Uncertainty Sampling Cheatsheet

When a Supervised Machine Learning model makes a prediction, it often gives a confidence in that prediction. If the model is uncertain (low confidence), then human feedback can help. Getting human feedback when a model is uncertain is a type of *Active Learning* known as *Uncertainty Sampling*.
This cheatsheet has four common ways to calculate uncertainty, with examples, equations and python code.

Example Machine Learning Prediction (x):

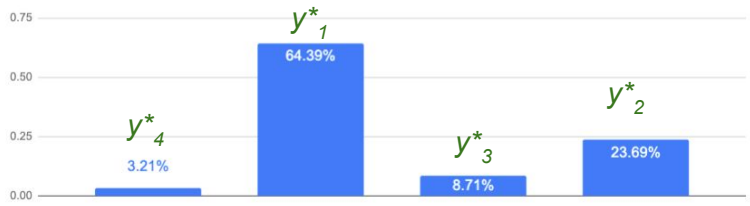

number of labels (n) = 4

The predictions are a probability distribution (x), meaning that every prediction is between 0 and 1 and the predictions add to 1. y^*_1 is the most confident, y^*_2 is the second most confident, etc. for n predicted labels.

This example can be expressed as a NumPy array:

```
prob_dist = np.array([0.0321, 0.6439, 0.0871, 0.2369])
```

Least Confidence: *difference between the most confident prediction and 100% confidence*

$$\frac{n\,(\,1 - P_\theta\,(\,y^*_1 \mid x\,)\,)}{n - 1}$$

```
most_conf = np.nanmax(prob_dist)
num_labels = prob_dist.size
numerator = (num_labels * (1 - most_conf))
denominator = (num_labels - 1)

least_conf = numerator / denominator
```

Margin of Confidence: difference between the top two most confident predictions

$$1 - (\,P_\theta(\,y^*_1 \mid x\,) - P_\theta(\,y^*_2 \mid x\,)\,)$$

```
prob_dist[::-1].sort()
difference = (prob_dist[0] - prob_dist[1])

margin_conf = 1 - difference
```

Ratio of Confidence: ratio between the top two most confident predictions

$$\frac{P_\theta(\,y^*_2 \mid x\,)}{P_\theta(\,y^*_1 \mid x\,)}$$

```
prob_dist[::-1].sort()

ratio_conf  = (prob_dist[1] / prob_dist[0])
```

Entropy: difference between all predictions, as defined by information theory

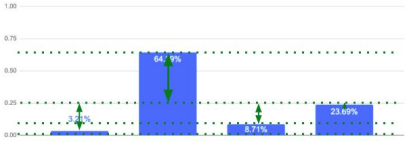

$$\frac{-\sum_y P_\theta(\,y \mid x\,)\,\log_2 P_\theta(\,y \mid x\,)\,)}{\log_2(n)}$$

```
prbslogs = prob_dist * np.log2(prob_dist)
numerator = 0 - np.sum(prbslogs)
denominator = np.log2(prob_dist.size)

entropy = numerator / denominator
```

Robert Munro. **Human-in-the-Loop Machine Learning**, *Manning Publications*. http://bit.ly/huml_book
See the book for more details on each method and for more sophisticated problems like sequence models and semantic segmentation, plus other sampling strategies like Diversity Sampling. robertmunro.com | @WWRob