MIT-IBM Watson AI Lab

Carnegie Mellon University

NUS National University of Singapore

Collective Model Fusion for Multiple Black-Box Experts

MINH HOANG, NGHIA HOANG, BRYAN LOW, CARL KINGSFORD

Collaborative AI: A health-care scenario

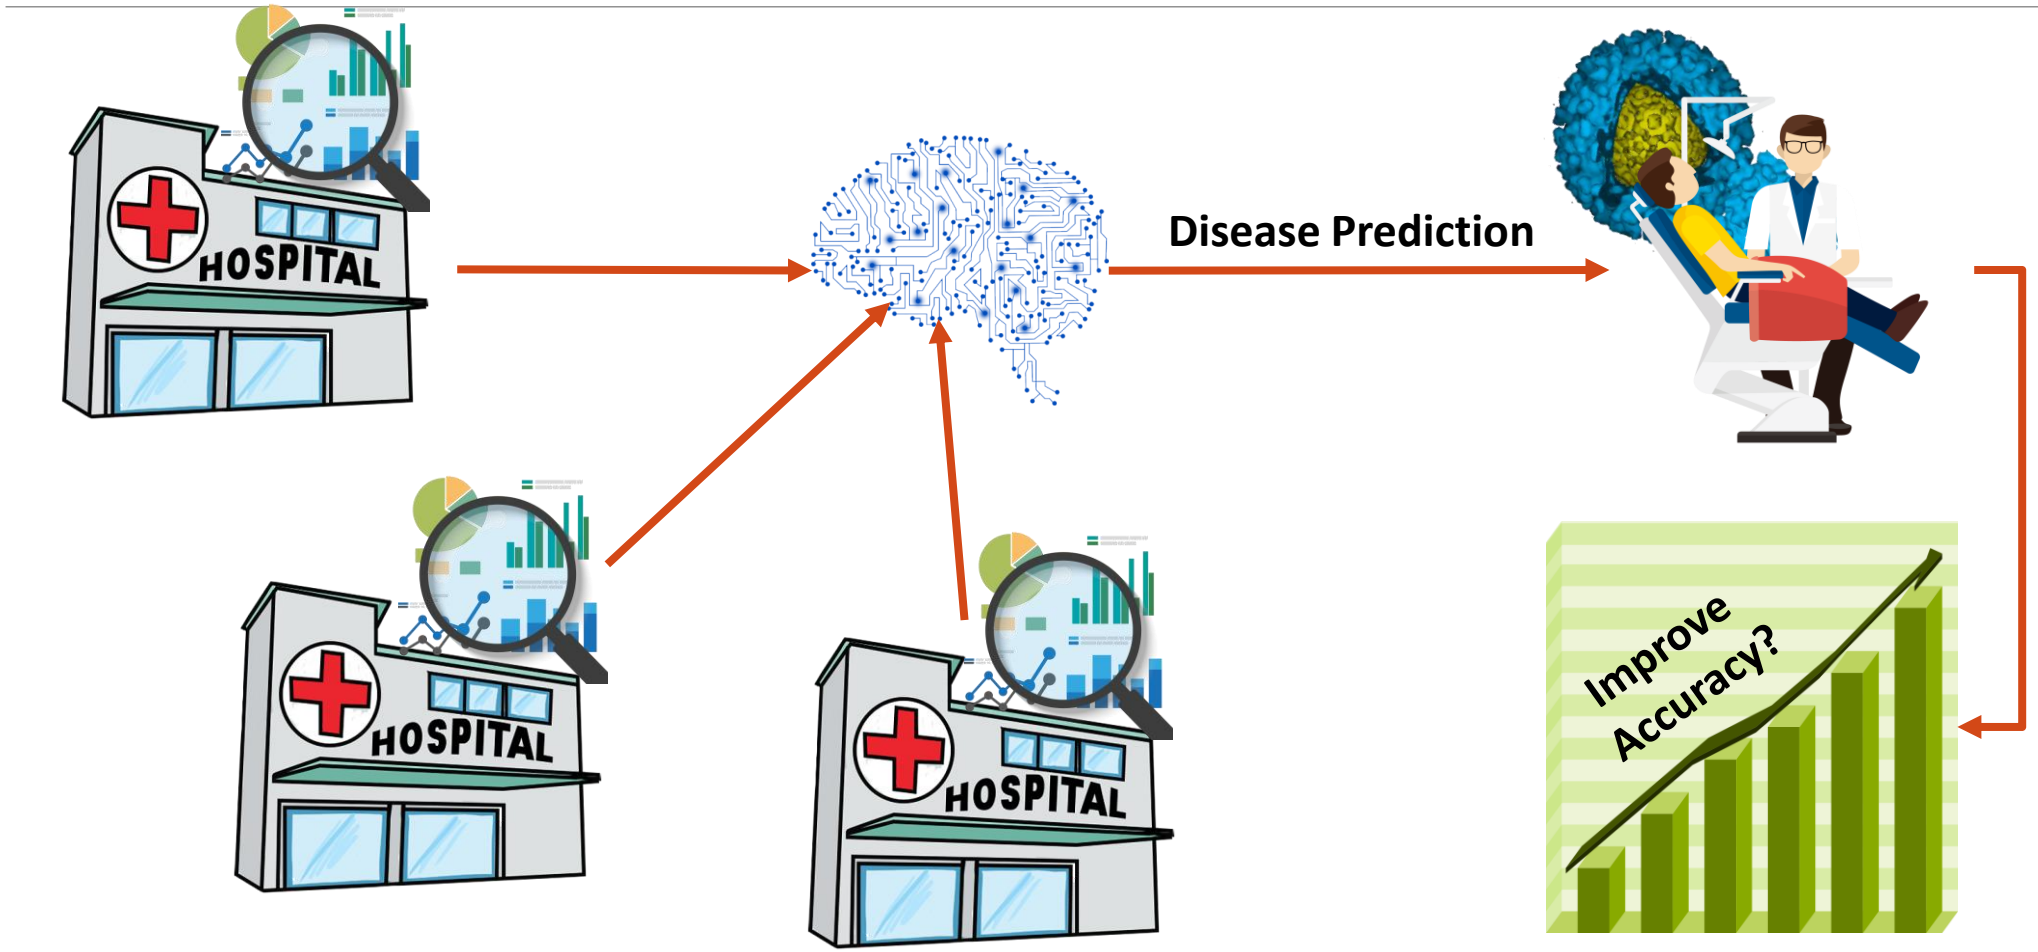

Related work: Data Fusion

Clinical Notes

Medical Codes

Vital Signs over time

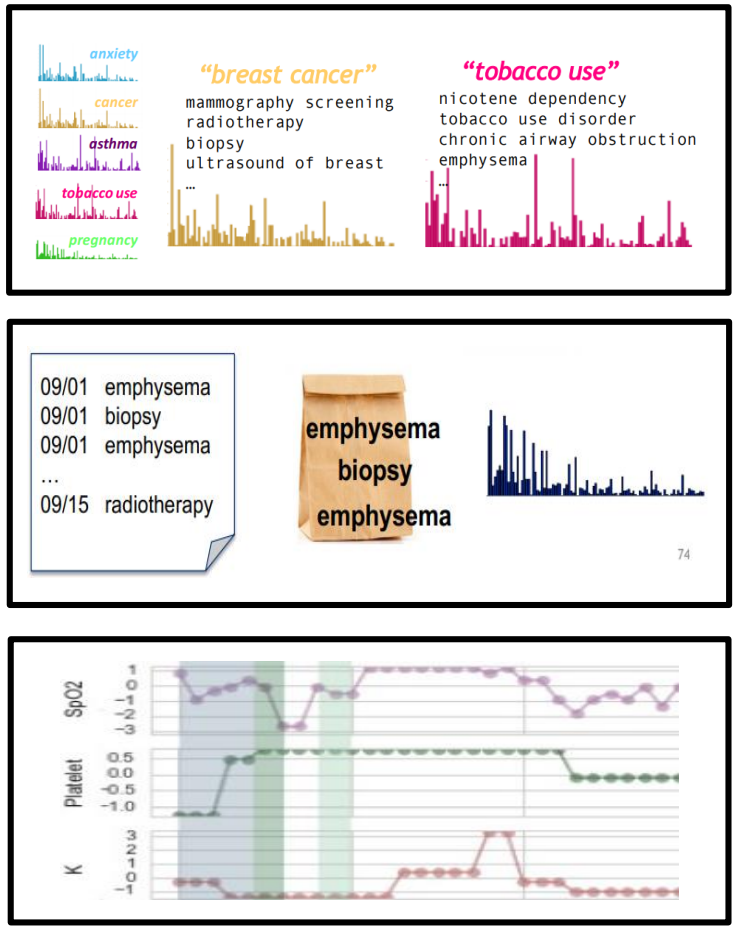

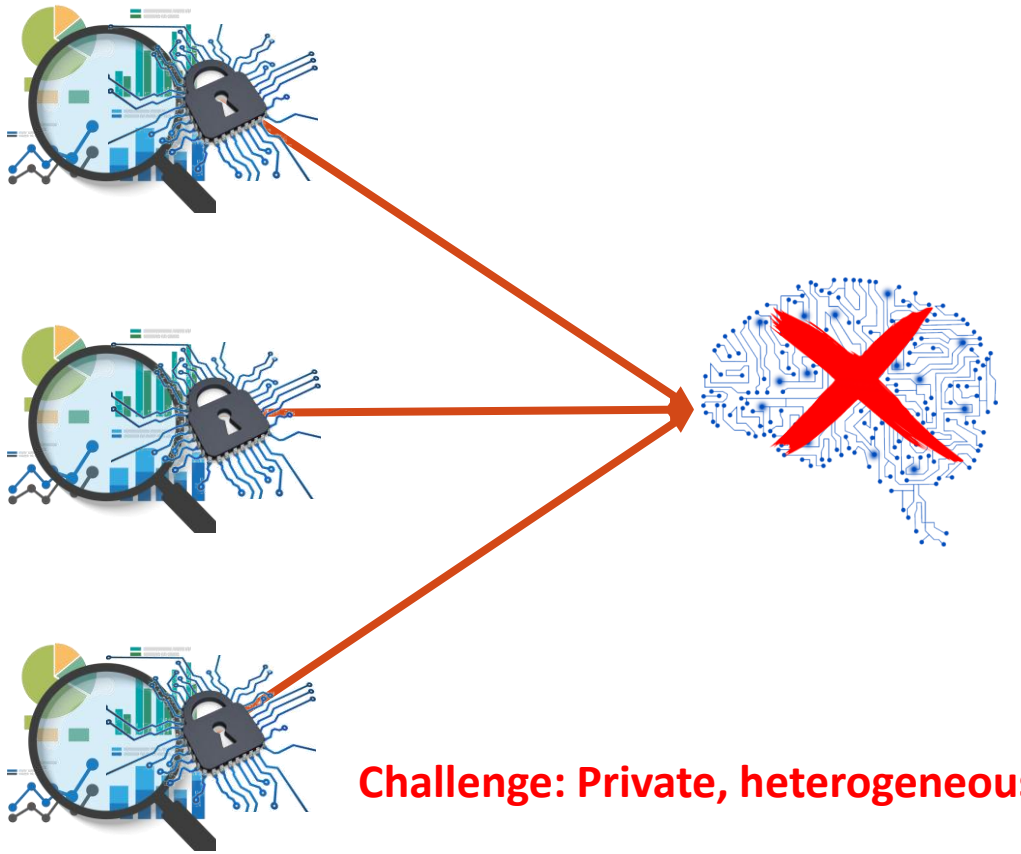

Challenge: Private, heterogeneous data

Related work: White-Box Homogeneous Model Fusion

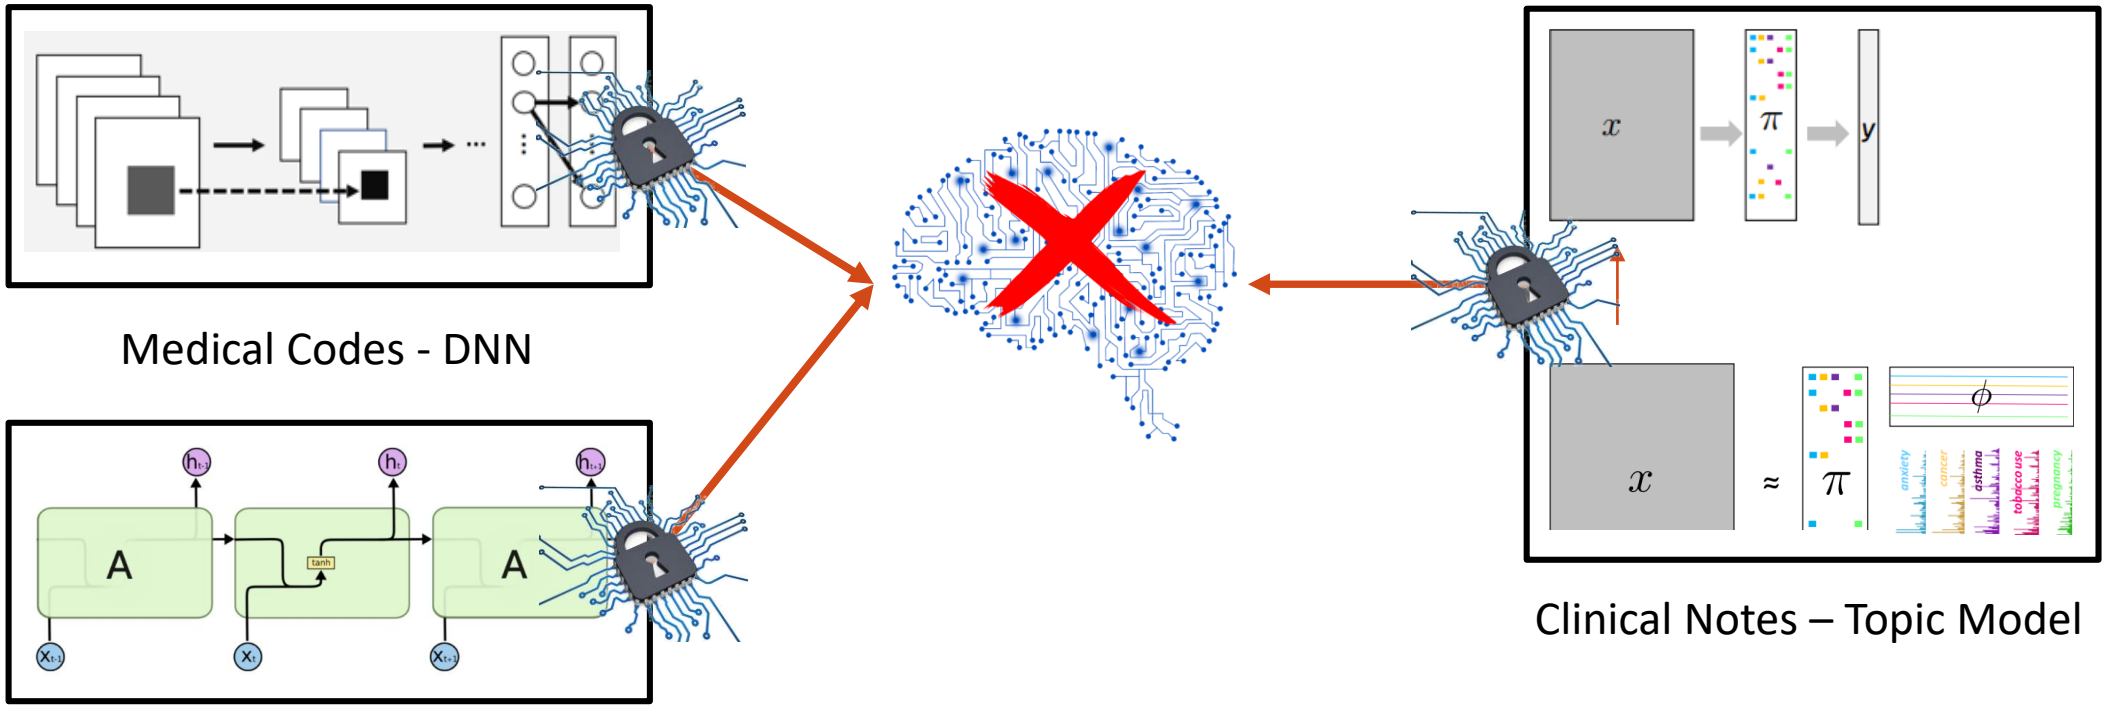

Challenge: Private, heterogeneous model architecture

A real-world setting: Black-Box Model Fusion

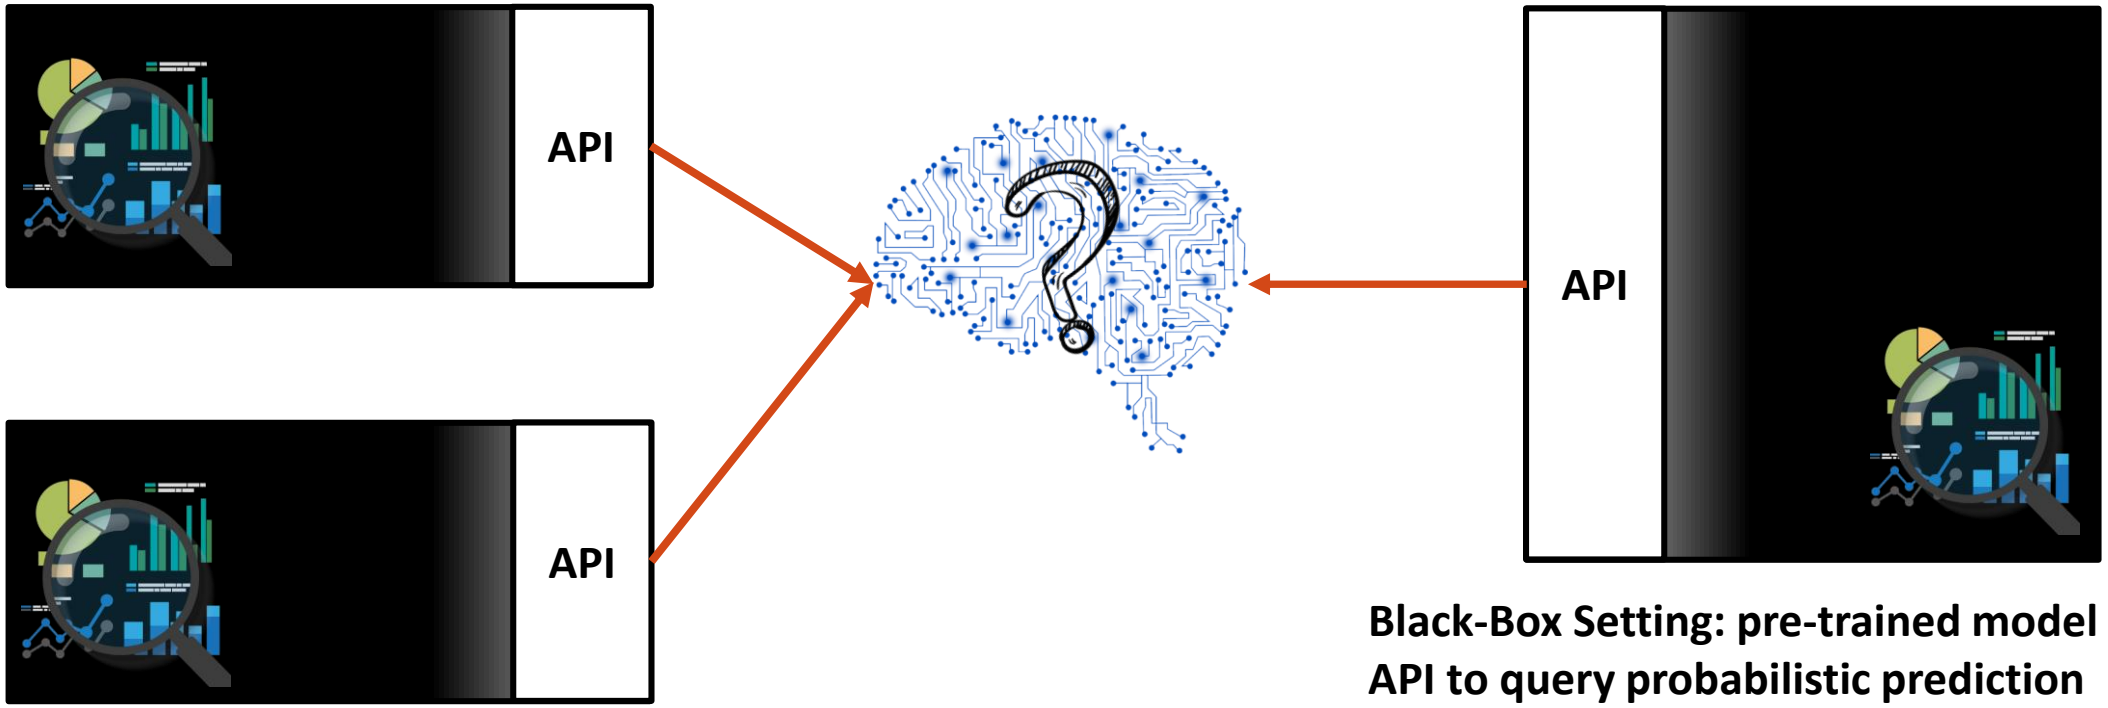

Black-Box Setting: pre-trained model API to query probabilistic prediction

Collective Inference via Gradient Aggregation (CIGAR)

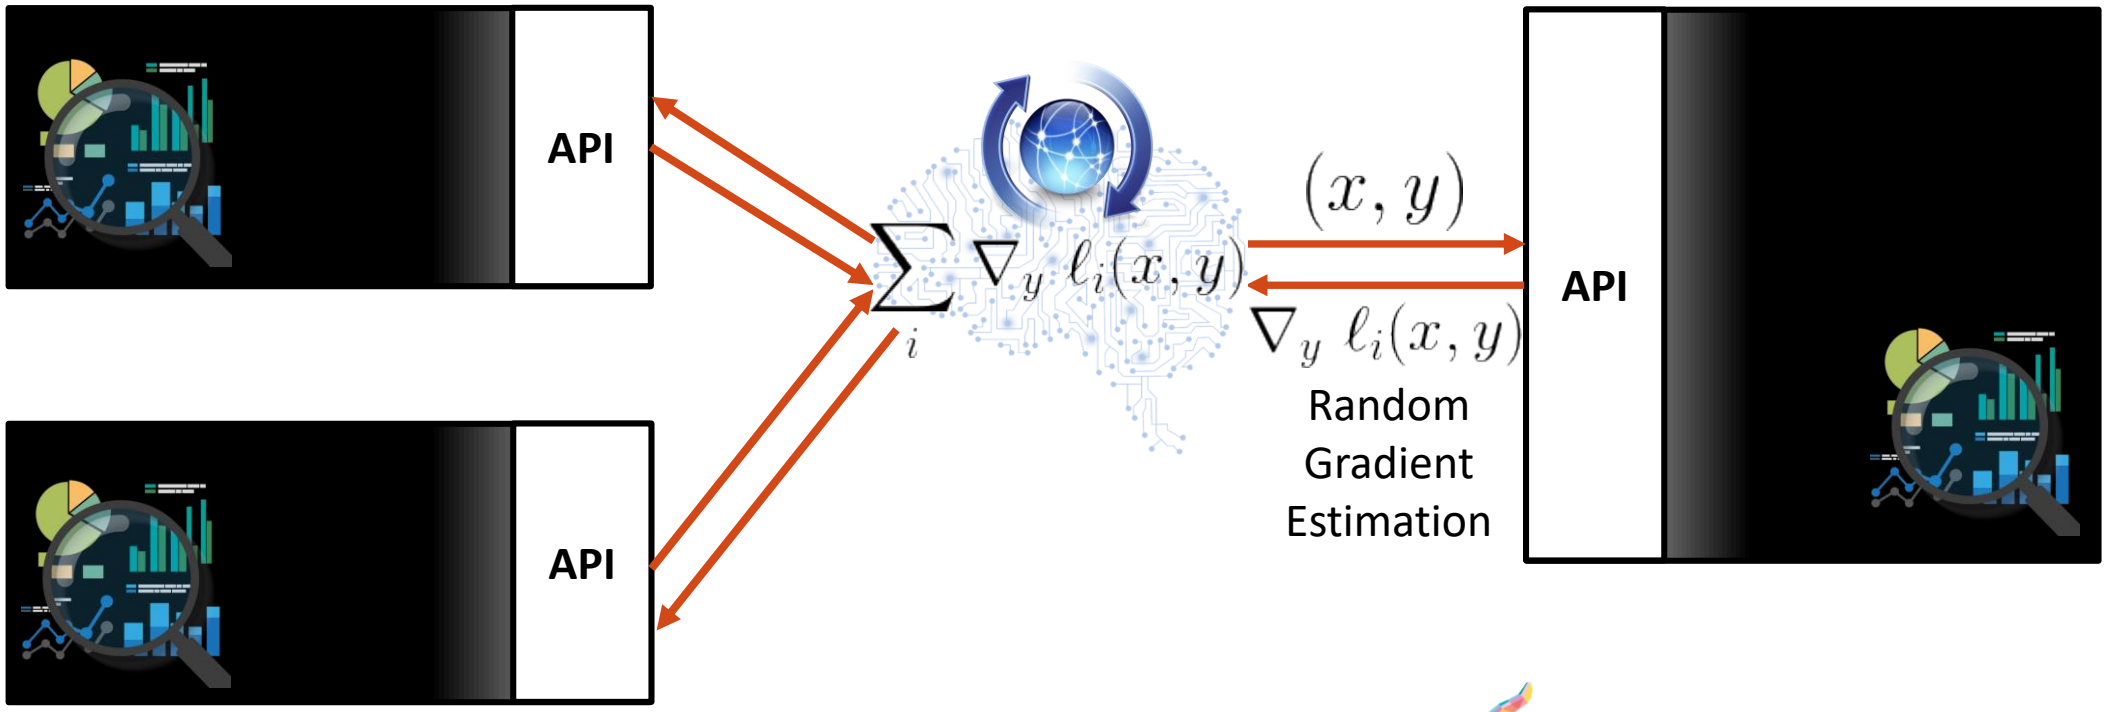

Light-weight Fusion

Collective Learning via Black-Box Imitation (COLBI)

$$\underset{w_i}{\text{minimize}} \; \mathbf{D}_{\mathrm{KL}}(q_i \;||\; p_i(y|x; w_i))$$

Guarantee: Disagreement rate is **upper-bounded by a constant** given sufficient training data

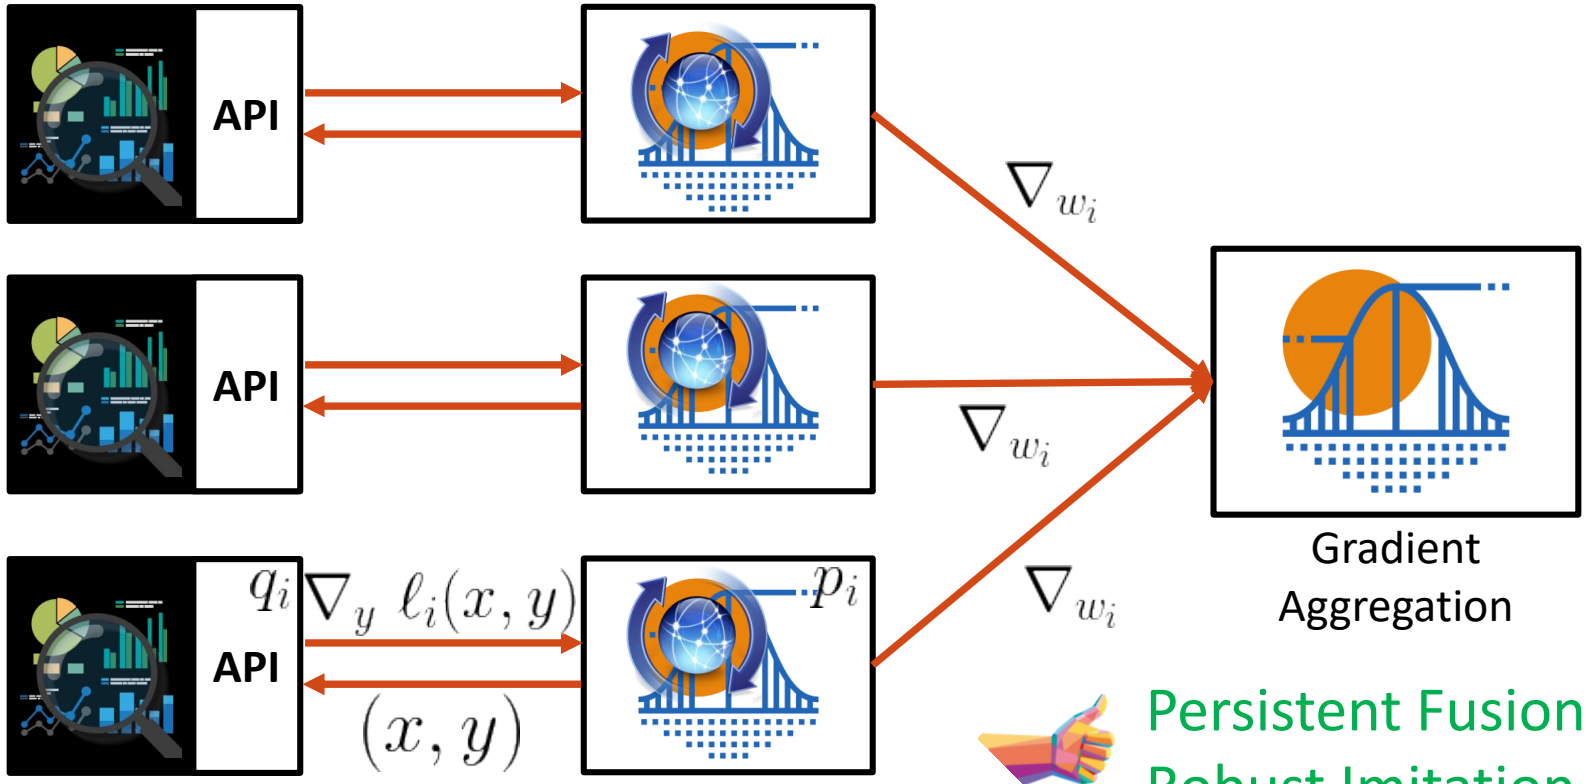

Persistent Fusion
Robust Imitation

CIGAR fusion improves performance

More accurate prediction with more fusion iterations

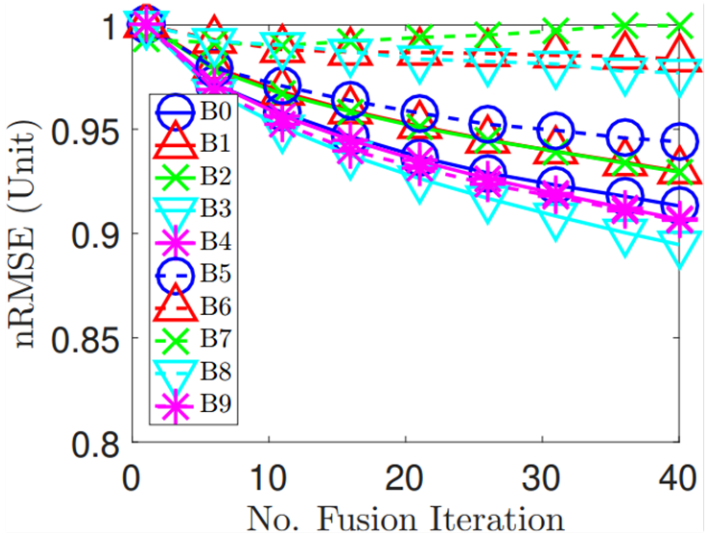

Up to 10% decrease in error for all black-box experts

High prediction variance PRE-FUSION → **Low prediction variance POST-FUSION**

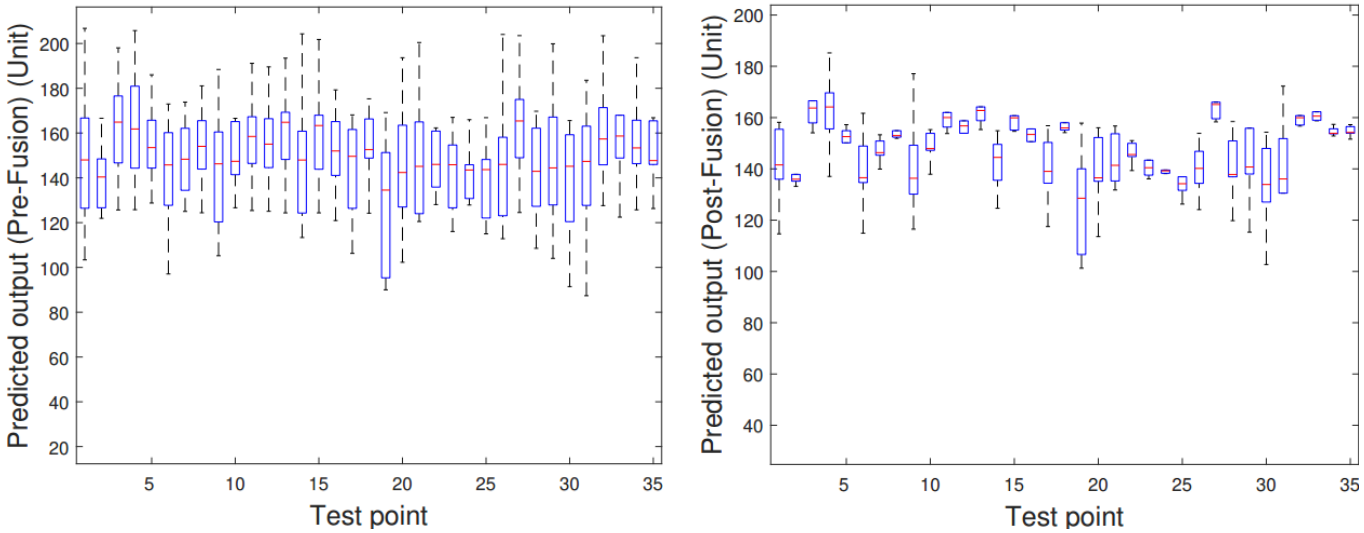

Before: Poor agreement
After: Better consensus

COLBI fusion improves performance

More accurate prediction with more fusion iterations

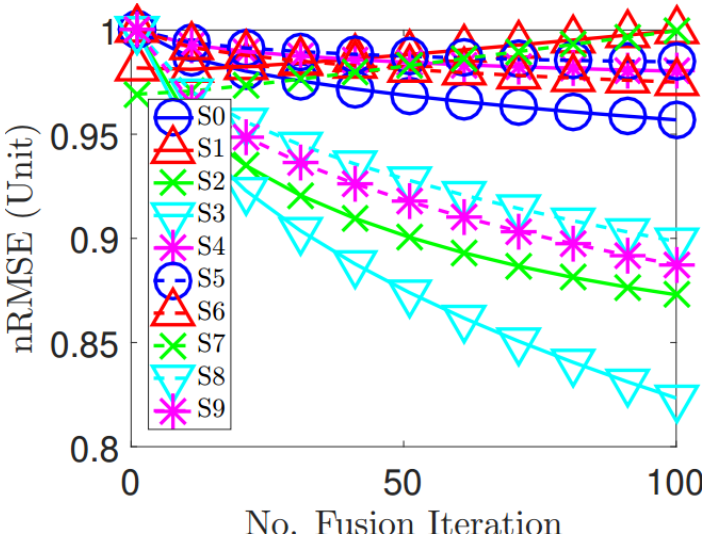

Up to 18% decrease in error for all black-box experts

High prediction variance PRE-FUSION → **Low prediction variance POST-FUSION**

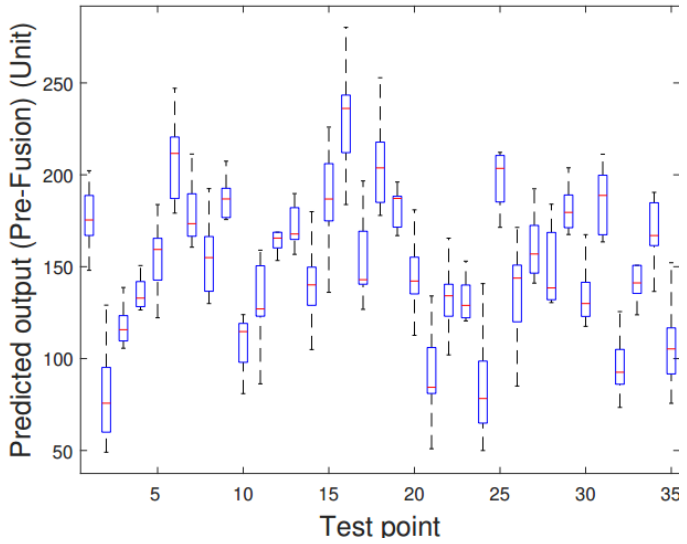

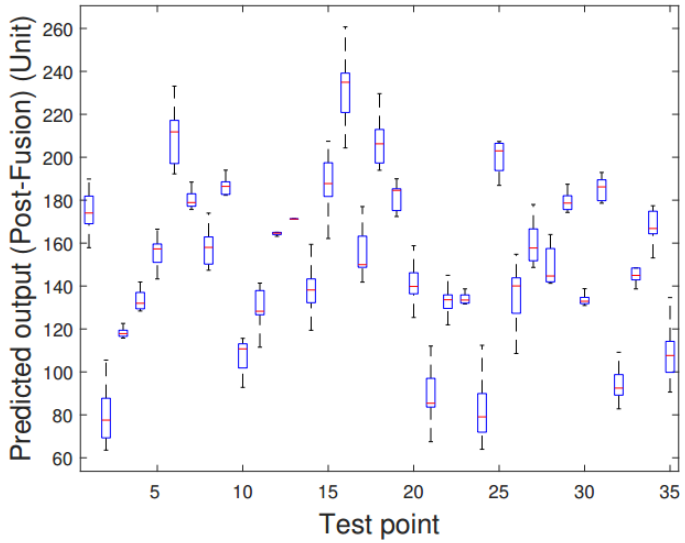

Before: Poor agreement
After: Better consensus

Thank you for listening!

Our poster session:

6:30pm Wednesday, Jun 12, 2019

Pacific Ballroom #184

Paper - Collective Model Fusion for Multiple Black-box Experts