

Synergizing Variability Modeling with Machine Learning: A Journey of Possibilities

Dr. Maxime Cordy

Interdisciplinary Center on Security and Trust (SnT), University of Luxembourg

It all started with an email...



Leopoldo Motta Teixeira

Invitation as Keynote Speaker at VAMOS 2024, Bern, Feb 7-9

To: Maxime CORDY, Cc: Marianne Huchard, Timo Kehrer

5 June 2023, 16:30

[Details](#)



Dear Maxime,

On behalf of the committees of the 18th International Working Conference on Variability Modelling of Software-Intensive Systems (VaMoS 2024), we are very pleased to invite you to deliver a keynote speech at the event.

We are familiar with your previous work on variability analysis and verification, and would like to know whether you could provide a perspective on the interaction of variability, testing and machine learning, given that you have been working around these topics in the recent years, and they directly relate to the VaMoS topics.

VaMoS 2024 will be held Feb 7-9, in Bern, Switzerland (<https://vamos2024.inf.unibe.ch/>).

If you accept our invitation, your keynote will be scheduled on either of the conference days (exact day to be confirmed). You are, of course, also invited to attend the other sessions during the week. All your travel expenses will be covered by VaMoS and you will get a free registration to the whole event.

Over the last years, VaMoS has already counted on renowned scientists or entrepreneurs as keynote speakers, such as Yves Bossu, Marsha Chechik, Mathieu Acher, Alfonso Pierantonio, Manuel Wimmer, Herwig Schreiner, Andrzej Wasowski, Nelly Bencomo, Norbert Siegmund and Matthias Galster.

We hope you will be interested and available for coming to Bern. It will be our great pleasure to have you as one of our keynote speakers.

Best regards,





Timo Kehrer, Marianne Huchard, Leopoldo Teixeira
VaMoS 2024 General chair and PC chairs

12 years ago...




Verification of variability-intensive systems

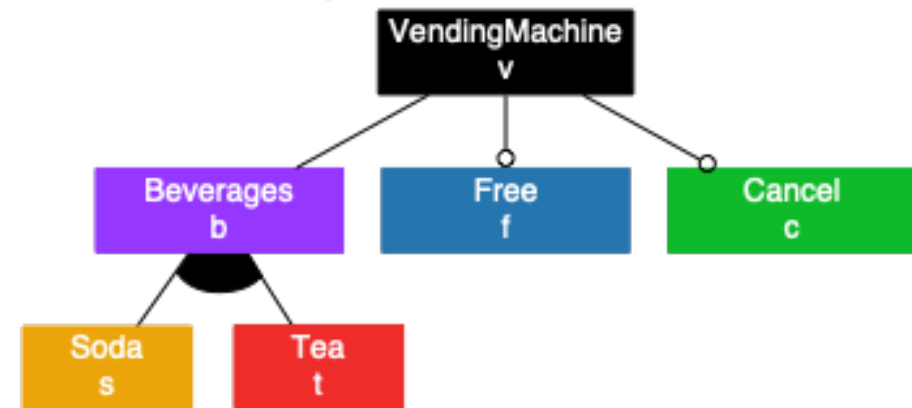
Instead of:



Write: **feature expressions**




+ companion feature model:



Uniform random sampling

Given a feature model
(a Boolean formula)...



... sample a valid variant (a solution to the formula) with uniform probability:

{v, b, s}
{v, b, s, f}
{v, b, s, c}
{v, b, s, f, c}

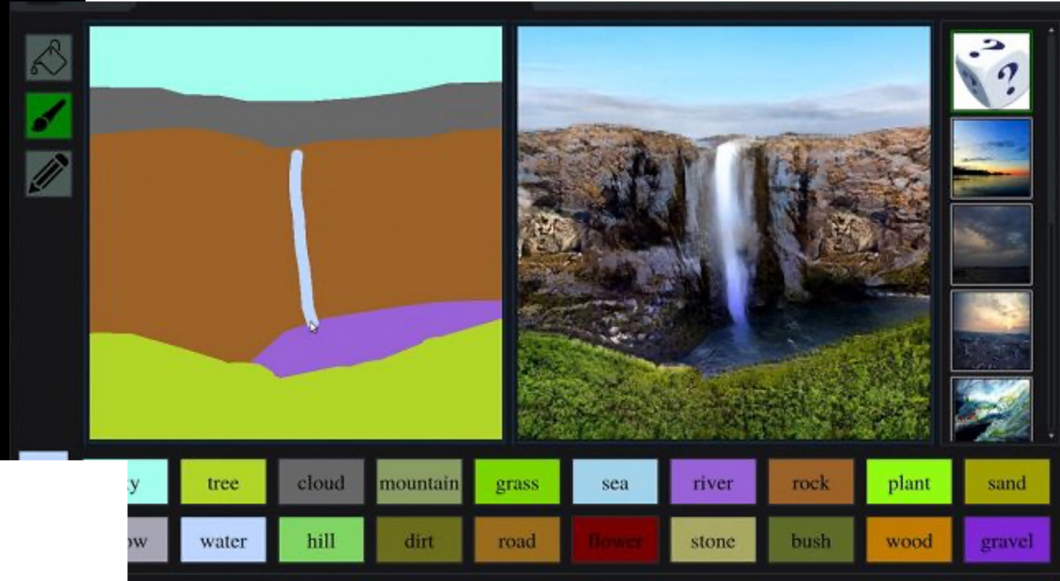
{v, b, t}
{v, b, t, f}
{v, b, t, c}
{v, b, t, f, c}

{v, b, s, t}
{v, b, s, t, f}
{v, b, s, t, c}
{v, b, s, t, f, c}

Deep learning: a new world of possibilities

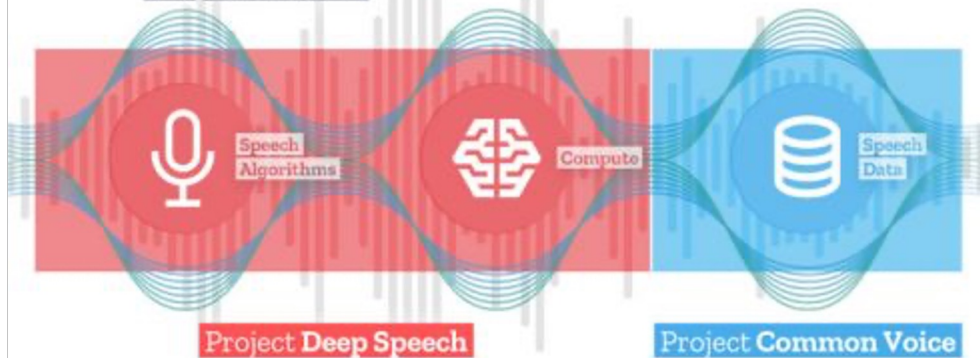
Deep Fake Nicolas Cage

Uploaded Mar 22 2018

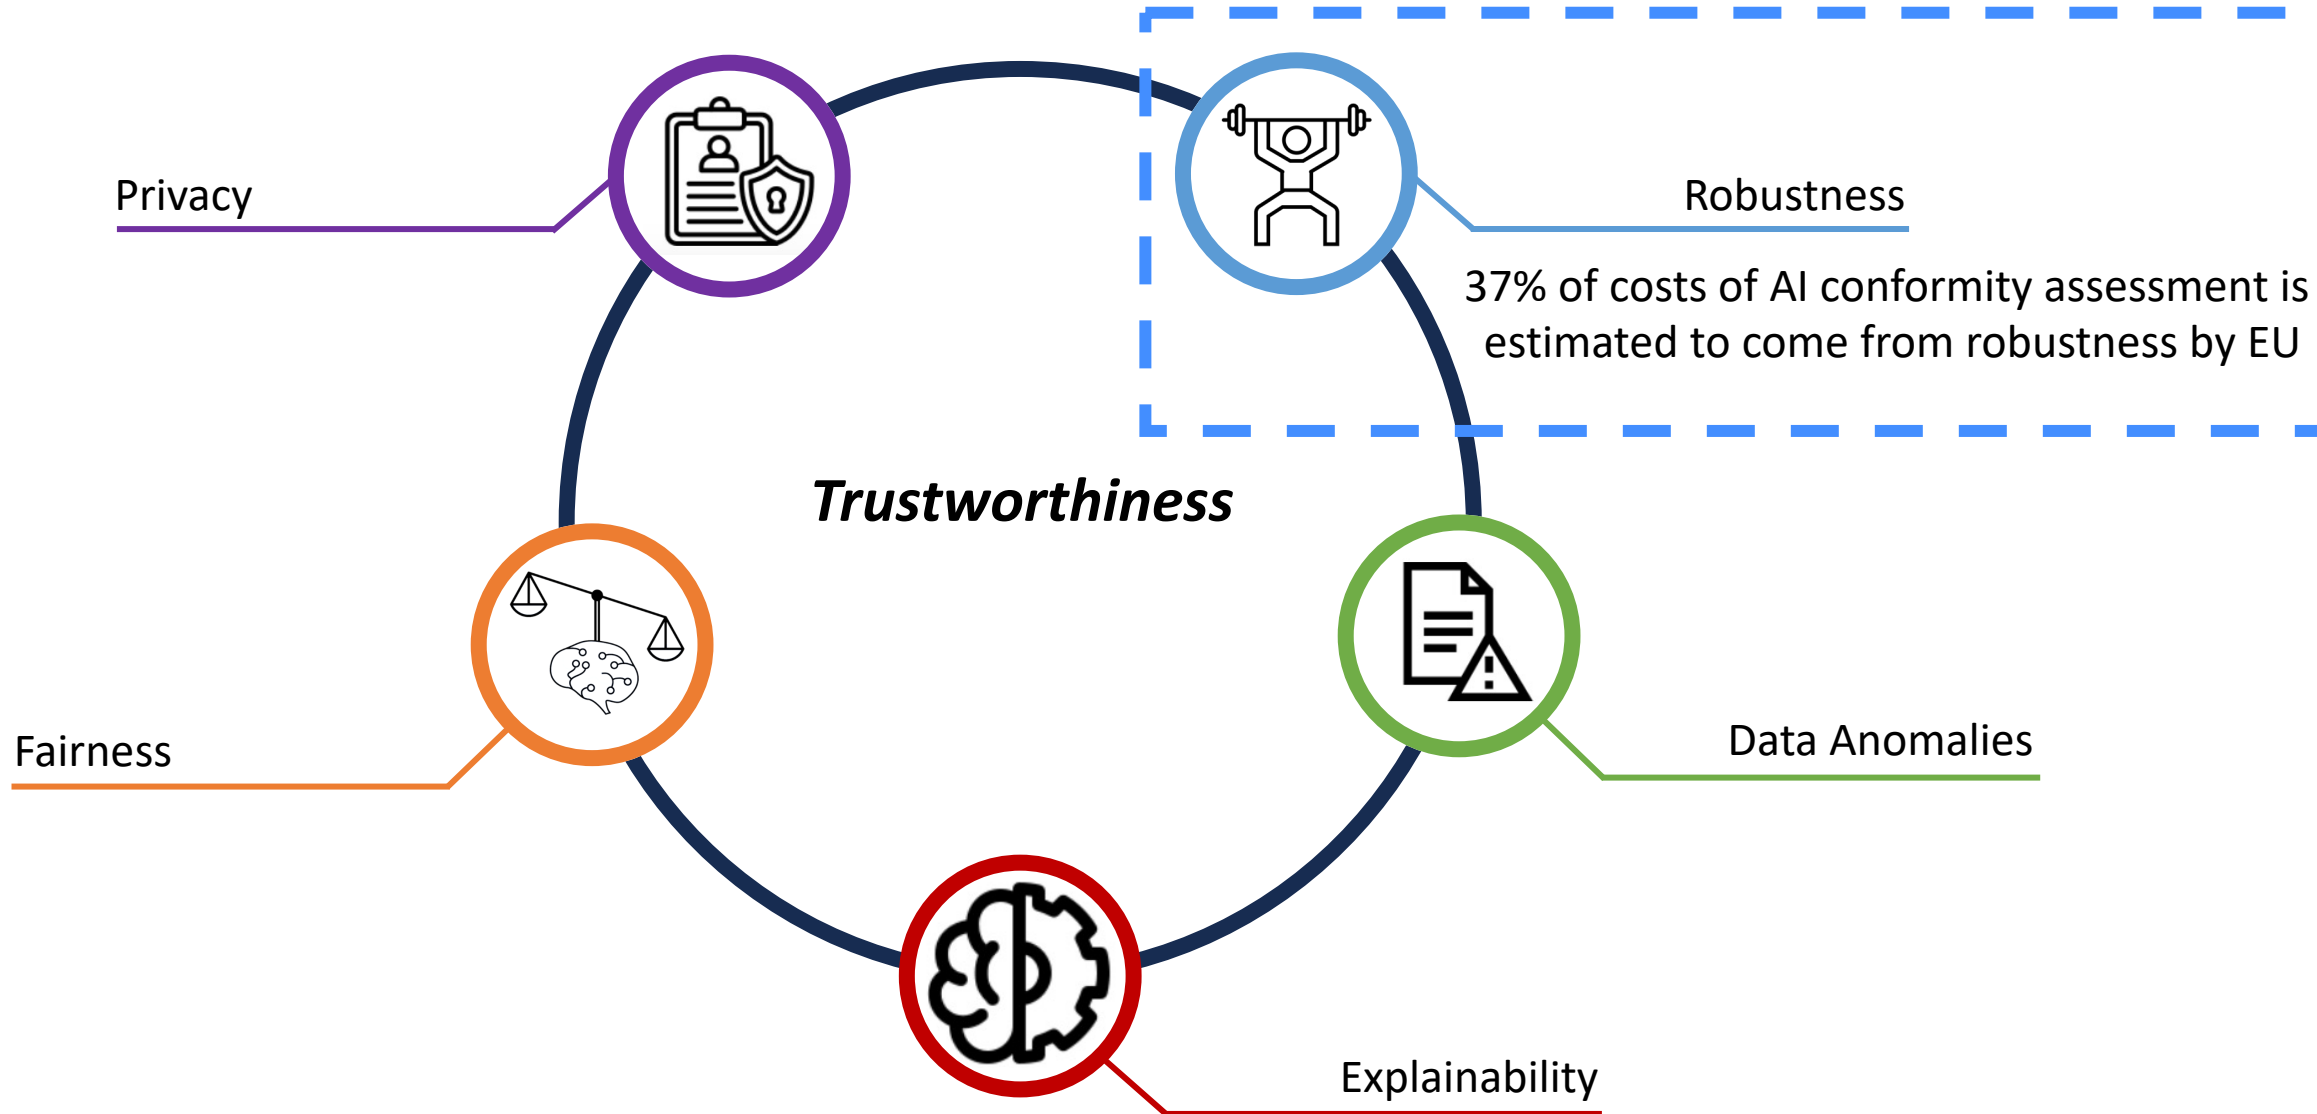


moz://a Speech Recognition

Ingredients



Different aspects of Trustworthy AI



Robustness

“the degree to which a model’s performance changes when confronted to data unseen during training”




“Natural Robustness”

Accuracy once running in production



Distribution Drift

“...an evolution of data that invalidates the data model. It happens when the statistical properties of the target variable, which the model is trying to predict, change over time in unforeseen ways.”



Security Threats

*Evasion/adversarial attacks
Inference attacks
Poisoning attacks
Model theft
Etc.*

Machine learning robustness (evasion attacks)

Original example



Small adversarial noise



Adversarial example



What humans still see

ML predicts:
"Panda"
(80% confidence)



What ML predicts: "Gibbon"
(99% confidence)

Gibbon

Machine learning robustness (evasion attacks)

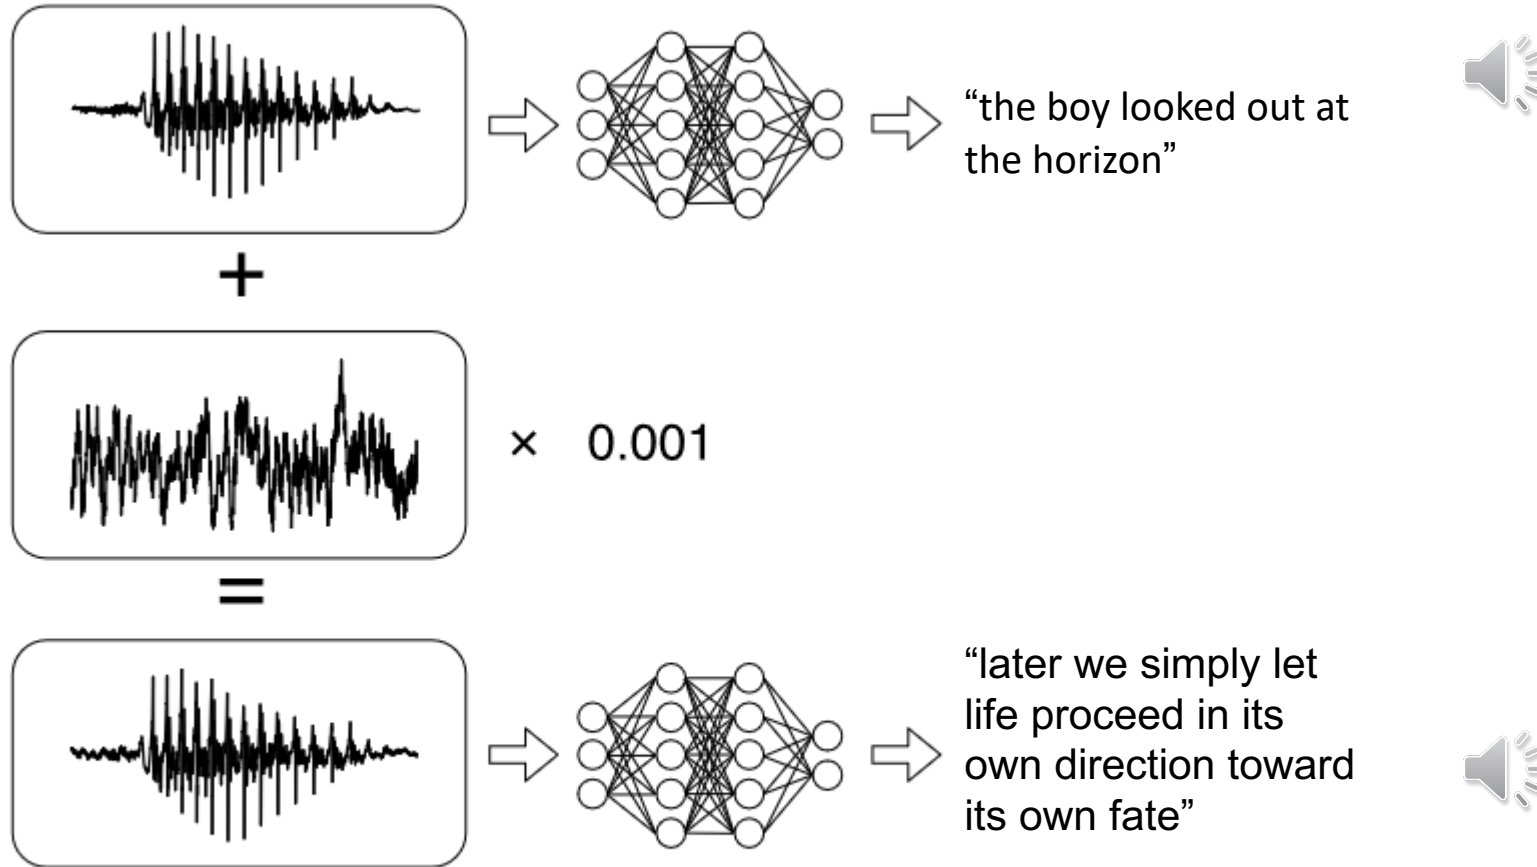



Figure 1. Illustration of our attack: given any waveform, adding a small perturbation makes the result transcribe as any desired target phrase.

Machine learning robustness (evasion attacks)



Crowd face recognition system



(a)

(b)

Journal of Information Security and Applications 60 (2021) 102874



Contents lists available at [ScienceDirect](#)

Journal of Information Security and Applications

journal homepage: www.elsevier.com/locate/jisa

Adversarial attacks by attaching noise markers on the face against deep face recognition

Gwangang Ryu^a, Hosung Park^b, Daeseon Choi^{c,*}

^a Department of Software Convergence, Graduate School of Soongsil University, Seoul, 07027, South Korea

^b Department of Cyber Security and Police, Busan University of Foreign Studies, Busan, 46234, South Korea

^c Department of Software, Soongsil University, Seoul, 07027, South Korea

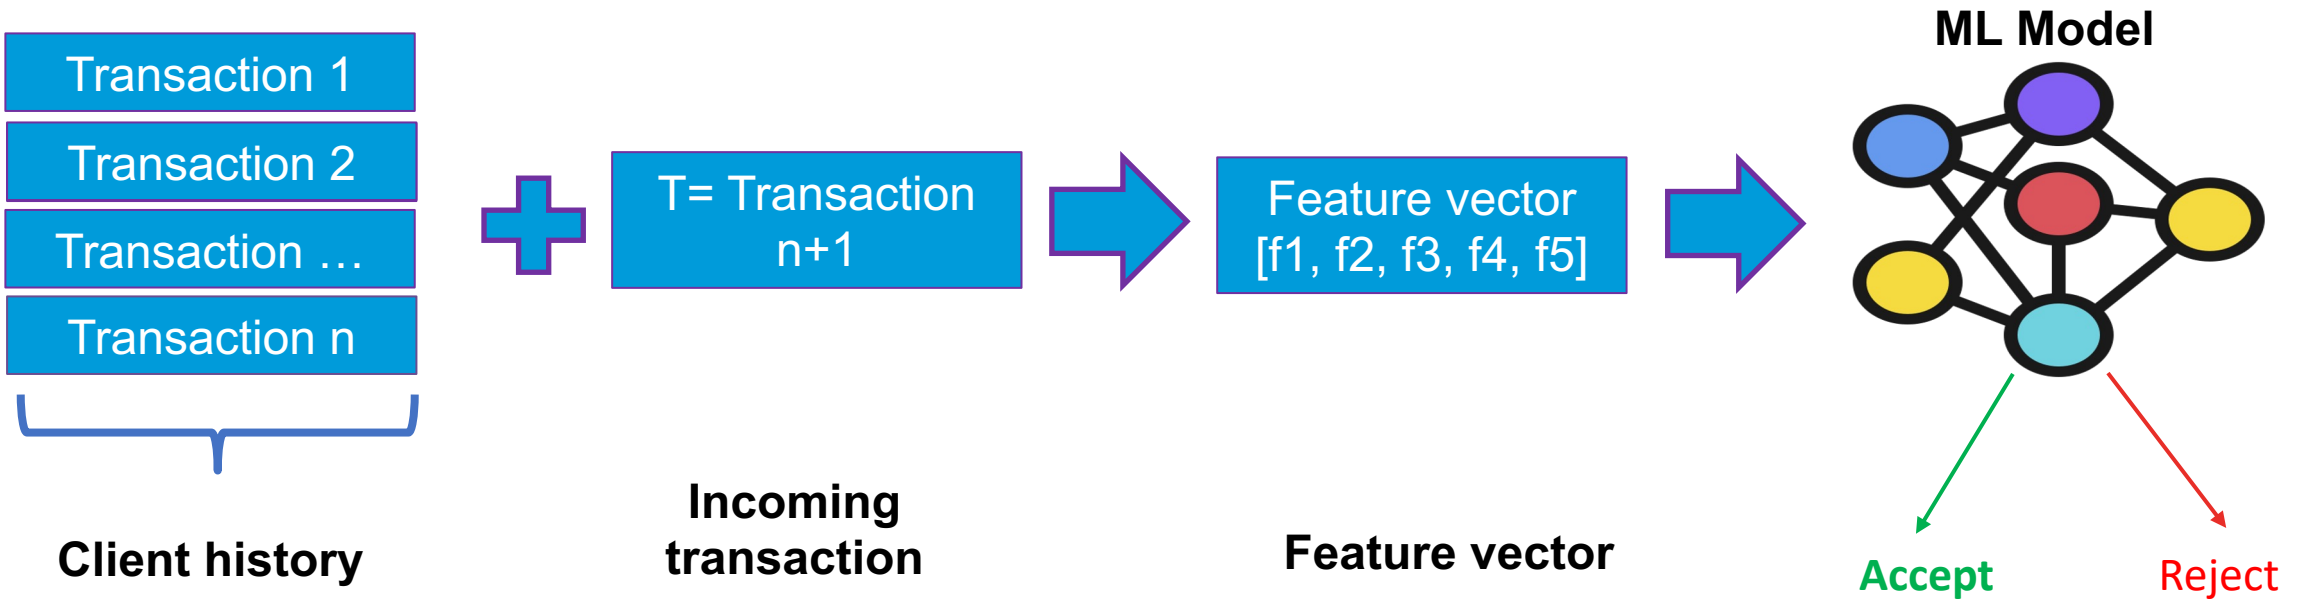
ARTICLE INFO

ABSTRACT



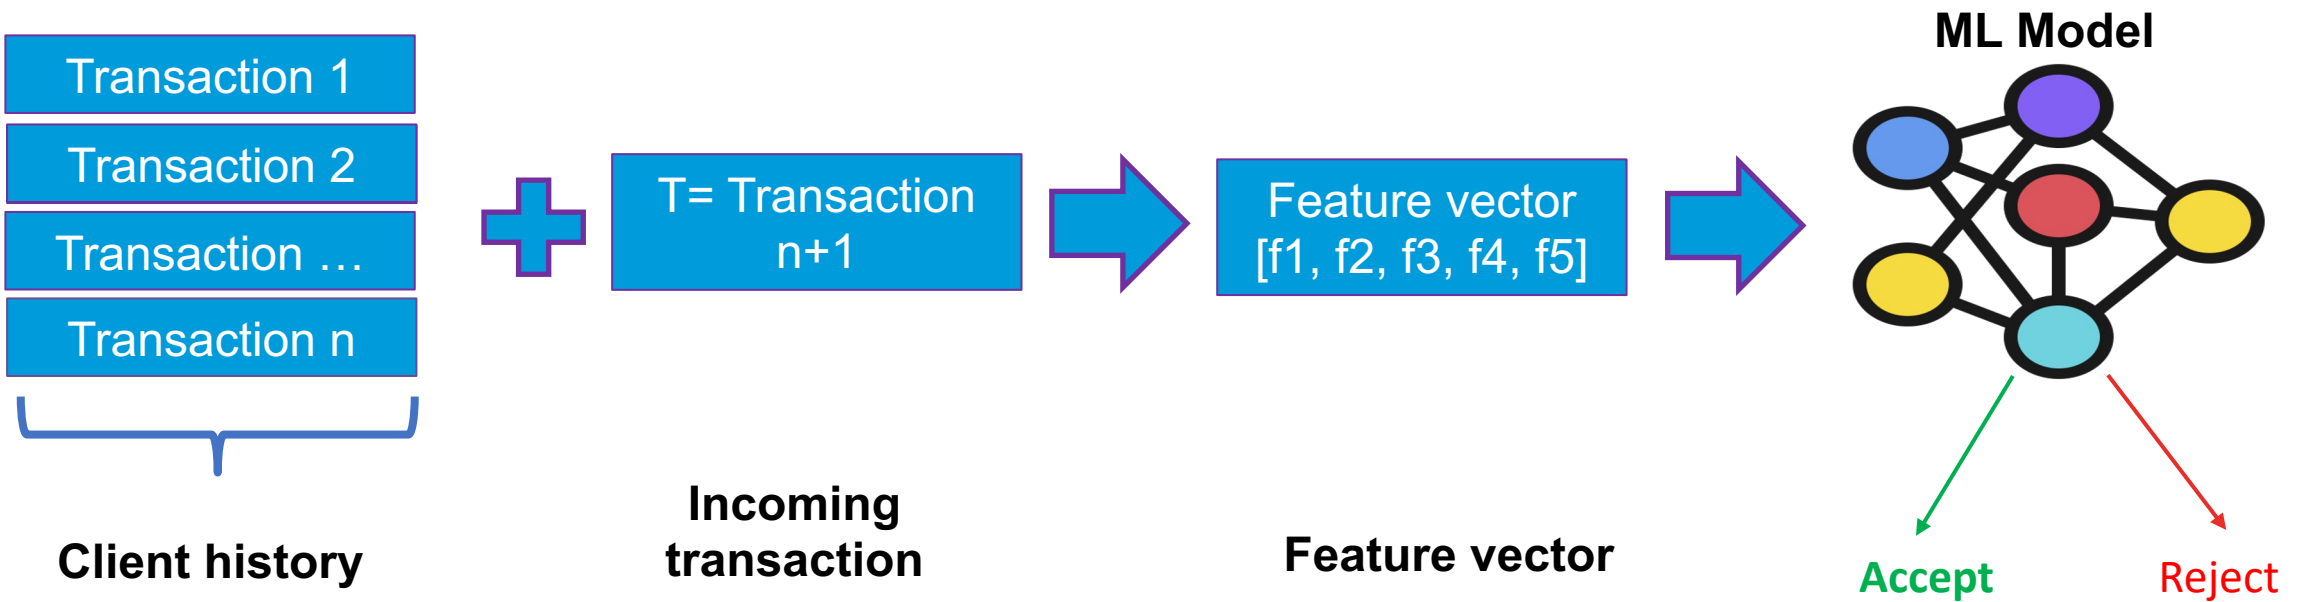
Machine learning robustness (evasion attacks)

Automated decision software in finance



Machine learning robustness (evasion attacks)

Automated decision software in finance



Evasion attack goal

Make the smallest change in transaction n+1
Such that the decision changes from reject to accept

Intersecting variability, machine learning and testing?

We are familiar with your previous work on variability analysis and verification, and would like to know whether you could provide a perspective on the interaction of variability, testing and machine learning, given that you have been working around these topics in the recent years, and they directly relate to the VaMoS topics.



Intersecting variability, machine learning and testing?



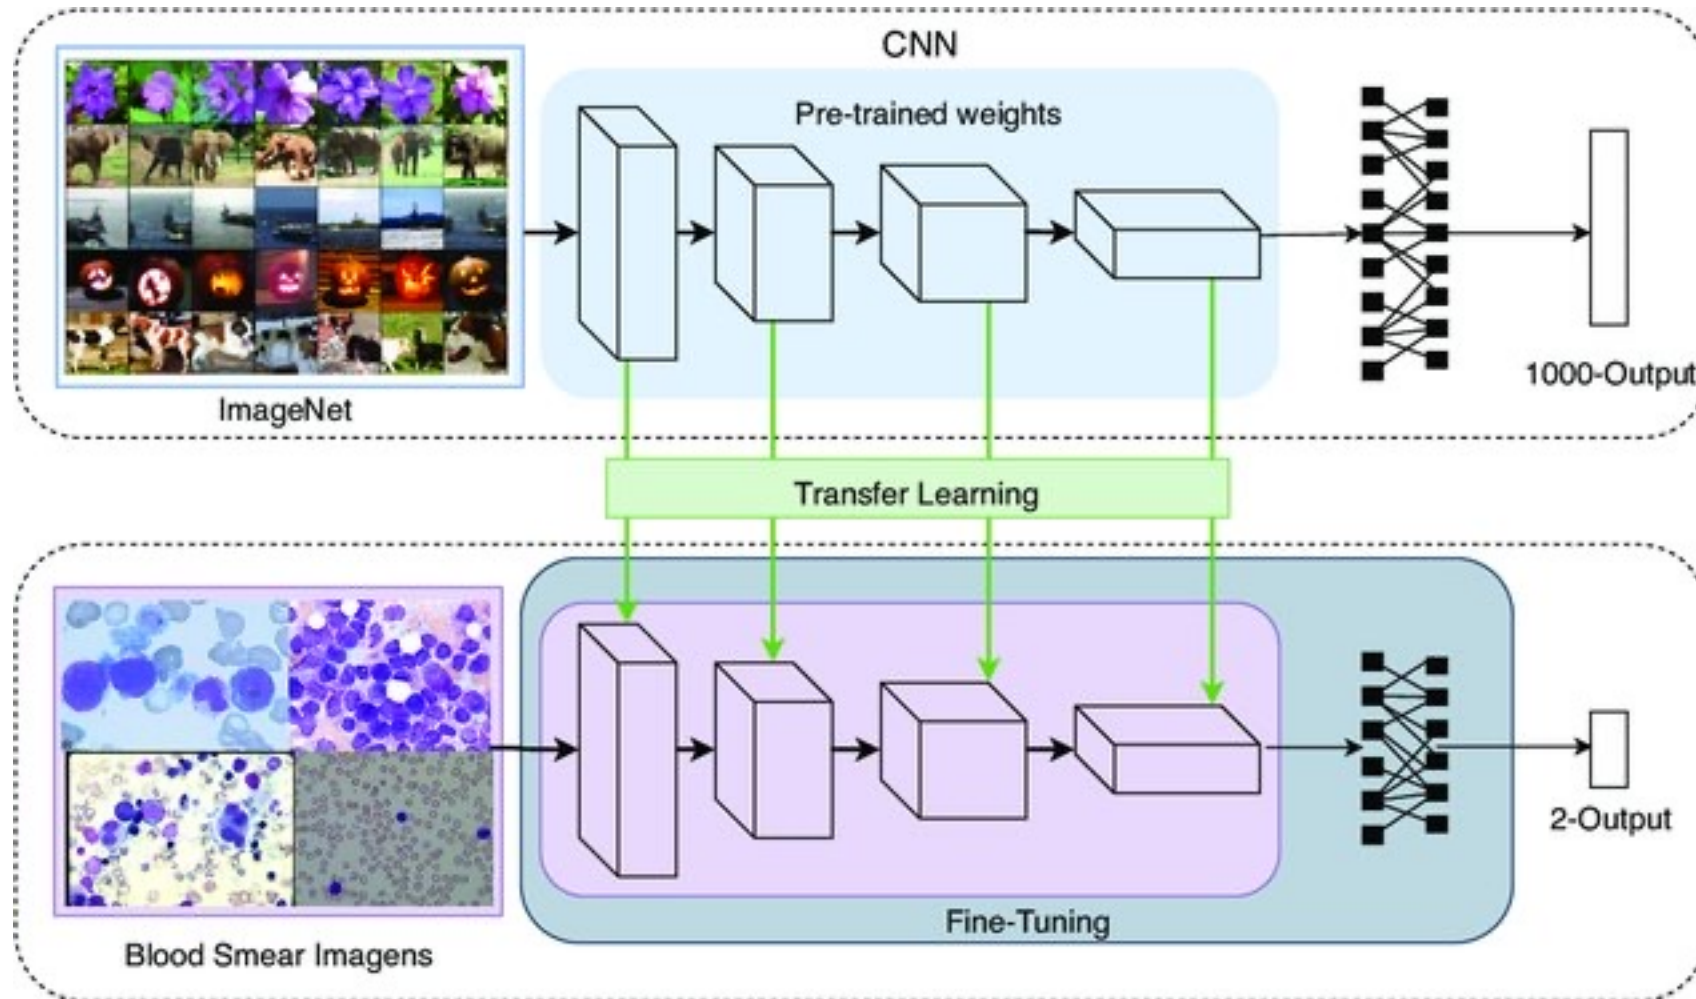
Early-stage research ahead!

Please change room or come back later for solidly established research ☺

Reusability, adaptability, and exploration (in ML)


Reusability, adaptability, and exploration (in ML)

Transfer learning and fine-tuning



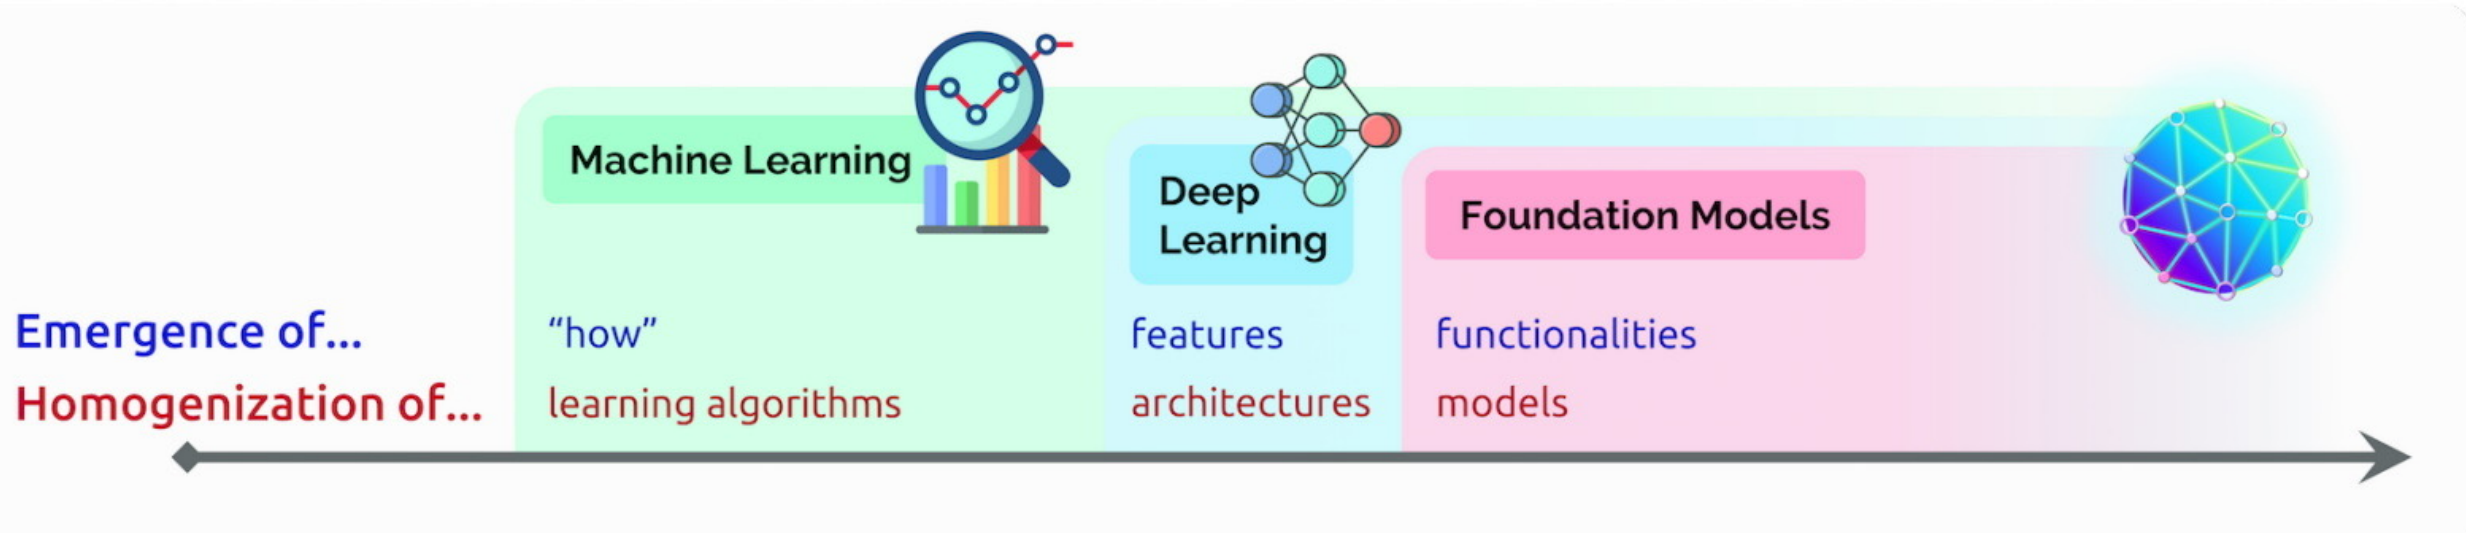
Reusability, adaptability, and exploration (in ML)

Multi-task learning



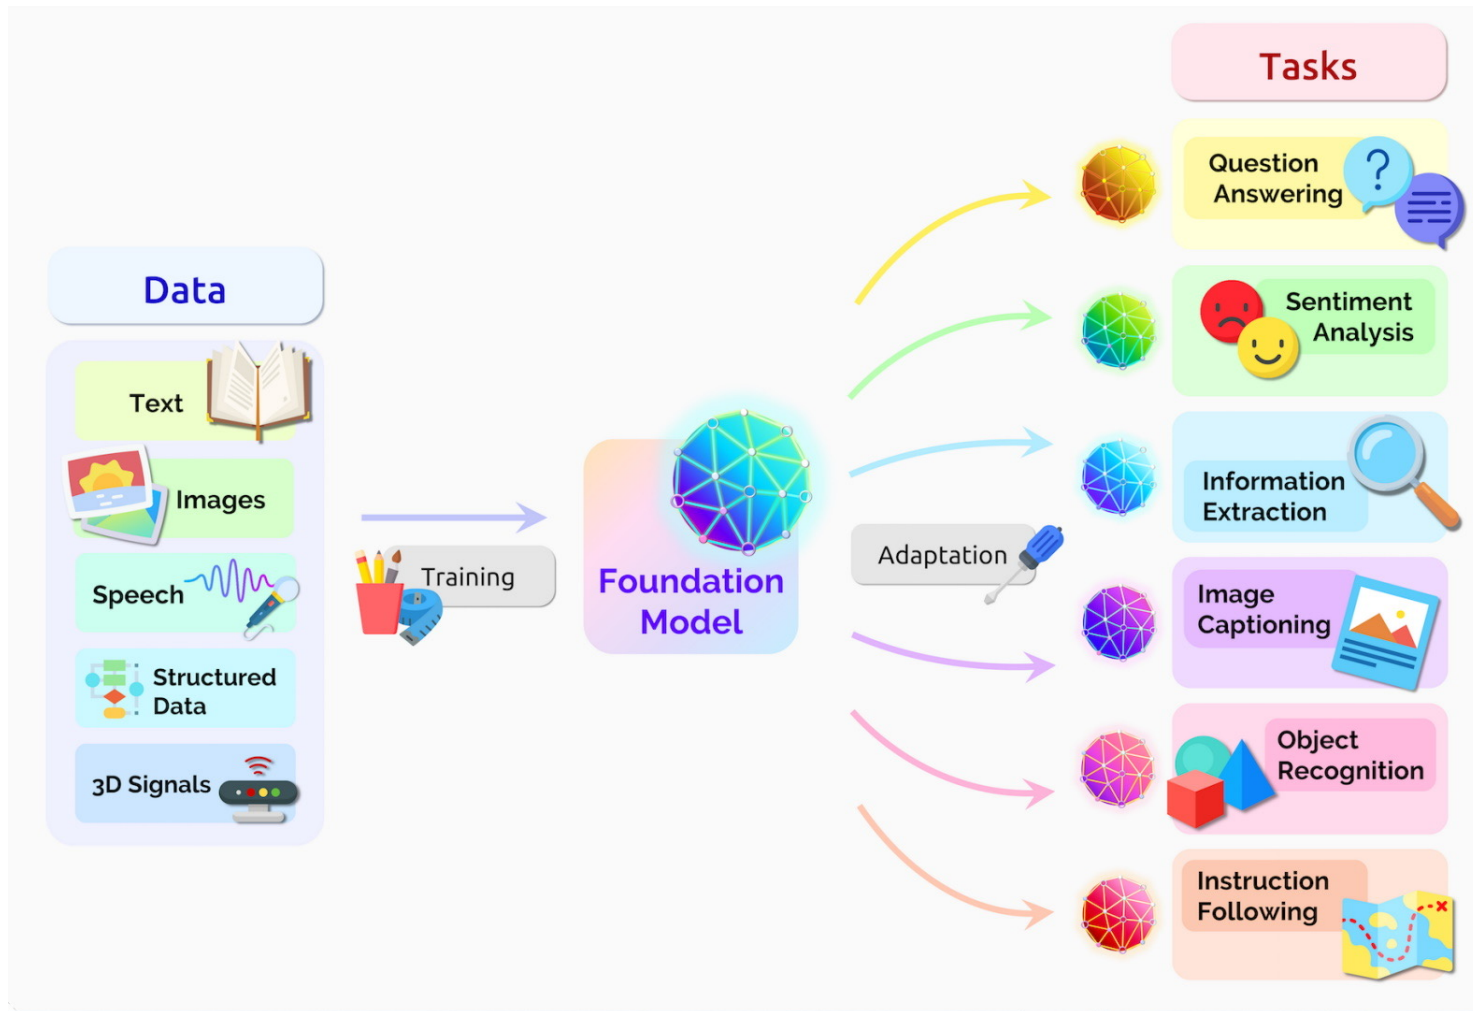
Reusability, adaptability, and exploration (in ML)

Foundation Models



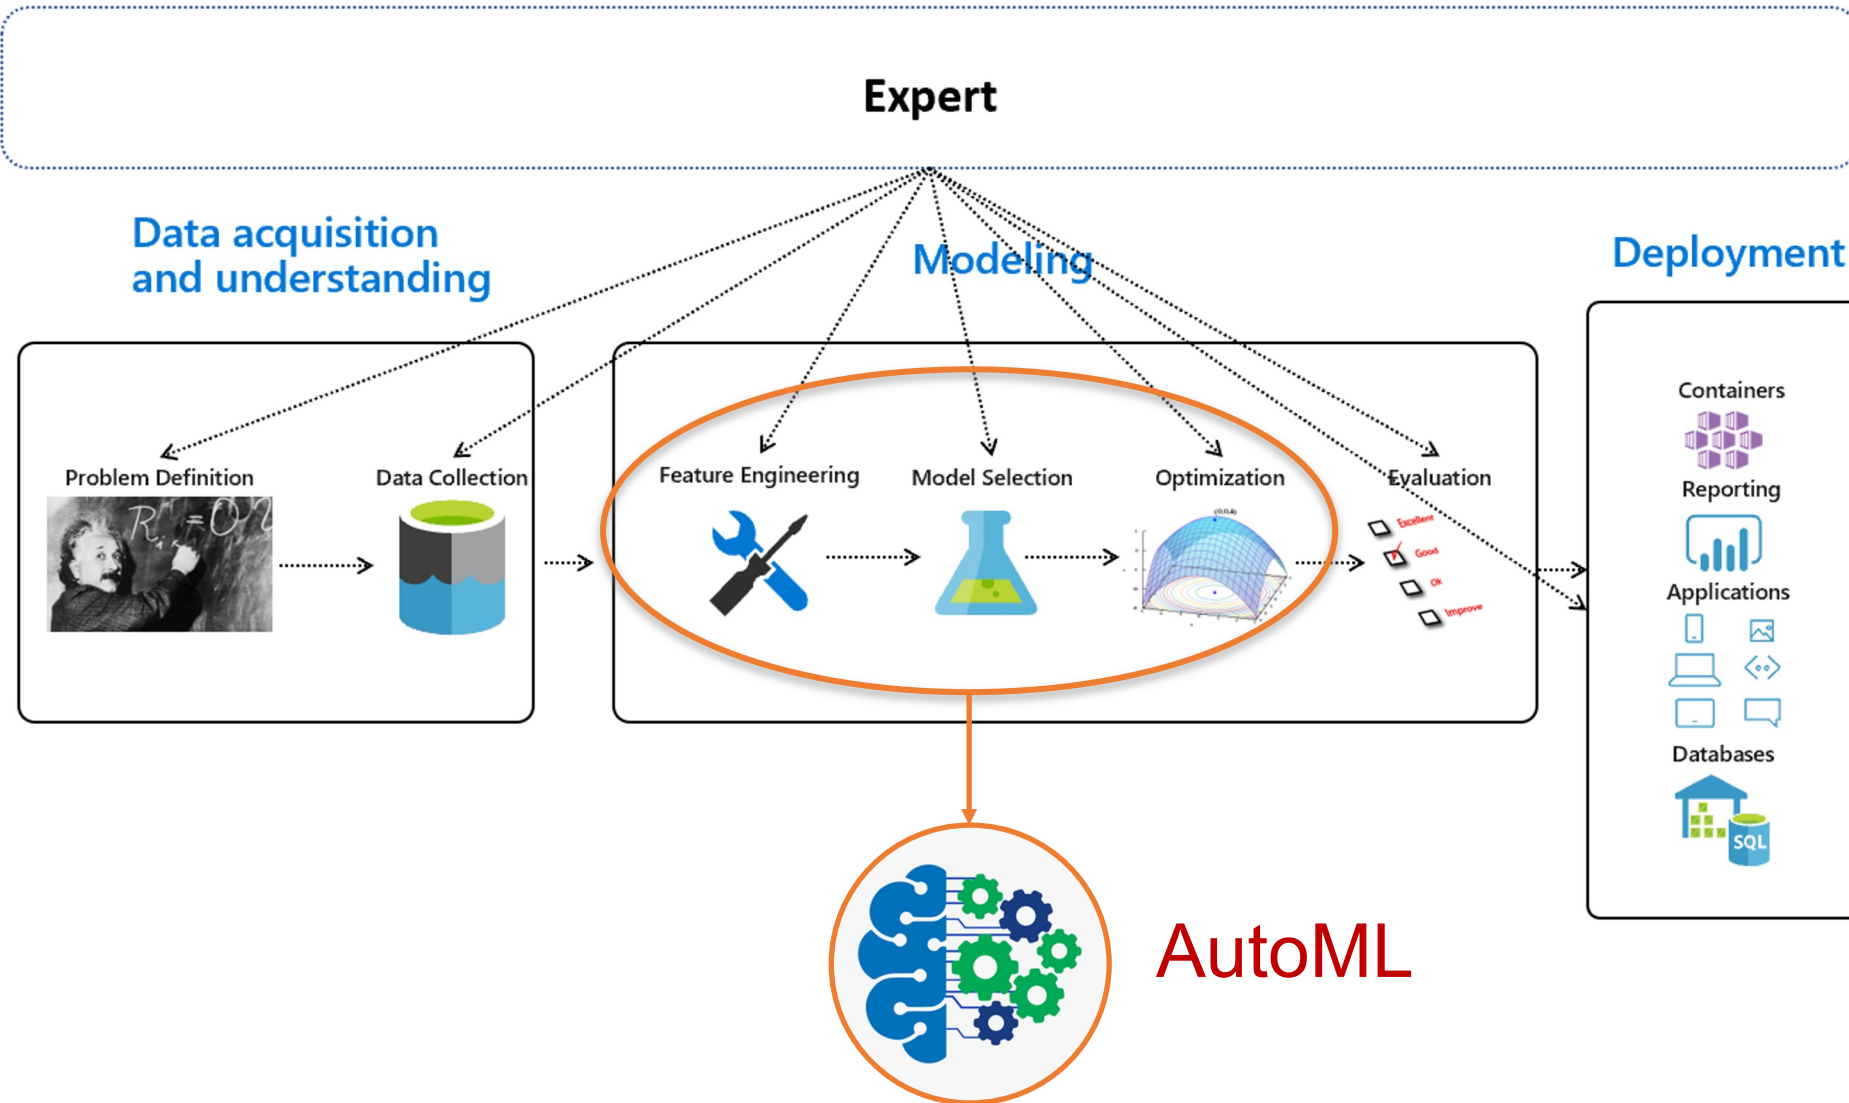
Reusability, adaptability, and exploration (in ML)

Foundation Models



Reusability, adaptability, and exploration (in ML)

AutoML



Automated Search for Configurations of Convolutional Neural Network Architectures

Salah Ghamizi, Maxime Cordy,
Mike Papadakis, Yves Le Traon

<https://github.com/yamizi/FeatureNet>

AutoML: exponential growth since 2017!

TOM SIMONITE BUSINESS 10.13.17 07:00 AM
GOOGLE'S LEARNING SOFTWARE LEARNS TO WRITE LEARNING SOFTWARE

Google's AutoML lets you train custom machine learning models without having to code

Frederic Lardinois @frederic / Jan 17, 2018

Comment

Forbes

Billionaires

Innovation

Leaders

AD



Web Content Management



Online Marketing

Google's self-training AI turns coders into machine-learning masters

Automating the training of machine-learning systems could make AI much more accessible.

by Will Knight January 17, 2018

Google has started using AI to build more advanced AI

Follow @BiNordic

Follow @BINordic

2,708 followers

David Nield

ScienceAlert

22 May 2017 11:44 AM

1227

3,139 views | May 28, 2019, 12:20pm

Google's AutoML And BigQuery ML: The Rise Of One-Click Hyperscale Machine Learning

Our research questions

RQ1: Can we develop a variability model that represents all possible DNN architectures?


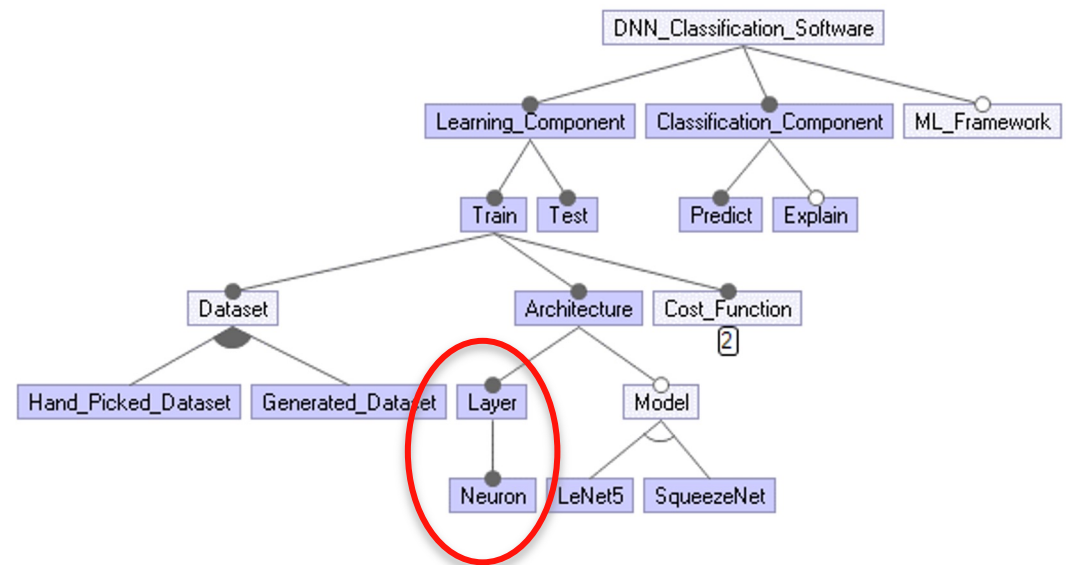

RQ2: Can we effectively search the configuration space and identify well-performing DNN architectures?

RQ3: Does our technique finds DNN architectures that outperform the state of the art?

A straightforward variability model for DNNs

a simple (dense) DL model

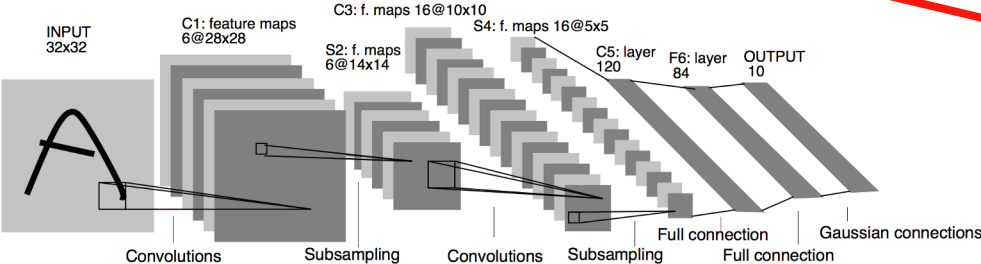
a naive feature model




insufficient to specify
newer architectures

Various architectures and hyperparameters

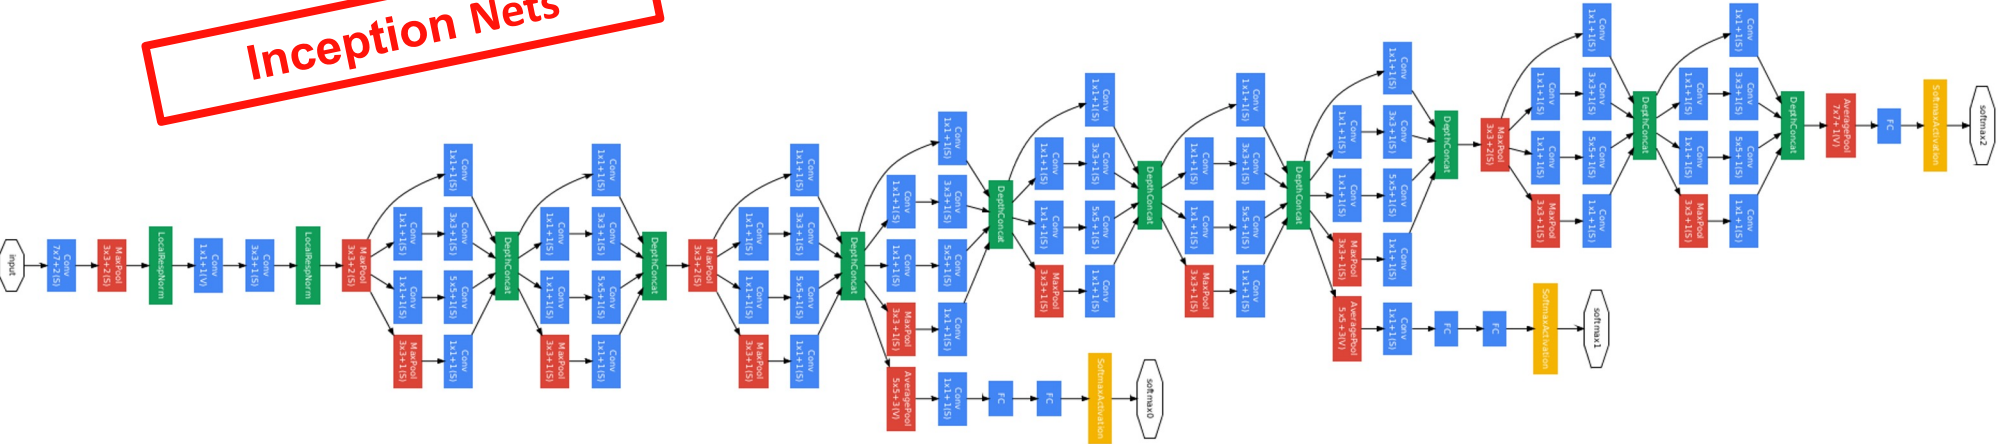
Convolutional Nets




Residual Nets



Inception Nets



Our DNN configuration pipeline

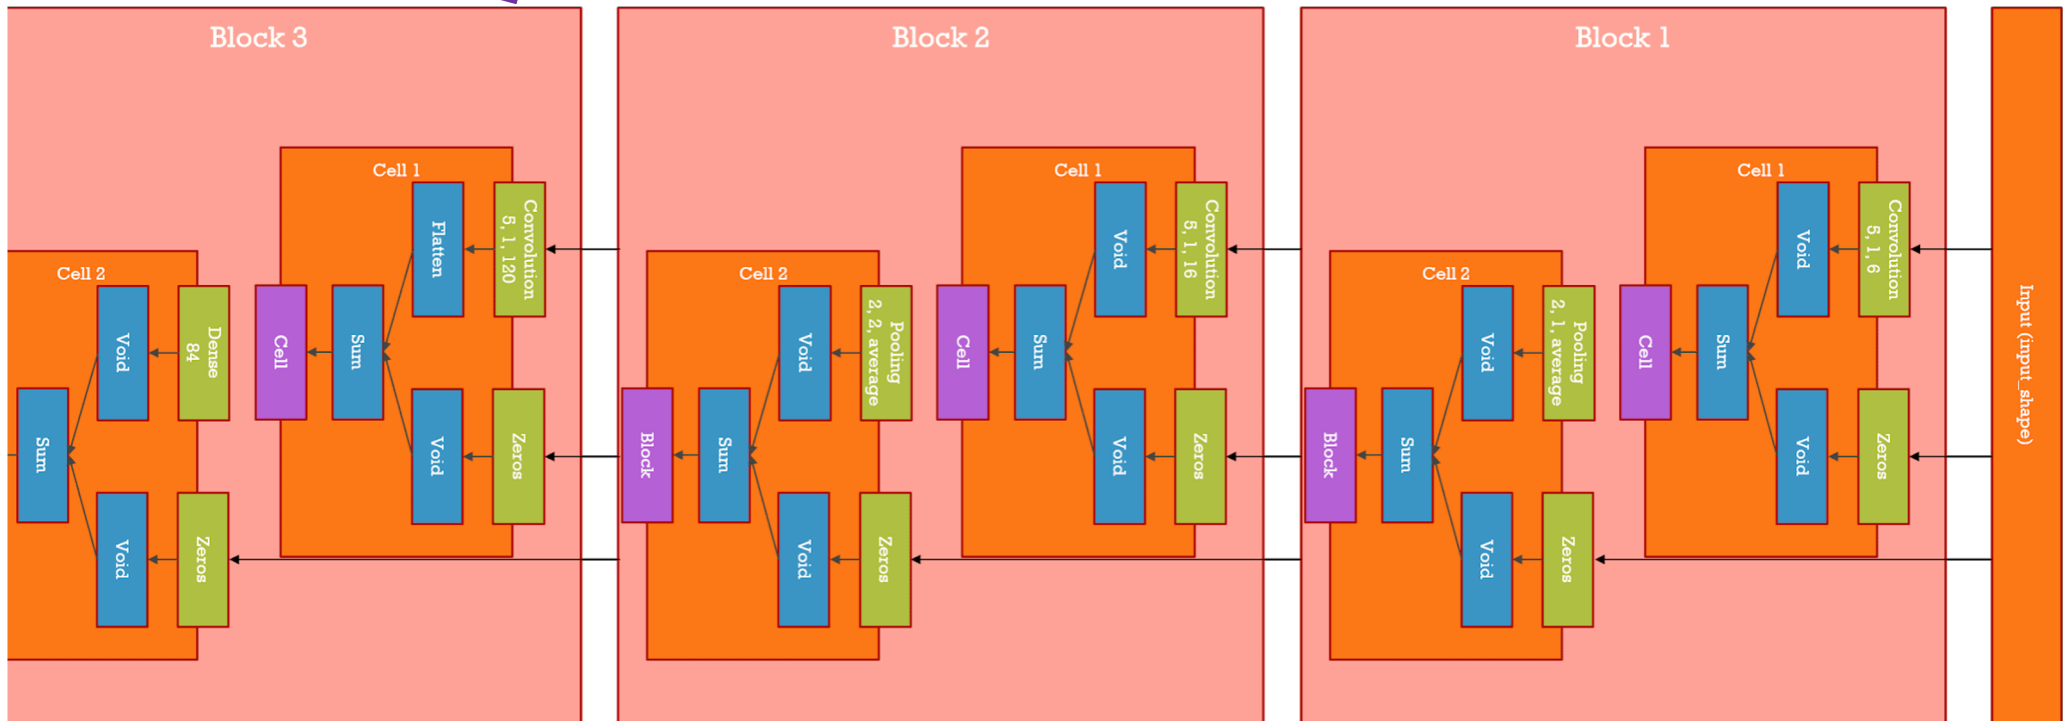


Our Block/Cell Meta-Model

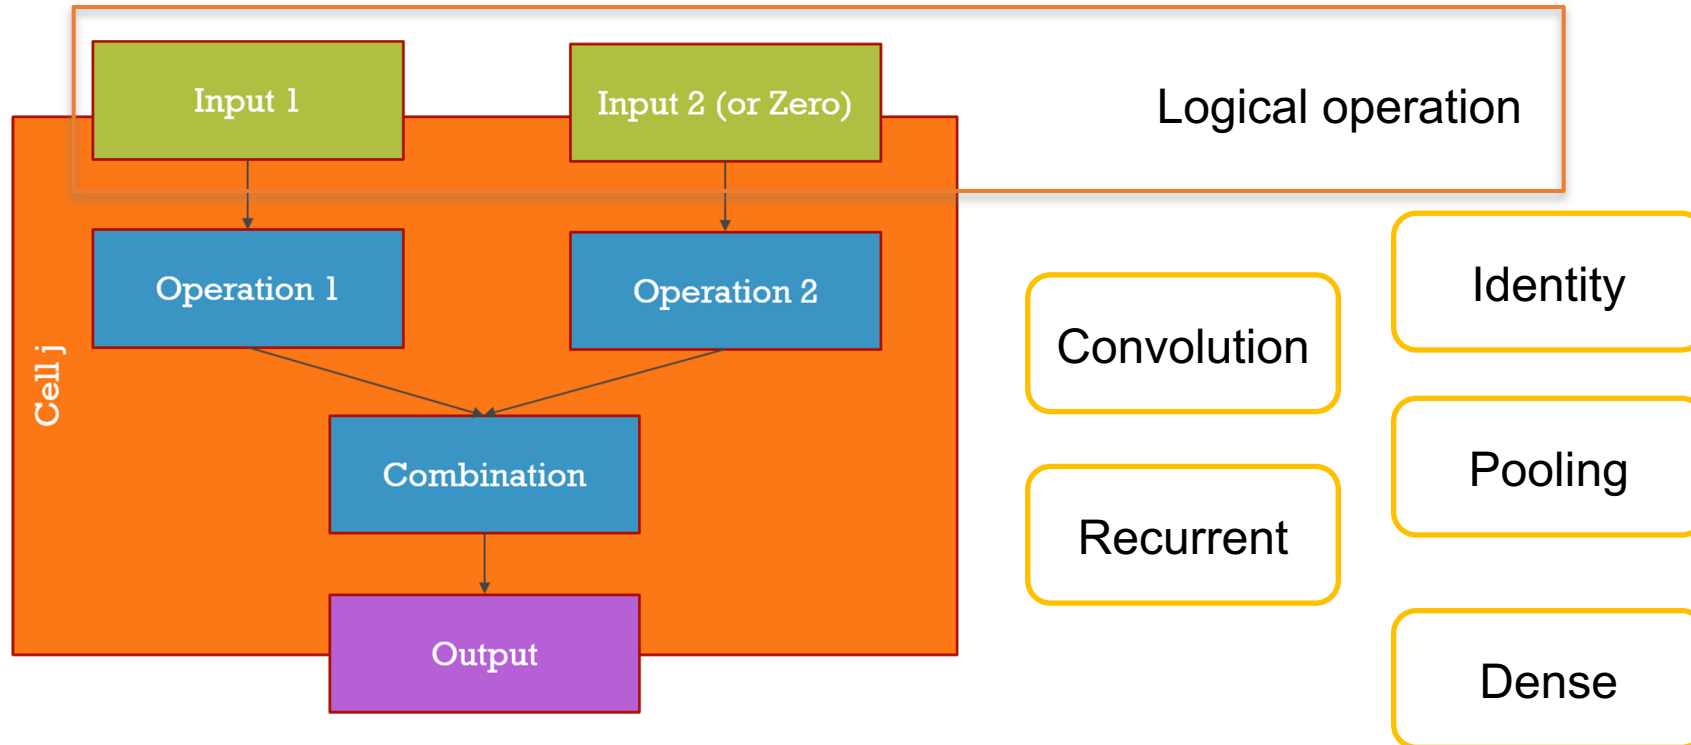
Any configuration sampler

Our Block/Cell Parser

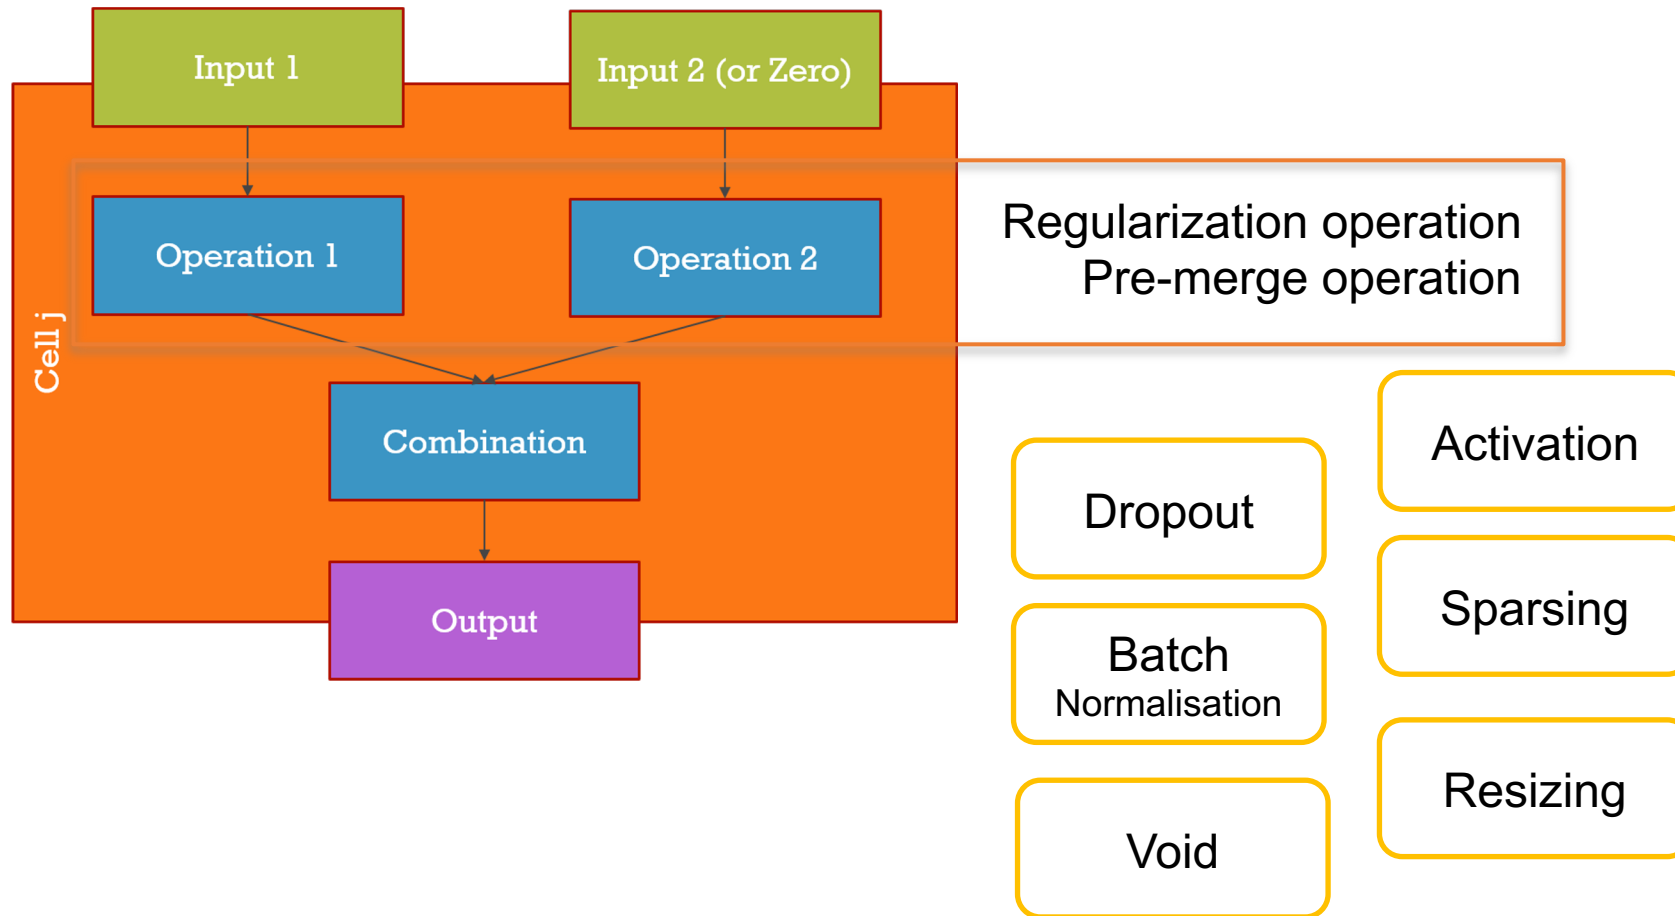
Trained models



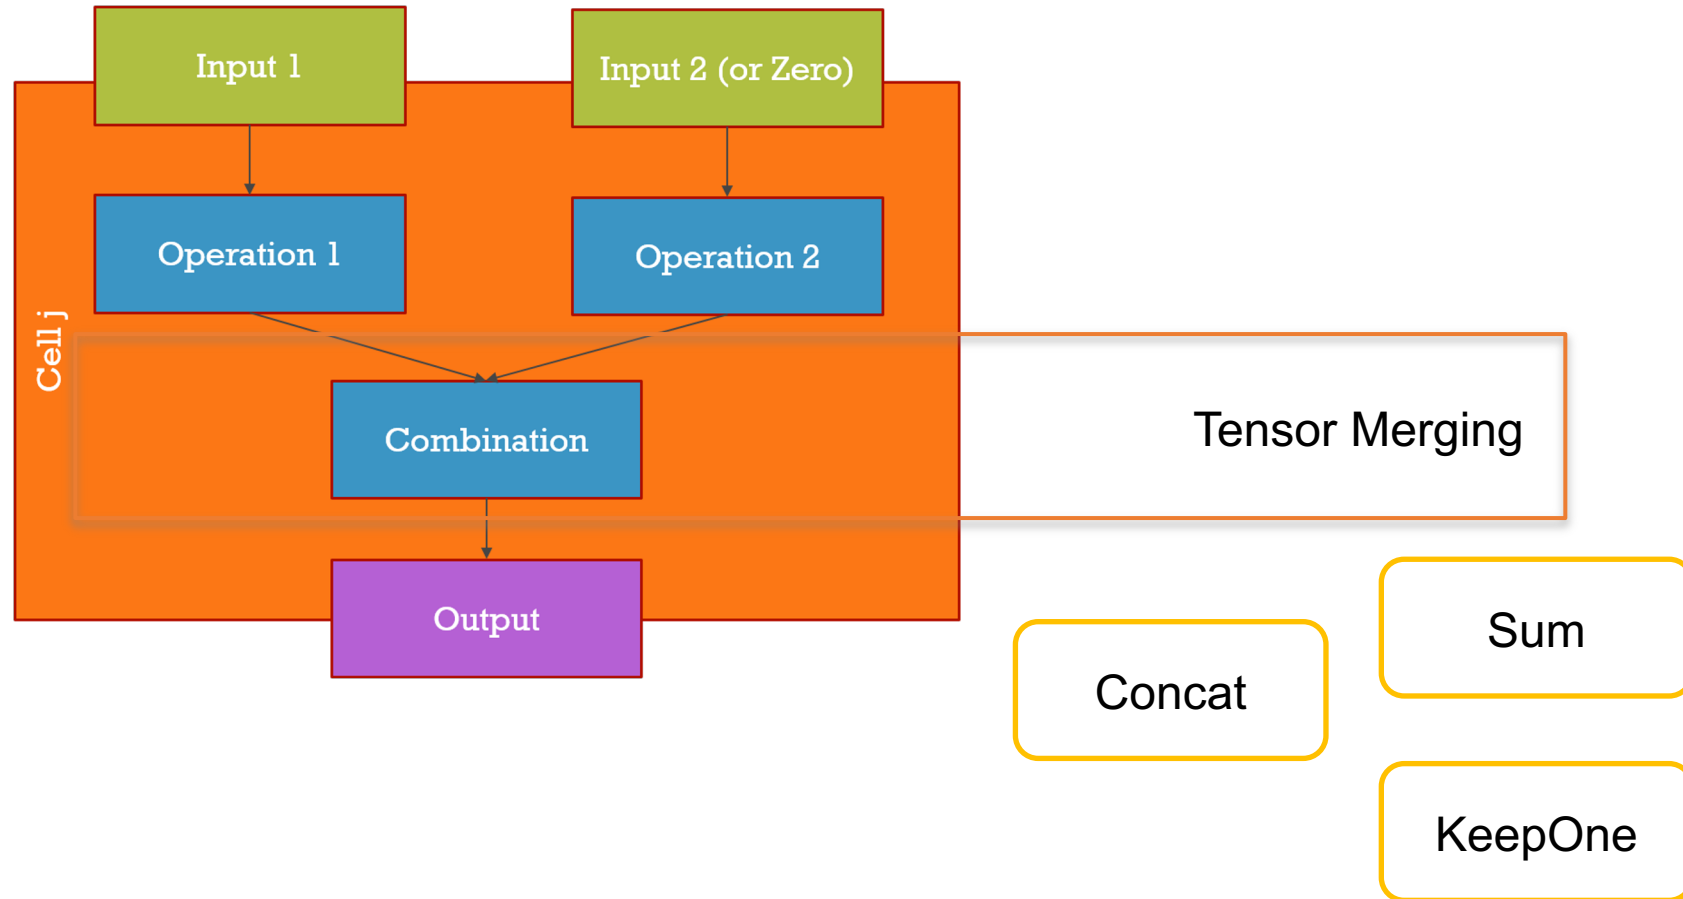
Meta-modelling DNNs with “Cells”



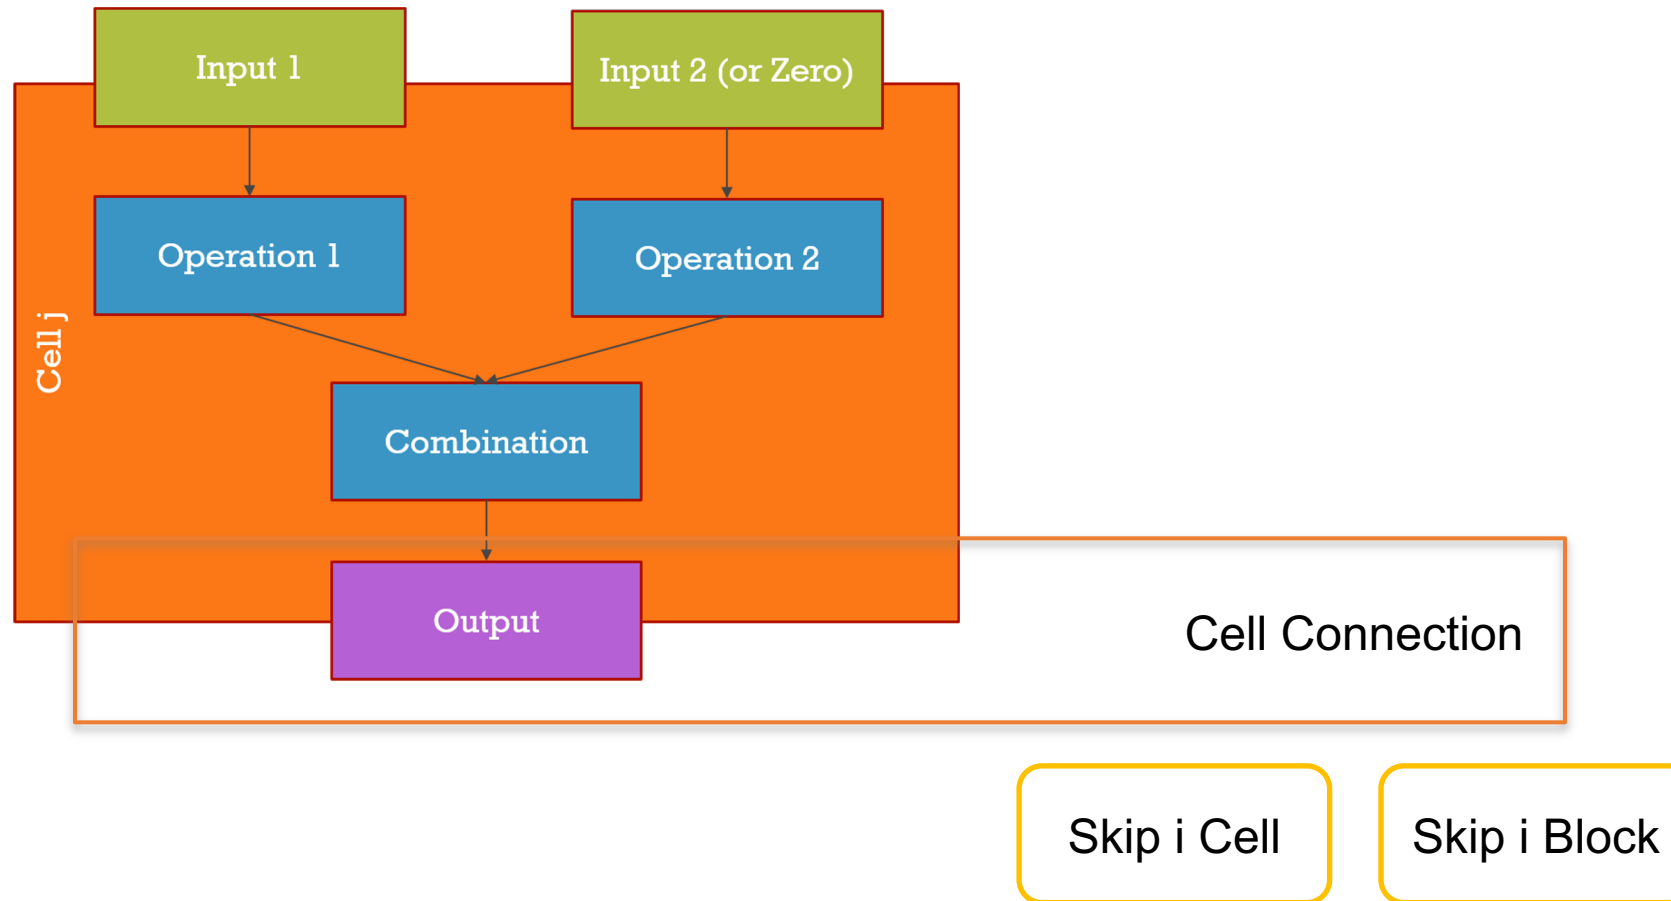
Meta-modelling DNNs with “Cells”



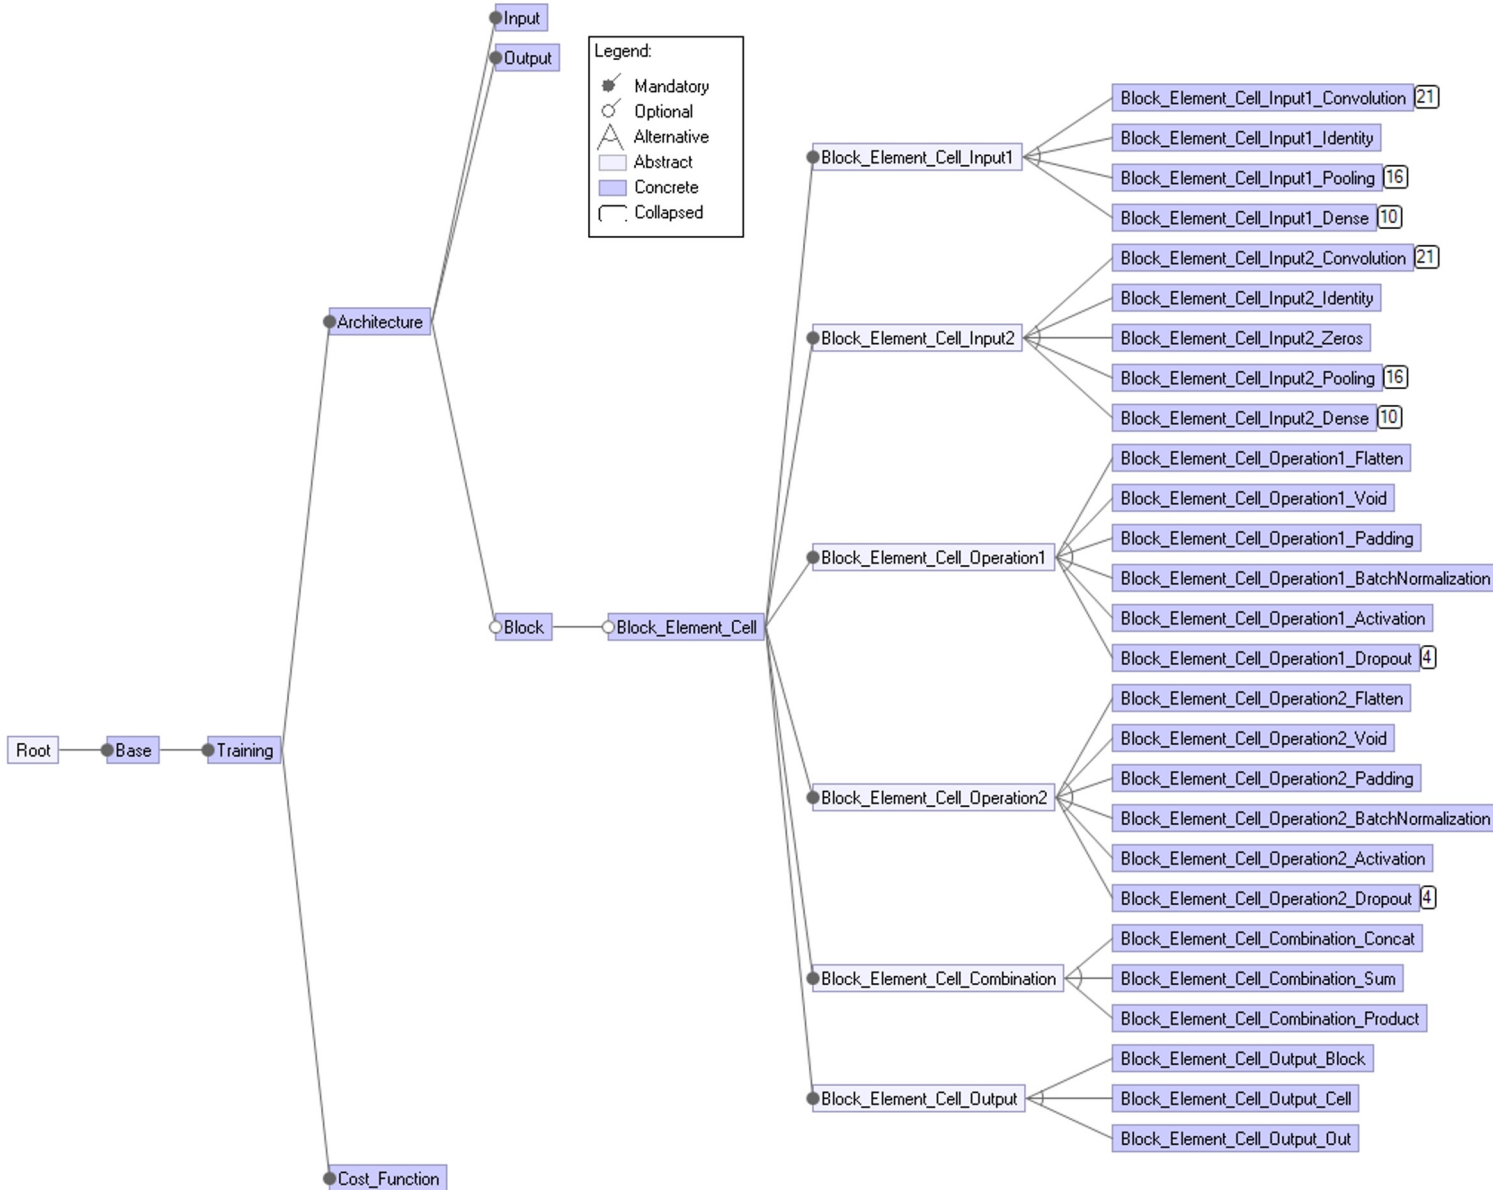
Meta-modelling DNNs with “Cells”



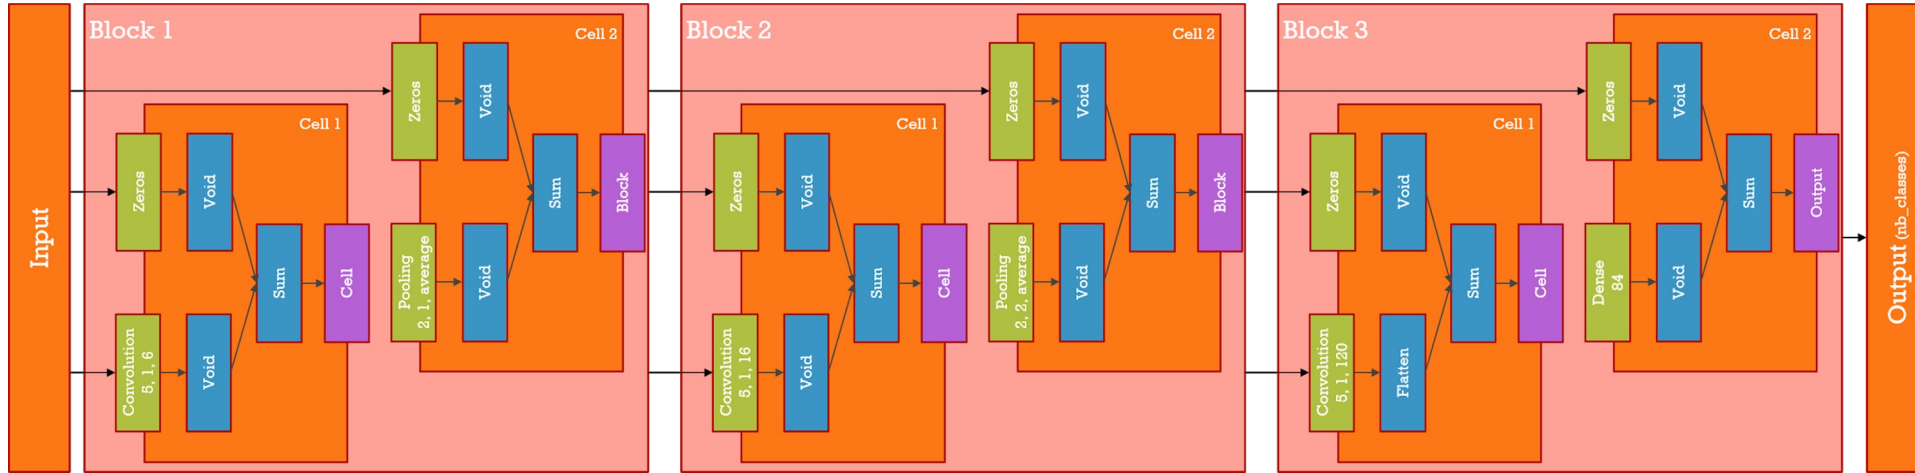
Meta-modelling DNNs with “Cells”



The corresponding variability model

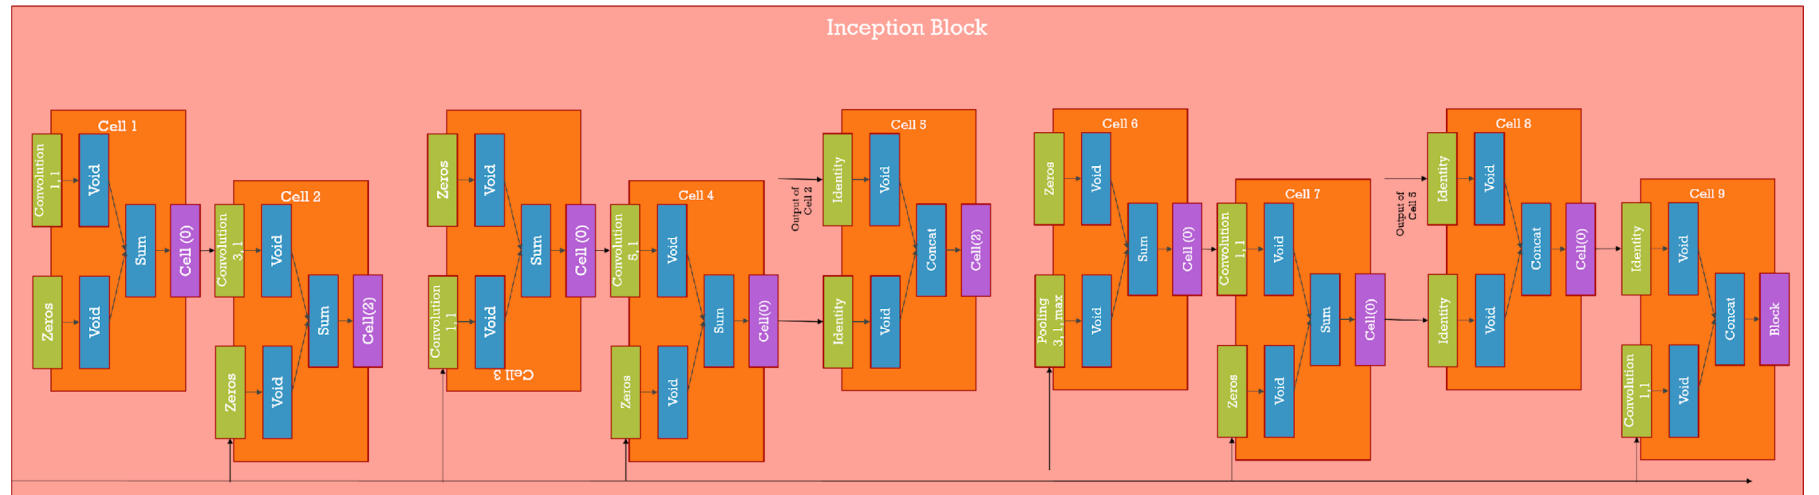


Recreating state of the art architecture



LeNet5 architecture


Inception Block



Experimental settings

- Computational limitations
 - Maximum blocks and cells per block set to 5
 - Limitation on convolution layers' filters (128) and dense layers' neurons (512)
 - No batch normalization, padding, concatenation and product operations
 - Only one output
- SGD with 0.01 learning rate, no decay, no momentum, batch size of 128 (no optimization)
- Datasets: MNIST (60k+10k images) and CIFAR-10 (50k + 10k images)
- Quality metrics: classification accuracy and *efficiency* (accuracy divided by number of weights)

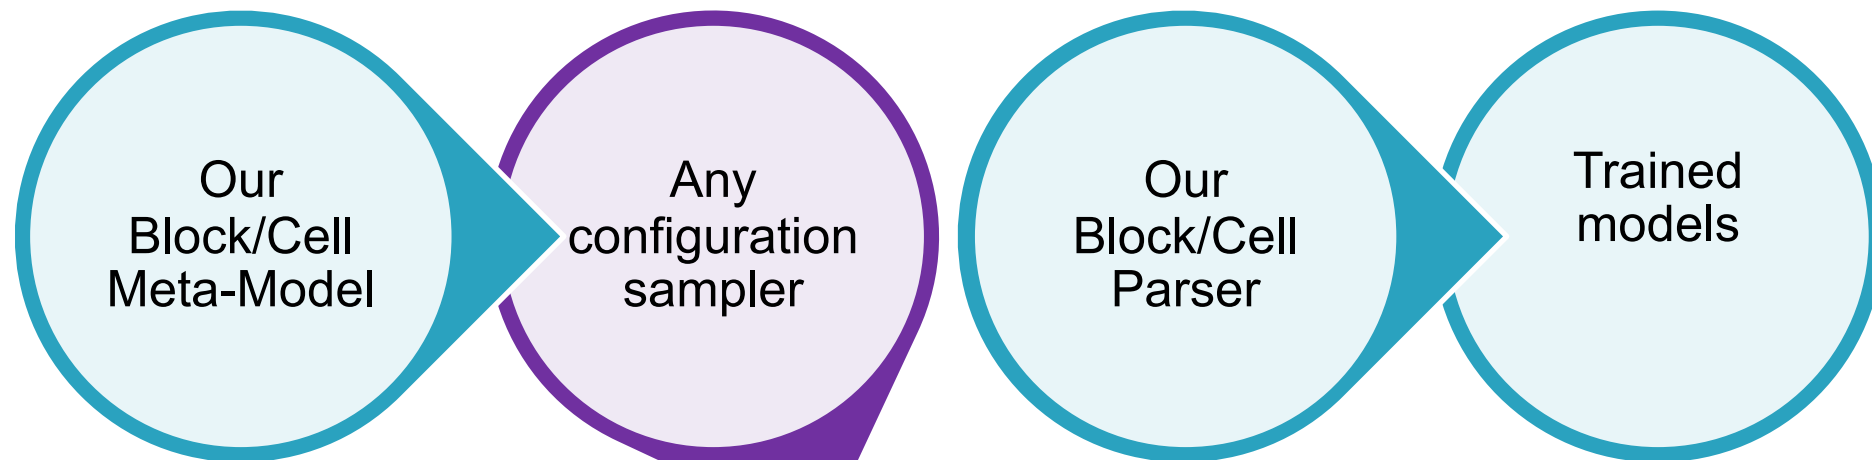
Original implementation vs ours



Our research questions

RQ1: Can we develop a variability model that represents all possible DNN architectures?

RQ2: Can we effectively search the configuration space and identify well-performing DNN architectures?



PLEDGE: Tool for Diversity based sampling

Christopher Henard, Mike Papadakis, Gilles Perrouin, Jacques Klein, and Yves Le Traon.
2013. PLEDGE: a product line editor and test generation tool (SPLC '13 Workshops).

PLEDGE - A Product Line Editor and tests GENERATION tool

Configuration Help

Feature Model Information

Name: Cellphone Number of CNF constraints: 22 Number of core features: 3
 Format: SPLOT Number of features: 11 Number of dead features: 0

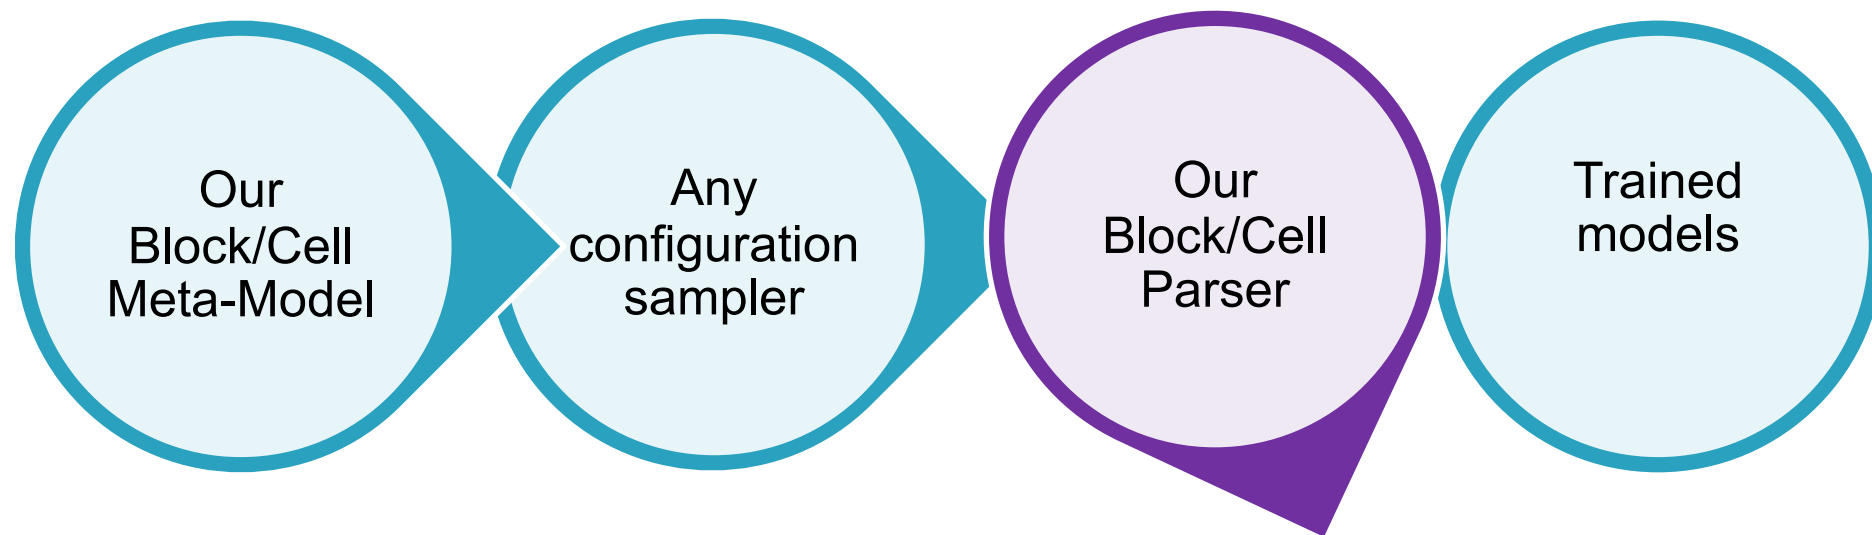
Features		
Number	Name	Type
1	cellphone	Core
2	wireless	Free
3	accu_cell	Core
4	display	Core
5	infrared	Free
6	bluetooth	Free
7	li_ion	Free
8	ni_mh	Free
9	ni_ca	Free
10	color	Free
11	monochrome	Free

CNF Constraints

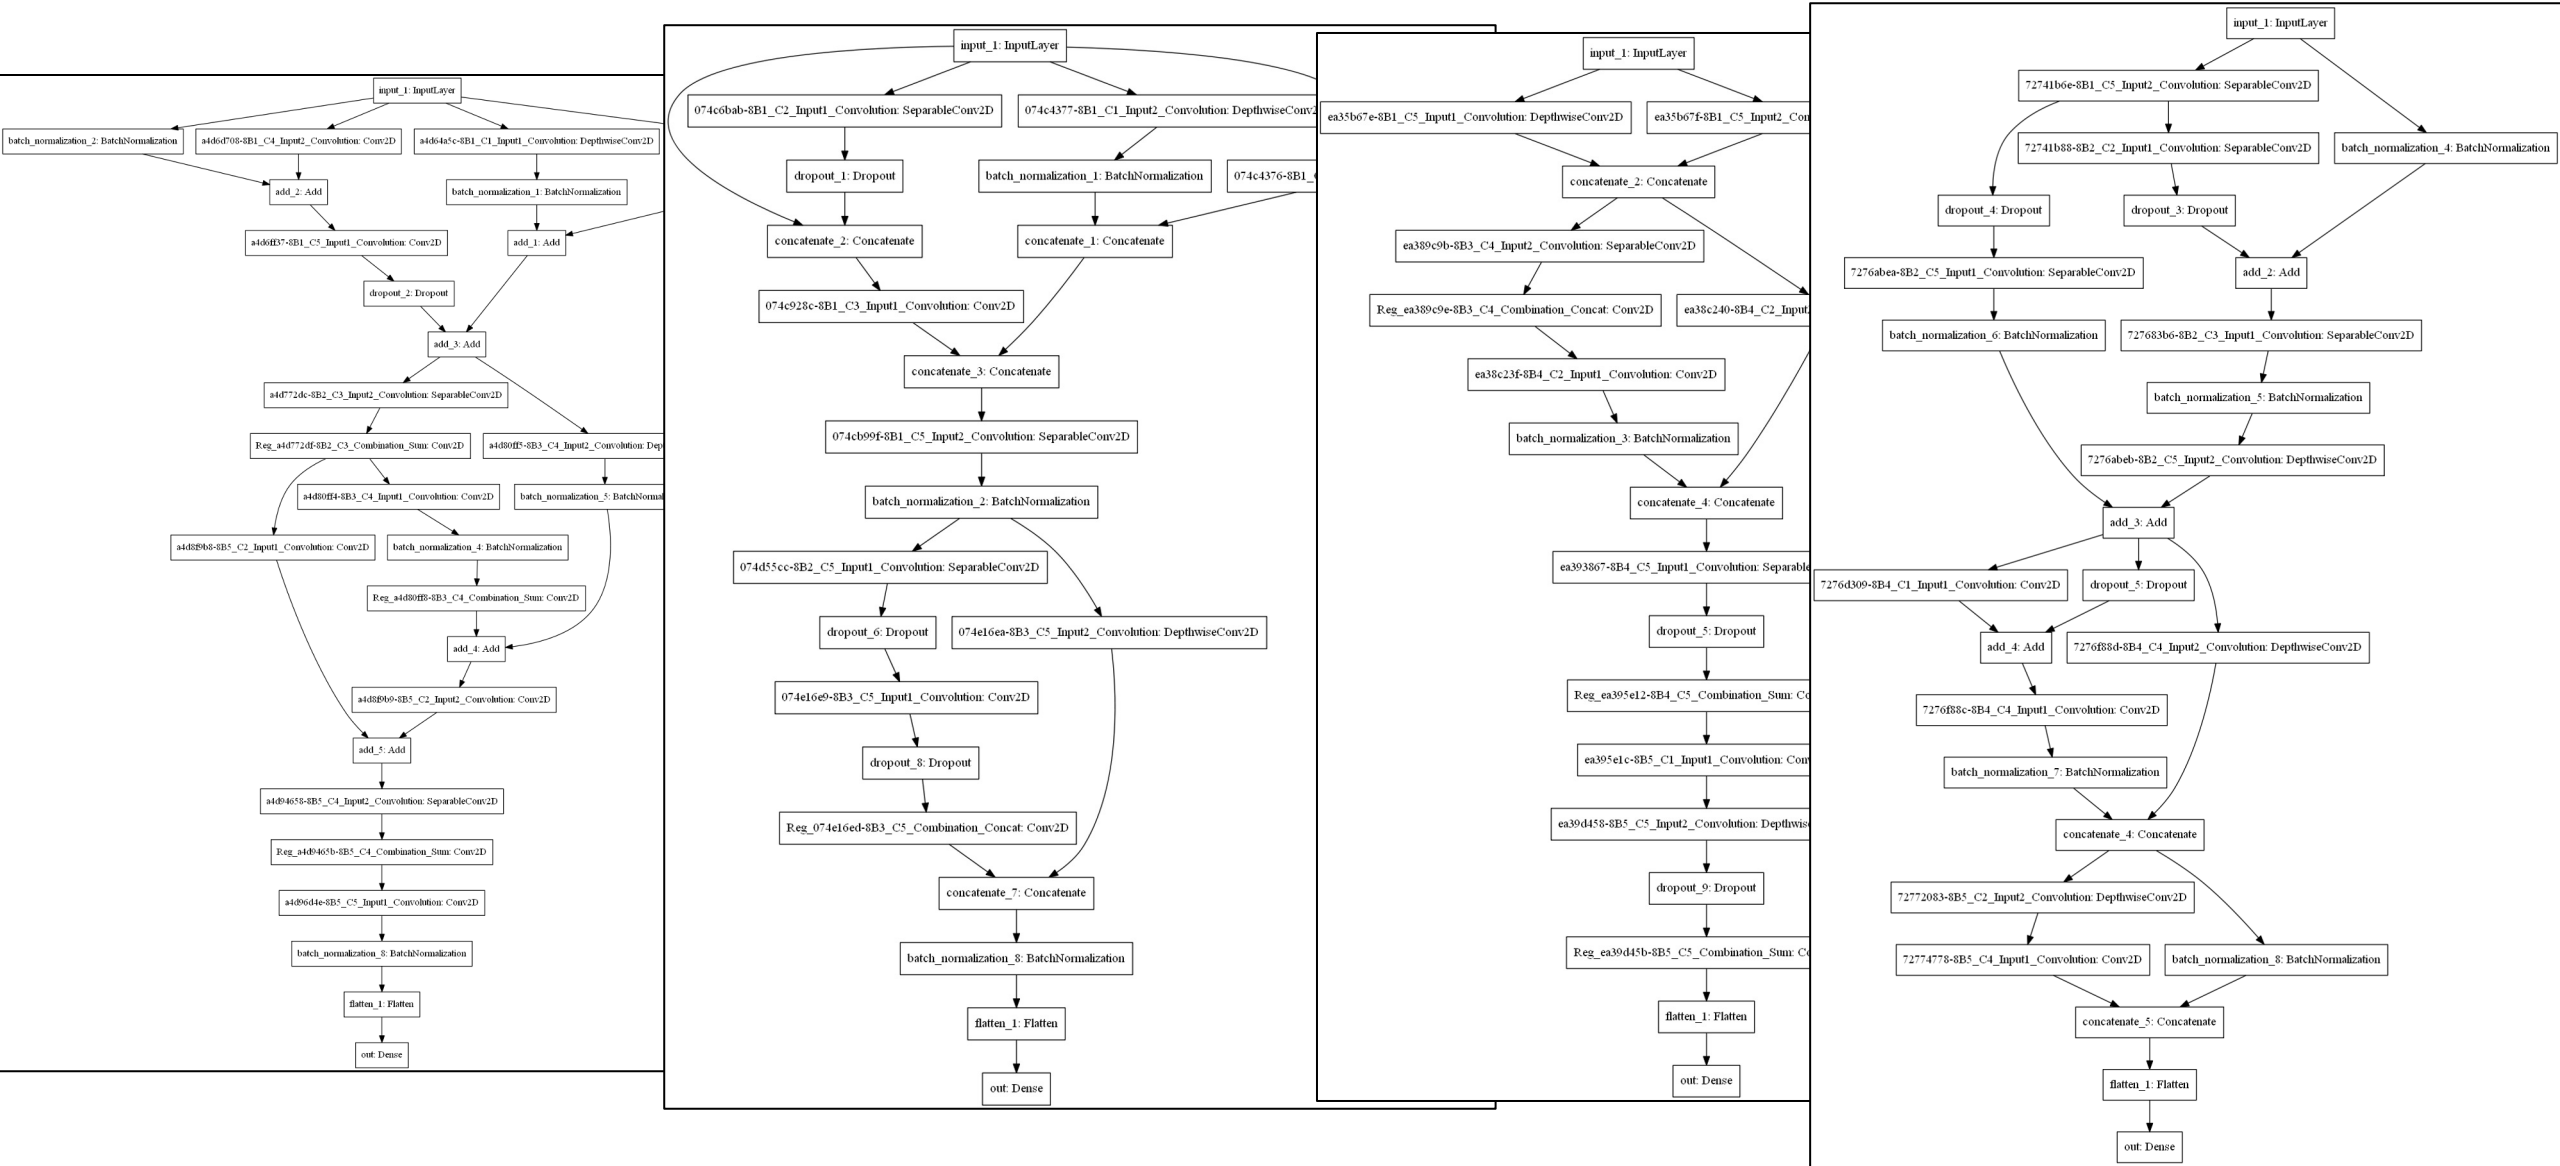
```

cellphone
!wireless OR cellphone
!accu_cell OR cellphone
!cellphone OR accu_cell
!display OR cellphone
!cellphone OR display
!infrared OR wireless
!bluetooth OR wireless
!wireless OR infrared OR bluetooth
!li_ion OR accu_cell
!ni_mh OR accu_cell
!ni_ca OR accu_cell
!accu_cell OR li_ion OR ni_mh OR
!li_ion OR !ni_mh
!li_ion OR !ni_ca
!ni_mh OR !ni_ca
!color OR display
!monochrome OR display
!display OR color OR monochrome
!color OR !monochrome
!bluetooth OR li_ion
  
```


Generating products Iteration number 4722 74%



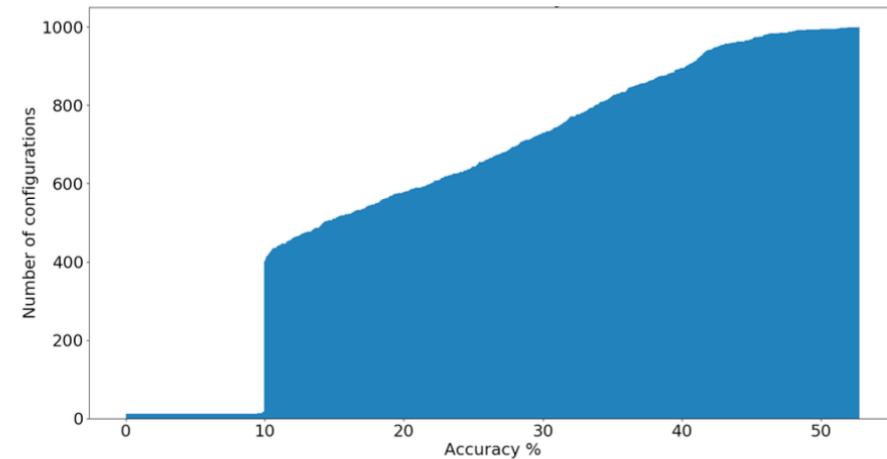
Our technique generates a wide range of architectures



... with diverse accuracy



(a) On MNIST



(b) On Cifar-10

Fig. 7. Distribution of the 1,000 generated architectures over all percentages of accuracy on two datasets. Any point (x, y) of the graph denotes that y architectures achieve an accuracy lower than $x\%$.

... and with diverse size




Fig. 8. Distribution of the accuracy and size of 1,000 generated architectures. The size is given in terms of number of trainable weights of the architectures. By sampling a diverse set of configurations (in terms of feature differences), our technique is able to generate architectures with a wide range of sizes an few thousands to millions of weights. Moreover, we observe that architectures with high accuracy are not necessarily the ones with the largest size, while architectures performing poorly are found in all size ranges.

Our research questions

RQ1: Can we develop a variability model that represents all possible DNN architectures?

RQ2: Can we effectively search the configuration space and identify well-performing DNN architectures?

RQ3: Does our technique find DNN architectures that outperform the state of the art?

FeatureNet vs. SotA architectures * (MNIST)

Table 1: Accuracy, size, and efficiency of LeNet5, SqueezeNet and our best sampled architectures, with a 12-epoch training, on MNIST (top part) and CIFAR-10 (bottom part).

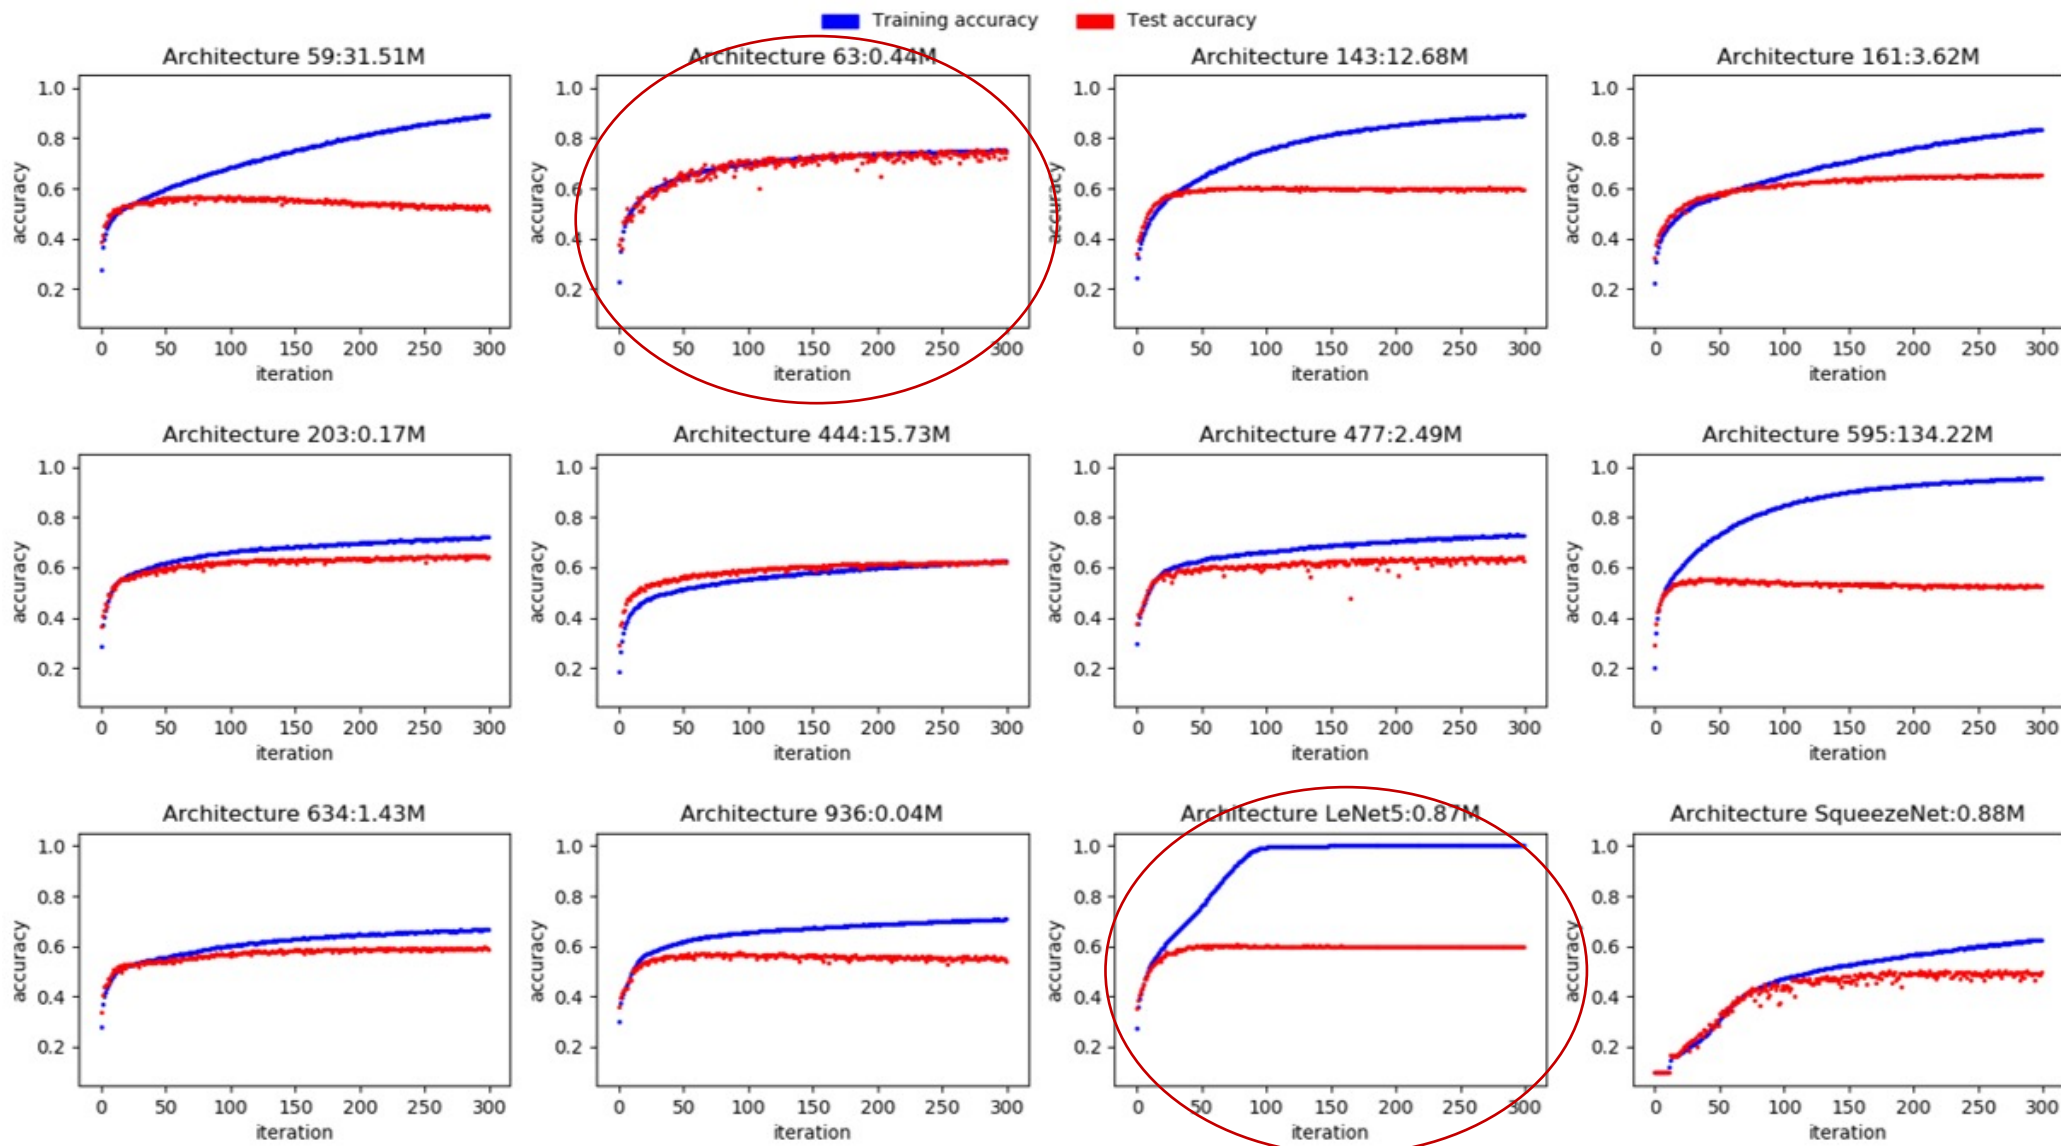
Dataset	Architecture	Accuracy	Size	Efficiency
MNIST 1,000 architectures with previous restrictions 100 LeNet5 architectures 100 smaller architectures	LeNet5	97.14%	545546	1.78
	Top S1	97.74%	365194	2.68
	Top S2	97.65%	570218	1.71
	Top S3	92.31%	43578	21.18
	SqueezeNet	43.67%	858154	0.51
CIFAR-10	LeNet5	49.13%	868406	0.57
	Top S1	52.77%	2494858	0.21
	Top S2	57.79%	862646	0.67
	Top S3	37.44%	38842	9.64
	SqueezeNet	17.96%	876970	0.51

FeatureNet vs. SotA architectures * (CIFAR-10)

Table 2: Accuracy of the 10 best architectures from S1, LeNet5 and SqueezeNet on CIFAR-10 and at 12, 300 and 600 training epochs. Architectures are ordered by descending order of accuracy at 600 epochs. * indicates shortened training.

Architecture	Size	(12)	(300)	(600)
#063	0.45M	48.25%	74.28%	74.74%
#203	0.17M	52.64%	64.55%	65.25%
#161	3.62M	48.57%	64.46%	65.24%
#477	2.49M	52.77%	63.43%	64.25%
#444	15.73M	49.97%	62.40%	62.80%
#143	12.68M	51.37%	60.46%	60.17%
LeNet5	0.87M	49.13%	59.42%	59.26%
#634	1.43M	51.11%	59.76%	* 59.06%
#936	0.04M	48.92%	54.26%	56.33%
#595	134.22M	52.16%	52.64%	* 53.64%
SqueezeNet	0.88M	17.96%	46.17%	49.69%
#059	31.51M	51.9%	52.10%	49.47%

Variability in performance



vs. other NAS methods

On Cifar-10:


NAS technique	Error
FeatureNet	25%
AutoKeras (TexasU,18')	15%
NasNet (Google, 18')	2.4%

But FeatureNet:

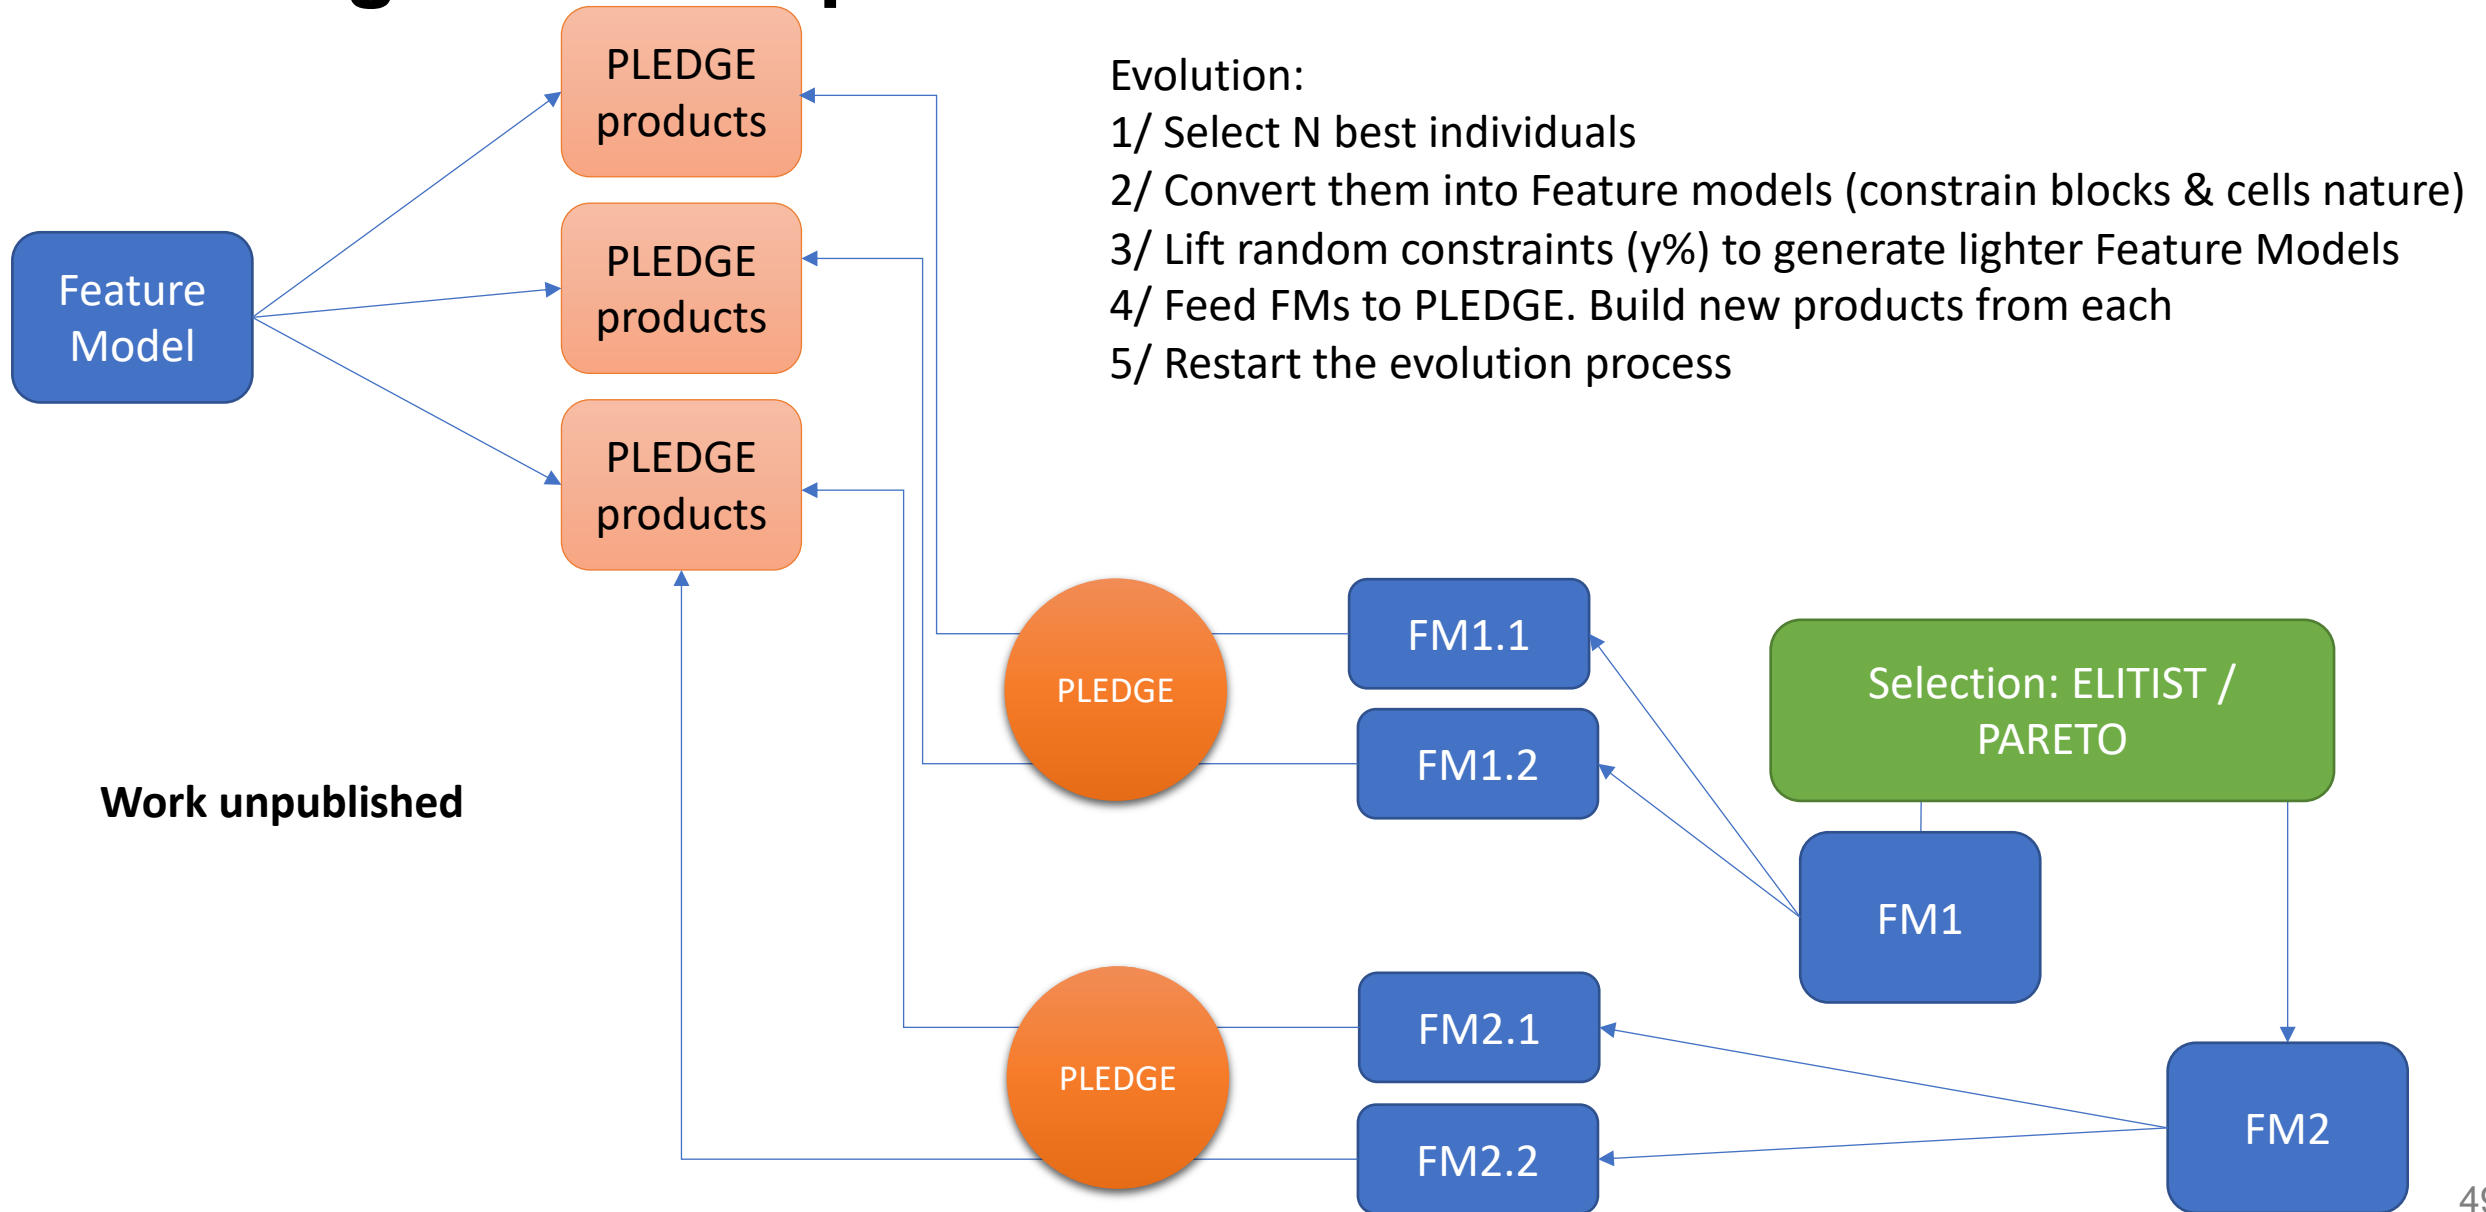
- Deals only with the architectures, others also consider:
 - Training optimization
 - Data Augmentation
- Is more efficient:
 - our full search took less than 1 GPU day
 - NasNet requires hundreds to thousands of GPU days

Moving ahead

Moving ahead: improved search



Moving ahead: improved search




Adversarial Robustness in Multi-Task Learning: Promises and Illusions

Salah Ghamizi, Maxime Cordy,
Mike Papadakis, Yves Le Traon



Reusability, adaptability, and exploration (in ML)

Multi-task learning



Machine learning robustness (evasion attacks)

Original example



Small adversarial noise



Adversarial example



What humans still see

ML predicts:
"Panda"
(80% confidence)



What ML predicts: "Gibbon"
(99% confidence)

Gibbon

Principles of evasion attacks




Image modified from <https://towardsdatascience.com/know-your-enemy-7f7c5038bdf3>

Objective: for x find δ

✓ With $f(x) \neq f(x + \delta)$

✓ With $L_p(x, x + \delta) < \epsilon$

By solving the optimization problem:

$$\operatorname{argmax}_{\delta} \mathcal{L}(x + \delta, y)$$

$$\text{s.t. } \|\delta\|_p \leq \epsilon$$

Evasion attacks in multi-task learning

An attacker seeks the perturbation δ that will maximize the joint loss function of the attacked tasks – i.e. the summed loss, within a p-norm bounded distance ϵ .

$$\operatorname{argmax}_{\delta} \mathcal{L}_i(x + \delta, y_i) \text{ s.t. } \|\delta\|_p \leq \epsilon$$

Single-task adversarial attacks

$$\operatorname{argmax}_{\delta} \mathcal{L}(x + \delta, \bar{y}) \text{ s.t. } \|\delta\|_p \leq \epsilon$$

Multi-task adversarial attacks

Does learning multiple tasks increase robustness?

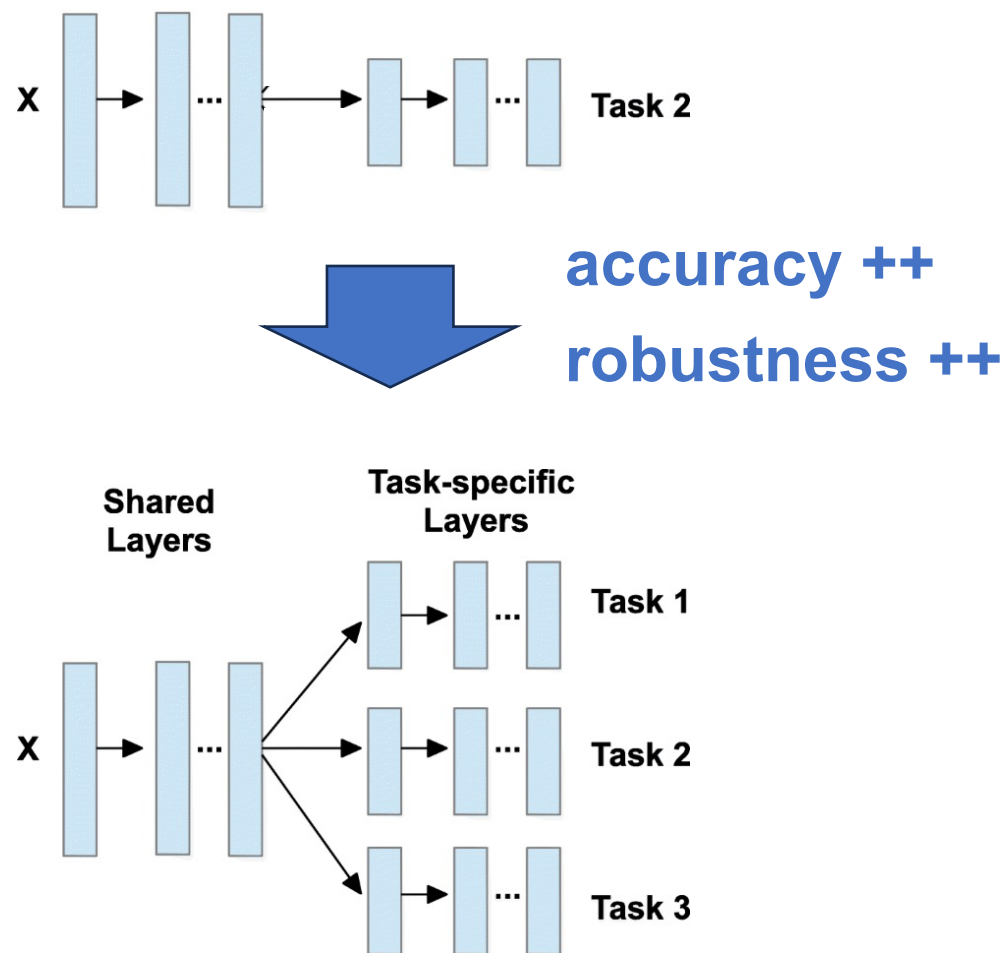
Multitask Learning Strengthens Adversarial Robustness

Chengzhi Mao, Amogh Gupta*, Vikram Nitin*,
Baishakhi Ray, Shuran Song, Junfeng Yang, and Carl Vondrick

Columbia University, New York, NY, USA
mcz,rayb,shurans,junfeng,vondrick@cs.columbia.edu,
ag4202,vikram.nitin@columbia.edu

Abstract. Although deep networks achieve strong accuracy on a range of computer vision benchmarks, they remain vulnerable to adversarial attacks, where imperceptible input perturbations fool the network. We present both theoretical and empirical analyses that connect the adversarial robustness of a model to the number of tasks that it is trained on. Experiments on two datasets show that attack difficulty increases as the number of target tasks increase. Moreover, our results suggest that when models are trained on multiple tasks at once, they become more robust to adversarial attacks on individual tasks. While adversarial defense remains an open challenge, our results suggest that deep networks are vulnerable partly because they are trained on too few tasks.

Keywords: Multitask Learning, Adversarial Robustness



Does learning multiple tasks increase robustness?

$$\mathbb{E}_x[\delta\mathcal{L}'] \approx K \cdot \sqrt{\frac{1 + \frac{2}{M} \sum_{i=1}^M \sum_{j=1}^{i-1} \frac{\text{Cov}(\mathbf{r}_i, \mathbf{r}_j)}{\text{Cov}(\mathbf{r}_i, \mathbf{r}_i)}}{M}}$$

constant wrt. ϵ

covariance

number of tasks

$\partial_x \mathcal{L}(x, y_i)$

average increase of task losses

Conclusion: Robustness increases more as

- more tasks are learnt
- tasks are less correlated.

Experimental settings

- Dataset: **Taskonomy**
- Tasks: Semantic Segmentation (**s**), Depth z-buffer Estimation (**d**), Depth euclidian Estimation (**D**), Surface Normal Prediction (**n**), SURF Keypoint Detection in 2D (**k**) and 3D (**K**), Canny Edge Detection (**e**), Edge Occlusion (**E**), Principal Curvature (**p**), Reshading (**r**) and Auto-Encoders (**A**)
- Models: ResNet18 + custom task decoder (same as Taskonomy paper)
- Metrics: cross-entropy (segmentation)
- Evasion attack: **PGD** (25 steps, step size = $2/255$, $e = 8/255$)


Is the robustness increase generalized?

Auxiliary →		s	d	D	n	E
Single	s	0.82	0.86	0.97	0.96	0.93
	d	5.74	5.61	5.28	6.88	6.41
	D	5.93	6.14	6.4	7.12	8.31
	n	7.43	9.48	8.93	10.82	9.08
	E	12.93	19.29	18.44	15.16	22.57


Adversarial vulnerability of a single main task after adding an auxiliary task (lower is better).

Diagonal elements represent models with a single task.

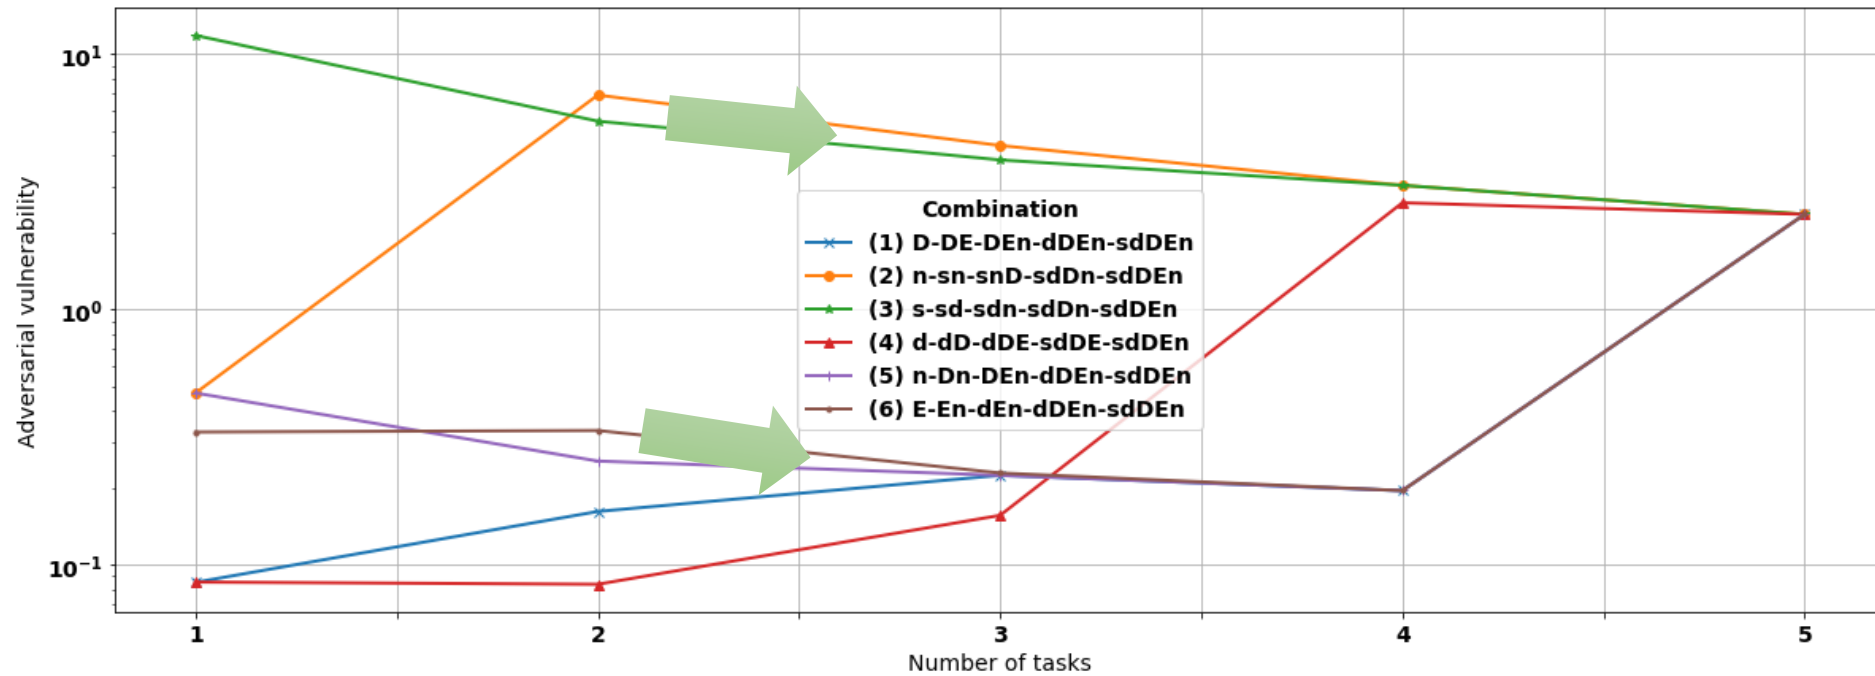
Marginal vulnerability of multi-task models



Marginal vulnerability of multi-task models



Marginal vulnerability of multi-task models



New theorem

vulnerability of the new task

$$\Delta_N \overbrace{\mathbb{E}_x[\delta \mathcal{L}']} \leq \epsilon \cdot (\mathbb{E}_x[\|\mathbf{r}_{N+1}\|] + \max_{i < N+1} \mathbb{E}_x[\|\mathbf{r}_i\|])$$

average loss increase
after adding task N+1

vulnerability of the
most vulnerable previous task

Conclusion: Robustness after adding a new task depends on

- the intrinsic vulnerability of the new task
- the most vulnerable previous task

Generalization to weighted tasks

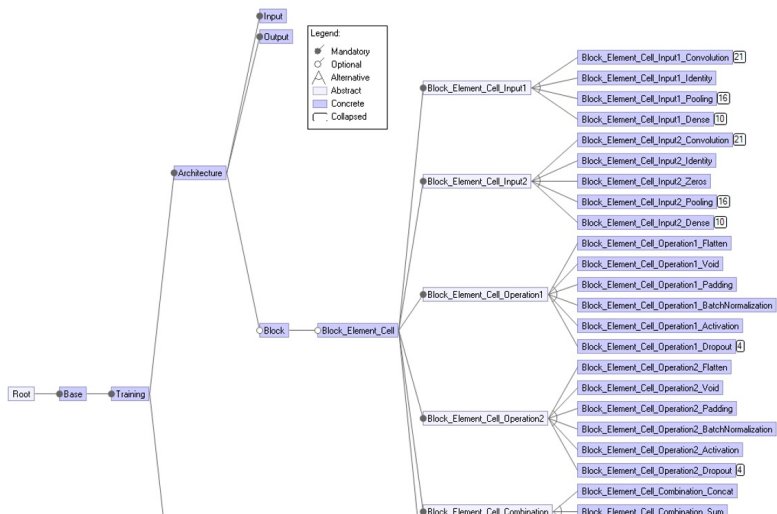
$$\Delta_N \widetilde{\mathbb{E}_x[\delta \mathcal{L}']} \leq \epsilon \cdot ((N + 1) \cdot w_{N+1} \mathbb{E}_x[\|\mathbf{r}_{N+1}\|] + N \cdot \max_{i < N+1} w_i \mathbb{E}_x[\|\mathbf{r}_i\|])$$

 learning weight of task i

Conclusion: appropriate task weighting can increase robustness!


Conclusion

The corresponding variability model

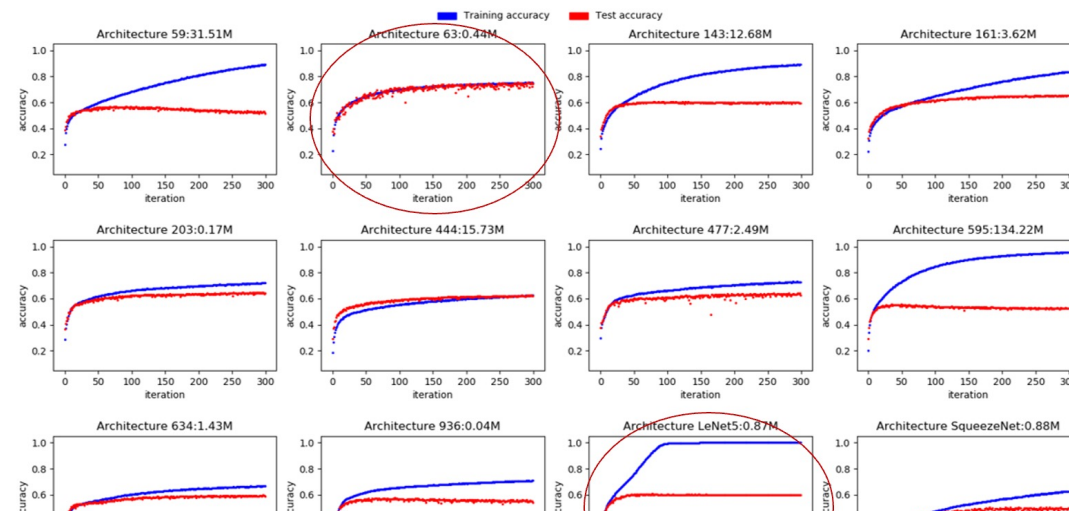


Reusability, adaptability, and exploration (in ML)

Multi-task learning



Variability in performance



Is the robustness increase generalized?

Auxiliary →	s	d	D	n	E
s	0.82	0.86	0.97	0.96	0.93
d	5.74	5.61	5.28	6.88	6.41
D	5.93	6.14	6.4	7.12	8.31
n	7.43	9.48	8.93	10.82	9.08
E	12.93	19.29	18.44	15.16	22.57

Adversarial vulnerability of a single main task after adding an auxiliary task (lower is better).

Diagonal elements represent models with a single task.

THANK YOU