CADET: Debugging and Fixing Misconfigurations using Counterfactual Reasoning

Md Shahriar Iqbal*
University of South Carolina
miqbal@email.sc.edu

Rahul Krishna*
Columbia University
rahul.krishna@columbia.edu

Mohammad Ali Javidian
Purdue University
mjavidia@purdue.edu

Baishakhi Ray
Columbia University
rayb@cs.columbia.edu

Pooyan Jamshidi
University of South Carolina
pjamshid@cse.sc.edu

Abstract

Modern computing platforms are highly-configurable with thousands of interacting configuration options. However, configuring these systems is challenging and misconfigurations can cause unexpected non-functional faults. This paper proposes CADET (short for Causal Debugging Toolkit) that enables users to *identify, explain*, and *fix* the root cause of non-functional faults early and in a principled fashion. CADET builds a causal model by observing the performance of the system under different configurations. Then, it uses casual path extraction followed by counterfactual reasoning over the causal model to (a) identify the root causes of non-functional faults, (b) estimate the effects of various configuration options on the performance objective(s), and (c) prescribe candidate repairs to the relevant configuration options to fix the non-functional fault. We evaluated CADET on 5 highly-configurable systems by comparing with state-of-the-art configuration optimization and ML-based debugging approaches. The experimental results indicate that CADET can find effective repairs for faults in multiple non-functional properties with (at most) 13% more accuracy, 32% higher gain, and 13× speed-up than other ML-based performance debugging methods. Compared to multi-objective optimization approaches, CADET can find fixes (at most) 8× faster with comparable or better performance gain. Our study of non-functional faults reported in NVIDIA's forum shows that CADET can find 14% better repairs than the experts' advice in less than 30 minutes.

1 Introduction

Modern computing systems are highly configurable and can seamlessly be deployed on various hardware platforms and under different environmental settings. The configuration space is combinatorially large with 100s if not 1000s of software and hardware configuration options that interact non-trivially with one another [38, 49, 99]. Unfortunately, configuring these systems to achieve specific goals is challenging and error-prone.

Incorrect configuration (*misconfiguration*) elicits unexpected interactions between software and hardware resulting *non-functional faults*, i.e., faults in *non-functional* system properties such as latency and energy consumption. These non-functional faults—unlike regular software bugs—do not cause the system to crash or exhibit an obvious misbehavior [75, 82, 94]. Instead, misconfigured systems remain operational while being compromised, resulting severe performance degradation in latency, energy consumption, and/or
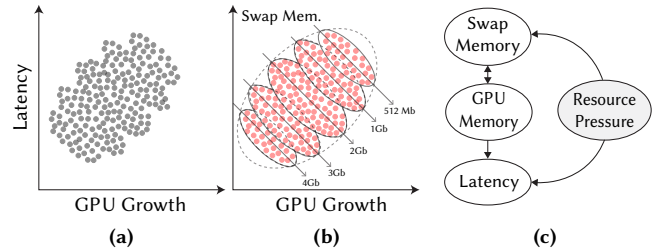
Figure 1: Observational data (in Fig. 1a) (incorrectly) shows that high GPU memory growth leads to high latency. The trend is reversed when the data is segregated by swap memory.

heat dissipation [15, 71, 74, 84]. The sheer number of modalities of software deployment is so large that exhaustively testing every conceivable software and hardware configuration is impossible. Consequently, identifying the root cause of non-functional faults is notoriously difficult [35] with as much as 99% of them going unnoticed or unreported for extended durations [4]. This has tremendous monetary repercussions costing companies worldwide an estimated $5 trillion in 2018 and 2019 [34]. Further, developers on online forums are quite vocal in expressing their dissatisfaction. For example, one developer on NVIDIA's developer forum bemoans: *"I am quite upset with CPU usage on TX2 [8],"* while another complained, *"I don't think it [the performance] is normal and it gets more and more frustrating [7]."* Crucially, these exchanges provoke other unanswered questions, such as, *"what would be the effect of changing another configuration 'X'? [2]."* Therefore, we seek methods that can *identify, explain, and fix* the root cause of non-functional faults early in a principled fashion.

Existing work. Much recent work has focused on configuration optimization, which are approaches aimed at finding a configuration that optimizes a performance objective [29, 72, 86, 100, 105]. Finding the optimum configuration using push-button optimization approaches is not applicable here because they do not give us any information about the underlying interactions between the faulty configuration options that caused the non-functional fault. This information is sought after by developers seeking to address these non-functional faults [82, 93].

Some previous work has used machine learning-based performance modeling approaches [36, 86, 87, 96]. These approaches are adept at inferring the correlations between certain configuration

---
*Joint First Author

options and performance, however, they lack the mathematical language to express *how* or *why* the configuration options affect performance. Without this, we risk drawing misleading conclusions. They also require significant time to gather the training samples, which grows exponentially with the number of configurations [56, 103].

Limitations of existing work. In Fig. 1, we present an example to help illustrate the limitations with current techniques. Here, the observational data gathered so far indicates that GPU memory growth is positively correlated with increased latency (as in Fig. 1a). An ML-model built on this data will, with high confidence, predict that larger GPU memory growth leads to larger latency. However, this is counter-intuitive because higher GPU memory growth should, in theory, reduce latency not increase it. When we segregate the same data on swap memory (as in Fig. 1b), we see that there is indeed a general trend downward for latency, i.e., within each group of swap memory, as GPU memory growth increases the latency decreases. We expect this because GPU memory growth controls how much memory the GPU can "borrow" from the swap memory. Depending on resource pressure imposed by other host processes, a resource manager may arbitrarily re-allocate some swap memory; this means the GPU borrows proportionally more/less swap memory thereby affecting the latency correspondingly. This is reflected by the data in Fig. 1b. If the ML-based model were to consult the available data (from Fig. 1a) unaware of such underlying factors, these models would be incorrect. With thousands of configurations, inferring such nuanced information from optimization or ML-based approaches would require a considerable amount of measurements and extensive domain expertise which can be impractical, if not impossible, to possess in practice.

Our approach. In this paper, we design, implement, and evaluate CADET (short for <u>Ca</u>usal <u>De</u>bugging <u>T</u>oolkit). CADET uses causal structural discovery algorithms [91] to construct a causal model using observational data [78]. Then, it uses counterfactual reasoning over the causal model to (a) identify the root causes of non-functional faults, (b) estimate the effects of various configurable parameters on the non-functional properties, and (c) prescribe candidate repairs to the relevant configuration options to fix the non-functional fault. For example, in Fig. 1, CADET constructs a *causal model* from observational data (as in Fig. 1c). This causal model indicates that GPU memory growth indirectly influences latency (via a swap memory) and that the configuration options may be affected by certain other factors, e.g., resource manager allocating resources for other processes running on the host. CADET uses counterfactual questions such as, "what is the effect of GPU memory growth on latency if the available swap memory was 2Gb?" to diagnose the faults and recommend changes to the configuration options to mitigate these faults.

Evaluation. We evaluate CADET on 5 real-world highly configurable systems (three machine learning systems, a video encoder, and a database system) deployed on 3 architecturally different NVIDIA Jetson systems-on-chip. We compare CADET with state-of-the-art configuration optimization and ML-based performance debugging approaches. Overall, we find that CADET is (at most) 13× faster with 13% better accuracy and 32% higher gain than the next best ML-based approaches for single-objective faults. Compared to single-objective optimization approaches, CADET can find repairs to misconfigurations 9× faster while performing as well

as (or better than) the best configuration discovered by the optimization techniques. Further, CADET can find effective repairs for faults in multiple performance objectives (i.e., latency and energy consumption) with accuracy and gain as high as 16% and 36%, respectively, than other ML-based performance debugging methods. Compared to multi-objective optimization approaches, CADET can find repairs 5× faster while having similar performance gain. Finally, with a case study, we demonstrate that CADET finds 14% better repairs than the experts' advice in 24 minutes.

Contributions. In summary, our contributions are as follows.

• We propose CADET, a novel causal diagnostics, and fault mitigation tool that identifies the root causes of non-functional faults and recommends repairs to resolve non-functional faults.

• Our empirical results, conducted on on 5 highly configurable systems deployed on 3 architecturally different hardware, CADET outperforms current state-of-the-art ML-based diagnostics and optimization approaches in terms of efficacy (+16% accuracy and +36% gain) and speed (up to 13× faster). CADET was also adept at handling both single- and multi-objective performance faults.

• We offer a manually curated performance fault dataset (called Jetson Faults) and accompanying code required to reproduce our findings at https://git.io/JtFNG.

2 Motivation

2.1 A real-world example

This section illustrates the challenges both users and component developers face when configuring and composing complex systems with a concrete issue report from the NVIDIA forum [6]. Here, a user notices some strange behavior when trying to transplant their code for real-time computation of depth information from stereo-cameras from NVIDIA Jetson TX1 to TX2. Since TX2 has twice the computational power, they expected at least 30% lower latency, but observed 4 times higher latency. The user indicates working with a reliable reference implementation of their code and is puzzled by the occurrence of this fault.

To solve this, the user solicits advice from the NVIDIA forums. After discussions spanning two days, they learn that they have made several misconfigurations: (1) **Wrong compilation flags**: Their compilation does not take into account the microarchitectural differences between the two platforms. These may be fixed by setting the correct microarchitecture with -gencode=arch parameter and compiling the code dynamically by disabling the CUDA_USE_STATIC flag. (2) **Wrong CPU/GPU clock frequency**: Their hardware configuration is set incorrectly. These may be fixed by setting the configuration -nvpmodel=MAX-N which changes the CPU and GPU clock settings. The Max-N setting in TX2 provides almost twice the performance of TX1 [30] due to a number of factors including increased clock speeds and TX2's use of 128-bit memory bus width versus the 64-bit in TX1 [30]. If TX2 is not configured to leverage these differences, it will face high latency. (3) **Wrong fan modes**: Their fan needs to be configured correctly to achieve higher CPU/GPU clock speeds. If the fan modes are not configured to be high, TX2 will thermal throttle the CPU and GPU to prevent overheating [3] and invariably increasing latency [17].

2.2 Non-functional faults (NF-FAULTS)

The previous example is one of many examples wherein misconfigurations severely impact performance and energy consumption. Such examples abound in many other systems and domains, including IoT ubiquitous systems (e.g. Amazon Alexa) [13, 25] and production-scale cloud-based systems [35]. We term such anomalies as non-functional faults (hereafter, NF-FAULTS), and they represent a mismatch between the expected performance of a non-functional

Figure 2: An example of multi-objective NF-FAULTS.

property (such as the latency or energy consumption) and what is actually observed when a system is deployed. Specifically, for certain configurations the software system may experience *tail values* of latency, energy consumption, or heat dissipation, i.e., they take values worse than predetermined SLA (we choose performance worse than 99th percentile). In certain cases, the software system might suffer from tail values in multiple non-functional properties simultaneously, i.e., multi-objective NF-FAULTS.

Fig. 2 illustrates the distribution of two non-functional properties (latency and energy consumption) measured for 2000 different configuration setting for an NVIDIA TX2 SoC running an image recognition task. There exist two types of non-functional faults that may be discerned from Fig. 2: *(1)* configurations whose latency *or* energy is worse that 99.99% of other configurations (shown in •); *(2)* configurations where *both latency and energy* is worse that 99.99% of other configurations (shown in •). The remaining configurations (shown in •) are not considered to be faulty.

2.3 Manual Study

To further understand the nature of NF-FAULTS, we explored the NVIDIA developer forums and analyzed 47 reported issues for different ML systems (e.g., image recognition, self-driving cars, IoT, etc.). The reports and corresponding discussions highlight the challenges that *developers* and *users* face when building complex systems and how they are remedied. Specifically, we studied composed systems deployed on NVIDIA's hardware platforms for various applications such as image processing, voice recognition, and speech recognition deployed on heterogeneous hardware platforms (TX1, TX2, XAVIER) allowing users to compose a custom system to meet specific functional and non-functional needs. The extensibility of composed systems exposes them to frequent misconfigurations resulting in non-functional faults. Analyzing the issues and discussions around these misconfigurations, we found the following: (1) **There are different types of NF-FAULTS across all systems as a result of misconfigurations** related to (i) latency (42%), (ii) energy consumption (36%), and (iii) thermal problems (22%). (2) **Configuration options interact with one another in different ways across environments.** 25% of issues report multiple types of NF-FAULTS simultaneously, highlighting challenging
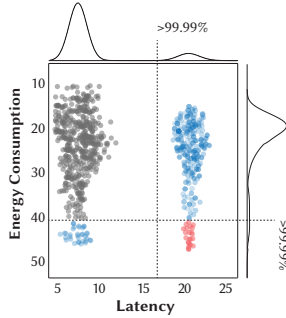
tradeoffs users need to make between, for example, reducing latency by increasing core frequency results in higher energy consumption. (3) **Non-functional faults take a long time to resolve.** Most performance issues (i.e., 61%) took more than 1 week (average 5 weeks, max 11 months) to resolve in comparison to other types of non-performance issues and require frequent back and forth between users and support teams, e.g. [1, 5]. Most of these issues are environment-specific and since the developers do not have access to the user's environment, they usually cannot replicate them.

To address these issues, developers manually inspect trace logs obtained by utilities such as perf. These traces present a large amount of information pertaining to the current state of the system. The volume of information presented makes it extremely burdensome for developers to decipher the actual root-cause of non-functional faults. To overcome this challenge, we propose using causal inference accompanied by counterfactuals to reason about and address these NF-FAULTS.

3 Background: Causal Inference

Causal inference is a systematic study to understand the interactions between variables in a system in terms of their cause-effect relationships between one another. This section provides a brief overview of how one may reason about causal relationships using the running example from example from Fig. 1. Accordingly, let us assume that we have gathered several samples of GPU memory growth, swap memory, and latency. If we are interested in how latency behaves given GPU memory growth, then, from a causal perspective, this can be formulated in three ways: observational, interventional, and counterfactual [78, 79].

3.1 Observation and Intervention

In the *observational* formulation, we measure the distribution of a target variable (e.g., latency Y) given that we *observe* another variable GPU memory growth (X) takes a certain value 'x' (i.e., $X = x$), denoted by $Pr(Y \mid X = x)$. This is a familiar expression that, given adequate data, modern supervised ML algorithms can estimate.

The *interventional* inference tackles a harder task of estimating the *effects of deliberate actions.* For example, in Fig. 1, we measure how the distribution of Latency (Y) would change if we (artificially) intervened during the data gathering process by forcing the variable GPU memory growth (X) to take a certain value 'x', i.e., $X = x$, but otherwise retain the other variables (i.e., swap memory) as is. This is done by building causal graphs and modifying them to reflect our intervention. Then, using the modified causal graph and *do-causal calculus* rules [78] we may estimate the outcome of the artificial intervention. This process is computationally efficient since we reason over existing data instead of gathering additional measurements.

3.2 Counterfactual Inference

These problems involve probabilistic answers to "what if?" questions. They pertain to alternate possibilities that are counter to the current facts (hence counterfactual) and their consequences. For example, in Fig. 1, a counterfactual question may ask:

> *Given that we observed a high latency, and given that* GPU memory growth *was set to (say) 33%, and given everything else we know about the circumstances of the observation (i.e., available* swap memory*), what is the probability of decreasing latency, had we set the* GPU memory growth *to 66%?*

In other words, we are interested in a scenario where:
• We *hypothetically* have low latency;
Conditioned on the following events:
• We *actually observed* high latency;
• GPU memory growth was initially set to 33%;
• We *hypothetically* set the new GPU memory growth to 66%; and
• Other circumstances (available swap memory, resource load, etc) remain the same;
Formally we represent this, by abusing notations, as:

$$Pr(\hat{Y} = y_{low} \mid \hat{X} = 0.66, \ X = 0.33, \ Y = y_{high}, \ Z) \qquad (1)$$

Here, Y stands for observed latency, X stands for current GPU memory growth, and Z stands for current swap memory. The variables $\hat{X}$ and $\hat{Y}$ are as yet unobserved (or hypothetical), i.e., these are variables that are either being predicted for ($\hat{Y} = y_{low}$) or being *hypothetically* changed to a new value ($\hat{X} = 0.66$).

Questions of this nature require precise mathematical language lest they will be misleading. For example, in Eq. (1), we are simultaneously conditioning on two values of GPU memory growth (i.e., $\hat{X} = 0.66$ and $X = 0.33$). Traditional machine learning approaches cannot handle such expressions. Instead, we must resort to causal models to compute them.

Computing counterfactual questions using causal models involves three steps [78, Section 7]: (i) *Abduction*: where we update our model of the world using observational data; (ii) *Action*: where we modify the model to reflect the counterfactual assumption being made; and (iii) *Prediction*: where we use the obtained modified model from the previous step to compute the value of the consequence of the counterfactual.

4 CADET: <u>Ca</u>usal Performance <u>De</u>bugging

Leveraging the Causal Analysis, described in Section 3, we build CADET, to debug the root causes and fix NF-FAULTS. This section presents a detailed description of CADET (outlined in Fig. 3).
Usage Scenario. In a typical scenario, CADET can be used when a developer experiences NF-FAULTS induced by misconfigurations to (a) identify which configuration options are the *root-cause* of the faults, and (b) prescribe how to set the root-cause configuration values to fix the non-functional fault. To do this, the user queries CADET with their questions about the non-functional fault (e.g., "what is the root cause of the non-functional fault?", or "how to improve the latency by 90%?", *etc*).
Overview. CADET works in three phases: (I) causal structure discovery, (II) causal path extraction, and (III) debug and repair using a counterfactual query. For Phase-I, we first generate a few dozen observed data, i.e., dynamic traces for measuring the non-functional properties of the system (e.g., latency, energy consumption, etc.) under different configuration settings. Using these traces, we construct a causal graph that captures the causal relationships between various configuration options and the system's non-functional properties. For Phase-II, we use the causal graph to identify *causal*
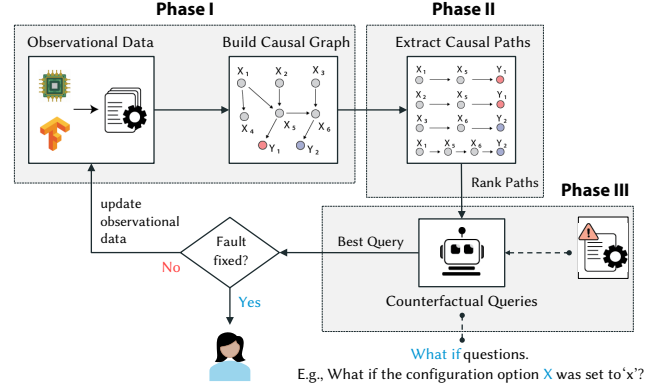

Figure 3: Overview of CADET

paths—paths that lead from configuration options to the non-functional properties. Next, in Phase-III, given an observed NF-FAULT, we aim to debug and repair the misconfigurations. In particular, we generate a series of counterfactual queries, i.e., what-if questions about specific changes to the values of each configuration option. Given a NF-FAULT as input, we check which of the counterfactual queries has the highest causal effect on remedying the NF-FAULT, and we generate a new configuration using that query. Finally, we evaluate the new configuration to assert if the newly generated configuration mitigates the fault. If not, we add this new configuration to the observational data of Phase-I and repeat the process until the NF-FAULT is fixed.

4.1 Phase-I: Causal Graph Discovery

In this stage, we express the relationships between configuration options (e.g., CPU freq, *etc*.) and the non-functional properties (e.g., latency, *etc*) using a causal model. A causal model is an *acyclic directed mixed graph* (hereafter, ADMG), i.e., an acyclic graph consisting of directed and bidirected edges representing the causal direction and existence of latent common cause(s), respectively [26, 83]. The nodes of the ADMG have the configuration options and the non-functional properties (e.g., latency, *etc*). Additionally, we enrich the causal graph by including nodes that represent the status of internal systems events, e.g., resource pressure (as in Fig. 1). Unlike configuration options, these system events cannot be modified. However, they can be observed and measured to explain how the causal-effect of changing configurations propagates to latency or energy consumption, e.g., in Fig. 1 resource pressure is a system event that determines how GPU memory growth affects latency.
Data collection. To build the causal model we gather a set of initial observational data resembling Table 4a[1]. It consists of configuration options set to randomly chosen values; various system events occurring when these configurations were used; and the measure of latency and energy consumption.
Fast Causal Inference. To convert observational data into a causal graph, we use a prominent structure discovery algorithm called *Fast Causal Inference* (hereafter, FCI) [91]. We picked FCI

---

[1]For our systems, we found that 25 initial samples were adequate to build an initial version of a causal model. However, this is a parameter that may change with the system.
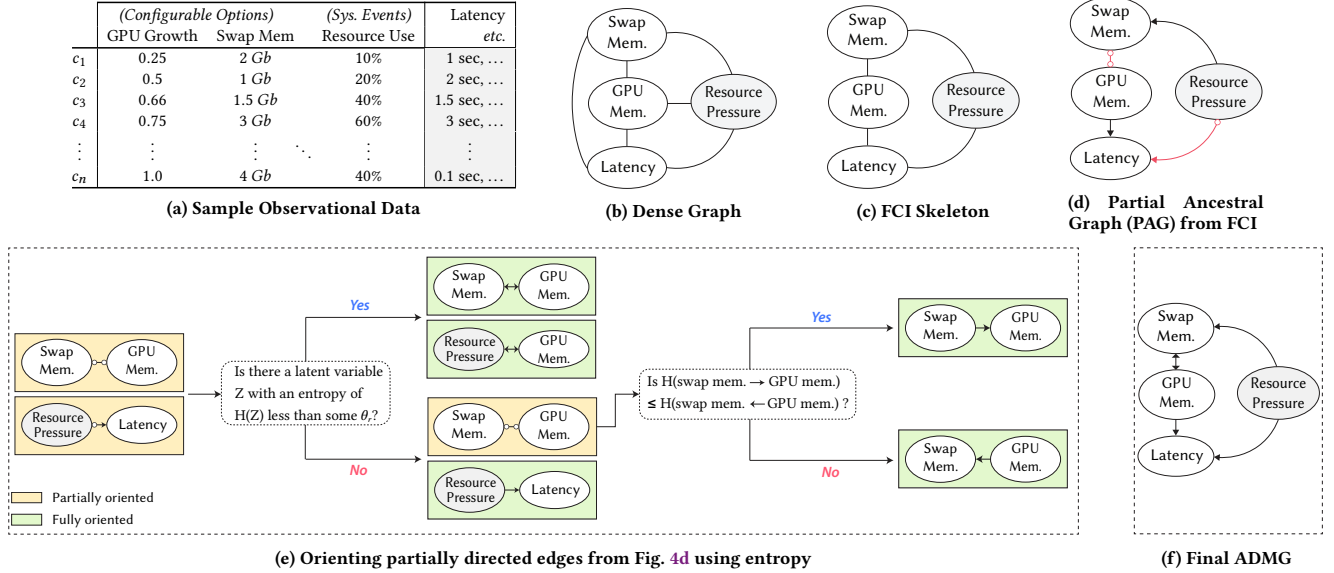
| | (Configurable Options) | | (Sys. Events) | Latency |
	GPU Growth	Swap Mem	Resource Use	etc.
c_1	0.25	2 Gb	10%	1 sec, …
c_2	0.5	1 Gb	20%	2 sec, …
c_3	0.66	1.5 Gb	40%	1.5 sec, …
c_4	0.75	3 Gb	60%	3 sec, …
⋮	⋮	⋮	⋱ ⋮	⋮
c_n	1.0	4 Gb	40%	0.1 sec, …

(a) Sample Observational Data

(b) Dense Graph

(c) FCI Skeleton

(d) Partial Ancestral Graph (PAG) from FCI

(e) Orienting partially directed edges from Fig. 4d using entropy

(f) Final ADMG

Figure 4: From observational data to a fully connected graph, a skeleton graph, and finally to partial ancestral graph (PAG).

because it has several useful properties. First, it accommodates for the existence of unobserved confounders [32, 76, 91], i.e., it operates even when there are unknown variables that have not been, or cannot be, measured. This is important because we do not assume absolute knowledge about configuration space, hence there could be certain configurations we could not modify or system events we have not observed. Second, it accommodates variables that belong to various data types such as nominal, ordinal, and categorical data common across the system stack.

FCI operates in three stages. First, we construct a fully connected undirected graph where each variable is connected to every other variable. Second, we use statistical independence tests to prune away edges between variables that are independent of one another. Finally, we orient undirected edges using prescribed edge orientation rules [19, 20, 32, 76, 91] to produce a *partial ancestral graph* (or PAG). A PAG contains the following types of (partially) directed edges:

- $X \longrightarrow Y$ indicating that vertex X causes Y.
- $X \longleftrightarrow Y$ which indicates that there are unmeasured confounders between vertices X and Y.

In addition, FCI produces two types of partially directed edges:

- $X \circ\!\!\longrightarrow Y$ indicating that either X causes Y, or that there are *unmeasured confounders* that cause both X and Y.
- $X \circ\!\!-\!\!\circ Y$ which indicates that either: (a) vertices X causes Y, or (b) vertex Y causes X, or (c) there are *unmeasured confounders* that cause both X and Y.

In the last two cases, the circle (○) indicates that there is an ambiguity in the edge type. In other words, given the current observational data, the circle can indicate an arrow head ($\longrightarrow$) or no arrow head ($-$), i.e., for $X \circ\!\!-\!\!\circ Y$, all three of $X \longrightarrow Y$, $Y \longrightarrow X$, and $X \longleftrightarrow Y$ might be compatible with current data, i.e., the current data could

be faithful to each of these statistically equivalent causal graphs inducing the same conditional independence relationships.

Resolving partially directed edges. For subsequent analyses over the causal graph, the PAG obtained must be fully resolved (directed with no ○ ended edges) in order to generate an ADMG, i.e., we must fully orient partially directed edges by replacing the circles in $\circ\!\!\longrightarrow$ and $\circ\!\!-\!\!\circ$ with the correct edge direction.

Resolving these partial edges from PAG is an open problem. This is tackled in two ways: (1) Solicit expert advice to orient edges [42], but this may be too cumbersome especially when the graph is complex [42]; or (2) Use alternative information-theoretic approaches to discover the correct orientations [58, 102]. In this paper, we use the information-theoretic approach, enabling us to orient all edges automatically without any expert intervention. Specifically, this paper uses an entropic approach proposed in [59, 60] to discover the true causal direction between two variables. Entropic causal discovery is inspired by Occam's razor, and the key intuition is that, among the possible orientations induced by partially directed edges (i.e., $\circ\!\!\longrightarrow$ and $\circ\!\!-\!\!\circ$), the most plausible orientation is that which has the lowest entropy.

Our work extends the theoretic underpinnings of entropic causal discovery to generate a fully directed causal graph by resolving the partially directed edges produced by FCI. For each partially directed edge, we follow two steps: (1) establish if we can generate a latent variable (with low entropy) to serve as a common cause between two vertices; (2) if such a latent variable does not exist, then pick the causal direction which has the lowest entropy.

For the first step, we assess if there could be an unmeasured confounder (say Z) that lies between two partially oriented nodes (say X and Y). For this, we use the *LatentSearch* algorithm proposed by Kocaoglu *et al.* [60]. *LatentSearch* outputs a joint distribution $q(X, Y, Z)$ of the variables X, Y, and Z which can be used to compute

the entropy $H(Z)$ of the unmeasured confounder Z. Following the guidelines of Kocaoglu *et al.*, we set an entropy threshold $\theta_r = 0.8 \times min\{H(X), H(Y)\}$. If the entropy $H(Z)$ of the unmeasured confounder falls *below* this threshold, then we declare that there is a simple unmeasured confounder Z (with a low enough entropy) to serve as a common cause between X and Y and accordingly, we replace the partial edge with a bidirected (i.e., $\longleftrightarrow$) edge.

When there is no latent variable with a sufficiently low entropy, there are two possibilities: *(a)* variable X causes Y; then, there is an arbitrary function $f(\cdot)$ such that $Y = f(X, E)$, where E is an exogenous variable (independent of X) that accounts for noise in the system; or *(b)* variable Y causes X; then, there is an arbitrary function $g(\cdot)$ such that $X = g(Y, \tilde{E})$, where $\tilde{E}$ is an exogenous variable (independent of Y) that accounts for noise in the system. The distribution of E and $\tilde{E}$ can be inferred from the data [59, see §3.1]. With these distributions, we measure the entropies $H(E)$ and $H(\tilde{E})$. If $H(E) < H(\tilde{E})$, then, it is simpler to explain the $X \longrightarrow Y$ (i.e., the entropy is lower when $Y = f(X, E)$) and we choose $X \longrightarrow Y$ as the causal direction. Otherwise, we choose $Y \longrightarrow X$.

Example. Fig. 4 shows the steps involved in generating the final ADMG from Fig. 1c. First, we build a dense graph by connecting all pairs of variables with an undirected egde (see Fig. 4b). Next, we use Fisher's exact test [21] to evaluate the independence of all pairs of variables conditioned on all remaining variables. Pruning edges between the independent variables results in skeleton graph as shown in Fig. 4c. Next, we orient undirected edges using edge orientation rules [20, 32, 76, 91] to produce a partial ancestral graph (as in Fig. 4d). In our example, we identify that there are two edges that are partially oriented (outlined in red in Fig. 4d): *(i)* swap memory ∘—∘ GPU memory growth; and *(ii)* resource pressure ∘—➛ latency. To resolve these two edge, we use the entropic orientation strategy. The steps involved to orient these edges are illustrated in Fig. 4e. After this step, we get the final ADMG shown in Fig. 4f.

4.2 Phase-II: Causal Path Extraction

In this stage, we extract paths from the causal graph (referred to as *causal paths*) and rank them from highest to lowest based on their average causal effect on latency, energy consumption, and heat dissipation (our three non-functional properties). Using path extraction and ranking, we reduce the complex causal graph into a small number of useful causal paths for further analyses. The configurations in this path are more likely to be associated with the root cause of the fault.

Extracting causal paths with backtracking. A causal path is a directed path originating from either the configuration options or the system event and terminating at a non-functional property (i.e., latency and/or energy consumption). To discover causal paths, we backtrack from the nodes corresponding to each non-functional property until we reach a node with no parents. If any intermediate node has more than one parent, then we create a path for each parent and continue backtracking on each parent.

Ranking causal paths. A complex causal graph can result in a large number of causal paths. It is not practical to reason over all possible paths as it may lead to a combinatorial explosion. Therefore, we rank the paths in descending order from ones having the highest

causal effect to ones having the lowest causal effect on each non-functional property. For further analysis, we use paths with the highest causal effect. To rank the paths, we measure the causal effect of changing the value of one node (say GPU memory growth or X) on its successor in the path (say swap memory or Z). We express this with the *do-calculus* notation: $\mathbb{E}[Z \mid do(X = x)]$. This notation represents the expected value of Z (swap memory) if we set the value of the node X (GPU memory growth) to x. To compute the *average causal effect* (ACE) of $X \rightarrow Z$ (i.e., GPU memory growth $\longrightarrow$ swap memory), we find the average effect over all permissible values of X (GPU memory growth), i.e.,

$$\text{ACE}(Z, X) = \frac{1}{N} \cdot \sum_{\forall a, b \in X} \mathbb{E}[Z \mid do(X = b)] - \mathbb{E}[Z \mid do(X = a)] \quad (2)$$

Here, N represents the total number of values X (GPU memory growth) can take. If changes in GPU memory growth result in a large change in swap memory, then the ACE (Z, X) will be larger, indicating that GPU memory growth on average has a large causal effect on swap memory. For the entire path, we extend Eq. (2) as:

$$\text{Path}_{ACE} = \frac{1}{K} \cdot \sum \text{ACE}(Z, X) \qquad \forall X, Z \in path \quad (3)$$

Eq. (3) represents the average causal effect of the causal path. The configuration options the lie in paths with larger P_{ACE} tend to have a greater causal effect on the corresponding non-functional properties in those paths. We select the top K paths with the largest P_{ACE} values, for each non-functional property. In this paper, we use K=3, however, this may be modified in our replication package.

4.3 Phase-III: Debugging non-functional faults

In this stage, we use the top K paths to (a) identify the root cause of non-functional faults; and (b) prescribe ways to fix the non-functional faults. When experiencing non-functional faults, a developer may ask specific queries to CADET and expect an actionable response. For this, we translate the developer's queries into formal probabilistic expressions that can be answered using causal paths. We use counterfactual reasoning to generate these probabilistic expressions. To understand query translation, we use the example causal graph of Fig. 4f where a developer observes high latency, i.e., a latency fault, and has the following questions:

❓ **"What is the root cause of my latency fault?"** To identify the root cause of a non-functional fault we must identify which configuration options have the most causal effect on the performance objective. For this, we use the steps outlined in §4.2 to extract the paths from the causal graph and rank the paths based on their average causal effect (i.e., Path$_{ACE}$ from Eq. (3)) on latency. We return the configurations that lie on the top K paths. For example, in Fig. 4f we may return the path (say) GPU memory growth $\longrightarrow$ swap memory $\longrightarrow$ Latency and the configuration options GPU memory growth *and* swap memory both being probable root causes.

❓ **"How to improve my latency?"** To answer this query, we first find the root cause as described above. Next, we discover what values each of the configuration options must take in order that the new latency is better (low latency) than the fault (high latency). For example, we consider the causal path GPU memory growth $\longleftrightarrow$ swap memory $\longrightarrow$ Latency, we identify the permitted values for the configuration options GPU memory growth and swap memory that

can result in a low latency (Y^{LOW}) that is better than the fault (Y^{HIGH}). For this, we formulate the following counterfactual expression:

$$\Pr(Y^{\text{LOW}}_{repair} | \neg repair, Y^{\text{HIGH}}_{\neg repair}) \qquad (4)$$

Eq. (4) measures the probability of "fixing" the latency fault with a "repair" (Y^{LOW}_{repair}) given that with no repair we observed the fault ($Y^{\text{HIGH}}_{\neg repair}$). In our example, the repairs would resemble GPU memory growth =0.66 or swap memory =4Gb, *etc*. We generate a *repair set* ($\mathcal{R}$) where the configurations GPU memory growth and swap memory are set to all permissible values, i.e.,

$$\mathcal{R} \equiv \bigcup \ \{\text{GPU memory growth} = x, \ \text{swap memory} = z, ...\}$$
$$\forall x \in \text{GPU memory growth}, \ z \in \text{swap memory}, \ \ldots \quad (5)$$

Next, we compute the *individual treatment effect* (ITE) on the latency (Y) for each repair in the repair set $\mathcal{R}$. In our case, for each repair $r \in \mathcal{R}$, ITE is given by:

$$\text{ITE}(r) = \Pr(Y^{\text{LOW}}_r | \neg r, Y^{\text{HIGH}}_{\neg r}) - \Pr(Y^{\text{HIGH}}_r | \neg r, Y^{\text{HIGH}}_{\neg r}) \quad (6)$$

ITE measures the difference between the probability that the latency *is low* after a repair r and the probability that the latency is *still high* after a repair r. If this difference is positive, then the repair has a higher chance of fixing the fault. In contrast, if the difference is negative then that repair will likely worsen the latency. To find the most useful repair ($\mathcal{R}_{best}$), we find a repair with the largest (positive) ITE, i.e.,

$$\mathcal{R}_{best} = \underset{\forall r \in \mathcal{R}}{\text{argmax}}[\text{ITE}(r)]$$

This provides the developer with a possible repair for the configuration options that can fix the latency fault.
Remarks. The ITE computation of Eq. (6) occurs *only* on the observational data. Therefore we may generate any number of repairs and reason about them without having to deploy those interventions and measuring their performance in the real-world. This offers significant monetary and runtime benefits.

4.4 Incremental Learning

Using the output of **Phase-III**, the system is reconfigured with the new configuration. If the new configuration addresses the NF-FAULT, we return the recommended repairs to the developer. Since the causal model uses limited observational data, there may be a discrepancy between the actual performance of the system after the repair and the value of the estimation derived from the current version of the causal graph. The more accurate the causal graph, the more accurate the proposed intervention will be [19, 20, 32, 76, 91]. Therefore, in case our repairs do not fix the faults, we update the observational data with this new configuration and repeat the process. Over time, the estimations of causal effects will become more accurate. We terminate the incremental learning once we achieve the desired performance.

5 Case Study: Latency Fault in TX2

This section revisits the real-world latency fault previously discussed in §2.1. For this study, we reproduce the developers' setup to assess how effectively CADET can diagnose the root-cause of the misconfigurations and fix them. For comparison, we use SMAC (an optimization approach) and BugDoc (an ML-based diagnosis

Problem [6]: For a real-time scene detection task, TX2 (faster platform) only processed 4 frames/sec whereas TX1 (slower platform) processed 17 frames/sec, i.e., the latency is 4× worse on TX2.
Observed Latency (frames/sec): 4 FPS
Expected Latency (frames/sec): 22-24 FPS *(30-40% better)*

Configuration Options	CADET	SMAC	BugDoc	Forum	ACE†
CPU Cores	✔	✔	✔	✔	3%
CPU Frequency	✔	✔	✔	✔	6%
EMC Frequency	✔	✔	✔	✔	13%
GPU Frequency	✔	✔	✔	✔	22%
Scheduler Policy	.	✔	✔	.	.
Sched rt runtime	.	.	✔	.	.
Sched child runs	.	.	✔	.	.
Dirty bg. Ratio	.	.	.	.	.
Dirty Ratio	.	.	✔	.	.
Drop Caches	.	✔	✔	.	.
CUDA_STATIC	✔	✔	✔	✔	55%
Cache Pressure	.	.	.	.	.
Swappiness	.	✔	✔	.	1%
Latency (TX2 frames/sec)	**26**	24	20	23	
Latency Gain (over TX1)	**53%**	42%	21%	39%	
Latency Gain (over default)	**6.5×**	6×	5×	5.75×	
Resolution time	**24 mins**	4 hrs	3.5 hrs	2 days	
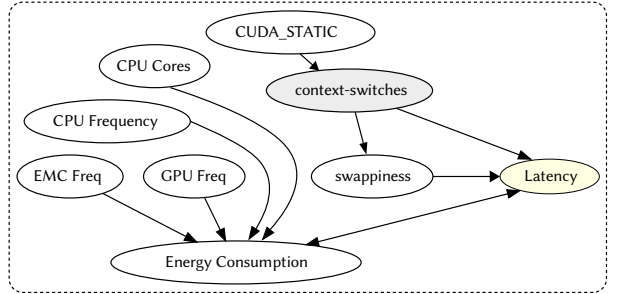
Figure 5: Using CADET on the real-world example from §2.1. CADET is better and faster than other methods.

tool). We evaluate the methods against the recommendations by the domain experts on the forum.
Findings. Fig. 5 illustrates our findings. We find that:
• CADET could diagnose the root-cause of the misconfiguration and recommends a fix within 24 minutes. Using the recommended configuration fixes from CADET, we achieved a frame rate of 26 FPS (53% better than TX1 and 6.5× better than the fault). This exceeds the developers' initial expectation of $30 - 40\%$ improvement.
• Using configuration optimization (SMAC), we auto-tune the system to find a near-optimal configuration. This (near-optimal) configuration had a latency of 24 FPS (which was 42% better than TX1 and 6× better than the fault). While SMAC also meets the developer's expectations, it performed slightly worse than CADET, and took 4 hours (6× longer than CADET). Further, the near-optimal configuration found by SMAC, change several unrelated configurations which were not recommended by the experts (✔ in Fig. 5).
• BugDoc (ML-based approach) has the least improvement compared to other approaches (21% improvement over TX1) while taking 3.5 hours (mostly spent on collecting training samples to train

internal the decision tree.). BugDoc failed to meet the developer's expectation. BugDoc also changed several unrelated configurations (depicted by ✔) not endorsed by the domain experts.

Why CADET works better (and faster)? CADET discovers the misconfigurations by constructing a causal model (a simplified version of this is shown in Fig. 5). This causal model rules out irrelevant configuration options and focuses on the configurations that have the highest (direct or indirect) causal effect on latency, e.g., we found the root-cause CUDA STATIC in the causal graph which indirectly affects latency via context-switches (an intermediate system event); this is similar to other relevant configurations that indirectly affected latency (via energy consumption). Using counterfactual queries, CADET can reason about changes to configurations with the highest average causal effect (last column in Fig. 5). The counterfactual reasoning occurs no additional measurements, significantly speeding up inference.

Together, the causal model and the counterfactual reasoning enable CADET to pinpoint the configuration options that were misconfigured and recommend a fix for them in a timely manner. As shown in Fig. 5, CADET accurately finds all the configuration options recommended by the forum (depicted by ✔ in Fig. 5). Further, CADET recommends fixes to these options that result in 24% better latency than the recommendation by domain experts in the forum. More importantly, CADET takes only 24 minutes (vs. 2 days of forum discussion) without modifying unrelated configurations.

6 Experimental Setup

6.1 Study subjects

Hardware Systems. This study uses three NVIDIA Jetson Platforms: TX1, TX2, and Xavier. Each platform has different hardware specifications, e.g., different CPU micro-architectures (ARM Caramel, Cortex), GPU micro-architectures (Maxwell, Pascal, Volta), energy requirements (5W–30W).

Software systems. We deploy five software systems on each NVIDIA Jetson Platform: (1) CNN based Image recognition with Xception to classify 5000 images from the CIFAR10 dataset [18]; (2) BERT (a transformer-based model) to perform sentiment analysis on 10000 reviews from the IMDb dataset [24]; (3) DeepSpeech an RNN based voice recognition on 5sec long audio files [41]; (4) SQLite, a database management system, to perform read, write, and insert operations; and (5) x264 video encoder to encode a video file of size 11MB with resolution 1920 x 1080.

Configuration Options. This work uses 28 configuration options including 10 software configurations, 8 OS/Kernel configurations, and 10 hardware configurations (c.f. Fig. 5). We also record the status of 19 non-intervenable system events. We choose these configurations and system events based on NVIDIA's configuration guides/tutorials and other related work [37].

6.2 Benchmark Dataset

We curate a non-functional faults dataset, called the Jetson Faults dataset, for each of the software and hardware system used in our study. This dataset and the data collection scripts are available in our replication package.

Data Collection. We exhaustively set each configuration option to all permitted values. Note, for continuous configuration options we choose 10 equally spaced values between the minimum and maximum permissible values, e.g., for GPU memory growth, we vary the value between 0% and 100% in steps of 10%. Next, we measure the latency, energy consumption, and heat dissipation for every configuration. We repeat each measurement 5 times and record the average to handle system noise and other variabilities [47].

Labeling misconfigurations. By definition, non-functional faults have latency, energy consumption, and heat dissipation that take tail values [35, 57], i.e., they are worse than the 99^{th} percentile. We filter our data set to find the configurations that result in tail values for latency and/or energy consumption, and label these configurations as '*faulty*'.

Ground Truth. We create a *ground-truth data* by inspecting each configuration labeled faulty and identifying their root-causes manually. We also manually reconfigure each faulty configuration to achieve the best possible latency, energy consumption, and/or heat dissipation thereby dispensing with the non-functional fault . For each reconfiguration, we record the changed configuration options and the values; these are used as a reference for evaluation.

6.3 Baselines

CADET is compared against four alternative state of-the-art ML-based methods for fault diagnostic and mitigation: (1) **CBI** [90]: A correlation based feature selection algorithm for identifying and fixing fault inducing configurations; (2) **Delta Debugging** [9]: A debugging technique that minimizes the difference between a pair of configurations, where one configuration causes a fault while the other does not; (3) **EnCore** [106]: A rule based technique that learns correlational information about misconfigurations from a given set of sample configurations. (4) **BugDoc** [66]: A Debugging Decision Tree based approach to automatically infer the root causes and derive succinct explanations of failures.

We also compare CADET with the near-optimal configurations from two state-of-the-art configuration optimization methods: (1) **SMAC** [46]: A sequential model based optimization technique to auto-tune software systems; (2) **PESMO** [43]: A multi-objective bayesian optimization to find Pareto-optimal configurations.

6.4 Evaluation Metrics

Relevance scores. We evaluate the predicted root-causes in terms of (a) the true positive rate (recall), (b) the true discovery rate (precision), and (c) accuracy (jaccard similarity). We prefer high accuracy, precision, and recall. To compute these metrics, we compare the set of configuration options identified by CADET to be the root cause with the true root-cause from the ground truth (§6.2).

Repair quality. To assess the quality of fixes, we measure the percentage improvement (gain %) after applying the recommended repairs using Δ_{gain} defined as: $\Delta_{gain} = \frac{\text{NFP}_{\text{FAULT}} - \text{NFP}_{\text{NOFAULT}}}{\text{NFP}_{\text{FAULT}}} \times 100$. Here $\text{NFP}_{\text{FAULT}}$ is the value of the faulty non-functional property (latency,*etc.*) and $\text{NFP}_{\text{NO FAULT}}$ is the the value of the faulty non-functional property after applying the repairs recommended by CADET. The larger the Δ_{gain}, the better the recommended fix.

Table 1: Efficiency of CADET compared to other approaches. Cells highlighted in blue indicate improvement over faults.

(a) Single objective performance fault in latency and energy consumption.

			Accuracy					Precision					Recall					Gain					Time[†]	
			CADET	CBI	δ-Debug	EnCore	BugDoc	**CADET**	CBI	δ-Debug	EnCore	BugDoc	**CADET**	CBI	δ-Debug	EnCore	BugDoc	**CADET**	CBI	δ-Debug	EnCore	BugDoc	**CADET**	Others
TX2	Latency	Image	**84**	66	65	68	71	**86**	67	61	63	67	**83**	64	68	69	62	**82**	48	42	57	59	**0.6**	4
		NLP	**75**	65	60	66	66	**76**	57	55	61	73	**71**	74	68	67	65	**74**	54	59	62	58	**0.4**	4
		Speech	**76**	64	63	63	72	**76**	58	69	61	71	**81**	73	61	63	69	**76**	59	53	55	66	**0.7**	4
		x264	**81**	67	60	61	70	**82**	69	58	65	66	**78**	64	67	63	72	**85**	69	72	68	71	**1.4**	4
Xavier	Energy	Image	**77**	63	55	63	64	**78**	56	58	66	65	**80**	69	55	63	68	**83**	59	50	35	51	**0.4**	4
		NLP	**76**	60	63	66	64	**70**	62	64	64	65	**79**	61	54	63	66	**62**	49	36	49	53	**0.5**	4
		Speech	**76**	66	65	61	71	**75**	55	59	54	68	**78**	53	52	59	71	**78**	64	48	65	63	**1.2**	4
		x264	**78**	62	57	59	67	**83**	63	53	61	66	**78**	67	53	54	72	**87**	73	71	76	76	**0.3**	4

(b) Multi-objective non-functional faults in *Energy, Latency*.

			Accuracy				Precision				Recall				Gain (Latency)				Gain (Energy)				Time[†]	
			CADET	CBI	EnCore	BugDoc	**CADET**	CBI	EnCore	BugDoc	**CADET**	CBI	EnCore	BugDoc	**CADET**	CBI	EnCore	BugDoc	**CADET**	CBI	EnCore	BugDoc	**CADET**	Others
Energy + Latency		Image	**80**	54	55	65	**77**	53	54	62	**81**	59	59	62	**84**	53	61	65	**75**	38	46	44	**0.9**	4
		NLP	**76**	51	56	65	**77**	42	56	63	**79**	59	62	65	**84**	53	59	61	**67**	41	27	48	**0.5**	4
		Speech	**81**	50	61	65	**80**	44	53	62	**81**	51	59	64	**88**	55	55	62	**77**	43	43	41	**1.1**	4
		x264	**81**	54	55	66	**83**	50	54	67	**80**	63	62	61	**75**	62	64	66	**76**	64	66	64	**1**	4

[†] Wallclock time in hours

7 Experimental Results

RQ1. CADET vs. Model-Based Diagnostics

Motivation. Machine learning approaches are commonly used in model-based fault diagnostics [9, 66, 90, 106]. These models are trained to learn the correlation between the configuration options and the non-functional properties (e.g, latency, *etc*) which are then extrapolated to diagnose and fix faults. This research question compares CADET and four state-of-the-art ML-based methods (described in §6.3) for single objective fault diagnostics and multi-objective fault diagnostics.

Note. All approaches require some initial observational data to operate. For CADET we begin with 25 initial samples to let CADET incrementally generate, evaluate, and update the causal model with candidate repairs. Other methods require a larger pool of training data. However, collecting the training data is expensive and time-consuming. Therefore, we use a budget of 3 hours to generate random configuration samples to train the other methods.

Approach. We assess the effectiveness of diagnostics for two types of non-functional faults: (a) "single-objective" faults that occur either in latency or in energy consumption, and (b) "multi-objective" faults where misconfigurations affect multiple non-functional properties simultaneously, i.e., in latency *and* energy consumption. For brevity, we evaluate latency faults in TX2 and energy consumption faults in Xavier. Our findings generalize over other hardware.

Approach. Tables 1a and 1b compare the effectiveness of CADET over other ML-based fault diagnosis approaches. We observe:

• ***Accuracy, precision, and recall:*** CADET *significantly outperforms ML-based methods in all cases*. For example, in image recognition with Xception on TX2, CADET achieves 13% more accuracy, 19% larger Precision compared to BugDoc and 14% more recall compared to EnCore. We observe similar trends in energy faults and multi-objective faults, i.e., CADET outperforms other methods.

• ***Gain.*** CADET *can recommend repairs for faults that significantly improve latency and energy usage*. Applying the changes to the configurations recommended by CADET increases the performance drastically. We observed latency gains as high as 83% (23% more than BugDoc) on TX2 and energy gain of 83% (32% more than BugDoc) on Xavier for image recognition.

• ***Wallclock time.*** CADET *can resolve misconfiguration faults significantly faster than ML-based approaches*. In Tables 1a and 1b, the last two columns indicate the time taken (in hours) by each approach to diagnosing the root cause. For all methods, we set a maximum budget of 4 hours. We find that, while other approaches use the entire budget to diagnose and resolve the faults, CADET can do so significantly faster, e.g., CADET is 13× faster in diagnosing and resolving faults in energy usage for x264 deployed on Xavier and 10× faster for latency faults for NLP task on TX2.

Discussion. Based on the above, we answer the following additional questions:

1. Why is CADET better than model-based approaches? Model-based methods rely only on correlation, and this can be misleading since they cannot incorporate the intrinsic complex causal structure of the underlying configuration space. In contrast, CADET relies on causal inference (instead of correlation) to model the configuration space thereby overcoming the limitations with current ML-models.

2. Why is CADET faster than model-based approaches? ML-based methods require a large number of initial observational data for training. They spend most of their allocated 4-hour budget on gathering these training samples. In contrast, CADET starts with only 25 samples and uses incremental learning (§4.4) to judiciously
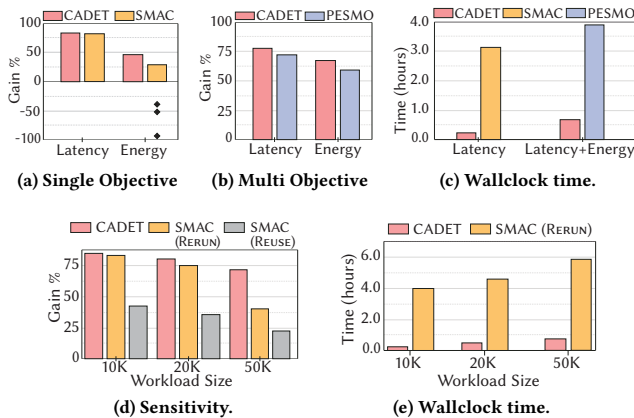
(a) Single Objective   **(b) Multi Objective**   **(c) Wallclock time.**

(d) Sensitivity.   **(e) Wallclock time.**

Figure 6: CADET vs. optimization with SMAC and PESMO.

update the casual graph with new configurations until a repair has been found. This drastically reduces the inference time.

RQ2: CADET vs. Search-Based Optimization

Motivation. Search based methods use techniques such as Bayesian optimization [85, 89] to explore the relationships between configuration options and corresponding non-functional properties. Albeit recent approaches seem promising [23, 45, 70], these methods are not viable for diagnostics since our objective is not to find an optimal configuration, rather it is to fix an already encountered fault. In this RQ, we explore the challenges of using optimization for fault diagnostics in terms of their efficacy (RQ2-A), their analysis time (RQ2-B), and their sensitivity to external changes (RQ2-C), by comparing CADET with two popular optimization approaches: (a) SMAC and (b) PESMO (see §6.3).

RQ2-A. How effective is optimization? In this RQ, we compare CADET with SMAC for fixing latency faults in TX2 (for Image recognition). In general, the fixes recommended by CADET were observed to have better latency gains than the near-optimal configuration discovered by SMAC (see Fig. 6a). Most notably, in numerous cases, SMAC generated "near-optimal" configurations which in fact lead to significant deterioration in other non-functional properties such as energy consumption and heat dissipation (denoted by ♦ in Fig. 6a). Out of 21 latency faults in image recognition on TX2, we found 7 cases (33%) where the near-optimal configuration caused a significant deterioration in energy consumption. In one case the optimal configuration increased the energy consumption by 96%! To overcome the above problem, we use a multi-objective optimization such as PESMO [43] to find a configuration that optimizes for energy and latency simultaneously. CADET can natively accommodate such multi-objective by slightly altering the counterfactual query (c.f. §4.3).

Fig. 6b compares CADET with PESMO for fixing energy and latency faults in TX2 (for Image recognition). As with the previous instance, we find that CADET performs better than PESMO, i.e., CADET recommends fixes misconfigurations that result in better gain than the Pareto-optimal configuration found by PESMO.

RQ2-B. What is the analysis time of optimization? The analysis time of all approaches are reported in Fig. 6c. Optimization

methods run until they converge, however it is a common practice to assign a maximum budget (we use 4 hours) to limit inference time and cost. CADET was found to have comparatively lesser inference time compared to the other methods. While SMAC takes slightly more than 3 hours to find a near-optimum in TX2 (image recognition), CADET resolved the faults in 20 minutes on average (9× faster). Likewise, PESMO takes close 4 hours compared to 45 minutes with CADET (5× faster).

Since CADET uses causal models to infer candidate fixes using only the available observational data, it tends to be much faster than SMAC and PESMO. During incremental learning, CADET judiciously evaluates the most promising fixes until the fault is resolved. This limits the number of times the system is reconfigured and profiled. In contrast, optimization techniques sample the configuration space exhaustively to find an optimal configuration. Every sampled configuration is deployed and profiled. Optimization techniques sample several sub-optimal configurations in their search process and they are all evaluated. This issue is compounded when performing multi-objective optimization.

RQ2-C. How sensitive are optimization techniques to external changes in the software? Configuration optimization methods are prone to be very sensitive to minor changes to the software stack. To demonstrate this, we use three larger additional image recognition workloads: 10K, 20K, and 50K test images (previous experiments used 5K test images). We evaluate two variants of SMAC: (1) SMAC (REUSE) where we *reuse* the near-optimum found with a 5K tests image on the larger workloads; and (2) SMAC (RERUN) where we *rerun* SMAC afresh on each workload.

Our experimental results demonstrate that CADET performs better than the two variants of SMAC (c.f. Fig. 6d). SMAC (REUSE) performs the worst when the workload changes. With 10K images, reusing the near-optimal configuration from 5K images results in latency gain of 39%, compared to 79% with CADET. With a workload size of 50K images, SMAC (REUSE) only achieves a latency improvement of 24% over the faulty configurations whereas CADET finds a fix that improves latency by 72%. SMAC (RERUN) performs better than SMAC (REUSE) but worse than CADET.

In Fig. 6e we report the inference times for SMAC (RERUN) and CADET. In keeping with the previous results, SMAC (RERUN) takes significantly longer than CADET, i.e., SMAC (RERUN) exceeds the 4 hour budget in all workloads whereas CADET takes at most 30 minutes for the largest workload to diagnose and fix the latency faults. We have to rerun SMAC every time the workload changes, and this limits its practical usability. In contrast, CADET incrementally updates the internal causal model with new samples from the larger workload to learn new relationships from the larger workload. Therefore, it is less sensitive and much faster.

RQ3. Scalability of CADET

Motivation. To examine the scalability of CADET, we attempt to remedy latency issues in SQLite database deployed on XAVIER by running the speedtest workload. SQLite was chosen because it offers a large number of configurable options, much more than DNN-based applications. These configurations control the compute, memory, data reuse operations of SQLite. Further, each of these

Table 2: Scalability of CADET for SQLite on Xavier.

Configs	Events	Paths	Queries	Degree	Gain (%)	Time/Fault (in sec.)		Total
						Discovery	Query Eval	
28	19	32	191	3.6	93	9	14	**291**
242	88	181	3136	1.8	94	68	141	**1692**
242	288	441	22372	1.6	92	111	854	**5312**

options can take on a large number of permitted values making SQLite a useful candidate to study scalability of CADET.

Approach. The causal graphs of previous experiments used 47 variables with 28 hardware, kernel/OS, and software configuration options in addition to 19 system events (c.f. §6.2). With SQLite, there are two additional scenarios: (a) 28 configuration options and 19 system events when we manually select the most relevant software/hardware options and events, (b) 242 configuration options and 88 events when we select all modifiable software and hardware options and system events, and (c) 242 configuration options and 288 events when we select not only all modifiable software and hardware options and system events but also intermediate `tracepoint events`. For each scenario, we track the time required to discover the causal structure, evaluate the counterfactual queries, and total time to resolve latency faults.

Results. Our findings, tabulated in Table 2, show that, with a larger number of configuration options, there are significantly more paths and queries from CADET. This adds to the causal graph discovery and query evaluation time. However, with as much as 242 configuration options and 88 events (Table 2, row 2), causal graph discovery takes roughly one minute, evaluating all 3136 queries takes roughly two minutes, and the total time to diagnose and fix a fault is roughly 28 minutes. This trend is observed even when the with 242 configuration options, 288 events (Table 2, row 3), and finer granularity of configuration values—the time required to causal model recovery is less than two minutes and the total time to diagnose and fix a fault is less than 2 hours. This indicates that CADET can scale to a much larger configuration space without an exponential increase in runtime for any of the intermediate stages. This property can be attributed to the sparsity of the causal graph (average degree of a node in Table 2 is at most 3.6 and it reduces to 1.6 when the number of configurations increase). This makes sense because not all variables (i.e., configuration options and/or system events) affect non-functional properties and a high number of variables in the graph end up as isolated nodes. Therefore, the number of paths and consequently the evaluation time do not grow exponentially as the number of variables increase.

Finally, the latency gain associated with repairs from larger configuration space with configurations was similar to the original space of 28 configurations (both were 93%). This indicates that: (a) imparting domain expertise to select most important configuration options can speedup the inference time of CADET, and (b) if the user chooses instead to use more configuration options (perhaps to avoid initial feature engineering), CADET can still diagnose and fix faults satisfactorily within a reasonable time.

8 Related Work

Performance Faults in Configurable Systems. Previous empirical studies have shown that a majority of performance issues are due to misconfigurations [39], with severe consequences in production environments [67, 92], and configuration options that cause such performance faults force the users to tune the systems themselves [108]. However, the configuration options keep increasing over time [103] and make the configuration-aware testing a difficult endeavor, where testing can be only practical by evaluating a representative sample of all possible configurations [81], by employing sampling strategies to discover potential (performance) faults [68]. Previous works have used static and dynamic program analysis to identify the influence of configuration options on performance [65, 97, 98] and to detect and diagnose misconfigurations [10, 11, 104, 107]. None of the white-box analysis approaches target configuration space across system stack, where it limits their applicabilities in identifying the true causes of a performance fault. The evaluation results indicate that by targeting the whole system stack, CADET goes beyond the existing approaches in terms of detecting and fixing performance faults.

Performance Modeling and Optimization. Performance behavior of configurable systems is complex (non-linear, non-convex, and multi-modal) [50, 99], where such behavior becomes even more intricate when multiple objectives are needed to be traded-off for performance-related tasks [48, 61, 73]. A common strategy to understand such complex performance behavior is to use machine-learning by building a black box model that characterizes system performance [36, 86, 96]. The learned performance model can then be used for performance debugging [36, 86, 99] and tuning [44]. While the sheer size of the configuration space and complex performance behavior prohibit any guarantee of finding a globally optimal configuration, optimization algorithms attempt to find near-optimal configurations under a limited sampling budget using a clever combination of sampling, model construction, and search. While there is no silver bullet [33], several methods have been attempted including hill climbing [101], optimization via guessing [77], Bayesian optimization [49], and multi-objective optimization [48, 100]. However, if one of the contextual/environmental conditions (e.g., hardware, workload) of the system changes, one needs to rerun the optimization approach again [52], therefore, transfer learning in terms of transferring measurement data or extracting some types of knowledge across environments has been used [47, 51, 53, 54, 63, 95, 96]. For a more comprehensive treatment of the literature, we refer to [80, 88]. Although black-box models have used extensively for optimization, they cannot diagnose non-functional faults as they are unaware of underlying causal factors due to their blind exploration without having access to the underlying causal relationships.

Statistical and Model-based Debugging. Debugging approaches such as Statistical Debugging [90], HOLMES [16], XTREE [62], BugDoc [66], EnCore [66], Rex [69], and PerfLearner [40] have been proposed to detect root causes of system faults. These methods make use of statistical diagnosis and pattern mining to rank the probable causes based on their likelihood of being the root causes of faults. However, these approaches may produce correlated predicates that lead to incorrect explanations.

Causal Testing and Profiling. Causal inference has been used for fault localization [12, 28], resource allocations in cloud systems [31], and causal effect estimation for advertisement recommendation systems [14]. More recently, AID [27] detects root causes of a intermittent software failure using fault injection as interventions and utilizing a guided search to pinpoint the root causes and generate an explanation how the root cause triggers the failure. Similarly LDX [64] determines causal dependencies of events for detection of information leak and security attacks. CAUSAL TESTING and HOLMES [55] modifies the system inputs to observe behavioral changes in order to build a causal model and then utilizes counterfactual reasoning to find the root causes of bugs. Causal profiling approaches like CoZ [22] points developers where optimizations will improve performance and quantifies their potential impact, where the interventions happen by virtual program speedups. Causal inference methods like X-RAY [10] and CONFAID [11] had previously been applied to analyze program failures using run-time control and data flow behavior. All approaches above are either orthogonal or complimentary to CADET, mostly they focus on functional bugs (e.g., CAUSAL TESTING) or if they are performance related, they are not configuration-aware (e.g., CoZ).

9 Conclusion

Modern computer systems are highly-configurable with thousands of interacting configurations with a complex performance behavior. Misconfigurations in these systems can elicit complex interactions between software and hardware configuration options resulting in non-functional faults. We propose CADET, a novel approach for diagnostics that learns and exploits the causal structure of configuration options, system events, and performance metrics. Our evaluation shows that CADET effectively and quickly diagnoses the root cause of non-functional faults and recommends high-quality repairs to mitigate these faults.

Acknowledgements

References

[1] 2017. PCIE x4, only 658MB/s - Jetson & Embedded Systems / Jetson TX2. In NVIDIA Developer Forums: https://forums.developer.nvidia.com/t/pcie-x4-only-658mb-s/55357.

[2] 2017. Slow Image Classification with Tensorflow on TX2. In NVIDIA developer forums: https://forums.developer.nvidia.com/t/54307.

[3] 2018. Question about thermal management. In NVIDIA developer forums: https://forums.developer.nvidia.com/t/59855.

[4] 2019. 99% of misconfiguration incidents in the cloud go unnoticed. In Help Net Security: https://www.helpnetsecurity.com/2019/09/25/cloud-misconfiguration-incidents/.

[5] 2019. Jetson Nano not responding or responding very slow over SSH. In NVIDIA Developer Forums: https://forums.developer.nvidia.com/t/jetson-nano-not-responding-or-responding-very-slow-over-ssh/79530/3.

[6] 2020. CUDA performance issue on TX2. In NVIDIA developer forums: https://forums.developer.nvidia.com/t/50477.

[7] 2020. General Performance Problems. In NVIDIA developer forums: https://forums.developer.nvidia.com/t/111704.

[8] 2020. High CPU usage on Jetson TX2 with GigE fully loaded. In NVIDIA developer forums: https://forums.developer.nvidia.com/t/124381.

[9] Cyrille Artho. 2011. Iterative delta debugging. International Journal on Software Tools for Technology Transfer 13, 3 (2011), 223–246.

[10] Mona Attariyan, Michael Chow, and Jason Flinn. 2012. X-ray: Automating root-cause diagnosis of performance anomalies in production software. In Presented as part of the 10th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 12). 307–320.

[11] Mona Attariyan and Jason Flinn. 2010. Automating Configuration Troubleshooting with Dynamic Information Flow Analysis.. In OSDI, Vol. 10. 1–14.

[12] George K Baah, Andy Podgurski, and Mary Jean Harrold. 2010. Causal inference for statistical fault localization. In Proceedings of the 19th international symposium on Software testing and analysis. 73–84.

[13] Dikla Barda, Roman Zaikin, and Yaara Shriki. 2020. Keeping the gate locked on your IoT devices: Vulnerabilities found on Amazon's Alexa. By Check Point Research: https://research.checkpoint.com/2020/amazons-alexa-hacked/.

[14] Léon Bottou, Jonas Peters, Joaquin Quiñonero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. 2013. Counterfactual reasoning and learning systems: The example of computational advertising. The Journal of Machine Learning Research 14, 1 (2013), 3207–3260.

[15] Randal E Bryant, O'Hallaron David Richard, and O'Hallaron David Richard. 2003. Computer systems: a programmer's perspective. Vol. 2.

[16] Trishul M Chilimbi, Ben Liblit, Krishna Mehra, Aditya V Nori, and Kapil Vaswani. 2009. HOLMES: Effective statistical debugging via efficient path profiling. In 2009 IEEE 31st International Conference on Software Engineering. IEEE, 34–44.

[17] Chia-Chang Chiu and Yi-Sheng Chueh. 2012. Method and computer system for thermal throttling protection. US Patent 8,301,873.

[18] François Chollet. 2017. Xception: Deep learning with depthwise separable convolutions. In Proceedings of the IEEE conference on computer vision and pattern recognition. 1251–1258.

[19] Diego Colombo and Marloes H Maathuis. 2014. Order-independent constraint-based causal structure learning. The Journal of Machine Learning Research 15, 1 (2014), 3741–3782.

[20] Diego Colombo, Marloes H Maathuis, Markus Kalisch, and Thomas S Richardson. 2012. Learning high-dimensional directed acyclic graphs with latent and selection variables. The Annals of Statistics (2012), 294–321.

[21] Lynne M Connelly. 2016. Fisher's exact test. Medsurg Nursing 25, 1 (2016).

[22] Charlie Curtsinger and Emery D Berger. 2015. Coz: Finding code that counts with causal profiling. In Proceedings of the 25th Symposium on Operating Systems Principles. 184–197.

[23] Valentin Dalibard, Michael Schaarschmidt, and Eiko Yoneki. 2017. BOAT: Building auto-tuners with structured Bayesian optimization. In Proceedings of the 26th International Conference on World Wide Web. 479–488.

[24] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018).

[25] Andrie Ene. 2020. Alexa Hack Jeopardized Echo Users Network. http://bit.ly/AlexaHack2020.

[26] Robin J Evans and Thomas S Richardson. 2014. Markovian acyclic directed mixed graphs for discrete data. The Annals of Statistics (2014), 1452–1482.

[27] Anna Fariha, Suman Nath, and Alexandra Meliou. 2020. Causality-Guided Adaptive Interventional Debugging. In Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data. 431–446.

[28] Farid Feyzi and Saeed Parsa. 2019. Inforence: effective fault localization based on information-theoretic analysis and statistical causal inference. Frontiers of Computer Science 13, 4 (2019), 735–759.

[29] Antonio Filieri, Henry Hoffmann, and Martina Maggio. 2015. Automated multi-objective control for self-adaptive software design. In Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering. 13–24.

[30] Dustin Franklin. 2017. NVIDIA Jetson TX2 Delivers Twice the Intelligence to the Edge. In NVIDIA Developer Blog: https://developer.nvidia.com/blog/jetson-tx2-delivers-twice-intelligence-edge/.

[31] Philipp Geiger, Lucian Carata, and Bernhard Schölkopf. 2016. Causal models for debugging and control in cloud computing. arXiv preprint arXiv 1603 (2016).

[32] Clark Glymour, Kun Zhang, and Peter Spirtes. 2019. Review of causal discovery methods based on graphical models. Frontiers in genetics 10 (2019), 524.

[33] Alexander Grebhahn, Norbert Siegmund, and Sven Apel. 2019. Predicting Performance of Software Configurations: There is no Silver Bullet. arXiv preprint arXiv:1911.12643 (2019).

[34] Jonathan Greig. 2020. Cloud misconfigurations cost companies nearly $5 trillion. In TechRepublic: https://www.techrepublic.com/article/cloud-misconfigurations-cost-companies-nearly-5-trillion/.

[35] Haryadi S Gunawi et al. 2018. Fail-slow at scale: Evidence of hardware performance faults in large production systems. ACM Transactions on Storage (TOS) 14, 3 (2018), 1–26.

[36] Jianmei Guo, Krzysztof Czarnecki, Sven Apel, Norbert Siegmund, and Andrzej Wąsowski. 2013. Variability-aware performance prediction: A statistical learning approach. In 2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE). IEEE.

[37] Hassan Halawa, Hazem A. Abdelhafez, Andrew Boktor, and Matei Ripeanu. 2017. NVIDIA jetson platform characterization. *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)* 10417 LNCS (2017), 92–105. https://doi.org/10.1007/978-3-319-64203-1_7

[38] Axel Halin, Alexandre Nuttinck, Mathieu Acher, Xavier Devroey, Gilles Perrouin, and Benoit Baudry. 2019. Test them all, is it worth it? Assessing configuration sampling on the JHipster Web development stack. *Empirical Software Engineering* 24, 2 (2019), 674–717.

[39] Xue Han and Tingting Yu. 2016. An empirical study on performance bugs for highly configurable software systems. In *Proceedings of the 10th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement.*

[40] Xue Han, Tingting Yu, and David Lo. 2018. PerfLearner: learning from bug reports to understand and generate performance test frames. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE).*

[41] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al. 2014. Deep speech: Scaling up end-to-end speech recognition. *arXiv preprint arXiv:1412.5567* (2014).

[42] Yang-Bo He and Zhi Geng. 2008. Active learning of causal networks with intervention experiments and optimal designs. *Journal of Machine Learning Research* 9, Nov (2008), 2523–2547.

[43] Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. 2016. Predictive entropy search for multi-objective bayesian optimization. In *International Conference on Machine Learning.* 1492–1501.

[44] Holger H Hoos. 2011. Automated algorithm configuration and parameter tuning. In *Autonomous search.* Springer, 37–71.

[45] Chin-Jung Hsu, Vivek Nair, Vincent W Freeh, and Tim Menzies. 2018. Arrow: Low-level augmented bayesian optimization for finding the best cloud vm. In *IEEE 38th International Conference on Distributed Computing Systems.*

[46] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2011. Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization.* Springer, 507–523.

[47] Md Shahriar Iqbal, Lars Kotthoff, and Pooyan Jamshidi. 2019. Transfer Learning for Performance Modeling of Deep Neural Network Systems. In *2019 {USENIX} Conference on Operational Machine Learning (OpML 19).* 43–46.

[48] Md Shahriar Iqbal, Jianhai Su, Lars Kotthoff, and Pooyan Jamshidi. 2020. FlexiBO: Cost-Aware Multi-Objective Optimization of Deep Neural Networks. *arXiv preprint arXiv:2001.06588* (2020).

[49] Pooyan Jamshidi and Giuliano Casale. 2016. An Uncertainty-Aware Approach to Optimal Configuration of Stream Processing Systems. In *Proc. Int'l Symp. on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)* (London, UK). IEEE, 10. https://doi.org/10.1109/MASCOTS.2016.17

[50] Pooyan Jamshidi and Giuliano Casale. 2016. An uncertainty-aware approach to optimal configuration of stream processing systems. In *2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS).* IEEE, 39–48.

[51] Pooyan Jamshidi, Norbert Siegmund, Miguel Velez, Christian Kästner, Akshay Patel, and Yuvraj Agarwal. 2017. Transfer Learning for Performance Modeling of Configurable Systems: An Exploratory Analysis. In *Proc. Int'l Conf. Automated Software Engineering (ASE).* ACM.

[52] Pooyan Jamshidi, Miguel Velez, Christian Kästner, and Norbert Siegmund. 2018. Learning to Sample: Exploiting Similarities Across Environments to Learn Performance Models for Configurable Systems. In *Proc. Int'l Symp. Foundations of Software Engineering (FSE).* ACM.

[53] Pooyan Jamshidi, Miguel Velez, Christian Kästner, Norbert Siegmund, and Prasad Kawthekar. 2017. Transfer Learning for Improving Model Predictions in Highly Configurable Software. In *Proc. Int'l Symp. Soft. Engineering for Adaptive and Self-Managing Systems (SEAMS)* (Buenos Aires, Argentina). IEEE, 11.

[54] Mohammad Ali Javidian, Pooyan Jamshidi, and Marco Valtorta. 2019. Transfer learning for performance modeling of configurable systems: A causal analysis. *arXiv preprint arXiv:1902.10119* (2019).

[55] Brittany Johnson, Yuriy Brun, and Alexandra Meliou. 2020. Causal Testing: Understanding Defects' Root Causes. In *Proceedings of the 2020 International Conference on Software Engineering.*

[56] Konstantinos Kanellis, Ramnatthan Alagappan, and Shivaram Venkataraman. 2020. Too Many Knobs to Tune? Towards Faster Database Tuning by Pre-selecting Important Knobs. In *12th {USENIX} Workshop on Hot Topics in Storage and File Systems (HotStorage 20).*

[57] Martin Kleppmann. 2017. *Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems.* " O'Reilly Media, Inc.".

[58] Rex B Kline. 2015. *Principles and practice of structural equation modeling.* Guilford publications.

[59] Murat Kocaoglu, Alexandros G. Dimakis, Sriram Vishwanath, and Babak Hassibi. 2017. Entropic Causal Inference. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (San Francisco, California, USA). 1156–1162.

[60] Murat Kocaoglu, Sanjay Shakkottai, Alexandros Dimakis, Constantine Caramanis, and Sriram Vishwanath. 2020. Applications of Common Entropy for Causal Inference. NeurIPS (2020).

[61] Sergiy Kolesnikov, Norbert Siegmund, Christian Kästner, Alexander Grebhahn, and Sven Apel. 2019. Tradeoffs in modeling performance of highly configurable software systems. *Software & Systems Modeling* 18, 3 (2019), 2265–2283.

[62] Rahul Krishna, Tim Menzies, and Lucas Layman. 2017. Less is more: Minimizing code reorganization using XTREE. *Information and Software Technology* 88 (2017), 53–66.

[63] Rahul Krishna, Vivek Nair, Pooyan Jamshidi, and Tim Menzies. 2020. Whence to Learn? Transferring Knowledge in Configurable Systems using BEETLE. *IEEE Transactions on Software Engineering* (2020).

[64] Yonghwi Kwon, Dohyeong Kim, William Nick Sumner, Kyungtae Kim, Brendan Saltaformaggio, Xiangyu Zhang, and Dongyan Xu. 2016. Ldx: Causality inference by lightweight dual execution. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems.* 503–515.

[65] Chi Li, Shu Wang, Henry Hoffmann, and Shan Lu. 2020. Statically inferring performance properties of software configurations. In *Proceedings of the Fifteenth European Conference on Computer Systems.* 1–16.

[66] Raoni Lourenço, Juliana Freire, and Dennis Shasha. 2020. BugDoc: A System for Debugging Computational Pipelines. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data.* 2733–2736.

[67] Ben Maurer. 2015. Fail at scale: Reliability in the face of rapid change. *Queue* 13, 8 (2015), 30–46.

[68] Flávio Medeiros, Christian Kästner, Márcio Ribeiro, Rohit Gheyi, and Sven Apel. 2016. A comparison of 10 sampling algorithms for configurable systems. In *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE).* IEEE.

[69] Sonu Mehta, Ranjita Bhagwan, Rahul Kumar, Chetan Bansal, Chandra Maddila, B Ashok, Sumit Asthana, Christian Bird, and Aditya Kumar. 2020. Rex: Preventing bugs and misconfiguration in large services using correlated change analysis. In *17th {USENIX} Symposium on Networked Systems Design and Implementation.*

[70] Takashi Miyazaki, Issei Sato, and Nobuyuki Shimizu. 2018. Bayesian optimization of hpc systems for energy efficiency. In *International Conference on High Performance Computing.* Springer, 44–62.

[71] I Molyneaux. 2009. The Art of Application Performance Testing: Help for Programmers and Quality Assurance.[Sl]:" O'Reilly Media.

[72] Vivek Nair, Tim Menzies, Norbert Siegmund, and Sven Apel. 2018. Faster discovery of faster system configurations with spectral learning. *Automated Software Engineering* 25, 2 (2018), 247–277.

[73] Luigi Nardi, Artur Souza, David Koeplinger, and Kunle Olukotun. 2019. Hyper-Mapper: a Practical Design Space Exploration Framework. In *2019 IEEE 27th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS).* IEEE, 425–426.

[74] Adrian Nistor, Po-Chun Chang, Cosmin Radoi, and Shan Lu. 2015. Caramel: Detecting and fixing performance problems that have non-intrusive fixes. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering.*

[75] Adrian Nistor, Tian Jiang, and Lin Tan. 2013. Discovering, reporting, and fixing performance bugs. In *10th working conference on mining software repositories.*

[76] Juan Miguel Ogarrio, Peter Spirtes, and Joe Ramsey. 2016. A hybrid causal search algorithm for latent variable models. In *Conference on Probabilistic Graphical Models.* 368–379.

[77] Takayuki Osogami and Sei Kato. 2007. Optimizing system configurations quickly by guessing at the performance. In *Int'l Conference on Measurement and Modeling of Computer Systems (SIGMETRICS).*

[78] Judea Pearl. 2009. *Causality.* Cambridge university press.

[79] Judea Pearl and Dana Mackenzie. 2018. *The book of why: the new science of cause and effect.* Basic Books.

[80] Juliana Alves Pereira, Hugo Martin, Mathieu Acher, Jean-Marc Jézéquel, Goetz Botterweck, and Anthony Ventresque. 2019. Learning Software Configuration Spaces: A Systematic Literature Review. *arXiv preprint arXiv:1906.03018* (2019).

[81] Xiao Qu, Myra B Cohen, and Gregg Rothermel. 2008. Configuration-aware regression testing: an empirical study of sampling and prioritization. In *Proceedings of the 2008 international symposium on Software testing and analysis.*

[82] C Mylara Reddy and N Nalini. 2016. Fault Tolerant Cloud Software Systems Using Software Configurations. In *2016 IEEE International Conference on Cloud Computing in Emerging Markets (CCEM).* IEEE, 61–65.

[83] Thomas Richardson, Peter Spirtes, et al. 2002. Ancestral graph Markov models. *The Annals of Statistics* 30, 4 (2002), 962–1030.

[84] Ana B Sánchez, Pedro Delgado-Pérez, Inmaculada Medina-Bulo, and Sergio Segura. 2020. TANDEM: A Taxonomy and a Dataset of Real-World Performance Bugs. *IEEE Access* 8 (2020), 107214–107228.

[85] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P Adams, and Nando De Freitas. 2015. Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* 104, 1 (2015), 148–175.

[86] Norbert Siegmund, Alexander Grebhahn, Sven Apel, and Christian Kästner. 2015. Performance-influence models for highly configurable systems. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering.* 284–294.

[87] Norbert Siegmund, Sergiy S Kolesnikov, Christian Kästner, Sven Apel, Don Batory, Marko Rosenmüller, and Gunter Saake. 2012. Predicting performance via automated feature-interaction detection. In *2012 34th International Conference*

on Software Engineering (ICSE). IEEE, 167–177.

[88] Norbert Siegmund, Nicolai Ruckel, and Janet Siegmund. 2020. Dimensions of software configuration: on the configuration context in modern software development. In *Proceedings of the 28th ESEC/FSE*. 338–349.

[89] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*. 2951–2959.

[90] Linhai Song and Shan Lu. 2014. Statistical debugging for real-world performance problems. *ACM SIGPLAN Notices* 49, 10 (2014), 561–578.

[91] Peter Spirtes, Clark N Glymour, Richard Scheines, and David Heckerman. 2000. *Causation, prediction, and search*. MIT press.

[92] Chunqiang Tang, Thawan Kooburat, Pradeep Venkatachalam, Akshay Chander, Zhe Wen, Aravind Narayanan, Patrick Dowell, and Robert Karl. 2015. Holistic configuration management at facebook. In *Proceedings of the 25th Symposium on Operating Systems Principles*. 328–343.

[93] Ferdian Thung, Shaowei Wang, David Lo, and Lingxiao Jiang. 2012. An empirical study of bugs in machine learning systems. In *2012 IEEE 23rd International Symposium on Software Reliability Engineering*. IEEE, 271–280.

[94] Sokratis Tsakiltsidis, Andriy Miranskyy, and Elie Mazzawi. 2016. On automatic detection of performance bugs. In *2016 IEEE international symposium on software reliability engineering workshops (ISSREW)*. IEEE, 132–139.

[95] Pavel Valov, Jianmei Guo, and Krzysztof Czarnecki. 2020. Transferring pareto frontiers across heterogeneous hardware environments. In *Proceedings of the ACM/SPEC International Conference on Performance Engineering*. 12–23.

[96] Pavel Valov, Jean-Christophe Petkovich, Jianmei Guo, Sebastian Fischmeister, and Krzysztof Czarnecki. 2017. Transferring performance prediction models across different hardware platforms. In *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering*. 39–50.

[97] Miguel Velez, Pooyan Jamshidi, Florian Sattler, Norbert Siegmund, Sven Apel, and Christian Kastner. 2019. Configcrusher: White-box performance analysis for configurable systems. *arXiv preprint arXiv:1905.02066* (2019).

[98] Miguel Velez, Pooyan Jamshidi, Norbert Siegmund, Sven Apel, and Christian Kästner. 2021. White-Box Analysis over Machine Learning: Modeling Performance of Configurable Systems. *arXiv preprint arXiv:2101.05362* (2021).

[99] Shu Wang, Chi Li, Henry Hoffmann, Shan Lu, William Sentosa, and Achmad Imam Kistijantoro. 2018. Understanding and auto-adjusting

performance-sensitive configurations. *ACM SIGPLAN Notices* 53, 2 (2018).

[100] Fan Wu, Westley Weimer, Mark Harman, Yue Jia, and Jens Krinke. 2015. Deep parameter optimisation. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. 1375–1382.

[101] Bowei Xi, Zhen Liu, Mukund Raghavachari, Cathy H Xia, and Li Zhang. 2004. A smart hill-climbing algorithm for application server configuration. In *13th International Conference on World Wide Web (WWW)*. ACM, 287–296.

[102] Shen Xinpeng, Ma Sisi, Vemuri Prashanthi, Gyorgy Simon, Michael W Weiner, Paul Aisen, Ronald Petersen, Clifford R Jack, Andrew J Saykin, William Jagust, et al. 2020. Challenges and Opportunities with Causal Discovery Algorithms: Application to Alzheimer's Pathophysiology. *Scientific Reports (Nature Publisher Group)* 10, 1 (2020).

[103] Tianyin Xu, Long Jin, Xuepeng Fan, Yuanyuan Zhou, Shankar Pasupathy, and Rukma Talwadker. 2015. Hey, you have given me too many knobs!: understanding and dealing with over-designed configuration in system software. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*.

[104] Tianyin Xu, Xinxin Jin, Peng Huang, Yuanyuan Zhou, Shan Lu, Long Jin, and Shankar Pasupathy. 2016. Early detection of configuration errors to reduce failure damage. USENIX Association, 619–634.

[105] Nezih Yigitbasi, Theodore L Willke, Guangdeng Liao, and Dick Epema. 2013. Towards machine learning-based auto-tuning of mapreduce. In *2013 IEEE 21st International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems*. IEEE, 11–20.

[106] Jiaqi Zhang, Lakshminarayanan Renganarayana, Xiaolan Zhang, Niyu Ge, Vasanth Bala, Tianyin Xu, and Yuanyuan Zhou. 2014. Encore: Exploiting system environment and correlation information for misconfiguration detection. In *Proceedings of the 19th international conference on Architectural support for programming languages and operating systems*. 687–700.

[107] Sai Zhang and Michael D Ernst. 2013. Automated diagnosis of software configuration errors. In *2013 35th International Conference on Software Engineering*.

[108] Yuanliang Zhang, Haochen He, Owolabi Legunsen, Shanshan Li, Wei Dong, and Tianyin Xu. 2021. An Evolutionary Study of Configuration Design and Implementation in Cloud Systems. In *Proceedings of International Conference on Software Engineering (ICSE'21)*.