

Research Article

A Dutch Financial Large Language Model

Sander Noels¹, Jorne De Blaere², Tijl De Bie¹

1. Ghent University, Gent, Belgium; 2. Independent researcher

This paper presents FinGEITje, the first Dutch financial Large Language Model (LLM) specifically designed and optimized for various financial tasks. Together with the model, we release a specialized Dutch financial instruction tuning dataset with over 140,000 samples, constructed employing an automated translation and data processing method. The open-source data construction method is provided, facilitating the creation of financial instruction datasets in different languages. To evaluate model performance, the study introduces the first Dutch financial evaluation benchmark, along with an automated evaluation method that utilizes an LLM as an independent evaluator, reducing manual intervention in performance evaluation. The experimental results highlight the superior performance of FinGEITje across five critical Dutch and English financial tasks.

Corresponding authors: Sander Noels, sander.noels@ugent.be; Jorne De Blaere, jorne.deblaere@silverfin.com; Tijl De Bie, tijl.debie@ugent.be

CCS Concepts

- Computing methodologies → Knowledge representation and reasoning;
- Computing methodologies → Natural language generation;
- Information systems → Information retrieval.

1. Introduction

Recent advances in Natural Language Processing (NLP) have significantly transformed the way we analyze unstructured financial data^{[1][2]}. The advent of LLMs has further expanded the capabilities of NLP, demonstrating exceptional performance across various financial tasks. However, substantial challenges remain, particularly with non-English languages such as Dutch, which are underrepresented by current financial models^[3].

The highly technical nature of financial texts requires domain-specific LLMs to understand complex financial language and concepts effectively^{[1][2]}. The financial sector struggles to create these models due to the scarcity of annotated datasets and high data annotation costs^[4]. In addition, existing financial LLMs mainly focus on English, creating a knowledge gap in localized financial contexts in languages such as Dutch. As a result, Dutch financial

documents are often misinterpreted, causing inefficiencies and inaccuracies when using general models for Dutch financial tasks.

Despite approximately 24 million Dutch speakers worldwide, existing Dutch LLMs have not focused on financial documents. Tailored financial models have consistently shown superior performance in domain-specific tasks compared to generalized LLMs. This emphasizes the necessity of developing specialized models, such as an open-source, transparent, and accessible financial LLM for the Dutch language. Such a model could democratize financial data analysis and foster advances in both academia and industry^{[4][5]}.

Our main contributions are summarized as follows:

- We introduce FinGEITje, the first open-source, transparent, and accessible Dutch financial LLM.
- We provide the first extensive Dutch financial instruction tuning dataset and a comprehensive evaluation benchmark for Dutch financial LLMs.
- We release a methodology for constructing financial instruction tuning datasets in multiple languages, supporting the development of financial LLMs across linguistic contexts.
- We introduce an automated evaluation method using an LLM as an independent evaluator, enabling large-scale evaluations.

FinGEITje offers valuable applications for financial professionals and researchers, including sentiment classification of financial news and social media posts, key entity identification in financial documents, and news headline classification to validate claims about price movements. It also supports financial relation extraction and can answer specific financial queries, thus streamlining information retrieval for financial decision-making processes.

The remainder of this paper is as follows. Section 2 reviews the related work. Section 3 details our proposed model, FinGEITje, followed by an empirical evaluation in Section 4. Section 5 provides the conclusions of our study. Section 6 explores potential limitations and future research directions. Finally, Section 7 discusses the ethical considerations associated with our work.

2. Related Work

To understand the context of our work, we review the literature on financial NLP models. We explore the challenges these models face, the necessity of financially-tailored LLMs, the scarcity of non-English language financial models, and the importance of Dutch-focused financial LLMs. Lastly, we consider the need for open-source and transparent financial LLMs to foster progress in this specialized field.

2.1. Challenges of Financial NLP Models

NLP has made significant strides in handling unstructured financial data, turning complex documents into clear insights and valuable market information. These advances have opened up various capabilities, ranging from stock

movement prediction to more advanced financial tasks^[3]. However, traditional NLP models face notable challenges in the financial sector.

One primary issue is the intrinsic complexity of financial language and jargon, which often results in a gap in understanding domain-specific documents^{[1][2][6]}. This complexity requires sophisticated models capable of accurately interpreting specialized terminology. Another significant barrier is the scarcity of annotated datasets, exacerbated by the high cost associated with data annotation in this specialized field^[4]. A third challenge lies in the limited inferential capabilities of traditional models, particularly in critical tasks such as making informed investment decisions^[4]. Lastly, the widespread adaptability of many NLP models remains constrained, as they are often optimized for a single task and lack the ability to generalize across multiple tasks^[7].

Due to these challenges, it is essential to develop more advanced, versatile, and robust NLP models tailored to the dynamic and complex requirements for understanding financial documents. In this evolving landscape, the need for financially tailored LLMs has become increasingly evident.

2.2. The Success of Financially-tailored LLMs

The financial industry has recently seen a surge in the application of LLMs due to their remarkable capabilities in performing various tasks using natural language instructions, which often do not require extensive retraining on specific data^{[8][9][10][11][12]}. Unlike traditional financial NLP models that typically excel only in narrowly defined tasks, LLMs exhibit substantial versatility across various financial applications.

This progress witnessed a significant boost with FinBERT, a domain-specific adaptation of BERT, showing that fine-tuning on financial documents greatly increases performance in financial tasks^[1]. The success of FinBERT spurred a series of innovative models, notably BloombergGPT, FinGPT, FinMA, and FinTral, each advancing the integration of financial data and improving model performance^{[2][13][5][14]}.

Earlier models were limited by their relatively small parameter sizes, typically below one billion parameters, which constrained their generalization capabilities. The push towards multi-billion parameter models has significantly improved their ability to generalize and perform across various financial tasks^[5]. This scalability plays a crucial role in the effectiveness of models in handling complex financial documents and generating precise and insightful outputs.

2.3. The Scarcity of Non-English Financial LLMs

One major issue in the current state of financial LLMs is the lack of models for languages other than English. Because English is a globally prominent language for international finance as well as machine learning research, it has been the focus for the majority of the existing work on financial LLMs. However, the comparative lack of models available in other languages represents a significant knowledge gap, preventing an in-depth understanding of localized financial contexts and nuances. The development of LLMs in these other languages is dependent on the availability of

sufficient datasets and appropriate evaluation methods. Although extensive datasets are available for English, there is a notable scarcity of such datasets for other languages^[3].

The development of financial LLMs in languages other than English is still at an early stage. Notable examples include the Mengzi-fin^[15] and BBT-FinT5^[16] models, both uniquely designed for the Chinese language. In addition to non-English language-specific models, multilingual financial LLMs, such as FinMA-ES^[17] and ICE-PIXIU^[18], have been developed, which have been fine-tuned on Spanish and English data and Chinese and English data, respectively. These models attempt to capture the linguistic nuances of the financial domain in various languages.

Language-specific financial benchmarks are a crucial factor for evaluating the impact and effectiveness of these financial LLMs in diversified linguistic contexts^[3]. The gap between English and other languages in the domain of financial language models is the absence of financial benchmarks tailored to those languages. A worthy mention is a newly established benchmark for Japanese financial LLMs, although it does not accompany a specified LLM^[19].

Table 1 compares the different financial LLMs available today across several parameters, including language compatibility. The addition of Chinese-specific models and multilingual models indicate a diversification in language applicability beyond English. For a complete analysis of the existing financial LLMs, see^[3].

Model	Backbone	Size	Open Source		Instruct	Language	Release Date
			Model	Data			
FinBERT ^[1]	BERT	110M	✓	✓	✗	English	08/2019
Mengzi-fin ^[15]	RoBERTa	103M	✓	✗	✗	Chinese	10/2021
BBT-FinT5 ^[16]	T5	220M	✓	✓	✗	Chinese	02/2023
BloombergGPT ^[2]	BLOOM	50B	✗	✗	✗	English	03/2023
FinGPT ^[4]	LLaMA	7/13B	✓	✓	✓	English	10/2023
FinMA ^[5]	LLaMA	7/30B	✓	✓	✓	English	06/2023
FinMA-ES ^[17]	LLaMA2	7B	✓	✓	✓	English, Spanish	02/2024
FinTral ^[12]	Mistral	7B	✗	✗	✓	English	02/2024
ICE-PIXIU ^[18]	LLaMA2	7B	✓	✓	✓	English, Chinese	03/2024

Table 1. Comparison of pre-existing LLMs in the financial domain.

2.4. The Importance of Dutch-focused Financial LLMs

The implementation of financial LLMs is largely centered around a few global languages, mainly English and Chinese, with little attention to languages like Dutch. This lacks a sufficient focus on understanding financial contexts where the Dutch language is dominant. An estimated 24 million people converse in Dutch around the world according to the Dutch Taalunie (Language Union), suggesting a considerable demand in developing a Dutch-specific financial LLM.¹

Generalized models, particularly those predominantly trained on English data, face a significant obstacle in their inability to generate Dutch sentences. They often revert to English after a few sentences or provide inaccurate translations from English to Dutch^[20]. This tendency leads to a loss of the finer nuances and subtleties inherent in the Dutch language. For a Dutch-specific LLM to retain its credibility and applicability, it is important to preserve linguistic subtleties and maintain the natural flow of the language, while avoiding the interweaving of English terms.

Moreover, the scarcity of Dutch-centered LLMs is not only due to the deficiency of pre-established, pre-trained models. It is also a reflection of the significant gap in critical supporting infrastructure, which includes specialized datasets, benchmarks, and leaderboards^[20]. In an attempt to address this gap, several Dutch pre-trained models have been introduced, such as GEITje-7B², GEITje-7B-ultra-sft³, GEITje-7B-ultra⁴, fietje-2b⁵, and Reynaerde-7B-Chat⁶. However, their focus does not specifically fall on Dutch financial documents, thus limiting their effectiveness in financial applications.

In summary, the need for LLMs proficient in Dutch financial contexts cannot be overstated. Developing Dutch-centric financial LLMs is crucial for generating insightful and accurate interpretations from financial documents, capturing the richness and specificity of Dutch financial language.

2.5. The Need for Open-Source Financial LLMs and Their Transparency

The rise of proprietary financial LLMs, such as BloombergGPT^[2], have utilized their exclusive access to specialized and nuanced financial data, leading to a pressing need for greater accessibility, transparency, and democratization of financial data^[1].

Open-source alternatives, such as FinGPT^[4], have shown the value of democratization in this domain, demonstrating similar or even superior performance to proprietary models with a lower investment in resources^[13]. However, the progress made in creating open-source financial LLMs, which primarily focus on English, has overlooked the development of Dutch or other non-English specific financial models.

The creation of a larger suite of open-source Dutch LLMs will accelerate both academic and industrial research in automated language generation, empowering non-expert users to employ generative AI tools in their native language. The open-source approach creates an honest culture, supports fair practices, and allows for clear checks on how the model works and behaves.

3. Methodology

In this section, we provide a detailed explanation of the methodology adopted for constructing FinGEITje, the first Dutch financial LLM. We begin by outlining the dataset preparation, proceed with the instruction tuning process, and end with a description of our Dutch financial evaluation benchmark.

3.1. Dataset

The dataset forms the backbone of our financial LLM, FinGEITje. Here, we detail the process of collecting raw data, constructing instruction datasets, and translating them into Dutch. We also discuss deduplication and filtering processes to maintain data quality.

3.1.1. Raw Data

The development of a high-quality Dutch financial LLM requires an effective dataset. We constructed an instruction tuning dataset specifically for performing Dutch financial tasks. Our dataset creation methodology draws much of its inspiration from FinGPT^[4]. Following FinGPT, we mainly rely on open-source data instead of depending on the self-instruct method^[21], such as Alpaca, due to several advantages: high-quality annotations by domain experts, minimal costs, no restrictions on commercial use, and a diverse range of text types and modalities. These sources include news articles, financial reports, tweets, and multi-modal data such as time series data and tables. We transform the English datasets into a Dutch instruction tuning dataset using an automated translation and data processing method, explained in Section 3.1.3.

Data	Task	Raw Data	Instruction Data	Data Types	Modalities	Source
FPB ^[22]	sentiment analysis	4,845	4,845	news	text	FinGPT/ fingpt-sentiment-train
FiQA-SA ^[23]	sentiment analysis	1,213	1,213	news headlines, tweets	text	FinGPT/ fingpt-sentiment-train
TFNS ^[24]	sentiment analysis	9,543	9,543	tweets	text	FinGPT/ fingpt-sentiment-train
NWGI	sentiment analysis	16,184	16,184	news	text	FinGPT/ fingpt-sentiment-train
NER ^[25]	named entity recognition	609	609	financial agreements	text	FinGPT/fingpt-ner
NER (CLS) ^[25]	named entity recognition	609	17,051	financial agreements	text	FinGPT/fingpt-ner-cls
Headline ^[26]	news headline classification	11,412	102,708	news headlines	text	FinGPT/fingpt-headline
FinRed ^[27]	relationship extraction	6,768	32,670	news, earning call transcripts	text	FinGPT/fingpt-finred
Finance Alpaca	question answering	68,912	68,912	news headlines, tweets, earnings reports	text	gbharti/finance-alpaca
FiQA-QA ^[23]	question answering	17,110	17,110	earnings reports	text, table	FinGPT/fingpt-fiqa_qa
ConvFinQA ^[28]	question answering	12,594	12,594	earnings reports	text, table	FinGPT/fingpt-convfinqa

Table 2. Summary of datasets used for training the Dutch financial LLM.

Table 2 provides an overview of the key datasets used, encompassing various financial tasks:

Sentiment analysis evaluates the sentiment of financial texts such as news articles and tweets, classifying them as positive, negative, or neutral. **Named entity recognition** identifies and classifies key financial entities, such as persons, organizations, and locations, facilitating the construction of financial knowledge graphs. **News headline classification** examines financial headlines, determining the accuracy of statements about price movements. **Relation extraction** seeks to identify and extract various financial relationships present within textual data, while the **question**

answering task automatically provides answers to financial questions based on text-based information. For more details on our dataset construction, please refer to our GitHub repository.⁷

3.1.2. Instruction Dataset Construction

The construction of an English financial instruction tuning dataset serves as the starting point in this process, aligning each raw data sample with one or more instruction prompts relevant to specific financial tasks. The creation of additional instruction prompts is crucial as it effectively increases the overall number of data samples^[13].

Our prompt construction is based on an innovative method tailored to the needs of financial LLMs. An illustration of how our prompting method operates is depicted in Figure 1. In our methodology, we attribute a memetic proxy^[22] to the model, simulating the persona of a financial expert. Following the memetic proxy, we present the task-linked instruction prompts, which further constrain the model to refine its output. The input provides additional financial context, and the response completes the instruction output.

The proper formulation of these instructions is pivotal in instruction fine-tuning, as they act as guides in training the model to perform specific tasks. From FinGPT, we adopt the instruction formulations for all tasks^[13], except Finance Alpaca, which derives instructions based on the self-instruct method^[21]. Table 2 provides information on how many instruction tuning samples are generated for each dataset. For a more detailed explanation of how the English financial instruction tuning datasets were construed for the English FinGPT model, we refer to the FinGPT repository.⁸

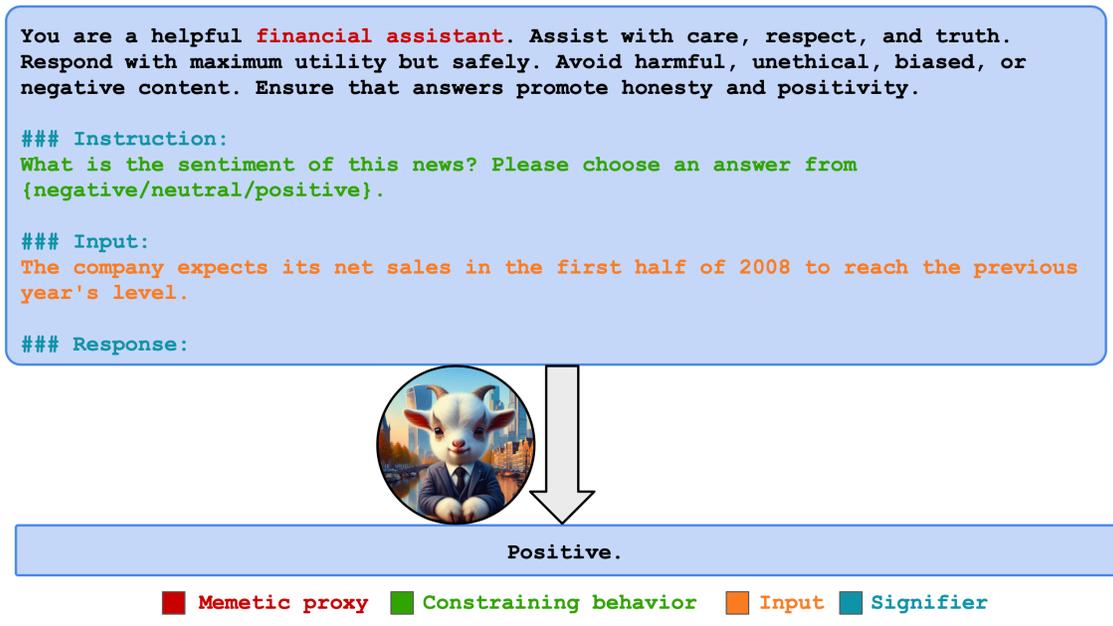


Figure 1. The prompting method of FinGPT.

3.1.3. Translation Approach

To translate the instruction datasets to Dutch, we made use of gpt-3.5-turbo-0125⁹ of OpenAI. It should be noted that due to the extensive size of the datasets, only limited manual verification was performed on the translated texts. Consequently, some translated content might contain inaccuracies or non-factual details. We provide a script, inspired by^[20], that uses OpenAI API services or open-source models to translate instruction tuning datasets.¹⁰ This approach can be adapted to translate datasets into any language supported by the translation model. This systematic but flexible translation approach has the potential to develop a diverse range of LLMs that surpass the boundaries of the English language. For more details on the translation process, see the Appendix A.

Given this context, it is crucial to acknowledge the potential drawbacks associated with automated translation. For instance, "translationese" effects might occur in the datasets, where the translated versions retain properties of the source text, such as word order and literal translation. This translation issue might lead to a biased advantage for non-Dutch models when evaluating on benchmarks. If these benchmarks were manually translated or originated in Dutch, it could limit the performance of non-Dutch models^[20].

3.1.4. Deduplication and Filtering

Data duplication is a common issue in collected datasets, which can obstruct the efficiency and effectiveness of further model training. To mitigate this, we performed an exact deduplication process designed to identify and remove exact duplicate samples within our dataset.

Dataset	Instructions
Total English Instructions	283,439
Total Translated Instructions	246,883
Total Filtered and Deduplicated Instructions	147,788

Table 3. The count of instruction tuning samples throughout the data cleaning pipeline.

After deduplication, we proceeded with additional filtering. We applied a language identification filter, which excluded all non-Dutch documents post-translation. This method ensured that our dataset mainly contains Dutch language samples, which is necessary for developing a Dutch-specific financial LLM. To identify and filter out potential poor translations, an additional script was used. The language identification tool *fasttext-language-identification*¹¹ was used to confirm if the instruction tuning samples were correctly translated into Dutch. Additional checks were carried out to ensure the presence of 'Signifiers' and a minimum length of three tokens in the translations. This filtering process

is mainly inspired by^[20]. Table 3 shows the number of instruction tuning samples throughout the pipeline. For a more comprehensive understanding of our deduplication and filtering pipeline, please refer to our repository.¹²

3.2. FinGEITje

In this section, we present FinGEITje, our Dutch financial LLM. We describe the model architecture, training details, and future directions to enhance its capability in the Dutch financial sector.

For FinGEITje, we use Mistral-7B-v0.1^[30] as our base model. This model is known for its superior performance, outperforming the LLaMA 13B model on numerous benchmarks. Most financial LLMs are built on smaller LLaMA models, but the robust Mistral BPE tokenizer, which efficiently segments numbers into single digits, makes it especially suitable for numerical tasks in the financial domain^[3].

Domain-Specific Pretraining. To implement language-specific adaptations, we use the foundation model GEITje-7B¹³, a Dutch base model capable of handling sequences up to 8,192 tokens, and trained on over 10 billion tokens of Dutch text. This model serves as the linguistic backbone of FinGEITje.

Financial Instruction Tuning. Instruction tuning^[31] aligns a model more closely with specific tasks while requiring less computational overhead compared to complete retraining. To efficiently handle instruction tuning, we employ FlashAttention^[32]. Additionally, we use QLoRA^[33] to perform instruction fine-tuning, focusing on all linear layers, known to provide performance similar to full fine-tuning but with markedly reduced resource demands. We set a batch size of 8 and a learning rate of $2.0e-04$. A 10% warmup phase is applied, ensuring smooth adaptation to the new data. FinGEITje undergoes a single epoch of training on 2 NVIDIA A100 40GB GPUs. Detailed parameters and training configurations are available on the model card [snoels/FinGEITje-7B-sft](#).¹⁴

Preference Optimization. While fine-tuning FinGEITje has provided a strong performance baseline, future iterations will incorporate preference optimization techniques, which are aimed at better aligning the model outputs with user expectations and improving response quality in interactive applications^[34].

3.3. Dutch Financial Evaluation Benchmark

Evaluation benchmarks provide valuable information on the performance of a model, allowing us to identify key areas for improvement. Unfortunately, the availability of benchmarks for Dutch LLMs is disproportionately low^[20]. There are none existing within the financial context. Currently, evaluation datasets for Dutch LLM have been created by translating English benchmarks into Dutch^[35], which lacks a genuine representation of the Dutch language.

To evaluate the performance of a Dutch financial LLM, a comprehensive and diverse benchmark is required that encompasses various key financial tasks. To this end, we introduce the first Dutch financial evaluation benchmark. It is designed based on our instruction tuning dataset construction method (see Section 3.1.2). This benchmark covers critical tasks essential for effective assessment, such as Sentiment Analysis (SA), News Headline Classification (HC), Named Entity Recognition (NER), Relationship Extraction (RE), and Question Answering (QA).

Dataset	Task	Test Samples	Evaluation Metrics
FPB ^[22] , FIQA-SA ^[23] , TFNS ^[24] , NWGI	SA	5,788	F1, Accuracy
NER (CLS) ^[25]	NER	1959	F1, Accuracy
Headline ^[26]	HC	9,094	F1, Accuracy
FinRed ^[27]	RE	1,790	F1, Accuracy
ConvFinQA ^[28]	QA	1,453	Accuracy

Table 4. Summary of datasets used in the Dutch financial evaluation benchmark.

Table 4 provides an overview of the datasets included in the benchmark and the respective key evaluation metrics. Along with our instruction dataset, we release the Dutch financial evaluation benchmark and the corresponding code for automated model evaluation on Github¹⁵ and Huggingface¹⁶. By providing a comprehensive, task-specific benchmark, we aim to encourage more robust and meaningful comparisons and evaluations of Dutch financial LLMs.

4. Experiments and Results

In this section, we discuss the experiments conducted to evaluate the performance of various LLMs, including our proposed FinGEITje, against our newly developed Dutch financial benchmark. We also discuss the methodology used for answer extraction and present detailed benchmarking results for both Dutch and English datasets.

4.1. Experiments

This subsection outlines the experimental setup and methodologies used to evaluate FinGEITje and other LLMs on our newly developed Dutch financial evaluation benchmark, providing insights into benchmarking and answer extraction methods.

4.1.1. Benchmarking Methods

We benchmark the performance of several LLMs, including FinGEITje, on our evaluation benchmark. We utilize two general Dutch instruction fine-tuned models: GEITje-7B-ultra-sft^{17[20]} and GEITje-7B-ultra¹⁸. Furthermore, two English-based LLaMa-models^[36] designed explicitly for financial applications, FinMA^[5] and FinGPT^[4], are selected. Furthermore, we evaluate a closed source, general-purpose GPT-3.5 model from OpenAI, gpt-3.5-turbo-0125¹⁹. We assess the zero-shot performance of all the models against 500 randomly chosen test samples to decrease the computational cost.

4.1.2. Answer Extraction

Some LLMs without task-specific fine-tuning fail to generate answers pre-defined in the instructions^{[5][14]}. Models might deviate from the task or fail to follow instructions completely. This discrepancy becomes problematic during model performance evaluation, particularly when outputs in a specific format are required, such as for QA and NER tasks. To address this, we propose applying an independent LLM as an evaluator tasked with extracting critical information from the generated responses to a task. We provide a script that uses the OpenAI API or open-source models to post-process the generated answers in the appropriate format as specified by the instruction.²⁰ This method could be instrumental in conducting large-scale evaluations where human intervention is not viable. Details on the answer extraction process can be found in Appendix B.

4.2. Benchmark Results

In this subsection, we present the comprehensive results of different LLMs against our Dutch-specific financial benchmark. Additionally, we evaluate the adaptability of FinGEITje by testing its performance against an English financial benchmark. This comparative analysis highlights the effectiveness of our Dutch financial LLM across different languages and tasks.

4.2.1. Results on Dutch Data

In this section, we present the results of different LLMs against our Dutch-specific financial benchmark. As shown in Table 5, our proposed LLM, FinGEITje, outperforms the other models across all tasks, indicating the effectiveness of language-specific instruction tuning for domain-specific tasks. Interestingly, the Dutch-specific models, which have been fine-tuned for general Dutch tasks, struggle when applied to financially-focused benchmarks. Additionally, OpenAI's gpt-3.5-turbo-0125 demonstrates strong zero-shot performance, highlighting its ability to perform tasks without detailed fine-tuning.

Dataset	Metrics	FinGEITje-7B-sft		geitje-7b-ultra		geitje-7b-chat		FinMA		FinGPT		gpt-3.5-turbo-0125	
		raw	extracted	raw	extracted	raw	extracted	raw	extracted	raw	extracted	raw	extracted
SA	Acc	0.790	0.790	0.564	0.674	0.454	0.540	0.276	0.632	0.268	0.350	<u>0.742</u>	<u>0.752</u>
	F1	0.790	0.790	0.555	0.662	0.466	0.520	0.136	0.645	0.113	0.276	<u>0.729</u>	<u>0.740</u>
NER (CLS)	Acc	0.840	0.912	<u>0.238</u>	0.662	<u>0.238</u>	0.670	0.236	<u>0.776</u>	0.222	0.386	<u>0.238</u>	0.690
	F1	0.851	0.916	<u>0.092</u>	0.684	<u>0.092</u>	0.693	0.091	<u>0.792</u>	0.089	0.357	<u>0.092</u>	0.712
HC	Acc	0.920	0.920	0.082	0.064	0.314	0.298	0.000	0.670	0.000	<u>0.696</u>	<u>0.606</u>	0.640
	F1	0.836	0.836	0.069	0.025	0.215	0.166	0.000	0.459	0.000	0.000	<u>0.466</u>	<u>0.486</u>
RE	Acc	0.569	0.569	0.000	0.111	0.000	0.123	0.003	0.045	<u>0.021</u>	<u>0.259</u>	0.000	0.157
	F1	0.560	0.560	0.000	0.121	0.000	0.148	0.005	0.057	<u>0.015</u>	<u>0.232</u>	0.000	0.156
QA	Acc	0.324	0.324	0.016	0.036	0.056	0.056	<u>0.286</u>	<u>0.294</u>	0.011	0.008	0.196	0.244

Table 5. Performance comparison across different models and datasets in Dutch. Best results are bold, second best are underlined.

4.2.2. Results on English Data

In order to test the adaptability of the Dutch financial language model, FinGEITje, we conducted a comparison of its performance against English financial evaluation data. The results are summarized in Table 6. Although English financial LLMs FinGPT and FinMA unsurprisingly perform best overall, FinGEITje, trained specifically on Dutch financial tasks, demonstrates notable competence. It emerges as the best model for NER (CLS), and manages to secure second highest performance for three tasks including SA, RE and QA.

Dataset	Metrics	FinGEITje-7B-sft		FinMA		FinGPT	
		raw	extracted	raw	extracted	raw	extracted
SA	Acc	<u>0.684</u>	<u>0.692</u>	0.728	0.728	0.402	0.402
	F1	<u>0.683</u>	<u>0.693</u>	0.729	0.729	0.247	0.248
NER (CLS)	Acc	0.806	0.880	0.402	0.484	<u>0.432</u>	<u>0.492</u>
	F1	0.847	0.883	0.309	0.442	<u>0.365</u>	<u>0.464</u>
HC	Acc	0.366	0.638	0.884	0.884	<u>0.798</u>	<u>0.798</u>
	F1	0.189	<u>0.290</u>	0.797	0.797	<u>0.205</u>	0.205
RE	Acc	<u>0.105</u>	<u>0.309</u>	0.004	0.088	0.589	0.614
	F1	<u>0.115</u>	<u>0.333</u>	0.007	0.085	0.613	0.630
QA	Acc	<u>0.408</u>	<u>0.420</u>	0.466	0.495	0.004	0.004

Table 6. Performance comparison across different models and datasets in English. Best results are bold, second best are underlined.

4.3. Discussion

FinGEITje is the first Dutch financial language model, optimized with a specialized instruction tuning dataset. Our benchmarks show that FinGEITje outperforms general Dutch instruction-tuned models, general-purpose models, and even English finance-specific models in all Dutch financial tasks. Surprisingly, FinGEITje also performs well on English financial tasks.

Our answer extraction process effectively improves performance by converting raw outputs to the correct format. For instance, when FinGEITje received English instructions, its Dutch outputs were correctly translated back into English, demonstrating the process’s reliability.

FinGEITje highlights the advantages of customizing language models for specific domains and languages. This approach is not only useful for Dutch but can also be applied to other languages. Our open-source instruction dataset translation method can help overcome the lack of datasets in non-English languages.

5. Conclusion

In this paper, we introduce FinGEITje, the first Dutch financial LLM fine-tuned on multiple financial tasks. We release a Dutch financial instruction tuning dataset comprising over more than 140,000 samples and an open-source data construction method that enables the creation of financial instruction datasets in various languages. In addition to

the dataset, we provide the first Dutch financial evaluation benchmark. In addition, we present an open-source automated evaluation method, which employs an LLM as an independent evaluator, reducing the need for human intervention in LLM performance evaluation.

Our extensive experimental evaluation demonstrates FinGEITje's superior performance on 5 critical Dutch and English financial tasks. Trained with limited computational resources and maximum transparency, our approach sets a new standard for Dutch financial NLP.

Our open-source contributions aim to promote further research and innovation in financial natural language understanding, facilitating the development of useful and safe LLMs in finance.

6. Future Work and Limitations

We reflect on possible future research directions and the current limitations of FinGEITje and similar financial LLMs, outlining potential enhancements.

6.1. Future Work

Our work on FinGEITje opens several promising avenues for future research. First, integrating advanced features such as vision, retrieval methods, and reinforcement learning approaches into the current model could potentially improve performance. Additionally, pre-training the Dutch foundational LLM on finance-specific texts may further improve domain-specific adaptation, and the application of domain-based filtering could also yield performance gains.

Moreover, expanding the scope of instruction tuning tasks is crucial for future improvements. This would allow the model to handle a wider range of financial tasks, thereby deepening its understanding of the financial domain. Challenging tasks, in particular, could benefit from increased instruction data volume.

Finally, the translation and instruction fine-tuning approach presented could be used as a framework for creating other domain-specific instruction tuning datasets and language-specific LLMs, thus improving the adaptability of LLMs across diverse realms.

6.2. Limitations

FinGEITje, alongside current state-of-the-art financial open-source LLMs, has several limitations. One significant constraint is the requirement for substantial computational resources, which poses a considerable barrier for researchers with limited computational capabilities. Additionally, since FinGEITje is specifically tailored for the finance domain, its performance may degrade when applied to tasks outside its specialized scope, potentially limiting its generalizability across diverse fields.

The model's capacity to handle complex reasoning tasks is also limited by the current architectural state of the 7B parameter models. Enhancing the reasoning abilities of the model remains a promising area for future research.

7. Ethical Considerations

It is important to address ethical considerations when conducting research involving large language models in the financial domain. This section discusses key ethical aspects related to energy consumption, data privacy and copyright, bias and fairness, and the responsible use of FinGEITje.

Energy Consumption. Training large language models requires significant computational resources, leading to considerable energy consumption and associated environmental impacts. We recognize the importance of sustainability in AI research and encourage ongoing efforts within the community to develop more energy-efficient algorithms and training techniques. Promoting sustainability is crucial for balancing technological advancement with ecological responsibility.

Data Privacy and Copyright. Our datasets are constructed from publicly available financial texts, including news articles, reports, and social media posts. We have ensured compliance with copyright laws, and no proprietary or confidential data has been used in training FinGEITje. By utilizing publicly available data, we aim to respect the intellectual property rights of content creators while providing a valuable resource for financial natural language processing.

Bias and Fairness. Language models can inadvertently learn and perpetuate biases present in their training data. FinGEITje is trained on a diverse range of financial documents to capture a broad spectrum of the financial domain. We acknowledge the possibility of residual biases and encourage users to critically assess the model's outputs, especially in contexts where fairness and impartiality are paramount. Continuous evaluation and feedback from the community are essential for addressing and mitigating potential biases.

Responsible Use. FinGEITje is intended to support education and research. It is important that users employ the model ethically and responsibly, adhering to all relevant laws and regulations. We emphasize the significance of considering the societal impacts of deploying language models in financial contexts. The potential misuse of financial LLMs raises concerns about the dissemination of financial misinformation or unethical market influence. Users should avoid applications that could lead to unethical practices, such as market manipulation or the spread of misinformation. This underlines the need for rigorous mechanisms that ensure the responsible and ethical application of financial LLMs in real-world scenarios.

Acknowledgements

This research received funding from the Flemish government, through Flanders Innovation & Entrepreneurship (VLAIO, project HBC.2020.2883) and from the Flemish government under the program “Onderzoekprogramma Artificiële Intelligentie (AI) Vlaanderen”.

References

A. Instruction Dataset Translation Details

All English instruction datasets were translated using *gpt-3.5-turbo-0125*, with *max_tokens* set to 1024 and *temperature* to 0.0. For all tasks, the following system message and prompt were used:

You are a helpful assistant that translates English to **<target language>** to the requirements that are given to you.

System message:

Prompt:

You are asked to translate a task's instruction, optional input to the task, and the output of the task, from English into **<target language>**. Here are the requirements that you should adhere to: 1. do not translate the identifiers 'instruction: ', 'input: ', and 'response: ' but instead copy them to your output; 2. make sure that text is fluent to read and does not contain grammatical errors. Use standard **<target language>** without regional bias; 3. translate the instruction and input text using informal, but standard, language; 4. make sure to avoid biases (such as gender bias, grammatical bias, social bias); 5. if the instruction is to correct grammar mistakes or spelling mistakes then you have to generate a similar mistake in the input in **<target language>**, and then also generate a corrected output version in the output in **<target language>**; 6. if the instruction is to translate text from one language to another, then you do not translate the text that needs to be translated in the instruction or the input, nor the translation in the output (just copy them as-is); 7. do not translate code fragments but copy them to your output. If there are English examples, variable names or definitions in code fragments, keep them in English; 8. maintain the format: the task consists of a task instruction (marked 'instruction: '), optional input to the task (marked 'input: ') and output for the task marked with 'response: '. Now translate the following task with the requirements set out above. Do not provide an explanation and do not add anything else.

B. Answer Extraction Prompt

The answers for all LLMs were extracted using *gpt-3.5-turbo-0125* with *max_tokens* set to 1024 and *temperature* to 0.0. Here we provide the general system and user prompt; visit our GitHub repository for further information.

System message:

You are a helpful assistant specialized in extracting the label of a message: The possible labels are **<scale_string>**. If none of the labels apply reply with 'unknown'.

<p>Prompt:</p> <p>### Instruction:</p> <p>Determine the label of the message.</p> <p>Options: <scale_string>.</p> <p>No other options may be given.</p> <p>### Input:</p> <p><prediction></p> <p>### Response:</p>

The *scale_string* signifies the set of possible outputs for a given financial task in a certain language, and the *prediction* represents the raw prediction of an LLM for a specified financial instruction task.

Footnotes

¹ <https://taalunie.org/informatie/24/feiten-cijfers>

² <https://huggingface.co/Rijgersberg/GEITje-7B>

³ <https://huggingface.co/Rijgersberg/BramVanroy/GEITje-7B-ultra-sft>

⁴ <https://huggingface.co/BramVanroy/GEITje-7B-ultra>

⁵ <https://huggingface.co/spaces/BramVanroy/fietje-2b>

⁶ <https://huggingface.co/ReBatch/Reynaerde-7B-Chat>

⁷ <https://github.com/snoels/fingeit>

⁸ <https://github.com/AI4Finance-Foundation/FinGPT>

⁹ <https://platform.openai.com/docs/models/gpt-3-5-turbo>

¹⁰ <https://github.com/snoels/fingeit>

¹¹ <https://huggingface.co/facebook/fasttext-language-identification>

¹² <https://github.com/snoels/fingeit>

¹³ <https://huggingface.co/Rijgersberg/GEITje-7B>

¹⁴ <https://huggingface.co/snoels/FinGEITje-7B-sft>

¹⁵ <https://github.com/snoels/fingeit>

¹⁶ <https://huggingface.co/snoels/FinGEITje-7B-sft>

¹⁷ <https://huggingface.co/BramVanroy/GEITje-7B-ultra-sft>

¹⁸ <https://huggingface.co/BramVanroy/GEITje-7B-ultra>

¹⁹ <https://platform.openai.com/docs/models/gpt-3-5-turbo>

²⁰ <https://github.com/snoels/fingeit>

References

1. ^a ^b ^c ^d ^e ^f Dogu Araci. 2019. *Finbert: Financial sentiment analysis with pre-trained language models*. arXiv preprint arXiv:1908.10063 (2019).
2. ^a ^b ^c ^d ^e ^f Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann. 2023. *Bloomberggpt: A large language model for finance*. arXiv preprint arXiv:2303.17564 (2023).
3. ^a ^b ^c ^d ^e ^f Yuqi Nie, Yaxuan Kong, Xiaowen Dong, John M Mulvey, H Vincent Poor, Qingsong Wen, and Stefan Zohren. 2024. *A Survey of Large Language Models for Financial Applications: Progress, Prospects and Challenges*. arXiv preprint arXiv:2406.11903 (2024).
4. ^a ^b ^c ^d ^e ^f ^g ^h Neng Wang, Hongyang Yang, and Christina Dan Wang. 2023. *Fingpt: Instruction tuning benchmark for open-source large language models in financial datasets*. arXiv preprint arXiv:2310.04793 (2023).
5. ^a ^b ^c ^d ^e ^f Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and Jimin Huang. 2023. *Pixiu: A large language model, instruction data and evaluation benchmark for finance*. arXiv preprint arXiv:2306.05443 (2023).
6. ^a Sander Noels, Sébastien Viaene, and Tijl De Bie. 2024. *TopoLedgerBERT: Topological Learning of Ledger Description Embeddings using Siamese BERT-Networks*. arXiv:2407.05175 [cs.CE]. Available at: <https://arxiv.org/abs/2407.05175>
7. ^a Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2021. *Cross-task generalization via natural language crowdsourcing instructions*. arXiv preprint arXiv:2104.08773 (2021).
8. ^a Moreno La Quatra and Luca Cagliero. 2020. *End-to-end training for financial report summarization*. In *Proceedings of the 1st Joint Workshop on Financial Narrative Processing and MultiLing Financial Summarisation*, pages 118–123.
9. ^a Xiang Deng, Vasilisa Bashlovkina, Feng Han, Simon Baumgartner, and Michael Bendersky. 2023. *What do llms know about financial markets? a case study on reddit market sentiment analysis*. In *Companion Proceedings of the ACM Web Conference 2023*, pages 107–110.
10. ^a Sander Noels, Adriaan Merlevede, Andrew Fecheyr, Maarten Vanhalst, Nick Meerlaen, Sébastien Viaene, and Tijl De Bie. 2023. *Automated Financial Analysis Using GPT-4*. In *Machine Learning and Knowledge Discovery in Databases: Applied Data Science and Demo Track*, Gianmarco De Francisci Morales, Claudia Perlich, Natali Ruchansky, Nicolas Kourtellis, Elena Baralis, and Francesco Bonchi (Eds.). Springer Nature Switzerland, Cham, 345–349. ISBN 978-3-031-43430-3.
11. ^a 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 (2020), 1877–1901.

12. ^aVictor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stieglar, Teven Le Scao, Arun Raja, et al. 2021. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207* (2021).
13. ^{a, b, c, d}Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. 2023. Fingpt: Open-source financial large language models. *arXiv preprint arXiv:2306.06031* (2023).
14. ^{a, b}Gagan Bhatia, El Moatez Billah Nagoudi, Hasan Cavusoglu, and Muhammad Abdul-Mageed. 2024. FinTral: A Family of GPT-4 Level Multimodal Financial Large Language Models. *arXiv preprint arXiv:2402.10986* (2024).
15. ^{a, b}Zhuosheng Zhang, Hanqing Zhang, Keming Chen, Yuhang Guo, Jingyun Hua, Yulong Wang, and Ming Zhou. 2021. Mengzi: Towards lightweight yet ingenious pre-trained models for chinese. *arXiv preprint arXiv:2110.06696* (2021).
16. ^{a, b}Dakuan Lu, Hengkui Wu, Jiaqing Liang, Yipei Xu, Qianyu He, Yipeng Geng, Mengkun Han, Yingsi Xin, and Yanghua Xiao. 2023. Bbt-fin: Comprehensive construction of chinese financial domain pre-trained language model, corpus and benchmark. *arXiv preprint arXiv:2302.09432* (2023).
17. ^{a, b, c}Xiao Zhang, Ruoyu Xiang, Chenhan Yuan, Duanyu Feng, Weiguang Han, Alejandro Lopez-Lira, Xiao-Yang Liu, Sophia Ananiadou, Min Peng, Jimin Huang, et al. 2024. D'olares or Dollars? Unraveling the Bilingual Prowess of Financial LLMs Between Spanish and English. *arXiv preprint arXiv:2402.07405* (2024).
18. ^{a, b}Gang Hu, Ke Qin, Chenhan Yuan, Min Peng, Alejandro Lopez-Lira, Benyou Wang, Sophia Ananiadou, Wanlong Yu, Jimin Huang, and Qianqian Xie. 2024. No Language is an Island: Unifying Chinese and English in Financial Large Language Models, Instruction Data, and Benchmarks. *arXiv preprint arXiv:2403.06249* (2024).
19. ^aMasanori Hirano. 2024. Construction of a japanese financial benchmark for large language models. *arXiv preprint arXiv:2403.15062* (2024).
20. ^{a, b, c, d, e, f, g}Bram Vanroy. 2023. Language resources for Dutch large language modelling. *arXiv preprint arXiv:2312.12852* (2023).
21. ^{a, b}Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560* (2022).
22. ^{a, b}Pekka Malo, Ankur Sinha, Pekka Korhonen, Jyrki Wallenius, and Pyry Takala. 2014. Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the Association for Information Science and Technology* 65, 4 (2014), 782--796.
23. ^{a, b, c}Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk, and Alexandra Balahur. 2018. Www'18 open challenge: financial opinion mining and question answering. In *Companion proceedings of the the web conference 2018*, pages 1941--1942.
24. ^{a, b}Neural Magic. 2022. Twitter financial news sentiment.
25. ^{a, b, c}Julio Cesar Salinas Alvarado, Karin Verspoor, and Timothy Baldwin. 2015. Domain adaption of named entity recognition to support credit risk assessment. In *Proceedings of the Australasian Language Technology Association Workshop 2015*, pages 84--90.

26. ^{a, b}Ankur Sinha and Tanmay Khandait. 2021. *Impact of news on the commodity market: Dataset and results*. In *Advances in Information and Communication: Proceedings of the 2021 Future of Information and Communication Conference (FI CC)*, Volume 2. Springer, pages 589--601.
27. ^{a, b}Soumya Sharma, Tapas Nayak, Arusarka Bose, Ajay Kumar Meena, Koustuv Dasgupta, Niloy Ganguly, and Pawan Goyal. 2022. *FinRED: A dataset for relation extraction in financial domain*. In *Companion Proceedings of the Web Conference 2022*, pages 595--597.
28. ^{a, b}Zhiyu Chen, Shiyang Li, Charese Smiley, Zhiqiang Ma, Sameena Shah, and William Yang Wang. 2022. *Convfinqa: Exploring the chain of numerical reasoning in conversational finance question answering*. *arXiv preprint arXiv:2210.03849* (2022).
29. ^ΔLaria Reynolds and Kyle McDonell. 2021. *Prompt programming for large language models: Beyond the few-shot paradigm*. In *Extended abstracts of the 2021 CHI conference on human factors in computing systems*, 1--7.
30. ^ΔAlbert 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* (2023).
31. ^ΔWei J, Bosma M, Zhao VY, Guu K, Yu AW, Lester B, Du N, Dai AM, Le QV (2021). "Finetuned language models are zero-shot learners". *arXiv preprint arXiv:2109.01652* (2021).
32. ^ΔDao T, Fu D, Ermon S, Rudra A, Ré C (2022). "Flashattention: Fast and memory-efficient exact attention with io-awareness". *Advances in Neural Information Processing Systems* 35 (2022): 16344--16359.
33. ^ΔDettmers T, Pagnoni A, Holtzman A, Zettlemoyer L (2024). "Qlora: Efficient finetuning of quantized llms". *Advances in Neural Information Processing Systems* 36 (2024).
34. ^ΔRafailov R, Sharma A, Mitchell E, Manning CD, Ermon S, Finn C (2024). "Direct preference optimization: Your language model is secretly a reward model". *Advances in Neural Information Processing Systems* 36 (2024).
35. ^ΔLai VD, Nguyen CV, Ngo NT, Nguyen T, Deroncourt F, Rossi RA, Nguyen TH (2023). "Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback". *arXiv preprint arXiv:2307.16039* (2023).
36. ^ΔTouvron H, Martin L, Stone K, Albert P, Almahairi A, Babaei Y, Bashlykov N, Batra S, Bhargava P, Bhosale S, Bikel D, Blecher L, Canton Ferrer C, Chen M, Cucurull G, Esiobu D, Fernandes J, Fu J, Fu W, Fuller B, Gao C, Goswami V, Goyal N, Hartshorn A, Hosseini S, Hou R, Inan H, Kardas M, Kerkez V, Khabsa M, Kloumann I, Korenev A, Koura PS, Lachaux MA, Lavril T, Lee J, Liskovich D, Lu Y, Mao Y, Martinet X, Mihaylov T, Mishra P, Molybog I, Nie Y, Poulton A, Reizenstein J, Rungta R, Saladi K, Schelten A, Silva R, Smith EM, Subramanian R, Tan XE, Tang B, Taylor R, Williams A, Kuan JX, Xu P, Yan Z, Zarov I, Zhang Y, Fan A, Kambadur M, Narang S, Rodriguez A, Stojnic R, Edunov S, Scialom T (2023). "Llama 2: Open Foundation and Fine-Tuned Chat Models". *arXiv:2307.09288 [cs.CL]*. <https://arxiv.org/abs/2307.09288>.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.