Conformal Prediction approach of Confident Learning

Christian Ngnie [1]

[1]Code For Africa

What is Confident Learning

How to confidently learn in noisy settings?

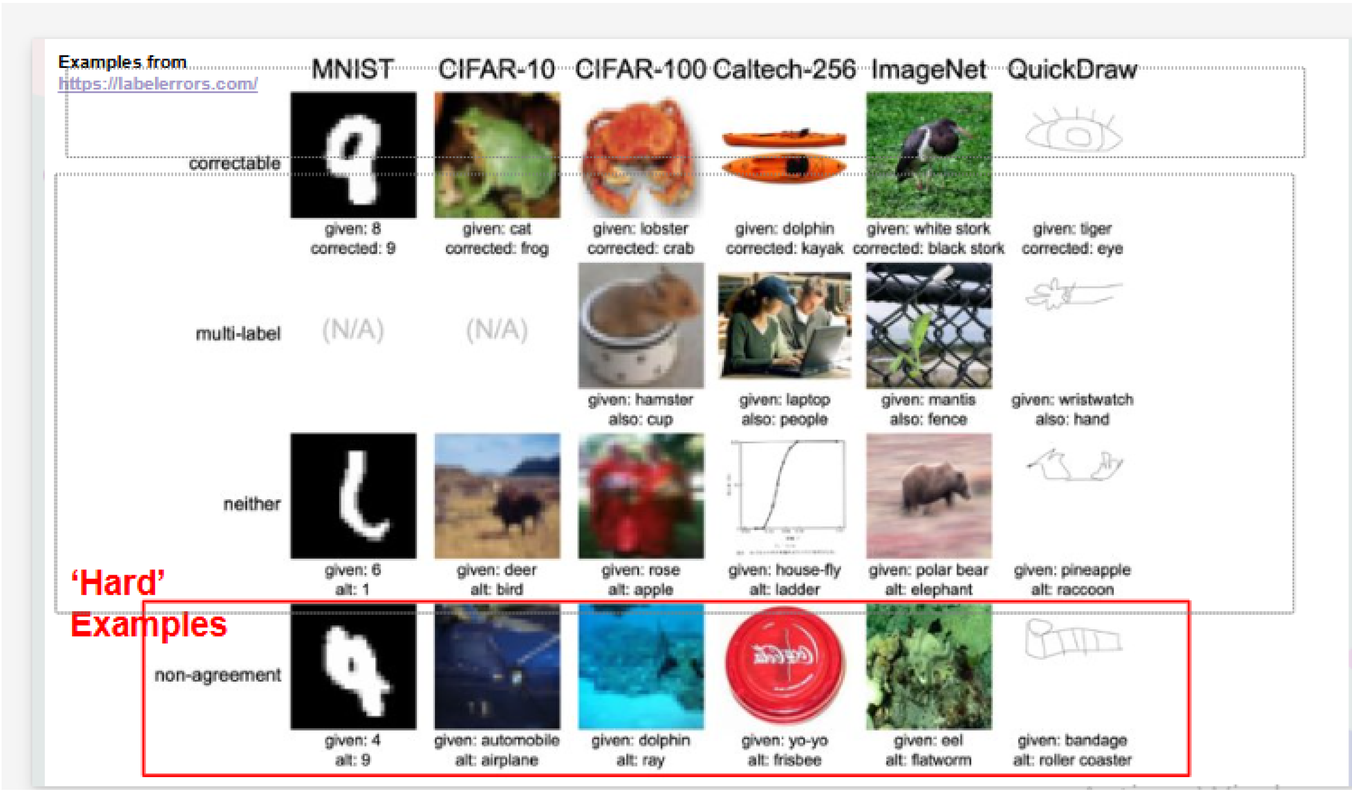

Figure 1. Some factors introducing noise in the labels

Confident Learning is a framework of theory and algorithm for:

- **Finding** label errors in a dataset
- **Ranking** data by likelihood of being label issue
- **Learning** with noisy labels
- **Complete** characterization of label noise in a dataset

Key idea: With confident Learning you can use any ML model's predicted probabilities to find label errors

- Uniform/symmetric class-conditional label noise: $P(\tilde{y} = i | y^* = j) = \epsilon, \forall i \neq j$
- Systematic/Asymmetric class-conditional label noise: $P(\tilde{y} = i | y^* = j)$ any valid distribution (the focus of CL)
- Instance-Dependent label noise: $P(\tilde{y} = i | y^* = j, x)$

How does confident learning work

Key idea: First we find thresholds as a proxy for the machine's self-confidence, on average for each task/class j

$$t_j = \frac{1}{|X_{\tilde{y}=j}|} \sum_{x \in X_{\tilde{y}=j}} \hat{P}(\tilde{y} = j; x; \theta)$$

$$\hat{X}_{\tilde{y}=i, y^*=j} := \left\{ x \in X_{\tilde{y}=i} : \hat{P}(\tilde{y} = j; x; \theta) \geq t_j, j = \underset{k, \hat{P}(\tilde{y}=k; x; \theta) \geq t_k}{\operatorname{argmax}} \hat{P}(\tilde{y} = k; x; \theta) \right\}$$

Sufficient condition

If $\hat{P}_{x,\tilde{y}=j} = P^*_{x,\tilde{y}=j} + \epsilon_{x,\tilde{y}=j}$, $\hat{X}_{\tilde{y}=i,y^*=j} \approxeq X_{\tilde{y}=i,y^*=j}$

where,

$$\epsilon_{x,\tilde{y}=j} \sim \begin{cases} U(\epsilon_j + t_j - P^*_{x,\tilde{y}=j}, \epsilon_j - t_j + P^*_{x,\tilde{y}=j}), \; P^*_{x,\tilde{y}=j} \geq t_j \\ U(\epsilon_j - t_j + P^*_{x,\tilde{y}=j}, \epsilon_j + t_j - P^*_{x,\tilde{y}=j}), \; P^*_{x,\tilde{y}=j} < t_j \end{cases} .$$

When CL fails

When the error in $\hat{P}(\tilde{y} = i; x; \theta)$ exceeds the threshold margins.

When might this happen?

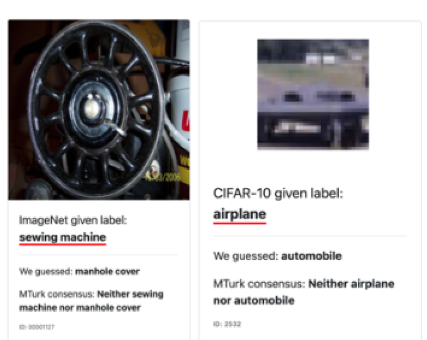

(really) hard examples

70%			
0	0.2	0.4	0.6
31.5	39.3	33.7	30.6
33.7	40.7	35.1	31.4
32.4	**41.8**	34.4	34.5
41.1	41.7	39.0	32.9
41.0	**41.8**	**39.1**	**36.4**

Acc. of CL-based methods for 70% noise for various settings.

too much (70+%) noise

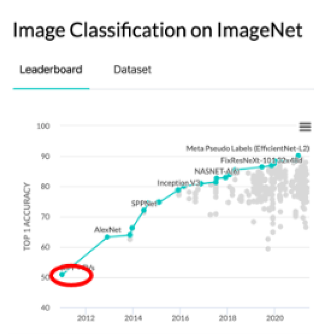

inappropriate model

Figure 2. Some cases where CL fails

Conformal prediction

1. **Theoretical guarantees** with little assumptions.
2. **Model-agnostic** and post processing
3. **Create the prediction set** $C_\alpha(X_{n+1})$ containing Y_{n+1} with statistical guarantees: $\mathbb{P}(Y_{n+1} \in C_\alpha(X_{n+1})) \geq 1 - \alpha$

Contribution: Conformalize CL with Jackknife+

idea: we threshold using empirical quantiles of class-based residuals $(q_j)_j$

$$t_j = \frac{1}{|X_{\tilde{y}=j}|} \sum_{x \in X_{\tilde{y}=j}} \hat{P}(\tilde{y} = j; x; \theta)$$

$$\hat{X}_{\tilde{y}=i, y^*=j} := \left\{ x \in X_{\tilde{y}=i} : \hat{P}(\tilde{y} = j; x; \theta) \geq 1 - q_j \right\}$$

Statistical guarantees

The set satisfies

$$\mathbb{P}(\tilde{y} = i \in \hat{X}_{\tilde{y}=i,y^*=j}) \geq 1 - 2\alpha$$

the framework identifies with at least $1 - 2\alpha$ guarantee an observed label as an error.

References

[1] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. Algorithmic learning in a random world. Springer, 2005.

[2] Anastasios N Angelopoulos, Stephen Bates, Clara Fannjiang, Michael I. Jordan, and Tijana Zrnic. Prediction-powered inference. arXiv preprint arXiv:2301.09633, 2023

[3] Yaniv Romano, Evan Patterson, and Emmanuel Candes. Conformalized quantile regression. In Advances in Neural Information Processing Systems, 2019b

[4] Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In International Conference on Machine Learning, pages 25834–25866. PMLR, 2022

[5] Curtis G. Northcutt, Lu Jiang, Isaac L. Chuang. Confident Learning: Estimating Uncertainty in Dataset Labels

[6] Anastasios N. Angelopoulos, Stephen Bates. A Gentle Introduction to Conformal Prediction and Distribution-Free Uncertainty Quantification

[7] Rina Foygel Barber, Emmanuel J. Candes, Aaditya Ramdas, Ryan J. Tibshirani. Predictive inference with the jackknife+