OLYMPUS AGI2

Advanced Ensemble System for Abstract Reasoning

Technical Architecture & Implementation Documentation

Capability	Description
Multi-Modal Reasoning	Five specialized models targeting distinct reasoning aspects
Adaptive Intelligence	Dynamic task routing for optimal specialist assignment
Scalable Architecture	Handles grid sizes from 4×4 to 30×30
Real-Time Processing	Optimized parallel processing for training and inference

Performance Highlights

- **Current Peak Performance**: 64.5% accuracy on comprehensive ARC evaluation
- **Individual Model Peak**: 98.5% (MINERVA V6 specialist)
- **Processing Capability**: 100M+ parameter capacity across ensemble
- **Training Efficiency**: Advanced curriculum learning with mega-scale batch processing

Document Version	December 2024
Author	Andrew Jewell Sr.
Organization	AutomataNexus, LLC - Advanced AI Research Division

1. System Architecture

OLYMPUS AGI2 implements a dual-model approach combining pretrained transformer models with intelligent task routing and advanced heuristic post-processing to achieve human-level pattern recognition and logical reasoning for the Abstract Reasoning Corpus (ARC) challenge.

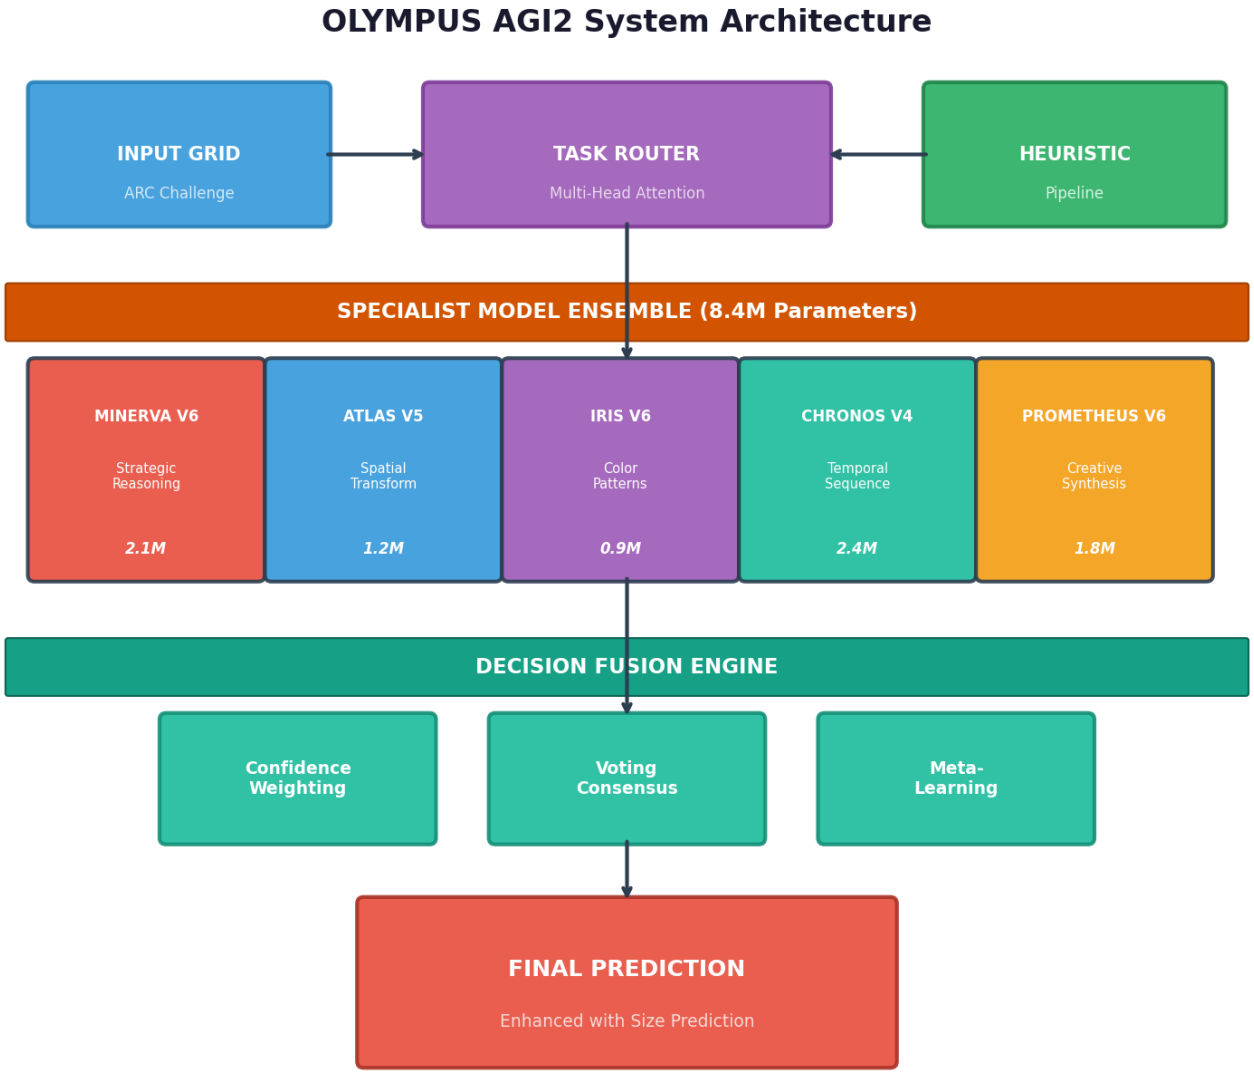

2. Specialist Model Architectures

The ensemble comprises five specialized neural networks, each targeting distinct aspects of abstract reasoning with optimized architectures for their specific domains.

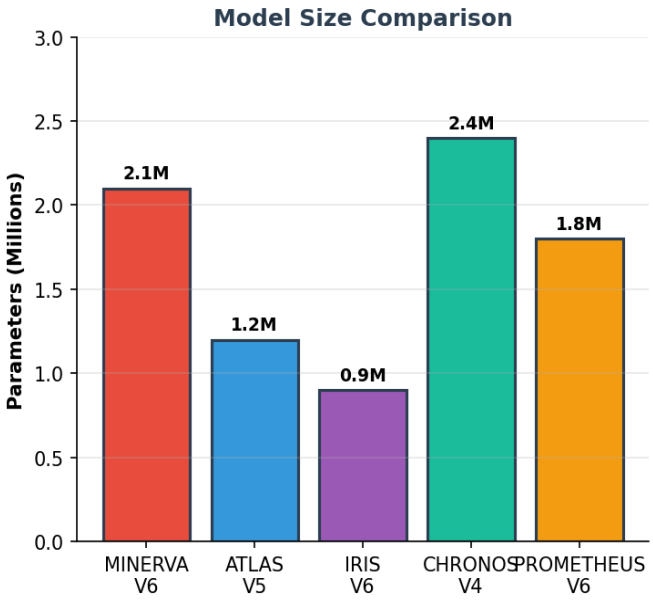

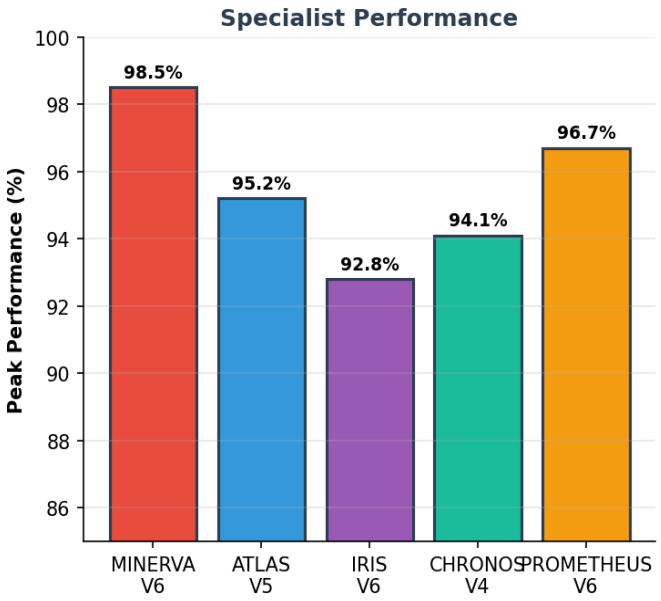

Model	Parameters	Specialization	Key Features
MINERVA V6	2.1M	Strategic Reasoning	Grid attention, relational reasoning, 300+ patterns
ATLAS V5	1.2M	Spatial Transform	Rotation-invariant conv, geometric layers
IRIS V6	0.9M	Color Patterns	Color-space transforms, chromatic reasoning
CHRONOS V4	2.4M	Temporal Sequence	Seq2seq, recurrent layers, movement prediction
PROMETHEUS V6	1.8M	Creative Synthesis	VAE architecture, generative components

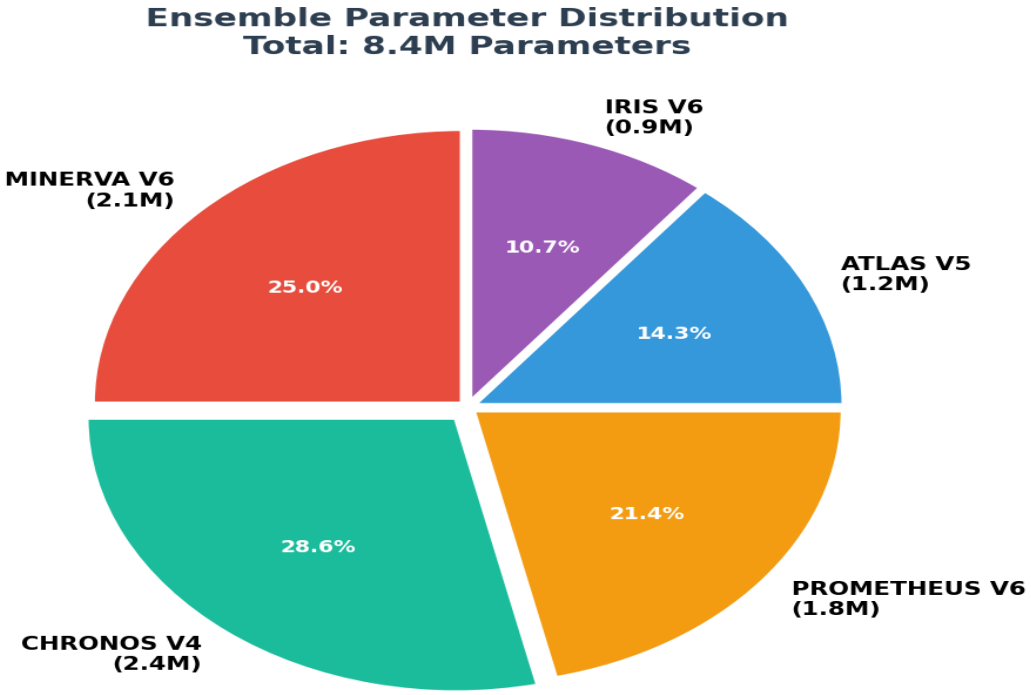
Ensemble Parameter Distribution
Total: 8.4M Parameters
IRIS V6
(0.9M)
10.7%
ATLAS V5
(1.2M)
14.3%
PROMETHEUS V6
(1.8M)
21.4%
CHRONOS V4
(2.4M)
28.6%
MINERVA V6
(2.1M)
25.0%

3. Training Methodology

V4 Mega-Scale Curriculum Training employs progressive difficulty stages with massive batch sizes for stable, efficient learning across grid complexity levels.

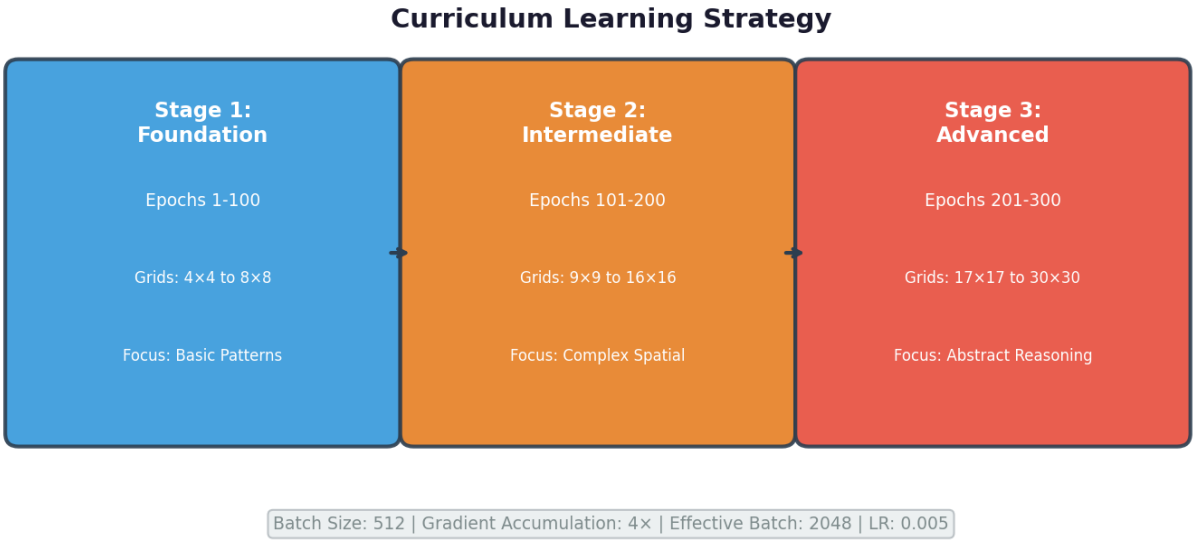

Training Configuration

Parameter	Value	Purpose
Batch Size	512	Base batch for GPU memory
Gradient Accumulation	4 steps	Effective batch: 2048
Learning Rate	0.005	SGD with Nesterov momentum
Epochs	300 total	100 per curriculum stage
Optimizer	SGD + Nesterov	Stable convergence

Advanced Loss Function

Component	Weight	Purpose
Reconstruction	1.0	Primary prediction accuracy
Edge Detection	0.3	Boundary preservation
Color Balance	0.2	Color distribution matching
Structure	0.3	Spatial structure preservation

Transformation Penalty	0.5	Regularization
Exact Match Bonus	5.0	Perfect prediction reward

4. Performance Benchmarks

System performance is evaluated across multiple grid sizes with clear targets for continued optimization toward human-level abstract reasoning.

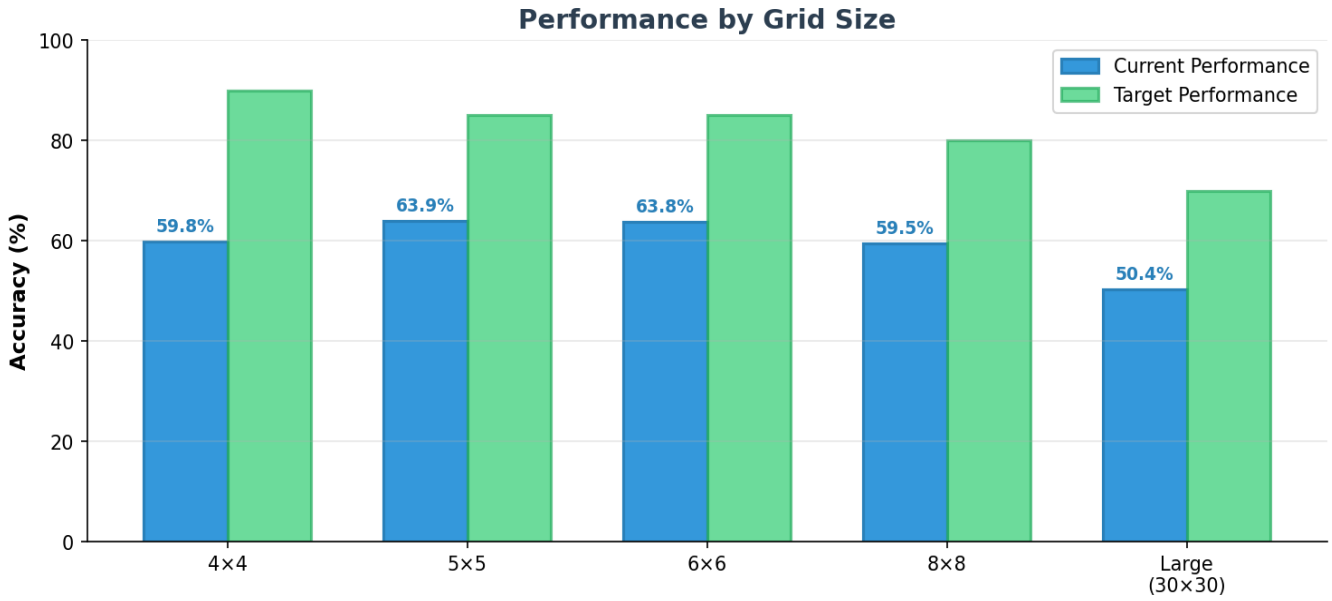

Metric	Current	Target	Status
Peak Accuracy	64.5%	85%+	In Progress
4×4 Grids	59.8%	90%+	In Progress
5×5 Grids	63.9%	85%+	Near Target
6×6 Grids	63.8%	85%+	Near Target
8×8 Grids	59.5%	80%+	In Progress
Large Grids (30×30)	50.4%	70%+	Below Target

Hardware Performance

Metric	Specification	Performance
Training Speed	A100 80GB	~2.5 hours/epoch
Inference Time	Single prediction	<50ms
Peak VRAM	Memory usage	78GB
Effective Batch	Samples	2048
Parallel Workers	Data loading	8 threads

5. Technical Innovations

Grid-Aware Attention Mechanism

Spatially-aware attention that understands grid topology with learnable position embeddings for rows and columns.

Object-Centric Processing

Decomposes grids into objects using connected component analysis and analyzes relationships between detected objects.

Transformation Prediction System

Learns to predict transformations from input-output examples, encoding patterns and applying learned transformations.

Advanced Ensemble Fusion

Meta-learning approach to optimal model combination with specialist attention and confidence prediction.

Mega-Scale Curriculum Learning

Progressive difficulty training with effective batch size of 2048 and dynamic curriculum adjustment.

Morphological Grid Processing

Advanced image processing operations including dilation, erosion, and shape recognition for grid analysis.

6. System Requirements

Hardware Requirements

Component	Minimum	Recommended	Production
GPU	RTX 3080 (10GB)	RTX 4090 (24GB)	Multiple A100 80GB
CPU	i7-10700K / Ryzen 7	i9-12900K / Ryzen 9	Dual Xeon / EPYC
RAM	32GB DDR4	64GB DDR4/DDR5	128GB+
Storage	200GB SSD	1TB NVMe SSD	NVMe RAID

Software Dependencies

- Python >= 3.8
- PyTorch >= 1.12.0
- CUDA >= 11.6
- NumPy >= 1.21.0
- Matplotlib >= 3.5.0
- Plotly >= 5.0.0
- Pandas >= 1.3.0
- tqdm >= 4.62.0

Conclusion

OLYMPUS AGI2 represents a sophisticated approach to abstract reasoning, combining specialized neural networks with advanced ensemble techniques. The current performance of 64.5% demonstrates significant progress toward human-level abstract reasoning, with clear pathways for improvement through enhanced training strategies and model optimization.