

Understanding the importance of enhanced logical reasoning among large language models with the help of hybrid symbolic architectures

MSc Research Project
Masters in Artificial Intelligence (MSCAIB)

Surya Prakash Chinnameda
Student ID: x23343699

School of Computing
National College of Ireland

Supervisor: Lavish Thomas

National College of Ireland
MSc Project Submission Sheet



School of Computing

Student Name: Surya prakash Chinnameda
Student ID: X23343699
Programme: Masters in Artificial Intelligence (MSCAIB) **Year:** 2025
Module: Practicum
Lecturer: Lavish Thomas
Submission Due Date: 15-09-2025
Project Title: Understanding the importance of enhanced logical reasoning among large language models with the help of hybrid symbolic architectures
Word Count: 855 **Page Count:** 4

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project. ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: Surya Prakash Chinnameda
Date: 15-09-2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	<input type="checkbox"/>
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	<input type="checkbox"/>
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	<input type="checkbox"/>

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Section 1 – Model Development Setup

A hybrid symbolic-neural architecture was mounted within a Python GPU-enabled environment (Python 3.11) through PyTorch as the deep learning framework to be used primarily. Some major dependencies include Hugging Face Transformers Library that will support our DistilBERT-based classifier, scikit-learn for preprocessing of data, TF-IDF Vectorizer, and Traditional Machine Learning Classifiers implementation among other uses. Pandas and NumPy were also used for normal numerical operations with Matplotlib plus Seaborn undertaking the duty of showing visuals from datasets.

The data comes from Kaggle’s Logical Reasoning Improvement Dataset. It is split into 80% training, 10% validation, and 10% test sets with a stratified approach so that the proportions of categories are kept. Preprocessing includes lowercasing, removing punctuation, tokenization, and minor normalization— χ^2 becomes “chi2”. WeightedRandomSampler, were used to fight the imbalance of categories representing learning equally across all classes

Section 2 – Model Training and Symbolic Rule Integration

TF-IDF vectorization max applied to text data used to trained Logistic Regression (rules-only and TF-IDF), plus DistilBERT Classifier models. They have generally emerged as the top performers when dealing with high-dimensional text classification challenges.

The DistilBERT transformer was fine-tuned with max tokens set to 288, five epochs for training, and validation-based early stopping applied. AdamW optimiser also used with its learning rate warm-up then decay to further enhance convergence. Rule-based features were created over the templates of logic patterns that domain-symbolic knowledge typically expresses — probability, combinatorics, and indicators of logic puzzles — among others. These symbolic features are induced as separate interpretability baseline such that the model is enabled to statistical learning with explicit reasoning signals before making a classification decision.

Section 3 – Evaluation and Deployment

TF-IDF vectorization was applied to the text data. The models that were trained included rules-only Logistic Regression, TF-IDF + Logistic Regression, and DistilBERT. These models have proven to be the most effective where high dimensions are involved in text classification tasks.

The DistilBERT transformer was fine-tuned for deep learning at max token length 288, trained across five epochs with early validation-based stopping, under AdamW optimiser accompanied by a scheduler for the learning rate (warm-up and decay). It injects symbolic reasoning through the generation of rule-based features that express domain-specific logic patterns - probability indicators, combinatorics indicators, and logic puzzle indicators. These symbolic features are induced as separate interpretability baseline of comparison to the neural model rather than being directly fused with its embeddings.

References

- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT*, 4171–4186. <https://doi.org/10.48550/arXiv.1810.04805>
- Liu, Z., Zhang, Y., & Huang, J. (2025). Transformer-based reasoning with symbolic integration. *Knowledge-Based Systems*, 289, 111410. <https://doi.org/10.1016/j.knosys.2024.111410>
- Mars, B. (2022). Word2Vec revisited: Modern applications for semantic representation in NLP. *Journal of Language Technology*, 14(3), 45–59.
- Nagahisarchoghaei, S., Hashemi, S., & Farhadi, H. (2023). Random forest enhancements for high-dimensional sparse text features. *Expert Systems with Applications*, 221, 119806. <https://doi.org/10.1016/j.eswa.2023.119806>
- Shari, R., Saleh, M., & Odabaş, M. (2021). Gradient boosting for text classification: Comparative analysis and practical recommendations. *Applied Soft Computing*, 107, 107381. <https://doi.org/10.1016/j.asoc.2021.107381>