

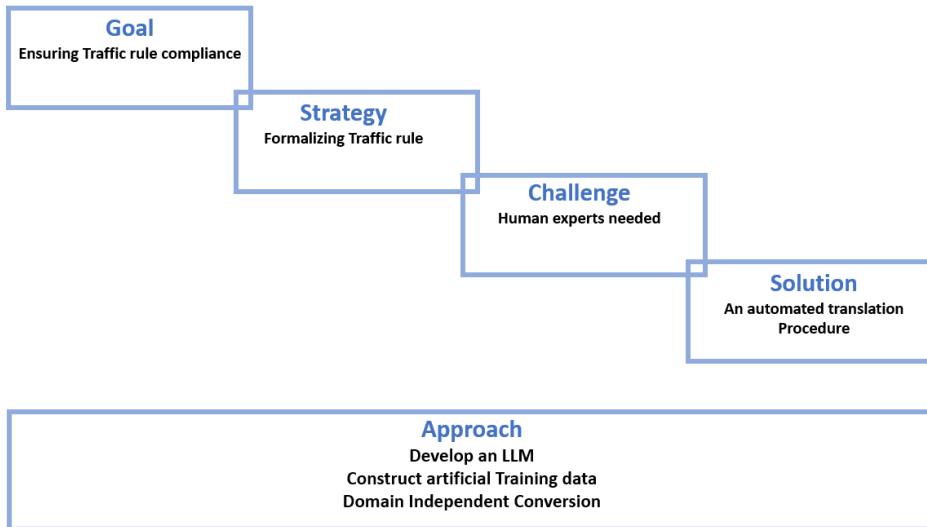
# Automated Formalization of Traffic Rules Using Natural Language Processing

## Master's Thesis of Mudita Kamble

### Mentoring:

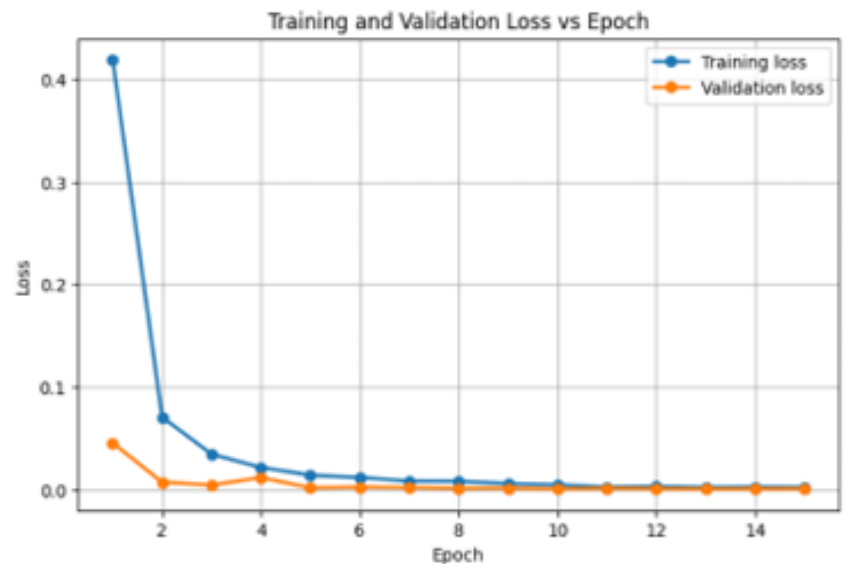
Evald Nexhipi, M.Sc

Athanasia Karalaku, M.Eng



By translating human-readable traffic regulations into formal specifications that machines can understand, the study aims to facilitate compliance with traffic rules in autonomous systems. A workable way to eliminate uncertainty and enable rule-based planning, monitoring, and verification is to formalize traffic regulations. The main issue, however, is that this procedure is typically dependent on human experts with expertise in both legal interpretation and formal methodologies. The thesis suggests an automated translation process based on large language models as a solution to this problem. The method comprises three primary components: developing an LLM-based rule-to-logic framework, creating artificial training data from existing traffic-rule formalizations, and utilizing pseudo-predicate placeholders to facilitate a more structured conversion process.

The loss curve shows how efficiently the selected model learns the mapping from structured natural-language input to formal logic output. A significant decrease in training loss during the early epochs indicates that the model quickly grasped the core rule-to-logic structure, whilst the low and stable validation loss demonstrates that this learning translated well to previously unseen validation samples. Stable convergence and little overfitting are indicated by the proximity of the two curves in the later epochs. Such behavior is especially crucial for this task as a usable translation model needs to be able to both fit the training data and maintain its dependability when presented with new traffic rule formulations. Thus, loss-based analysis shows that the selected model's architecture is appropriate for structured formalization and that the fine-tuning setup was successful.



		TESTING RULES AND SENTENCES			
		Rule 4	Rule 7	Unseen	Full Rule 4
MODELS / FRAMEWORKS	M1_FT1 (T5)	No	Yes	No	No
	M1_FT2 (T5)	No	Yes	No	Yes
	M1_FT3 (T5)	Yes	Yes	No	Yes
	M2_FT1 (Qwen)	Yes	Yes	Yes	Yes
	M3_FT1 (TinyLlama)	No	Yes	No	No
	M3_FT2 (TinyLlama)	No	Yes	No	No
	M4_FT1 (Phi)	No	NA	No	No

The final model comparison was carried out on rule-level prediction tasks involving held-out rules, an unseen sentence, and an extended full-rule formulation. Practical correctness, or whether the created logical formula fits the intended target output, is the main focus of this evaluation. According to the comparison, Qwen performed the best overall, properly anticipating every test case that was chosen. While it continued to perform poorly on the unknown test sentence, the optimal FLAN-T5 configuration also demonstrated good performance and significantly improved over subsequent fine-tuning tests. TinyLlama and Phi, on the other hand, generated less dependable results, indicating that not all open-weight language models are equally appropriate for this purpose. Overall, the findings show that automated traffic-rule formalization is possible, with T5 continuing to be a very promising substitute for small, structured datasets and Qwen emerging as the best current performer.