LCHAIM - Investigating Long Context Reasoning in Hebrew

Anonymous ACL submission

Abstract

Natural Language Inference (NLI) has gained significant attention recently due to its importance in understanding how machines comprehend and reason about language. While English has received tremendous interest, Morphologically Rich Languages (MRLs) like Hebrew, require more research. In this paper, we address the evaluation of Hebrew NLI models by introducing LCHAIM, a dataset specifically designed to evaluate Hebrew NLI models on tasks involving long premises and complex reasoning. The dataset, created by translating and validating the English ConTRoL dataset, consists of 8,325 context-hypothesis pairs that require advanced reasoning such as coreferential, temporal, and analytical reasoning. Our findings highlight the difficulty of contextual reasoning in Hebrew, as evidenced by the performance of different models. Fine-tuning the LongHero model on both the shorter premise Hebrew NLI and the LCHAIM datasets yielded an accuracy and an F1-Score of 52%, that is 35% less than human performance. Similarly, Large language Models (LLMs) like Gemma-9B and Dicta-LM-2.0-7B achieved mean accuracies of 57% and 39% respectively in zero-shot setting. Additionally, our results indicate that fine-tuning or providing a few shots from a different Hebrew NLI dataset (HebNLI) generally does not improve the performance on the more complex task of LCHAIM.

1 Introduction

NLI, also called Textual Entailment (TE) is a key task in natural language processing (Dagan et al., 2013; Bowman et al., 2015). It involves determining whether a given hypothesis can be logically inferred from a premise. This task has shown to be very helpful in various applications, including text classification, event extraction, and summarization evaluation, (Yin et al., 2019; Sainz et al., 2022; Scirè et al., 2024). NLI has gained significant attention recently due to its importance in understanding how machines comprehend and reason about language. This is largely because almost any task can be generalized to determining entailment or contradiction in context between texts (Liu et al., 2021a).

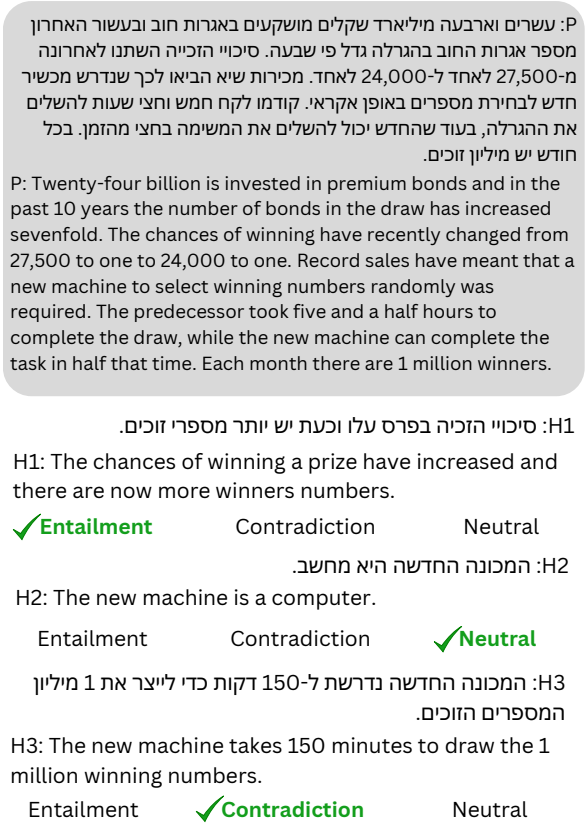

Figure 1: A sample of NLI task from the proposed LCHAIM dataset and CoNTRoL dataset.

Although many datasets have been developed to train and evaluate NLI models (Dagan et al., 2005; Giampiccolo et al., 2007; Bowman et al., 2015; Williams et al., 2018; Welleck et al., 2019, See Figure 2), these datasets primarily focus on sentence-level examples, which do not fully capture the complexity of real-world language understanding.

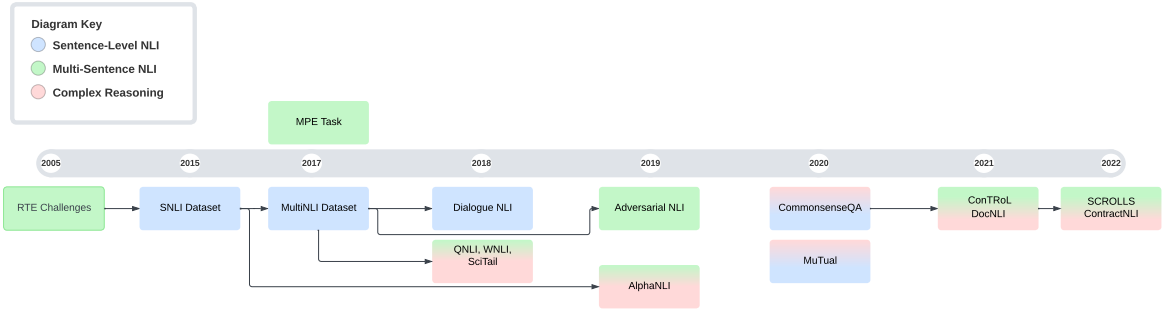

Figure 2: Timeline of related datasets and works published in single-sentence, and multi-sentence NLI.

To address this issue, Liu et al. (2021a) proposed ConTRoL, a long context (premise) NLI dataset that enabled investigating contextual reasoning for NLI [1]. The example in Figure 1 demonstrates Mathematical reasoning, which requires an understanding of calculus to solve it.

However, despite large advances in NLP, research and datasets predominantly focus on English, leaving rich morphological languages like Hebrew underrepresented. On the other hand, the rich morphological structure of Hebrew presents significant challenges for Natural Language Understanding (NLU) tasks, leading for example to inconsistencies and errors in evaluation metrics (Cohen et al., 2023).

Recent studies have increased interest in Hebrew (Seker et al., 2022; Gueta et al., 2022), but the emphasis remains largely on morpho-syntactic tasks. Despite the existence of a few NLI datasets for MRLs (Klemen et al., 2024; Halat and Atlamaz, 2024; Jallad and Ghneim, 2022), to the best of our knowledge, no existing dataset specifically addresses the challenges associated with long premises that require complex reasoning. Furthermore, none of the aforementioned datasets is in Hebrew.

To address these gaps, we present in this paper a first investigation of long context reasoning in Hebrew. As part of this investigation, we compile LCHAIM - **L**ong **C**ontext **H**ebrew with **A**dvanced reasoning **I**nference **M**odel Benchmark - a Hebrew version of ConTRoL (Liu et al., 2021a). LCHAIM contains 8,325 premise-hypothesis pairs in Hebrew, labeled for contradiction, entailment, or neutral relationship between premise p and hypothesis h, and is obtained from the ConTRoL dataset via automatic translation and human validation. Then, we provide an evaluation of Hebrew Language Models for the task of NLI, using the resulting dataset. We show that the task of NLI with contextual reasoning is challenging even for larger language models, which shows the need for the construction of complex NLU benchmarks in Hebrew.

The research contribution is three-fold:

1. We present LCHAIM, a dataset for training, and validating models on complex reasoning in Hebrew.

2. We propose, validate, and document A translation methodology for translating datasets from English to other languages.

3. We perform detailed experiments to evaluate Hebrew models on the task of long-premise NLI with complex reasoning.

2 LCHAIM Dataset

In this section we describe the LCHAIM dataset. Since we closely follow the ConTRoL paper, most of the key statistics and descriptions can be found in detail in the corresponding paper. We address the dataset as LCHAIM, but being a translated version of ConTRoL, the basic properties also hold for the original dataset.

The LCHAIM dataset is compiled from publicly available online verbal reasoning tests, including those used in police initial recruitment exams, medical college admissions, and university clinical aptitude tests, as well as corporate verbal aptitude assessments. These tests are structurally similar to NLI tasks, presenting a premise and a hypothesis with three answer choices: true, false, or cannot say. These choices correspond to the NLI labels ENTAILMENT, CONTRADICTION, and NEUTRAL.

[1] Other English datasets with long context NLI include (Yin et al., 2021; Koreeda and Manning, 2021; Shaham et al., 2022), but these do not specifically include complex reasoning.

The data format of LCHAIM aligns with existing NLI benchmarks (Bowman et al., 2015; Williams et al., 2018), where each instance includes a premise, a hypothesis, and a label from ENTAILMENT, NEUTRAL, or CONTRADICTION. A notable difference from existing datasets is that ConTRoL features longer premises, often spanning one or more passages. Additionally, each premise is paired with three or more hypotheses, distinguishing the original ConTRoL, and therefore LCHAIM from traditional NLI datasets that typically associate each premise with a single hypothesis.

The verbal reasoning tests require candidates to comprehend meaning, evaluate logical strength, make valid inferences, and identify appropriate conclusions. The passages cover various topics, including current affairs, business, science, the environment, economics, history, meteorology, health, and education. These questions are of high quality and are used in rigorous assessments, reflecting a high level of difficulty. After removing duplicates, the LCHAIM dataset consists of 8,325 context-hypothesis pairs. Lexical overlap analysis shows only 4.87% overlap for ENTAILMENT pairs and 5.49% for CONTRADICTION pairs, indicating that the dataset presents significant challenges for simple lexical matching techniques.(Liu et al., 2021a)

Inferring from the original ConTRoL dataset, and with the hypothesis translation does not affect the reasoning types present in the dataset at large. LCHAIM contains various reasoning tasks such as Coreferential Reasoning (Ye et al., 2020), Verbal Logical Reasoning (Liu et al., 2021b), Temporal and Mathematical Reasoning (Nakhimovsky, 1987), Information Integration over Paragraphs (Welbl et al., 2018), and Analytical Reasoning (Williams et al., 2019).

In the remaining of this section, we provide a detailed account of the data preparation and validation procedures. More specifically, we describe the automatic translation process, the comparison between different translations, and the manual validation performed.

2.1 Translation

The LCHAIM dataset was created via translation of the previously mentioned ConTRoL dataset, a large NLI dataset with especially long premises that aims to address contextual reasoning. To ensure the optimal translation we employ two translation models, Amazon Translate[2], a Neural Network-based machine translation service, and DictaLM Translation, a Hebrew-LLM-based translation service (Shmidman et al., 2024). We translate the entire ConTRoL dataset using both translation models using the API of both services.

2.2 Human Validation

To determine which translation is superior for this task, we perform a manual evaluation of the translation.

To evaluate the performance of Amazon Translate and DictaLM, we employed a stratified random sampling approach. From our dataset of 8,325 samples, we randomly selected 200 samples, maintaining the proportional representation of the original train, test, and development splits.

Two human annotators (two authors of the paper) independently reviewed and ranked each translated sample based on three ordinal scale parameters and one binary parameter:

1. Translation Accuracy: The degree of precision in translating from English to Hebrew (Popović et al., 2006).

2. Fluency: The naturalness and coherence of the translated premise and hypothesis (Graham et al., 2013).

3. Cultural Fit: The extent to which the translation preserves the cultural nuances and appropriateness of the original English text in Hebrew (Nida, 1964).

4. Label Validity: Binary, whether the label is still valid after the translation of the premise and hypothesis.

Each ordinal parameter was scored on a scale of 1 to 5, with 5 representing the highest quality and 1 the lowest. Additionally, the annotators assessed whether the original label (contradiction, entailment, or neutral) remained valid after translation, ensuring the preservation of the logical relationship between the premise and hypothesis in the target language.

Following this validation process, we calculated the average annotators' score for each of the three parameters (Translation Accuracy, Fluency, and Cultural Fit). We also computed the percentage of cases where the original label was maintained

[2] https://aws.amazon.com/translate/

after translation. These results are presented in Table 1, providing a comprehensive overview of the performance of both Amazon Translate and DictaLM across all evaluated aspects.

Amazon Translate was selected as the better translation in the context of this dataset, with a mean of accuracy, fluency, and cultural fit 85.57%, 82.53% ,74.49% respectively. The percentage of samples with original labels remaining valid after translation was 97.8% for Amazon Translate and 94% for DictaLM Translation. This evaluation methodology allowed us to comprehensively assess the quality and reliability of both automatic translation models in the context of our specific task, with quantitative measures to support our choice.

We believe this translation and validation methodology can be replicated for other NLI datasets and languages.

3 Experiments

3.1 Hebrew Pre-trained Language Models

We implement and evaluate the performance of SOTA pre-trained language models in Hebrew to demonstrate their NLU capabilities, proxied by their performance on LCHAIM.

AlephBERT AlephBERT (Seker et al., 2022) is a Hebrew version of the classic BERT (Devlin et al., 2019) pre-trained language model with a large vocabulary. It was trained on three datasets: the Hebrew portion of the OSCAR corpus cleaned Twitter texts, and the entire Hebrew Wikipedia. Both variants utilize wordpieces with a vocabulary size of 52,000. AlephBERT was trained only using the masked token prediction objective, excluding next-sentence prediction.

LongHero LongHero (Shalumov and Haskey, 2023) is a variant of the HeRo model [3], designed to handle long sequences more effectively. It uses a BPE tokenizer with a vocabulary size of 50,265 tokens, trained on the HeDC4 dataset.

We implement both models following the original ConTRoL paper implementation using the last hidden state and a classification head (For more information see Appendix A). For AlephBERT, we use the proposed methodology by Devlin et al. (2019). We adhere to the original implementation details, using the same dataset split into training, development, and test sets with an 8:1:1 ratio (6,692:837:836 samples for training, validation and test respectively) and training all models for 10 epochs (see Appendix B for hyperparameter optimization details). The maximum sequence length is set to 512 tokens for AlephBERT and 4,096 for LongHero.

To further investigate the performance of these models in different settings, we employ fine-tuning in 3 ways. Firstly, we fine-tune both AlephBERT and LongHero using the HebNLI[4] dataset, a translated version of the original SNLI dataset (Bowman et al., 2015). This is done to investigate the effect of fine-tuning using a short-premise NLI dataset for the task of long-premise NLI in LCHAIM. Second, we perform further fine-tuning of the same models using the train set (80% of samples) of the LCHAIM dataset. We report Accuracy over the test set as the main metric, and Precision (P), Recall (R), and F1-Score (F1) for each class.

3.2 In-Context Learning with LLMs

We further evaluate SOTA Hebrew and multilingual LLMs for the LCHAIM NLI task.

Dicta-LM-2.0-7B Dicta LM 2.0 (Shmidman et al., 2024) is a Hebrew LLM with a custom tokenizer, continuously pre-trained from Mistral-7B-v0.1 (Jiang et al., 2023). The pertaining is done using a Hebrew corpus, and further Supervised Fine Tuning is performed using custom Hebrew datasets.

Gemma-9B The Gemma-9B model (Team et al., 2024) is a multilingual LLM, designed with a transformer decoder architecture and trained on a context length of 8192 tokens. It incorporates several advanced techniques to enhance its performance. Multi-Query Attention (Shazeer, 2019) is utilized, featuring multi-head attention to improve efficiency. Rotary Positional Embeddings (RoPE Su et al., 2024) are employed instead of absolute positional embeddings, with shared embeddings across inputs and outputs to reduce model size. The standard ReLU activation function is replaced with GeGLU (Shazeer, 2020) and RMSNorm (Zhang and Sennrich, 2019) is applied to normalize the inputs to each transformer sub-layer, contributing to more stable training.

Inference with LLMs was performed using the Huggingface Transformers package, by parsing

[3] HeRo (Shalumov and Haskey, 2023) is a Hebrew language model based on the RoBERTa architecture (Liu et al., 2019).

[4] https://github.com/NNLP-IL/HebNLI

Model	Accuracy	Fluency	Cult. Fit	Label Validity
Amazon Translate	**85.57**%	**82.53**%	75.49%	**97.8**%
DICTA-LM-2.0	81.13%	81.92%	**78.53**%	94%

Table 1: Translation evaluation metrics for Amazon Translate and DICTA-LM-2.0.

the LLM's response, results are reported over the test set only. During the prompt engineering process, several methodologies were investigated to enhance the model's performance. Initially, we employed a zero-shot approach, wherein only free-text instructions were provided to the model to guide the task. However, this method proved insufficient, as it was overly simplistic for the models to accurately predict the specific labels, resulting in incorrect outputs.

We then implemented a prompt that included detailed instructions regarding the structure of the desired output (see Appendix C). This adjustment led to an improvement in the model's performance. Finally, we introduced examples, known as the "few-shot"(Brown, 2020) approach or in-context learning, which involved providing the model with several instances of NLI tasks along with the correct labels.

Additional Experiments To obtain a human performance baseline ,we perform a manual evaluation, where 4 of the authors, Hebrew native speakers, test 200 random samples from the LCHAIM dataset.

To further investigate the complexity of the various types of reasoning, and the increased context length available in LCHAIM, we test AlephBERT, LongHero, and Gemma-9B using the shorter-premise HebNLI dataset, a translated version of the original Stanford NLI (SNLI) (Bowman et al., 2015), for comparison. We initially ran models without fine-tuning, to highlight the impact of fine-tuning and compare baseline performance between models.

4 Results

Table 2 presents the results for the models tested. AlephBERT, achieved an overall mean accuracy of 39.5%, with an F1-Score of 36%. However, the model showed relatively uneven calibration, consistently predicting the entailment class in the test set, after accounting for overfitting through class balance, vanishing gradient, and varying model complexity, we assume this behavior indicates the difficulty of the task at hand. Fine-tuning AlephBERT using the HebNLI dataset achieved 37.5% mean accuracy and 35.42% F1 Score resulting in a -2%, and -1.5% absolute difference respectively, compared to the non fine-tuned model.

Fine-tuning only using LCHAIM, without the HebNLI step resulted in a mean accuracy of 38.4% and an F1 Score of 32.4%, which indicates an absolute difference of -1%, and -3% respectively. Continuous fine-tuning on both HebNLI and LCHAIM resulted in a mean accuracy of 42.6% and an F1 Score of 41.3% or an overall 3%, and 5% absolute increase respectively. This means for AlephBERT the best setting is fine-tuning on both the short premise HebNLI and the long premise LCHAIM dataset, these settings also showed better calibration than the original pre-trained version, with relatively balanced class prediction.

LongHero, which has a long enough context to be able to accommodate the long premises, was able to achieve a mean accuracy and an F1-Score of 34.5%, and 33.7% respectively without fine-tuning. The model displayed some imbalance towards the entailment class. Fine-tuning LongHero using the shorter premise HebNLI resulted in an accuracy of 32% and an F1 Score of 25%, or an absolute difference of -2% and -8% compared to the non-fine-tuned version. Fine-tuning only using the LCHAIM training set resulted in an accuracy of 41% and an F1 Score of 34% or an absolute difference of 7% and 1% respectively. Furthermore, fine-tuning LongHero on both HebNLI, and then on LCHAIM, resulted in the best overall performance of 52% mean accuracy and F1 Score, an absolute performance difference of 18% and 19% respectively. This can also be viewed as a relative increase of an impressive 35% increase in mean accuracy and F1 Score, which indicates a significant increase in performance. This model also showed the overall best calibration between the 3 predicted classes. The best model for the task was a fine-tuned LongHero using both HebNLI and LCHAIM.

For the LLMs evaluated, Gemma-9B obtained

Model	Acc.	F1	Entailment			Contradiction			Neutral		
			P	R	F1	P	R	F1	P	R	F1
Human (English)	87.06	93.15	94.83	95.65	95.24	93.33	91.21	92.26	93.02	91.91	92.95
Human (Hebrew)	84.97	84.26	77.05	87.04	81.74	91.23	78.79	84.55	85.33	87.67	86.49
AlephBERT	39.50	35.94	42.47	68.19	52.34	31.38	24.38	27.44	39.13	15.25	21.95
AlephBERT-HebNLI	37.51	35.42	42.43	56.57	48.49	28.92	14.46	19.28	33.06	34.74	33.88
AlephBERT-LCHAIM	38.38	32.39	42.61	60.85	50.12	32.82	44.21	37.67	25.00	0.01	0.02
AlephBERT-ALL	42.60	41.32	46.26	56.88	51.02	36.45	43.38	39.62	45.21	22.03	29.63
LongHero	34.53	33.74	39.42	46.17	42.53	28.7	33.47	30.91	32.85	19.49	24.46
LongHero-HebNLI	31.55	24.92	50.77	10.09	16.84	34.48	12.40	18.24	29.25	**80.93**	42.97
LongHero-LCHAIM	41.49	34.68	41.94	**81.96**	55.49	44.12	12.40	19.35	36.73	15.25	21.56
LongHero-ALL	**52.17**	**52.20**	**59.74**	57.19	**58.44**	**47.56**	**44.22**	**45.82**	**47.19**	53.39	**50.10**

Table 2: Performance results for different models over the LCHAIM test set. Metrics include overall accuracy, F1 score, Precision (P), Recall (R), and F1 scores for each class: Entailment, Contradiction, and Neutral. -HebNLI indicates fine-tuning using HebNLI, -LCHAIM indicates fine-tuning using LCHAIM, and -ALL indicates sequential fine-tuning on HebNLI and then LCHAIM. Human (English) and Human (Hebrew) indicate performance of four educated testees of 300, and 200 samples from ConTRoL and LCHAIM respectively. Bold indicates the best performing model.

Model	Shots' Dataset	0-shot	1-shot	2-shot	3-shot
Gemma-9B	HebNLI	**56.15%**	50.93%	49.94%	48.45%
Gemma-9B	LCHAIM	**56.15%**	55.65%	49.57%	49.32%
Dicta-LM-2.0-7B	HebNLI	39.38 %	40.45%	40.99%	40.62%
Dicta-LM-2.0-7B	LCHAIM	39.38%	42.36%	40.00%	41.37%

Table 3: Mean accuracy of LLMs Gemma-9B and Dicta-LM-2.0-7B over the LCHAIM test set, under 0-shot, 1-shot, 2-shot, and 3-shot settings. The 'Shots' Dataset column indicates the dataset used for gathering samples for the shots. Best performance model and setting in bold. Gemma-9B scored the best result, 56.15% for the 0-shot setting.

a mean accuracy of 56.15% for the zero-shot settings, and performed relatively worse for the few-shot setting with 50.93%, 49.94%, and 48.45% for 1,2,and 3-shot prompting respectively. This is measured when considering few-shots sampled from the shorter premises HebNLI dataset. Randomly sampling from the longer premise LCHAIM dataset, was unattainable due to context length restrictions of the model, and the long premises in LCHAIM, therefore, we randomly sampled from shorter premise instances from the LCHAIM dataset to serve as few-shot evaluation for the Gemma-9B model. Zero-shot setting resulted in 56.15% accuracy, and performance worsened with few-shot prompting to 55.65%, 49.57% and 49.32% for 1, 2, and 3 shot respectively. Dicta-lm-2.0-7B obtained a mean accuracy of 39% for zero-shot setting, and a mean accuracy of 40.45%, 40.99% and 40.62% in the 1,2, and 3-shot prompting, indicating a mere 1.5% absolute increase in accuracy for few-shot prompting compared to zero-shot. This is measured for shots randomly sampled from HebNLI. For shots sampled from LCHAIM, Dicta-lm-2.0-7B obtained a 42.36%, 40% and 41.37% for 1,2, and 3-shots prompting respectively, indicating a slight increase from zero-shot prompting. Results are presented in Table 3.

As per the manual evaluation, human performance are calculated and displayed in Table 2. The testers were able to achieve a mean accuracy of 84.97% with an F1 Score of 84.26%, showcasing a significant difference of 33% absolute accuracy and F1 Score compared to the best LongHero model. Human performance also show a difference of 28% mean accuracy compared to the best performing LLM, Gemma-9B.

Result for the shorter HebNLI are presented in Table 4. The base AlephBERT model achieved a mean accuracy and F1 Score of 53.50% and 53.42% respectively without fine-tuning. While the LongHero base model achieved a 47.17% and 46.74% mean accuracy and F1 Score respectively. Fine-tuning using the HebNLI datasets resulted in a 15% and 27% absolute increase in accuracy for AlephBERT and LongHero respectively. However, fine-tuning using the LCHAIM dataset, resulted in

a 17% and 10% absolute **decrease** in performance for both models respectively, consistently showing that fine-tuning using LCHAIM reduces the model's ability when considering the HebNLI task.

Fine-tuning using first the LCHAIM and then HebNLI dataset resulted in a mean accuracy and F1 Score of of 78.28% and 78.24% for AlephBERT, offering similar results to fine-tuning only with HebNLI . The same approach resulted in a mean accuracy and F1 Score of 83.93% and 83.78% respectively for the LongHero model, a 5.5% increase in performance compared to fine-tuning only with HebNLI.

The best model for the HebNLI test set (LongHero), achieved a 31.76% absolute higher mean accuracy compared to the best model for the LCHAIM test set (also LongHero). Both models showed good calibration over both test sets. For comparison, SOTA performance for HebNLI is reported in the Hebrew LLM leaderboard, with a mean accuracy of 95.48% for Qwen2.5-72B [5].

5 Error Analysis

To further explore the performance of the tested models, we perform an error analysis. With inspiration from Liu et al. (2021a), we investigate the percentage of errors across premise lengths. This is done with the hypothesis that longer texts usually require a higher degree of reasoning, therefore encompassing a harder task. Results are presented in Figure 3. The error analysis showed no significant change in mean accuracy as text lengths increase for LongHero. Therefore, we hypothesise the complexity of the task is not only correlated to the text length, rather it can be correlated to the vocabulary, and the reasoning types (coreferential, verbal logical, temporal and mathematical, information integration, and analytical) present in the premise-hypothesis pair. This is contrary to the English equivalent, which showcased lower performance across several models, as premise length increased, which could possibly indicate the differences between English and Hebrew.

6 Discussion

The results of our evaluation of the LCHAIM dataset shed light on the challenges and performance of models dealing with long-premise NLI tasks in Hebrew. Our findings show both the strengths and limitations of the current models and how dataset characteristics and translation quality impact their performance. To the best of our knowledge, LCHAIM is the first long premise, complex reasoning, NLI dataset available in Hebrew.

[5] https://huggingface.co/spaces/hebrew-llm-leaderboard/leaderboard

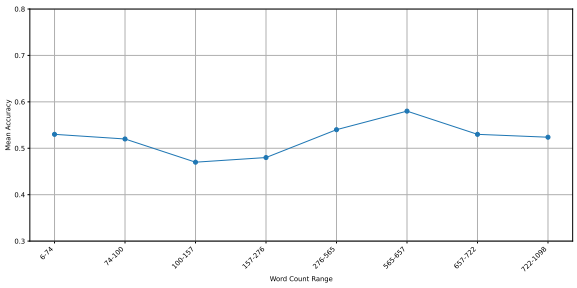

Figure 3: Mean accuracy for LongHero model, across different premise word length bins, over the LCHAIM test set. 100 samples minimum in each bin.

Model Performance The AlephBERT model had difficulties with long-premise NLI tasks. It achieved a mean accuracy of 39.50% and an F1-Score of 35.94%, showing that it struggles with extended contexts and tends to over-predict entailment. Fine-tuning AlephBERT on the HebNLI dataset and on the LCHAIM train set decreased its performance slightly. However, an improvement in performance was seen when further fine-tuning the model with LCHAIM after the initial HebNLI fine-tuning, resulting in a mean accuracy and F1-Score of 42.60% and 41.32% respectively. This suggests the shorter, more simple task of HebNLI does not contribute towards the harder, longer-premise LCHAIM test set performance on its own, but can provide a good foundation for further fine-tuning. Therefore, including the long-premise data from LCHAIM is essential for better performance in the longer premise tasks. With that said, performance was still considerably lower then the 84.97% mean accuracy human performance.

LongHero, which is designed to handle longer contexts, achieved a mean accuracy and F1-Score of 34.53% and 33.74% respectively, without fine-tuning. This model had also calibration issues similar to AlephBERT. It is worth noting the English equivalent models (BERT-base, and Longformer) did not experience these calibration issues when considering the English version of this dataset (Liu et al., 2021a). The best results were seen with the fine-tuned LongHero using both HebNLI and then LCHAIM, with a mean accuracy and F1 Score of 52.17% and 52.20% respectively, indicating that

Model	Acc.	F1	Entailment			Contradiction			Neutral		
			P	R	F1	P	R	F1	P	R	F1
AlephBERT	53.50	53.42	56.14	55.22	55.31	56.11	59.38	57.84	46.23	45.57	45.40
AlephBERT-HebNLI	78.73	78.68	78.21	83.88	80.95	80.81	78.17	79.47	77.01	73.62	75.28
AlephBERT-LCHAIM	36.42	31.02	36.68	61.18	45.86	36.08	41.36	38.54	36.00	0.03	0.06
AlephBERT-ALL	78.28	78.24	77.71	82.56	80.01	80.39	78.82	79.60	76.54	72.89	74.67
LongHero	47.17	46.74	48.22	58.22	52.75	50.51	47.55	49.99	41.22	34.43	37.52
LongHero-HebNLI	74.52	73.86	76.68	81.13	81.30	72.36	74.46	76.97	72.80	62.57	67.25
LongHero-LCHAIM	37.30	31.46	37.22	79.80	51.12	38.77	14.10	21.66	32.06	15.04	21.60
LongHero-ALL	**83.93**	**83.78**	**87.41**	**86.84**	**87.13**	**84.19**	**85.01**	**84.60**	**79.78**	**79.48**	**79.63**

Table 4: Performance results for different models using the shorter-premise HebNLI test set. Metrics include overall accuracy, F1 score, Precision (P), Recall (R), and F1 scores for each class: Entailment, Contradiction, and Neutral. -HebNLI indicates fine-tuning using HebNLI, -LCHAIM indicates fine-tuning using LCHAIM, and -ALL indicates sequential fine-tuning with LCHAIM and then HebNLI. Numbers in bold indicate the best performing model setting.

using a combination of datasets is key to getting the best performance.

For large language models like Gemma-9B and Dicta-LM-2.0-7B, performance varied across zero-shot and few-shot settings. Gemma-9B did well in the zero-shot setting but declined with more shots from the shorter HebNLI or the longer LCHAIM, possibly due to the context length required. Dicta-LM-2.0-7B however, showed slight improvements with few-shot prompting from HebNLI or LCHAIM.

Performance over the shorter HebNLI dataset were significantly higher than the performance over the LCHAIM test set considering the sequentially fine-tuned AlephBERT and LongHero. These models showed a 25-30% higher mean accuracy and F1 Score considering the HebNLI test set, which further shows the complexity of the LCHAIM dataset. With caution, we hypothesise this is an indicator of the complexity of the reasoning in the LCHAIM dataset. LongHero achieved slightly better performance than AlephBERT.

Complex Reasoning in LCHAIM We hypothesised that the LCHAIM dataset is a long reach to Hebrew models because, unlike previous NLI benchmarks, it contains tasks that require complex reasoning. One example shows a premise and hypothesis pair which require temporal reasoning. To infer the hypothesis entails the premise, the model needs to understand the 17th century is previous to the 18th century (see Figure 4).

Like ConTRol, Hebrew models struggle to achieve human performance with LCHAIM. Unlike ConTRoL, Hebrew models show bad calibration, which could indicate specific complexities in Hebrew.

P: **ההיסטוריונים** של ימינו שואפים לבנות **תיעוד של פעילויות אנושיות ולהשתמש בתיעוד זה כדי להשיג הבנה עמוקה יותר של האנושות**. תפיסה זו של משימתם היא עדכנית למדי, המתוארכת להתפתחות **מהמאה ה -18 ותחילת המאה ה -19** של ההיסטוריה המדעית, וטופחה במידה רבה על ידי היסטוריונים מקצועיים שאימצו את ההנחה כי חקר הפעילות האנושית הטבעית והבלתי נמנעת. לפני סוף המאה ה -18, ההיסטוריה לא נלמדה כמעט בשום בתי ספר, והיא לא ניסתה לספק פרשנות לחיי האדם בכללותם. זה מתאים יותר לתפקידה של הדת, של הפילוסופיה, או אולי אפילו של השירה.

P: Today's **historians** aim to build **a record of human activities and use this documentation to gain a deeper understanding of humanity**. This conception of their mission is fairly recent, dating back to **the 18th and early 19th** century development of scientific history, and has been largely fostered by professional historians who have adopted the assumption that the study of natural and inevitable human activity. Before the end of the 18th century, history was not taught in almost any schools, and it did not attempt to provide an interpretation of human life as a whole. This is more suited to the role of religion, of philosophy, or perhaps even of poetry.

H1: במאה ה -17 ההיסטוריה לא הייתה נחשבת כדרך להבנת האנושות

H1: In the 17th century history would not have been considered as a way of understanding humanity.

✓**Entailment** Contradiction Neutral

Figure 4: Sample from the LCHAIM dataset which requires temporal reasoning to solve.

7 Conclusion

This study introduced the LCHAIM dataset, a benchmark for long-premise NLI task in Hebrew, that focuses on complex reasoning types like coreferential, temporal, and analytical reasoning. Our experiments show that current state-of-the-art models perform far below human levels. The LCHAIM dataset provides a tough challenge for future research, underscoring the need for better models and methods. By making LCHAIM and our code available, we hope to encourage further exploration and advancements in Hebrew NLU research.

Limitations

While this study advances one's knowledge about Hebrew NLI, several limitations should be acknowledged. First, our analysis did not explore the performance across different reasoning types, such as coreferential, temporal, and analytical reasoning. A more granular evaluation could offer deeper insights into how models handle specific reasoning challenges in Hebrew NLI tasks. Second, in the scope of this study we used four models, primarily Hebrew ones. While these models represent strong baselines, they may not fully capture the capabilities of larger and more advanced, multilingual LLMs, such as Command-R-plus, which has shown proficiency in Hebrew. Addressing these limitations in future work could enhance the robustness and generalizability of our findings. Future work should focus on testing larger models like Aya-35B, Command-R-Plus-104B, GPT-4o, and longer context oriented models like those based on the Mamba architecture (Gu and Dao, 2023). These advanced models may offer better performance and a deeper understanding of complex reasoning.

References

Samuel Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642.

Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

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.

Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine learning challenges workshop*, pages 177–190. Springer.

Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzoto. 2013. *Recognizing Textual Entailment: Models and Applications*. Morgan and Claypool.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1, page 2. Minneapolis, Minnesota.

Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and William B Dolan. 2007. The third pascal recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, pages 1–9.

Yvette Graham, Timothy Baldwin, Alistair Moffat, and Justin Zobel. 2013. Continuous measurement scales in human evaluation of machine translation. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 33–41.

Albert Gu and Tri Dao. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.

Eylon Gueta, Avi Shmidman, Shaltiel Shmidman, Cheyn Shmuel Shmidman, Joshua Guedalia, Moshe Koppel, Dan Bareket, Amit Seker, and Reut Tsarfaty. 2022. Large pre-trained models with extra-large vocabularies: A contrastive analysis of hebrew bert models and a new one to outperform them all. *arXiv preprint arXiv:2211.15199*.

Mustafa Halat and Ümit Atlamaz. 2024. Implicatr: A granular dataset for natural language inference and pragmatic reasoning in turkish. In *Proceedings of the First Workshop on Natural Language Processing for Turkic Languages (SIGTURK 2024)*, pages 29–41.

Khloud Al Jallad and Nada Ghneim. 2022. Arnli: Arabic natural language inference for entailment and contradiction detection. *arXiv preprint arXiv:2209.13953*.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Matej Klemen, Aleš Žagar, Jaka Čibej, and Marko Robnik-Šikonja. 2024. Si-nli: A slovene natural language inference dataset and its evaluation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14859–14870.

Yuta Koreeda and Christopher D Manning. 2021. Contractnli: A dataset for document-level natural language inference for contracts. *arXiv preprint arXiv:2110.01799*.

Hanmeng Liu, Leyang Cui, Jian Liu, and Yue Zhang. 2021a. Natural language inference in context-investigating contextual reasoning over long texts. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 15, pages 13388–13396.

Jian Liu, Leyang Cui, Hanmeng Liu, Dandan Huang, Yile Wang, and Yue Zhang. 2021b. Logiqa: a challenge dataset for machine reading comprehension with logical reasoning. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3622–3628.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Alexander Nakhimovsky. 1987. Temporal reasoning in natural language understanding: The temporal structure of the narrative. In *Third Conference of the European Chapter of the Association for Computational Linguistics*.

Eugene Albert Nida. 1964. *Toward a science of translating: with special reference to principles and procedures involved in Bible translating*. Brill Archive.

Maja Popović, Adrià de Gispert, Deepak Gupta, Patrik Lambert, Hermann Ney, José B Mariño, Marcello Federico, and Rafael E Banchs. 2006. Morpho-syntactic information for automatic error analysis of statistical machine translation output. In *Proceedings on the Workshop on Statistical Machine Translation*, pages 1–6.

Oscar Sainz, Itziar Gonzalez-Dios, Oier Lopez de Lacalle, Bonan Min, and Eneko Agirre. 2022. Textual entailment for event argument extraction: Zero-and few-shot with multi-source learning. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2439–2455.

Alessandro Scirè, Karim Ghonim, and Roberto Navigli. 2024. Fenice: Factuality evaluation of summarization based on natural language inference and claim extraction. In *Findings of the Association for Computational Linguistics: ACL 2024*.

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.

Uri Shaham, Elad Segal, Maor Ivgi, Avia Efrat, Ori Yoran, Adi Haviv, Ankit Gupta, Wenhan Xiong, Mor Geva, Jonathan Berant, et al. 2022. Scrolls: Standardized comparison over long language sequences. *arXiv preprint arXiv:2201.03533*.

Vitaly Shalumov and Harel Haskey. 2023. Hero: Roberta and longformer hebrew language models. *arXiv preprint arXiv:2304.11077*.

Noam Shazeer. 2019. Fast transformer decoding: One write-head is all you need. *arXiv preprint arXiv:1911.02150*.

Noam Shazeer. 2020. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*.

Shaltiel Shmidman, Avi Shmidman, Amir DN Cohen, and Moshe Koppel. 2024. Adapting llms to hebrew: Unveiling dictalm 2.0 with enhanced vocabulary and instruction capabilities. *arXiv preprint arXiv:2407.07080*.

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. 2024. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063.

Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.

Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. 2018. Constructing datasets for multi-hop reading comprehension across documents. *Transactions of the Association for Computational Linguistics*, 6:287–302.

Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2019. Dialogue natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3731–3741.

Adina Williams, Nikita Nangia, and Samuel R Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of NAAfCL-HLT*, pages 1112–1122.

Chad C Williams, Mitchel Kappen, Cameron D Hassall, Bruce Wright, and Olave E Krigolson. 2019. Thinking theta and alpha: Mechanisms of intuitive and analytical reasoning. *NeuroImage*, 189:574–580.

Deming Ye, Yankai Lin, Jiaju Du, Zhenghao Liu, Peng Li, Maosong Sun, and Zhiyuan Liu. 2020. Coreferential reasoning learning for language representation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7170–7186.

Wenpeng Yin, Jamaal Hay, and Dan Roth. 2019. Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3914–3923.

Wenpeng Yin, Dragomir Radev, and Caiming Xiong. 2021. Docnli: A large-scale dataset for document-level natural language inference. *arXiv preprint arXiv:2106.09449*.

Biao Zhang and Rico Sennrich. 2019. Root mean square layer normalization. *Advances in Neural Information Processing Systems*, 32.

A NLI Task with Pre-trained Language Models

Given a premise p and a hypothesis h, we concatenate them into a new sequence: $[\texttt{CLS}] + p + [\texttt{SEP}] + h + [\texttt{SEP}]$, where $[\texttt{CLS}]$ denotes the classification token and $[\texttt{SEP}]$ serves as the separator token. After encoding this sequence with the pre-trained model, the hidden representation of the $[\texttt{CLS}]$ token from the final layer is fed into a custom Multi-Layer Perceptron (MLP) followed by a softmax function for classification. The custom MLP consists of three hidden layers followed by an output layer. The first hidden layer performs a linear transformation from the input size to 300 units, followed by ReLU activation, layer normalization, and dropout with a rate of 0.3. The second hidden layer reduces this to 100 units with similar operations, and the third hidden layer further reduces it to 50 units. The final output layer maps the 50 units to 3 output classes for classification. Each hidden layer includes ReLU activation and layer normalization.

B Hyperparameters Optimization Details

This section outlines the hyperparameter optimization process employed during model training. Hyperparameters were optimized through an extensive grid search across a variety of configurations to achieve optimal model performance. Batch sizes of 8, 16, and 32 were tested, while learning rates ranging from $1e^{-5}$ to $5e^{-5}$ (specifically $1e^{-5}$, $2e^{-5}$, $3e^{-5}$, $4e^{-5}$, $5e^{-5}$) were explored to determine the optimal convergence rate. Gradient accumulation steps of 1, 2, and 4 were also evaluated. These combinations enabled us to systematically test different hyperparameter settings and identify the best configuration for each model.

C Prompt Structure and Classification Guidelines

This section provides details of the Hebrew prompt used in our study for NLI tasks. In our study, we provided the models with a Hebrew prompt for classifying the relationship between a premise and a hypothesis (illustrated in Figure 5). The English translation mirrors the Hebrew prompt and was included for clarity in the paper.

In the Hebrew version provided to the models, the prompt instructs them to classify the premise-hypothesis relationship as follows:

c (Contradiction): Hypothesis contradicts the premise.

English version	Hebrew version
I will provide a premise and an hypothesis, classify whether the hypothesis contradicts the premise (c), entails the premise (e) or is neutral towards it (n). You must answer with the one letter representing the relation between the premise and hypothesis only. Examples of a response: Answer: e Answer: c Answer: n Classify for the following premise and hypothesis. Premise: {p} Hypothesis: {h} Answer:	אני אתן לך פסקה ומשפט ואתה תצטרך לסווג האם המשפט סותר את הנאמר הפסקה (ס) , נגזר מהנאמר מהפסקה (מ) או לא סותר את הנאמר בפסקה ולא נגזר מהנאמר בפסקה (נ) עליך לענות באמצעות אחת האותיות המייצגות את הקשר של המשפט לטקסט בלבד דוגמאות לתשובות: תשובה: מ תשובה: ס תשובה: נ ענה על הפסקה והמשפט הבאים פסקה: {p} משפט: {h} תשובה:

Figure 5: The prompt used in our research for NLI tasks. The left side shows the English translation, while the right side displays the original Hebrew version provided to the models. The prompt instructs the model to classify the relationship between a given premise and hypothesis as either contradiction (c), entailment (e), or neutral (n), using a single letter for the response.

e (Entailment): Hypothesis follows from the premise.

n (Neutral): Hypothesis neither contradicts nor follows.

The model responded with a single letter (c, e, or n), and examples were provided to guide the responses.