

Aligning Code-Generating Models with Formal Specifications

UC San Diego

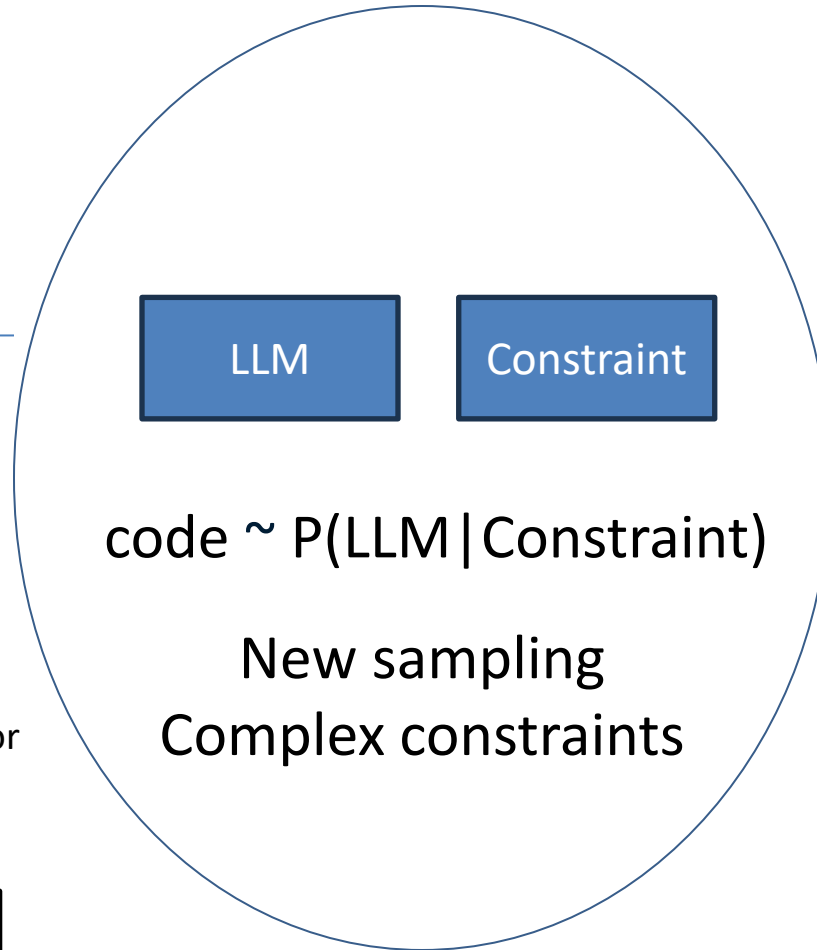
Challenge:

- LLMs are not guaranteed to produce code that satisfies a user given specification.

Solution:

- Constrained Decoding can force LLM output to satisfy specification
- New sampling approaches for constrained decoding
- Constrained decoding for complex specs (attribute grammars)

Project info (2422214 UCSD, D'Antoni, Berg-Kirkpatrick, Polikarpova)



Scientific Impact:

- New left-to-right static analysis algorithms that operate at the string level
- The work will allow the LLM community to understand how formal specs can be pushed to left-to-right (programming) language generation

Broader Impact and Broader Participation:

- LLMs for code are now adopted by millions of people and need to be aligned with specs
- We plan to build tools that integrate with ANY LLM through known APIs
- Already covering this topic in Graduate Courses

