AEX CONTINUUM
doi:10.xxxx/ctm.pp.2026.0001

Title of the Research Paper in Sentence or Title Case

An optional one-line subtitle describing the question or the result

First Author[1,*] **Second Author**[1,*] **Third Author**[2] **Fourth Author**[1,3,†]

[1] Open Systems Lab, Address Line, City, Country

[2] Department of Computer Science, Name of a Rather Long University Institute for Advanced Studies in Artificial Intelligence, City, Country

[3] Third Collaborating Centre for Machine Learning and Robotics, Faculty of Engineering, Name of Another University, City, Country

[*] Equal contribution [†] Corresponding author: fourth.author@example.com

Abstract

State the problem, the approach and the headline result in five or six sentences. A reader should be able to decide from this paragraph alone whether the rest of the report is relevant to them. Quantify the main claim where you can, and say what is new compared with prior work. Placeholder text: replace it with your own.

Keywords: keyword one, keyword two, keyword three

Received 12 October 2026 • Accepted 30 October 2026 • Published 6 November 2026

1 Introduction

Open with the motivation and the gap in current understanding. Cite prior work in brackets [1, 2] or in running text, as Brown et al. [3] do. State the research question precisely, and say why existing approaches do not answer it.

Our contributions are as follows:

- the first contribution, stated as a testable claim;
- the second contribution, ideally with a number attached;
- a released artefact (code, data or model), if there is one.

2 Related Work

Summarise what the reader needs to know and position your work against it. Group related work by theme rather than by paper, and end each group by saying how your approach differs. Cross-reference with `\cref`, which names the target for you: Figure 1, Section 3, Table 1, Equation (1) and Algorithm 1.

3 Method

Describe the approach in enough detail to be reproduced. Equations are numbered and set in the Times-compatible maths font:

$$\mathcal{L}(\theta) = -\frac{1}{N}\sum_{i=1}^{N} \log p_\theta(y_i \mid x_i) + \lambda \lVert \theta \rVert_2^2. \tag{1}$$

Definition 1. Definitions, theorems and proofs use the same restrained style: a bold sans-serif run-in heading, with upright text for definitions and italic text for results.

Theorem 1. *State the result precisely, with all assumptions.*

Proof. Give the argument, or point to Appendix A. □

Algorithm 1 Placeholder training procedure

Input: dataset $\mathcal{D}$, initial parameters θ_0, learning rate η
Output: trained parameters θ_T

```
for t = 0, ..., T − 1 do
    sample a mini-batch B ~ D
    θ_{t+1} ← θ_t − η ∇_θ L(θ_t; B)
    if t mod k = 0 then
        evaluate on the validation set
return θ_T
```

3.1 Implementation details

Configuration and code can be shown with the listing style:

```python
# placeholder example
def train(model, data, epochs=3):
    for _ in range(epochs):
        loss = model.step(data)
    return model
```

Hyper-parameters. Run-in paragraph headings suit short, parallel details such as optimiser, learning rate and batch size.

4 Experiments

Report the setup, then the results. Keep tables light: a heavy top and bottom rule, a single rule under the header, and numbers aligned to the right.

Key finding. Use this box sparingly, for the single result you most want a skimming reader to remember.

© 2026 AEX CONTINUUM and the authors.
This work is licensed under a Creative Commons "Attribution 4.0 International" license.

Cite as: First Author, Second Author, Third Author and Fourth Author (2026). *Title of the Research Paper in Sentence or Title Case.* AEX CONTINUUM. https://doi.org/10.xxxx/ctm.pp.2026.0001

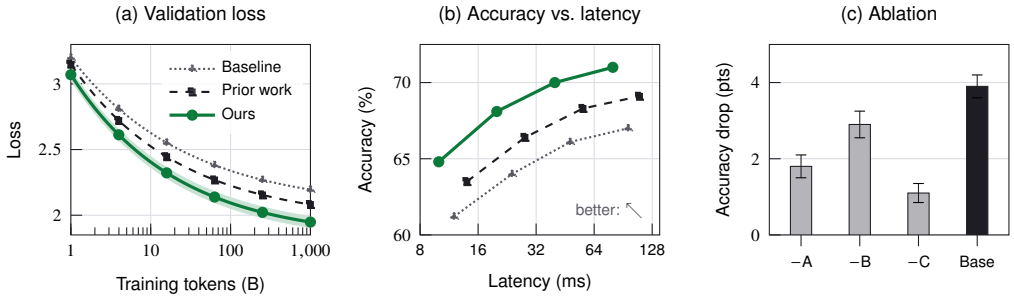


Figure 1. Analysis of the proposed method. (a) Validation loss as the training budget grows; the shaded band is ±1 standard deviation over five seeds. (b) Accuracy against inference latency; up and to the left is better. (c) Ablation: accuracy lost, relative to the full model, when each component (A, B, C) or all of them (Base) is removed; error bars show the standard deviation. *All data shown here are synthetic placeholders.*

Table 1. Placeholder results table. Bold marks the best result in each column. Replace the entries with real measurements.

Method	Metric A ↑	Metric B ↑	Latency (ms) ↓	Params
Baseline	—	—	—	—
Prior work	—	—	—	—
Ours	**—**	**—**	**—**	**—**

4.1 Ablations

Describe which components matter and by how much, and refer to the table that supports each claim.

5 Discussion

Interpret the results, state the limitations honestly, and flag threats to validity and known failure cases. Say what you would do next.

6 Conclusion

Restate the contribution and its significance in two or three sentences.

Acknowledgements

Thank collaborators, reviewers and funders here.

Declarations

Author contributions

First Author and Second Author contributed equally to the methodology and experiments. Third Author provided the data. Fourth Author supervised the work. All authors wrote and approved the manuscript. (Placeholder; the CRediT role names are a good vocabulary here.)

Funding

Placeholder: name the funders and grant numbers, or state that there was none.

Competing interests

Placeholder: the authors declare no competing interests (or list them).

Data and code availability

Placeholder: say where the code and data can be found, with a persistent identifier, and note any access restrictions.

References

[1] A. Vaswani, N. Shazeer, N. Parmar, et al. “Attention Is All You Need”. In: *Advances in Neural Information Processing Systems*. Vol. 30. 2017. arXiv: 1706.03762.

[2] Y. LeCun, Y. Bengio, and G. Hinton. “Deep Learning”. In: *Nature* 521.7553 (2015), pp. 436–444. doi: 10.1038/nature14539.

[3] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, et al. “Language Models are Few-Shot Learners”. In: *Advances in Neural Information Processing Systems*. Vol. 33. 2020. arXiv: 2005.14165.

A Additional Details

Long derivations, extra tables, prompts and hyperparameters belong here.