暨南大學 JINAN UNIVERSITY

Fongwell 芳禾数据

Fusion: Efficient and Secure Inference Resilient to Malicious Servers

Caiqin Dong, Jian Weng, Jia-Nan Liu, Yue Zhang,

Yao Tong, Anjia Yang, Yudan Cheng, and Shun Hu

Outline

- Background
- Design Goals
- Our Solution: *Fusion*
- Performance

Real-World Applications

Machine Learning as a Service

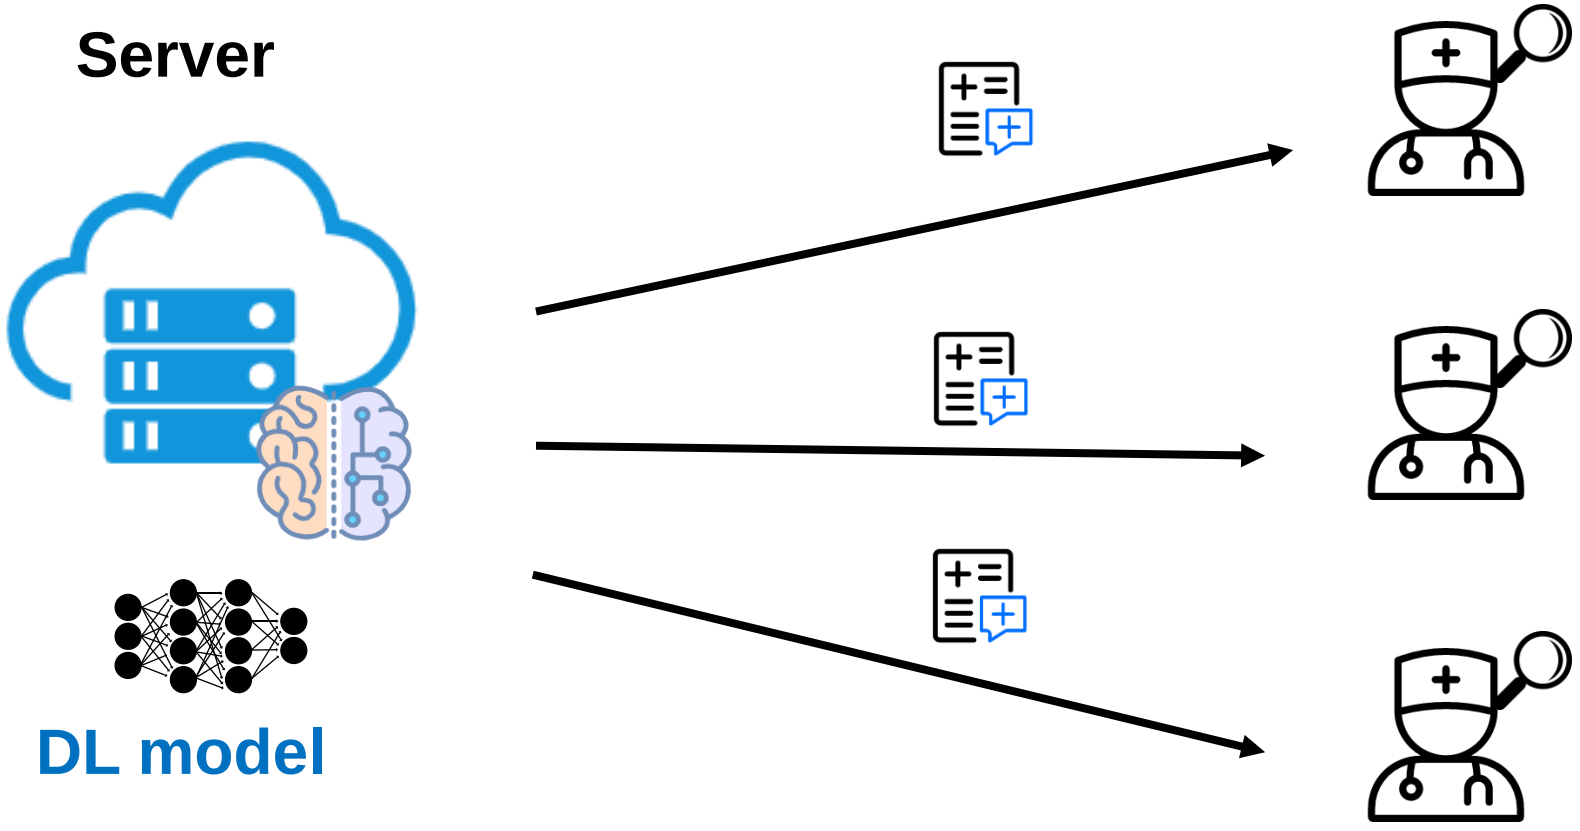

Motivation and Design Goals

Security Requirements

- Privacy Preservation
- Model Accuracy
- Computation Correctness

Malicious

Client

Server

Accurate inference results

Design Challenges

Possible Solutions

Model Accuracy

Computation Correctness

Privacy

Zero-knowledge proof

Maliciously secure 2PC framework

Need complex and careful design

Our Key Insight

An important observation

Client can know some computation results in advance

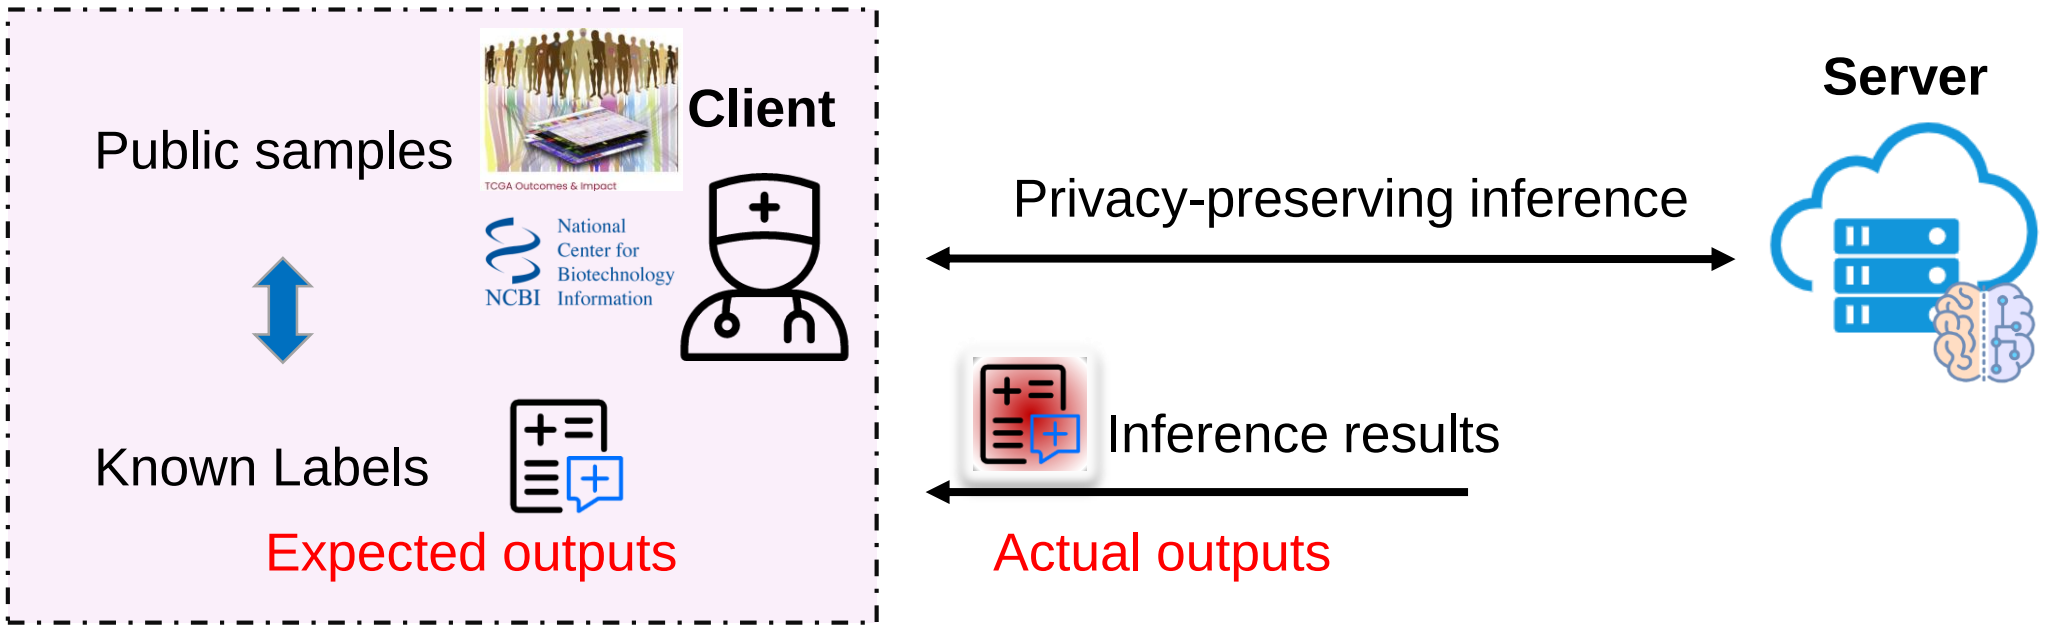

Our Key Insight

Mix-and-Check

Client

(1) Prepare public samples

(2) Duplicate each query sample

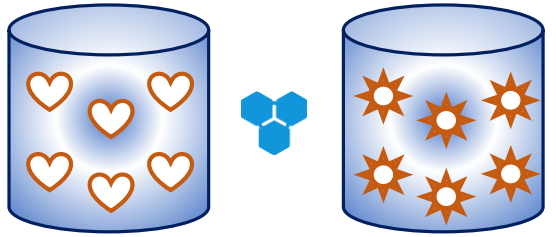

Model Accuracy

Computation Correctness

Client Detects Server's Malicious Behaviors

Server

- Low-Quality model
- Incorrect computations for some samples

Inconsistent results

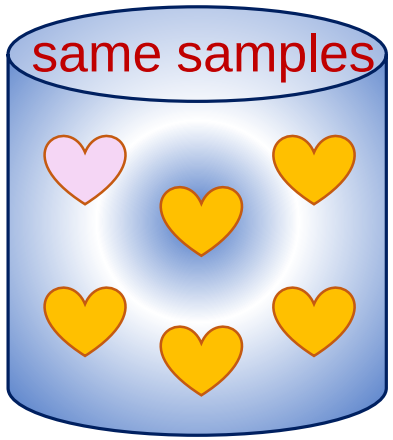

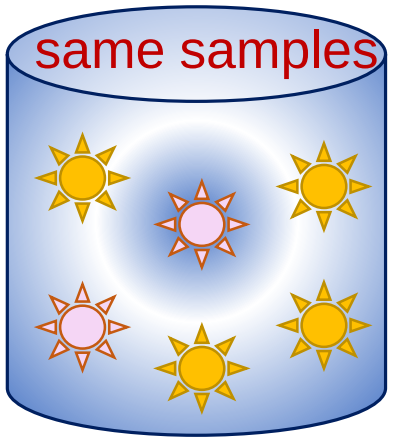

Inaccurate results

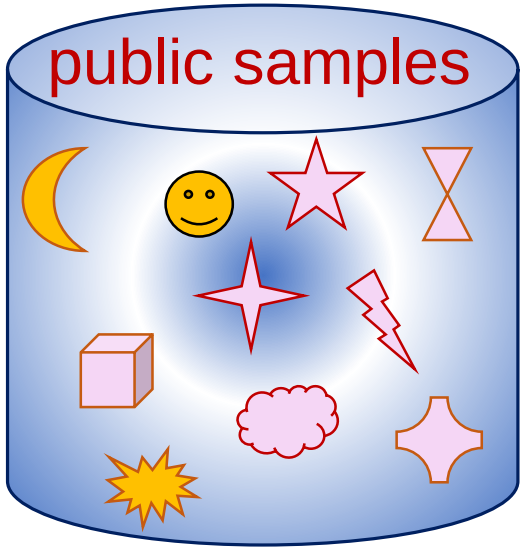

Solution: *Fusion*

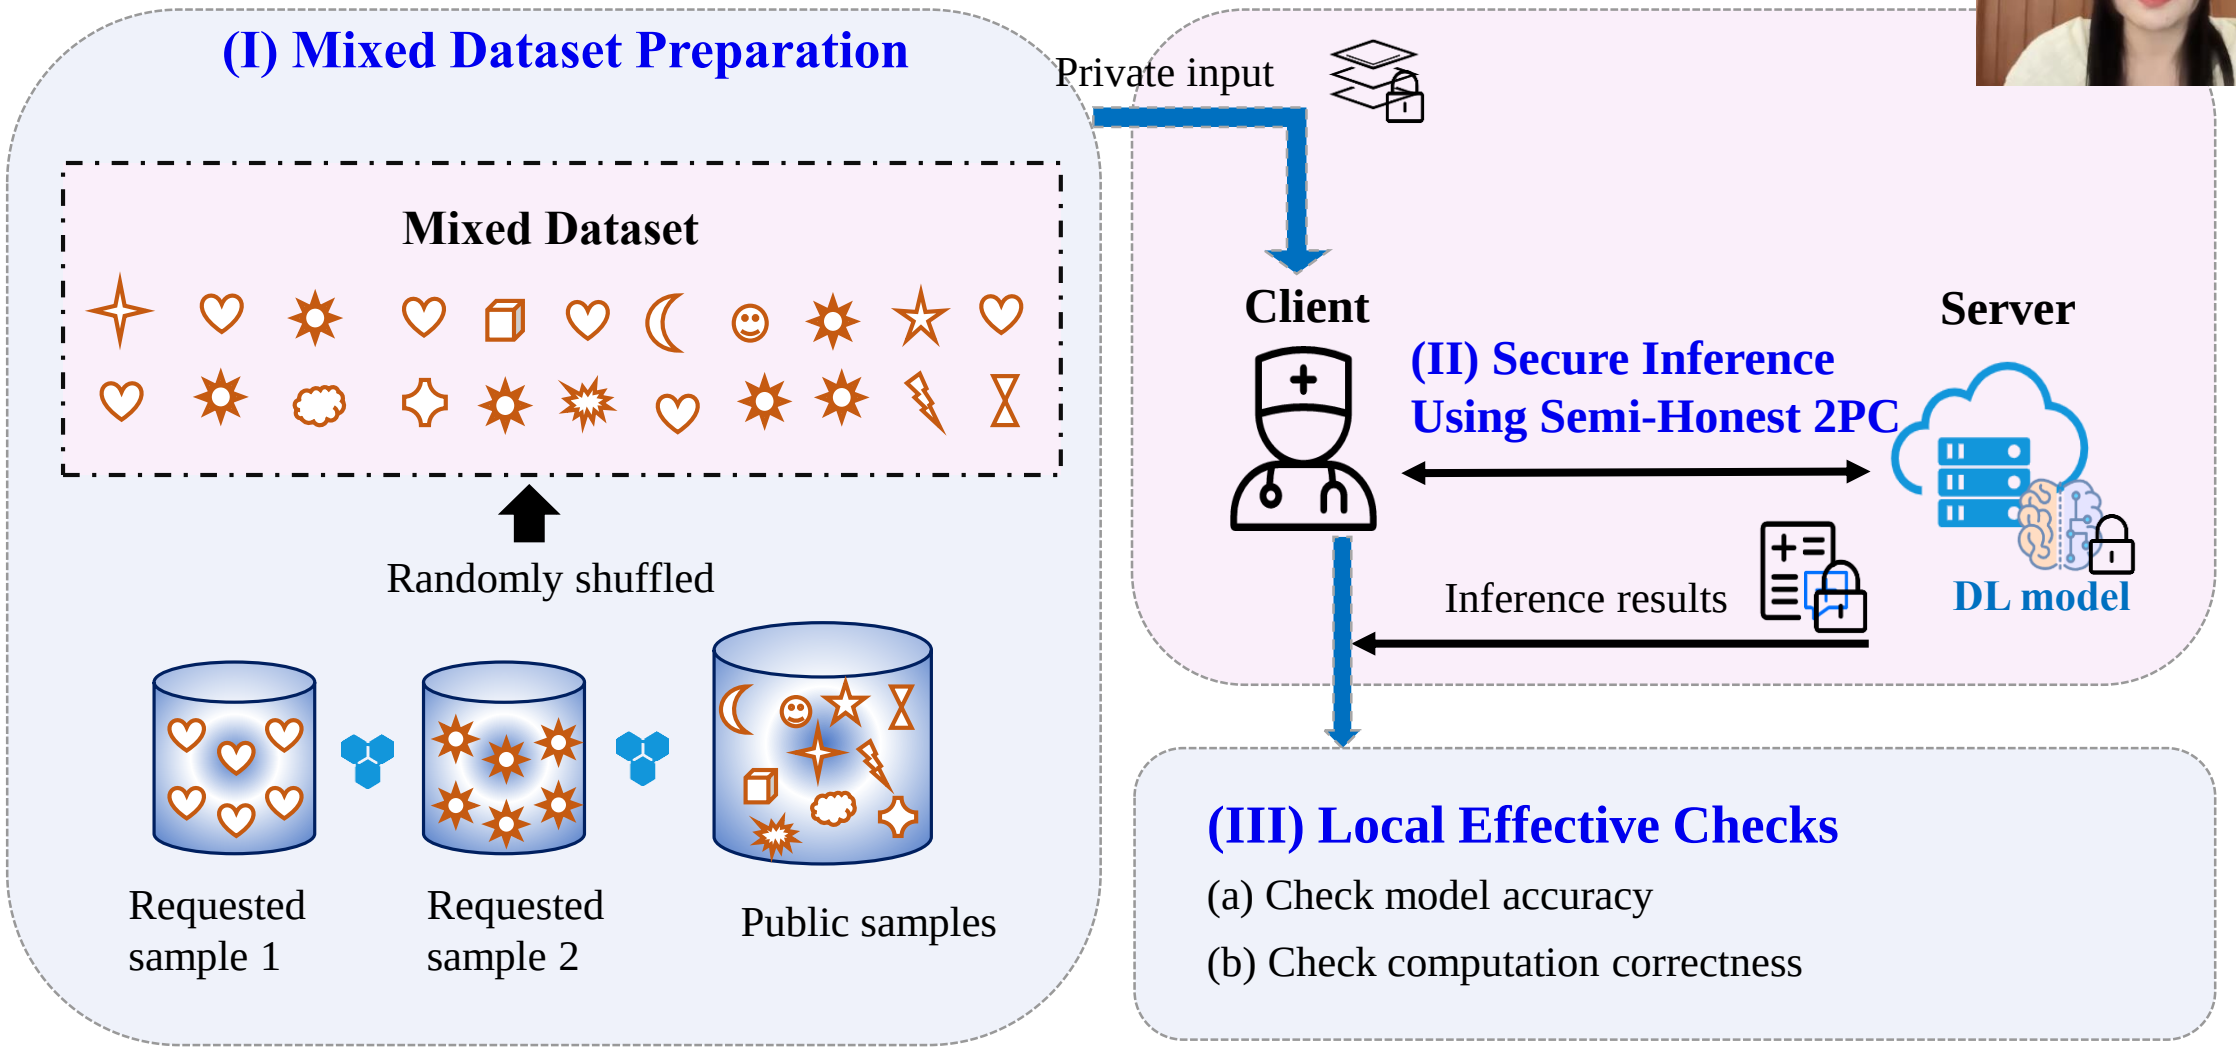

Solution: *Fusion*

Client

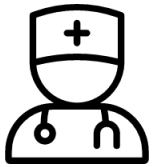

R query samples

T public samples

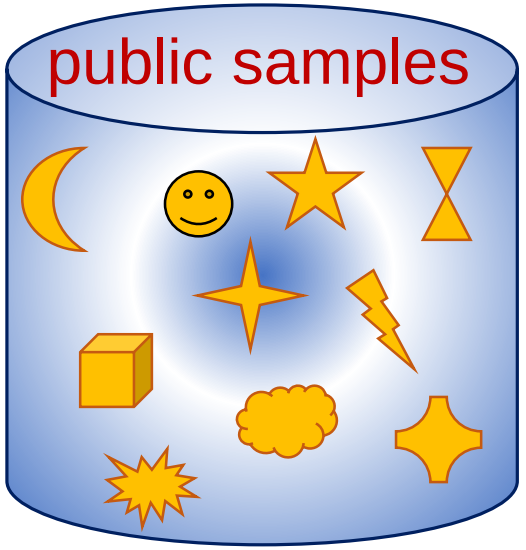

Server

A trained model

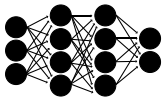

Client Prepares a Mixed Dataset

- (1) Prepare Mixed Dataset

Client

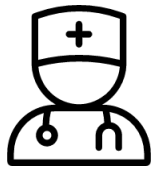

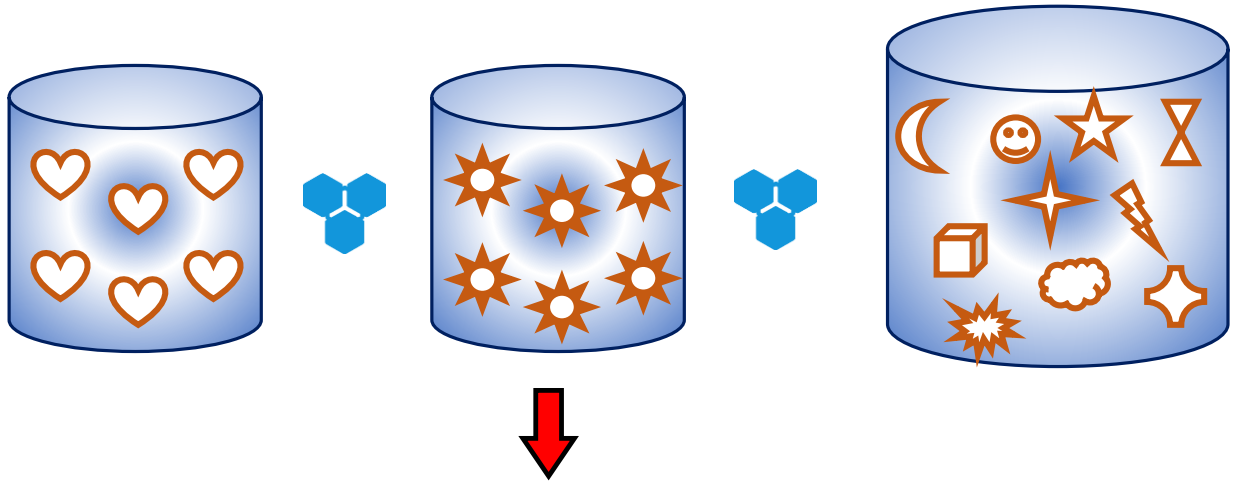

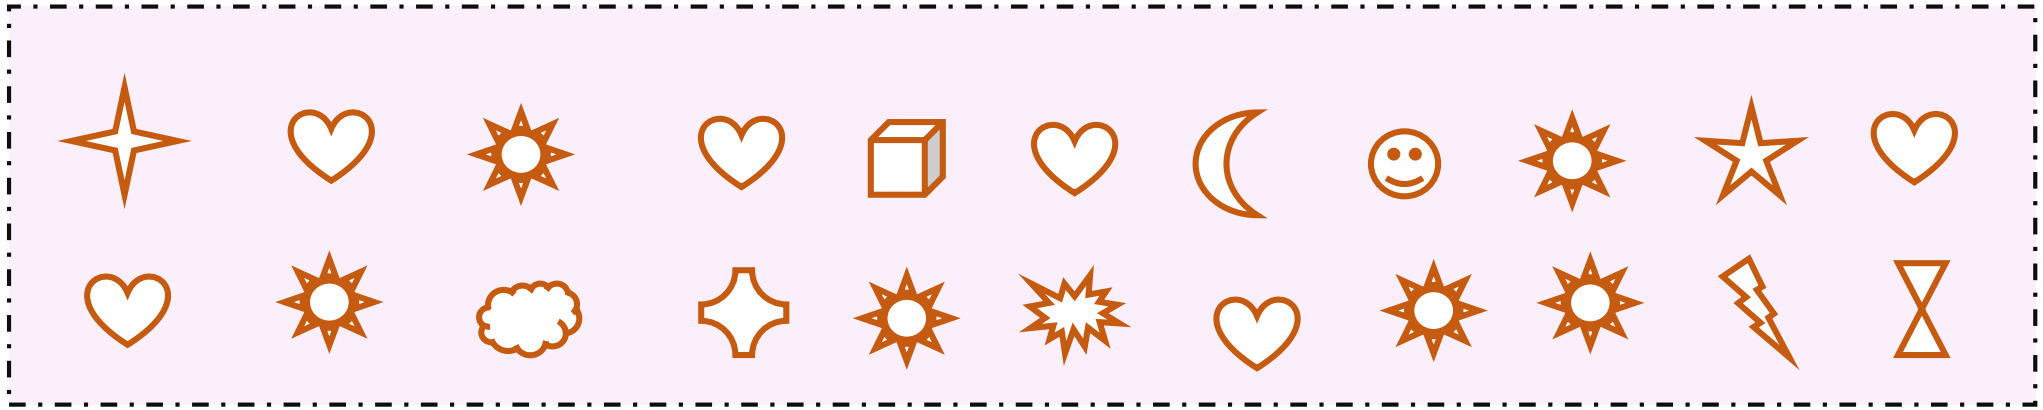

Privacy-Preserving Inference Execution

- (2) Obtain Inference Results

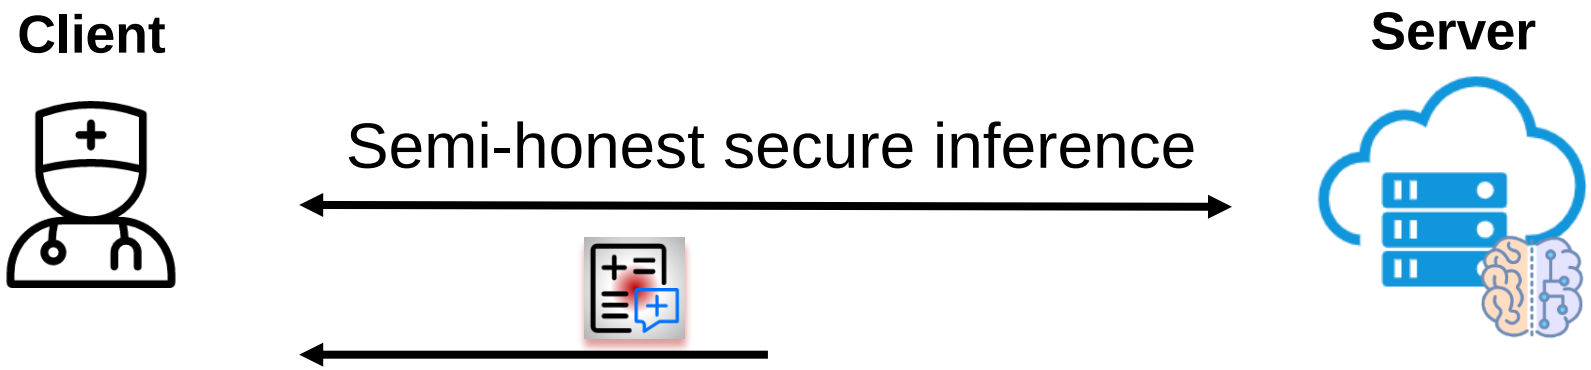

Client Checks Inference Results

- (3) Simple-but-Effective Local Checks

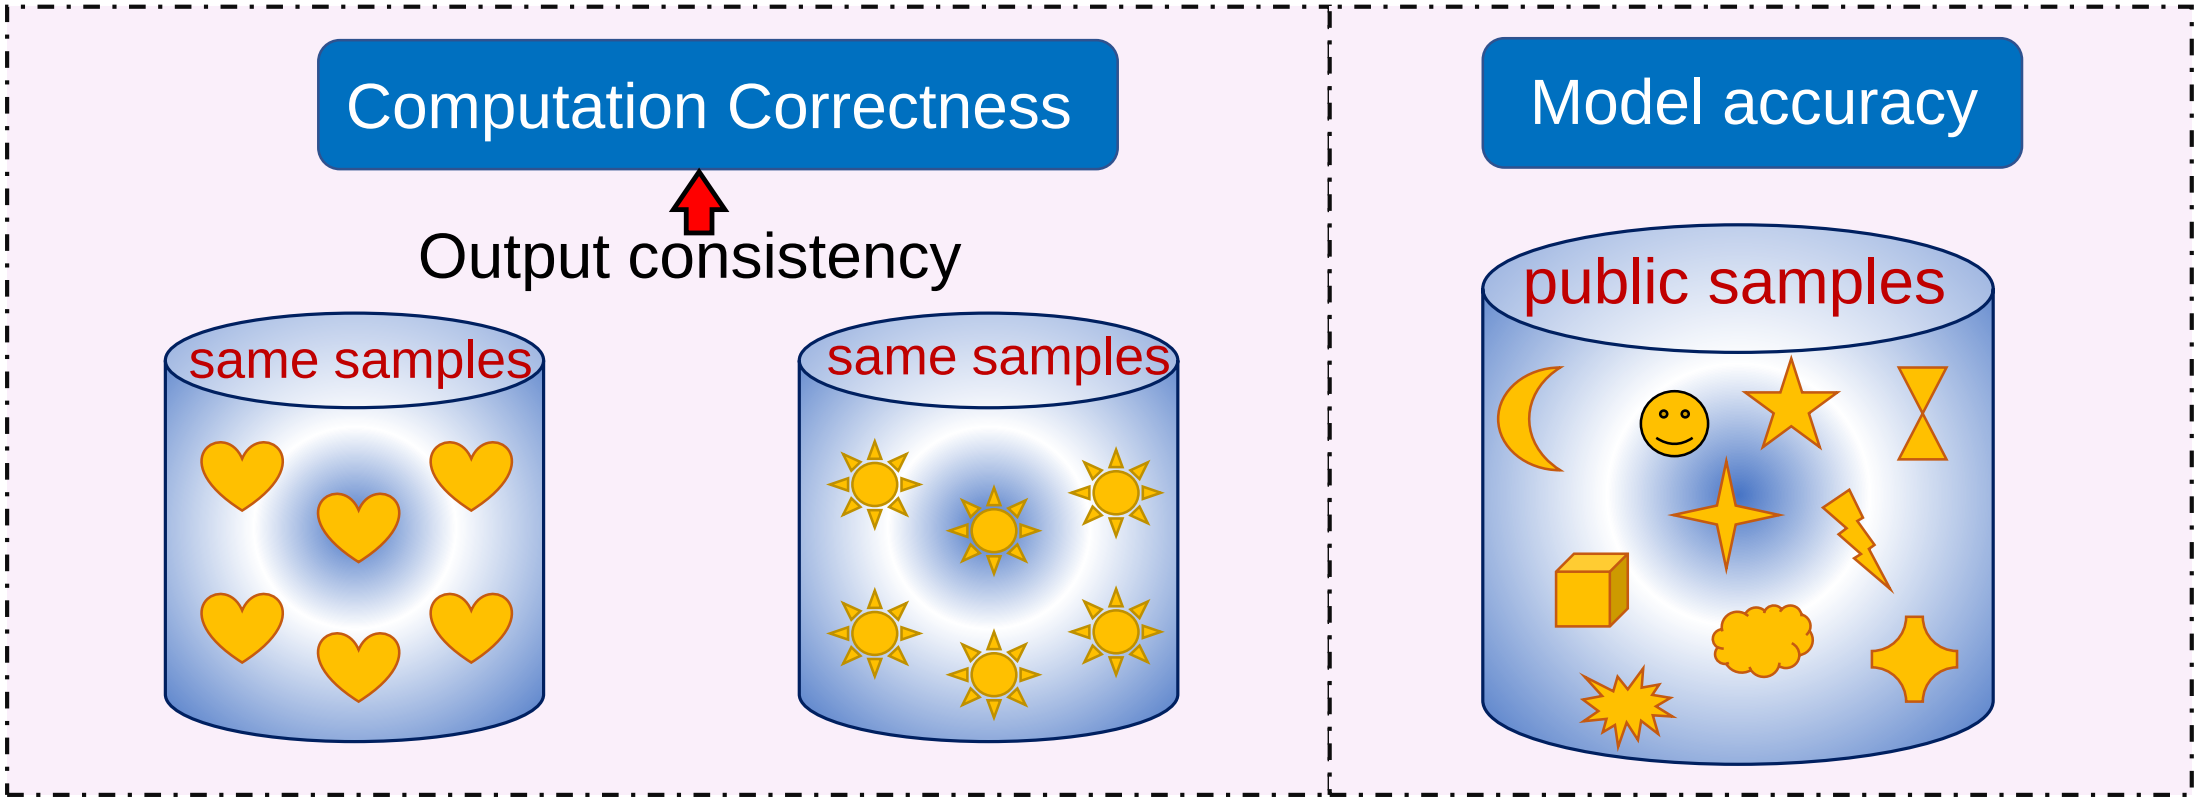

Optimal Number Selection

Client

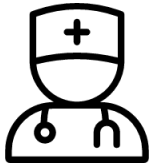

Given R, select appropriate B, T

Security Requirement

- Detect server's cheating

Cost Requirement

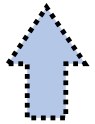

- Decrease the average cost

Security Requirement

Server succeeds in cheating

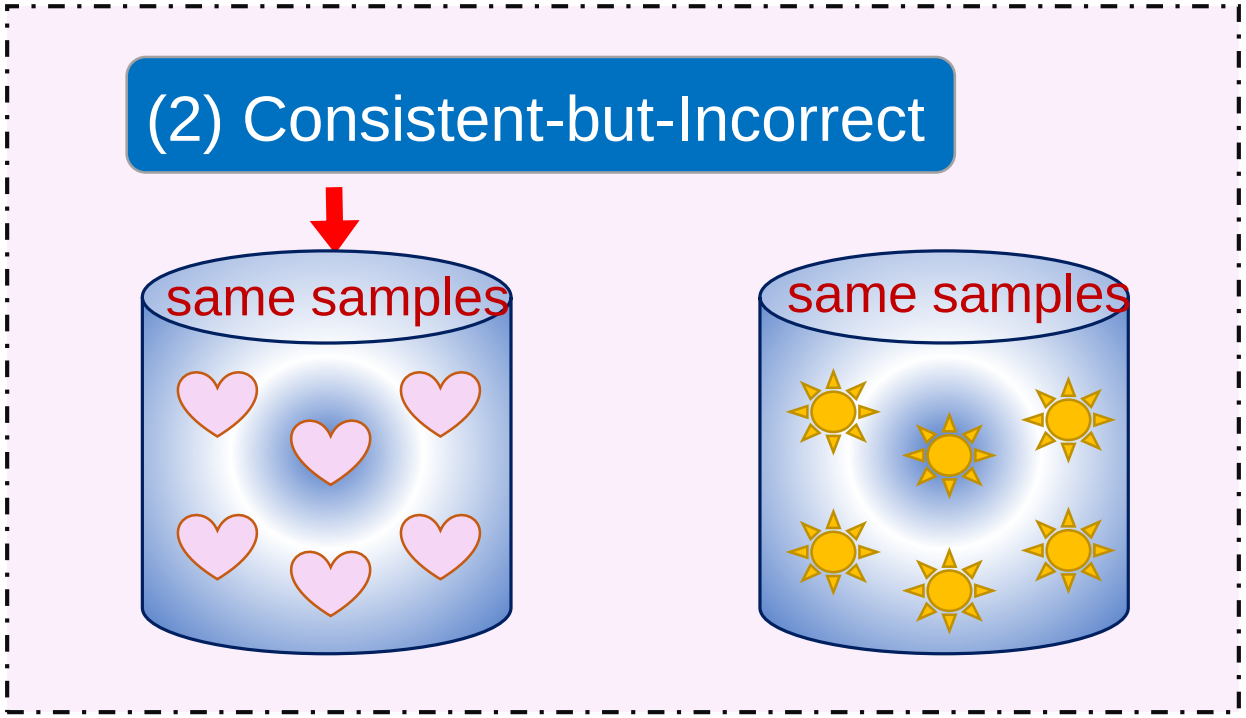

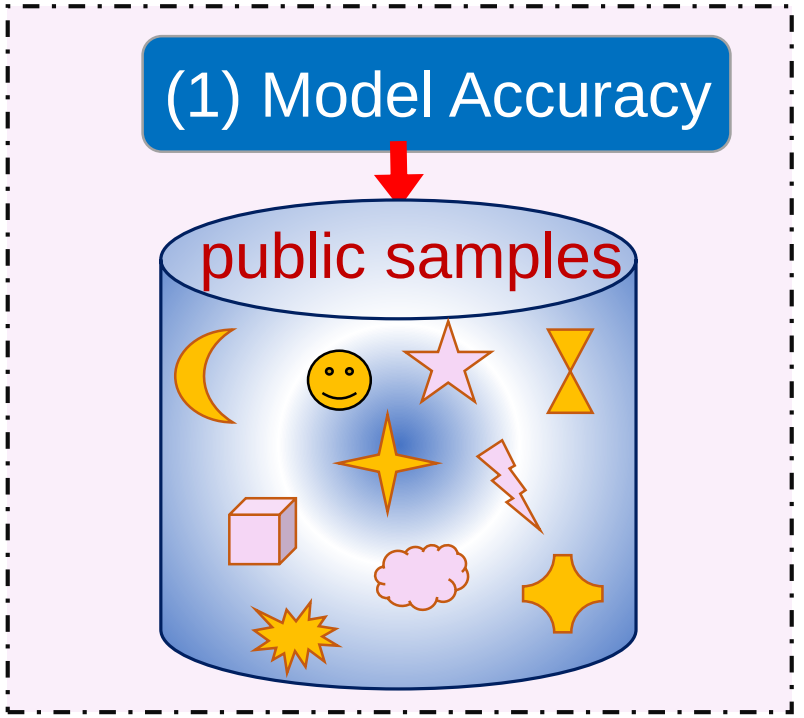

$$\Pr[E_B] = \frac{\binom{R}{i}(iB)!\,(RB-iB)!}{(RB)!} \qquad (2)$$

$$\Pr[E_T] = \frac{\binom{RB+T-iB}{T}}{\binom{RB+T}{T}} \qquad (1)$$

Client Selects Numbers Ensuring Security

Client

Search for the optimal numbers ensuring security

Security Requirements → Cost Optimization

$$\Pr_{success} \leq 2^{-\lambda}$$

$$\Pr_{success} = \Pr[E_T] \times \Pr[E_B]$$

satisfy

$$\underset{B,T}{arg\ min}\ \text{Cost(B,T,R)} = \frac{RB + T}{R}$$

Related Works

Popular related works

Threat Models
- Semi-Honest Security
- Malicious Security

Secret Sharing

NDSS15' ABY

CCS19' LevioSA (IPS compiler)

S&P20' Cryptflow (TEE)

S&P21' SIRNN

Usenix21' ABY2.0

CCS17' MiniONN

Usenix18' Gazelle

Usenix20' DELPHI

CCS20' CrypTFlow2

Usenix21' MUSE (MAC)

Usenix22' Cheetah

Usenix22' SIMC (MAC)

Homomorphic Encryption + Garbled Circuits/Secret Sharing

Model Quality

Performance

Table I: Comparison between Cheetah-based *Fusion* and *LevioSA* (CCS19')

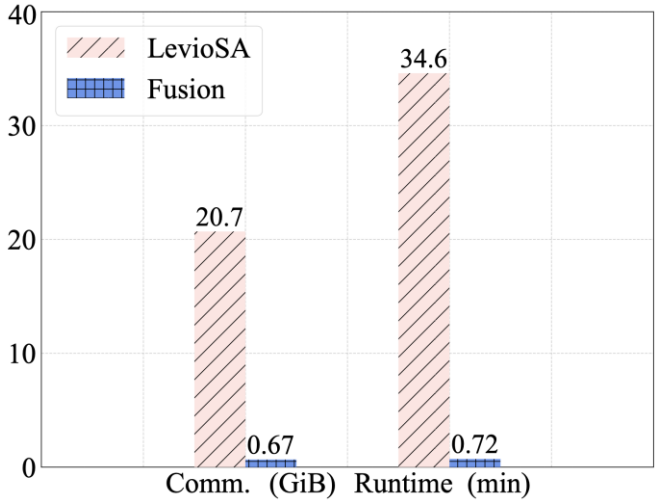

Runtime: 48.06× faster

Communication: 30.90× less

Table II: Performance of Fusion using different semi-honest inference protocols

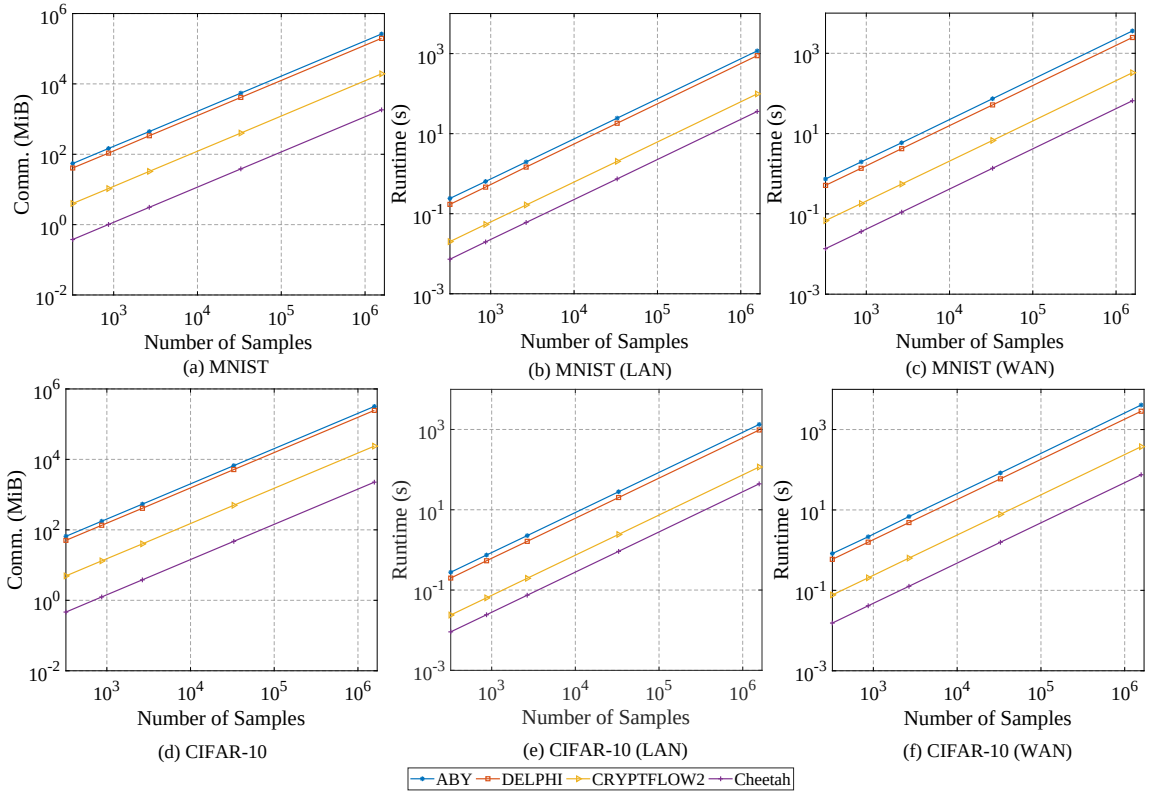

Performance

Table III: Performance of Cheetah-based Fusion and comparison with semi-honest inference protocols

	Scheme	*MNIST*			*CIFAR-10*		
		Comm.	**LAN**	**WAN**	**Comm.**	**LAN**	**WAN**
Fusion	(2^3, 8, 100)	24.499	487.500	850.000	30.080	575.000	975.000
	(2^5, 7, 100)	12.102	228.125	425.000	14.838	284.375	481.250
	(2^7, 6, 100)	8.106	154.688	283.594	9.949	189.844	323.438
	(2^9, 5, 100)	6.210	118.164	215.820	7.617	145.508	246.289
	(2^{13}, 4, 100)	4.799	90.686	166.968	5.886	112.341	191.724
	(2^{19}, 3, 100)	3.580	68.204	125.570	4.407	84.297	143.249
	CrypTFlow2 [59]	12.591	62.499	208.392	15.473	73.736	238.012
	DELPHI [46]	128.412	563.924	1573.818	160.079	617.572	1814.989
	ABY [13]	170.980	741.813	2293.657	207.421	850.366	2591.182

3.5X 0.9X 1.7X
35.9X 8.3X 12.5X
47.8X 10.9X 18.3X

Table IV: Performance on ResNet50

Scheme	Comm.	LAN	WAN
Fusion (2^3, 8, 100)	39.921	20.410	34.241
Fusion (2^5, 7, 100)	19.714	10.082	16.912
Fusion (2^7, 6, 100)	13.205	6.750	11.326
Fusion (2^9, 5, 100)	10.117	5.173	8.678
CrypTFlow2 [59] (SCI_{HE})	26.742	3.988	10.204
CrypTFlow2 [59] (SCI_{OT})	281.497	4.795	39.466

2.64X 1.30X slower 1.18X faster

Conclusion

Strong Security

- Model accuracy
- Computation correctness
- Privacy preservation

High Efficiency

- Low average overhead
- Comparable efficiency with semi-honest protocols