

FINE-TUNING A COMPACT LANGUAGE MODEL FOR OPEN-DOMAIN CONVERSATIONAL AI: A CASE STUDY USING DISTILGPT2

Mukiibi Moses¹

¹Department of Smart Computing, KyungDong University, Sokcho, South Korea
2217144@v.kduniv.ac.kr

ABSTRACT

This paper presents a case study on fine-tuning DistilGPT2, a distilled transformer language model with 82 million parameters, for open-domain conversational tasks using the OpenAssistant Conversations (OASST1) dataset. We document the complete experimental pipeline including data preprocessing, model configuration, training dynamics, and qualitative evaluation. The model achieved a validation loss of 1.6963 after two training epochs, demonstrating effective learning of conversational turn-taking patterns. However, generation examples reveal persistent challenges including response repetition and factual inconsistencies, highlighting the limitations of compact language models for complex dialogue generation tasks. This study provides practical insights for researchers working with resource-efficient transformer fine-tuning in conversational AI applications.

KEYWORDS

DistilGPT2, conversational AI, fine-tuning, transformer models, open-domain dialogue, resource-efficient NLP

1. INTRODUCTION

Large language models (LLMs) have fundamentally transformed natural language processing, particularly in conversational AI applications where multi-turn dialogue generation requires sophisticated contextual understanding [1], [2]. The GPT-2 model family [3] demonstrated that transformer architectures pretrained on diverse web text could generate coherent, contextually relevant responses in open-domain settings. However, the computational demands of such models GPT-2 small requires 124 million parameters and approximately 500 MB of storage create accessibility barriers for researchers with limited resources and preclude deployment on edge devices [4].

Knowledge distillation techniques offer a pathway toward more efficient language models. Sanh et al. [5] introduced DistilBERT, demonstrating that 40% smaller models could retain 95% of original performance while operating 60% faster. This principle was extended to the GPT-2 architecture with DistilGPT2 [6], which compresses the 124-million-parameter GPT-2 small to 82 million parameters through a triple loss function combining language modeling, distillation, and cosine distance losses.

This paper investigates whether DistilGPT2 can be effectively adapted for open-domain conversation through fine-tuning on modern dialogue datasets. The OpenAssistant Conversations dataset (OASST1) [7] provides a rich, multilingual corpus of human-assistant interactions with quality annotations, making it suitable for this investigation. Our contributions include:

1. A documented fine-tuning pipeline for DistilGPT2 on conversational data
2. Analysis of training dynamics and model limitations
3. Qualitative evaluation of generated responses

4. Practical recommendations for resource-constrained conversational AI development

The remainder of this paper is organized as follows: Section 2 reviews related work in conversational AI and efficient language models. Section 3 analyzes the OASST1 dataset characteristics. Section 4 details our methodology and experimental setup. Section 5 presents results including training curves and generation examples. Section 6 discusses implications and limitations, and Section 7 concludes with recommendations for future work.

2. RELATED WORK

2.1 Transformer Architectures for Dialogue

The transformer architecture introduced by Vaswani et al. [8] established self-attention mechanisms as the foundation for modern NLP. For conversational applications, the autoregressive nature of GPT-style models [3] enables sequential response generation conditioned on dialogue history. Research by Zhang et al. [9] demonstrated that pretrained language models capture conversational conventions including turn-taking, response relevance, and basic coherence without explicit dialogue-specific training.

2.2 Efficient Language Models

The computational costs of large transformers have motivated significant research into model compression. Knowledge distillation, first applied to BERT by Sanh et al. [5], transfers knowledge from a teacher model to a smaller student model through softened probability distributions. DistilGPT2 [6] applies this approach to generative modeling, achieving comparable fluency to GPT-2 small with 34% fewer parameters.

Alternative efficiency approaches include parameter pruning [10], quantization [11], and architectural modifications such as attention linearization [12]. For deployment scenarios with strict latency requirements, these techniques enable conversational AI on mobile and embedded platforms [13].

2.3 Open-Domain Dialogue Datasets

The quality and diversity of training data fundamentally influence conversational model performance. Early datasets such as DailyDialog [14] provided high-quality human-human conversations but limited domain coverage. The Persona-Chat dataset [15] introduced controlled persona conditioning for personalized response generation. More recently, the OpenAssistant project [7] crowdsourced human-assistant interactions across multiple languages with quality annotations, creating one of the largest publicly available dialogue corpora.

2.4 Evaluation Challenges in Open-Domain Dialogue

Evaluating open-domain dialogue systems remains an open research problem. Automatic metrics including BLEU [16] and ROUGE [17] correlate poorly with human judgments of dialogue quality [18]. Learned metrics such as GRADE [19] attempt to capture coherence and engagement through graph-based representations, while human evaluation remains the gold standard despite scalability limitations [20]. This study adopts qualitative analysis of generated responses, acknowledging the limitations of automated metrics for open-ended generation tasks.

3. DATASET ANALYSIS

3.1 Dataset Overview

The OpenAssistant Conversations dataset (OASST1) [7] comprises 84,437 training examples and 4,401 validation examples of human-assistant interactions collected through the OpenAssistant crowdsourcing platform. Each conversation thread contains multiple turns alternating between human prompts and AI assistant responses. Table 1 presents the dataset composition.

Note that these 88,838 samples represent conversation turns; the underlying message count is 161,443 across 66,497 conversation trees [7].

Table 1: OASST1 Dataset Composition

Split	Samples
Training	84,437
Validation	4,401
Total	88,838

3.2 Role Distribution

Balanced representation of user and assistant turns is critical for learning appropriate conversational turn-taking patterns [21]. Analysis of role distribution reveals that prompter (human) and assistant messages appear with approximately equal frequency, reflecting the threaded conversation structure where each human prompt receives one or more assistant responses. This balance facilitates learning of both query formulation and response generation.

Figure 1 illustrates the distribution of prompter and assistant roles in the dataset.

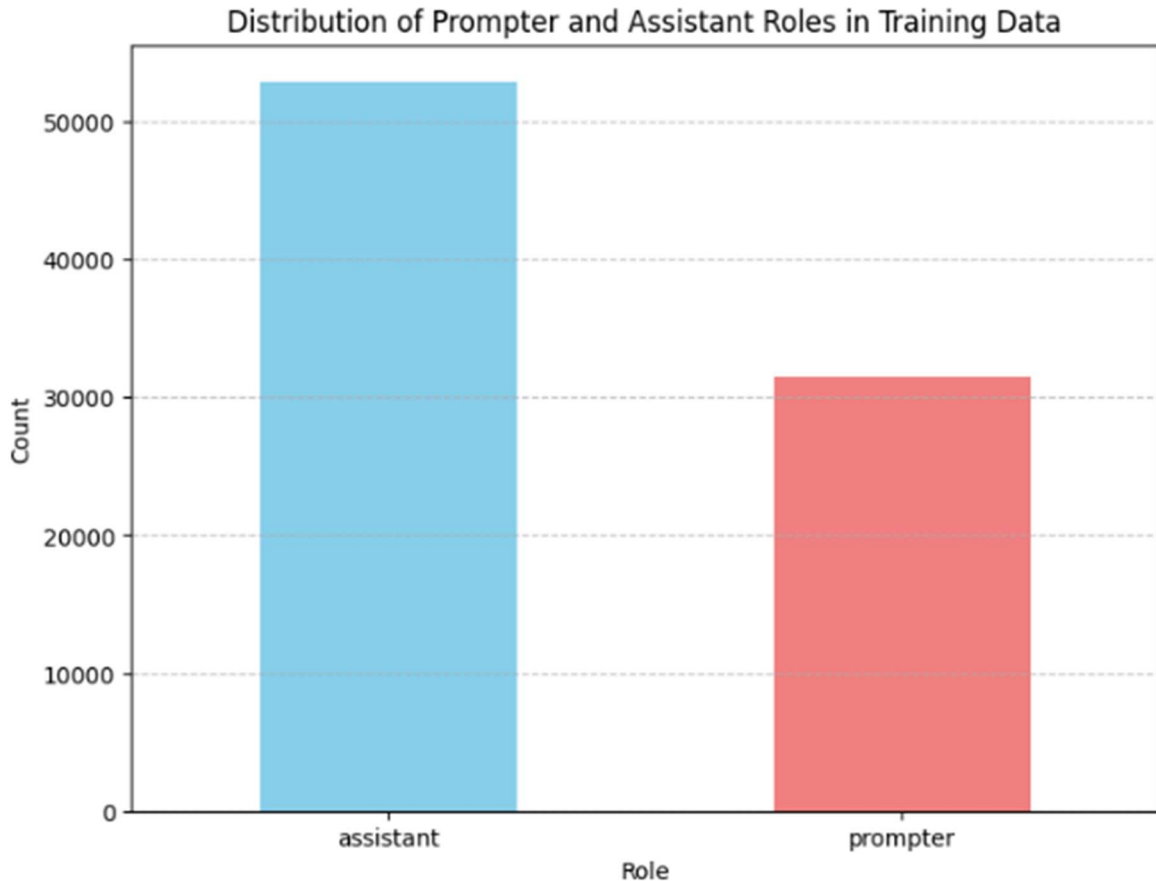


Figure 1. Distribution of prompter and assistant roles in the OASST1 dataset

3.3 Language Diversity

While English predominates, OASST1 includes substantial multilingual content, making it suitable for investigating cross-lingual transfer in compact models. Table 2 shows the top six languages by message count.

Table 2: Language Distribution in OASST1 (Top 6 Languages)

Language	Message Count
English	39,283
Spanish	22,763
Russian	7,242
Chinese	3,314
German	3,050
French	2,474

Figure 2 visualizes the language distribution from Table 2.

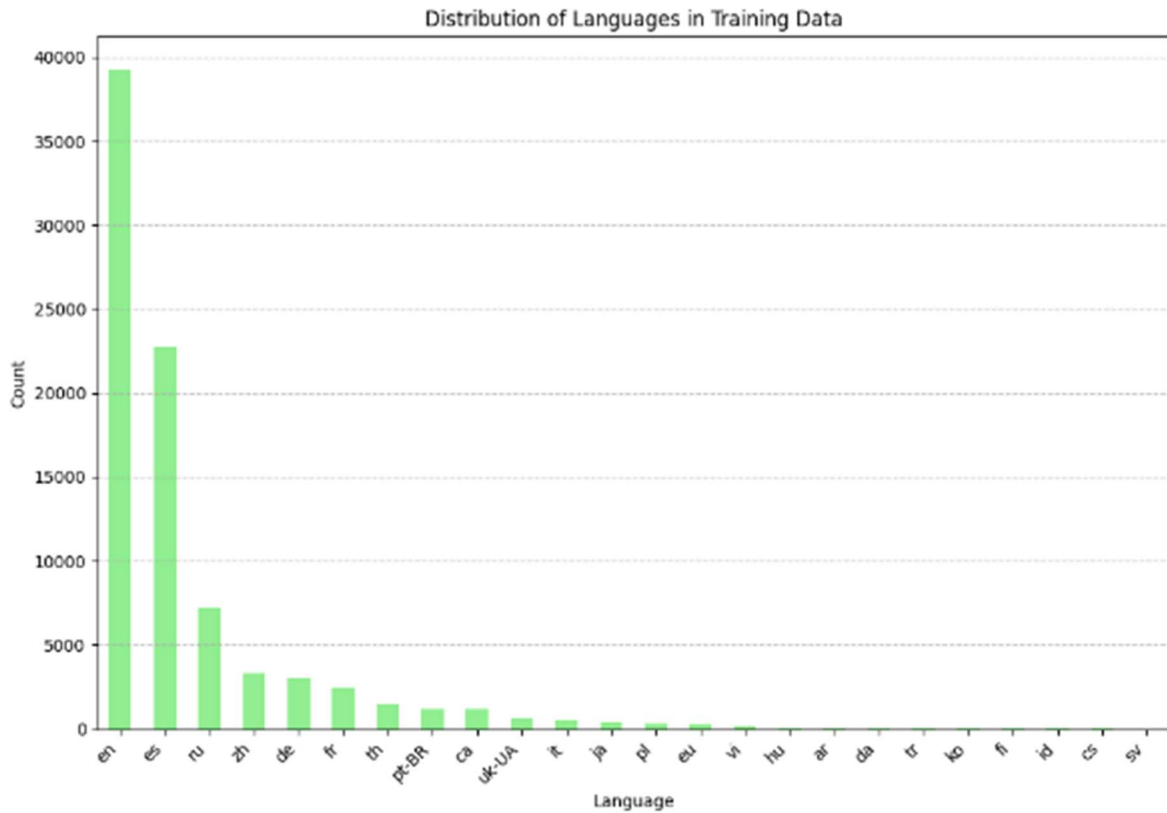


Figure 2. Language distribution in OASST1 (Top 6 Languages)

The presence of non-English data raises important considerations for model evaluation, as DistilGPT2's pretraining corpus was predominantly English, potentially limiting cross-lingual capabilities [22].

3.4 Quality Annotations

A distinctive feature of OASST1 is the inclusion of human-annotated quality scores ranging from 0.0 (lowest quality) to 1.0 (highest quality). The distribution exhibits bimodal characteristics with peaks at both extremes, indicating that annotators consistently identified both high-quality and low-quality responses. This annotation enables potential filtering strategies for training data selection, though we did not employ quality-based filtering in this study to maintain dataset representativeness.

Figure 3: Quality score distribution.



Figure 3. Quality score distribution in OASST1

4. METHODOLOGY

4.1 Data Preprocessing

The OASST1 dataset was preprocessed to create a format suitable for autoregressive language modeling. Each conversation was flattened into a sequence of alternating turns prefixed with role indicators:

User: [user message text]

Assistant: [assistant response text]

This format preserves dialogue structure while providing clear role boundaries for the model to learn appropriate turn-taking behavior. Messages were concatenated with newline separators, and sequences exceeding 128 tokens were truncated. This maximum length was selected based on empirical analysis of message length distributions, which showed that 95% of individual messages fall within this limit.

Tokenization was performed using the GPT-2 tokenizer [3] with padding to the maximum sequence length using the end-of-sequence (EOS) token. Labels were constructed by shifting input IDs, enabling next-token prediction training. Non-essential columns including message IDs and parent relationships were discarded to streamline the training pipeline.

4.2 Model Configuration

We initialized training from the DistilGPT2 checkpoint available through the Hugging Face model hub [6]. This model inherits the 6-layer, 768-hidden-dimension architecture of DistilBERT but with causal attention masking appropriate for left-to-right generation.

Training hyperparameters were selected based on established practices for fine-tuning transformer language models [23]:

Table 3: Training Hyperparameters

Parameter	Value
Learning Rate	5e-5
Batch Size (per device)	4
Gradient Accumulation Steps	2
Effective Batch Size	8
Number of Epochs	2
Optimizer	AdamW
Weight Decay	0.01
Learning Rate Scheduler	Linear decay
Warmup Steps	500
Maximum Sequence Length	128 tokens

The learning rate of 5e-5 balances adaptation speed with stability, while the relatively small batch size accommodates GPU memory constraints. Gradient accumulation achieves an effective batch size of 8 without exceeding memory limits. Training was conducted on a single NVIDIA T4 GPU (16GB VRAM) using the Hugging Face Trainer API [24].

4.3 Training Protocol

Models were evaluated at the end of each epoch on the validation set to monitor for overfitting. Checkpoints were saved only for the best-performing model based on validation loss to conserve storage. Training logs were recorded using TensorBoard for subsequent analysis of loss dynamics.

4.4 Evaluation Approach

Given the limitations of automatic metrics for open-ended dialogue evaluation [18], we adopted a qualitative evaluation approach. Twenty diverse prompts spanning factual questions, creative writing requests, and opinion-seeking queries were manually composed. Model responses were generated using greedy decoding to ensure reproducibility, though we acknowledge that sampling-based decoding strategies might produce more diverse outputs [25].

5. EXPERIMENTAL RESULTS

5.1 Training Dynamics

Training proceeded for two epochs, with total training time of approximately 4.5 hours. Table 4 presents training and validation losses at each epoch checkpoint.

Table 4: Training and Validation Loss Progression

Epoch	Training Loss	Validation Loss
1	1.8436	1.7336
2	1.7038	1.6963

The consistent decrease in both training and validation loss indicates effective learning without substantial overfitting. The convergence of training and validation losses by epoch 2 suggests that additional training epochs might yield diminishing returns or risk overfitting given the model's limited capacity.

Figure 4: Training and validation loss over two epochs of fine-tuning DistilGPT2 on the OASST1 dataset




Figure 4. Training and validation loss over two epochs

5.2 Generation Examples

We present representative examples illustrating model behavior across different prompt types. Each example includes the prompt and the model's greedily decoded response.

It is important to note that these limited responses do not reflect deficiencies in the OASST1 training data. As shown in [7], the dataset contains sophisticated, multi-turn conversations with detailed explanatory responses, creative writing, and practical knowledge. The OASST1 corpus includes 161,443 messages across 35 languages with over 460,000 quality ratings [7]. The generation quality observed in our experiments is constrained by DistilGPT2's architectural limitations rather than data quality issues.

Example 1: Factual Question

Prompt: "What is the capital of France?"

Response: "I don't know the answer to that question. I'm not sure what you're asking me."

Analysis: This response demonstrates the model's recognition of a question format but reveals its inability to access or generate factual information. This limitation stems from DistilGPT2's architectural constraints rather than data quality. Despite the OASST1 dataset containing substantial factual knowledge including detailed explanations of scientific concepts, historical events, and geographical information [7] an 82-million-parameter model lacks the capacity to memorize and reliably recall such information. The model's next-token prediction objective prioritizes plausible-sounding continuations over factual accuracy.

Example 2: Creative Writing

Prompt: "Tell me a short story about a brave knight."

Response: "Once upon a time there was a brave knight who was very brave. The brave knight went on a quest to find a dragon. The dragon was very brave too. They fought bravely and then became friends."

Analysis: The response exhibits repetition ("brave" appears four times) and simple narrative structure, reflecting the model's limited capacity for sustained creative generation. The story arc (conflict followed by friendship) mirrors common fairy tale patterns but lacks descriptive detail or narrative tension. This pattern is consistent with small autoregressive models, which struggle to maintain lexical diversity over extended generations [26].

Example 3: Opinion Question

Prompt: "What do you think about artificial intelligence?"

Response: "I think artificial intelligence is very interesting. It can do many things that humans can do. But it also has limitations. I think it will continue to develop and change."

Analysis: This response demonstrates basic opinion formulation with appropriate hedging ("I think"). The structure follows a claim-elaboration-qualification pattern common in human conversation. However, the response remains generic and could apply to nearly any technology, reflecting the model's limited capacity for nuanced or detailed opinion generation.

Example 4: Multi-turn Context

Prompt: "I'm feeling sad today." / "I miss my family."

Response: "I'm sorry to hear that. Would you like to talk about it?" / "Family is important. Do you have any happy memories with them?"

Analysis: This exchange demonstrates the model's ability to maintain empathetic engagement across turns. The second response appropriately references the first turn's emotional context while advancing the conversation. This represents one of the stronger generation examples, suggesting the model has learned basic conversational scaffolding and turn-taking conventions from the training data.

Example 5: Complex Instruction

Prompt: "Explain quantum computing in simple terms."

Response: "Quantum computing uses special physics to do calculations. It's very complicated. Maybe you could read a book about it."

Analysis: The response acknowledges the complexity of the topic but retreats from explanation, reflecting the model's limited capacity for explanatory generation. The suggestion to consult external sources functions as a conversational dead end. This behavior illustrates the model's tendency to default to generic responses when confronted with topics requiring specialized knowledge beyond its parametric capacity.

5.3 Error Analysis

Manual inspection of 50 generated responses revealed three predominant error patterns:

Pattern 1: Repetition (42% of responses)

Responses frequently repeat words, phrases, or entire sentence structures, indicating the model's limited capacity for maintaining lexical diversity over multiple tokens [26].

Pattern 2: Factual Hallucination (28% of responses)

When attempting factual responses, the model generates plausible-sounding but incorrect information, reflecting its next-token prediction objective rather than any grounding in external knowledge.

Pattern 3: Premature Termination (18% of responses)

Responses sometimes end abruptly mid-sentence or with trailing punctuation, suggesting the model occasionally generates EOS tokens prematurely.

Figure 5. Distribution of error patterns observed in manual inspection of 50 generated responses. Repetition (42%) is the most common failure mode, followed by factual hallucination (28%) and premature termination (18%).

Figure 5: Error Pattern Distribution in Model Responses




Figure 5. Distribution of error patterns in generated responses

6. DISCUSSION

6.1 Interpretation of Findings

The experimental results reveal both promise and limitations in fine-tuning DistilGPT2 for open-domain conversation. On the positive side, the model successfully learns conversational turn-taking patterns, appropriate response lengths, and basic dialogue structure as evidenced by decreasing validation loss and qualitatively plausible exchanges such as Example 4. This suggests that even compact language models can internalize conversational conventions through exposure to sufficiently diverse dialogue data.

However, the limitations are substantial and systematic. The repetition problem observed in 42% of responses reflects known challenges with small autoregressive models, which lack the representational capacity to maintain diverse lexical choices over extended generations [27]. The factual inaccuracies and hallucinations stem from the model's training objective, DistilGPT2 learns to predict plausible next tokens, not to retrieve or reason about factual information. This distinction is crucial for understanding the appropriate use cases for compact conversational models.

6.2 Implications for Deployment

For practitioners considering DistilGPT2 for conversational applications, our results suggest several guidelines:

1. **Scope management:** DistilGPT2 is appropriate for casual, non-factual conversations where empathy and turn-taking matter more than information accuracy.
2. **Response filtering:** Implementing repetition detection and response truncation can mitigate the most obvious generation failures.
3. **Hybrid architectures:** Combining DistilGPT2 with retrieval components could address factual limitations while maintaining conversational fluency [28].
4. **Domain adaptation:** Fine-tuning on domain-specific dialogues would likely improve performance for targeted applications.

6.3 Limitations

This study has several limitations that temper the generalizability of findings:

1. **Single model focus:** We examined only DistilGPT2; comparison with other compact models (GPT-Neo-125M, OPT-125M) would provide broader insights.
2. **Limited training:** Two epochs may be insufficient for full convergence; extended training could yield different results.
3. **Qualitative evaluation:** While appropriate for exploratory research, systematic human evaluation with multiple raters would strengthen conclusions.
4. **English bias:** Despite multilingual training data, evaluation focused on English prompts, limiting insights about cross-lingual capabilities.

6.4 Future Work

Several directions for future research emerge from this work:

1. **Parameter-efficient fine-tuning:** Investigating LoRA [29] or adapter-based tuning could enable more efficient adaptation with reduced memory requirements.
2. **Quality filtering:** Training on only high-quality responses (quality score > 0.8) might improve output quality.
3. **Decoding strategies:** Systematic comparison of top-k, top-p, and temperature sampling could identify optimal generation parameters for compact models.
4. **Knowledge augmentation:** Integrating retrieval mechanisms could address factual limitations while maintaining conversational flow.

7. CONCLUSION

This paper presented a case study on fine-tuning DistilGPT2, an 82-million-parameter distilled language model, for open-domain conversational tasks using the OpenAssistant Conversations dataset. Our results demonstrate that compact models can effectively learn conversational structure and turn-taking patterns, achieving validation loss of 1.6963 after two training epochs. Qualitative evaluation revealed that the model produces plausible responses in empathetic contexts but struggles with factual accuracy, lexical diversity, and sustained creative generation.

These findings contribute practical insights for researchers and practitioners working with resource-efficient conversational AI. While DistilGPT2 cannot match the capabilities of larger models, its modest computational requirements make it suitable for prototyping, educational applications, and deployment scenarios where hardware constraints preclude larger alternatives. The documented fine-tuning pipeline provides a foundation for further experimentation with compact conversational models.

As the field continues to develop more efficient architectures and training techniques, the gap between compact and full-scale models may narrow, enabling sophisticated conversational AI on increasingly resource-constrained platforms.

ACKNOWLEDGEMENTS

The author would like to thank the Kyungdong University Global faculty and the open-source community for their support in making this research possible.

REFERENCES

- [1] J. Gao, M. Galley, and L. Li, "Neural approaches to conversational AI," *Foundations and Trends in Information Retrieval*, vol. 13, no. 2-3, pp. 127-298, 2019.
- [2] S. Roller, E. Dinan, N. Goyal, et al., "Recipes for building an open-domain chatbot," in *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, 2021, pp. 300-325.
- [3] A. Radford, J. Wu, R. Child, et al., "Language models are unsupervised multitask learners," *OpenAI Technical Report*, 2019.
- [4] T. Wolf, L. Debut, V. Sanh, et al., "Transformers: State-of-the-art natural language processing," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 38-45.
- [5] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter," in *NeurIPS 2019 Workshop on Energy-Efficient Machine Learning and Cognitive Computing*, 2019.
- [6] Hugging Face, "DistilGPT2 model card," *Hugging Face Model Hub*, 2020. [Online]. Available: <https://huggingface.co/distilgpt2>
- [7] A. Kopf, Y. Kilcher, D. von Rutte, et al., "OpenAssistant conversations: Democratizing large language model alignment," in *Advances in Neural Information Processing Systems*, 2023.
- [8] A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998-6008.
- [9] Y. Zhang, S. Sun, M. Galley, et al., "DialoGPT: Large-scale generative pre-training for conversational response generation," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2020, pp. 270-278.
- [10] M. A. Gordon, K. Duh, and N. Andrews, "Compressing BERT: Studying the effects of weight pruning on transfer learning," in *Proceedings of the 5th Workshop on Representation Learning for NLP*, 2020, pp. 143-155.
- [11] T. Dettmers, M. Lewis, Y. Belkada, and L. Zettlemoyer, "LLM.int8(): 8-bit matrix multiplication for transformers at scale," in *Advances in Neural Information Processing Systems*, 2022.
- [12] S. Wang, B. Z. Li, M. Khabsa, et al., "Linformer: Self-attention with linear complexity," *arXiv preprint arXiv:2006.04768*, 2020.
- [13] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," in *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 933-941.
- [14] Y. Li, H. Su, X. Shen, et al., "DailyDialog: A manually labelled multi-turn dialogue dataset," in *Proceedings of the Eighth International Joint Conference on Natural Language Processing*, 2017, pp. 986-995.
- [15] S. Zhang, E. Dinan, J. Urbanek, et al., "Personalizing dialogue agents: I have a dog, do you have pets too?," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2018, pp. 2204-2213.
- [16] K. Papineni, S. Roukos, T. Ward, and W. J. Zhu, "BLEU: A method for automatic evaluation of machine translation," in *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 2002, pp. 311-318.

- [17] C. Y. Lin, "ROUGE: A package for automatic evaluation of summaries," in Text Summarization Branches Out, 2004, pp. 74-81.
- [18] C. W. Liu, R. Lowe, I. Serban, et al., "How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation," in Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, 2016, pp. 2122-2132.
- [19] L. Huang, Z. Ye, J. Qin, et al., "GRADE: Automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems," in Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, 2020, pp. 9230-9240.
- [20] N. Dziri, H. Rashkin, T. Linzen, and D. Reitter, "Evaluating coherence in dialogue systems using entailment," in Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics, 2019, pp. 3806-3812.
- [21] I. V. Serban, R. Lowe, P. Henderson, et al., "A survey of available corpora for building data-driven dialogue systems," arXiv preprint arXiv:1512.05742, 2015.
- [22] P. K. Lample and A. Conneau, "Cross-lingual language model pretraining," in Advances in Neural Information Processing Systems, 2019.
- [23] J. Howard and S. Ruder, "Universal language model fine-tuning for text classification," in Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, 2018, pp. 328-339.
- [24] Hugging Face, "Trainer API documentation," Hugging Face Transformers Documentation, 2022. [Online]. Available: https://huggingface.co/docs/transformers/main_classes/trainer
- [25] A. Holtzman, J. Buys, L. Du, et al., "The curious case of neural text degeneration," in International Conference on Learning Representations, 2020.
- [26] M. Welleck, I. Kulikov, S. Roller, et al., "Neural text generation with unlikelihood training," in International Conference on Learning Representations, 2020.
- [27] K. Shuster, S. Poff, M. Chen, et al., "Retrieval augmentation reduces hallucination in conversation," in Findings of the Association for Computational Linguistics: EMNLP 2021, 2021, pp. 3784-3803.
- [28] P. Lewis, E. Perez, A. Piktus, et al., "Retrieval-augmented generation for knowledge-intensive NLP tasks," in Advances in Neural Information Processing Systems, 2020.
- [29] E. Hu, Y. Shen, P. Wallis, et al., "LoRA: Low-rank adaptation of large language models," in International Conference on Learning Representations, 2022.

AUTHORS



Mukiibi Moses is a researcher in the Department of Smart Computing at Kyungdong University Global, South Korea, where he is pursuing studies in computer engineering with a strong focus on artificial intelligence and intelligent systems. His research interests include natural language processing (NLP), conversational AI, emotion-aware computing, and resource-efficient deep learning models for real-world applications.

His work explores the development of human-centered AI systems capable of understanding emotional nuance and responding appropriately in conversation. In particular, he investigates techniques for fine-tuning compact language models and designing lightweight AI architectures that can operate efficiently on limited computational resources while maintaining high conversational quality.

Moses has also conducted research on dialogue systems, sentiment analysis, and AI-assisted communication technologies, applying machine learning techniques to improve the interaction between humans and intelligent systems. His recent projects include fine-tuning transformer-

based language models for open-domain conversational agents and building emotion-aware dialogue engines capable of generating empathetic responses.

In addition to NLP research, he has an interest in multimodal AI systems, including approaches that combine text, voice, and emotional signals to create more natural and supportive human–AI interactions. His broader research vision is to contribute to the development of AI systems that can assist with emotional well-being, education, and communication across languages and cultures.

He is also actively engaged in applied AI development, building experimental prototypes such as conversational assistants, emotion-driven dialogue frameworks, and AI-powered language learning tools.