

Continual Machine Learning

Summer 2023

Teacher

Dr. Martin Mundt,

hessian.AI-DEPTH junior research group leader on Open World Lifelong Learning (OWLL)

& researcher in the Artificial Intelligence and Machine Learning (AIML) group at TU Darmstadt


Time

Every Friday 14:25 - 16:05 CEST

Course Homepage

http://owll-lab.com/teaching/cl_lecture_23

<https://www.youtube.com/playlist?list=PLm6QXeaB-XkA5-IVBB-h7XeYzFzgSh6sk>




Continual **AI**



hessian.AI



TECHNISCHE
UNIVERSITÄT
DARMSTADT





Week 8: Open world learning - learning & prediction in the presence of the unknown

Recall sequences so far




We've discussed various ways to measure + assumptions, but so far it was always clear what to test on



Recall: the tasks we considered



What if we don't know the boundary & aren't constrained on our testing examples?



Recall: the tasks we considered



What if we don't know the boundary & aren't constrained on our testing examples?

What if future or unrelated data is in the test set?





Figure 1: Schematic of split MNIST task protocol.

Recall: distribution shifts



--- Ideal reproducibility ● Model accuracy — Linear fit


Recall: natural data distributions are complex & can easily shift!

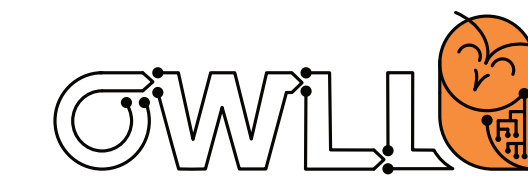
Performance loss even happens if we recollect another “test set” with the same instructions a second time!

Recall: noisy oracles

Recall our active learning assumptions:

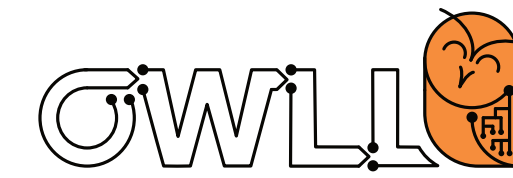
- *Oracle is infallible:*
the teacher/labeler does not make mistakes!
- *Pool belongs to task:*
we will cover this in our lecture on “learning and the unknown”





Perspectives to address these challenges

More than known vs. unknown



1. Known knowns:

Do you have an intuition what these 4 categories could represent?

2. Known unknowns:

3. Unknown unknowns:

4. Unknown knowns:

More than known vs. unknown



1. **Known knowns:**

Examples belong to the distribution from training set was drawn. Assumption of an accurate & confident prediction.

2. **Known unknowns:**

3. **Unknown unknowns:**

4. **Unknown knowns:**

More than known vs. unknown



1. **Known knowns:**

Examples belong to the distribution from training set was drawn. Assumption of an accurate & confident prediction.

2. **Known unknowns:**

Unknown examples where models are not confident or uncertainty is high. Can be optionally “negatively” labelled examples used in training.

3. **Unknown unknowns:**

4. **Unknown knowns:**

More than known vs. unknown



1. **Known knowns:**

Examples belong to the distribution from training set was drawn. Assumption of an accurate & confident prediction.

2. **Known unknowns:**

Unknown examples where models are not confident or uncertainty is high. Can be optionally “negatively” labelled examples used in training.

3. **Unknown unknowns:**

Unseen instances belonging to unexplored & unknown data distributions. Predictions generally overconfident & by definition false.

4. **Unknown knowns:**

More than known vs. unknown



1. **Known knowns (or simply knowns):**

Examples belong to the distribution from training set was drawn. Assumption of an accurate & confident prediction.

2. **Known unknowns:**

Unknown examples where models are not confident or uncertainty is high. Can be optionally “negatively” labelled examples used in training.

3. **Unknown unknowns:**

Unseen instances belonging to unexplored & unknown data distributions. Predictions generally overconfident & by definition false.

4. **Unknown knowns:**

Usually not considered: we know the concept but choose to treat it as unknown (willful ignorance?) or our ML system cannot represent the concept + structure altogether

Three types of approaches



What do you think: how can we solve our challenge?

Three types of approaches



Anomalies in predictions:

The *unsuspecting angle*, where out-of-distribution are hopefully separable through anomalous output values.

Incorporating prior knowledge:

The *intuitive idea* to include “background” or “non-example” data population explicitly.

Open Set recognition:

The more *formal approach* ensures that we only rely on predictions from our “covered space”; we create bounds.

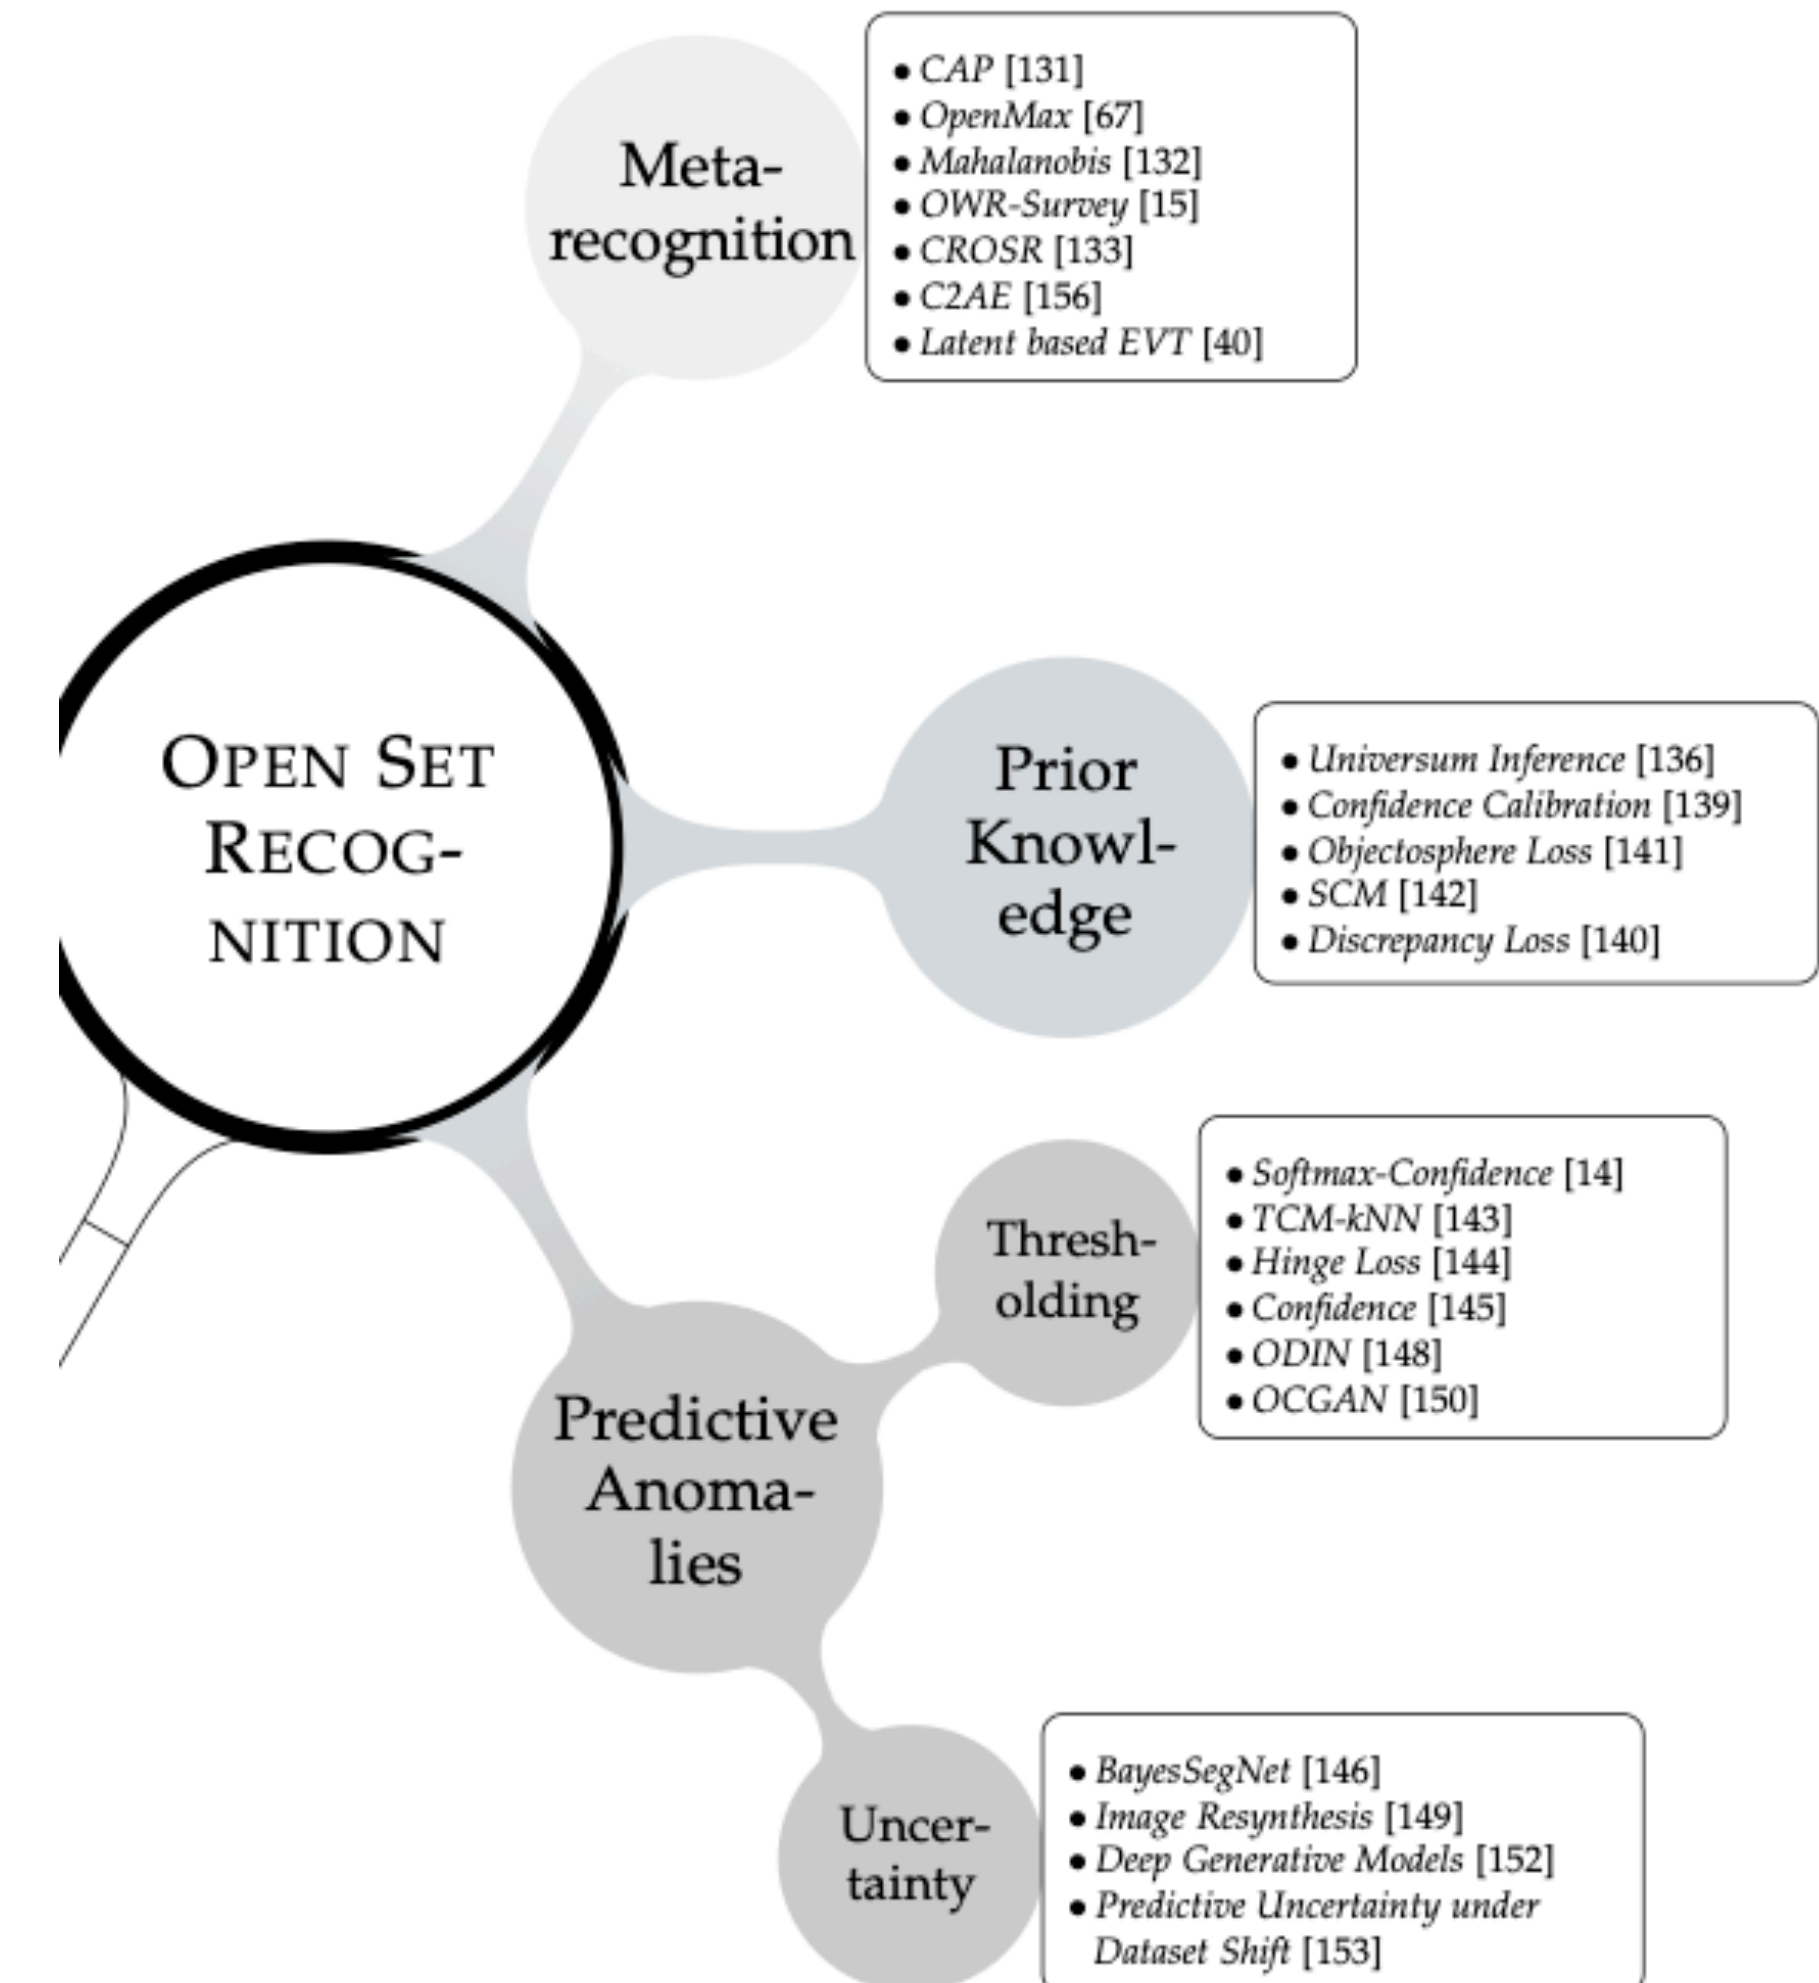



Figure from “A Wholistic View of Deep Neural Networks: Forgotten Lessons and the Bridge to Active and Open World Learning”, Mundt et al, Neural Networks, 2023



Predictive anomalies: the unfortunate part of the story

**Disclaimer: I'll use my many figures from our papers for convenience,
without trying to imply that we discovered these phenomena**

Recall lecture 1: overconfidence



Recall the quantitative example:


1. Train a neural network classifier on a dataset (here fashion items)
2. Log predictions for arbitrary other datasets
3. Observe that majority of misclassifications happen with large output “probability”

Overconfidence & uncertainty



Unfortunately uncertainty is not a necessarily a “fix”

Standard neural network classifier



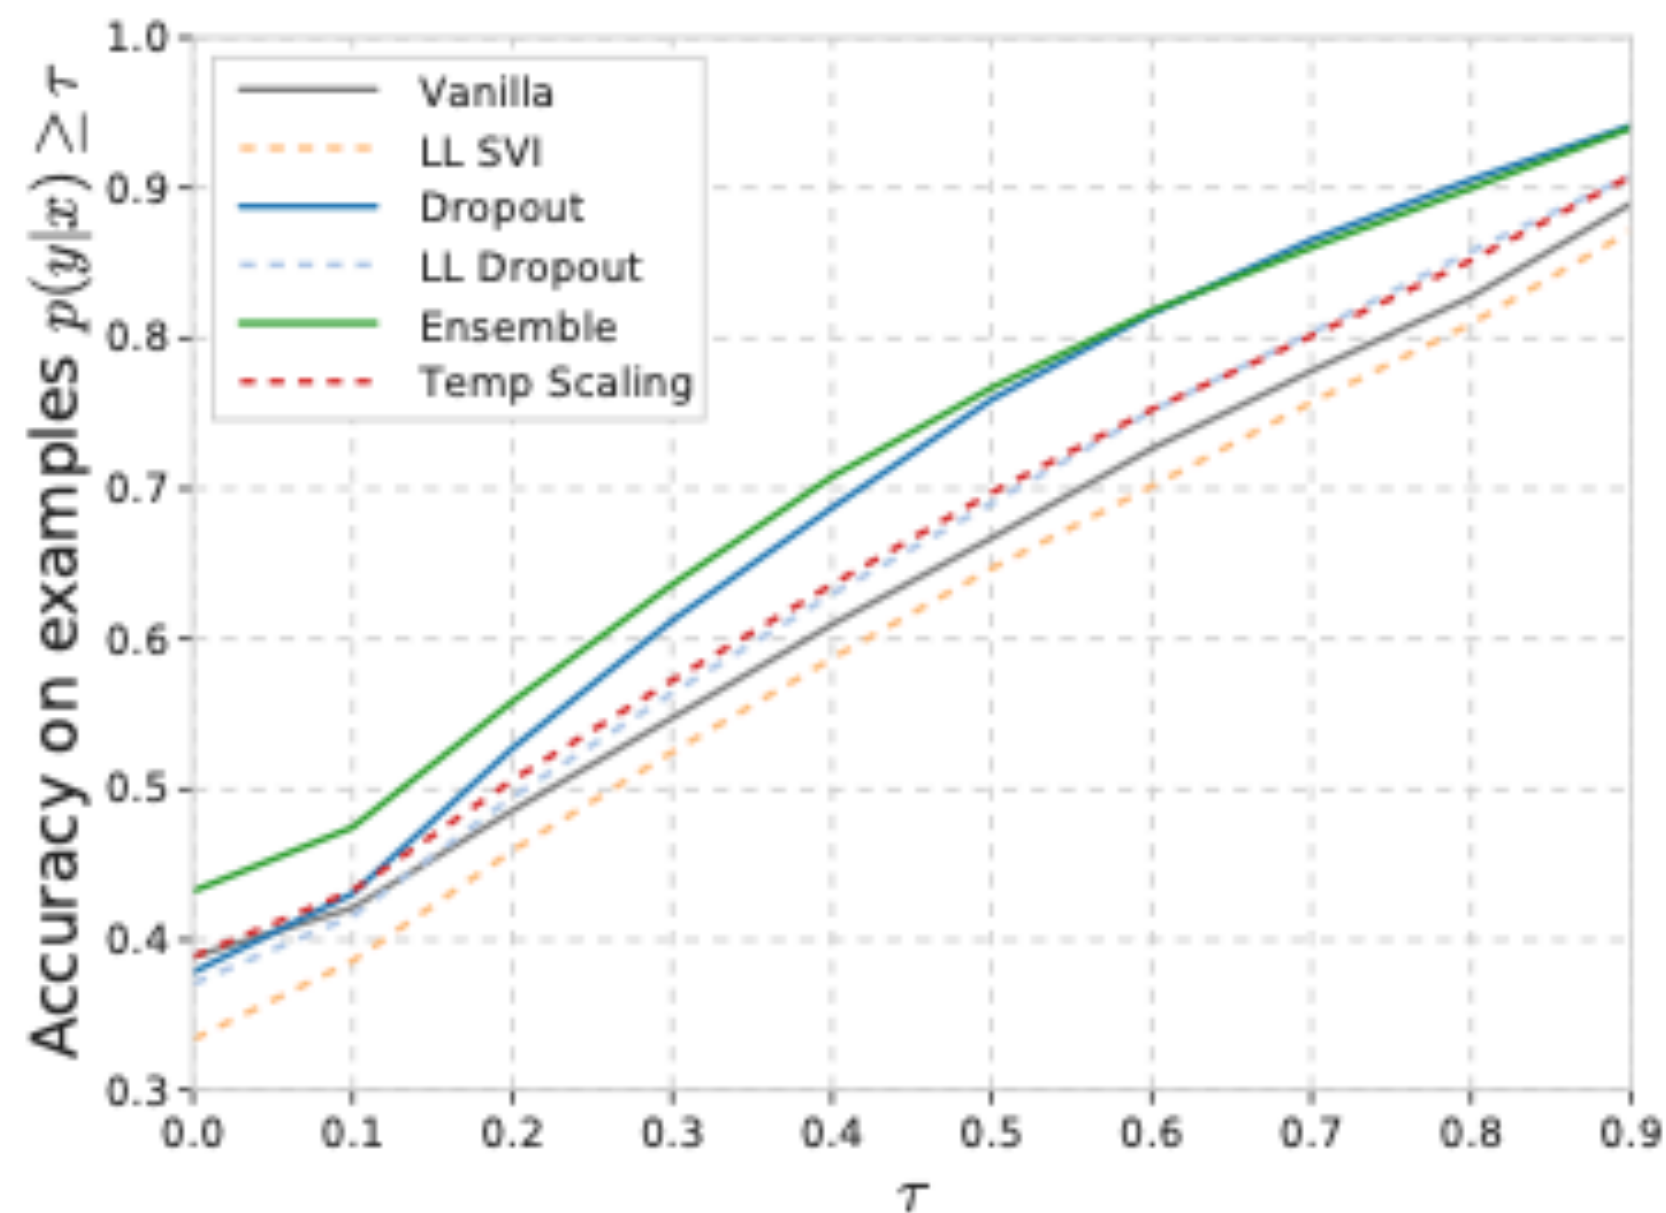
Average over 50 MCD stochastic forward passes




Overconfidence & uncertainty



Unfortunately uncertainty is not a necessarily a “fix”
& it get’s even harder when we try to select a threshold



(d) ImageNet: Confidence vs Acc

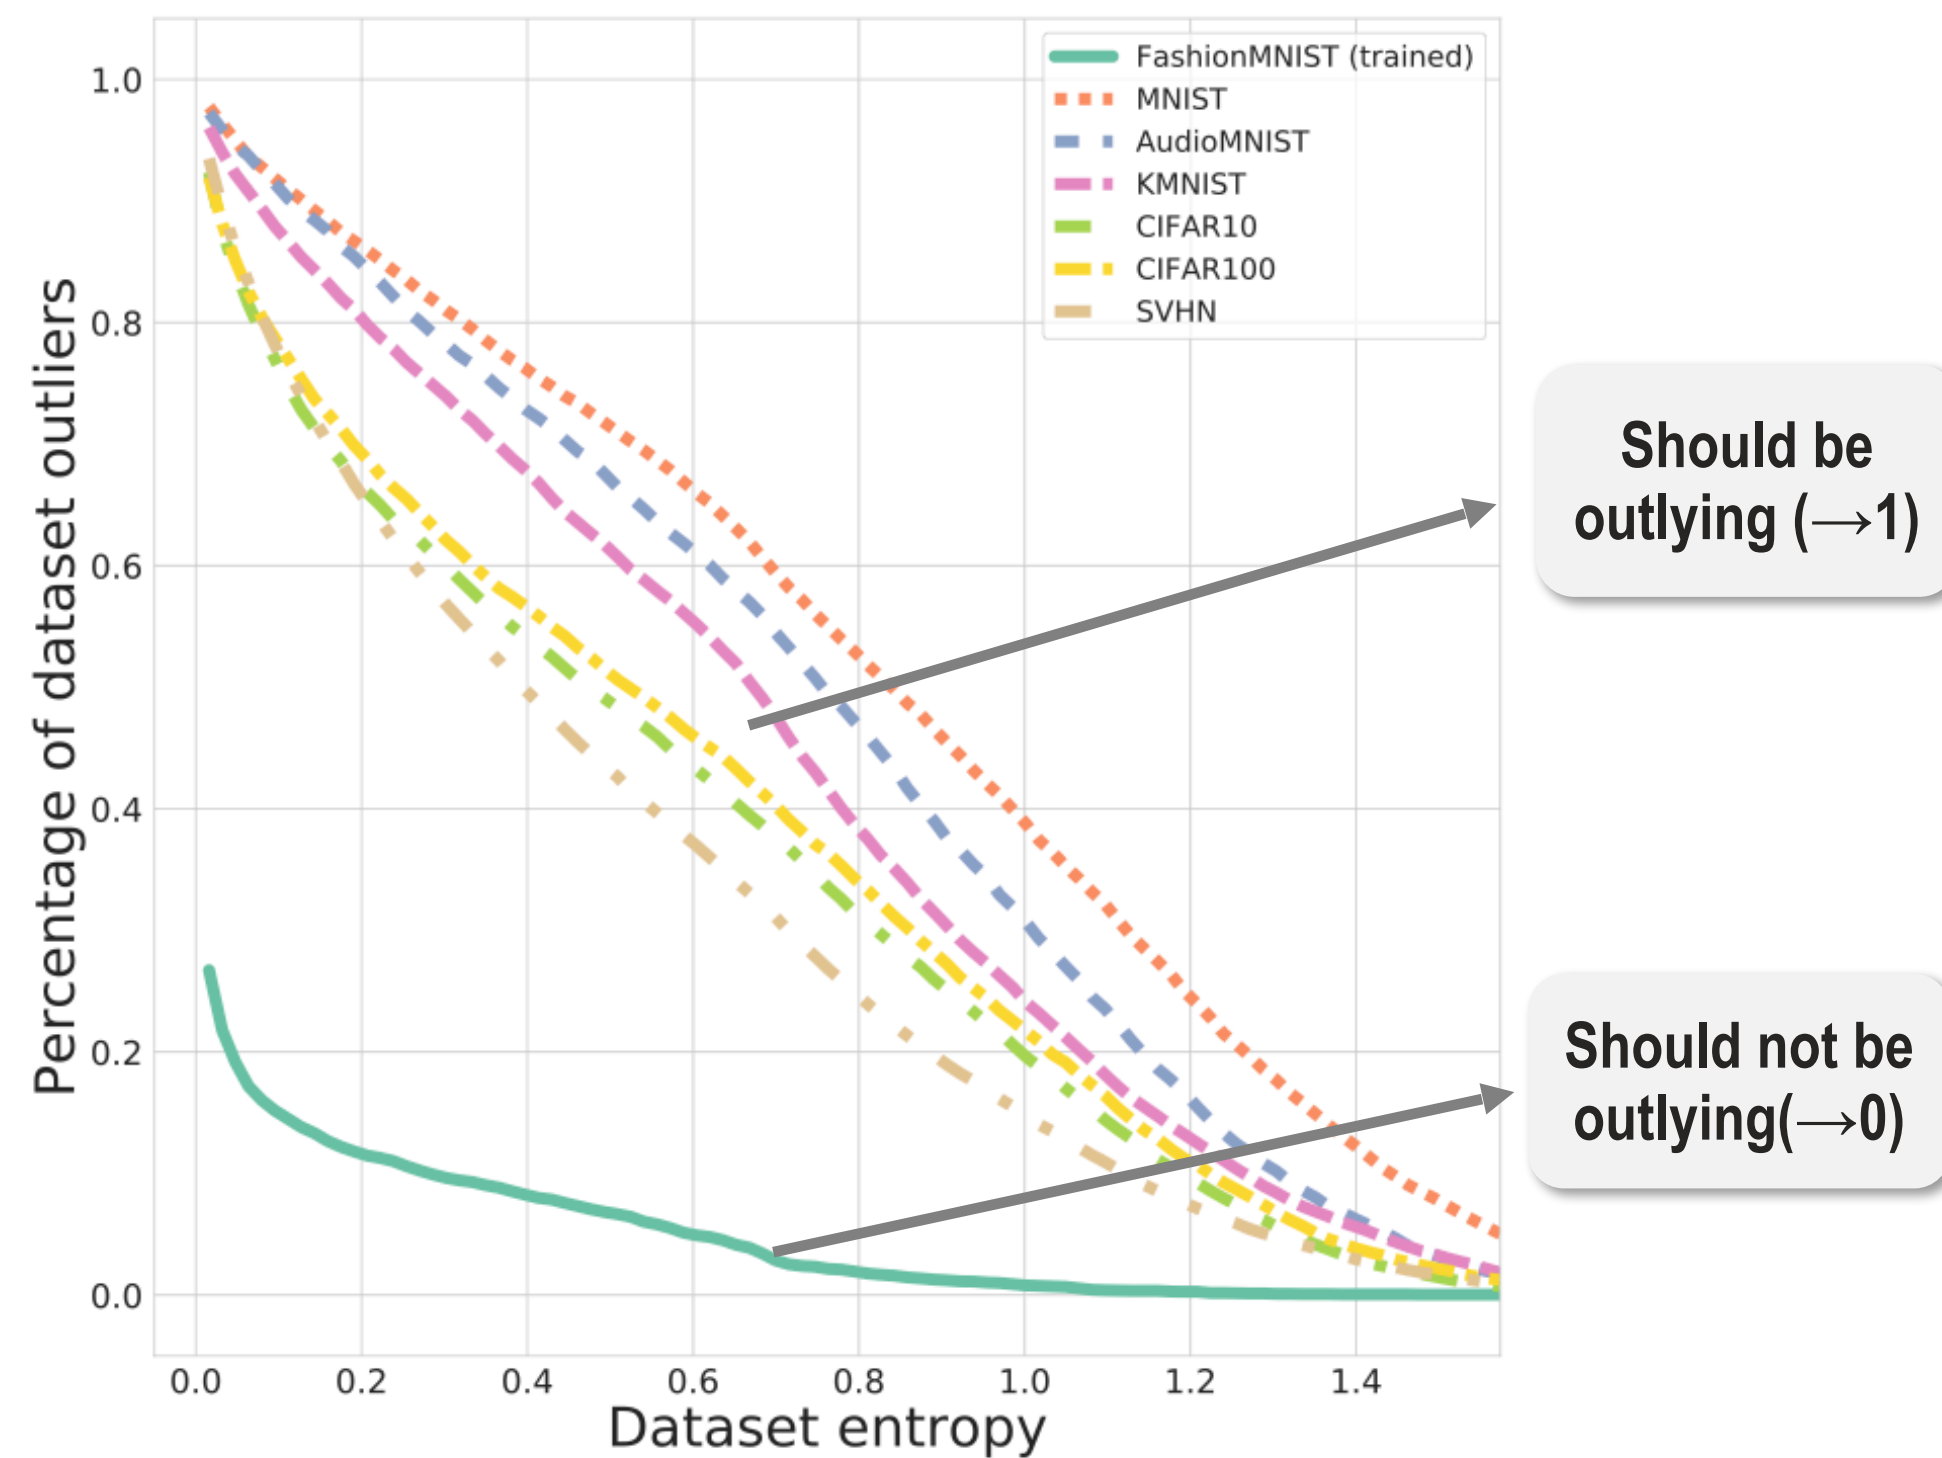


(f) CIFAR: Confidence on OOD

Overconfidence & gen. models



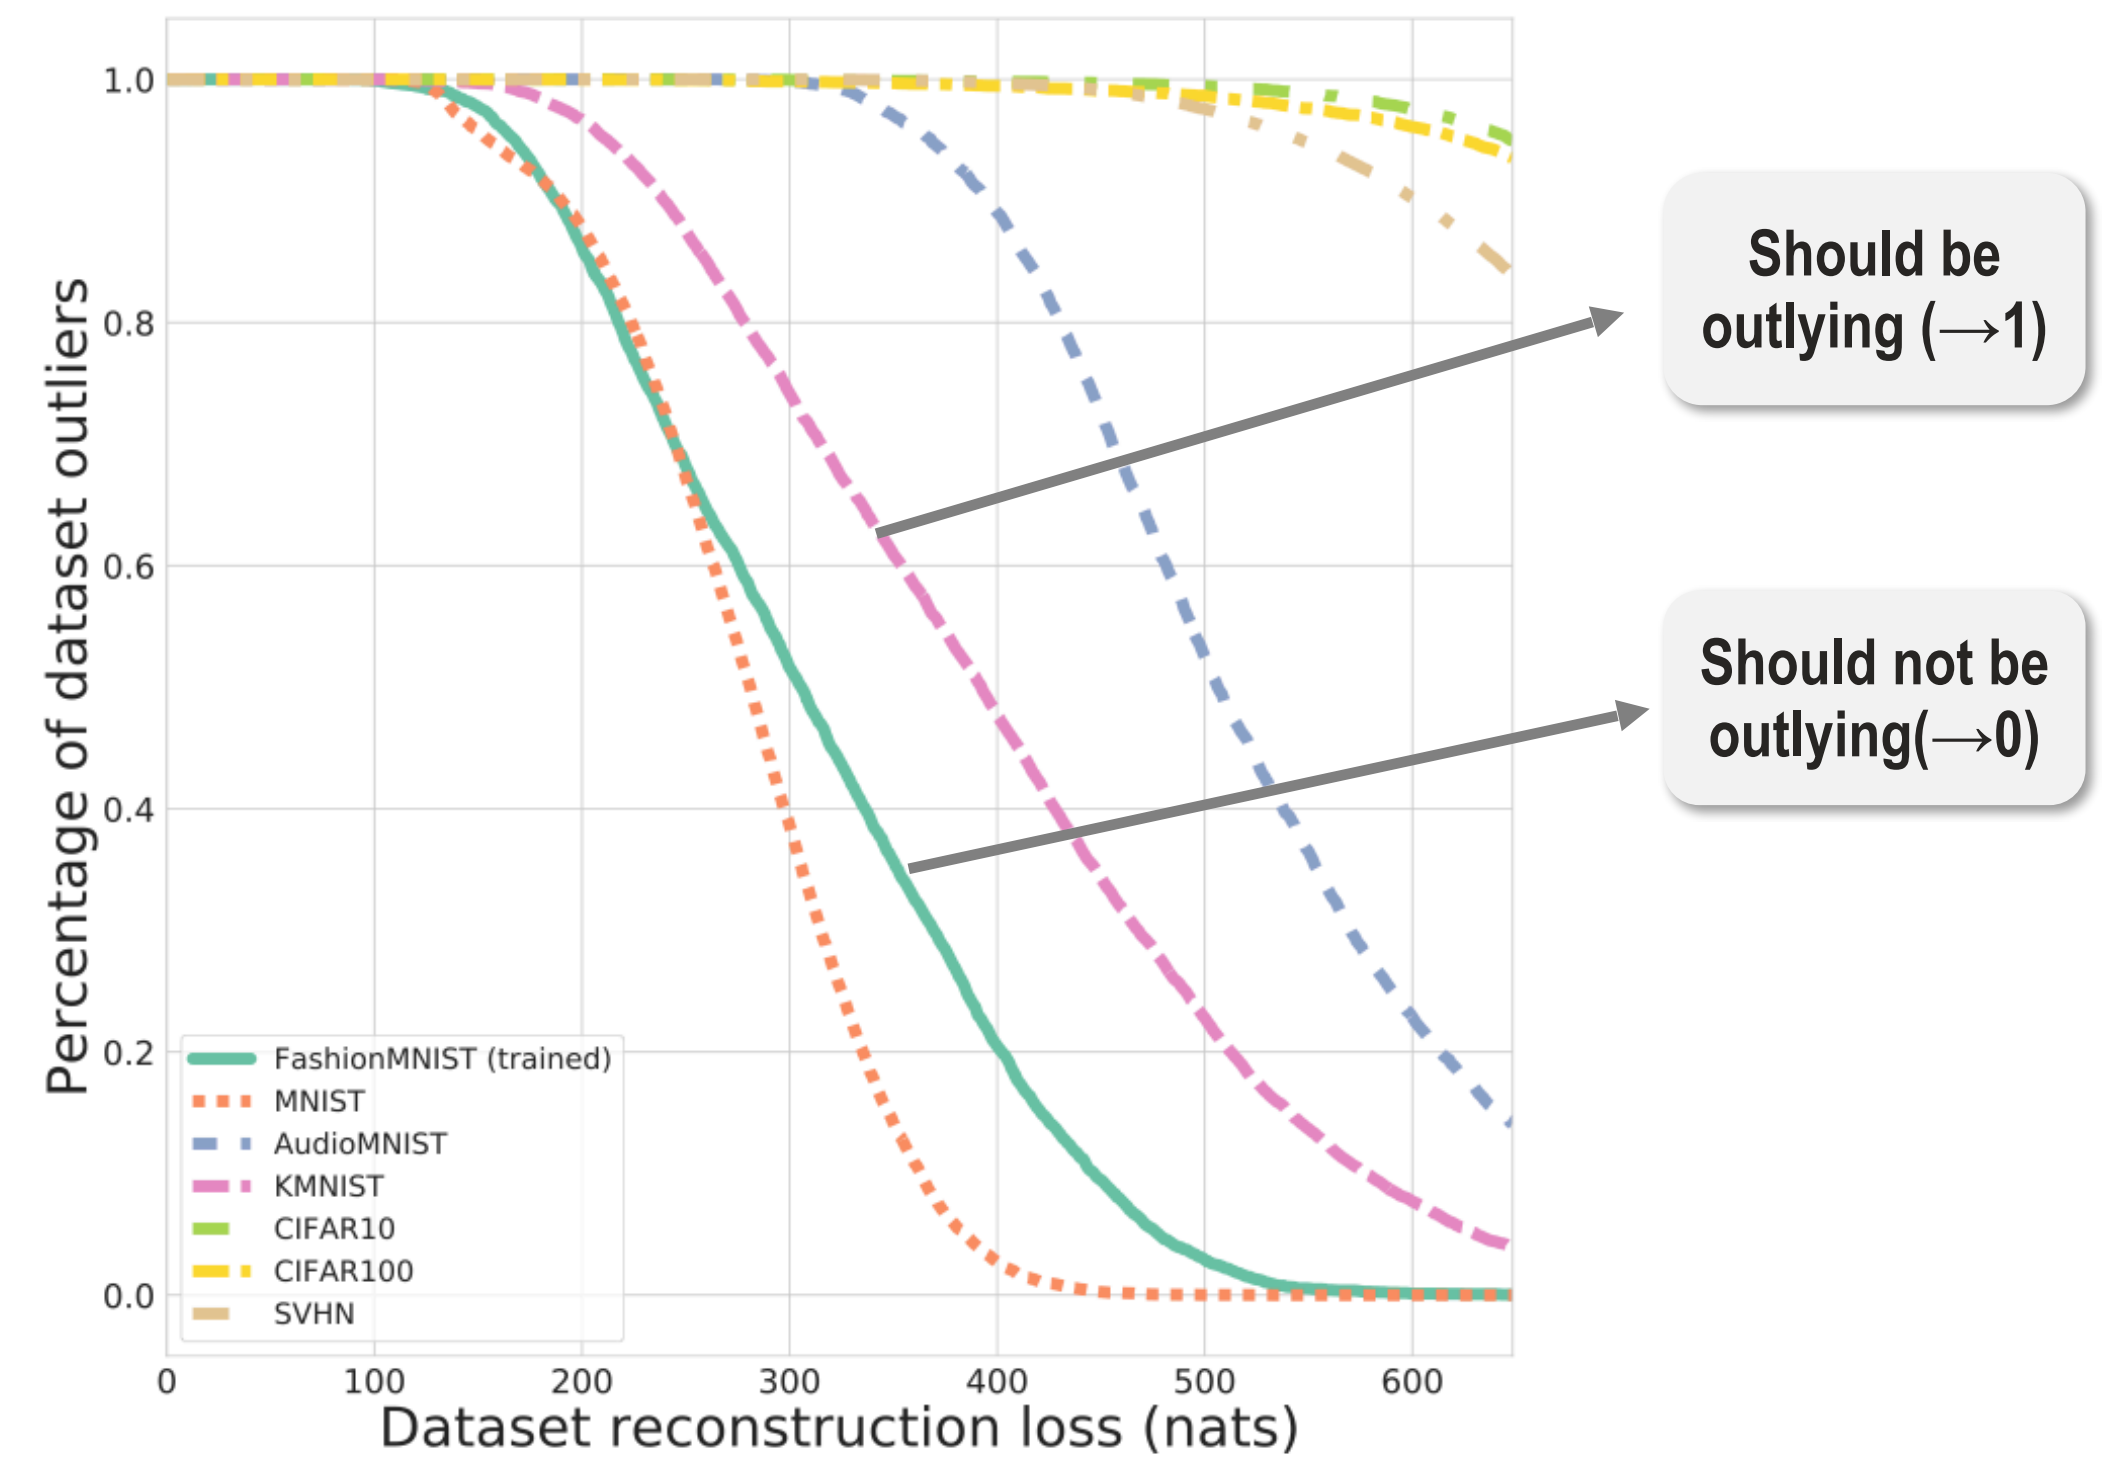

**Unfortunately uncertainty is not a necessarily a “fix”
& it get’s even harder when we try to select a threshold**



Overconfidence & uncertainty




Overconfidence is not exclusive to discriminative models




Overconfidence & gen. models




Overconfidence is not exclusive to discriminative models



Glow



PixelCNN

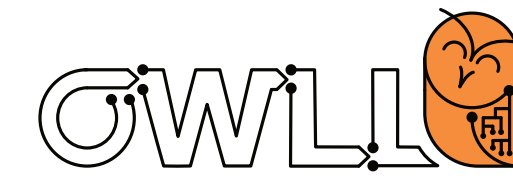


VAE

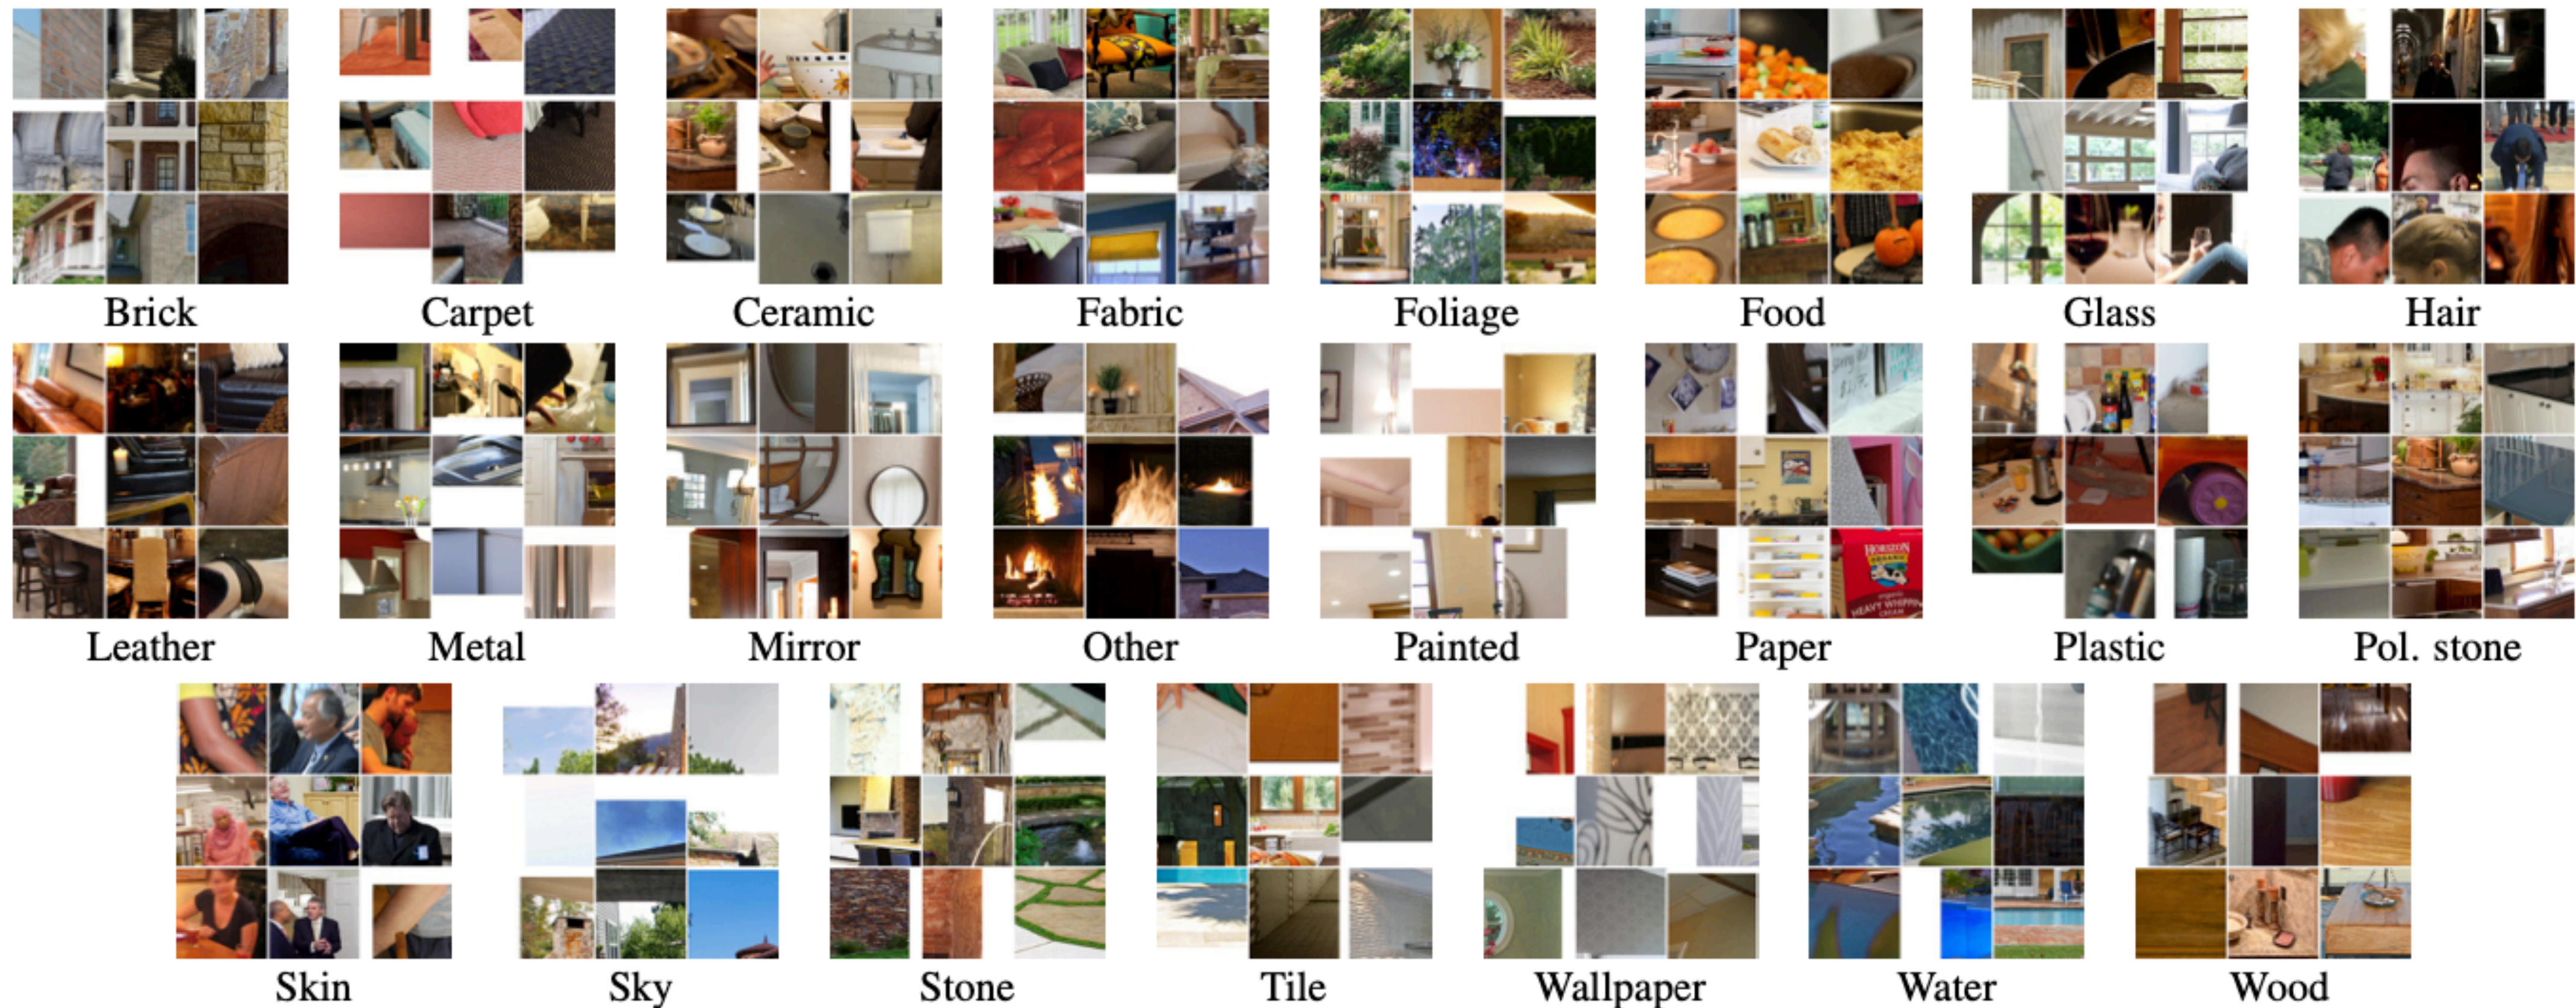


Including prior knowledge: an alternative?


The intuitive idea



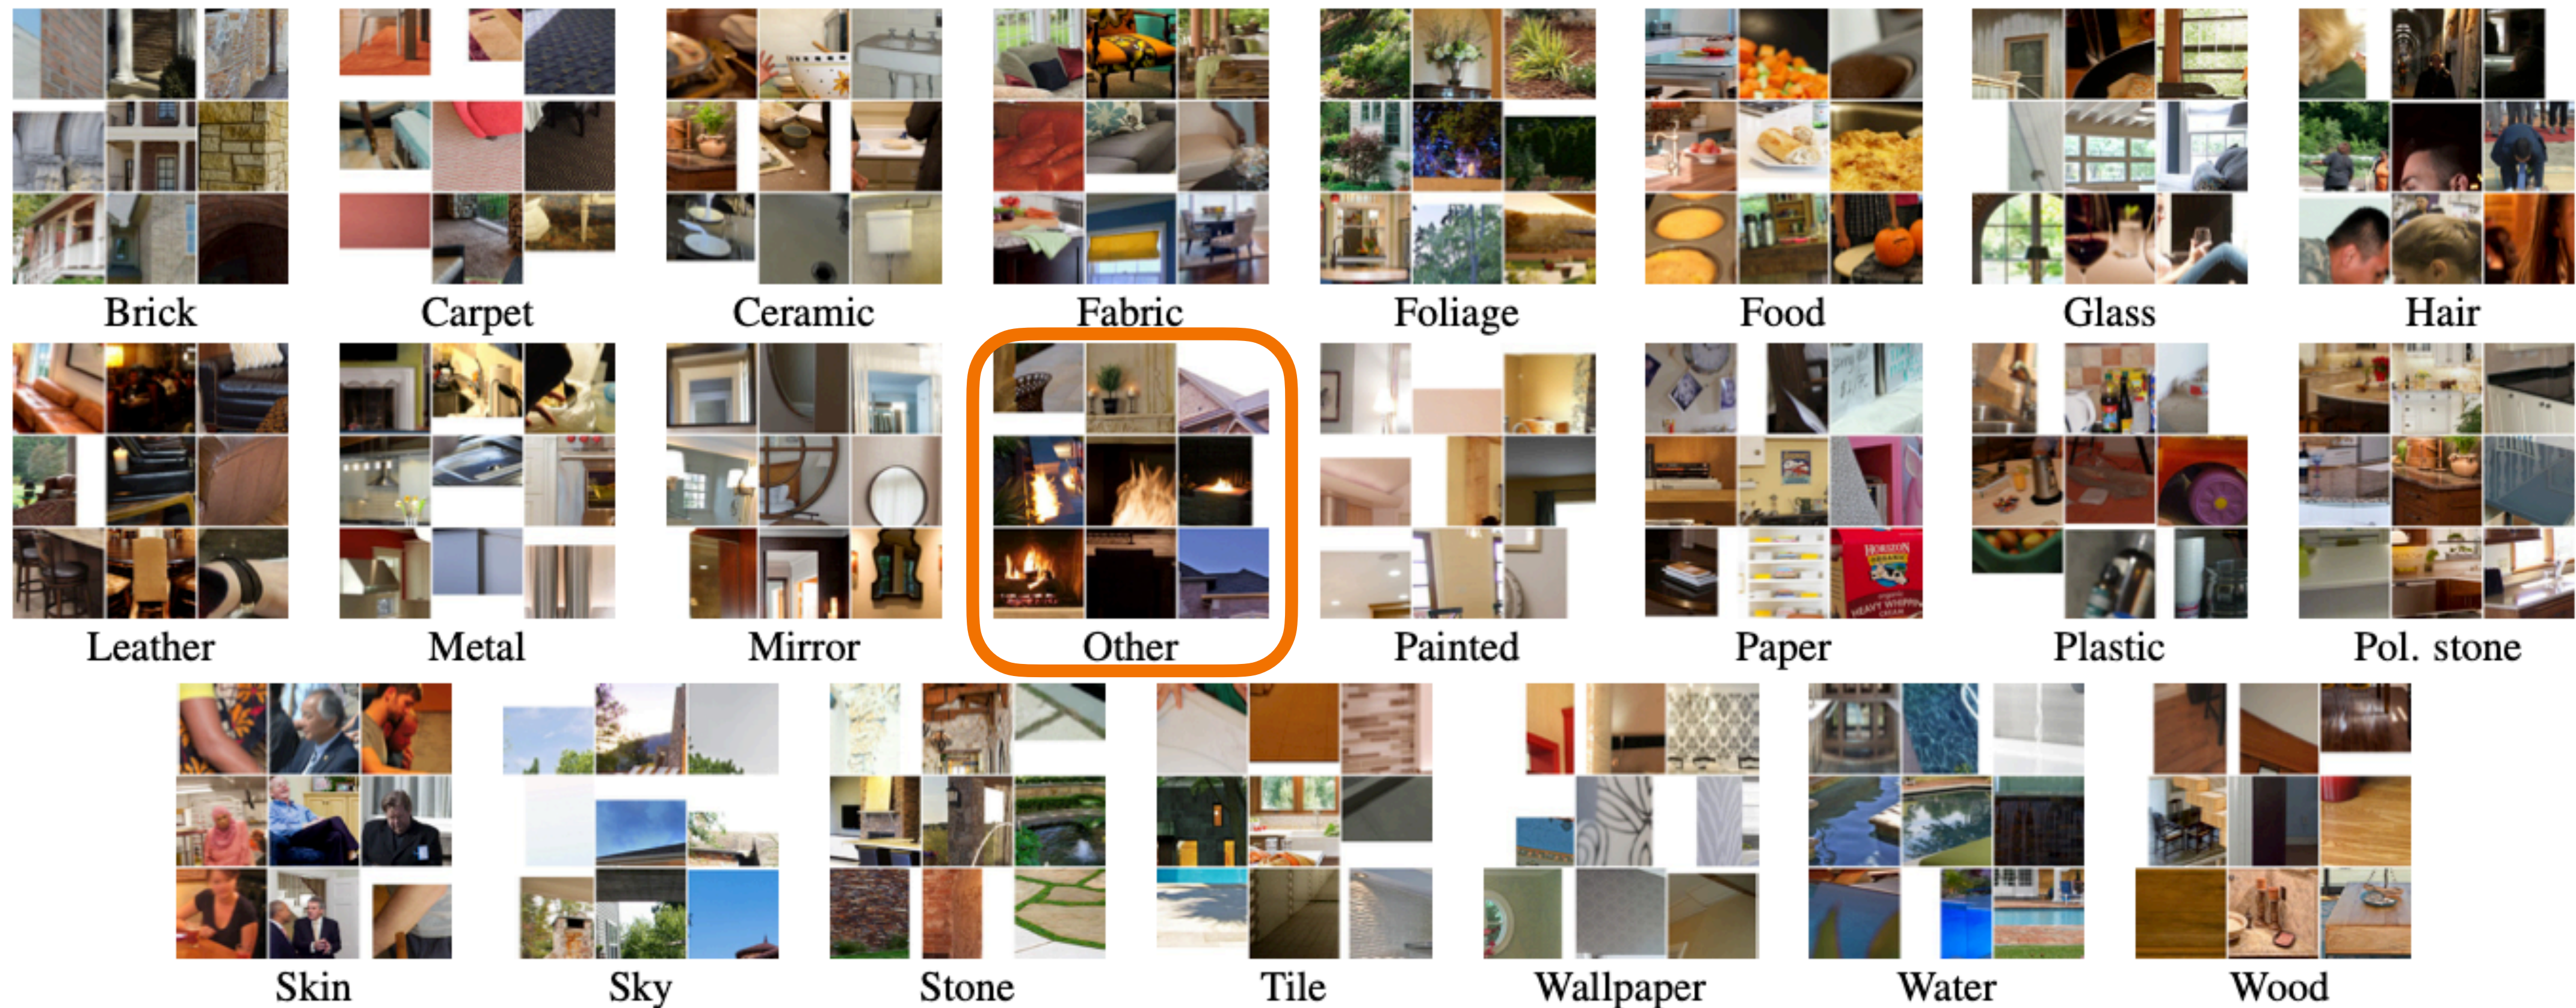
Take a look at the below Materials in Context (MINC) dataset: what do you notice?



The intuitive idea



An intuitive idea is to incorporate everything we know that does not belong to our task(s)



Inference with the universum



In essence: include **background class / “non-examples”** that aren't of interest

Key questions:

- How to implement the loss: many many conceivable conceivable
(Disclaimer: possibly *uncountable* amount of works)
- “what part of the universum is useful” (“Inference with the universum”, Weston et al, ICML 2006)
- "what are we expected to see during prediction later"? (Noise? Other concepts? Etc.)

Calibration: some examples



1. We could let our predictions (classifier) explicitly follow a uniform distribution for “out” data (Kimin Lee et al, “Training confidence-calibrated classifiers for detecting out-of-distribution samples”, ICLR 2018)

$$\min_{\theta} \mathbb{E}_{P_{\text{in}}(\hat{\mathbf{x}}, \hat{y})} \left[-\log P_{\theta}(y = \hat{y} | \hat{\mathbf{x}}) \right] + \beta \mathbb{E}_{P_{\text{out}}(\mathbf{x})} \left[KL(\mathcal{U}(y) \parallel P_{\theta}(y | \mathbf{x})) \right]$$

Calibration: some examples



1. We could let our predictions (classifier) explicitly follow a uniform distribution for “out” data (Kimin Lee et al, “Training confidence-calibrated classifiers for detecting out-of-distribution samples”, ICLR 2018)

$$\min_{\theta} \mathbb{E}_{P_{\text{in}}(\hat{\mathbf{x}}, \hat{y})} \left[-\log P_{\theta}(y = \hat{y} | \hat{\mathbf{x}}) \right] + \beta \mathbb{E}_{P_{\text{out}}(\mathbf{x})} \left[KL(\mathcal{U}(y) \parallel P_{\theta}(y | \mathbf{x})) \right]$$

2. We could calibrate our outputs, e.g. by scaling a temperature parameter later (Liang et al, “Enhancing the reliability of out-of-distribution image detection in neural networks”, ICLR 2018)

$$S_i(\mathbf{x}; T) = \frac{\exp(f_i(\mathbf{x})/T)}{\sum_{j=1}^N \exp(f_j(\mathbf{x})/T)}$$

Calibration: some examples



1. We could let our predictions (classifier) explicitly follow a uniform distribution for “out” data
(Kimin Lee et al, “Training confidence-calibrated classifiers for detecting out-of-distribution samples”, ICLR 2018)

$$\min_{\theta} \mathbb{E}_{P_{\text{in}}(\hat{\mathbf{x}}, \hat{y})} \left[-\log P_{\theta}(y = \hat{y} | \hat{\mathbf{x}}) \right] + \beta \mathbb{E}_{P_{\text{out}}(\mathbf{x})} \left[KL(\mathcal{U}(y) \parallel P_{\theta}(y | \mathbf{x})) \right]$$

2. We could calibrate our outputs, e.g. by scaling a temperature parameter later
(Liang et al, “Enhancing the reliability of out-of-distribution image detection in neural networks”, ICLR 2018)

$$S_i(\mathbf{x}; T) = \frac{\exp(f_i(\mathbf{x})/T)}{\sum_{j=1}^N \exp(f_j(\mathbf{x})/T)}$$


3. And many other versions to modify our loss, e.g.:
(Dhamija et al, “Reducing network agnostophobia”, NeurIPS 2018)

$$J_E(x) = \begin{cases} -\log S_c(x) & \text{if } x \in \mathcal{D}'_c \text{ is from class } c \\ -\frac{1}{C} \sum_{c=1}^C \log S_c(x) & \text{if } x \in \mathcal{D}'_b \end{cases}$$


Background & Objectosphere




We could also think about encouraging features to be zero for OOD data



(a) Softmax



(b) Background



(c) Objectosphere

Figure 1: LENET++ RESPONSES TO KNOWN AND UNKNOWN. The network in (a) was only trained to classify the 10 MNIST classes (\mathcal{D}'_c) using softmax, while the networks in (b) and (c) added NIST letters [15] as known unknowns (\mathcal{D}'_b) trained with softmax or our novel Objectosphere loss.




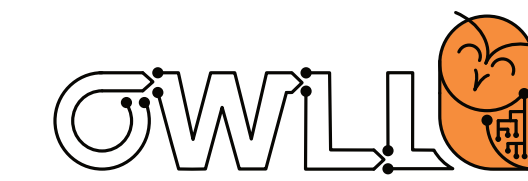
What do you think are the up & downsides so far?

Closed -open world assumption



We may need a different approach: as the world grows more “open” we move from **known unknowns** to **unknown unknowns**. Our two perspectives only handle the former






Open set recognition & explicit bounds

Intuition behind open space



Intuitively: we could take into account distances from the known data points



Intuition behind open space

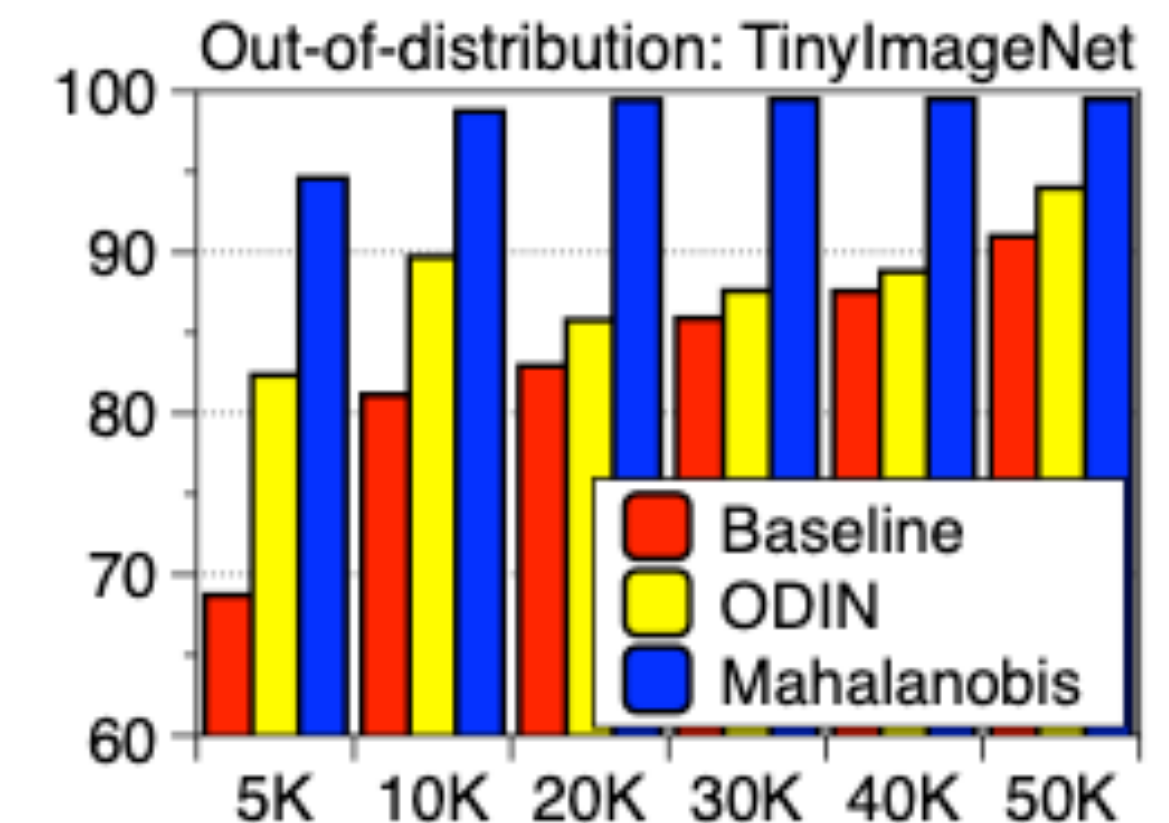




Intuitively: we could take into account distances from the known data points

Example 1 : we could make assumptions like every class being Normal distributed & then calculate distances to our existing data points, e.g. Mahalanobis distance


$$\hat{\mu}_c = \frac{1}{N_c} \sum_{i:y_i=c} f(\mathbf{x}_i), \quad \hat{\Sigma} = \frac{1}{N} \sum_c \sum_{i:y_i=c} (f(\mathbf{x}_i) - \hat{\mu}_c)(f(\mathbf{x}_i) - \hat{\mu}_c)^\top$$

$$M(\mathbf{x}) = \max_c - (f(\mathbf{x}) - \hat{\mu}_c)^\top \hat{\Sigma}^{-1} (f(\mathbf{x}) - \hat{\mu}_c)$$



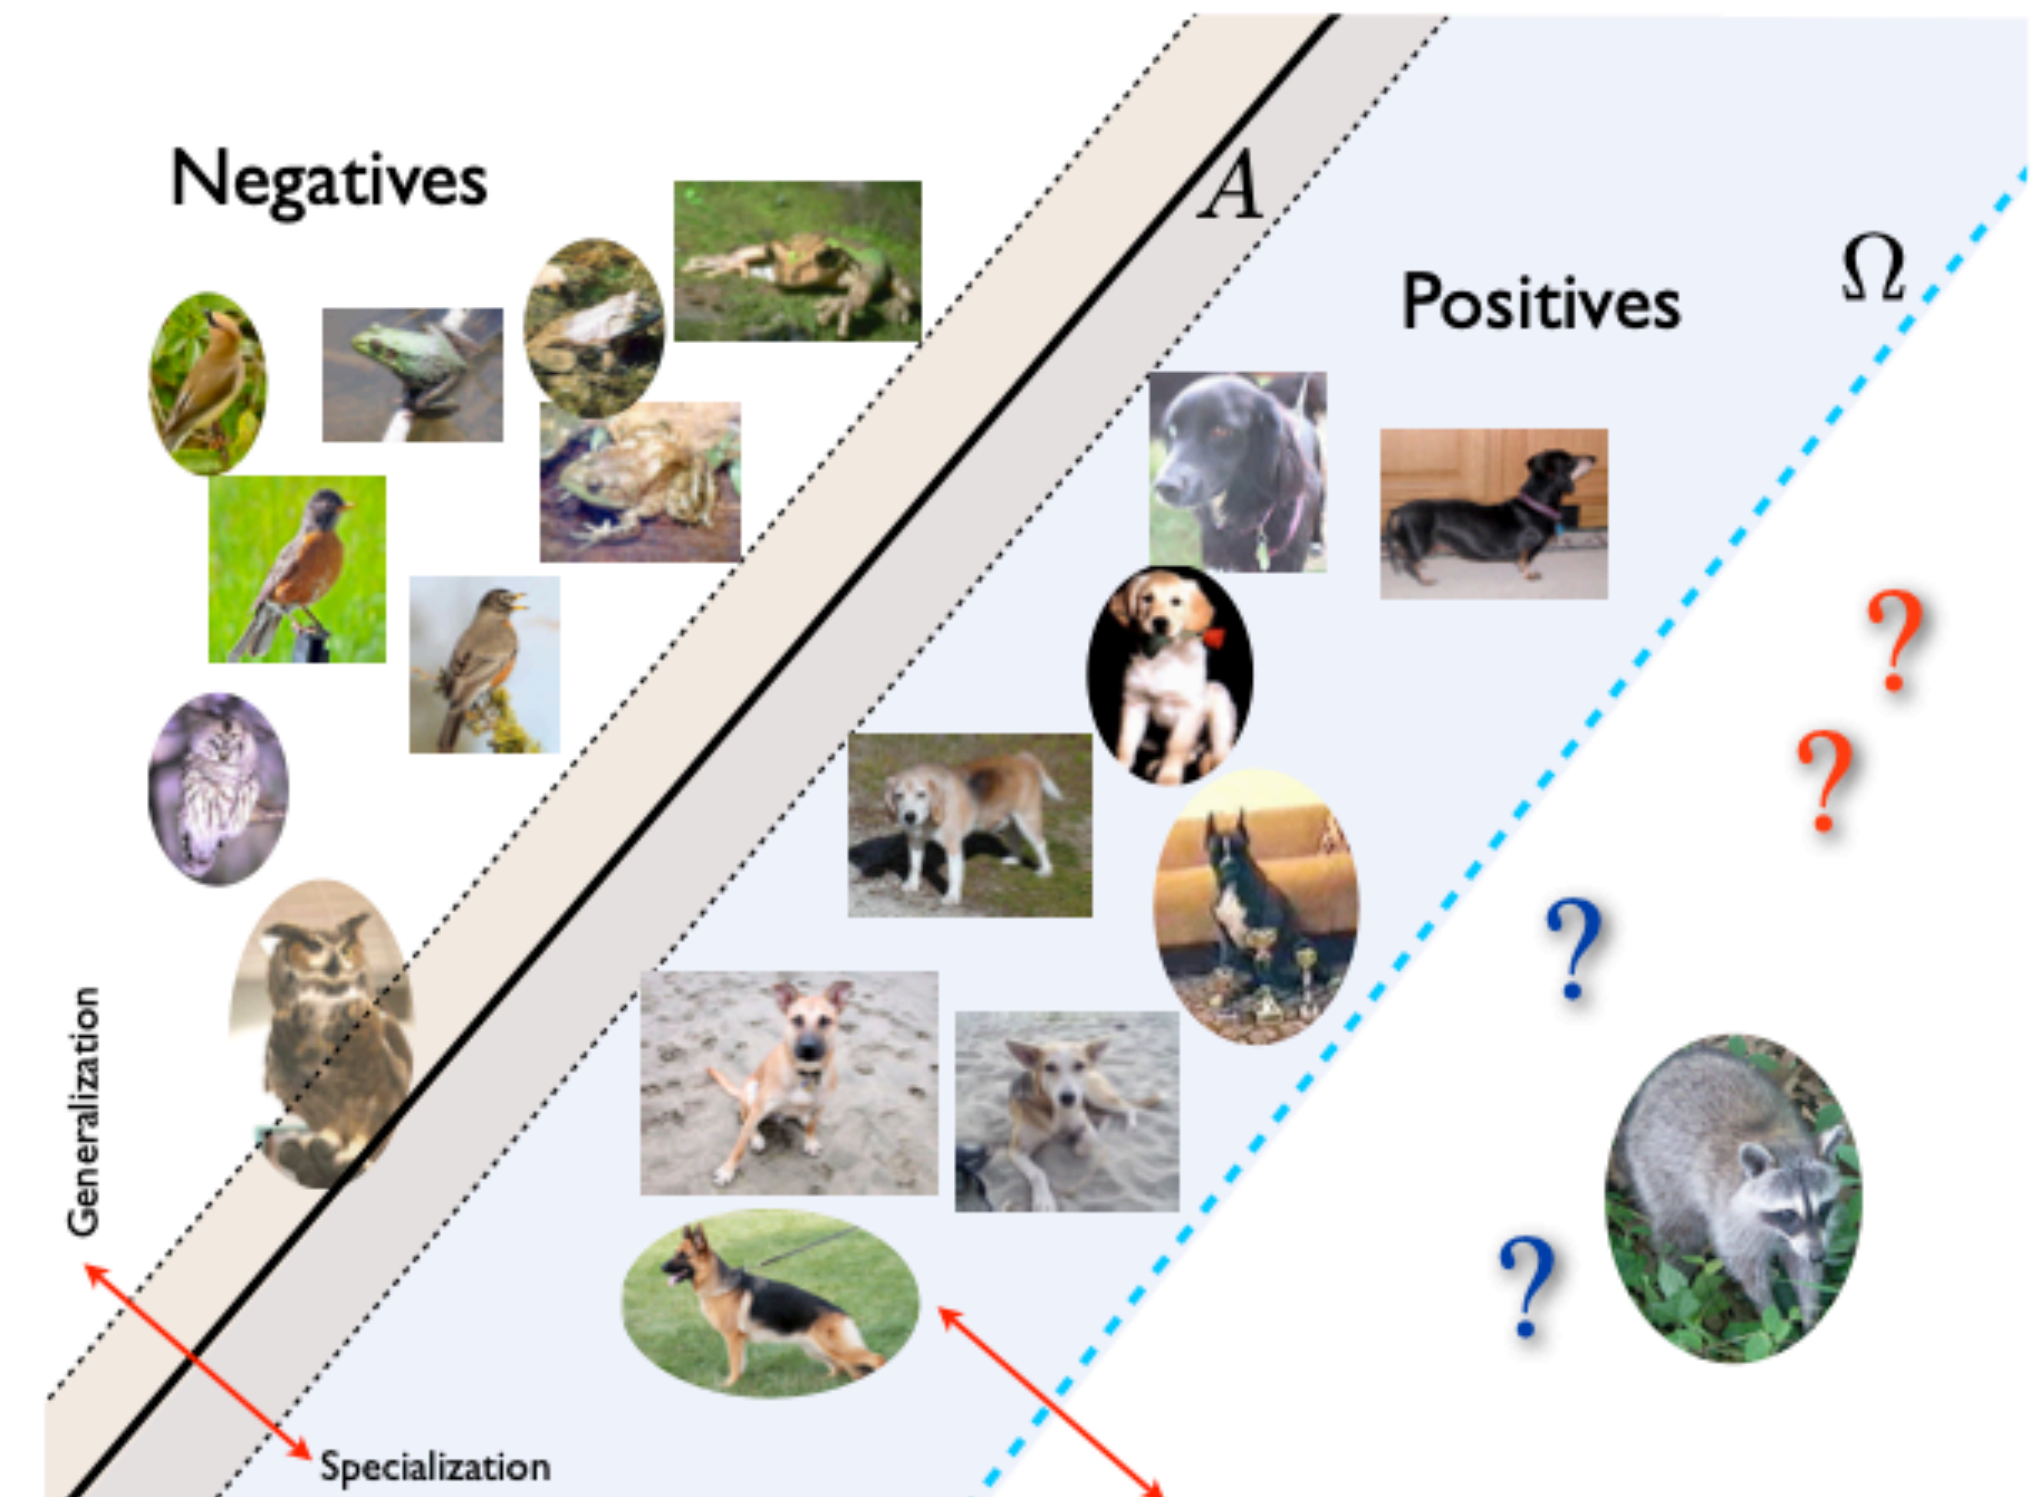
(E.g. Kimin Lee et al, “A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks”, NeurIPS 2018)

Intuition behind open space



Intuitively: we could take into account distances from the known data points

Example 2: we could fit another parallel plane in an SVM, for a reject option, based on the support set with large distances



Formalizing open space/sets




Intuitively: open space is what we have not covered with known data

Formally: (see e.g. “Learning and the Unknown”, Boult et al, AAI 2019)

For a recognition function f over space \mathcal{X} & a union of balls with radius r that includes all known training examples:

$$\mathcal{O} = \mathcal{X} - \bigcup_{i \in N} B_r(x_i)$$



Formalizing open space/sets




For a recognition function f over space \mathcal{X} & a union of balls with radius r that includes all known training examples:

$$\mathcal{O} = \mathcal{X} - \bigcup_{i \in N} B_r(x_i)$$

Can now define open space risk as a relative measure of open space to the full space, but see the survey for the full math


For now: the aim would be to decay the probability away from supporting evidence



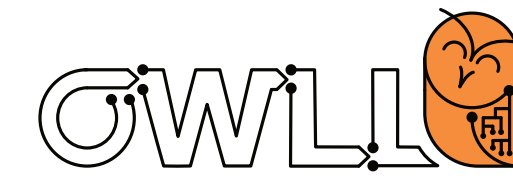
Bounds with extreme values



In other words, we could fit a distance based model (following the radius idea), e.g. here based on the mean activations of training data in a deep net



Bounds with extreme values



In other words, we could fit a distance based model (following the radius idea), e.g. here based on the mean activations of training data in a deep net

Algorithm 1 EVT Meta-Recognition Calibration for Open Set Deep Networks, with per class Weibull fit to η largest distance to mean activation vector. Returns libMR models ρ_j which includes parameters τ_i for shifting the data as well as the Weibull shape and scale parameters: κ_i, λ_i .

Require: FitHigh function from libMR

Require: Activation levels in the penultimate network layer $\mathbf{v}(\mathbf{x}) = v_1(x) \dots v_N(x)$

Require: For each class j let $S_{i,j} = v_j(x_{i,j})$ for each correctly classified training example $x_{i,j}$.

1: **for** $j = 1 \dots N$ **do**

2: **Compute mean AV**, $\mu_j = \text{mean}_i(S_{i,j})$

3: **EVT Fit** $\rho_j = (\tau_j, \kappa_j, \lambda_j) = \text{FitHigh}(\|\hat{S}_j - \mu_j\|, \eta)$

4: **end for**

5: **Return** means μ_j and libMR models ρ_j

Bounds with extreme values



In other words, we could fit a distance based model (following the radius idea), e.g. here based on the mean activations of training data in a deep net

But which distribution should we choose?

- We are mainly interested in the extreme distances, as we want to make a decision of when to reject
- Extreme value theory may provide an answer for us

Algorithm 1 EVT Meta-Recognition Calibration for Open Set Deep Networks, with per class Weibull fit to η largest distance to mean activation vector. Returns libMR models ρ_j which includes parameters τ_i for shifting the data as well as the Weibull shape and scale parameters: κ_i, λ_i .

Require: FitHigh function from libMR

Require: Activation levels in the penultimate network layer $\mathbf{v}(\mathbf{x}) = v_1(x) \dots v_N(x)$

Require: For each class j let $S_{i,j} = v_j(x_{i,j})$ for each correctly classified training example $x_{i,j}$.

- 1: **for** $j = 1 \dots N$ **do**
 - 2: **Compute mean AV**, $\mu_j = \text{mean}_i(S_{i,j})$
 - 3: **EVT Fit** $\rho_j = (\tau_j, \kappa_j, \lambda_j) = \text{FitHigh}(\|\hat{S}_j - \mu_j\|, \eta)$
 - 4: **end for**
 - 5: **Return** means μ_j and libMR models ρ_j
-


Bounds with extreme values




Extreme value theory is interested in the probability of events that are more extreme than any previously observed

$$f(x; \lambda, k) = \begin{cases} \frac{k}{\lambda} \left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k} & x \geq 0 \\ 0 & x < 0 \end{cases}$$

Regardless of the overall distribution, if the data is bounded, EVT tells us that sampling the tail/the extrema away from the median of our distribution results in an EVT distribution: Weibull, Gumbel or Fréchet



Bounds with extreme values



We can use the cumulative distribution function (CDF) to either reject right away, because we exceed our extremely observed distances, or use the value to modify our prediction score (Referred to as OpenMax here)

Algorithm 2 OpenMax probability estimation with rejection of unknown or uncertain inputs.

Require: Activation vector for $\mathbf{v}(\mathbf{x}) = v_1(x), \dots, v_N(x)$

Require: means μ_j and libMR models $\rho_j = (\tau_i, \lambda_i, \kappa_i)$

Require: α , the number of “top” classes to revise

1: Let $s(i) = \text{argsort}(v_j(x))$; Let $\omega_j = 1$

2: **for** $i = 1, \dots, \alpha$ **do**

3: $\omega_{s(i)}(x) = 1 - \frac{\alpha - i}{\alpha} e^{-\left(\frac{\|x - \tau_{s(i)}\|}{\lambda_{s(i)}}\right)^{\kappa_{s(i)}}$

4: **end for**

5: Revise activation vector $\hat{v}(x) = \mathbf{v}(\mathbf{x}) \circ \omega(\mathbf{x})$

6: Define $\hat{v}_0(x) = \sum_i v_i(x)(1 - \omega_i(x))$.


7:

$$\hat{P}(y = j | \mathbf{x}) = \frac{e^{\hat{v}_j(\mathbf{x})}}{\sum_{i=0}^N e^{\hat{v}_i(\mathbf{x})}}$$

8: Let $y^* = \text{argmax}_j P(y = j | \mathbf{x})$



9: Reject input if $y^* == 0$ or $P(y = y^* | \mathbf{x}) < \epsilon$

OpenMax in a generative variant



OpenMax seem to improve a lot!

But why is there still so much room for improvement?






Do we need generative models on top?


Overconfidence & gen. models




Recall earlier: overconfidence is not exclusive to discriminative models, but what if it's only about predictive values again?



Glow



PixelCNN

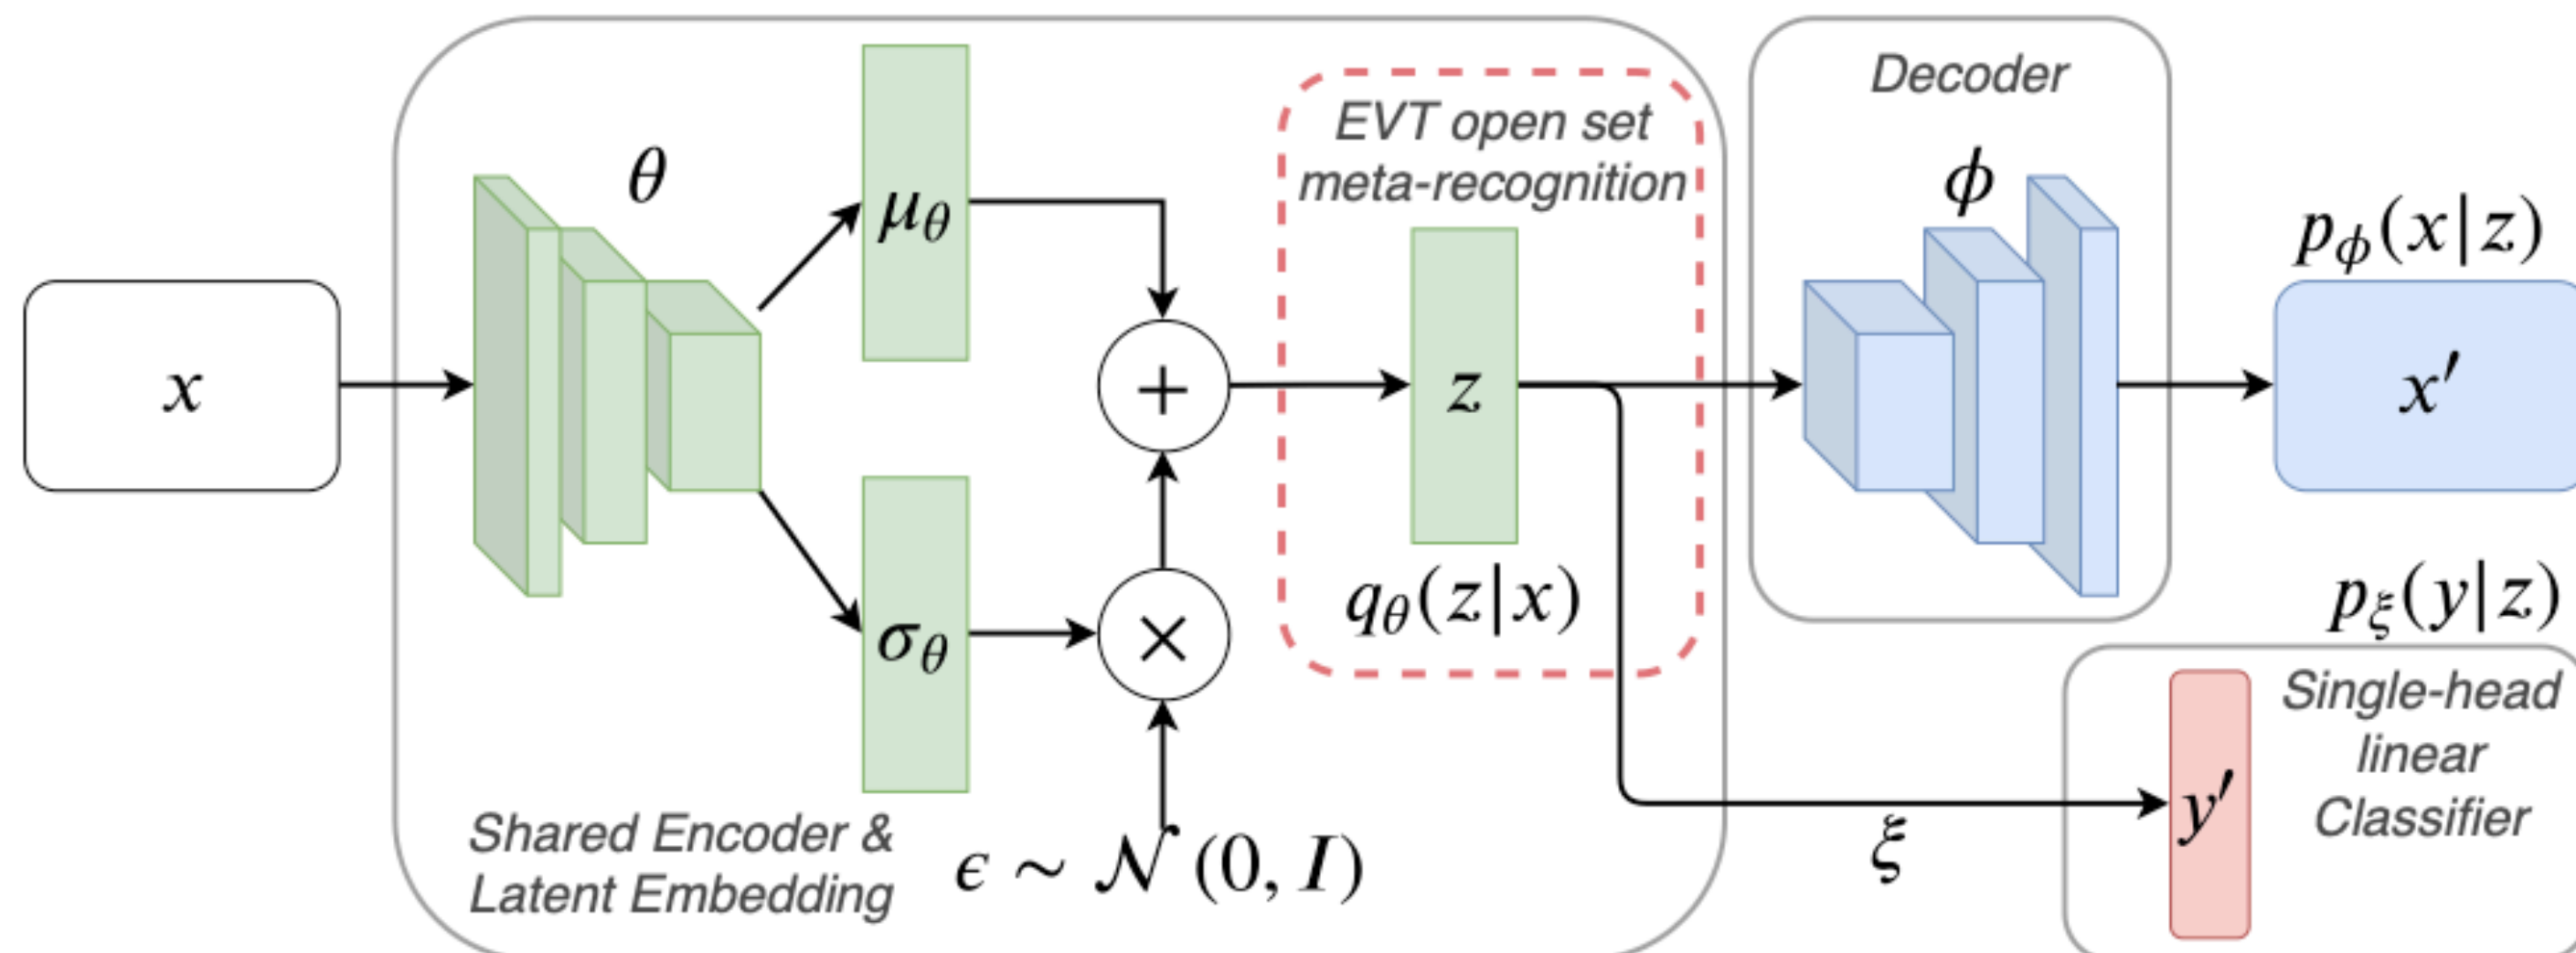


VAE

OpenMax in a generative variant



We could formulate an OpenMax variant based on a VAE, based on generative factors



Algorithm 1 Open set recognition calibration for deep variational neural networks. A Weibull model fit of tail-size η is conducted to bound the per class approximate posterior. Per class c Weibull models ρ_c with their respective shift τ_c , shape κ_c and scale λ_c parameters are returned.

Require: Trained encoder $q_\theta(z|x)$ and classifier $p_\xi(y|z)$
Require: Classifier probabilities $p_\xi(y|z)$ and samples from the approximate posterior $z(x^{(i)}) \sim q_\theta(z|x^{(i)})$ for each training dataset example $x^{(i)}$

Require: For each class c , let $S_c^{(i)} = z(x_c'^{(i)})$ for each correctly classified training example $x_c'^{(i)}$


- 1: **for** $c = 1 \dots C$ **do**
- 2: **Get per class latent mean** $\bar{S}_c = \text{mean}(S_c^{(i)})$
- 3: **Weibull model** $\rho_c = \text{Fit Weibull}(\|S_c - \bar{S}_c\|, \eta)$
- 4: **Return** means \bar{S} and Weibull models ρ

OpenMax in a generative variant

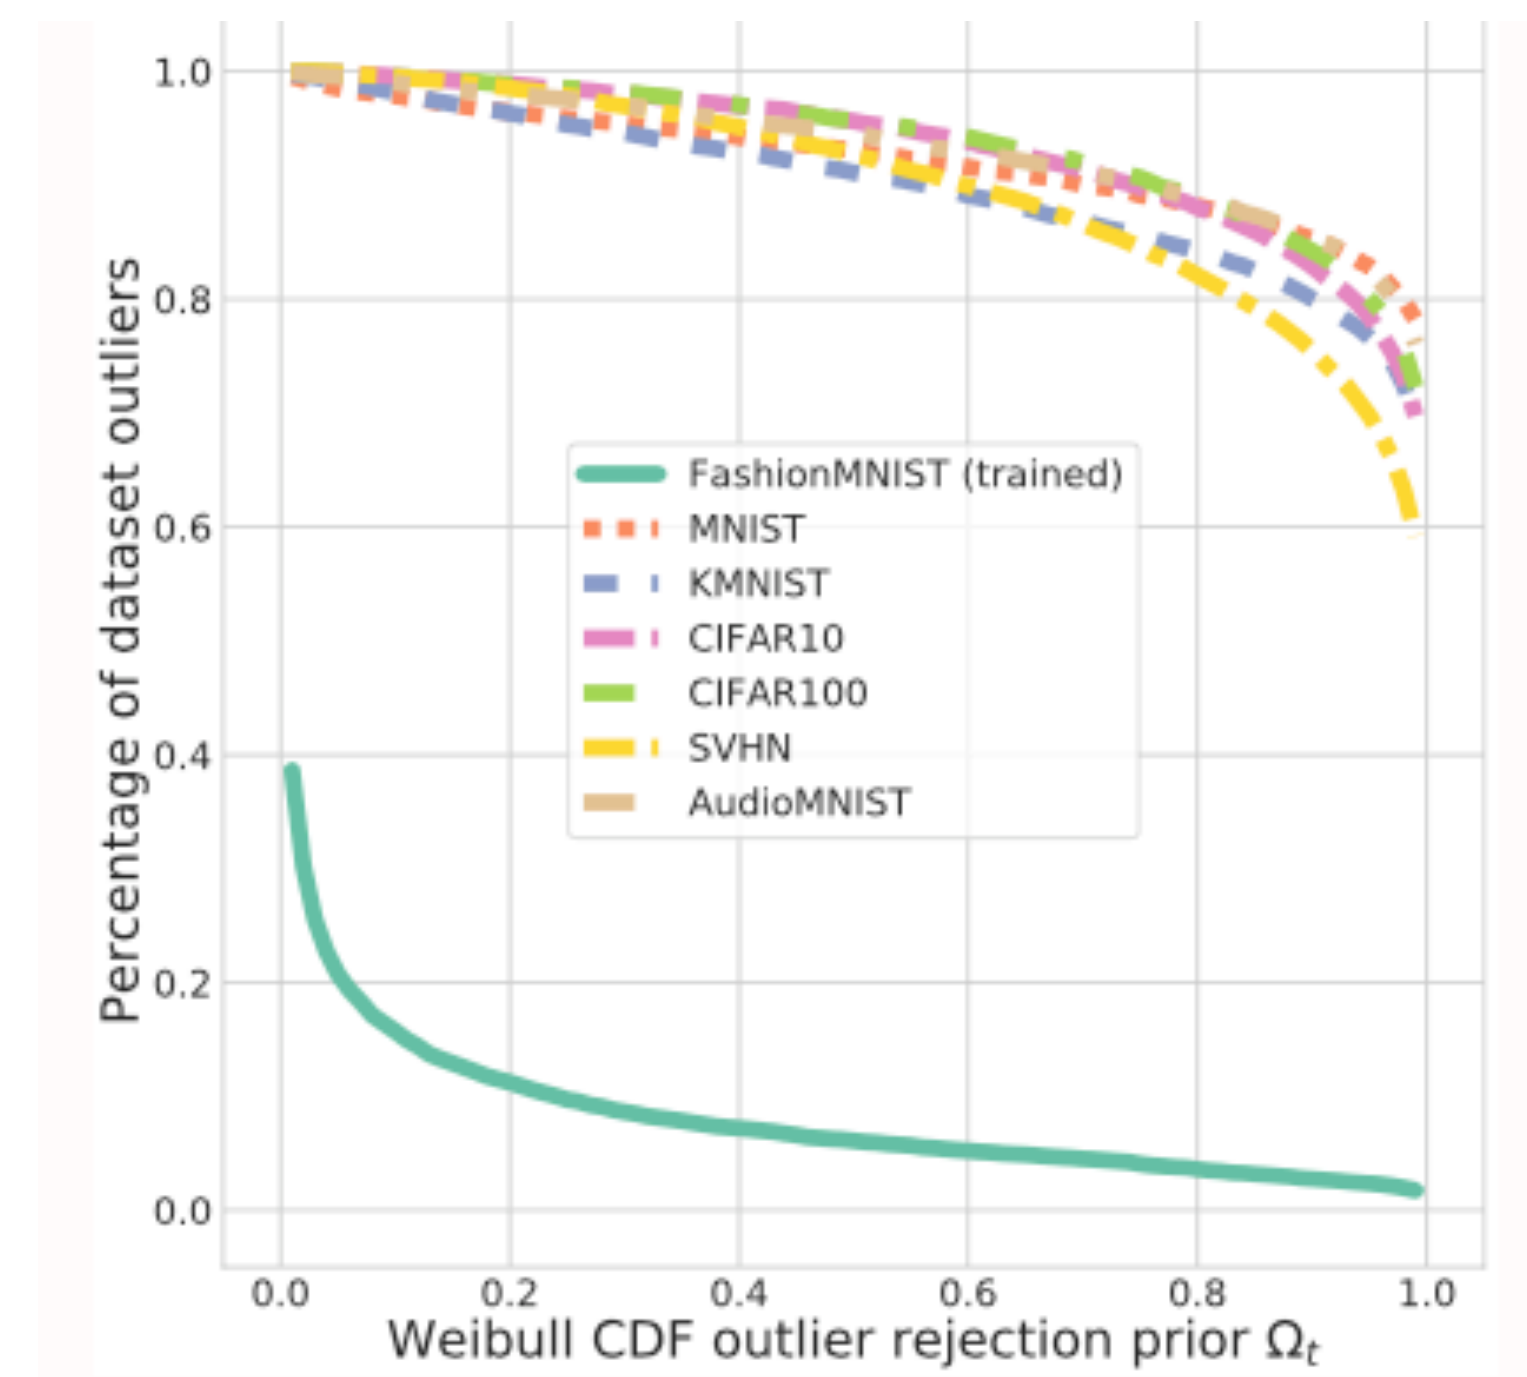



It may indeed be a question of the learned representations

Standard classifier $p(y|x)$ with OpenMax



“Open”VAE approach: $p(x,y)$






As an alternative/auxiliary approach, we could also take a direct look at the functions that we use in our model

An alternative/auxiliary view



Hypothesis: specific functions in our ML models, like ReLU in NNs are (at least in parts) the culprit - they always produce high confidence far away from the data (Hein et al, “Why ReLU networks yield high confidence predictions far away from the training data and how to mitigate the problem”, CVPR 2019)

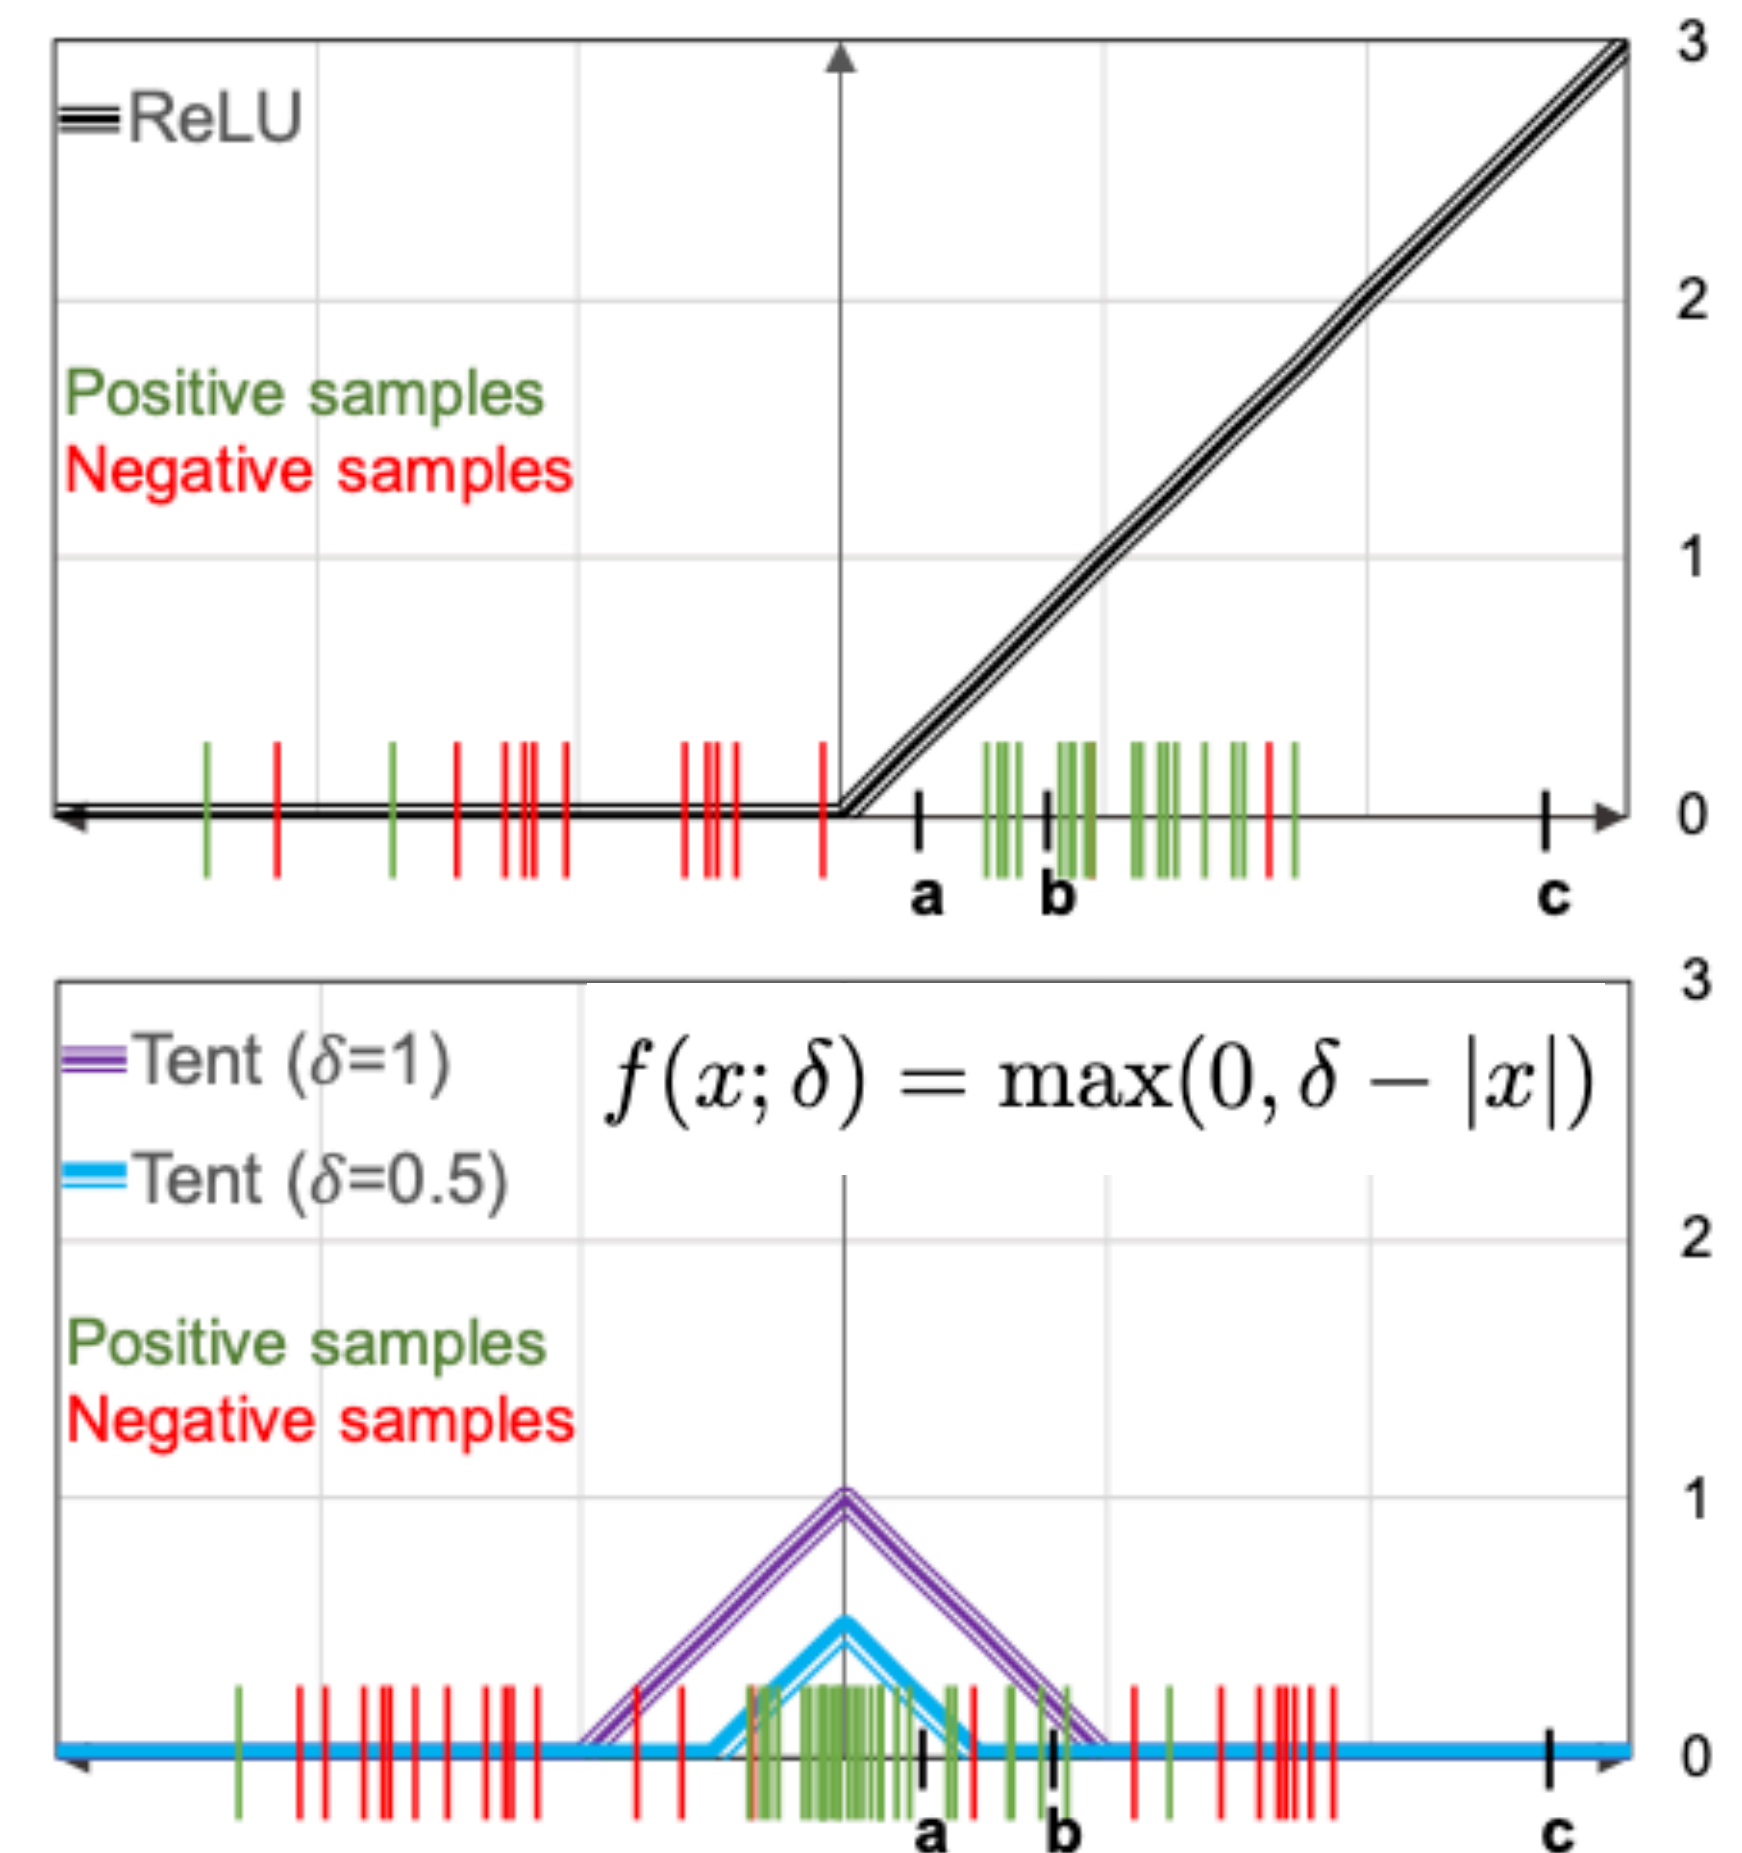


An alternative/auxiliary view



Hypothesis: specific functions in our ML models, like ReLU in NNs are (at least in parts) the culprit - they always produce high confidence far away from the data (Hein et al, “Why ReLU networks yield high confidence predictions far away from the training data and how to mitigate the problem”, CVPR 2019)

Alternative idea: use functions that are bounded and try to determine their “extent” based on the observed data (Rozsa & Boulton, “Improved Adversarial Robustness by Reducing Open Space Risk via Tent Activations”, 2019)



An alternative/auxiliary view



Hypothesis: specific functions in our ML models, like ReLU in NNs are (at least in parts) the culprit - they always produce high confidence far away from the data (Hein et al, “Why ReLU networks yield high confidence predictions far away from the training data and how to mitigate the problem”, CVPR 2019)

Alternative idea: use functions that are bounded and try to determine their “extent” based on the observed data (Rozsa & Boulton, “Improved Adversarial Robustness by Reducing Open Space Risk via Tent Activations”, 2019)




Figure produced by Quentin Delfosse, illustrating ReLU vs Tent activations



Open world learning: combining ideas

Open world learning



In retrospect: although there have been increments, the types of continual learning we have seen so far were indeed in a closed world





Figure from CVPR16 “Statistical Methods for Open Set Recognition” by Scheirer & Boulton, <https://www.wjscheirer.com/misc/openset/cvpr2016-open-set-part3.pdf>

Open world learning



Open world learning tries to “puzzle together” the pieces we have seen so far

“An effective open world recognition system must efficiently perform four tasks: detect unknown, choose which points to label for addition to the model, label the points, and update the model” (Boult et al, “Learning and the Unknown”, AAAI 2019)

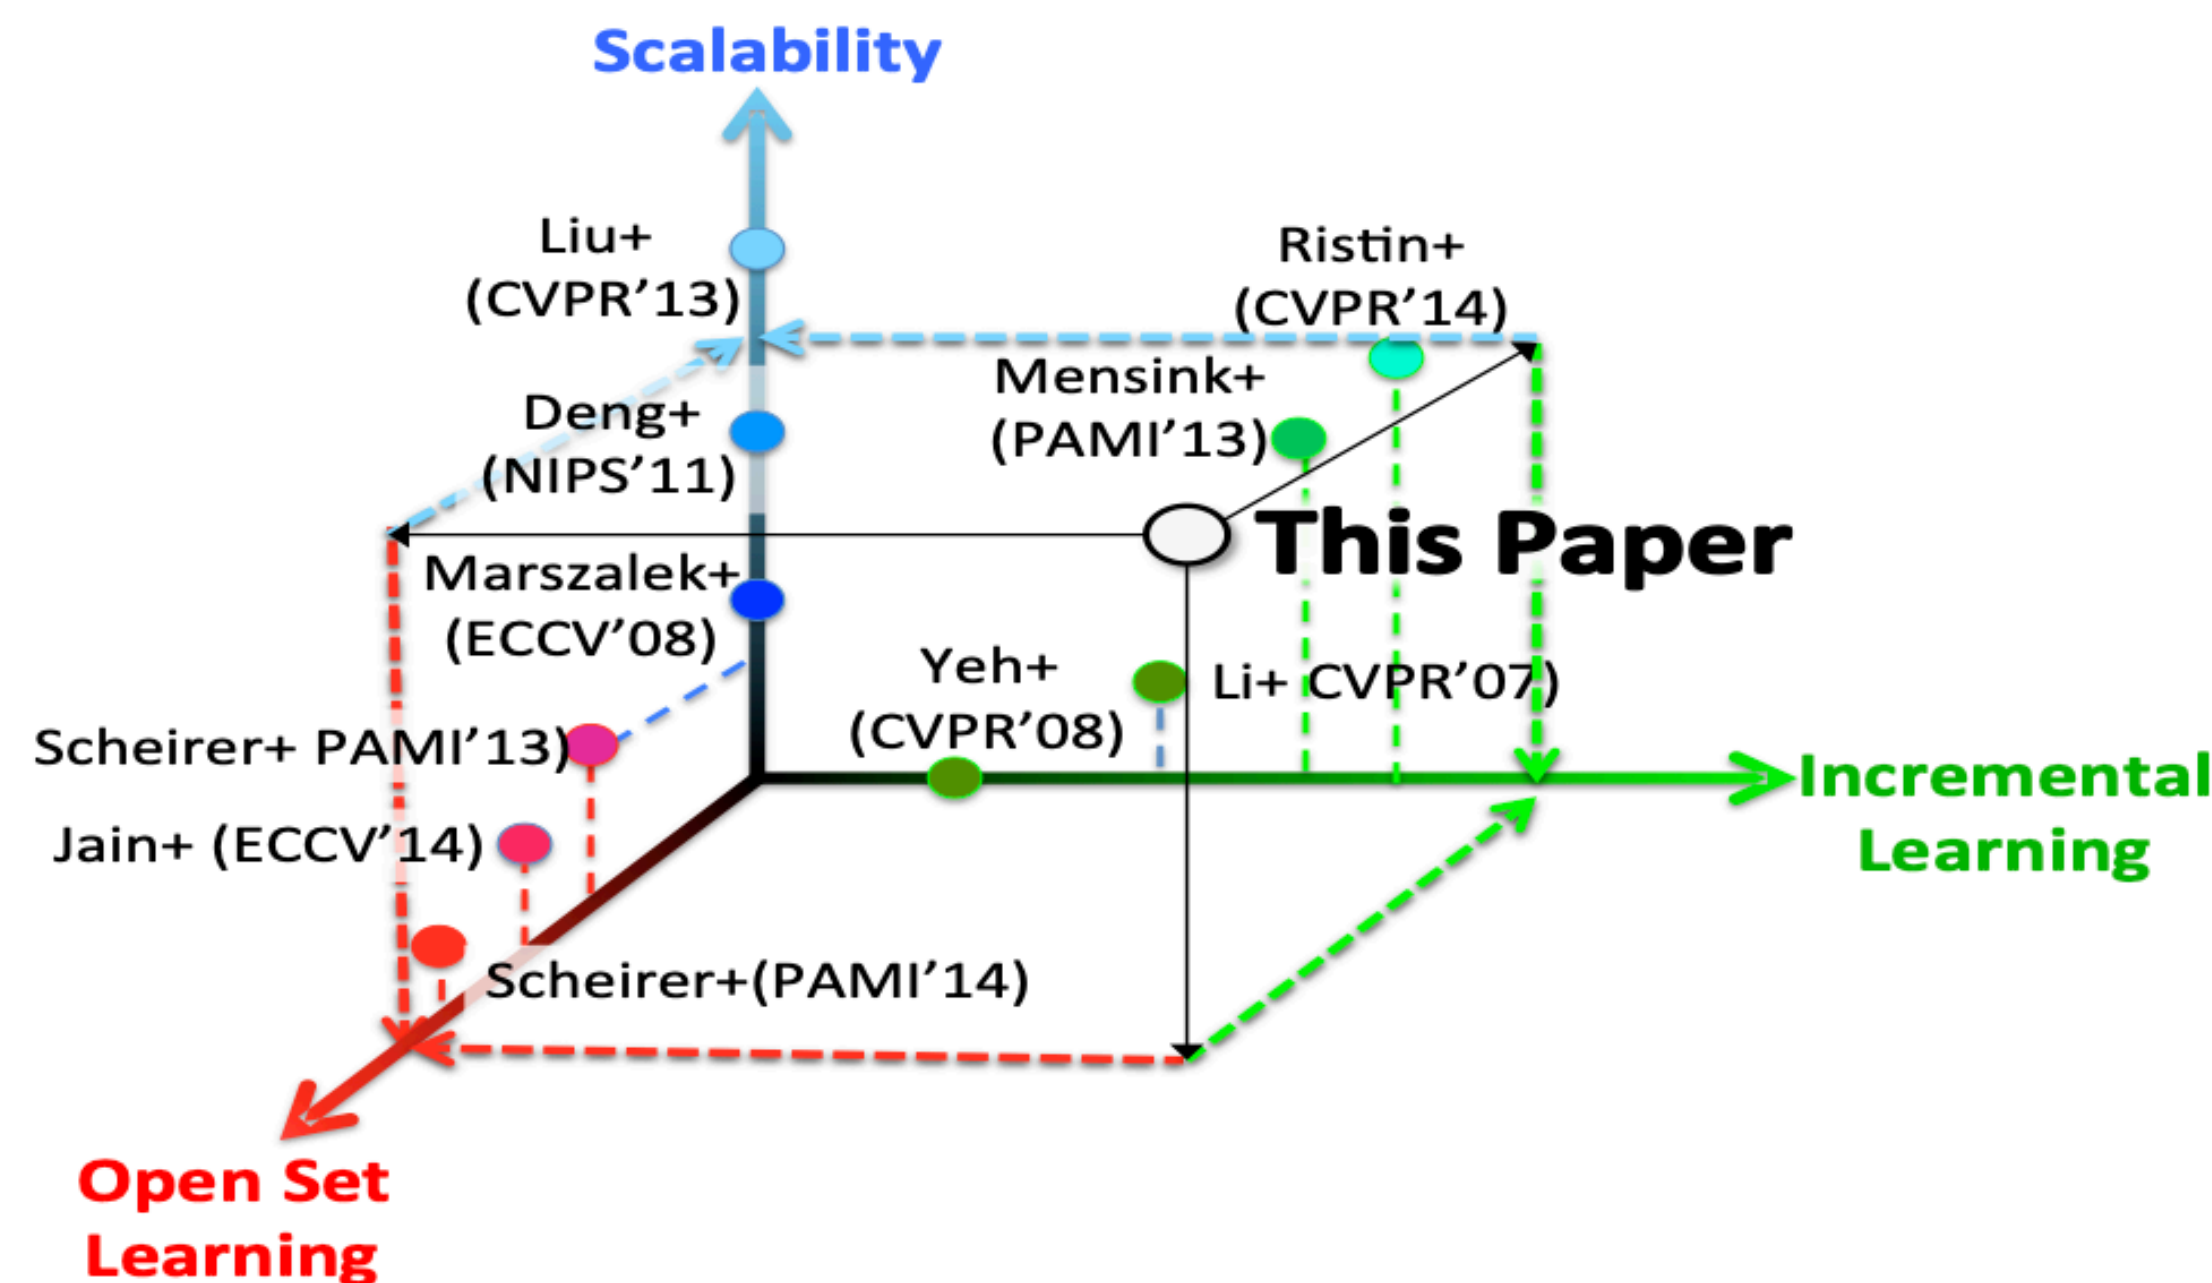


Open world learning




Open world learning tries to “puzzle together” the pieces we have seen so far

“An effective open world recognition system must efficiently perform four tasks: detect unknown, choose which points to label for addition to the model, label the points, and update the model” (Boult et al, “Learning and the Unknown”, AAAI 2019)

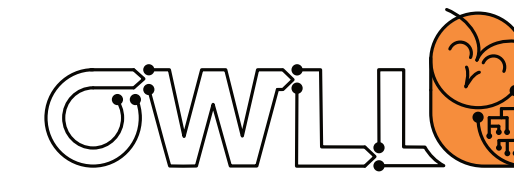


Bendale & Boult, “Towards Open World Recognition”, CVPR 2015



Joseph et al, “Towards Open World Object Detection”, CVPR 2021

Open world learning



We can try to puzzle the pieces together now. As it is very much a cutting-edge research frontier, let's talk about it more in the "frontiers" lecture

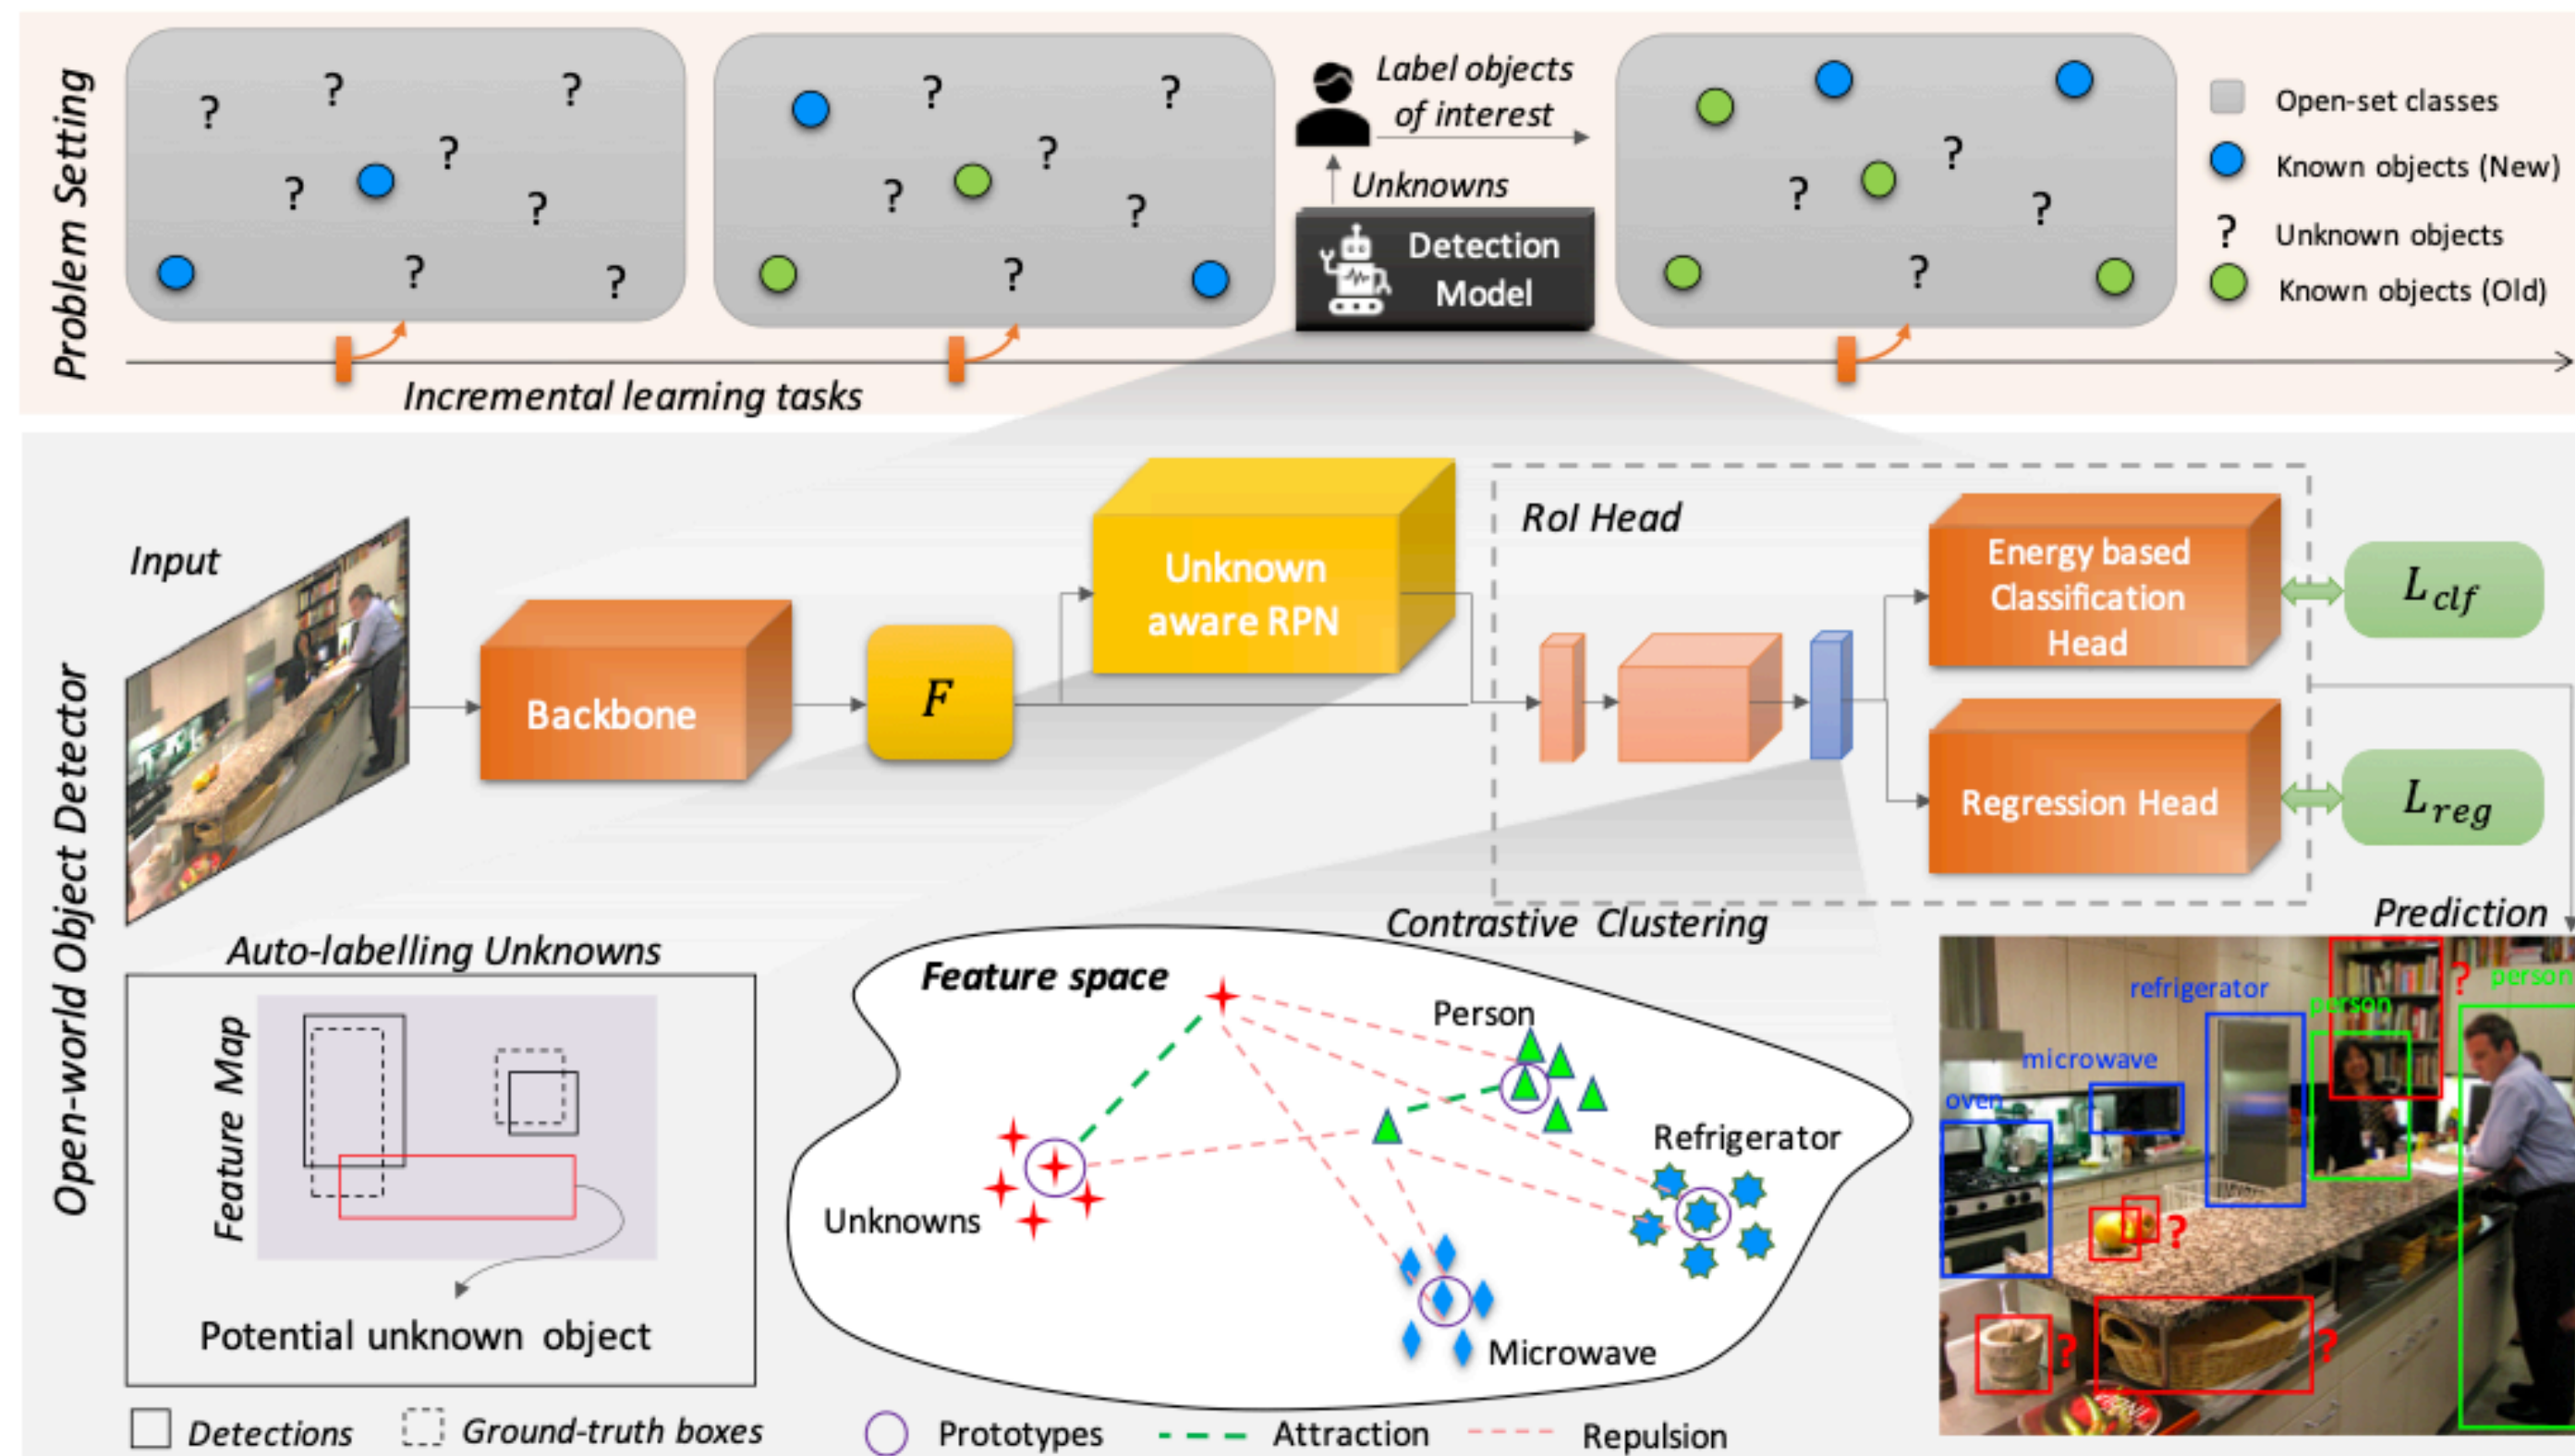


Figure 2: *Approach Overview:* *Top row:* At each incremental learning step, the model identifies unknown objects (denoted by “?”), which are progressively labelled (as blue circles) and added to the existing knowledge base (green circles). *Bottom row:* Our open world object detection model identifies potential unknown objects using an energy-based classification head and the unknown-aware RPN. Further, we perform contrastive learning in the feature space to learn discriminative clusters and can flexibly add new classes in a continual manner without forgetting the previous classes.

Ending on some open questions & a disclaimer:

- Note the “towards” in many of the paper titles
- There is much to be done still: what about avoiding forgetting in addition now?
- Naturally, evaluation gets even more complicated now!
- It’s no longer a question of ML algorithms, perhaps it already was a systems question beforehand, but now it definitely is

Corruptions, adversarial etc.



What about natural corruptions, adversarial attacks etc.?



x

“panda”

57.7% confidence

+ .007 ×



$\text{sign}(\nabla_x J(\theta, x, y))$

“nematode”

8.2% confidence

=



$x +$

$\epsilon \text{sign}(\nabla_x J(\theta, x, y))$

“gibbon”

99.3 % confidence

Corruptions, adversarial etc.



What about natural corruptions, adversarial attacks etc.?


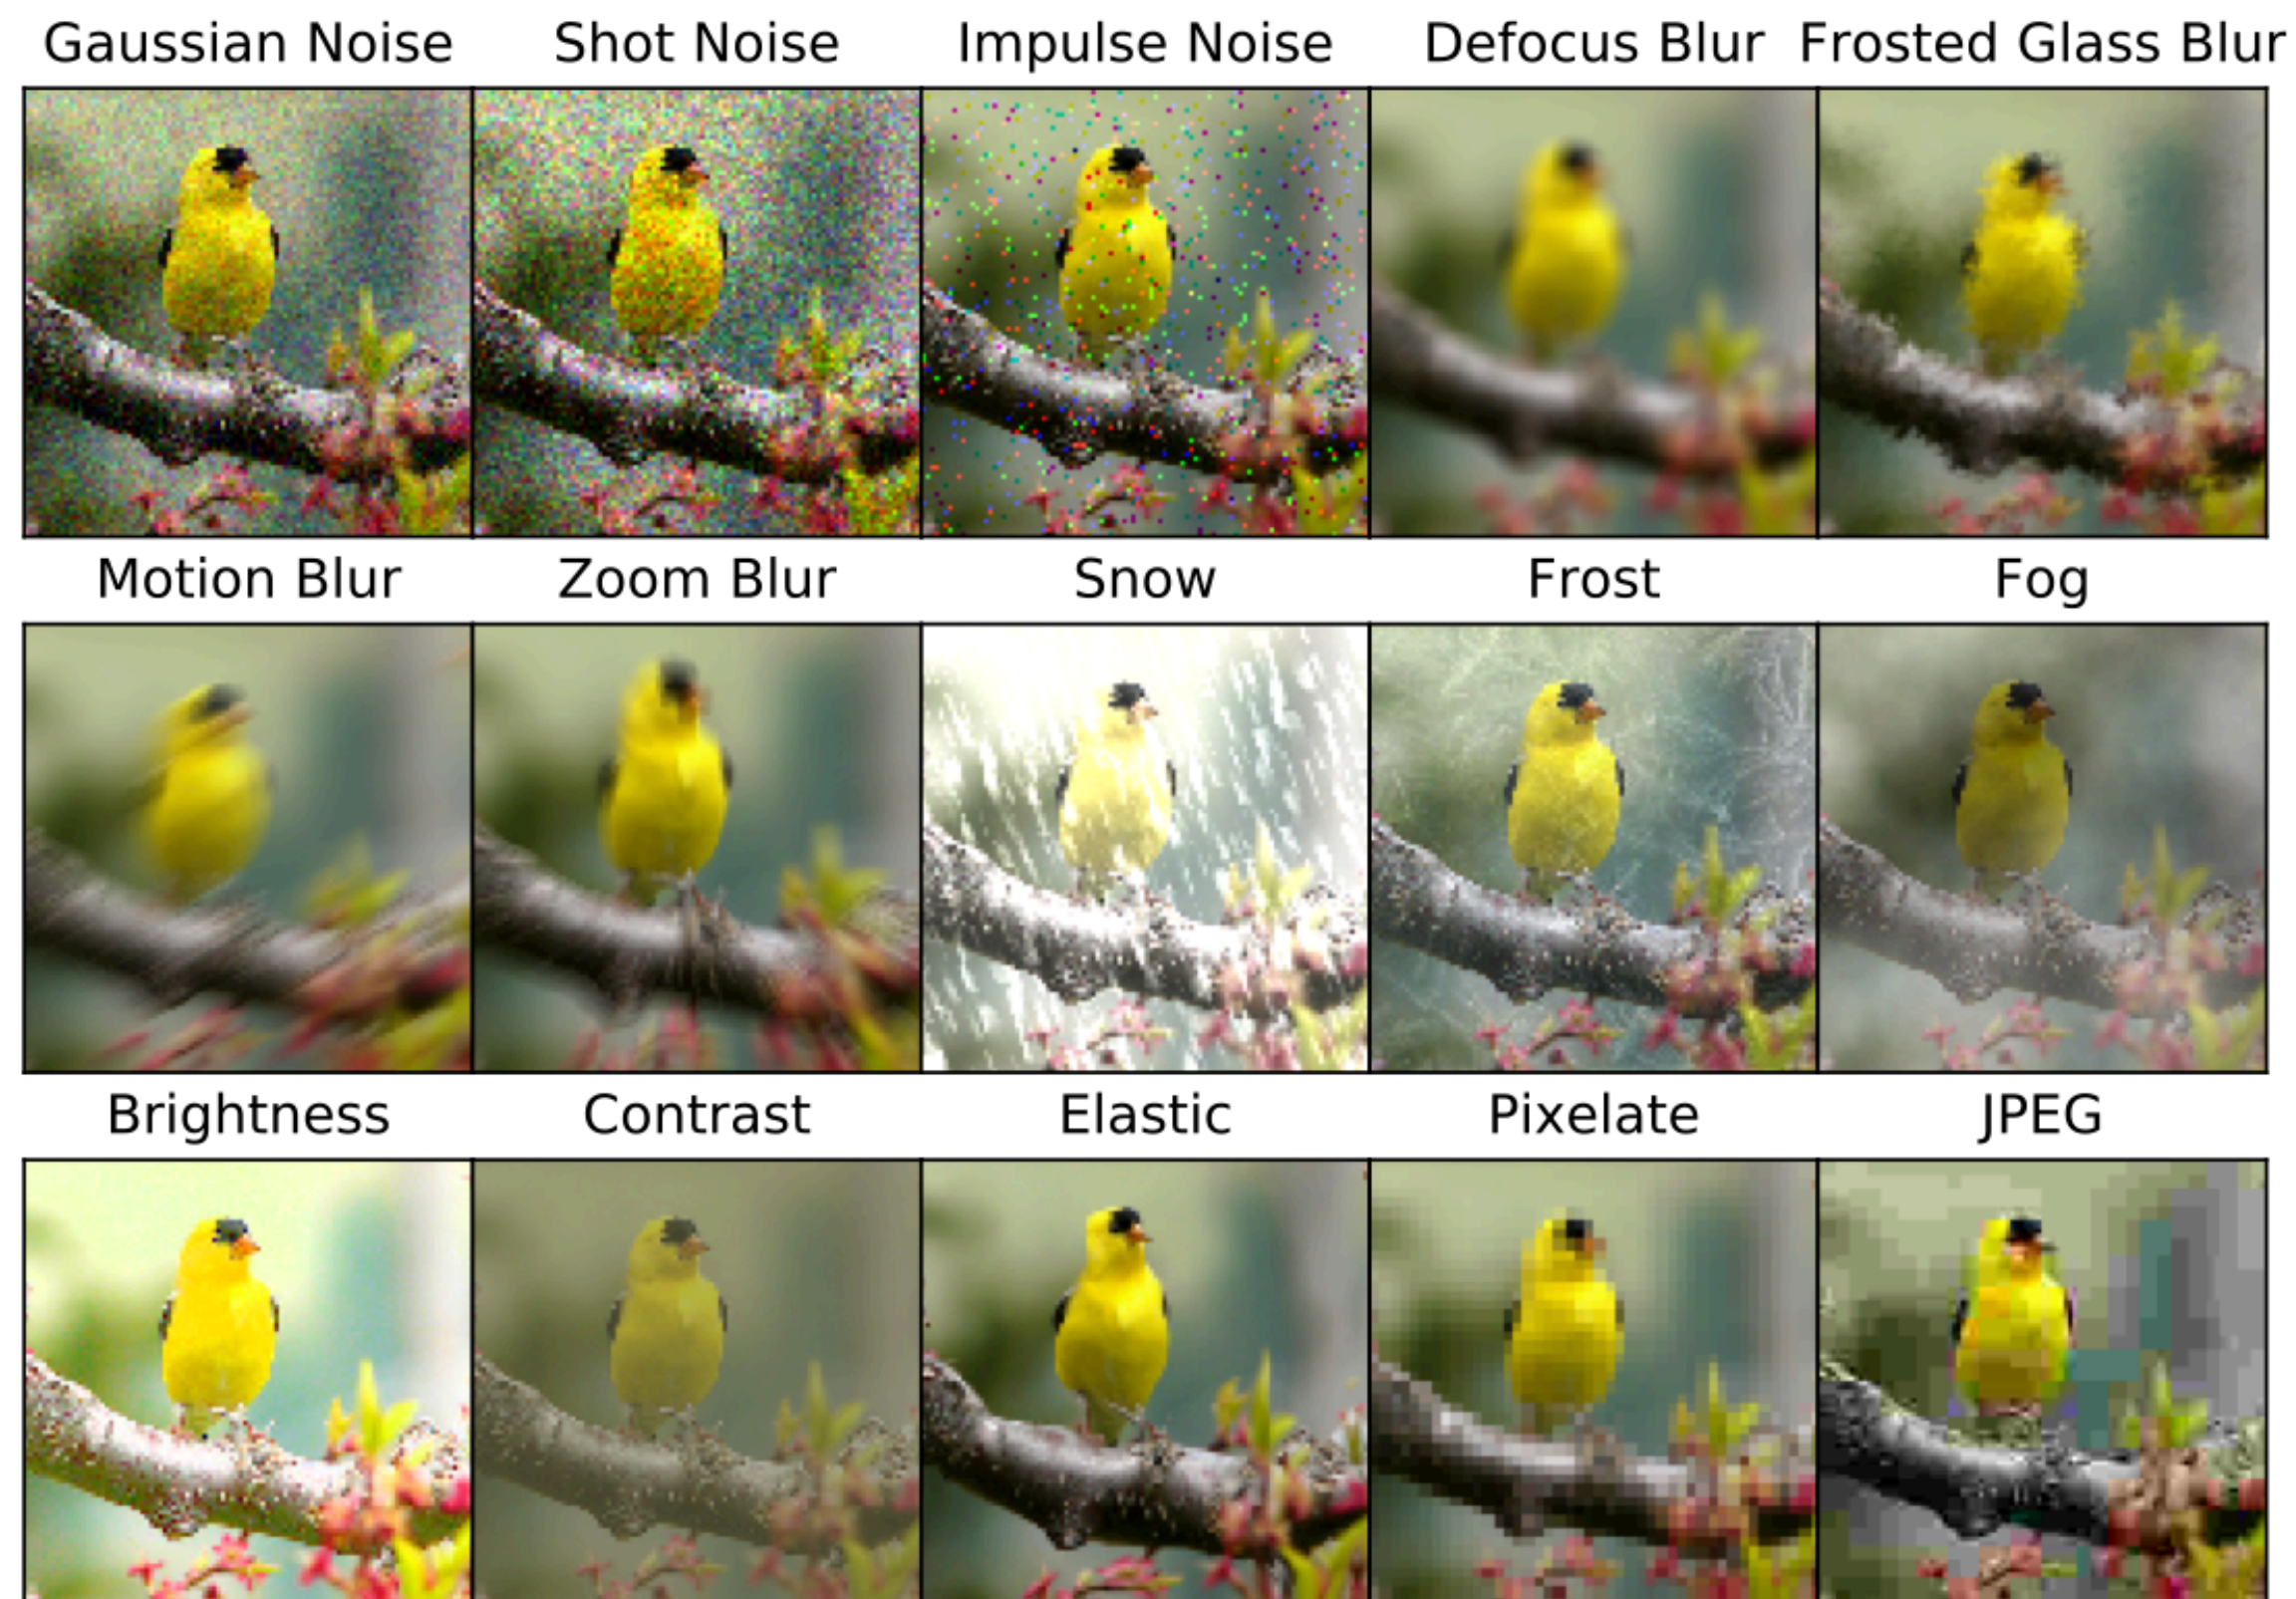


Figure 3: Robustness (mCE) and Relative mCE IMAGENET-C values. Relative mCE values suggest robustness in itself declined from AlexNet to ResNet. “BN” abbreviates Batch Normalization.