Squinting Vision Pipelines: Detecting and Correcting Errors in Vision Models at Runtime

Ken Wenger

Chief Technology Officer

Squint AI Inc

www.squintinc.com

SQUINT AI™

Squint AI Inc

Humans squint when we are not sure.

AI CAN SQUINT TOO.

SQUINT AI™

Squint AI Inc

What problems are we solving?

Fundamental question

Can we trust the predictions of AI models ?

Model expression

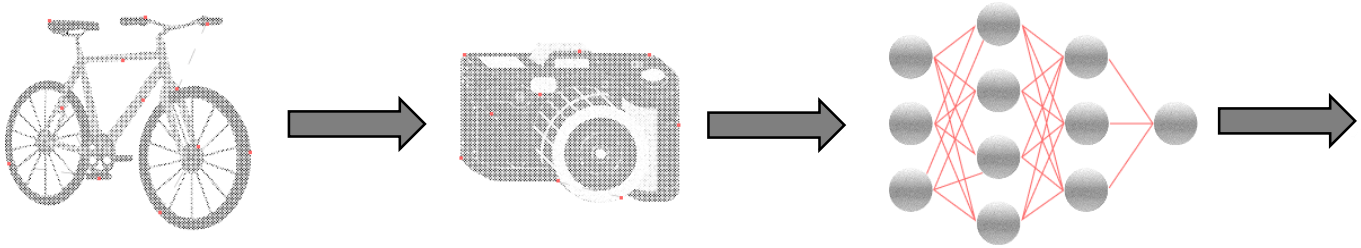

Challenges we must address

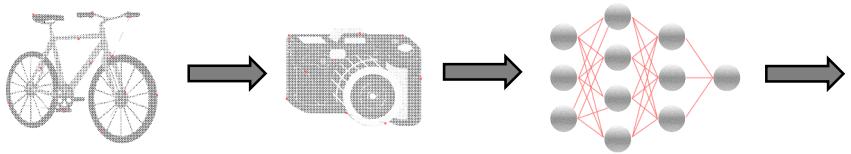

- Models often make mistakes.
- It is difficult to assess an accurate expectation of performance in a production setting.
- It is difficult to assess what the models are learning.

Models often make mistakes

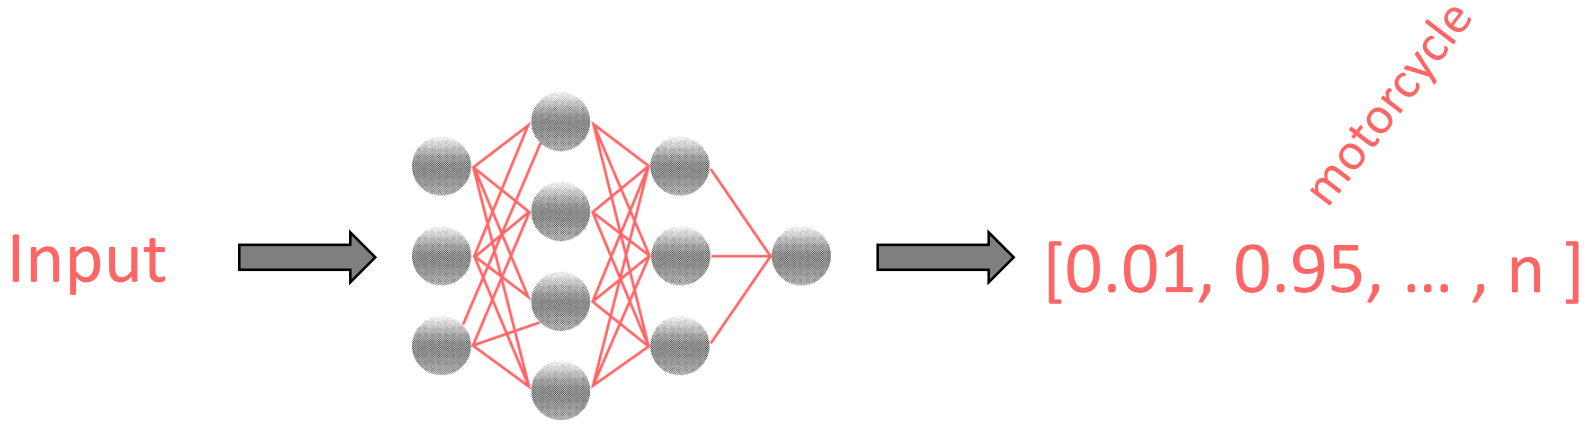

Challenge addressing performance expectation

What can cause models to get it wrong?

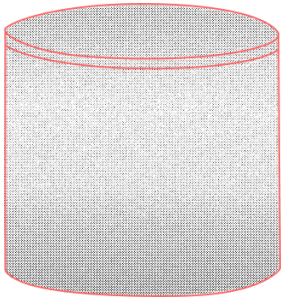

Dataset

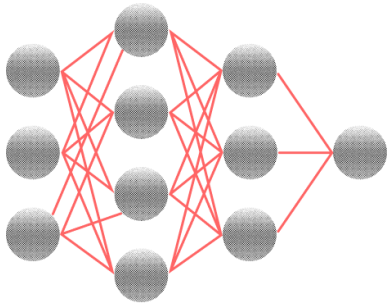

Model

Addressing dataset problems

Why is it difficult to get the dataset right?

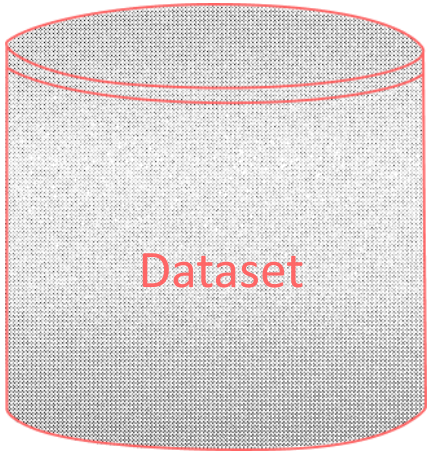

ENTROPY

The variance that each unit of information can experience (before information is completely lost).

COMPLEXITY

A measure of how many categories and data samples are required to accurately describe the operational domain.

AMBIGUITY

How much overlap exists between the datapoints in the different categories.

Consequence of a poor dataset on performance expectation

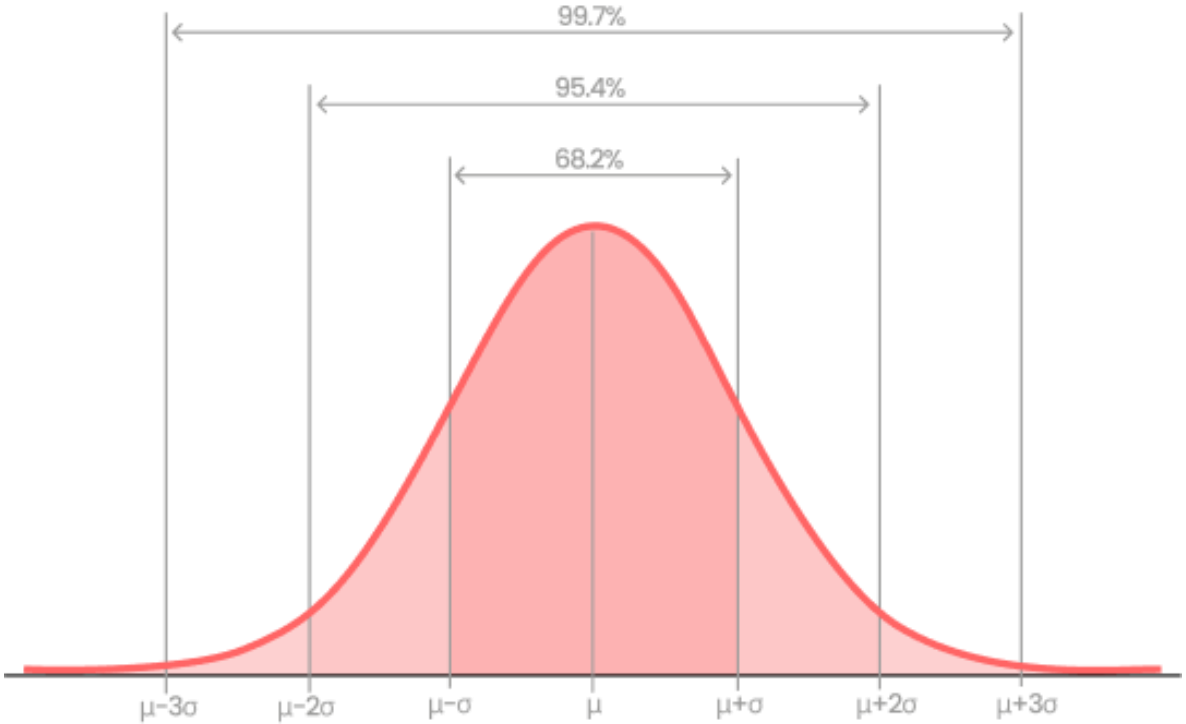

SQUINT AI™

Bounding trust

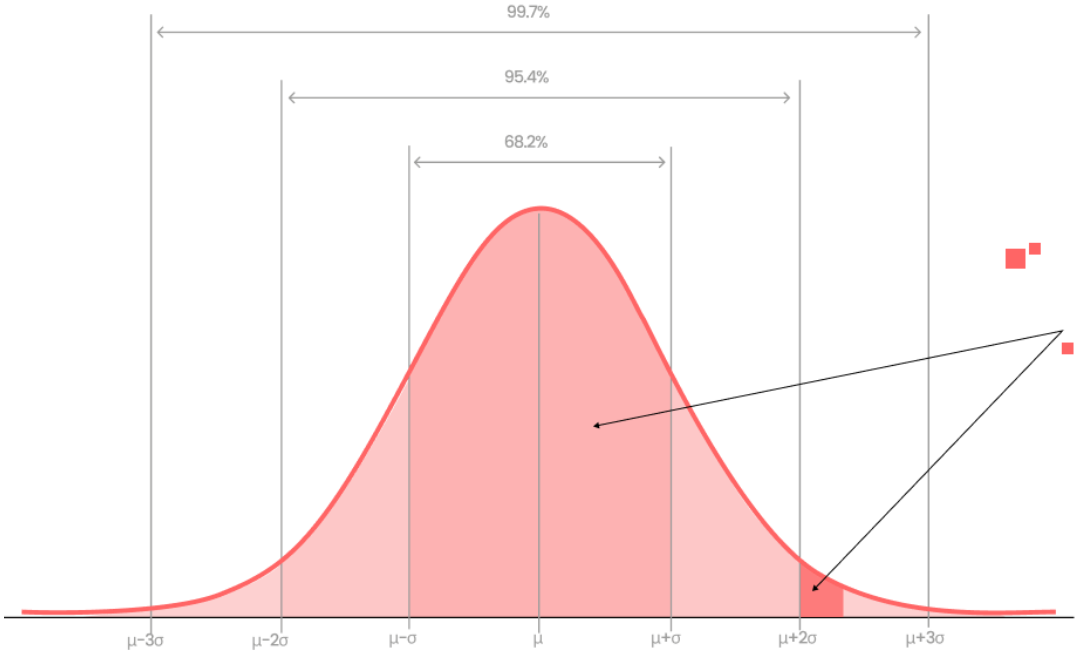

Can we detect this for each prediction?

Explainable AI / Slice discovery

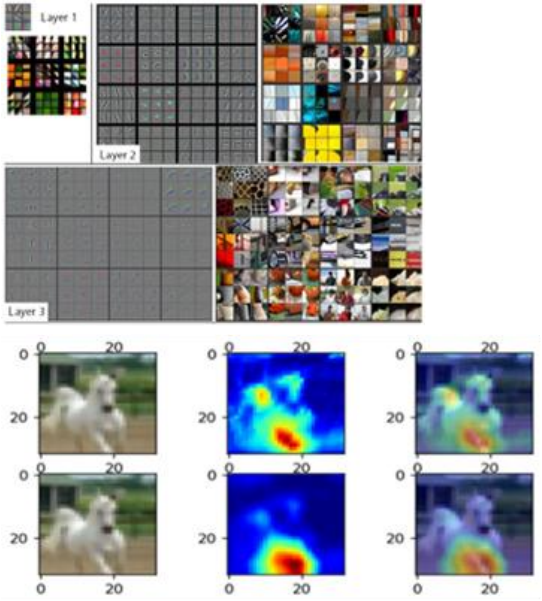

[A novel application of XAI in squinting models: A position paper - ScienceDirect]

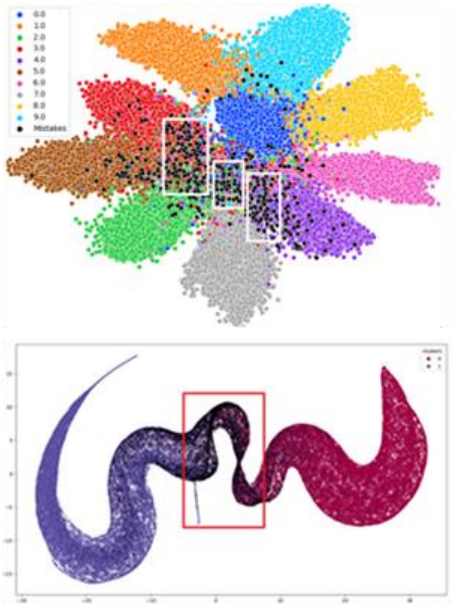

[Visualization and Model Explanations in Convolutional Neural Networks | by Kumar Devesh | Medium]

Explainable AI (XAI) is not only useful during model development.

Explainable AI can be used to build runtime watchdogs to actively guarantee the integrity of the AI application.

SQUINT AI™

Bounding confidence

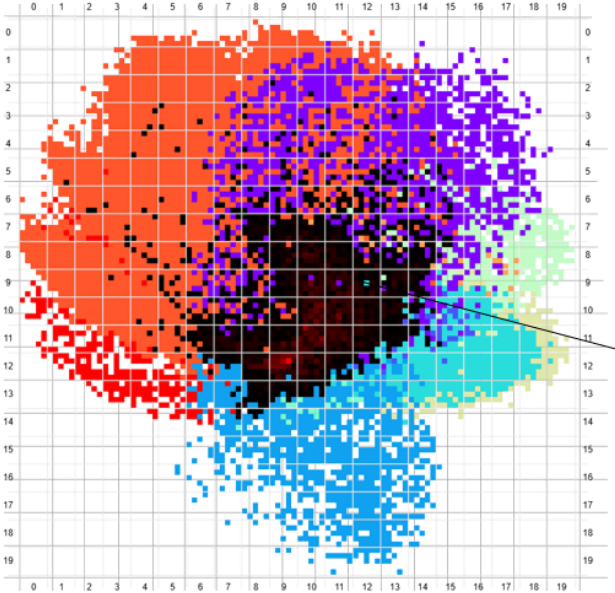

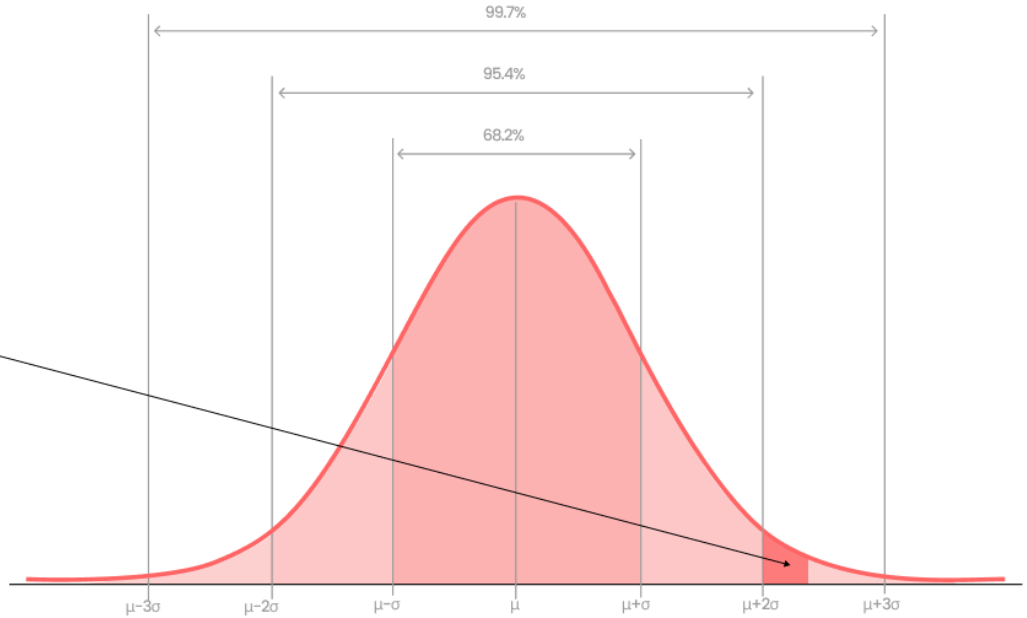

Adding contextual information

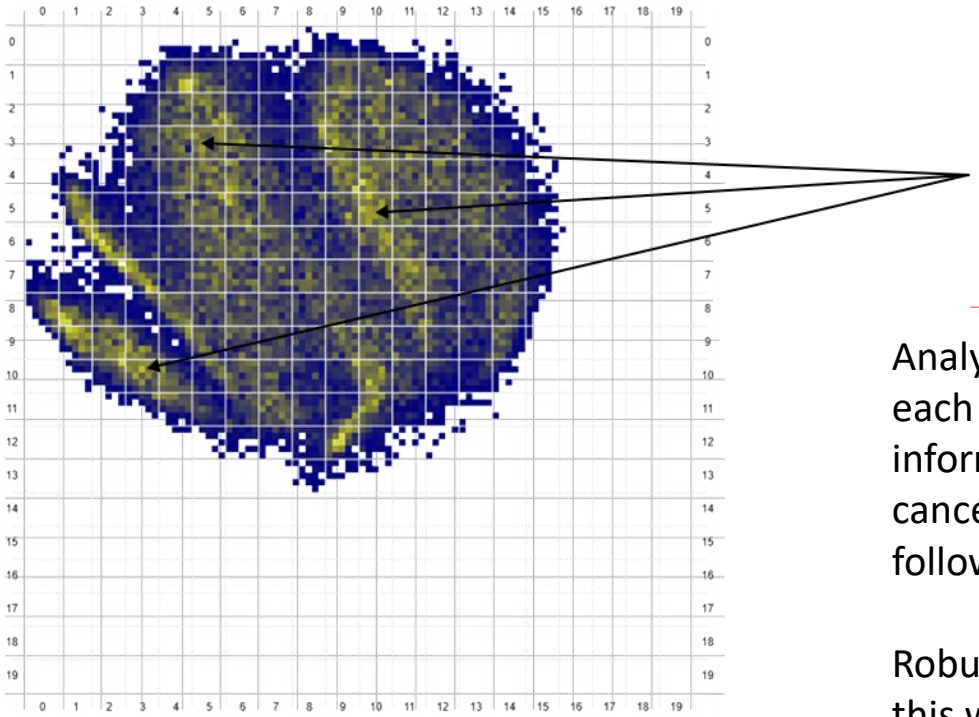

Different regions within a cluster contain different semantic information

Analyzing the semantic information within each cluster can provide contextual information. E.g., this isn't just "grade 1 cancer" this is "grade 1 cancer" with the following cell morphology…

Robust, context-aware pipelines can be built this way.

SQUINT AI™

Bounding a robust pipeline

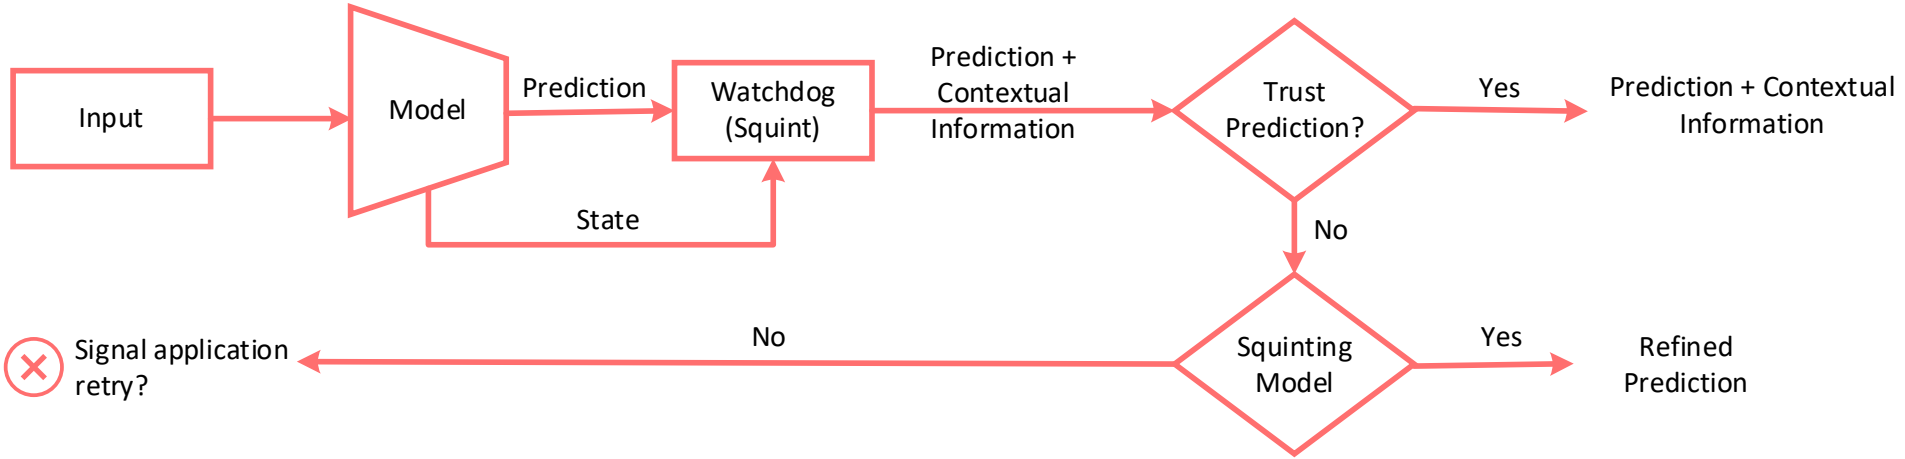

Summary

- Models often make mistakes.
  - Geofence the ambiguous region

- It is difficult to assess an accurate <u>expectation</u> of performance in a production setting.
  - Bound performance to trusted regions

- It is difficult to assess what the models are learning.
  - Use XAI, input saliency, semantic clustering analysis, etc.

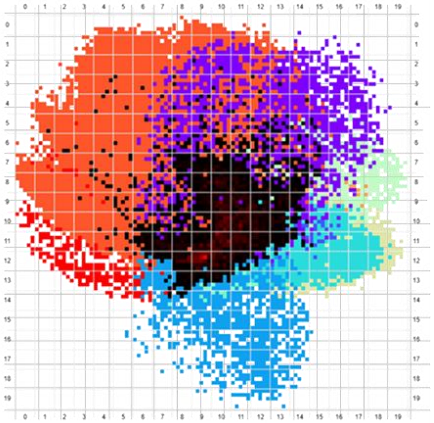

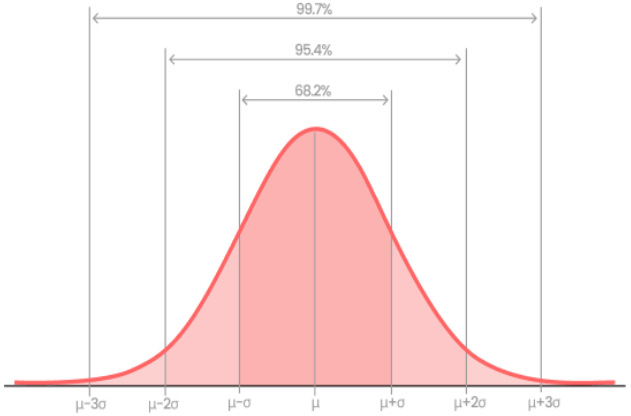

Results – Squint pipeline with human in the loop

Breast Cancer Experiment:

- Dataset: 126,056 images (62,901 pos, 63,155 neg)
- ResNet-based classifier with state of the art 89.37% accuracy

Baseline model	Squint watchdog + Human in the loop	Squint watchdog + Squinting model (Vision Transformer)
10.63% error rate	**2.02% error rate**	**1.79% error rate**
89.38% accuracy	**97.98% accuracy**	**98.21% accuracy**

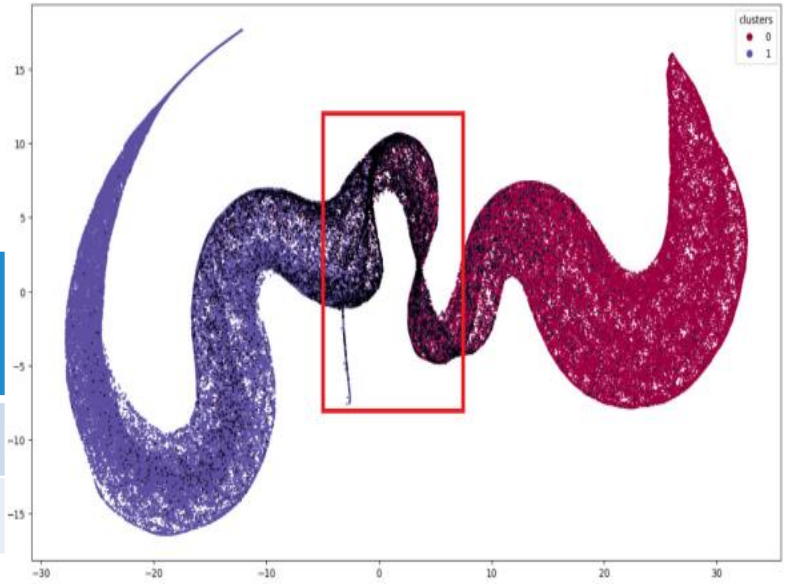

Squint AI Inc

How do we build this framework today?

Booth: 619 (Come see us!)

SQUINT INSIGHTS STUDIO™

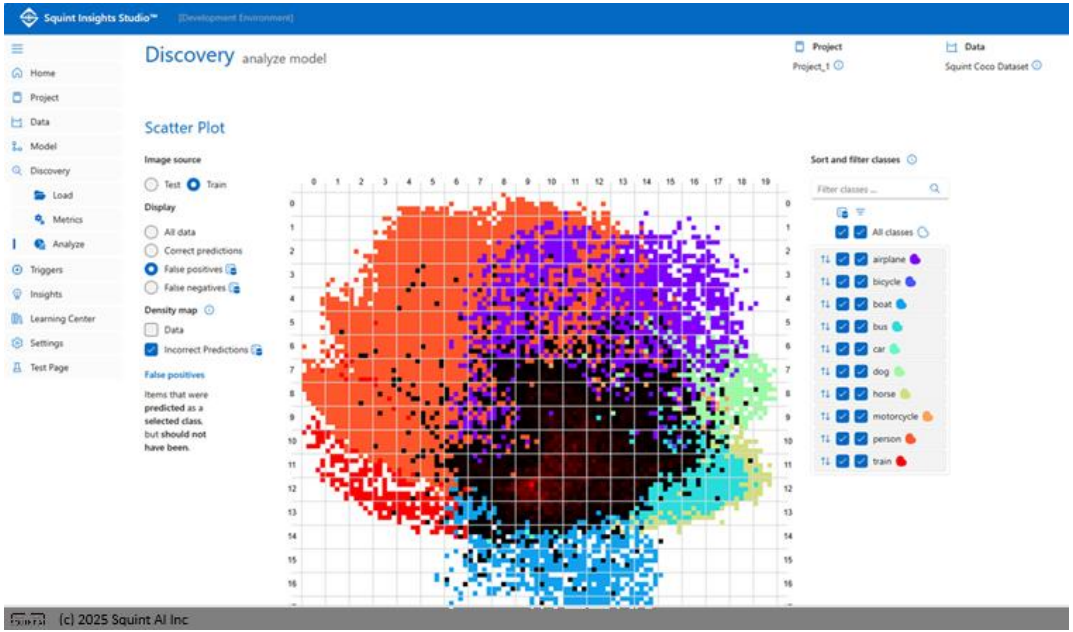

SQUINT AI™

Squint AI’s explainable AI paper:
A novel application of XAI in squinting models

Thank you!

www.squintinc.com

SQUINT AI™