AS² - Attention-Based Soft Answer Sets: An End-to-End Differentiable Neuro-Soft-Symbolic Reasoning Architecture

Wael Abd-Almageed • Clemson University • wabdalm@clemson.edu

zeteo.sites.clemson.edu

NeurIPS 2026

AS² code
huggingface.co/wabdalm/AS2

THE PROBLEM

Disconnected reasoning and perception. Constraint satisfaction.

How can the rules teach *both* perception and reasoning?

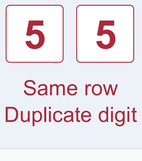

Discrete pipeline — Perception → hard symbols → external solver
Exact training inference can be costly; early choices lose uncertainty.

AS² — **Soft beliefs + declarative rules → differentiable feedback**
End-to-end learning without an external solver.

Rules become differentiable computation

Sudoku instantiation

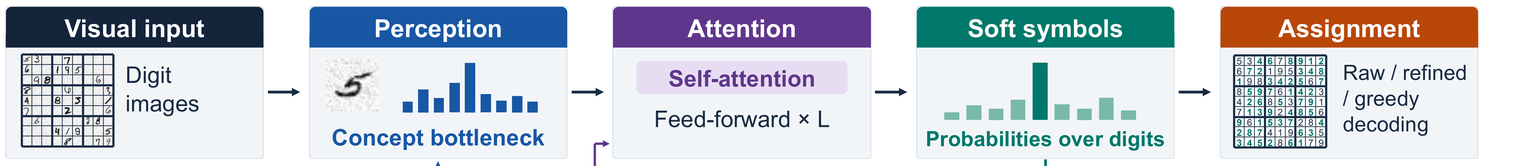

Structure, not sequence index

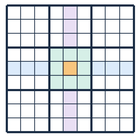

row + column + box membership embeddings
The constraints define structure.

Predict ≈ what the rules permit

$L = \sum_G \|p_G - T_{P,G}(p)\|^2$

T suppresses competing claims within a group.
Zero residual ↔ valid one-hot Sudoku groups.



Beyond a single puzzle

Visual Sudoku	Mutual exclusivity
CLEVR-Hans	Existential conjunction
MNIST Addition	Arithmetic sums

Why this matters

Materials science · **Drug discovery**

Visual Sudoku

Read handwritten clues. Infer every blank.

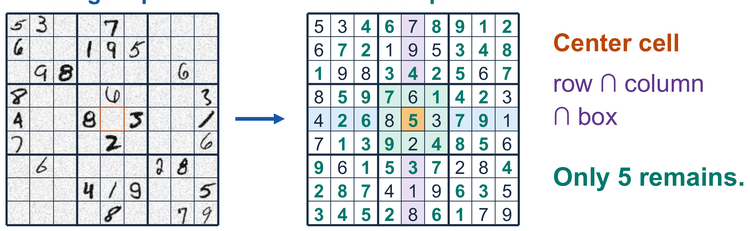

Matched illustration, not a measured prediction.

95.60%	98.70%	100%
Raw argmax	10 soft steps	Greedy decode

Board accuracy • 1,000 test boards • ≈45 clues/board
Greedy decoding is discrete, with no backtracking.

CLEVR-Hans: true rules or visual shortcuts?

Classify 3-D scenes of 3–10 objects into 3 or 7 rule-defined classes.

Class C_0 requires a large cube AND a large cylinder. Color is irrelevant.

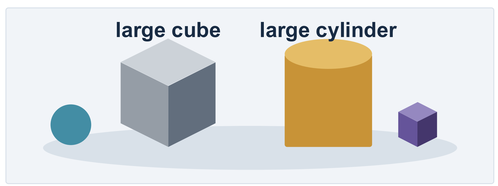

TRAIN / VAL • rule satisfied

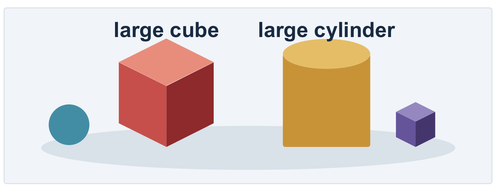

TEST • same rule, different color

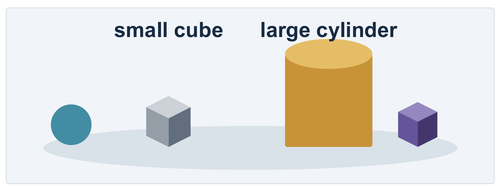

NEGATIVE • required large cube absent

Train/val tie large cubes to gray; test removes that shortcut. Scenes above are schematic.

Hans3 • preserving the rules helps



Hans7 • a harder generalization test



Test scene accuracy • 3-seed means • AS² uses Crop perception

Sources: camera-ready paper §3–4 & Appendix A • CLEVR-Hans: Stammer et al., CVPR 2021 • MNIST imagery: LeCun et al.