Robust probabilistic inference engines for autonomous agents

Stefano Ermon

Stanford University

Problem Solving in AI

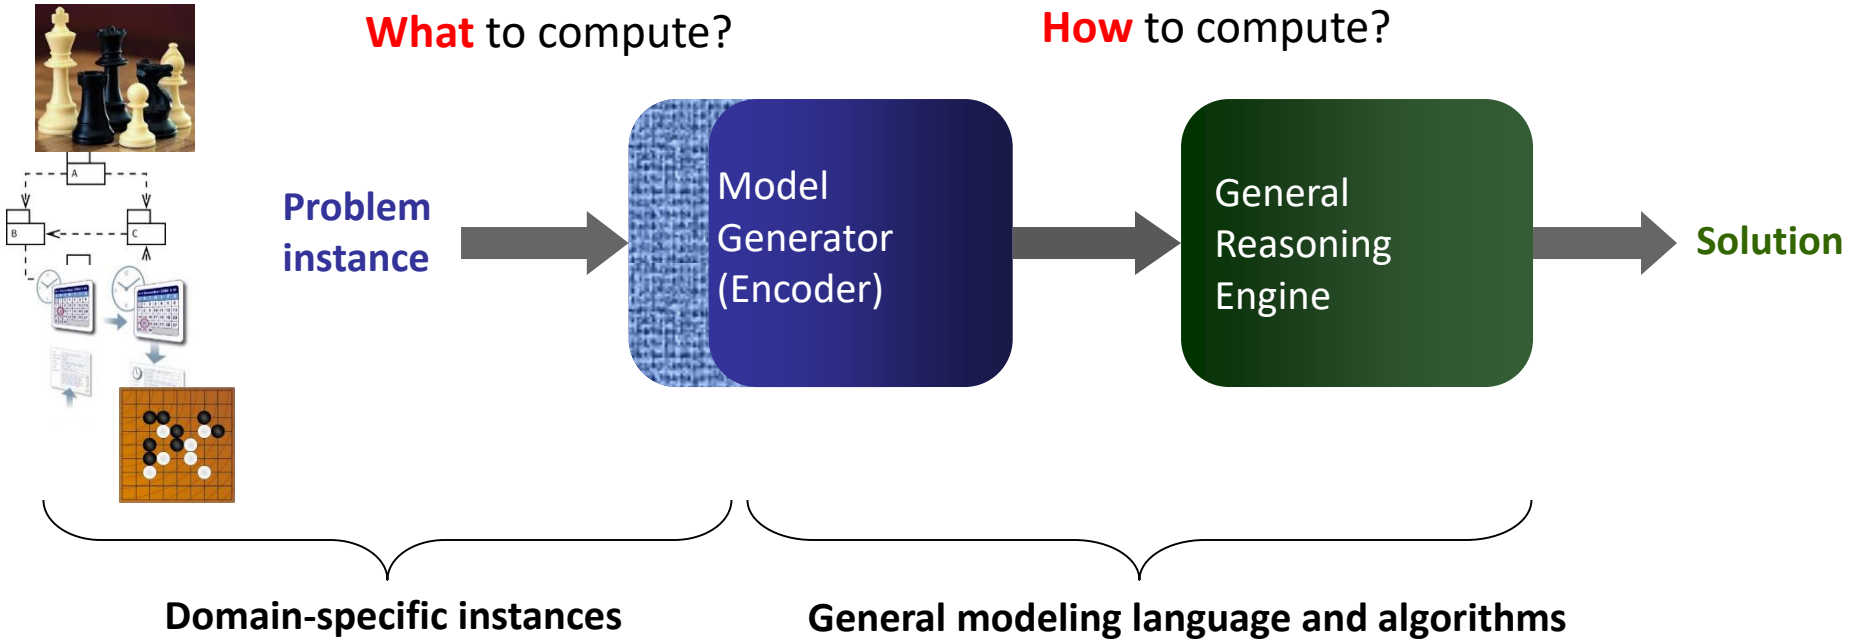

Problem solving in AI:

Separate modeling from algorithms

Problem Solving in AI

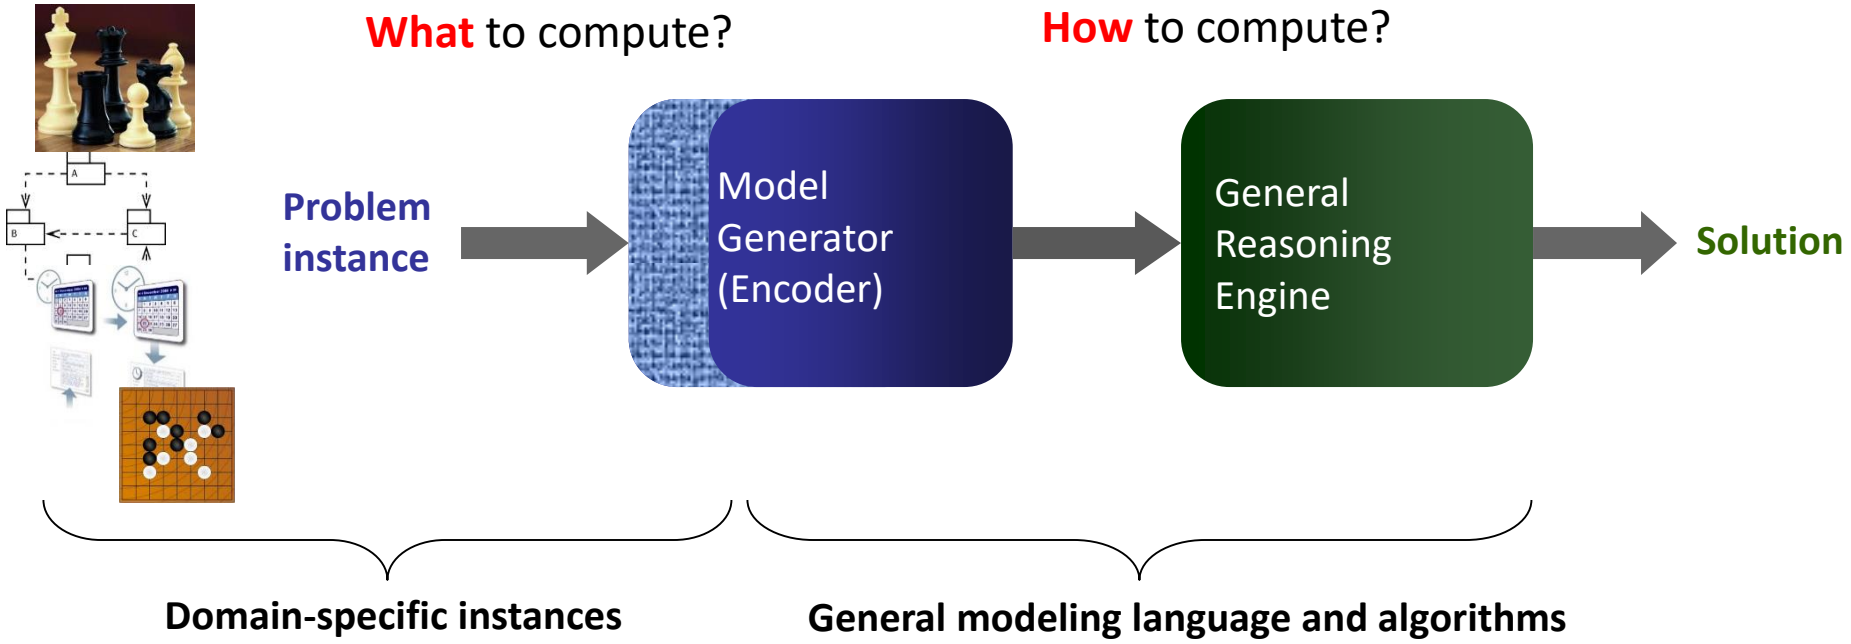

Safety and **reliability** require:

1. **precise models**
2. **accurate reasoning techniques**

Challenges in reasoning about complex systems

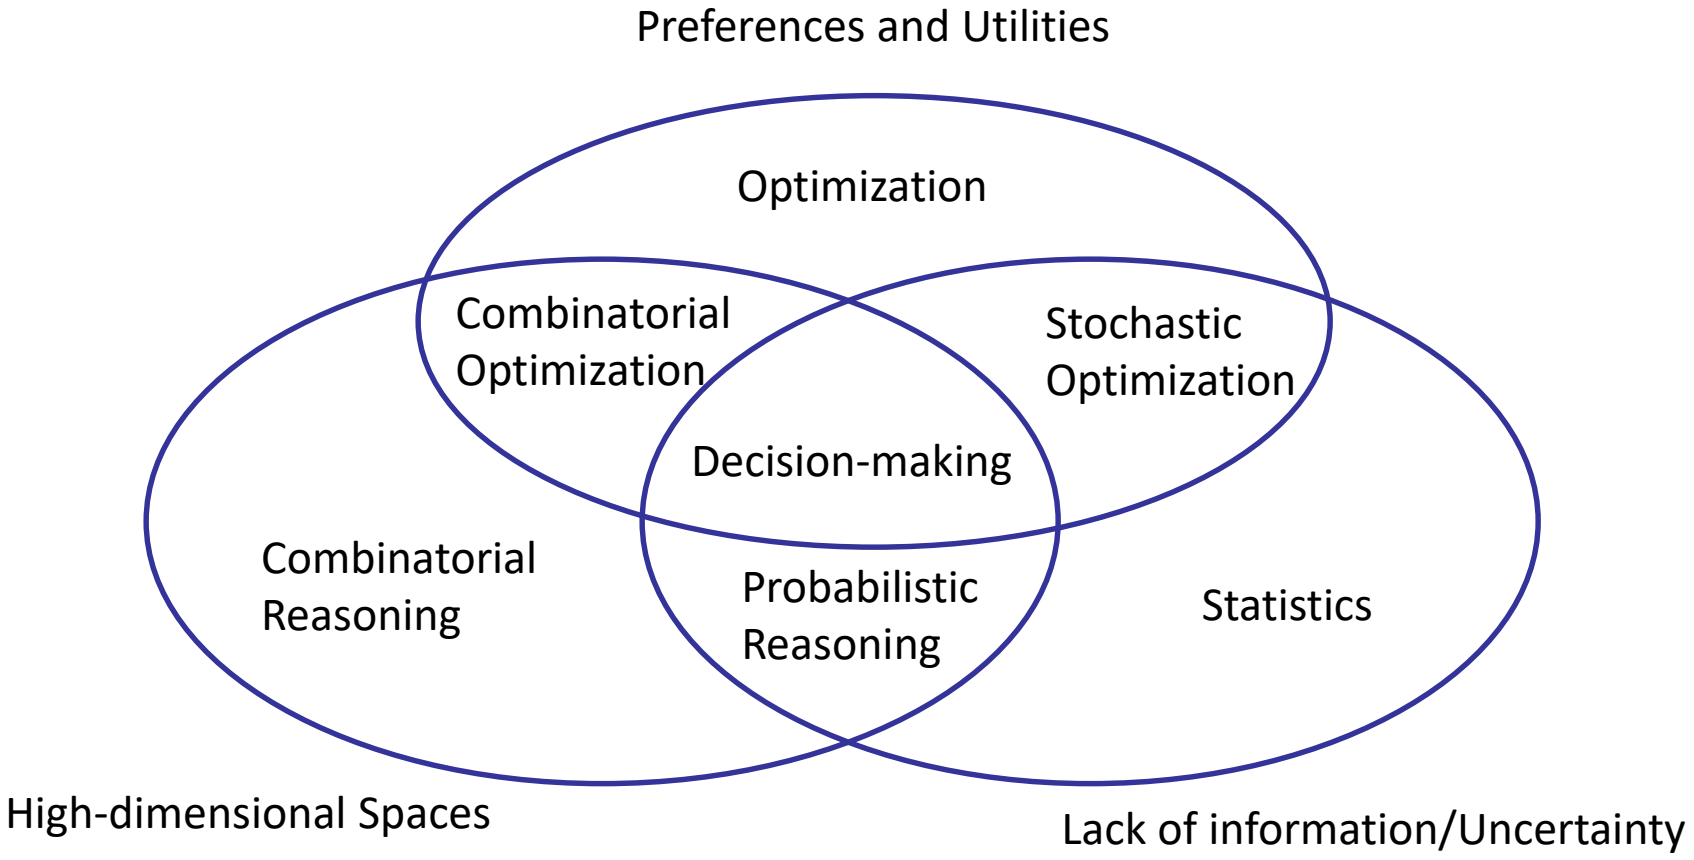

Challenges in reasoning about complex systems

Theorem provers, logical reasoning, SAT solvers, soundness, certificates of optimality, …

hardware verification

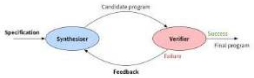

Program synthesis

Preferences and Utilities

Optimization

Combinatorial Optimization

Stochastic Optimization

Decision-making

?

Combinatorial Reasoning

Probabilistic Reasoning

Statistics

High-dimensional Spaces

Lack of information/Uncertainty

Proposal and recent results

Proposal: use *combinatorial reasoning/optimization* techniques (logic, verification, synthesis) for *probabilistic reasoning* tasks (machine learning)

- Algorithms that can provide **certificates/proofs of accuracy**
- Handle **extreme (unsafe) events**
- Can support deterministic + probabilistic dependencies

- Some recent results:
  - Satisfiability Modulo Theory solvers for statistical hypothesis testing (Zhao et al., AAAI-2016)
  - Integer Linear Programming for sampling (Kim et al., AAAI-2016)
  - Integer Linear Programming and SAT for decision making under uncertainty (Xue et al., NIPS-2016)
  - Variational methods with guarantees (Achim et al., AISTATS-2016)

Thanks!