

# Self-Repetition in Abstractive Neural Summarizers

Nikita Salkar<sup>1</sup>, Thomas Trikalinos<sup>2</sup>, Byron C. Wallace<sup>1</sup>, Ani Nenkova<sup>3</sup>

<sup>1</sup>Khoury College of Computer Sciences, Northeastern University, USA

<sup>2</sup>Health Services, Policy and Practice, Brown University, USA

<sup>3</sup>Adobe Research, USA

{salkar.n,b.wallace}@northeastern.edu, thomas\_trikalinos@brown.edu, nenkova@adobe.com

## Abstract

We provide a quantitative and qualitative analysis of self-repetition in the output of neural summarizers. We measure self-repetition as the number of  $n$ -grams of length four or longer that appear in multiple outputs of the same system. We analyze the behavior of three popular architectures (BART, T5 and Pegasus), fine-tuned on five datasets. In a regression analysis, we find that the three architectures have different propensities for repeating content across output summaries for inputs, with BART being particularly prone to self-repetition. Fine-tuning on more abstractive data, and on data featuring formulaic language, is associated with a higher rate of self-repetition. In qualitative analysis we find systems produce artefacts such as ads and disclaimers unrelated to the content being summarized, as well as formulaic phrases common in the fine-tuning domain. Our approach to corpus level analysis of self-repetition may help practitioners clean up training data for summarizers and ultimately support methods for minimizing the amount of self-repetition.

## 1 Introduction

Sequence-to-sequence neural models for conditional text generation such as BART (Lewis et al., 2019), T5 (Raffel et al., 2020), and Pegasus (Zhang et al., 2020) achieve strong empirical results on abstractive summarization tasks. The summaries that such systems output often appear to be novel, in that they repeat text verbatim from inputs sparingly or not at all. Here, we set out to study the novelty of models with respect to their own outputs, by measuring the extent to which the content a model generates is formulaic repetition produced *across inputs*.

More specifically, we analyze how often long  $n$ -grams (length  $\geq 4$ ) appear in at least two summaries for different inputs. Repetition of some such  $n$ -grams may be natural, for example in news covering the same type of event, or in academic

papers with accepted formulaic descriptions of research questions and findings. To contextualize our measurements, we therefore contrast repetition in summaries written by humans with what we observe in system outputs. The former provides a baseline expectation regarding how much repetition is normal in a particular domain. In three out of the five domains we study we find that long  $n$ -gram repetition is considerably higher in automatically produced summaries than in human-written summaries. In the fourth domain, scientific papers, self-repetition even in human summaries is so high that the measure we use may not be sensitive enough to distinguish differences in repetition at this range.

We hypothesized that such undesirable behavior can be easier to quantify when we evaluate systems across domains, tasking a system trained in one domain to generate summaries in another. The intuition was that the repeated  $n$ -grams will be typical for the fine-tuning domain but rare in the test domain, so problematic repetitions may be easier to detect. This setting leads to clear cases of hallucinations reflecting the training data, e.g., fine-tuning BART (Lewis et al., 2019) on an academic paper summarization dataset and then applying it to a news summarization task yields hundreds of generated summaries that contain the phrase *this paper reports the results of an investigation*. Further, the phrase *The past few years have seen a dramatic increase* appears in a dozen news summaries, as do slight variations. Table 1 shows more examples of self-repetition and Section 5 describes the details of our qualitative analysis of  $n$ -grams identified by manually scanning repeated  $n$ -grams that clearly do not match the domain of text for which the summaries were generated.

To characterize this repetition behavior quantitatively, we perform a regression analysis in which we include as predictors system architecture, as well as training and test datasets (Section 6). We find that BART (Lewis et al., 2019) is especially

Repeating n-gram	Freq
click here for all the latest transfer news	73/11490
<b>Example:</b> Moha El Ouriachi is set to sign for Stoke City, according to his agent. The 19-year-old Barcelona B player is keen to seek first-team action. Stoke have already signed Bojan Krkic and Marc Muniesa from Barcelona. <a href="#">Click here for all the latest transfer news.</a>	
this paper reports the results of an investigation	143/11490
<b>Example:</b> schoolgirl killer Zbigniew Huminski was arrested for a range of crimes which are likely to see him jailed for life . <a href="#">this paper reports the results of an investigation</a> into the circumstances under which he was arrested in the northern port city of Calais	
In our series of letters from African-American journalists, film-maker and columnist Farai Sevenzo considers	16/11490
<b>Example:</b> <a href="#">In our series of letters from African-American journalists, film-maker and columnist Farai Sevenzo considers</a> the lessons learned from the 2013 Boston Marathon bombings.	
however, there is insufficient evidence to	1086/6440
<b>Example:</b> @xmath3 is an effective solution for the vacuum state of qcd . <a href="#">However, there is insufficient evidence to</a> support or refute the use of lattice simulations with @xmath3.	
but there is a lack of evidence to support	103/6440
<b>Example:</b> The Apple Watch is officially going on sale - <a href="#">but there is a lack of evidence to support</a> its decision to make it available through online orders.	

Table 1: Examples of self-repetition.

prone to self-repetition, more so than the other architectures we consider, and that the type of training data used to fine tune the sequence-to-sequence model for summarization has a considerable impact on the propensity of models to repeat themselves.

Our work highlights a dimension of repetition and novelty in summarization that, to our knowledge, has not been explored previously. The repetition metrics we introduce may be broadly useful in characterizing the performance of new abstractive summarization systems, as we show that models differ markedly with respect to these measures.

## 2 Related work

Prior work in abstractive neural summarization has focused on phrases repeated *within a given output*, and proposed various means for mitigating this problem (See et al., 2017; Paulus et al., 2018; Fu et al., 2021; Nair and Singh, 2021). By contrast, our work quantifies the extent to which systems produce *the same n-grams across different inputs*, and the factors that correlate with this behavior.

Research in text generation has documented that

systems often self-repeat and have quantified how much models repeat content from their pre-training data (McCoy et al., 2021; Carlini et al., 2022). We provide some puzzling examples where we are unable to trace the origin of repeated content<sup>1</sup>. We also recognize a portion of the repetitions as hallucinations that are influenced by the training data. Oftentimes, the hallucinations are stylistic, similar to the formulaic phrases from academic papers that we mentioned in the introduction. Prior work has shown that neural summarization systems are capable of choosing important content across domains but need in-domain data to faithfully reproduce the style of a given domain (Hua and Wang, 2017). In our work, we find that once systems pick up stylistic templates from one domain, they are likely to reuse them in other domains, where the formulaic phrases look out of place.

Self-repetition is well-documented in dialog systems research. Dialog systems often produce generic formulaic responses regardless of the preceding utterance (Li et al., 2016): in one of the reported experiments, four generic responses (*I don't know*, *I don't know what you are talking about*, *I don't think this is a good idea*, *Oh my god*) constitute 32% of system generated responses. These phrases were common in the training data, with 0.4% of training data sentences containing the phrase *I don't know*, even though overall the training data was diverse. Our findings for summarization are similar, as we discover in our regression analysis that training on data with higher incidence of formulaic phrases, like academic papers and summaries of medical evidence, results in a summarizer that is overall more likely to repeat content across inputs, at rates markedly higher than done by humans.

Human summaries are typically considered an appropriate reference while enhancing abstractive text summarization models (Yang et al., 2019, 2020). For our analysis too, we contrast model generated summaries against the human summaries as baseline to determine the threshold over which self-repetition is considered anomalous.

## 3 Defining Self-Repetition

We introduce a *repetition score* to measure how often systems repeat themselves. The score is a

<sup>1</sup>Recently developed techniques for attributing content in a summary to the language model or the input (Xu and Durrett, 2021) would be more powerful than the manual inspection we carried out and will support future work on self-repetition.

function of  $n$ -grams of length four and longer in different summaries, which is indicative of text similarity and potential plagiarism (Lyon et al., 2001). We consider an  $n$ -gram to be *repeating* when it appears in two or more summaries in a dataset. The repetition score can be computed at the *dataset* and *individual summary* level.

At the dataset level, we count the number of summaries that contain at least one  $n$ -gram ( $n \geq 4$ ) that also appears in another summary. We define the repetition score for a dataset as the number of summaries containing repeating  $n$ -grams divided by the total number of summaries in that dataset. We divide by the total in order to normalize the values allowing for meaningful comparison between datasets of different sizes.

For an individual summary, we define the repetition score as:

$$R_i = \log\left(\sum_{k=1}^m N_k + 1\right). \quad (1)$$

Where  $i$  indexes summaries,  $m$  is the number of repeating  $n$ -grams in summary  $i$ , and  $N_k$  denotes the count of summaries that contain the  $k$ th repeating  $n$ -gram found within summary  $i$ . We take the log to this value to produce the final score, to make the repetition score less sensitive to outliers.

## 4 Models and Datasets

We consider three models: BART (Lewis et al., 2019), T5 (Raffel et al., 2020), and Pegasus (Zhang et al., 2020), each fine-tuned on five summarization datasets: CNN/DailyMail (Hermann et al., 2015), BBC XSum (Narayan et al., 2018), Scientific Papers (SP; Cohan et al. 2018), Reddit (Völske et al., 2017) and a corpus of Randomized Controlled Trials (RCTs; Wallace et al. 2021). We evaluate each model on the five datasets, yielding 75 (3·5·5) combinations of architectures, train, and test datasets.

Table 2 reports repetition scores for each architecture on the datasets considered. To contextualize these, we also report repetition scores for the reference (i.e., human-written) summaries. Reddit shows the least amount of human repetition; only 27% of summaries contain at least one  $n$ -gram of length four or greater that also appears in another Reddit summary. Scientific Papers are the most formulaic: 99% of abstracts contain such repetition. The RCTs data (also scientific in nature) is similarly repetitive. News—from both CNN/Daily

Dataset	Human	BART	T5	Pegasus
CNN/DailyMail	0.69	0.96	0.90	0.80
XSum	0.60	0.85	0.70	0.81
Reddit	0.27	0.26	0.28	0.29
Scientific Papers	0.99	0.99	0.99	0.99
RCT	0.88	1.0	0.96	1.0

Table 2: Repetition scores for human and in-domain system summaries produced with different architectures.

Dataset	Unigram	Bigram	Trigram	4-gram
CNN/DailyMail	30.20	54.40	71.53	79.99
XSum	40.40	81.47	91.47	93.64
Reddit	9.50	2.71	2.53	2.77
SP	48.41	49.99	70.08	81.48
RCT	52.56	77.87	92.02	96.08

Table 3: Percent abstractiveness of human summaries.

Mail and XSum—is somewhere in-between: 60–70% of human summaries contain a long repeated  $n$ -gram.

In model outputs we observe a level of repetition similar to what is seen in the references on the Reddit and Scientific Papers dataset. For news corpora (CNN and XSum) and the medical evidence summarization task (RCTs) however, system repetition scores are markedly higher than the scores for the human-written summaries. BART seems particularly prone to repetition.

We contrast the repetition score of the human summaries in each domain with their level of *abstractiveness*, defined as the fraction of  $n$ -grams of a given size that *do not* appear in the input (and so are “novel”). As pointed out in (Narayan et al., 2018), reference summaries in XSum are more abstractive than those in the CNN/Daily Mail dataset. Table 3 also highlights that Reddit summaries are particularly extractive, e.g., bi-grams in references almost always appear in the corresponding inputs. Aside from Reddit, the number of novel with respect to the input  $n$ -grams increases with  $n$ .

## 5 Qualitative Analysis

To glean a qualitative view of repetition behavior, we randomly sampled 20 long  $n$ -grams that appeared in more than 10 summaries. These  $n$ -grams often do not appear in the corresponding inputs.

We show examples in Table 1. The first  $n$ -gram is generated in 73 out of 11,490 summaries by a Pegasus model fine-tuned on CNN/Daily Mail and applied to test instances from the same domain; there is no domain shift here. This  $n$ -gram does not occur in the train or the test set.

Repetition is particularly pronounced when the

**Summary:** In our series of letters from African journalists, filmmaker and columnist Ahmedou Ould-Abdallah reflects on his time at the University of Cape Town.

**Input:** This is great. I hope you discuss some of these points in the next episode of TLDR: CLG.

**Summary:** In our series of letters from African journalists, filmmaker and columnist GustavoM looks at the relationship between humans and animals.

**Input:** Listen to GustavoM, a friendship could never come close to that of a companion.tl;dr: GustavoM is right.

**Summary:** In our series of letters from African journalists, filmmaker and columnist Ahmedou Ould-Abdallah reflects on his time at the University of Cape Town.

**Input:** We had to take business writing classes to graduate in finance. It was essentially a class on how to do effective TL;DRs.

Table 4: Examples of hallucinations in summaries.

model is trained to summarize data for one domain and then applied to another. For example, the second  $n$ -gram shown (*"this paper reports the results of an investigation"*) was repeated in 143/11,490 summaries generated by a BART model trained on Scientific Papers and then applied to CNN/DailyMail inputs. This  $n$ -gram also appears in two out of 203,037 training inputs of Scientific Papers with its sub- $n$ -grams appearing with even greater frequency.

The next  $n$ -gram is found in 16 out of 11,490 summaries produced by a BART model trained on XSum and applied to CNN/Daily Mail. This  $n$ -gram does not appear in the XSum train set; moreover, there is no mention of "Farai Sevenzo" in the CNN dataset at all (inputs or outputs). While these examples contain summaries that are at least related to the input, Table 4 shows examples of hallucinating summaries generated by Pegasus trained on XSum and then applied to Reddit. The  $n$ -gram "In our series of letters from African journalists, filmmaker and columnist" occurs in the generated summaries without having any relevance to the input. These examples indicate that models sometimes produce formulaic content unrelated to inputs, which may not even have been encountered in the training data.

## 6 Regression Analysis

We next quantify the association between self-repetition and factors that might influence this, including system architecture and pre-training, and the datasets used for training and testing. We would also expect that repetition would be proportional

	Coef	P>  t	[0.025	0.975]
<b>Intercept</b>	1.94	0.00	1.91	1.97
<b>Length of Summary</b>	0.35	0.00	0.34	0.36
<b>BART</b>	1.79	0.00	1.77	1.82
<b>T5</b>	-0.11	0.00	-0.13	-0.09
<b>Pegasus</b>	-0.02	0.07	-0.05	0.00
<b>Train SP</b>	1.43	0.00	1.40	1.46
<b>Train RCT</b>	2.28	0.00	2.25	2.31
<b>Train Reddit</b>	-0.37	0.00	-0.40	-0.34
<b>Train XSum</b>	0.24	0.00	0.20	0.27
<b>Test SP</b>	0.55	0.00	0.52	0.59
<b>Test RCT</b>	-0.95	0.00	-1.06	-0.84
<b>Test Reddit</b>	-0.52	0.00	-0.55	-0.49
<b>Test XSum</b>	-0.37	0.00	-0.40	-0.34
<b>RCT - SP</b>	2.90	0.00	2.85	2.95
<b>RCT - RCT</b>	2.41	0.00	2.25	2.56
<b>RCT - Reddit</b>	0.40	0.00	0.36	0.44
<b>RCT - XSum</b>	-0.07	0.00	-0.11	-0.03
<b>Reddit - SP</b>	0.60	0.00	0.54	0.65
<b>Reddit - RCT</b>	0.40	0.00	0.24	0.56
<b>Reddit - Reddit</b>	-0.71	0.00	-0.75	-0.67
<b>Reddit - XSum</b>	0.33	0.00	0.29	0.38
<b>SP - SP</b>	0.51	0.00	0.45	0.56
<b>SP - RCT</b>	-0.45	0.00	-0.61	-0.29
<b>SP - Reddit</b>	0.49	0.00	0.45	0.53
<b>SP - XSum</b>	0.15	0.00	0.11	0.20
<b>XSum - SP</b>	0.66	0.00	0.61	0.71
<b>XSum - RCT</b>	0.81	0.00	0.65	0.97
<b>XSum - Reddit</b>	0.44	0.00	0.40	0.48
<b>XSum - XSum</b>	0.11	0.00	0.07	0.16

Table 5: Regression results; detailed descriptions of predictors are in the Appendix.

to summary length: More words naturally afford more opportunities for repetition, even if by chance. And indeed we observe that the repetition scores of human summaries are proportional to their average lengths. We report summary lengths in Appendix Table A1 which can be compared to the repetition scores in Table 2. Model generated summaries exhibit a similar correlation.

We also hypothesized that domain shift — e.g., testing a model trained to summarize scientific texts on news articles — would increase repetition across summaries (the model may default to stock phrases in such cases). We provide qualitative examples of this in Section 5.

We fit a regression model to 731,406 summaries generated by 75 combinations of architecture, train and test data, along with the reference summaries for all datasets. We have multiple one-hot encoded categorical variables, which means we must select reference categories for these (effectively the intercept term). We use human generated summaries as the reference architecture and the CNN/Daily Mail as the reference train and test sets.

This model treats the repetition observed in a given summary as defined in Equation 1 as a lin-



ear function of predictors including: the length of the generated summary in number of white space delimited tokens (Length of Summary); the model architecture used to generate the summary. Differences in pre-training data will be folded in the behavior due to architecture (BART, T5, Pegasus); the training data to which this model was fit; the test data for which a summary is produced; and interaction terms between train and test datasets. The latter we denote by “TRAIN - TEST”, e.g., “XSum - Reddit” indicates a summary produced by a model fine-tuned on XSum given an input drawn from the Reddit corpus. This is a *cross-domain* model. By contrast, “XSum - XSum” is an *in-domain* example of a summary produced on an XSum test instance by a model fine-tuned using the XSum training data. Table 5 enumerates all covariates (more details in the Appendix).

Table 5 reports results from this analysis. We make a few key observations here. First, it would seem BART is most prone to repetition of the models considered. From the average summary lengths reported in Appendix Table A1, we observe that the BART summaries on CNN/DailyMail are almost double the length of human summaries. This suggests the possibility that the observed tendency of BART to disproportionately produce repetitions may owe to the fact that it is prone to producing lengthier summaries in general. To investigate this, we imposed a restriction on the *max-length* while decoding — specifically we set this to 50, which falls between the average lengths of T5 and Pegasus of each corresponding model (Appendix Table A2). This resulted in BART yielding summaries that are shorter (on average) than those of T5 and Pegasus. Table A2 shows the regression results when the analysis performed with these shortened BART summaries. This does shrink the coefficient for BART by a small amount, but it remains by far the largest (compared to T5 and Pegasus). This indicates that while the summary length may somewhat influence the overall repetition, BART seems prone to this behavior independent of its tendency to produce lengthier outputs.

In Table 5, among the source data, RCT has the maximum amount of repetition in comparison to the baseline CNN DailyMail followed by Scientific Papers and XSum, which aligns with the results of Table 2. Among the test set, Scientific Papers is the only corpus to have an influence on the repetition. The interaction terms yield higher coefficients

when the training data is Scientific Papers or Randomized Controlled Trials in comparison to when the train source is XSum or Reddit. Further, for all the training datasets, the highest values are for when the test data is Scientific Papers or RCT.

To ascertain whether domain shift (in general) is indeed a significant factor associated with repetition, we perform a likelihood ratio test with the interaction terms. Specifically we use as our nested model a regression with all interaction terms omitted, and compare this to the full model with all factors. We choose 0.001 as the critical value. The likelihood ratio test results in a p-value of  $\ll 0.001$ . This implies that the domain interactions do impart information in terms of quantifying the self-repetition, i.e., applying a summarization model to data from a domain that differs from its training source correlates with increased repetition.

## 7 Conclusions

We evaluated the tendency of neural summarization models to *repeat* themselves *across outputs* on five datasets. To our knowledge this is the first analysis of this phenomenon. Our results indicate that BART has the greatest tendency to self-repeat, and that the training source is a significant factor which may lead to this repetition behavior. Adapting a summarization model trained on one domain to another (distinct) domain also correlates significantly with repetition; the model may “not know what to say” in such cases, and default to stock phrases from the training data. We also found that models sometimes repeat long strings of text that do not contain any references in the corresponding inputs or even the training sets. These may originate in pre-training data, but more research into such hallucinations is warranted. We hope this analysis will encourage development of methods for mitigating the repetition across summaries and for controlling hallucinations in abstractive neural summarizers.

## Acknowledgements

We are grateful to Tracy King for her careful reading of and detailed comments on an earlier version of this paper.

This research was supported in part by the National Institutes of Health (NIH) under the National Library of Medicine (NLM) grant 2R01LM012086, and in part by the National Science Foundation (NSF) under grant 2211954.

## References

- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O'Reilly Media, Inc."
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2022. Quantifying memorization across neural language models. In *CoRR*, vol. *abs/2202.07646*.
- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*.
- Zihao Fu, Wai Lam, Anthony Man-Cho So, and Bei Shi. 2021. A theoretical analysis of the repetition problem in text generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28.
- Xinyu Hua and Lu Wang. 2017. [A pilot study of domain adaptation effect for neural abstractive summarization](#). In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 100–106, Copenhagen, Denmark. Association for Computational Linguistics.
- 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. In *Association for Computational Linguistics (ACL)*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Caroline Lyon, James Malcolm, and Bob Dickerson. 2001. [Detecting short passages of similar text in large document collections](#). In *Proceedings of the 2001 Conference on Empirical Methods in Natural Language Processing*.
- R Thomas McCoy, Paul Smolensky, Tal Linzen, Jianfeng Gao, and Asli Celikyilmaz. 2021. How much do language models copy from their training data? evaluating linguistic novelty in text generation using raven. In *CoRR*, *abs/2111.09509*.
- Pranav Nair and Anil Kumar Singh. 2021. On reducing repetition in abstractive summarization. In *Proceedings of the Student Research Workshop Associated with RANLP 2021*, pages 126–134.
- 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. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *6th International Conference on Learning Representations, Vancouver, BC, Canada*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Skipper Seabold and Josef Perktold. 2010. statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*.
- Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. [TL;DR: Mining Reddit to learn automatic summarization](#). In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.
- Byron C. Wallace, Sayantan Saha, Frank Soboczenski, and Iain J. Marshall. 2021. Generating (Factual?) Narrative Summaries of RCTs: Experiments with Neural Multi-Document Summarization. In *Proceedings of AMIA Informatics Summit*.
- Jiacheng Xu and Greg Durrett. 2021. [Dissecting generation modes for abstractive summarization models via ablation and attribution](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6925–6940, Online. Association for Computational Linguistics.
- Min Yang, Chengming Li, Ying Shen, Qingyao Wu, Zhou Zhao, and Xiaojun Chen. 2020. Hierarchical human-like deep neural networks for abstractive text summarization. *IEEE Transactions on Neural Networks and Learning Systems*, 32(6):2744–2757.
- Min Yang, Qiang Qu, Wenting Tu, Ying Shen, Zhou Zhao, and Xiaojun Chen. 2019. Exploring human-like reading strategy for abstractive text summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.

## A Appendix

### A.1 Regression Model Details

The dataset for the regression model comprises 731406 summaries, generated by the 75 (3·5·5) combinations of architectures, train and test datasets. The predictors corresponding to each summary  $i$  and the observed repetition score  $R_i$  constitutes an  $(x_i, y_i)$  pair. More specifically, “ $x_i$ ” is composed of the features of the summary we use in our analysis, which we describe individually below. Note that some of our predictors (those related to architectures and datasets) are categorical, and so need to be “one-hot” encoded. In such cases, one option must serve as a reference category with respect to which the remaining coefficients can be interpreted.

Regarding these categorical variables: We analyze four architectures for producing summaries — including “Human” in addition to BART, T5 and Pegasus. “Human” serves as our reference architecture, so we do not have an explicit coefficient for this. Similarly, we include five datasets in our analysis; for any summary one dataset will have served as the training source and another as the source of test inputs. We use CNN/Daily Mail as the reference category for both of these categorical predictors.

Because we are interested in the effects of applying models trained on one summarization domain to another, we also include “interaction terms” that encode pairs of train/test datasets via indicators. As such, we one-hot encode all pairwise interactions between our four datasets.

We estimate coefficients to these predictors given the observed summary data in an Ordinary Least Squares (OLS) linear regression model, as implemented the `statsmodels` (v0.12.2) Python module (Seabold and Perktold, 2010).

**Details about regression predictors** We discuss the individual terms in our regression (coefficients for which are reported in Table 5) in greater detail below.

- **Length of Summary** This is the number of words in a summary extracted by the NLTK word tokenizer (Bird et al., 2009). Because

lengths are quite variable, we standardize the length using the Z-score normalization. The value of 0.35 in the analysis suggests a positive correlation between the length of a summary and the amount of repetition which also corroborates our observations from Table A1 and Table 2.

- **Human** This denotes the special neural “architecture” responsible for generating the reference summaries: Humans. Recall that “humans” serve as our reference architecture category for one-hot encoding, so are folded into the intercept term.
- **BART** This denotes the summaries generated by the BART architecture (Lewis et al., 2019). The somewhat large positive coefficient (1.79) indicates BART is particularly prone to generating repetitions across its outputs.
- **T5** This denotes the summaries generated by the T5 architecture (Raffel et al., 2020). Overall, our regression results suggest that in aggregate T5 is about comparable to humans in terms of its tendency to repeat itself in general, although it is also subject to this in domain adaptation settings (as are all models considered).
- **Pegasus** This denotes the summaries generated by the Pegasus architecture (Zhang et al., 2020). The interpretation of the corresponding coefficient here is similar to for T5.
- **Train CNN/Daily Mail** This indicates summaries produced by models *trained* on the CNN/Daily Mail dataset (Hermann et al., 2015). CNN/Daily Mail serves as our reference for this categorical feature, and so we do not have an explicit coefficient for it.
- **Train SP** Indicates a summary produced by a model *trained* on the Scientific Papers dataset (Cohan et al., 2018). The positive coefficient (1.43) suggests that in aggregate models trained on Scientific Papers are more prone to repeat than those trained on CNN/Daily Mail dataset.
- **Train RCT** Indicates a summary produced by a model *trained* on the Randomized Controlled Trials (RCTs) dataset (Wallace et al.,

Dataset	Train / Val / Test	Input Document	Human Summary	BART	T5	Pegasus
CNN/Daily Mail	287113 / 11338 / 11490	683.51	52.12	103.37	58.01	53.16
XSum	204045 / 11332 / 11334	360.58	21.09	19.34	20.13	17.86
Reddit	67198 / 16800 / 16000	222.66	21.06	19.54	21.69	22.98
Scientific Papers	203037 / 6436 / 6440	5702.14	163.13	96.52	81.09	97.37
RCT	3721 / 464 / 466	2689.83	68.15	22.68	58.75	39.64

Table A1: Average lengths of test inputs, the corresponding human summaries, and model-generated summaries.

	Coef	P>  t	[0.025	0.975]
<b>Intercept</b>	1.64	0.01	122.87	0.00
<b>Length of Summary</b>	0.43	0.00	129.30	0.00
<b>BART</b>	1.60	0.01	129.28	0.00
<b>T5</b>	-0.13	0.01	-10.52	0.00
<b>Pegasus</b>	-0.04	0.01	-3.29	0.00
<b>Train SP</b>	1.69	0.02	105.39	0.00
<b>Train RCT</b>	2.60	0.02	165.89	0.00
<b>Train Reddit</b>	-0.09	0.02	-5.40	0.00
<b>Train XSum</b>	0.65	0.02	41.28	0.00
<b>Test SP</b>	0.52	0.02	27.68	0.00
<b>Test RCT</b>	-0.77	0.06	-13.17	0.00
<b>Test Reddit</b>	-0.51	0.01	-35.13	0.00
<b>Test XSum</b>	-0.28	0.02	-17.84	0.00
<b>RCT - SP</b>	2.97	0.03	110.00	0.00
<b>RCT - RCT</b>	2.24	0.08	28.60	0.00
<b>RCT - Reddit</b>	0.41	0.02	19.50	0.00
<b>RCT - XSum</b>	-0.13	0.02	-5.76	0.00
<b>Reddit - SP</b>	0.63	0.03	23.42	0.00
<b>Reddit - RCT</b>	0.25	0.08	3.05	0.00
<b>Reddit - Reddit</b>	-0.60	0.02	-28.07	0.00
<b>Reddit - XSum</b>	0.26	0.02	11.21	0.00
<b>SP - SP</b>	0.45	0.03	16.97	0.00
<b>SP - RCT</b>	-1.03	0.08	-12.41	0.00
<b>SP - Reddit</b>	0.51	0.02	24.37	0.00
<b>SP - XSum</b>	0.11	0.02	4.61	0.00
<b>XSum - SP</b>	0.69	0.03	25.60	0.00
<b>XSum - RCT</b>	0.63	0.08	7.61	0.00
<b>XSum - Reddit</b>	0.43	0.02	20.63	0.00
<b>XSum - XSum</b>	0.01	0.02	0.41	0.68

Table A2: Regression results after restricting length of BART summaries.

2021). The positive coefficient (2.28) suggests that training on this dataset results in comparatively large amount of repetition.

- **Train Reddit** Indicates a summary produced by a model *trained* on the Reddit dataset (Völske et al., 2017). The small negative coefficient value of -0.37 indicates that models trained on Reddit are somewhat less prone to repetition, on average.
- **Train XSum** Indicates a summary produced by a model *trained* on the XSum dataset (Narayan et al., 2018). The small positive coefficient estimate of 0.24 implies that models trained on XSum may repeat slightly

more than those trained on CNN/Daily Mail, on average.

- **Test CNN/Daily Mail** Indicates that the corresponding summary was generated for an instance drawn from the CNN/Daily Mail test set. We again treat this as the reference category.
- **Test SP** Indicates that the corresponding summary was generated for an instance drawn from the Test SP test set. The small positive value of 0.55 suggests that evaluating models on Scientific Paper instances correlates with a greater amount of repetition.
- **Test RCT** Indicates that the corresponding summary was generated for an instance drawn from the Test RCT test set. The negative value of -0.95 indicates that when tested on RCT instances, models are slightly less prone to repetition.
- **Test Reddit** Indicates that the corresponding summary was generated for an instance drawn from the Reddit test set. The small negative value of -0.52 implies that when models are evaluated on Reddit instances they may tend to repeat themselves across summaries comparatively less frequently.
- **Test XSum** Indicates that the corresponding summary was generated for an instance drawn from the XSum test set. The negative coefficient of -0.37 implies a slightly lower tendency for repetition when models are tested on instances from the XSum test set.
- **RCT - SP** This denotes a summary produced by a model trained on the RCTs train set and evaluated on Scientific Papers test set; a cross-domain scenario. The estimate coefficient of



Train	Test	Bart	T5	Pegasus
CNN/Daily Mail	CNN/Daily Mail	103.37	58.00	53.16
	XSum	65.2	45.59	44.26
	SP	92.33	63.43	45.05
	Reddit	91.48	44.92	52.35
	RCT	77.63	39.78	43.99
XSum	CNN/Daily Mail	21.69	23.06	19.11
	XSum	19.34	20.13	17.86
	SP	22.62	25.2	20.24
	Reddit	20.12	19.83	17.61
	RCT	22.72	20.13	19.32
SP	CNN/Daily Mail	69.51	83.39	95.33
	XSum	58.42	78.28	66.72
	SP	96.51	81.09	97.37
	Reddit	56.53	71.79	78.41
	RCT	83.46	46.69	66.12
Reddit	CNN/Daily Mail	54.06	84.04	78.96
	XSum	53.89	78.92	69.51
	SP	62.07	69.14	83.60
	Reddit	19.53	21.69	22.98
	RCT	44.15	46.41	92.51
RCT	CNN/Daily Mail	35.16	61.95	73.53
	XSum	28.92	62.61	48.73
	SP	28.71	45.38	49.48
	Reddit	24.60	60.40	62.59
	RCT	22.68	58.75	39.64

Table A3: The Average Lengths of Systems before restricting the max-length during BART decoding.

2.90 indicates that this combination of interaction yields a comparatively high amount of repetition.

- **RCT – RCT** This denotes a summary generated by a model trained and tested on the Randomized Controlled Trials. This is an in-domain scenario. A coefficient of 2.41 indicates that this combination of interaction also yields a much higher amount of repetition than the baseline train - test combination.
- **RCT – Reddit** This denotes a summary produced by a model trained on Randomized Controlled Trials and evaluated on Reddit. This is again a cross-domain scenario. A coefficient of 0.40 means that this combination has a negligibly higher self-repetition than the baseline. Similarly, for the rest.
- **RCT – XSum** Denotes a summary generated by a model trained on Randomized Controlled Trials and tested on XSum.
- **SP – SP** Denotes a summary generated by an in-domain model trained and tested on Scientific Papers.
- **SP – RCT** Denotes a summary produced by

Train	Test	Bart	T5	Pegasus
CNN/Daily Mail	CNN/Daily Mail	36.67	58.00	53.16
	XSum	36.61	45.59	44.26
	SP	38.36	63.43	45.05
	Reddit	38.52	44.92	52.35
	RCT	32.17	39.78	43.99
XSum	CNN/Daily Mail	21.69	23.06	19.11
	XSum	19.34	20.13	17.86
	SP	22.62	25.2	20.24
	Reddit	20.12	19.83	17.61
	RCT	22.72	20.13	19.32
SP	CNN/Daily Mail	69.51	83.39	95.33
	XSum	58.42	78.28	66.72
	SP	96.51	81.09	97.37
	Reddit	56.53	71.79	78.41
	RCT	35.33	46.69	66.12
Reddit	CNN/Daily Mail	54.06	84.04	78.96
	XSum	53.89	78.92	69.51
	SP	62.07	69.14	83.60
	Reddit	19.53	21.69	22.98
	RCT	44.15	46.41	92.51
RCT	CNN/Daily Mail	35.16	61.95	73.53
	XSum	28.92	62.61	48.73
	SP	28.71	45.38	49.48
	Reddit	24.60	60.40	62.59
	RCT	22.68	58.75	39.64

Table A4: The Average Lengths of Systems after restricting the max-length during BART decoding.

a trained on Scientific Papers and tested on Randomized Controlled Trials.

- **SP – Reddit** Denotes a summary produced by a model trained on Scientific Papers and tested on Reddit.
- **SP – XSum** Denotes a summary produced by a model trained on Scientific Papers and tested on XSum.
- **Reddit – SP** Denotes a summary generated by a model trained on Reddit and tested on Scientific Papers.
- **Reddit – RCT** Denotes a summary produced by a model trained on Reddit and tested on Scientific Papers.
- **Reddit – Reddit** Denotes a summary produced by an in-domain model trained and tested on Reddit.
- **Reddit – XSum** Denotes a summary produced by a model trained on Reddit and tested on XSum.
- **XSum – SP** Denotes a summary generated by a model trained on XSum and tested on Scientific Papers.

- **XSum – RCT** Denotes a summary produced by a model trained on XSum and tested on RCT.
- **XSum – Reddit** Denotes a summary produced by a model trained on XSum and tested on Reddit.
- **XSum – XSum** Denotes a summary produced by an in-domain model trained and tested on XSum.

Table A3 reports the average lengths of summaries generated by each system. We can see that when the training data is CNN/Daily Mail, BART has the highest average lengths. Further BART trained on Scientific Papers and applied to RCTs also have lengths higher than corresponding models.

We restrict the max-lengths of these systems to 50 which lies between the corresponding T5 and Pegasus models' average lengths. Table A4 depicts the average lengths after imposing the restrictions. From A2 we can see that shortening the lengths of BART summaries does not mitigate its tendency to repeat the most of all the models.