

# Performance Analysis for Highly-Configurable Systems

Christian Kästner, Carnegie Mellon University, @SPLC'19



# Collaborators



Norbert Siegmund  
University of Weimar



Sven Apel  
Saarland University



Miguel Velez  
Carnegie Mellon Univ.



Pooyan Jamshidi  
Univ. South Carolina



David Garlan  
Carnegie Mellon Univ.

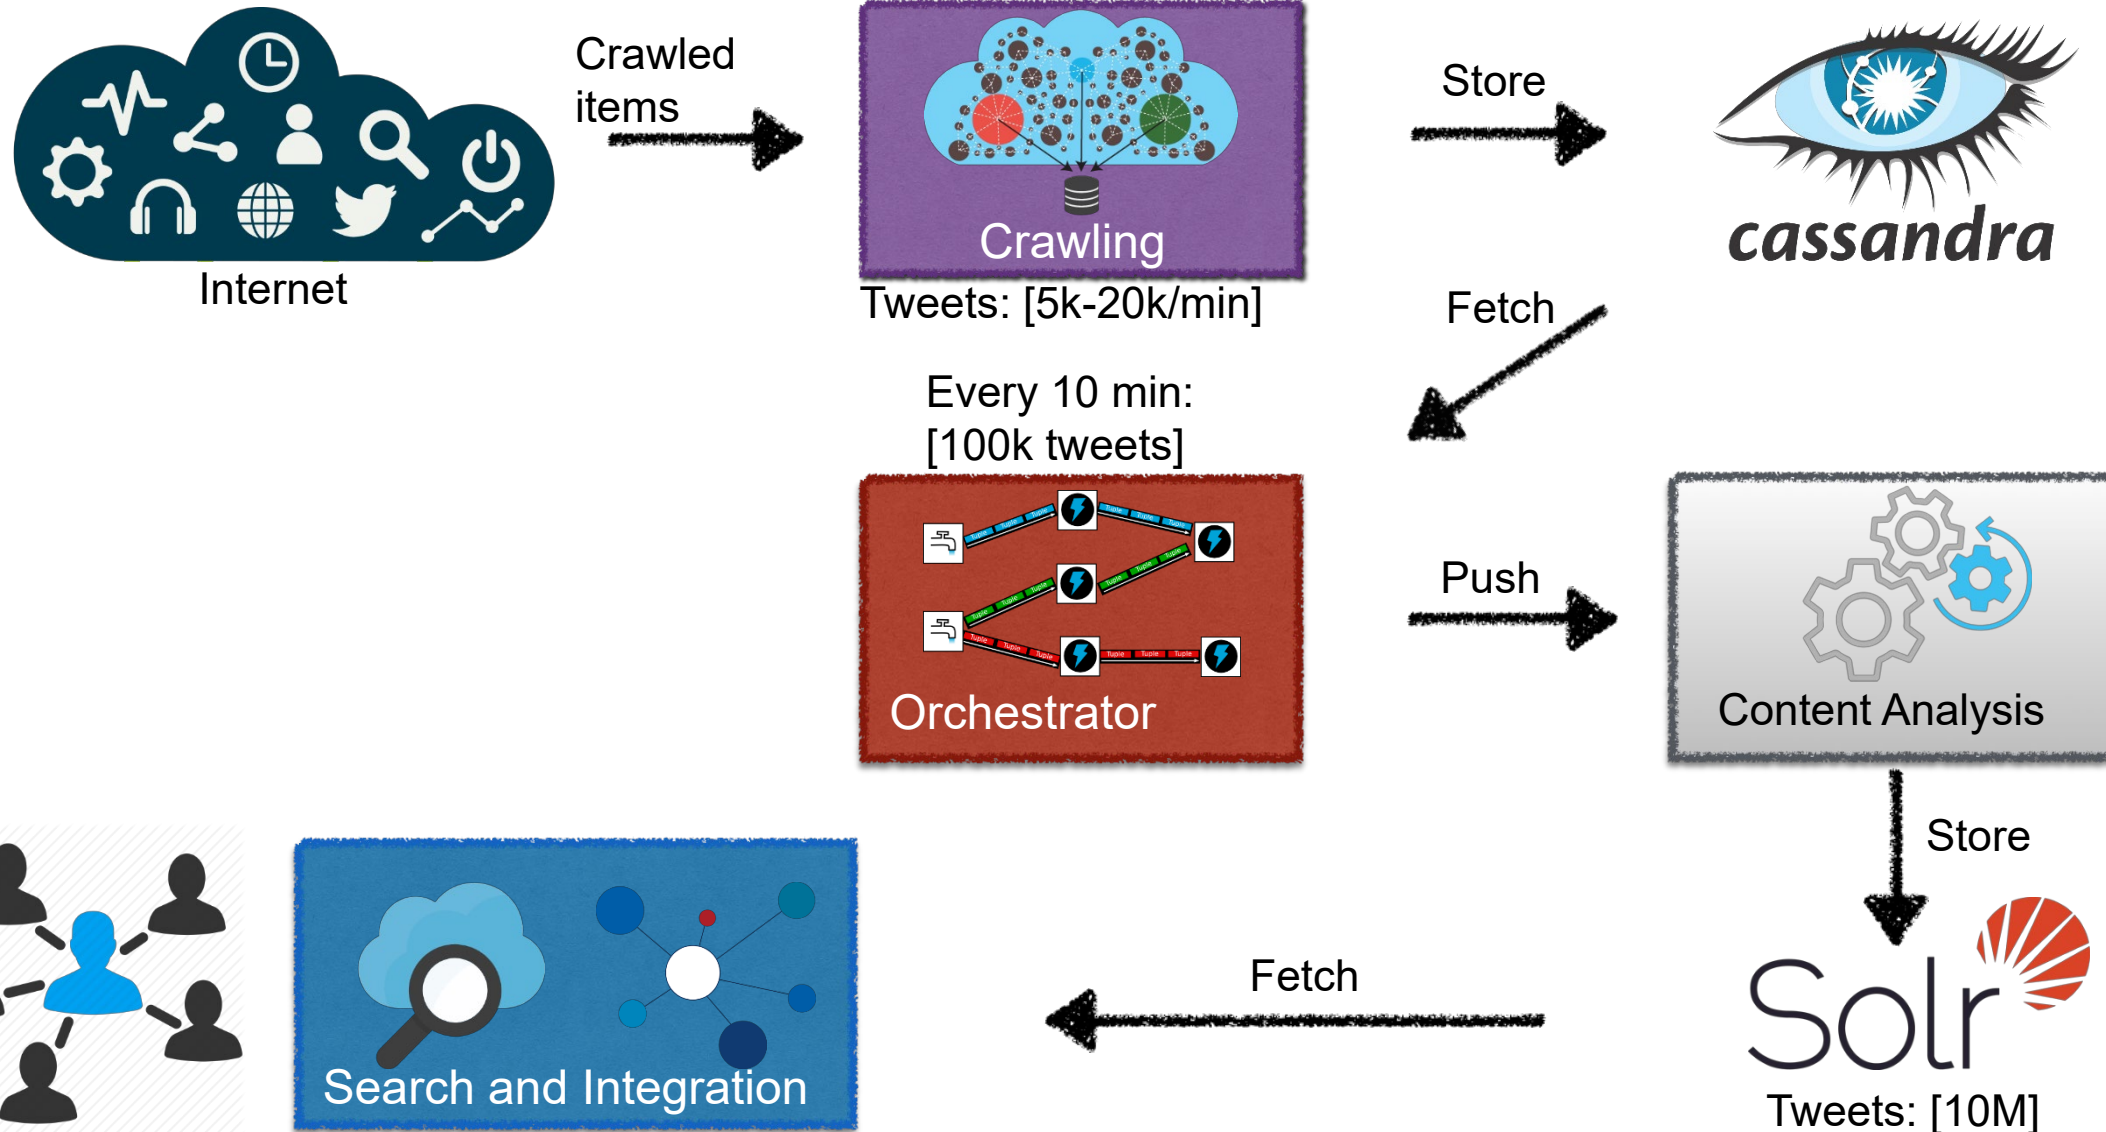


Sergiy Kolesnikov  
University of Passau

A grayscale, blurred office environment with several computer monitors and an office chair in the background. In the foreground, a silver laptop is open on a wooden desk, its screen displaying a dark, blank page. The text "Case Study" is overlaid in a large, bold, black font across the bottom of the laptop.

# Case Study

“**SocialSensor** quickly surfaces trusted and relevant material from social media - with context.”



# Changing Requirements



Internet

Crawled items


**10X**



Crawling

Tweets: [5k-20k/min]

Store




*cassandra*

Fetch

**Real time**

Every 10 min:  
[100k tweets]



Orchestrator

Push



Content Analysis

Store



Tweets: [10M]

Fetch

**100X**



Search and Integration

# Opportunity



**Data processing engines highly configurable**



> 100 options



> 100 options



> 100 options

```
102
103 drpc.port: 3772
104 drpc.worker.threads: 64
105 drpc.max_buffer_size: 1048576
106 drpc.queue.size: 128
107 drpc.invocations.port: 3773
108 drpc.invocations.threads: 64
109 drpc.request.timeout.secs: 600
110 drpc.childopts: "-Xmx768m"
111 drpc.http.port: 3774
112 drpc.https.port: -1
113 drpc.https.keystore.password: ""
114 drpc.https.keystore.type: "JKS"
115 drpc.http.creds.plugin: org.apache.storm.security.auth.DefaultHttpCredentialsPlugin
116 drpc.authorizer.acl.filename: "drpc-auth-acl.yaml"
117 drpc.authorizer.acl.strict: false
118
119 transactional.zookeeper.root: "/transactional"
120 transactional.zookeeper.servers: null
121 transactional.zookeeper.port: null
122
123 ## blobstore configs
124 supervisor.blobstore.class: "org.apache.storm.blobstore.NimbusBlobStore"
125 supervisor.blobstore.download.thread.count: 5
126 supervisor.blobstore.download.max_retries: 3
127 supervisor.localizer.cache.target.size.mb: 10240
128 supervisor.localizer.cleanup.interval.ms: 600000
129
```


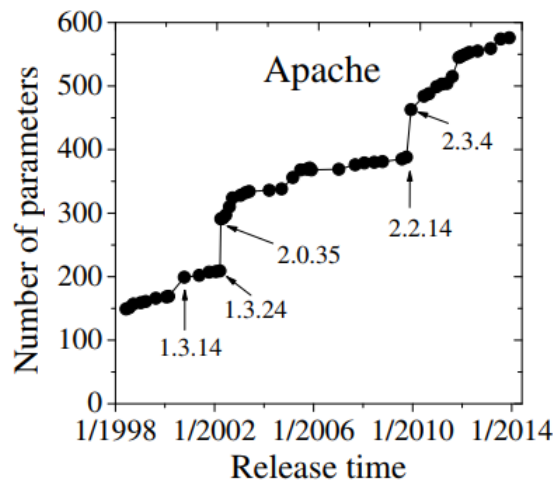




320<sup>optional, independent</sup> options

more combinations than estimated  
atoms in the universe



# Developers and users are overwhelmed with configuration options [Xu et al. FSE'15]



|   |  |
|---|--|
| <b>Parameter:</b> optimizer_prune_level (Boolean) | /*MySQL*/  |
| <b>Desc.:</b>                                     | Controls the heuristics applied during query optimization to prune less-promising partial plans from the optimizer search space. |
| <b>Values:</b>                                    | 0 or 1   |
| <b>Usage:</b>                                     | No user set the parameter in our dataset.  |

(a) Empirical, heuristic usages

|  |  |
|--|--|
| <b>Parameter:</b> key_cache_block_size (Numeric) | /*MySQL*/  |
| <b>Desc.:</b>                                    | The size in bytes of blocks in the key cache.                  |
| <b>Values:</b>                                   | [512, 16384]   |
| <b>Usage:</b>                                    | All the users stay with the default value 1024 in our dataset. |

Substantial increase in configurability

80% options ignored

# Changing Requirements



Internet

Crawled items


**10X**



Crawling

Tweets: [5k-20k/min]

Store




*cassandra*

Fetch

**Real time**

Every 10 min:  
[100k tweets]



Orchestrator

Push



Content Analysis

Store



Tweets: [10M]


Fetch

**100X**

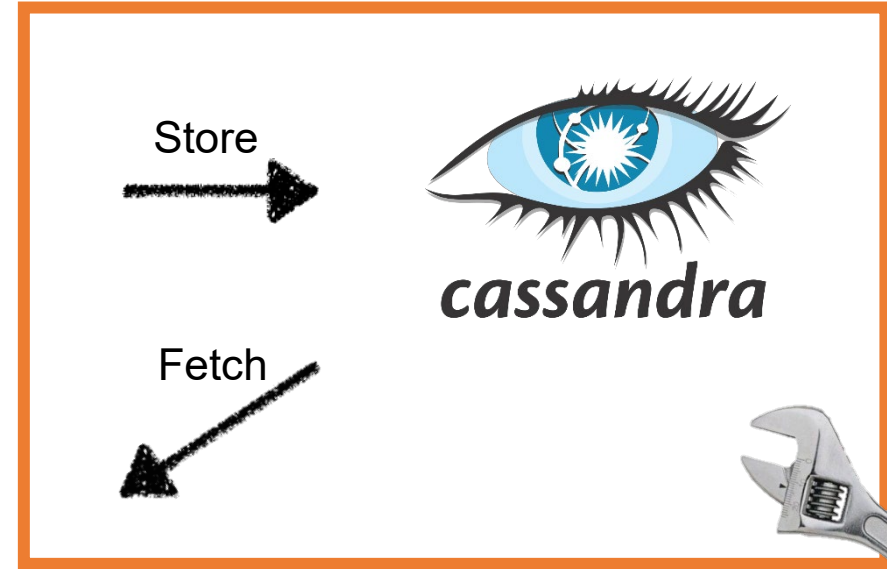


Search and Integration

# Experiment



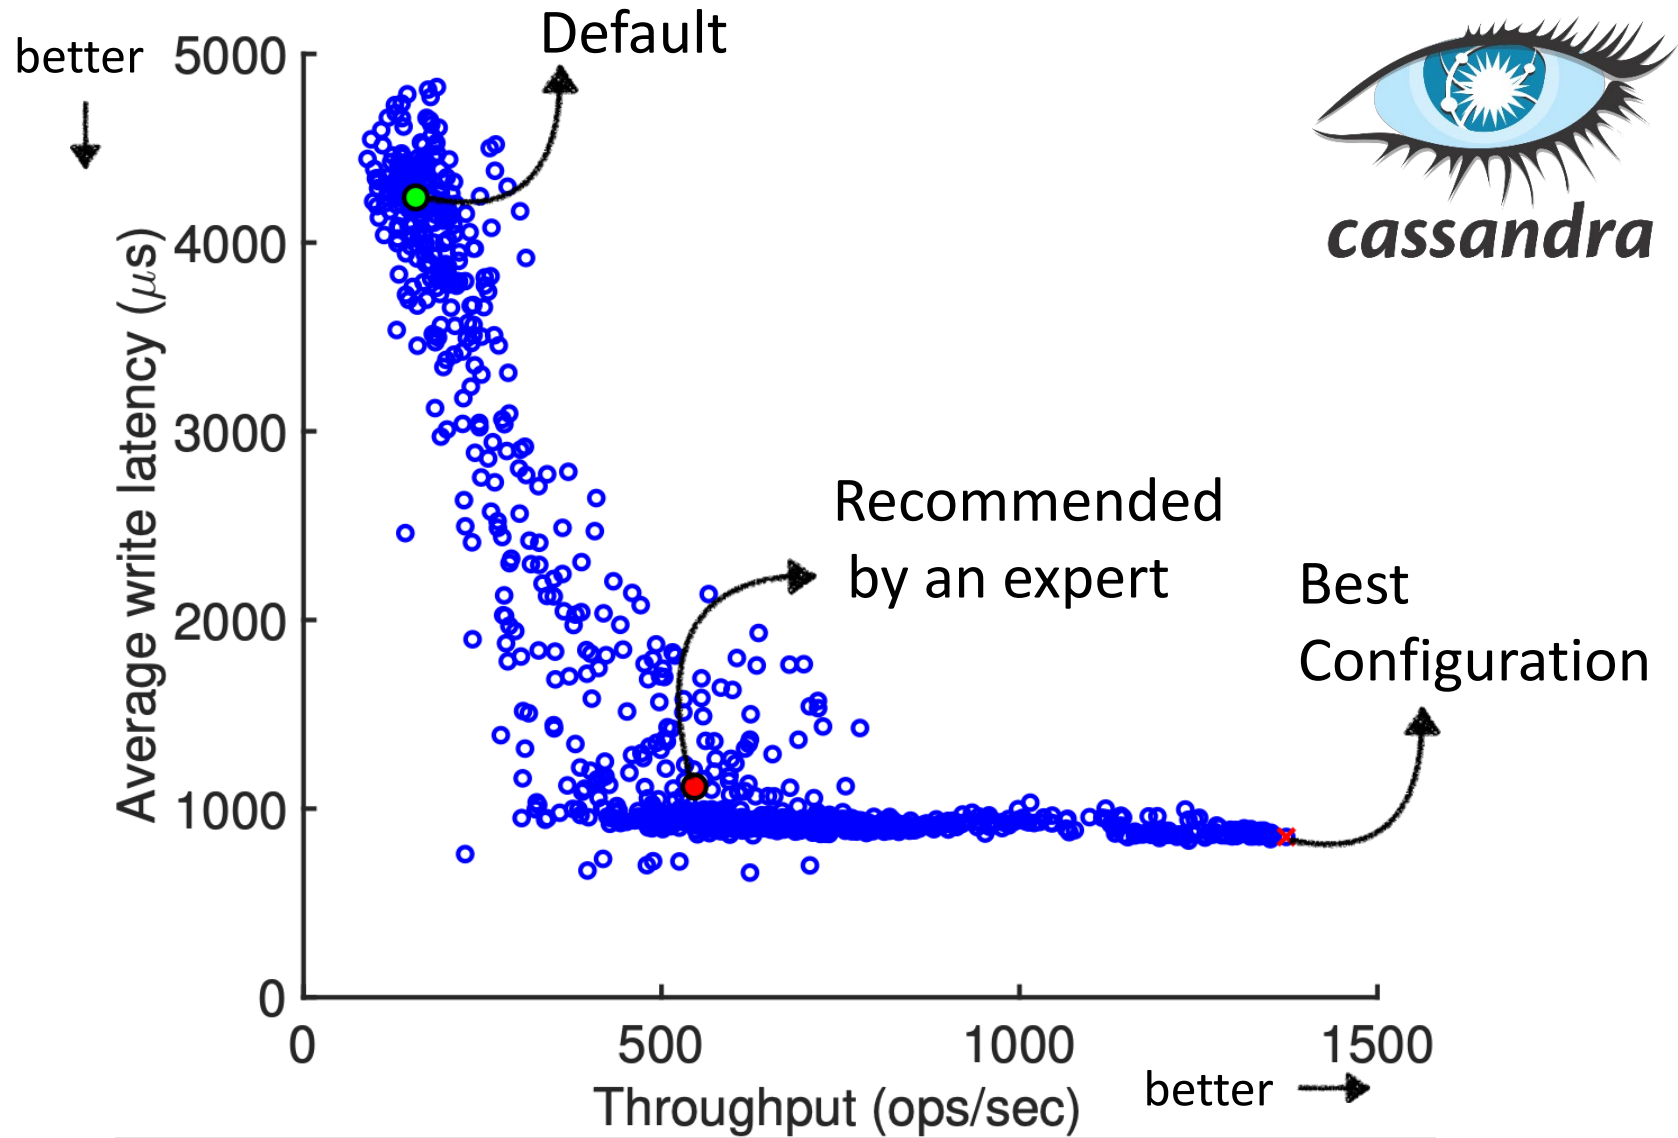
Representative  
workload

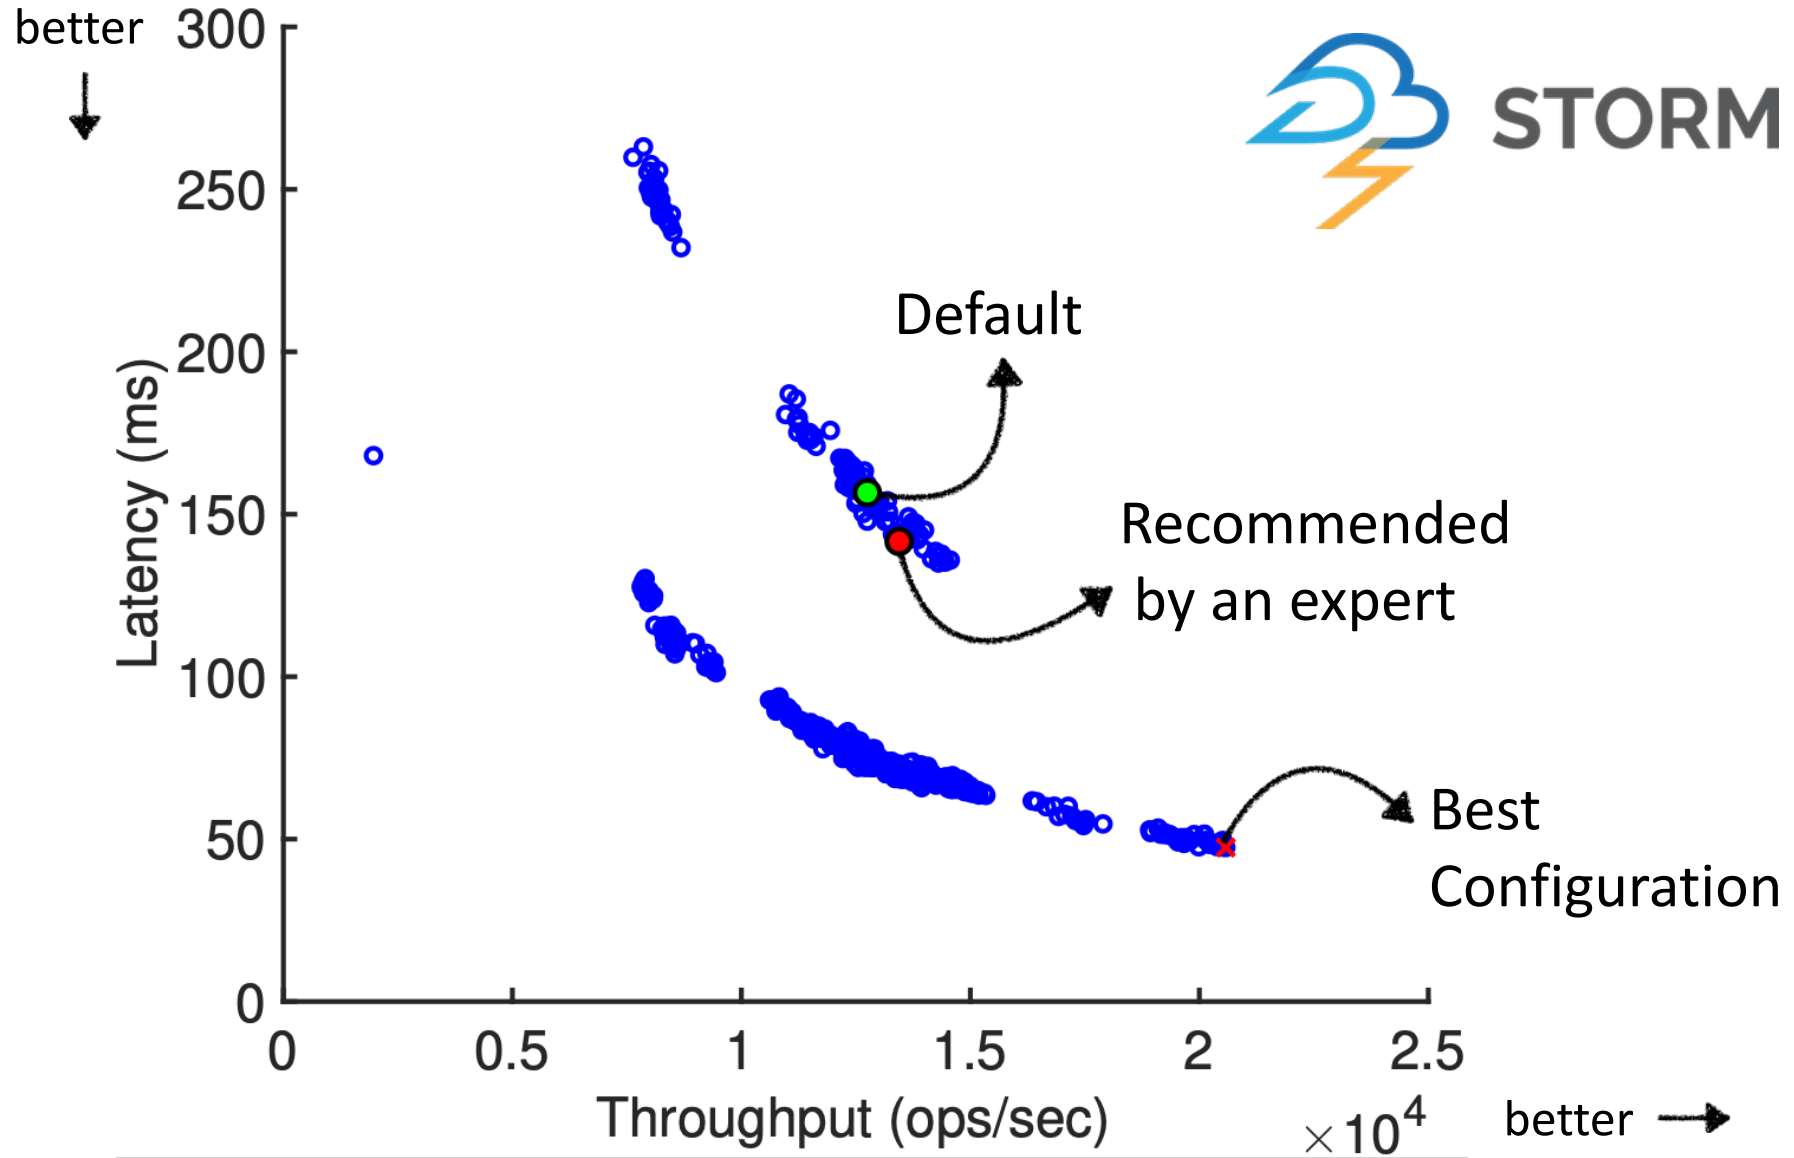


Change  
configuration



Performance  
measure  
of interest







**100X** more user  
cloud resources reduced 20%  
outperform expert recommendation



# Why Options?

Design for performance vs configuration

No one size fits all

Workload, hardware differs

Different users make different tradeoffs

Options are deferred design decisions



Software Config. | Workload C... | Hardware C..

Page Size

- 1K
- 2K
- 4K
- 8K

Index Structures

- Btree
- Queue
- Hash

Functional Features

- Encryption
- Diagnostic
- Transactions

Cache Size

### Binary Option: Encryption

Description: Encrypts data in the database file via a user-defined password.

#### Influence on Configuration

- Energy ↘ +5% (+ 50W/h; 131\$/year)
- Performance ↘ -15% (-530 Transactions/s)
- Footprint ↘ +18% (+329KB)
- Memory (peak) → +0,5% (250MB)
- Quality ↗ + Security

Configurations:  
 - New  
 - Last  
 - Default

#### Interacts with:

Database size

CPU frequency

# Example

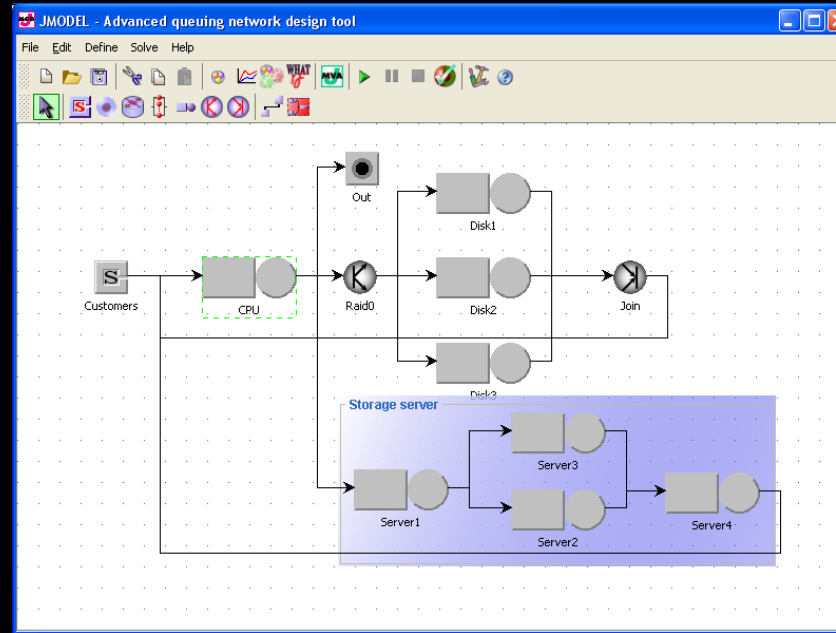
- Index
- Encryption
- Compression
- Logging

Write performance  
 Query performance  
 Scalability  
 Security

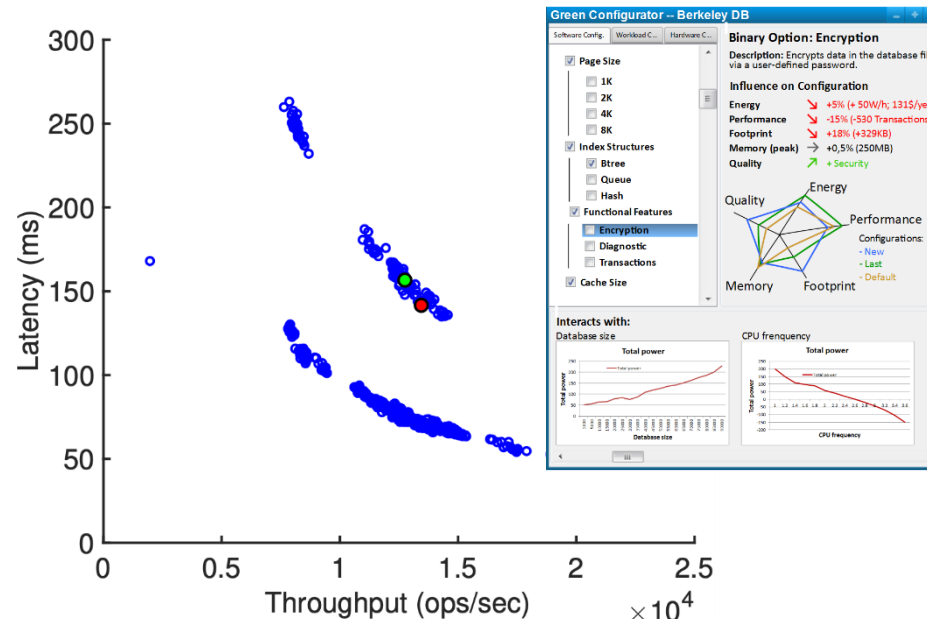


# Performance Analysis: Goals

# Performance Modeling for Design




# Optimization & Model Inference for Existing Code

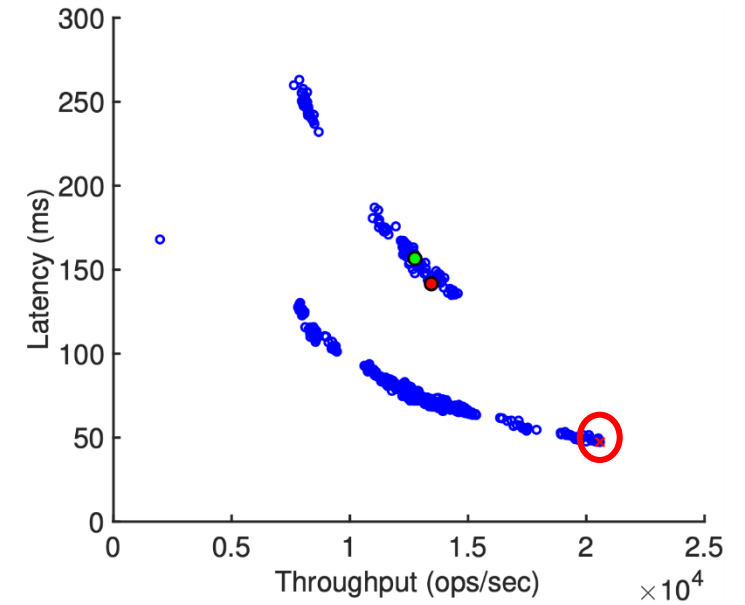




**Optimization**



# Optimization



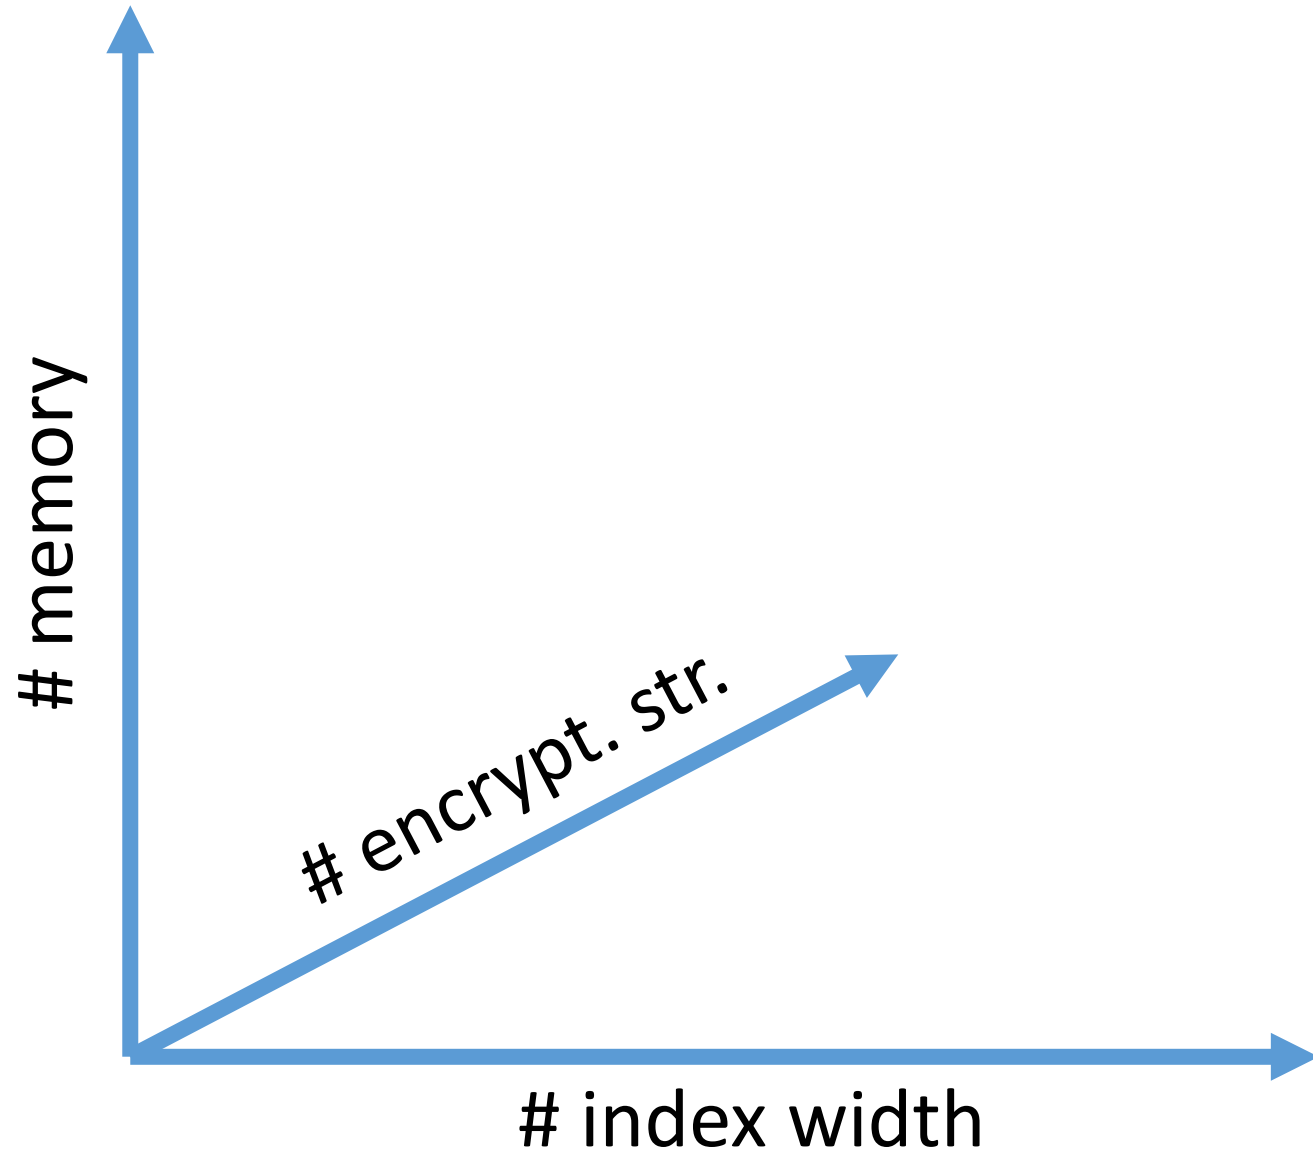
Goal: Find the fastest configuration or optimize a fitness function


Search problem, implementation given

Explored in many communities (databases, AI, systems, SE, ...)

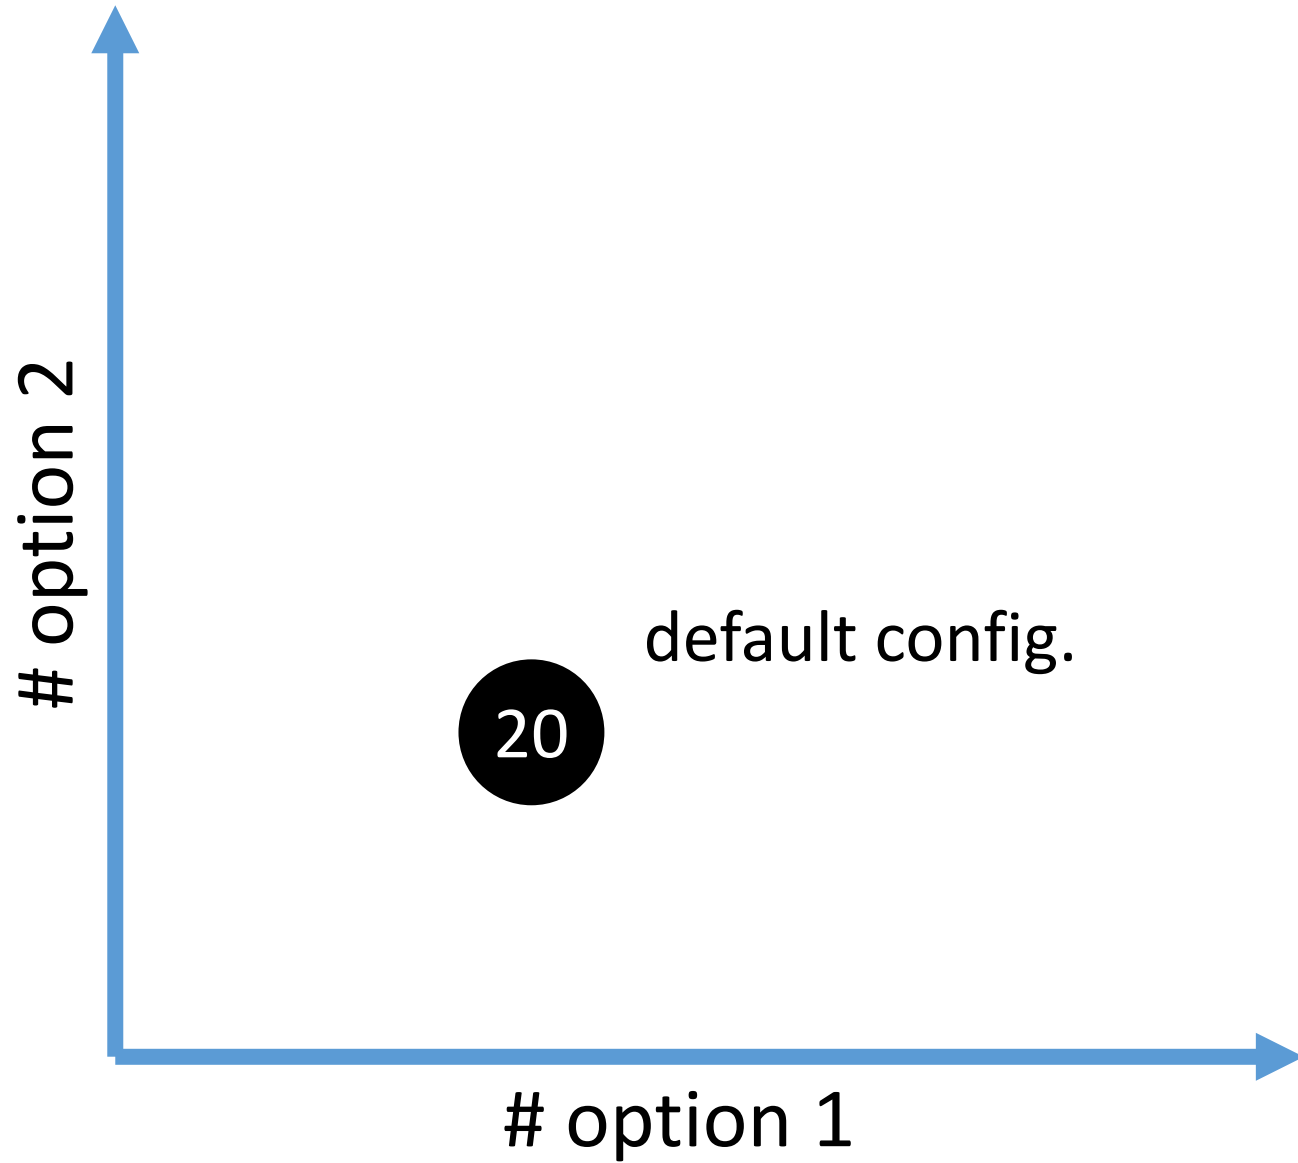
Auto-tuning, hyperparameter optimization, algorithm selection


Evaluated by performance improvement or distance to optimum

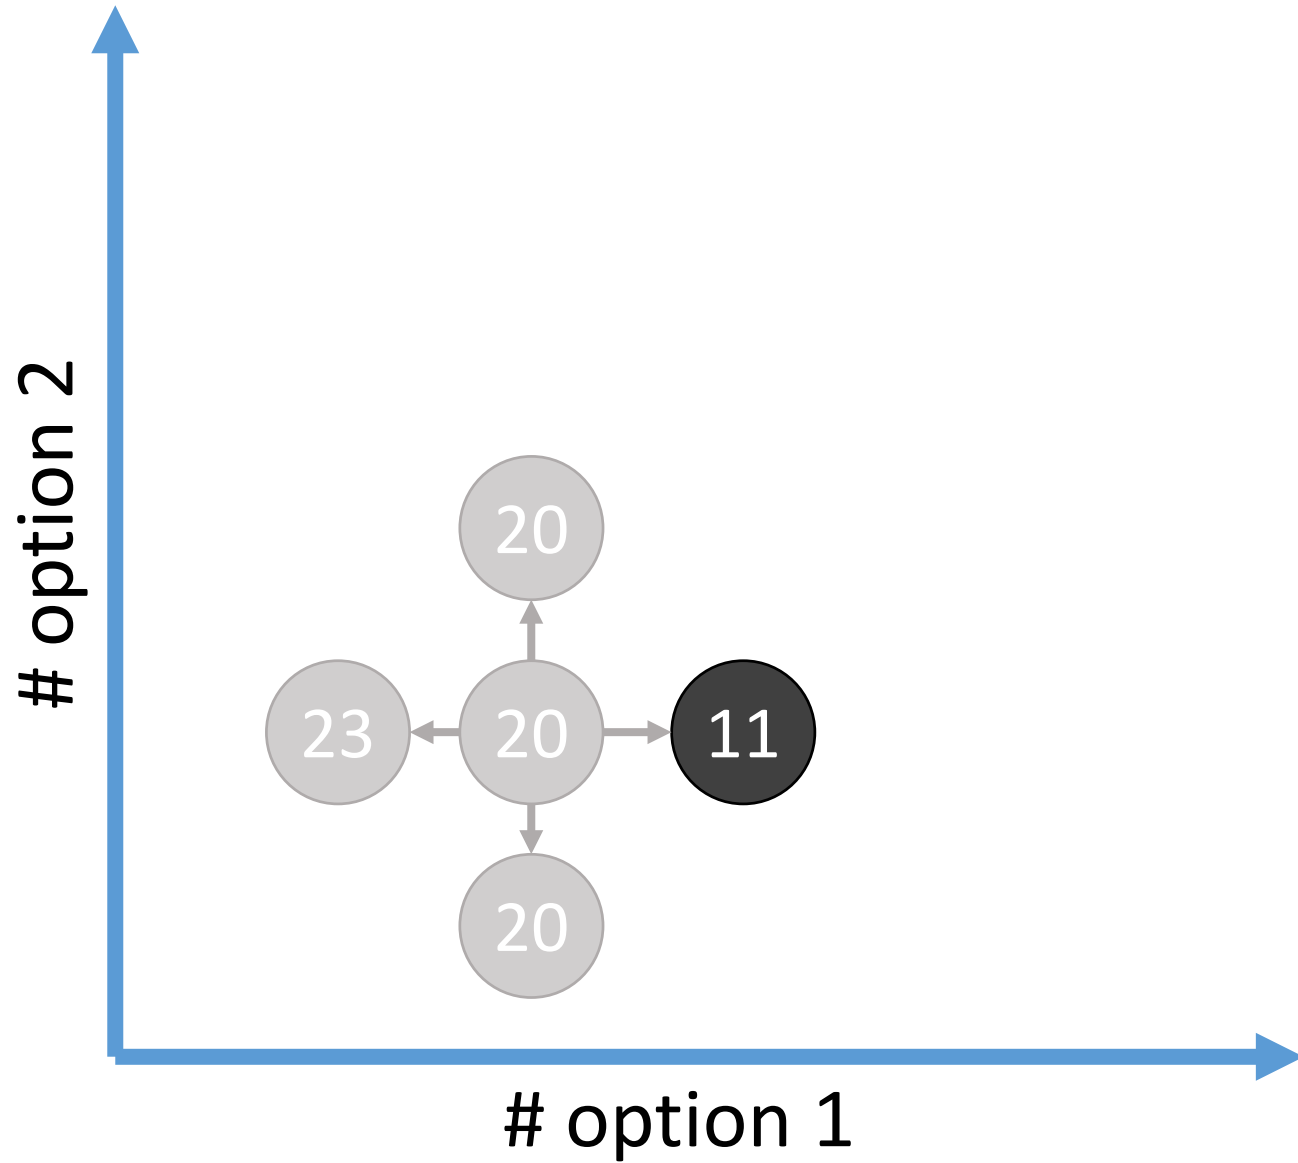


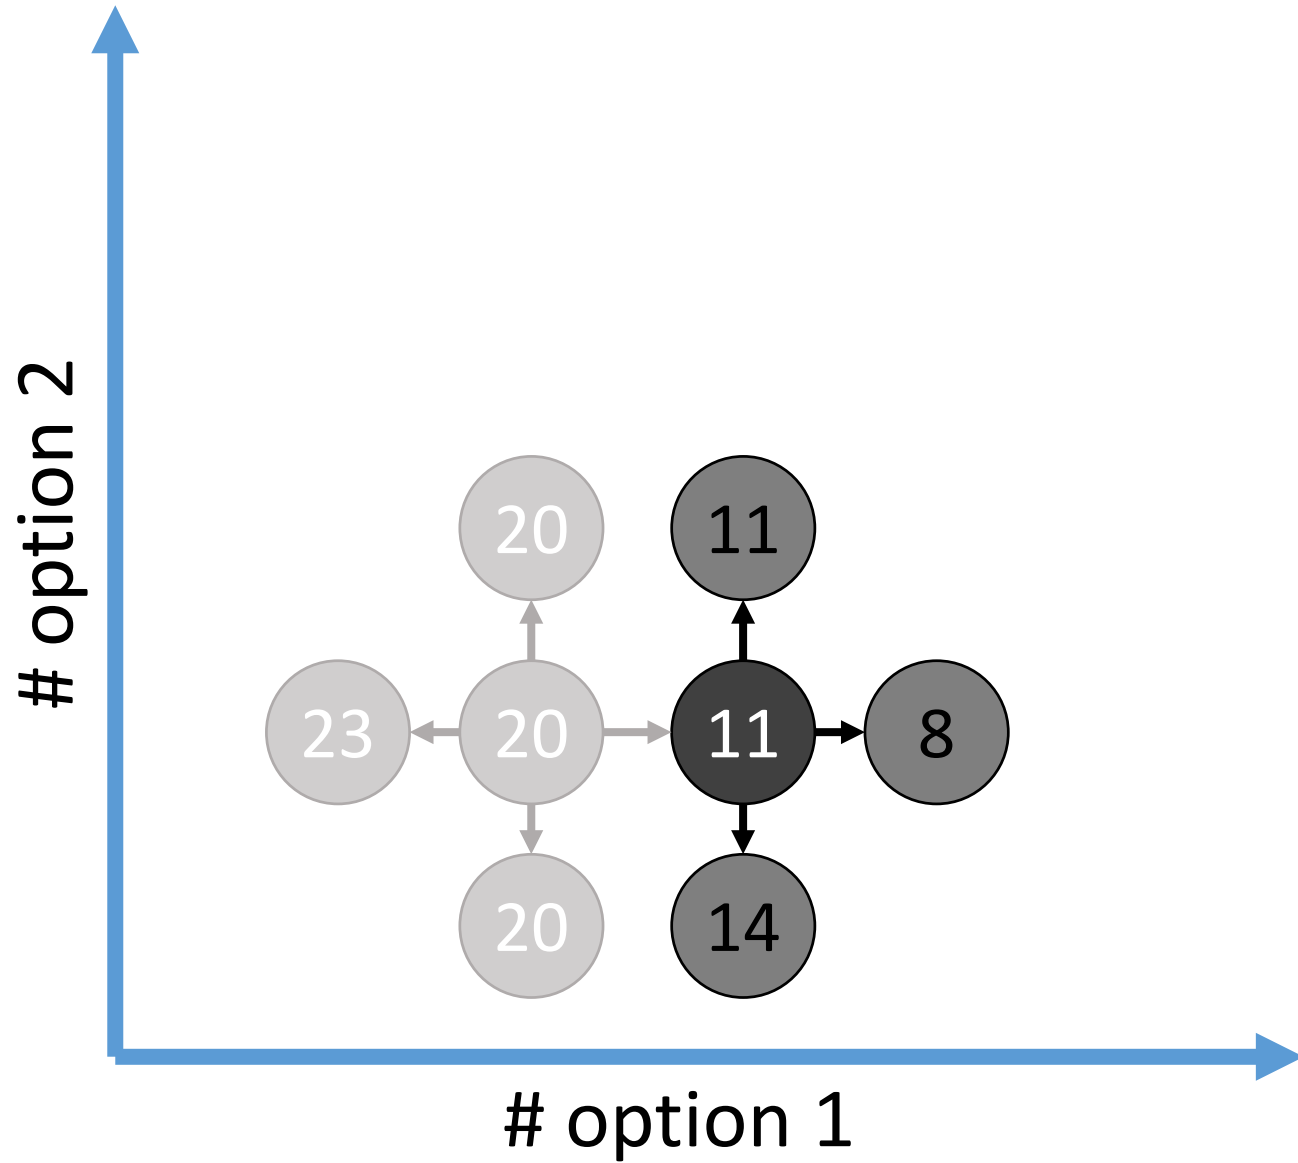




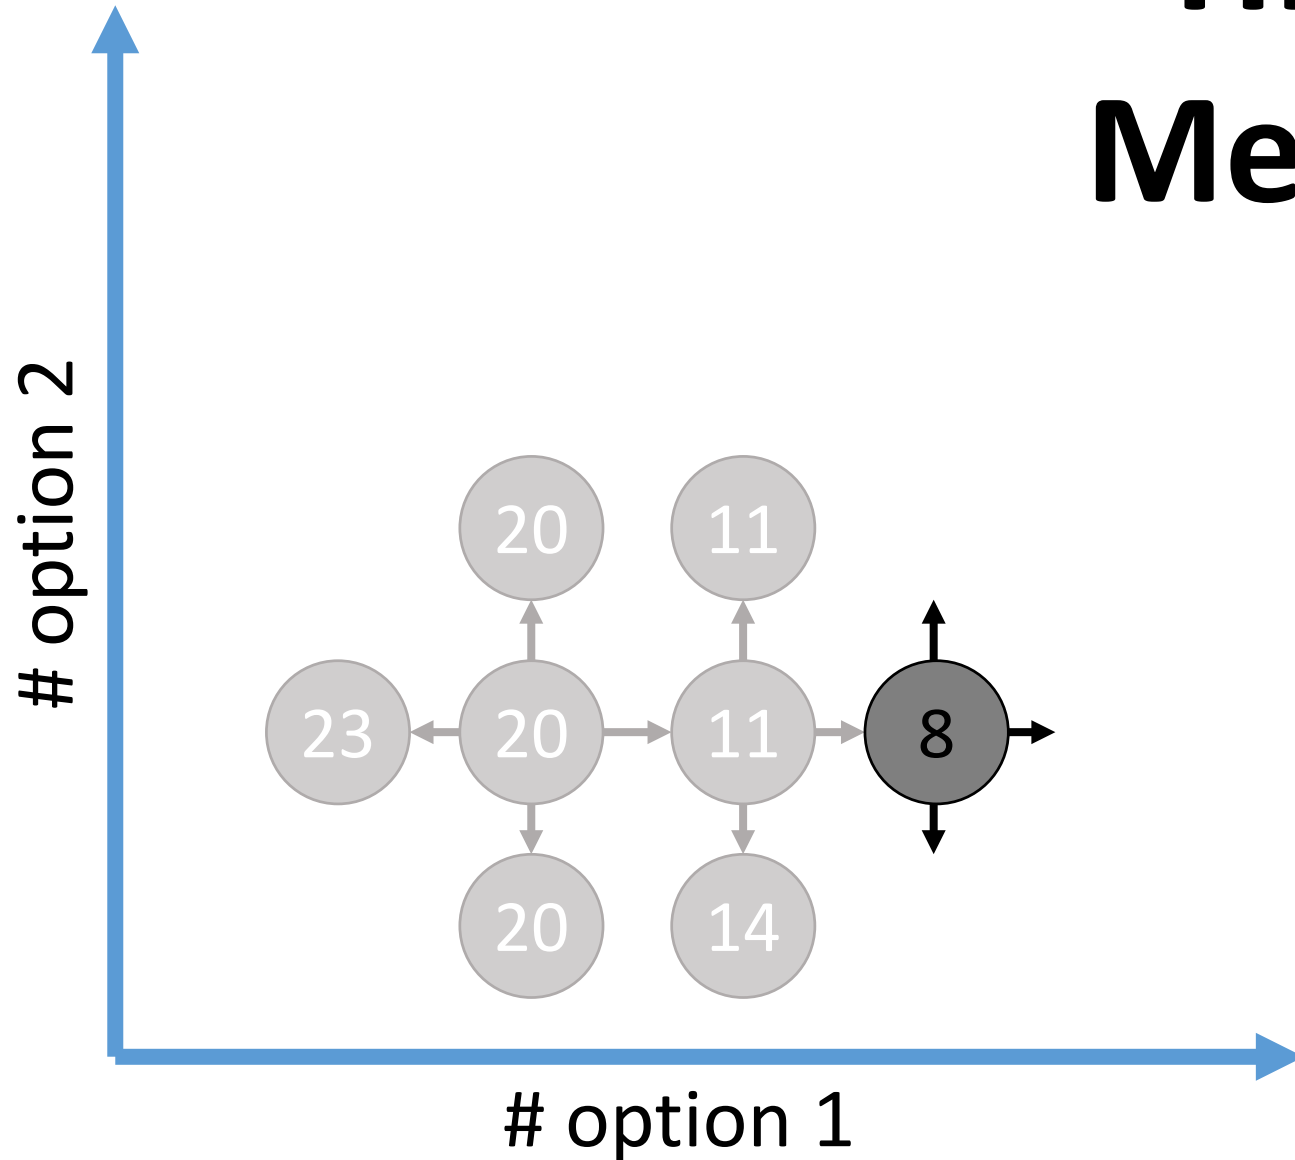


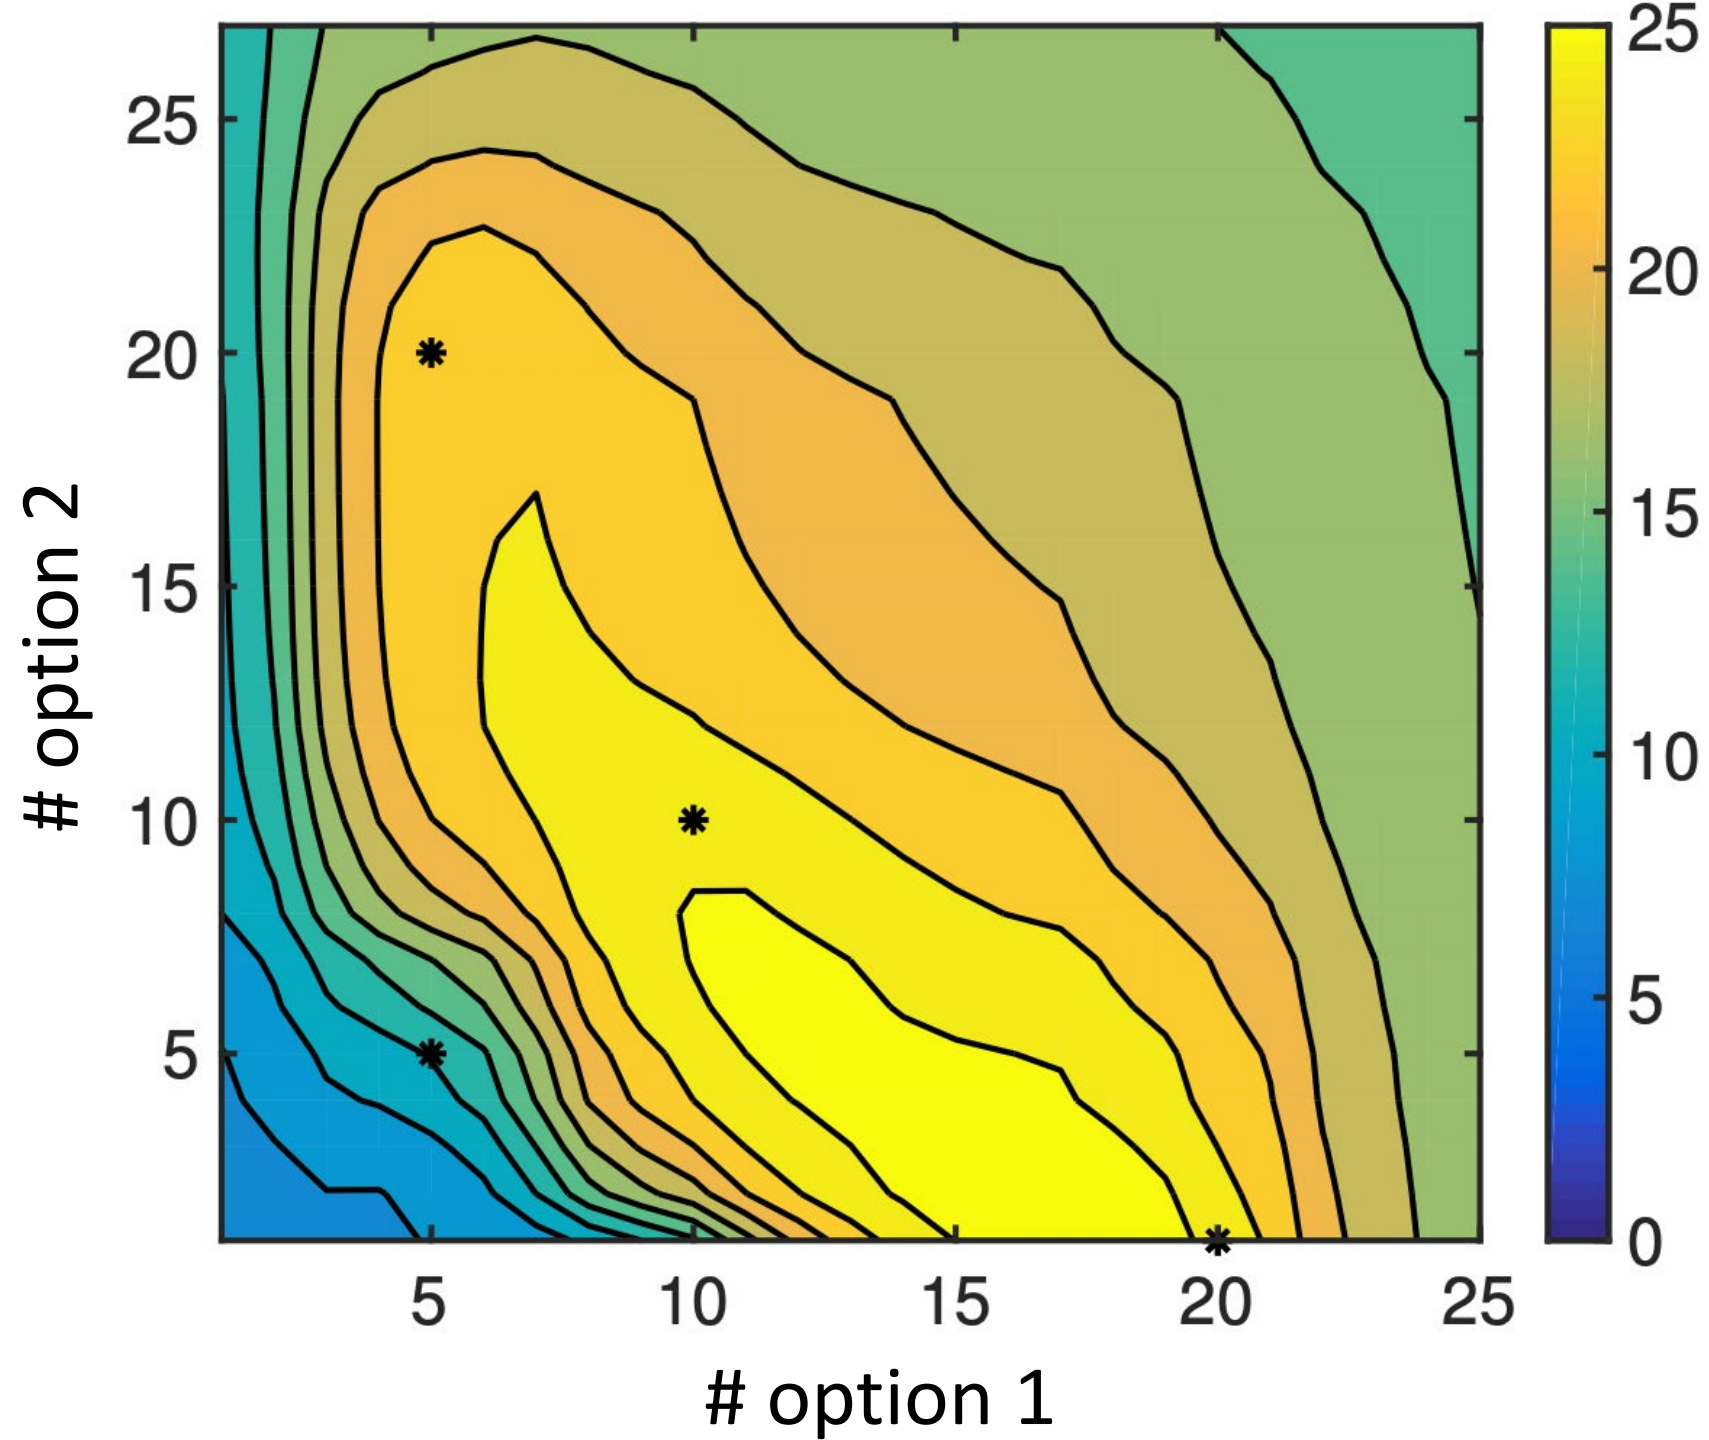






# Hill Climbing Metaheuristic





## Optimizing Selection of Competing Features via Feedback-Directed Evolutionary Algorithms



Tian Huat Tan<sup>†</sup> Yinxing Xue<sup>\*</sup> Manman Chen<sup>\*</sup> Jun Sun<sup>†</sup> Yang Liu<sup>‡</sup> Jin Song Dong<sup>\*</sup>  
<sup>†</sup>Singapore University of Technology and Design, Singapore  
<sup>\*</sup>National University of Singapore, Singapore  
<sup>‡</sup>Nanyang Technological University, Singapore

## Faster discovery of faster system configurations with spectral learning

Vivek Nair<sup>1</sup>  · Tim Menzies<sup>1</sup> ·  
Norbert Siegmund<sup>2</sup> · Sven Apel<sup>3</sup>

## Combining Evolutionary Algorithms with Constraint Solving for Configuration Optimization

## Sequential Model-Based Optimization for General Algorithm Configuration

Frank Hutter, Holger H. Hoos, and Kevin Leyton-Brown

## SIP: Optimal Product Selection from Feature Models Using Many-Objective Evolutionary Optimization

ROBERT M. HIERONS, MIQING LI, and XIAOHUI LIU, Brunel University London, UK  
SERGIO SEGURA, University of Seville, Spain  
WEI ZHENG, Northwestern Polytechnical University, China

## Comparison of Exact and Approximate Multi-Objective Optimization for Software Product Lines

Rafael Olaechea, Derek Rayside, Jianmei Guo, Krzysztof Czarnecki  
University of Waterloo  
Waterloo, Ontario  
{rolaechea, gjm,kczarnec}@gsd.uwaterloo.ca, {drayside}@uwaterloo.ca

## Random Search for Hyper-Parameter Optimization

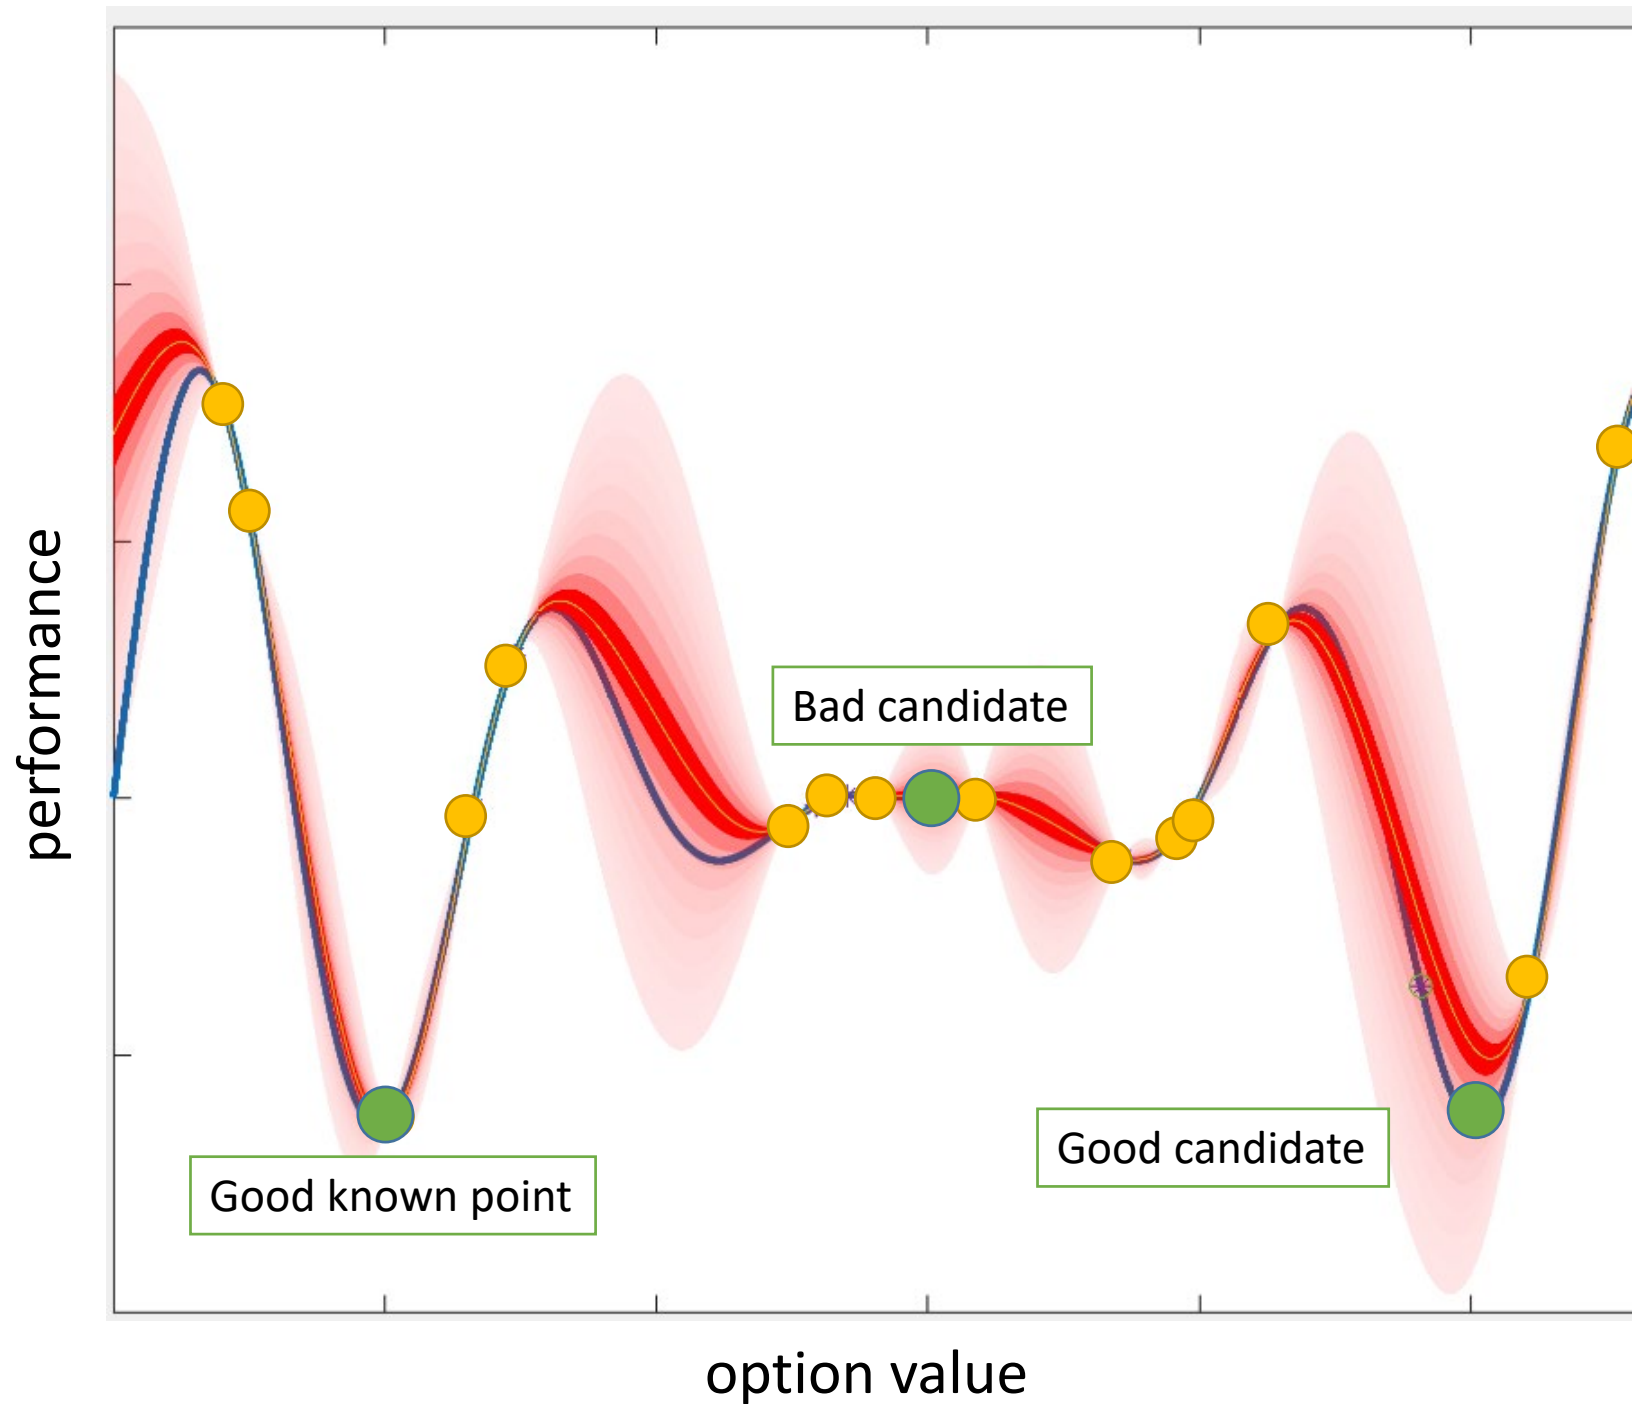
James Bergstra  
Yoshua Bengio

JAMES.BERGSTRA@UMONTREAL.CA  
YOSHUA.BENGIO@UMONTREAL.CA


## Hyperopt: A Python Library for Optimizing the Hyperparameters of Machine Learning Algorithms

James Bergstra<sup>\*†</sup>, Dan Yamins<sup>‡</sup>, David D. Cox<sup>§</sup>

# Gaussian processes to model uncertainty







# Model Inference

# Configuration

## Green Configurator -- Berkeley DB

Software Config. Workload C.. Hardware C..

**Page Size**

- 1K
- 2K
- 4K
- 8K

**Index Structures**

- Btree
- Queue
- Hash

**Functional Features**

- Encryption
- Diagnostic
- Transactions

**Cache Size**

### Binary Option: Encryption

Description: Encrypts data in the database file via a user-defined password.

#### Influence on Configuration

|               |                              |
|---------------|------------------------------|
| Energy        | ↘ +5% (+ 50W/h; 131\$/year)  |
| Performance   | ↘ -15% (-530 Transactions/s) |
| Footprint     | ↘ +18% (+329KB)              |
| Memory (peak) | → +0,5% (250MB)              |
| Quality       | ↗ + Security                 |

Configurations:  
- New  
- Last  
- Default

### Interacts with:

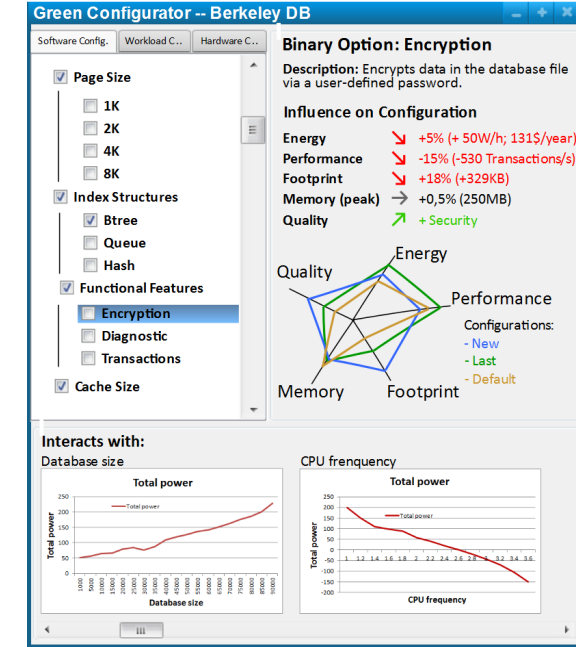
#### Database size

| Database size | Total power |
|---------------|-------------|
| 1,000         | 50          |
| 5,000         | 60          |
| 10,000        | 70          |
| 15,000        | 80          |
| 20,000        | 85          |
| 25,000        | 90          |
| 30,000        | 100         |
| 35,000        | 110         |
| 40,000        | 120         |
| 45,000        | 130         |
| 50,000        | 140         |
| 55,000        | 150         |
| 60,000        | 160         |
| 65,000        | 170         |
| 70,000        | 180         |
| 75,000        | 190         |
| 80,000        | 200         |
| 85,000        | 210         |
| 90,000        | 230         |

#### CPU frequency

| CPU frequency | Total power |
|---------------|-------------|
| 1.0           | 200         |
| 1.2           | 150         |
| 1.4           | 120         |
| 1.6           | 100         |
| 1.8           | 90          |
| 2.0           | 80          |
| 2.2           | 70          |
| 2.4           | 60          |
| 2.6           | 50          |
| 2.8           | 40          |
| 3.0           | 30          |
| 3.2           | 20          |
| 3.4           | 10          |
| 3.6           | -180        |

# Model Inference





Goal: Predict configuration's performance or explain performance influence of options

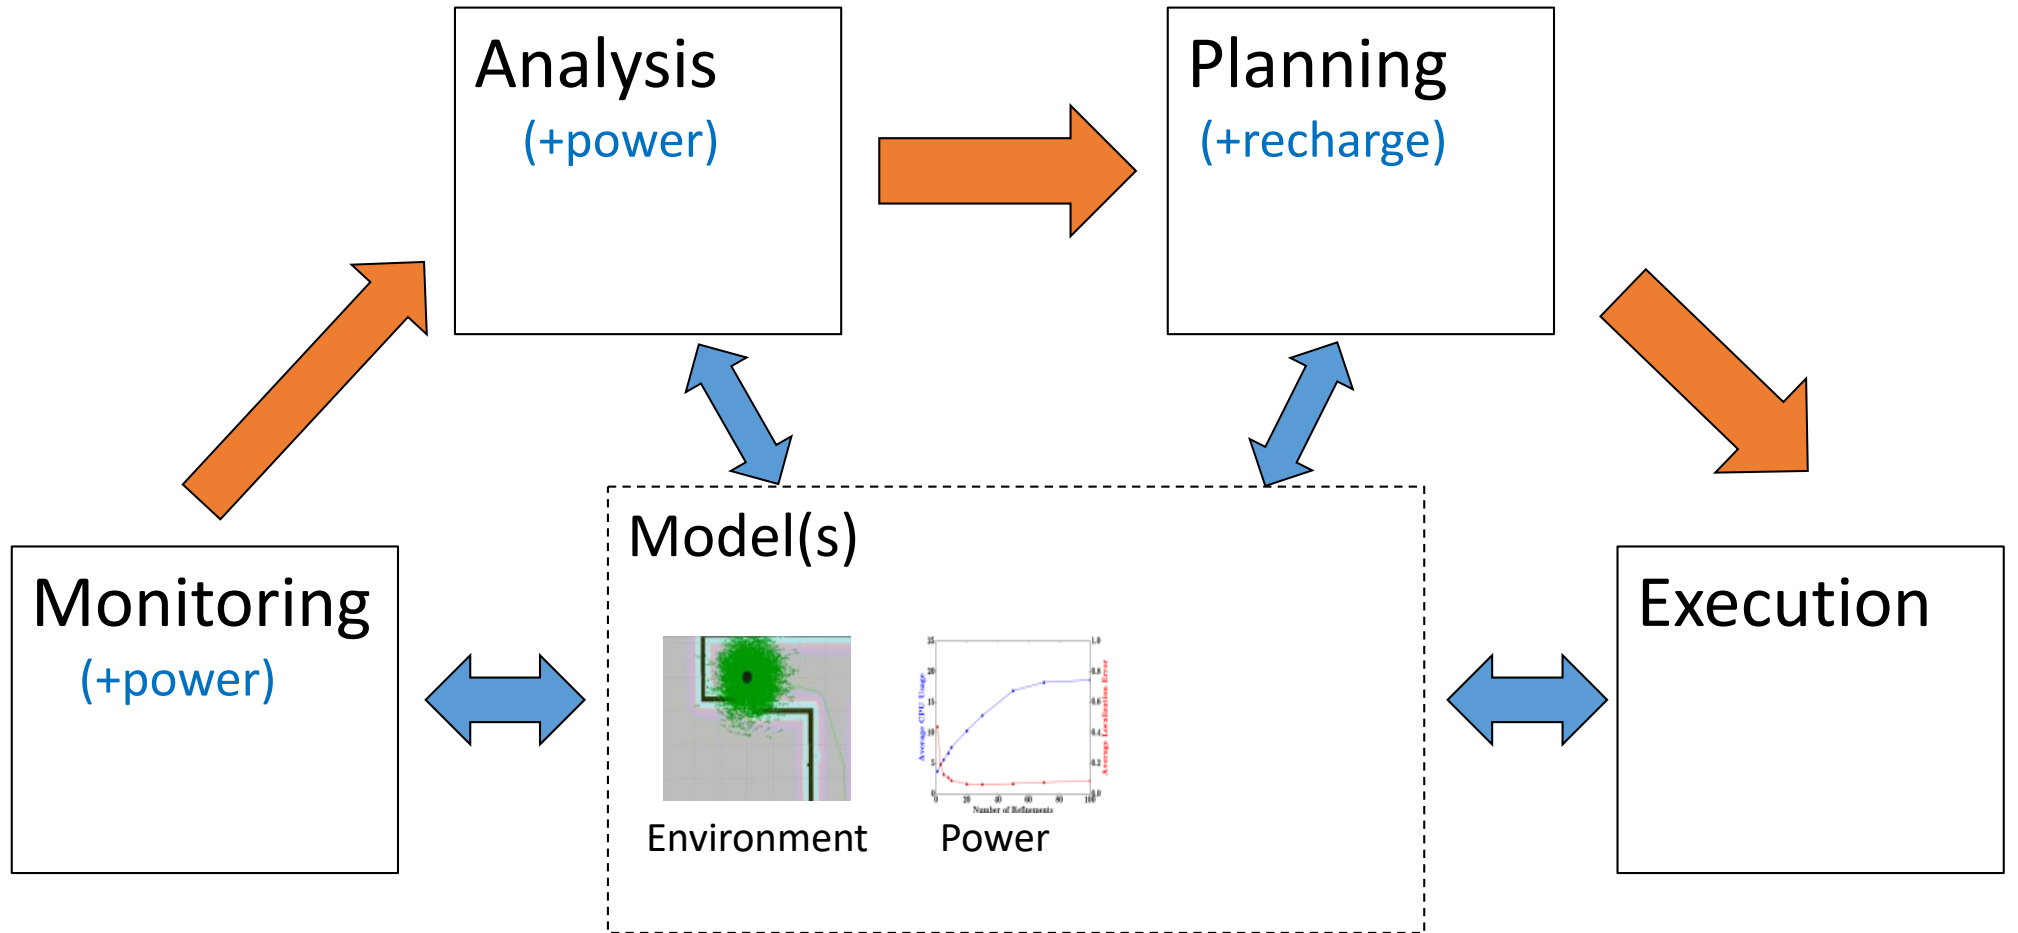
Given implementation, use cases: Understanding, debugging, planning, ...

Typically sampling + learning, various techniques

Evaluated by prediction accuracy



# Planning



# Debugging

**Binary Configuration Options**

**Coarse-grid solver**

- IP\_CG
- RED\_AMG
- IP\_AMG


**Smoother**

- Jac
- GS
- GSAC
- RBGS
- RBGSAC
- BS

**Numeric Configuration Options**

**Pre-smoothing Steps:** 0 6

**Post-smoothing Steps:** 0 6






“the HIPA<sup>CC</sup> expert was surprised to see that *pixelsPerThread* configuration option had only a small influence on system performance”

# Sensitivity Analysis



# One-At-A-Time Sensitivity Analysis

Options

|  |  |  |  |  | Perf. |
|---|---|---|---|---|-------|
| 1   | 0   | 0   | 1   | 0   | 10.3  |





Configurations




# Random Sampling + Regression



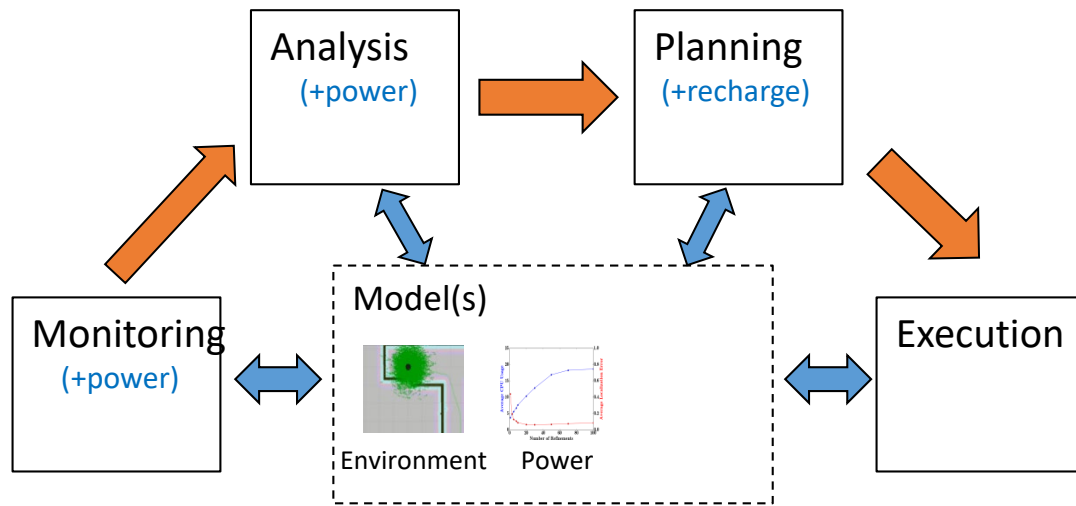
Individual Options

|                |  |  |  |  |  | Perf. |
|----------------|---|---|---|---|---|-------|
| Configurations | 1   | 0   | 0   | 1   | 0   | 10.3  |
|                | 0   | 1   | 1   | 0   | 1   | 55.8  |
|                | 1   | 0   | 1   | 1   | 1   | 21.9  |
|                | 1   | 1   | 0   | 1   | 1   | 60.1  |
|                | 1   | 1   | 1   | 0   | 0   | 59.3  |
|                | 0   | 0   | 0   | 1   | 0   | 29.3  |
|                | 1   | 1   | 1   | 1   | 1   | 45.6  |

$\beta_0$     $\beta_1$     $\beta_2$     $\beta_3$     $\beta_4$     $\beta_5$    Error:

 36.5   -24.3   15.3   -2.5   5.4   0.3   21%

$$p = \beta_0 + \beta_1 * \text{index} + \beta_2 * \text{encr} + \beta_3 * \text{compr} + \dots$$



### Green Configurator -- Berkeley DB

Software Config. | Workload C... | Hardware C...

#### Binary Option: Encryption

Description: Encrypts data in the database file via a user-defined password.

#### Influence on Configuration

|               |                              |
|---------------|------------------------------|
| Energy        | ↘ +5% (+ 50W/h; 131\$/year)  |
| Performance   | ↘ -15% (-530 Transactions/s) |
| Footprint     | ↘ +18% (+329KB)              |
| Memory (peak) | → +0,5% (250MB)              |
| Quality       | ↗ + Security                 |

Configurations:  
 - New (blue)  
 - Last (green)  
 - Default (orange)

**Configuration Options:**

- Page Size
  - 1K
  - 2K
  - 4K
  - 8K
- Index Structures
  - Btree
  - Queue
  - Hash
- Functional Features
  - Encryption
  - Diagnostic
  - Transactions
- Cache Size

#### Interacts with:

**Database size**

**CPU frequency**

$$p = 36.5 - 24.3 * \text{index} + 15.3 * \text{encr} - 2.5 * \text{compr} + \dots$$

# Interactions?





Benchmark execution: 30 sec



+ Index

Benchmark execution: 10 sec

Index reduces by 20 sec



+ Encryption

Benchmark execution: 35 sec

Encryption adds 5 sec


$$p = 30 - 20 * \text{index} + 5 * \text{encr}$$

Benchmark execution: 30 sec



+ Index

Benchmark execution: 10 sec

Index reduces by 20 sec



+ Encryption

Benchmark execution: 35 sec

Encryption adds 5 sec



+ Index  
+ Encryption

Benchmark execution: 50 sec


$$p = 30 - 20 * \text{index} + 5 * \text{encr} + 25 * \text{index} * \text{encr} + \dots$$



+ Index

Benchmark execution: 10 sec

Index reduces by 20 sec



+ Encryption

Benchmark execution: 35 sec

Encryption adds 5 sec



+ Index  
+ Encryption

Benchmark execution: 50 sec

# Random Sampling + Regression



| Configurations | Individual Options |           |           |           |           | Interactions |           |           | Perf.     |        |
|----------------|--------------------|-----------|-----------|-----------|-----------|--------------|-----------|-----------|-----------|--------|
|                |                    |           |           |           |           |              |           |           |           |        |
|                | 1                  | 0         | 0         | 1         | 0         | 0            | 0         | 10.3      |           |        |
|                | 0                  | 1         | 1         | 0         | 1         | 0            | 0         | 55.8      |           |        |
|                | 1                  | 0         | 1         | 1         | 1         | 0            | 0         | 21.9      |           |        |
|                | 1                  | 1         | 0         | 1         | 1         | 1            | 0         | 60.1      |           |        |
|                | 1                  | 1         | 1         | 0         | 0         | 1            | 0         | 59.3      |           |        |
|                | 0                  | 0         | 0         | 1         | 0         | 0            | 0         | 29.3      |           |        |
|                | 1                  | 1         | 1         | 1         | 1         | 1            | 1         | 45.6      |           |        |
|                | $\beta_0$          | $\beta_1$ | $\beta_2$ | $\beta_3$ | $\beta_4$ | $\beta_5$    | $\beta_6$ | $\beta_7$ | $\beta_8$ | Error: |
|                | 36.5               | -24.3     | 15.3      | -2.5      | 5.4       | 0.3          |           |           |           | 21%    |
|                | 30.4               | -21.1     | 5.3       | -2.4      | 3.5       | 0.2          | 26.1      | -5.4      | 0.01      | 6%     |

$$p = \beta_0 + \beta_1 * \text{index} + \beta_2 * \text{encr} + \beta_3 * \text{compr} + \dots$$

# Random Sampling + Regression



Individual Options

|           |           |           |           |           |           |           | Perf.  |
|-----------|-----------|-----------|-----------|-----------|-----------|-----------|--------|
|           | 1         | 0         | 0         | 1         | 0         | 1         | 10.3   |
|           | 0         | 1         | 1         | 0         | 1         | 0         | 55.8   |
|           | 1         | 0         | 1         | 1         | 1         | 1         | 21.9   |
|           | 1         | 1         | 0         | 1         | 1         | 1         | 60.1   |
|           | 2         | 1         | 1         | 0         | 0         | 4         | 59.3   |
|           | 0         | 0         | 0         | 1         | 0         | 0         | 29.3   |
|           | 1         | 1         | 1         | 1         | 1         | 1         | 45.6   |
| $\beta_0$ | $\beta_1$ | $\beta_2$ | $\beta_3$ | $\beta_4$ | $\beta_5$ | $\beta_6$ | Error: |

$$p = \beta_0 + \beta_1 * \text{index} + \dots + \beta_6 * \text{index}^2$$



# Random Sampling + Regression



| Configurations | Individual Options |           |           |           |           | Interactions |           |           | Perf.     |        |
|----------------|--------------------|-----------|-----------|-----------|-----------|--------------|-----------|-----------|-----------|--------|
|                |                    |           |           |           |           |              |           |           |           |        |
|                | 1                  | 0         | 0         | 1         | 0         | 0            | 0         | 0         | 10.3      |        |
|                | 0                  | 1         | 1         | 0         | 1         | 0            | 0         | 0         | 55.8      |        |
|                | 1                  | 0         | 1         | 1         | 1         | 0            | 0         | 0         | 21.9      |        |
|                | 1                  | 1         | 0         | 1         | 1         | 1            | 0         | 1         | 60.1      |        |
|                | 1                  | 1         | 1         | 0         | 0         | 1            | 1         | 0         | 59.3      |        |
|                | 0                  | 0         | 0         | 1         | 0         | 0            | 0         | 0         | 29.3      |        |
|                | 1                  | 1         | 1         | 1         | 1         | 1            | 1         | 1         | 45.6      |        |
|                | $\beta_0$          | $\beta_1$ | $\beta_2$ | $\beta_3$ | $\beta_4$ | $\beta_5$    | $\beta_6$ | $\beta_7$ | $\beta_8$ | Error: |
|                | 36.5               | -24.3     | 15.3      | -2.5      | 5.4       | 0.3          |           |           |           | 21%    |
|                | 30.4               | -21.1     | 5.3       | -2.4      | 3.5       | 0.2          | 26.1      | -5.4      | 0.01      | 6%     |

$$p = \beta_0 + \beta_1 * \text{index} + \beta_2 * \text{encr} + \beta_3 * \text{compr} + \dots$$

320<sup>optional, independent</sup> options

more combinations than estimated  
atoms in the universe




# Random Sampling + Regression



| Configurations | Individual Options |           |           |           |           | Interactions |           |           | Perf.     |        |
|----------------|--------------------|-----------|-----------|-----------|-----------|--------------|-----------|-----------|-----------|--------|
|                |                    |           |           |           |           |              |           |           |           |        |
|                | 1                  | 0         | 0         | 1         | 0         | 0            | 0         | 0         | 10.3      |        |
|                | 0                  | 1         | 1         | 0         | 1         | 0            | 0         | 0         | 55.8      |        |
|                | 1                  | 0         | 1         | 1         | 1         | 0            | 0         | 0         | 21.9      |        |
|                | 1                  | 1         | 0         | 1         | 1         | 1            | 0         | 1         | 60.1      |        |
|                | 1                  | 1         | 1         | 0         | 0         | 1            | 1         | 0         | 59.3      |        |
|                | 0                  | 0         | 0         | 1         | 0         | 0            | 0         | 0         | 29.3      |        |
|                | 1                  | 1         | 1         | 1         | 1         | 1            | 1         | 1         | 45.6      |        |
|                | $\beta_0$          | $\beta_1$ | $\beta_2$ | $\beta_3$ | $\beta_4$ | $\beta_5$    | $\beta_6$ | $\beta_7$ | $\beta_8$ | Error: |
|                | 36.5               | -24.3     | 15.3      | -2.5      | 5.4       | 0.3          |           |           |           | 21%    |
|                | 30.4               | -21.1     | 5.3       | -2.4      | 3.5       | 0.2          | 26.1      | -5.4      | 0.01      | 6%     |


$$p = \beta_0 + \beta_1 * \text{index} + \beta_2 * \text{encr} + \beta_3 * \text{compr} + \dots$$

# Feature Selection



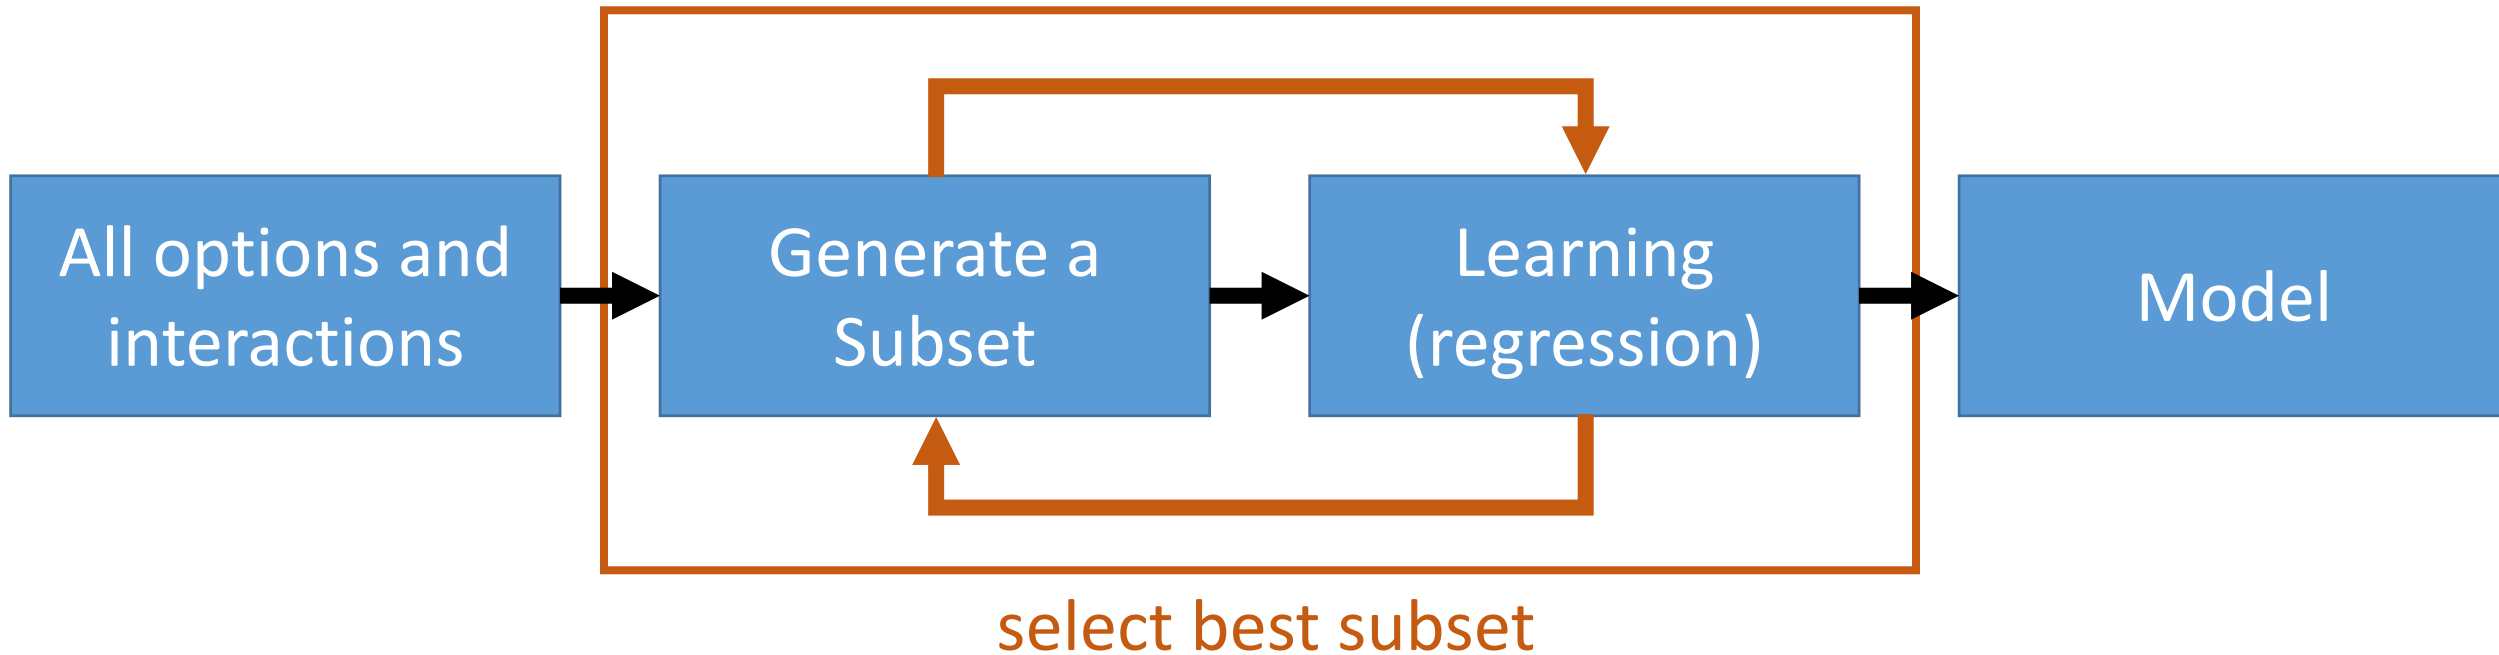
Goal: Select small set of features that best explain observations

# Feature Selection



Goal: Select small set of features that best explain observations

# Feature Selection



heuristics: try individual options first, then combinations of selected options


Goal: Select small set of features that best explain observations

# Round 1



# Round 1

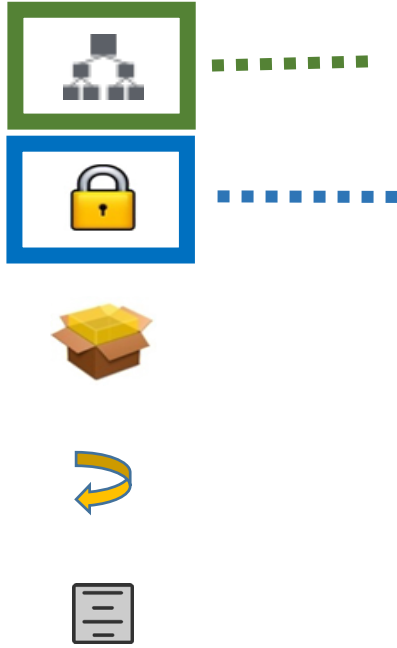


|           |  | Perf.  |
|-----------|---|--------|
|           | 1   | 10.3   |
|           | 0   | 55.8   |
|           | 1   | 21.9   |
|           | 1   | 60.1   |
|           | 1   | 59.3   |
|           | 0   | 29.3   |
|           | 1   | 45.6   |
| $\beta_0$ | $\beta_1$   | Error: |

**$p = 36.5 - 24.3 * \text{index}$**



Round 1



|           |   | Perf. |
|-----------|---|-------|
| $\beta_1$ | 1 | 10.3  |
| $\beta_0$ | 0 | 55.8  |
|           | 1 | 21.9  |
|           | 1 | 60.1  |
|           | 1 | 59.3  |
|           | 0 | 29.3  |
|           | 1 | 45.6  |

Error:

$p = 36.5 - 24.3 * \text{index}$

|           |   | Perf. |
|-----------|---|-------|
| $\beta_2$ | 0 | 10.3  |
| $\beta_0$ | 1 | 55.8  |
|           | 0 | 21.9  |
|           | 1 | 60.1  |
|           | 1 | 59.3  |
|           | 0 | 29.3  |
|           | 1 | 45.6  |

Error:

$p = 26.5 + 8.4 * \text{encr}$

Round 1



Round 1






Round 2



Round 1

Round 2




|           |  |  | Perf. |
|-----------|---|---|-------|
| $\beta_0$ | 1   | 0   | 10.3  |
| $\beta_1$ | 0   | 1   | 55.8  |
| $\beta_2$ | 1   | 0   | 21.9  |
|           | 1   | 1   | 60.1  |
|           | 1   | 1   | 59.3  |
|           | 0   | 0   | 29.3  |
|           | 1   | 1   | 45.6  |

Error:

$p = 32.5 - 19.3*i + \dots$

Round 1

Round 2



|   |   |   | Perf. |
|---|---|---|-------|
| 1 | 0 | 0 | 10.3  |
| 0 | 1 | 0 | 55.8  |
| 1 | 0 | 1 | 21.9  |
| 1 | 1 | 1 | 60.1  |
| 1 | 1 | 0 | 59.3  |
| 0 | 0 | 0 | 29.3  |
| 1 | 1 | 1 | 45.6  |

$\beta_0$     $\beta_1$     $\beta_2$    Error:

**$p = 32.5 - 19.3*i + \dots$**

|   |   |   | Perf. |
|---|---|---|-------|
| 1 | 0 | 0 | 10.3  |
| 0 | 1 | 0 | 55.8  |
| 1 | 0 | 1 | 21.9  |
| 1 | 1 | 1 | 60.1  |
| 1 | 1 | 0 | 59.3  |
| 0 | 0 | 0 | 29.3  |
| 1 | 1 | 1 | 45.6  |


$\beta_0$     $\beta_1$     $\beta_2$    Error:

**$p = 26.5 + 8.4*enc - \dots$**


Round 1



Round 2



Round 3



Round 4



heuristics: try individual options first, then combinations of selected options

$3.55 \cdot \text{ref} + 0.01 \cdot \text{keyint} + 0.05 \cdot \text{bframes}$

$1.87 \cdot \text{ref} + 0.10 \cdot \text{keyint} + 0.5 \cdot \text{bframes} + -$

$0.63 \cdot \text{ref} \cdot \text{no\_asm} + 0.19 \cdot \text{keyint} \cdot \text{no\_asm} + -$

$0.02 \cdot \text{bframes} \cdot \text{crfRatio} + 1.8\text{E}-$

$0.8 \cdot \text{keyint} \cdot \text{rc\_lookahead} \cdot \text{rc\_lookahead} +$

$0.0003 \cdot \text{bframes} \cdot \text{crfRatio} \cdot \text{crfRatio} +$

$0.0002 \cdot \text{qpSetting} \cdot \text{qpSetting} + 20.92 \cdot \text{no\_asm} + -$

$0.0004 \cdot \text{no\_asm} \cdot \text{keyint} \cdot \text{keyint} + 5.80\text{E}-$

$0.7 \cdot \text{bframes} \cdot \text{keyint} \cdot \text{keyint} + 0.50 \cdot \text{no\_asm} \cdot \text{crfRatio}$

# Empirical Observations




Few options responsible for most of the variation

Interactions are important to explain variation

Few options interact





**Performance  
behavior has  
“structure”**




# LLVM

computation time (minutes)

0.04 0.05 0.06 0.07 0.09 0.17 0.26 0.39 0.75 2.59

number of model terms

21  
18  
15  
12  
9  
6  
3  
0



term size

4  
3  
2  
1

prediction error (%)


4.4 4.0 3.6 3.2 2.8 2.4 2.0 1.6 1.2 0.8



# Apache

computation time (minutes)

0.004 0.01 0.010 0.015 0.021 0.029 0.042 0.084 0.2



$3.55 \cdot \text{ref} + 0.01 \cdot \text{keyint} + 0.05 \cdot \text{bframes}$

$1.87 \cdot \text{ref} + 0.10 \cdot \text{keyint} + 0.5 \cdot \text{bframes} + -$

$0.63 \cdot \text{ref} \cdot \text{no\_asm} + 0.19 \cdot \text{keyint} \cdot \text{no\_asm} + -$

$0.02 \cdot \text{bframes} \cdot \text{crfRatio} + 1.8\text{E}-$

$0.8 \cdot \text{keyint} \cdot \text{rc\_lookahead} \cdot \text{rc\_lookahead} +$

$0.0003 \cdot \text{bframes} \cdot \text{crfRatio} \cdot \text{crfRatio} +$

$0.0002 \cdot \text{qpSetting} \cdot \text{qpSetting} + 20.92 \cdot \text{no\_asm} + -$

$0.0004 \cdot \text{no\_asm} \cdot \text{keyint} \cdot \text{keyint} + 5.80\text{E}-$

$0.7 \cdot \text{bframes} \cdot \text{keyint} \cdot \text{keyint} + 0.50 \cdot \text{no\_asm} \cdot \text{crfRatio}$

# So Far: Model Inference

Sensitivity analysis: detect performance interactions from small samples

Blackbox approach, simple, quite accurate, works surprisingly well

Data and machine learning tools available

e.g.,  <http://fospd.net/SPLConqueror>



# Open Challenges

# Challenge: Finding Interactions




# Challenge: Measurement Costs






# Challenge: Workload/Environment Changes






# Recent Advances

# Better Learning



Various learning approaches explored:  
linear regression, Gaussian processes, decision trees, DNN, ...  
+ tuning of those


# Better Sampling



## PLACKETTE BURMAN

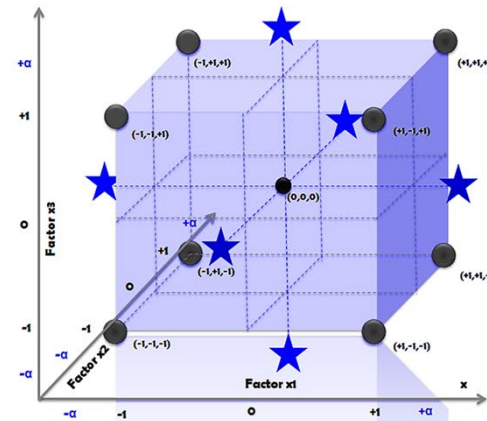
|    | X <sub>1</sub> | X <sub>2</sub> | X <sub>3</sub> | X <sub>4</sub> | X <sub>5</sub> | X <sub>6</sub> | X <sub>7</sub> | X <sub>8</sub> | X <sub>9</sub> | X <sub>10</sub> | X <sub>11</sub> |
|----|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|----------------|-----------------|-----------------|
| 1  | +1             | +1             | +1             | +1             | +1             | +1             | +1             | +1             | +1             | +1              | +1              |
| 2  | -1             | +1             | -1             | +1             | +1             | +1             | -1             | -1             | -1             | +1              | -1              |
| 3  | -1             | -1             | +1             | -1             | +1             | +1             | +1             | -1             | -1             | -1              | +1              |
| 4  | +1             | -1             | -1             | +1             | -1             | +1             | +1             | +1             | -1             | -1              | -1              |
| 5  | -1             | +1             | -1             | -1             | +1             | -1             | +1             | +1             | +1             | -1              | -1              |
| 6  | -1             | -1             | +1             | -1             | -1             | +1             | -1             | +1             | +1             | +1              | -1              |
| 7  | -1             | -1             | -1             | +1             | -1             | -1             | +1             | -1             | +1             | +1              | +1              |
| 8  | +1             | -1             | -1             | -1             | +1             | -1             | -1             | +1             | -1             | +1              | +1              |
| 9  | +1             | +1             | -1             | -1             | -1             | +1             | -1             | -1             | -1             | +1              | -1              |
| 10 | +1             | +1             | +1             | -1             | -1             | -1             | -1             | -1             | -1             | +1              | -1              |
| 11 | -1             | +1             | +1             | +1             | -1             | -1             | -1             | +1             | -1             | -1              | +1              |
| 12 | +1             | -1             | +1             | +1             | +1             | -1             | -1             | -1             | +1             | -1              | -1              |

## FRACTIONAL FACTORIAL




• Number of Runs:  $L^{F \pm f}$

## CENTRAL COMPOSITE



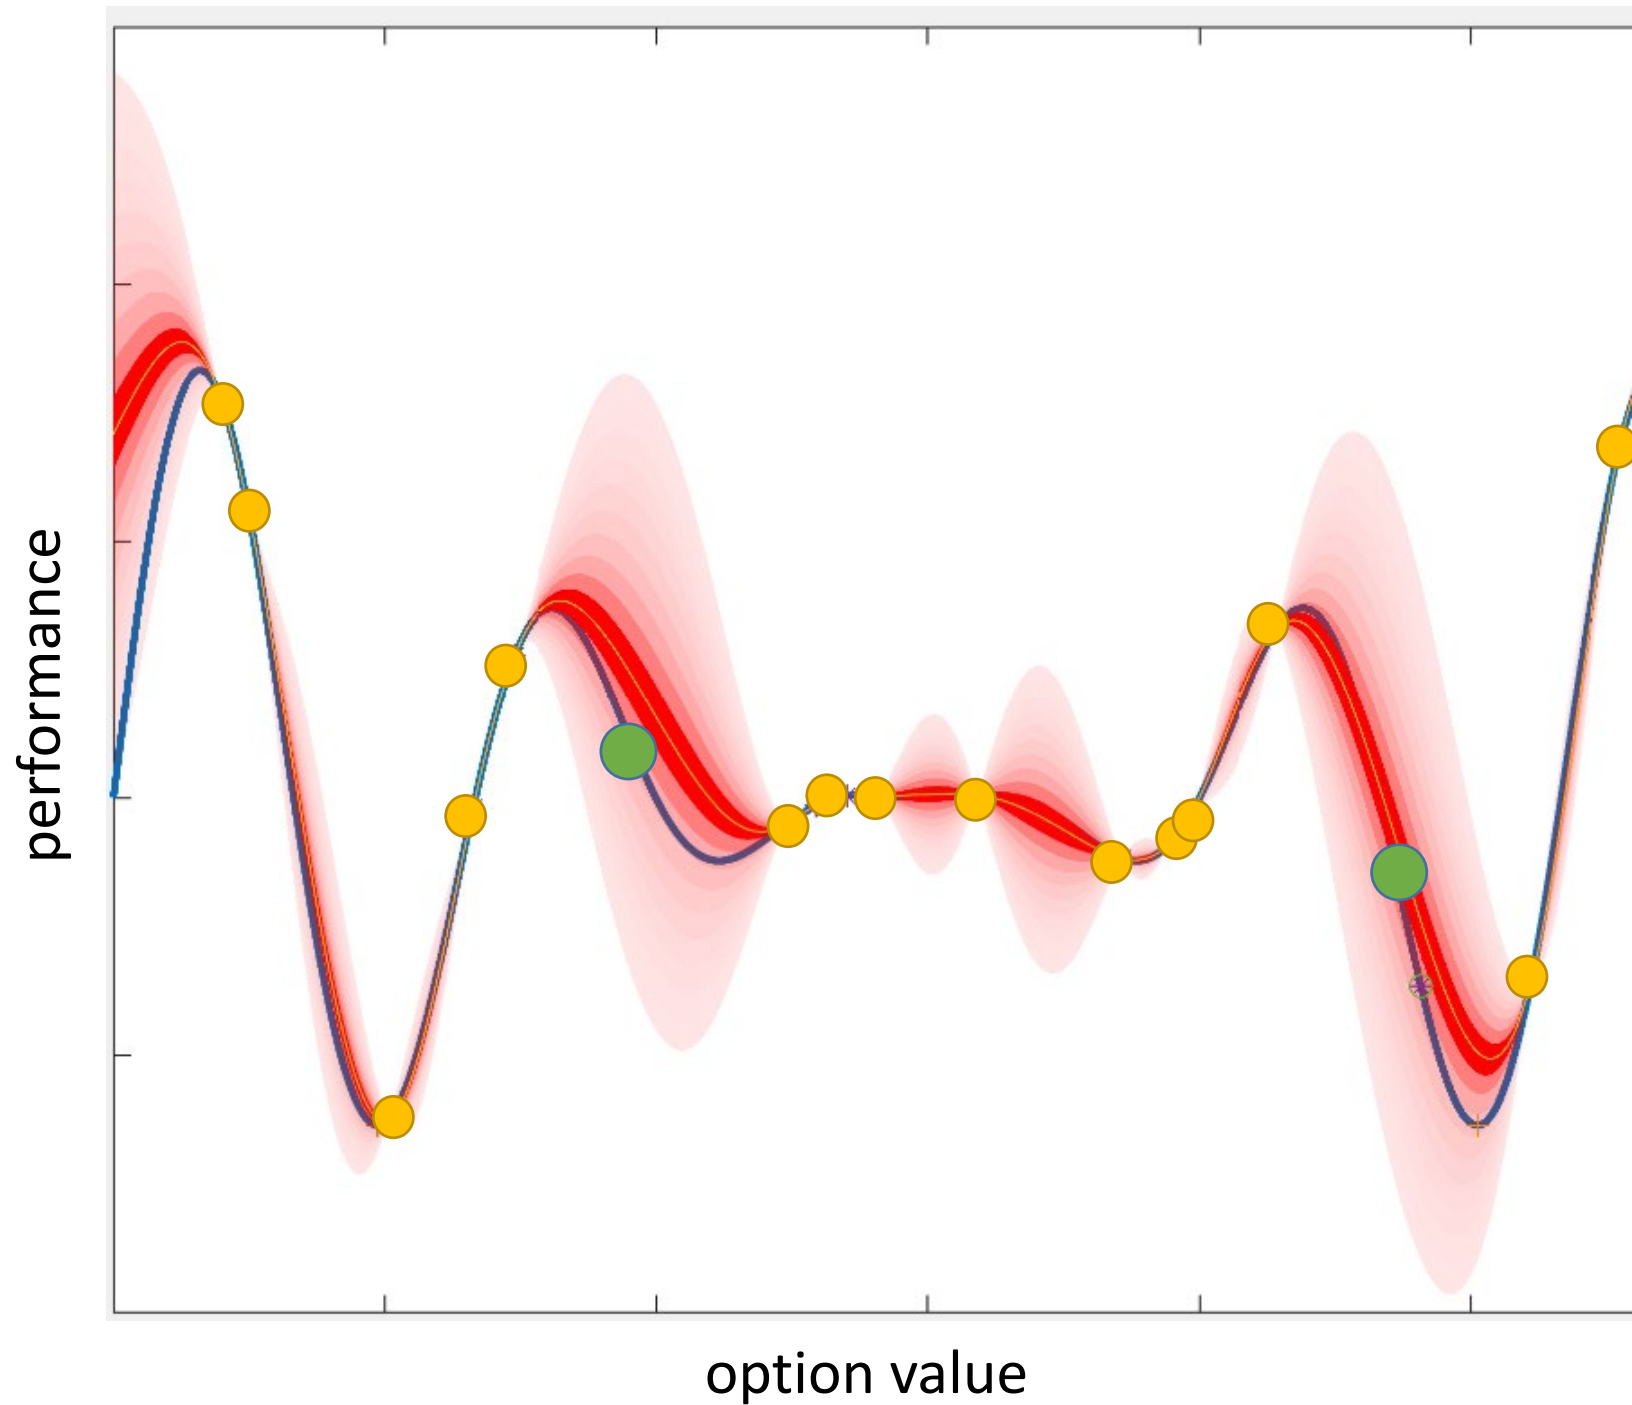
- Levels: 5  
[“-α”, -1, “0” and +1, +α’]
- No of Runs:  $2^{fp} + 2SP + CP$

## BOX BEHNKEN



- Levels: 3 levels per factor
- Design Points: at the “mid” points of edges of the process space & center



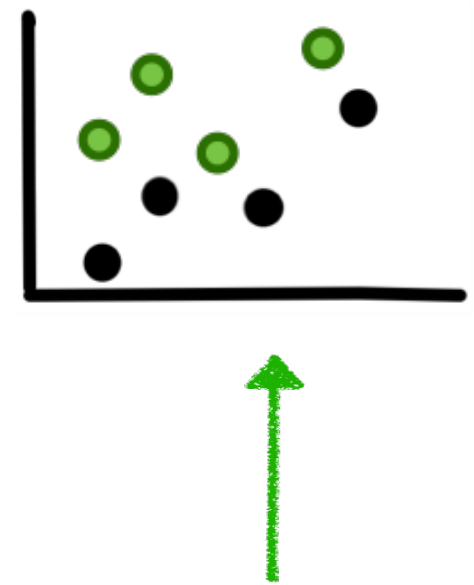
# Active Sampling




# Reducing Measurement Costs



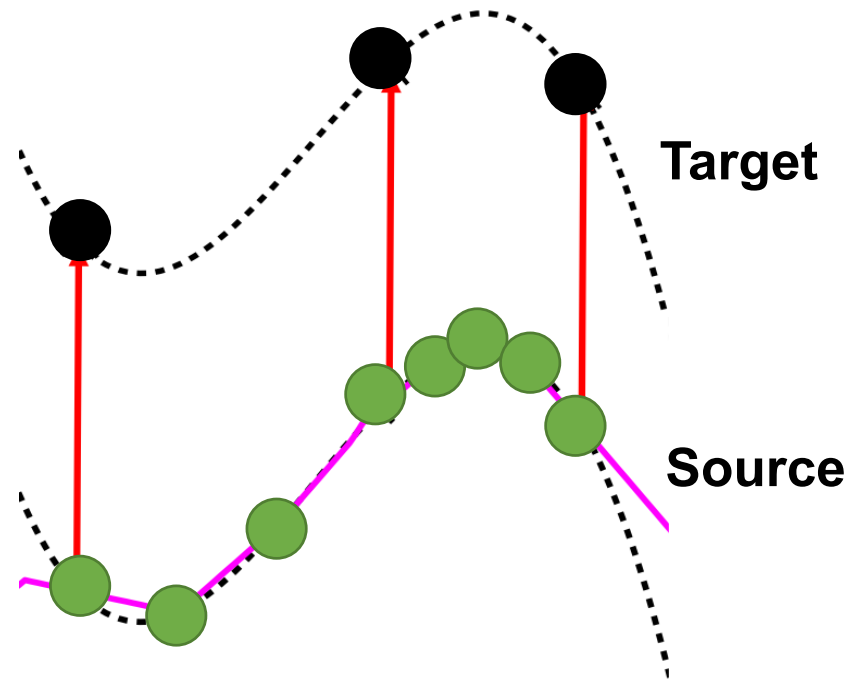
# Performance of real system is “similar” to performance in simulators



# Transfer Learning



Throughput ↑








# Workload and Hardware Changes

Most approaches assume fixed workload,  
learned models specific to benchmark

Model workload and hardware variability as additional options



# Transfer Learning





**TABLE II:** Results indicate that there exist several forms of knowledge that can be transferred across environments and can be used in transfer learning.

| Environment  | ES | RQ1         |             |             |             |             | RQ2  |    |      |    |             | RQ3       |           |      |             | RQ4         |             |          |             |      |
|--|----|-------------|-------------|-------------|-------------|-------------|------|----|------|----|-------------|-----------|-----------|------|-------------|-------------|-------------|----------|-------------|------|
|  |    | H1.1        |             | H1.2        | H1.3        |             | H1.4 |    | H2.1 |    |             | H2.2      |           | H3.1 |             | H3.2        |             | H4.1     |             | H4.2 |
|  |    | M1          | M2          | M3          | M4          | M5          | M6   | M7 | M8   | M9 | M10         | M11       | M12       | M13  | M14         | M15         | M16         | M17      | M18         |      |
| <b>SPEAR— Workload (#variables/#clauses): <math>w_1 : 774/5934, w_2 : 1008/7728, w_3 : 1554/11914, w_4 : 978/7498</math>; Version: <math>v_1 : 1.2, v_2 : 2.7</math></b> |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_1 : [h_2 \rightarrow h_1, w_1, v_2]$   | S  | <b>1.00</b> | 0.22        | <b>0.97</b> | <b>0.92</b> | <b>0.92</b> | 9    | 7  | 7    | 0  | 1           | 25        | 25        | 25   | 1.00        | 0.47        | 0.45        | 1        | 1.00        |      |
| $ec_2 : [h_4 \rightarrow h_1, w_1, v_2]$   | L  | 0.59        | 24.88       | <b>0.91</b> | <b>0.76</b> | <b>0.86</b> | 12   | 7  | 4    | 2  | 0.51        | <b>41</b> | <b>27</b> | 21   | <b>0.98</b> | 0.48        | 0.45        | 1        | 0.98        |      |
| $ec_3 : [h_1, w_1 \rightarrow w_2, v_2]$   | L  | <b>0.96</b> | 1.97        | 0.17        | 0.44        | 0.32        | 9    | 7  | 4    | 3  | 1           | 23        | 23        | 22   | 0.99        | 0.45        | 0.45        | 1        | 1.00        |      |
| $ec_4 : [h_1, w_1 \rightarrow w_3, v_2]$   | M  | <b>0.90</b> | 3.36        | -0.08       | 0.30        | 0.11        | 7    | 7  | 4    | 3  | 0.99        | 22        | 23        | 22   | 0.99        | 0.45        | 0.49        | 1        | 0.94        |      |
| $ec_5 : [h_1, w_1, v_2 \rightarrow v_1]$   | S  | 0.23        | <b>0.30</b> | 0.35        | 0.28        | 0.32        | 6    | 5  | 3    | 1  | 0.32        | <b>21</b> | <b>7</b>  | 7    | 0.33        | 0.45        | 0.50        | 1        | 0.96        |      |
| $ec_6 : [h_1, w_1 \rightarrow w_2, v_1 \rightarrow v_2]$   | L  | -0.10       | <b>0.72</b> | -0.05       | 0.35        | 0.04        | 5    | 6  | 1    | 3  | <b>0.68</b> | <b>7</b>  | <b>21</b> | 7    | 0.31        | <b>0.50</b> | <b>0.45</b> | <b>1</b> | <b>0.96</b> |      |
| $ec_7 : [h_1 \rightarrow h_2, w_1 \rightarrow w_4, v_2 \rightarrow v_1]$   | VL | -0.10       | 6.95        | 0.14        | 0.41        | 0.15        | 6    | 4  | 2    | 2  | <b>0.88</b> | <b>21</b> | <b>7</b>  | 7    | -0.44       | <b>0.47</b> | <b>0.50</b> | <b>1</b> | <b>0.97</b> |      |
| <b>x264— Workload (#pictures/size): <math>w_1 : 8/2, w_2 : 32/11, w_3 : 128/44</math>; Version: <math>v_1 : r2389, v_2 : r2744, v_3 : r2744</math></b>                   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_1 : [h_2 \rightarrow h_1, w_3, v_3]$   | SM | <b>0.97</b> | 1.00        | <b>0.99</b> | <b>0.97</b> | <b>0.92</b> | 9    | 10 | 8    | 0  | 0.86        | 21        | 33        | 18   | 1.00        | 0.49        | 0.49        | 1        | 1           |      |
| $ec_2 : [h_2 \rightarrow h_1, w_1, v_3]$   | S  | <b>0.96</b> | 0.02        | <b>0.96</b> | <b>0.76</b> | <b>0.79</b> | 9    | 9  | 8    | 0  | 0.94        | 36        | 27        | 24   | 1.00        | 0.49        | 0.49        | 1        | 1           |      |
| $ec_3 : [h_1, w_1 \rightarrow w_2, v_3]$   | M  | 0.65        | <b>0.06</b> | 0.63        | 0.53        | 0.58        | 9    | 11 | 8    | 1  | <b>0.89</b> | 27        | 33        | 22   | <b>0.96</b> | 0.49        | 0.49        | 1        | 1           |      |
| $ec_4 : [h_1, w_1 \rightarrow w_3, v_3]$   | M  | 0.67        | <b>0.06</b> | 0.64        | 0.53        | 0.56        | 9    | 10 | 7    | 1  | <b>0.88</b> | 27        | 33        | 20   | <b>0.96</b> | 0.49        | 0.49        | 1        | 1           |      |
| $ec_5 : [h_1, w_3, v_2 \rightarrow v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_6 : [h_1, w_3, v_1 \rightarrow v_2]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_7 : [h_1, w_1 \rightarrow w_3, v_2]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_8 : [h_2 \rightarrow h_1, w_1, v_3]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| <b>SQLite— Workload</b>  |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_1 : [h_3 \rightarrow h_2, w_1, v_1]$   | S  | <b>0.96</b> | 1.27        | <b>0.83</b> | 0.40        | 0.35        | 2    | 3  | 1    | 0  | 1           | 9         | 9         | 7    | 0.99        | N/A         | N/A         | N/A      | N/A         |      |
| $ec_2 : [h_3 \rightarrow h_2, w_2, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_3 : [h_2, w_1 \rightarrow w_2, v_1]$   | M  | 0.50        | <b>1.24</b> | 0.43        | 0.17        | 0.43        | 1    | 1  | 0    | 0  | 1           | 4         | 2         | 2    | <b>1.00</b> | N/A         | N/A         | N/A      | N/A         |      |
| $ec_4 : [h_2, w_3 \rightarrow w_4, v_1]$   | M  | <b>0.95</b> | 1.00        | 0.79        | 0.24        | 0.29        | 2    | 4  | 1    | 0  | 1           | 12        | 11        | 7    | 0.99        | N/A         | N/A         | N/A      | N/A         |      |
| $ec_5 : [h_1, w_1, v_1 \rightarrow v_2]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_6 : [h_1, w_2 \rightarrow w_1, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_7 : [h_2 \rightarrow h_1, w_2, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| <b>SaC— Workload: <math>w_1 : 1/1, w_2 : 1/1, w_3 : 1/1, w_4 : 1/1, w_5 : 1/1, w_6 : 1/1, w_7 : 1/1, w_8 : 1/1, w_9 : 1/1</math></b>                                     |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_1 : [h_1, w_1 \rightarrow w_2, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_2 : [h_1, w_1 \rightarrow w_3, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_3 : [h_1, w_1 \rightarrow w_4, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_4 : [h_1, w_1 \rightarrow w_5, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_5 : [h_1, w_2 \rightarrow w_3, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_6 : [h_1, w_2 \rightarrow w_4, v_1]$   | S  | <b>0.91</b> | 5.54        | <b>0.80</b> | 0.00        | <b>0.91</b> | 14   | 11 | 8    | 0  | 0.85        | 64        | 65        | 31   | -0.40       | 0.13        | 0.15        | 0.12     | 0.64        |      |
| $ec_7 : [h_1, w_2 \rightarrow w_5, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_8 : [h_1, w_3 \rightarrow w_4, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_9 : [h_1, w_3 \rightarrow w_5, v_1]$   |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_{10} : [h_1, w_4 \rightarrow w_5, v_1]$  |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_{11} : [h_1, w_6 \rightarrow w_7, v_1]$  |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_{12} : [h_1, w_6 \rightarrow w_8, v_1]$  |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_{13} : [h_1, w_8 \rightarrow w_9, v_1]$  |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |
| $ec_{14} : [h_1, w_9 \rightarrow w_8, v_1]$  |    |             |             |             |             |             |      |    |      |    |             |           |           |      |             |             |             |          |             |      |

**Insight 1. Performance distributions can be transferred: Potential for learning a non-linear transfer function.**

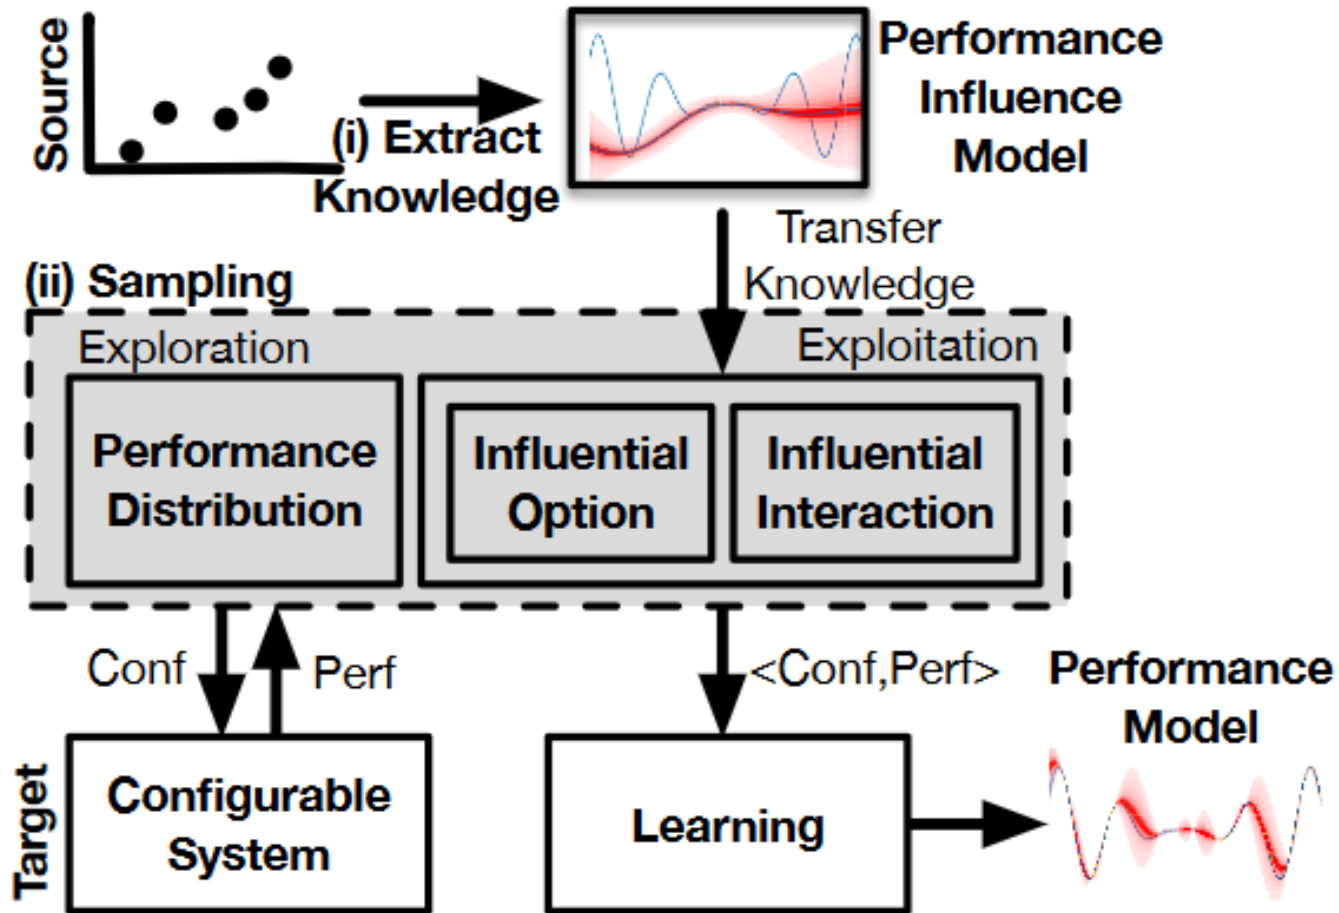
**Insight 2. Configuration ranks can be transferred: Good configurations stay good for changing hardware.**

**Insight 3. Influential options and interactions can be transferred: Relevant options in one environment stay relevant in other environments.**

ES: Expected severity  
SaC workload descri  
tbody: simulation of

**h1:** NUC4/1.30/15/SSD; **h2:** NUC12/2.15/1/SSD; **h3:** Station/1/2.8/3/SSD; **h4:** Amazon/1/2.4/1/SSD; **h5:** Amazon/1/2.4/0.5/SSD; **h6:** Azure/1/2.4/3/SSD  
**Metrics:** **M1:** Pearson correlation; **M2:** Kullback-Leibler (KL) divergence; **M3:** Spearman correlation; **M4/M5:** Perc. of top/bottom conf.; **M6/M7:** Number of influential options; **M8/M9:** Number of options agree/disagree; **M10:** Correlation btw importance of options; **M11/M12:** Number of interactions; **M13:** Number of interactions agree on effects;

# Learn to Sample




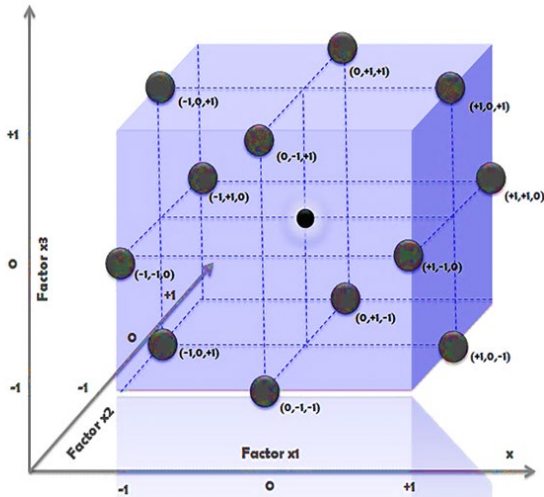

## Challenge: Finding Interactions



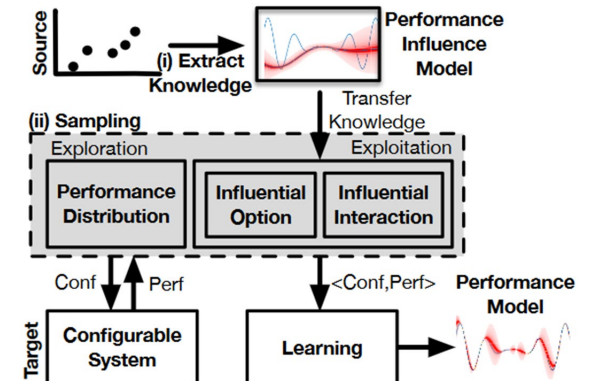
## Challenge: Measurement Costs



## Challenge: Workload/Environment Changes



## Learn to Sample





# **Future Directions**

# Exploiting Code Structure

```
#ifdef LOCKING
list *locks;
void lock() { /*...*/ }
void unlock() { /*...*/ }
#endif

void put(Object key, Object data){
#ifdef LOCKING
    lock();
#endif
    /*...*/
}

#ifdef STATISTICS
int getDbSize() {
    return calculateDbSize();
}
int calculateDbSize() {
#ifdef LOCKING
    lock();
#endif
    /*...*/
}
#endif
```

# Tracking Load-Time Options

```
class ProxyService {
    static boolean NATIVE_PROXY_SUPPORTED
        = Build.VERSION.SDK_INT >= 12;
    public void onSharedPreferencesChanged() {
        boolean enableSharing = false;
        if (!NATIVE_PROXY_SUPPORTED)
            if (Context.getSystemService("bluetooth"))
                enableSharing = true;

        if (enableSharing) {
            // slow code
        }
    }
}
```



# Whitebox Performance Analysis



Static/dynamic tracking of options w/ slicing & taint tracking

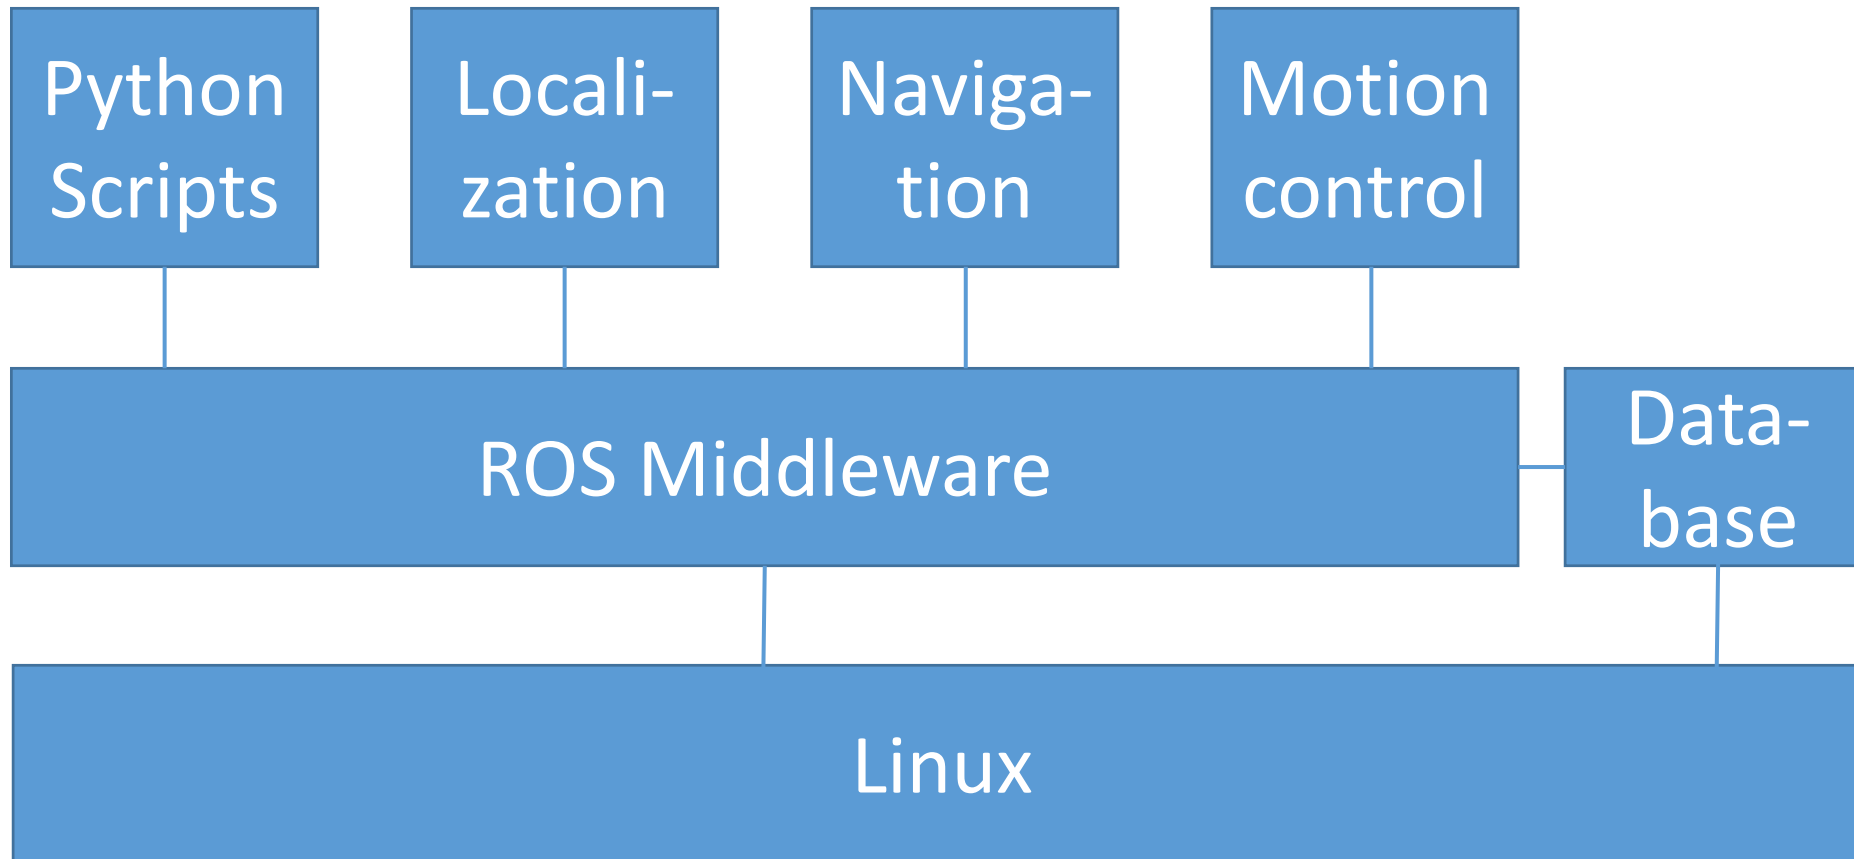
Which options influence performance-critical regions?

Which options interact?

-> Reduce number of measurements

-> Target measurements for likely interactions

# Graybox Performance Analysis?




## Challenge: Finding Interactions



## Challenge: Measurement Costs

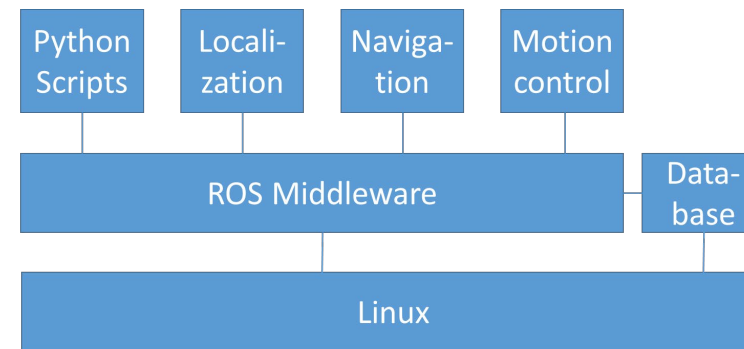


## Challenge: Workload/Environment Changes



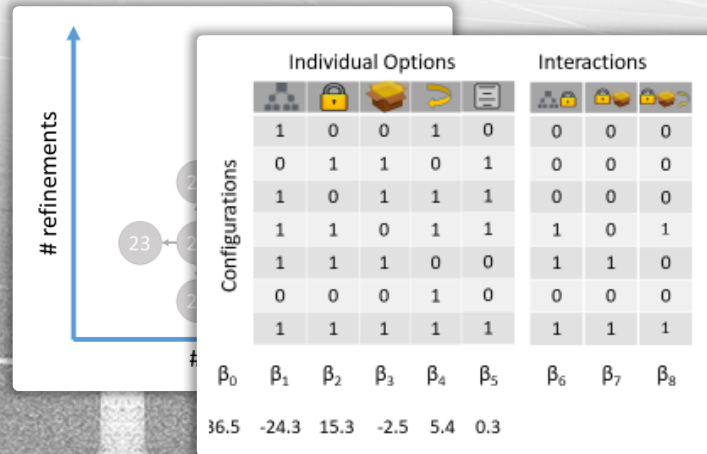
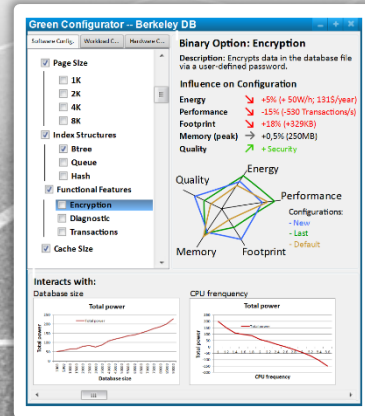



```
class ProxyService {
  static boolean NATIVE_PROXY_SUPPORTED
    = Build.VERSION.SDK_INT >= 12;
  public void onSharedPreferenceChanged() {
    boolean enableSharing = false;
    if (!NATIVE_PROXY_SUPPORTED)
      if (Context.getSystemService("bluetooth"))
        enableSharing = true;

    if (enableSharing) {
      // slow code
    }
  }
}
```



# Performance Analysis for Highly-Configurable Systems



```
class ProxyService {
  static boolean NATIVE_PROXY_SUPPORTED
    = Build.VERSION.SDK_INT >= 12;
  public void onSharedPreferencesChanged() {
    boolean enableSharing = false;
    if (!NATIVE_PROXY_SUPPORTED)
      if (Context.getSystemService("bluetooth"))
        enableSharing = true;
    if (enableSharing) {
```