Building a Large Language Model from Scratch

A Step-by-Step Guide Using Python and PyTorch

Dr. Yves J. Hilpisch[1]

February 1, 2026

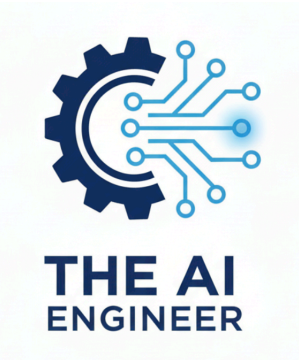

[1] Get in touch: https://linktr.ee/dyjh. Web page: https://theaiengineer.dev. Research, structuring, drafting, and visualizations were assisted by GPT 5.x as a co-writing tool under human direction. Comments and feedback are welcome.

Preface

Small models, big clarity. This book is a hands-on journey to demystify large language models by building a tiny, readable GPT from first principles — one piece at a time. You'll see every moving part, run every example, and finish with code you understand well enough to modify with confidence.

Why start small? Because understanding scales. A compact transformer you can hold in your head is a better teacher than a giant you can only run. Once the ideas click — tokens to vectors, attention as selective mixing, training as gentle iteration — the bigger systems look less like magic and more like engineering.

How we'll work together: narrative first, then action. You'll encounter short, runnable snippets you can try immediately and small scripts you can reuse. Figures come from code so they stay honest. The workflow is deliberately light — venv, a Makefile, and fast edit→run loops — so progress feels steady.

What to expect along the way: a steady rhythm of ideas and practice. You'll measure the right things (shape sanity, learning signals, perplexity), make deliberate trade-offs (size, speed, quality), and ship your work as a tidy bundle with a CLI and a pocket UI. Most of all, you'll build the habit of keeping models legible — clean inputs, explicit masks, and parameters you can explain.

Bring curiosity and patience. Type some code. Read error messages. Ask what could break. The goal is not to memorize an API; it's to learn how to think with tensors, one transparent step at a time. When you reach the end, you'll have a miniature model with a clear voice — and the confidence to push further.

Let's begin.

Disclaimer

This book and its code are for educational purposes only. No warranty is given; use at your own risk. For full details, see Appendix F.

Contents

List of Figures

Contact

Building a Large Language Model from Scratch

A Step-by-Step Guide Using Python and PyTorch

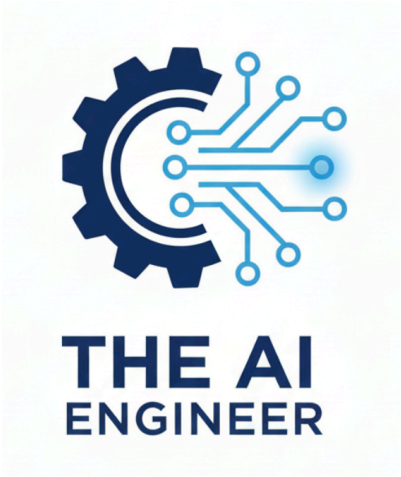

Get in touch:

https://linktr.ee/dyjh

https://theaiengineer.dev