KD-Crowd: A Knowledge Distillation Framework for Learning from Crowds

Shao-Yuan LI, Yu-Xiang ZHENG, Ye SHI, Sheng-Jun HUANG, Songcan CHEN

Frontiers of Computer Science, DOI: 10.1007/s11704-023-3578-7

Problems & Ideas

- Problems of existing model worker expertise methods :
  - Transition matrix is oversimplified for complex noise patterns.
  - Estimating the unknown worker expertise and latent true labels interact, for which estimation errors occur and propagate easily.
- Ideas: A knowledge distillation framework combining the complementary strength of noise-model-free robust learning techniques and transition matrix-based worker expertise modeling.

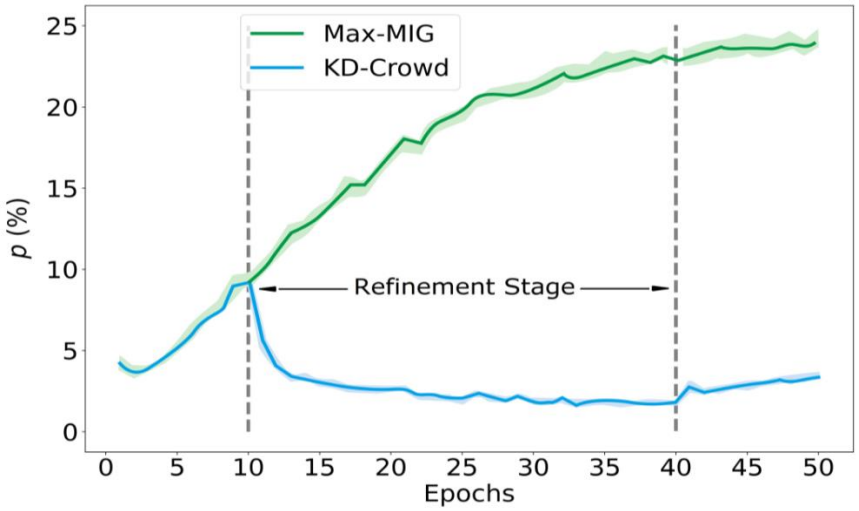

Illustrative example for dynamic memorization of annotation noise. The vertical axis means the fraction of wrong annotations that are memorized by each worker's true label predictions. The shaded regions for the curves encapsulate the range within one standard deviation of the respective means.

Main Contributions

- Contributions:
  - We are possibly the first to show that existing crowdsourcing methods easily overfit the annotation noise;
  - We propose KD-Crowd, a knowledge distillation framework that refines crowdsourcing learning using noise-free techniques. This includes an f-mutual information gain-based loss to avoid the student model learning the teacher's significant errors.;
  - We validate the effectiveness and generality of KD-Crowd over two different crowdsourcing models on synthetic and real-world data, showing that KD-Crowd achieves significant performance gains.

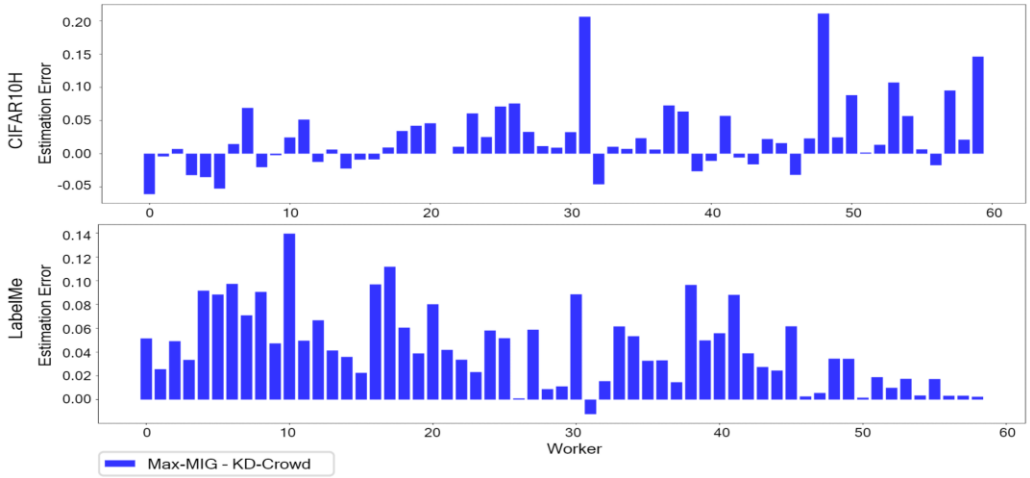

The gap between the overall error $\boldsymbol{E_m}$ of Max-MIG and KD-Crowd for all workers on CIFAR10H and LabelMe. $\boldsymbol{E_m}$ is calculated as the mean value of entries in $T_e^m = |T_m - \hat{T}_m|/T_m$