Privately Evaluating Decision Trees and Random Forests

David J. Wu, Tony Feng, Michael Naehrig, and Kristin Lauter

July, 2016

Machine Learning as a Service

Big Data + Machine Learning = New Applications

patient profile and
symptoms

recommended
treatment plan

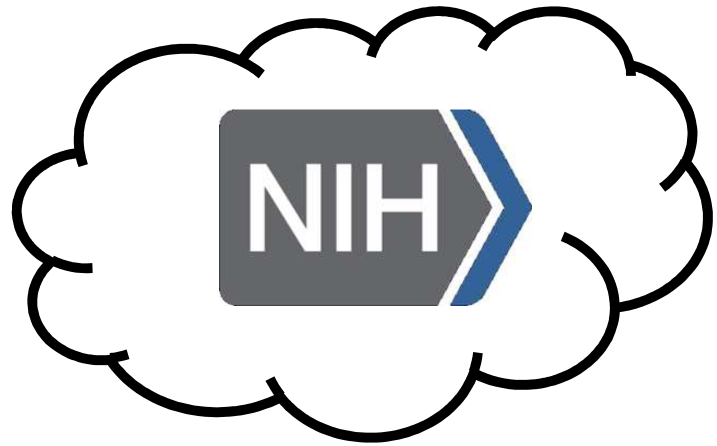

Machine Learning as a Service

Big Data + Machine Learning = New Risks

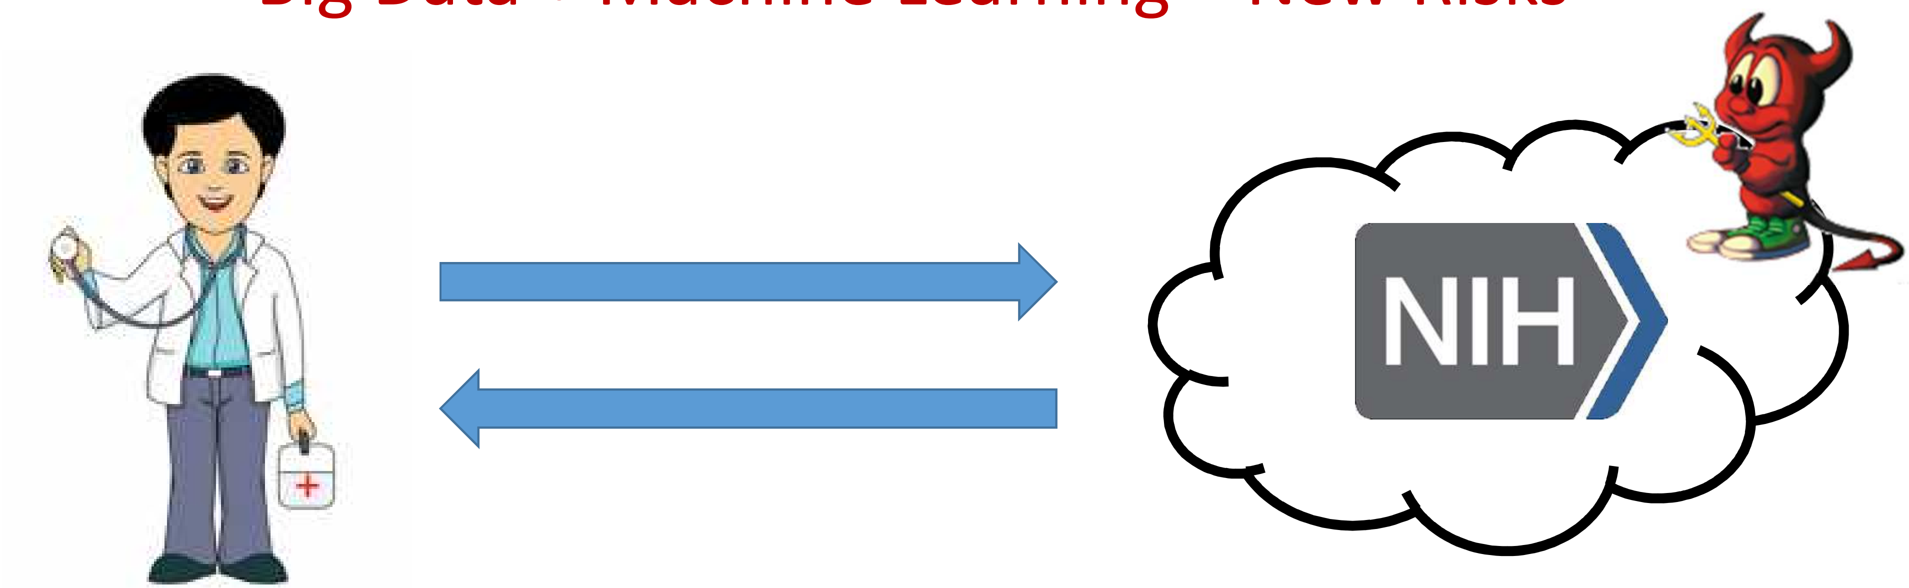

adversary that compromises cloud
service learns patient profile

Machine Learning as a Service

Big Data + Machine Learning = New Risks

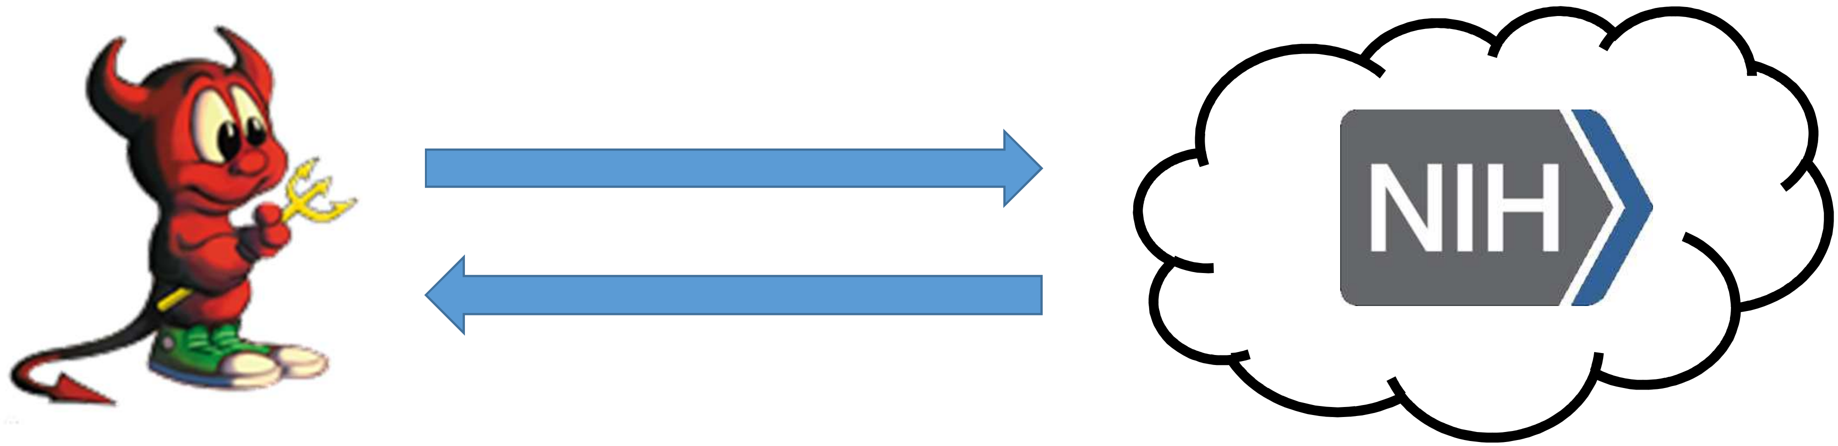

malicious client might recover
information about the model

Our Work: Decision Trees

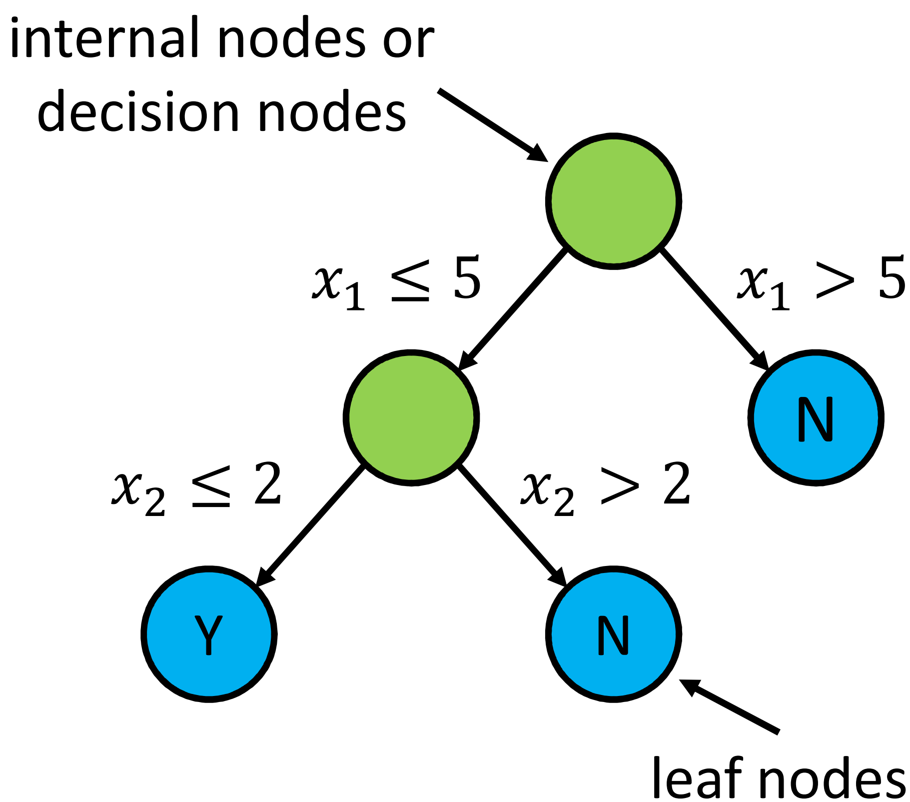

- Nonlinear models for regression or classification
- Consists of a series of decision variables (tests on the feature vector)
- Evaluation corresponds to tree traversal

Input: feature vector $[x_1, \ldots, x_n]$

Fully Private Decision Tree Evaluation

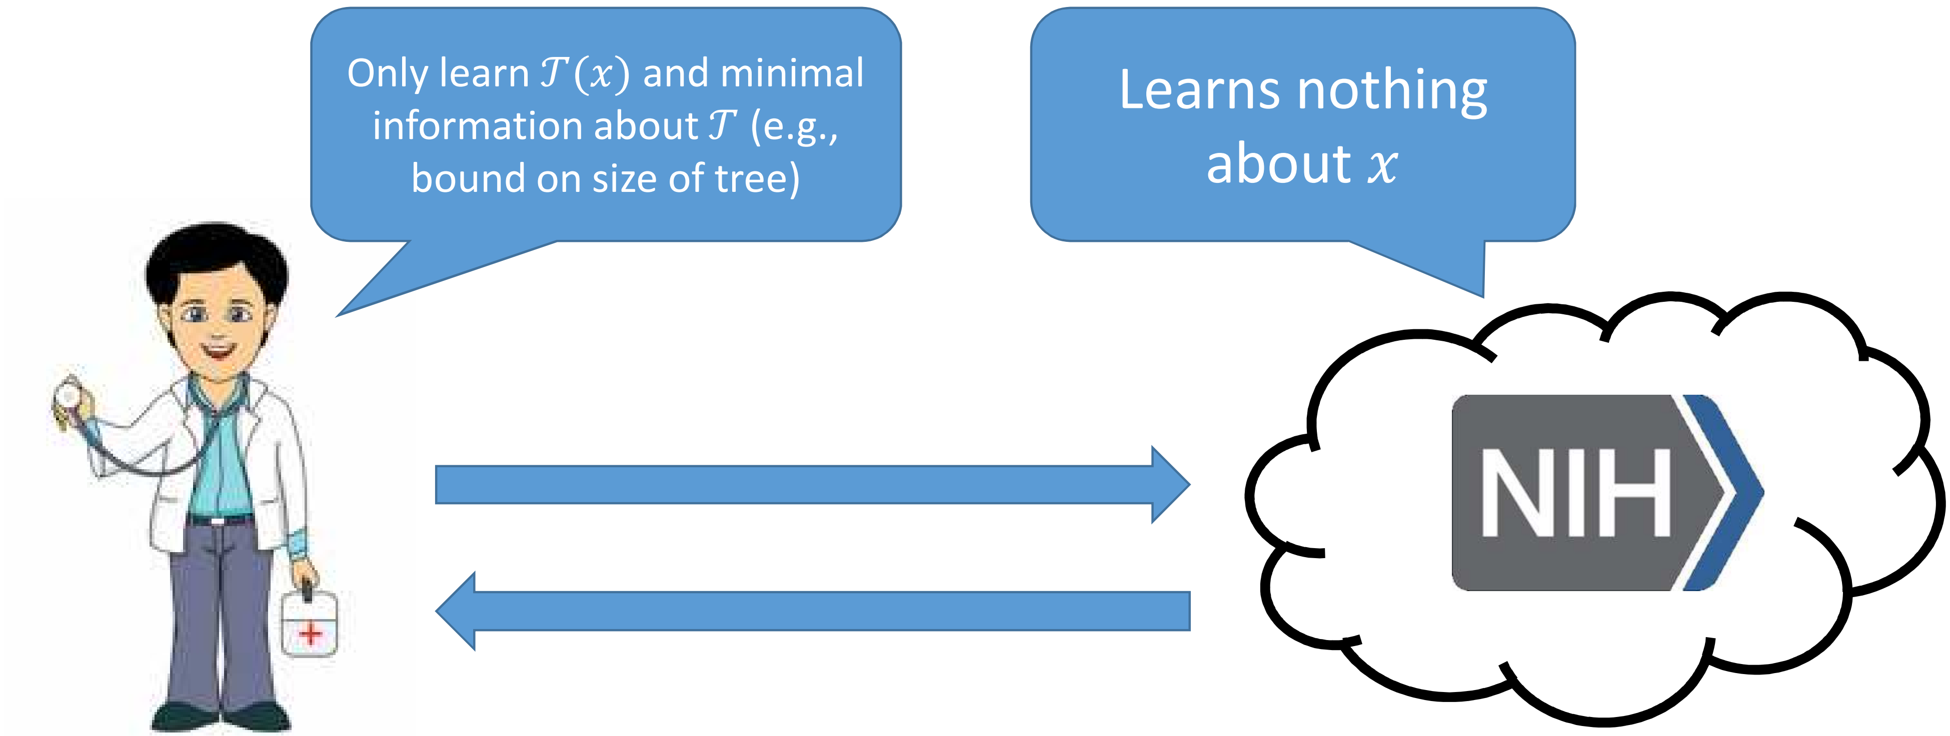

input: feature vector x

input: decision tree $\mathcal{T}$

Fully Private Decision Tree Evaluation

Focus on model evaluation –
assume server already has model

input: feature vector x

input: decision tree $\mathcal{T}$

Comparison of Approaches

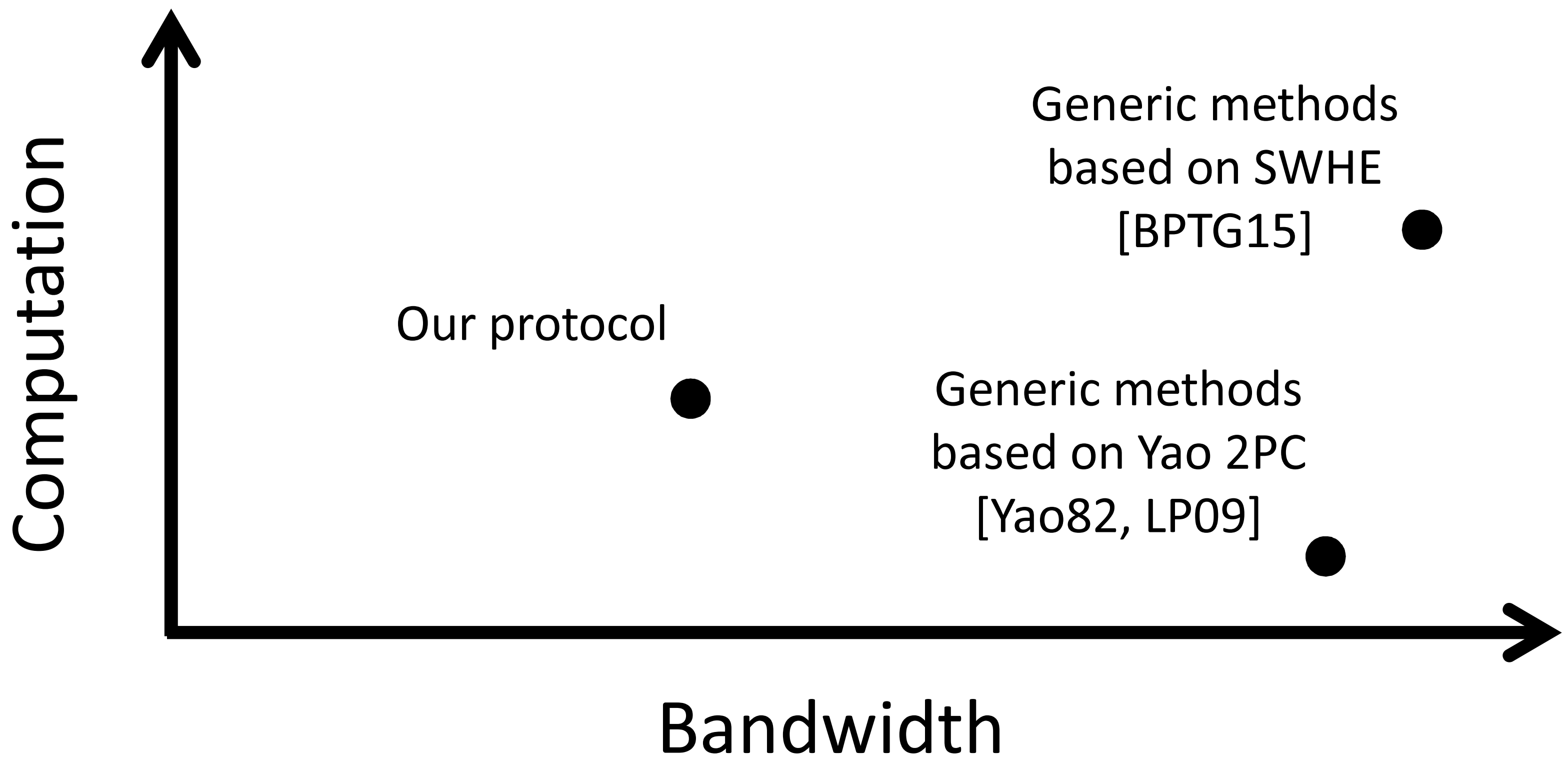

Not drawn to scale

Comparison of Approaches

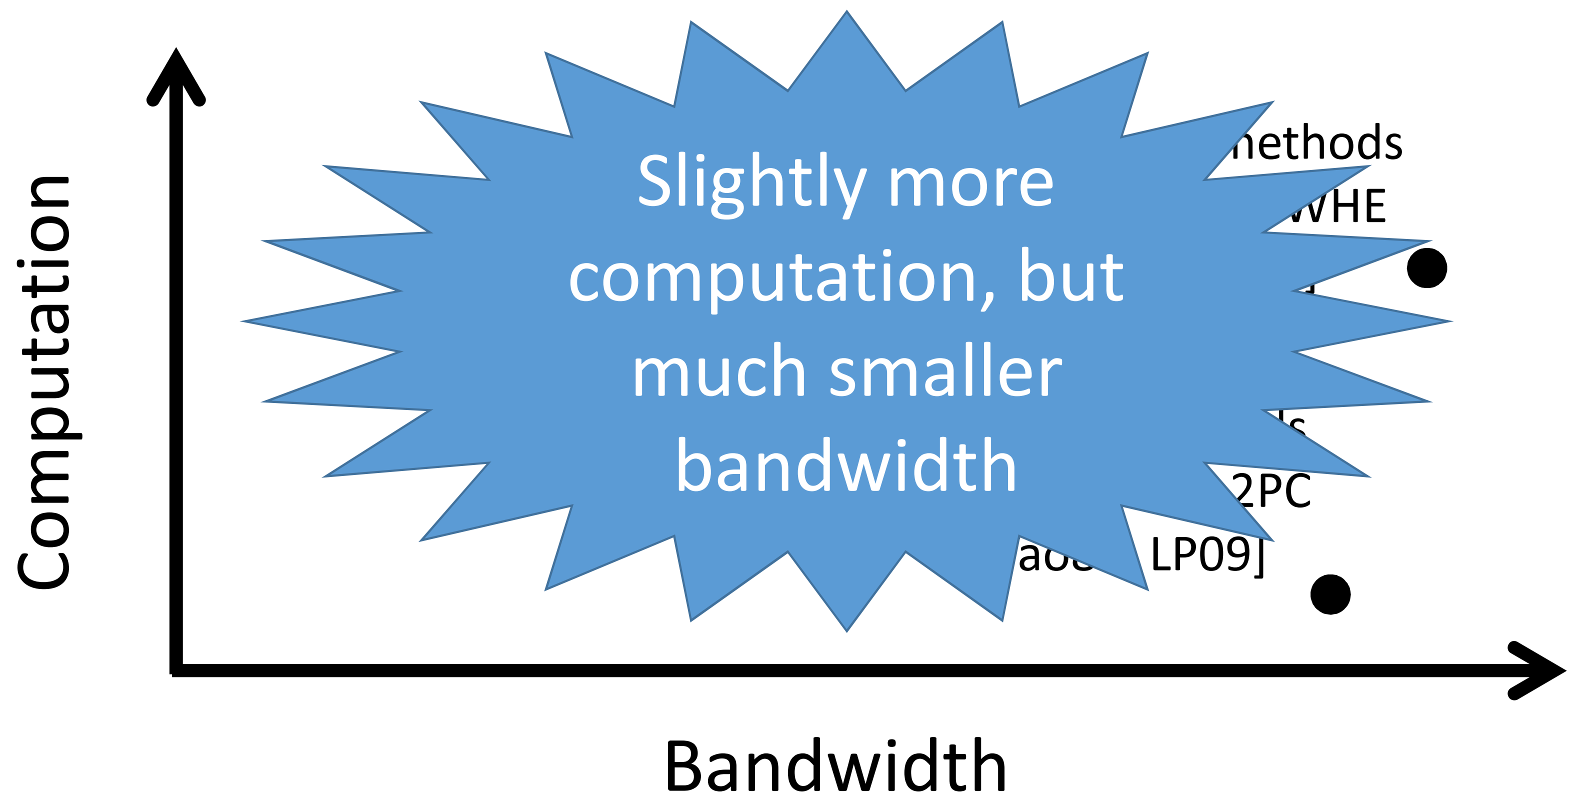

Not drawn to scale

Protocol Building Blocks: Comparisons

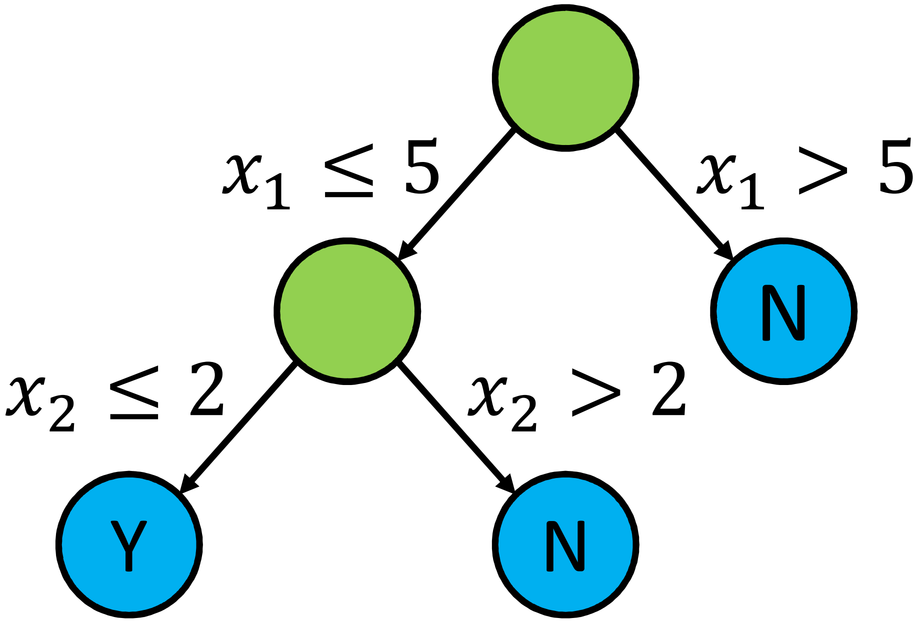

Require protocol to compare components of client's feature vector with thresholds

Comparison Protocol [DGK07, BPGT15]

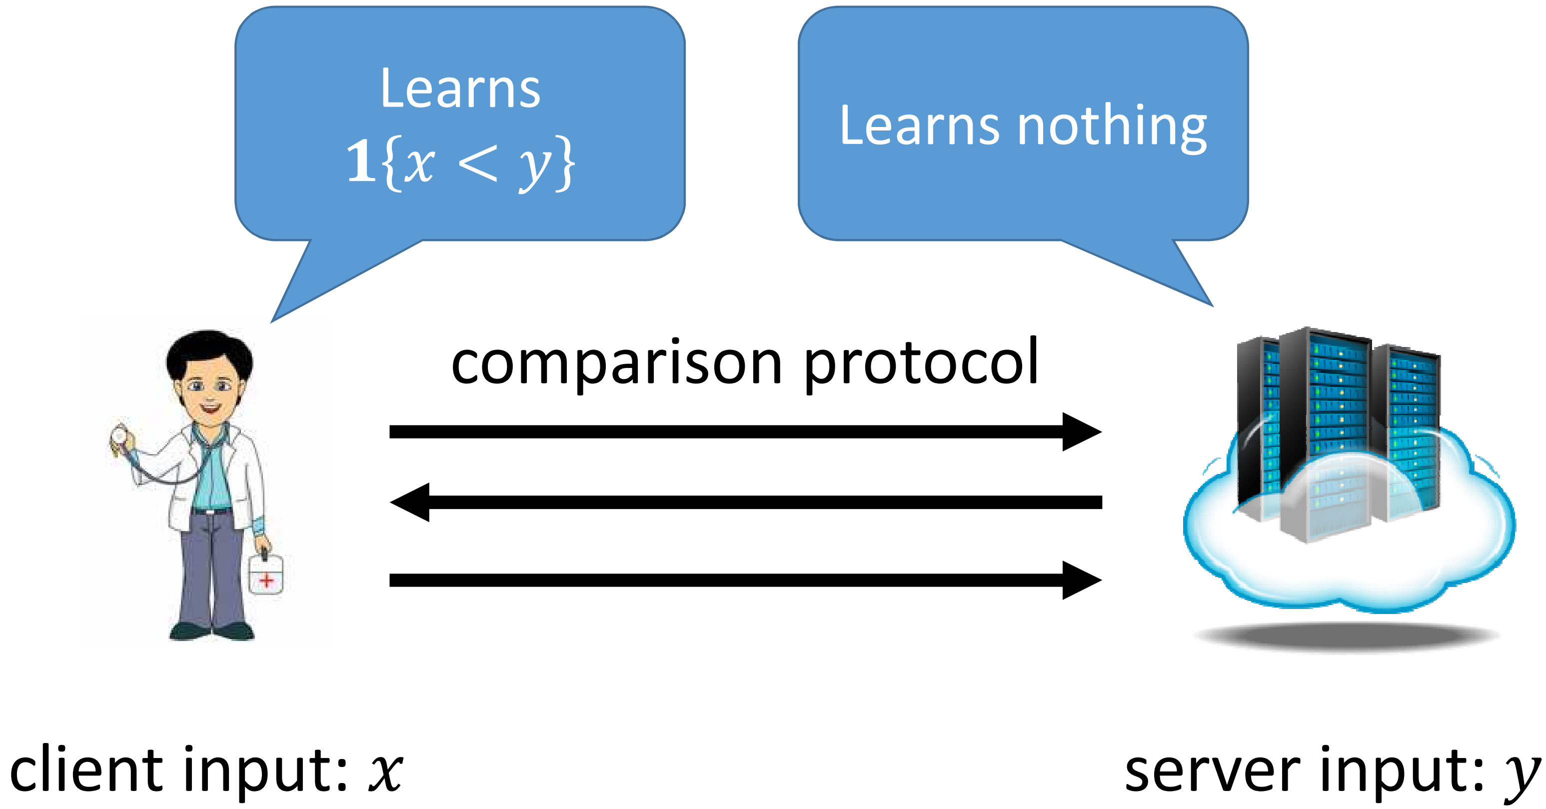

Private Decision Tree Evaluation

Suppose client knows b_1, b_2, and the structure of the tree

Then, client can compute the *index* of the outcome

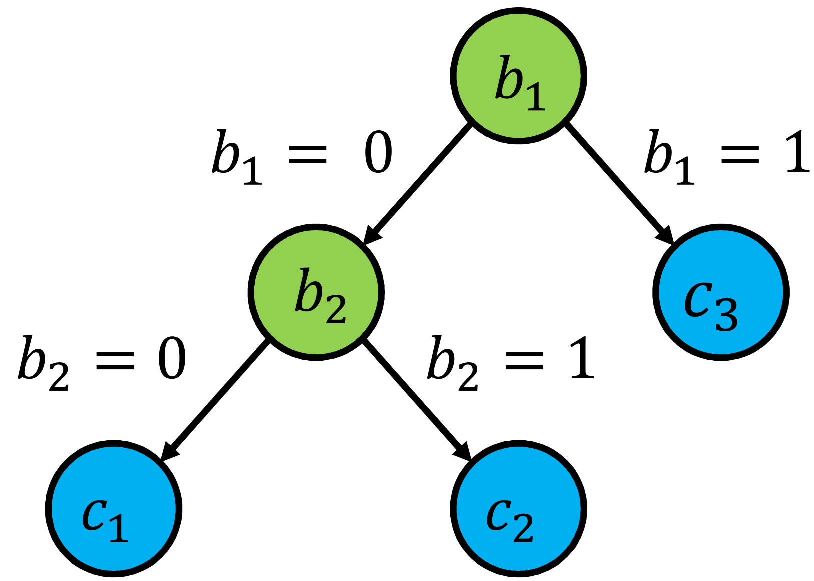

Private Decision Tree Evaluation

Suppose client knows the index of the outcome

Problem reduces to oblivious transfer: treat leaves as database, client knows index

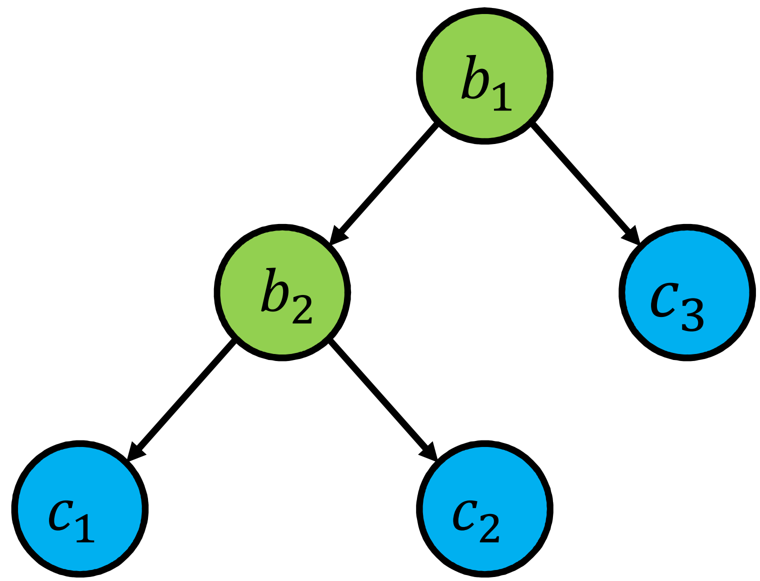

Oblivious Transfer (OT) [Kil88, NP99, NP01]

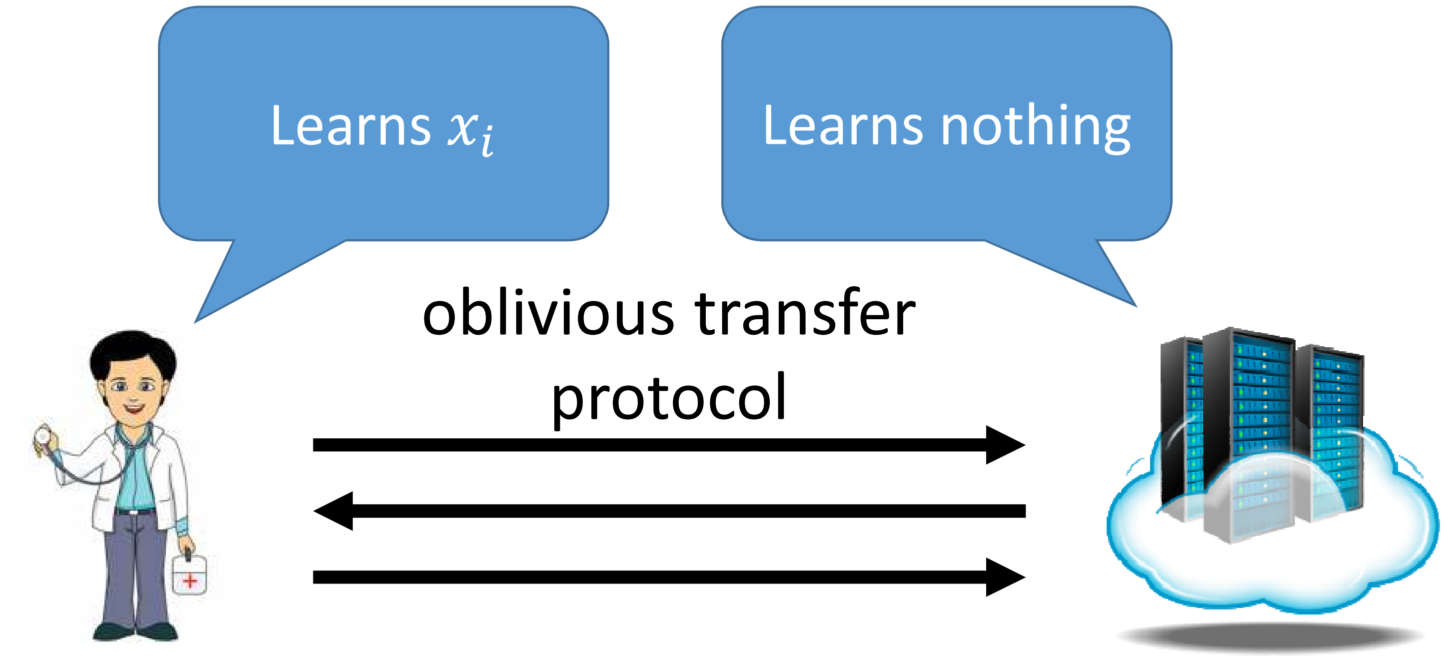

client input: index i

server input: database $\{x_1, \ldots, x_n\}$

Private Decision Tree Evaluation

Suppose client knows the index of the outcome

Problem reduces to oblivious transfer: treat leaves as database, client knows index

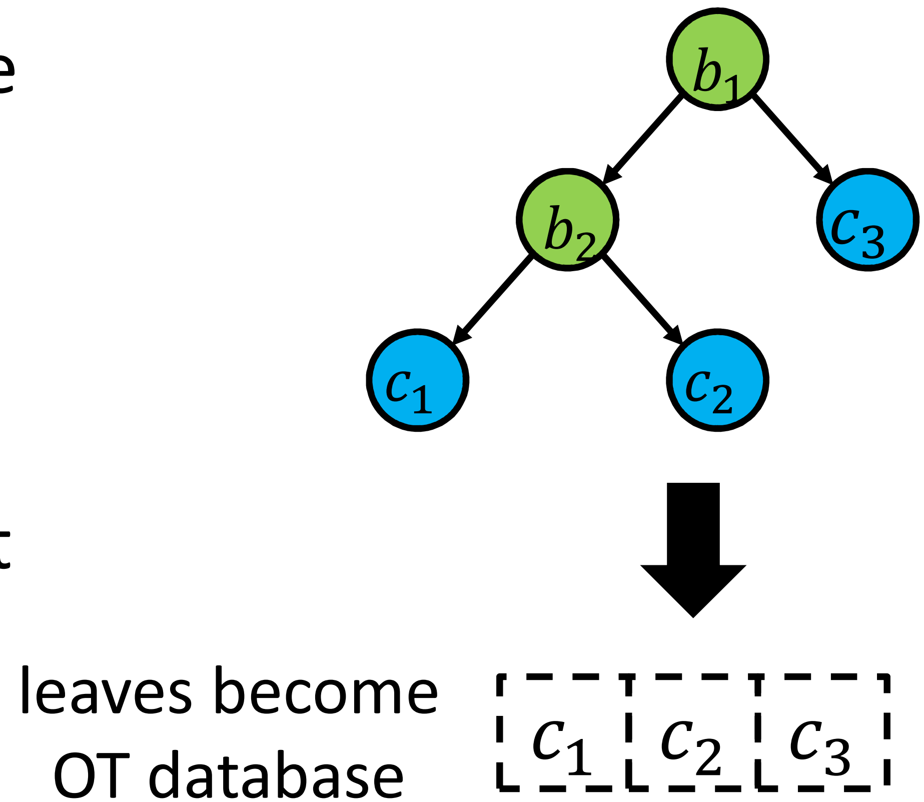

Private Decision Tree Evaluation

1. Client obtains b_1, b_2 using comparison protocol
2. Client uses OT to retrieve classification value

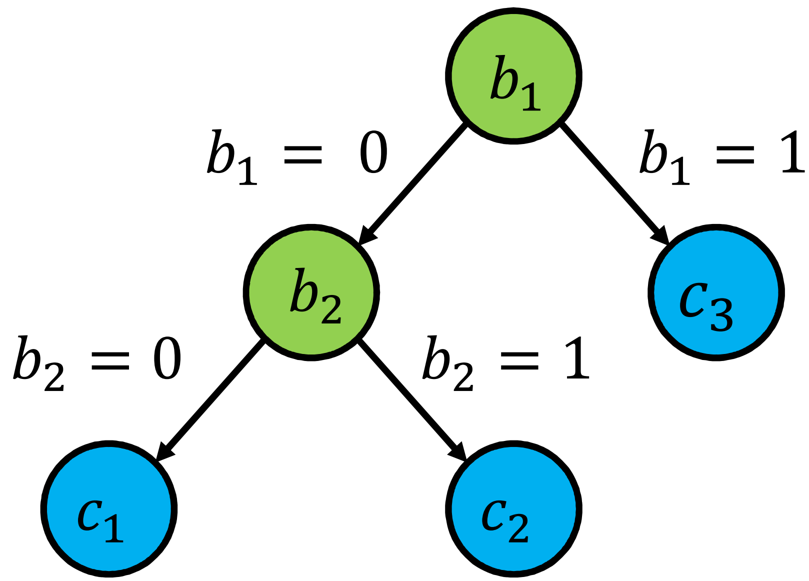

Problem: Requires client to learn/know structure of the tree

Hiding the Structure

1. **Padding:** Insert "dummy" nodes to obtain complete tree

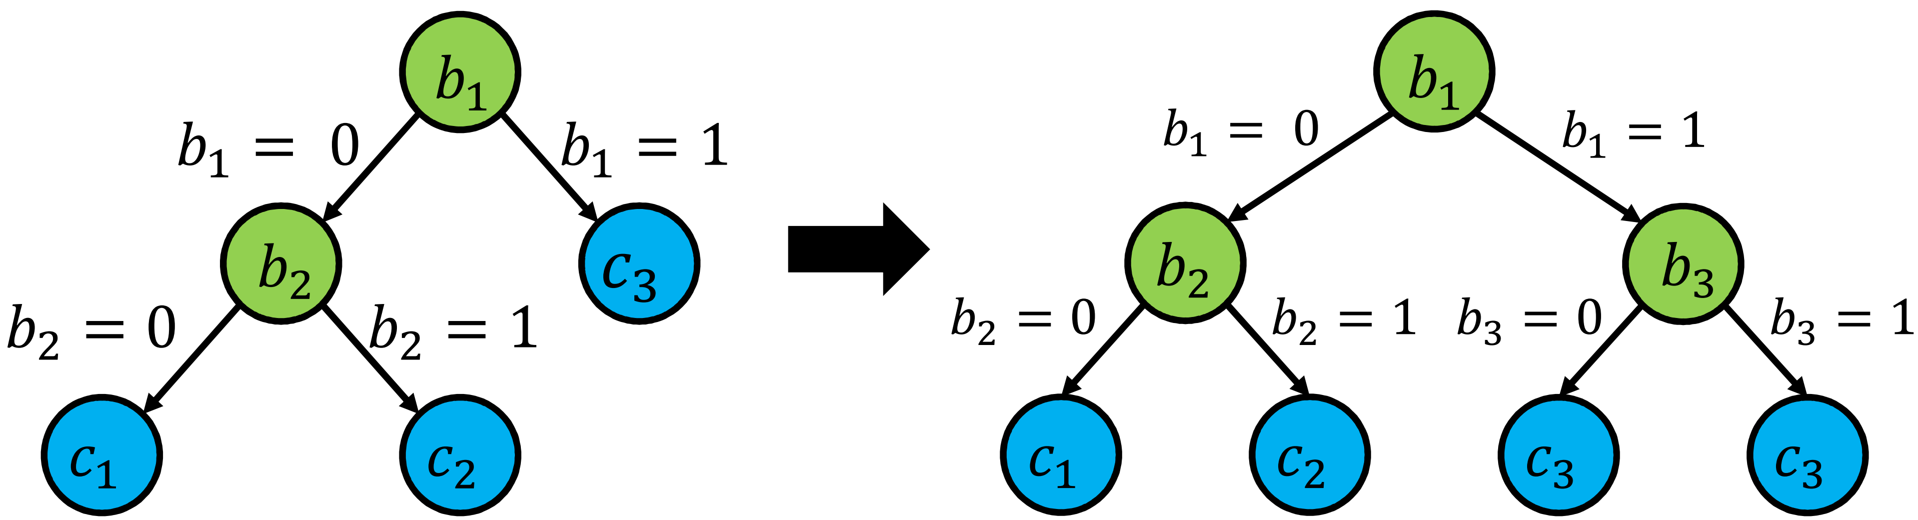

Hiding the Structure

2. Randomization: Randomly flip decision variables:

$$\bar{b}_i \coloneqq 1 - b_i$$

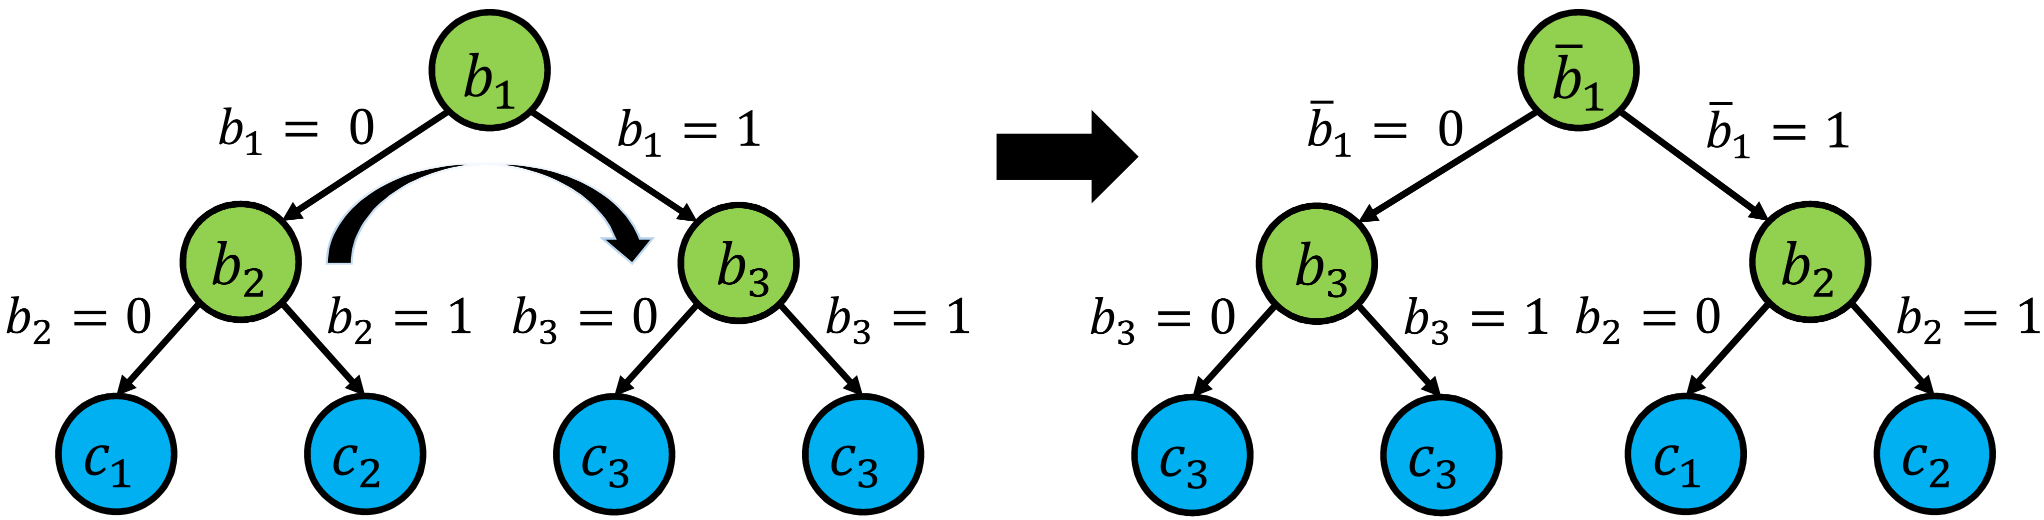

Private Decision Tree Evaluation

1. **Server:** Pad and permute the decision tree
2. **Server & Client:** Comparison protocol to compute b_i in permuted tree
3. **Client:** Compute the index j of the leaf node
4. **Client & Server:** Engage in OT to obtain c_j

Theorem. This protocol is secure against *semi-honest* adversaries.

Further Extensions

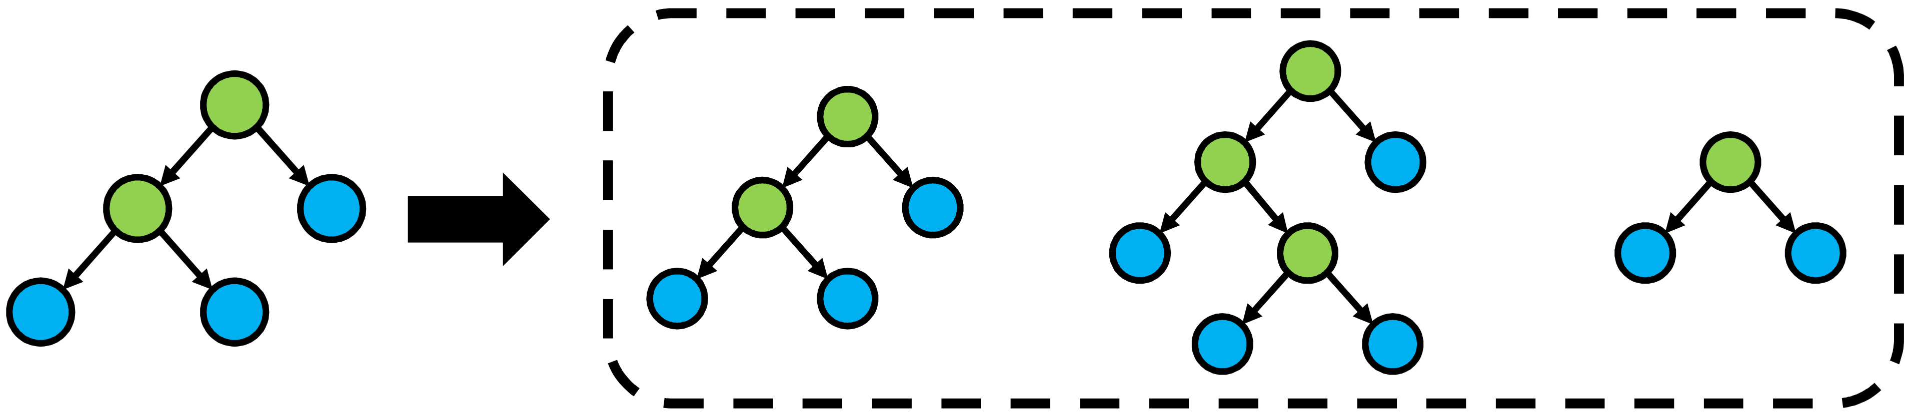

evaluating random forests without
revealing individual classifications

Ensuring security against **malicious** adversaries

See paper for details!

Experiments

Implemented private decision tree + random forest protocol

Benchmarks taken between a laptop client and an EC2 server

Decision Tree Evaluation on ECG Data

	Security Level	Computation (s)		Bandwidth (KB)
		Client	Server	
[BFK+09]	80	2.609	6.260	112.2
[BPGT14]	80	2.297	1.723	3555
Generic 2PC (Estimated)	128	-	-	$\geq$ 180.5
This work	**128**	**0.091**	**0.188**	**101.9**

Experimental Parameters:

- Data Dimension: 6
- Depth of Decision Tree: 4
- Number of Comparisons: 6

Decision Tree Evaluation on ECG Data

	Security Level	Computation (s)		Bandwidth (KB)
[BFK+				
[BPGT				
Generic (Estima				
This w				

10x faster than previous protocols

- Data Dimension: 6
- Depth of Decision Tree: 4
- Number of Comparisons: 6

Performance for Complete Decision Trees

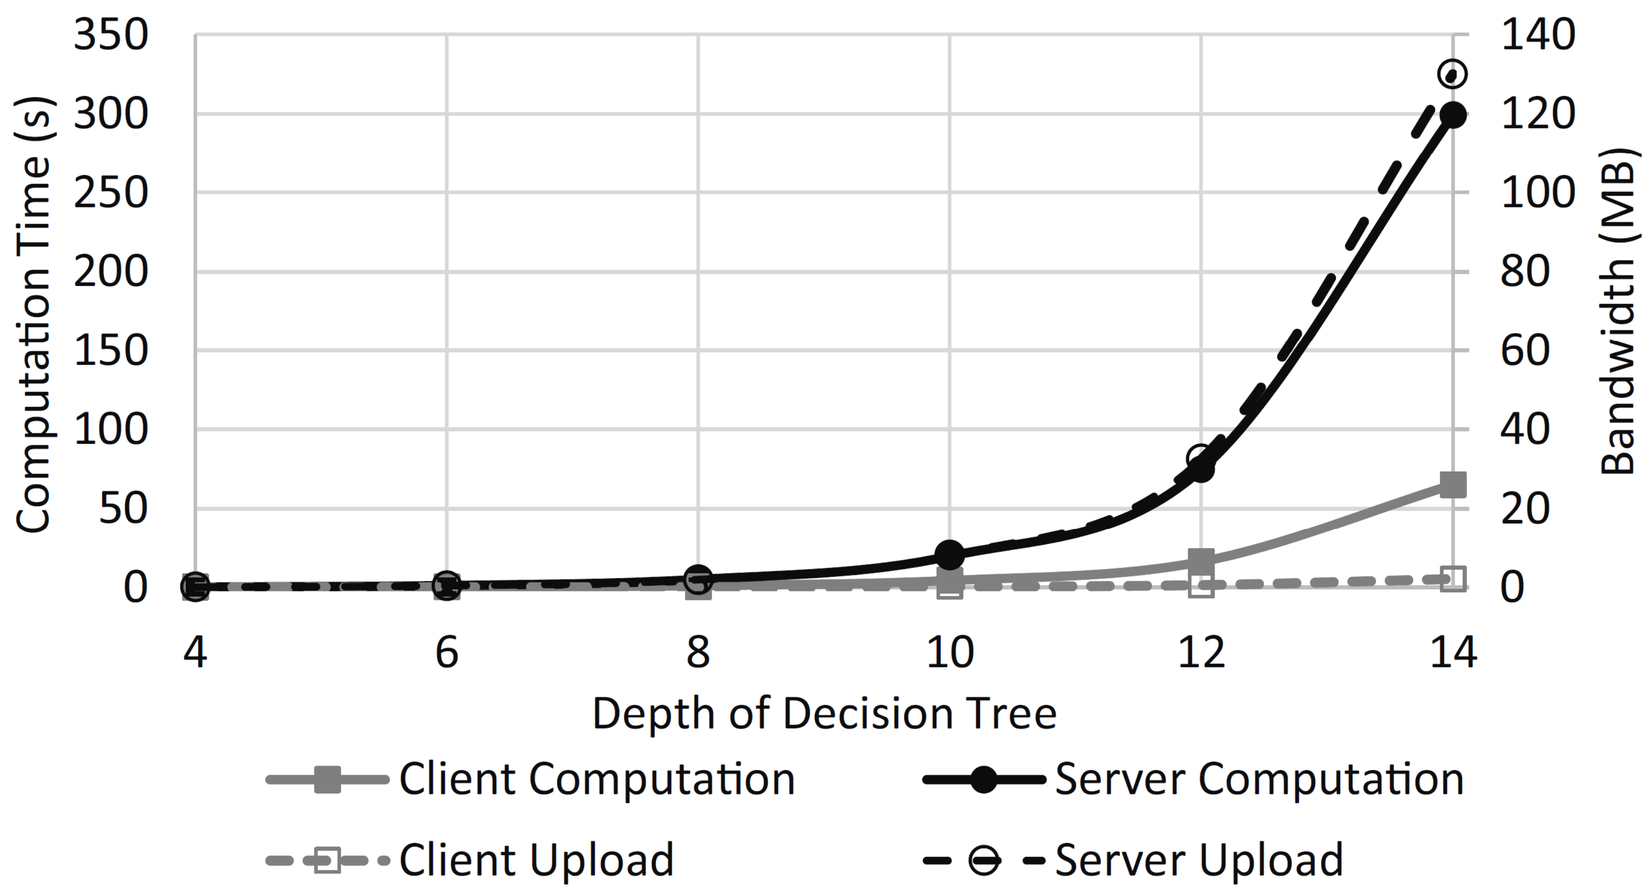

Conclusions

Simple protocols for decision tree evaluation in both semi-honest (and malicious) setting

Semi-honest (and malicious-secure) decision tree protocols provide new computation/communication tradeoffs

Thanks!

http://eprint.iacr.org/2015/386