Neurosymbolic AI (NSAI): Learning Meets Reasoning

Kamal Acharya, Ph.D.
Postdoctoral Research Associate
Aerospace Vehicle Intelligence and Autonomy (AVIA) Laboratory
Department of Mechanical Engineering
School of Engineering and Computer Science
Baylor University

Why accurate prediction is not enough

WHY NEUROSYMBOLIC AI?

Neural AI predicts
Route A: 18 minutes
Landing reserve: 12%

Symbolic safety rule
Required reserve >20%
Reject Route A; revise the plan

NSAI conclusion: Route A is the fastest but it violates the safety rule

The neural model estimates outcomes and symbolic reasoning enforces a non-negotiable constraint.

AI evolved through rules, learning and their reunion

Each wave solved problems the previous one could not.

1 SYMBOLIC AI	2 NEURAL AI	3 NEUROSYMBOLIC AI
1950s–1980s	**1980s–2010s**	**2000s–present**
Rules, logic and search make reasoning explicit, but systems can be brittle.	Backpropagation, data and computing make pattern learning practical.	Learning is connected with rules, knowledge or reasoning.

Neurosymbolic AI is a convergence: learning supplies adaptability; symbols supply structure.

Neurosymbolic AI grew from two research traditions

The integration is new in momentum, not in intellectual origin.

1956	1986	2005	2022
Dartmouth workshop helped establish AI as a field	Backpropagation renewed learning in multilayer networks	Structured survey mapped neural–symbolic integration	Kautz framed a “Third AI Summer” including hybrid systems

The NSAI journey: explicit reasoning → trainable learning → systematic integration → scalable hybrid intelligence

The two paradigms solve different parts of intelligence

NEURAL AI

Learns patterns from examples

- Handles images, sensor streams and complex data
- Adapts when sufficient examples are available
- Often difficult to inspect or constrain directly

Recognizes what usually happens

SYMBOLIC AI

Represents knowledge explicitly

- Encodes rules, relations and constraints
- Supports transparent reasoning and checking
- Can be brittle and expensive to construct manually

Specifies what must or may happen

Neurosymbolic AI connects statistical pattern recognition with explicit, inspectable knowledge.

Neurosymbolic AI connects a learner with a reasoner

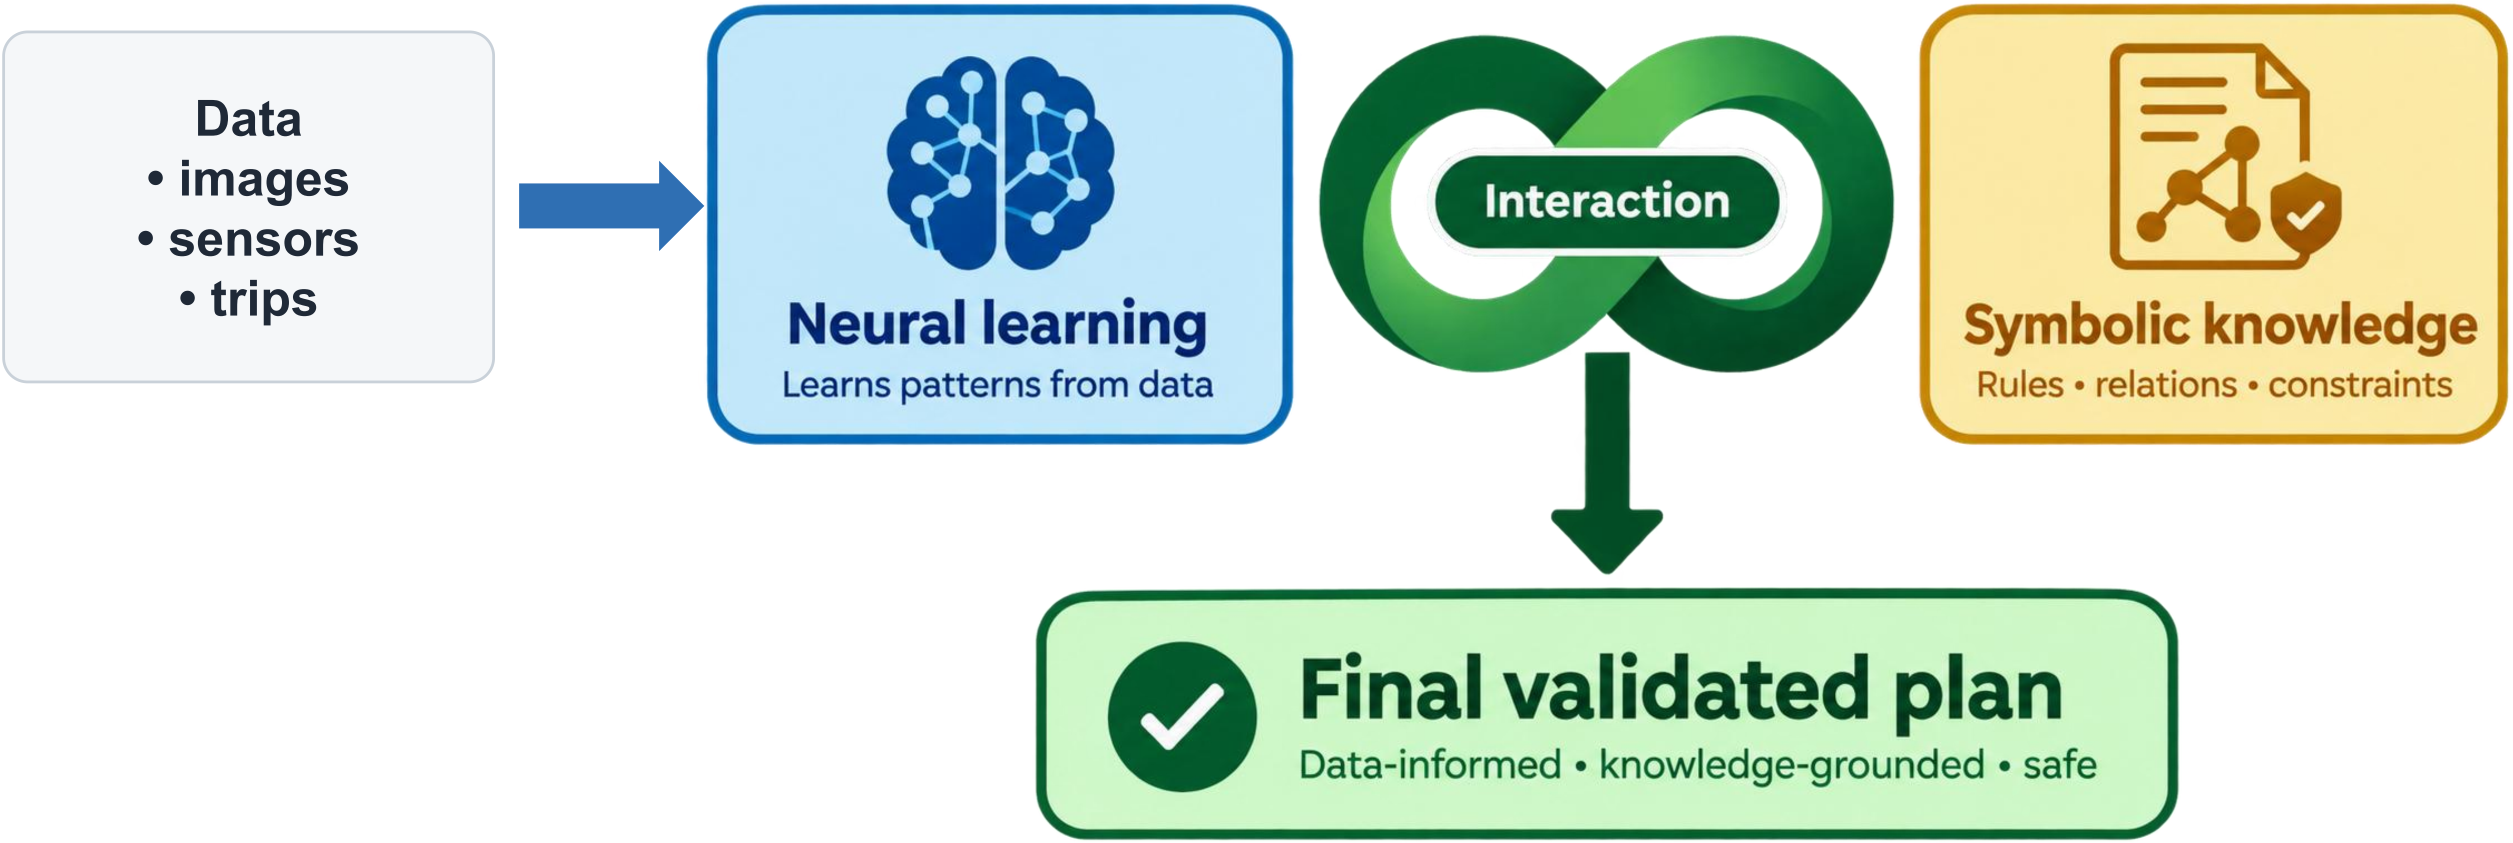

Symbolic knowledge makes assumptions inspectable

A "symbol" is a named, inspectable representation that can participate in reasoning.

IF–THEN RULE	CONSTRAINT	KNOWLEDGE GRAPH	PROGRAM OR PLAN
If reserve < minimum, reject the route	No two aircraft may occupy the same protected region	Aircraft → uses → vertiport Vertiport → has → charger	Take off → cruise → land with ordered preconditions

Symbols may come from experts, data, or a neural model and then participate in explicit reasoning.

Information-flow patterns organize Neurosymbolic AI

1 LEARNING FOR REASONING

Neural → Symbolic

A learned model creates symbols or representations that a reasoner uses.

2 REASONING FOR LEARNING

Symbolic → Neural

Rules or constraints shape what and how the neural model learns.

3 LEARNING–REASONING

Neural ↔ Symbolic

Learning and reasoning exchange information repeatedly in one system.

LLMs can turn implicit knowledge into explicit symbols

LEARNING FOR REASONING

Large language model implicit knowledge

Knowledge distillation extracts facts, rules or graphs

Symbolic reasoner checks, combines and applies them

Inspectable knowledge for downstream reasoning

Source: Acharya, Velasquez, and Song, IEEE Transactions on Artificial Intelligence, 2024.
DOI: 10.1109/TAI.2024.3428519

Symbolic rules shapes what neural model learns

REASONING FOR LEARNING

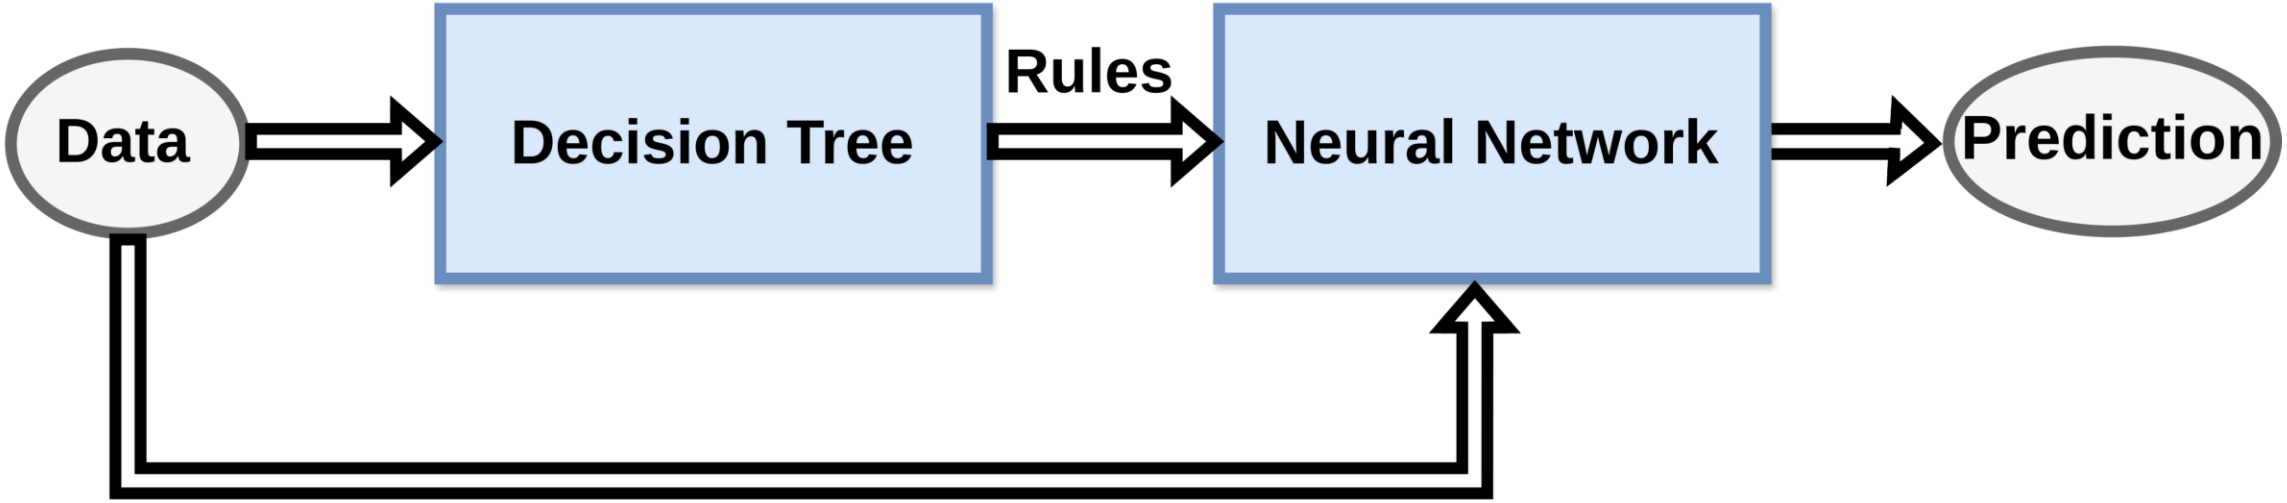

Source: Acharya et al., IEEE IWCMC, 2025. DOI: 10.1109/IWCMC65282.2025.11059465

Learning–reasoning systems close the loop

LEARNING–REASONING

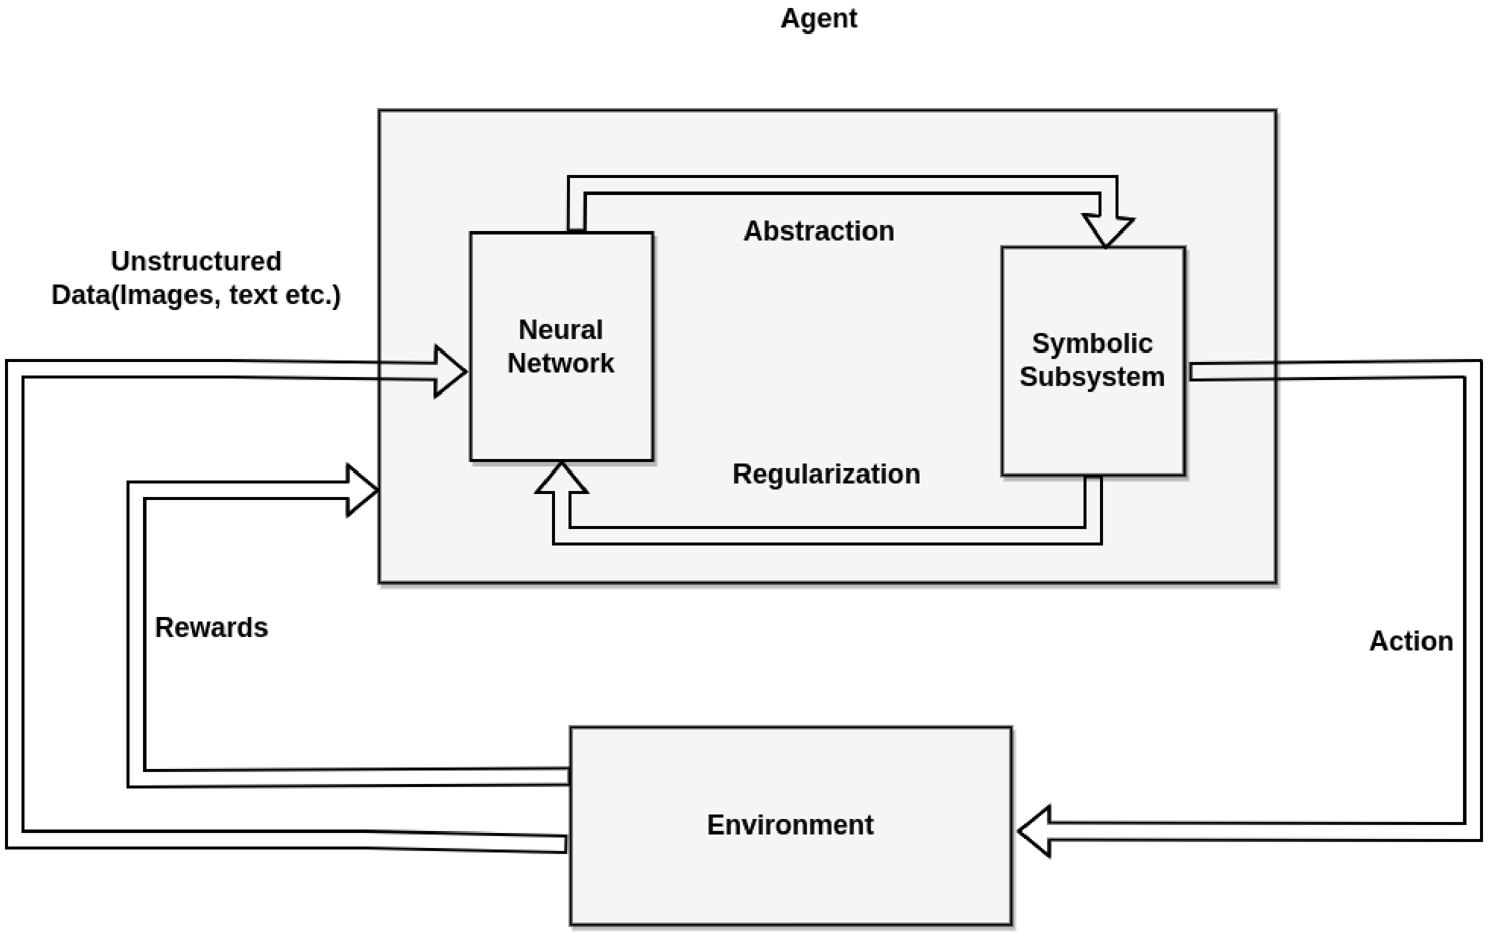

Source: Acharya, Raza, Dourado, Velasquez, and Song, IEEE Transactions on Artificial Intelligence, 2024. DOI: 10.1109/TAI.2023.3311428

Integration alone does not guarantee trustworthy AI

A hybrid system can still fail in its data, knowledge or neural–symbolic interface.

NEURAL FAILURE	SYMBOLIC FAILURE	INTERFACE FAILURE
Noisy data, unfamiliar conditions or poorly calibrated predictions	Incorrect, incomplete, outdated or conflicting rules	Information may be lost, misaligned or applied at the wrong stage

Trustworthy NSAI requires explicit evaluation not just architectural integration.

Trustworthiness needs three properties

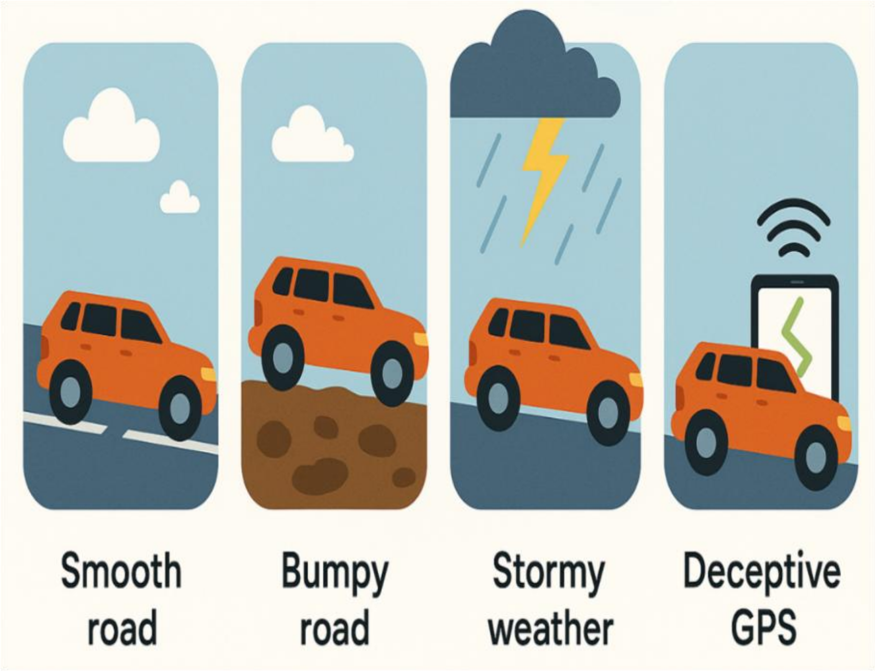

Stable Performance Under Varied Condition
Robustness

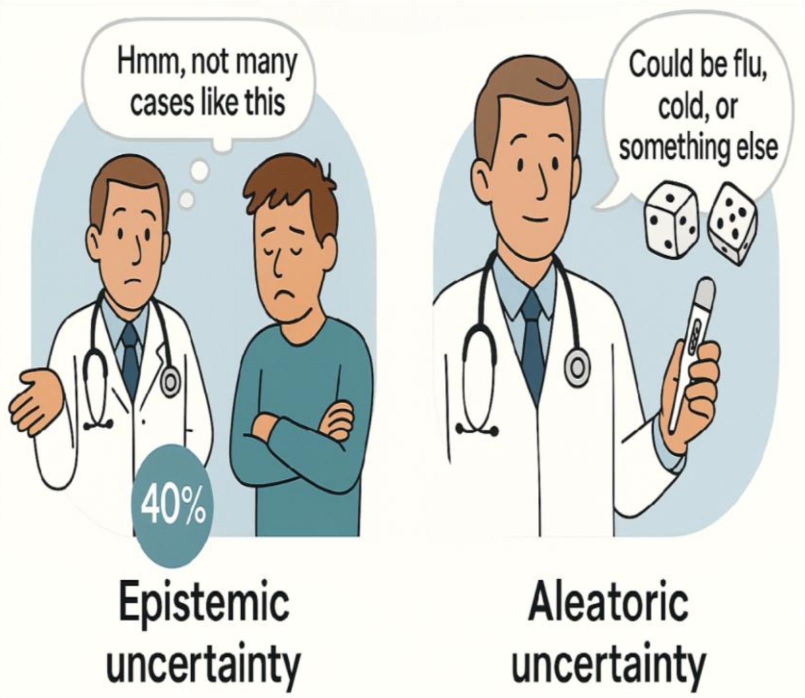

Measuring Confidence in Prediction
Uncertainty Quantification

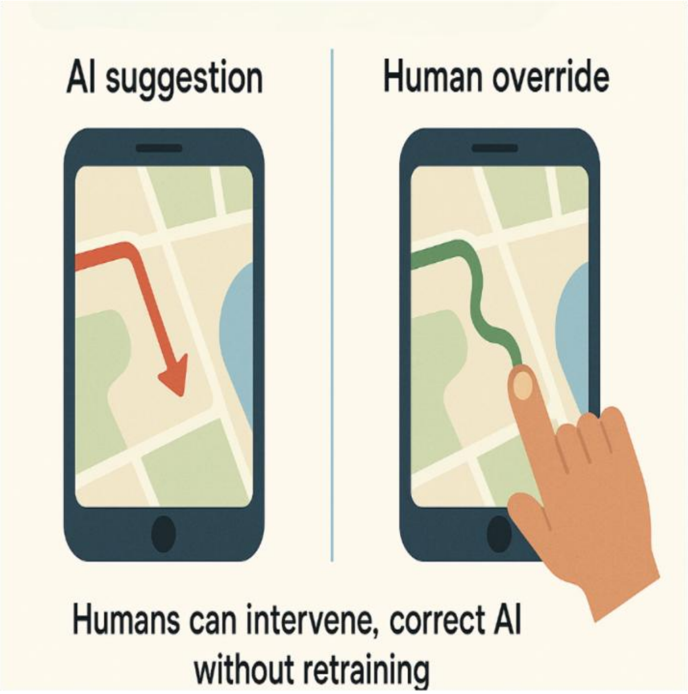

Human Ability to Modify or Influence AI
Intervenability

Neurosymbolic Robustness Pipeline

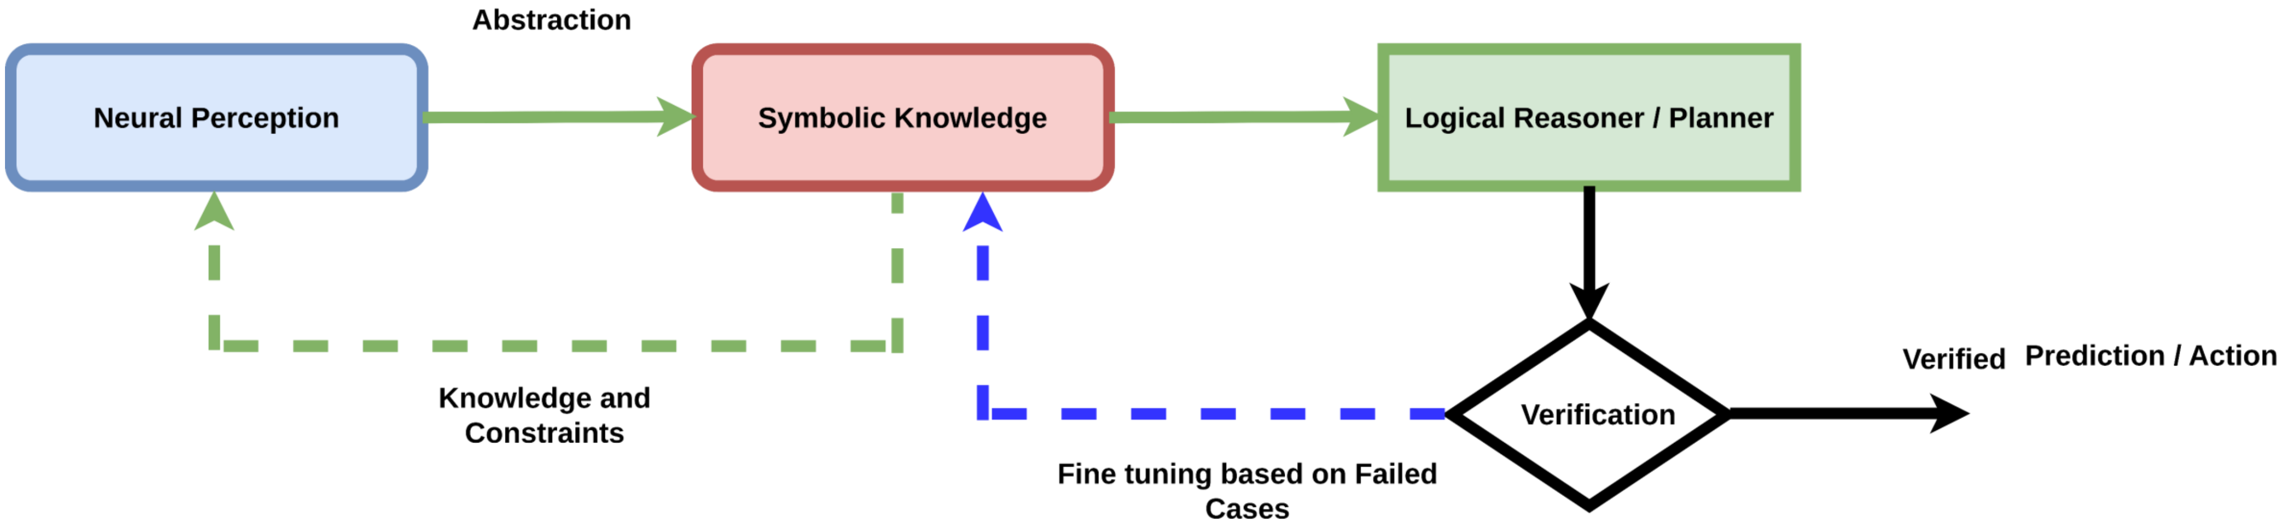

Source: Acharya and Song, Arabian Journal for Science and Engineering, 2026. DOI: 10.1007/s13369-025-10887-3

Neurosymbolic Uncertainty Quantification Pipeline

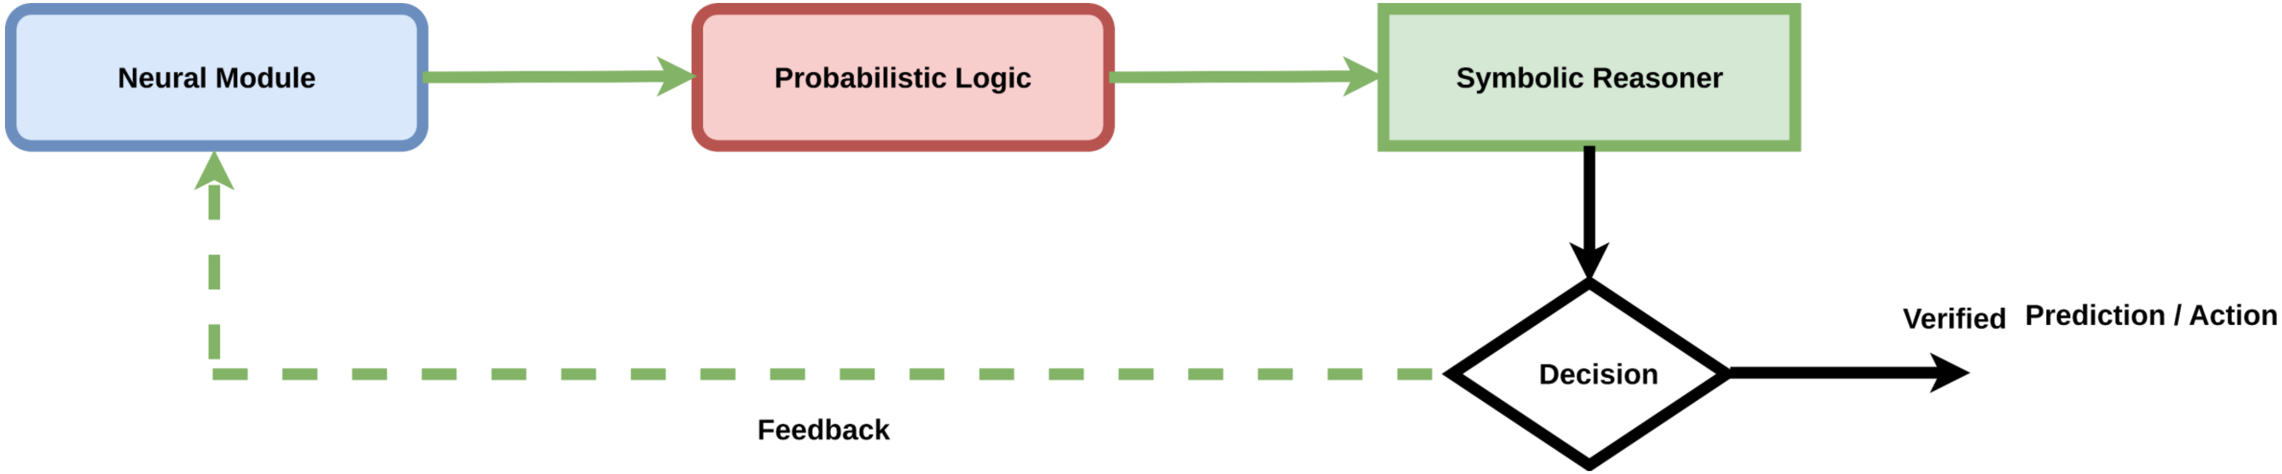

Source: Acharya and Song, Arabian Journal for Science and Engineering, 2026. DOI: 10.1007/s13369-025-10887-3

Human Intervention in a Neurosymbolic Pipeline

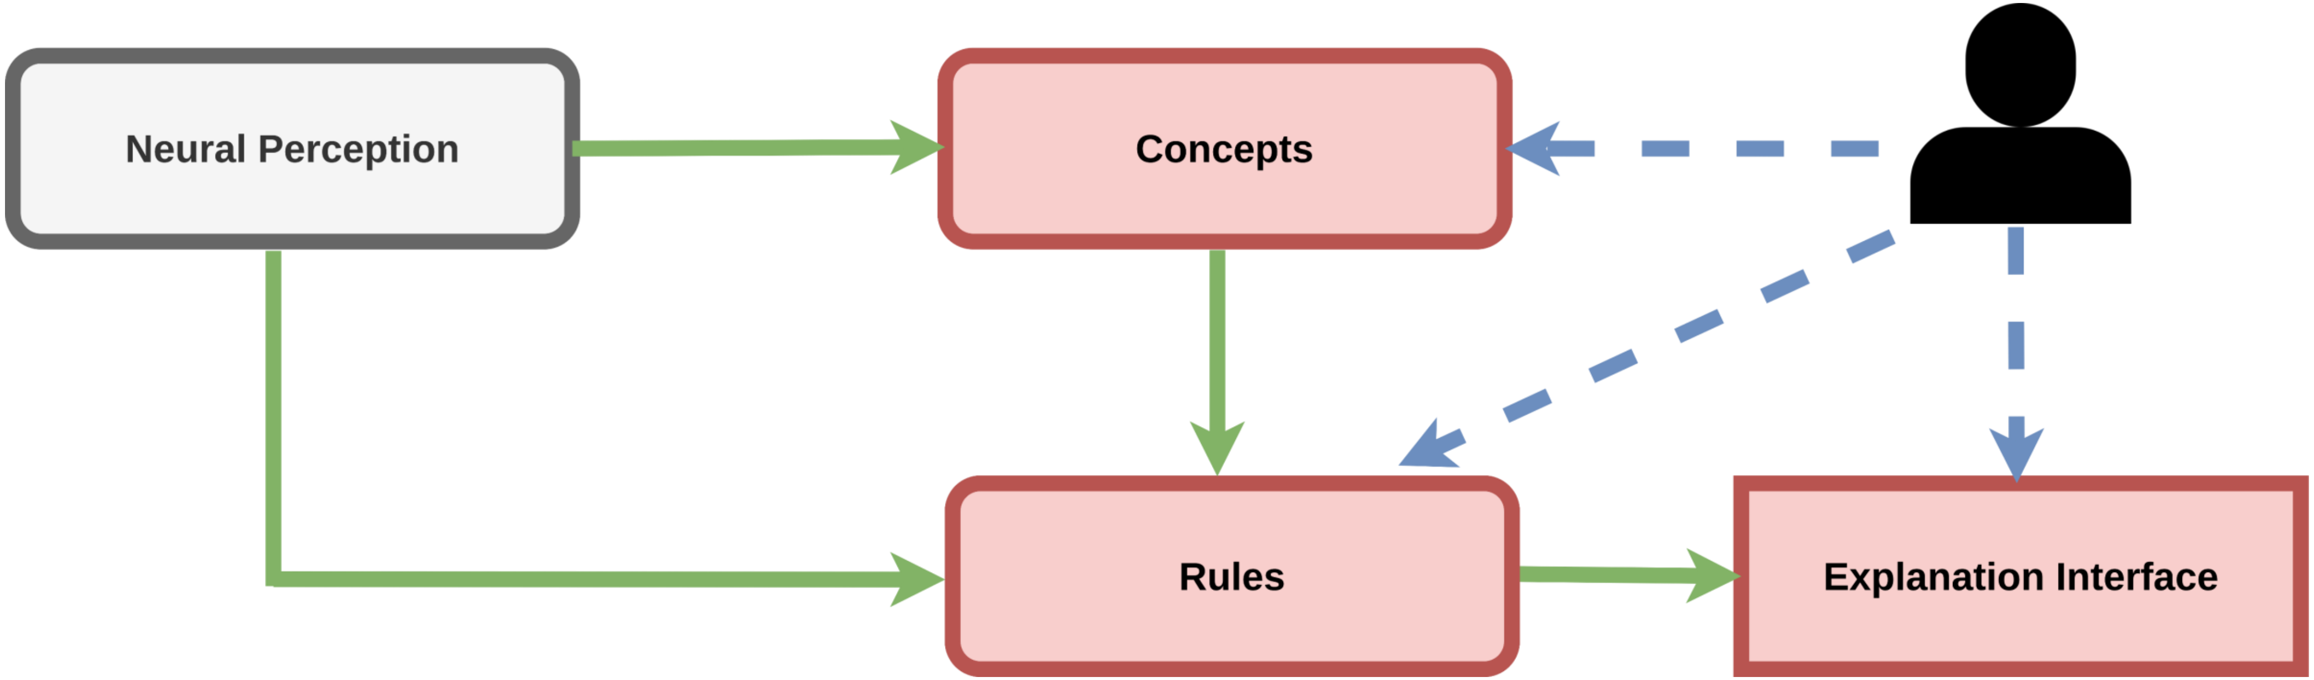

Source: Acharya and Song, Arabian Journal for Science and Engineering, 2026. DOI: 10.1007/s13369-025-10887-3

AAM and NSAI

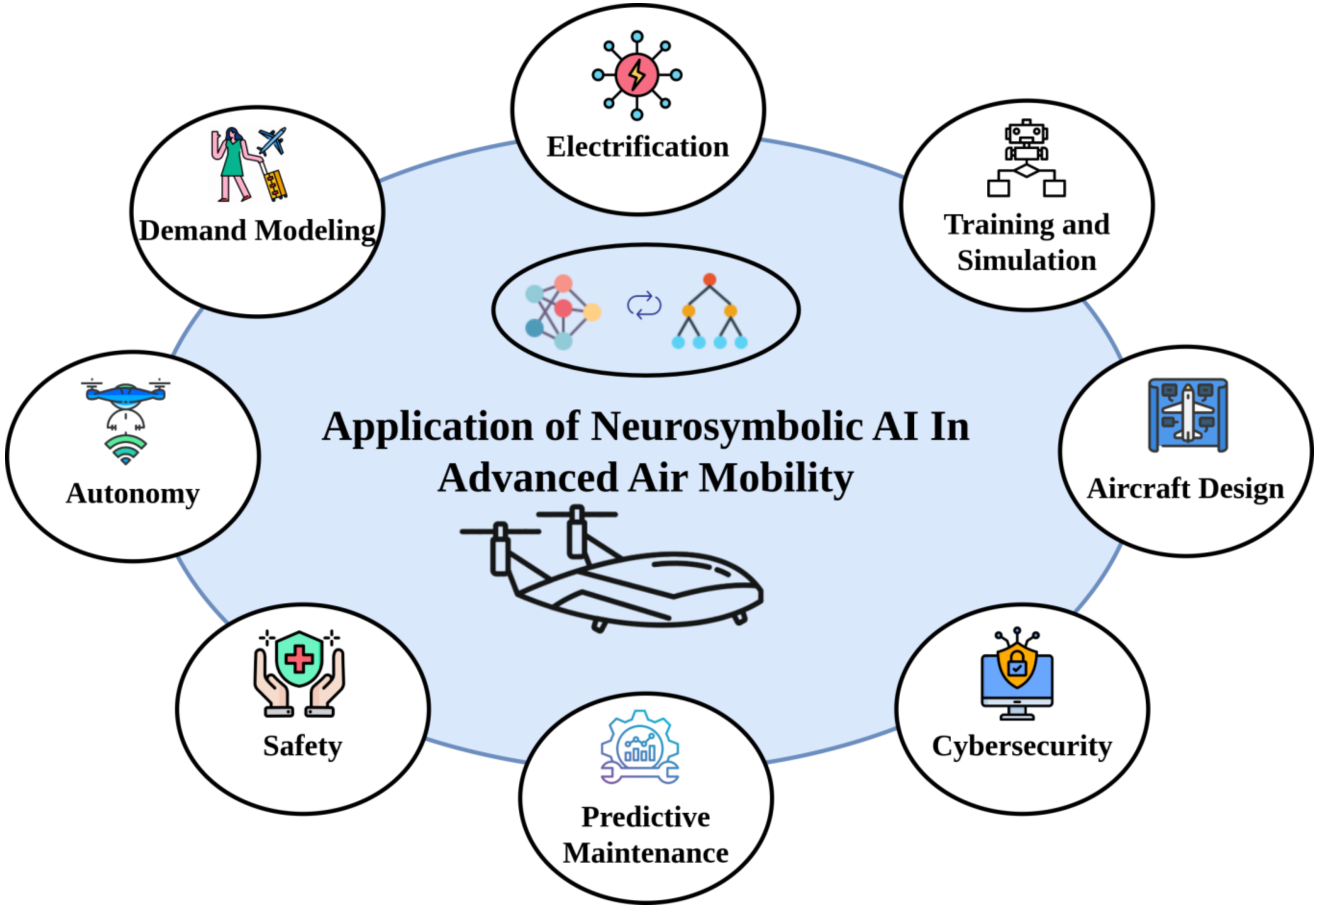

Source: Acharya et al. (IJCAI 2025)

OPERATIONS

Demand modeling
Autonomy
Predictive maintenance

ENGINEERING

Electrification
Aircraft design
Training and simulation

ASSURANCE

Safety • Cybersecurity

Next step: a neurosymbolic gravity model

DATA-DRIVEN GRAVITY BASELINE

We built data-driven gravity model that learns how O-D, geospatial and socioeconomic data shape its trip demand.

Acharya et al., IEEE CAI (2025)

THEORY-CONSISTENT EXTENSION

Add gravity knowledge as soft constraints:

. More opportunities should not reduce flow

COUNTERFACTUAL EXPLANATIONS

Ask: “What would change the predicted flow?”

Return plausible changes while keeping the explanation consistent with gravity relations.

Goal: extend our tested baseline into an interpretable model with theory-consistent predictions and what-if explanations.

NSAI connects our research program

1

Learn from data

Neural models discover complex patterns in LLMs, mobility data and AAM demand.

2

Make knowledge explicit

Rules, relations and domain theory make predictions inspectable and constrainable.

3

Combine, evaluate and extend

Trustworthiness guides deployment; the gravity model carries the architecture forward.

The research direction: accurate predictions that can be checked, explained and revised.

References

K. Acharya, W. Raza, C. Dourado, A. Velasquez and H. H. Song, "Neurosymbolic Reinforcement Learning and Planning: A Survey," in *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 5, pp. 1939-1953, May 2024, doi: 10.1109/TAI.2023.3311428

K. Acharya, A. Velasquez and H. H. Song, "A Survey on Symbolic Knowledge Distillation of Large Language Models," in *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 12, pp. 5928-5948, Dec. 2024, doi: 10.1109/TAI.2024.3428519

K. Acharya, M. Lad, L. Sun and H. Song, "Neurosymbolic Approach for Travel Demand Prediction: Integrating Decision Tree Rules into Neural Networks," *2025 International Wireless Communications and Mobile Computing (IWCMC)*, Abu Dhabi, United Arab Emirates, 2025, pp. 600-605, doi: 10.1109/IWCMC65282.2025.11059465.

K. Acharya, I. Sharifi, M. Lad, L. Sun, and H. H. Song, “Integrating Neurosymbolic AI in Advanced Air Mobility: A Comprehensive Survey,” in *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-25)*, pp. 10362–10370, 2025, doi: 10.24963/ijcai.2025/1151

K. Acharya and H. H. Song, “A Comprehensive Review of Neuro-symbolic AI for Robustness, Uncertainty Quantification, and Intervenability,” *Arabian Journal for Science and Engineering*, vol. 51, pp. 35–67, 2026, doi: 10.1007/s13369-025-10887-3.

K. Acharya, M. Lad, L. Sun and H. Song, "A Data-Driven Approach to Enhancing Gravity Models for Trip Demand Prediction," *2025 IEEE Conference on Artificial Intelligence (CAI)*, Santa Clara, CA, USA, 2025, pp. 815-820, doi: 10.1109/CAI64502.2025.00145.