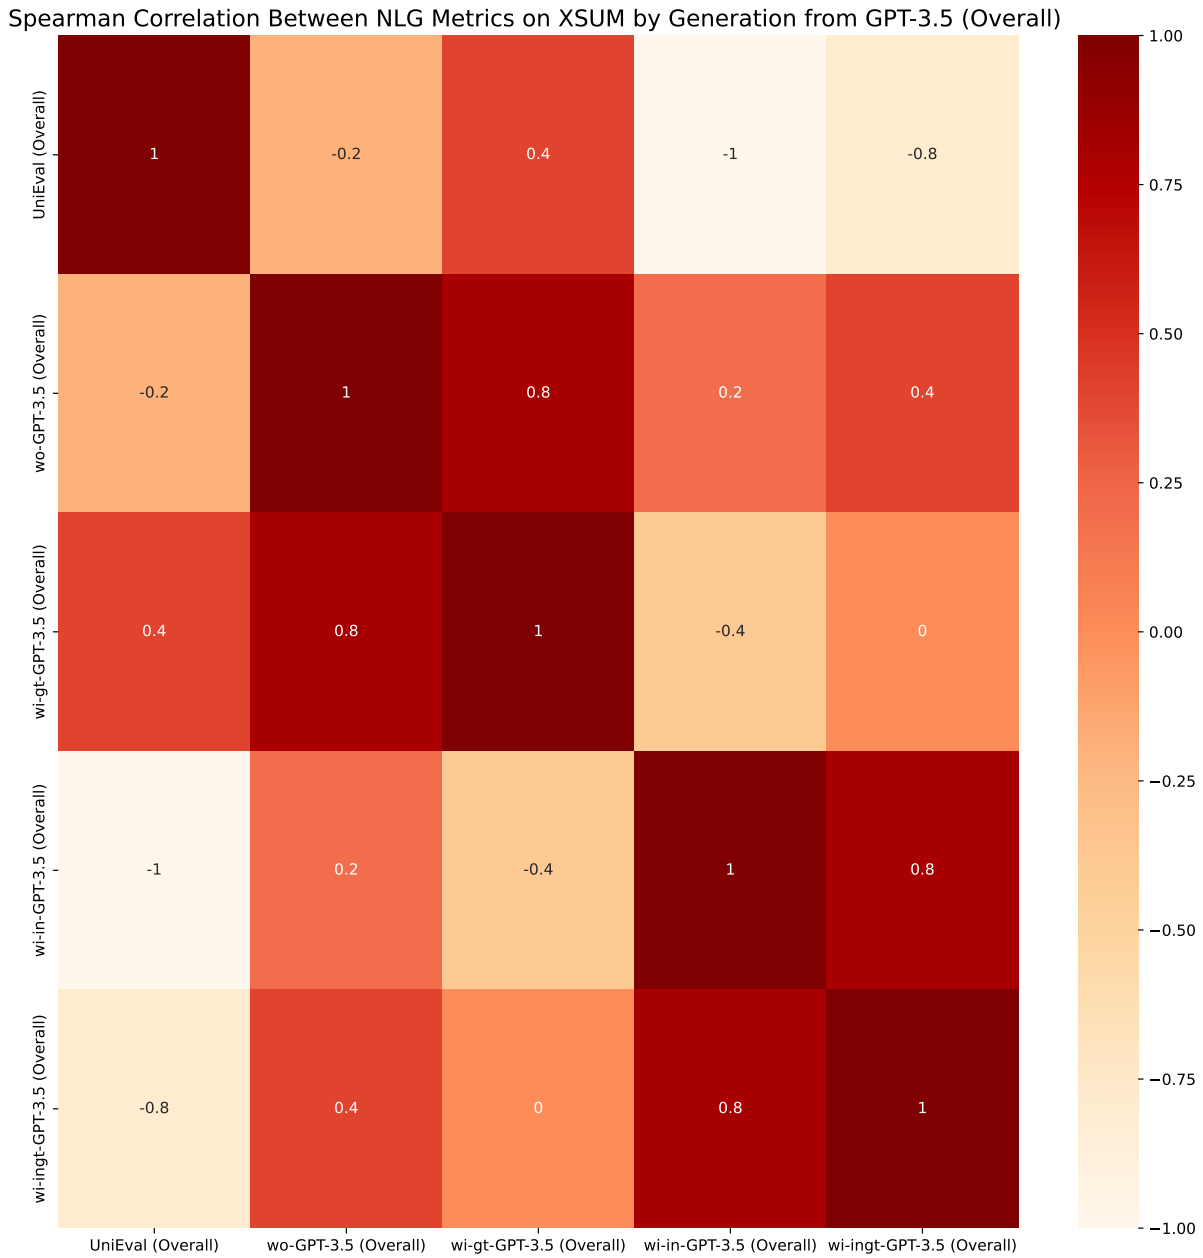


Figure 37: Diagram of Spearman correlation in terms of overall between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 13. The generated summaries are from GPT-3.5. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

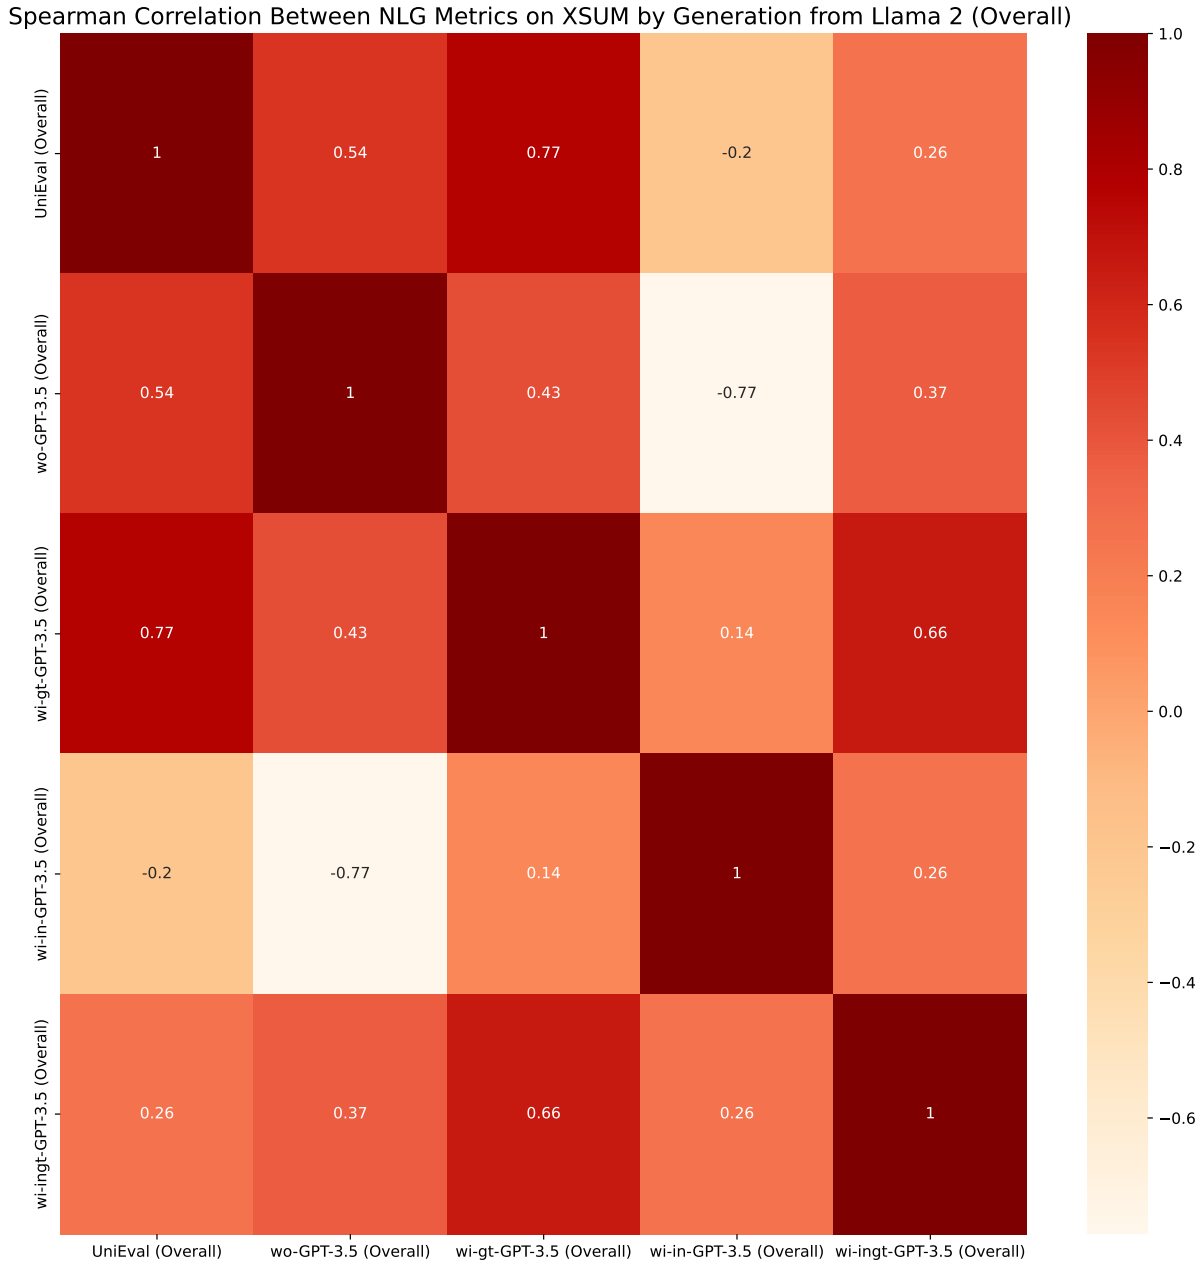


Figure 38: Diagram of Spearman correlation in terms of overall between NLG metrics on XSUM dataset from the view of uncertainty estimation methods used in Fig. 14. The generated summaries are from Llama 2. For the GPT-3.5-based NLG metrics, we only draw wi-ingt-GPT-3.5 results to save space.

NLG Metrics	NumSets	ECC	LexSim	EigV
ROUGE-L	-0.1869	0.0045	0.3268	0.0035
SummaC	0.1602	-0.1263	-0.1231	-0.0938
CTC	-0.2152	0.095	0.419	0.0952
Spearman	0.1283	-0.0595	0.1782	-0.0114
Kendall-Tau	0.1287	-0.055	0.1783	-0.0116
UniEval (Relevance)	-0.0173	-0.2415	-0.2378	-0.3188
UniEval (Consistency)	0.2134	-0.2464	-0.2447	-0.0808
UniEval (Coherence)	-0.0798	-0.1396	-0.2953	-0.0963
UniEval (Fluency)	0.2574	-0.7312	-0.6689	-0.6715
UniEval (Overall)	0.0921	-0.3211	-0.3598	-0.2563
wi-ingt-GPT-3.5 (Relevance)	0.2014	-0.0067	0.2394	0.1356
wi-ingt-GPT-3.5 (Consistency)	-0.1509	0.3037	0.4858	0.4395
wi-ingt-GPT-3.5 (Coherence)	0.0309	0.1237	0.3326	0.2971
wi-ingt-GPT-3.5 (Fluency)	-0.1544	0.2456	0.4483	0.4712
wi-ingt-GPT-3.5 (Overall)	0.1455	0.152	0.3106	0.3046
Col Mean	0.0369	-0.0669	0.066	0.0137

Table 10: Main results of the relationship between the uncertainty estimation methods and NLG metrics on TofuEval dataset using generation from GPT-3.5.

Uncertainty Estimation Methods	EI	MR	SOAF	RE	TME	CO	NMS	Mean
NumSets	0.1929	0.0414	-0.1774	-0.3319	-0.2410	0.0287	-0.0487	-0.0766
ECC	0.1619	-0.4640	-0.7079	0.2717	0.0072	-0.1679	-0.2993	-0.1712
LexSim	0.3119	0.1317	0.5493	0.3579	0.2672	-0.1781	0.0264	0.2095
EigV	0.2887	-0.2198	-1.1169	0.1539	0.1343	-0.1724	-0.3462	-0.1826

Table 11: Main results of the relationship between the uncertainty estimation methods and human annotations on TofuEval dataset using generation from GPT-3.5.

because no row consistently shows positivity across all dimensions.

- UniEval (Coherence) shows a large positive PRR (0.6794) with SOAF. However, it also shows a negative PRR (-0.4737) with RE.
- All NLG metrics show negative PRR values in terms of row mean. This means that the NLG metrics do not show good consistency with the human annotation.

From Table 11, we can obtain below conclusions.

- The ROUGE-L and CTC exhibit similar performance based on human annotation, while SummaC demonstrates dissimilar performance compared to these two.
- The Spearman and Kendall-Tau perform the same from the perspective of human annotation.
- The five dimensions (relevance, consistency, coherence, fluency, overall) of UniEval perform weak correlations among these five dimensions from the perspective of human annotation. However, the four dimensions (consistency, coherence, fluency, overall) of wi-ingt-GPT-3.5 show positive correlations from the perspective of human annotation.

C.2.4 UE-NLG Experimental Results

Table 10 and Figure 39 show the UE-NLG experimental results.

From Table 4, we can conclude below.

- It is hard for any uncertainty estimation method to achieve positive PRR for all the NLG metric methods.
- LexSim performs the best among the NLG metrics, consistent with the UE-HUM experimental results. Thus, NLG metrics remain useful for identifying the optimal uncertainty estimation method, which may align with human annotation.

Based on Figures 11 and 39, we found that rankings of certain uncertainty estimation methods differ between NLG metrics and human annotation. For instance, NumSets shows a negative correlation with EigV according to NLG metrics, while a positive correlation is observed based on human annotation. However, utilizing NLG metrics as an oracle remains meaningful, as it sometimes yields similar results to using human annotation as the oracle.

Spearman Correlation Between Uncertainty Estimation Methods on TofuEval by Generation from GPT-3.5

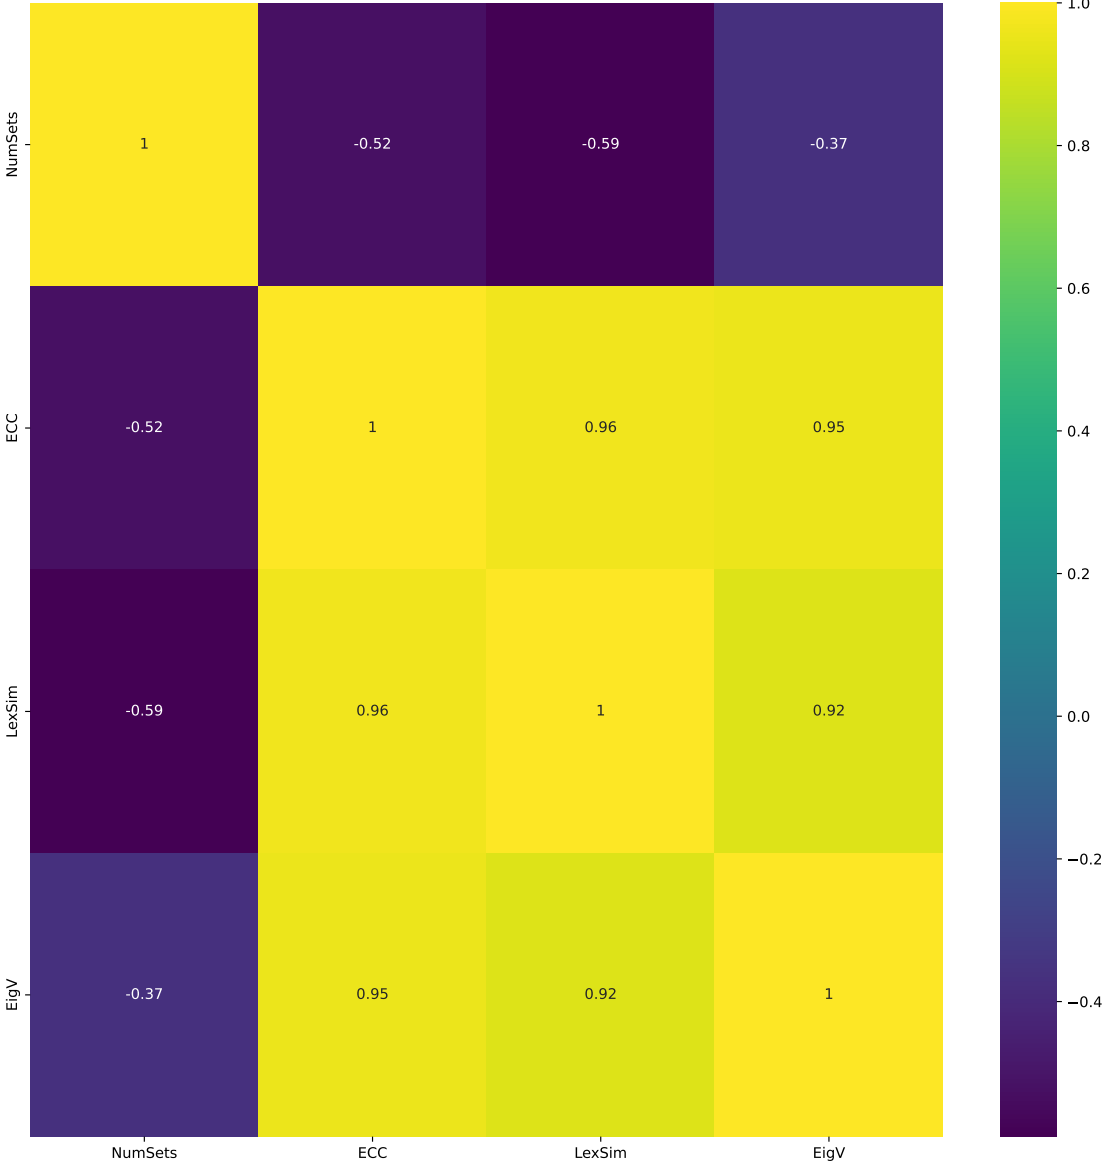


Figure 39: Diagram of Spearman correlation between uncertainty estimation on TofuEval dataset from the view of NLG metrics. The generated summaries are from GPT-3.5.

Spearman Correlation Between Uncertainty Estimation Methods on TofuEval by Generation from GPT-3.5

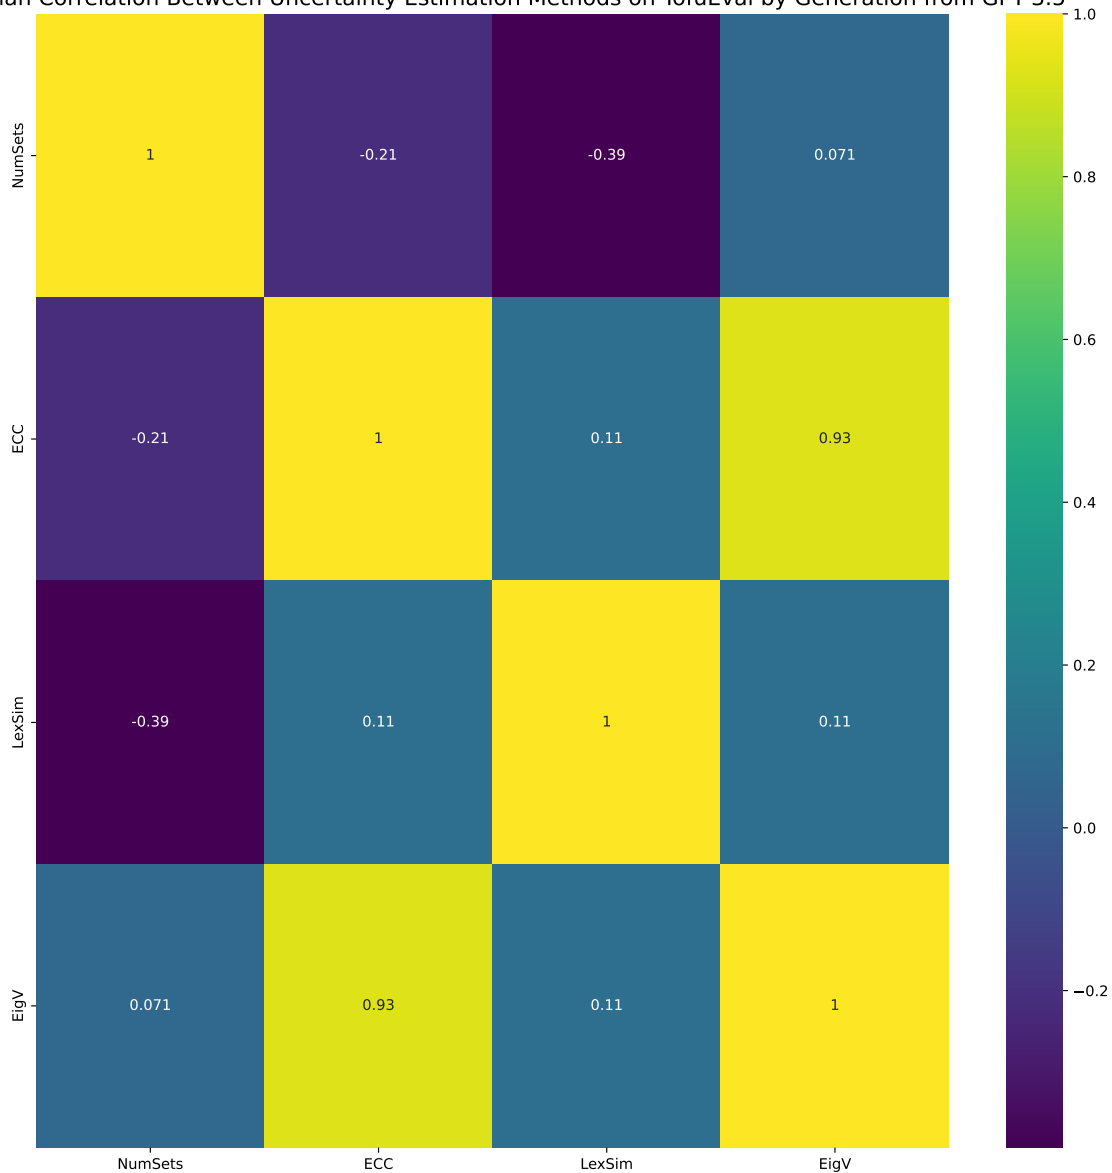


Figure 40: Diagram of Spearman correlation between uncertainty estimation on TofuEval dataset from the view of human annotation. The generated summaries are from GPT-3.5.

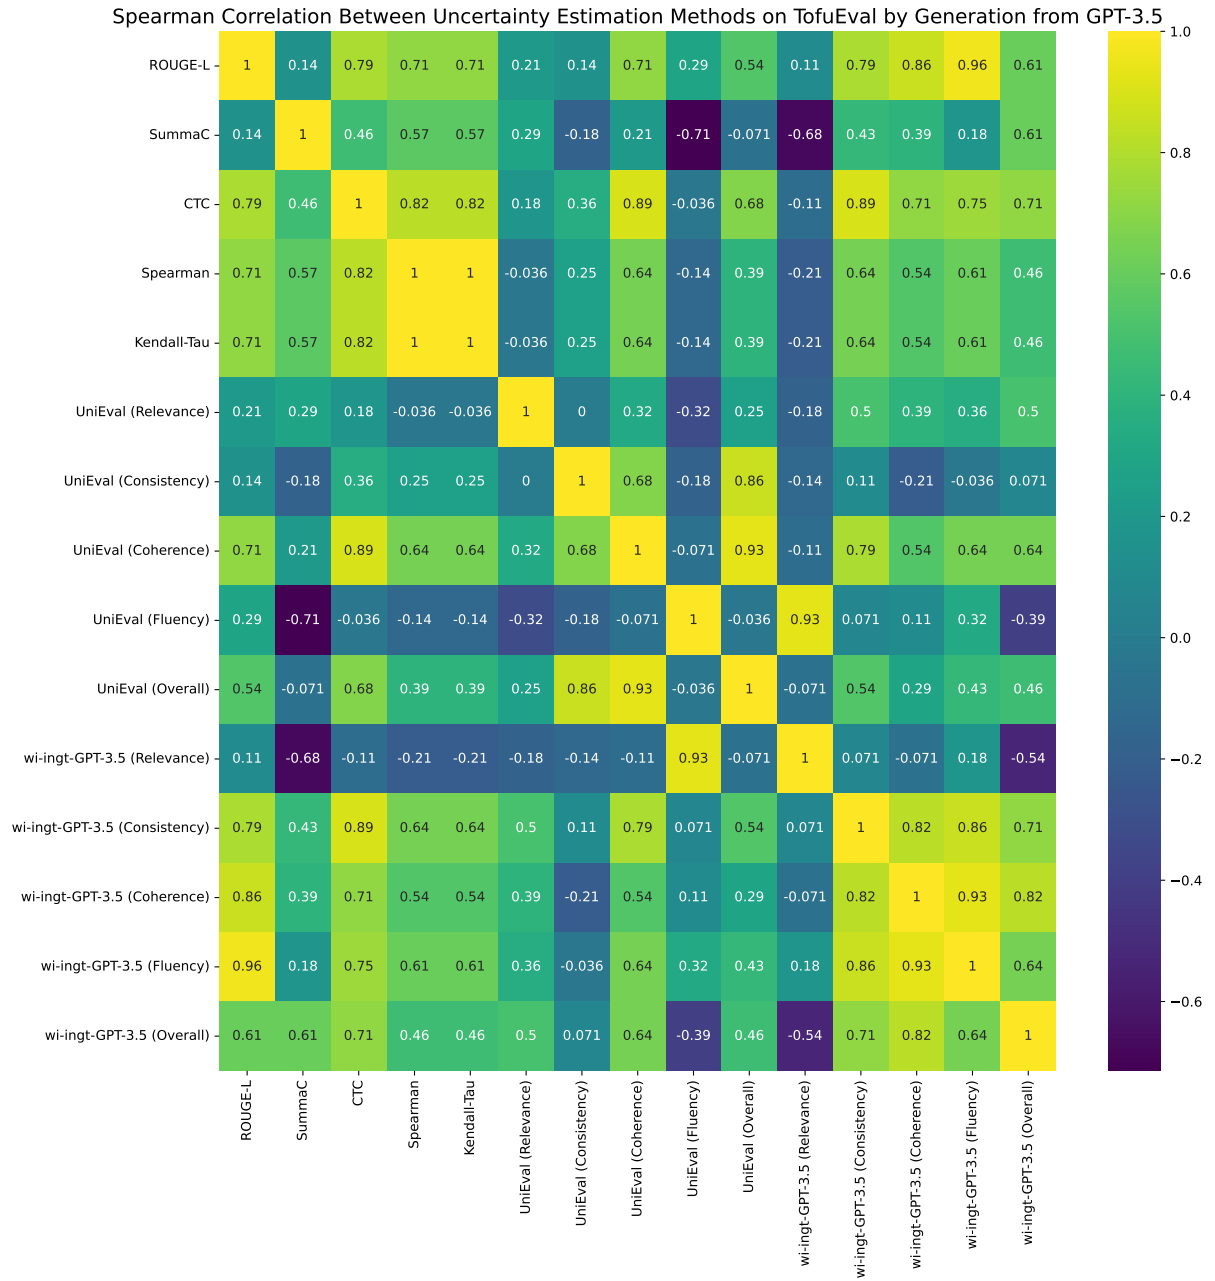


Figure 41: Diagram of Spearman correlation between NLG metrics on TofuEval dataset from the view of human annotation. The generated summaries are from GPT-3.5.