Thesis Reviewer´s Report

Student:	Godfred Enim Douglas
Title:	Temporal Fusion Transformers for Traffic Flow Prediction in Smart Cities
Supervisor:	Petr Hájek, prof. Ing. Ph.D.
Reviewer:	Komárková Jitka, prof. Ing. Ph.D.
Reviewer's job title:	professor, Institute of Systems Engineering and Informatics, University of Pardubice

Assessment criteria

	excellent	very good	acceptable	unacceptable	N/A
Achievement of the aims of the thesis		x	☐	☐	☐
Use of appropriate methods		x	☐	☐	☐
Depth of analysis (in relation to the topic)		☐	x	☐	☐
Structure and extent of the thesis		x	☐	☐	☐
Use of Czech and foreign sources (including references)		x	☐	☐	☐
Formal aspects (text, diagrams, charts)		**x**	☐	☐	☐
Quality of language (style, grammar, terminology)		**x**	☐	☐	☐

Usability of the results

	high	medium	low	N/A
In theory			x	☐
In practice		x	☐	☐

Other comments

The main aim of the thesis is to introduce temporal fusion transformers, characterise approaches to traffic flow prediction in smart cities, propose a traffic flow prediction system using temporal fusion transformers, pre-process traffic flow datasets, validate the proposed prediction system using the datasets, and discuss implications for smart cities. The thesis defines six specific objectives that logically and systematically break down the main aim. The research questions are well-aligned with these objectives. While the stated objectives form a coherent framework, their fulfilment is uneven, with stronger emphasis on implementation than on critical analysis, validation breadth, or practical deployment considerations.

The thesis represents a comprehensive piece of work. The thesis meets the given aim and objectives. The author uses appropriate methods and a systematic approach, but the methodological scope is narrow. The research relies on a single dataset from one city, limiting the generalisability of findings. The absence of a deeper examination of how the model would perform in different urban contexts weakens the practical value of the results. The spatial component is handled in a simplified way—each of the 36 detector sites is treated as a static covariate, without modelling the actual spatial relationships or dependencies that could significantly influence traffic flow predictions.

Last date and time of printout 11.8.2025 15:14

The literature review identifies a clear literature gap: while TFT has been successful in other domains, its application to traffic flow prediction in a smart city context remains underexplored. This forms a clear motivation for the thesis.

The author uses a real-world dataset comprising traffic measurements from 36 detector sites. The problem is formally defined as a multi-horizon forecasting task, which aligns with the capabilities of the TFT model. The evaluation uses standard metrics (MAE, RMSE), which are appropriate, but the analysis is largely limited to reporting numerical results without deeper interpretation of model behaviour, limitations, or error sources. There is no comparative analysis beyond a simple baseline, and no ablation or sensitivity study to understand the contribution of individual features or model components.

The spatial component of the data is analysed in a simplified manner. The model treats each of the 36 detector sites as a simple static covariate, meaning it can distinguish between different locations but does not explicitly model or analyse the spatial relationships and dependencies between them.

The results presented in Chapter 5 are convincing and supported by data and visualisations. The chapter tends toward repetition rather than synthesis or critical reflection. While the discussion of limitations mentions data dependency, computational cost, and covariate drift, it remains brief and does not explore how these issues could be mitigated in practice.

The author could be more consistent while using some of the terms. Throughout the thesis, the author inconsistently uses "spatiotemporal," "spatio-temporal," and "spatial-temporal" to describe the same type of data. Similarly, slight inconsistencies in hyphenation (e.g., "hyperparameters" vs. "hyper-parameters," "Deep Learning" vs. "Deep-Learning") and the use of synonyms like "prediction" vs. "forecasting" could be improved for greater precision.

The thesis is clearly organised, although the balance between technical detail, critical evaluation, and practical implications could be improved. The thesis is clearly written. Only minor grammatical and typographical errors can be found, which do not detract from the overall quality of the work.

Questions and suggestions for the defence

1. In the thesis, the spatial component is addressed by treating each detector site as a static covariate. This approach does not explicitly model the relationships between different locations. How could the spatial relationships between the detector sites be more deeply analysed?
2. Given that your experiments are based on data from a single city, what challenges and adaptations would be necessary to apply your TFT-based forecasting system to multiple cities with varying traffic patterns, infrastructure, and data availability?"

Overall evaluation

I **recommend** the thesis for defence.
The proposed grade for the thesis: C

In Pardubice on 11.8.2025

Signature: Jitka Komárková