

Two-Party Decision Tree Training from Updatable Order-Revealing Encryption

Robin Berger¹, Felix Dörre¹, Alexander Koch²
2024-03-07 @ACNS 2024

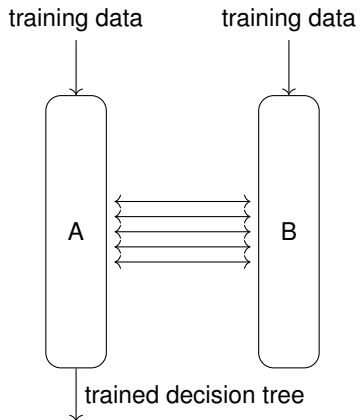
¹: KASTEL, KIT



²: CNRS & IRIF



Our Setting



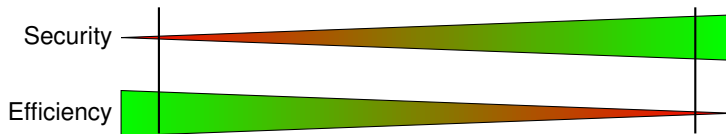
Theory/Practice

Practice:

- Focus on efficiency
- Computation on plaintext data

Theory:

- Focus on security
- Computation using generic
Multiparty Computation



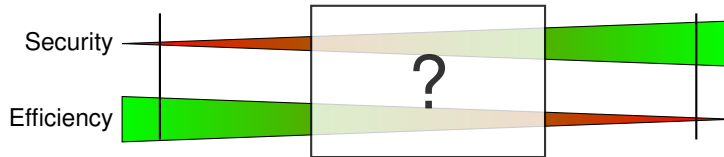
Theory/Practice

Practice:

- Focus on efficiency
- Computation on plaintext data

Theory:

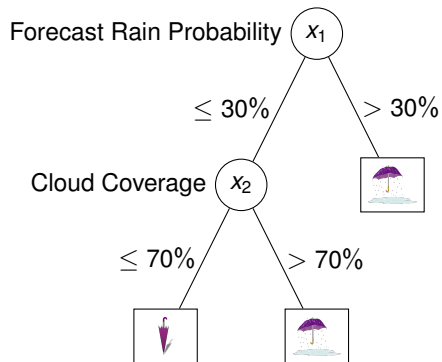
- Focus on security
- Computation using generic Multiparty Computation



Can we achieve a speedup by allowing some leakage?
 (Specifically for decision tree training)

Decision Trees

Should I bring an umbrella?

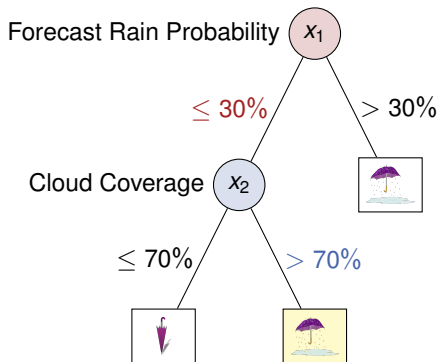


Example datapoint:

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 27\% \\ 91\% \end{pmatrix}$$

Decision Trees

Should I bring an umbrella?



Example datapoint:

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 27\% \\ 91\% \end{pmatrix}$$

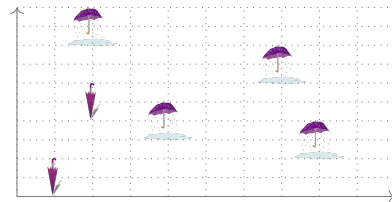
How to Train a Decision Tree?

Decision Tree Training

```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



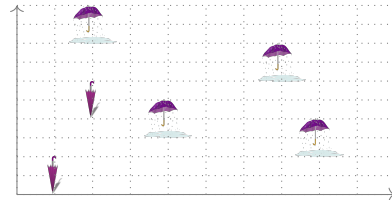
How to Train a Decision Tree?

Decision Tree Training

```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



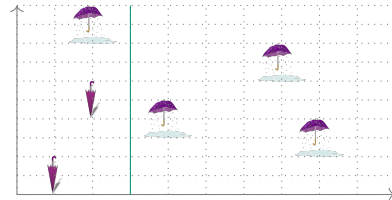
How to Train a Decision Tree?

Decision Tree Training

```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



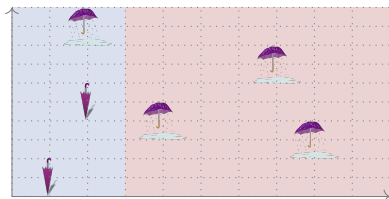
How to Train a Decision Tree?

Decision Tree Training

```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



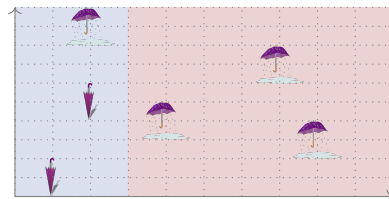
How to Train a Decision Tree?

Decision Tree Training

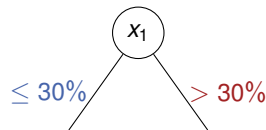
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



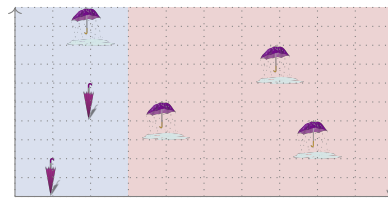
How to Train a Decision Tree?

Decision Tree Training

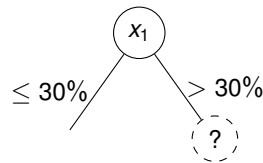
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



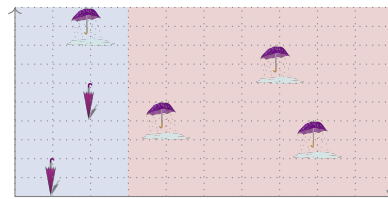
How to Train a Decision Tree?

Decision Tree Training

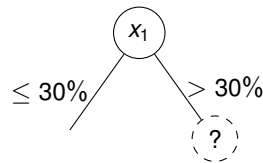
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



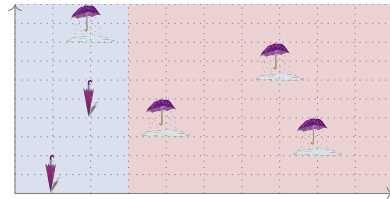
How to Train a Decision Tree?

Decision Tree Training

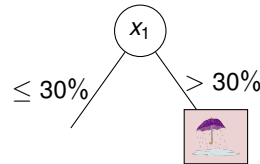
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



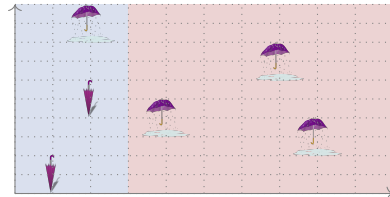
How to Train a Decision Tree?

Decision Tree Training

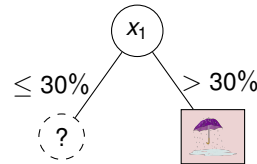
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



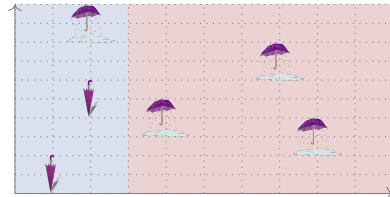
How to Train a Decision Tree?

Decision Tree Training

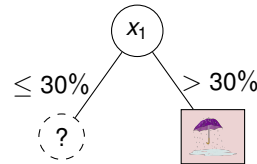
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



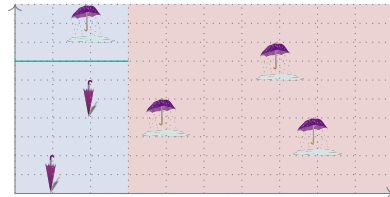
How to Train a Decision Tree?

Decision Tree Training

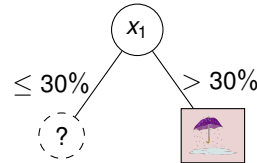
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



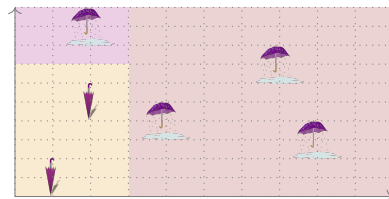
How to Train a Decision Tree?

Decision Tree Training

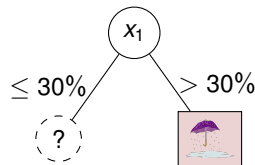
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



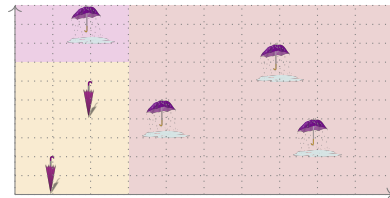
How to Train a Decision Tree?

Decision Tree Training

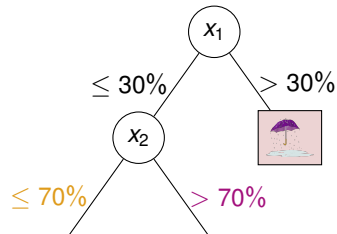
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



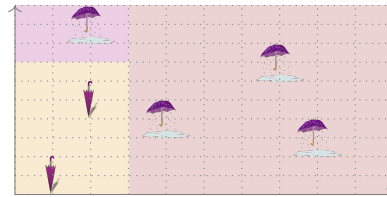
How to Train a Decision Tree?

Decision Tree Training

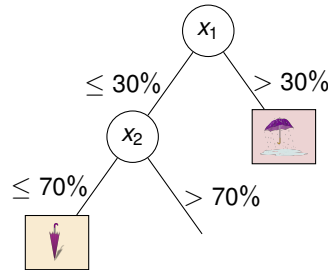
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



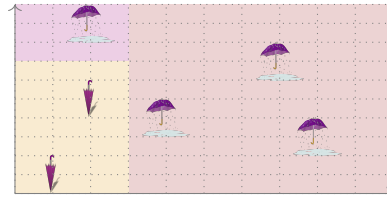
How to Train a Decision Tree?

Decision Tree Training

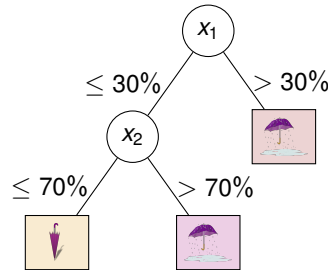
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



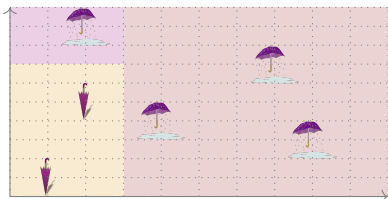
How to Train a Decision Tree?

Decision Tree Training

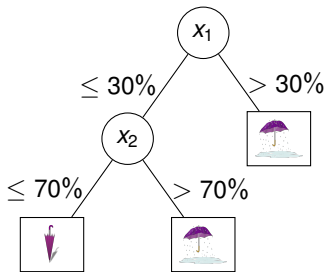
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

x_2 : Cloudyness



x_1 : Forecast Rain Probability



How to Train a Decision Tree?

Decision Tree Training

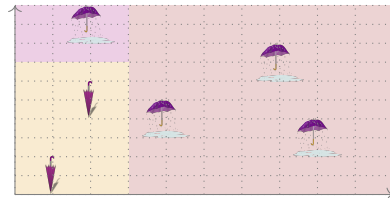
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

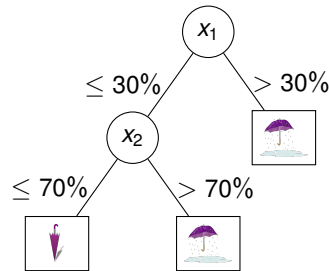
Required operations:

- Equality checks
- Comparisons

x_2 : Cloudyness



x_1 : Forecast Rain Probability



How to Train a Decision Tree?

Decision Tree Training

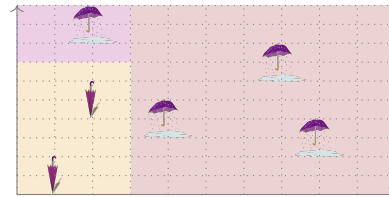
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

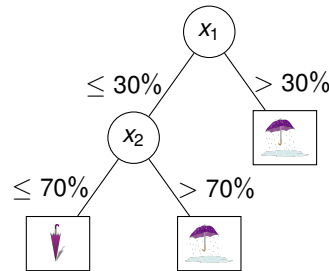
Required operations:

- Equality checks
- Comparisons

x_2 : Cloudyness



x_1 : Forecast Rain Probability



How to Train a Decision Tree?

Decision Tree Training

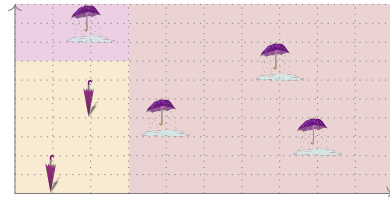
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

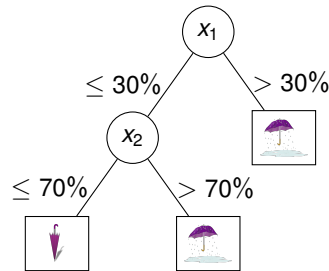
Required operations:

- Equality checks
- Comparisons

x_2 : Cloudyness



x_1 : Forecast Rain Probability



How to Train a Decision Tree?

Decision Tree Training

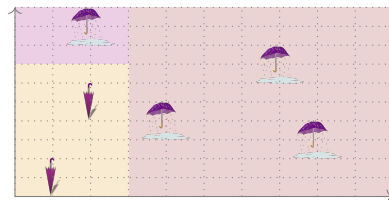
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return Leaf(label)
4:   end if
5:   split = FINDBESTSPLIT(data)
6:   (X, Y) := split(data)
7:   return InnerNode(split, TRAIN(X), TRAIN(Y))
8: end function
  
```

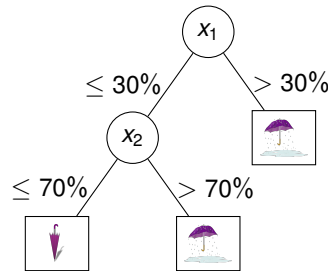
Required operations:

- Equality checks
- Comparisons

x_2 : Cloudyness



x_1 : Forecast Rain Probability

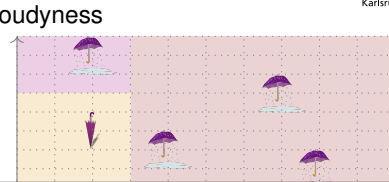


How to Train a Decision Tree?

Decision Tree Training

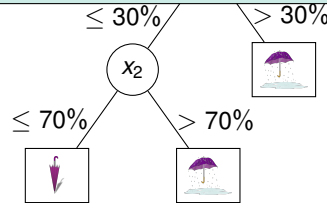
```

1: function TRAIN(data)
2:   if all datapoints in data have the same label then
3:     return leaf(label)
4:
5:
6:
7:
8: end function
  
```



Idea: Encrypt data with an encryption scheme that permits comparisons/equality checks.

ability








Required operations:

- Equality checks
- Comparisons


Order-Revealing Encryption



An Order-Revealing Encryption scheme consists of the following algorithms:


- Gen \rightarrow 
- Enc(, m) \rightarrow 
- Cmp(, ) $\rightarrow m_0 \stackrel{?}{<} m_1$


Order-Revealing Encryption

An Order-Revealing Encryption scheme consists of the following algorithms:

■ Gen \rightarrow 

■ Enc(, m) \rightarrow 

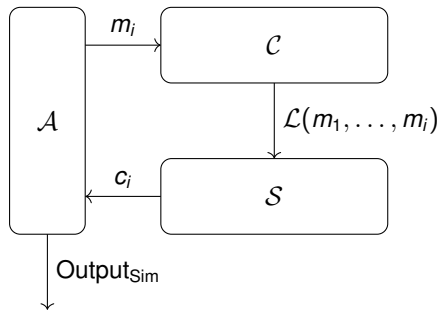
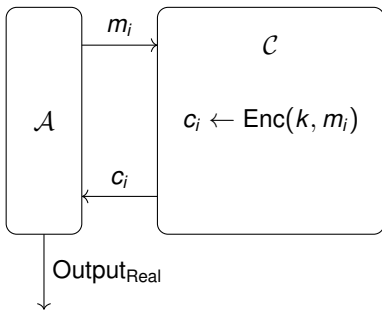
■ Cmp(, ) $\rightarrow m_0 \stackrel{?}{<} m_1$

* Dec(, ) $\rightarrow m$

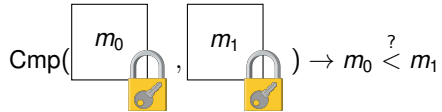
Security of Order-Revealing Encryption

ORE security w.r.t. leakage \mathcal{L} from [Che+16]

There exists a PPT \mathcal{S} , such that for all PPT \mathcal{A} : $\text{Output}_{\text{Real}} \stackrel{c}{\approx} \text{Output}_{\text{Sim}}$

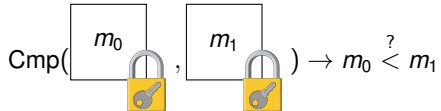


Comparisons on Ciphertexts

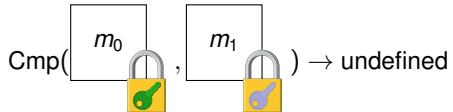


Comparisons on Ciphertexts

$\text{Cmp}(\text{[} m_0 \text{] } \text{[} m_1 \text{]}) \rightarrow m_0 < m_1$



$\text{Cmp}(\text{[} m_0 \text{] } \text{[} m_1 \text{]}) \rightarrow \text{undefined}$




Comparisons on Ciphertexts

$$\text{Cmp}(\underbrace{m_0}_{\text{lock}}, \underbrace{m_1}_{\text{lock}}) \rightarrow m_0 <^? m_1$$

$$\text{Cmp}(\underbrace{m_0}_{\text{lock}}, \underbrace{m_1}_{\text{lock}}) \rightarrow \text{undefined}$$

Problem when applied to decision tree training:

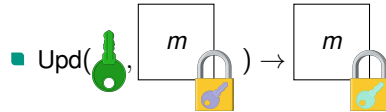
- All parties need to use the same .

Updatable Order-Revealing Encryption

Updatable Order-Revealing Encryption adds the following algorithms:

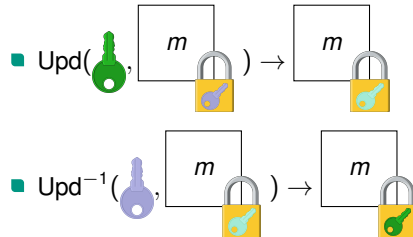
Updatable Order-Revealing Encryption

Updatable Order-Revealing Encryption adds the following algorithms:



Updatable Order-Revealing Encryption

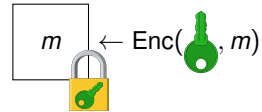
Updatable Order-Revealing Encryption adds the following algorithms:



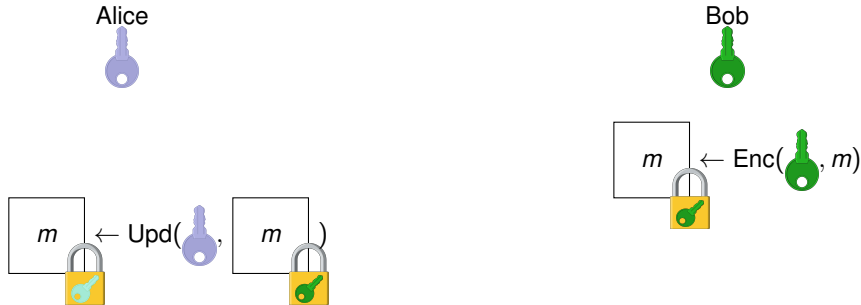
Interactive Encryption



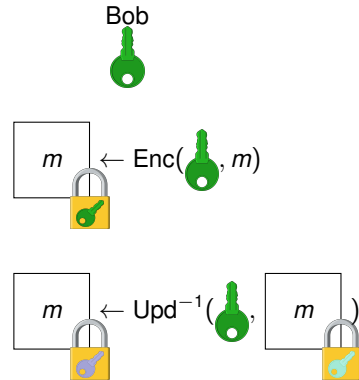
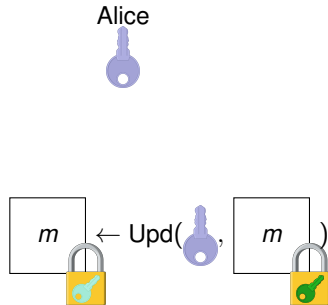
Interactive Encryption



Interactive Encryption



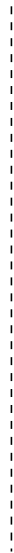
Interactive Encryption

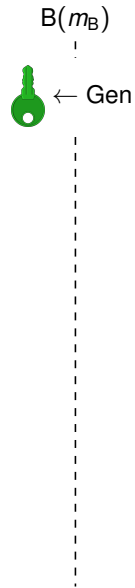
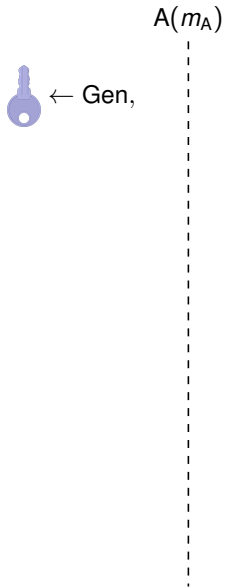


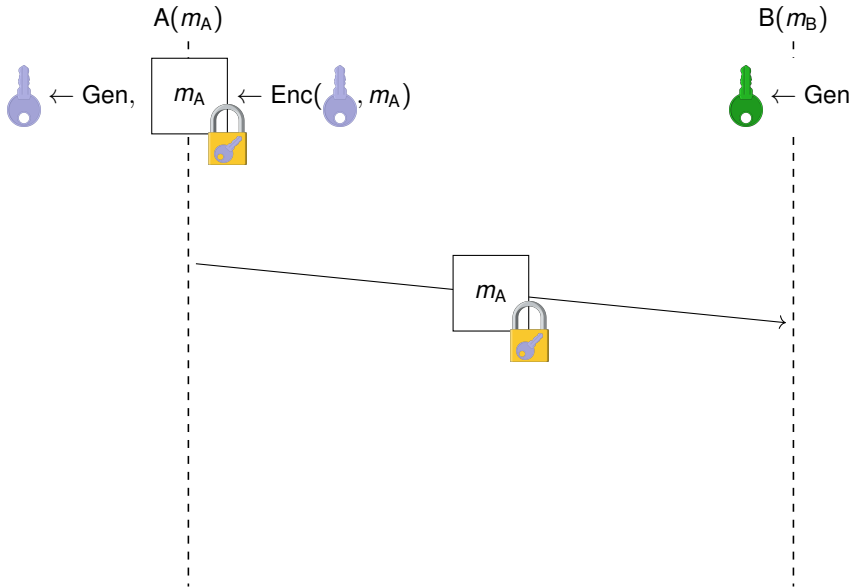
$A(m_A)$

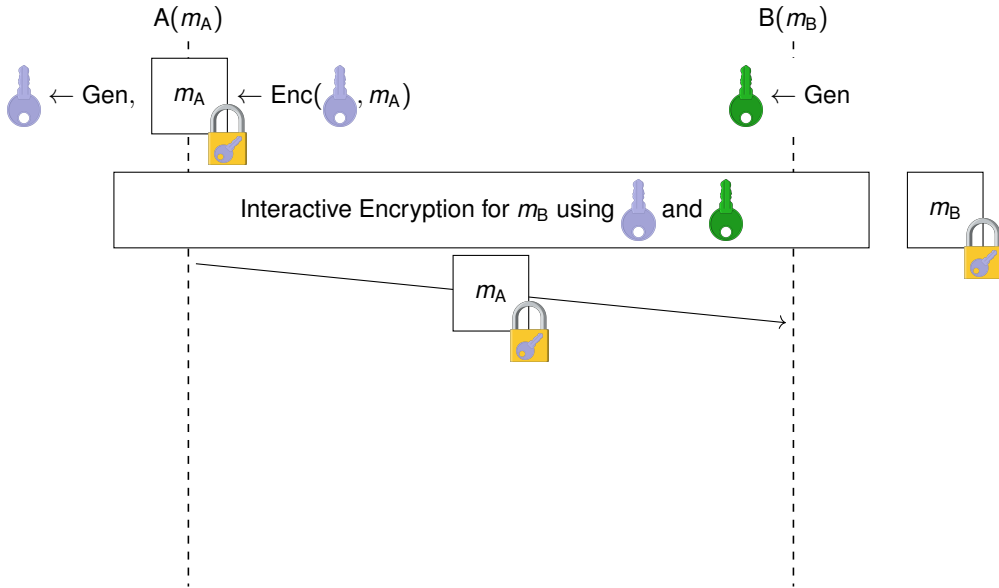


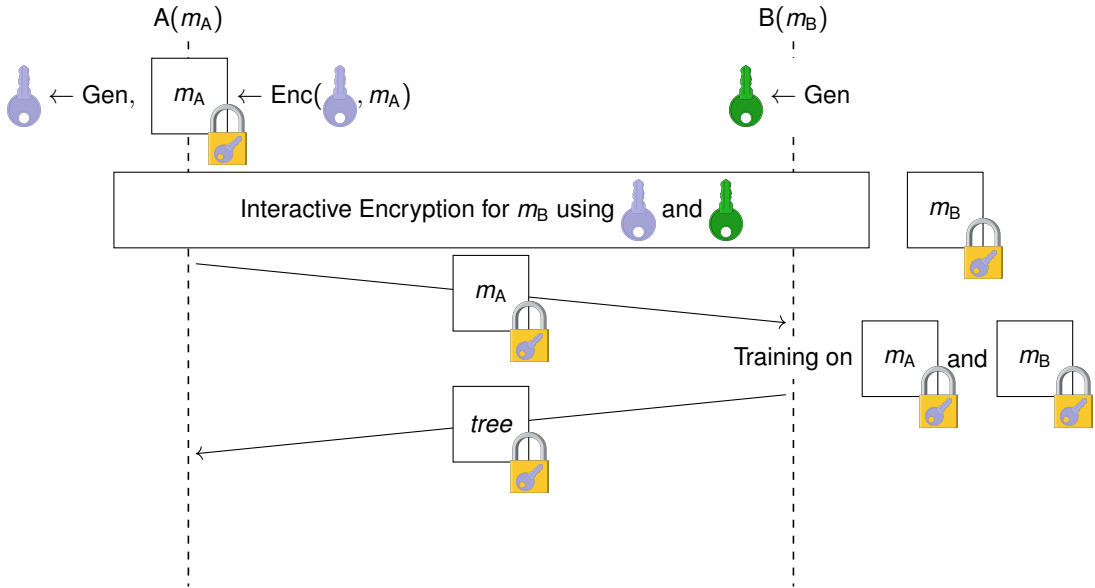
$B(m_B)$

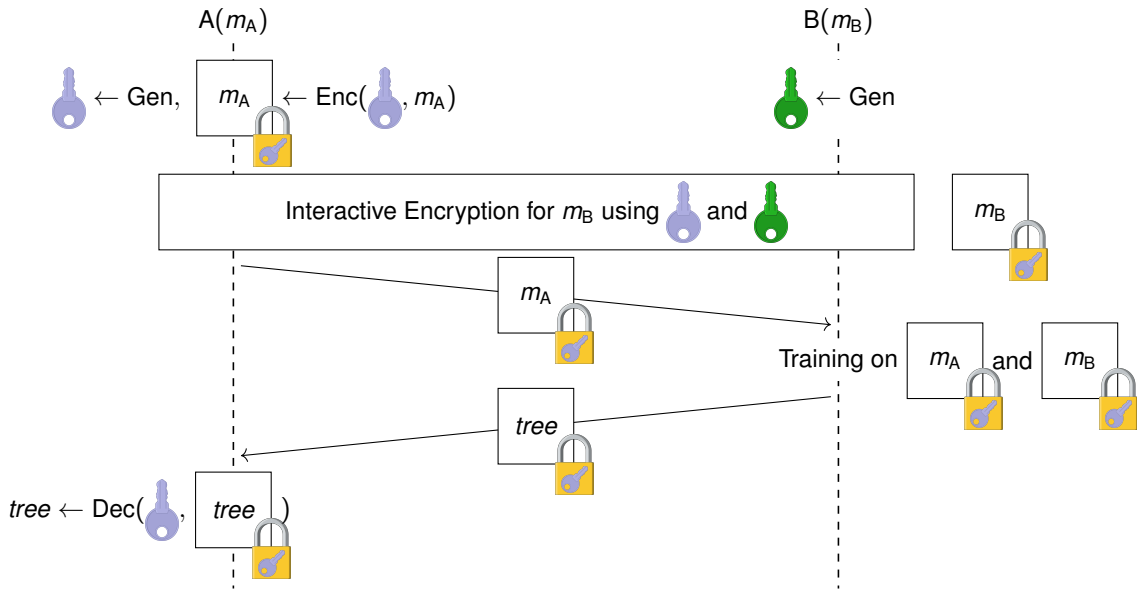




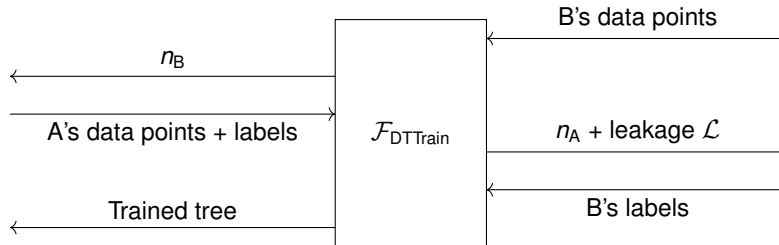








Semi-Honest Security of the Protocol



Performance evaluation

Benchmarks on

- Synthetic dataset
 - 2^{13} samples
 - 11 attributes
- 50ms network latency

Training Protocol	Runtime
Plaintext training	0.4s
Our work	20.1s
Hamada et al. [Ham+23]	4821.6s

Conclusion

Efficient approach for decision tree training if

- Training algorithm only requires comparisons and equality checks
- Leakage is acceptable
 - Case-by-case decision

References

- [Che+16] Nathan Chenette, Kevin Lewi, Stephen A. Weis, and David J. Wu. “Practical Order-Revealing Encryption with Limited Leakage”. In: 2016, pp. 474–493. DOI: 10.1007/978-3-662-52993-5_24.
- [Ham+23] Koki Hamada, Dai Ikarashi, Ryo Kikuchi, and Koji Chida. “Efficient decision tree training with new data structure for secure multi-party computation”. In: 2023.1 (Jan. 2023), pp. 343–364. DOI: 10.56553/popets-2023-0021.

Leakage

Leakage of our scheme:

$$\mathcal{L}(m_1, \dots, m_N) = \{(i, j, \text{hsb}(m_i \oplus m_j)) \mid m_i < m_j\}$$

Example:

message	leakage
00100	00???
01110	01???
10000	1??0?
10011	1==1?

Maximum leakage: $\mathcal{O}(\log N)$ bits per message