## Active Learning \& Beyond

## Stefano Teso

University of Trento, Italy

## Preliminaries

Strategies

## Extensions

Conclusion and Further Reading

## Preliminaries

"Imagine that you are the leader of a colonial expedition from Earth to an extrasolar planet. Luckily, this planet is habitable and has a fair amount of vegetation suitable for feeding your group. Impor- tantly, the most abundant source of food comes from a plant whose fruits are sometimes smooth and round, but sometimes bumpy and irregular."


Figure 1.1: Several alien fruits, which vary in shape from round to irregular.
"The physicians assure you that the shape of a fruit is the only feature that seems related to its safety. The problem, though, is that a wide variety of fruit shapes from these plants exist: almost a continuous range from round to irregular. Since the colony has essential uses for both safe and noxious fruits, you want to be able to classify them as accurately as possible."

Source: [Settles, 2012].

- We know that smoother fruits are (monotonically) safer, but we don't know where to set the threshold.
$\square$ We know that smoother fruits are (monotonically) safer, but we don't know where to set the threshold.

■ In other words, we want to learn a threshold function:

$$
f_{\theta}(x)= \begin{cases}1 & \text { if } x_{\text {irreg }}<\theta \\ -1 & \text { otherwise }\end{cases}
$$

where x are measurements of fruit features and $x_{\text {irreg }}$ captures its shape "irregularity".

Idea: use regular supervised learning

- Collect a large enough training set $S=\{(\mathrm{x}, y)\}$, fit threshold classifier $f_{\theta}$ on $S$
- If maximum \% errors is $\epsilon \in(0,1)$, enough to collect $O\left(\frac{1}{\epsilon}\right)$ examples [Shalev-Shwartz and Ben-David, 2014]


Figure 1.2: Supervised learning for the alien fruits example. Given a set of $\langle x, y\rangle$ instance-label pairs, we want to choose the threshold $\theta^{*}$ that classifies them most accurately.

## Idea: use regular supervised learning

- Collect a large enough training set $S=\{(\mathrm{x}, y)\}$, fit threshold classifier $f_{\theta}$ on $S$
- If maximum \% errors is $\epsilon \in(0,1)$, enough to collect $O\left(\frac{1}{\epsilon}\right)$ examples [Shalev-Shwartz and Ben-David, 2014]


Figure 1.2: Supervised learning for the alien fruits example. Given a set of $\langle x, y\rangle$ instance-label pairs, we want to choose the threshold $\theta^{*}$ that classifies them most accurately.
$\square$ We want to find $\theta$ as quickly and as economically as possible, by requiring fewer tests.Can we do better?

## Key features:

- Fruits x are plentiful and easy to harvest and measure
- Obtaining $y$ incurs a cost: person that eats the fruit may get sick

So we definitely want to minimize the number of needed labels.

Idea: gather large set of unlabeled fruits $U=\left\{\mathbf{x}_{i}\right\}$ and arrange them by roughness.


Figure 1.3: A binary search through the set of ordered, untested alien fruits. By only testing this subset of fruits, we can exponentially reduce costs while achieving the same result. The labels shown in light blue can be inferred, and therefore do not need to be tested.

Then use binary search (as in the illustration) to figure out the threshold $\theta$. This only requires $O\left(\log _{2} \frac{1}{\epsilon}\right)$ tests!

Idea: gather large set of unlabeled fruits $U=\left\{\mathbf{x}_{i}\right\}$ and arrange them by roughness, then use binary search:

| $\epsilon$ | $\frac{1}{\epsilon}$ | $\log _{2} \frac{1}{\epsilon}$ |
| :---: | :---: | :---: |
| 0.1 | 10 | 3.321 |
| 0.001 | 1000 | 9.966 |
| 0.00001 | 100000 | 16.610 |

$\square$ In this (cleverly designed illustrative) scenario, there is an exponential improvement in sample complexity

## Active vs Passive

"The key hypothesis is that if the learner is allowed to choose the data from which it learns - to be active, curious, or exploratory, if you will - it can perform better with less training." [Settles, 2012]

## Preconditions:

- Collecting unlabelled instances x is cheap
- Obtaining their labels $y$ is expensive


## Preconditions:

- Collecting unlabelled instances x is cheap
- Obtaining their labels $y$ is expensive


## Example: Citizen Science

There are tons of images of celestial bodies (think sky surveys). However, in order to undestand what's in an image (is it a spiral galaxy? is it a gravitational lensing effect?) you have to ask a human expert.

## Preconditions:

- Collecting unlabelled instances x is cheap
- Obtaining their labels $y$ is expensive


## Example: Citizen Science

There are tons of images of celestial bodies (think sky surveys). However, in order to undestand what's in an image (is it a spiral galaxy? is it a gravitational lensing effect?) you have to ask a human expert.

## Example: Recommendation

There are millions of products on online catalogues (think Amazon), but in order to discover what are the tastes of a user, you have to actually convince them to score the items. This information is personalized, so this is the only way to obtain supervision.

## Example: Scientific Discovery

■ Adam, the "robot scientist" [King et al., 2009]

## The Automation of Science

Ross D. King, ${ }^{1}{ }^{1 *}$ Jem Rowland, ${ }^{1}$ Stephen G. Oliver, ${ }^{2}$ Michael Young, ${ }^{3}$ Wayne Aubrey, ${ }^{1}$ Emma Byrne, ${ }^{1}$ Maria Liakata, ${ }^{1}$ Magdalena Markham, ${ }^{1}$ Pinar Pir, ${ }^{2}$ Larisa N. Soldatova, ${ }^{1}$ Andrew Sparkes, ${ }^{1}$ Kenneth E. Whelan, ${ }^{1}$ Amanda Clare ${ }^{1}$

The basis of science is the hypothetico-deductive method and the recording of experiments in sufficient detail to enable reproducibility. We report the development of Robot Scientist "Adam," which advances the automation of both. Adam has autonomously generated functional genomics hypotheses about the yeast Saccharomyces cerevisiae and experimentally tested these hypotheses by using laboratory automation. We have confirmed Adam's conclusions through manual experiments. To describe Adam's research, we have developed an ontology and logical language. The resulting formalization involves over 10,000 different research units in a nested treelike structure, 10 levels deep, that relates the 6.6 million biomass measurements to their logical description. This formalization describes how a machine contributed to scientific knowledge.

The learner obtains labels by operating an automated testing machine.

## Example: Scientific Discovery



Similar strategies used in chemical engineering, material engineering, etc.

A summary of frequently used terms:

- Instances $\mathbf{x} \in \mathbb{R}^{d}$ are unlabelled d-dimensional vectors of observations
- Examples $z=(x, y)$ are instances annotated by a label $y \in\{0,1\}$ or $y \in\{1, \ldots, c\}$
- A classifier $f: \mathbb{R}^{d} \rightarrow\{0,1\}$ maps instances to labels
- $\mathcal{F}=\left\{f_{\theta}\right\}$ is a family of classifiers parameterized by $\theta$
$\square$ The meaning of $\theta$ depends on the model class, e.g., for neural nets with a fixed architecture, $\theta$ represents their weights; for random forests, $\theta$ represents the structure and leaves of all trees.


## Assumptions

- We assume the data to be distributed according to a ground-truth data generating process $p^{*}(Y, \mathbf{X})$, and use $p^{*}(Y \mid \mathbf{X})$ and $p^{*}(\mathbf{X})$ to indicate the prior and class likelihood of this process, i.e.,

$$
\begin{equation*}
p^{*}(Y, \mathbf{X}) \equiv p^{*}(Y \mid \mathbf{X}) \cdot p^{*}(\mathbf{X}) \tag{1}
\end{equation*}
$$

The data may be corrupted by label noise, but we explicitly avoid adversarial settings.

- We also assume to deal with probabilistic classifiers:

$$
\begin{equation*}
p_{\theta}(Y=y \mid \mathbf{X}=\mathbf{x}) \tag{2}
\end{equation*}
$$

Predictions are computed as:

$$
\begin{equation*}
f_{\theta}(\mathrm{x}):=\underset{y=1, \ldots, c}{\operatorname{argmax}} p_{\theta}(Y=y \mid \mathbf{X}=\mathrm{x}) \tag{3}
\end{equation*}
$$

This includes, e.g., logistic regression, Gaussian processes, and neural nets with a top softmax activation.
$\square$ Predictors that only output a score (e.g., Support Vector Machines output the distance to the separating hyperplane) can often be adapted (via, e.g., calibration or Platt scaling). This is not necessary to leverage AL.

## The Labeling Oracle

$\square$ The labeling oracle label : $\mathbb{R}^{d} \rightarrow\{0,1\}$ returns:

$$
\begin{equation*}
\operatorname{label}(\mathrm{x}):=\underset{y \in\{0,1\}}{\operatorname{argmax}} p^{*}(Y=y \mid \mathbf{X}=\mathrm{x}) \tag{4}
\end{equation*}
$$

where $p^{*}(Y \mid \mathbf{X})$ is the ground-truth distribution.
$\square$ Invoking the oracle comes at a cost, which is usually non-negligible, instance- and class-dependent, and possibly unknown [Herde et al., 2021]. ${ }^{1}$

For simplicity, we assume the cost to be unitary: querying any instance $\mathbf{x} \in \mathbb{R}^{d}$ costs the same.

[^0]
(b) pool-based sampling


Active Learning (Pool-based). Given:

- a family of classifiers $\mathcal{F}$,
- a set of unlabelled instances $U=\left\{\mathrm{x}_{1}, \ldots, \mathrm{x}_{m}\right\} \subseteq \mathbb{R}^{d}$ sampled i.i.d. from $p^{*}(\mathbf{X})$,
- a (costly) labeling oracle label : $\mathbb{R}^{d} \rightarrow\{0,1\}$,

Find a classifier $\widehat{f} \in \mathcal{F}$ that achieves low risk on $p^{*}(\mathbf{X}, Y)$ while keeping annot. cost low

(a) query synthesis


Active Learning (Query Synthesis). Given:

- a family of classifiers $\mathcal{F}$,
- a generator of instances synthesize(region) $\mapsto \mathbf{x}$,
- a (costly) labeling oracle label : $\mathbb{R}^{d} \rightarrow\{0,1\}$,

Find a classifier $\widehat{f} \in \mathcal{F}$ that achieves low risk on $p^{*}(\mathbf{X}, Y)$ while keeping annot. cost low


## Active Learning (Selective Sampling). Given:

- a family of classifiers $\mathcal{F}$,
- a sequence of unlabelled instances $\mathrm{x}_{1}, \mathrm{x}_{2}, \mathrm{x}_{3}, \ldots$,
- a (costly) labeling oracle label : $\mathbb{R}^{d} \rightarrow\{0,1\}$

Find a classifier $\widehat{f}_{t} \in \mathcal{F}$ that achieves low risk on future data $\mathrm{x}_{t+1}, \mathrm{x}_{t+2}, \ldots$ while keeping annot. cost low

## Query Sampling vs. Query Synthesis



- Left to right:
- Pool-based: moderate control over queries, requires memory to store $U$
- Query synthesis: maximum control over queries, can generate uninterpretable queries [Baum and Lang, 1992], although deep generative models can help somehow [Nguyen et al., 2016].
- Selective sampling: little control over the distribution of queries, often solved under tight memory constraints (online learning)

■ We will focus on pool-based AL.

## Strategies

## Template

```
Input: models \(\mathcal{F}\), examples \(L\), pool \(U\), query budget \(T \geq 1\)
Output: selected model \(f \in \mathcal{F}\)
    1: \(f \leftarrow \operatorname{fit}(\mathcal{F}, L) \quad \triangleright\) initialize the model
    2: for \(t=1,2, \ldots, T\) do \(\quad \triangleright\) until the budget is exhausted
    \(\mathbf{x} \leftarrow \operatorname{argmax}_{x \in U} \operatorname{acq}(f, \mathbf{x}) \quad \triangleright\) select a query instance
    obtain label \(y\) of \(x\) from annotator
        \(U \leftarrow U \backslash\{\mathbf{x}\} \quad \triangleright\) remove unlabeled instance from pool
        \(L \leftarrow L \cup\{(\mathrm{x}, y)\} \quad \triangleright\) update training set
\(\underset{\sim}{f} \leftarrow \operatorname{fit}(\mathcal{F}, L) \quad \triangleright\) update the model
```

$\square$ fit performs training (e.g., trains for a fixed \# of epochs)
$\square$ acq scores instances based on their "informativeness"
$\square$ What instance $\mathrm{x} \in U$ should be selected so to convey as much information as possible to $f$ ?

## Uncertainty Sampling

What's the point of asking the label of instances on which the classifier is already certain? ${ }^{2}$

■ Left: two Gaussians (40 points each)
■ Middle: picking points completely at random (ignoring the class label!)

- Right: picking points based on uncertainty

[^1]
## Uncertainty Sampling

$\square$ Idea: pick $\mathrm{x} \in U$ on which the classifier is most uncertain.

## Uncertainty Sampling

$\square$ Idea: pick $\mathrm{x} \in U$ on which the classifier is most uncertain.


Figure 2.1: The binary search from Figure 1.3, re-interpreted as an uncertainty sampling approach. The best instance to query is deemed to be the one closest to the threshold $\theta$.

## Uncertainty Sampling

$\square$ Idea: pick $\mathrm{x} \in U$ on which the classifier is most uncertain.


Figure 2.1: The binary search from Figure 1.3, re-interpreted as an uncertainty sampling approach. The best instance to query is deemed to be the one closest to the threshold $\theta$.

- How should uncertainty be defined?


## Uncertainty Sampling

■ Define uncertainty using the confidence, i.e., distance from certainty:

$$
\begin{equation*}
\operatorname{acq}(\theta, \mathbf{x}):=1-p_{\theta}(\hat{y} \mid \mathbf{x}) \tag{5}
\end{equation*}
$$

where $\hat{y}$ is the predicted label:

$$
\begin{equation*}
\hat{y}:=f_{\theta}(\mathrm{x})=\underset{y}{\operatorname{argmax}} p_{\theta}(y \mid \mathrm{x}) \tag{6}
\end{equation*}
$$

## Uncertainty Sampling

- Define uncertainty using the margin, i.e., difference in (conditional) likelihood:

$$
\begin{equation*}
\operatorname{acq}(\theta, \mathbf{x}):=p_{\theta}\left(\hat{y}^{\prime} \mid \mathbf{x}\right)-p_{\theta}(\hat{y} \mid \mathbf{x}) \tag{7}
\end{equation*}
$$

where $\hat{y}$ is the predicted label and $\hat{y}^{\prime}$ is the 2nd best label:

$$
\begin{gather*}
\hat{y}=\underset{y}{\operatorname{argmax}} p_{\theta}(y \mid \mathbf{x})  \tag{8}\\
\hat{y}^{\prime}:=\underset{y \neq \hat{y}}{\operatorname{argmax}} p_{\theta}(y \mid \mathbf{x}) \tag{9}
\end{gather*}
$$

## Uncertainty Sampling

$\square$ Define uncertainty using the Shannon entropy of the label:

$$
\begin{equation*}
\operatorname{acq}(\theta, \mathbf{x}):=H_{\theta}(Y \mid \mathbf{X}=\mathbf{x}) \tag{10}
\end{equation*}
$$

where $H_{\theta}$ is defined as:

$$
\begin{equation*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x}):=-\sum_{y \in[c]} p_{\theta}(y \mid \mathbf{x}) \log _{2} p_{\theta}(y \mid \mathbf{x}) \tag{11}
\end{equation*}
$$

- It achieves a minimum on dead certain distributions:

$$
p_{\theta}(Y \mid \mathrm{x})=(0,1,0, \ldots, 0)
$$

- and a maximum on the uniform distribution:

$$
p_{\theta}(Y \mid \mathbf{x})=\left(\frac{1}{c}, \cdots, \frac{1}{c}\right)
$$

## Confidence vs. Margin vs. Entropy

- Left: confidence considers prob. of top class only
- Middle: margin considers prob. of top \& runner up classes
- Right: entropy considers prob. of all classes

If $c=2$, they are equivalent. If $c>2$, no obvious best choice, it really depends on the task and loss (e.g., crossentropy vs.accuracy)

## Confidence vs. Margin vs. Entropy

Example: for classifiers with a sigmoid top layer:

uncertainty depends on distance from separating hyperplane of predicted vs. top two vs. all classes

## Uncertainty Sampling

- Uncertainty sampling is very easy to implement.

■ Margin \& Confidence can be defined even in terms of unnormalized scores.
■ Usually performs reasonably well (though not optimally) in practice: a useful baseline/starting point.

## Example: Structured Output

Consider an LSTM that takes a sequence of MNIST images $X=\left[\mathrm{x}_{1}, \ldots, \mathrm{x}_{n}\right]$ that composes a word and outputs the word itself $\mathbf{y}=\left(y_{1}, \ldots, y_{n}\right)$.

- Computing the most likely output $\hat{\mathbf{y}}$ can be done efficiently.
- Computing the entropy amounts to:

$$
\begin{equation*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x}):=-\sum_{\mathbf{y} \in\{1, \ldots, 26\}^{n}} p_{\theta}(\mathbf{y} \mid X) \log _{2} p_{\theta}(\mathbf{y} \mid X) \tag{12}
\end{equation*}
$$

This involves summing over $26^{n}$ possible outputs, which takes time exponential in $n$.

■ Computing the most likely output can be NP-hard. For instance, if $y$ is molecular structure that mast satisfy specific hard constraints (chemical validity), then finding the best structure amounts to solving a hard combinatorial problem.

Hence, the confidence and margin can also be very hard.

## Uncertainty in Regression

$\square$ When considering regression models with $Y \in \mathbb{R}$, uncertainty at x can be implemented as differential entropy:

$$
\begin{align*}
H_{\theta}(Y \mid \mathbf{X}=\mathrm{x}) & :=\mathbb{E}\left[-\log _{2} p_{\theta}(y \mid \mathrm{x}) \mid \mathrm{x}\right]  \tag{13}\\
& =-\int_{\mathbb{R}} p_{\theta}(y \mid \mathrm{x}) \log _{2} p_{\theta}(y \mid \mathrm{x}) \tag{14}
\end{align*}
$$

## Uncertainty in Regression

$\square$ When considering regression models with $Y \in \mathbb{R}$, uncertainty at x can be implemented as differential entropy:

$$
\begin{align*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x}) & :=\mathbb{E}\left[-\log _{2} p_{\theta}(y \mid \mathbf{x}) \mid \mathbf{x}\right]  \tag{13}\\
& =-\int_{\mathbb{R}} p_{\theta}(y \mid \mathbf{x}) \log _{2} p_{\theta}(y \mid \mathbf{x}) \tag{14}
\end{align*}
$$

- As an alternative heuristic, use the variance:

$$
\begin{align*}
\operatorname{Var}_{\theta}(Y \mid \mathbf{x}) & :=\mathbb{E}[(Y-\underbrace{\mathbb{E}[Y \mid \mathrm{x}]}_{\mu_{\theta}(Y \mid \mathbf{x}):=})^{2} \mid \mathbf{x}]  \tag{15}\\
& =\int_{\mathbb{R}}\left(y-\mu_{\theta}(Y \mid \mathbf{x})\right)^{2} p_{\theta}(y \mid \mathbf{x}) d y  \tag{16}\\
\mu_{\theta}(Y \mid \mathbf{x}) & =\int_{\mathbb{R}} y p_{\theta}(y \mid \mathbf{x}) d y \tag{17}
\end{align*}
$$

## Uncertainty in Regression

$\square$ When considering regression models with $Y \in \mathbb{R}$, uncertainty at x can be implemented as differential entropy:

$$
\begin{align*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x}) & :=\mathbb{E}\left[-\log _{2} p_{\theta}(y \mid \mathbf{x}) \mid \mathbf{x}\right]  \tag{13}\\
& =-\int_{\mathbb{R}} p_{\theta}(y \mid \mathbf{x}) \log _{2} p_{\theta}(y \mid \mathbf{x}) \tag{14}
\end{align*}
$$

$\square$ As an alternative heuristic, use the variance:

$$
\begin{align*}
\operatorname{Var}_{\theta}(Y \mid \mathbf{x}) & :=\mathbb{E}[(Y-\underbrace{\mathbb{E}[Y \mid \mathbf{x}]}_{\mu_{\theta}(Y \mid \mathbf{x}):=})^{2} \mid \mathbf{x}]  \tag{15}\\
& =\int_{\mathbb{R}}\left(y-\mu_{\theta}(Y \mid \mathbf{x})\right)^{2} p_{\theta}(y \mid \mathbf{x}) d y  \tag{16}\\
\mu_{\theta}(Y \mid \mathbf{x}) & =\int_{\mathbb{R}} y p_{\theta}(y \mid \mathbf{x}) d y \tag{17}
\end{align*}
$$

How to compute them?

## Uncertainty in Regression

$\square$ Differential entropy and variance:

$$
\begin{equation*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x})=-\int_{\mathbb{R}} p_{\theta}(y \mid \mathbf{x}) \log _{2} p_{\theta}(y \mid \mathbf{x}) \quad \operatorname{Var}_{\theta}(Y \mid \mathbf{x})=\int_{\mathbb{R}}\left(p_{\theta}(y \mid \mathbf{x})-\mu_{\theta}(Y \mid \mathbf{x})\right) d y \tag{18}
\end{equation*}
$$

[^2]
## Uncertainty in Regression

$\square$ Differential entropy and variance:

$$
\begin{equation*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x})=-\int_{\mathbb{R}} p_{\theta}(y \mid \mathbf{x}) \log _{2} p_{\theta}(y \mid \mathbf{x}) \quad \operatorname{Var}_{\theta}(Y \mid \mathbf{x})=\int_{\mathbb{R}}\left(p_{\theta}(y \mid \mathbf{x})-\mu_{\theta}(Y \mid \mathbf{x})\right) d y \tag{18}
\end{equation*}
$$

Both expensive to compute for general models, can approximate via quadrature or sampling, but closed-form solutions exist for some models (e.g., Gaussian Processes and NNs with a Gaussian output)

[^3]
## Uncertainty in Regression

Differential entropy and variance:

$$
\begin{equation*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x})=-\int_{\mathbb{R}} p_{\theta}(y \mid \mathbf{x}) \log _{2} p_{\theta}(y \mid \mathbf{x}) \quad \operatorname{Var}_{\theta}(Y \mid \mathbf{x})=\int_{\mathbb{R}}\left(p_{\theta}(y \mid \mathbf{x})-\mu_{\theta}(Y \mid \mathbf{x})\right) d y \tag{18}
\end{equation*}
$$

Both expensive to compute for general models, can approximate via quadrature or sampling, but closed-form solutions exist for some models (e.g., Gaussian Processes and NNs with a Gaussian output)

## Example: 1-dimensional Gaussian Output

Consider one-dimensional output $y \in \mathbb{R}$ and a neural net:

$$
\begin{equation*}
n n: \mathrm{x} \mapsto(\mu, \sigma), \quad y \sim \mathcal{N}(\mu, \sigma) \tag{19}
\end{equation*}
$$

In this case, it is well known ${ }^{3}$ that:

$$
\begin{equation*}
\operatorname{Var}_{\theta}(Y \mid \mathbf{x})=\sigma^{2}, \quad H_{\theta}(Y \mid \mathbf{x})=\frac{1}{2} \log \left(2 \pi \sigma^{2}\right)+\frac{1}{2} \tag{20}
\end{equation*}
$$

Notice that $\operatorname{Var}_{\theta}(Y \mid x) \propto \exp H_{\theta}(Y \mid x)$, so they change monotonically.

[^4]
## Uncertainty in Regression

Differential entropy and variance:

$$
\begin{equation*}
H_{\theta}(Y \mid \mathbf{X}=\mathbf{x})=-\int_{\mathbb{R}} p_{\theta}(y \mid \mathbf{x}) \log _{2} p_{\theta}(y \mid \mathbf{x}) \quad \operatorname{Var}_{\theta}(Y \mid \mathbf{x})=\int_{\mathbb{R}}\left(p_{\theta}(y \mid \mathbf{x})-\mu_{\theta}(Y \mid \mathbf{x})\right) d y \tag{21}
\end{equation*}
$$

Both expensive to compute for general models, can approximate via quadrature or sampling, but closed-form solutions exist for some models (e.g., Gaussian Processes and NNs with a Gaussian output)

## Example: k-dimensional Gaussian Output

Consider one-dimensional output $y \in \mathbb{R}^{k}$ and a neural net:

$$
\begin{equation*}
n n: \mathbf{x} \mapsto(\boldsymbol{\mu}, S), \quad \Sigma \leftarrow S S^{T}, \quad \mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma) \tag{22}
\end{equation*}
$$

with $\Sigma$ PSD by construction. In this case, it is well known ${ }^{4}$ that:

$$
\begin{equation*}
\operatorname{Var}_{\theta}(Y \mid \mathbf{x}) \propto \operatorname{tr} \Sigma \quad H_{\theta}(Y \mid \mathbf{x}) \propto \log \operatorname{det} \Sigma \tag{23}
\end{equation*}
$$

where the trace is cheap to compute but the determinant is more challenging.
${ }^{4}$ See https://en.wikipedia.org/wiki/Multivariate_normal_distribution.

## Illustration



- Synthetic dataset: 25 clusters of red points arranged in a $5 \times 5$ grid, surrounded by a sea of blue points

$\square$ After $\mathbf{1 0}$ iterations of uncertainty sampling.

$\square$ After $\mathbf{7 0}$ iterations of uncertainty sampling.

$\square$ After $\mathbf{1 4 0}$ iterations of uncertainty sampling. Not nice!
- Discriminative models are over-confident:


Uncertainty does not decrease with distance from the training set.

■ Bayesian generative models not so much:

(a), posterior

Uncertainty does decrease with distance from the training set.

## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.

## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.

- Aleatoric uncertainty ("random") captures how much we can trust the supervision itself. It cannot be decreased. (left)


## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.
$\square$ Aleatoric uncertainty ("random") captures how much we can trust the supervision itself. It cannot be decreased. (left)

■ Epistemic uncertainty ("relating to knowledge") captures how little we know about the world. This reflects on uncertainty on the choice of $\theta$. It decreases by acquiring more data. (right)

## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.

■ Aleatoric uncertainty ("random") captures how much we can trust the supervision itself. It cannot be decreased. (left)

■ Epistemic uncertainty ("relating to knowledge") captures how little we know about the world. This reflects on uncertainty on the choice of $\theta$. It decreases by acquiring more data. (right)

- There isn't much point in trying to reduce aleatoric uncertainty in AL [Sharma and Bilgic, 2017]


## Uncertainty Sampling for Streaming Data

```
Input: models \mathcal{F}}\mathrm{ , bootstrap training set L, threshold }
Output: selected model }f\in\mathcal{F
:f\leftarrow\operatorname{fit}(\mathcal{F},L)\quad\triangleright initialize the model
for }t=1,2,3,\ldots,\mathrm{ do
    receive instance x
        if unc(f,\mathbf{x})>\tau\mathrm{ then }\quad\triangleright\mathrm{ if }f\mathrm{ is uncertain about }\mathbf{x}
            obtain label y of x from annotator
            L\leftarrowL\cup{(\mathbf{x},y)}\quad\triangleright update training set
        f\leftarrow\operatorname{fit}(\mathcal{F},L)\quad\triangleright update the model
return f
```

- The tricky bit is setting $\tau$. Many algorithms update it dynamically by, e.g, starting from a large $\tau$ and lowering it as new data is received and the model improves

■ For some problems, US converges to the right thing - because it is uncertain enough

■ If you are unluckly, US becomes over-confident: in this example, the model becomes confident that the regions inside the black blob cannot be white, so it does not sample them and converges to the wrong shape.

- Uncertainty sampling is quite heuristic. Are there more principled approaches?


## Version Space

■ Consider a hypothesis space $\mathcal{F}=\left\{f_{\theta}: \mathbb{R}^{d} \rightarrow[c]\right\}$ and a data set $S=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}$

## Version Space

- Consider a hypothesis space $\mathcal{F}=\left\{f_{\theta}: \mathbb{R}^{d} \rightarrow[c]\right\}$ and a data set $S=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}$


## Consistency

A hypothesis $f \in \mathcal{F}$ is consistent with $S$, written $f \models S$, iff it makes zero mistakes on it, that is:

$$
\begin{equation*}
(f \models S) \quad \Longleftrightarrow \quad\left(\sum_{(\mathrm{x}, y) \in S} \mathbb{1}\{f(\mathrm{x}) \neq y\}\right)=0 \tag{24}
\end{equation*}
$$

## Version Space

■ Consider a hypothesis space $\mathcal{F}=\left\{f_{\theta}: \mathbb{R}^{d} \rightarrow[c]\right\}$ and a data set $S=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}$

## Consistency

A hypothesis $f \in \mathcal{F}$ is consistent with $S$, written $f \models S$, iff it makes zero mistakes on it, that is:

$$
\begin{equation*}
(f \models S) \quad \Longleftrightarrow \quad\left(\sum_{(\mathrm{x}, y) \in S} \mathbb{1}\{f(\mathrm{x}) \neq y\}\right)=0 \tag{24}
\end{equation*}
$$

## Version Space

The version space $V S_{\mathcal{F}}(S)$ of $\mathcal{F}$ given $S$ is the set of hypotheses $f \in \mathcal{F}$ that is consistent with $S$, that is:

$$
\begin{equation*}
V S_{\mathcal{F}}(S)=\{f \in \mathcal{F}: f \models S\} \tag{25}
\end{equation*}
$$

## Version Space

■ Consider a hypothesis space $\mathcal{F}=\left\{f_{\theta}: \mathbb{R}^{d} \rightarrow[c]\right\}$ and a data set $S=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}$

## Consistency

A hypothesis $f \in \mathcal{F}$ is consistent with $S$, written $f \models S$, iff it makes zero mistakes on it, that is:

$$
\begin{equation*}
(f \models S) \quad \Longleftrightarrow \quad\left(\sum_{(\mathrm{x}, y) \in S} \mathbb{1}\{f(\mathrm{x}) \neq y\}\right)=0 \tag{24}
\end{equation*}
$$

## Version Space

The version space $V S_{\mathcal{F}}(S)$ of $\mathcal{F}$ given $S$ is the set of hypotheses $f \in \mathcal{F}$ that is consistent with $S$, that is:

$$
\begin{equation*}
V S_{\mathcal{F}}(S)=\{f \in \mathcal{F}: f \models S\} \tag{25}
\end{equation*}
$$

$\square V S_{\mathcal{F}}(L)$ contains those hypotheses that have not yet been ruled out by the acquired examples $L$ : they are all just as good w.r.t. $S$

## Version Space

## Version Space

The version space $V S_{\mathcal{F}}(S)$ of $\mathcal{F}$ given $S$ is the set of hypotheses $f \in \mathcal{F}$ that is consistent with $S$, that is:

$$
\begin{equation*}
V S_{\mathcal{F}}(S)=\{f \in \mathcal{F}: f \models S\} \tag{26}
\end{equation*}
$$

■ If $S$ is not separable w.r.t. $\mathcal{F}$, i.e., if there is no hypothesis $f \in \mathcal{F}$ that is constent with it, then the version space is empty $\left(V S_{\mathcal{F}}(S)=\varnothing\right)$. This happens because:

## Version Space

## Version Space

The version space $V S_{\mathcal{F}}(S)$ of $\mathcal{F}$ given $S$ is the set of hypotheses $f \in \mathcal{F}$ that is consistent with $S$, that is:

$$
\begin{equation*}
V S_{\mathcal{F}}(S)=\{f \in \mathcal{F}: f \models S\} \tag{26}
\end{equation*}
$$

■ If $S$ is not separable w.r.t. $\mathcal{F}$, i.e., if there is no hypothesis $f \in \mathcal{F}$ that is constent with it, then the version space is empty $\left(V S_{\mathcal{F}}(S)=\varnothing\right)$. This happens because:

- $\mathcal{F}$ is not expressive enough. Example: no $f^{*} \in \mathcal{F}$ that correctly labels the data.


## Version Space

## Version Space

The version space $V S_{\mathcal{F}}(S)$ of $\mathcal{F}$ given $S$ is the set of hypotheses $f \in \mathcal{F}$ that is consistent with $S$, that is:

$$
\begin{equation*}
V S_{\mathcal{F}}(S)=\{f \in \mathcal{F}: f \models S\} \tag{26}
\end{equation*}
$$

$\square$ If $S$ is not separable w.r.t. $\mathcal{F}$, i.e., if there is no hypothesis $f \in \mathcal{F}$ that is constent with it, then the version space is empty $\left(V S_{\mathcal{F}}(S)=\varnothing\right)$. This happens because:

- $\mathcal{F}$ is not expressive enough. Example: no $f^{*} \in \mathcal{F}$ that correctly labels the data.
- $S$ is noisy. Example: $S$ contains the same instance twice but annotated with different labels - e.g., $(\mathrm{x}, 1)$ and ( $\mathrm{x}, 3$ ) - so no $f \in \mathcal{F}$ can classify both correctly.


## Version Space

## Version Space

The version space $V S_{\mathcal{F}}(S)$ of $\mathcal{F}$ given $S$ is the set of hypotheses $f \in \mathcal{F}$ that is consistent with $S$, that is:

$$
\begin{equation*}
V S_{\mathcal{F}}(S)=\{f \in \mathcal{F}: f \models S\} \tag{26}
\end{equation*}
$$

$\square$ If $S$ is not separable w.r.t. $\mathcal{F}$, i.e., if there is no hypothesis $f \in \mathcal{F}$ that is constent with it, then the version space is empty $\left(V S_{\mathcal{F}}(S)=\varnothing\right)$. This happens because:

- $\mathcal{F}$ is not expressive enough. Example: no $f^{*} \in \mathcal{F}$ that correctly labels the data.
- $S$ is noisy. Example: $S$ contains the same instance twice but annotated with different labels - e.g., $(\mathrm{x}, 1)$ and ( $\mathrm{x}, 3$ ) - so no $f \in \mathcal{F}$ can classify both correctly.
$\square$ We assume the realizable case: $\exists f^{*} \in \mathcal{F}$ s.t. $y=f^{*}(\mathrm{x})$ for all x and no noise.
This implies that $f^{*} \in V S_{\mathcal{F}}(L)$ for all choices of labeled examples $L$, because the supervision $(\mathrm{x}, y)$ is always consistent with $f^{*}$. In addition, $V S_{\mathcal{F}}(L) \neq \varnothing$.


## Version Space $\leftrightarrow$ Disagreement Region

## Disagreement Region

Given $\mathcal{F}$ and $S$, the disagreement region is the set of points $\mathrm{x} \in \mathbb{R}^{d}$ such that there exist two classifiers $f, f^{\prime}$ in the version space $V S_{\mathcal{F}}(S)$ that produce different predictions for them:

$$
\begin{equation*}
D I S_{\mathcal{F}}(S)=\left\{\mathrm{x} \in \mathbb{R}^{d}: \exists f, f^{\prime} \in V S_{\mathcal{F}}(S) \cdot f(\mathrm{x}) \neq f^{\prime}(\mathrm{x})\right\} \tag{27}
\end{equation*}
$$

## Version Space $\leftrightarrow$ Disagreement Region

## Disagreement Region

Given $\mathcal{F}$ and $S$, the disagreement region is the set of points $\mathrm{x} \in \mathbb{R}^{d}$ such that there exist two classifiers $f, f^{\prime}$ in the version space $V S_{\mathcal{F}}(S)$ that produce different predictions for them:

$$
\begin{equation*}
D I S_{\mathcal{F}}(S)=\left\{\mathrm{x} \in \mathbb{R}^{d}: \exists f, f^{\prime} \in V S_{\mathcal{F}}(S) \cdot f(\mathrm{x}) \neq f^{\prime}(\mathrm{x})\right\} \tag{27}
\end{equation*}
$$

■ If $\mathrm{x} \notin D I S_{\mathcal{F}}(S)$, then all candidate classifiers $f$ in the version space classify it the same: acquiring its label is pointless.

## Version Space $\leftrightarrow$ Disagreement Region

## Disagreement Region

Given $\mathcal{F}$ and $S$, the disagreement region is the set of points $\mathrm{x} \in \mathbb{R}^{d}$ such that there exist two classifiers $f, f^{\prime}$ in the version space $V S_{\mathcal{F}}(S)$ that produce different predictions for them:

$$
\begin{equation*}
D I S_{\mathcal{F}}(S)=\left\{\mathrm{x} \in \mathbb{R}^{d}: \exists f, f^{\prime} \in V S_{\mathcal{F}}(S) \cdot f(\mathrm{x}) \neq f^{\prime}(\mathrm{x})\right\} \tag{27}
\end{equation*}
$$

■ If $\mathrm{x} \notin D I S_{\mathcal{F}}(S)$, then all candidate classifiers $f$ in the version space classify it the same: acquiring its label is pointless.
$\square$ If $\mathrm{x} \in \operatorname{DIS}_{\mathcal{F}}(S)$, then at least one $f$ in the version space classifies it differently: acquiring its label is useful.

Left: input space $\mathbb{R}^{d}$, data set $S$ of red crosses vs blue circles. Right: hypothesis space $\mathcal{F}$, each $f$ is a point; the ground-truth $f^{*}$ is in red.

Left: input space $\mathbb{R}^{d}$, data set $S$ of red crosses vs blue circles. Right: hypothesis space $\mathcal{F}$, each $f$ is a point; the ground-truth $f^{*}$ is in red.
$\mathcal{F}$ is the set of $2 D$ rectangles. Rectangles in instance space (left) are points in hypothesis space (right), as shown by the arrows.

Left: input space $\mathbb{R}^{d}$, data set $S$ of red crosses vs blue circles. Right: hypothesis space $\mathcal{F}$, each $f$ is a point; the ground-truth $f^{*}$ is in red.
$\mathcal{F}$ is the set of $2 D$ rectangles. Rectangles in instance space (left) are points in hypothesis space (right), as shown by the arrows.

The version space $V S_{\mathcal{F}}(S)$ contains all the rectangles (pale gray) between inner \& outer rectangles (darker gray)

Left: input space $\mathbb{R}^{d}$, data set $S$ of red crosses vs blue circles. Right: hypothesis space $\mathcal{F}$, each $f$ is a point; the ground-truth $f^{*}$ is in red.
$\mathcal{F}$ is the set of $2 D$ rectangles. Rectangles in instance space (left) are points in hypothesis space (right), as shown by the arrows.

The version space $V S_{\mathcal{F}}(S)$ contains all the rectangles (pale gray) between inner \& outer rectangles (darker gray) The disagreement region $\operatorname{DIS}_{\mathcal{F}}(S)$ is the space enclosed between these two rectangles.

## Streaming AL

Input: models $\mathcal{F}$
Output: selected model $f \in \mathcal{F}$
1: $L \leftarrow \varnothing$
2: $\mathcal{V} \leftarrow \mathcal{F} \quad \triangleright$ implements the version space $V S_{\mathcal{F}}(L)$
3: for $t=1,2,3, \ldots$, do
4: receive instance x
5: if $\mathrm{x} \in \operatorname{DIS}(\mathcal{V})$ then $\quad \triangleright$ if x falls in the disagreement region
6: obtain label $y$ of $x$
7: update $\mathcal{V} \leftarrow\{f \in \mathcal{V}: f(x)=y\} \quad \triangleright$ update version space
: return any $f \in \mathcal{V}$

## Streaming AL

Input: models $\mathcal{F}$
Output: selected model $f \in \mathcal{F}$
1: $L \leftarrow \varnothing$
2: $\mathcal{V} \leftarrow \mathcal{F} \quad \triangleright$ implements the version space $V S_{\mathcal{F}}(L)$
3: for $t=1,2,3, \ldots$, do
4: receive instance x
5: $\quad$ if $\mathrm{x} \in \operatorname{DIS}(\mathcal{V})$ then obtain label $y$ of $x$ update $\mathcal{V} \leftarrow\{f \in \mathcal{V}: f(\mathrm{x})=y\} \quad \triangleright$ update version space
return any $f \in \mathcal{V}$
$\square$ If $\mathrm{x} \in \operatorname{DI} S_{\mathcal{F}}(L)$, then there are two $f, f^{\prime} \in V S_{\mathcal{F}}(L)$ that disagree on how x should be labeled. Getting its label allows us to get rid of either $f$ or $f^{\prime}$, so $V S_{\mathcal{F}}(S)$ and $D I S_{\mathcal{F}}(S)$ both shrink.

## Streaming AL

Input: models $\mathcal{F}$
Output: selected model $f \in \mathcal{F}$
1: $L \leftarrow \varnothing$
$\mathcal{V} \leftarrow \mathcal{F} \quad \triangleright$ implements the version space $V S_{\mathcal{F}}(L)$
: for $t=1,2,3, \ldots$, do receive instance x if $\mathrm{x} \in \operatorname{DIS}(\mathcal{V})$ then $\quad \triangleright$ if x falls in the disagreement region obtain label $y$ of $x$ update $\mathcal{V} \leftarrow\{f \in \mathcal{V}: f(\mathrm{x})=y\} \quad \triangleright$ update version space
return any $f \in \mathcal{V}$
$\square$ If $\mathrm{x} \in D I S_{\mathcal{F}}(L)$, then there are two $f, f^{\prime} \in V S_{\mathcal{F}}(L)$ that disagree on how x should be labeled. Getting its label allows us to get rid of either $f$ or $f^{\prime}$, so $V S_{\mathcal{F}}(S)$ and $D I S_{\mathcal{F}}(S)$ both shrink.
$\square$ Recall that $f^{*}$ is always compatible with examples $(x, y)$, so it is always in $\mathcal{V} \rightarrow$ algorithm zooms into it!

## Streaming AL

```
Input: models \(\mathcal{F}\)
Output: selected model \(f \in \mathcal{F}\)
    1: \(L \leftarrow \varnothing\)
    \(\mathcal{V} \leftarrow \mathcal{F} \quad \triangleright\) implements the version space \(V S_{\mathcal{F}}(L)\)
    : for \(t=1,2,3, \ldots\), do
        receive instance \(\mathbf{x}\)
        if \(\mathrm{x} \in \operatorname{DIS}(\mathcal{V})\) then \(\quad \triangleright\) if x falls in the disagreement region
            obtain label \(y\) of \(x\)
            update \(\mathcal{V} \leftarrow\{f \in \mathcal{V}: f(\mathbf{x})=y\} \quad \triangleright\) update version space
    return any \(f \in \mathcal{V}\)
```

$\square$ If $\mathrm{x} \in D I S_{\mathcal{F}}(L)$, then there are two $f, f^{\prime} \in V S_{\mathcal{F}}(L)$ that disagree on how x should be labeled. Getting its label allows us to get rid of either $f$ or $f^{\prime}$, so $V S_{\mathcal{F}}(S)$ and $D I S_{\mathcal{F}}(S)$ both shrink.
$\square$ Recall that $f^{*}$ is always compatible with examples $(\mathrm{x}, y)$, so it is always in $\mathcal{V} \rightarrow$ algorithm zooms into it!

- This algorithm makes no useless queries!

Streaming AL

```
Input: models \(\mathcal{F}\)
Output: selected model \(f \in \mathcal{F}\)
    1: \(L \leftarrow \varnothing\)
    \(\mathcal{V} \leftarrow \mathcal{F} \quad \triangleright\) implements the version space \(V S_{\mathcal{F}}(L)\)
    : for \(t=1,2,3, \ldots\), do
        receive instance \(\mathbf{x}\)
        if \(\mathrm{x} \in \operatorname{DIS}(\mathcal{V})\) then \(\quad \triangleright\) if x falls in the disagreement region
            obtain label \(y\) of \(x\)
            update \(\mathcal{V} \leftarrow\{f \in \mathcal{V}: f(\mathbf{x})=y\} \quad \triangleright\) update version space
    return any \(f \in \mathcal{V}\)
```

$\square$ If $\mathrm{x} \in D I S_{\mathcal{F}}(L)$, then there are two $f, f^{\prime} \in V S_{\mathcal{F}}(L)$ that disagree on how x should be labeled. Getting its label allows us to get rid of either $f$ or $f^{\prime}$, so $V S_{\mathcal{F}}(S)$ and $D I S_{\mathcal{F}}(S)$ both shrink.
$\square$ Recall that $f^{*}$ is always compatible with examples $(x, y)$, so it is always in $\mathcal{V} \rightarrow$ algorithm zooms into it!

■ This algorithm makes no useless queries! ...can we do better if we can choose x ?

## Pool-based AL

```
Input: models \(\mathcal{F}\)
Output: selected model \(f \in \mathcal{F}\)
: \(L \leftarrow \varnothing\)
2: \(\mathcal{V} \leftarrow \mathcal{F} \quad \triangleright\) implements the version space \(V S_{\mathcal{F}}(L)\)
3: for \(t=1,2, \ldots, T\) do
        \(\mathrm{x} \leftarrow \operatorname{argmax}_{\mathrm{x} \in U} \operatorname{acq}_{V S}(\mathcal{V}, \mathcal{F}, \mathrm{x})\)
        obtain label \(y\) of \(x\)
        update \(\mathcal{V} \leftarrow\{f \in \mathcal{V}: f(\mathrm{x})=y\} \quad \triangleright\) update version space
    return any \(f \in \mathcal{V}\)
```

$\square$ We can always ensure that $\mathrm{x} \in D S_{\mathcal{F}}(L)$ - as long as $V S_{\mathcal{F}}(S) \neq \varnothing$, in which case we can simply terminate.

Problem: how do we define the acquisition function?

Consider the homogeneous linear classifiers:

$$
\begin{equation*}
\mathcal{F}=\left\{f_{\theta}(\mathrm{x})=\mathbb{1}\left\{\boldsymbol{\theta}^{\top} \mathrm{x}>0\right\}: \boldsymbol{\theta} \in \mathbb{R}^{d},\|\boldsymbol{\theta}\|_{2}=1\right\} \tag{28}
\end{equation*}
$$

The version space of $S$ is essentially the set of direction vectors $\boldsymbol{\theta}$ that classify all points correctly.

- Classifiers are hyperplanes in instance space and instances are hyperplanes in hypothesis space (duality)
$\square$ Pick the point $\mathrm{x} \in U$ that (greedily) restricts the version space as much as possible. In this special case, x it passes close to the center of $V S_{\mathcal{F}}(S)$.

Idea: pick $\mathrm{x} \in U$ that reduces the volume of the version space $V S_{\mathcal{F}}(S)$ as much as possible.

Idea: pick $\mathrm{x} \in U$ that reduces the volume of the version space $V S_{\mathcal{F}}(S)$ as much as possible.

The volume of a region $A \subseteq \mathcal{F}$ is:

$$
\operatorname{Vol}(A)=\int_{A} d \boldsymbol{\theta}
$$

Idea: pick $\mathrm{x} \in U$ that reduces the volume of the version space $V S_{\mathcal{F}}(S)$ as much as possible.

The volume of a region $A \subseteq \mathcal{F}$ is:

$$
\begin{equation*}
\operatorname{Vol}(A)=\int_{A} d \boldsymbol{\theta}=\int_{\boldsymbol{\theta} \in \mathbb{R}^{|\boldsymbol{\theta}|}} \delta\{\boldsymbol{\theta} \in A\} d \boldsymbol{\theta} \tag{29}
\end{equation*}
$$

So computing a volume in general requires integration.

Idea: pick $\mathrm{x} \in U$ that reduces the volume of the version space $V S_{\mathcal{F}}(S)$ as much as possible.

The volume of a region $A \subseteq \mathcal{F}$ is:

$$
\begin{equation*}
\operatorname{Vol}(A)=\int_{A} d \boldsymbol{\theta}=\int_{\boldsymbol{\theta} \in \mathbb{R}^{|\boldsymbol{\theta}|}} \delta\{\boldsymbol{\theta} \in A\} d \boldsymbol{\theta} \tag{29}
\end{equation*}
$$

So computing a volume in general requires integration.
$\square$ Pick x that minimizes $\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathrm{x}, \mathrm{y})\})\right)$. However, we don't know the label y of x

Idea: pick $\mathrm{x} \in U$ that reduces the volume of the version space $V S_{\mathcal{F}}(S)$ as much as possible.
$\square$ The volume of a region $A \subseteq \mathcal{F}$ is:

$$
\begin{equation*}
\operatorname{Vol}(A)=\int_{A} d \boldsymbol{\theta}=\int_{\boldsymbol{\theta} \in \mathbb{R}^{|\boldsymbol{\theta}|}} \delta\{\boldsymbol{\theta} \in A\} d \boldsymbol{\theta} \tag{29}
\end{equation*}
$$

So computing a volume in general requires integration.
$\square$ Pick x that minimizes $\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathrm{x}, \mathrm{y})\})\right)$. However, we don't know the label y of $\mathrm{x} \rightarrow$ marginalize it:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmin}} \mathbb{E}_{y \sim p_{\theta}(Y \mid \mathbf{x})}\left[\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathbf{x}, y)\})\right)\right] \tag{30}
\end{equation*}
$$

Question: how to encode the version space?

Question: how to encode the version space?

- If $\mathcal{F}$ is finite, can explicitly store $f \models S$. Bonus: computing expected volume is doable (integral becomes sum).

Question: how to encode the version space?
$\square$ If $\mathcal{F}$ is finite, can explicitly store $f \models S$. Bonus: computing expected volume is doable (integral becomes sum).

■ If $\mathcal{F}$ is infinite, cannot store explicitly. However, we only need to compute its volume:

$$
\begin{equation*}
\mathbb{E}_{y \sim p_{\theta}(Y \mid \mathrm{x})}\left[\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathrm{x}, y)\})\right)\right]=\sum_{y \in[c]} p_{\theta}(y \mid \mathrm{x}) \cdot \underbrace{\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathrm{x}, y)\})\right)}_{\text {this is the difficult bit }} \tag{31}
\end{equation*}
$$

Question: how to encode the version space?
$\square$ If $\mathcal{F}$ is finite, can explicitly store $f \models S$. Bonus: computing expected volume is doable (integral becomes sum).

■ If $\mathcal{F}$ is infinite, cannot store explicitly. However, we only need to compute its volume:

$$
\begin{equation*}
\mathbb{E}_{y \sim p_{\theta}(Y \mid \mathbf{x})}\left[\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathbf{x}, y)\})\right)\right]=\sum_{y \in[c]} p_{\theta}(y \mid \mathbf{x}) \cdot \underbrace{\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathbf{x}, y)\})\right)}_{\text {this is the difficult bit }} \tag{31}
\end{equation*}
$$

- If $\mathcal{F}$ is "simple" and/or $S$ is small, volume can be approximated cheaply using Monte Carlo techniques. For instance with rejection sampling, let $B \subseteq V S_{\mathcal{F}}(S)$ of known volume:

$$
\begin{equation*}
\left\{\widetilde{\boldsymbol{\theta}}_{i} \sim \operatorname{Uniform}(B): i \in[s]\right\}, \quad \operatorname{Vol}\left(V S_{\mathcal{F}}\left(S^{\prime}\right)\right) \approx \frac{1}{\operatorname{Vol}(B)} \cdot \frac{1}{s} \sum_{i \in[s]} \mathbb{1}\left\{\widetilde{\boldsymbol{\theta}}_{i} \in V S_{\mathcal{F}}\left(S^{\prime}\right)\right\} \tag{32}
\end{equation*}
$$

To check, $\mathbb{1}\left\{\boldsymbol{\theta}_{\boldsymbol{i}} \in V S_{\mathcal{F}}\left(S^{\prime}\right)\right\}$, check that $f_{\boldsymbol{\theta}}$ classifies all examples in $S$ correctly.

Question: how to encode the version space?
$\square$ If $\mathcal{F}$ is finite, can explicitly store $f \models S$. Bonus: computing expected volume is doable (integral becomes sum).

■ If $\mathcal{F}$ is infinite, cannot store explicitly. However, we only need to compute its volume:

$$
\begin{equation*}
\mathbb{E}_{y \sim p_{\theta}(Y \mid \mathbf{x})}\left[\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathbf{x}, y)\})\right)\right]=\sum_{y \in[c]} p_{\theta}(y \mid \mathbf{x}) \cdot \underbrace{\operatorname{Vol}\left(V S_{\mathcal{F}}(S \cup\{(\mathbf{x}, y)\})\right)}_{\text {this is the difficult bit }} \tag{31}
\end{equation*}
$$

- If $\mathcal{F}$ is "simple" and/or $S$ is small, volume can be approximated cheaply using Monte Carlo techniques. For instance with rejection sampling, let $B \subseteq V S_{\mathcal{F}}(S)$ of known volume:

$$
\begin{equation*}
\left\{\widetilde{\boldsymbol{\theta}}_{i} \sim \operatorname{Uniform}(B): i \in[s]\right\}, \quad \operatorname{Vol}\left(V S_{\mathcal{F}}\left(S^{\prime}\right)\right) \approx \frac{1}{\operatorname{Vol}(B)} \cdot \frac{1}{s} \sum_{i \in[s]} \mathbb{1}\left\{\widetilde{\boldsymbol{\theta}}_{i} \in V S_{\mathcal{F}}\left(S^{\prime}\right)\right\} \tag{32}
\end{equation*}
$$

To check, $\mathbb{1}\left\{\boldsymbol{\theta}_{\boldsymbol{i}} \in V S_{\mathcal{F}}\left(S^{\prime}\right)\right\}$, check that $f_{\boldsymbol{\theta}}$ classifies all examples in $S$ correctly.

- Otherwise (think CNN on ImageNet), can be extremely challenging
- These approaches make two assumptions:
- These approaches make two assumptions:
- Disagreement is measures using all hypotheses in the version space $V S_{\mathcal{F}}(S)$.
- These approaches make two assumptions:
- Disagreement is measures using all hypotheses in the version space $V S_{\mathcal{F}}(S)$.
- Disagreement is binary: it is only 0 if all hypotheses fully agree on $\mathrm{x} \in U$.

■ Let's relax both of them $\rightarrow$ speed-up!

- These approaches make two assumptions:
- Disagreement is measures using all hypotheses in the version space $V S_{\mathcal{F}}(S)$.
- Disagreement is binary: it is only 0 if all hypotheses fully agree on $x \in U$.

■ Let's relax both of them $\rightarrow$ speed-up!

- Moreover, version space is only non-empty in the realizable case. How do we deal with this?


## Query By Committee (QBC)

Idea:

- Subsample $k$ representatives $C \subset V S_{\mathcal{F}}(S), k \gg 1$
- Measure disagreement efficiently between representatives in $C$


## Query By Committee (QBC)

## Idea:

- Subsample $k$ representatives $C \subset V S_{\mathcal{F}}(S), k \gg 1$
- Measure disagreement efficiently between representatives in $C$

How to generate $C$ ?

- Sample $\mathcal{C}:=\left\{c_{j}: j \in[k]\right\} \stackrel{k}{\sim} V S_{\mathcal{F}}(S)$, as done previously (i.e., we don't gain much)


## Query By Committee (QBC)

## Idea:

- Subsample $k$ representatives $C \subset V S_{\mathcal{F}}(S), k \gg 1$
- Measure disagreement efficiently between representatives in $C$
- How to generate $C$ ?
- Sample $\mathcal{C}:=\left\{c_{j}: j \in[k]\right\} \stackrel{k}{\sim} V S_{\mathcal{F}}(S)$, as done previously (i.e., we don't gain much)
- Define a distribution $p(\theta \mid L)$, very Bayesian $\rightarrow$ can be challenging to sample from the posterior


## Query By Committee (QBC)

## Idea:

- Subsample $k$ representatives $C \subset V S_{\mathcal{F}}(S), k \gg 1$
- Measure disagreement efficiently between representatives in $C$
- How to generate $C$ ?
- Sample $\mathcal{C}:=\left\{c_{j}: j \in[k]\right\} \stackrel{k}{\sim} V S_{\mathcal{F}}(S)$, as done previously (i.e., we don't gain much)
- Define a distribution $p(\theta \mid L)$, very Bayesian $\rightarrow$ can be challenging to sample from the posterior
- Bagging: sample $k$ subsets of $L$, train one classifier $c_{j}$ on each.


## Query By Committee (QBC)

## Idea:

- Subsample $k$ representatives $C \subset V S_{\mathcal{F}}(S), k \gg 1$
- Measure disagreement efficiently between representatives in $C$

■ How to generate $C$ ?

- Sample $\mathcal{C}:=\left\{c_{j}: j \in[k]\right\} \stackrel{k}{\sim} V S_{\mathcal{F}}(S)$, as done previously (i.e., we don't gain much)
- Define a distribution $p(\theta \mid L)$, very Bayesian $\rightarrow$ can be challenging to sample from the posterior
- Bagging: sample $k$ subsets of $L$, train one classifier $c_{j}$ on each.
- Boosting: randomly reweight $L$, sequentially train $k$ classifiers by repeatedly reweighting examples by mistakes made by previous classifier.
$\square$ How to measure disagreement of $\mathcal{C}$ on $\mathrm{x} \in U$ ?
$\square$ How to measure disagreement of $\mathcal{C}$ on $\mathrm{x} \in U$ ?

■ "Hard" Voting + Entropy:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}}-\sum_{y} \frac{n(y, \mathbf{x})}{k} \log \frac{n(y, \mathbf{x})}{k}, \quad n(y, \mathbf{x}):=\sum_{c \in \mathcal{C}} \mathbb{1}\{c(\mathbf{x})=y\} \tag{33}
\end{equation*}
$$

Each classifiers votes either 0 or 1.
$\square$ How to measure disagreement of $\mathcal{C}$ on $\mathrm{x} \in U$ ?

■ "Hard" Voting + Entropy:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}}-\sum_{y} \frac{n(y, \mathbf{x})}{k} \log \frac{n(y, \mathbf{x})}{k}, \quad n(y, \mathbf{x}):=\sum_{c \in \mathcal{C}} \mathbb{1}\{c(\mathbf{x})=y\} \tag{33}
\end{equation*}
$$

Each classifiers votes either 0 or 1.

■ "Soft" Voting + Entropy:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}}-\sum_{y} p_{\mathcal{C}}(y \mid \mathbf{x}) \log p_{\mathcal{C}}(y \mid \mathbf{x}), \quad p_{\mathcal{C}}(y \mid \mathbf{x}):=\frac{1}{k} \sum_{c \in \mathcal{C}} p_{c}(y \mid \mathbf{x}) \tag{34}
\end{equation*}
$$

Output probabilities of each $c \in \mathcal{C}$ taken into account.
$\square$ How to measure disagreement of $\mathcal{C}$ on $\mathrm{x} \in U$ ?

■ "Hard" Voting + Entropy:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}}-\sum_{y} \frac{n(y, \mathbf{x})}{k} \log \frac{n(y, \mathbf{x})}{k}, \quad n(y, \mathbf{x}):=\sum_{c \in \mathcal{C}} \mathbb{1}\{c(\mathbf{x})=y\} \tag{33}
\end{equation*}
$$

Each classifiers votes either 0 or 1 .

■ "Soft" Voting + Entropy:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}}-\sum_{y} p_{\mathcal{C}}(y \mid \mathrm{x}) \log p_{\mathcal{C}}(y \mid \mathrm{x}), \quad p_{\mathcal{C}}(y \mid \mathrm{x}):=\frac{1}{k} \sum_{c \in \mathcal{C}} p_{c}(y \mid \mathrm{x}) \tag{34}
\end{equation*}
$$

Output probabilities of each $c \in \mathcal{C}$ taken into account.

- Average Kullback-Liebler divergence:

$$
\begin{align*}
& \underset{\mathrm{x} \in U}{\operatorname{argmax}} \frac{1}{k} \sum_{c \in \mathcal{C}} \mathrm{KL}\left(p_{c}(Y \mid \mathrm{x}) \| p_{\mathcal{C}}(Y \mid \mathrm{x})\right)  \tag{35}\\
& \quad \mathrm{KL}(p(Y \mid \mathrm{x}) \| q(Y \mid \mathrm{x})):=\sum_{y} p(y \mid \mathrm{x}) \log \frac{p(y \mid \mathrm{x})}{q(y \mid \mathrm{x})} \tag{36}
\end{align*}
$$

Very expressive, measures difference between whole distributions.

## Model Improvement

Idea: pick the point that gives the maximal improvement in model quality

## Model Improvement

Idea: pick the point that gives the maximal improvement in model quality

## Useful Concepts

The loss of $p_{\theta}$ on example $z=(\mathrm{x}, y)$ is denoted $\ell(\theta, z)$. For instance, cross-entropy loss:

$$
\begin{equation*}
\ell(\theta, z):=-\sum_{j} \mathbb{1}\{j=y\} \log p_{\theta}(j \mid \mathbf{x})=-\log p_{\theta}(y \mid \mathbf{x}) \tag{37}
\end{equation*}
$$

## Model Improvement

Idea: pick the point that gives the maximal improvement in model quality

## Useful Concepts

The loss of $p_{\theta}$ on example $z=(\mathrm{x}, y)$ is denoted $\ell(\theta, z)$. For instance, cross-entropy loss:

$$
\begin{equation*}
\ell(\theta, z):=-\sum_{j} \mathbb{1}\{j=y\} \log p_{\theta}(j \mid \mathbf{x})=-\log p_{\theta}(y \mid \mathbf{x}) \tag{37}
\end{equation*}
$$

The true risk $\mathcal{L}^{*}$ of $\theta$ w.r.t. the ground-truth distribution $p^{*}(\mathbf{X}, Y)$ is:

$$
\begin{equation*}
\mathcal{L}^{*}(\theta):=\mathbb{E}_{z \sim p^{*}}[\ell(\theta, z)] \tag{38}
\end{equation*}
$$

It measures the true quality of the model, unobserved.

## Model Improvement

Idea: pick the point that gives the maximal improvement in model quality

## Useful Concepts

The loss of $p_{\theta}$ on example $z=(\mathrm{x}, y)$ is denoted $\ell(\theta, z)$. For instance, cross-entropy loss:

$$
\begin{equation*}
\ell(\theta, z):=-\sum_{j} \mathbb{1}\{j=y\} \log p_{\theta}(j \mid \mathbf{x})=-\log p_{\theta}(y \mid \mathbf{x}) \tag{37}
\end{equation*}
$$

The true risk $\mathcal{L}^{*}$ of $\theta$ w.r.t. the ground-truth distribution $p^{*}(\mathbf{X}, Y)$ is:

$$
\begin{equation*}
\mathcal{L}^{*}(\theta):=\mathbb{E}_{z \sim p^{*}}[\ell(\theta, z)] \tag{38}
\end{equation*}
$$

It measures the true quality of the model, unobserved.
The empirical risk $\widehat{\mathcal{L}}_{S}$ of $\theta$ w.r.t. data set $S=\left\{z_{1}, \ldots, z_{m}\right\}$ sampled i.i.d. from $p^{*}$ is:

$$
\begin{equation*}
\widehat{\mathcal{L}}_{S}(\theta):=\frac{1}{|S|} \sum_{z \in S} \ell(\theta, z) \tag{39}
\end{equation*}
$$

It estimates the quality of the model from a sample $S$, optimized during training.

## Model Improvement

Let $\widehat{\theta}$ be the parameters obtained by training on $S$ and $\widehat{\theta}^{+z}$ those obtained by training on $S \cup\{z\}$, i.e.,

$$
\begin{equation*}
\widehat{\theta}:=\underset{\theta}{\operatorname{argmin}} \widehat{\mathcal{L}}_{S}(\theta) \quad \widehat{\theta}^{+z}:=\underset{\theta}{\operatorname{argmin}} \widehat{\mathcal{L}}_{S \cup\{z\}}(\theta) \tag{40}
\end{equation*}
$$

where optimization is possibly approximate, e.g., based on SGD.

## Model Improvement

Let $\widehat{\theta}$ be the parameters obtained by training on $S$ and $\widehat{\theta}^{+z}$ those obtained by training on $S \cup\{z\}$, i.e.,

$$
\begin{equation*}
\widehat{\theta}:=\underset{\theta}{\operatorname{argmin}} \widehat{\mathcal{L}}_{S}(\theta) \quad \widehat{\theta}^{+z}:=\underset{\theta}{\operatorname{argmin}} \widehat{\mathcal{L}}_{S \cup\{z\}}(\theta) \tag{40}
\end{equation*}
$$

where optimization is possibly approximate, e.g., based on SGD.

## Model Improvement

The model improvement (MI) given by a new example $z \notin S$ is the decrease in true risk:

$$
\begin{equation*}
\operatorname{acq}(\mathrm{x}):=\mathcal{L}^{*}(\widehat{\theta})-\mathcal{L}^{*}\left(\widehat{\theta}^{+\mathrm{z}}\right) \tag{41}
\end{equation*}
$$

The higher, the better $\longrightarrow$ pick the $\mathrm{x} \in U$ that maximizes the improvement.

## Model Improvement as Greedy Optimization

$\square \mathrm{MI}$ amounts to solving:

$$
\begin{equation*}
\underset{\mathrm{x} \in U}{\operatorname{argmax}} \mathcal{L}^{*}(\widehat{\theta})-\mathcal{L}^{*}\left(\widehat{\theta}^{+\mathrm{z}}\right)=\underset{\mathrm{x} \in U}{\operatorname{argmin}} \mathcal{L}^{*}\left(\widehat{\theta}^{+\mathrm{z}}\right) \tag{42}
\end{equation*}
$$

It is guaranteed to find the best next candidate!

[^5]
## Model Improvement as Greedy Optimization

$\square \mathrm{MI}$ amounts to solving:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}} \mathcal{L}^{*}(\widehat{\theta})-\mathcal{L}^{*}\left(\widehat{\theta}^{+\mathrm{z}}\right)=\underset{\mathbf{x} \in U}{\operatorname{argmin}} \mathcal{L}^{*}\left(\widehat{\theta}^{+\mathrm{z}}\right) \tag{42}
\end{equation*}
$$

It is guaranteed to find the best next candidate!
$\square \mathrm{MI}$ is essentially a greedy strategy for solving: ${ }^{5}$

$$
\begin{align*}
\underset{S \subseteq U}{\operatorname{argmin}} & \mathcal{L}^{*}(\widehat{\theta})  \tag{43}\\
\text { s.t. } & |S| \leq \text { query budget } \tag{44}
\end{align*}
$$

In this view, $A L$ is a subset optimization problem, and $M I$ solves it directly.

[^6]
## Model Improvement as Greedy Optimization

$\square \mathrm{MI}$ amounts to solving:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}} \mathcal{L}^{*}(\widehat{\theta})-\mathcal{L}^{*}\left(\widehat{\theta}^{+\mathrm{z}}\right)=\underset{\mathbf{x} \in U}{\operatorname{argmin}} \mathcal{L}^{*}\left(\widehat{\theta}^{+\mathrm{z}}\right) \tag{42}
\end{equation*}
$$

It is guaranteed to find the best next candidate!
$\square \mathrm{MI}$ is essentially a greedy strategy for solving. ${ }^{5}$

$$
\begin{align*}
\underset{S \subseteq U}{\operatorname{argmin}} & \mathcal{L}^{*}(\widehat{\theta})  \tag{43}\\
\text { s.t. } & |S| \leq \text { query budget } \tag{44}
\end{align*}
$$

In this view, AL is a subset optimization problem, and MI solves it directly.
$\square$ Compare this to uncertainty sampling, which is not as sound

[^7]```
\underset{x\inU}{\operatorname{argmin}}\mp@subsup{\mathcal{L}}{}{*}(\mp@subsup{\widehat{0}}{}{+z})
```

[^8]We want to solve:

$$
\begin{equation*}
\underset{\mathrm{x} \in U}{\operatorname{argmin}} \mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \tag{45}
\end{equation*}
$$

Problem: $\mathcal{L}^{*}(\cdot)$ is an integral over $\mathrm{x}^{\prime} \in \mathbb{R}^{d}:$

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right)=\mathbb{E}_{z^{\prime} \sim p^{*}}\left[\ell\left(\widehat{\theta}^{+z}, z^{\prime}\right)\right]=\int_{\mathbb{R}^{d}} \ell\left(\widehat{\theta}^{+z},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right) d \mathrm{x}^{\prime} \tag{46}
\end{equation*}
$$

which is intractable

[^9]We want to solve:

$$
\begin{equation*}
\underset{\mathrm{x} \in U}{\operatorname{argmin}} \mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \tag{45}
\end{equation*}
$$

Problem: $\mathcal{L}^{*}(\cdot)$ is an integral over $\mathrm{x}^{\prime} \in \mathbb{R}^{d}:$

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right)=\mathbb{E}_{z^{\prime} \sim p^{*}}\left[\ell\left(\widehat{\theta}^{+z}, z^{\prime}\right)\right]=\int_{\mathbb{R}^{d}} \ell\left(\widehat{\theta}^{+z},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right) d \mathrm{x}^{\prime} \tag{46}
\end{equation*}
$$

which is intractable $\rightarrow$ approximate using empirical average over $U:{ }^{6}$

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \approx \widehat{\mathcal{L}}_{U}\left(\widehat{\theta}^{+z}\right)=\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+z},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right) \tag{47}
\end{equation*}
$$

[^10]We want to solve:

$$
\begin{equation*}
\underset{\mathrm{x} \in U}{\operatorname{argmin}} \mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \tag{45}
\end{equation*}
$$

Problem: $\mathcal{L}^{*}(\cdot)$ is an integral over $\mathrm{x}^{\prime} \in \mathbb{R}^{d}:$

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right)=\mathbb{E}_{z^{\prime} \sim p^{*}}\left[\ell\left(\widehat{\theta}^{+z}, z^{\prime}\right)\right]=\int_{\mathbb{R}^{d}} \ell\left(\widehat{\theta}^{+z},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right) d \mathrm{x}^{\prime} \tag{46}
\end{equation*}
$$

which is intractable $\rightarrow$ approximate using empirical average over $U:{ }^{6}$

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \approx \widehat{\mathcal{L}}_{U}\left(\widehat{\theta}^{+z}\right)=\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+z},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right) \tag{47}
\end{equation*}
$$

Example: if $\ell$ is the $0-1$ loss, then this amounts to:

$$
\begin{equation*}
\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{1}\left\{f_{\widehat{\theta}^{+}}\left(\mathbf{x}^{\prime}\right) \neq y^{\prime}\right\} \tag{48}
\end{equation*}
$$

[^11]■ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \widehat{\mathcal{L}}_{U}\left(\widehat{\theta}^{+z}\right)=\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+z},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right) \tag{49}
\end{equation*}
$$

We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \widehat{\mathcal{L}}_{U}\left(\widehat{\theta}^{+z}\right)=\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+z},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right) \tag{49}
\end{equation*}
$$

Problem: we don't have access to the ground-truth label $z=(x, y)$

We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \widehat{\mathcal{L}}_{U}\left(\widehat{\theta}^{+z}\right)=\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+z},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right) \tag{49}
\end{equation*}
$$

Problem: we don't have access to the ground-truth label $z=(x, y) \rightarrow$ marginalize w.r.t. $p^{*}:$

$$
\begin{equation*}
\mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right] \tag{50}
\end{equation*}
$$

This averages over alternative future models $\widehat{\theta}^{+(x, y)}$ obtained after retraining on $L \cup(\mathrm{x}, y)$.
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \widehat{\mathcal{L}}_{U}\left(\widehat{\theta}^{+z}\right)=\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+z},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right) \tag{49}
\end{equation*}
$$

Problem: we don't have access to the ground-truth label $z=(x, y) \rightarrow$ marginalize w.r.t. $p^{*}$ :

$$
\begin{equation*}
\mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+(\mathbf{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right] \tag{50}
\end{equation*}
$$

This averages over alternative future models $\widehat{\theta}^{+(x, y)}$ obtained after retraining on $L \cup(\mathbf{x}, y)$.

Problem: we don't have access to the ground-truth label $y^{\prime}$ either
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \widehat{\mathcal{L}}_{U}\left(\widehat{\theta}^{+z}\right)=\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+z},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right) \tag{49}
\end{equation*}
$$

Problem: we don't have access to the ground-truth label $z=(x, y) \rightarrow$ marginalize w.r.t. $p^{*}$ :

$$
\begin{equation*}
\mathbb{E}_{y \sim p^{*}}(Y \mid \mathrm{x})\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right] \tag{50}
\end{equation*}
$$

This averages over alternative future models $\widehat{\theta}^{+(x, y)}$ obtained after retraining on $L \cup(\mathbf{x}, y)$.

Problem: we don't have access to the ground-truth label $y^{\prime}$ either $\rightarrow$ marginalize w.r.t. $p^{*}$ :

$$
\begin{equation*}
\mathbb{E}_{y \sim p^{*}(Y \mid \mathbf{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p^{*}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{51}
\end{equation*}
$$

This averages over the unknown labels $y^{\prime}$ of the instances in $x^{\prime} \in U$.
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p^{*}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{52}
\end{equation*}
$$

$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p^{*}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{52}
\end{equation*}
$$

Problem: we don't have access to $p^{*}$ at all
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p^{*}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{52}
\end{equation*}
$$

Problem: we don't have access to $p^{*}$ at all $\rightarrow$ estimate using model's distribution:

where $\widehat{\theta}^{+}:=\widehat{\theta}^{+(x, y)}$. If $p_{\theta}$ is "good enough", then the approximation is good.
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p^{*}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{52}
\end{equation*}
$$

Problem: we don't have access to $p^{*}$ at all $\rightarrow$ estimate using model's distribution:

where $\widehat{\theta}^{+}:=\widehat{\theta}^{+(x, y)}$. If $p_{\theta}$ is "good enough", then the approximation is good.
(a) Is the expected loss of the updated model on $\mathbf{x}^{\prime} \in U$,
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p^{*}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+(\mathrm{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{52}
\end{equation*}
$$

Problem: we don't have access to $p^{*}$ at all $\rightarrow$ estimate using model's distribution:

where $\widehat{\theta}^{+}:=\widehat{\theta}^{+(x, y)}$. If $p_{\theta}$ is "good enough", then the approximation is good.
(a) Is the expected loss of the updated model on $\mathbf{x}^{\prime} \in U$,
(b) Is the average expected oss of the updated model on all of $U$,

We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p^{*}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p^{*}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+(\mathbf{x}, y)},\left(\mathbf{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{52}
\end{equation*}
$$

Problem: we don't have access to $p^{*}$ at all $\rightarrow$ estimate using model's distribution:

where $\widehat{\theta}^{+}:=\widehat{\theta}^{+(x, y)}$. If $p_{\theta}$ is "good enough", then the approximation is good.
(a) Is the expected loss of the updated model on $\mathbf{x}^{\prime} \in U$,
(b) Is the average expected oss of the updated model on all of $U$,
(c) Is the above averaged over the possible updated models $\widehat{\theta}^{+(x, y)}$.

- We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{54}
\end{equation*}
$$

- We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{54}
\end{equation*}
$$



We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}\left(Y \mid \mathrm{x}^{\prime}\right)}\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{54}
\end{equation*}
$$



$$
\left.\left.\begin{array}{rl}
\mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)
\end{array}\right] \mathbb{1}\left\{f_{\widehat{\theta}^{+}}\left(\mathrm{x}^{\prime}\right) \neq y^{\prime}\right\}\right]=p_{\widehat{\theta}^{+}}\left(\hat{y}^{\prime} \neq y^{\prime} \mid \mathrm{x}^{\prime}\right), \quad \hat{y}^{\prime}:=f_{\widehat{\theta}^{+}}\left(\mathrm{x}^{\prime}\right)
$$

$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{54}
\end{equation*}
$$



$$
\left.\left.\begin{array}{rl}
\mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)
\end{array}\right] \mathbb{1}\left\{f_{\widehat{\widehat{\theta}}^{+}}\left(\mathrm{x}^{\prime}\right) \neq y^{\prime}\right\}\right]=p_{\widehat{\theta}^{+}}\left(\hat{y}^{\prime} \neq y^{\prime} \mid \mathrm{x}^{\prime}\right), \quad \hat{y}^{\prime}:=f_{\widehat{\theta}^{+}}\left(\mathrm{x}^{\prime}\right)
$$

Hence, the above can be rewritten as ( $\frac{1}{|U|}$ doesn't matter because it is independent of $\mathbf{x}$ ):

$$
\begin{equation*}
\mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U}\left(1-p_{\widehat{\theta}^{+}}\left(\hat{y}^{\prime} \mid \mathrm{x}^{\prime}\right)\right)\right] \propto \sum_{y \in[c]} p_{\widehat{\theta}}(y \mid \mathbf{x}) \sum_{\mathrm{x}^{\prime} \in U}\left(1-p_{\widehat{\theta}^{+}}\left(\hat{y}^{\prime} \mid \mathrm{x}^{\prime}\right)\right) \tag{57}
\end{equation*}
$$

$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathbf{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{54}
\end{equation*}
$$



$$
\left.\left.\begin{array}{rl}
\mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)
\end{array}\right] \mathbb{1}\left\{f_{\widehat{\widehat{\theta}}^{+}}\left(\mathrm{x}^{\prime}\right) \neq y^{\prime}\right\}\right]=p_{\widehat{\theta}^{+}}\left(\hat{y}^{\prime} \neq y^{\prime} \mid \mathrm{x}^{\prime}\right), \quad \hat{y}^{\prime}:=f_{\widehat{\theta}^{+}}\left(\mathrm{x}^{\prime}\right)
$$

Hence, the above can be rewritten as ( $\frac{1}{|U|}$ doesn't matter because it is independent of $\mathbf{x}$ ):

$$
\begin{equation*}
\mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathbf{x})}\left[\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U}\left(1-p_{\widehat{\theta}^{+}}\left(\hat{y}^{\prime} \mid \mathbf{x}^{\prime}\right)\right)\right] \propto \sum_{y \in[c]} p_{\widehat{\theta}}(y \mid \mathbf{x}) \sum_{\mathbf{x}^{\prime} \in U}\left(1-p_{\widehat{\theta}^{+}}\left(\hat{y}^{\prime} \mid \mathbf{x}^{\prime}\right)\right) \tag{57}
\end{equation*}
$$

$\square$ We pick $\mathrm{x} \in U$ that minimizes the above $\rightarrow$ minimizes expected future confidence on $U$

- We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{58}
\end{equation*}
$$

- We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{58}
\end{equation*}
$$

Example: consider the negative $\log$-likelihood $\ell(\theta,(\mathbf{x}, y))=-\log p_{\theta}(y \mid \mathbf{x})$.

- We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{58}
\end{equation*}
$$

Example: consider the negative $\log$-likelihood $\ell(\theta,(\mathbf{x}, y))=-\log p_{\theta}(y \mid \mathbf{x})$. Then:

$$
\begin{align*}
\mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[-\log p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathrm{x}^{\prime}\right)\right] & =-\sum_{y^{\prime} \in[c]} p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathbf{x}^{\prime}\right) \log p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathbf{x}^{\prime}\right)  \tag{59}\\
& =H_{\widehat{\theta}^{+}}(Y \mid \mathbf{x}) \tag{60}
\end{align*}
$$

- We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{58}
\end{equation*}
$$

Example: consider the negative $\log$-likelihood $\ell(\theta,(\mathbf{x}, y))=-\log p_{\theta}(y \mid \mathbf{x})$. Then:

$$
\left.\left.\begin{array}{rl}
\mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid x^{\prime}\right)
\end{array}\right]-\log p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid x^{\prime}\right)\right]=-\sum_{y^{\prime} \in[c]} p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathbf{x}^{\prime}\right) \log p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathbf{x}^{\prime}\right)
$$

Hence, the above can be rewritten as:

$$
\begin{equation*}
\mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U}\left(H_{\widehat{\theta}^{+}}(Y \mid \mathrm{x})\right)\right] \propto \sum_{y \in[c]} p_{\widehat{\theta}}(y \mid \mathrm{x}) \sum_{\mathrm{x}^{\prime} \in U}\left(H_{\widehat{\theta}^{+}}(Y \mid \mathrm{x})\right) \tag{61}
\end{equation*}
$$

$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{58}
\end{equation*}
$$

Example: consider the negative $\log$-likelihood $\ell(\theta,(\mathbf{x}, y))=-\log p_{\theta}(y \mid \mathbf{x})$. Then:

$$
\begin{align*}
\mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[-\log p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathrm{x}^{\prime}\right)\right] & =-\sum_{y^{\prime} \in[c]} p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathbf{x}^{\prime}\right) \log p_{\widehat{\theta}^{+}}\left(y^{\prime} \mid \mathbf{x}^{\prime}\right)  \tag{59}\\
& =H_{\widehat{\theta}^{+}}(Y \mid \mathbf{x}) \tag{60}
\end{align*}
$$

Hence, the above can be rewritten as:

$$
\begin{equation*}
\mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U}\left(H_{\widehat{\theta}^{+}}(Y \mid \mathrm{x})\right)\right] \propto \sum_{y \in[c]} p_{\widehat{\theta}}(y \mid \mathrm{x}) \sum_{\mathrm{x}^{\prime} \in U}\left(H_{\widehat{\theta}^{+}}(Y \mid \mathrm{x})\right) \tag{61}
\end{equation*}
$$

$\square$ We pick $\mathrm{x} \in U$ that minimizes the above $\rightarrow$ minimizes expected future entropy on $U$

■ In uncertainty sampling, we pick x that minimizes model's estimate of current uncertainty w.r.t. itself, this is myopic

■ In expected model improvement, we pick x that minimizes model's estimate of expected future uncertainty w.r.t. unlabeled set, this is less myopic
comp.graphics vs. comp.windows.x

comp.sys.ibm.pc.hardware vs. comp.os.ms-windows.misc


Figure 4.1: Learning curves showing that expected error reduction can outperform QBC and uncertainty sampling for two binary text classification tasks. Source: Adapted from Roy and McCallum (2001), with kind permission of the authors.
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{62}
\end{equation*}
$$

$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{62}
\end{equation*}
$$

Problem: computing $\widehat{\theta}^{+}$requires to fit model on $L \cup\{(x, y)\}$ (slow)
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{62}
\end{equation*}
$$

Problem: computing $\widehat{\theta}^{+}$requires to fit model on $L \cup\{(x, y)\}$ (slow)

Problem: this has to be done $|U| \times[c]$ times.
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{62}
\end{equation*}
$$

Problem: computing $\widehat{\theta}^{+}$requires to fit model on $L \cup\{(\mathrm{x}, y)\}$ (slow)

Problem: this has to be done $|U| \times[c]$ times.

Problem: this has to be done in each iteration of active learning.
$\square$ We already decided on this approximation:

$$
\begin{equation*}
\mathcal{L}^{*}\left(\widehat{\theta}^{+z}\right) \quad \longrightarrow \quad \mathbb{E}_{y \sim p_{\widehat{\theta}}(Y \mid \mathrm{x})}\left[\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} \in U} \mathbb{E}_{y^{\prime} \sim p_{\widehat{\theta}^{+}}}\left(Y \mid \mathrm{x}^{\prime}\right)\left[\ell\left(\widehat{\theta}^{+},\left(\mathrm{x}^{\prime}, y^{\prime}\right)\right)\right]\right] \tag{62}
\end{equation*}
$$

Problem: computing $\widehat{\theta}^{+}$requires to fit model on $L \cup\{(x, y)\}$ (slow)

Problem: this has to be done $|U| \times[c]$ times.

Problem: this has to be done in each iteration of active learning.
$\square$ Only practical for classes of models that support closed-form updates (e.g., Gaussian Processes) or stable incremental learning (e.g., perceptron-like learning algorithms).

## Expected Model Change

- Unless a candidate ( $\mathrm{x}, y$ ) induces a large change in the model $\widehat{\theta}$ upon retraining, then it cannot reduce the model's risk by much: change is a prerequisite for improvement.


## Expected Model Change

- Unless a candidate ( $\mathrm{x}, \mathrm{y}$ ) induces a large change in the model $\widehat{\theta}$ upon retraining, then it cannot reduce the model's risk by much: change is a prerequisite for improvement.

Inituition:

$$
\begin{equation*}
\ell\left(\widehat{\theta}, z^{\prime}\right)-\ell\left(\widehat{\theta}^{+z}, z^{\prime}\right) \leq\left|\ell\left(\widehat{\theta}, z^{\prime}\right)-\ell\left(\widehat{\theta}^{+z}, z^{\prime}\right)\right| \leq c \cdot\left\|\widehat{\theta}-\widehat{\theta}^{+z}\right\|, \quad c>0 \tag{63}
\end{equation*}
$$

where $\|\cdot\|$ is, e.g., the Euclidean norm. This formally holds for all $c$-Lipshitz loss functions $\ell$.

## Expected Model Change

- Unless a candidate ( $\mathrm{x}, \mathrm{y}$ ) induces a large change in the model $\widehat{\theta}$ upon retraining, then it cannot reduce the model's risk by much: change is a prerequisite for improvement.


## Inituition:

$$
\begin{equation*}
\ell\left(\widehat{\theta}, z^{\prime}\right)-\ell\left(\widehat{\theta}^{+z}, z^{\prime}\right) \leq\left|\ell\left(\widehat{\theta}, z^{\prime}\right)-\ell\left(\widehat{\theta}^{+z}, z^{\prime}\right)\right| \leq c \cdot\left\|\widehat{\theta}-\widehat{\theta}^{+z}\right\|, \quad c>0 \tag{63}
\end{equation*}
$$

where $\|\cdot\|$ is, e.g., the Euclidean norm. This formally holds for all $c$-Lipshitz loss functions $\ell$.

■ Large change also occurs when the loss increases - hence the absolute value in the second step of Eq. 63.
All in all, EMC looks for examples $\mathrm{x} \in U$ that "make a difference" one way or the other.
But once ( $\mathrm{x}, y$ ) is acquried it is added to the training set $L$ on which $\widehat{\theta}$ is fit, so loss is likely to go down rather than up.

## Expected Model Change

The trick is that if $\widehat{\theta}$ is obtained via gradient descent, the difference $\widehat{\theta}-\widehat{\theta}^{+z}$ is easy to compute:

$$
\begin{equation*}
\widehat{\theta}-\widehat{\theta}^{+z}=\eta \cdot \nabla_{\theta} \ell(\theta, z) \tag{64}
\end{equation*}
$$

where $\eta$ is the learning rate. This gives expected gradient length:

$$
\begin{equation*}
\operatorname{acq}_{\mathrm{EGL}}(\mathrm{x}):=\mathbb{E}_{y \sim p_{\theta}(Y \mid \mathrm{x})}\left[\left\|\nabla_{\theta} \ell(\widehat{\theta},(\mathrm{x}, y))\right\|^{2}\right] \tag{65}
\end{equation*}
$$

The square does not change ranking of examples \& avoids computing a square root.

- Quite cheap to compute using automatic differentiation packages (using Jacobian to parallelize over $U$ )
- Assuming $\eta$ is constant across examples and GD, the computation is exact. For other optimizers, it is an approximation


## Are Uncertain Points Representative?

## Diversity-based Selection

Idea: pick instances $\mathrm{x} \in U$ that are both locally informative and also similar to as many other unlabeled points as possible:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}} \operatorname{acq}(f, \mathbf{x}) \cdot\left(\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \operatorname{sim}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right)^{\beta} \tag{66}
\end{equation*}
$$

where:

- $\operatorname{acq}(f, x)$ is a "standard" acquisition function based on, e.g., pointwise uncertainty.
- $\operatorname{sim}\left(\mathrm{x}, \mathrm{x}^{\prime}\right)$ measures the similarity between x and $\mathrm{x}^{\prime}$, e.g., a Gaussian kernel, Pearson's correlation coefficient, Spearman's rank correlation. Application specific.
- $\beta>0$ is a hyper-parameter

Intuitively, x's label conveys information about the label on the other points in $U$

## Example



Figure 5.2: Learning curves showing that, by explicitly weighting queries by their representativeness among the input instances, information density can yield better results than the base uncertainty sampling heuristic by itself.

■ We optimize:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}} \operatorname{acq}(f, \mathrm{x}) \cdot\left(\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \operatorname{sim}\left(\mathrm{x}, \mathrm{x}^{\prime}\right)\right)^{\beta} \tag{67}
\end{equation*}
$$

## Properties:

- Tends to work better than pure more "local" acquisition functions [Settles, 2012]
- Even when uncertainty sampling is worse than random, information density performs well
- We optimize:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}} \operatorname{acq}(f, \mathrm{x}) \cdot\left(\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \operatorname{sim}\left(\mathrm{x}, \mathrm{x}^{\prime}\right)\right)^{\beta} \tag{67}
\end{equation*}
$$

## Properties:

- Tends to work better than pure more "local" acquisition functions [Settles, 2012]
- Even when uncertainty sampling is worse than random, information density performs well
- Similarity computation can be sped-up using caching: "simply" store similarity matrix $S_{i j}=\left[\operatorname{sim}\left(\mathbf{x}_{i}, \mathrm{x}_{j}\right)\right]$ for all $\mathrm{x}_{i}, \mathrm{x}_{j} \in U$ (only needs to be done once)

We optimize:

$$
\begin{equation*}
\underset{\mathbf{x} \in U}{\operatorname{argmax}} \operatorname{acq}(f, \mathbf{x}) \cdot\left(\frac{1}{|U|} \sum_{\mathbf{x}^{\prime} \in U} \operatorname{sim}\left(\mathbf{x}, \mathrm{x}^{\prime}\right)\right)^{\beta} \tag{67}
\end{equation*}
$$

## Properties:

- Tends to work better than pure more "local" acquisition functions [Settles, 2012]
- Even when uncertainty sampling is worse than random, information density performs well
- Similarity computation can be sped-up using caching: "simply" store similarity matrix $S_{i j}=\left[\operatorname{sim}\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)\right]$ for all $\mathrm{x}_{i}, \mathrm{x}_{j} \in U$ (only needs to be done once)
- Approximate using clustering: cluster $U$ so that points within cluster are similar and points across clusters are not $\rightarrow$ block-diagonal similarity matrix, lowers storage requirement from $O\left(|U|^{2}\right)$ to $O\left(\sum_{i} \mid\right.$ cluster $\left.\left._{i}\right|^{2}\right)$

■ Do we gain anything by "summarizing" the data using clustering?

■ Do we gain anything by "summarizing" the data using clustering?

## Idea:

- Cluster unlabeled data set $U \rightarrow\left\{C_{i} \subset U: i \in[k]\right\}$
- Treat each $C_{i}$ as a separate problem, e.g., query cluster centroidsDo we gain anything by "summarizing" the data using clustering?


## Idea:

- Cluster unlabeled data set $U \rightarrow\left\{C_{i} \subset U: i \in[k]\right\}$
- Treat each $C_{i}$ as a separate problem, e.g., query cluster centroids


## Problems:

- U may not have a good clustering structure or $\operatorname{sim}(\cdot, \cdot)$ may not be able to capture it
- How many clusters and at what granularity?
- Clusters of $x$ 's may not correlate well with label $y$.

Figure: the swiss roll dataset has no obvious clustering structure.

## Unified Derivation

WRITEME

Extensions
$\square$ Consider a neural network $f_{\theta}: \mathbb{R}^{d} \rightarrow[c]$ :

$$
\begin{aligned}
f_{\theta}(\mathbf{x}) & =\underset{y \in[c]}{\operatorname{argmax}} p_{\theta}(y \mid \mathbf{x}) \\
p_{\theta}(y \mid \mathbf{x}) & =\operatorname{softmax}\left(W \phi_{\omega}(\mathbf{x})\right)_{y}
\end{aligned}
$$

where:

- $\theta=\{W, \omega\}$ are parameters
- $\phi_{\omega}: \mathbb{R}^{d} \rightarrow \mathbb{R}^{k}$ is an embedding function (e.g., convolutions + pooling layers)
- $W \in \mathbb{R}^{c \times k}$ are the parameters of the top dense layer


## Deep Architectures

- Deep NNs have a number of quirks:


## Deep Architectures

■ Deep NNs have a number of quirks:

- Very overconfident even away from the training set: their uncertainty cannot be trusted $\rightarrow$ strategies based on confidence, margin, entropy will underperform (including uncertainty sampling, model improvement, density-aware sampling, etc.)


## Deep Architectures

- Deep NNs have a number of quirks:
- Very overconfident even away from the training set: their uncertainty cannot be trusted $\rightarrow$ strategies based on confidence, margin, entropy will underperform (including uncertainty sampling, model improvement, density-aware sampling, etc.)
- Expensive to fit on data: training a ResNet on a realistic data set can take minutes to days $\rightarrow$ hard to ensure responsivity


## Deep Architectures

- Deep NNs have a number of quirks:
- Very overconfident even away from the training set: their uncertainty cannot be trusted $\rightarrow$ strategies based on confidence, margin, entropy will underperform (including uncertainty sampling, model improvement, density-aware sampling, etc.)
- Expensive to fit on data: training a ResNet on a realistic data set can take minutes to days $\rightarrow$ hard to ensure responsivity
- Quite insensitive to the addition of a single example $\rightarrow$ what's the point of querying individual instances?


## Deep Architectures

■ Deep NNs have a number of quirks:

- Very overconfident even away from the training set: their uncertainty cannot be trusted $\rightarrow$ strategies based on confidence, margin, entropy will underperform (including uncertainty sampling, model improvement, density-aware sampling, etc.)
- Expensive to fit on data: training a ResNet on a realistic data set can take minutes to days $\rightarrow$ hard to ensure responsivity
- Quite insensitive to the addition of a single example $\rightarrow$ what's the point of querying individual instances?
- Training is stochastic (i.e., not $100 \%$ stable) $\rightarrow$ changes in performance can depend on factors other than new labeled examples, high variance


## Overconfidence

Problem: Deep NNs tend to be very overconfident even away from the training set $\rightarrow$ their uncertainty cannot be trusted

## Overconfidence

Problem: Deep NNs tend to be very overconfident even away from the training set $\rightarrow$ their uncertainty cannot be trusted


Figure 1. Binary classification on a toy dataset using a MAP estimate, temperature scaling, and both last-layer and all-layer Gaussian approximations over the weights which are obtained via Laplace approximations. Background color and black line represent confidence and decision boundary, respectively. Bottom row shows a zoomed-out view of the top row. The Bayesian approximations-even in the last-layer case-give desirable uncertainty estimates: confident close to the training data and uncertain otherwise. MAP and temperature scaling yield overconfident predictions. The optimal temperature is picked as in Guo et al. (2017).

Credit: [Kristiadi et al., 2020].

## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.

## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.

- Aleatoric uncertainty ("random") captures how much we can trust the supervision itself. It cannot be decreased. (left)


## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.
$\square$ Aleatoric uncertainty ("random") captures how much we can trust the supervision itself. It cannot be decreased. (left)

■ Epistemic uncertainty ("relating to knowledge") captures how little we know about the world. This reflects on uncertainty on the choice of $\theta$. It decreases by acquiring more data. (right)

## Aleatoric vs Epistemic [Hüllermeier and Waegeman, 2021]



Figure 5: Left: Even with precise knowledge about the optimal hypothesis, the prediction at the query point (indicated by a question mark) is aleatorically uncertain, because the two classes are overlapping in that region. Right: A case of epistemic uncertainty due to a lack of knowledge about the right hypothesis, which is in turn caused by a lack of data.

■ Aleatoric uncertainty ("random") captures how much we can trust the supervision itself. It cannot be decreased. (left)

■ Epistemic uncertainty ("relating to knowledge") captures how little we know about the world. This reflects on uncertainty on the choice of $\theta$. It decreases by acquiring more data. (right)

- There isn't much point in trying to reduce aleatoric uncertainty in AL [Sharma and Bilgic, 2017]


## Bayesian NNs

- The problem with NNs is that uncertainty depends on a single model:
- This gives poor epistemic uncertainty


## Bayesian NNs

- The problem with NNs is that uncertainty depends on a single model:
- This gives poor epistemic uncertainty
- Using ensambles of NNs is computationally challenging: training one $N N$ is expensive, training $k$ even more so


## Bayesian NNs

- The problem with NNs is that uncertainty depends on a single model:
- This gives poor epistemic uncertainty
- Using ensambles of NNs is computationally challenging: training one NN is expensive, training $k$ even more so
- Using Bayesian techniques - i.e., maintaining a distribution over alternative NNs - is also challenging.


## Bayesian NNs

The problem with NNs is that uncertainty depends on a single model:

- This gives poor epistemic uncertainty
- Using ensambles of NNs is computationally challenging: training one NN is expensive, training $k$ even more so
- Using Bayesian techniques - i.e., maintaining a distribution over alternative NNs - is also challenging.

Idea of Bayesian NNs:

- Replace parameters $\theta$ with distribution over alternative parameters $p(\theta \mid L)$


## Bayesian NNs

The problem with NNs is that uncertainty depends on a single model:

- This gives poor epistemic uncertainty
- Using ensambles of NNs is computationally challenging: training one $N N$ is expensive, training $k$ even more so
- Using Bayesian techniques - i.e., maintaining a distribution over alternative NNs - is also challenging.

Idea of Bayesian NNs:

- Replace parameters $\theta$ with distribution over alternative parameters $p(\theta \mid L)$
- Compute predictions by marginalizing over $\theta$ :

$$
\begin{equation*}
p(y \mid \mathbf{x})=\int \underbrace{p(y \mid \mathbf{x}, \theta)}_{\text {NN with params } \theta} \cdot \underbrace{p(\theta \mid L)}_{\text {posterior over params }} d \theta \tag{68}
\end{equation*}
$$

## Bayesian NNs

- The problem with NNs is that uncertainty depends on a single model:
- This gives poor epistemic uncertainty
- Using ensambles of NNs is computationally challenging: training one $N N$ is expensive, training $k$ even more so
- Using Bayesian techniques - i.e., maintaining a distribution over alternative NNs - is also challenging.


## Idea of Bayesian NNs:

- Replace parameters $\theta$ with distribution over alternative parameters $p(\theta \mid L)$
- Compute predictions by marginalizing over $\theta$ :

$$
\begin{equation*}
p(y \mid \mathbf{x})=\int \underbrace{p(y \mid \mathbf{x}, \theta)}_{\text {NN with params } \theta} \cdot \underbrace{p(\theta \mid L)}_{\text {posterior over params }} d \theta \tag{68}
\end{equation*}
$$

- Learn by updating distribution:

$$
\begin{equation*}
p(\theta \mid L) \quad \rightarrow \quad p(\theta \mid L \cup\{(\mathrm{x}, y)\}) \tag{69}
\end{equation*}
$$

Not trivial! Is there an efficient approximation?

## Dropout

- Randomly set nodes to 0 with a fixed probability.


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.

- Used as a regularization technique: by randomly removing neurons, prevents them from relying on each other "too much"


## Dropout as Bayesian Approximation

■ Computing class probabilities:

$$
\begin{align*}
p(y \mid \mathbf{x}, L) & =\int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta}  \tag{70}\\
& \approx \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {dropout }}(\boldsymbol{\theta}) d \boldsymbol{\theta}  \tag{71}\\
& \approx \frac{1}{R} \sum_{r=1}^{R} p\left(y \mid \mathbf{x}, \widehat{\boldsymbol{\theta}}_{r}\right), \quad \widehat{\boldsymbol{\theta}}_{r} \sim p_{\text {dropout }}(\boldsymbol{\theta}) \tag{72}
\end{align*}
$$

In other words, run NN $R$ times with dropout enabled (during inference!) then average the $R$ vectors of class probabilities.

## Dropout as Bayesian Approximation

■ Computing class probabilities:

$$
\begin{align*}
p(y \mid \mathbf{x}, L) & =\int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta}  \tag{70}\\
& \approx \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {dropout }}(\boldsymbol{\theta}) d \boldsymbol{\theta}  \tag{71}\\
& \approx \frac{1}{R} \sum_{r=1}^{R} p\left(y \mid \mathbf{x}, \widehat{\boldsymbol{\theta}}_{r}\right), \quad \widehat{\boldsymbol{\theta}}_{r} \sim p_{\text {dropout }}(\boldsymbol{\theta}) \tag{72}
\end{align*}
$$

In other words, run NN $R$ times with dropout enabled (during inference!) then average the $R$ vectors of class probabilities.

■ Dropout can be viewed as variational Bayesian approximation where the approximating distribution is a mixture of two Gaussians [Gal and Ghahramani, 2016].

## Dropout as Bayesian Approximation

■ Computing class probabilities:

$$
\begin{align*}
p(y \mid \mathbf{x}, L) & =\int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta}  \tag{70}\\
& \approx \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {dropout }}(\boldsymbol{\theta}) d \boldsymbol{\theta}  \tag{71}\\
& \approx \frac{1}{R} \sum_{r=1}^{R} p\left(y \mid \mathbf{x}, \widehat{\boldsymbol{\theta}}_{r}\right), \quad \widehat{\boldsymbol{\theta}}_{r} \sim p_{\text {dropout }}(\boldsymbol{\theta}) \tag{72}
\end{align*}
$$

In other words, run NN $R$ times with dropout enabled (during inference!) then average the $R$ vectors of class probabilities.

■ Dropout can be viewed as variational Bayesian approximation where the approximating distribution is a mixture of two Gaussians [Gal and Ghahramani, 2016]. The approximation is independent of $L \rightarrow$ no training required.

## Dropout as Bayesian Approximation

■ Computing class probabilities:

$$
\begin{align*}
p(y \mid \mathbf{x}, L) & =\int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta}  \tag{70}\\
& \approx \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {dropout }}(\boldsymbol{\theta}) d \boldsymbol{\theta}  \tag{71}\\
& \approx \frac{1}{R} \sum_{r=1}^{R} p\left(y \mid \mathbf{x}, \widehat{\boldsymbol{\theta}}_{r}\right), \quad \widehat{\boldsymbol{\theta}}_{r} \sim p_{\text {dropout }}(\boldsymbol{\theta}) \tag{72}
\end{align*}
$$

In other words, run NN $R$ times with dropout enabled (during inference!) then average the $R$ vectors of class probabilities.

■ Dropout can be viewed as variational Bayesian approximation where the approximating distribution is a mixture of two Gaussians [Gal and Ghahramani, 2016]. The approximation is independent of $L \rightarrow$ no training required.

■ Immediately leads to more calibrated output probabilities!

## BALD [Gal et al., 2017]

Question: does dropout help with query selection too?

## BALD [Gal et al., 2017]

Question: does dropout help with query selection too? Yes.

- Uncertainty sampling:

$$
\begin{equation*}
\operatorname{acq}_{U N C}(\mathrm{x})=-\sum_{y \in[c]} p(Y=y \mid \mathrm{x}, L) \log p(Y=y \mid \mathrm{x}, L) \tag{73}
\end{equation*}
$$

Application is immediate.

## BALD [Gal et al., 2017]

Question: does dropout help with query selection too? Yes.

■ Uncertainty sampling:

$$
\begin{equation*}
\operatorname{acquNC}(\mathrm{x})=-\sum_{y \in[c]} p(Y=y \mid \mathrm{x}, L) \log p(Y=y \mid \mathrm{x}, L) \tag{73}
\end{equation*}
$$

Application is immediate.

- Mutual information between predictions and model posterior (BALD):

$$
\begin{equation*}
\operatorname{acq}_{B A L D}(\mathbf{x})=H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})] \tag{74}
\end{equation*}
$$

## BALD [Gal et al., 2017]

Question: does dropout help with query selection too? Yes.

■ Uncertainty sampling:

$$
\begin{equation*}
\operatorname{acquNC}(\mathrm{x})=-\sum_{y \in[c]} p(Y=y \mid \mathrm{x}, L) \log p(Y=y \mid \mathrm{x}, L) \tag{73}
\end{equation*}
$$

Application is immediate.

- Mutual information between predictions and model posterior (BALD):

$$
\begin{equation*}
\operatorname{acq}_{B A L D}(\mathrm{x})=H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})] \tag{74}
\end{equation*}
$$

- Left: entropy of the prediction $\rightarrow$ high when the model's prediction is uncertain


## BALD [Gal et al., 2017]

Question: does dropout help with query selection too? Yes.

■ Uncertainty sampling:

$$
\begin{equation*}
\operatorname{acquNC}(\mathrm{x})=-\sum_{y \in[c]} p(Y=y \mid \mathrm{x}, L) \log p(Y=y \mid \mathrm{x}, L) \tag{73}
\end{equation*}
$$

Application is immediate.

- Mutual information between predictions and model posterior (BALD):

$$
\begin{equation*}
\operatorname{acq}_{B A L D}(\mathrm{x})=H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})] \tag{74}
\end{equation*}
$$

- Left: entropy of the prediction $\rightarrow$ high when the model's prediction is uncertain
- Right: expected entropy of the prediction over the posterior of the model parameters $\rightarrow$ low when the model is overall certain for each draw of model parameters from the posterior.


## BALD [Gal et al., 2017]

Question: does dropout help with query selection too? Yes.

■ Uncertainty sampling:

$$
\begin{equation*}
\operatorname{acq}_{U N C}(\mathrm{x})=-\sum_{y \in[c]} p(Y=y \mid \mathrm{x}, L) \log p(Y=y \mid \mathrm{x}, L) \tag{73}
\end{equation*}
$$

Application is immediate.

- Mutual information between predictions and model posterior (BALD):

$$
\begin{equation*}
\operatorname{acq}_{B A L D}(\mathrm{x})=H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})] \tag{74}
\end{equation*}
$$

- Left: entropy of the prediction $\rightarrow$ high when the model's prediction is uncertain
- Right: expected entropy of the prediction over the posterior of the model parameters $\rightarrow$ low when the model is overall certain for each draw of model parameters from the posterior.

High when model has many possible ways to explain the data, i.e., the posterior draws are disagreeing among themselves.

- BALD can be computed as:

$$
\begin{aligned}
& H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})] \\
&=-\sum_{y \in[c]} p(y \mid \mathbf{x}, L) \log p(y \mid \mathbf{x}, L)+\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
&=-\sum_{y \in[c]} \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta} \cdot \log \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta}+\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
& \approx-\sum_{y \in[c]} \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {drop }}(\theta) d \boldsymbol{\theta} \cdot \log \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {drop }}(\theta) d \boldsymbol{\theta}+\mathbb{E}_{\boldsymbol{\theta} \sim p_{\text {drop }}(\theta)}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
& \approx-\sum_{y \in[c]}\left(\frac{1}{R} \sum_{r} \hat{p}_{y}^{r}\right) \log \left(\frac{1}{R} \sum_{r} \hat{p}_{y}^{r}\right)+\frac{1}{R} \sum_{y, r} \hat{p}_{y}^{r} \log \hat{p}_{y}^{r},
\end{aligned}
$$

where $\hat{p}_{y}^{r}=p\left(y \mid \mathbf{x}, \boldsymbol{\theta}_{r}\right)$ is the output of the NN's softmax for class $y$ and randomized parameters $\boldsymbol{\theta}_{r} \sim p_{\text {drop }}(\boldsymbol{\theta})$.

- BALD can be computed as:

$$
\begin{aligned}
& H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})] \\
&=-\sum_{y \in[c]} p(y \mid \mathbf{x}, L) \log p(y \mid \mathbf{x}, L)+\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L L}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
&=-\sum_{y \in[c]} \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta} \cdot \log \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta}+\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
& \approx-\sum_{y \in[c]} \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {drop }}(\theta) d \boldsymbol{\theta} \cdot \log \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {drop }}(\theta) d \boldsymbol{\theta}+\mathbb{E}_{\boldsymbol{\theta} \sim p_{\text {drop }}(\theta)}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
& \approx-\sum_{y \in[c]}\left(\frac{1}{R} \sum_{r} \hat{p}_{y}^{r}\right) \log \left(\frac{1}{R} \sum_{r} \hat{p}_{y}^{r}\right)+\frac{1}{R} \sum_{y, r} \hat{p}_{y}^{r} \log \hat{p}_{y}^{r},
\end{aligned}
$$

where $\hat{p}_{y}^{r}=p\left(y \mid \mathbf{x}, \boldsymbol{\theta}_{r}\right)$ is the output of the NN's softmax for class $y$ and randomized parameters $\boldsymbol{\theta}_{r} \sim p_{\text {drop }}(\boldsymbol{\theta})$.

■ Only need to compute $\hat{\mathbf{p}}^{r}$ once per $r \in[R]$, for a modest $R$ times increase in runtime; could be parallelized.

- BALD can be computed as:

$$
\begin{aligned}
& H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})] \\
&=-\sum_{y \in[c]} p(y \mid \mathbf{x}, L) \log p(y \mid \mathbf{x}, L)+\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L L}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
&=-\sum_{y \in[c]} \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta} \cdot \log \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid L) d \boldsymbol{\theta}+\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
& \approx-\sum_{y \in[c]} \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {drop }}(\theta) d \boldsymbol{\theta} \cdot \log \int p(y \mid \mathbf{x}, \boldsymbol{\theta}) p_{\text {drop }}(\theta) d \boldsymbol{\theta}+\mathbb{E}_{\boldsymbol{\theta} \sim p_{\text {drop }}(\theta)}\left[\sum_{y \in[c]} p(y \mid \mathbf{x}, \boldsymbol{\theta}) \log p(y \mid \mathbf{x}, \boldsymbol{\theta})\right] \\
& \approx-\sum_{y \in[c]}\left(\frac{1}{R} \sum_{r} \hat{p}_{y}^{r}\right) \log \left(\frac{1}{R} \sum_{r} \hat{p}_{y}^{r}\right)+\frac{1}{R} \sum_{y, r} \hat{p}_{y}^{r} \log \hat{p}_{y}^{r},
\end{aligned}
$$

where $\hat{p}_{y}^{r}=p\left(y \mid \mathbf{x}, \boldsymbol{\theta}_{r}\right)$ is the output of the NN's softmax for class $y$ and randomized parameters $\boldsymbol{\theta}_{r} \sim p_{\text {drop }}(\boldsymbol{\theta})$.

■ Only need to compute $\hat{\mathbf{p}}^{r}$ once per $r \in[R]$, for a modest $R$ times increase in runtime; could be parallelized.
$\square$ As $R \rightarrow \infty$, this approximates the original mutual information.

## Illustration



■ For all choices of acquisition function, the dropout-based uncertainty helps!
$\square$ Let us look at batch-based active learning.

## Batch Selection

Given $L, U$ and a classifier $f \in \mathcal{F}$ trained on $L$, find a batch $B \subseteq U$ of $b \gg 1$ unlabeled instances that brings maximal information to the model:

$$
\begin{gather*}
\underset{B \subseteq U}{\operatorname{argmax}} \operatorname{acq}_{B A L D}(f, B)  \tag{75}\\
\text { s.t. }|B|=b
\end{gather*}
$$

Let us look at batch-based active learning.

## Batch Selection

Given $L, U$ and a classifier $f \in \mathcal{F}$ trained on $L$, find a batch $B \subseteq U$ of $b \gg 1$ unlabeled instances that brings maximal information to the model:

$$
\begin{align*}
& \underset{B \subseteq U}{\operatorname{argmax}} \operatorname{acq}_{B A L D}(f, B)  \tag{75}\\
& \quad \text { s.t. }|B|=b
\end{align*}
$$

## Advantages:

- Only retrain the model after ever $b$ examples, meaning that supervision has an effect.
- Retraining is less frequent, leading to faster overall execution (at the expense of possibly instance selection, because $b$ examples depend on a fixed $f$ ).
- Supports parallel annotation for, e.g., crowd-sourcing scenarios.

Question: can regular acquisition function (like BALD) be extended to this setting?

- Natural generalization of instance-level strategies:

$$
\begin{equation*}
\operatorname{acq}(f, B)=\sum_{\mathrm{x} \in B} \operatorname{acq}(f, \mathrm{x}) \tag{77}
\end{equation*}
$$

How well does this work?

- Natural generalization of instance-level strategies:

$$
\begin{equation*}
\operatorname{acq}(f, B)=\sum_{\mathrm{x} \in B} \operatorname{acq}(f, \mathrm{x}) \tag{77}
\end{equation*}
$$

How well does this work?
$\square$ This ignores correlation between instances in $\mathbf{x}$ :

- Even if all of them are informative, they may carry the same information
- We want $B$ to be informative as a whole!


## Illustration



Figure 1: Idealised acquisitions of BALD and Batch$B A L D$. If a dataset were to contain many (near) replicas for each data point, then BALD would select all replicas of a single informative data point at the expense of other informative data points, wasting data efficiency.


Figure 2: Performance on Repeated MNIST with acquisition size 10 . See section 4.1 for further details. BatchBALD outperforms BALD while BALD performs worse than random acquisition due to the replications in the dataset.

## BatchBALD

■ The problem with the "natural generalization":

$$
\begin{equation*}
\operatorname{acq}(f, B)=\sum_{\mathbf{x} \in B} \operatorname{acq}(f, \mathbf{x}) \tag{78}
\end{equation*}
$$

is that the sum doesn't consider the overlap between the information carried by different $\mathrm{x} \in b$.

## BatchBALD

The problem with the "natural generalization":

$$
\begin{equation*}
\operatorname{acq}(f, B)=\sum_{\mathbf{x} \in B} \operatorname{acq}(f, \mathbf{x}) \tag{78}
\end{equation*}
$$

is that the sum doesn't consider the overlap between the information carried by different $\mathrm{x} \in b$.

Idea: don't break the acquisition function into a sum! For BALD, this means replacing:

$$
\begin{equation*}
\sum_{\mathbf{x} \in B}\{\underbrace{H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})]}_{M I(Y, \Theta \mid \mathbf{x}, L)}\} \tag{79}
\end{equation*}
$$

with

$$
\begin{equation*}
\operatorname{MI}\left(\left\{Y_{1}, \ldots, Y_{b}\right\}, \Theta \mid\left\{\mathrm{x}_{1}, \ldots, \mathrm{x}_{b}\right\}, L\right) \tag{80}
\end{equation*}
$$

## BatchBALD

- The problem with the "natural generalization":

$$
\begin{equation*}
\operatorname{acq}(f, B)=\sum_{\mathbf{x} \in B} \operatorname{acq}(f, \mathbf{x}) \tag{78}
\end{equation*}
$$

is that the sum doesn't consider the overlap between the information carried by different $\mathrm{x} \in b$.

Idea: don't break the acquisition function into a sum! For BALD, this means replacing:

$$
\begin{equation*}
\sum_{\mathbf{x} \in B}\{\underbrace{H(Y \mid \mathbf{x}, L)-\mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta} \mid L)}[H(Y \mid \mathbf{x}, \boldsymbol{\theta})]}_{M I(Y, \Theta \mid \mathbf{x}, L)}\} \tag{79}
\end{equation*}
$$

with

$$
\begin{equation*}
\operatorname{MI}\left(\left\{Y_{1}, \ldots, Y_{b}\right\}, \Theta \mid\left\{\mathrm{x}_{1}, \ldots, \mathrm{x}_{b}\right\}, L\right) \tag{80}
\end{equation*}
$$

$\square$ In other words, don't assume independence between the elements of $B$ !

## Illustration



(a) BALD


(b) BatchBALD

Figure 3: Intuition behind BALD and BatchBALD using I-diagrams [30]. BALD overestimates the joint mutual information. BatchBALD, however, takes the overlap between variables into account and will strive to acquire a better cover of $\omega$. Areas contributing to the respective score are shown in grey, and areas that are double-counted in dark grey.
(Credit: [Kirsch et al., 2019].)

Problem: batch selection amounts to solving

$$
\begin{equation*}
\underset{B \subset U \cup|B|=b}{\operatorname{argmax}} \operatorname{MI}\left(\left\{Y_{1}, \ldots, Y_{b}\right\}, \Theta \mid\left\{\mathbf{x}_{1}, \ldots, \mathbf{x}_{b}\right\}, L\right) \tag{81}
\end{equation*}
$$

How can we solve this?

Problem: batch selection amounts to solving

$$
\begin{equation*}
\underset{B \subseteq U:|B|=b}{\operatorname{argmax}} \operatorname{MI}\left(\left\{Y_{1}, \ldots, Y_{b}\right\}, \Theta \mid\left\{\mathbf{x}_{1}, \ldots, \mathrm{x}_{b}\right\}, L\right) \tag{81}
\end{equation*}
$$

How can we solve this?

## Submodular Function [Krause and Guestrin, 2008]

Let $S$ be a set. A function $f$ that maps subsets of $S$ to real values is submodular if for every $B \subset A \subseteq S$ and any $x \in S \backslash A$ it holds that:

$$
\begin{equation*}
f(A \cup\{x\})-f(A) \leq f(B \cup\{x\})-f(B) \tag{82}
\end{equation*}
$$

$f$ enjoys a diminishing returns property: adding an element $x$ to a smaller set $B$ "adds more" than adding the same element to a superset $A \supset B$.

Problem: batch selection amounts to solving

$$
\begin{equation*}
\underset{B \subseteq U:|B|=b}{\operatorname{argmax}} \operatorname{MI}\left(\left\{Y_{1}, \ldots, Y_{b}\right\}, \Theta \mid\left\{\mathbf{x}_{1}, \ldots, \mathrm{x}_{b}\right\}, L\right) \tag{81}
\end{equation*}
$$

How can we solve this?

## Submodular Function [Krause and Guestrin, 2008]

Let $S$ be a set. A function $f$ that maps subsets of $S$ to real values is submodular if for every $B \subset A \subseteq S$ and any $x \in S \backslash A$ it holds that:

$$
\begin{equation*}
f(A \cup\{x\})-f(A) \leq f(B \cup\{x\})-f(B) \tag{82}
\end{equation*}
$$

$f$ enjoys a diminishing returns property: adding an element $x$ to a smaller set $B$ "adds more" than adding the same element to a superset $A \supset B$.

## Maximizing a Submodular Functions

Let $f(A)$ be submodular and $S$ the domain. Then, the greedy algorithm:

- $A_{1} \leftarrow \varnothing$.
- $A_{t+1} \leftarrow \operatorname{argmax}_{x \in\left(S \backslash A_{t}\right)} f\left(A_{t} \cup\{x\}\right)$.
- Stop when budget $T$ is exhausted.
finds $A_{T} \subseteq S$ that has score $\left(1-\frac{1}{e}\right) \approx 67 \%$ as good as the score of the global optimum $A^{*} \subseteq S$ of $f$.

Problem: batch selection amounts to solving

$$
\begin{equation*}
\underset{B \subseteq U:|B|=b}{\operatorname{argmax}} \operatorname{MI}\left(\left\{Y_{1}, \ldots, Y_{b}\right\}, \Theta \mid\left\{\mathbf{x}_{1}, \ldots, \mathbf{x}_{b}\right\}, L\right) \tag{83}
\end{equation*}
$$

How can we solve this?

Problem: batch selection amounts to solving

$$
\begin{equation*}
\underset{B \subseteq U:|B|=b}{\operatorname{argmax}} \operatorname{MI}\left(\left\{Y_{1}, \ldots, Y_{b}\right\}, \Theta \mid\left\{\mathbf{x}_{1}, \ldots, \mathrm{x}_{b}\right\}, L\right) \tag{83}
\end{equation*}
$$

How can we solve this?

- The mutual information is submodular! $\rightarrow$ apply greedy optimization:
- Pick $x_{1}$ to optimize $\operatorname{MI}\left(Y_{1}, \Theta \mid \mathrm{x}_{1}, L\right)$ (BALD)
- Pick $x_{t+1}$ to optimize $M I\left(Y_{t+1} \cup Y_{1: t}, \Theta \mid \mathrm{x}_{t+1} \cup \mathrm{x}_{1: t}, L\right)$ (BALD over updated MI)
where $B=\left\{\mathrm{x}_{1}, \ldots, \mathrm{x}_{b}\right\}$.

(Credit: [Kirsch et al., 2019].)

(Credit: [Kirsch et al., 2019].)


## Discriminative Active Learning

Idea: we want instances in $L$ to be distributed as ground-truth

- This involves getting $p^{*}(\mathbf{X})$ right
- It also involves getting all the modes in it right


## Discriminative Active Learning

Idea: we want instances in $L$ to be distributed as ground-truth

- This involves getting $p^{*}(\mathbf{X})$ right
- It also involves getting all the modes in it right
$\square$ Highly non-trivial: we could train a generative model $\hat{p}_{\theta}(X, Y)$ using density estimation and use that to guide query selection


## Discriminative Active Learning

Idea: we want instances in $L$ to be distributed as ground-truth

- This involves getting $p^{*}(\mathbf{X})$ right
- It also involves getting all the modes in it right
- Highly non-trivial: we could train a generative model $\hat{p}_{\theta}(X, Y)$ using density estimation and use that to guide query selection $\rightarrow$ hard to train, break down in high dimension (in general)


## Discriminative Active Learning

Idea: we want instances in $L$ to be distributed as ground-truth

- This involves getting $p^{*}(\mathbf{X})$ right
- It also involves getting all the modes in it right

■ Highly non-trivial: we could train a generative model $\hat{p}_{\theta}(\mathbf{X}, Y)$ using density estimation and use that to guide query selection $\rightarrow$ hard to train, break down in high dimension (in general)
$\square$ Is there any alternative?

## Discriminative Active Learning

Idea: if we cannot find a classifier $f \in \mathcal{F}$ that tells $L$ from $U$ apart, and the latter is large enough (i.e., it can be used to approximate the ground-truth distribution $p^{*}(\mathbf{X})$ ), then $L$ is high-quality.

## Discriminative Active Learning

Idea: if we cannot find a classifier $f \in \mathcal{F}$ that tells $L$ from $U$ apart, and the latter is large enough (i.e., it can be used to approximate the ground-truth distribution $p^{*}(\mathbf{X})$ ), then $L$ is high-quality.

■ Given $\mathrm{x} \in U$, how certain are we that it comes from $U$ rather than from $L$ ?

- if indistinguishable, we represented the true distribution using $L$
- if distinguishable, it looks different from $L$ so labeling it should be informative


## $\mathcal{F}$-divergence

Given two distribution $p_{S}(X)$ and $p_{T}(X)$ on $X \in \mathcal{X}$ and a hypothesis class $\mathcal{F}$ also on $\mathcal{X}$, the $\mathcal{F}$-divergence between $p_{S}$ and $p_{T}$ is:

$$
\begin{equation*}
d_{\mathcal{F}}\left(p_{S}, p_{T}\right)=2 \cdot \sup _{f \in \mathcal{F}}\left|p_{S}(\{x: f(x)=1\})-p_{T}(\{x: f(x)=1\})\right| \tag{84}
\end{equation*}
$$

## $\mathcal{F}$-divergence

Given two distribution $p_{S}(X)$ and $p_{T}(X)$ on $X \in \mathcal{X}$ and a hypothesis class $\mathcal{F}$ also on $\mathcal{X}$, the $\mathcal{F}$-divergence between $p_{S}$ and $p_{T}$ is:

$$
\begin{equation*}
d_{\mathcal{F}}\left(p_{S}, p_{T}\right)=2 \cdot \sup _{f \in \mathcal{F}}\left|p_{S}(\{x: f(x)=1\})-p_{T}(\{x: f(x)=1\})\right| \tag{84}
\end{equation*}
$$

$\square$ Measures how different two domains $p_{S}$ and $p_{T}$ are from the perspective of a model class $\mathcal{F}$ : the larger the difference, the more different they look.
$\square \mathcal{F}$-divergence used to identify concept drift and - symmetrically - to estimate how well a classifier trained on one task will perform on a different, related task

Idea: use $\mathcal{F}$-divergence to discriminate between $L$ and $U$

Idea: use $\mathcal{F}$-divergence to discriminate between $L$ and $U$

The $\mathcal{F}$-divergence is:

$$
\begin{equation*}
d_{\mathcal{F}}\left(p_{S}, p_{T}\right)=2 \cdot \sup _{f \in \mathcal{F}}\left|p_{S}(\{x: f(x)=1\})-p_{T}(\{x: f(x)=1\})\right| \tag{85}
\end{equation*}
$$

Idea: use $\mathcal{F}$-divergence to discriminate between $L$ and $U$
$\square$ The $\mathcal{F}$-divergence is:

$$
\begin{equation*}
d_{\mathcal{F}}\left(p_{S}, p_{T}\right)=2 \cdot \sup _{f \in \mathcal{F}}\left|p_{S}(\{x: f(x)=1\})-p_{T}(\{x: f(x)=1\})\right| \tag{85}
\end{equation*}
$$

- Approximate as follows:
- $p_{S}(\mathrm{x}):=\frac{1}{|L|} \sum_{\mathrm{x}^{\prime} \in L} \delta\left\{\mathrm{x}^{\prime}=\mathrm{x}\right\}$
- $p_{T}(\mathrm{x}):=\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} i n U} \delta\left\{\mathrm{x}^{\prime}=\mathrm{x}\right\}$.

Then $\mathcal{F}$-divergence becomes:

$$
\begin{equation*}
2 \sup _{f \in \mathcal{F}}\left|\frac{1}{|L|} \sum_{\mathrm{x} \in L} f(\mathrm{x})-\frac{1}{|U|} \sum_{\mathrm{x} \in U} f(\mathrm{x})\right| \tag{86}
\end{equation*}
$$

Idea: use $\mathcal{F}$-divergence to discriminate between $L$ and $U$
$\square$ The $\mathcal{F}$-divergence is:

$$
\begin{equation*}
d_{\mathcal{F}}\left(p_{S}, p_{T}\right)=2 \cdot \sup _{f \in \mathcal{F}}\left|p_{S}(\{x: f(x)=1\})-p_{T}(\{x: f(x)=1\})\right| \tag{85}
\end{equation*}
$$

- Approximate as follows:
- $p_{S}(\mathrm{x}):=\frac{1}{|L|} \sum_{\mathrm{x}^{\prime} \in L} \delta\left\{\mathrm{x}^{\prime}=\mathrm{x}\right\}$
- $p_{T}(\mathrm{x}):=\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} i n U} \delta\left\{\mathrm{x}^{\prime}=\mathrm{x}\right\}$.

Then $\mathcal{F}$-divergence becomes:

$$
\begin{equation*}
2 \sup _{f \in \mathcal{F}}\left|\frac{1}{|L|} \sum_{\mathrm{x} \in L} f(\mathrm{x})-\frac{1}{|U|} \sum_{\mathrm{x} \in U} f(\mathrm{x})\right| \tag{86}
\end{equation*}
$$

How to compute this?

Idea: use $\mathcal{F}$-divergence to discriminate between $L$ and $U$
$\square$ The $\mathcal{F}$-divergence is:

$$
\begin{equation*}
d_{\mathcal{F}}\left(p_{S}, p_{T}\right)=2 \cdot \sup _{f \in \mathcal{F}}\left|p_{S}(\{x: f(x)=1\})-p_{T}(\{x: f(x)=1\})\right| \tag{85}
\end{equation*}
$$

- Approximate as follows:
- $p_{S}(\mathrm{x}):=\frac{1}{|L|} \sum_{\mathrm{x}^{\prime} \in L} \delta\left\{\mathrm{x}^{\prime}=\mathrm{x}\right\}$
- $p_{T}(\mathrm{x}):=\frac{1}{|U|} \sum_{\mathrm{x}^{\prime} i n U} \delta\left\{\mathrm{x}^{\prime}=\mathrm{x}\right\}$.

Then $\mathcal{F}$-divergence becomes:

$$
\begin{equation*}
2 \sup _{f \in \mathcal{F}}\left|\frac{1}{|L|} \sum_{x \in L} f(\mathrm{x})-\frac{1}{|U|} \sum_{\mathrm{x} \in U} f(\mathrm{x})\right| \tag{86}
\end{equation*}
$$

- How to compute this? $\sup _{f \in \mathcal{F}}$ can be implemented by learning $f$ from data set:


## $\mathcal{F}$-divergence as learning

$\square$ Given $L$ and $U$, define binary classification task:

## $\mathcal{F}$-divergence as learning

$\square$ Given $L$ and $U$, define binary classification task:

- For all $\mathrm{x} \in L$, add $(\phi(\mathrm{x})$, labeled $)$ to dataset


## $\mathcal{F}$-divergence as learning

$\square$ Given $L$ and $U$, define binary classification task:

- For all $x \in L$, add $(\phi(x)$, labeled $)$ to dataset
- For all $\mathbf{x} \in U$, add $(\phi(x)$, unlabeled $)$ to dataset


## $\mathcal{F}$-divergence as learning

$\square$ Given $L$ and $U$, define binary classification task:

- For all $x \in L$, add $(\phi(x)$, labeled $)$ to dataset
- For all $\mathrm{x} \in U$, add $(\phi(\mathrm{x})$, unlabeled) to dataset
- Train (simple) classifier $p_{\psi}(\cdot \mid \phi(\mathrm{x}))$ to distinguish between the two sources by optimizing cross-entropy loss


## $\mathcal{F}$-divergence as learning

$\square$ Given $L$ and $U$, define binary classification task:

- For all $\mathrm{x} \in L$, add $(\phi(\mathrm{x})$, labeled) to dataset
- For all $\mathrm{x} \in U$, add $(\phi(\mathrm{x})$, unlabeled $)$ to dataset
- Train (simple) classifier $p_{\psi}(\cdot \mid \phi(\mathrm{x}))$ to distinguish between the two sources by optimizing cross-entropy loss

This implicitly makes use of the labels $y$ in $L$ through $\phi(x)$.

## $\mathcal{F}$-divergence as learning

$\square$ Given $L$ and $U$, define binary classification task:

- For all $\mathrm{x} \in L$, add $(\phi(\mathrm{x})$, labeled $)$ to dataset
- For all $x \in U$, add $(\phi(x)$, unlabeled $)$ to dataset
- Train (simple) classifier $p_{\psi}(\cdot \mid \phi(\mathrm{x}))$ to distinguish between the two sources by optimizing cross-entropy loss

This implicitly makes use of the labels $y$ in $L$ through $\phi(x)$.

■ Pick those instances $\mathrm{x} \in U$ that have lowest probability $p_{\psi}$ (labeled $\left.\mid \boldsymbol{\phi}(\mathrm{x})\right)$

## $\mathcal{F}$-divergence as learning

$\square$ Given $L$ and $U$, define binary classification task:

- For all $\mathrm{x} \in L$, add $(\phi(\mathrm{x})$, labeled $)$ to dataset
- For all $x \in U$, add $(\phi(x)$, unlabeled $)$ to dataset
- Train (simple) classifier $p_{\psi}(\cdot \mid \phi(\mathrm{x}))$ to distinguish between the two sources by optimizing cross-entropy loss

This implicitly makes use of the labels $y$ in $L$ through $\phi(x)$.
$\square$ Pick those instances $\mathrm{x} \in U$ that have lowest probability $p_{\psi}$ (labeled $\left.\mid \boldsymbol{\phi}(\mathrm{x})\right)$

Remark: this is closely related to GANs!

## Conclusion and Further Reading

## Take-away

$\square$ AL useful when supervision is expensive high $\rightarrow$ choose it wisely
■ Many variants: pool-based, streaming, and query synthesis
■ Many practical approaches: uncertainty-based (uncertainty sampling, QBC, expected gradient length), diversity-based (information density).

Some can be derived from version spaces and model improvement.

- Deep variants select entire batches and often rely on Bayesian techniques

■ Critique \& realistic annotators, costs, etc.: [Herde et al., 2021] [Settles, 2011]
■ Plenty of room for new research ;-)

Baum, E. B. and Lang, K. (1992)
Query learning can work poorly when a human oracle is used.
In International joint conference on neural networks, volume 8, page 8.Gal, Y. and Ghahramani, Z. (2016).
Dropout as a bayesian approximation: Representing model uncertainty in deep learning.
In international conference on machine learning, pages 1050-1059. PMLR.Gal, Y., Islam, R., and Ghahramani, Z. (2017).
Deep bayesian active learning with image data.
In International Conference on Machine Learning, pages 1183-1192. PMLR
Herde, M., Huseljic, D., Sick, B., and Calma, A. (2021).
A survey on cost types, interaction schemes, and annotator performance models in selection algorithms for active learning in classification.
arXiv preprint arXiv:2109.11301.
Hüllermeier, E. and Waegeman, W. (2021).
Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. Machine Learning, 110(3):457-506.

King, R. D., Rowland, J., Oliver, S. G., Young, M., Aubrey, W., Byrne, E., Liakata, M., Markham, M., Pir, P., Soldatova, L. N., et al. (2009).
The automation of science.

Science, 324(5923):85-89.
Kirsch, A., Van Amersfoort, J., and Gal, Y. (2019).
Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning.
Advances in neural information processing systems, 32:7026-7037.Krause, A. and Guestrin, C. (2008).
Beyond convexity: Submodularity in machine learning.
ICML Tutorials.
Kristiadi, A., Hein, M., and Hennig, P. (2020).
Being bayesian, even just a bit, fixes overconfidence in relu networks.
In International Conference on Machine Learning, pages 5436-5446. PMLR.
Nguyen, A., Dosovitskiy, A., Yosinski, J., Brox, T., and Clune, J. (2016).
Synthesizing the preferred inputs for neurons in neural networks via deep generator networks.
Advances in neural information processing systems, 29:3387-3395.
Settles, B. (2011).
From theories to queries: Active learning in practice.
In Active Learning and Experimental Design workshop In conjunction with AISTATS 2010, pages 1-18.
JMLR Workshop and Conference Proceedings.
Settles, B. (2012).
Active learning.

Shalev-Shwartz, S. and Ben-David, S. (2014).
Understanding machine learning: From theory to algorithms.
Cambridge university press.
國 Sharma, M. and Bilgic, M. (2017).
Evidence-based uncertainty sampling for active learning.
Data Mining and Knowledge Discovery, 31(1):164-202.


[^0]:    ${ }^{1}$ Although it could be estimated from interaction data.

[^1]:    ${ }^{2}$ There is a point to doing so, as we will see later.

[^2]:    ${ }^{3}$ See https://en.wikipedia.org/wiki/Normal_distribution.

[^3]:    ${ }^{3}$ See https://en.wikipedia.org/wiki/Normal_distribution.

[^4]:    ${ }^{3}$ See https://en.wikipedia.org/wiki/Normal_distribution.

[^5]:    ${ }^{5}$ Note: MI is greedy, not optimal! Non-greedy alternatives are conceptually better, but they also computationally infeasible and for this reason they are ignored in the AL literature.

[^6]:    ${ }^{5}$ Note: MI is greedy, not optimal! Non-greedy alternatives are conceptually better, but they also computationally infeasible and for this reason they are ignored in the AL literature.

[^7]:    ${ }^{5}$ Note: MI is greedy, not optimal! Non-greedy alternatives are conceptually better, but they also computationally infeasible and for this reason they are ignored in the AL literature.

[^8]:    ${ }^{6}$ The unlabeled set $U$ is ideally pretty large, so the approximation is reasonable.

[^9]:    ${ }^{6}$ The unlabeled set $U$ is ideally pretty large, so the approximation is reasonable.

[^10]:    ${ }^{6}$ The unlabeled set $U$ is ideally pretty large, so the approximation is reasonable.

[^11]:    ${ }^{6}$ The unlabeled set $U$ is ideally pretty large, so the approximation is reasonable.

