

Task : Solve constraint-based logic puzzles by deriving a consistent sequence of atomic deductions.

- Summary:**
- COT & Logic-LM collapse on constraint-dense logic puzzles vs. SymStep+G reaches **97-100%** with same model and no additional training
 - LLMs don't just make wrong deductions, but they don't know **what to deduce next**

6 domains

119 instances

\$0.0013 per puzzle

1 Why do LLMs fail at multi-step logic?

Errors accumulation

Step 1

Step 2

Step 3

Step 4

Step 5

wrong answer

One bad step, silently accepted, poisons every step after it. No mechanism catches it.

Deductive deadlock

LLM

“re-checks clue 3 ... again”

“no obvious next move”

every step is correct — yet zero progress

Without direction, the model cycles in place. It needs direction.

2 The fix: one claim at a time — verified, then guided

Problem (natural language)

LLM M

MRV Hint (SymStep+G)
 $(p^*, a^*) = \arg \min_{|\Pi| > 1} |\Pi(p, a)|$

Accepted
update Π , cascade

Contradiction
exact error message

+G adds an MRV hint: “Blake's drink $\in$ {Coffee, Juice}”

DEDUCE: p, a, v

Symbolic constraint propagation
arc-consistency cascade

accept, next step

revise

MRV hint

No unverified fact enters the context.
~60% of assignments derive automatically — zero extra LLM calls.

SymStep in action

LLM DEDUCE: Alice, color, NOT Blue

OK. Hint: Alice.color $\in$ {Red, Green}

LLM DEDUCE: Bob, color, Green

OK. Cascade $\rightarrow$ Alice.color = Red (free)

LLM DEDUCE: Alice, pet, Dog

CONTRADICTION — Alice.pet is Cat. Revise.

DEDUCE: Alice, pet, Cat

Solved — all cells determined.

the whole protocol — regex-parseable, no formal logic, no training

3 Results: baselines collapse, SymStep+G doesn't

CoT

SymStep

SymStep+G

LGP-14

ZebraLogicBench*

AR-LSAT

0%

86%

100%

0%

80%

97%

87%

100%

100%

* 35 unseen puzzles from a published 1,000-puzzle benchmark · all runs on Claude Haiku

Generalizes across 6 reasoning domains!

Logic puzzles · scheduling · analytical reasoning · arithmetic · finance

	LGP-14	ZLB	SP-6	LSAT	MWP-8	FIN-6
Direct	14	0	0	100	88	67
CoT	0	0	0	87	100	83
SymStep	86	80	83	100	100	100
SymStep+G	100	97	67	100	100	100

4 Interesting: direction beats correction

The bottleneck is **direction**, not correctness.

Pointed at the most constrained variable, the LLM makes only correct deductions. MRV guidance prevents LLM from entering dead-end.

CoT (neither)

Verification only

Guidance only

Both (SymStep+G)

0%

72%

100%

100%

LGP-6 · 3 independent runs · zero variance

One claim, one check
deterministic verification; never LLM-judged

Direction > correction
MRV hint removes deadlocks before they happen

No logic training needed
regex-parseable protocol