How much compute did the key AI innovations need?

Advanced · A few weeks · Free · Project idea

Online version: https://ai-safety-sandbox.vercel.app/projects/compute-behind-key-innovations/

Summary

Could algorithmic progress continue under a compute halt? Catalogue the innovations in a major model paper and estimate the compute used in the experiments behind each, extending published MIRI TGT work.

Why it matters

If key innovations needed little compute, limits on big training runs would not stop algorithmic progress. This question shapes how pauses and agreements are designed.

What you will have at the end

A dataset of innovations with compute estimates.

How to start (tick each box as you go)

- [ ] **1. Read the original paper on compute used in experiments behind key innovations.**
- [ ] **2. Pick a different model's technical report than the one they used.**
- [ ] **3. List its innovations and find the experiments behind each.**
- [ ] **4. Estimate compute per experiment, with ranges.**
- [ ] **5. Compare your findings with the original paper.**

Key terms

Compute
The raw computing power (specialised chips running for a long time) used to train and run AI models. Easier to measure and track than data or ideas.

Fermi estimate
A quick, rough calculation that breaks a hard number into easier guesses multiplied together. Good enough to tell whether something is a big deal.

Source and links

Technical AI Governance resource pack (MIRI TGT reading lists)
https://www.lesswrong.com/posts/GRPswd5v7haob97fA/recommendations-for-people-getting-into-technical-ai

Compute Requirements for Algorithmic Innovation in Frontier AI Models (2025)
https://arxiv.org/abs/2507.10618

When you have finished

Write a few paragraphs on what you did, what you found and what you would try next, and share it with a friend, your society or online.

Courses, mentors, small grants and communities:
https://ai-safety-sandbox.vercel.app/next-steps/

Questions? Ask in AI Safety Cardiff University (AISCU): https://aiscu.org/

Notes