HOLMES:
Efficient Distribution Testing for Secure Collaborative Learning

Ian Chang UC Berkeley
Katerina Sotiraki Yale University
Weikeng Chen DZK Labs
Murat Kantarcioglu UT Dallas
Raluca Ada Popa UC Berkeley

USENIX Security ‘23

Secure Collaborative Learning

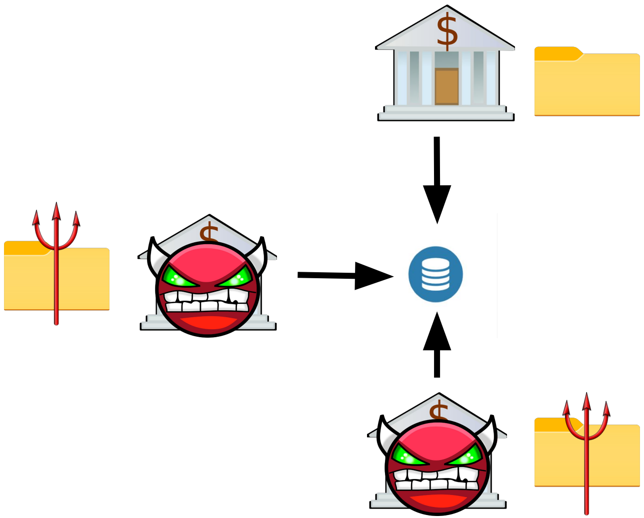

Multiple datasets lead to better accuracy

Privacy

- secure computation [GMW87, Y82]

Security

- malicious security [CLOS02, DPSZ12, WRK13]

Corrupted datasets can ruin model, e.g.[PY17, WRJI19, RSARRJ20]

- Privacy technique blinds parties' corrupted dataset

Attempt 1: range checks

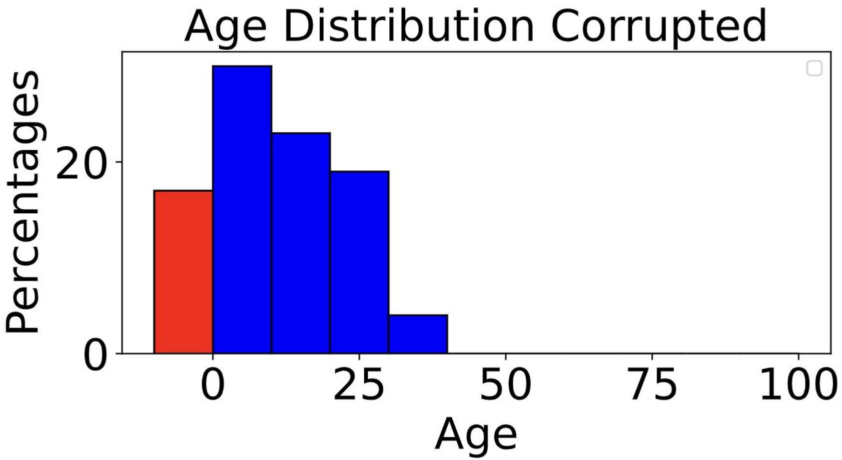

Corrupted Input:

- negative age (age < 0)
- too old age (age > 120)

Range checks

e.g. [BBBPWM18,CB17,AGJOP21]

- Enforce a range of values that each input can take
- Previously the only technique against malicious inputs

Are range checks enough?

- Introduce distribution testing (check properties of distribution)

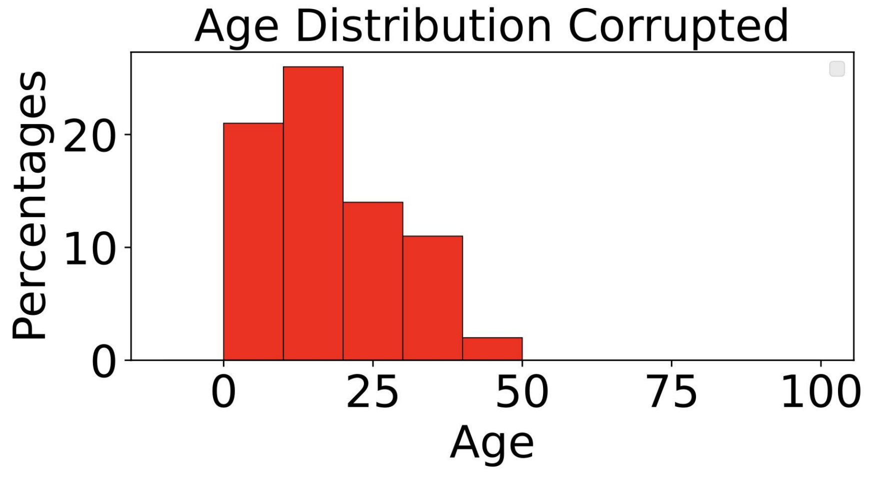

ages > 0 and ages < 120

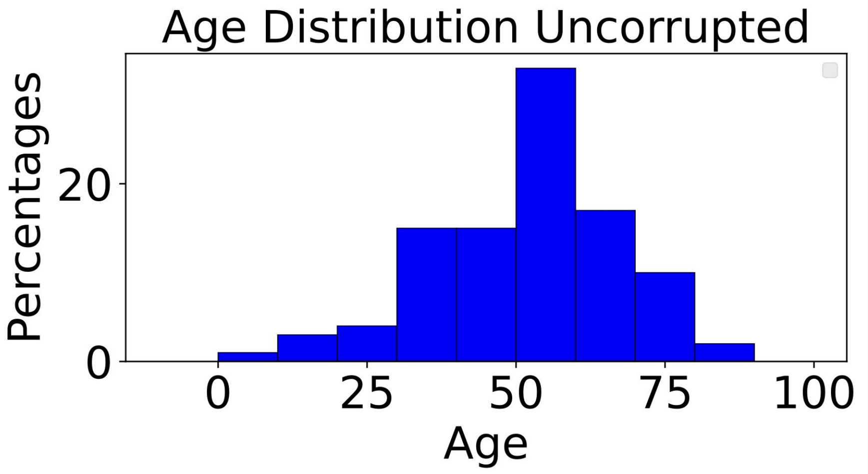

ages > 0 and ages < 120 **and μ ≈ 50**

- Distribution testing + range checks >>> range checks!

Our work: **HOLMES**

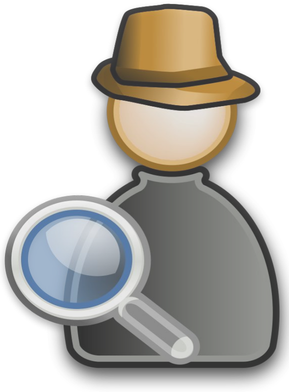

- Checks malicious input using **distribution testing**

- Operates in highest level of security
  - Malicious security (e.g. n - 1 out of n parties)

- Perform distribution testing efficiently
  - 10-10000x faster than baselines

Why distribution testing?

- Pragmatic Clinical Trials
  - Compare distributions of datasets to detect discrepancies
- Group fairness
  - Biased data => biased trained model
- Data quality
  - Model of joint dataset > models of individual datasets

Beyond Distribution Testing

- Distribution testing **cannot** detect input poisoning attacks
    - Input poisoning: small pertubations to inputs

- Input poisoning attacks are ineffective in certain cases
    - e.g., federated learning [SHKR21]

Roadmap

- Use zero-knowledge (ZK) for fast distribution testing
    - Offload and verify computation of local dataset using ZK
    - Refer to the paper for more details

- Design efficient multidimensional tests
    - 10000x times faster than strawman!

- Perform experimental evaluation
    - HOLMES distribution testing vs. Naive

Histogram goodness-of-fit

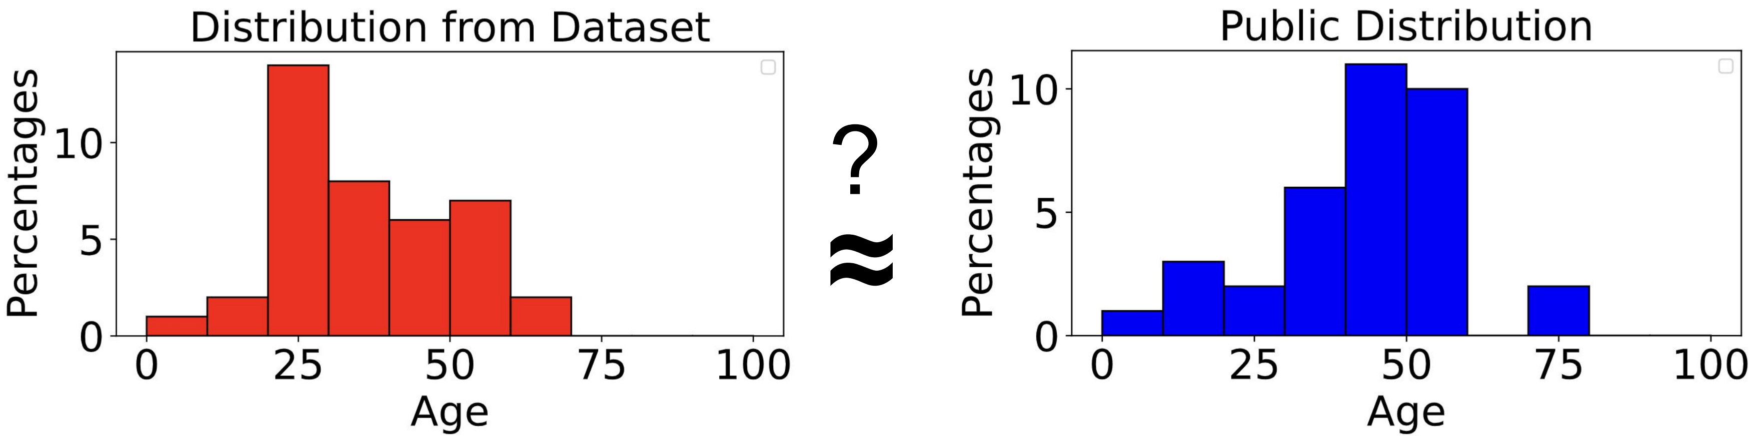

Classical histogram checks use Pearson's χ^2-test

Intuitively, check if $\Sigma_i(\text{count}_{\text{dataset}}[i] - \text{count}_{\text{public}}[i])^2$ is small

What happens in multidimensional data?

Multidimensional goodness-of-fit

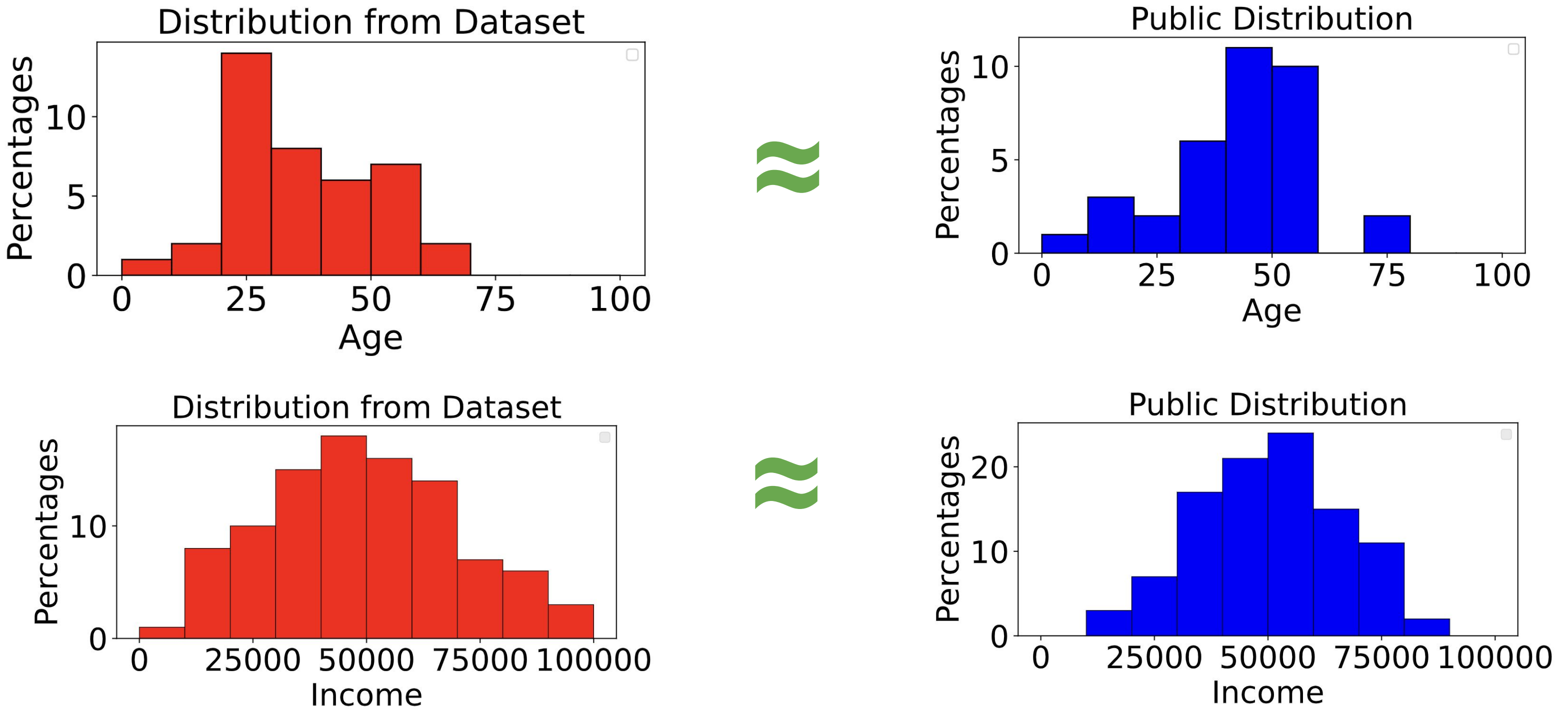

Perform histogram check for each attribute: age & income

Multidimensional goodness-of-fit

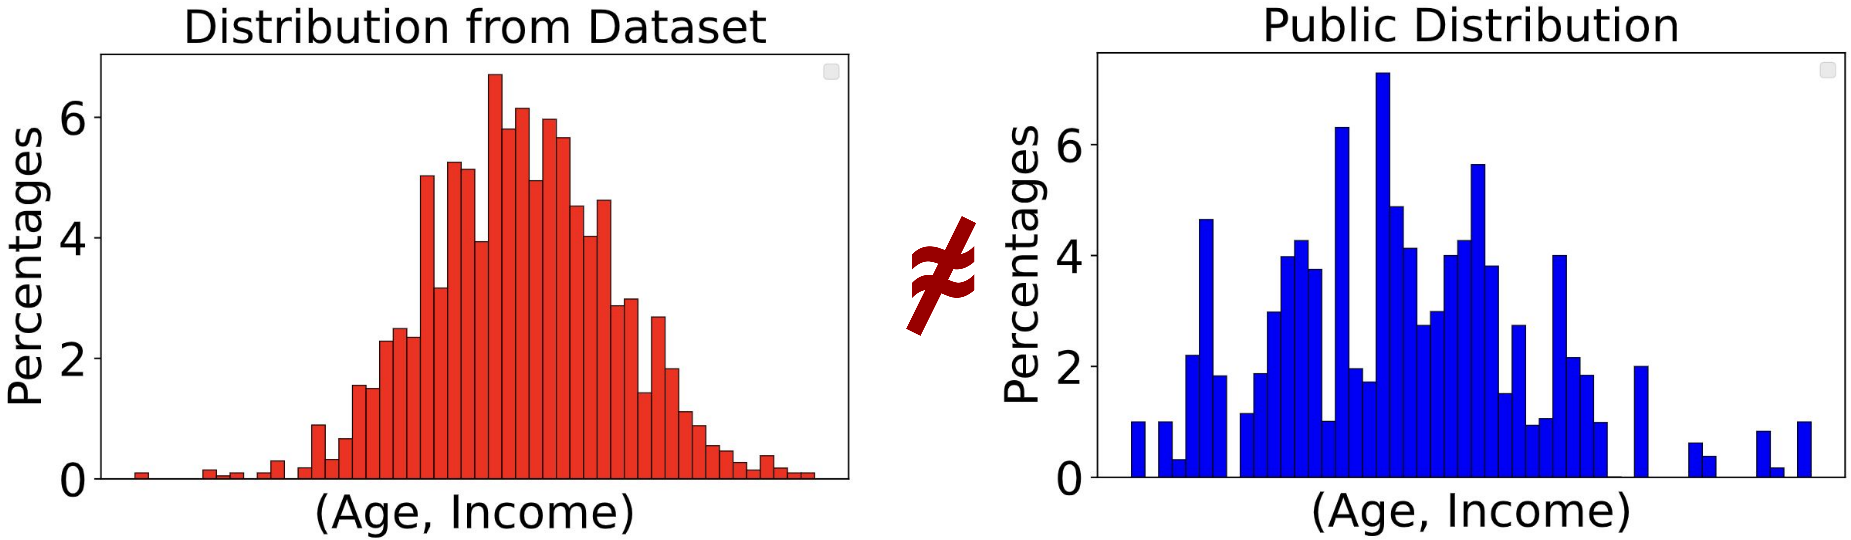

Checking histograms for individual attributes does not suffice

Number of histogram bins grows exponentially

Pearson's χ^2 test is prohibitively expensive

Our solution: efficient sketching

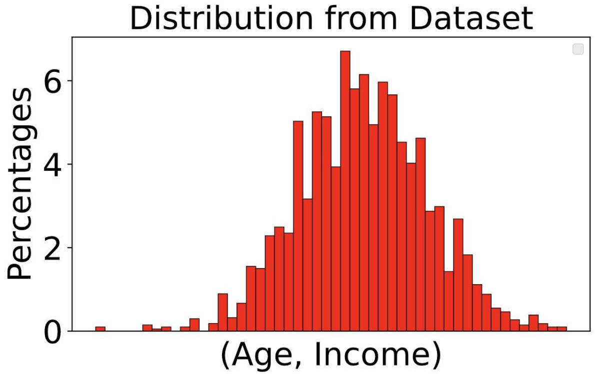

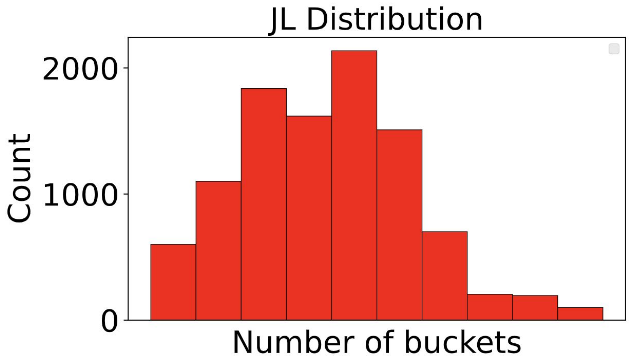

Johnson-Lindenstrauss Lemma [JL84,A03]:

For suitable random matrix A, $||\mathbf{x}||_2 \approx ||A\mathbf{x}||_2$

Only works when comparing to a public distribution

Experimental Evaluation

Setup:

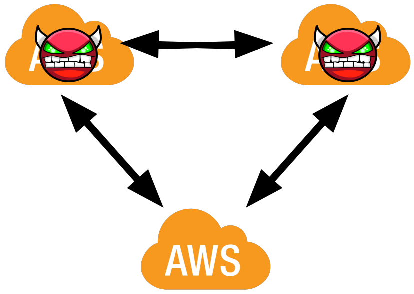

- QuickSilver for ZK, SCALE-MAMBA for MPC
- AWS c5.9xlarge instances, each containing 36 cores
  - Each instance is a different party
- Vary: 2 to 10 parties, input dataset size, real-world datasets

Highlights:

- 10 times speedup for classical distribution tests
- 10000 times speedup for multidimensional distribution tests

Single dimension histogram check w/ varying input size

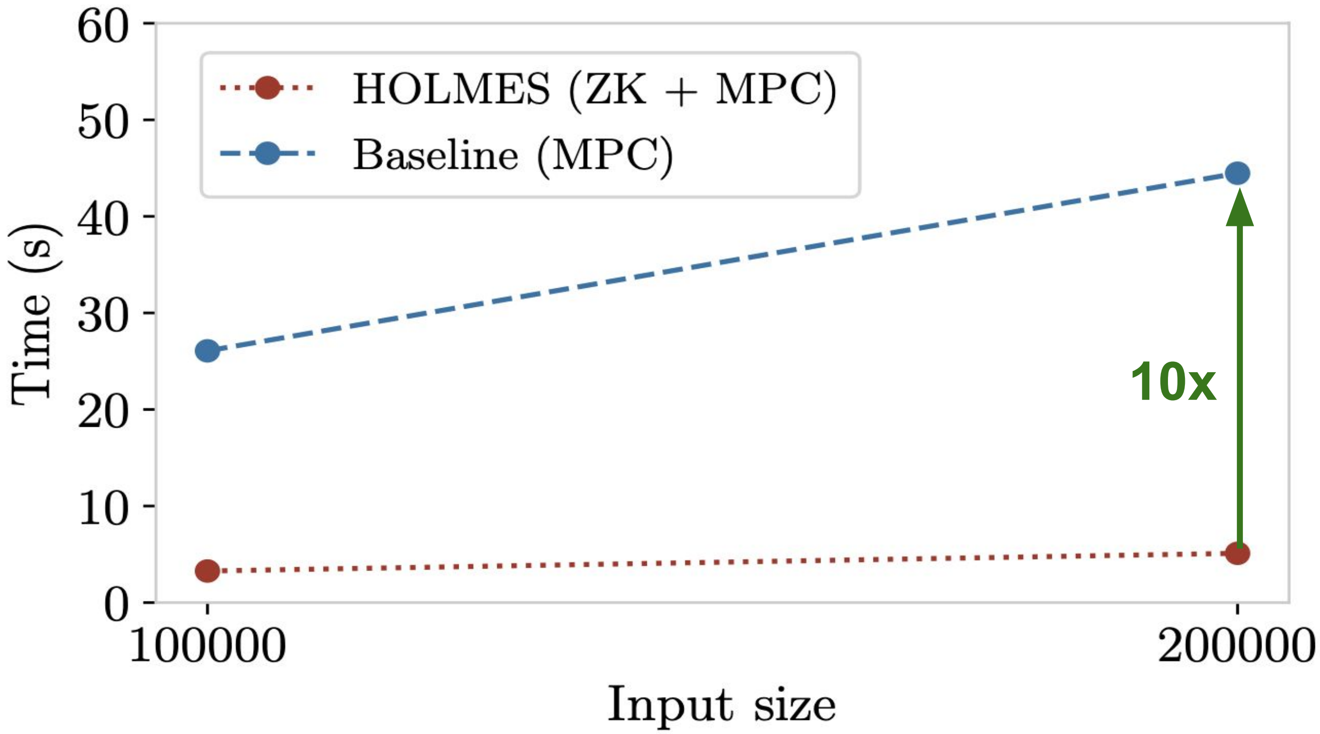

10x speedup with ZK at an input size of 200k entries

Histogram check w/ varying number of attributes

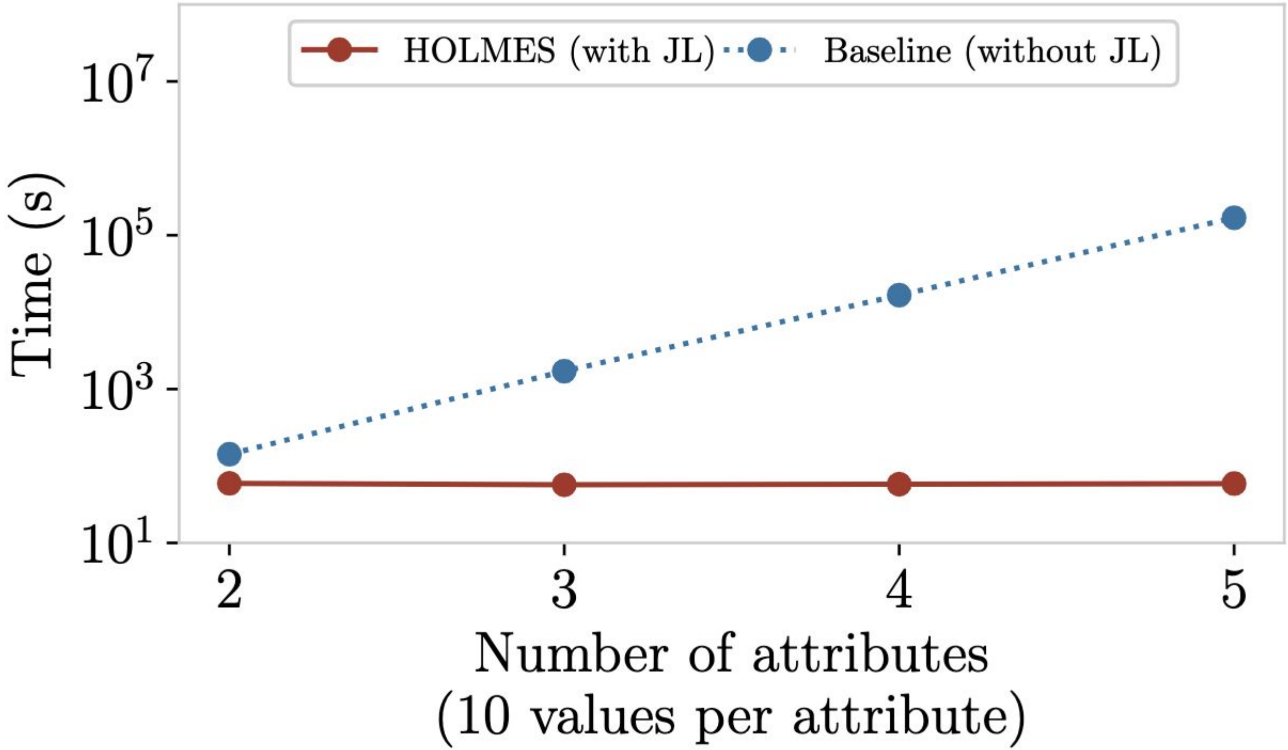

10000x speedup with JL at five attributes per input entry

Conclusion

- We present HOLMES, an efficient framework for distribution testing
- HOLMES is a lot more efficient than the baseline generic MPC
  - Combines MPC + ZK (10x speedup)
  - Sketching for multidimensional distribution tests (10000x speedup)
- E-print: https://eprint.iacr.org/2021/1517

Questions?