

2021 IWLS Programming Contest

Machine Learning + Logic Synthesis (Part II)

Sat Chatterjee (Google AI), Alan Mishchenko (UC Berkeley), and Gai Liu (Xilinx)

Agenda

1. Overview of the Contest
2. A Look at the Entries
3. Announcement of Winners

The Contest

Recap of Last Year's Contest (ML + LS Part I)

Goal: Learn an unknown boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ in the form of an AIG from a training set consisting of input-output pairs, that is, from a set of minterms

- Size restriction: Had to be under 5000 And nodes
- Generated many interesting ideas and approaches (see [DATE 21 paper](#))

This year, we wanted to extend it to

- Multi-output case, that is, learn a function $f: \{0, 1\}^n \rightarrow \{0, 1\}^m$
- Explore accuracy v/s area Pareto-optimality

The Contest This Year

Goal: Learn an And-Inverter Graph (AIG) which performs well on the CIFAR-10 dataset.

airplane




automobile



bird



cat



deer




dog




frog



horse



ship



truck



CIFAR-10 dataset is a standard image recognition dataset used in machine learning

- 50K training images belonging to 10 classes
- 10K test images
- Each image is 32x32x3 8-bit pixels

Learn an AIG with 24,576 binary inputs and 10 binary outputs (1-hot encoded).

3 size limits (with separate winner for each)

- Small no more than 10K And nodes
- Medium no more than 100K And nodes
- Large no more than 1 million And nodes

Evaluation

Based on Top-1 Accuracy on the CIFAR-10 Test Set

- Participants were prohibited from using the test set for training
- Asked to minimize looking at the test set even for hyper-parameter tuning
- Had to report how many times they looked at the test set

We also used an alternate test set ([CIFAR 10.1 v6](#)) that is less well known

- Based on Recht et al. [Do CIFAR-10 Classifiers Generalize to CIFAR-10?](#)

Evaluation Infrastructure

Command provided in ABC for contestants to evaluate an AIG on a given set of images

```
abc 01> &iwls21test Team-1/small.aig dataset/test_batch.bin
Successfully read 29.31 MB (10000 images) from file "dataset/test_batch.bin".
Summary: Total = 10000. Errors = 5894. Correct = 4106. ( 41.06 %) Naive guess = 1000. ( 10.00 %)
Total checking time = 0.15 sec
abc 01> █
```

```
abc 01> abc 01> &iwls21test Team-1/large.aig dataset/test_batch.bin
Successfully read 29.31 MB (10000 images) from file "dataset/test_batch.bin".
Summary: Total = 10000. Errors = 4321. Correct = 5679. ( 56.79 %) Naive guess = 1000. ( 10.00 %)
Total checking time = 4.15 sec
abc 01> █
```

The Entries

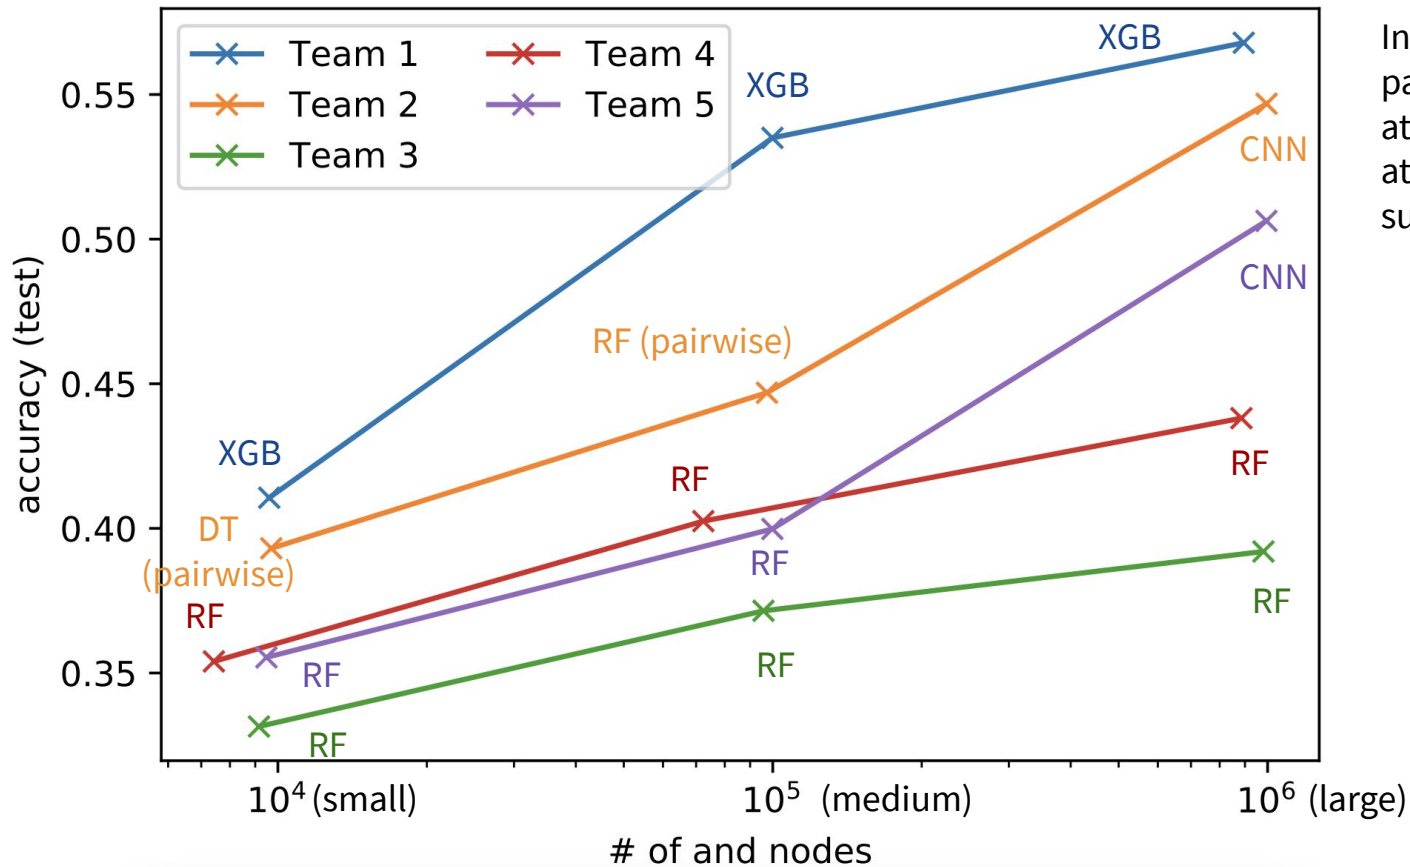
We got 5 entries (in random order below)

- UC Berkeley + U Tokyo + IIT KGP + IIT Bombay
- Universidade Federal de Pelotas (UFPel) + Universidade Federal do Rio Grande do Sul (UFRGS)
- Federal University of Rio Grande do Sul (UFRGS) and Federal University of Santa Catarina (UFSC)
- University of Wisconsin–Madison
- National Taiwan University

Techniques Used


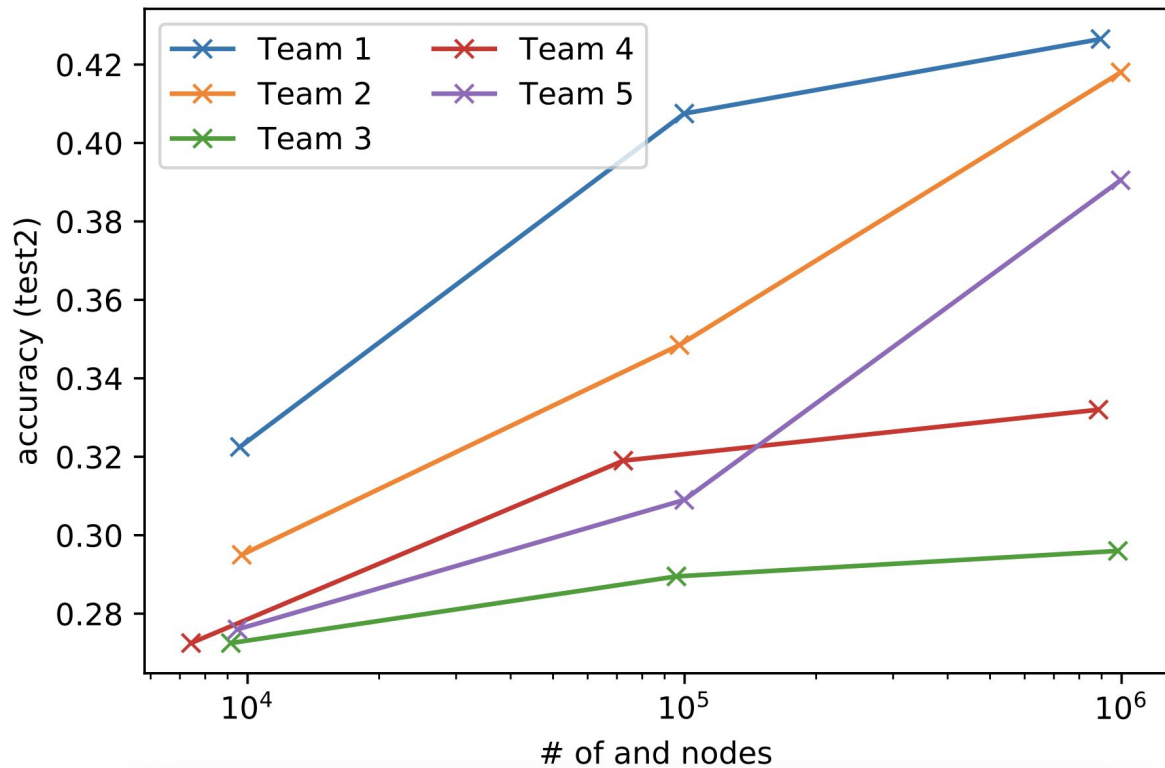
- Model Families:
 - Decision Trees, Random Forests, XGBoost, and Convolutional Neural Nets
 - Quantization, PO2 weights
- Problem Setup
 - Reduction to binary classification: one v/s rest, pairwise comparisons
 - Native multi-class
- Training Methods
 - Data augmentation
- Logic Optimization
 - ABC scripts to minimize area
 - Cartesian Genetic Programming to reduce AIG size
 - Rare pattern elimination

Results: CIFAR-10 Test Set



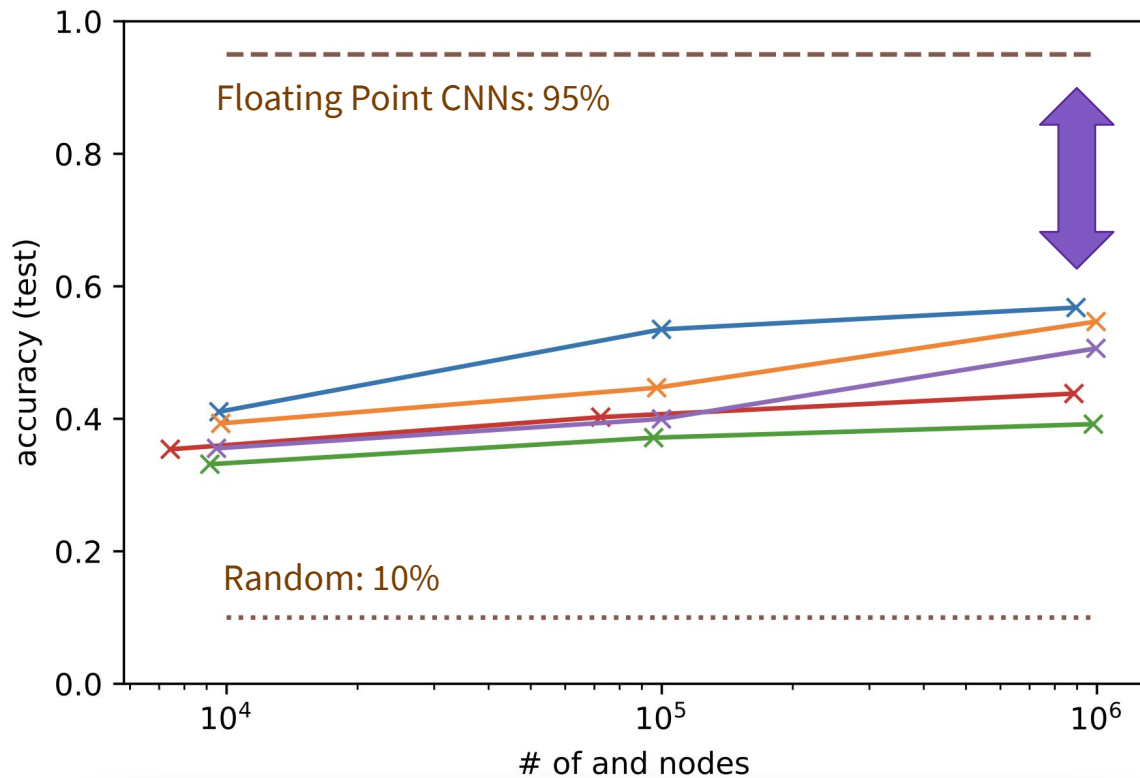
In almost all cases the participants did not look at the test set or looked at it once before final submission.

Results: CIFAR-10 Alternate Test Set (10.1 v6)



Same relative performance as on original test set though lower absolute numbers (in line with what Recht et al. found in their study).



Where next?



Can we
close this
gap?

And Finally ... Revealing The Winners

3rd: Team 5 (UC Berkeley + U Tokyo + IIT KGP + IIT B)


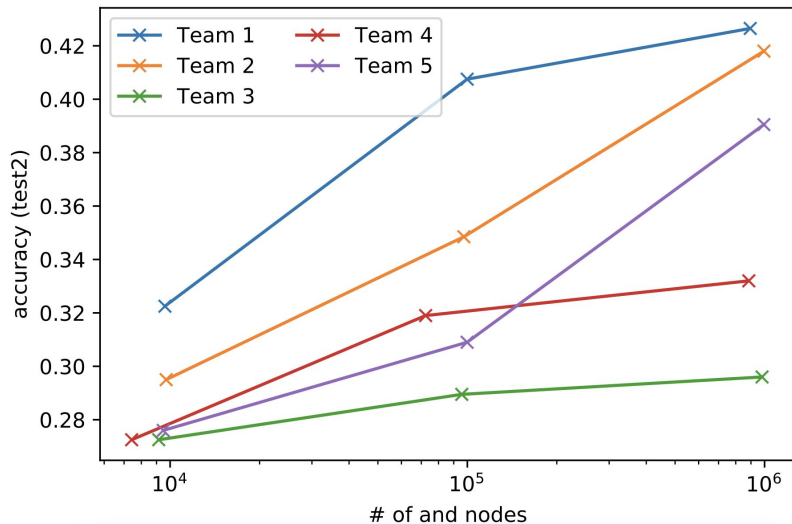


Yukio Miyasaka, Sai Sanjeet, Xinpei Zhang, Mingfei Yu, Qingyang Yi, Ryogo Koike, Takemaru Kadoi, Prof. Masahiro Fujita, Prof. Bidhu Datta Sahoo, Prof. Virendra Singh, and Prof. John Wawrzynek

[Presentation Video](#)

(No peeks for small and medium, 16 peeks for large)

2nd: Team 2 (National Taiwan University)


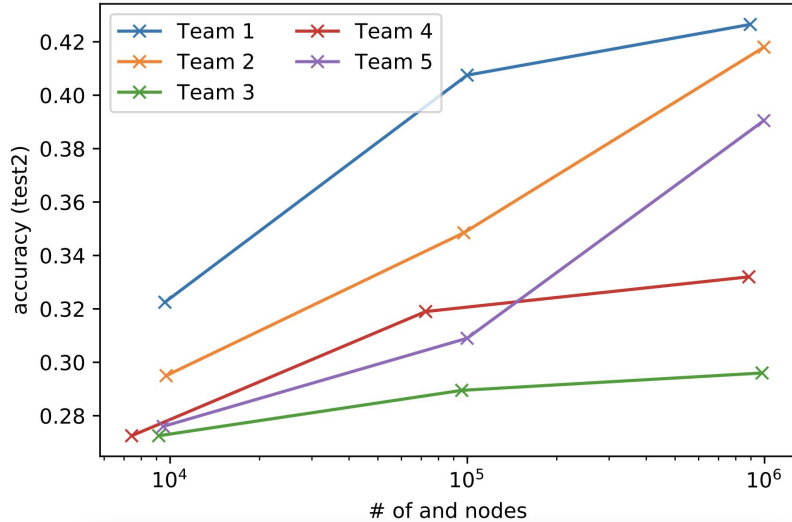


Po-Chun Chien, Yu-Shan Huang, Hao-Ren Wang, Nai-Ning Ji, and Prof. Jie-Hong Roland Jiang

[Presentation Video](#)

(One peek just before final submission)

1st: Team 1 (University of Wisconsin–Madison)



Wei Zeng, Tianen Chen, Younghyun Kim, and Azadeh Davoodi

[Presentation Video](#)

(No peeks)