

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.

Revision #3

Created 15 October 2025 08:59:02 by James

Updated 16 October 2025 14:58:10 by James