

SLMs

Specialised and Small Language Models (SLMs) including information on alternative approaches to transformers

- [Less is More: Recursive Reasoning with Tiny Networks](#)
- [Mistral Small 3.2](#)
- [Hierarchical Reasoning Model](#)

Less is More: Recursive Reasoning with Tiny Networks

Paper: <https://arxiv.org/html/2510.04871v1>

Abstract

[Hierarchical Reasoning Model \(HRM\)](#) is a novel approach using two small neural networks recursing at different frequencies. This biologically inspired method beats Large Language models (LLMs) on hard puzzle tasks such as Sudoku, Maze, and ARC-AGI while trained with small models (27M parameters) on small data (~ 1000 examples). HRM holds great promise for solving hard problems with small networks, but it is not yet well understood and may be suboptimal. We propose Tiny Recursive Model (TRM), a much simpler recursive reasoning approach that achieves significantly higher generalization than HRM, while using a single tiny network with only 2 layers. With only 7M parameters, TRM obtains 45% test-accuracy on ARC-AGI-1 and 8% on ARC-AGI-2, higher than most LLMs (e.g., Deepseek R1, o3-mini, Gemini 2.5 Pro) with less than 0.01% of the parameters.

Mistral Small 3.2

Mistral Small is a 24B param LLM that

Running in Ollama

```
ollama pull hf.co/gabriellarson/Mistral-Small-3.2-24B-Instruct-2506-GGUF:Q4_K_M
```

References

- <https://simonwillison.net/2025/Jun/20/mistral-small-32/>

Hierarchical Reasoning Model

Paper URL: <https://arxiv.org/pdf/2506.21734>

Code Repo: <https://github.com/sapientinc/HRM>

HRM is an alternative to transformer architecture that is better able to reason. It outperforms transformer-based LLMs at ARC-AGI2 with only 27M parameters.

Training a 27M Parameter Model with 1000 Examples

In the paper the authors refer to the fact that they only use between 1000 and 10,000 examples for specific problem domains:

- **Sudoku-Extreme:** 1000 training examples (used in main experiments)
- **Sudoku-Extreme-Full:** ~10,000 examples (used in analysis experiments for convergence guarantees)
- **ARC-AGI:** ~1000 examples from the official dataset, heavily augmented with translations, rotations, flips, and color permutations

This may seem quite low considering that this is a 27M parameter neural network and it seems likely that the network would be underfit after so few examples. The authors provide some additional clarifications around this point:

1. Data augmentation is used in order to functionally boost the size of the training set.
2. The authors use deep supervision to augment the training process (rather than relying on back-propagation alone).
3. The problem domain is simpler than for language - particularly for things like Sudoku and ARC-AGI - these are structured grid type problems.