

# MetaEase: Heuristic Analysis from Source Code via Symbolic-Guided Optimization

---

**Pantea Karimi<sup>1</sup>**

Siva Kesava Reddy Kakarla<sup>2</sup>, Ryan Beckett<sup>2</sup>, Santiago Segarra<sup>3</sup>,  
Pooria Namyar<sup>2</sup>, Mohammad Alizadeh<sup>1</sup>, Behnaz Arzani<sup>2</sup>



# Heuristics Are Everywhere, Yet Can Break Unexpectedly



Approximate  
Solution  
(Heuristic)

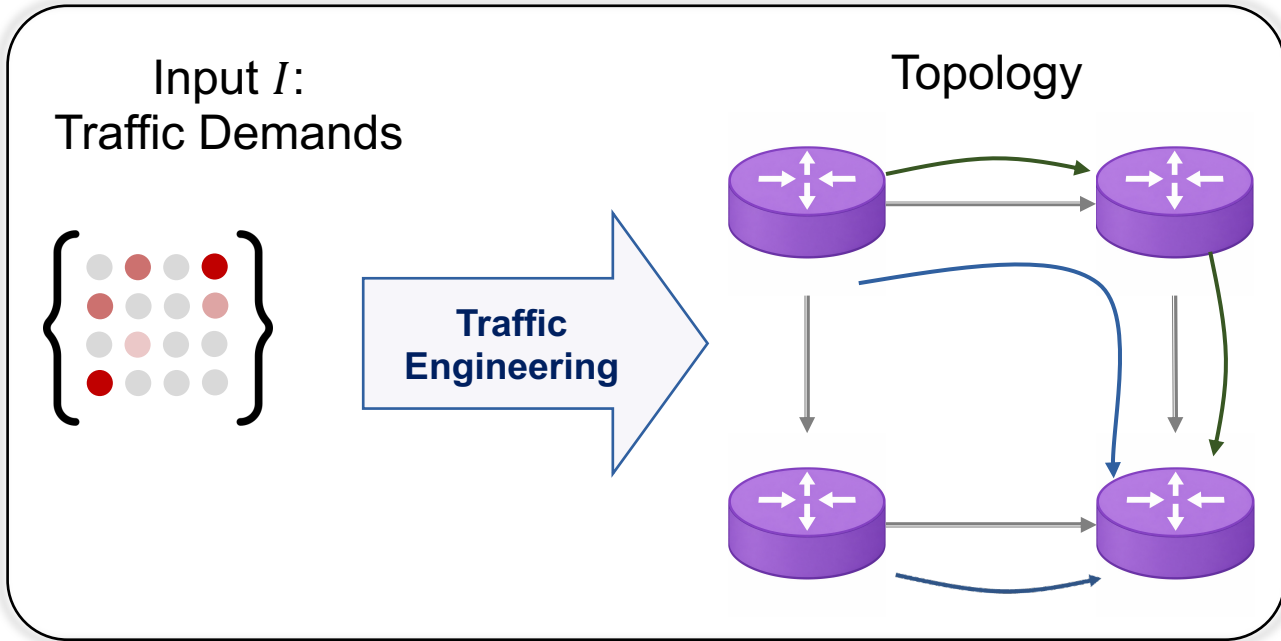
Example of heuristics:  
Demand Pinning in Traffic Engineering  
First-Fit-Decreasing in VM Placement



Risk of underperformance and inefficiency



# Traffic Engineering Example



Performance objective:  
Maximize total routed traffic

Performance Gap

*Opt(I)* = Optimal Performance scales input *I*

Heuristic

Demand Pinning (DP)

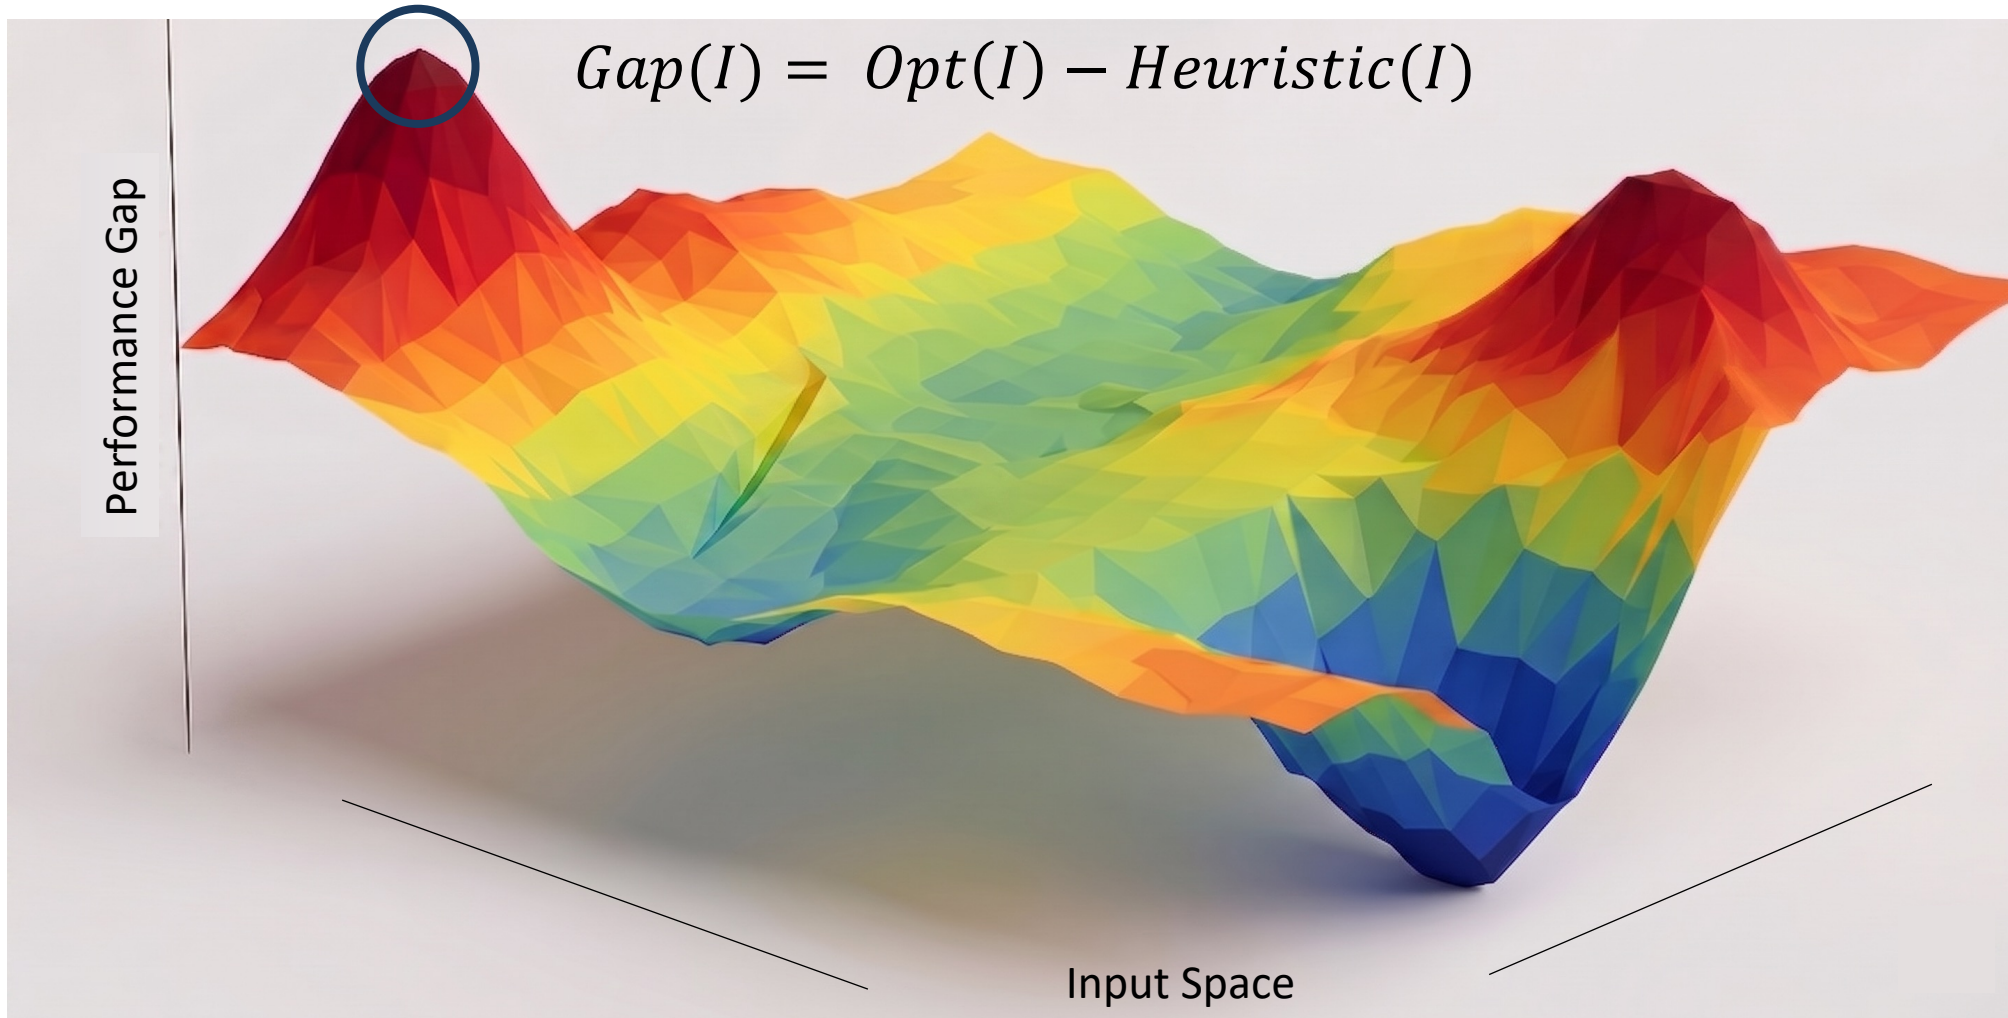
*Heuristic(I)* → DP's Performance on input *I*

*If demand ≤ P: pin to shortest path  
else: route optimally*

$$Gap(I) = Opt(I) - Heuristic(I)$$

**Demand Pinning *may* have 30% performance gap for some topologies!**

# Heuristic Analysis: Find Worst-Case Performance Gap



# State-of-the-Art Solutions for Heuristic Analysis

Example: MetaOpt [NSDI '24], FPerf[NSDI '23], Xplain [Hotnets'24]

Encode any new heuristic into a **mathematical format** (e.g., optimization, SAT)

The Heuristic (Demand Pinning) *If demand  $I_k \leq P$ : use shortest path; else: route optimally*

Demand Pinning

Outer Problem (Leader)

Hard to Use or Domain Specific

Inner Problem 1 — OPT (Max-Flow LP):

$$\begin{aligned} \text{OPT}(I) = \max_{f'} & \sum_k \sum_{p \in \mathcal{P}_k} f'_{k,p} \\ \text{s.t.} & \sum_k \sum_{\substack{p \in \mathcal{P}_k \\ e \in p}} f'_{k,p} \leq \text{Cap}_e \quad \forall e \\ & \sum_{p \in \mathcal{P}_k} f'_{k,p} \leq I_k \quad \forall k \\ & f'_{k,p} \geq 0 \end{aligned}$$


(Conditional + Max-Flow):

$$\begin{aligned} b_k & \in \{0, 1\} && \text{(binary indicator)} \\ b_k & \leq P + M(1 - b_k) \quad \forall k && \text{(big-}M\text{)} \\ b_k & \geq P + \epsilon - M b_k \quad \forall k && \text{(big-}M\text{)} \\ f_{k,\hat{p}_k} & = I_k \cdot b_k && \text{(pin if small)} \\ f_{k,p} & \leq M(1 - b_k) \quad \forall p \neq \hat{p}_k \\ \sum_k \sum_{\substack{p \in \mathcal{P}_k \\ e \in p}} f_{k,p} & \leq \text{Cap}_e \quad \forall e \\ \sum_{p \in \mathcal{P}_k} f_{k,p} & \leq I_k \quad \forall k \end{aligned}$$

Then rewrite to single-level via KKT:


- + dual variables  $\lambda_e, \mu_k$  for each constraint
- + stationarity:  $\frac{\partial \mathcal{L}}{\partial f_{k,p}} = 0 \quad \forall k, p$
- + complementary slackness:  $\lambda_e \cdot g_e = 0$
- + nonlinear terms requiring SOS / big- $M$

2 lines of logic  $\rightarrow$  bi-level optimization + binary variables + big- $M$ , ... **~1-2 days, domain expert.**



# MetaEase: The First Heuristic Analyzer from Code

Just bring *heuristic code*  
No modeling or encoding



```
for demand in demands:  
    if demand <= P:  
        # Pin to shortest path  
        allocate_shortest_path(demand)  
    else:  
        # Route remaining demands optimally  
        optimization_set.add(demand)  
solve_optimization(optimization_set)
```

*Heuristic Code*

*Opt* Formulation


Fix, usually a well-known optimization

# Key Challenges

*Gap(I)* sampling is expensive



Avoid bad local Maxima



# Challenge 1: Sampling Gap Function Is Expensive

$$Gap(I) = Opt(I) - Heuristic(I)$$

No encoding or model is available

*Opt* is expensive to run

$$Opt(I_1) - Heuristic(I_1)$$

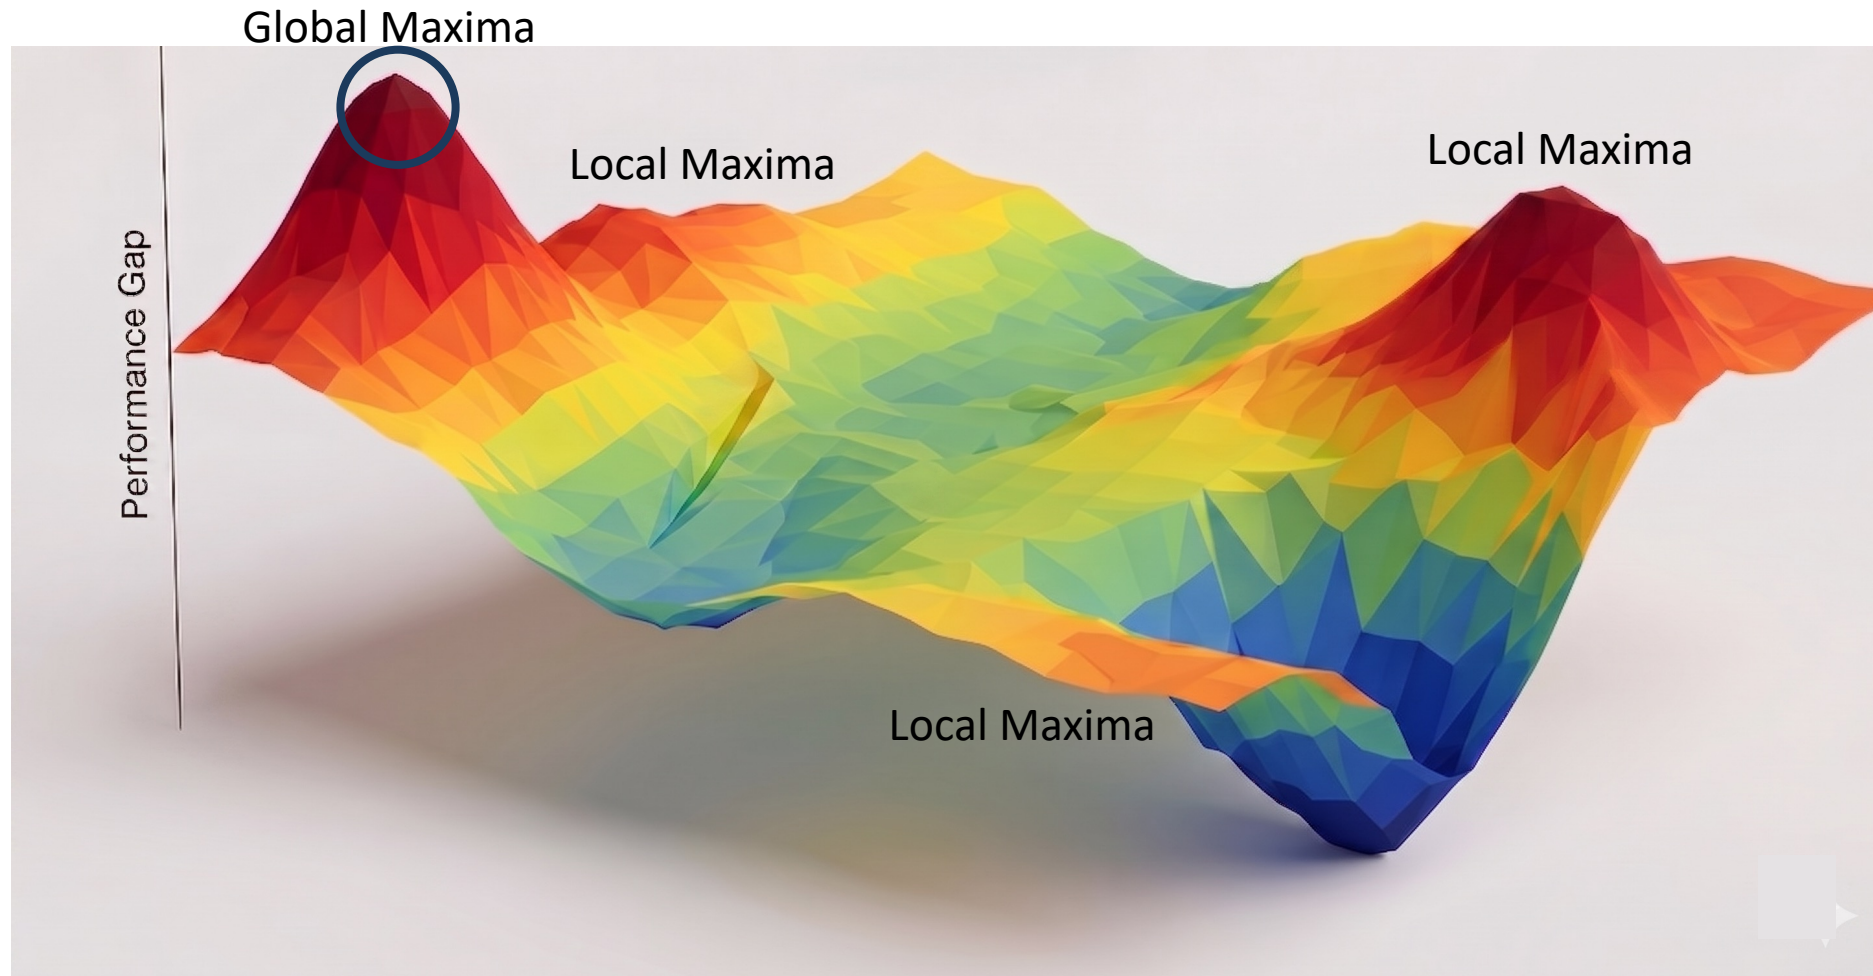
$I_1$

```
for demand in demands:  
    if demand <= P:  
        # Pin to shortest path  
        allocate_shortest_path(demand)  
    else:  
        # Route remaining demands optimally  
        optimization_set.add(demand)  
solve_optimization(optimization_set)
```



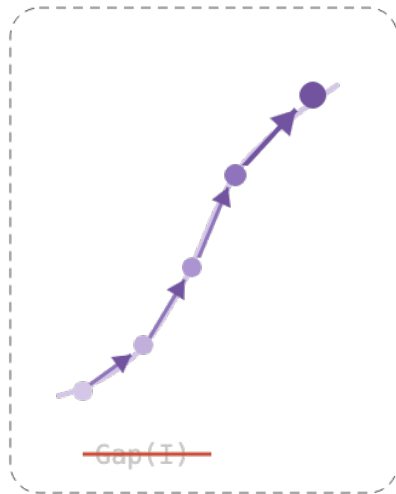
Efficient progress without repeated  $Gap(I)$  sampling?

# Challenge 2: Gap Function Can Have Local Maxima




# MetaEase Mechanisms

*Gap(I)* sampling is expensive







Gradient ascent

Avoid bad local Maxima




Smart seeds

# Gradient Ascent to Avoid Repeated Gap Sampling


$$Gap(I) = Opt(I) - Heuristic(I)$$

$$\nabla Gap(I) = \nabla Opt(I) - \nabla Heuristic(I)$$


**MetaEase: Compute efficiently separately, subtract, use as gradient direction**

Estimate  $\nabla Opt$  without **repeatedly solving**  $Opt$ , using duality theory.  
(💡 we can take a step toward the optimal direction, but finding the exact  $Opt$  value takes many more steps.)



Gradient ascent


# How MetaEase Computes $\nabla Heuristic(I)$

*Heuristic(I)*

```
for demand in demands:  
    if demand <= P:  
        # Pin to shortest path  
        allocate_shortest_path(demand)  
    else:  
        # Route remaining demands optimally  
        optimization_set.add(demand)  
solve_optimization(optimization_set)
```


This is arbitrary code, how do we efficiently get the gradient?

**Insight: Gaussian Process (GP) Surrogate can approximate it locally as a differentiable function.**




Gradient ascent

# MetaEase Uses GP Surrogate for $\nabla Heuristic(I)$



**Fast and Sample Efficient**


# MetaEase's Gradient Ascent Steps



Gradient ascent

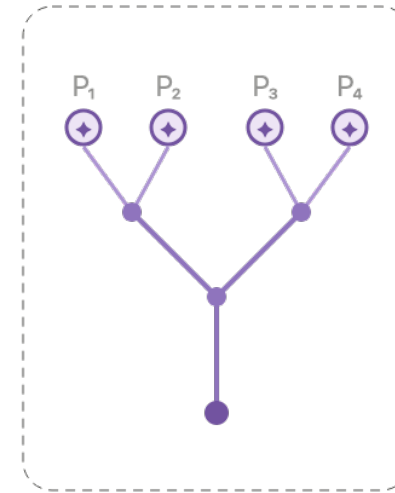
# MetaEase Mechanisms

No repeated  $Gap(I)$  sampling




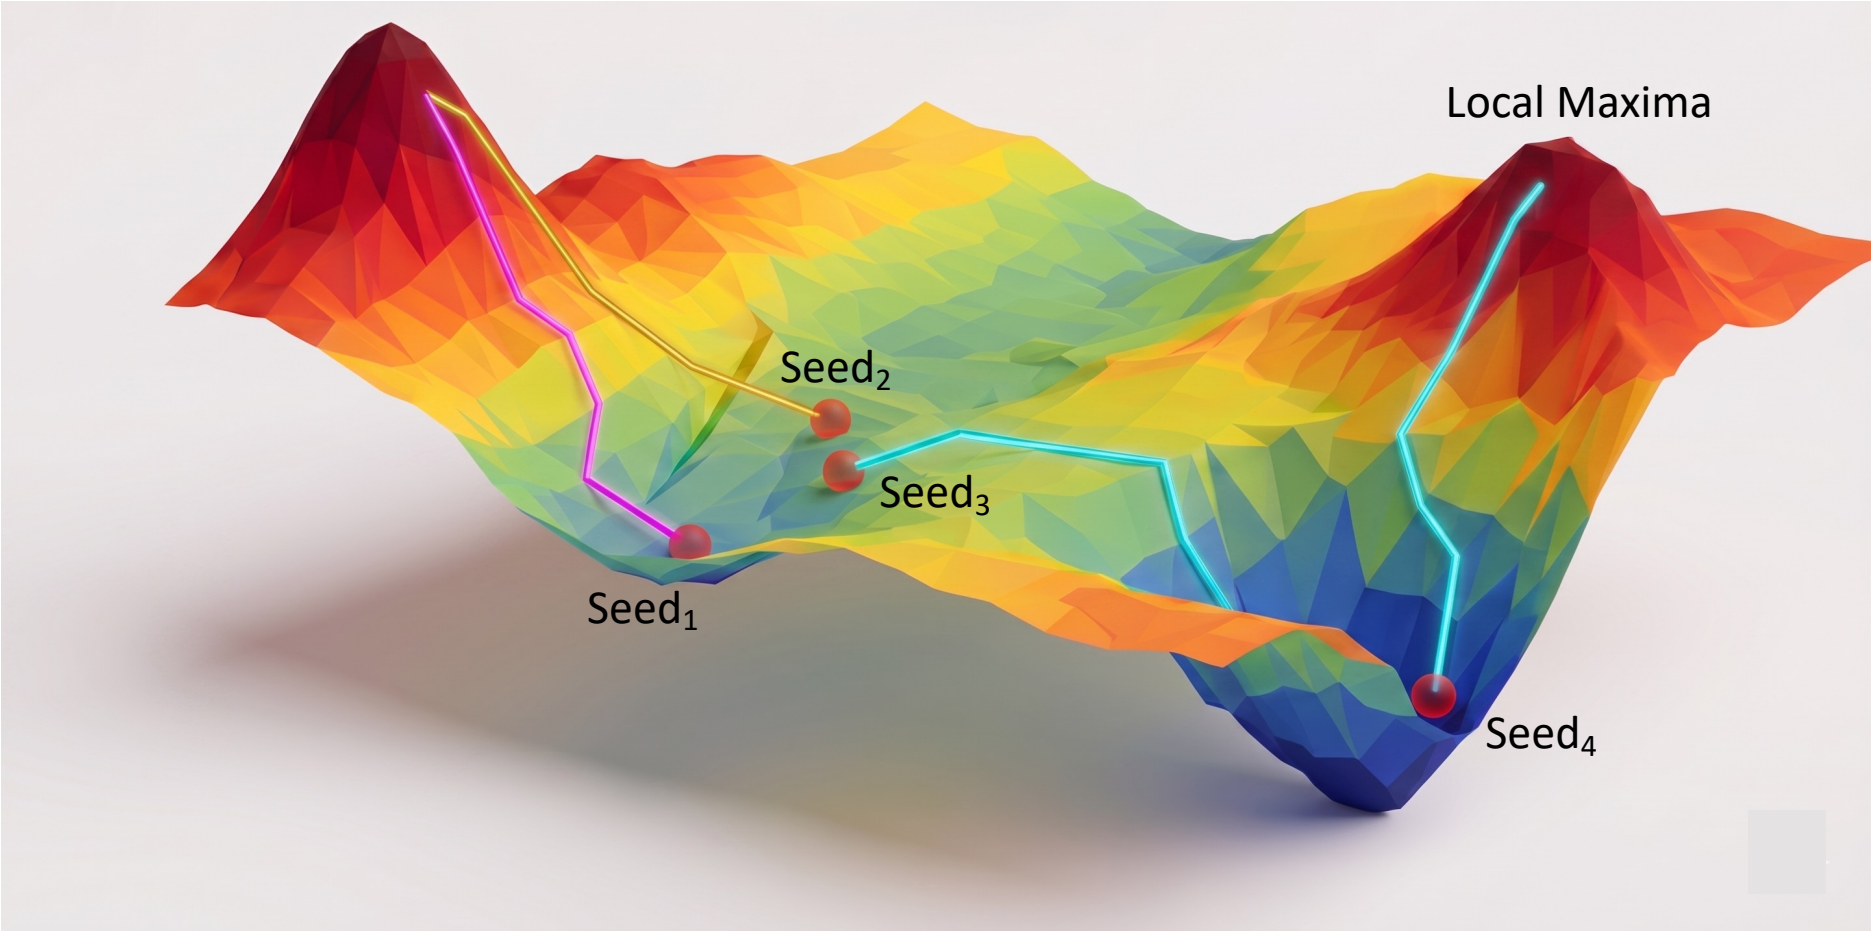
Gradient ascent

Avoid bad local Maxima



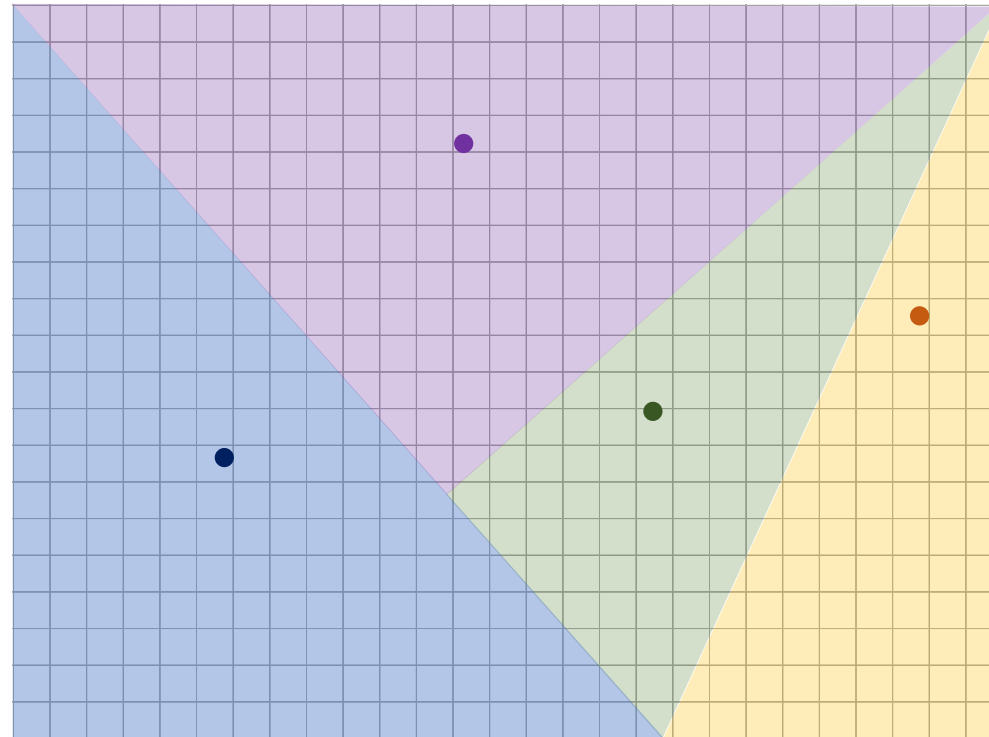
Smart seeds

# Selecting Smart Seeds Matters




Smart seeds

# Seed Different Heuristic Behavior to Avoid Local Maxima




In formal verification, these regions are equivalence classes.



Smart seeds

# Code Paths are Good Equivalence Classes for Heuristic



Smart seeds

**MetaEase uses Symbolic Execution (KLEE) to find Seeds.**

# What About the Cliffs?

# MetaEase: Putting it All Together




## Heuristic Code

```
for demand in demands:  
    if demand <= P:  
        # Pin to shortest path  
        allocate_shortest_path(demand)  
    else:  
        # Route remaining demands optimally  
        optimization_set.add(demand)  
solve_optimization(optimization_set)
```


*Opt* Formulation

## Symbolic Execution



For every seed

## MetaEase's Gradient Ascent



# Comparable Results to State-of-the-Art (MetaOpt)

## MetaEase: Demand Pinning as Code

```

for demand in demands:
    if demand <= P:
        # Pin to shortest path
        allocate_shortest_path(demand)
    else:
        # Route remaining demands optimally
        optimization_set.add(demand)

solve_optimization(optimization_set)
    
```

## MetaOpt: Demand Pinning as Optimization

### Outer Problem (Leader):

$$\begin{aligned} \max_{\mathbf{I} \in \mathbb{R}^n} \quad & \text{OPT}(\mathbf{I}) - \text{DP}(\mathbf{I}) \\ \text{s.t.} \quad & 0 \leq I_k \leq D_{\max} \quad \forall k \end{aligned}$$

### Inner Problem 1 — OPT (Max-Flow LP):

$$\begin{aligned} \text{OPT}(\mathbf{I}) = \max_{f'} \quad & \sum_k \sum_{p \in \mathcal{P}_k} f'_{k,p} \\ \text{s.t.} \quad & \sum_k \sum_{\substack{p \in \mathcal{P}_k \\ e \in p}} f'_{k,p} \leq \text{Cap}_e \quad \forall e \\ & \sum_{p \in \mathcal{P}_k} f'_{k,p} \leq I_k \quad \forall k \\ & f'_{k,p} \geq 0 \end{aligned}$$

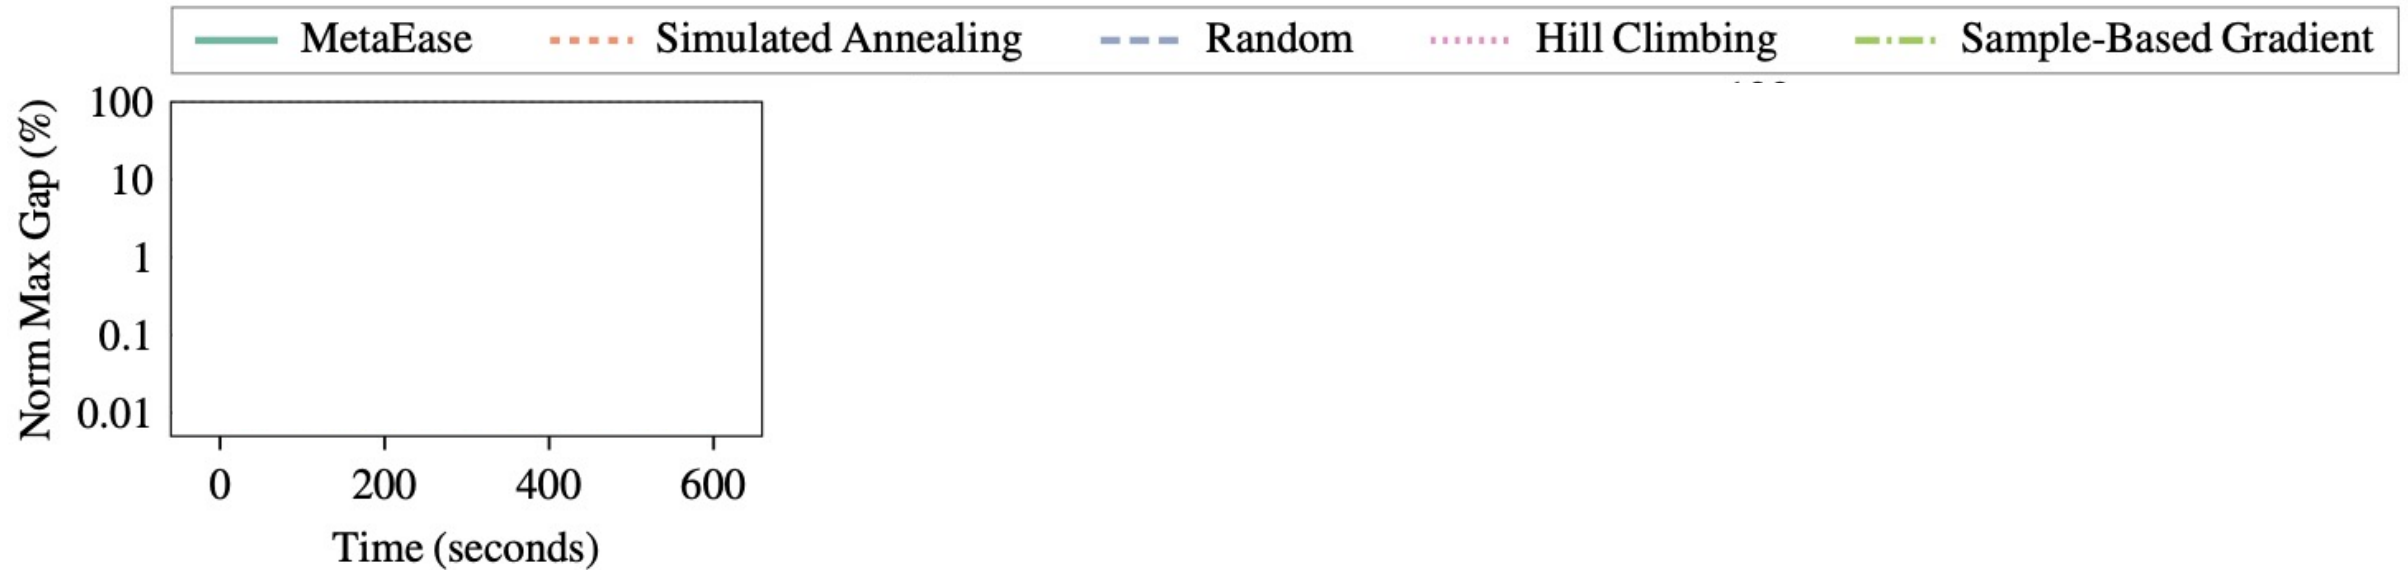
### Inner Problem 2 — DP (Conditional + Max-Flow):

$$\begin{aligned} b_k &\in \{0, 1\} && \text{(binary indicator)} \\ I_k &\leq P + M(1 - b_k) \quad \forall k && \text{(big-}M\text{)} \\ I_k &\geq P + \varepsilon - M b_k \quad \forall k && \text{(big-}M\text{)} \\ f_{k,\hat{p}_k} &= I_k \cdot b_k && \text{(pin if small)} \\ f_{k,p} &\leq M(1 - b_k) \quad \forall p \neq \hat{p}_k \\ \sum_k \sum_{\substack{p \in \mathcal{P}_k \\ e \in p}} f_{k,p} &\leq \text{Cap}_e \quad \forall e \\ \sum_{p \in \mathcal{P}_k} f_{k,p} &\leq I_k \quad \forall k \end{aligned}$$

### Then rewrite to single-level via KKT:

- + dual variables  $\lambda_e, \mu_k$  for each constraint
- + stationarity:  $\frac{\partial \mathcal{L}}{\partial f_{k,p}} = 0 \quad \forall k, p$
- + complementary slackness:  $\lambda_e \cdot g_e = 0$
- + nonlinear terms requiring SOS / big- $M$

# MetaEase Outperforms Black-Box Baselines




(a) Abilene

# Unique Capability: Analyzing any Performance Heuristic


**DNN-Based** Heuristic:

DOTE [NSDI'23] in Traffic Engineering




**Randomized** Heuristic:

Arrow [SIGCOMM'21] in Optical-IP TE

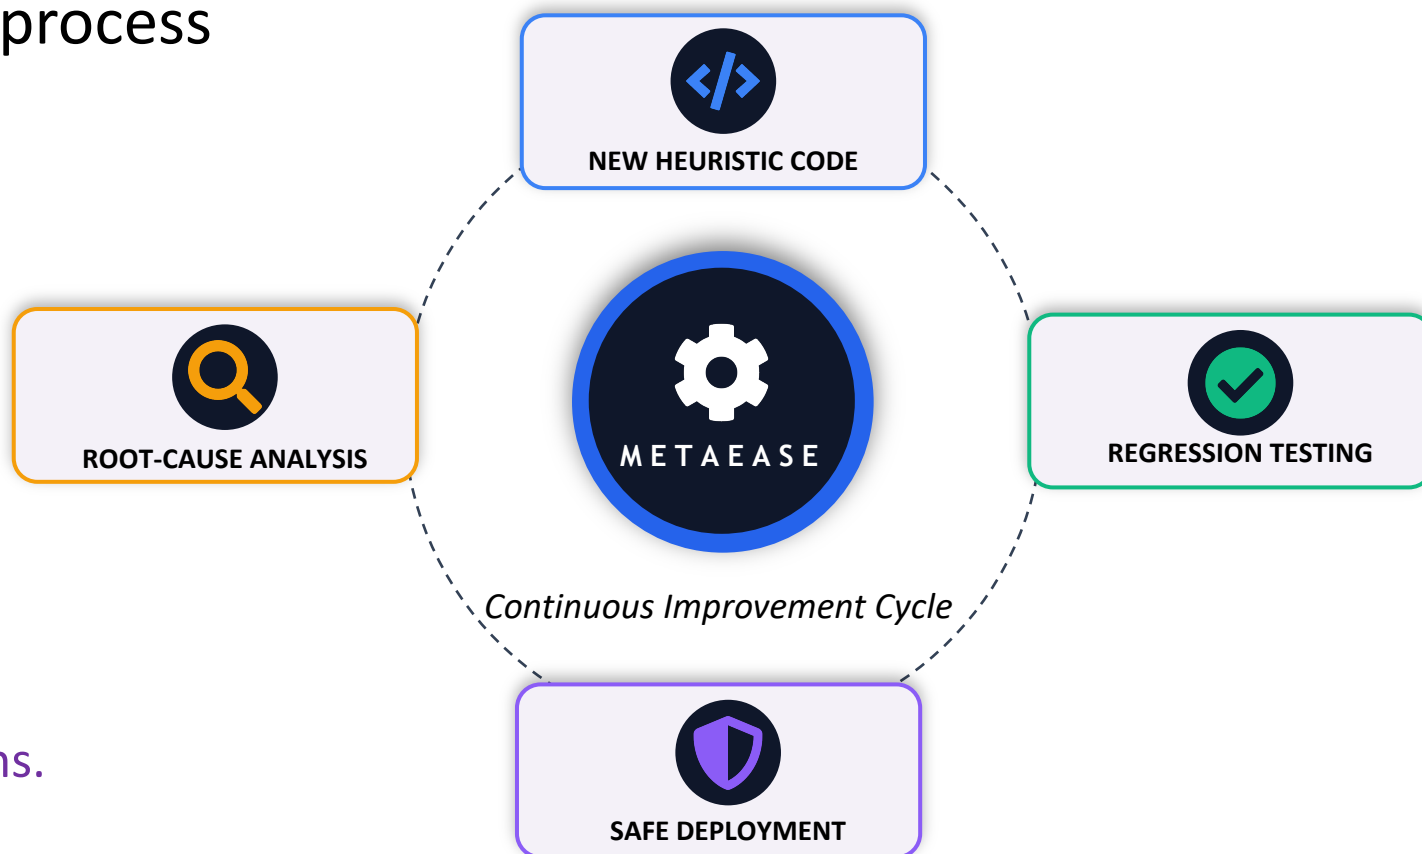


**MetaEase can analyze heuristics that other analyzers cannot analyze or require difficult encoding.**



# Conclusion

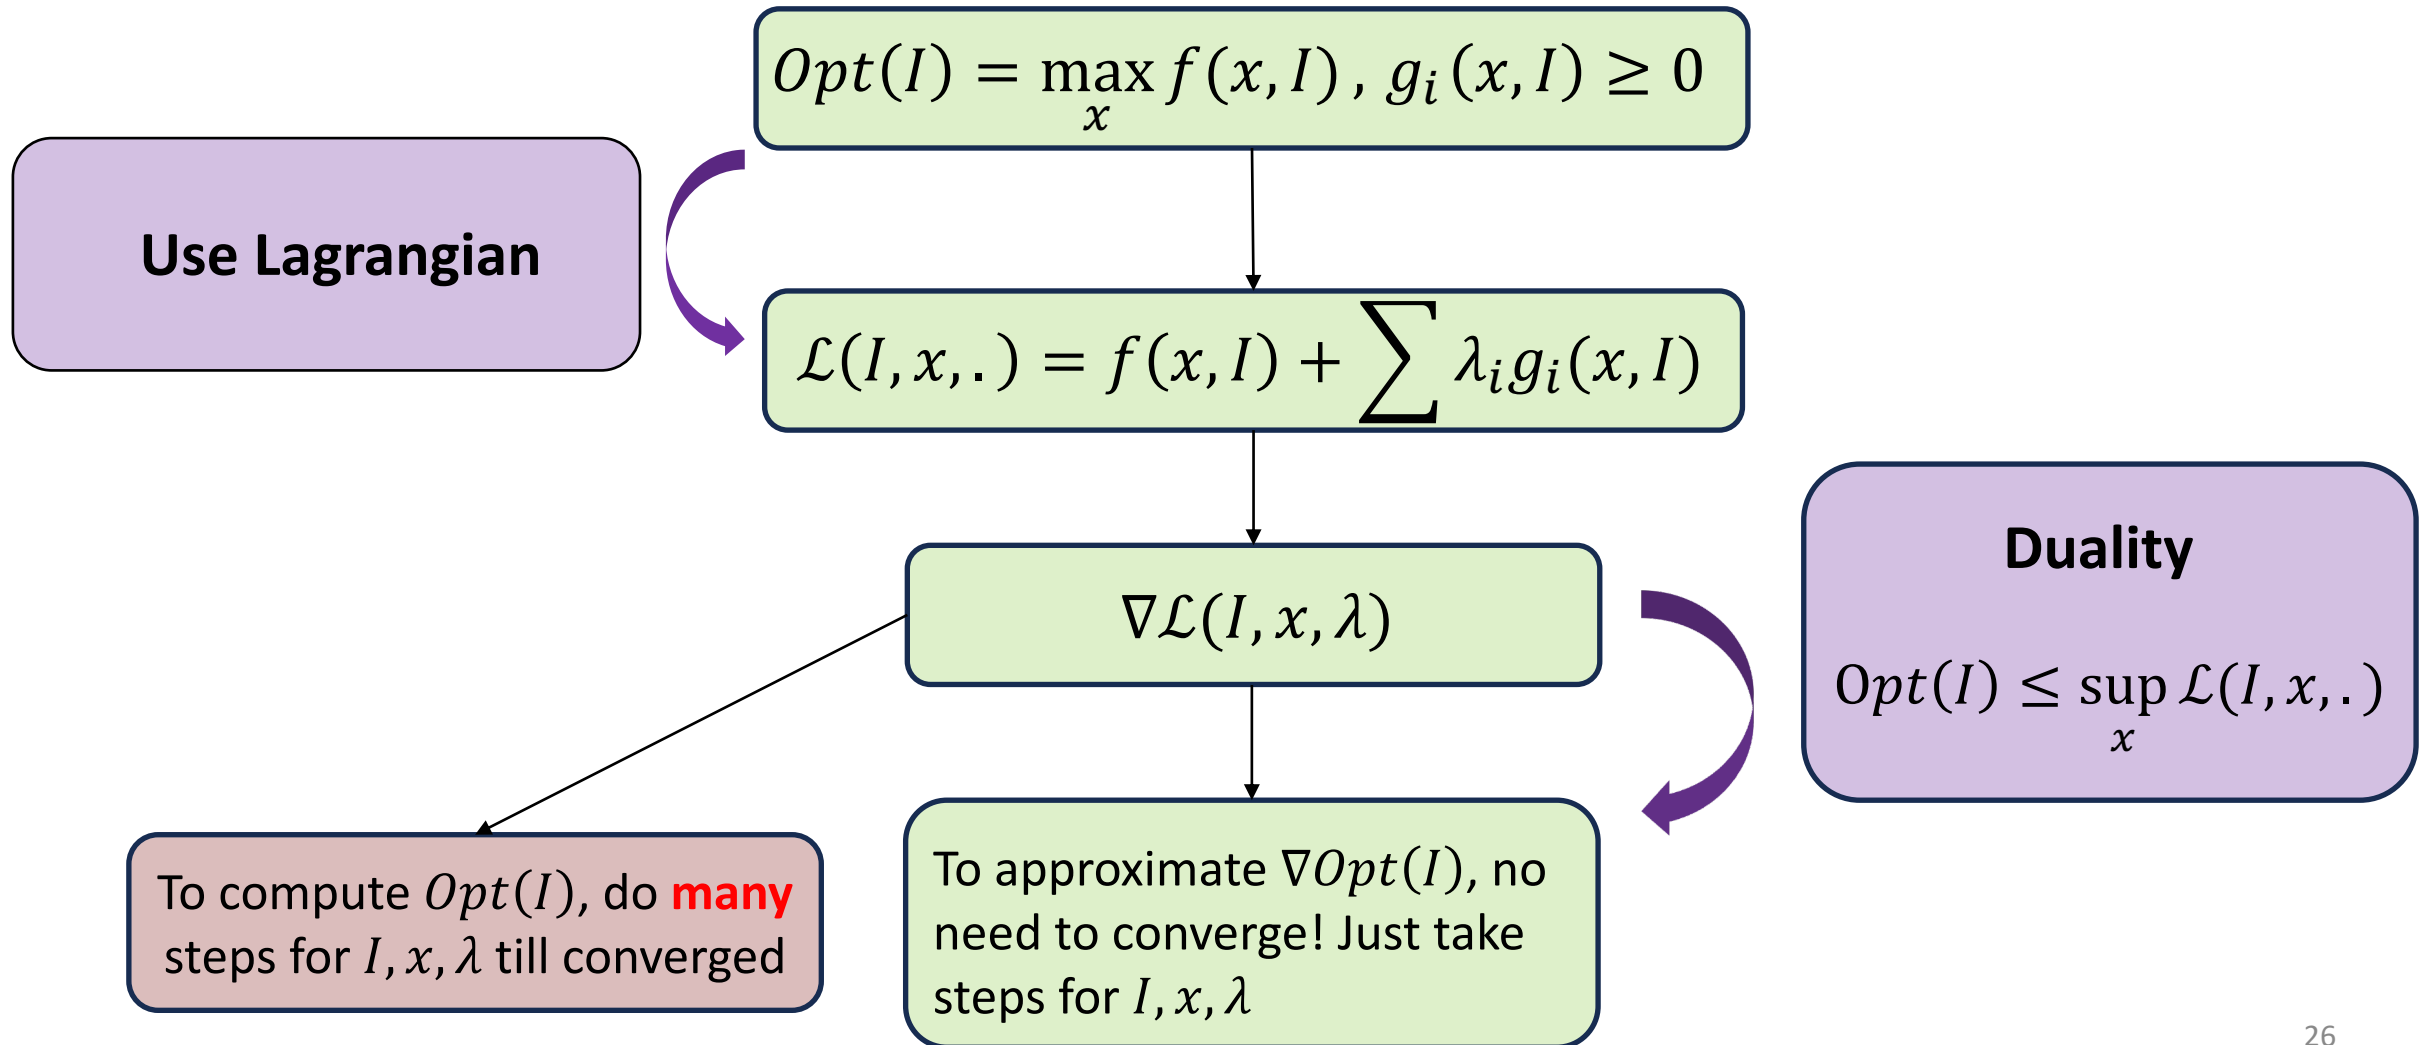
- First general-domain heuristic analyzer from code
- It matches or beats the SOTA across 5 domains and 8 heuristics
- MetaEase part of a design process



Thanks! I would love to answer your questions.  
Contact: [pkarimib@mit.edu](mailto:pkarimib@mit.edu)



# How MetaEase Computes $\nabla Opt(I)$



# How MetaEase Gets the Optimal

give me the traffic engineering optimal formulation in python

5:21 PM ↻ ✎ 📄 < 2/2 >

Architected network optimization solution using linear programming method... >

Here's a classic traffic engineering (TE) optimal formulation — a multi-commodity flow problem minimizing maximum link utilization, implemented with `scipy.optimize.linprog`.

Marshaled traffic engineering optimization framework for Python implement... >

The solver runs on a 4-node diamond topology and finds the optimal MLU of **87.5%** — the LP splits traffic across both paths (0→1→3 and 0→2→3, plus the cross-link 1→2) to balance utilization evenly across the bottleneck links.

The formulation covers the standard TE building blocks:

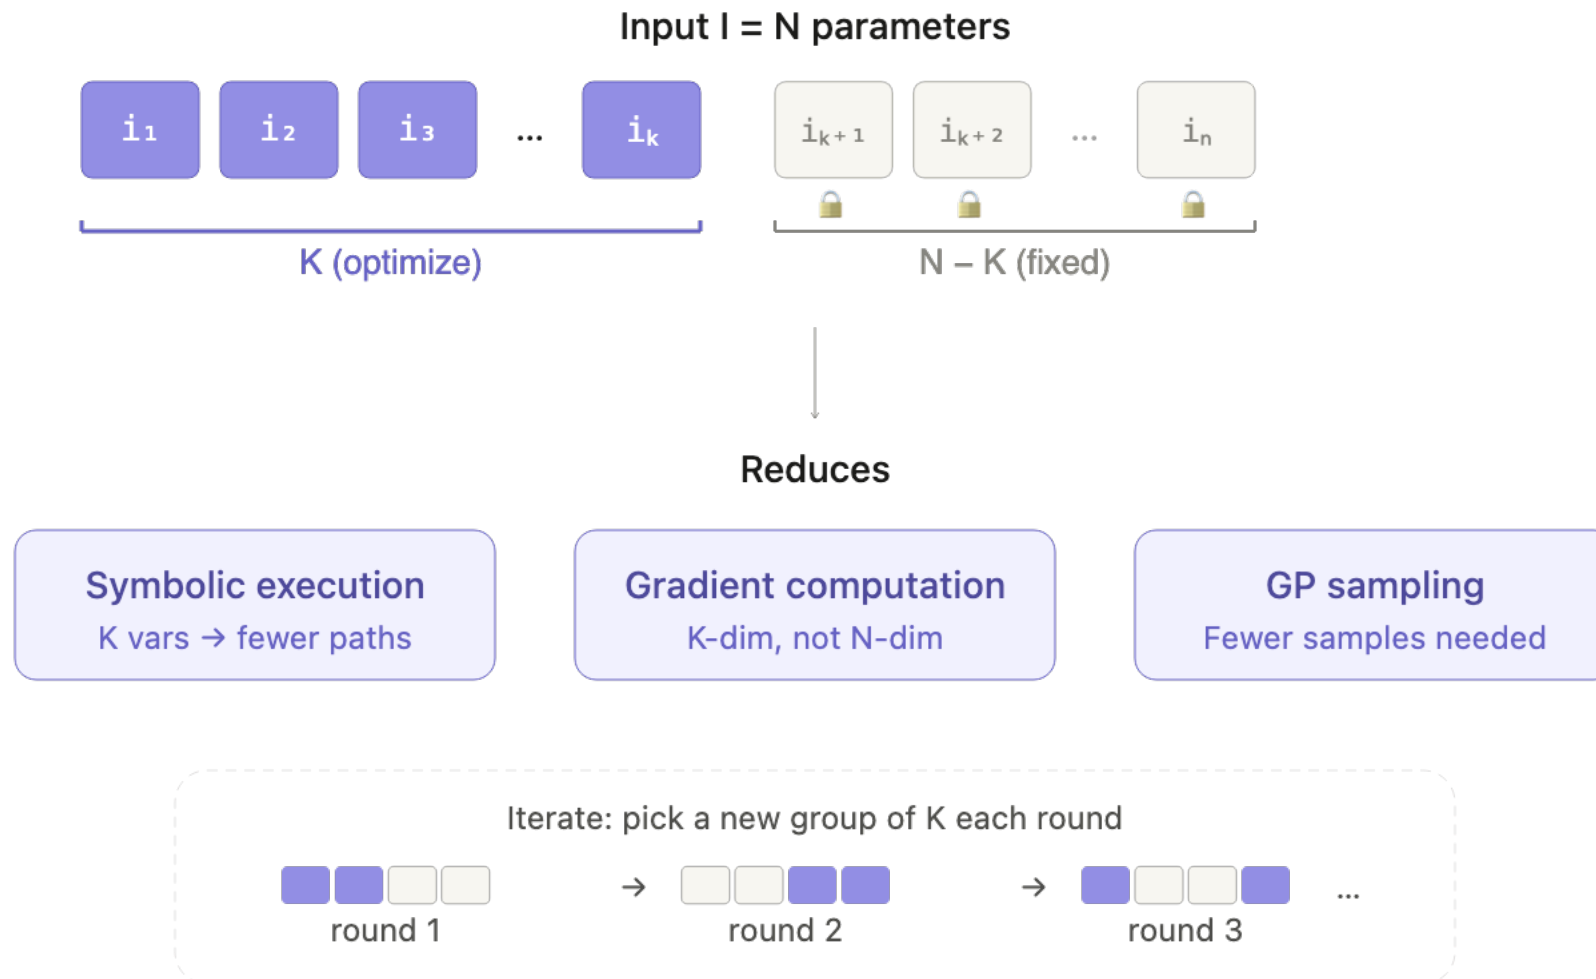
- **Objective:** minimize the maximum link utilization ( $u_{\max}$ ) across all links
- **Flow conservation:** at every node, incoming flow minus outgoing flow equals the demand (source/sink/transit)
- **Capacity coupling:** total flow on each link  $\leq$  capacity  $\times u_{\max}$ , which forces  $u_{\max}$  up to the worst-case link
- **Multi-commodity flow splitting:** each demand can split across

Write a message...

```
1 """
2 Traffic Engineering Optimal Formulation
3 =====
4 Multi-commodity flow LP that minimizes maximum link utilization (MLU).
5
6 Formulation:
7     minimize    u_max
8     subject to:
9         (1) Flow conservation at each node for each demand
10        (2) Link capacity: sum of flows on link (i,j) <= capacity(i,j) * u_max
11        (3) Demand satisfaction: total flow for demand k = demand_k
12        (4) Non-negativity: all flows >= 0, u_max >= 0
13
14 This is the standard Fortz-Thorup style TE objective used widely in WAN
15 and ISP backbone optimization.
16 """
17
18 import numpy as np
19 from scipy.optimize import linprog
20 from dataclasses import dataclass
21
22
23 # — Data Structures —————
24
25 @dataclass
26 class Link:
27     src: int
28     dst: int
29     capacity: float
30
31
```

# Optimizations for Scalability


- Projected gradient ascent: Update  $K$  parameters at a time



# Additional Optimizations for Scalability

- Adding domain customization “Hints”
  - Adding hints to prune and find harder samples
- Prune and prioritize equivalence classes

# Challenge 3: Gap Function Can Be Non-Smooth




# MetaEase Mechanisms


No repeated  $Gap(I)$  sampling

Avoid bad local Maxima

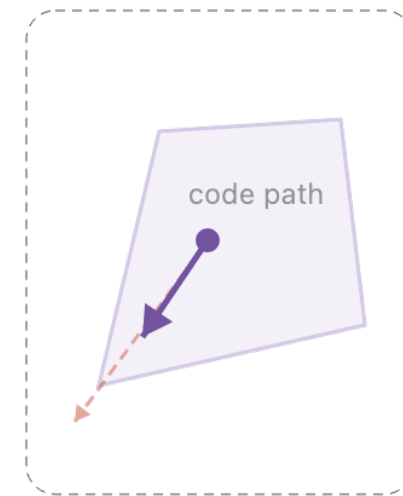
Handle non-differentiability



Gradient ascent



Smart seeds



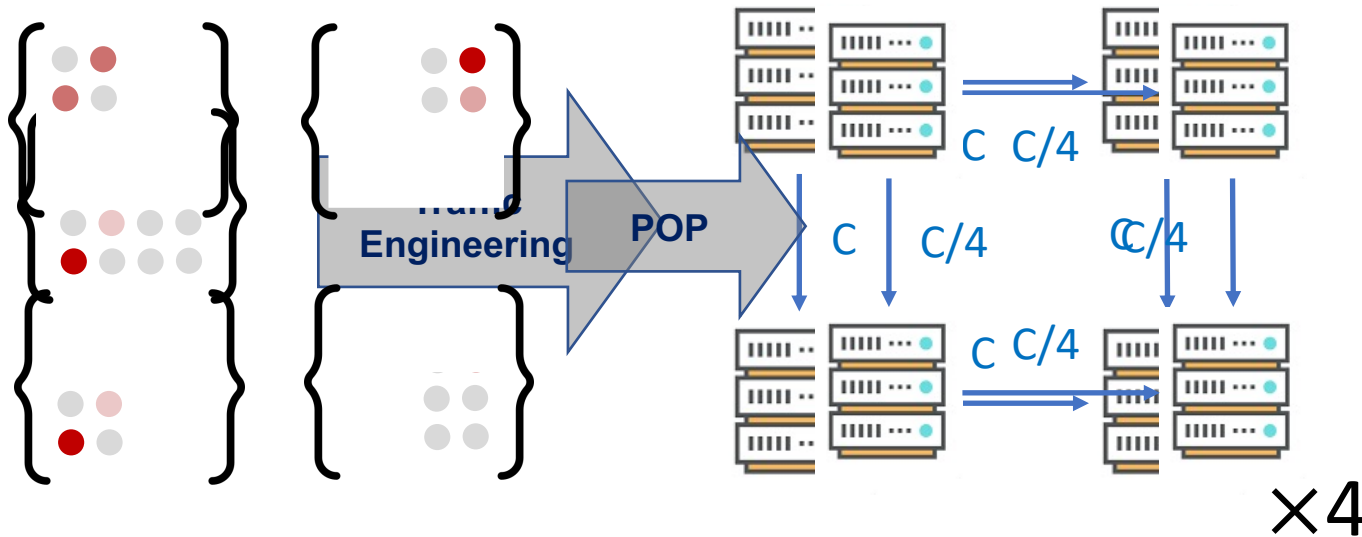
Path projection

# MetaEase Projects the Gradients Back to Code Path

# MetaEase vs MetaOpt Gap Ratio

$$\text{Gap ratio} = \frac{\text{MetaEase Max Gap}}{\text{MetaOpt Max Gap}}$$


Heuristic	Topo	Abilene	B4	Cogentco	Swan	Uninett2010
POP		98%	89%	88%	91%	65%



# MetaEase vs MetaOpt - Vector Bin Packing

$$\frac{\text{MetaEase Max Gap}}{\text{Metaopt Max Gap}}$$

Heuristic	#Balls, Dim	10, 1D	10, 2D	15-1D	15-2D	20-1D	20-2D
First-Fit Decreasing		100%	66%	100%	75%	100%	250%



# Hints In MetaEase

```
for (int k = 0; k < K; ++k) {
    // Restrict
    //     each demand to a bounded interval.
    klee_assume(0 <= demand[k] && demand[k] <= D_MAX);
}

bool any_edge_overloaded = false;

for (int e = 0; e < num_edges; ++e) {
    // Compute the total
    //     load induced on edge e by all demands.
    int edge_load
        = sum_demands_on_edge(demand, edges[e]);

    // Track whether
    //     this input stresses at least one edge.
    // We want: exists
    //     e such that edge_load(e) > capacity(e).
    any_edge_overloaded = any_edge_overloaded
        || (edge_load > edges[e].Capacity);
}

// Hardness predicate:
// If no edge is overloaded
//     , the instance is typically "easy",
// and is unlikely to expose a gap
//     between the heuristic and the benchmark.
klee_assume(any_edge_overloaded);
```

# Path-Aware Gradient Step

---

**Algorithm 1:** Path-aware gradient step in MetaEase

---

**Input:** Block size  $\Delta$ , samples  $N$ , step size  $\eta$

```
1  $h \leftarrow \text{path}(I)$  // current code path
2  $\mathbf{X} \leftarrow \{u \in \text{SAMPLEBOX}(I, \Delta) \mid \text{path}(u) = h\}$ 
3  $\mathbf{Y} \leftarrow [\text{Heuristic}(x) \mid x \in \mathbf{X}]$ 
4  $\text{Heuristic}_{GP} \leftarrow \text{FITGP}(\mathbf{X}, \mathbf{Y})$ 
5 Compute  $g_b \leftarrow \nabla \text{Benchmark}(I)$ 
6 Compute  $g_h \leftarrow \nabla \text{Heuristic}_{GP}(I)$ 
7  $g \leftarrow g_b - g_h$ 
8  $\tilde{I} \leftarrow I + \eta \cdot g$ 
9 if  $\text{path}(\tilde{I}) = h$  then
10 |  $I \leftarrow \tilde{I}$  // stay in same path
11 else
12 | Pick  $z \in \mathbf{X}$  most aligned with  $g$ 
13 |  $I \leftarrow z$  // project back into path
```

---

# Encoding LLM-Generated Heuristic In MetaEase

Code 1: New Sort-based Traffic Engineering Heuristic

```
double SortBasedHeuristic(Graph *G) {  
    // collect and sort demands  
    Demands D = positive_demands(G);  
    sort_desc(D);  
  
    // split top 20% as critical  
    int k = max(1, 0.2 * size(D));  
    Demands Critical = top_k(D, k);  
    Demands Remaining = rest(D, Critical);  
  
    // solve optimal on critical demands first  
    Routing  
        criticalRoutes = SolveOptimal(G, Critical);  
    // Adjust the capacities  
    Graph G = AdjustCapacities(G, criticalRoutes);  
    // solve optimal on non critical demands  
    Routing  
        nonCriticalRoutes = SolveOptimal(G, Remaining);  
  
    return G->total_met_demands;  
}
```

# Encoding LLM-Generated Heuristic In MetaOpt

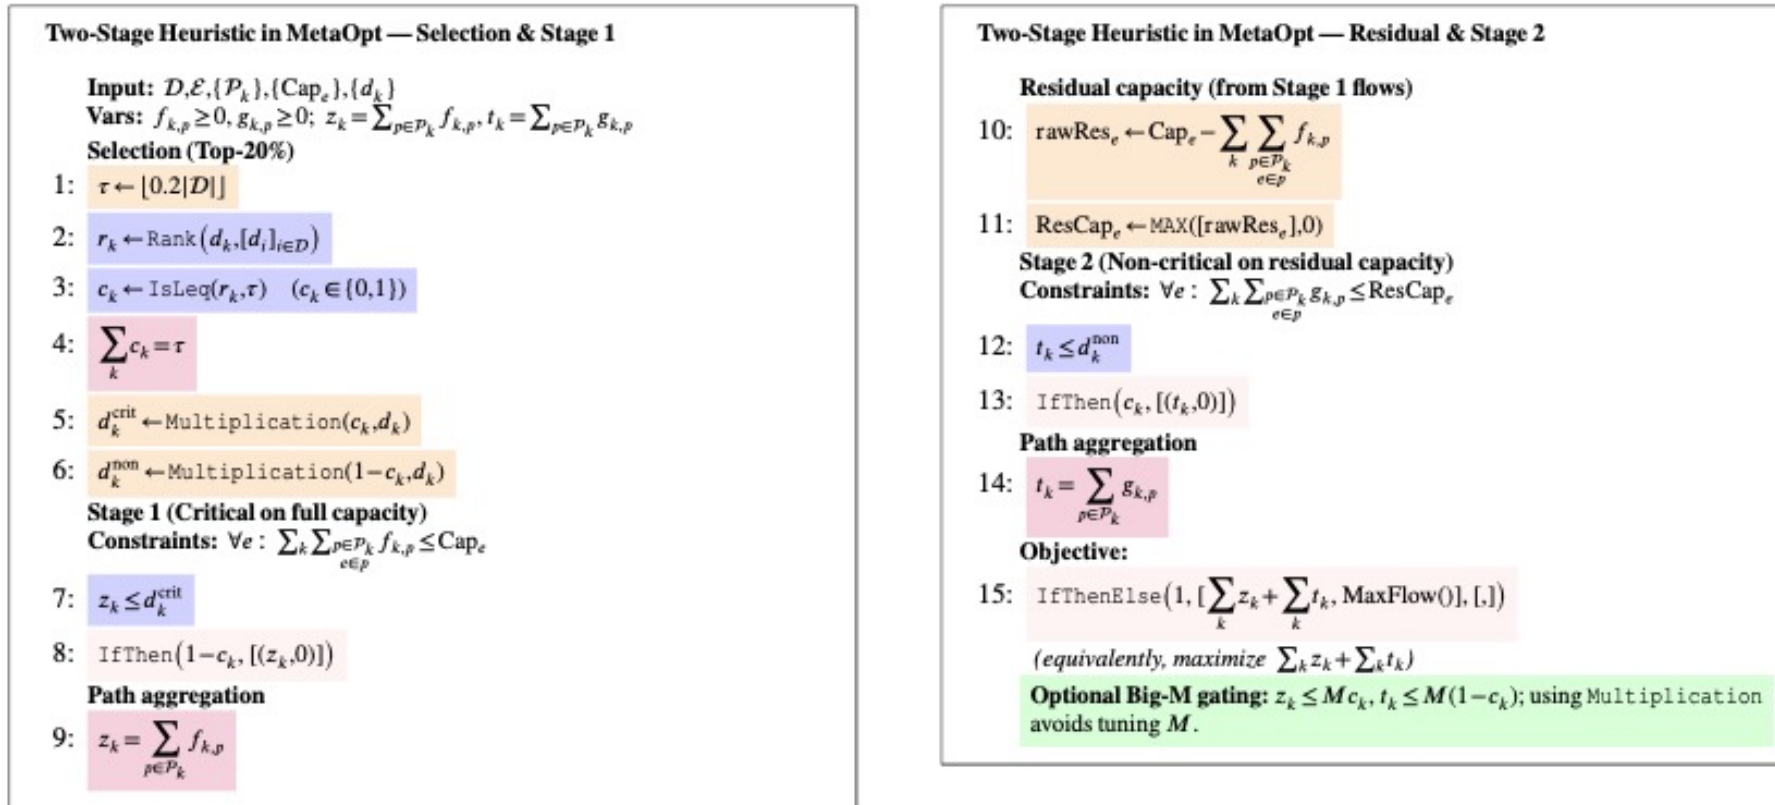



Figure 19: Encoding of Traffic Engineering heuristic in MetaOpt. This heuristic selects the top-20% of demands as critical (1) route critical set optimally, (2) then route remaining non-critical demands on residual capacities. This required a lot of attention and LLMs can't do it without human supervision to check for correctness.

# MetaEase vs MetaOpt Progress Over Time



# MetaEase vs Black-Box

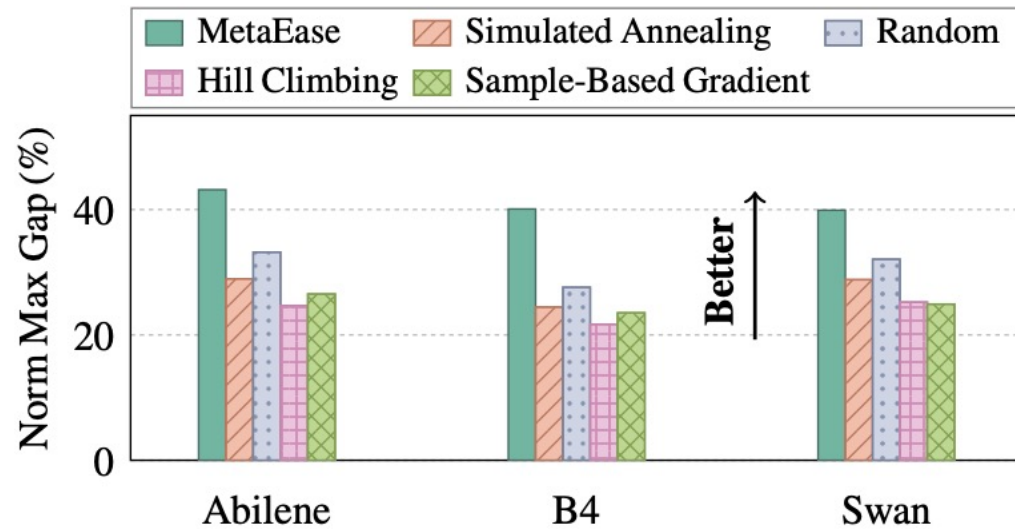


Figure 10: MetaEase finds a larger performance gap than black-box baselines for POP across production topologies under the same time budget.

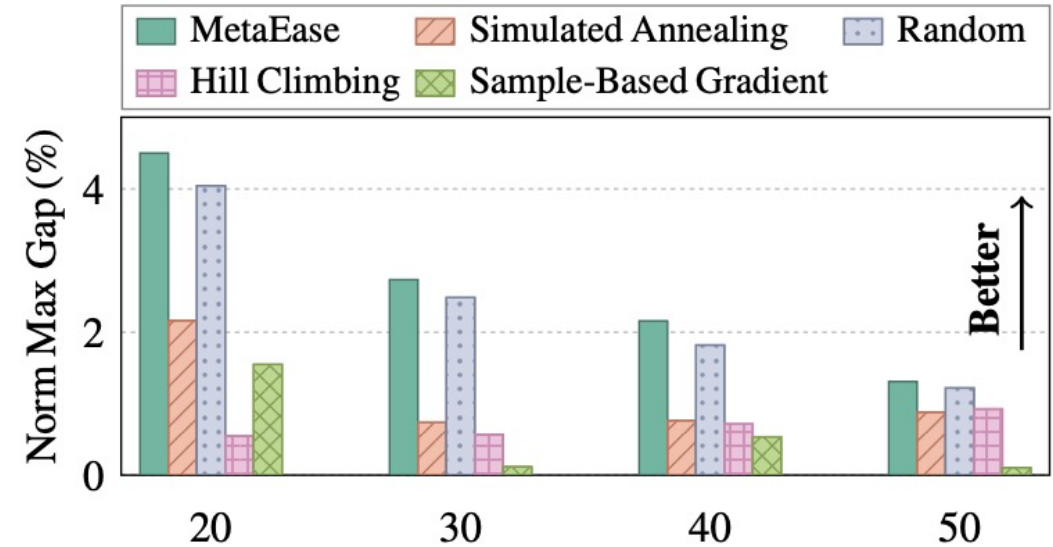


Figure 11: MetaEase outperforms all the baselines when it analyzes knapsack problems for 20, 30, 40, and 50 items. We normalized the gaps by number of items multiplied by the upper bound on each item's weight.

# MetaEase Ablations

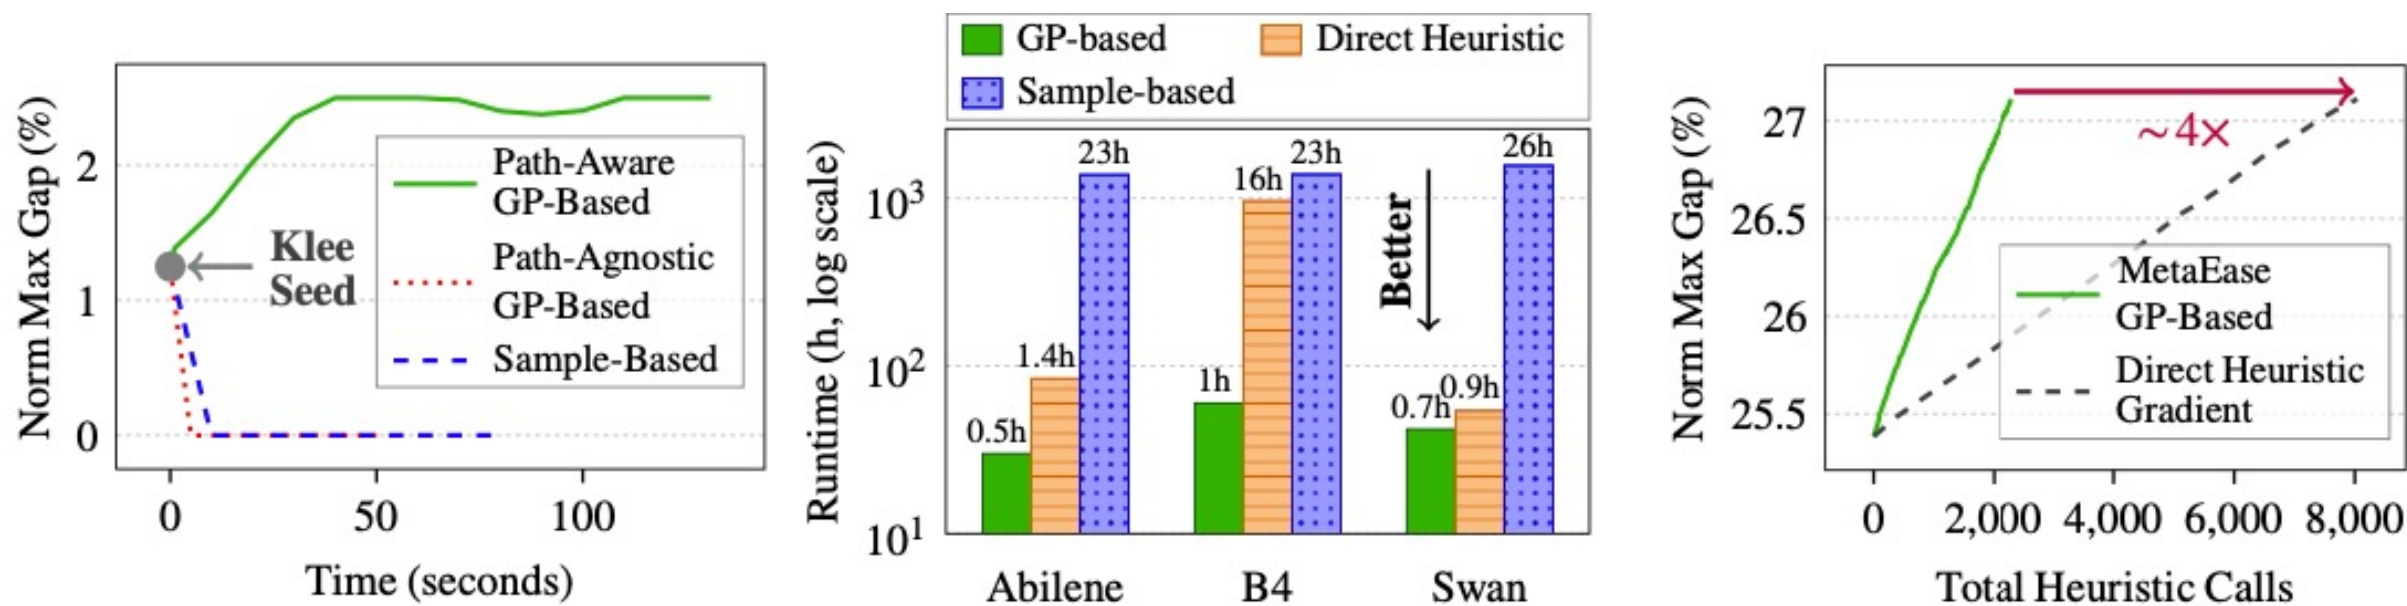
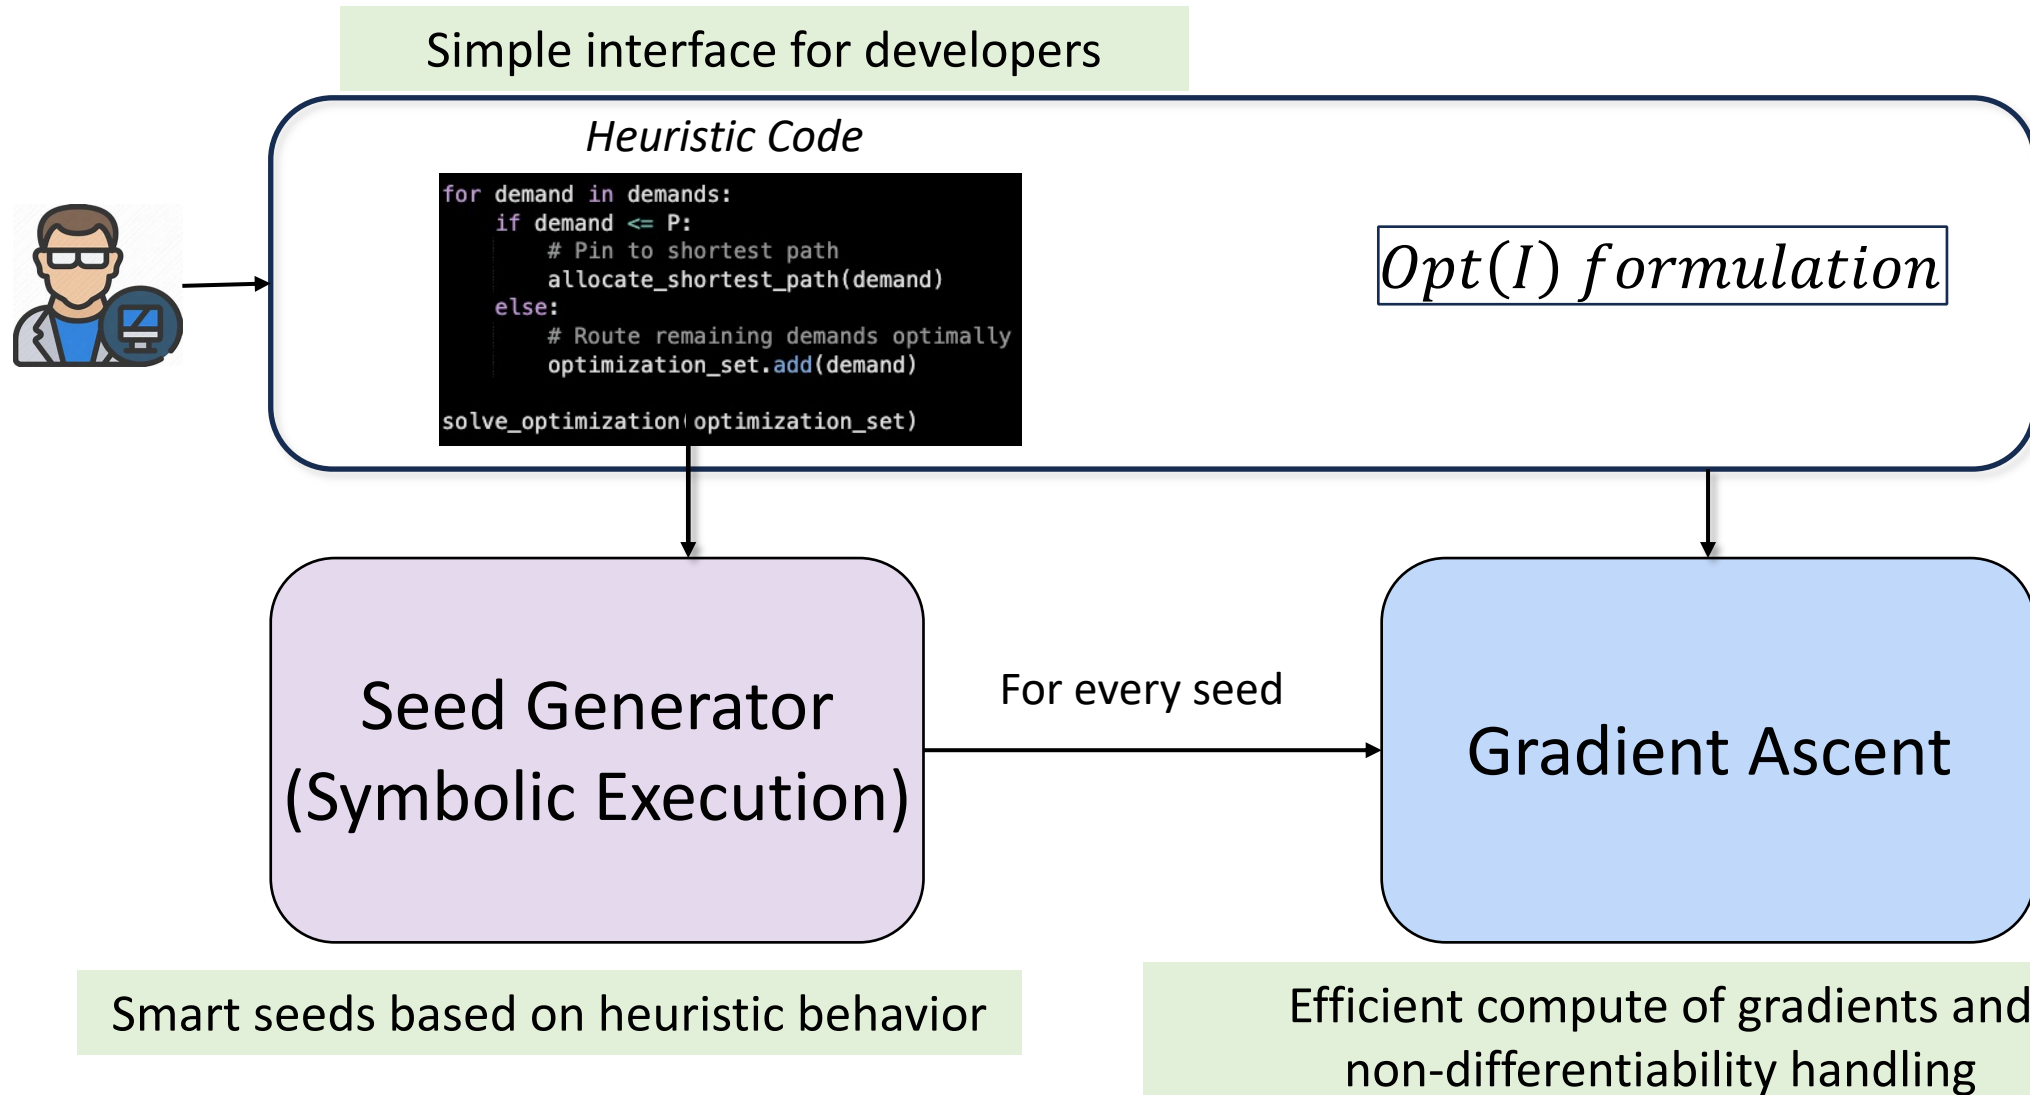


Figure 15: “Path-based” and “Surrogate-based” methods are effective: (left) MetaEase’s gradient ascent improves upon the initial point and handles the discontinuities; (middle) surrogate-based is faster than sample-based variants, including vanilla sample-based and one where it only estimates heuristic’s gradient; (right) MetaEase reduces the number of heuristic calls.

# MetaEase: Putting It All Together




# MetaEase Mechanisms


No repeated  $Gap(I)$  sampling

Avoid bad local Maxima


Handle non-differentiability



Gradient ascent

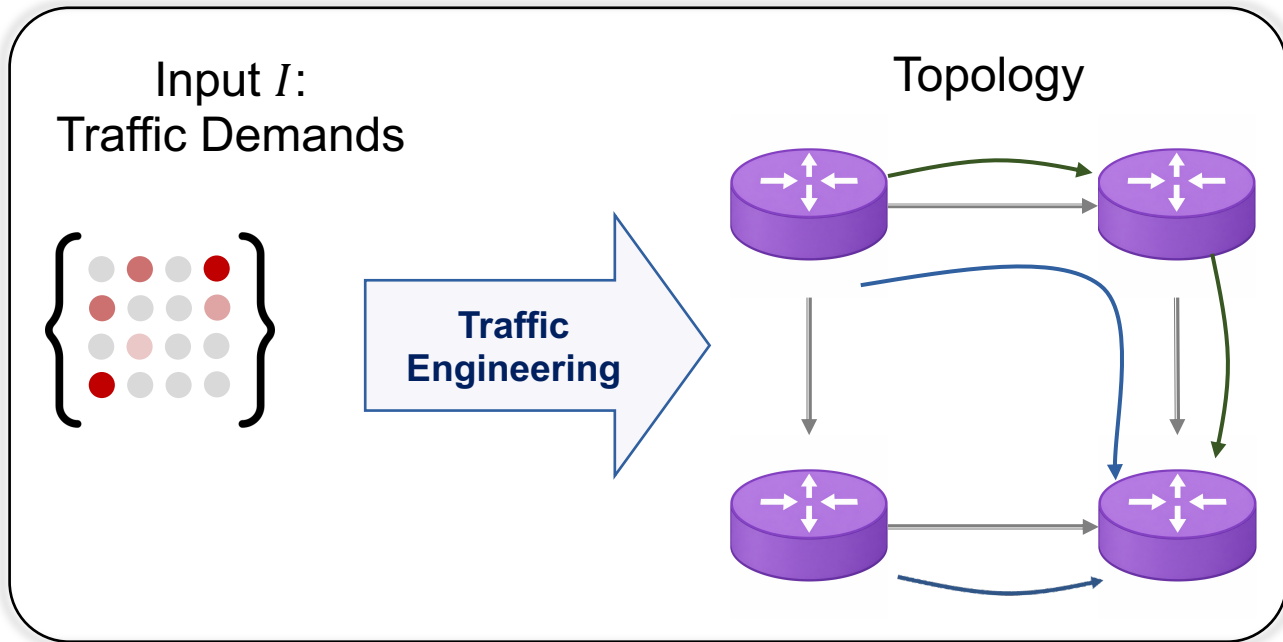


Smart seeds



Path projection

# Traffic Engineering Example



Performance Gap

$Opt(I) \rightarrow$  Optimal Performance on input  $I$


$Heuristic(I) \rightarrow$  DP's Performance on input  $I$

$$Gap(I) = Opt(I) - Heuristic(I)$$

**Demand Pinning *may* have 30% performance gap for some topologies!**


# MetaEase Mechanisms

No repeated  $Gap(I)$  sampling




Gradient ascent

Avoid bad local Maxima



Smart seeds

Handle non-differentiability



Path projection

# Unique Capability: Analyzing Any Numeric Heuristic


## DNN-Based Heuristic:

DOTE [NSDI'23] in Traffic Engineering

Method	Normalized Max Gap
MetaEase	73.13%
DNN Analyzer in [59]	71.84%
Random	63.02%
Simulated Annealing	61.39%
Hill Climbing	58.78%
Sample-based Gradient	58.78%

## Randomized Heuristic:

Arrow [SIGCOMM'21] in Optical-IP TE




MetaEase can analyze heuristics that other analyzers cannot analyze or require difficult encoding.

# Existing Solutions

## Model-based Tools

Encode any new heuristic into a mathematical format

Example: MetaOpt [NSDI '24], FPerf[NSDI '24] [Hotnets'24]



Encode

**OuterVar:**  $Y$ (size of balls)  
**Input:**  $C$ (capacity of bins)  
**for all ball  $i$  and bin  $j$  do**  
$$r_{ij} = C_j - Y_i - \sum_{\text{ball } u < i} x_{uj}^d$$
$$f_{ij} = \text{AllLeq}([-r_{ij}^d]_d, 0)$$
$$\gamma_{ij} = \text{AllEq}([x_{ik}^d]_{d,k < j}, 0)$$
$$\alpha_{ij} = \text{AND}(f_{ij}, \gamma_{ij})$$
$$\text{IfThenElse}(\alpha_{ij}, [(x_{ij}, Y_i)], [(x_{ij}, 0)])$$
**end for**

This is hard!

**X Hard To Use**  
**Domain Specific**

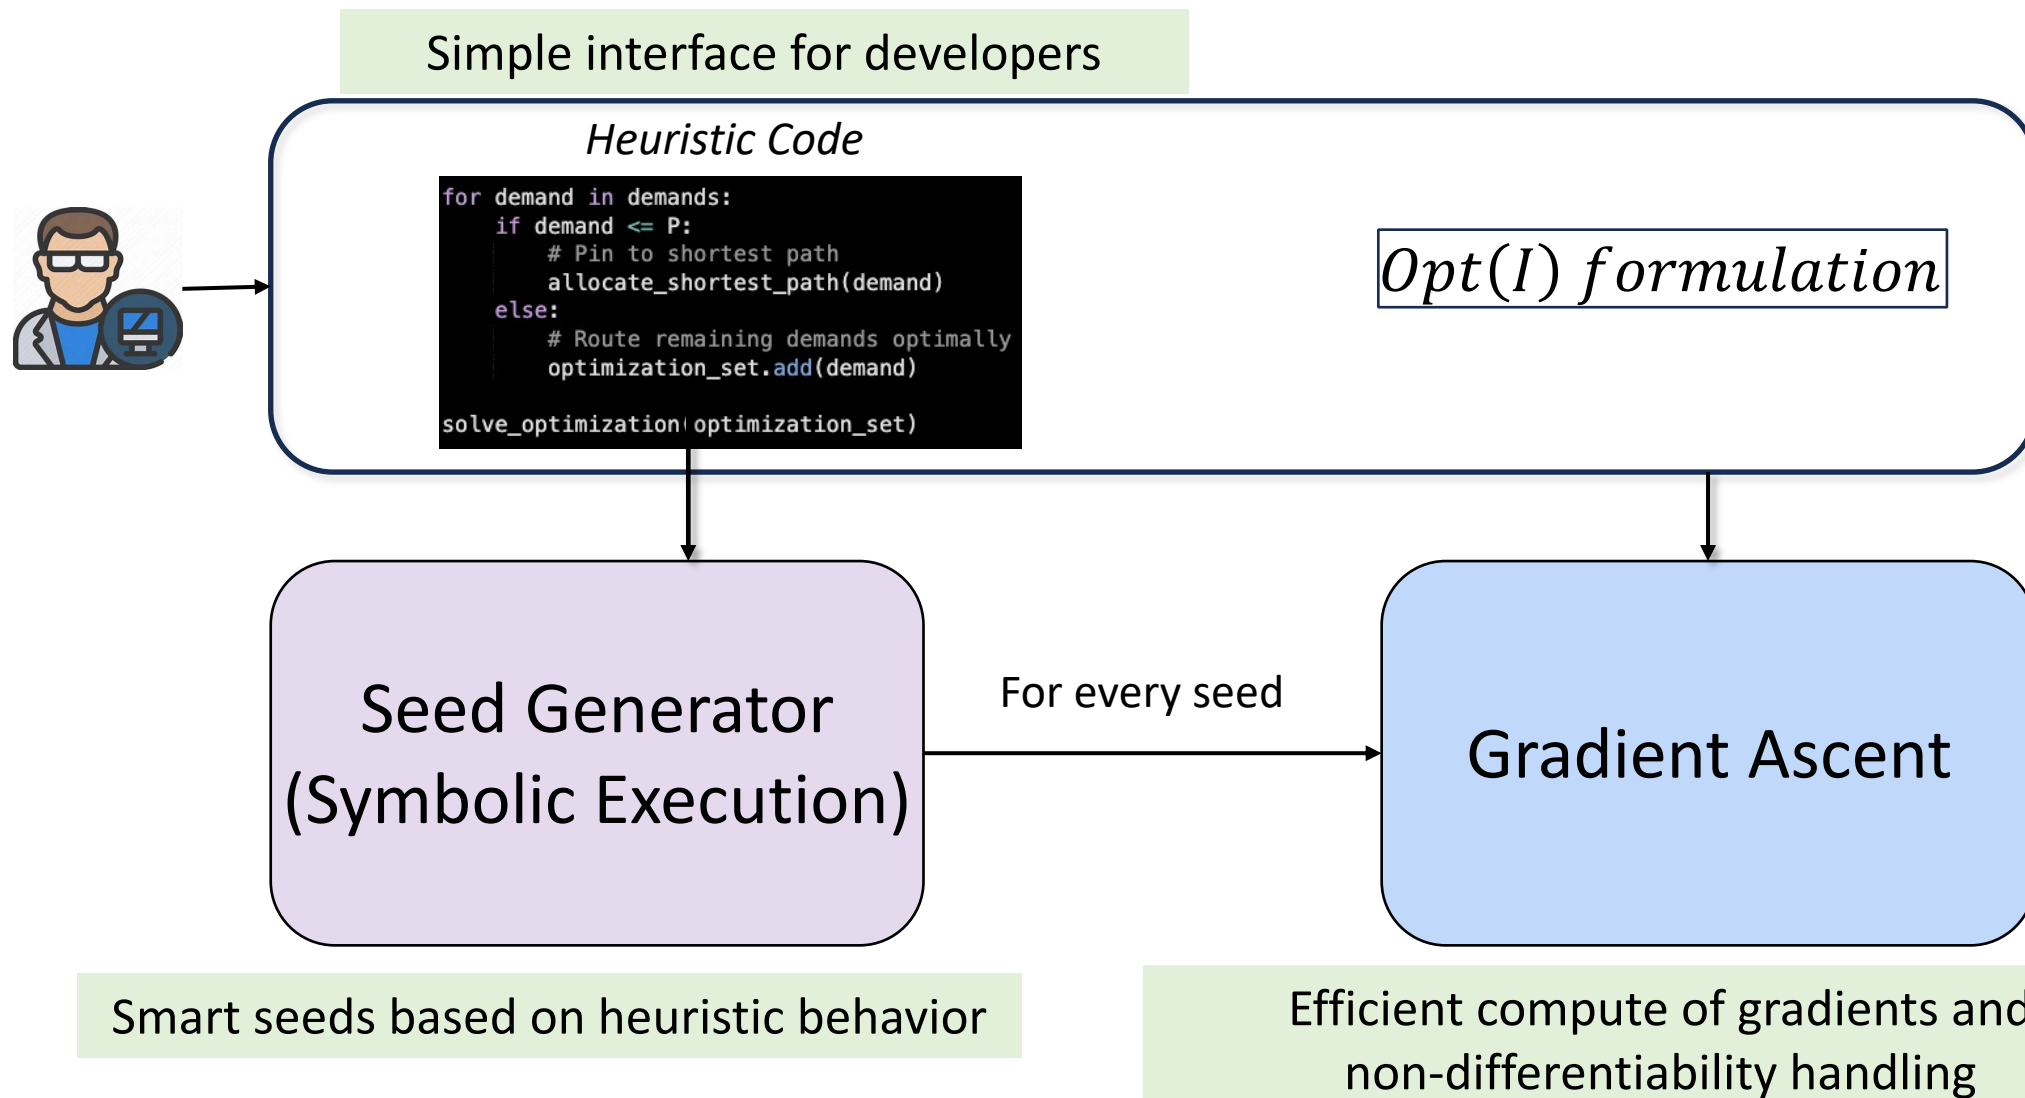
## Black-Box Tools

Example: Random Sampling, Hill Climbing, Simulated Annealing

Easy to use, progress by sampling many  $Gap(I)$  values


**X Inefficient**

# MetaEase Overview



# MetaEase Mechanism's For Gradient Ascent

$$\nabla Gap(I) = \nabla Opt(I) - \nabla Heuristic(I)$$



Compute  $\nabla Opt$  efficiently ✓

Compute  $\nabla Heuristic$  efficiently ✓

Smart gradient ascent seeds

Handling non-differentiability