Transfer Learning for Performance Analysis of Machine Learning Systems

Pooyan Jamshidi University of South Carolina









built Today's most popular systems are configurable

102		🚺 🛞 — 💷 rqt_reconfigureParam	jure_Param - rqt	
103	drpc.port: 3772	Dynamic Reconfigure		
104	drpc.worker.threads: 64		-	
105	drpc.max_buffer_size: 1048576	Filter key:		
106	drpc.queue.size: 128		1	
107	drpc.invocations.port: 3773	Collapse all Expand all	acc_lim_x	
108	drpc.invocations.threads: 64		acc_lim_y	
109	drpc.request.timeout.secs: 600	* move_base	acc lim theta	
110	drpc.childopts: "-Xmx768m"		max vel x	
111	drpc.http.port: 3774	▶ local costmap	IIIdx_vet_x	
112	drpc.https.port: -1		min_vel_x	
113	drpc.https.keystore.password: ""		max_vel_thet	
114	drpc.https.keystore.type: "JKS"		min vel theta	
115	drpc.http.creds.plugin: org.apache.storm.security.auth.DefaultHttpCredentialsPlugi			
116	drpc.authorizer.acl.filename: "drpc-auth-acl.yaml"	K (1)	mm_m_place_	
117	drpc.authorizer.acl.strict: false		sim_time	
118			sim_granulari	
119	transactional.zookeeper.root: "/transactional"		angular sim o	
120	transactional.zookeeper.servers: null		adiat acala	
121	transactional.zookeeper.port: null		pdist_scale	
122		1 E	gdist_scale	
123	## blobstore configs		occdist_scale	
124	supervisor.blobstore.class: "org.apache.storm.blobstore.NimbusBlobStore"		oscillation re-	
125	supervisor.blobstore.download.thread.count: 5		oscillation_re	
, 126	supervisor.blobstore.download.max_retries: 3		escape_reset	
127	supervisor.localizer.cache.target.size.mb: 10240		escape_reset	
128	supervisor.localizer.cleanup.interval.ms: 600000		wy camples	
179		8		





	/move_base/Trajectory	PlannerROS	3
	0.0 -	20.0	1.5
	0.0	20.0	0.0
eta	0.0 -	20.0	1.2
	0.0	20.0	0.15
	0.0	20.0	0.1
heta	0.0	20.0	1.0
neta	-20.0	0.0	-1.0
ice_vel_theta	0.0	20.0	0.4
	0.0 -	10.0	1.0
larity	0.0	5.0	0.05
m_granularity	0.0	1.57079632679	0.1
2	0.0 -	5.0	0.5
2	0.0	5.0	0.8
ale	0.0	5.0	0.05
_reset_dist	0.0	5.0	0.05
set_dist	0.0	5.0	0.1
set_theta	0.0	5.0	1.57079632679
c	1 -	300	14



. . .





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

drpc.http.creds.plugin: org.apache.storm.security.auth.DefaultHttpCredentialsPlugi



Empirical observations confirm that systems are becoming increasingly configurable



[Tianyin Xu, et al., "Too Many Knobs...", FSE'15]

Empirical observations confirm that systems are becoming increasingly configurable



Configurations determine the performance behavior

```
void Parrot setenv(. . . name, . . . value){
#ifdef PARROT HAS SETENV
 my_setenv(name, value, 1);
#else
  int name len=strlen(name);
  int val_len=strlen(value);
  char* envs=glob env;
  if(envs==NULL){
    return;
  strcpy(envs,name);
  strcpy(envs+name_len,"=");
  strcpy(envs+name_len + 1,value);
  putenv(envs);
```

#endif



How do we understand performance behavior of real-world highly-configurable systems that scale well...

... and enable developers/users to reason about qualities (performance, energy) and to make tradeoff?



SocialSensor

- Identifying trending topics
- Identifying user defined topics
- Social media search

SocialSensor



Internet

Crawled items









Challenges



Internet

Crawled items

10X





100X





using more resources?

How can we gain a better performance without

Let's try out different system configurations!

Opportunity: Data processing engines in the pipeline were all configurable





STORM

> 100



 2^{300}



Default configuration was bad, so was the expert'







The default configuration is typically bad and the optimal configuration is noticeably better than median **Default Configuration** better 5000 r



- 2X-10X faster than worst
- Noticeably faster than median



What did happen at the end?

- Achieved the objectives (100X user, same experience)
- Saved money by reducing cloud resources up to 20%
- Our tool was able to identify configurations that was consistently better than expert recommendation





A typical approach for understanding the performance behavior is sensitivity analysis $O_1 \times O_2 \times \cdots \times O_{19} \times O_{20}$ $C_{1} \quad 0 \times 0 \times \dots \times 0 \times 1$ $C_{2} \quad 0 \times 0 \times \dots \times 1 \times 0$ $C_{3} \quad 0 \times 0 \times \dots \times 1 \times 1$ $y_1 = f(c_1)$ $y_2 = f(c_2)$ $y_3 = f(c_3)$ $f \sim f(\cdot)$ Learn Training/Sample $\begin{array}{c} 1 \times 1 \times \cdots \times 1 \times 0 \\ c_n 1 \times 1 \times \cdots \times 1 \times 1 \end{array} \quad y_n = f(c_n) \end{array}$

Performance model could be in any appropriate form of black-box models

 $O_1 \times O_2 \times \cdots \times O_{19} \times O_{20}$ $y_1 = f(c_1)$ $C_1 0 \times 0 \times \cdots \times 0 \times 1$ $C_2 0 \times 0 \times \cdots \times 1 \times 0$ $y_2 = f(c_2)$ $C_3 0 \times 0 \times \cdots \times 1 \times 1$ $y_3 = f(c_3)$ Training/Sample $f(\cdot)$ $\begin{array}{c} 1 \times 1 \times \cdots \times 1 \times 0 \\ c_n \ 1 \times 1 \times \cdots \times 1 \times 1 \end{array} \quad y_n = f(c_n) \end{array}$





Evaluating a performance model





A performance model contain useful information about influential options and interactions

$f: \mathbb{C} \to \mathbb{R}$ $f(\cdot) = 1.2 + 3o_1 + 5o_3 + 0.9o_7 + 0.8o_3o_7 + 4o_1o_3o_7$

Performance model can then be used to reason about qualities

```
void Parrot setenv(. . . name, . . . value){
#ifdef PARROT HAS SETENV
 my_setenv(name, value, 1);
#else
  int name_len=strlen(name);
  int val_len=strlen(value);
  char* envs=glob env;
  if(envs==NULL){
    return;
  strcpy(envs,name);
  strcpy(envs+name_len,"=");
  strcpy(envs+name_len + 1,value);
  putenv(envs);
#endif
```



Insight: Performance measurements of the real system is "similar" to the ones from the simulators



We developed methods to make learning cheaper via transfer learning

Goal: Gain strength by transferring information across environments



What is transfer learning?





What is transfer learning?



Transfer learning is a machine learning technique, where knowledge gain during training in one type of problem is used to train in other similar type of problem







What is the advantage of transfer learning?

- During learning you may need thousands of rotten and fresh potato and hours of training to learn.
- But now using the same knowledge of rotten features you can identify rotten tomato with **less samples and training time**.
- You may have learned during daytime with enough light and exposure; but your present tomato identification job is at night.
- You may have learned sitting very close, just beside the box of potato; but now for tomato identification you are in the other side of the glass.

A simple transfer learning via model shift

Machine twice as fast







[Pavel Valov, et al. "Transferring performance prediction models...", ICPE'17]







[P. Jamshidi, et al., "Transfer learning for improving model predictions", SEAMS'17]









Gaussian Processes enables reasoning about performance

Step 1: Fit GP to the data seen so far

Step 2: Explore the model for regions of most variance

Step 3: Sample that region

Step 4: Repeat


The intuition behind our transfer learning approach

Intuition: Observations on the source(s) can affect predictions on the target

Example: Learning the chess game make learning the Go game a lot easier!



CoBot experiment: DARPA BRASS































Result: CoBot experiment



2500

Results: Other configurable systems



Details: [SEAMS '17]

Transfer Learning for Improving Model Predictions in Highly Configurable Software

Pooyan Jamshidi, Miguel Velez, Christian KästnerNorbert SiegmundCarnegie Mellon University, USABauhaus-University Weimar, Germany{pjamshid,mvelezce,kaestner}@cs.cmu.edunorbert.siegmund@uni-weimar.de

Abstract—Modern software systems are built to be used in dynamic environments using configuration capabilities to adapt to changes and external uncertainties. In a self-adaptation context, we are often interested in reasoning about the performance of the systems under different configurations. Usually, we learn a black-box model based on real measurements to predict the performance of the system given a specific configuration. However, as modern systems become more complex, there are many configuration parameters that may interact and we end up learning an exponentially large configuration space. Naturally, this does not scale when relying on real measurements in the actual changing environment. We propose a different solution: Prasad Kawthekar Stanford University, USA pkawthek@stanford.edu





Looking further: When transfer learning goes wrong Non-transfer-learning

Insight: Predictions become more accurate when the source is more related to the target.







Key question: Can we develop a theory to explain when transfer learning works?



Q1: How source and target are "related"?

Q2: What characteristics are preserved?

Q3: What are the actionable insights?

Mathematical Framework of configuration optimization

sinle casla ьg a <u>6</u> 6 a; $sin^2 d + cos^2 d = 1;$ $\frac{sin d}{cos d} = \frac{1}{2} \frac{sin d}{sin d}$ · d = <u>T</u> 180 d d \$ 180°=TL; sind cscd=1; 4=as Cosd = cbg d **△** >0 $tg \varphi = \pm a^2 (\frac{3}{2})^{\frac{3}{2}};$

We define configuration optimization as a multi-objective optimization problem with unknown feasibility constraints

$x^* = argmin_{c \in \mathscr{C}} f(x)$

 $\phi_i(\mathbf{c}) \leq b_i, i = 1, ..., q$

Configuration space may include real, ordinal, and categorical configuration options

- The variables defining the configuration space can be ordinal (real, integer), and categorical.
- Ordinal parameters have a domain of a finite set of values which are either integer and/or real values.
- Ordinal values must have an ordering by the less-than operator.
- Categorical parameters (Boolean) also have domains of a finite set of values but have no ordering requirement.

We assume the derivative of the optimization function is not available

- We assume that the **derivative** of f is **not available**.
- And that bounds, such as Lipschitz constants, for the derivative of f is also unavailable.
- Evaluating feasibility is often in the same order of expense as evaluating the objective function f.
- As for the objective function, **no particular assumptions** are made on the constraint functions.

We induce partial ordering between configurations in the configuration space



$\mathbb{R}^d: y \prec y' \iff \forall i \in [d] y_i \leq y'_i \& \exists j y_i < y'_i$

 $\mathscr{C}: \mathbf{c} \prec \mathbf{c}' \iff f(\mathbf{c}) \prec f(\mathbf{c}')$

Based on the induced ordering of configurations we can define the Pareto optimal set of configurations

 $\Gamma = \{\mathbf{c} \in \mathscr{C} : \nexists \mathbf{c}' < \mathbf{c}\}$

Applying these constraints gives the constrained Pareto-optimal set

$\Gamma = \{ \mathbf{c} \in \mathscr{C} : \nexists \mathbf{c}' < \mathbf{c} \And \phi_i(\mathbf{c}') \leq b_i \}$

The multi-objective function maps each point in the 3-dimensional configuration space on the left to the optimization space on the right



We hypothesized that we can exploit similarities across environments to learn "cheaper" performance models

Source Environment (Execution time of Program X) $O_1 \times O_2 \times \cdots \times O_{19} \times O_{20}$

$$C_{1} \quad 0 \times 0 \times \dots \times 0 \times 1 \quad y_{s1} = f_{s}(c_{1})$$

$$C_{2} \quad 0 \times 0 \times \dots \times 1 \times 0 \quad y_{s2} = f_{s}(c_{2})$$

$$C_{3} \quad 0 \times 0 \times \dots \times 1 \times 1 \quad y_{s3} = f_{s}(c_{3})$$

$$\dots$$

$$1 \times 1 \times \dots \times 1 \times 0 \quad y_{sn} = f_{s}(c_{n})$$

[P. Jamshidi, et al., "Transfer learning for performance modeling of configurable systems....", ASE'17]

Target Environment (Execution time of Program Y) $O_1 \times O_2 \times \cdots \times O_{19} \times O_{20}$



Similarity

Our empirical study: We looked at different highlyconfigurable systems to gain insights

$(A \vee B) \wedge (\neg A \vee \neg B \vee \neg C) \wedge (\neg A \vee B \vee C)$





SPEAR (SAT Solver) Analysis time

14 options

X264 (video encoder) **SQLite (DB engine) Encoding time** Query time 14 options

16 options

[P. Jamshidi, et al., "Transfer learning for performance modeling of configurable systems....", ASE'17]



SaC (Compiler) **Execution time** 50 options

Linear shift happens only in limited environmental changes

Soft	Environmental change	Sever
SPEAR	NUC/2 -> NUC/4	Small
	Amazon_nano -> NUC	Large
	Hardware/workload/version	V Lar
	and the second state and the second state second state second state second state and the second state second state second state second states and the second states	
v964	Version	Large
x264	Version Workload	Large Mediu
x264	Version Workload write-seq -> write-batch	Large Mediu Small

Implication: Simple transfer learning is limited to hardware changes in practice



Target

Source

Influential options and interactions are preserved across environments

Soft	Environmental change	Severity
x264	Version	Large
	Hardware/workload/ver	V Large
SQLite	write-seq -> write-batch	V Large
	read-rand -> read-seq	Medium
SaC	Workload	V Large

Implication: Avoid wasting budget on non-informative part of configuration space and focusing where it matters.



We only need to explore part of the space: = 0.00000000058

Transfer learning across environment

Source (Execution time of Program X) $O_1 \times O_2 \times \cdots \times O_{19} \times O_{20}$

$$C_{1} \quad 0 \times 0 \times \dots \times 0 \times 1 \qquad y_{s1} = f_{s}(c_{1})$$

$$C_{2} \quad 0 \times 0 \times \dots \times 1 \times 0 \qquad y_{s2} = f_{s}(c_{2})$$

$$C_{3} \quad 0 \times 0 \times \dots \times 1 \times 1 \qquad y_{s3} = f_{s}(c_{3}) \qquad \text{perf}$$

$$\dots$$

$$1 \times 1 \times \dots \times 1 \times 0 \qquad y_{sn} = f_{s}(c_{n})$$



 $\hat{f}_{S} \sim f_{S}(\cdot)$

and interactions degree between options are not high

$\mathbb{C} = O_1 \times O_2 \times O_3 \times O_4 \times O_5 \times O_6 \times O_7 \times O_8 \times O_9 \times O_{10}$ $\hat{f}_{s}(\cdot) = 1.2 + 3o_{1} + 5o_{3} + 0.9o_{7} + 0.8o_{3}o_{7} + 4o_{1}o_{3}o_{7}$

Observation 1: Not all options and interactions are influential



Observation 2: Influential options and interactions are preserved across environments

$$\hat{f}_s(\cdot) = 1.2 + 3o_1 + 5o_1$$

$$\hat{f}_t(\cdot) = 10.4 - 2.1o_1 + 1.2o_3$$



Transfer Learning for Performance Modeling of Configurable Systems: An Exploratory Analysis

Pooyan Jamshidi Carnegie Mellon University, USA

Norbert Siegmund Bauhaus-University Weimar, Germany

Abstract—Modern software systems provide many configuration options which significantly influence their non-functional properties. To understand and predict the effect of configuration options, several sampling and learning strategies have been proposed, albeit often with significant cost to cover the highly dimensional configuration space. Recently, transfer learning has been applied to reduce the effort of constructing performance models by transferring knowledge about performance behavior across environments. While this line of research is promising to learn more accurate models at a lower cost, it is unclear why and when transfer learning works for performance modeling. To shed light on when it is beneficial to apply transfer learning, we conducted an empirical study on four popular software systems, varying software configurations and environmental conditions, such as hardware, workload, and software versions, to identify the key knowledge pieces that can be exploited for transfer learning. Our results show that in small environmental changes (e.g., homogeneous workload change), by applying a linear transformation to the performance model, we can understand the performance behavior of the target environment, while for severe environmental changes (e.g., drastic workload change) we

Details: [ASE '17]

Miguel Velez, Christian Kästner Akshay Patel, Yuvraj Agarwal Carnegie Mellon University, USA



reliable, and less costly model for the target environment.



How to sample the configuration space to learn a "better" performance behavior? How to select the most informative configurations?



The similarity across environment is a rich source of knowledge for exploration of the configuration space

 $O_1 \times O_2 \times \cdots \times O_{19} \times O_{20}$ $\begin{array}{c} c_1 \\ 0 \times 0 \times \cdots \times 0 \times 1 \\ c_2 \\ 0 \times 0 \times \cdots \times 1 \times 0 \\ c_3 \\ 0 \times 0 \times \cdots \times 1 \times 1 \end{array}$ $1 \times 1 \times \cdots \times 1 \times 0$ $C_n 1 \times 1 \times \cdots \times 1 \times 1$

- We therefore end up blindly explore the configuration space
- That is essentially the key reason why "most" work in this area consider random sampling.

When we treat the system as black boxes, we cannot typically distinguish between different configurations



 $f_s(\cdot) = 1.2 + 3o_1 + 5o_3 + 0.9o_7 + 0.8o_3o_7 + 4o_1o_3o_7$ $\times 1 \times 0 \times 0 \times 0 \times 1 \times 0 \times 1 \times 0 \qquad \hat{f}_s(c_3) = 14.9$



Without knowing this knowledge, many blind/ random samples may not provide any additional information about performance of the system





In higher dimensional spaces, the blind samples even become less informative/effective



Learning to Sample (L2S)

Extracting the knowledge about influential options and interactions: Step-wise linear regression

Source (Execution time of Program X) $O_1 \times O_2 \times \cdots \times O_{19} \times O_{20}$

 $C_{1} \quad 0 \times 0 \times \dots \times 0 \times 1 \quad y_{s1} = f_{s}(c_{1})$ $C_{2} \quad 0 \times 0 \times \dots \times 1 \times 0 \quad y_{s2} = f_{s}(c_{2})$ $C_{3} \quad 0 \times 0 \times \dots \times 1 \times 1 \quad y_{s3} = f_{s}(c_{3})$ $\times 1 \times \cdots \times 1$ $C_n \quad 1 \times 1 \times \dots \times 1 \times 1 \quad y_{sn} = f_s(c_n)$

- 1. Fit an **initial model**
- 2. Forward selection: Add terms iteratively
- Learn performance model $\hat{f}_{s} \sim f_{s}(\cdot)$
- 3. Backward elimination: **Removes terms iteratively**
- 4. Terminate: When neither (2) or (3) improve the model

Build a performance distribution using kernel density estimation using the source data


L2S extracts the knowledge about influential options and interactions via performance models

$\hat{f}_{s}(\cdot) = 1.2 + 3o_1 + 5o_3 + 0.9o_7 + 0.8o_3o_7 + 4o_1o_3o_7$





 $\hat{f}_{s}(\cdot) = 1.2 + 3o_{1} + 5o_{3} + 0.9o_{7} + 0.8o_{3}o_{7} + 4o_{1}o_{3}o_{7}$ $\hat{f}_t(\cdot) = 10.4 - 2.1o_1 + 1.2o_3 + 2.2o_7 + 0.1o_1o_3 - 2.1o_3o_7 + 14o_1o_3o_7$ $\hat{f}_t(c_1) = 10.4$ $\hat{f}_t(c_2) = 8.1$ $\hat{f}_t(c_3) = 11.6$ $\hat{f}_t(c_4) = 12.6$ $\hat{f}_t(c_5) = 11.7$ $C_6 1 \times 0 \times 1 \times 0 \times 0 \times 0 \times 1 \times 0 \times 0 \times 0 = 23.7$



Exploration vs Exploitation

We also explore the configuration space using pseudo-random sampling to detect missing interactions



For capturing options and interactions that only appears in the target, L2S relies on exploration (random sampling)



L2S transfers knowledge about the structure of performance models from the source to guide the sampling in the target environment





















Evaluation: also exists

Other transfer learning approaches

"Model shift" builds a model in the source and uses the shifted model "directly" for predicting the target

Machine twice as fast







[Pavel Valov, et al. "Transferring performance prediction models...", ICPE'17]





"Data reuse" combines the data collected in the source with the ones in the target in a "multi-task" setting to predict the target



[P. Jamshidi, et al., "Transfer learning for improving model predictions", SEAMS'17]



"Data reuse" with guided sampling





Simulator (Gazebo)







Evaluation: Learning performance behavior of Machine Learning Systems

ML system: <u>https://pooyanjamshidi.github.io/mls</u>

Configurations of deep neural networks affect accuracy and energy consumption



Deep neural network as a highly configurable system





DNN measurements are costly Each sample cost ~1h

4000 * 1h ~= 6 months



L2S enables learning a more accurate model with less samples exploiting the knowledge from the source



Convolutional Neural Network

L2S may also help data-reuse approach to learn faster



XGBoost

Evaluation: Learning performance behavior of **Big Data Systems**

Some environments the similarities across environments may be too low and this results in "negative transfer"



Why performance models using L2S sample are more accurate?



The samples generated by L2S contains more information... "entropy <-> information gain"



Limitations

- Limited number of systems and environmental changes
 - Synthetic models
 - <u>https://github.com/pooyanjamshidi/GenPerf</u>
- Binary options
 - Non-binary options -> binary
- Negative transfer

Learning to Sample: Exploiting Similarities across Environments to Learn Performance Models for Configurable Systems

Pooyan Jamshidi University of South Carolina USA

ABSTRACT

Most software systems provide options that allow users to tailor the system in terms of functionality and qualities. The increased flexibility raises challenges for understanding the configuration space and the effects of options and their interactions on performance and other non-functional properties. To identify how options and interactions affect the performance of a system, several sampling and learning strategies have been recently proposed. However, existing approaches usually assume a fixed environment (hardware, workload, software release) such that learning has to be repeated once the environment changes. Repeating learning and measurement for each environment is expensive and often practically infeasible. Instead, we pursue a strategy that transfers knowledge across environments but sidesteps heavyweight and expensive transfer-

Details: [FSE '18]





What will the software systems of the future look like?



VISION Software 2.0

Increasingly customized and configurable

namic Reconfigure					D? - 0		
er key:		/move_base/TrajectoryPlar	nnerROS		x		11
acc acc acc acc acc acc acc acc	cc_lim_x cc_lim_y cc_lim_theta ax_vel_x in_vel_x ax_vel_theta in_vel_theta in_in_place_vel_theta m_time	0.0	20.0 20.0 20.0 20.0 20.0 20.0 20.0 20.0	1.5 0.0 1.2 0.15 0.1 1.0 -1.0 0.4 1.0		BN-Reby-Con	
sir an pd gd oc os es es	m_granularity ngular_sim_granularity dist_scale dist_scale ccdist_scale scillation_reset_dist scape_reset_dist scape_reset_theta	0.0	5.0 1.57079632679 5.0 5.0 5.0 5.0 5.0 5.0 5.0 5.0	0.05 0.1 0.5 0.8 0.05 0.05 0.1 1.570796326	79		

Increasingly competing objectives

Energy Model size



- Inference speed
 - **Training speed**
 - Accuracy

Deep neural network as a highly configurable system





We found many configuration with the same accuracy while having drastically different energy demand





Deep architecture design level variations

NAME	DESCRIPTION	ΤΥΡΕ	MIN	MAX
batch_size	SGD parameter	int	8	32
conv_1_filter_size	Architecture parameter	int	2	10
conv_1_num_filters	Architecture parameter	int	32	256
conv_2_filter_size	Architecture parameter	int	2	10
conv_2_num_filters	Architecture parameter	int	32	256
conv_3_filter_size	Architecture parameter	int	2	10
conv_3_num_filters	Architecture parameter	int	32	256
log_beta_1	Adam SGD parameter	real	-4.6	-0.7
log_beta_2	Adam SGD parameter	real	-13.8	-0.7
log_decay	Adam SGD parameter	real	-23	-2.3
log_epsilon	Adam SGD parameter	real	-23	-13.8
log_lr	Adam SGD parameter	real	-23	0

Deep architecture deployment variations















Deep architecture hardware-level variations





NVIDIA Xavier

Exploring the design space of deep networks





Optimal Architecture (Yesterday)

New Fraud Pattern



Optimal Architecture (Today)









Configuration options and interactions influence performance of DNNs



Insight: Learn a model on a cheaper workload to explore the expensive workload faster









Configuration errors are prevalent










Configuration errors are common





Configuration complexity and dependencies between options is a major source of configuration errors



rpc.port: 3772 drpc.worker.threads: 64 drpc.queue.size: 128







Operational context II



Operational context III







Configuration complexity and dependencies between options is a major source of configuration errors

drpc.port: 3772 drpc.worker.threads: 64 drpc.max_buffer_size: 1048576 drpc.queue.size: 128 drpc.invocations.port: 3773 drpc.invocations.threads: 64 drpc.request.timeout.secs: 600

Configuration Options



Apache Hadoop Architecture

Configurations are software too

We can find the repair patches faster with a lighter workload

- [localization]: Using transfer learning to derive the most likely configurations that manifest the bugs
- [repair]: Automatically prepare patches to fix the configuration bug

built Many systems are now configurable

	😸 — 🗇 rqt_reconfigure_Param - rqt	
	Dynamic Reconfigure	
	Filter key:	
	<u>C</u> ollapse all <u>E</u> xpan	d all acc_lim_x
▼ move_base TrajectoryPlanne ▶ global_costmap ▶ local_costmap	▼ move_base	acc_um_y acc_lim_theta
	 ▶ global_costmap ▶ local_costmap 	max_vel_x





Trained Network

Given the ever growing configurable systems, how can we enable learning practical models that scale well and provide reliable predictions for exploring the configuration space?







Challenges Crawled



Our empirical study: We looked at different highlyconfigurable systems to gain insights





Exploring the design space of deep networks

