



Domain Specific Model Working Group

Open Power AI Consortium



Ben Sooter
Director R&D

19 February 2026

Domain Specific Model Working Group Vision

Develop
Electric Sector
DSMs and
Provide to the
Industry



End-User
Deployment
Experiences



Benchmark
DSMs to
Understand
Performance



Technical
Transfer of
Lessons
Learned



Value through
Collaboration



Collecting and Sharing Lessons Learned through Developing and Using DSMs



Public Electric Sector Data Sources

Public Electric Sector Data Sources

- EPRI Knowledge Graph
- Other Electric Sector Data Source Opportunities



Synthetic Data

Synthetic Datasets for ML Object Detection Training

Concept and Pipeline

- Use **parametric 3D CAD models** rendered photorealistically (like Blender-based automations) to generate **labeled training data at scale**
- Automated variation across:
 - **Lighting**: intensity, angle, glare, color temperature
 - **Scene conditions**: occlusion, clutter, motion blur, partial obstruction
 - **Camera parameters**: distance, focal length, depth-of-field, sensor noise
- Output includes **perfect ground truth labels** (bounding boxes, masks) with zero annotation cost

Why Synthetic Data Improves Training


- **Data coverage & edge cases**: Generate rare or unsafe scenarios that are hard to capture in the real world
- **Label accuracy**: Eliminates human annotation error and inconsistency
- **Rapid iteration**: New objects, configurations, or failure modes can be introduced immediately
- **Cost & scalability**: Orders-of-magnitudes cheaper and faster than physical data collection

Common Pitfalls and Considerations


- **Domain gap risk**: Models may overfit to synthetic artifacts
- **Bias amplification**: Over-sampling ideal conditions can degrade real-world robustness
- **False confidence**: High synthetic validation accuracy may not reflect field performance

Best Practices to Mitigate Risk

- Mix **synthetic + real-world data** (synthetic for coverage, real for grounding)
 - Roughly 90% synthetic, 10% real data
- Apply **domain randomization** aggressively, not just photorealism
- **Validate** on real datasets early and continuously



Best Practice: Combine synthetic pre-training with fine-tuning on real-world samples (mixed-reality approach) for maximum model robustness.



Discussion





The Future of DSMs



'ENGINEERING DOMAIN SUPERIORITY'

Advanced Methodologies for
High-Performance
Specialized Language Models



★ Data-
Efficiency
Gap

THE 2025-2026 ARCHITECTURAL PIVOT

★ THE
PERFORMANCE

★ CEILING

GENERALIST
LLM

★ ? Deep
Parametric
Understanding

MEDICINE

LAW

FINANCE

ENGINEERING

'THE DISAPPOINTMENT GAP'

NAIVE RAG FAILURES




STANDARD SFT FAILURES



Pillar 1: Deep Parametric Intuition (CPT)

Self-Supervised
Magic



Source
Distribution

Scaling Law

Target
Loss

Data Scale


It's learning
the distribution!



Weights, not
just Prompts!

ICLR 2025
Breakthrough

Pillar 1.6: The Synthetic Power-Up



→ Turning Scarcity
into Abundance.

Large-Scale
Synthetic Corpus

Log-Linear
Accuracy
Gains!

Pillar 1.6: The Synthetic Power-Up

ICLR
2025
Breakthrough

The
EntiGraph
Synthesizer

Small
Real
Data

Entities

Relations

→ Turning
Scarcity into
Abundance.

Large-Scale
Synthetic Corpus

Log-Linear
Accuracy
Gains!

Pillar 2: The Open-Book Exam (RAFT)

Don't just memorize—
Use the evidence!



Real-world
production
recipe.

Ignore
the noise!




'Pillar 3: The Retrieval Evolution'

Standard Pipeline RAG




V
S
VERSUS


Knowledge Graph RAG





🌀 Pillar 3.1: Hierarchies & Communities (GraphRAG) 🌀




Pillar 3.2: Agentic Self-Correction




Pillar 4: The Long-Context Speedster (Mamba)




Pillar 4.1: The Specialist Brain (MoE)




Pillar 5: Precision Optimization (DoRA & GaLore)





Pillar 6: Alignment via Utility (KTO)



Pillar 7: The Tool-Augmented Advantage









Start by acknowledging that we are in the Era of Expertise. In 2020, the hype around one-size-fits-all generalist models was at its peak. The Goal: Today, we're looking at the "Living AI Stack"—a combination of synthetic training, structured retrieval, and prompt engineering.

Pillar 8: Enterprise AI Readiness



The Architecture of Expertise (2026)

The End of
One-Model-Fits-All



Next Steps

- ✓ Build the Eval
- ✓ Upgrade Retrieval
- ✓ Distill the Best





TOGETHER...SHAPING THE FUTURE OF ENERGY®