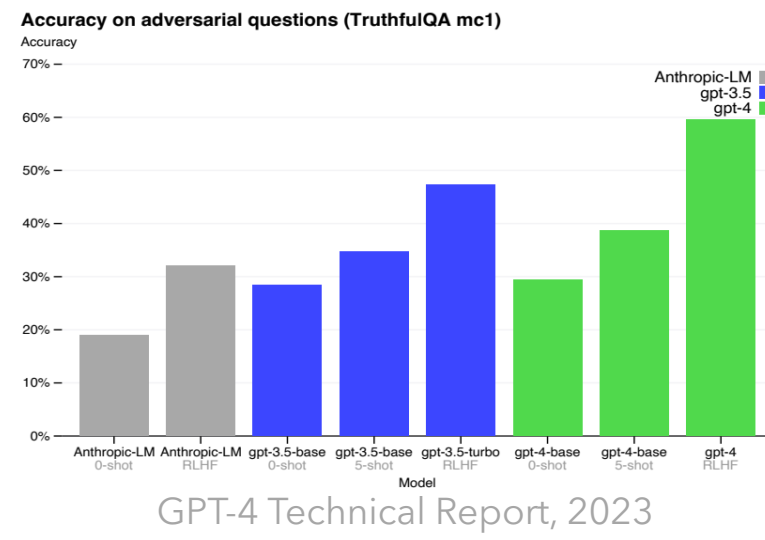


1. LLMs Hallucinate

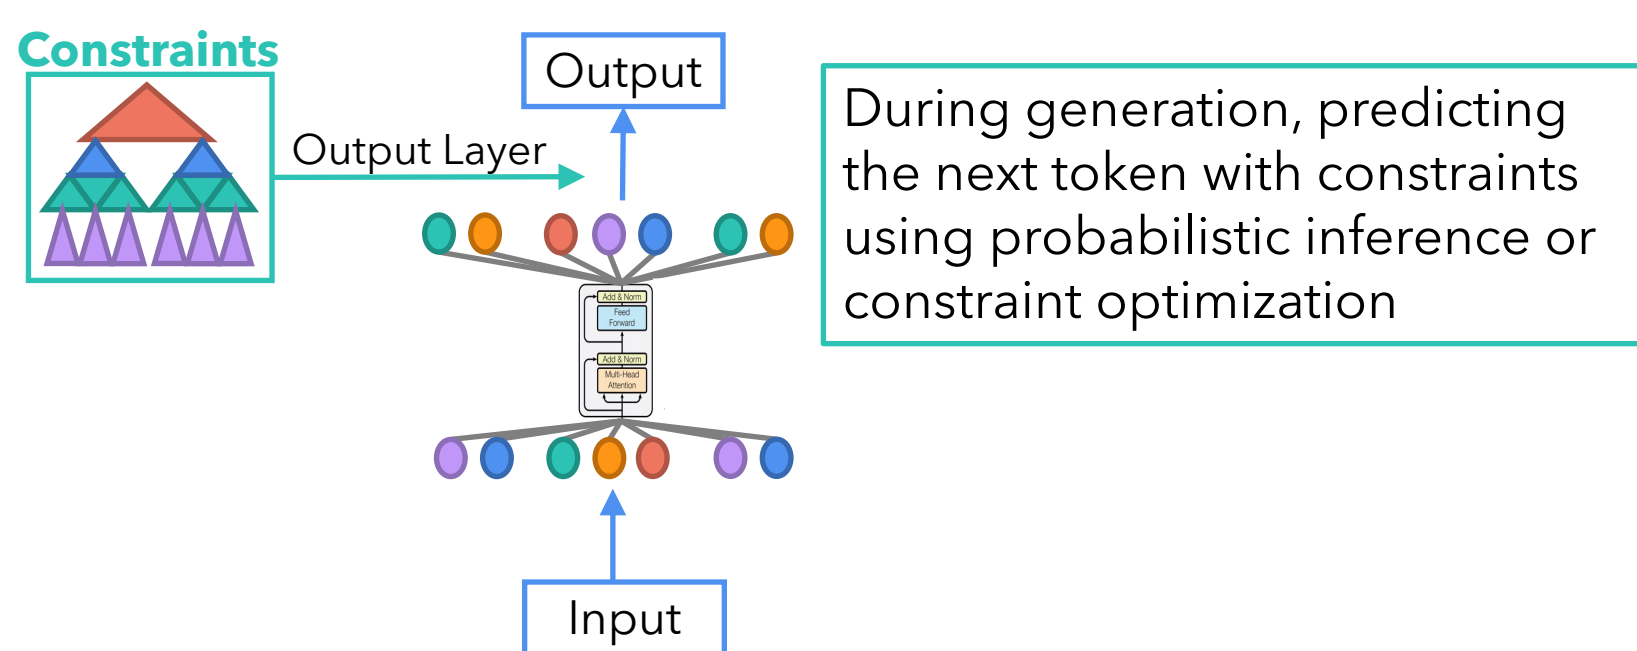
- Produces inconsistent information and inaccurate results
- LLMs *over-generalize patterns and relationships* from pretraining data
 - Causes inconsistent and inaccurate results

GPT-4: 40% hallucination rate




2. Current Efforts for Reducing Inconsistencies

- Ensuring the output of LLMs is consistent with domain constraints

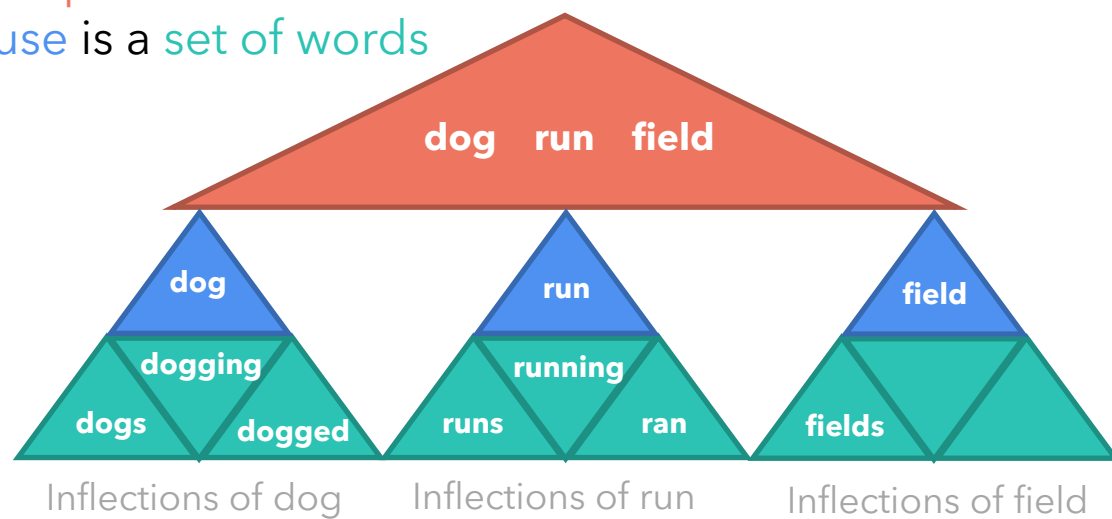


- CommonGen constraints = contains keywords or their inflections



3. Constraints in CommonGen

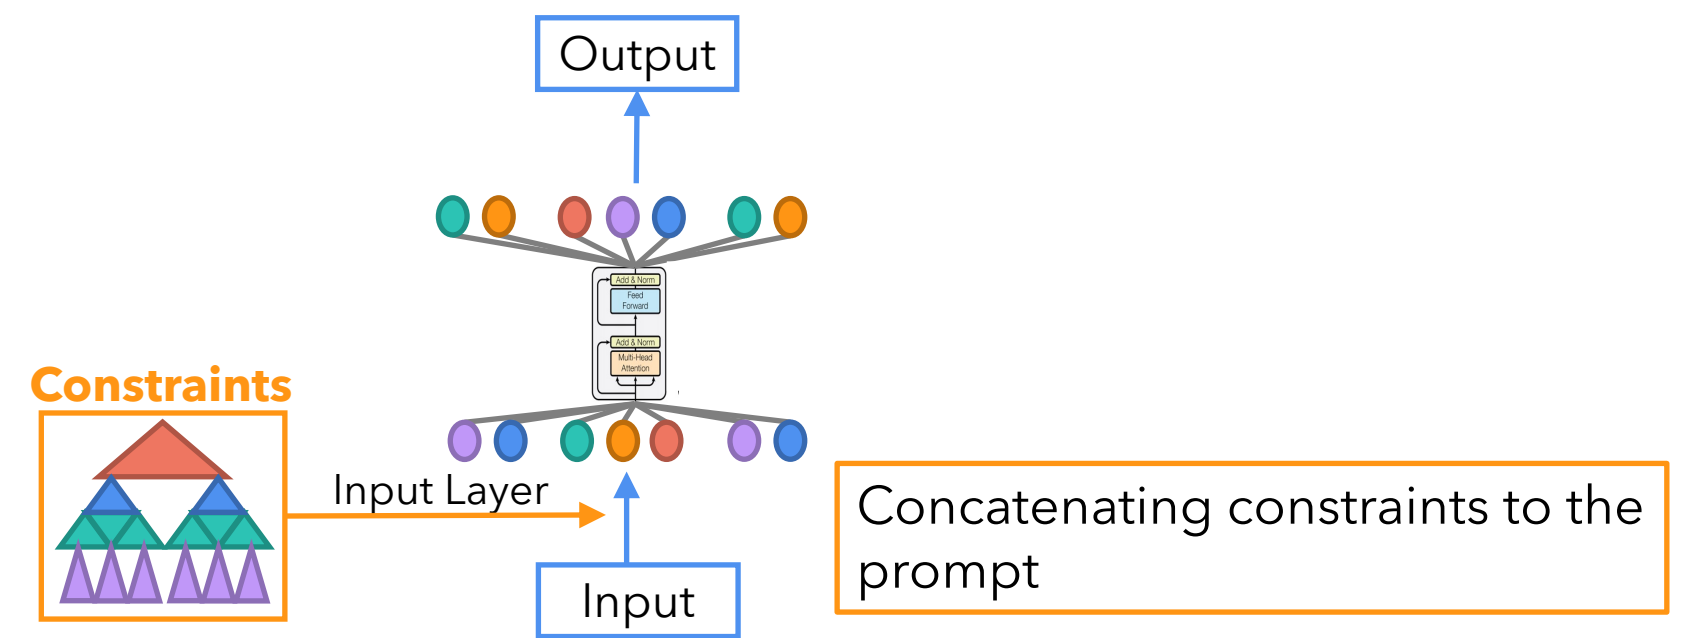
- **Given:** Concept word set = [dog, field, run]
 - **Goal:** Write a sentence using all concepts from the set
- Representing Constraints in Conjunctive Normal Form (CNF)
- Each **concept** is a **clause**
 - Each **clause** is a **set of words**



(dog V dogs V dogging V dogged) \wedge (run V runs V running V ran) \wedge (field V fields)

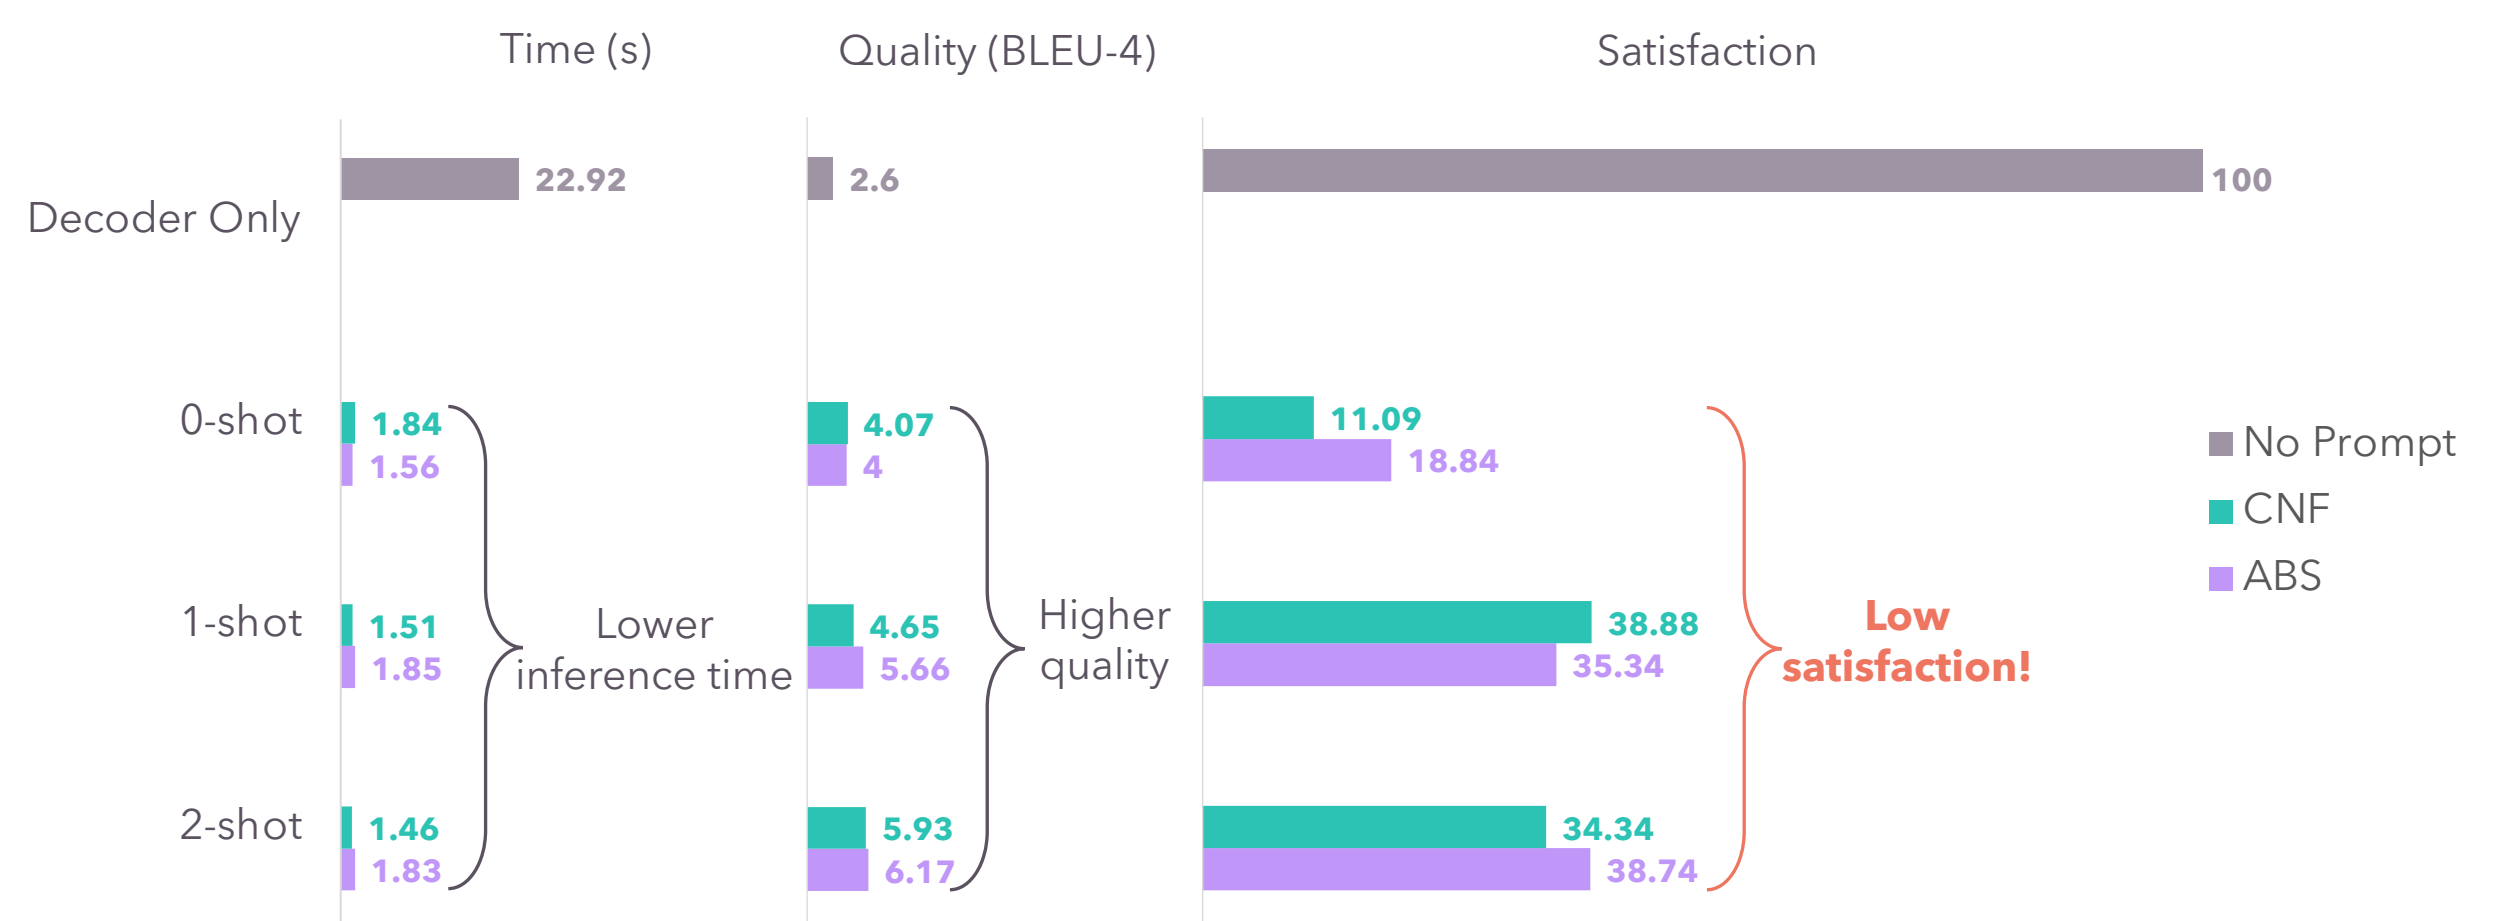
4. Can Prompting with Constraints Help Mitigate Decoder Limitations?

- Adding constraints to the input of LLMs

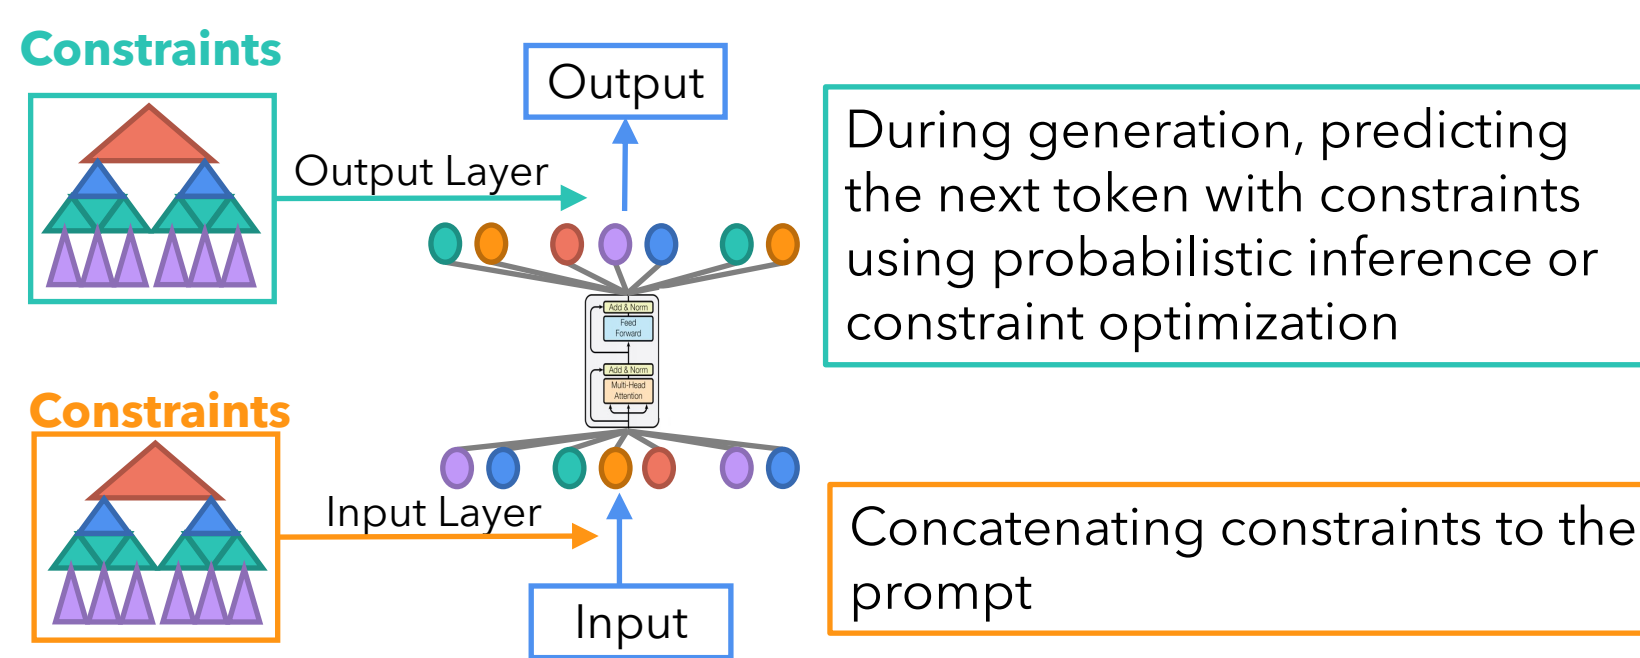


CNF **Conjunctive Normal Form**
CNF constraint as text
Write a sentence using the words (dog or dogs or ...) and (run or running or ...) and (field or fields)

ABS **Abstract**
Abstract instance of constraint
Given a set of words x, write a sentence using all words in x or inflections of x



5. Empirical Study



Metrics

1. Generation Quality (perfect score = 100.0)
 - Automated metrics: **ROUGE-L**, **BLEU-4**, **CIDEr**, **SPICE**
2. Satisfaction (perfect score = 100.0)
 - If all CNF clauses in constraint C are met: $\bigwedge_{i=1}^{|C|} C_i$
3. Time
 - Inference time (seconds)

Decoder Strategies

NL **NeuroLogic**

- Soft constraint decoder
- Beam-based look ahead strategy

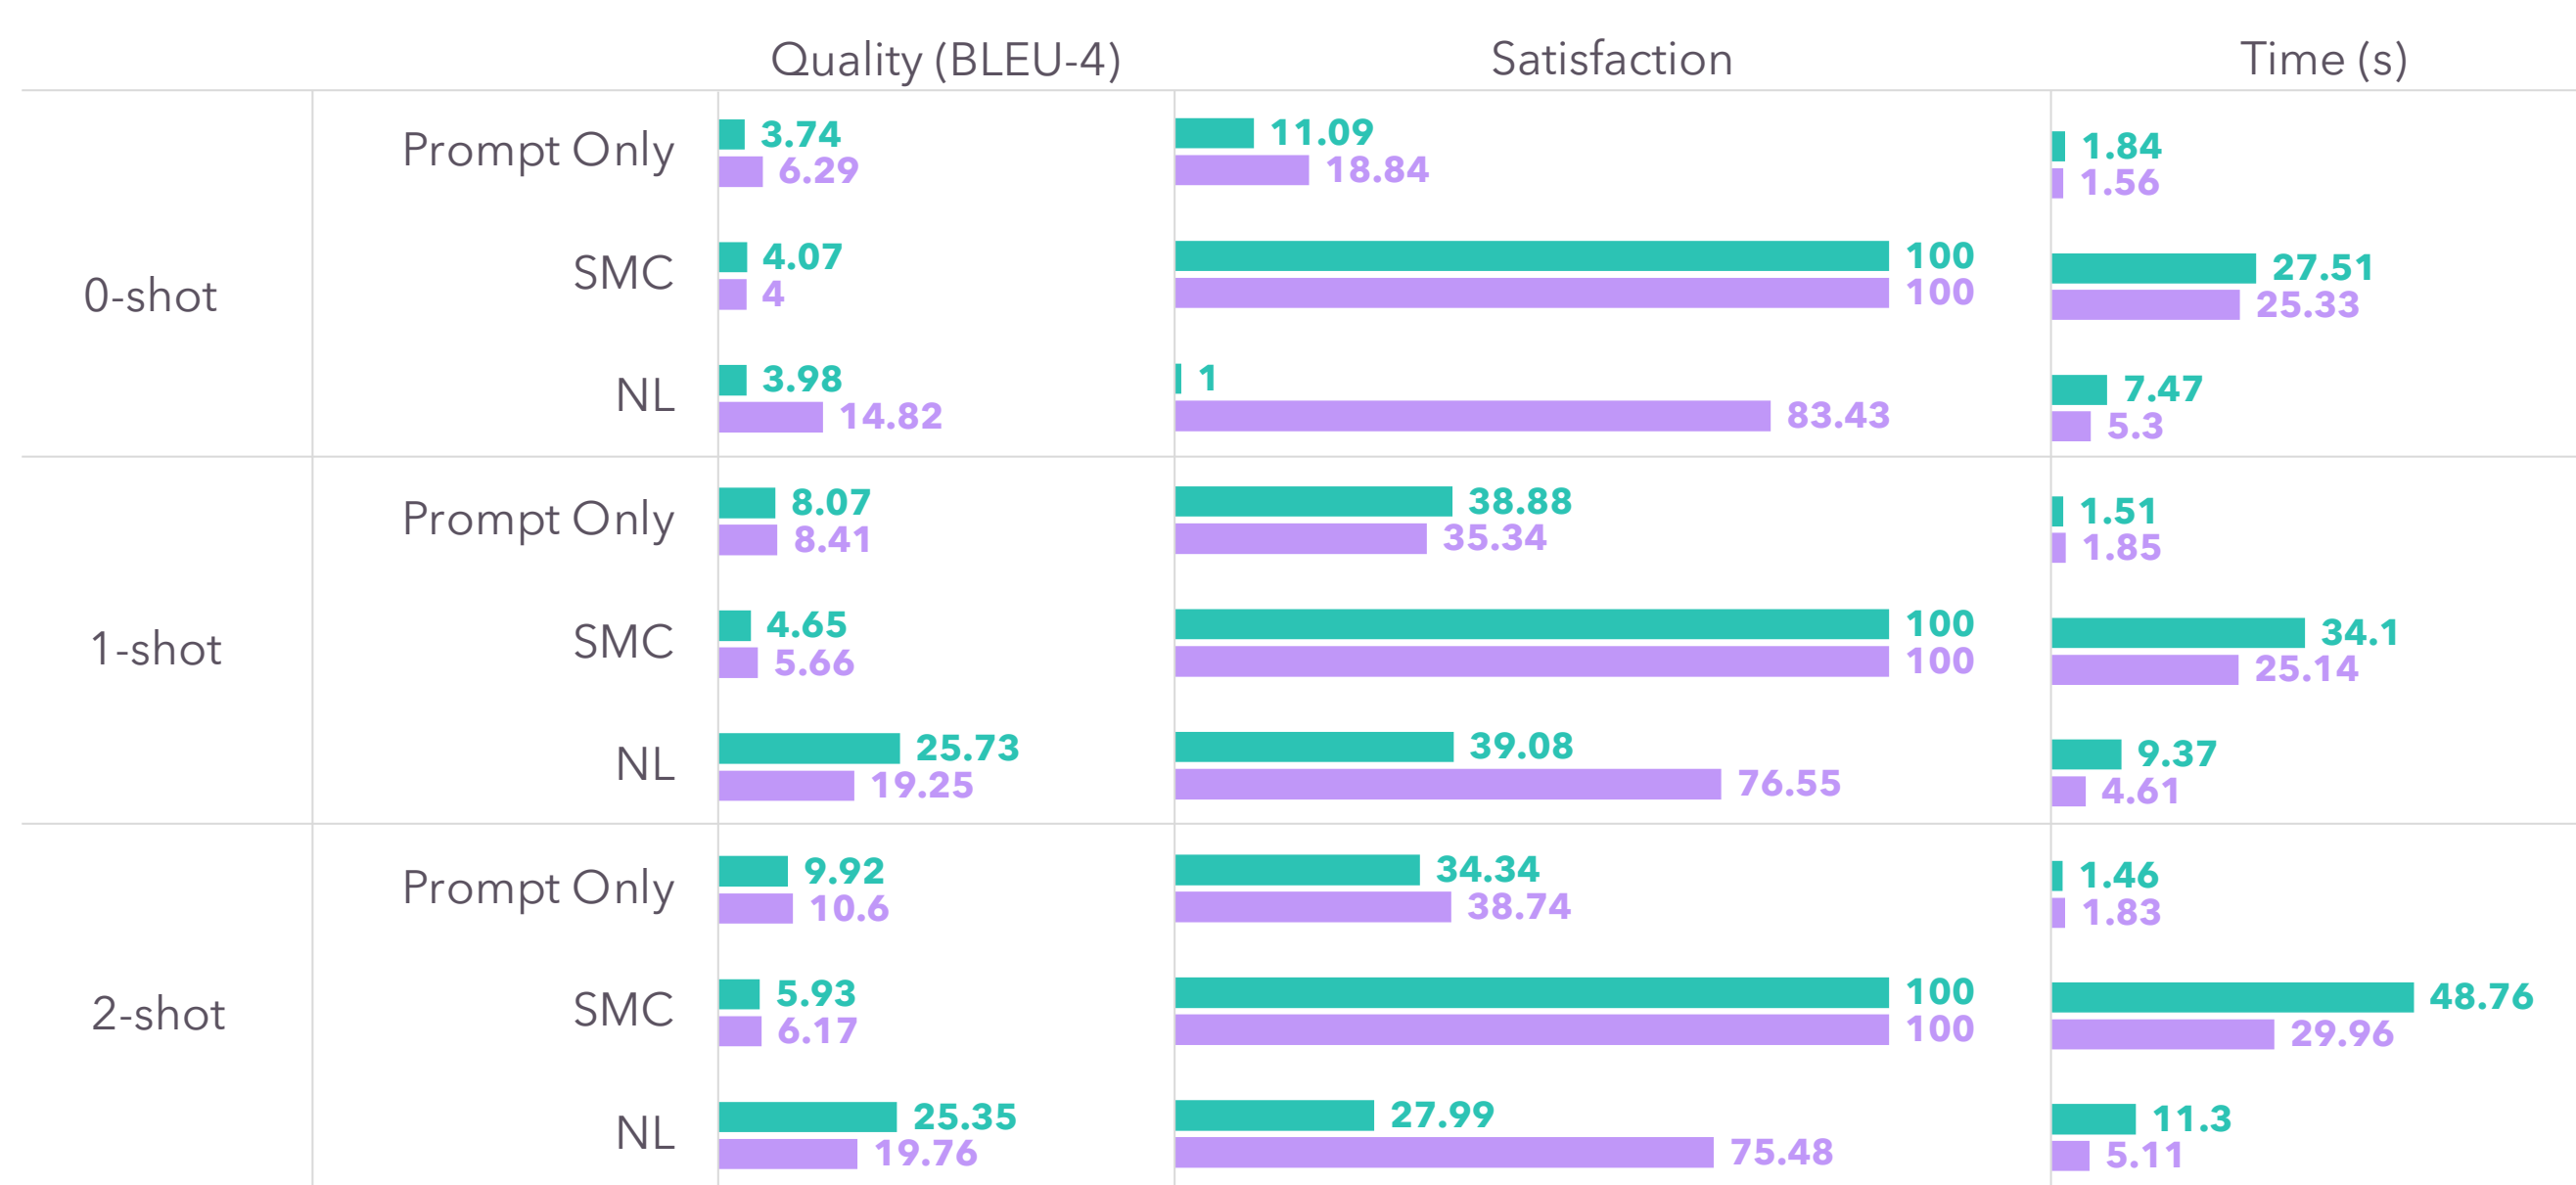
NeuroLogic Decoding: (Un)supervised Neural Text Generation with Predicate Logic Constraints, 2021

SMC **Sequential Monte Carlo**


- Hard constraint decoder
- Masked vocabulary strategy

Sequential Monte Carlo Steering of Large Language Models using Probabilistic Programs, 2023


6. Results & Key Findings




Prompting reduces the search space for decoders using search strategies



Irrelevant context hurts satisfaction



Tradeoff between constraint satisfaction and inference time + quality



Soft Constraint Decoders

- **Benefit:** Balances tradeoff
- **Cost:** \downarrow satisfaction, reliance on high quality distributions

Hard Constraint Decoders

- **Benefit:** 100% satisfaction
- **Cost:** \downarrow quality, \uparrow time