

# Minimizing Compute Costs: When Should We Run More Expensive Malware Analysis?

**Andre T. Nguyen**, Richard Zak, Luke E. Richards, Maya Fuchs, Fred Lu, Robert Brandon,  
Gary Lopez Munoz, Edward Raff, Charles Nicholas, James Holt



Booz | Allen | Hamilton

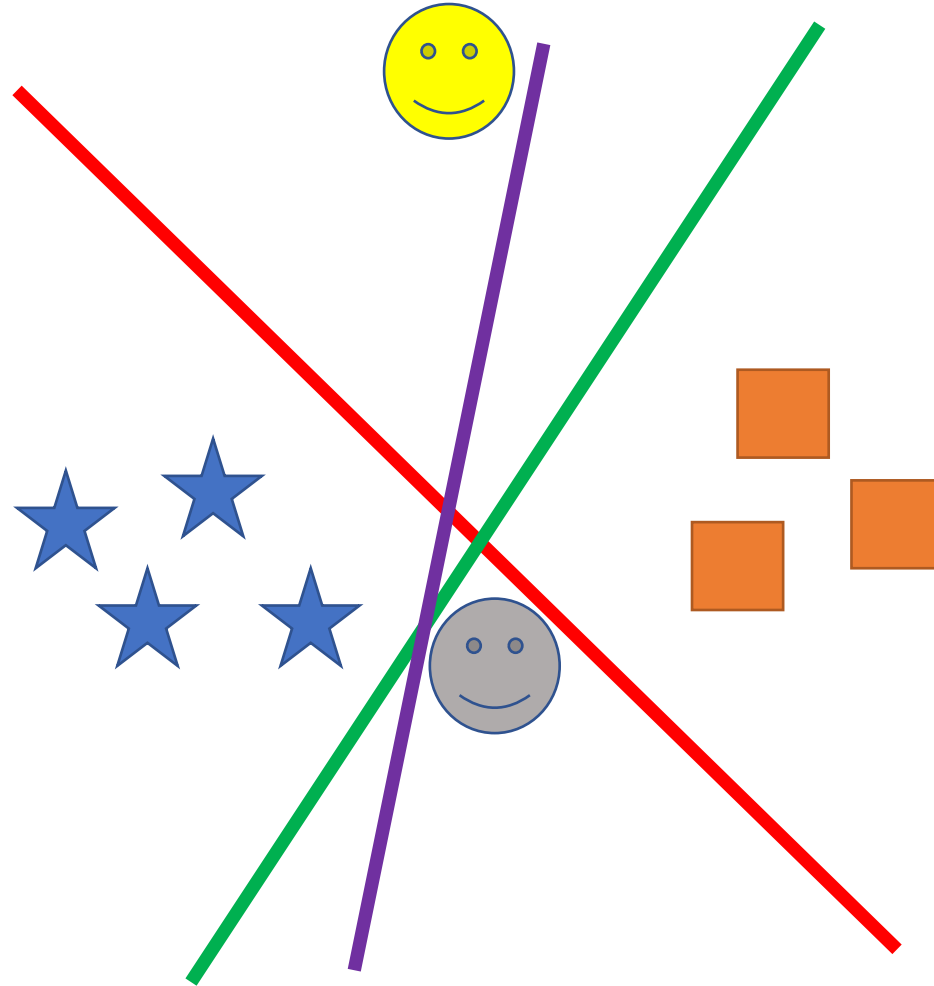
# Research Goals

Existing machine learning based approaches to malware detection have not yet leveraged uncertainty in a systematic manner.

Cyber security intrinsically requires operating under uncertain conditions, so uncertainty should not be ignored.

# Why can't we just use the softmax output probabilities?

[0.99, 0.01]  
[0.01, 0.99]  
[0.60, 0.40]



# The Quantification of Uncertainty

Suppose a cat/dog classifier trained on zoomed out images of dogs and cats...



## Epistemic Uncertainty

Uncertainty due to a lack of similar data.

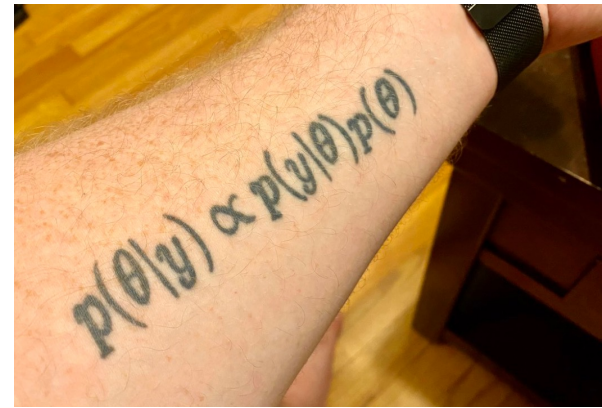
If you had more cat/dog training images focused on tails, this would be easier.



## Aleatoric Uncertainty

Inherently confusing because both classes are present.

# Bayes Rule in the context of ML



Posterior belief about model parameters given the data.

Likelihood of the data given the model parameters.

Prior belief about model parameters.

$$P(\theta|D, M) = \frac{P(D|\theta, M)P(\theta|M)}{P(D|M)}$$

Model Parameters    Data    Model Class

Likelihood of the data given the model class, with model parameters having been integrated out. Useful for model comparison.

# Bayesian inference operates on distributions to capture beliefs and uncertainty.




Figure 1. Left: each weight has a fixed value, as provided by classical backpropagation. Right: each weight is assigned a distribution, as provided by Bayes by Backprop.




Point estimates.


The proper quantification of uncertainty using distributions.

Usually you begin by selecting a model class. Then, suppose you were asked to draw what you think the right model is, without observing any data...

$$\kappa = \exp\left(\frac{-\|x-x'\|^2}{2l^2}\right)$$



$$\kappa = \min(x, x')$$



$$\kappa = (x^T x' + c)^2$$


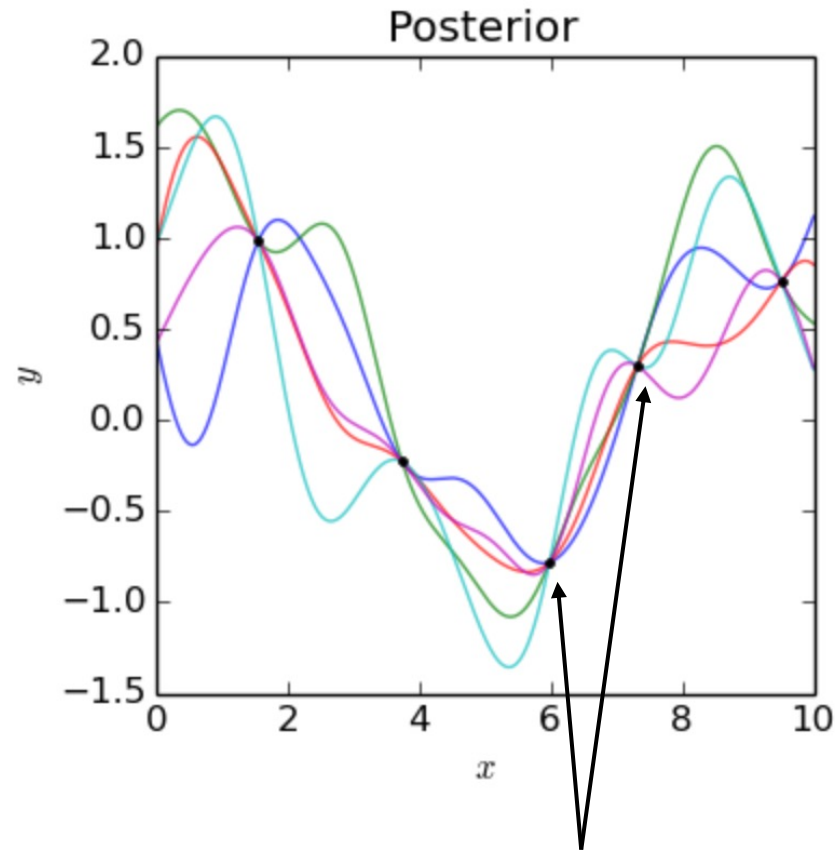



Image from:


[https://en.wikipedia.org/wiki/Gaussian\\_process](https://en.wikipedia.org/wiki/Gaussian_process)



Once you start observing data, the possible functions you can draw become constrained.



Data




Uncertainty is higher when far from the data.

Image from:

[https://en.wikipedia.org/wiki/Gaussian\\_process](https://en.wikipedia.org/wiki/Gaussian_process)





# Synergy

Uncertainty + Machine Learning Based Malware Detection



# Model Example: MalConv

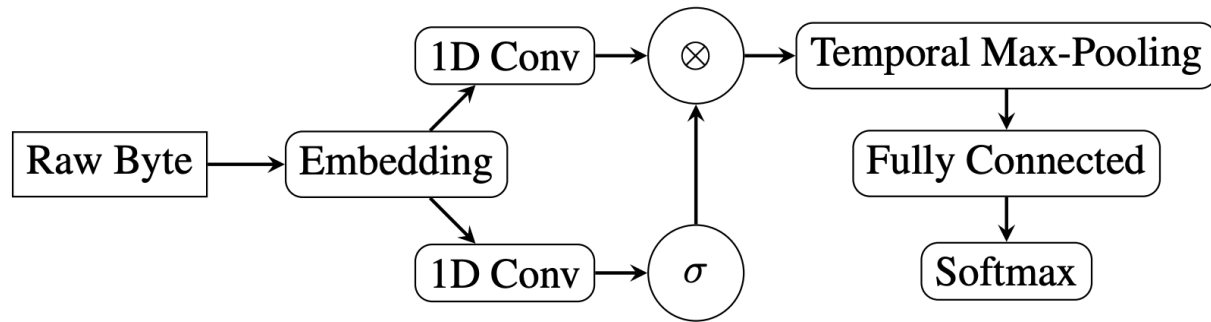



Figure 1. Architecture diagram of MalConv model.



# Bayesian Dropout



- Variational inference approach for Bayesian deep learning.
- <https://arxiv.org/pdf/1506.02142.pdf>
- Probably the easiest method to implement and deploy.
- Essentially, add dropout to all layers, and leave dropout on during prediction. A sample is run through the model multiple times to generate the predictive distribution.
- Getting well calibrated uncertainties is a bit trickier as dropout probabilities need to be tuned: <https://arxiv.org/pdf/1705.07832.pdf>

# Dropout




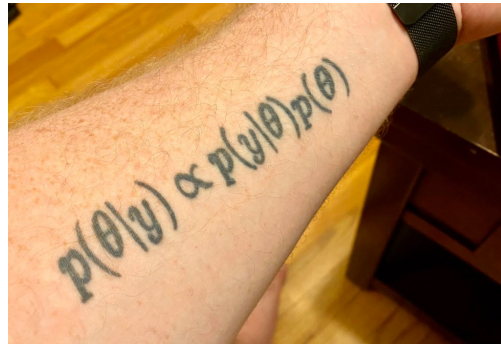
Figure 1: Dropout Neural Net Model. **Left:** A standard neural net with 2 hidden layers. **Right:** An example of a thinned net produced by applying dropout to the network on the left. Crossed units have been dropped.

<https://www.cs.toronto.edu/~hinton/absps/JMLRdropout.pdf>

# Bayesian MalConv (MalBayes)



MalConv



Dropout as a Bayesian Approximation



MalBayes



# Alternatively, we can ensemble...




Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.




Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.




Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.




Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.




Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.




Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.

Figure 1. Architecture diagram of MalConv model.





# Leveraging Uncertainty for Improved Static Malware Detection Under Extreme False Positive Constraints

Workshop on Adaptive Cyber Defense at IJCAI 2021.





# Uncertainty on Errors and New AV Classes





(a) EMBER, Bayesian MalConv




(a) Bayesian MalConv





(b) EMBER, Bayesian Logistic Regression



(b) Bayesian Logistic Regression



(c) EMBER, LGBM Ensemble



(c) LGBM

# Out of Distribution Data Detection Using Dropout Bayesian Neural Networks

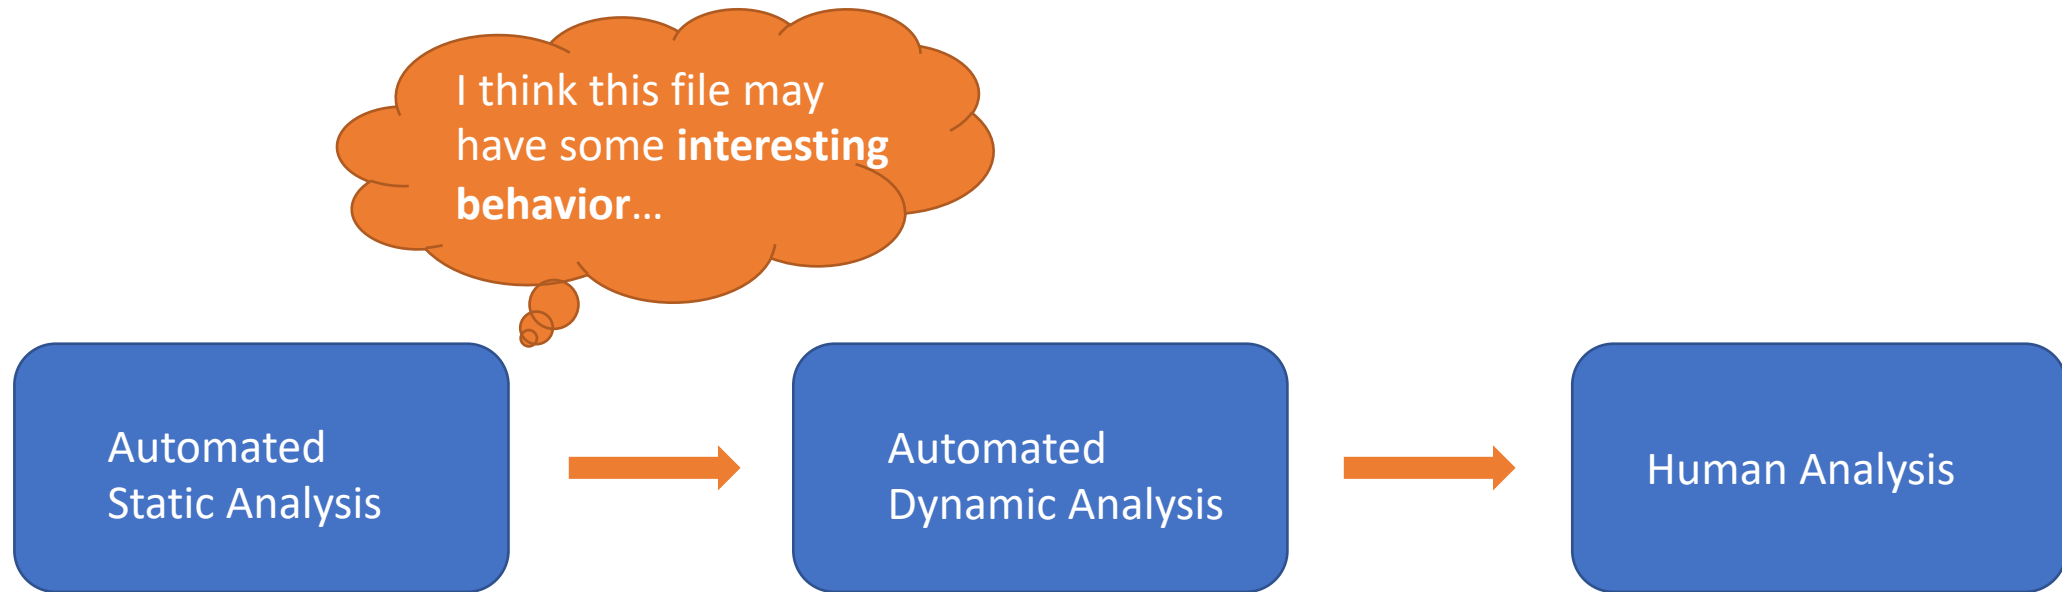
AAAI 2022.

OOD		Num/Class	n=100		n=50		n=25	
Experiment	Model	Metric Features	AUC	Recall	AUC	Recall	AUC	Recall
EMBER2018	LR	Last	0.789	0.704	<b>0.786</b>	0.682	<b>0.778</b>	0.650
		Last+Spread	<b>0.793</b>	<b>0.718</b>	0.783	<b>0.689</b>	0.766	<b>0.658</b>
	RF	Last	0.757	0.735	0.752	0.727	0.748	0.714
		Last+Spread	<b>0.791</b>	<b>0.784</b>	<b>0.782</b>	<b>0.764</b>	<b>0.770</b>	<b>0.743</b>
Brazilian	LR	Last	0.685	<b>0.645</b>	0.680	0.607	0.668	0.584
		Last+Spread	<b>0.741</b>	0.620	<b>0.734</b>	<b>0.617</b>	<b>0.712</b>	<b>0.605</b>
	RF	Last	0.724	0.693	0.705	0.674	0.679	0.652
		Last+Spread	<b>0.839</b>	<b>0.797</b>	<b>0.813</b>	<b>0.772</b>	<b>0.776</b>	<b>0.736</b>

# Minimizing compute costs: When should we run more expensive analysis?

CAMLIS 2022.

# Can we optimize the decision process?




# CAPA Rules

```
1 rule:
2   meta:
3     name: capture webcam image
4     namespace: collection/webcam
5     author: johnk3r
6     scope: function
7     att&ck:
8       - Collection::Video Capture [T1125]
9     examples:
10      - a30101595f6f28ab2f4b0b2cd177c3c4d2ab34a355ab7761a3795d0887c24ada:0x4011C0
11   features:
12     - or:
13       - and:
14         - api: capCreateCaptureWindow
15         - basic block:
16           - and:
17             - api: SendMessage
18             - number: 0x40a = WM_CAP_DRIVER_CONNECT
19         - optional:
20           - basic block:
21             - and:
22               - api: SendMessage
23               - number: 0x40B = WM_CAP_DRIVER_DISCONNECT
24         - basic block:
25           - and:
26             - api: SendMessage
27             - number: 0x419 = WM_CAP_FILE_SAVEDIB
```


```
1 rule:
2   meta:
3     name: encrypt data using Curve25519
4     namespace: data-manipulation/encryption/elliptic-curve
5     author: dimiter.andonov@mandiant.com
6     scope: basic block
7     att&ck:
8       - Defense Evasion::Obfuscated Files or Information [T1027]
9     examples:
10      - 0a0882b8da225406cc838991b5f67d11:0x4135f6
11      - 0a0882b8da225406cc838991b5f67d11:0x416f51
12      - 80372de850597bd9e7e021a94f13f0a1:0x406480
13      - 80372de850597bd9e7e021a94f13f0a1:0x4086f4
14   features:
15     # initializes a 32-byte array with
16     # array[0] = 0xf8,
17     # array[31] = array[31] & 0x3f | 0x40
18     - and:
19       - and:
20         - number: 0xf8
21         - mnemonic: and
22       - and:
23         - number: 0x3f
24         - mnemonic: and
25     - and:
26       - number: 0x40
27       - mnemonic: or
```

# Predicting CAPA Outputs

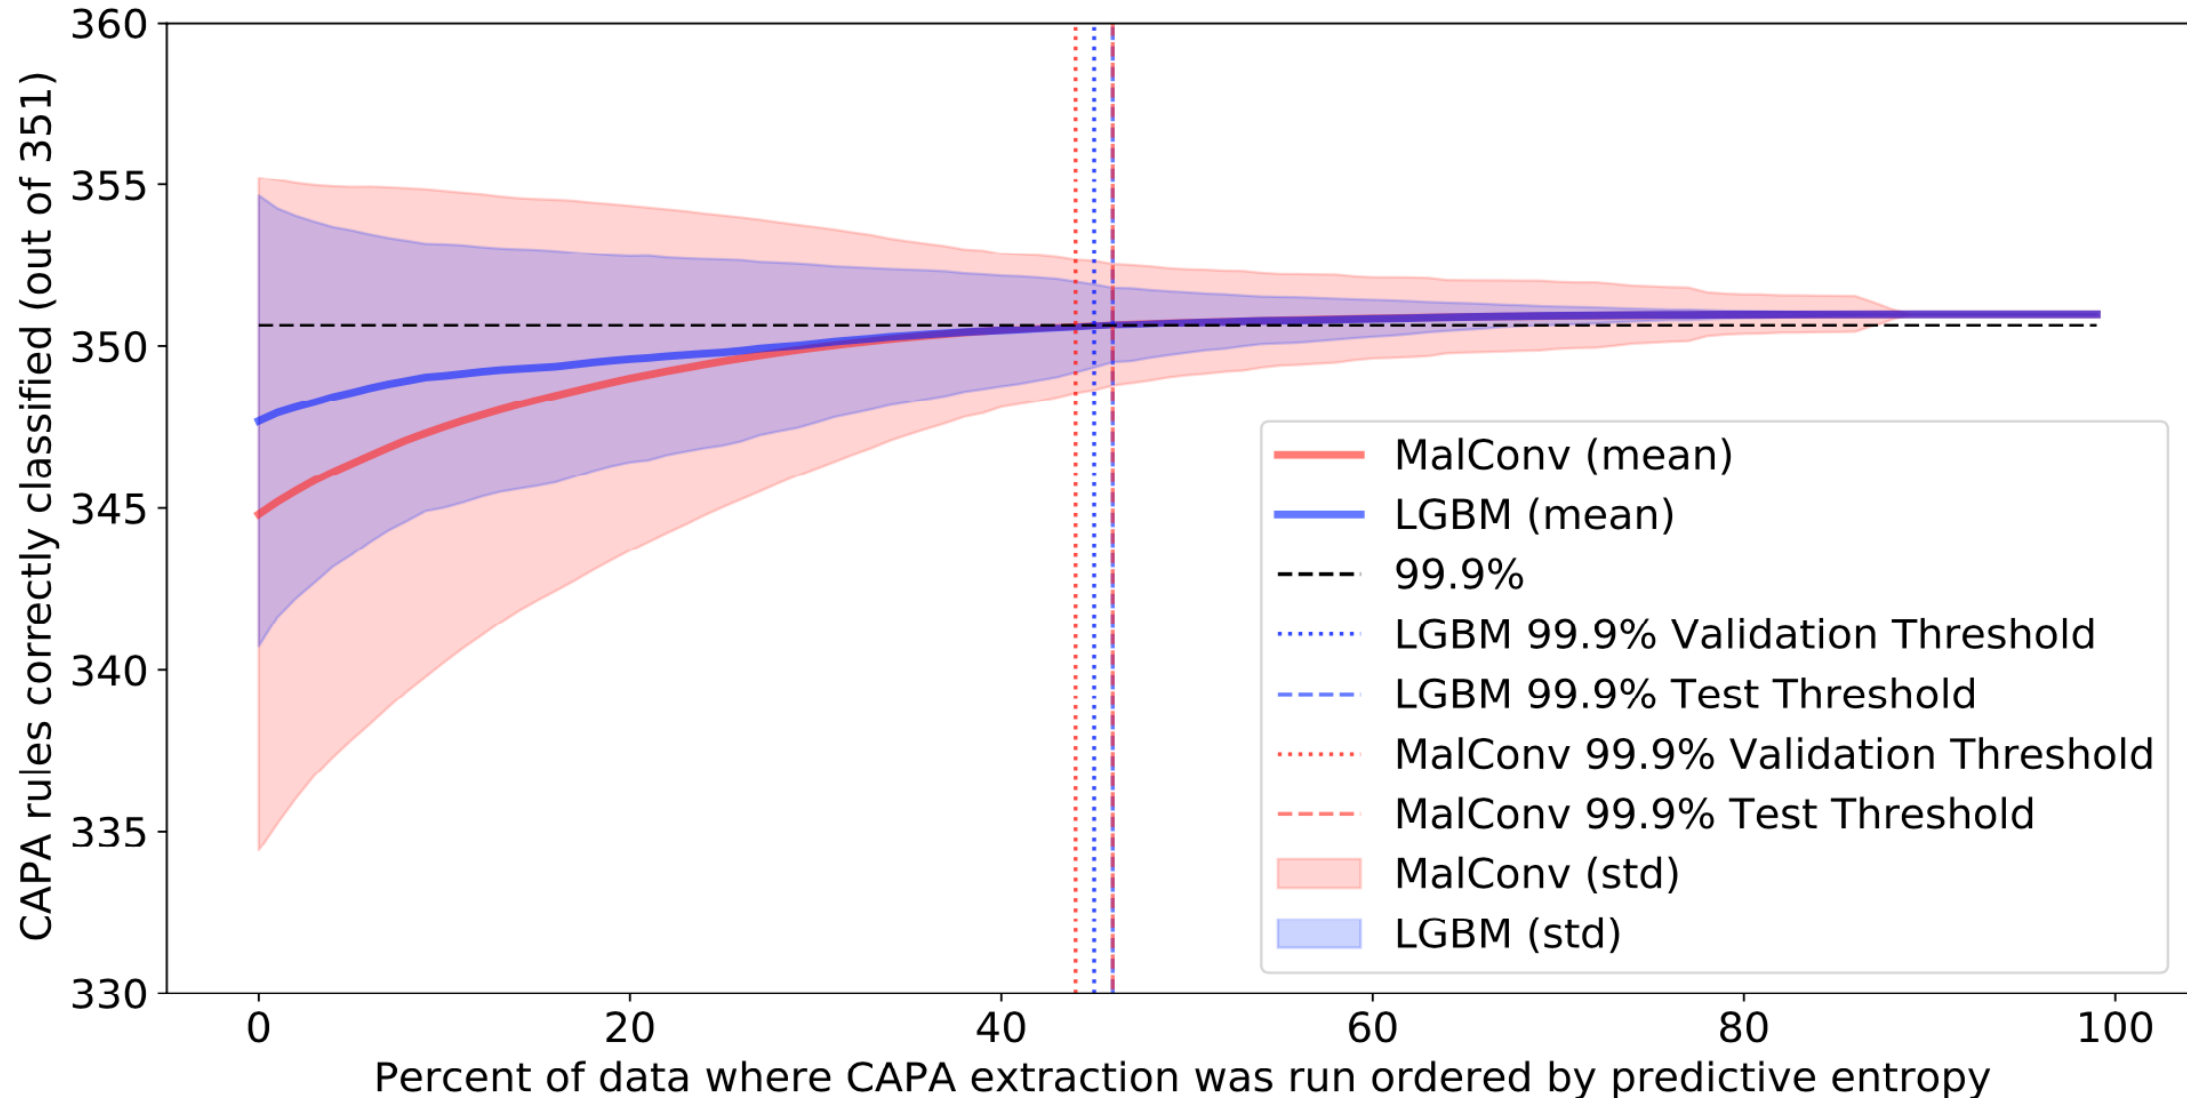
CAPA Prediction Accuracies By Rule



CAPA Prediction AUCs By Rule

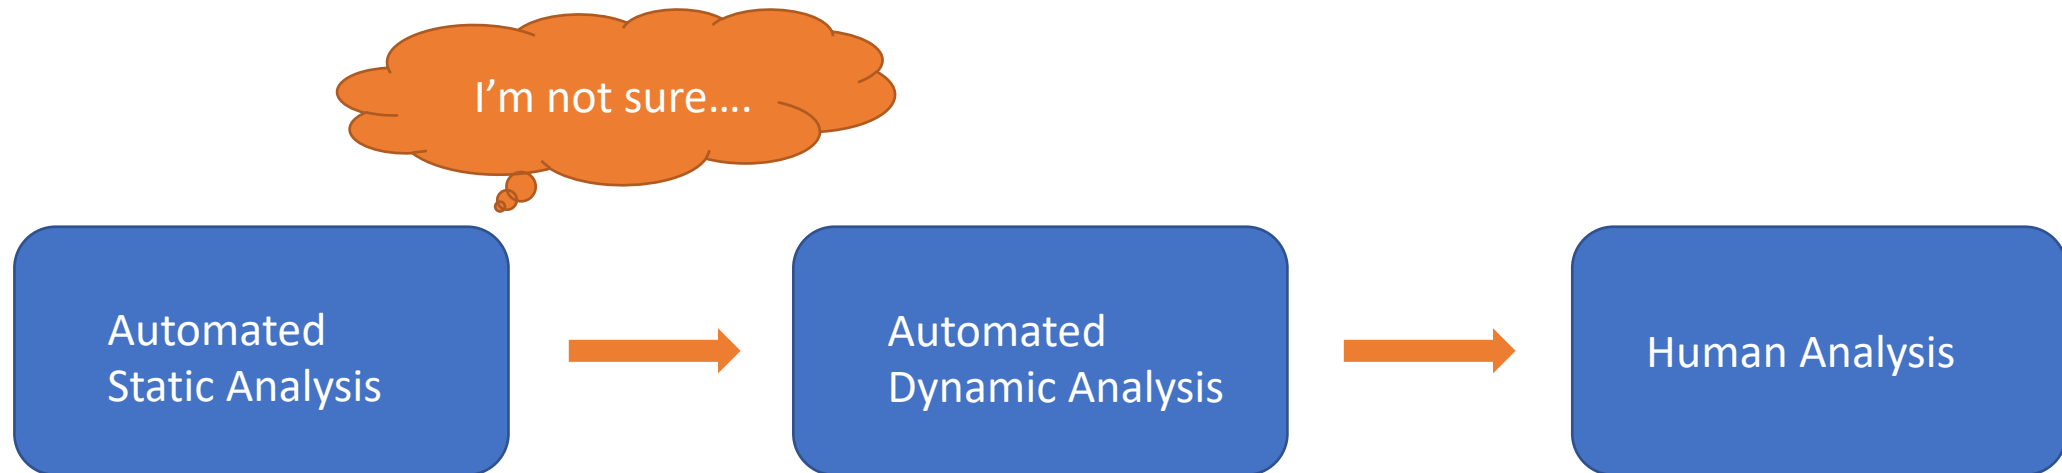


# How much expensive CAPA analysis do we need?





# Can we optimize the decision process?



Bayesian MalConv with sampling: 0.02 seconds per file

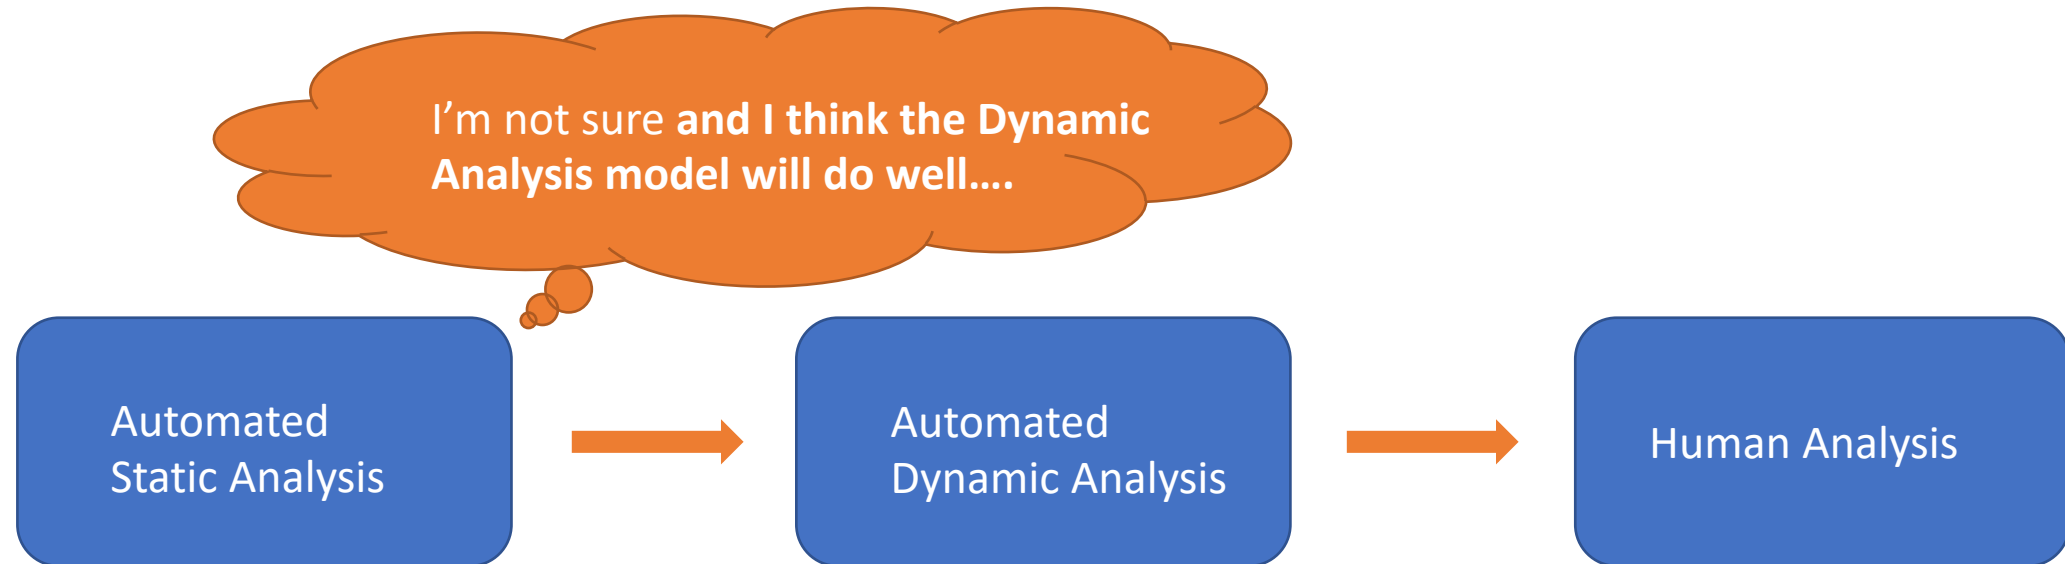
EMBER Feature extraction: 0.09 seconds per file

CAPA feature extraction: 45.75 seconds per file

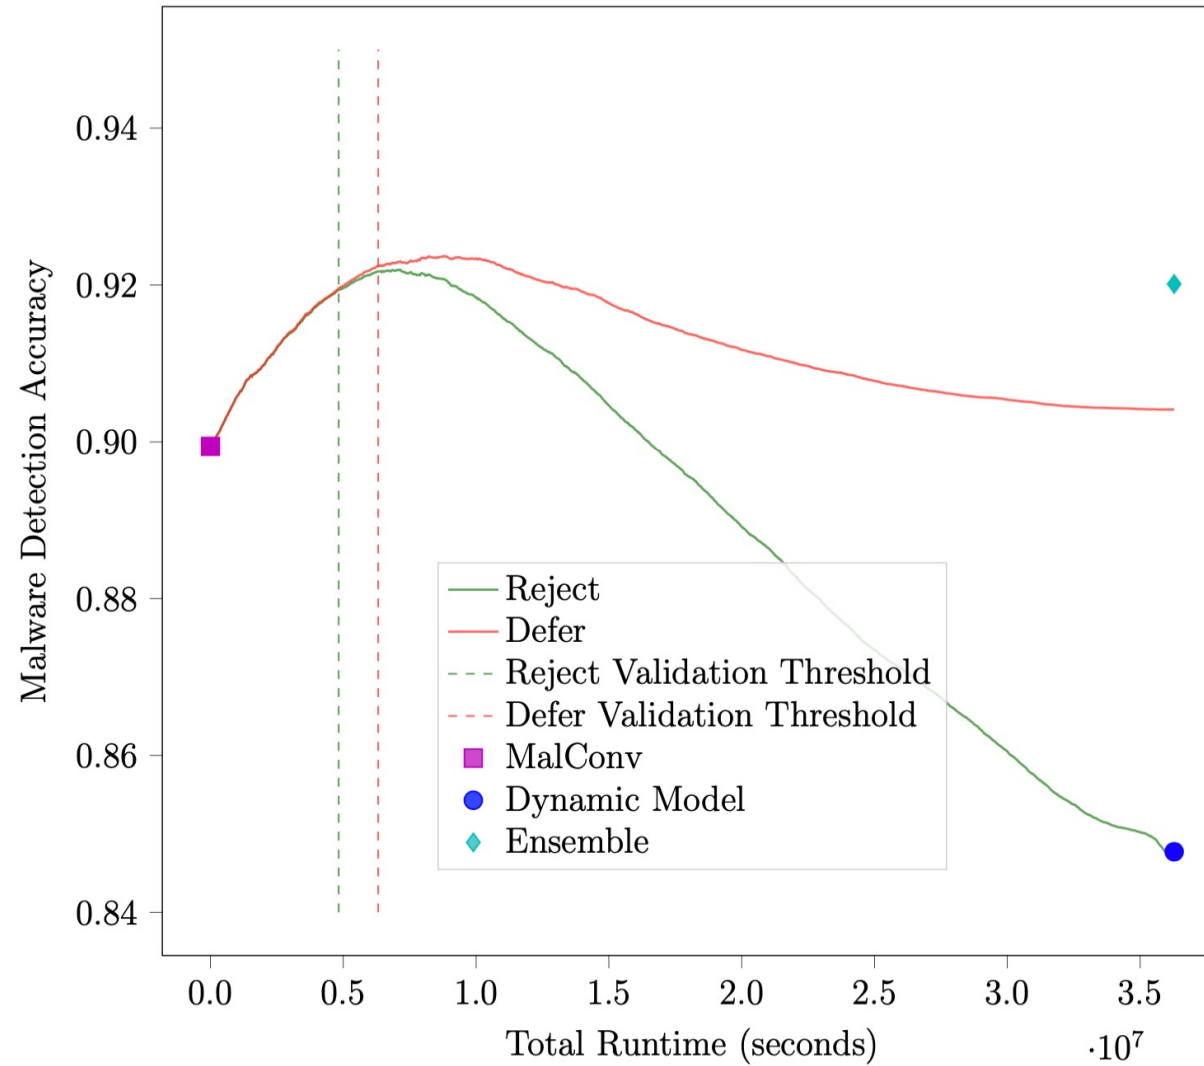
Running dynamic analysis: 526 seconds per file

Running Bayesian MalConv on a file is over 26,300 times faster than running dynamic analysis!

# Can we optimize the decision process?



# MalConv -> Dynamic



# Thanks!

**André T. Nguyen, Ph.D.**  
Director of Machine Learning

JURA Bio, Inc.  
Boston, MA  
[an@jurabio.com](mailto:an@jurabio.com)

