

HeSum: a Novel Dataset for Abstractive Text Summarization in Hebrew

Tzuf Paz-Argaman^{*,a}, Itai Mondshine^{*,a}, Asaf Achi Mordechai^a, and Reut Tsarfaty^a

^aBar-Ilan University, Israel,

{tzuf.paz-argaman, mondshi1, asaf.achimordechai, reut.tsarfaty}@biu.ac.il

Abstract

While large language models (LLMs) excel in various natural language tasks in English, their performance in lower-resourced languages like Hebrew, especially for generative tasks such as abstractive summarization, remains unclear. The high morphological richness in Hebrew adds further challenges due to the ambiguity in sentence comprehension and the complexities in meaning construction. In this paper, we address this resource and evaluation gap by introducing HeSum, a novel benchmark specifically designed for abstractive text summarization in Modern Hebrew. HeSum consists of 10,000 article-summary pairs sourced from Hebrew news websites written by professionals. Linguistic analysis confirms HeSum’s high abstractness and unique morphological challenges. We show that HeSum presents distinct difficulties for contemporary state-of-the-art LLMs, establishing it as a valuable testbed for generative language technology in Hebrew, and MRLs generative challenges in general.¹

1 Introduction

Recent advances with large language models (LLMs, Brown et al., 2020; Chowdhery et al., 2023) demonstrate impressive capabilities, encompassing diverse tasks such as natural language (NL) understanding and reasoning, including *classification* tasks such as commonsense reasoning (Bisk et al., 2020) and sentiment analysis (Liang et al., 2022), as well as *generative* tasks like summarization and dialogue systems (Cohen et al., 2022). However, these impressive achievements are primarily demonstrated for the English language. Our understanding of how these models perform on low-resource languages is limited, as current evaluations are primarily focused on languages with abundant data (Ahuja et al., 2023; Lai et al., 2023).

^{*}Equal contribution.

¹The dataset, code, and fine-tuned models are publicly available at <https://github.com/OnIplab/HeSum>

This concern is particularly relevant for *morphologically rich languages* (MRLs) such as Hebrew, which is known for their word complexity and ambiguity, leading to processing difficulty (Tsarfaty et al., 2019, 2020). Despite advances in natural language processing for Hebrew, which so far covered tasks as reading comprehension (Keren and Levy, 2021; Cohen et al., 2023), named entity recognition (Bareket and Tsarfaty, 2021), sentiment analysis (Chriqui and Yahav, 2022), and text-based geolocation (Paz-Argaman et al., 2023); a crucial gap persists in the ability to evaluate novel, human-like generated text, as in *abstractive text generation*.

Abstractive text-generation requires both natural language understanding and reasoning over the input, and the ability to generate grammatically, and in particular *morpho-syntactically*, correct text, as well as *semantically* and *morpho-semantically* coherent, fluent text that conveys consistent meanings. Notably, text-generation models are also prone to ‘hallucinations’ — generating factually incorrect content. These challenges are further amplified in Hebrew due to its morphological richness which leads to a complex realization of sentence structure and meaning.

In order to enable empirically quantified assessment of these aspects of text generation in MRLs, we present a novel benchmark dataset for **Hebrew** abstractive text **Summarization (HeSum)**. HeSum consists of 10,000 articles paired with their corresponding summaries, all of which have been sourced from three different Hebrew news websites, all written by professional journalists. This curated collection offers several key advantages: (i) *High Abstractness* – extensive linguistic analysis validates HeSum’s summaries as demonstrably more abstractive even when compared to English benchmarks. (ii) *Unique Hebrew Challenges* – meticulous linguistic analysis pinpoints the inherent complexities specific to Hebrew summarization, offering valuable insights into the nuanced

Set	Size	Vocabulary size (over Articles)		Avg. Document Length		Avg. Word Ambiguity	Avg. Morph Anaphors	Avg. Construct-State	BertScore
		Lemmas	Tokens	Article	Summary	Article	Article	Summary	Semantic Similarity Article-Summary
Train	8,000	47,903	269,168	1,427.4	33.2	58	98.8	2.4	76
Validation	1,000	23,134	104,383	1,410.0	33.8	90	87.9	2.5	76
Test	1,000	22,543	102,387	1,507.6	34.7	89	95.7	2.6	74

Table 1: Linguistic Analysis of the HeSum dataset.

characteristics that differentiate it from its English counterpart. And (iii) *Thorough LLM Evaluation* – we conducted a comprehensive empirical analysis using state-of-the-art LLMs, demonstrating that HeSum presents unique challenges even for these contemporary models. By combining high abstractness, nuanced morphological complexities, and a rigorous LLM evaluation, HeSum establishes itself as a valuable resource for advancing the frontiers of abstractive text summarization in MRL settings.

2 The Challenge

Linguistic Challenges in Hebrew Morphologically rich languages (MRLs) pose distinct challenges for generative tasks, above and beyond morphologically impoverished ones such as English.

In MRLs, each input token can be composed of multiple lexical and functional elements, each contributing to the overall structure and semantic meanings of the generated text. For instance, the Hebrew word ‘וכשמביהנו’ is composed of seven morphemes: ‘ו’ (‘and’), ‘כש’ (‘when’), ‘מ’ (‘from’), ‘ה’ (‘the’), ‘בית’ (‘house’), ‘של’ (‘of’), and ‘אנחנו’ (‘us’). This has ramifications for both the understanding of MRL texts, a process that necessitates morphological segmentation, and for generating MRL texts, requiring morphological fusion. At comprehension, Hebrew poses an additional challenge due to its inherent ambiguity, with many tokens admitting multiple valid segmentations, e.g., ‘הקפה’ could be interpreted as ‘ה’+‘קפה’ (‘the’+‘coffee’); as ‘הקפה’ (‘orbit’); or as ‘הקף’ + ‘של’ + ‘היא’ (‘perimeter’+‘of’+‘her’). During generation, the emergence of unseen morphological compositions, where unfamiliar morphemes combine in familiar ways, poses an additional challenge (Hofmann et al., 2021; Gueta et al., 2023). These challenges, coupled with inherent linguistic features like morphological inflections, construct-state nouns (*smixut*), and flexible word order, create a multifaceted challenge for LLMs in processing and generating Hebrew texts.

The HeSum Task We aim to unlock the comprehension-and-generation challenge in MRL

settings by first tackling the abstractive text summarization task (Moratanch and Chitrakala, 2016), here focusing on Modern Hebrew.

Given an input document in Hebrew, specifically a news article, our goal is to generate a short, clear, Hebrew summary of the key information in the article. In contrast to *extractive* summarization, here novel morphosyntactic structures need to be generated to communicate the summary.

3 Dataset, Statistics and Analysis

3.1 Data Collection

The HeSum dataset consists of article-and-summary pairs. The articles were collected from three Hebrew news websites: “Shakuf”,² “HaMakom”,³ and “The Seventh Eye”.⁴ These websites focus on independent journalism, providing articles on topics such as government accountability, corporate influence, and environmental issues. Each article on these websites is accompanied by an extended subheading written by a professional editor, that serves as a summary of the content. To ensure data quality, articles that were not in Hebrew, or ones that had particularly short summaries (i.e., the extended subheading was less than 10 tokens) were excluded from the dataset.

3.2 Linguistic Analysis

We examined the linguistic, morpho-syntactic and semantic, properties of the HeSum dataset. For the extraction of syntactic and semantic features, we used DictaBert (Shmidman et al., 2023). Additionally, AlephBert (Seker et al., 2022), a Hebrew monolingual BERT-based encoder model (Devlin et al., 2018), was employed to compute semantic similarity between articles and their corresponding summaries, leveraging the BertScore method (Zhang* et al., 2020). Notably, semantic similarity was performed only on article-summary pairs within the model’s 512-token limit.

²<https://shakuf.co.il>

³<https://www.ha-makom.co.il>

⁴<https://www.the7eye.org.il>

Dataset	novel n-grams				CMP	RED (n=1)	RED (n=2)
	n = 1	n = 2	n = 3	n = 4			
CNN/Daily Mail	13.20	52.77	72.20	81.40	90.90	13.73	1.10
XSum	35.76	83.45	95.50	98.49	90.40	5.83	0.16
HeSum	42.00	73.20	82.00	85.36	95.48	4.83	0.10

Table 2: HeSum’s Intrinsic Evaluation compared to English Benchmarks (CNN/Daily Mail and XSum).

Table 1 highlights the Hebrew language’s multifaceted complexities as reflected in this task. The notable disparity in the vocabulary size between token and lemma counts underscores extensive morphological richness, necessitating models adept at handling linguistic diversity. The abundance of morphological anaphoric expressions and numerous Hebrew construct-state constructions necessitate advanced models attuned to entity relations that are expressed via Hebrew’s unique morphological traits. Lattice analysis reveals a high degree of word ambiguity (numerous lattice paths), highlighting natural language understanding challenges and the consequent difficulty of accurate tokenization for downstream processing tasks. The substantial document length, necessitate the use of models adept at long-form text processing. Finally, the relatively high semantic similarity score indicates effective information distillation in the summaries.

3.3 Summarization Intrinsic Analysis

To assess the challenges of the HeSum summaries we used three established metrics: (i) *Abstractness*: the percentage of summary novel n-grams, unseen in the article (Narayan et al., 2018). (ii) *Compression Ratio (CMP)*: the word count in summary S divided by the word count in article A : $CMP_w(S, A) = 1 - \frac{|S|}{|A|}$. Higher compression ratios indicate greater word-level reduction and, subsequently, potentially a more challenging summarization task (Bommasani and Cardie, 2020). (iii) *Redundancy (RED)*: measures repetitive n-grams within a summary (S) using the form: $RED(S) = \frac{\sum_{i=1}^m (f_i - 1)}{\sum_{i=1}^m f_i}$ where m is the number of unique n-grams in the summary and $f_i \geq 1$ is the frequency count of a specific n-gram (Hasan et al., 2021).

Table 2 presents a quantitative analysis of HeSum’s summarization characteristics, underscoring its challenges. HeSum demonstrates a high degree of abstractness, with approximately half of its unique vocabulary and over 73% of bigrams unseen in the original articles. Furthermore, HeSum

presents a significant compression challenge, as summaries average less than 5% of the input article length. Additionally, the analysis reveals minimal redundancy within the summaries, with less than 5% repeated n-grams. These findings underscore HeSum’s efficacy in conveying the central ideas of the articles’ information in a novel, distillate, and non-redundant manner. Comparative analysis with established abstractive summarization benchmarks, CNN/Daily Mail (Nallapati et al., 2016) and XSum (Narayan et al., 2018), confirms HeSum’s high abstractness, compression ratio, and low redundancy, even when compared to these datasets.

4 Experiments

4.1 Experimental setup

Models To demonstrate the complexity of this task, we conducted an evaluation of two LLMs in a zero-shot setting: the GPT-4 model with 32K context window (version 0613), and GPT-3.5-turbo with 16K context (version 0613). To find the most effective prompt format, we tested on the HeSum validation set various prompting strategies, including translating parts of the prompt to English. Ultimately, we adopted the English-translated approach (Brown et al., 2020), where both the instruction and input were translated. The output summaries are strictly in Hebrew. Additionally, to address the limitations of available generative models for Hebrew, we fine-tuned the multilingual mLongT5 model (Uthus et al., 2023) on the HeSum training set with two versions, base (2.37 GB) and large (4.56 GB). mLongT5 is a sequence-to-sequence model based on mT5 (Xue et al., 2020) specifically designed for handling long sequences. Appendix B includes the GPT models’ prompting strategies experiments, and the mLongT5 training details.

Automatic Evaluation Metrics To evaluate the generated summaries with respect to the original texts, we used two standardly-used automatic metrics: ROUGE and BertScore.

Model	ROUGE			BertScore	Human Evaluation	
	ROUGE-1	ROUGE-2	ROUGE-L		Coherence	Completeness
GPT-4	13.59	3.70	10.39	77.3	4.48	4.14
GPT-3.5	13.69	3.84	10.55	77.0	4.38	3.98
mLongT5 (fine-tuned)	17.47	7.56	14.68	57.6	3.46	2.10

Table 3: Models’ performance on the HeSum test-set.

Phenomenon	GPT-4	GPT-3.5	mLongT5	Example error in Hebrew	Example error translated into English	Explanation
Repetition	0	0	1	האם הוא יכול להיות אליים? אם הוא יכול להיות אליים?...	Can he be violent? If he can be violent?	Duplication with subtle alterations.
Copy from Article	0	0	4	האם עיתונאי יכול להציע ביקורת פומבית על...?	Can a journalist make public criticism of...?	The model-generated summary replicates a section of the original article.
Non-alphabetic omission	1	0	14	צחצחים	Chachachim	Missing diacritic – it should be ‘צחצחים’ instead of ‘צחצחים’.
Incorrect disambiguation	1	1	2	נאומו של הרוד טופז	Uncle Topaz’s speech...	‘רודו’ is incorrectly interpreted as ‘רוד’ + ‘של’ (‘Uncle’ + ‘of’), instead of as a man’s name – ‘רודו’, which is why the model added a definite ‘ה’ to ‘רוד’.
Hallucination	3	3	2	...ערב את נח...	...involved Noah...	Noah is not a person mentioned in the article.
Culture transfer	1	1	0	...למנהיגת הקמפיין הנבחרת, נסי ברנדס...	...to the campaign leader-elect, Nancy Brands...	The article refers to Nancy as a ‘he’, but the summary uses feminine inflection (leader), probably due to Nancy being a common female name in English.
Incorrect gender	6	11	1	...חושפות בחקירתם...	...reveal in their investigation...	Gender inflection mismatch: ‘reveal’ (fem.) clashes with ‘their’ (masc.).
Incorrect definite (e.g., construct state)	3	2	3	...המשטר המשפטי פירסם...	The Ministry of the Justice published...	Definite articles on both words in ‘The Ministry of the Justice’ violate Hebrew construct state rules.

Table 4: Error analysis comparing generated summaries from GPT-4, GPT-3.5, and mLongT5 based on 30 inputs.

ROUGE (Lin, 2004) is a widely-used metric in summarization that measures n-gram overlap between generated summaries and human-written references. We calculated ROUGE-1 (unigrams), ROUGE-2 (bigrams), and ROUGE-L scores (longest common subsequence) to capture different levels of granularity. However, n-gram metrics such as ROUGE can struggle with capturing semantic similarity if paraphrases are used. To address this, we also employed BertScore (Zhang* et al., 2020) with AlephBert (Seker et al., 2022) as its backbone. BertScore leverages the pre-trained language model to provide a more semantically meaningful evaluation of the summary.

Human Evaluation To validate the quality of model-generated summaries for the HeSum task, seven independent expert annotators evaluated a total of 186 summaries (62 per model) based on the same set of 62 reference articles. Annotators evaluated each summary using a 1-5 Likert scale (Likert, 1932) based on two key criteria: *coherence*, which assessed the summaries’ grammaticality and readability, and *completeness*, which measured the degree to which they capture the main ideas of the articles. To measure the consistency of the annotators’ scores, we calculated Krippendorff’s α (Krippendorff, 2018) for an interval scale, and received an α score of 0.78 which indicates a good inter-annotator agreement rate.

4.2 Results

Quantitative Analysis Table 3 summarizes the quantitative evaluation results. While mLongT5 consistently achieved higher ROUGE scores, a metric focused on surface-level similarity, GPT-based models exhibited superior performance in BertScore, a semantic similarity metric, and in human evaluation scores that assess coherence and completeness. The consistently higher ROUGE scores of mLongT5 might be partially attributed to limitations in the ROUGE metric itself. ROUGE scores favor summaries that closely mimic the source text, even if they lack originality or fluency. Additionally, n-gram-based metrics like ROUGE may discount grammatically correct sentences that convey the required meaning even with morphological or lexical word variations, or changes in word order, compared to the source text.

Furthermore, when comparing ROUGE to human evaluation we found a negative correlation of human evaluation with ROUGE scores. We computed Pearson correlation coefficients (PCC) and found the coefficients to be around -0.16 with highly statistically significant p-values (less than 2.39×10^{-5}), indicating that higher ROUGE scores do not in actuality correspond to human evaluations of good summaries. Similarly, using Kendall’s τ correlations resulted in negative values. Further research is needed for developing automatic summarization metrics that correlate with human scores.

Qualitative Analysis Following the identification of key error categories, we conducted a comparative analysis by randomly selecting 30 summaries generated by each of the three models for the same set of 30 articles. For each model, we then quantified the occurrences of each identified phenomenon within the sampled summaries. The results in Table 4 reveal disparities between the errors of GPT-based models and those of the finetuned mLongT5 on various linguistic phenomena.

The finetuned mLongT5 exhibits pronounced disruptions like repetition (3.33%) and exact copy of sections from the articles (13.33%), which weren't observed in the GPT-based results. However, the GPT-based models demonstrate errors in morphological phenomena specific to Hebrew, such as incorrect gender and wrong definiteness marking on *smixut*, indicating that the morphological richness of the language remains a challenge for these LLMs. Additionally, known phenomena of GPT-based models such as "hallucinations" (Cui et al., 2023; Guerreiro et al., 2023) are also observed in our analysis, as is familiar from other languages.

5 Conclusion

This research seeks to fill a critical gap in the field of LLMs assessment for generative creative tasks in MRLs, by presenting HeSum, a new dataset for Hebrew abstractive summarization, that includes 10K article-summary pairs sourced from professional journalists on Hebrew news websites. HeSum offers three key advantages: high level of abstractness in summarization, distinct challenges specific to the Hebrew language, and a thorough empirical assessment of LLMs using this dataset. By integrating these aspects, HeSum establishes itself as a valuable resource for researchers striving to push the boundaries of generative tasks, and specifically abstractive text summarization in Hebrew.

Limitations

Evaluation Metrics based on n-gram matching, such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and Meteor (Denkowski and Lavie, 2014), are commonly used for evaluating summarization quality in English. However, these metrics can be problematic when applied to Hebrew text. Hebrew allows for more flexible word order compared to English. Additionally, its morphological richness entails that the same concept can be expressed in multiple ways due to variations in pre-

fixes, suffixes, and root conjugations. Furthermore, Hebrew has variations in spelling words due to missing vowels (*Ktiv haser* vs. *Ktiv male*). These factors can lead to n-gram-based metrics overlooking grammatically correct sentences in the generated summary that convey the full meaning even though they show slight differences. Our findings in Section 4.2, which demonstrate a negative correlation between ROUGE scores and human evaluation scores, highlight the limitations of ROUGE evaluation in the context of Hebrew summarization.

Subset of LLMs Although we aspired to evaluate HeSum on a broad range of large language models (LLMs), our current analysis is limited to only two generative models. This might overlook newer models offering potentially superior performance. Additionally, resource constraints prevented us from investigating the behavior of these models in few-shot settings. Having acknowledged that, the timeliness of this resource is uncompromised, as it can be used with contemporary and future models alike, to track advances on this challenge.

Open Access vs. Domain Focus HeSum predominantly comprises articles from news websites, which may bias models' success in this task towards news-style writing, and may not fully capture the linguistic diversity across different genres and domains. The reason for selecting these domains specifically stems from our ability to obtain a permissive license for the resource, allowing open and free access by the community. However, the websites we have chosen – "Shakuf", "HaMakom", and "The Seventh Eye" – deviate from typical news platforms, offering a diverse range of topics that go beyond the typical content found on many popular news websites in Hebrew. This variety ensures that our dataset reflects a broader spectrum of real-world topics.

Dataset Scale vs. Quality In the realm of abstractive summarization, datasets like CNN/Daily Mail (Nallapati et al., 2016) and XSum (Narayan et al., 2018) are commonly employed. These datasets utilize news articles from websites, treating the article content as the document and a corresponding field (often not explicitly intended as a summary) as the summary. However, this approach has faced criticism due to uncertainty about whether the chosen field truly represents a summary (Tejaswin et al., 2021). An alternative approach involves human

summarization, but this tends to result in smaller datasets (e.g., PriMock57, Papadopoulos Korfiatis et al., 2022). To advance Hebrew NLP (and the study of generation in MRLs in general) beyond traditional classification tasks, there is a need for extensive generative datasets. Given the current lack of viable alternatives within the NLP community, we have adopted a similar approach to XSUM, albeit with longer summaries. Additionally, collecting human-generated summaries in low-resource languages presents challenges, including the scarcity of crowdsourcing platforms that support Hebrew. To ensure quality, we meticulously reviewed 100 articles, their subheadings, and brief introductory sentences. Ultimately, we chose subheadings as our summary source because they provided more informative content, capturing additional details from the articles. Furthermore, we filtered out articles with subheadings containing very few tokens (10 or fewer) to ensure our summaries adequately represent the article content.

Ethics

Following the generous permission of “Shakuf”, “HaMakom”, and “The Seventh Eye” — organizations committed to independent journalism, media scrutiny, and transparency in Israel — we were granted the valuable opportunity not only to access and analyze their published articles but also to publish the data for broader research use. This unique collaboration fosters open access and empowers other researchers to build upon the data extracted from their articles and our findings within Hebrew abstraction summarization, expanding knowledge in this important field. Also, we are guaranteed not to have offensive language or hate speech in our data. It should be borne in mind, however, that the opinions or biases reflected in these data may differ from other sources of information (news websites, social media, non-Hebrew news reports, and the like). So, the deployment of technology trained on this resource should be done with care.

Acknowledgements

We wholeheartedly thank Naomi Niddam (Editor-in-Chief, “Shakuf”), Nir Ben-Zvi (CEO, “Shakuf” and “The Seventh Eye”), Shuki Tausig (Editor-in-Chief, “The Seventh Eye”), and Ohad Stone (Creative and digital projects, “Shakuf”) for granting us access to and permission to analyze and share their articles for broader research purposes. We

also thank Yoav Goldberg, Ido Dagan, Amir David Nissan Cohen, and the participants of the NLP seminar at Bar-Ilan University for their valuable feedback. This research has been funded by a grant from the European Research Council, ERC-StG grant number 677352, a grant by the Israeli Ministry of Science and Technology (MOST), and a KAMIN grant from the Israeli Innovation Authority, for which we are grateful.

References

- Kabir Ahuja, Rishav Hada, Millicent Ochieng, Prachi Jain, Harshita Diddee, Samuel Maina, Tanuja Ganu, Sameer Segal, Maxamed Axmed, Kalika Bali, et al. 2023. Mega: Multilingual evaluation of generative ai. *arXiv preprint arXiv:2303.12528*.
- Google Translate API. 2023. [Google translate api v2 documentation](#).
- Dan Bareket and Reut Tsarfaty. 2021. Neural modeling for named entities and morphology (nemo²). *Transactions of the Association for Computational Linguistics*, 9:909–928.
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.
- Rishi Bommasani and Claire Cardie. 2020. Intrinsic evaluation of summarization datasets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8075–8096.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira,

- Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2023. [Palm: Scaling language modeling with pathways](#). *Journal of Machine Learning Research*, 24(240):1–113.
- Avihay Chriqui and Inbal Yahav. 2022. Hebert and hebemo: A hebrew bert model and a tool for polarity analysis and emotion recognition. *INFORMS Journal on Data Science*, 1(1):81–95.
- Aaron Daniel Cohen, Adam Roberts, Alejandra Molina, Alena Butryna, Alicia Jin, Apoorv Kulshreshtha, Ben Hutchinson, Ben Zevenbergen, Blaise Hilary Aguera-Arcas, Chung ching Chang, Claire Cui, Cosmo Du, Daniel De Freitas Adiwardana, Dehao Chen, Dmitry (Dima) Lepikhin, Ed H. Chi, Erin Hoffman-John, Heng-Tze Cheng, Hongrae Lee, Igor Krivokon, James Qin, Jamie Hall, Joe Fenton, Johnny Soraker, Kathy Meier-Hellstern, Kristen Olson, Lora Moys Aroyo, Maarten Paul Bosma, Marc Joseph Pickett, Marcelo Amorim Menegali, Marian Croak, Mark Díaz, Matthew Lamm, Maxim Krikun, Meredith Ringel Morris, Noam Shazeer, Quoc V. Le, Rachel Bernstein, Ravi Rajakumar, Ray Kurzweil, Romal Thoppilan, Steven Zheng, Taylor Bos, Toju Duke, Tulsee Doshi, Vincent Y. Zhao, Vinodkumar Prabhakaran, Will Rusch, YaGuang Li, Yanping Huang, Yanqi Zhou, Yuanzhong Xu, and Zhifeng Chen. 2022. Lamda: Language models for dialog applications. In *arXiv*.
- Amir Cohen, Hilla Merhav-Fine, Yoav Goldberg, and Reut Tsarfaty. 2023. [HeQ: a large and diverse Hebrew reading comprehension benchmark](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13693–13705, Singapore. Association for Computational Linguistics.
- Chenhang Cui, Yiyang Zhou, Xinyu Yang, Shirley Wu, Linjun Zhang, James Zou, and Huaxiu Yao. 2023. Holistic analysis of hallucination in gpt-4v (ision): Bias and interference challenges. *arXiv preprint arXiv:2311.03287*.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the ninth workshop on statistical machine translation*, pages 376–380.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Nuno M Guerreiro, Duarte M Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André FT Martins. 2023. Hallucinations in large multilingual translation models. *Transactions of the Association for Computational Linguistics*, 11:1500–1517.
- Eylon Gueta, Omer Goldman, and Reut Tsarfaty. 2023. Explicit morphological knowledge improves pre-training of language models for hebrew. *arXiv preprint arXiv:2311.00658*.
- Tahmid Hasan, Abhik Bhattacharjee, Md Saiful Islam, Kazi Samin, Yuan-Fang Li, Yong-Bin Kang, M Sohel Rahman, and Rifat Shahriyar. 2021. Xl-sum: Large-scale multilingual abstractive summarization for 44 languages. *arXiv preprint arXiv:2106.13822*.
- Valentin Hofmann, Janet Pierrehumbert, and Hinrich Schütze. 2021. [Superbizarre is not superb: Derivational morphology improves BERT’s interpretation of complex words](#). 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 3594–3608, Online. Association for Computational Linguistics.
- Omri Keren and Omer Levy. 2021. [ParaShoot: A Hebrew question answering dataset](#). In *Proceedings of the 3rd Workshop on Machine Reading for Question Answering*, pages 106–112, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- Viet Dac Lai, Nghia Trung Ngo, Amir Pouran Ben Veysseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023. Chatgpt beyond english: Towards a comprehensive evaluation of large language models in multilingual learning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, Singapore. Association for Computational Linguistics.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. 2022. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*.
- Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Aviya Maimon and Reut Tsarfaty. 2023. Cohesentia: A novel benchmark of incremental versus holistic assessment of coherence in generated texts. *arXiv preprint arXiv:2310.16329*.
- N Moratanch and S Chitrakala. 2016. A survey on abstractive text summarization. In *2016 International Conference on Circuit, power and computing technologies (ICCPCT)*, pages 1–7. IEEE.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [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, Berlin, Germany. Association for Computational Linguistics.
- 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*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Alex Papadopoulos Korfiatis, Francesco Moramarco, Radmila Sarac, and Aleksandar Savkov. 2022. [Pri-Mock57: A dataset of primary care mock consultations](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 588–598, Dublin, Ireland. Association for Computational Linguistics.
- 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.
- Tzuf Paz-Argaman, Tal Bauman, Itai Mondshine, Itzhak Omer, Sagi Dalyot, and Reut Tsarfaty. 2023. [HeGeL: A novel dataset for geo-location from Hebrew text](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7311–7321, Toronto, Canada. Association for Computational Linguistics.
- Tanya Reinhart. 1980. Conditions for text coherence. *Poetics today*, 1(4):161–180.
- Amit Seker, Elron Bandel, Dan Bareket, Idan Brusilovsky, Refael Greenfeld, and Reut Tsarfaty. 2022. [AlephBERT: Language model pre-training and evaluation from sub-word to sentence level](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 46–56, Dublin, Ireland. Association for Computational Linguistics.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, et al. 2022. Language models are multilingual chain-of-thought reasoners. *arXiv preprint arXiv:2210.03057*.
- Shaltiel Shmidman, Avi Shmidman, and Moshe Koppel. 2023. Dictabert: A state-of-the-art bert suite for modern hebrew. *arXiv preprint arXiv:2308.16687*.
- Priyam Tejaswin, Dhruv Naik, and Pengfei Liu. 2021. [How well do you know your summarization datasets?](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3436–3449, Online. Association for Computational Linguistics.
- Reut Tsarfaty, Dan Bareket, Stav Klein, and Amit Seker. 2020. [From SPMRL to NMRL: What did we learn \(and unlearn\) in a decade of parsing morphologically-rich languages \(MRLs\)?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7396–7408, Online. Association for Computational Linguistics.
- Reut Tsarfaty, Shoval Sadde, Stav Klein, and Amit Seker. 2019. [What's wrong with Hebrew NLP? and how to make it right](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 259–264, Hong Kong, China. Association for Computational Linguistics.
- David Uthus, Santiago Ontañón, Joshua Ainslie, and Mandy Guo. 2023. mlongt5: A multilingual and efficient text-to-text transformer for longer sequences. *arXiv preprint arXiv:2305.11129*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mt5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934*.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

A The HeSum Dataset

Collection Protocol Since the websites we collected (Shakuf, HaMakom, and The Seventh Eye) lack archives or RSS feeds, we developed a crawler to systematically navigate through pages, beginning from the homepage and exploring various article links. Leveraging their shared HTML structure, we could efficiently scrape the sites. We excluded pages without textual content, such as multimedia pages or those not in Hebrew. Additionally, articles with summaries of less than 10 tokens were filtered out, as they often lack sufficient detail to be a summary. In addition, all the articles were cleaned from Unicode characters or unrelated content.

Dataset	novel n-grams				CMP	RED (n=1)	RED (n=2)
	n = 1	n = 2	n = 3	n = 4			
CNN/Daily Mail	13.20	52.77	72.20	81.40	90.90	13.73	1.10
XSum	35.76	83.45	95.50	98.49	90.90	5.83	0.16
HeSum	42.00	73.20	82.00	85.36	95.48	4.83	0.10
HeSum (morpheme-based)	17.41	48.02	67.51	76.86	95.57	25.90	2.72

Table 5: HeSum’s Intrinsic Evaluation compared to English Benchmarks (CNN/Daily Mail and XSum).

Coherence

1. Very Incoherent: The summary is extremely confusing and lacks any clear connection between sentences.
2. Incoherent: The summary is somewhat understandable.
3. Somewhat Coherent
4. Coherent
5. Very Coherent

Completeness

1. Very Incomplete: The summary lacks essential information and does not convey the main points effectively.
2. Incomplete: The summary provides some information but misses key details.
3. Somewhat Complete
4. Complete
5. Very Complete

Figure 1: Evaluation Criteria

Human Evaluation Details We collected annotations from seven volunteered participants aged 25 and above, all with at least one academic degree. The participants were instructed to rate two parameters – *coherence* and *completeness*, based on known criteria, as depicted in Figure 1. While completeness measures the extent to which the summary captures all the essential information from the source text, coherence is a more complex metric. According to Reinhart (1980), coherence encompasses three core aspects: (i) cohesion, (ii) consistency, and (iii) relevance. While metrics like BertScore can also assess completeness, automatic evaluation of coherence remains a challenge (Mai-

Model	Rouge1	Rouge2	RougeL	Epochs	Loss
mLongT5-Base	18.62	8.68	15.92	18	2.15
mLongT5-Large	20.22	9.66	18.12	12	1.92

Table 6: mLongT5 performance on validation set and training details.

mon and Tsarfaty, 2023). Therefore, the measurement of coherence is evaluated in this work solely by humans.

Data Analysis Table 5 provides a quantitative analysis of HeSum’s summarization characteristics, highlighting its challenges. The analysis utilizes two tokenization approaches: word-based (above the dashed line) and morpheme-based (below the dashed line). This distinction allows for a deeper examination of the dataset’s abstractness and the influence of morphological features. As the table demonstrates, the number of unique vocabulary items (novel n-grams) decreases when using morpheme tokenization. However, HeSum still exhibits a higher uni-gram count compared to CNN/Daily Mail. This suggests that the task itself inherently involves a high degree of abstractness, and the morphological nature of the data presents an additional challenge.

B Models

Fine-tuning mLongT5 The HeSum corpus exhibits characteristics of long-form text, with an average document length of 2,747 tokens and a 90th percentile reaching 5,276 tokens. This extensive content poses challenges for the vanilla mT5 model, whose capacity for processing such lengths may be limited. Consequently, we have fine-tuned the mLongT5 model (Uthus et al., 2023), which is suitable for handling long inputs. The paper presents results obtained with the base version of mLongT5, which is 2.37 GB. We are also releasing

Dataset	novel n-grams				CMP	RED (n=1)	RED (n=2)
	n = 1	n = 2	n = 3	n = 4			
HeSum	42.00	73.20	82.00	85.36	95.48	4.83	0.10
GPT-4	47.24	80.35	91.37	95.92	91.89	8.14	0.68
GPT-3.5	45.69	80.18	91.73	96.35	93.46	7.53	0.83
mLongT5-large	8.26	30.10	43.50	50.21	95.39	11.89	5.98
mLongT5-base	7.21	28.77	42.06	49.46	92.25	15.74	10.25

Table 7: Intrinsic Evaluation of Summarization. A Comparative Analysis of GPT-4, GPT-3.5, mT5 Models and the Hesum Dataset.

Model	prefix	input	output	ROUGE-1	ROUGE-2	ROUGE-L
GPT-3.5	E	E	E	16.10	4.06	11.43
GPT-3.5	H	H	H	16.34	4.26	11.69
GPT-3.5	E	E	H	12.80	2.30	11.00
mLongT5	—	H	H	17.47	7.56	14.68
GPT-3.5	E	H	H	17.08	4.95	12.46
GPT-3.5	H	H	E	14.35	3.13	9.89
GPT-3.5	E	H	E	14.31	3.11	10.40
GPT-3.5	H	E	H	15.90	4.23	10.80

Table 8: Testing different configurations of language prompting to find the best configuration to evaluate GPT-3.5. 'H' denotes Hebrew and 'E' denotes English. 'prefix' is the instruction to the model, 'input' is the article itself, and the output is the desired summarization language.

a larger model (4.56 GB).⁵ The training regimen employed an 8-GPU A100 cluster for 36 hours for the large model, while the base model leveraged a single A100 GPU with 40 GB of memory. Early stopping, utilizing ROUGE-1 as the metric, was implemented to optimize the training process. Further details regarding model performance and implementation specifics are provided in Table 6.

Prompting GPT-based models Here, we leverage the translate-English approach, suggested by (Shi et al., 2022) and (Ahuja et al., 2023), which translates instances from target languages into English before prompting. We decompose the prompt task into three parts: (i) the input article (ii) the instruction (prefix), and (iii) the output. All three parts could be done in either Hebrew or English for the HeSum task. In our experiment, Google Translate API (2023, API, 2023) handled the translation of prompts (input and/or prefix) from Hebrew to English and the translated outputs back to Hebrew for analysis. Testing GPT-3.5 on different configurations of language prompting in the HeSum validation set, we found that the best prompt-language

⁵<https://huggingface.co/biunlp/mT5LongHeSum-large>

configuration is English-English-English (Table 8). We then applied this prompting strategy to both GPT-3.5 and GPT-4 on the test set. The prompt we used depicted in Figure 2.

You are a genius summarizer. Your task is to summarize the main points of the following text. Please follow these instructions step by step:

1. The summary should be concise, consisting of up to 3 sentences.
2. If there are several main topics, create a separate sentence for each topic.
3. The output should be in English.

Figure 2: The prompt we used for the GPT-based models

C Implementation Details

For the intrinsic evaluation of the dataset, we created a Jupyter notebook which computes the different metrics. For computing the n-grams, we used the NLTK package,⁶ and for loading and processing the data, we used NumPy⁷ and Pandas.⁸ For evaluation of the different models, we used the most common ROUGE package for non-English papers,⁹ and the HuggingFace implementation of Transformers for BertScore.¹⁰

⁶<https://pypi.org/project/nltk/>

⁷<https://pypi.org/project/numpy/>

⁸<https://pypi.org/project/pandas/>

⁹https://github.com/csebuetnlp/xl-sum/tree/master/multilingual_rouge_scoring

¹⁰<https://pypi.org/project/transformers/>

D Additional Models Performance Analysis

Table 7 presents the intrinsic evaluation results for the models, corresponding to the metrics introduced in Section 3.3. Notably, GPT-based models generate text with greater abstractness, as evidenced by their higher count of novel n-grams compared to the fine-tuned mT5. This finding aligns with mT5’s tendency towards repetitive generation, which is further supported by its high RED score and by the qualitative analysis presented in Table 4.