

# Liputan6: A Large-scale Indonesian Dataset for Text Summarization

Fajri Koto    Jey Han Lau    Timothy Baldwin

School of Computing and Information Systems  
The University of Melbourne

ffajri@student.unimelb.edu.au, jeyhan.lau@gmail.com, tbaldwin@unimelb.edu.au

## Abstract

In this paper, we introduce a large-scale Indonesian summarization dataset. We harvest articles from *Liputan6.com*, an online news portal, and obtain 215,827 document–summary pairs. We leverage pre-trained language models to develop benchmark extractive and abstractive summarization methods over the dataset with multilingual and monolingual BERT-based models. We include a thorough error analysis by examining machine-generated summaries that have low ROUGE scores, and expose both issues with ROUGE itself, as well as with extractive and abstractive summarization models.

## 1 Introduction

Despite having the fourth largest speaker population in the world, with 200 million native speakers,<sup>1</sup> Indonesian is under-represented in NLP. One reason is the scarcity of large datasets for different tasks, such as parsing, text classification, and summarization. In this paper, we attempt to bridge this gap by introducing a large-scale Indonesian corpus for text summarization.

Neural models have driven remarkable progress in summarization in recent years, particularly for abstractive summarization. One of the first studies was Rush et al. (2015), where the authors proposed an encoder–decoder model with attention to generate headlines for English Gigaword documents (Graff et al., 2003). Subsequent studies introduced pointer networks (Nallapati et al., 2016b; See et al., 2017), summarization with content selection (Hsu et al., 2018; Gehrmann et al., 2018), graph-based attentional models (Tan et al., 2017), and deep reinforcement learning (Paulus et al., 2018). More recently, we have seen the widespread adoption

of pre-trained neural language models for summarization, e.g. BERT (Liu and Lapata, 2019), BART (Lewis et al., 2020), and PEGASUS (Zhang et al., 2020a).

Progress in summarization research has been driven by the availability of large-scale English datasets, including 320K *CNN/Daily Mail* document–summary pairs (Hermann et al., 2015) and 100k *NYT* articles (Sandhaus, 2008) which have been widely used in abstractive summarization research (See et al., 2017; Gehrmann et al., 2018; Paulus et al., 2018; Lewis et al., 2020; Zhang et al., 2020a). News articles are a natural candidate for summarization datasets, as they tend to be well-structured and are available in large volumes. More recently, English summarization datasets in other flavours/domains have been developed, e.g. *XSum* has 226K documents with highly abstractive summaries (Narayan et al., 2018), *BIGPATENT* is a summarization dataset for the legal domain (Sharma et al., 2019), *Reddit TIFU* is sourced from social media (Kim et al., 2019), and Cohan et al. (2018) proposed using scientific publications from arXiv and PubMed for abstract summarization.

This paper introduces the first large-scale summarization dataset for Indonesian, sourced from the *Liputan6.com* online news portal over a 10-year period. It covers various topics and events that happened primarily in Indonesia, from October 2000 to October 2010. Below, we present details of the dataset, propose benchmark extractive and abstractive summarization methods that leverage both multilingual and monolingual pre-trained BERT models. We further conduct error analysis to better understand the limitations of current models over the dataset, as part of which we reveal not just modelling issues but also problems with ROUGE.

To summarize, our contributions are: (1) we release a large-scale Indonesian summarization corpus with over 200K documents, an order of mag-

<sup>1</sup><https://www.visualcapitalist.com/100-most-spoken-languages/>.

Example-1	
<p><b>Dokumen:</b> Liputan6.com, Jakarta : Organisasi Negara-negara Pengekspor Minyak (OPEC) mengakui mengalami kesulitan untuk menjaga stabilitas harga minyak dunia. Itu lantaran harga minyak terus melonjak sepanjang tahun ini. Hingga kini harga minyak mentah dunia masih mencapai tingkat tertinggi sejak pecah perang teluk sepuluh tahun silam. [3 kalimat dengan 57 kata tidak ditampilkan] Padahal , sebelumnya OPEC telah merevisi produksi minyak sebanyak tiga kali dalam enam bulan terakhir. Pertama, April hingga Juni dengan kenaikan mencapai 500 ribu barel dan terakhir, September ini, OPEC kembali menaikkan produksi sebesar 800 ribu barel per hari. [5 kalimat dengan 96 kata setelahnya tidak ditampilkan] <b>Ringkasan:</b> OPEC kesulitan menjaga stabilitas harga minyak dunia lantaran harga minyak dipasaran terus melonjak. Padahal, OPEC telah tiga kali menaikkan produksi dalam enam bulan terakhir.</p>	<p><b>Document:</b> Liputan6.com, Jakarta: The Organization of Petroleum Exporting Countries (OPEC) has admitted that it is having difficulty maintaining the stability of world oil prices. That's because oil prices continue to soar this year. Until now world crude oil prices have still reached the highest level since the gulf war broke out ten years ago. [3 sentences with 57 words are abbreviated from here] In fact, OPEC had previously revised oil production three times in the last six months. First, April to June with an increase of 500 thousand barrels and last, this September, OPEC has again increased production by 800 thousand barrels per day. [5 sentences with 96 words are abbreviated from here] <b>Summary:</b> OPEC is struggling to maintain the stability of world oil prices because oil prices on the market continue to soar. In fact, OPEC has raised production three times in the past six months.</p>
Example-2	
<p><b>Dokumen:</b> Liputan6.com, Jakarta : Gara-gara berusaha kabur saat diminta menunjukkan barang hasil curian, Rosihan bin Usman, tersangka pencurian tas wisatawan asing, baru-baru ini, tersungkur ditembak aparat Kepolisian Resor Denpasar Barat, Bali. Sebelumnya, Rosihan ditangkap massa setelah mencuri tas Nicholas Dreyden, wisatawan asing asal Inggris. Tas yang berisi dokumen keimigrasian dan surat penting itu diambil Rosihan setelah mengelabui korban. [7 kalimat dengan 78 kata setelahnya tidak ditampilkan] <b>Ringkasan:</b> Seorang pencuri tas wisatawan asing ditembak polisi. Ia berusaha kabur saat diminta menunjukkan hasil curian. Karena itu, polisi menembaknya.</p>	<p><b>Document:</b> Liputan6.com, Jakarta: Because of trying to escape when asked to show stolen goods, Rosihan bin Usman, a suspect of the theft of a foreign tourist bag, recently fell down, shot by the West Denpasar Resort Police, Bali. Previously, Rosihan was arrested by the mob after stealing the bag of Nicholas Dreyden, a foreign tourist from England. The bag containing immigration documents and important letters was taken by Rosihan after tricking the victim. [7 sentences with 78 words are abbreviated from here] <b>Summary:</b> A foreign tourist bag thief was shot by police. He tried to run away when asked to show the loot. Because of this, the police shot him.</p>

Figure 1: Example articles and summaries from Liputan6. To the left is the original document and summary, and to the right is an English translation (for illustrative purposes). We additionally highlight sentences that the summary is based on (noting that such highlighting is not available in the dataset).

nitude larger than the current largest Indonesian summarization dataset and one of the largest non-English summarization datasets in existence;<sup>2</sup> (2) we present statistics to show that the summaries in the dataset are reasonably abstractive, and provide two test partitions, a standard test set and an extremely abstractive test set; (3) we develop benchmark extractive and abstractive summarization models based on pre-trained BERT models; and (4) we conduct error analysis, on the basis of which we share insights to drive future research on Indonesian text summarization.

## 2 Data Construction

*Liputan6.com* is an online Indonesian news portal which has been running since August 2000, and provides news across a wide range of topics including politics, business, sport, technology, health, and entertainment. According to the Alexa ranking of websites at the time of writing,<sup>3</sup> *Liputan6.com* is ranked 9th in Indonesia and 112th globally. The website produces daily articles along

<sup>2</sup>The data can be accessed at [https://github.com/fajri91/sum\\_liputan6](https://github.com/fajri91/sum_liputan6)

<sup>3</sup><https://www.alexa.com/topsites>

with a short description for its RSS feed. The summary is encapsulated in the javascript variable `window.kmklabs.article` and the key `shortDescription`, while the article is in the main body of the associated HTML page. We harvest this data over a 10-year window — from October 2000 to October 2010 — to create a large-scale summarization corpus, comprising 215,827 document–summary pairs. In terms of preprocessing, we remove formatting and HTML entities (e.g. “ and \_), lowercase all words, and segment sentences based on simple punctuation heuristics. We provide example articles and summaries, with English translations for expository purposes (noting that translations are not part of the dataset), in Figure 1.

As a preliminary analysis of the document–summary pairs over the 10-year period, we binned the pairs into 5 chronologically-ordered groups containing 20% of the data each, and computed the proportion of novel  $n$ -grams (order 1 to 4) in the summary (relative to the source document). Based on the results in Figure 2, we can see that the proportion of novel  $n$ -grams drops over time, implying that the summaries of more recent articles are less

Variant	#Doc			% of Novel $n$ -grams			
	Train	Dev	Test	1	2	3	4
Canonical	193,883	10,972	10,972	16.2	52.5	71.8	82.4
Xtreme	193,883	4,948	3,862	22.2	66.7	87.5	96.6

Table 1: Statistics for the canonical and Xtreme variants of our data. The percentage of novel  $n$ -grams is based on the combined Dev and Test set.

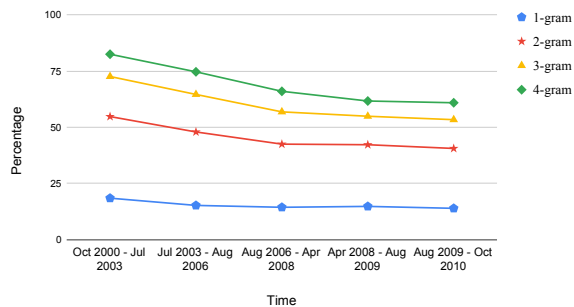


Figure 2: Proportion of novel  $n$ -grams over time in the summaries.

abstractive. For this reason, we decide to use the earlier articles (October 2000 to Jan 2002) as the development and test documents, to create a more challenging dataset. This setup also means there is less topic overlap between training and development/test documents, allowing us to assess whether the summarization models are able to summarize unseen topics.

For the training, development and test partitions, we use a splitting ratio of 90:5:5. In addition to this canonical partitioning of the data, we provide an “Xtreme” variant (inspired by *Xsum*; Narayan et al. (2018)) whereby we discard development and test document–summary pairs where the summary has fewer than 90% novel 4-grams (leaving the training data unchanged), creating a smaller, more challenging data configuration. Summary statistics for the “canonical” and “Xtreme” variants are given in Table 1.

We next present a comparison of Liputan6 (canonical partitioning) and IndoSum (the current largest Indonesian summarization dataset, as detailed in Section 6; Kurniawan and Louvan (2018)) in Table 2. In terms of number of documents, Liputan6 is approximately 11 times larger than IndoSum (the current largest Indonesian summarization dataset), although articles and summaries in Liputan6 are slightly shorter.

To understand the abstractiveness of the summaries in the two datasets, in Table 3 we present

ROUGE scores for the simple baseline of using the first  $N$  sentences as an extractive summary (“LEAD- $N$ ”), and the percentage of novel  $n$ -grams in the summary.<sup>4</sup> We use LEAD-3 and LEAD-2 for IndoSum and Liputan6 respectively, based on the average number of sentences in the summaries (Table 2). We see that Liputan6 has consistently lower ROUGE scores (R1, R2, and RL) for LEAD- $N$ ; it also has a substantially higher proportion of novel  $n$ -grams. This suggests that the summaries in Liputan6 are more abstractive than IndoSum.

To create a ground truth for extractive summarization, we follow Cheng and Lapata (2016) and Nallapati et al. (2016a) in greedily selecting the subset of sentences in the article that maximizes the ROUGE score based on the reference summary. As a result, each sentence in the article has a binary label to indicate whether they should be included as part of an extractive summary. Extractive summaries created this way will be referred to as “ORACLE”, to denote the upper bound performance of an extractive summarization system.

### 3 Summarization Models

We follow Liu and Lapata (2019) in building extractive and abstractive summarization models using BERT as an encoder to produce contextual representations for the word tokens. The architecture of both models is presented in Figure 3. We tokenize words with WordPiece, and append [CLS] (prefix) and [SEP] (suffix) tokens to each sentence. To further distinguish the sentences, we add even/odd segment embeddings ( $T_A/T_B$ ) based on the order of the sentence to the word embeddings. For instance, for a document with sentences  $[s_1, s_2, s_3, s_4]$ , the segment embeddings are  $[T_A, T_B, T_A, T_B]$ . Position embeddings ( $P$ ) are also used to denote the position of each token. The WordPiece, segment, and position embeddings are summed together and provided as input to BERT.

BERT produces a series of contextual representations for the word tokens, which we feed into a (second) transformer encoder/decoder for the extractive/abstractive summarization model. We detail the architecture of these two models in Sections 3.1 and 3.2. Note that this second transformer is initialized with random parameters (i.e. it is not pre-trained).

For the pre-trained BERT encoder, we use mul-

<sup>4</sup>All statistics are based on the entire dataset, encompassing the training, dev, and test data.

Dataset	#Doc			Article			Summary		
	Train	Dev	Test	$\mu(\text{Word})$	$\mu(\text{Sent})$	#Vocab	$\mu(\text{Word})$	$\mu(\text{Sent})$	#Vocab
IndoSum	14,252	750	3,762	347.23	18.37	117K	68.09	3.47	53K
Liputan6	193,883	10,972	10,972	232.91	12.60	311K	30.43	2.09	100K

Table 2: A comparison of IndoSum and Liputan6.  $\mu(\text{Word})$  and  $\mu(\text{Sent})$  denote the average number of words and sentences, respectively.

Dataset	LEAD- $N$			% of Novel $n$ -grams			
	R1	R2	RL	1	2	3	4
IndoSum	65.6	58.9	64.8	3.1	10.8	16.2	20.3
Liputan6	41.2	27.1	38.7	12.9	41.6	57.6	66.9

Table 3: Abtractiveness of the summaries in IndoSum and Liputan6.

tilingual BERT (mBERT) and our own IndoBERT (Koto et al., to appear).<sup>5</sup> IndoBERT is a BERT-Base model we trained ourselves using Indonesian documents from three sources: (1) Indonesian Wikipedia (74M words); (2) news articles (55M words) from Kompas,<sup>6</sup> Tempo (Tala et al., 2003),<sup>7</sup> and Liputan6,<sup>8</sup> and (3) the Indonesian Web Corpus (90M words; Medved and Suchomel (2017)). In total, the training data has 220M words. We implement IndoBERT using the Huggingface framework,<sup>9</sup> and follow the default configuration of BERT-Base (uncased): hidden size = 768d, hidden layers = 12, attention heads = 12, and feed-forward = 3,072d. We train IndoBERT with 31,923 Word-Pieces (vocabulary) for 2 million steps.

### 3.1 Extractive Model

After the document is processed by BERT, we have a contextualized embedding for every word token in the document. To learn inter-sentential relationships, we use the [CLS] embeddings ( $[x_{S_1}, x_{S_2}, \dots, x_{S_m}]$ ) to represent the sentences, to which we add a sentence-level positional embedding ( $P$ ), and feed them to a transformer encoder (Figure 3). An MLP layer with sigmoid activation is applied to the output of the transformer encoder to predict whether a sentence should be extracted (i.e.  $\tilde{y}_S \in \{0, 1\}$ ). We train the model with binary

<sup>5</sup>The pre-trained mBERT is sourced from: <https://github.com/google-research/bert>.

<sup>6</sup><https://kompas.com>

<sup>7</sup><https://koran.tempo.co>

<sup>8</sup>For Liputan6, we use only the articles from the training partition.

<sup>9</sup><https://huggingface.co/>

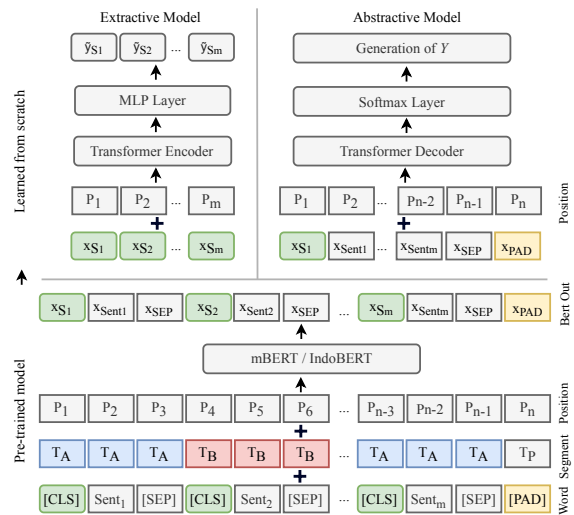


Figure 3: Architecture of the extractive and abstractive summarization models.

cross entropy, and update all model parameters (including BERT) during training. Note that the parameters in the transformer encoder and the MLP layer are initialized randomly, and learned from scratch.

The transformer encoder is configured as follows: layers = 2, hidden size = 768, feed-forward = 2,048, and heads = 8. In terms of training hyperparameters, we train using the Adam optimizer with learning rate  $lr = 2e^{-3} \cdot \min(\text{step}^{-0.5}, \text{step} \cdot \text{warmup}^{-1.5})$  where  $\text{warmup} = 10,000$ . We train for 50,000 steps on  $3 \times V100$  16GB GPUs, and perform evaluation on the development set every 2,500 steps. At test time, we select sentences for the extractive summary according to two conditions: the summary must consist of: (a) at least two sentences, and (b) at least 15 words. These values were set based on the average number of sentences and the minimum number of words in a summary. We also apply trigram blocking to reduce redundancy (Paulus et al., 2018). Henceforth, we refer to this model as “BERTEXT”.

### 3.2 Abstractive Model

Similar to the extractive model, we have a second transformer to process the contextualized embeddings from BERT. In this case, we use a transformer decoder instead (i.e. an attention mask is used to prevent the decoder from attending to future time steps), as we are learning to generate an abstractive summary. But unlike the extractive model, we use the BERT embeddings for all tokens as input to the transformer decoder (as we do not need sentence representations). We add to these BERT embeddings a second positional encoding before feeding them to the transformer decoder (Figure 3). The transformer decoder is initialized with random parameters (i.e. no pre-training).

The transformer decoder is configured as follows: layers = 6, hidden size = 768, feed-forward = 2,048, and heads = 8. Following Liu and Lapata (2019), we use a different learning rate for BERT and the decoder when training the model:  $lr = 2e^{-3} \cdot \min(\text{step}^{-0.5}, \text{step} \cdot 20,000^{-1.5})$  and  $0.1 \cdot \min(\text{step}^{-0.5}, \text{step} \cdot 10,000^{-1.5})$  for BERT and the transformer decoder, respectively. Both networks are trained with the Adam optimizer for 200,000 steps on  $4 \times V100$  16GB GPUs and evaluated every 10,000 steps. For summary generation, we use beam width = 5, trigram blocking, and a length penalty (Wu et al., 2016) to generate at least two sentences and at least 15 words (similar to the extractive model).

Henceforth the abstractive model will be referred to as “BERTABS”. We additionally experiment with a third variant, “BERTEXTABS”, where we use the weights of the fine-tuned BERT in BERTEXT for the encoder (instead of off-the-shelf BERT weights).

## 4 Experiment and Results

We use three ROUGE (Lin, 2004) F-1 scores as evaluation metrics: R1 (unigram overlap), R2 (bigram overlap), and RL (longest common subsequence overlap). In addition, we also provide BERTSCORE (F-1), as has recently been used for machine translation evaluation (Zhang et al., 2020b).<sup>10</sup> We use the development set to select the best checkpoint during training, and report the evaluation scores for the canonical and Xtreme test sets in Table 4. For both test sets, the summarization models are trained using the same training

<sup>10</sup>[https://github.com/Tiiiger/bert\\_score](https://github.com/Tiiiger/bert_score)

set, but they are tuned with a different development set (see Section 2 for details). In addition to the BERT models, we also include two pointer-generator models (See et al., 2017): (1) the base model (PTGEN); and (2) the model with coverage penalty (PTGEN+COV).<sup>11</sup>

We first look at the baseline LEAD- $N$  and ORACLE results. LEAD-2 is the best LEAD- $N$  baseline for Liputan6. This is unsurprising, given that in Table 2, the average summary length was 2 sentences. We also notice there is a substantial gap between ORACLE and LEAD-2: 12–15 points for R1 and 5–7 points for BERTSCORE, depending on the test set. This suggests that the baseline of using the first few sentences as an extractive summary is ineffective. Comparing the performance between the canonical and Xtreme test sets, we see a substantial drop in performance for both LEAD- $N$  and ORACLE, highlighting the difficulty of the Xtreme test set due to its increased abstractiveness.

For the pointer-generator models, we see little improvement when including the coverage mechanism (PTGEN+COV vs. PTGEN), implying that there is minimal repetition in the output of PTGEN. We suspect this is due to the Liputan6 summaries being relatively short (2 sentences with 30 words on average). A similar observation is reported by Narayan et al. (2018) for *XSum*, where the summaries are similarly short (a single sentence with 23 words, on average).

Next we look at the BERT models. Overall they perform very well, with both the mBERT and IndoBERT models outperforming the LEAD- $N$  baselines and PTGEN models by a comfortable margin. IndoBERT is better than mBERT (approximately 1 ROUGE point better on average over most metrics), showing that a monolingually-trained BERT is a more effective pre-trained model than the multilingual variant. The best performance is achieved by IndoBERT’s BERTEXTABS. In the canonical test set, the improvement over LEAD-2 is +4.4 R1, +2.62 R2, +4.3 R3, and +3.4 BERTSCORE points. In the Xtreme test set, BERTEXTABS suffers a substantial drop compared to the canonical test set (6–7 ROUGE and 2 BERTSCORE points), although the performance gap between it and LEAD-2 is about the same.

<sup>11</sup>We use the default hyper-parameter configuration recommended by the original authors for the pointer-generator models.

Model	Canonical Test Set				Xtreme Test Set			
	R1	R2	RL	BS	R1	R2	RL	BS
LEAD-1	32.67	18.50	29.40	72.62	27.27	11.56	23.60	71.19
LEAD-2	36.68	20.23	33.71	74.58	31.10	12.78	27.63	72.98
LEAD-3	34.49	18.84	32.06	74.31	29.54	12.05	26.68	72.78
ORACLE	51.54	30.56	47.75	79.24	43.69	18.57	38.84	76.75
PTGEN	36.10	19.19	33.56	75.92	30.41	12.05	27.51	74.10
PTGEN+COV	35.53	18.56	32.92	75.75	30.27	11.81	27.26	74.11
BERTEXT (mBERT)	37.51	20.15	34.57	75.22	31.83	12.63	28.37	73.62
BERTABS (mBERT)	39.48	21.59	36.72	77.19	33.26	13.82	30.12	75.40
BERTEXTABS (mBERT)	39.81	21.84	37.02	77.39	33.86	14.13	30.73	75.69
BERTEXT (IndoBERT)	38.03	20.72	35.07	75.33	31.95	12.74	28.47	73.64
BERTABS (IndoBERT)	40.94	<b>23.01</b>	37.89	77.90	34.59	<b>15.10</b>	31.19	75.84
BERTEXTABS (IndoBERT)	<b>41.08</b>	22.85	<b>38.01</b>	<b>77.93</b>	<b>34.84</b>	15.03	<b>31.40</b>	<b>75.99</b>

Table 4: ROUGE results for the canonical and Xtreme test sets. All ROUGE (“R1”, “R2”, and “RL”) scores have a confidence interval of at most  $\pm 0.3$ , as reported by the official ROUGE script. “BS” is BERScore computed with `bert-base-multilingual-cased` (layer 9), as suggested by Zhang et al. (2020b).

## 5 Error Analysis

In this section, we analyze errors made by the extractive (BERTEXT) and abstractive (BERTEXTABS) models to better understand their behaviour. We use the mBERT version of these models in our analysis.<sup>12</sup>

### 5.1 Error Analysis of Extractive Summaries

We hypothesized that the disparity between ORACLE and BERTEXT (14.03 point difference for R1 in the canonical test set) was due to the number of extracted sentences. To test this, when extracting sentences with BERTEXT, we set the total number of extracted sentences to be the same as the number of sentences in the ORACLE summary. However, we found minimal benefit using this approach, suggesting that the disparity is not a result of the number of extracted sentences.

To investigate this further, we present the frequency of *sentence positions* that are used in the summary in ORACLE and BERTEXT for the canonical test set in Figure 4a. We can see that BERTEXT tends to over-select the first two sentences as the summary. In terms of proportion, 65.47% of

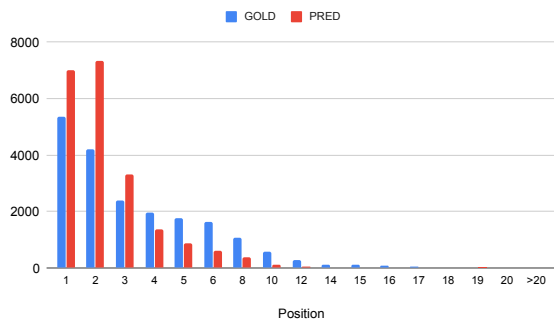
<sup>12</sup>The error analysis is based on mBERT rather than IndoBERT simply because this was the best-performing model at the time the error analysis was performed. While IndoBERT ultimately performed slightly better, given that the two models are structurally identical, we would expect to see a similar pattern of results.

BERTEXT summaries involve the first two sentences. In comparison, only 42.54% of ORACLE summaries use sentences in these positions. One may argue that this is because the training and test data have different distributions under our chronological partitioning strategy (recall that the test set is sampled from the earliest articles), but that does not appear to be the case: as Figure 4b shows, the distribution of sentence positions in the training data is very similar to the test data — 43.14% of ORACLE summaries involve the first two sentences.

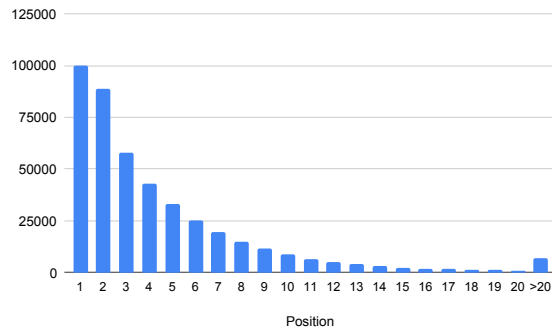
### 5.2 Error Analysis of Abstractive Summaries

To perform error analysis for BERTEXTABS, we randomly sample 100 documents with an R1 score  $< 0.4$  in the canonical test set (which accounts for nearly 50% of the test documents). Two native Indonesian speakers examined these 100 samples to manually assess the quality of the summaries, and score them on a 3-point ordinal scale: (1) *bad*; (2) *average*; and (3) *good*. Each annotator is presented with the source document, the reference summary, and the summary generated by BERTEXTABS. In addition to the overall quality evaluation, we also asked the annotators to analyze a number of (fine-grained) attributes in the summaries:

- *Abbreviations*: the system summary uses abbreviations that are different to the reference summary.



(a) Distribution of sentence positions for ORACLE and BERTEXT in the canonical test set.



(b) Distribution of sentence positions for ORACLE in the training set.

Figure 4: Position of ORACLE and/or Predicted Extractive Summaries

Category	Bad	Avg.	Good
#Samples (100)	32	8	60
Abbreviation (%)	21.9	25.0	40.0
Morphology (%)	12.5	25.0	36.7
Paraphrasing (%)	50.0	87.5	<b>86.7</b>
Lack of coverage (%)	<b>90.6</b>	<b>100.0</b>	40.0
Wrong focus (%)	68.8	0.00	8.3
Un. details (from doc) (%)	<b>90.6</b>	75.0	75.0
Un. details (not from doc) (%)	18.8	12.5	5.0

Table 5: Error analysis for 100 samples with R1 <0.4.

- *Morphology*: the system summary uses morphological variants of the same lemmas contained in the reference summary.
- *Synonyms/paraphrasing*: the system summary contains paraphrases of the reference summary.
- *Lack of coverage*: the system summary lacks coverage of certain details that are present in the reference summary.
- *Wrong focus*: the system summarizes a different aspect/focus of the document to the reference summary.
- *Unnecessary details (from document)*: the system summary includes unimportant but factually correct information.
- *Unnecessary details (not from document)*: the system summary includes unimportant and factually incorrect information (hallucinations).

We present a breakdown of the different error types in Table 5. Inter-annotator agreement for the overall quality assessment is high (Pearson’s  $r = 0.69$ ). Disagreements in the quality label (*bad*,

*average*, *good*) are resolved as follows: (1)  $\{bad, average\} \rightarrow bad$ ; and (2)  $\{good, average\} \rightarrow good$ . We only have four examples with  $\{bad, good\}$  disagreement, which we resolved through discussion. Interestingly, more than half (60) of our samples were found to have *good* summaries. The primary reasons why these summaries have low ROUGE scores are paraphrasing (86.7%), and the inclusion of additional (but valid) details (75.0%). Abbreviations and morphological differences also appear to be important factors. These results underline a problem with the ROUGE metric, in that it is unable to detect good summaries that use a different set of words to the reference summary. One way forward is to explore metrics that consider sentence semantics beyond word overlap such as METEOR (Banerjee and Lavie, 2005) and BERTSCORE,<sup>13</sup> and question-answering system based evaluation such as APES (Eyal et al., 2019) and QAGS (Wang et al., 2020). Another way is to create more reference summaries (which will help with the issue of the system summaries including [validly] different details to the single reference).

Looking at the results for *average* summaries (middle column), BERTEXTABS occasionally fails to capture salient information: 100% of the summaries have coverage issues, and 75.0% contain unnecessary (but valid) details. They also tend to use paraphrases (87.5%), which further impacts on a lower ROUGE score. Finally, the *bad* system summaries have similar coverage issues, and also tend to have a very different focus compared to the

<sup>13</sup>Indeed, we suggest that BERTSCORE should be used as the canonical evaluation metric for the dataset, but leave empirical validation of its superiority for Indonesian summarization evaluation to future work.

Example-1 of error analysis (Abbreviation, morphoplogy, synonyms/paraphrasing, and details from the document)	
<p><b>Dokumen:</b> Liputan6.com , Jakarta : Protes masih bergema menyambut Keputusan Menteri Tenaga Kerja dan Transmigrasi Nomor 78 Tahun 2001 . Kebijakan yang sengaja dikeluarkan sebagai wujud perubahan keputusan sebelumnya ini , sampai sekarang , masih mengundang kecaman keras dari pekerja di Indonesia . Itulah sebabnya , mereka menuntut Kepmenakertrans baru ini dicabut karena dinilai merugikan pekerja . [19 kalimat dengan 406 kata tidak ditampilkan] Sementara itu , SPSI secara tegas menolak segala bentuk negosiasi . [3 kalimat dengan 45 kata setelahnya tidak ditampilkan] <b>Ringkasan manusia:</b> pemberlakuan <b>kepmenakertrans</b> 78/2001 masih mengundang rasa tidak puas di dada sejumlah pekerja indonesia . maka , lahirlah <b>tuntutan</b> agar peraturan yang dinilai merugikan ini dicabut . <b>Ringkasan sistem [Good]:</b> <b>keputusan menteri tenaga kerja dan transmigrasi</b> nomor 78 tahun 2001 mengundang kecaman keras dari pekerja di indonesia . mereka menuntut kepmenakertrans dicabut karena dinilai merugikan pekerja . spsi menolak negosiasi .</p>	<p><b>Document:</b> Liputan6.com, Jakarta: Protests still resonate with welcoming Minister of Manpower and Transmigration Decree No. 78/2001. This policy, which was deliberately issued as an amendment to the previous decision, until now, still invites harsh criticism from workers in Indonesia. That is why they demand to revoke the new Kepmenakertrans because it is considered detrimental to workers. [19 sentences with 406 words are abbreviated from here] Meanwhile, SPSI firmly rejected all forms of negotiation. [3 sentences with 45 words are abbreviated from here] <b>Gold Summary:</b> The enactment of <b>Kepmenakertrans</b> 78/2001 still invites the dissatisfaction of Indonesian workers. hence, <b>demands</b> to revoke the regulation arose as it was considered to be detrimental. <b>System Summary [Good]:</b> <b>Minister of Manpower and Transmigration Decree</b> number 78 of 2001 invited strong criticism from workers in Indonesia. They <b>demand</b> to revoke Kepmenakertrans because it is considered detrimental to workers. <b>SPSI</b> rejects negotiations.</p>
Example-2 of error analysis (Lack of coverage, wrong focus, and details that are not from the document)	
<p><b>Dokumen:</b> Liputan6.com , Jakarta : Langkah reshuffle yang dilakukan Presiden Abdurrahman Wahid , agaknya tak mendapat restu . Buktinya , Wakil Presiden Megawati Sukarnoputri kembali tidak hadir dalam pelantikan tiga menteri bidang ekonomi , Rabu ( 13/6 ) . [8 kalimat dengan 113 kata setelahnya tidak ditampilkan] <b>Ringkasan manusia:</b> wapres megawati sukarnoputri , kembali tidak hadir dalam pelantikan tiga menteri baru . dalam reshuffle 1 juni , megawati juga tak muncul dalam pelantikan , karena merasa tak dilibatkan dalam reshuffle kabinet . <b>Ringkasan sistem [Bad]:</b> presiden abdurrahman wahid kembali tidak hadir dalam pelantikan tiga menteri bidang ekonomi . ketidakepakatan soal perombakan kabinet itu juga terjadi 1 juni silam . presiden meminta mereka lebih menjaga koordinasi antarmenteri .</p>	<p><b>Document:</b> Liputan6.com, Jakarta: The reshuffle step was taken by President Abdurrahman Wahid, apparently did not get the blessing. The proof, Vice President Megawati Sukarnoputri was again not present at the inauguration of three ministers in the economic sector, Wednesday (6/13). [8 sentences with 113 words are abbreviated from here] <b>Gold Summary:</b> Vice President Megawati Sukarnoputri, is not present at the inauguration of three new ministers again. In the reshuffle on June 1, Megawati also did not appear in the inauguration, because she felt not involved in the cabinet reshuffle. <b>System Summary [Bad]:</b> President Abdurrahman Wahid was again absent from the inauguration of three ministers in the economic sector. disagreement about the cabinet reshuffle also occurred 1 June ago. the president asked them to maintain more coordination between ministries.</p>

Figure 5: Two examples to highlight error categories used in our error analysis.

reference summary (90.6%).

In Figure 5 we show two representative examples from BERTEXTABS. The first example is considered *good* by our annotators, but due to abbreviations, morphological differences, paraphrasing, and additional details compared to the reference summary, the ROUGE score is  $<0.4$ . In this example, the gold summary uses the abbreviation *kepmenakertrans* while BERTEXTABS generates the full phrase *keputusan menteri tenaga kerja dan transmigrasi* (which is correct). The example also uses paraphrases (*invites strong criticism* to explain *dissatisfaction*), and there are morphological differences in words such as *tuntutan* (noun) vs. *menuntut* (verb). The low ROUGE score here highlights the fact that the bigger issue is with ROUGE itself rather than the summary.

The second example is considered to be *bad*, with the following issues: lack of coverage, wrong focus, and contains unnecessary details that are not from the article. The first sentence *President Abdurrahman Wahid was absent* has nothing to do

with the original article, creating a different focus (and confusion) in the overall summary.

To summarize, coverage, focus, and the inclusion of other details are the main causes of low quality summaries. Our analysis reveals that abbreviations and paraphrases are another cause of summaries with low ROUGE scores, but that is an issue with ROUGE rather than the summaries. Encouragingly, hallucination (generating details not in the original document) is not a major issue for these models (notwithstanding that almost 20% of *bad* samples contain hallucinations).

## 6 Related Datasets

Previous studies on Indonesian text summarization have largely been extractive and used small-scale datasets. Gunawan et al. (2017) developed an unsupervised summarization model over 3K news articles using heuristics such as sentence length, keyword frequency, and title features. In a similar vein, Najibullah (2015) trained a naive Bayes model to extract summary sentences in a 100-article dataset.



Aristoteles et al. (2012) and Silvia et al. (2014) apply genetic algorithms to a summarization dataset with less than 200 articles. These studies do not use ROUGE for evaluation, and the datasets are not publicly available.

Koto (2016) released a dataset for chat summarization by manually annotating chat logs from *WhatsApp*.<sup>14</sup> However, this dataset contains only 300 documents. The largest summarization data to date is *IndoSum* (Kurniawan and Louvan, 2018), which has approximately 19K news articles with manually-written summaries. Based on our analysis, however, the summaries of *IndoSum* are highly extractive.

Beyond Indonesian, there is only a handful of non-English summarization datasets that are of sufficient size to train modern deep learning summarization methods over, including: (1) LCSTS (Hu et al., 2015), which contains 2 million Chinese short texts constructed from the Sina Weibo microblogging website; and (2) ES-News (Gonzalez et al., 2019), which comprises 270k Spanish news articles with summaries. LCSTS documents are relatively short (less than 140 Chinese characters), while ES-News is not publicly available. Our goal is to create a benchmark corpus for Indonesian text summarization that is both large scale and publicly available.

## 7 Conclusion

We release Liputan6, a large-scale summarization corpus for Indonesian. Our dataset comes with two test sets: a canonical test set and an “Xtreme” variant that is more abstractive. We present results for several benchmark summarization models, in part based on IndoBERT, a new pre-trained BERT model for Indonesian. We further conducted extensive error analysis, as part of which we identified a number of issues with ROUGE-based evaluation for Indonesian.

## Acknowledgments

We are grateful to the anonymous reviewers for their helpful feedback and suggestions. In this research, Fajri Koto is supported by the Australia Awards Scholarship (AAS), funded by the Department of Foreign Affairs and Trade (DFAT), Australia. This research was undertaken using the LIEF HPC-GPGPU Facility hosted at The University of

Melbourne. This facility was established with the assistance of LIEF Grant LE170100200.

## References

- Aristoteles Aristoteles, Yeni Herdiyeni, Ahmad Ridha, and Julio Adisantoso. 2012. Text feature weighting for summarization of document Bahasa Indonesia using genetic algorithm. *IJCSI International Journal of Computer Science Issues*, 9(1):1–6.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72.
- Jianpeng Cheng and Mirella Lapata. 2016. Neural summarization by extracting sentences and words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 484–494.
- 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 *NAACL HLT 2018: 16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 2, pages 615–621.
- Matan Eyal, Tal Baumel, and Michael Elhadad. 2019. Question answering as an automatic evaluation metric for news article summarization. In *NAACL-HLT 2019: Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 3938–3948.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-up abstractive summarization. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 4098–4109.
- J.-A. Gonzalez, L.-F. Hurtado, E. Segarra, F. Garcia-Granada, and E. Sanchis. 2019. Summarization of Spanish talk shows with siamese hierarchical attention networks. *Applied Sciences*, 9(18).
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English Gigaword. Linguistic Data Consortium.
- D Gunawan, A Pasaribu, R F Rahmat, and R Budiarto. 2017. Automatic text summarization for Indonesian language using TextTeaser. *IOP Conference Series: Materials Science and Engineering*, 190(1):12048.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Neural Information Processing Systems*, pages 1693–1701.

<sup>14</sup><https://www.whatsapp.com/>.

- Wan Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. A unified model for extractive and abstractive summarization using inconsistency loss. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 132–141.
- Baotian Hu, Qingcai Chen, and Fangze Zhu. 2015. LCSTS: A large scale Chinese short text summarization dataset. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1967–1972.
- Byeongchang Kim, Hyunwoo Kim, and Gunhee Kim. 2019. Abstractive summarization of Reddit posts with multi-level memory networks. In *NAACL-HLT 2019: Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 2519–2531.
- Fajri Koto. 2016. A publicly available Indonesian corpora for automatic abstractive and extractive chat summarization. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*.
- Fajri Koto, Afshin Rahimi, Jey Han Lau, and Timothy Baldwin. to appear. IndoLEM and IndoBERT: A benchmark dataset and pre-trained language model for Indonesian NLP. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING 2020)*.
- Kemal Kurniawan and Samuel Louvan. 2018. Indosum: A new benchmark dataset for Indonesian text summarization. In *2018 International Conference on Asian Language Processing (IALP)*, pages 215–220.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *2019 Conference on Empirical Methods in Natural Language Processing*, pages 3728–3738.
- Marek Medved and Vít Suchomel. 2017. Indonesian web corpus (idWac). In *LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (FAL), Faculty of Mathematics and Physics, Charles University*.
- Ahmad Najibullah. 2015. Indonesian text summarization based on naive Bayes method. *Proceeding Of The International Seminar and Conference 2015*, 1(1).
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2016a. SummaRuNNer: A recurrent neural network based sequence model for extractive summarization of documents. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*, pages 3075–3081.
- Ramesh Nallapati, Bowen Zhou, Cicero Nogueira dos santos, Caglar Gulcehre, and Bing Xiang. 2016b. Abstractive text summarization using sequence-to-sequence RNNs and beyond. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290.
- 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 *EMNLP 2018: 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. A deep reinforced model for abstractive summarization. In *Proceedings of the 6th International Conference on Learning Representations*.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 379–389.
- Evan Sandhaus. 2008. The New York Times annotated corpus. Linguistic Data Consortium.
- 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*, pages 1073–1083.
- Eva Sharma, Chen Li, and Lu Wang. 2019. BIGPATENT: A large-scale dataset for abstractive and coherent summarization. In *ACL 2019: The 57th Annual Meeting of the Association for Computational Linguistics*, pages 2204–2213.
- Silvia, Pitri Rukmana, Vivi Regina Aprilia, Derwin Suhartono, Rini Wongso, and Meiliana. 2014. Summarizing text for Indonesian language by using latent Dirichlet allocation and genetic algorithm. In *1st International Conference on Electrical Engineering, Computer Science and Informatics 2014*, pages 148–153.
- F. Tala, J. Kamps, K.E. Miller, and M. de Rijke. 2003. The impact of stemming on information retrieval in Bahasa Indonesia. In *The 14th Meeting of Computational Linguistics in the Netherlands*.

- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. Abstractive document summarization with a graph-based attentional neural model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1171–1181.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. Asking and answering questions to evaluate the factual consistency of summaries. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, ukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020a. PEGASUS: Pre-training with extracted gap-sentences for abstractive summarization. In *ICML 2020: 37th International Conference on Machine Learning*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020b. BERTScore: Evaluating text generation with BERT. In *ICLR 2020: Eighth International Conference on Learning Representations*.