Mixed integer linear programming

Most popular tool for solving combinatorial (& nonconvex) problems

Mixed integer linear program (MILP)

maximize $\boldsymbol{c}^T\boldsymbol{x}$

subject to $A\boldsymbol{x} \leq \boldsymbol{b}$

$x_i \in \mathbb{Z}$ for all $i \in I$

Routing

Scheduling

Planning

- Solver performance heavily influenced by parameter configurations
- Configuration is challenging without expertise and computational resources

Conventional automated configuration pipeline

1. Gather **training set** of MILPs (historical problems from application domain)

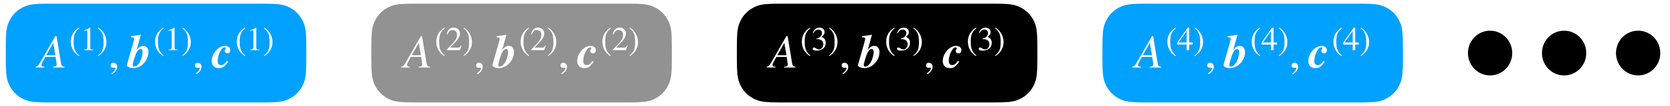

2. Find configuration leading to **good runtime** on average over training set

3. Hope for good runtime on **future problems** from the same application

Key scalability challenge:

Evaluating a configuration's average runtime over the training set…

Requires **solving** every MILP in the training set using that configuration!

🔭 Overview of contributions

Can we use LLMs to configure MILP solvers with minimal training data?

- New **LLM-based framework** to configure cutting plane separators
- Finds high-performing configuration by **solving only a few MILPs**
- **Ensembling strategy** to build portfolio of high-performing configurations
- Requires **no custom solver interface**
- Competitive with existing configuration approaches
  but only requires a **fraction of the training data and computation time**

Background: Branch-and-cut algorithm

Uses guidance from LP relaxations to guide search

MILP

maximize $\boldsymbol{c}^T\boldsymbol{x}$

subject to $A\boldsymbol{x} \leq \boldsymbol{b}$

$x_i \in \mathbb{Z}$ for all $i \in I$

Linear programming (LP) relaxation

maximize $\boldsymbol{c}^T\boldsymbol{x}$

subject to $A\boldsymbol{x} \leq \boldsymbol{b}$

~~$x_i \in \mathbb{Z}$ for all $i \in I$~~

Cutting planes (CPs) are additional constraints that:

- Separate LP optimal solution
- Don't separate any integer point

Many different families of CPs; which to use when?

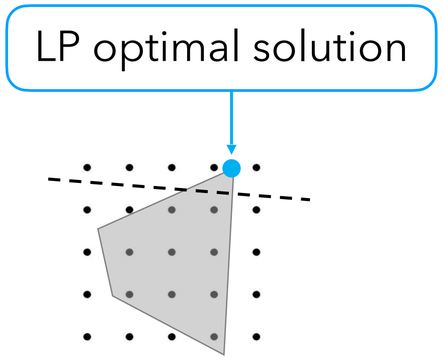

LLM configuration method

Prompt technique

Documentation: You're an optimization expert tasked with …

CP descriptions: Specify which of the following CPs to turn on:

- Clique cuts: Based on cliques in a conflict graph…
- …
- 2 means use aggressively, 1 means use normally, …

Problem description: Stable set problem: find a subset of vertices…

Output: Name: CliqueCut, value: 1; …

Ensembling strategy

LLMs are stochastic – generate many configurations

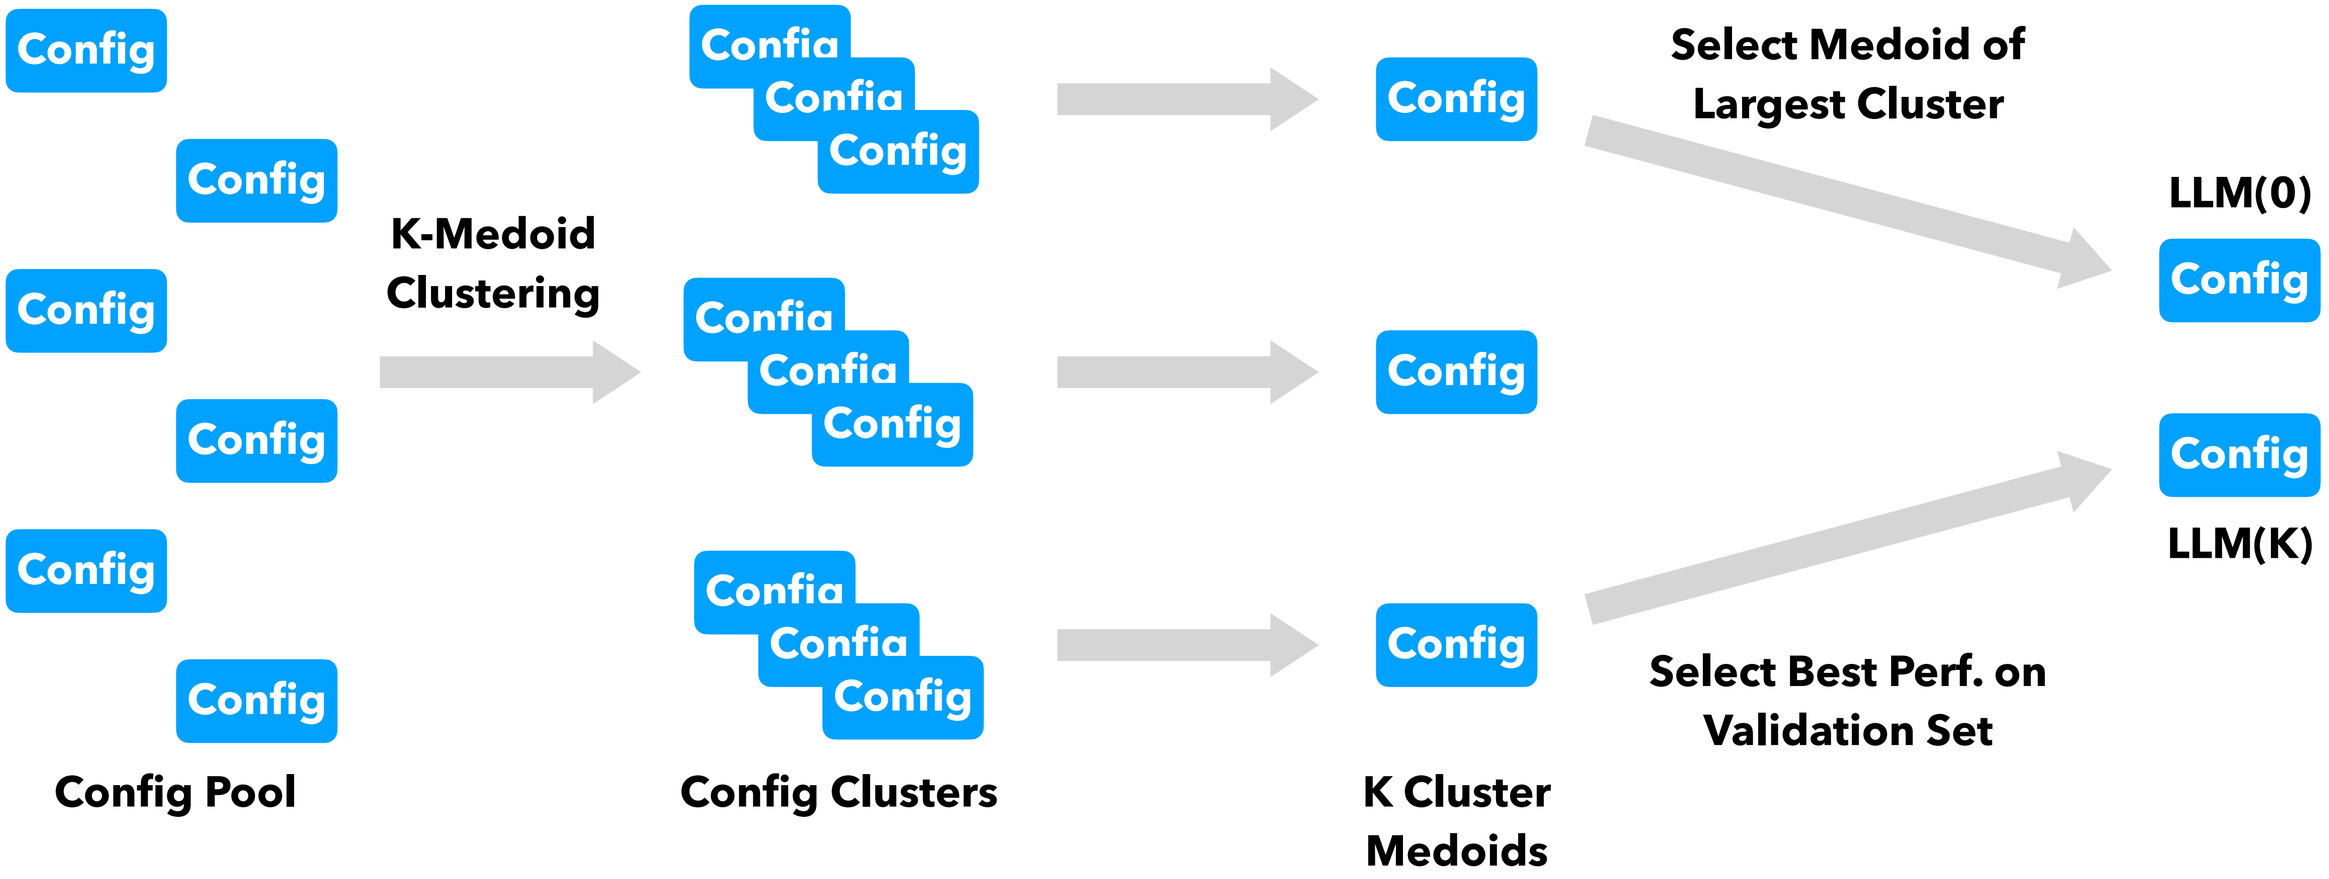

Experimental evaluation

Synthetic & real-world MILPs from Distributional MIPLIB [Huang et al., arXiv'24]

Name	# variables	# constraints
Binary packing	300	300
Capacitated facility location	100	100
Combinatorial auction	100	500
Max independent set	500	1088
Max cut	54	134
Packing	60	60
Set cover	500	250
Load balancing	64340	61000
Middle-mile consolidation network design	569	248

- 30 validation instances
- 100 test instances
  - Each solved with 3 random seeds
- SCIP: wall clock time (s)
- Gurobi: work units
- Relative improvement over default

Baselines

Pruning:

- Turns off CPs that weren't used on validation instances w/ default settings

Search(d):

- Samples d candidate configurations uniformly at random
- Applies the one with best median performance on validation set
- Instance agnostic configuration method by Li et al. [NeurIPS'24]

Computation vs. performance

LLM-based configuration is pareto optimal