

Zero-Knowledge Proofs for Secure and Private Machine Learning

Dario Fiore | IMDEA Software Institute



Foundations and Applications of Zero-Knowledge Proofs | Edinburgh, UK | Sep 6, 2024



European Research Council
Established by the European Commission

Agenda

Security of ML inference

How to use ZKPs for secure ML

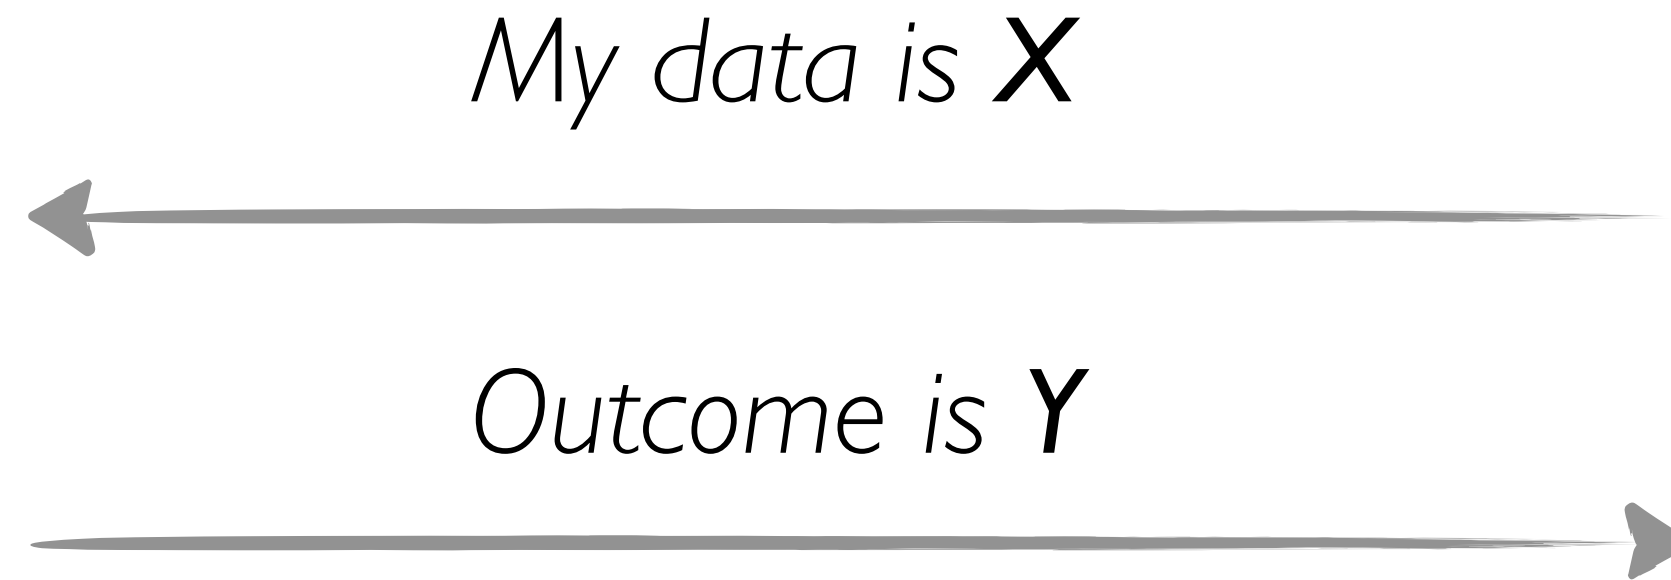
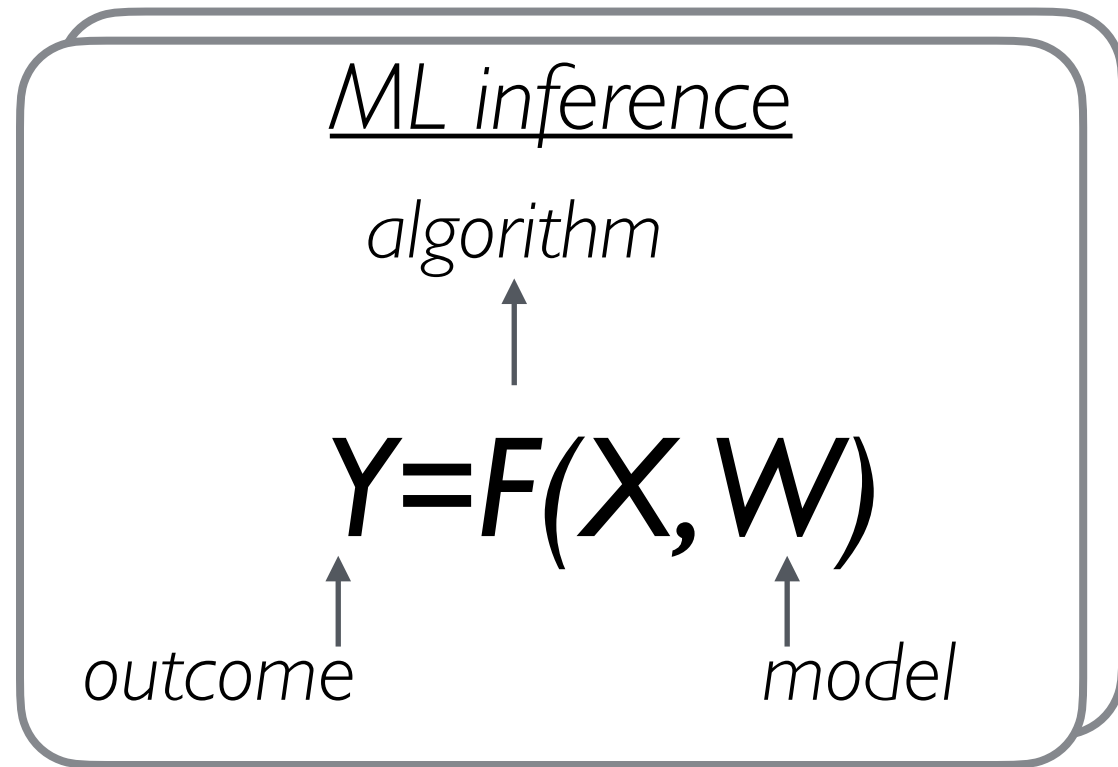
Efficiency challenges of ZKPs for ML

Efficient ZKPs for Neural Networks

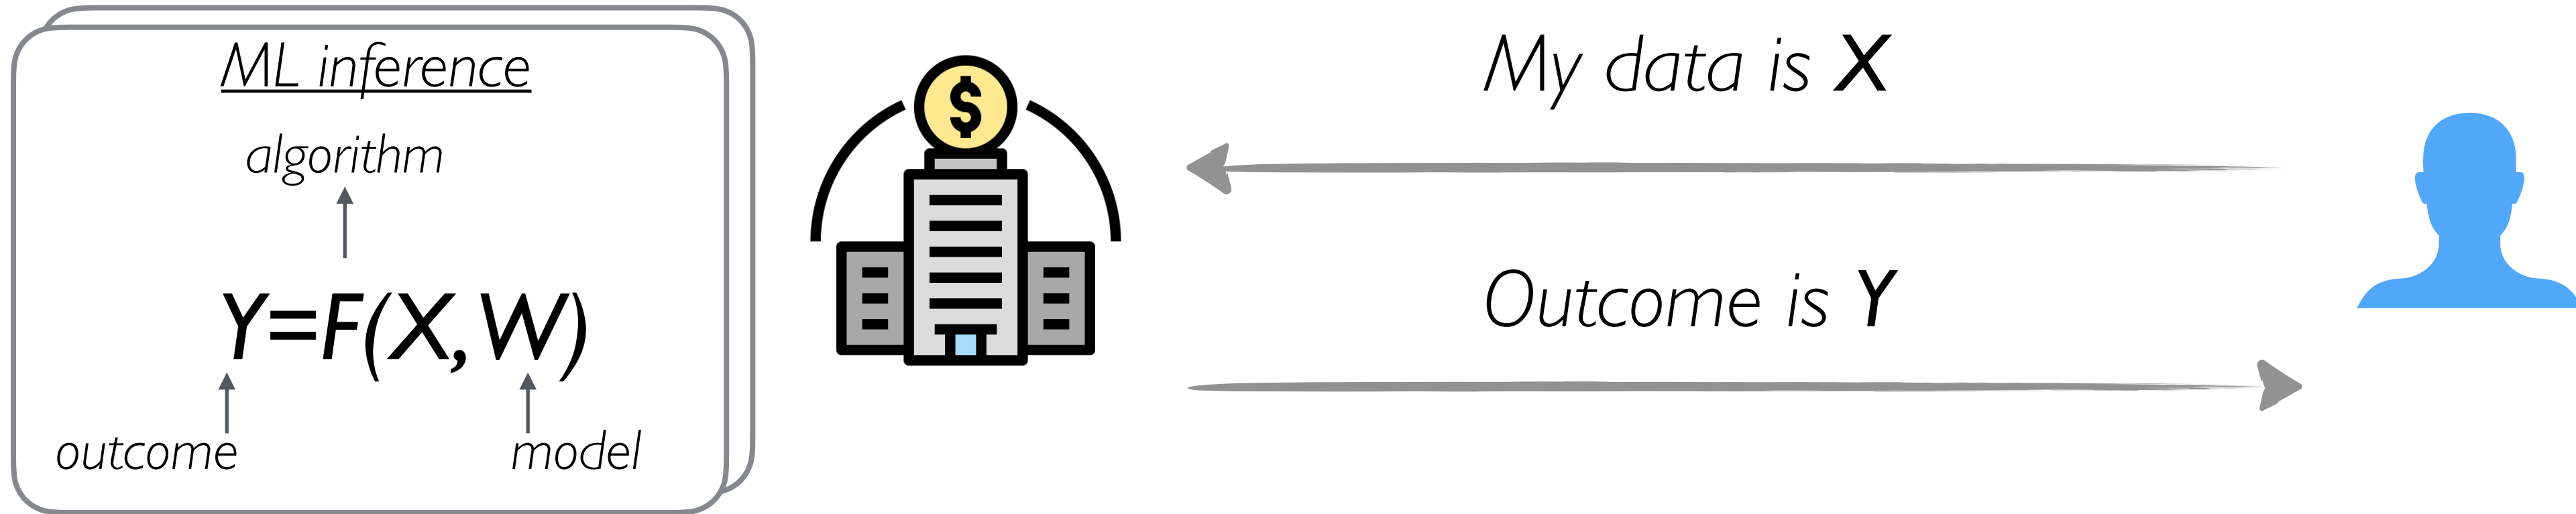
Efficient ZKPs for Decision Trees

Conclusions

Motivation: outsourcing machine learning



Motivation: outsourcing machine learning

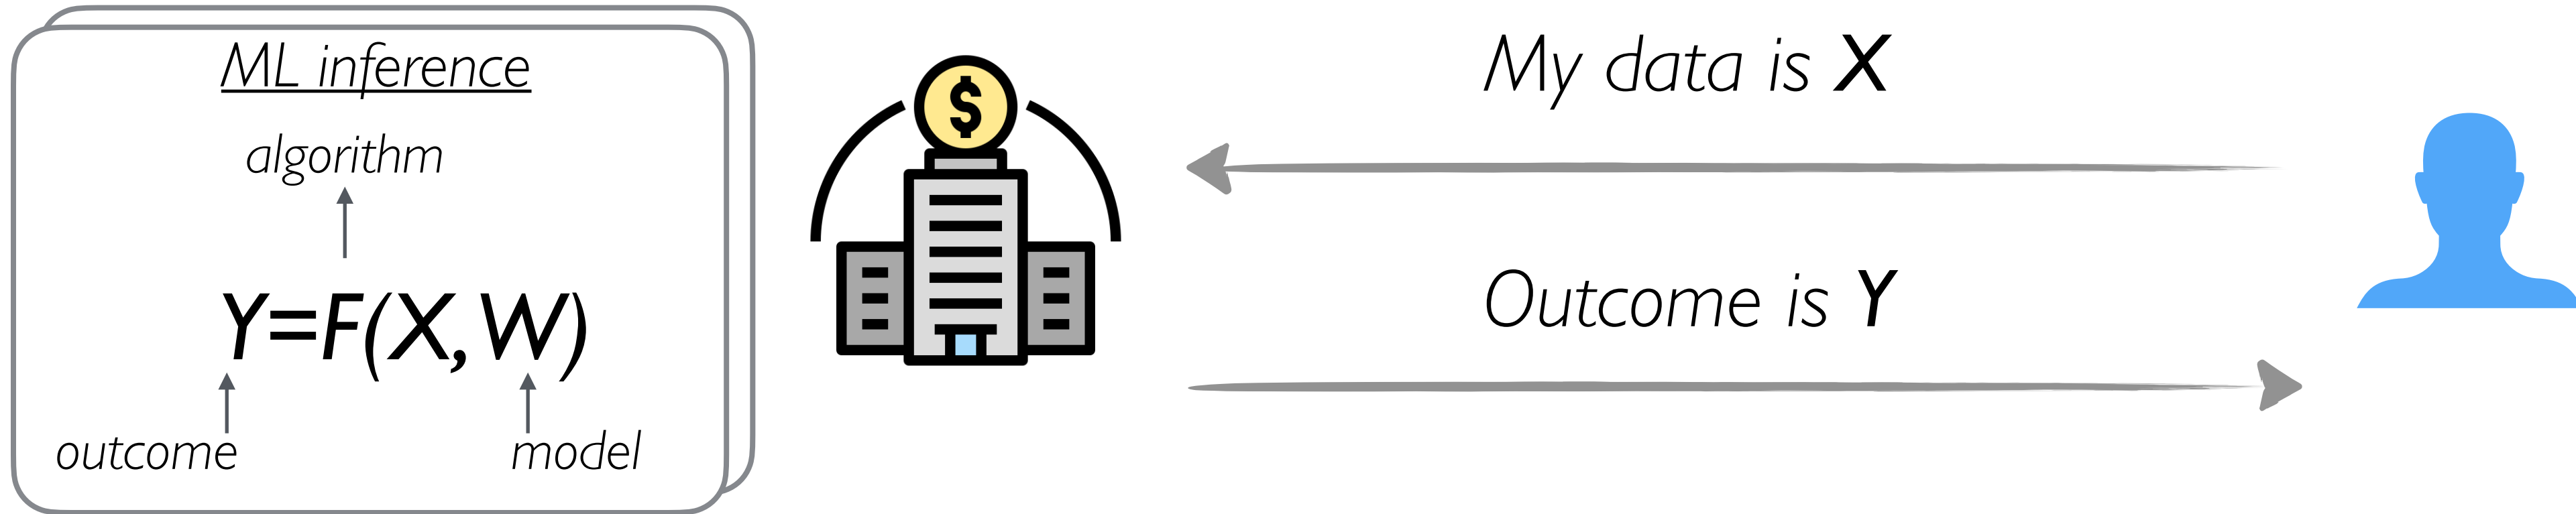


Banking & Finance



Can I have a loan?

Motivation: outsourcing machine learning

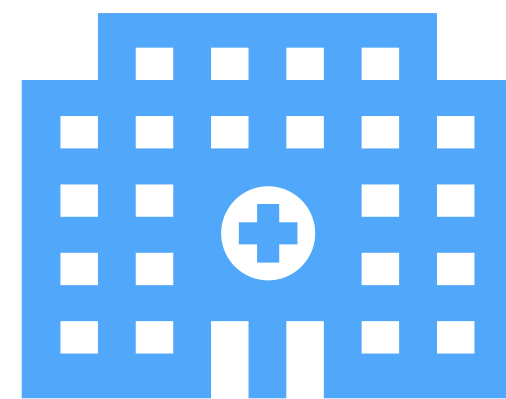


Banking & Finance



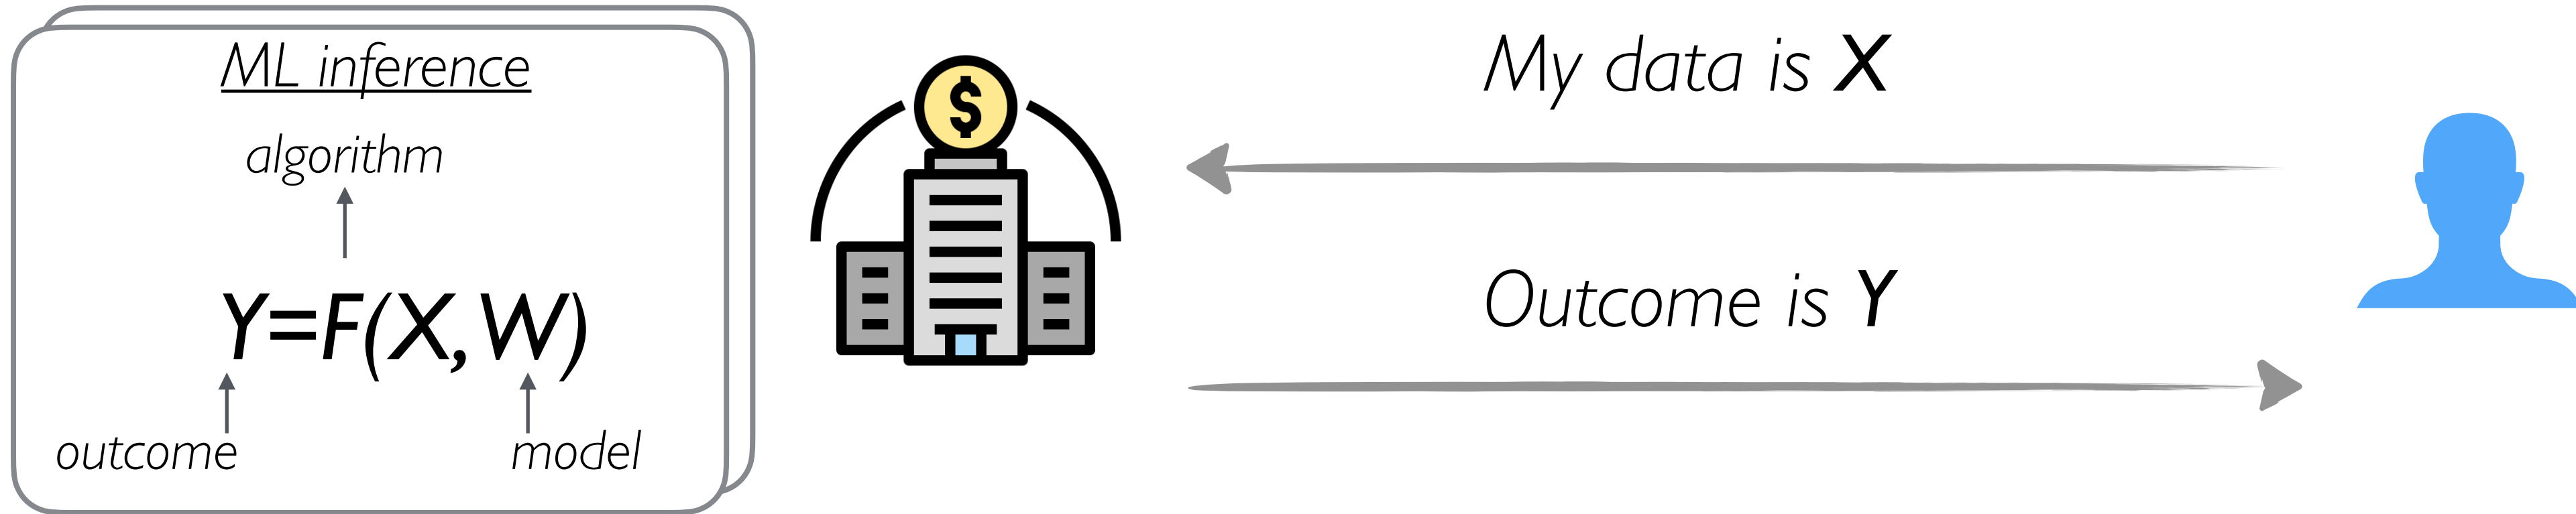
Can I have a loan?

Healthcare



Risk of a disease?

Motivation: outsourcing machine learning

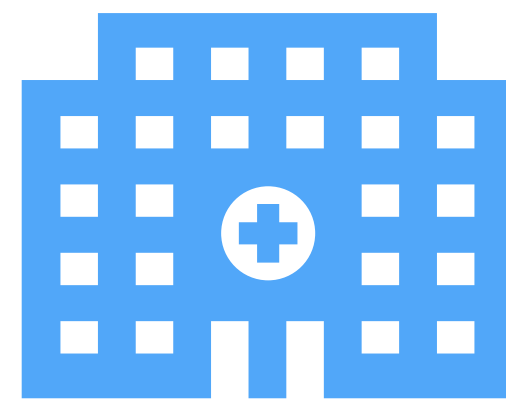


Banking & Finance



Can I have a loan?

Healthcare



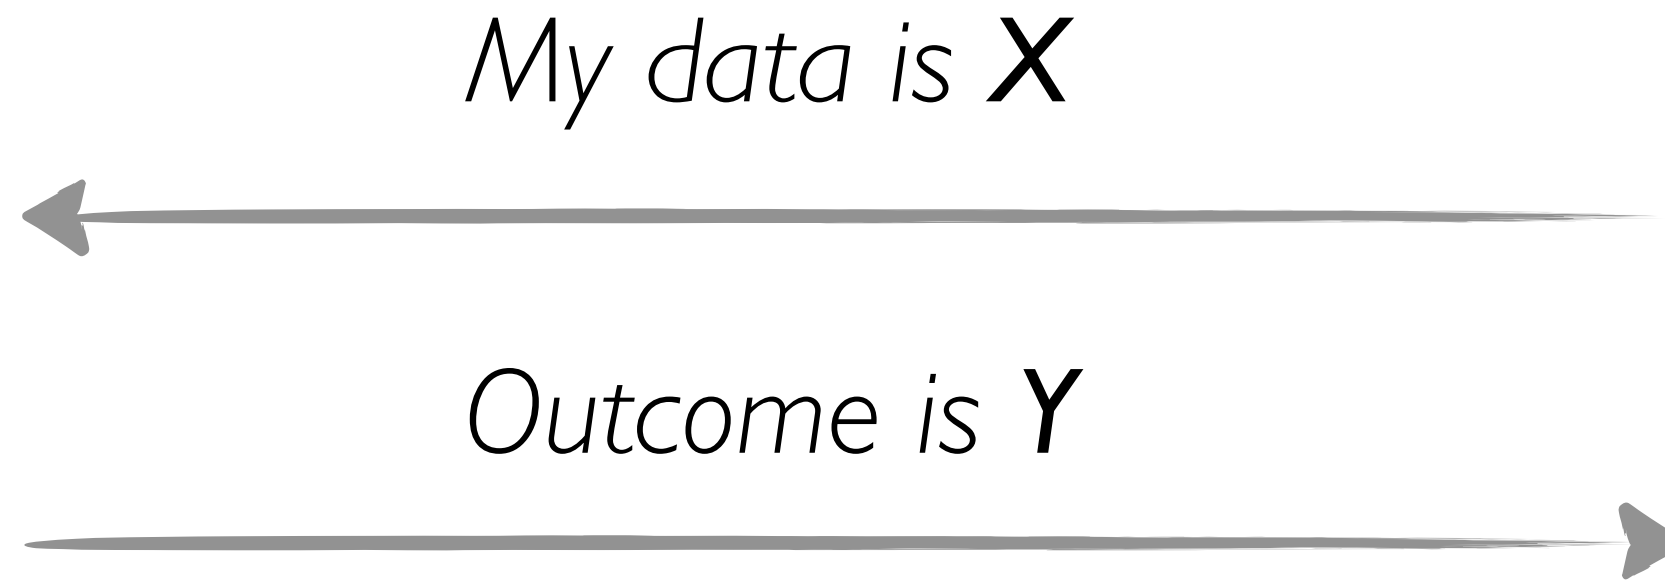
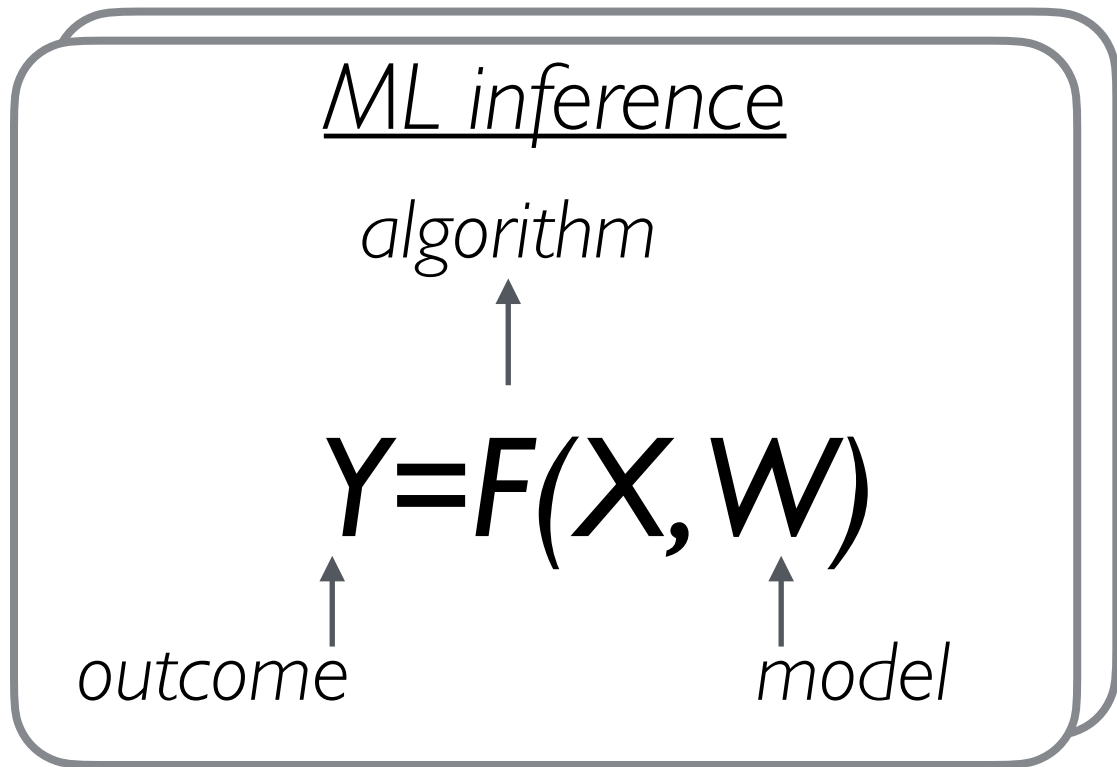
Risk of a disease?

Criminal justice



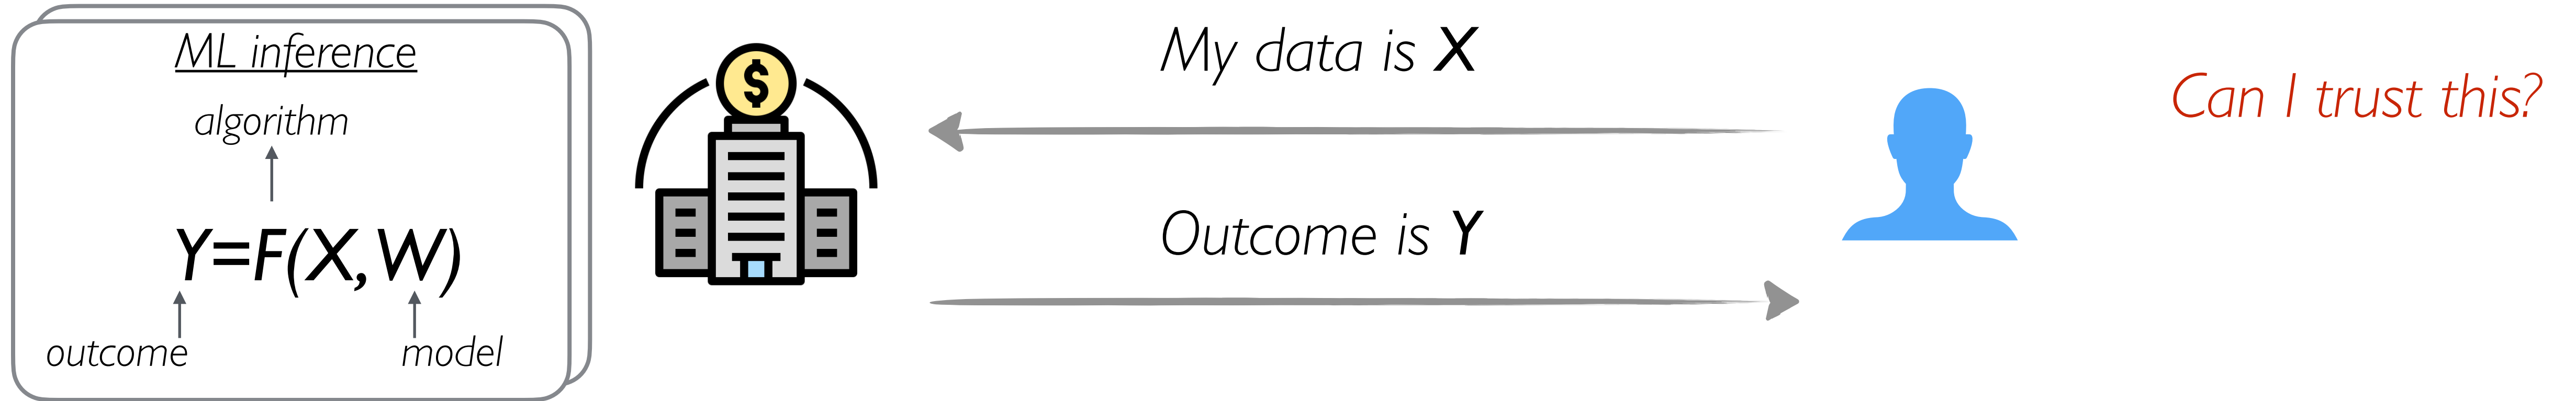
Released or retained?

Security of outsourced machine learning



Can I trust this?

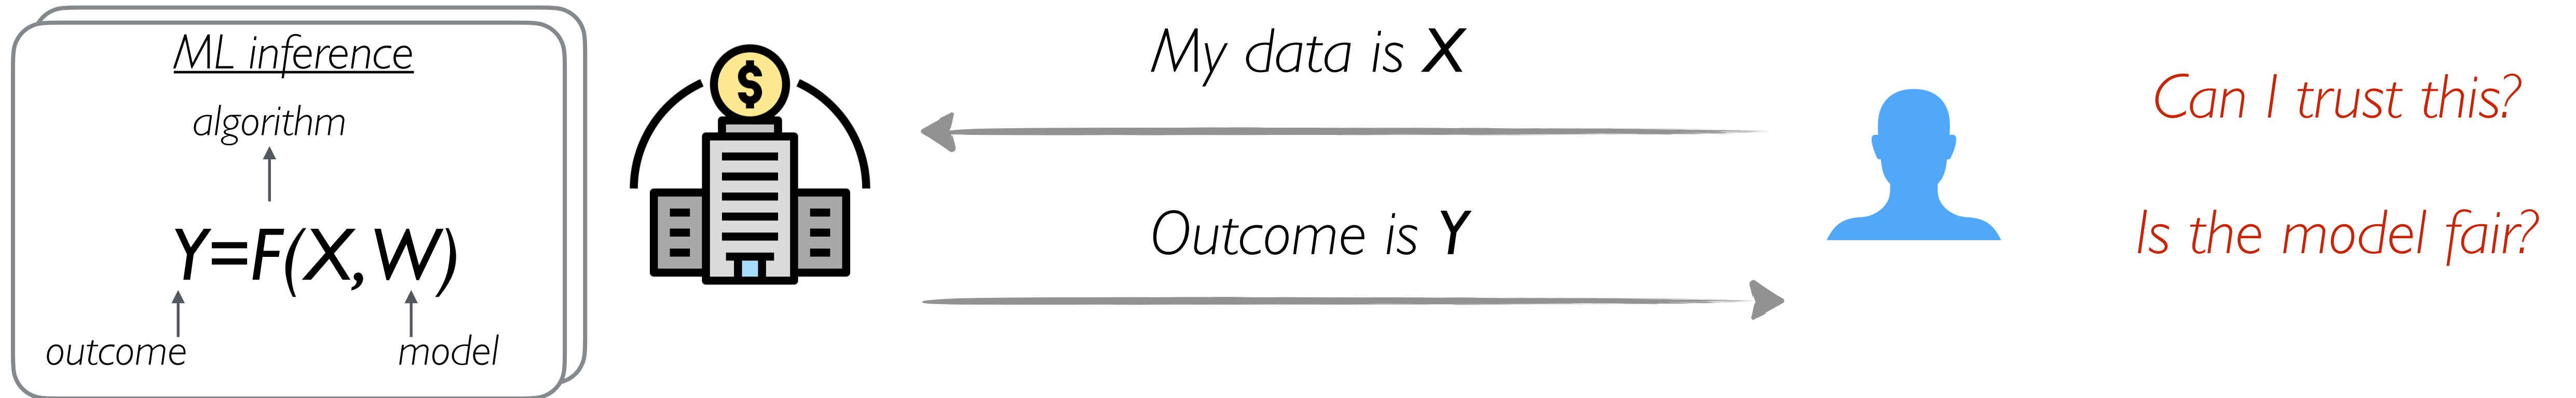
Security of outsourced machine learning



Goals

Integrity: detect tampered computations

Security of outsourced machine learning

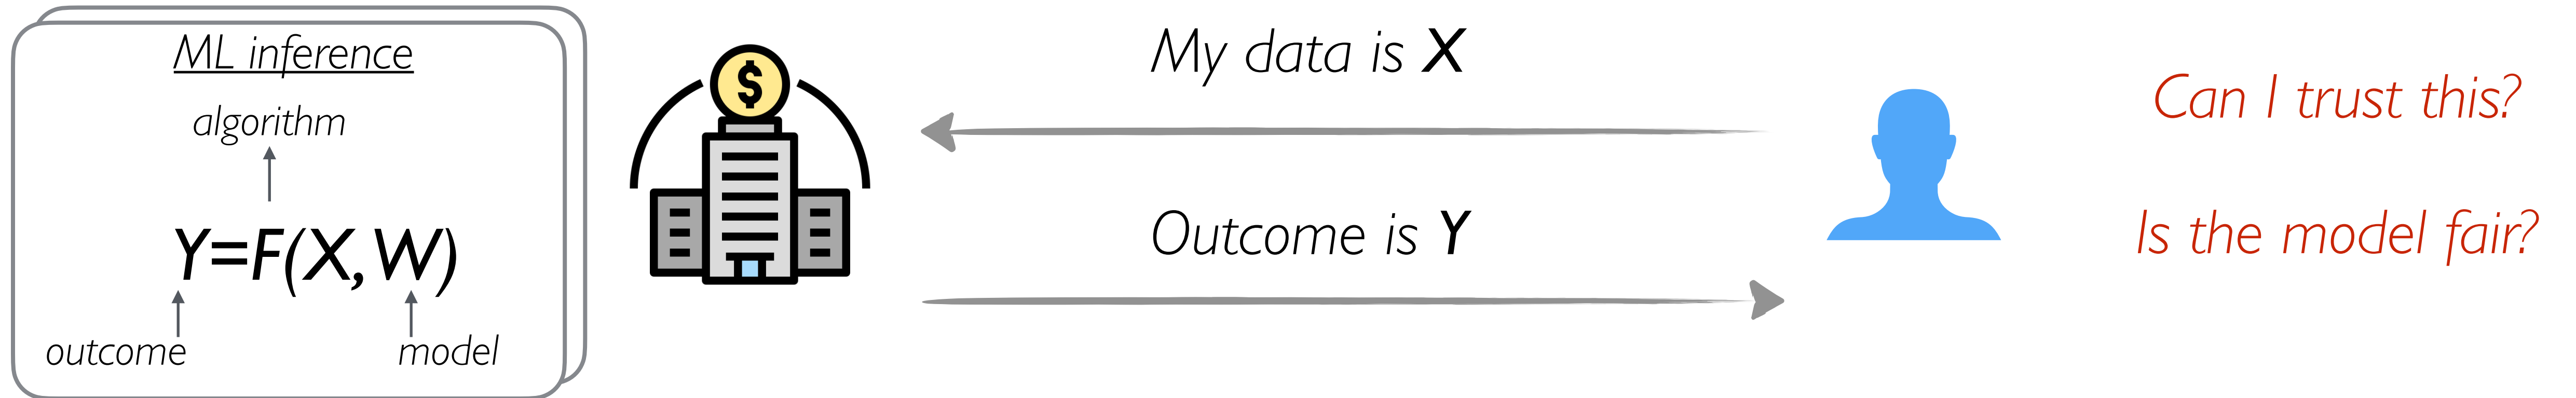


Goals

Integrity: detect tampered computations

Fairness: assuming the model is trusted, decision process is the same for all clients

Security of outsourced machine learning



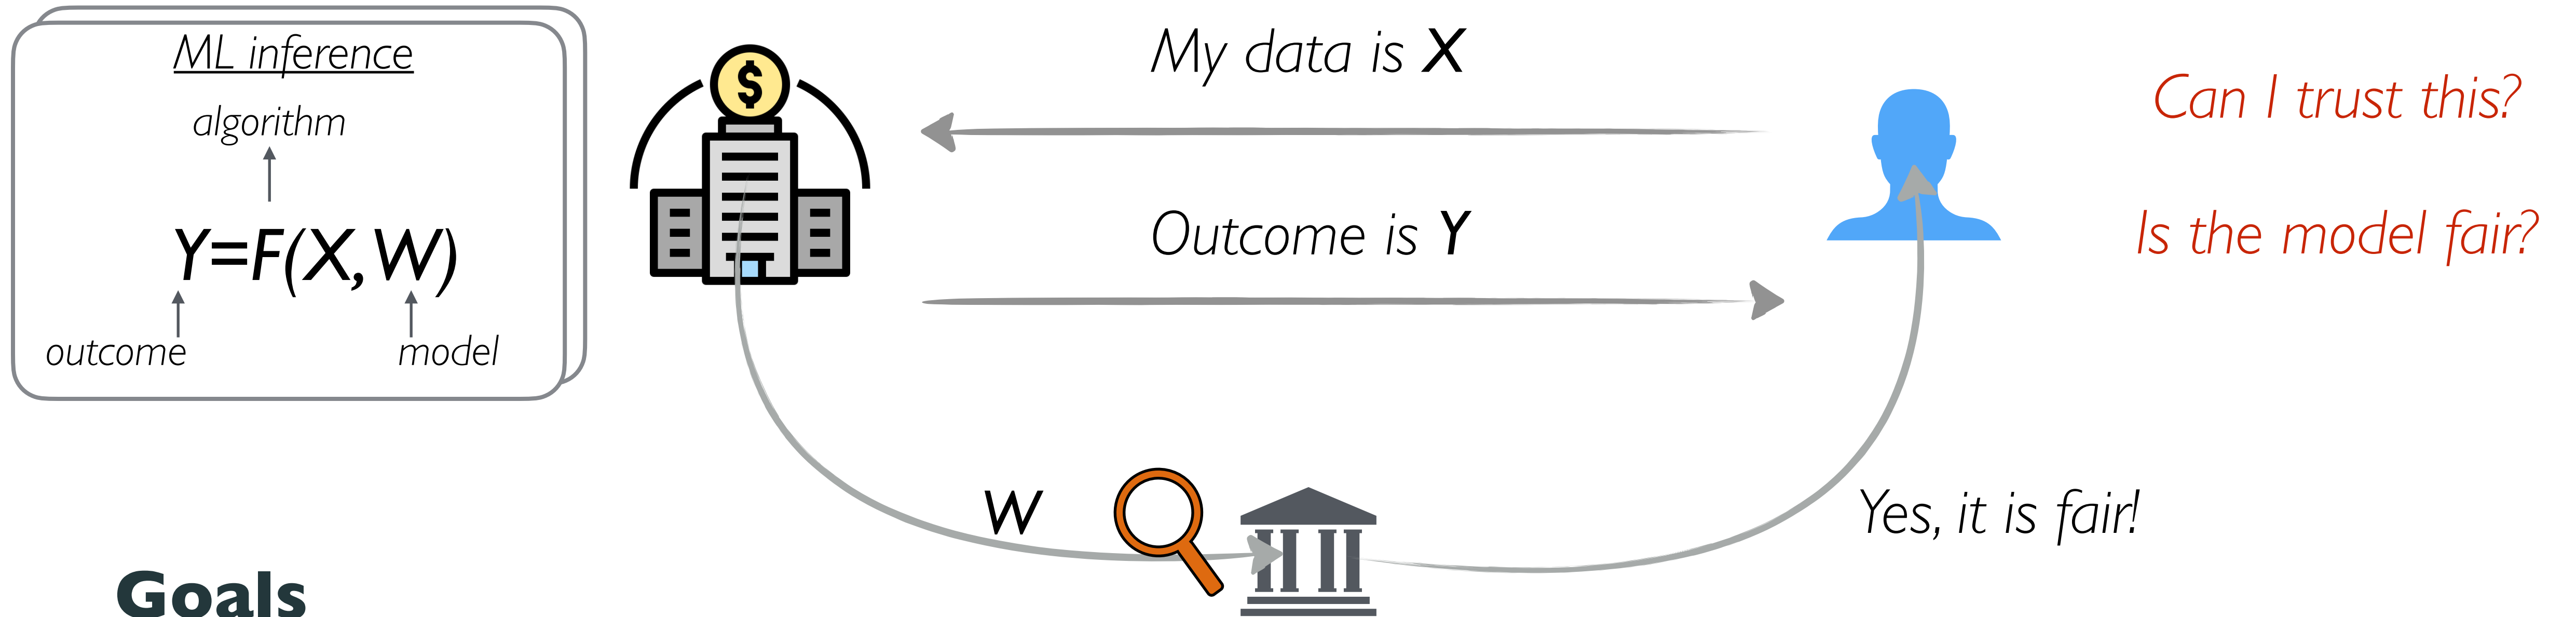
Goals

Integrity: detect tampered computations

Fairness: assuming the model is trusted, decision process is the same for all clients

Privacy: clients should not learn anything about the model W

Security of outsourced machine learning



Goals

Integrity: detect tampered computations

Fairness: assuming the model is trusted, decision process is the same for all clients

Privacy: clients should not learn anything about the model W

Secure and Private ML Inference



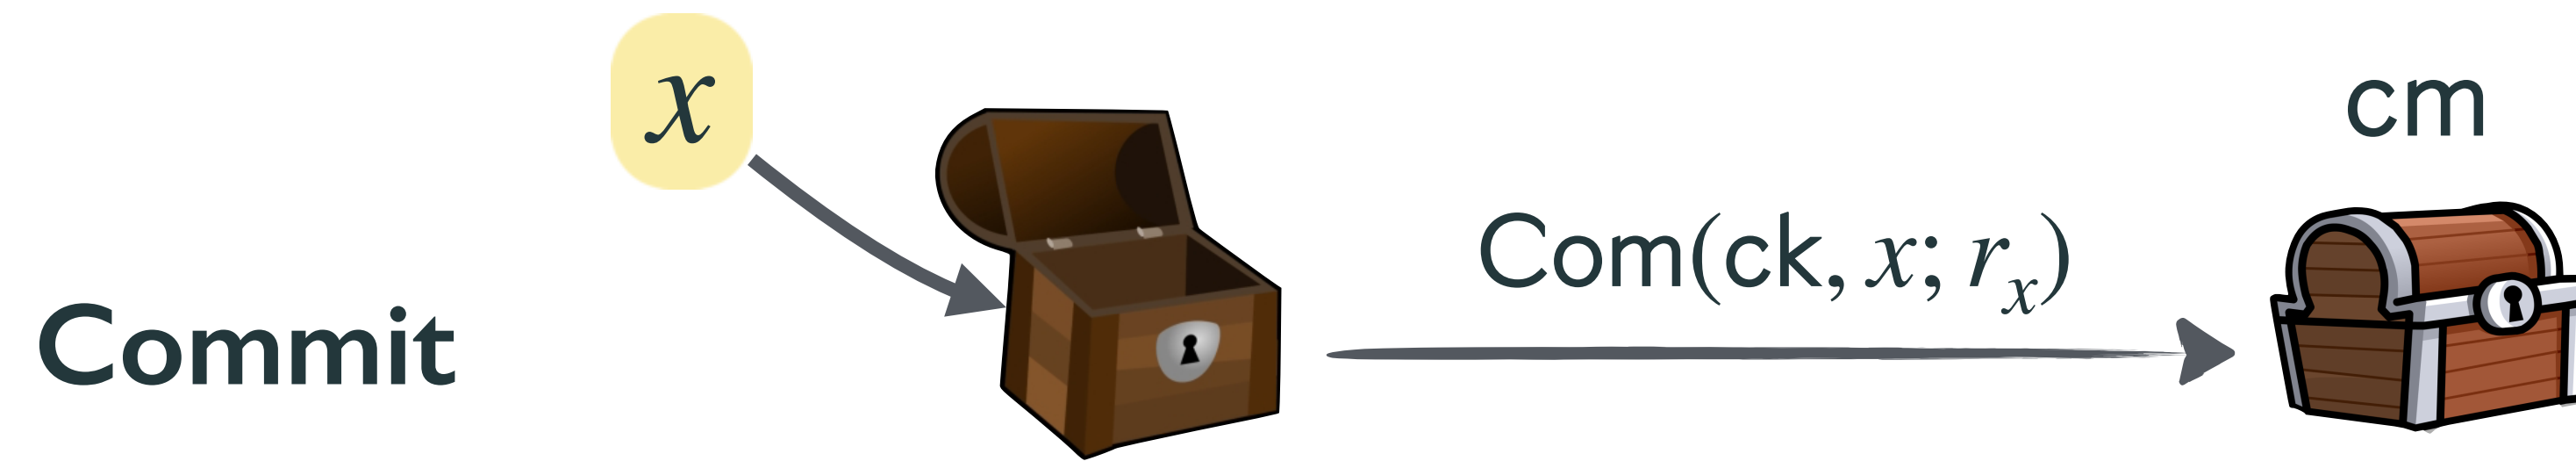
Ingredients

Zero Knowledge Proofs

Commitments

Commitments

[Blum, Even '81]



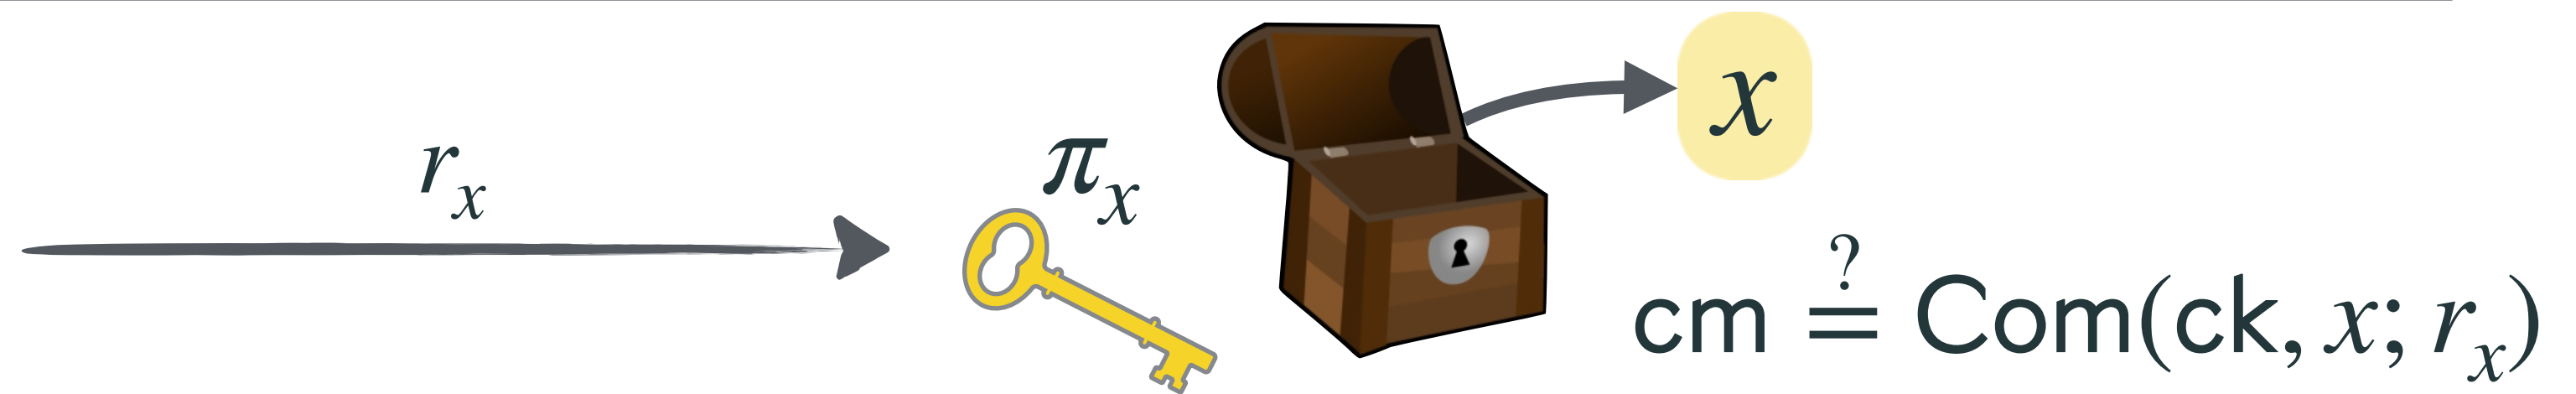
Commitments

[Blum, Even '81]

Commit

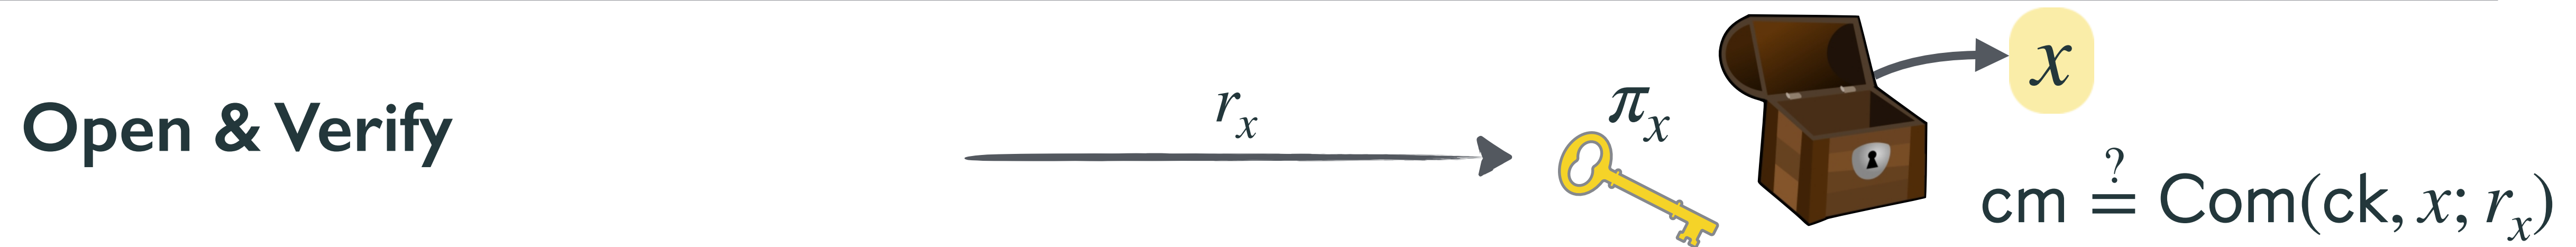
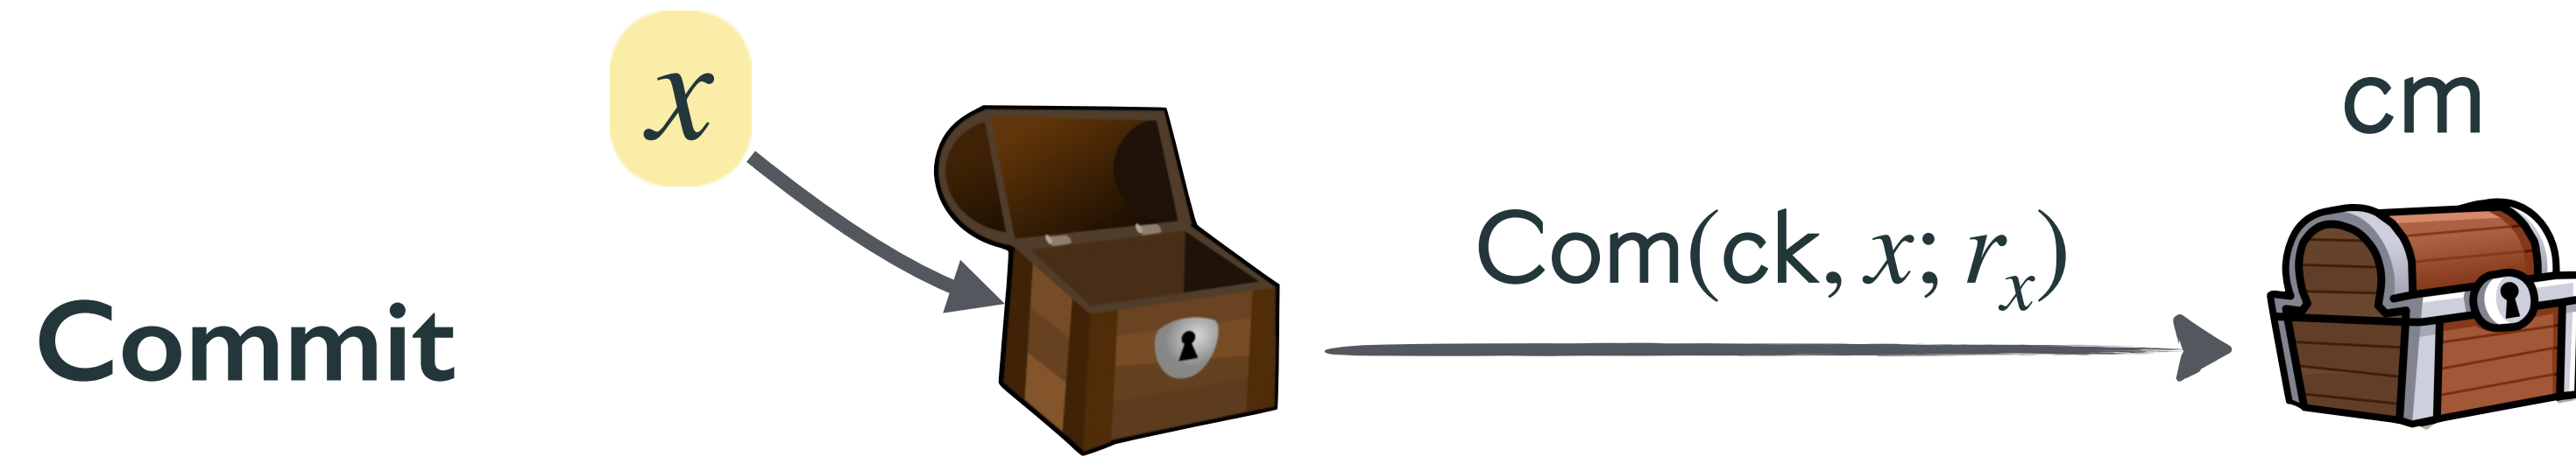


Open & Verify



Commitments

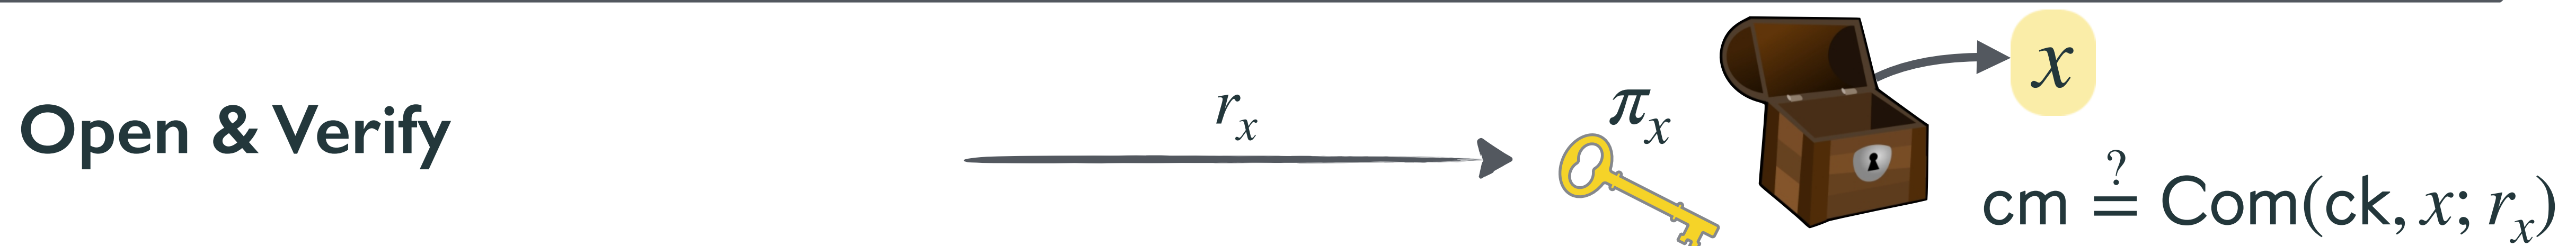
[Blum, Even '81]



- **Hiding:** $\text{Com}(x) \approx \text{Com}(x')$

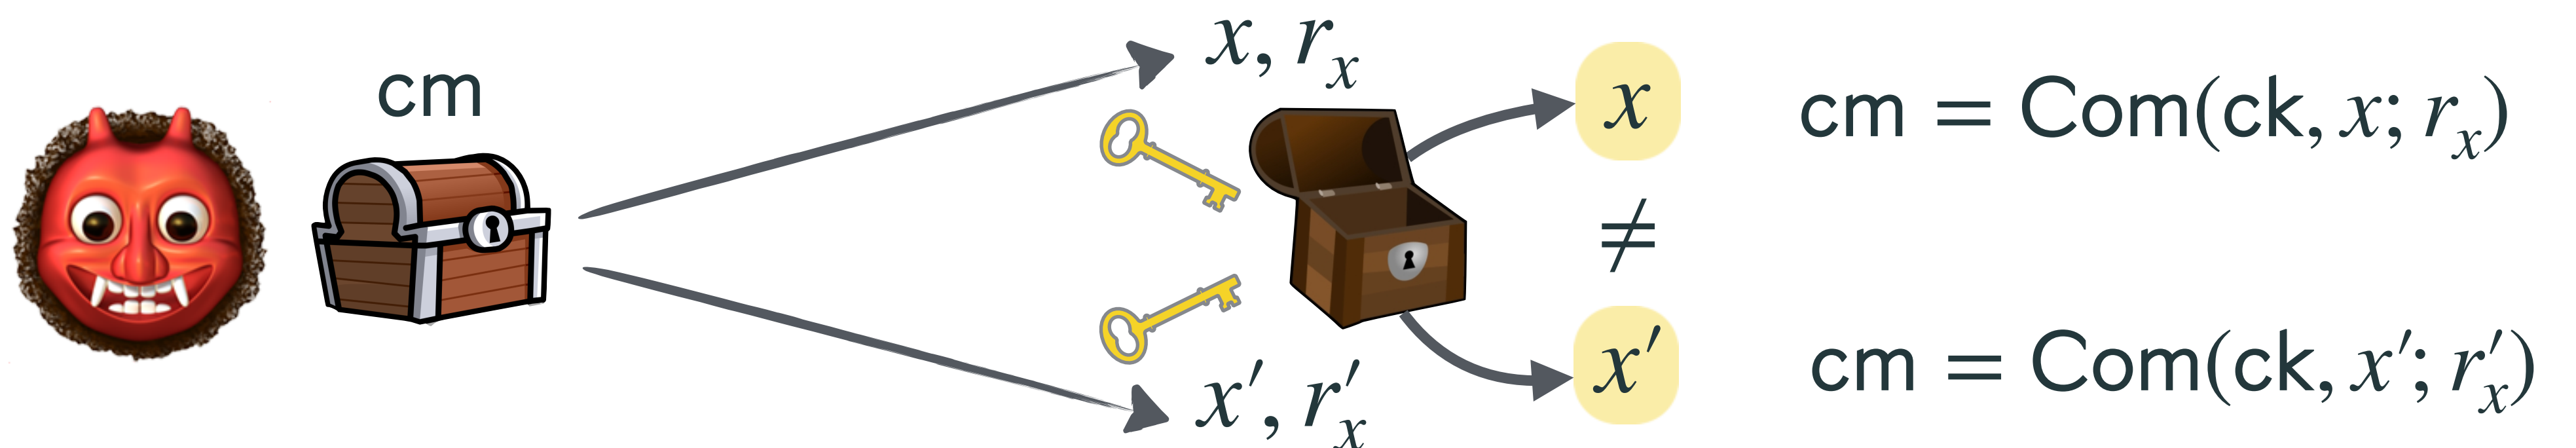
Commitments

[Blum, Even '81]

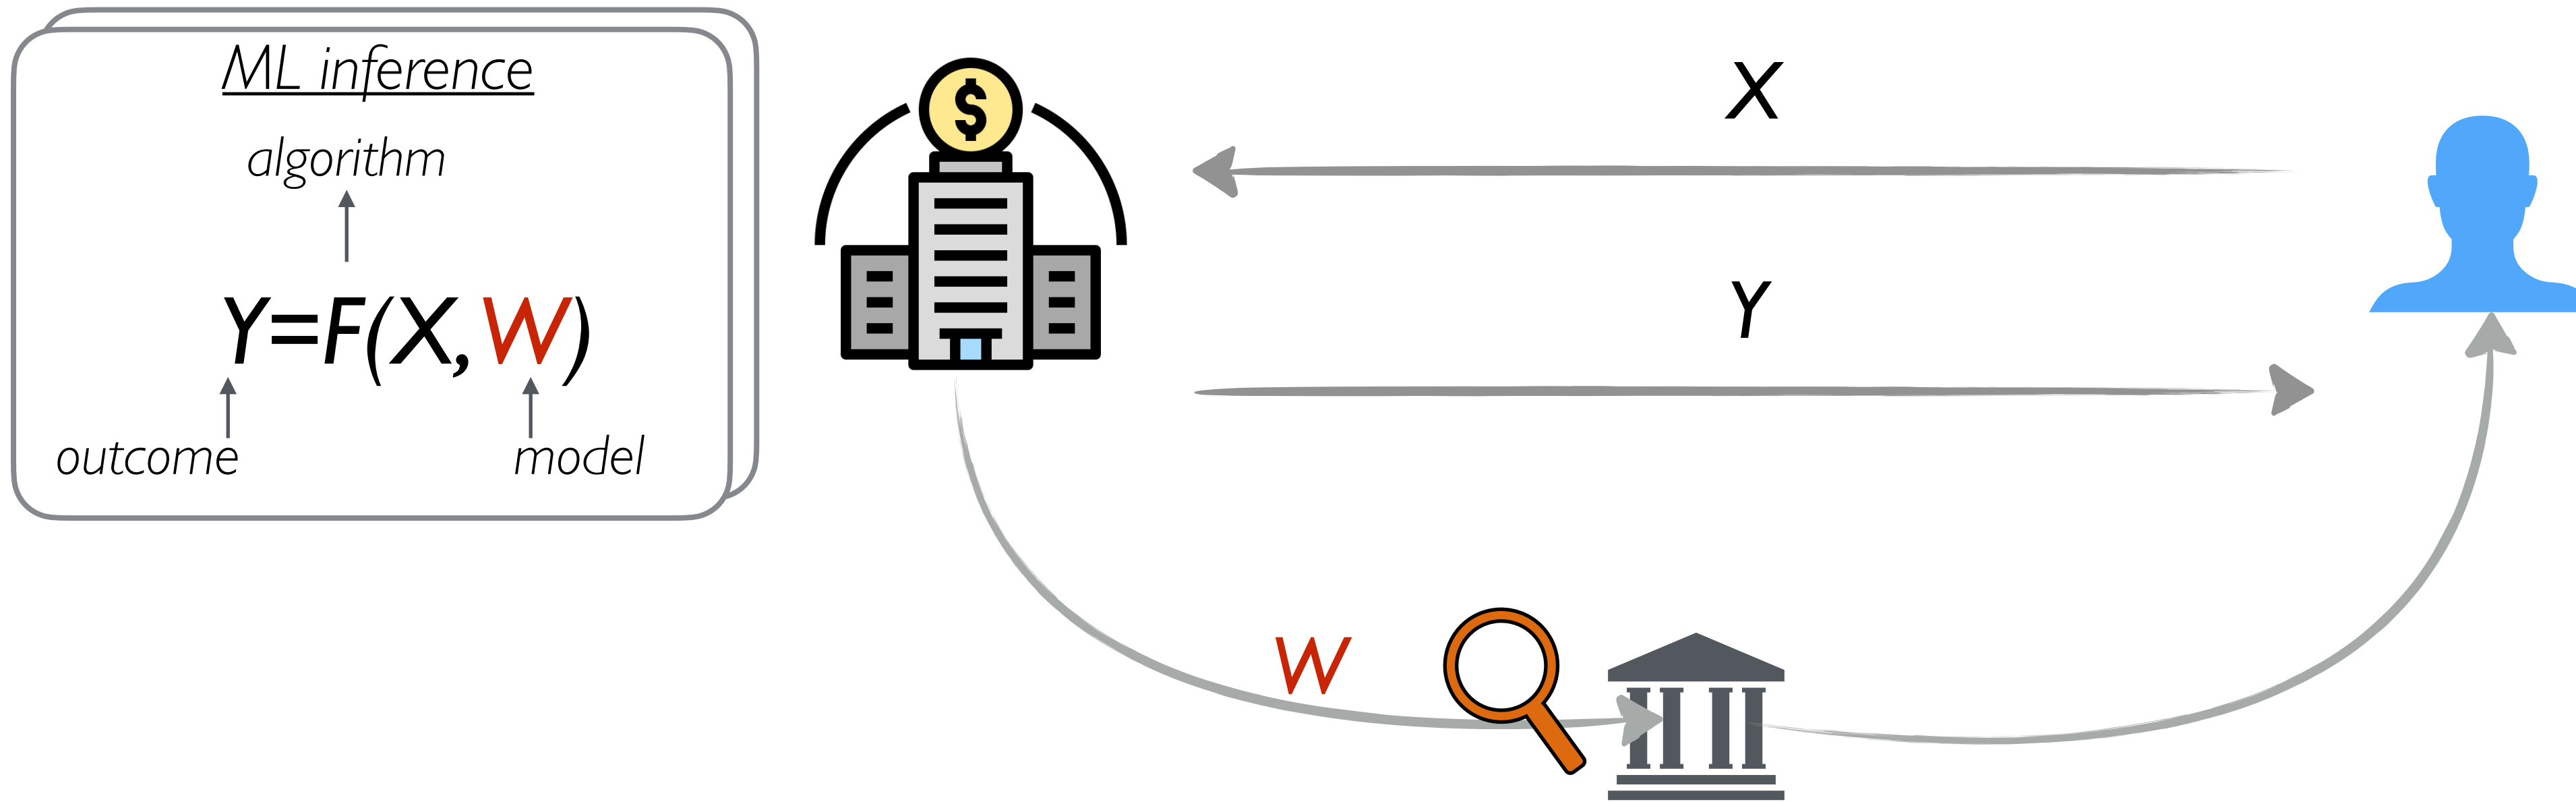


• **Hiding:** $\text{Com}(x) \approx \text{Com}(x')$

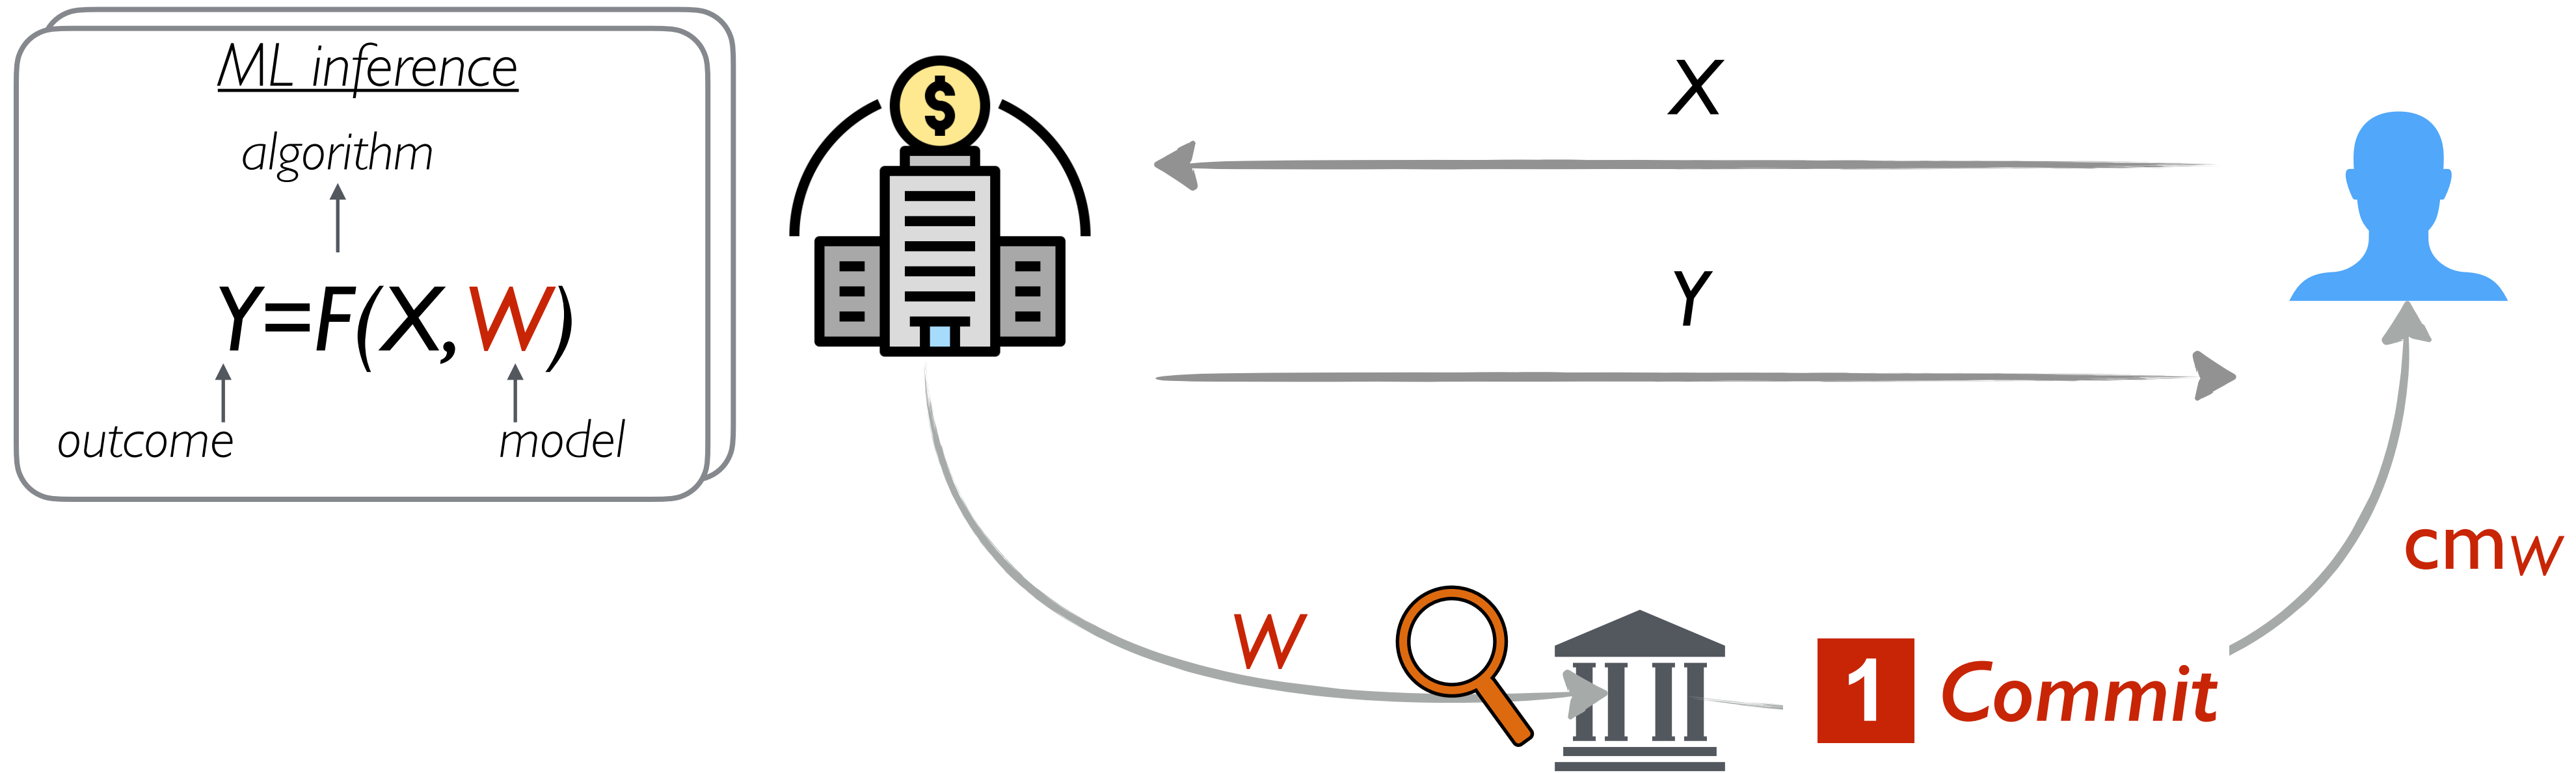
• **Binding:**



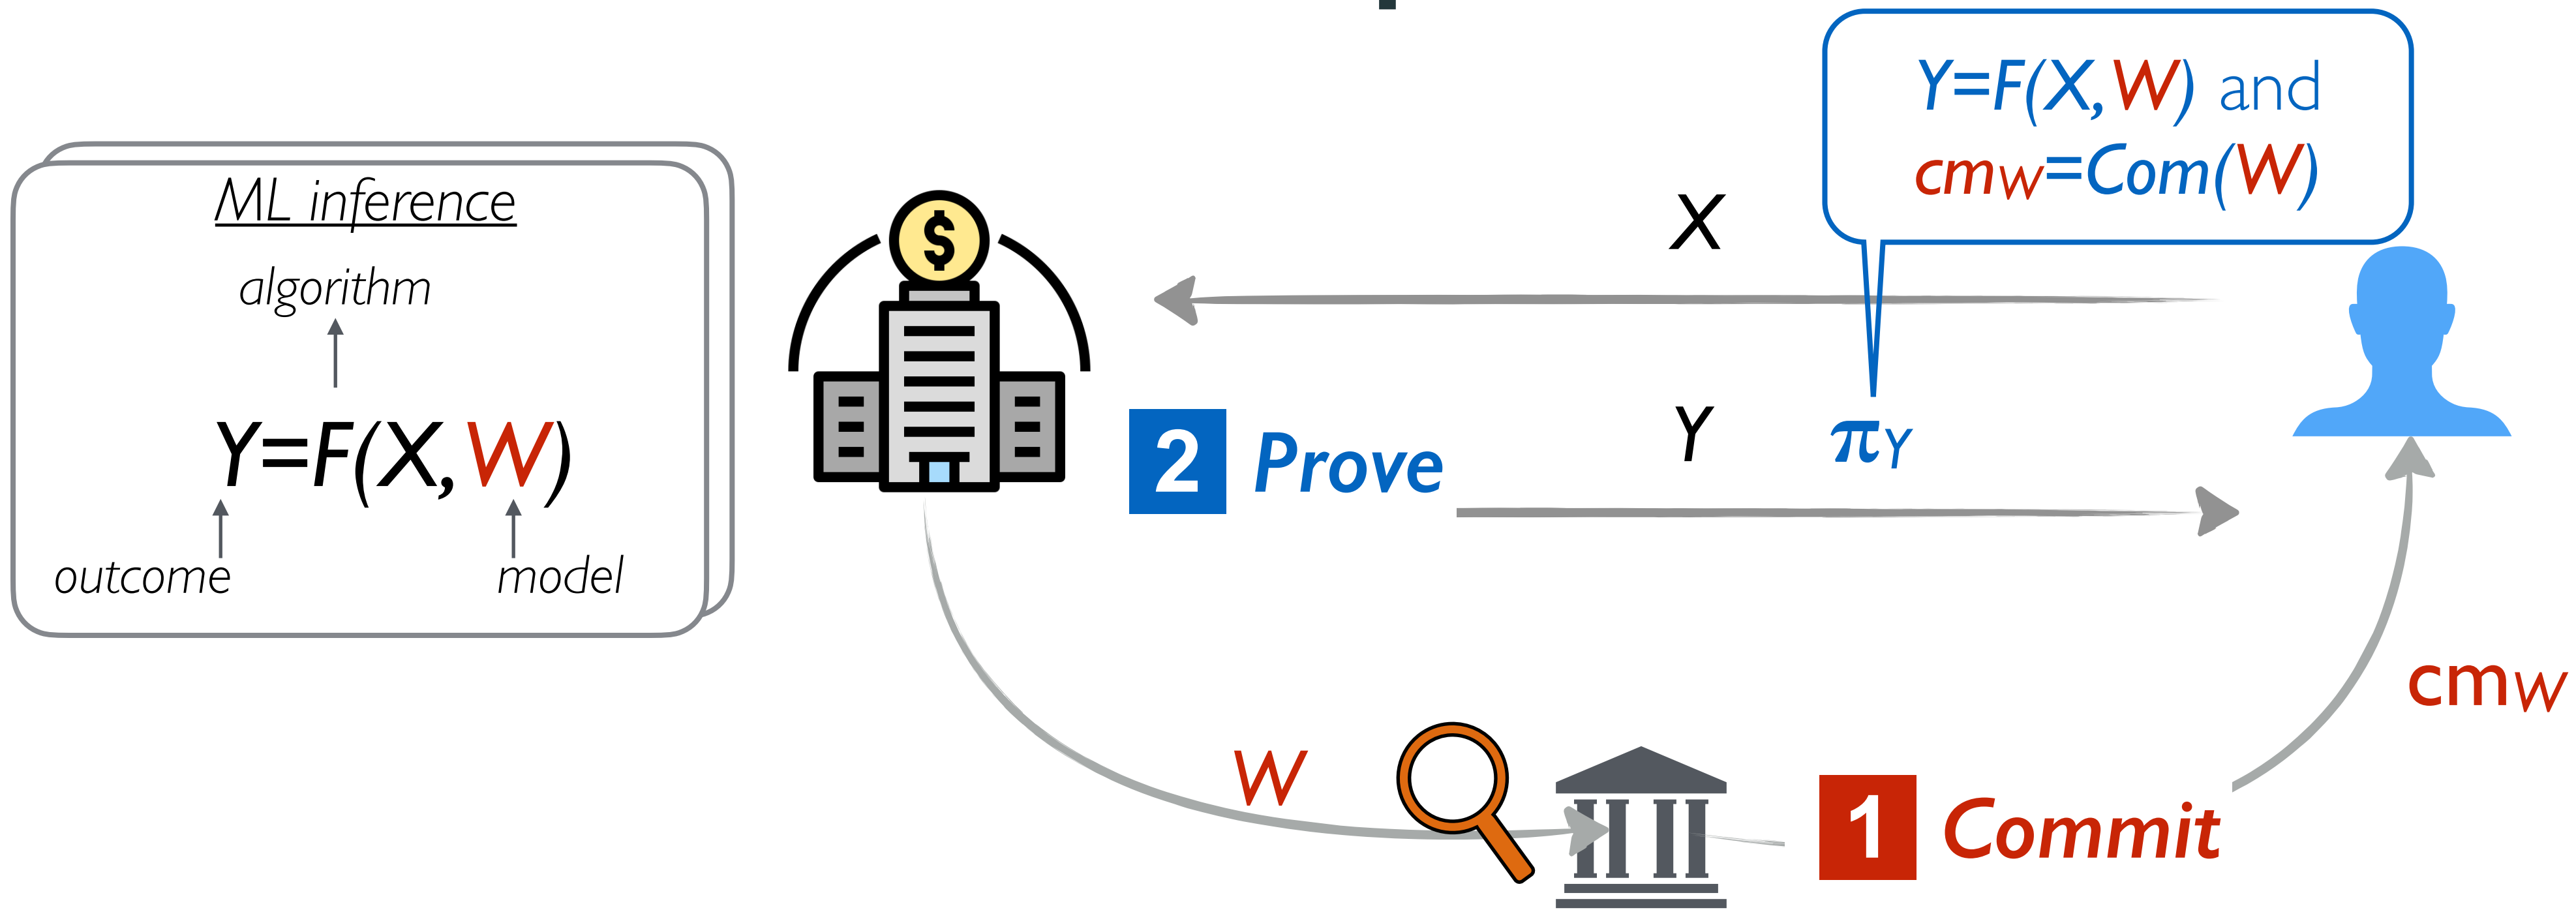
ZKPs for secure and private ML inference



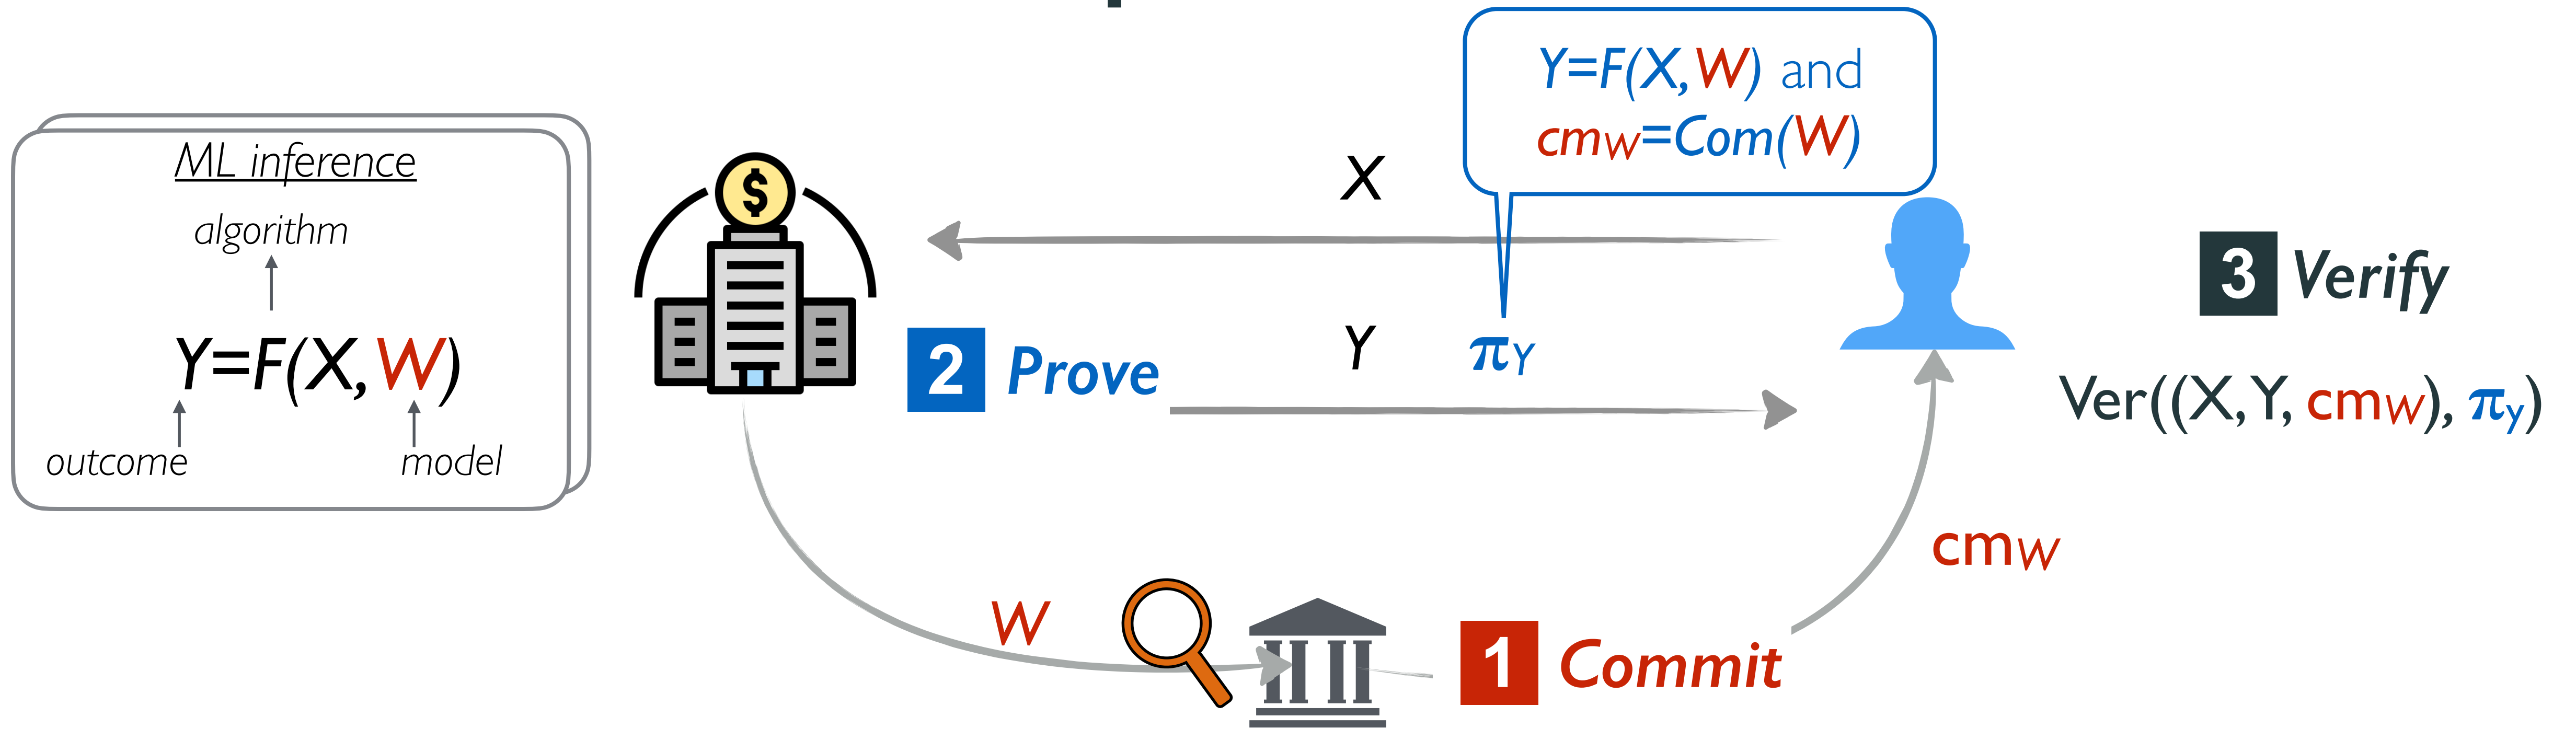
ZKPs for secure and private ML inference



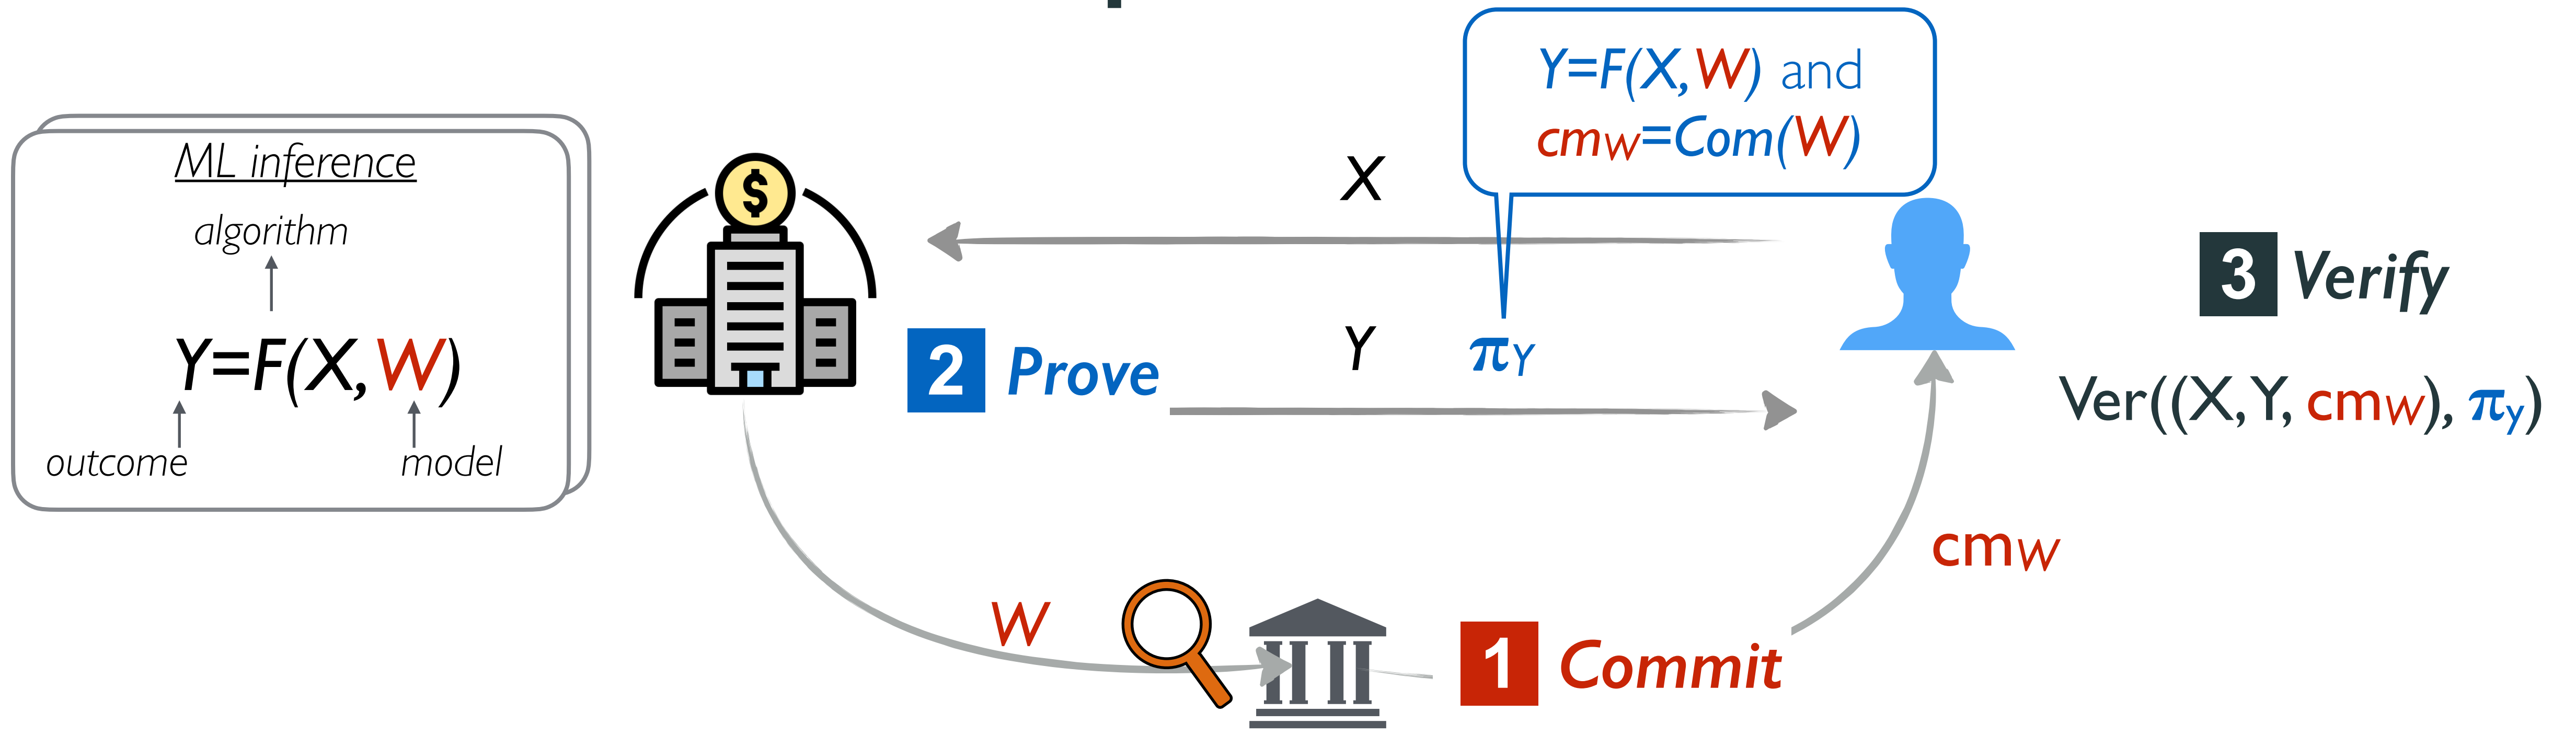
ZKPs for secure and private ML inference



ZKPs for secure and private ML inference

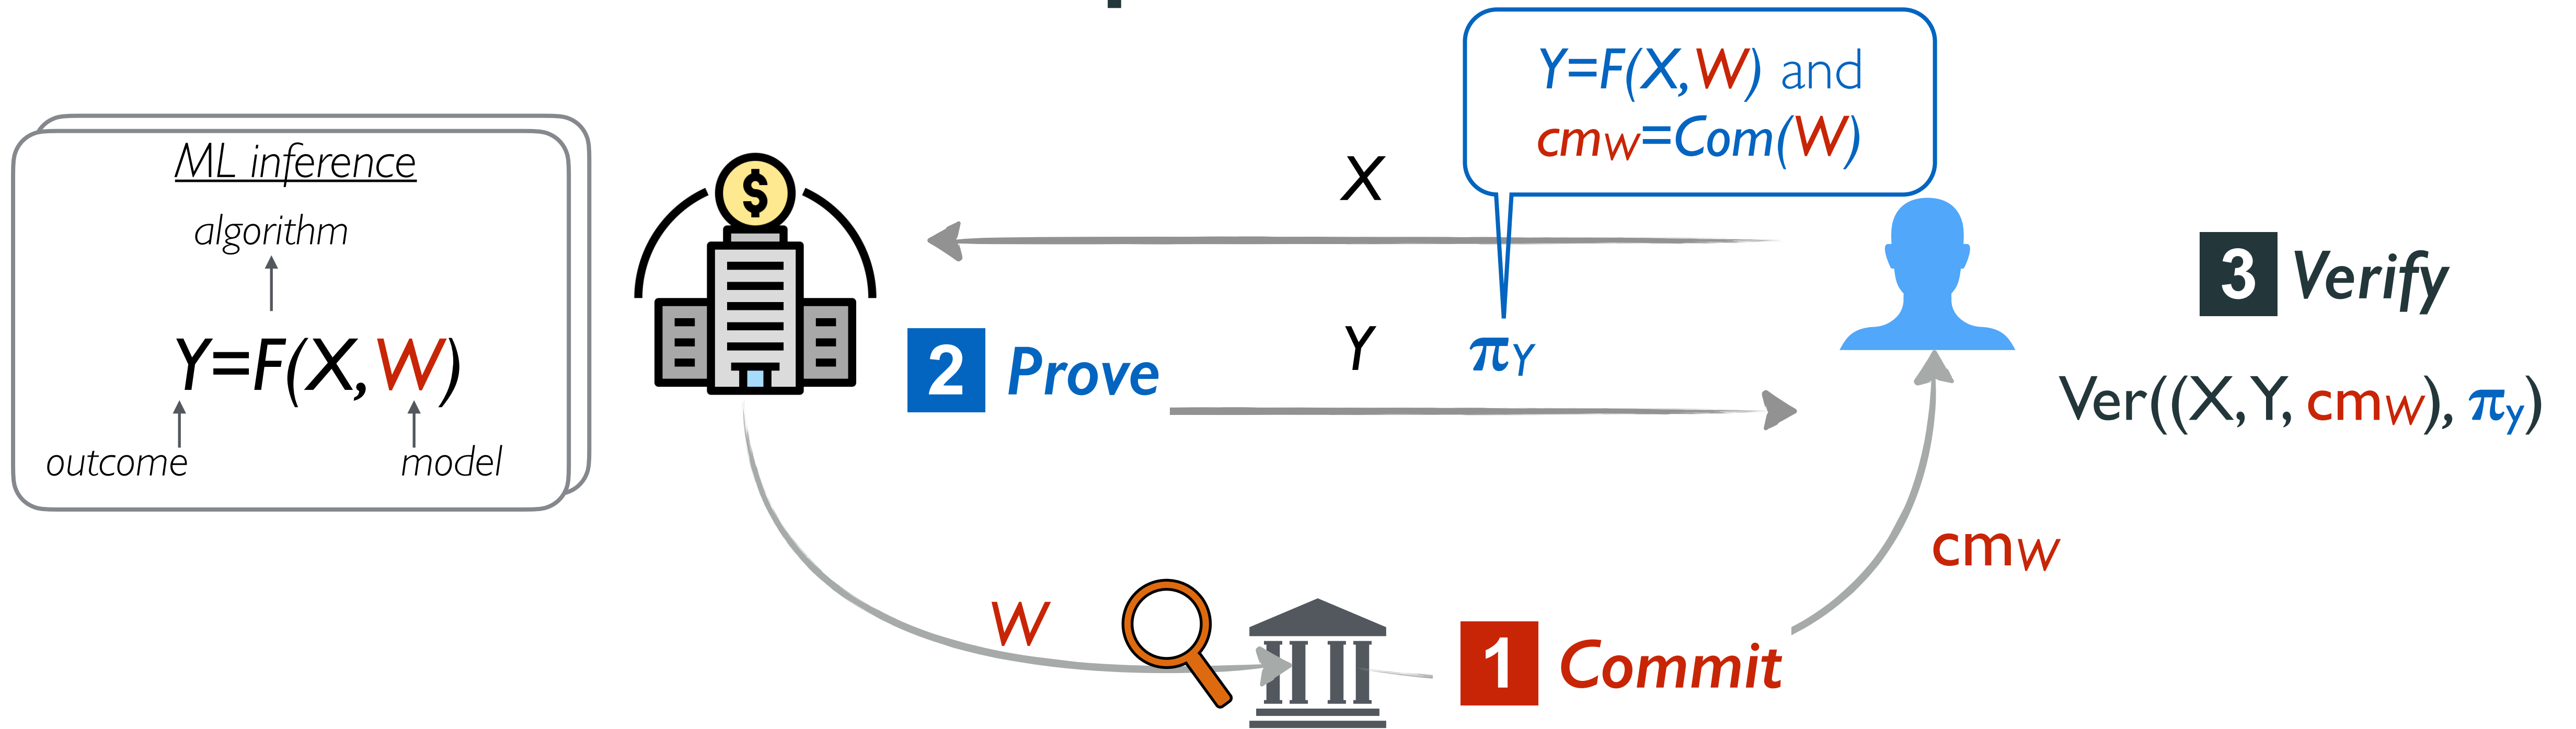


ZKPs for secure and private ML inference



Integrity: detect tampered computations ← ZKP Soundness + Com binding

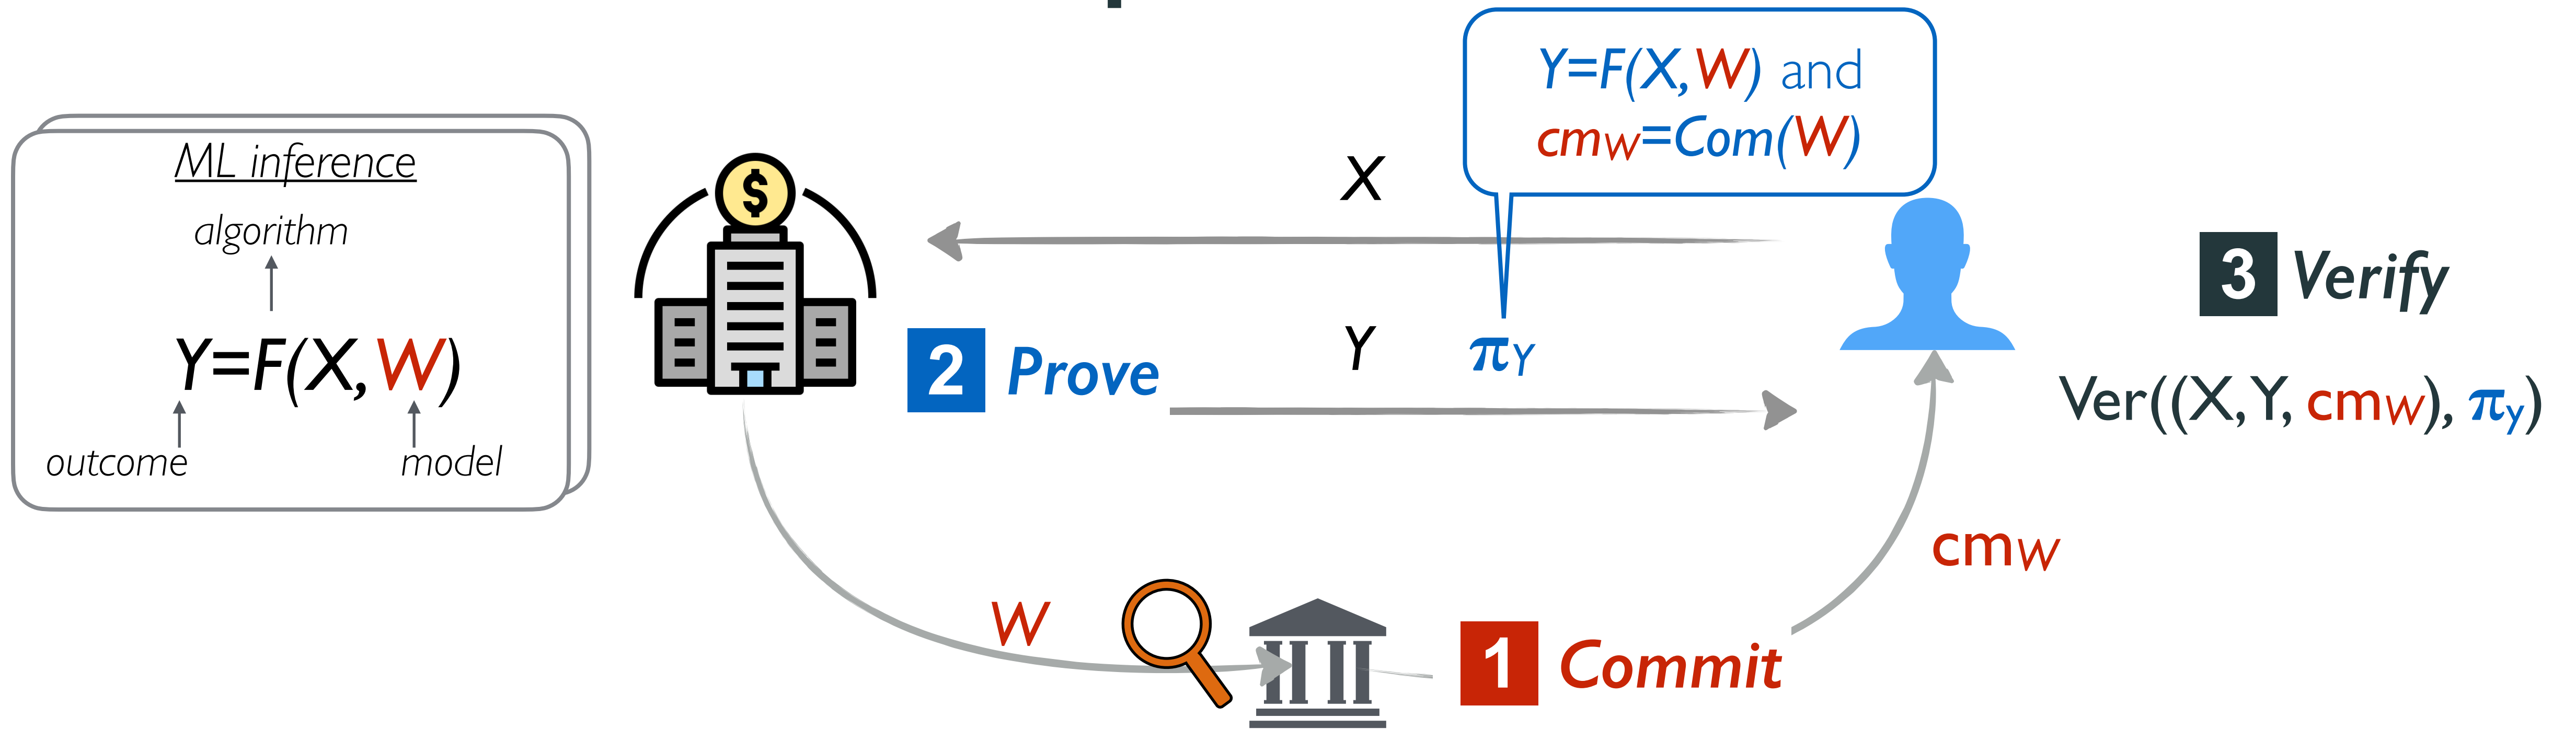
ZKPs for secure and private ML inference



Integrity: detect tampered computations ← ZKP Soundness + Com binding

Privacy: clients should not learn anything about the model W ← ZKP ZK + Com hiding

ZKPs for secure and private ML inference

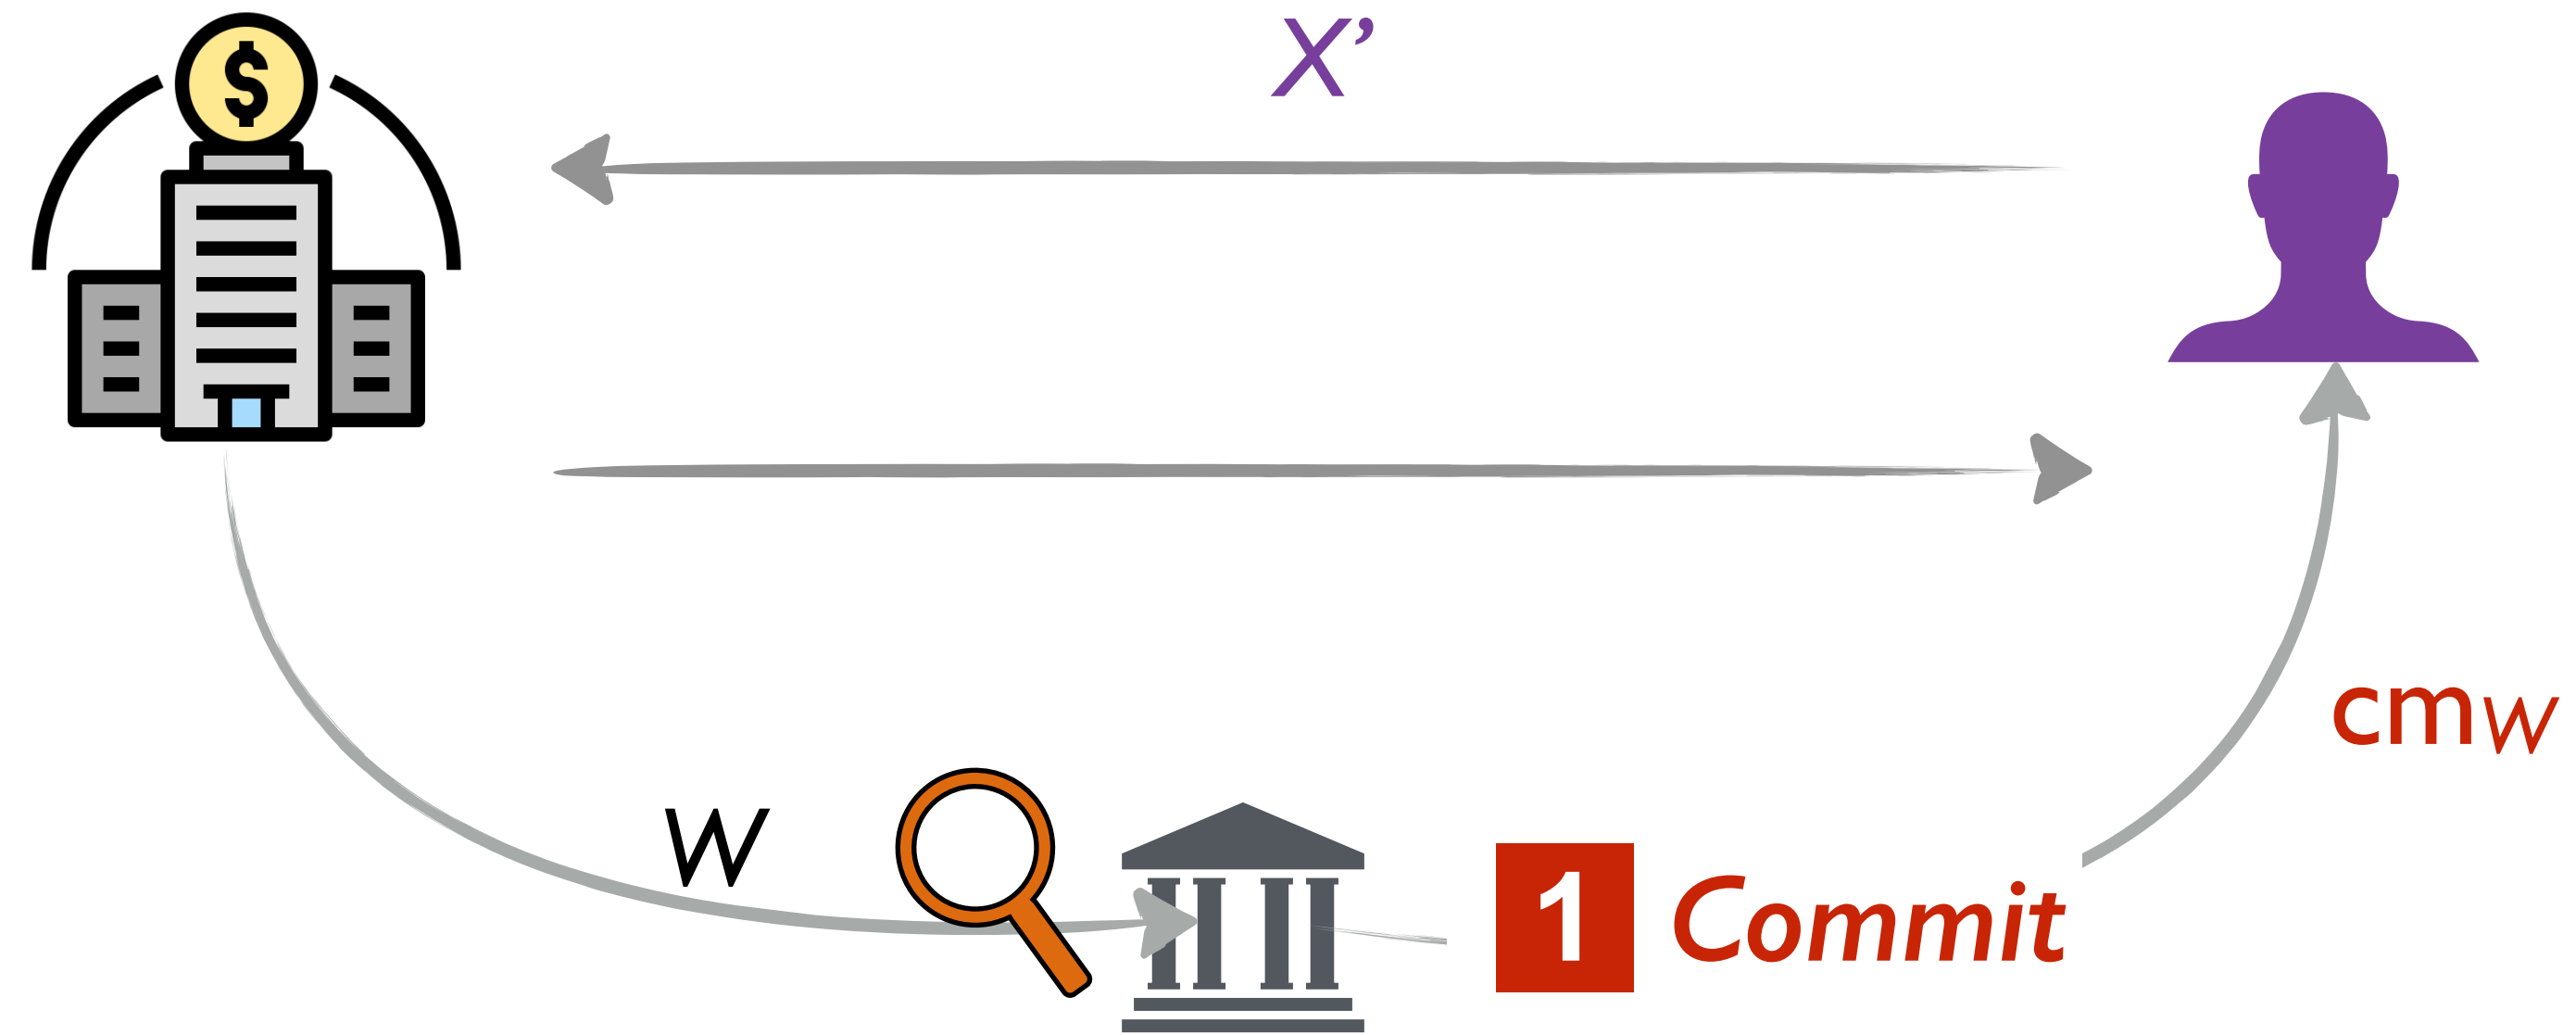


Integrity: detect tampered computations \leftarrow ZKP Soundness + Com binding

Privacy: clients should not learn anything about the model W \leftarrow ZKP ZK + Com hiding

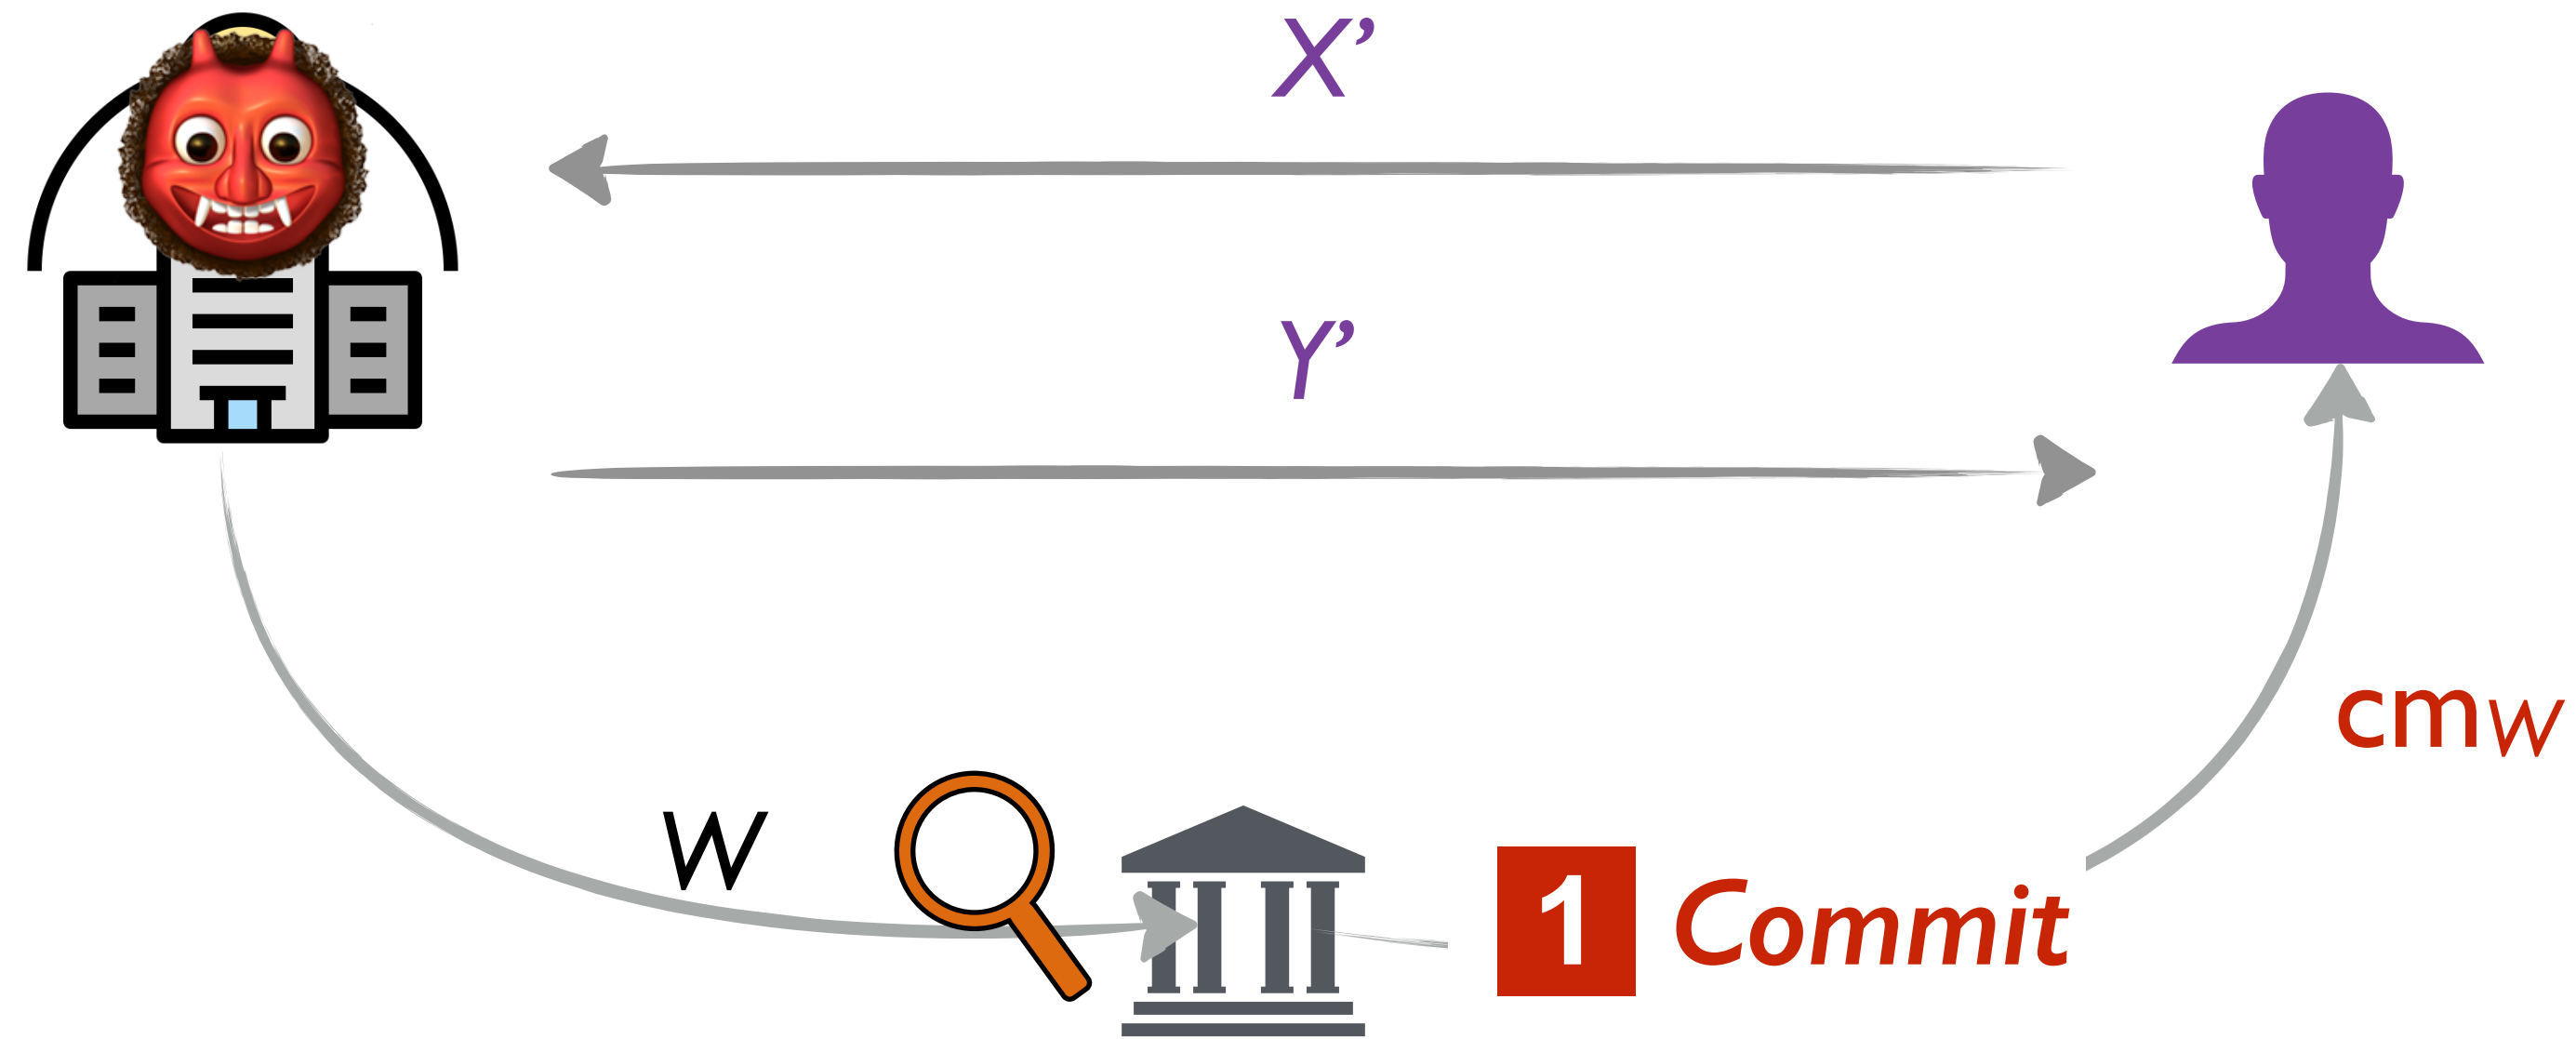
Fairness: decision process the same for all clients...

ZKPs for secure and private ML inference

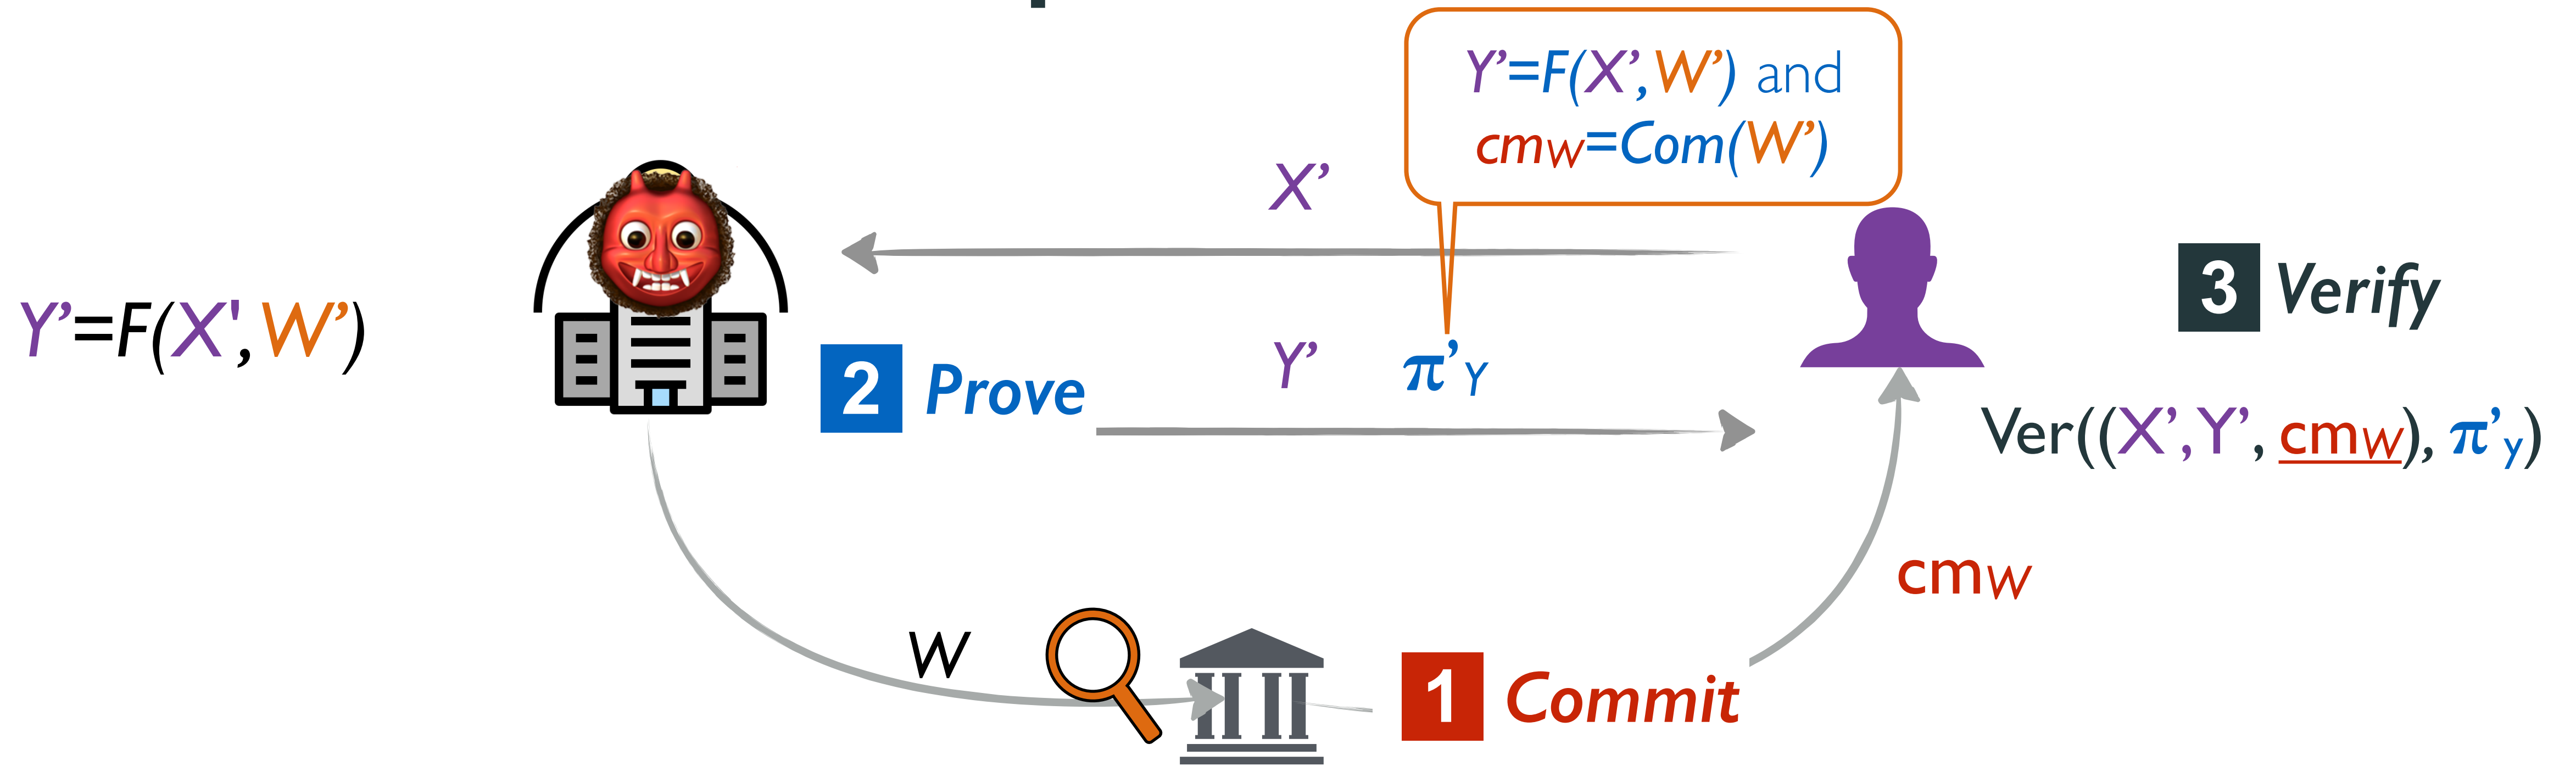


ZKPs for secure and private ML inference

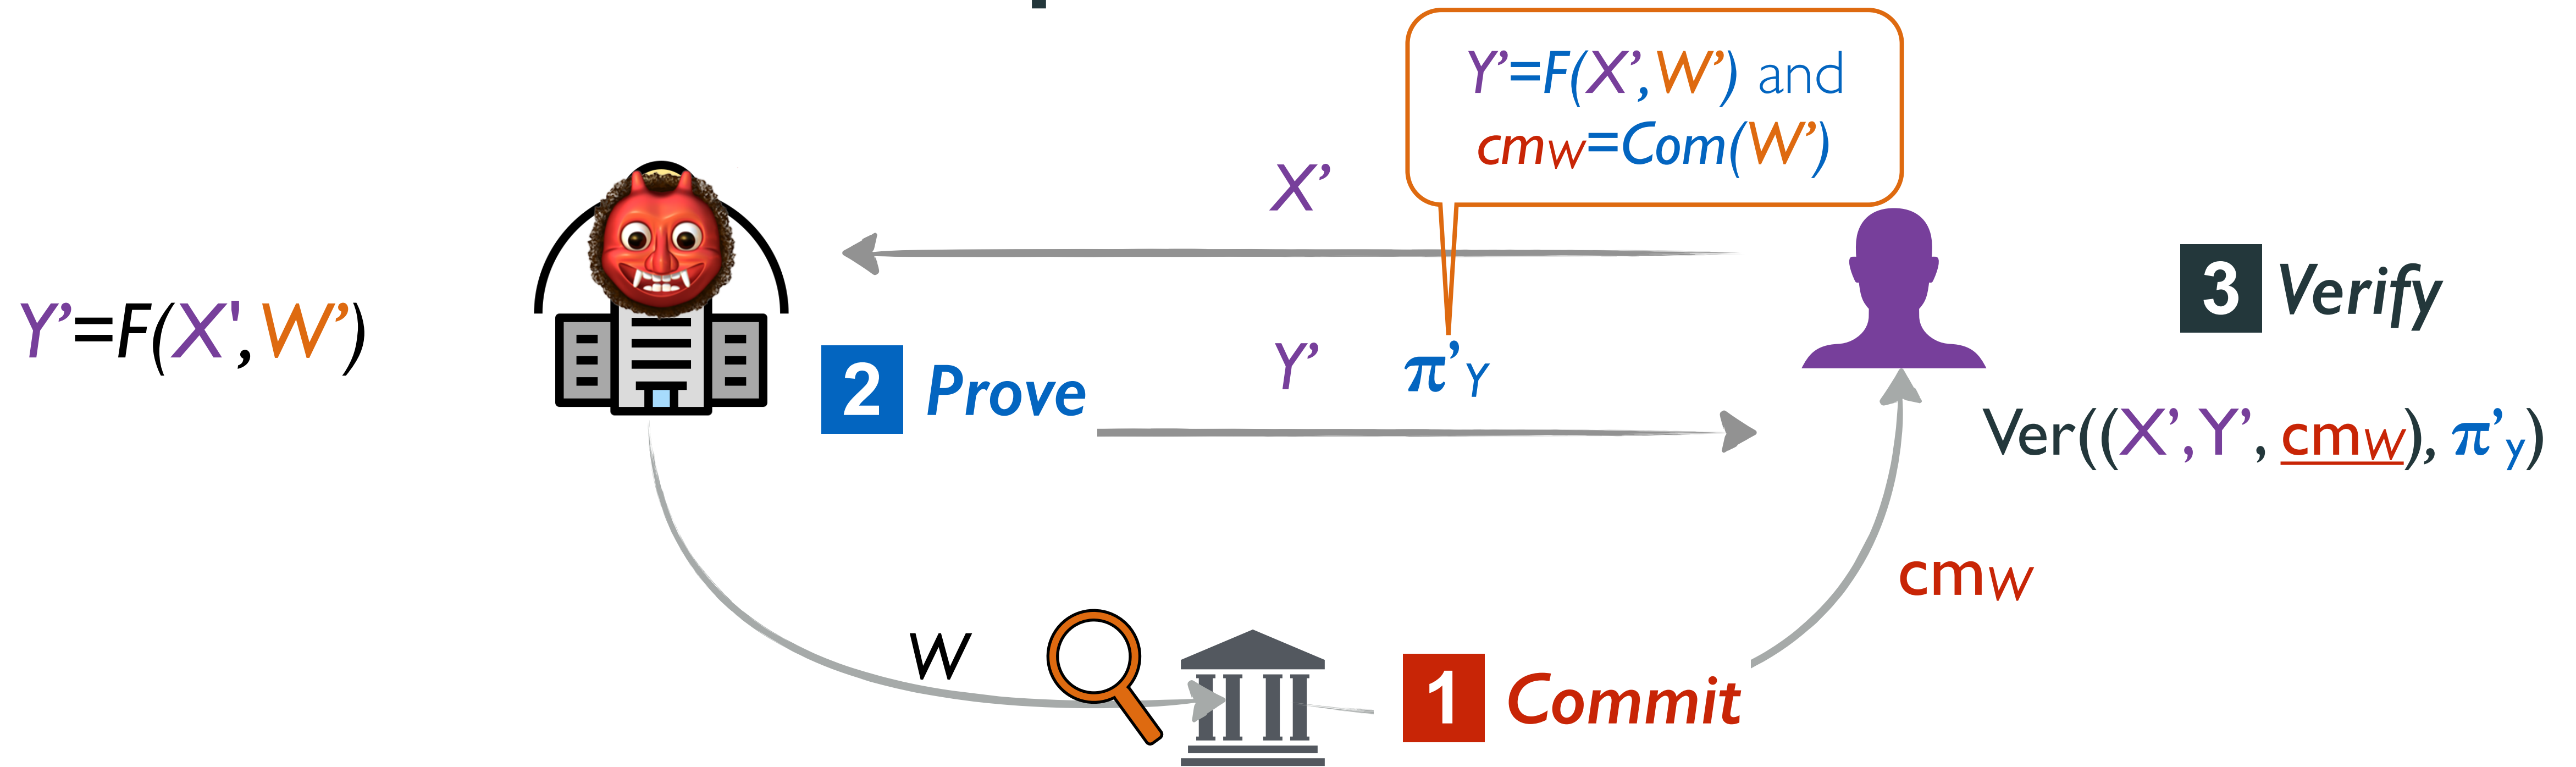
$$Y' = F(X', W')$$



ZKPs for secure and private ML inference



ZKPs for secure and private ML inference

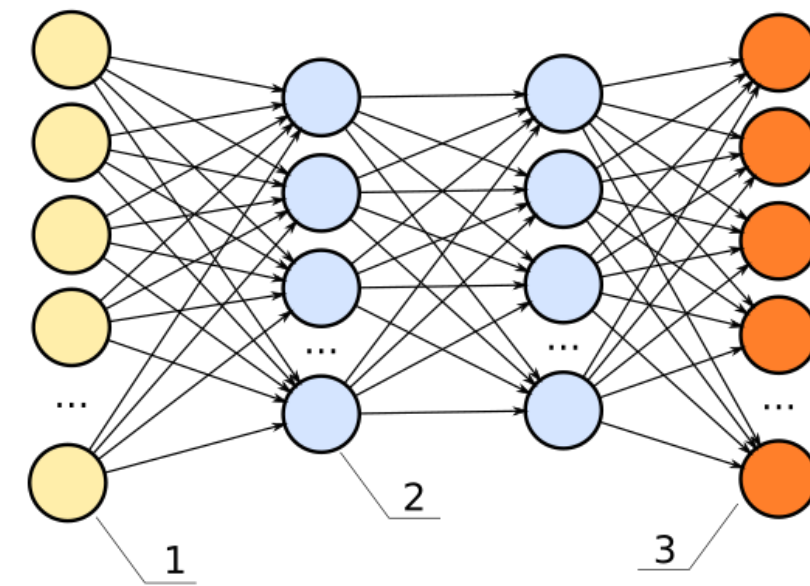
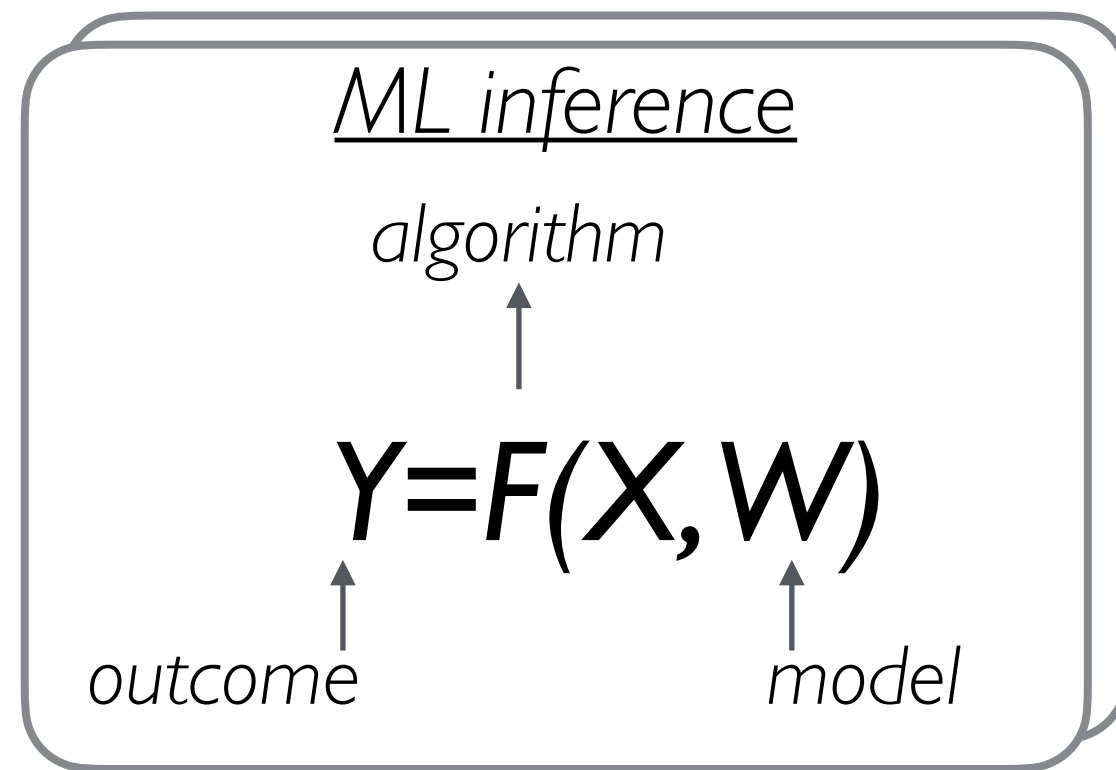


Integrity: detect tampered computations ← ZKP Soundness + Com binding

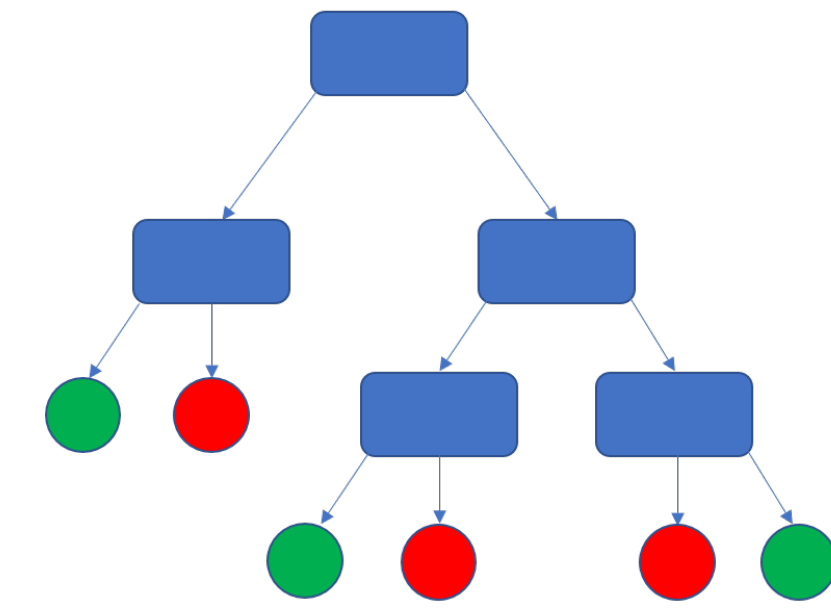
Privacy: clients should not learn anything about the model W ← ZKP ZK + Com hiding

Fairness: decision process the same for all clients ← ZKP Soundness + Com binding

Practical challenges of constructing ZKPs for ML



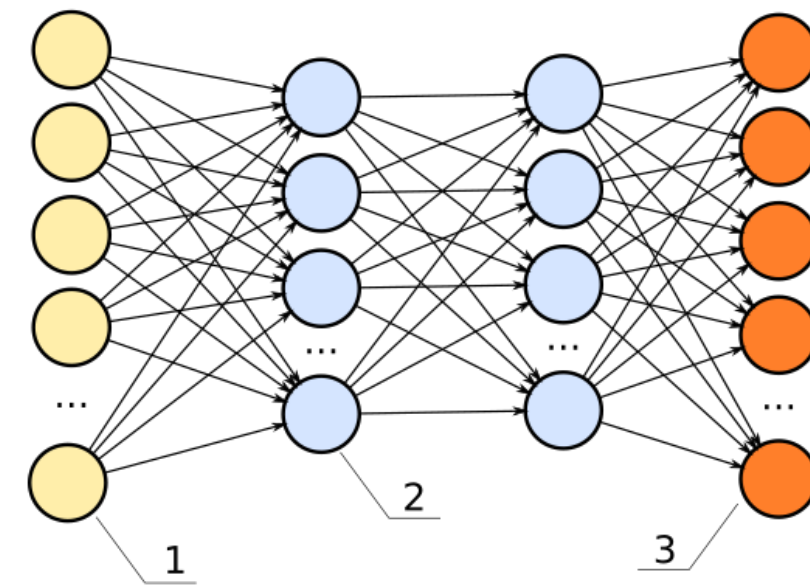
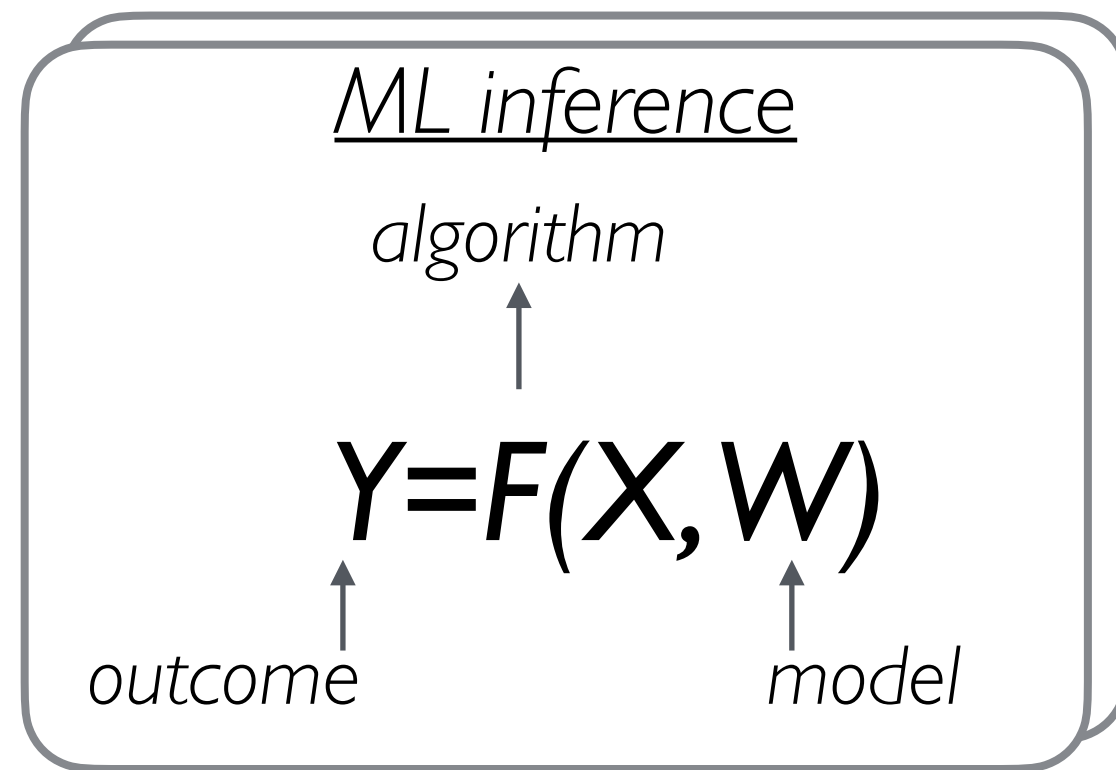
Neural Networks



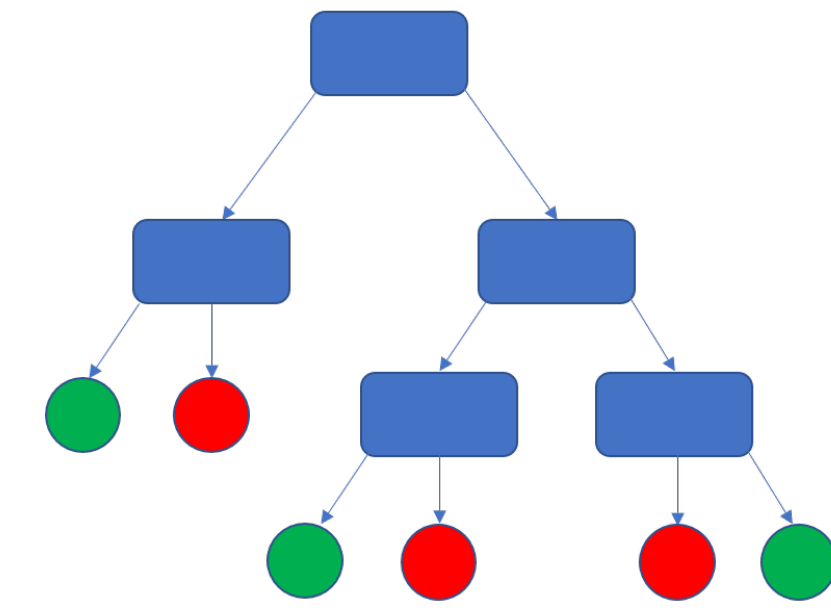
Decision Trees

Scale with **large models** and **not-that-ZKP-friendly** computations

Practical challenges of constructing ZKPs for ML



Neural Networks



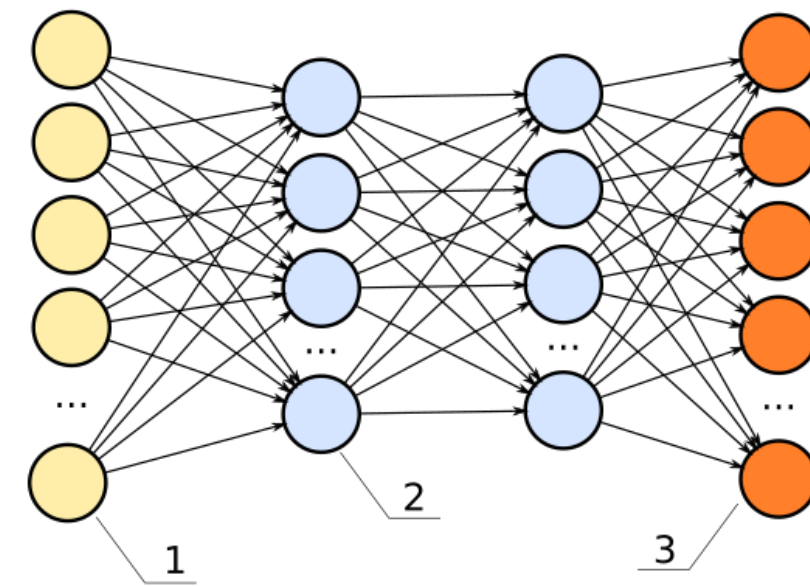
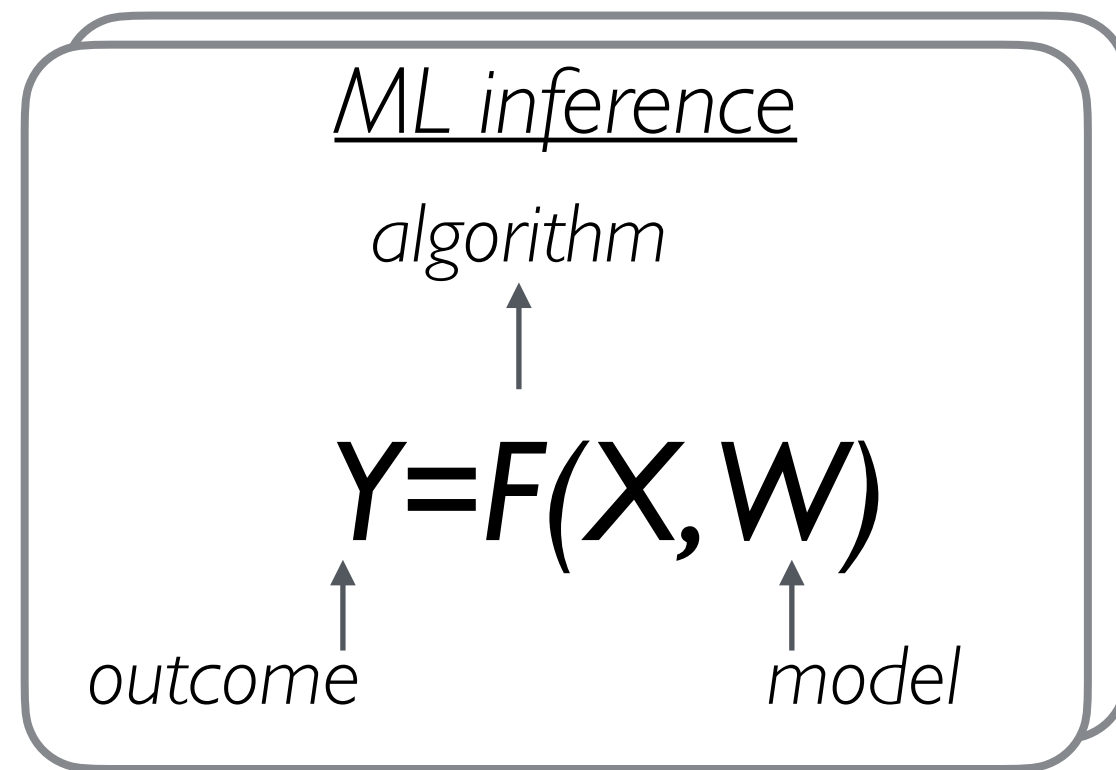
Decision Trees

Scale with **large models** and **not-that-ZKP-friendly** computations

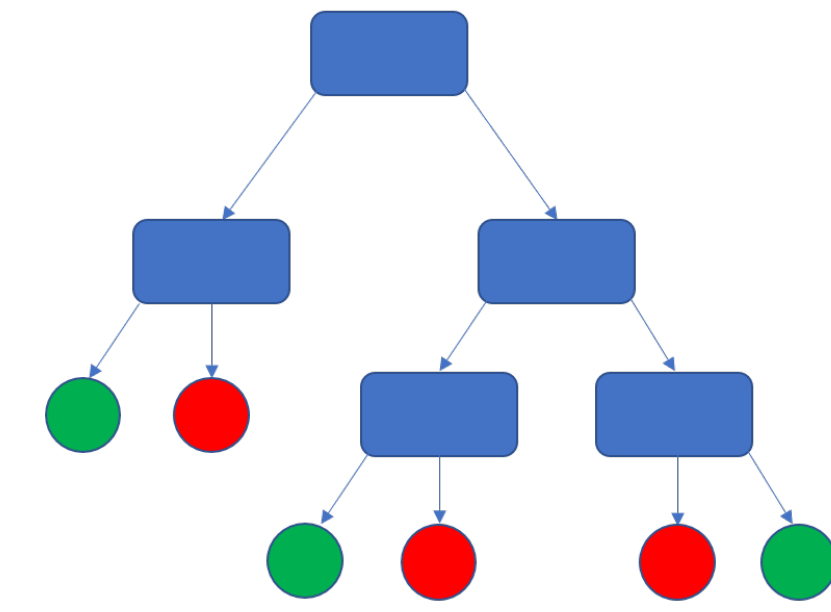
VGG16 (one of the best computer vision NN) has ~500MB parameters

Dense matrix operations, non-linear layers (NN), threshold&comparison gates (DT)

Practical challenges of constructing ZKPs for ML



Neural Networks



Decision Trees

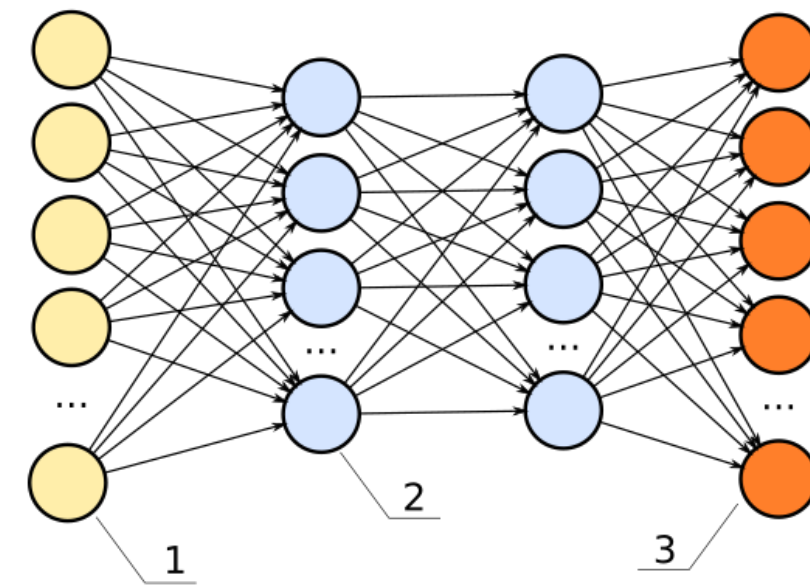
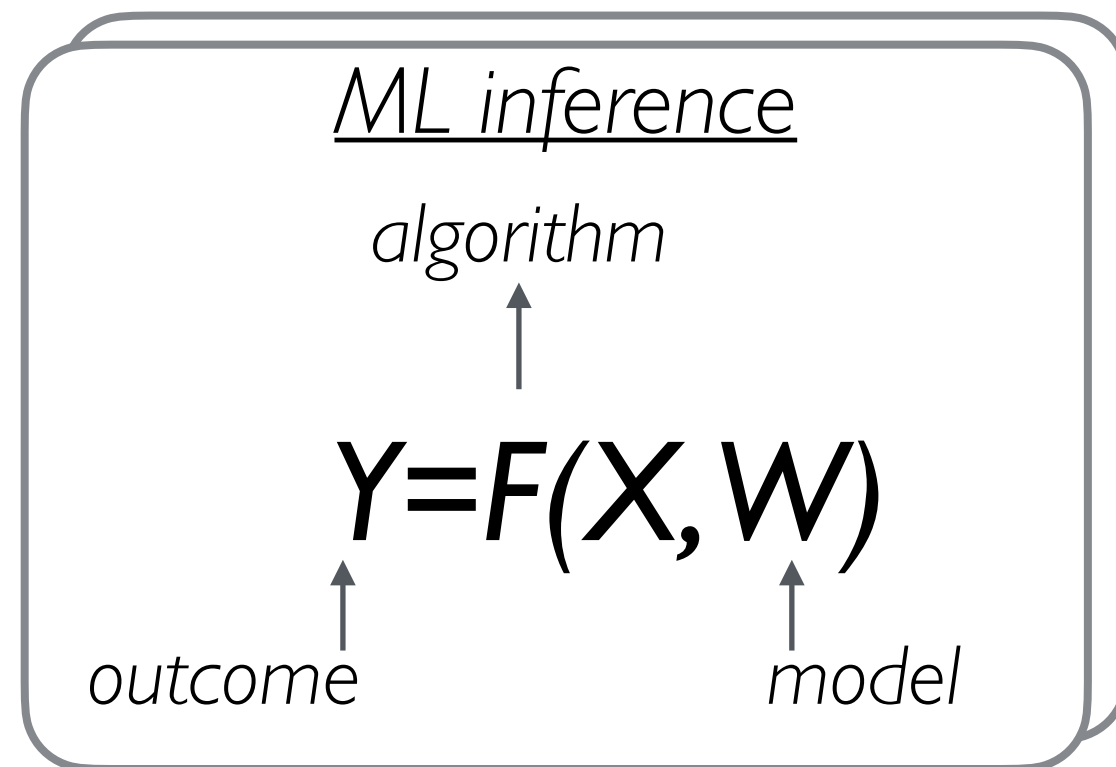
Scale with **large models** and **not-that-ZKP-friendly** computations

VGG16 (one of the best computer vision NN) has ~500MB parameters

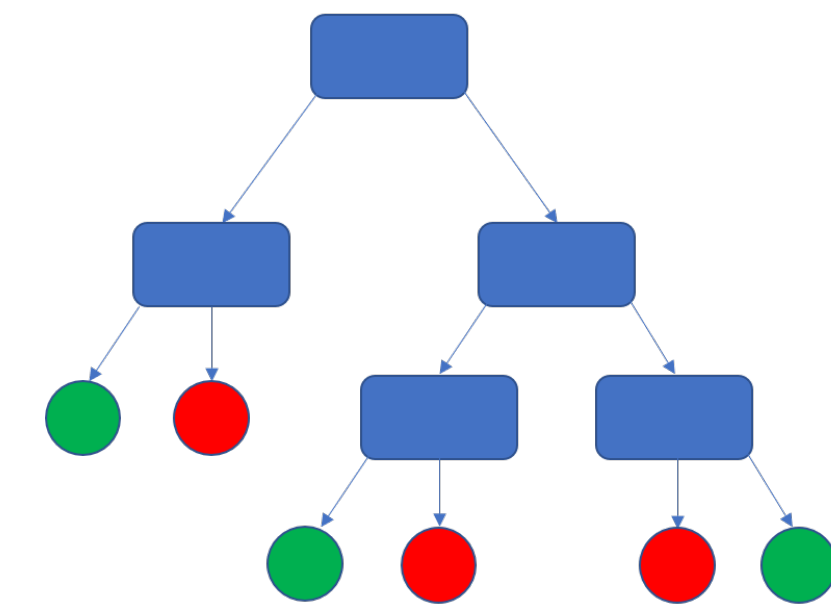
Dense matrix operations, non-linear layers (NN), threshold&comparison gates (DT)

Proof verification — can leverage SNARKs

Practical challenges of constructing ZKPs for ML



Neural Networks



Decision Trees

Scale with **large models** and **not-that-ZKP-friendly** computations

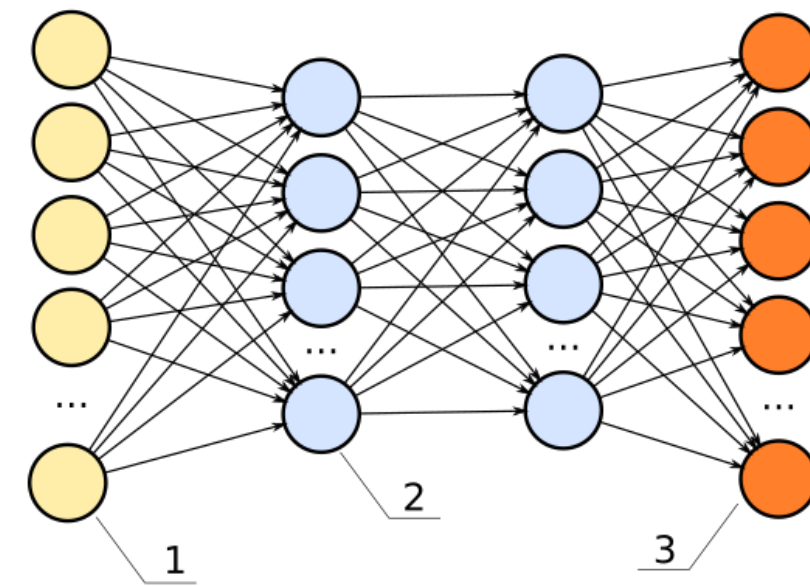
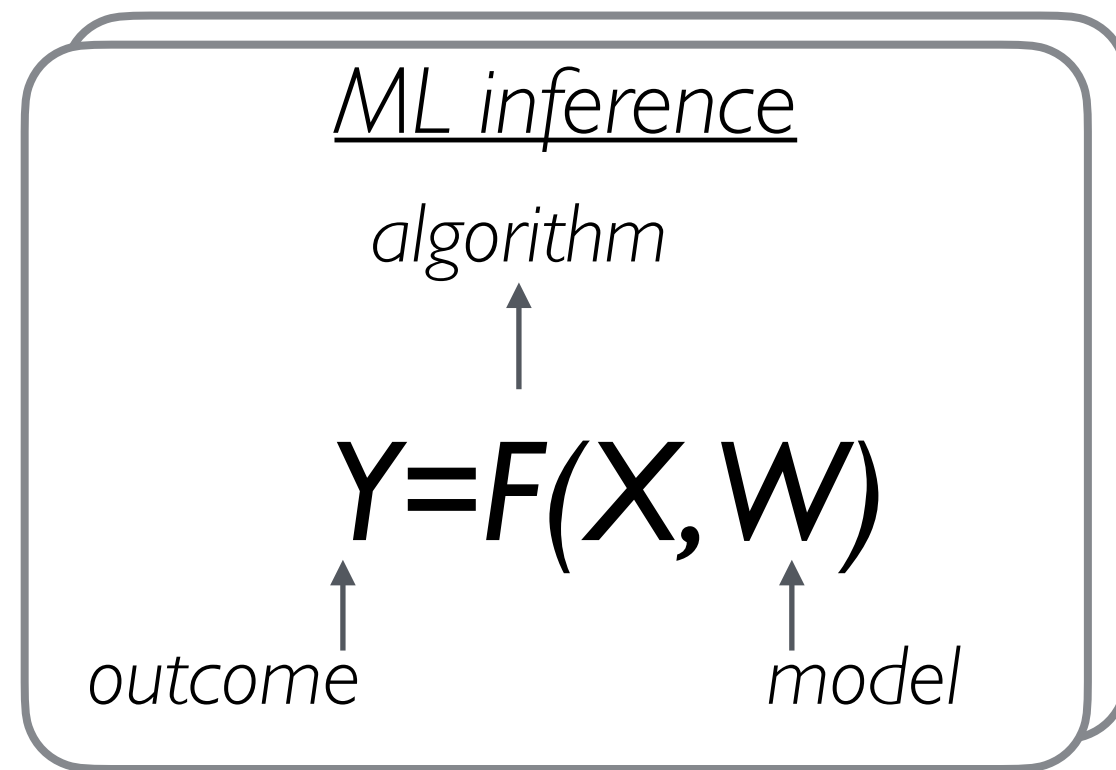
VGG16 (one of the best computer vision NN) has ~500MB parameters

Dense matrix operations, non-linear layers (NN), threshold&comparison gates (DT)

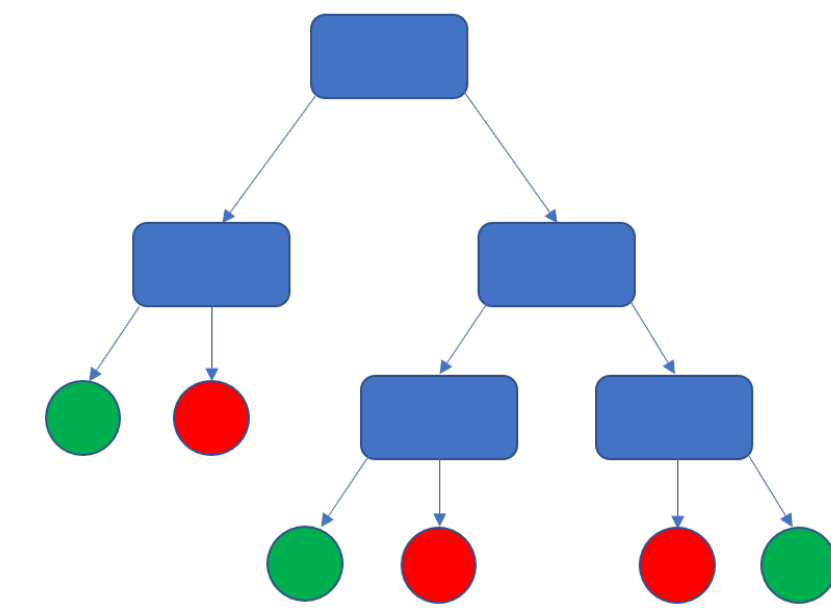
Proof verification — can leverage SNARKs

Proof generation — the most challenging, we'd like proving time $O(|F|)$ and concretely close to $|F|$

Practical challenges of constructing ZKPs for ML



Neural Networks



Decision Trees

Scale with **large models** and **not-that-ZKP-friendly** computations

VGG16 (one of the best computer vision NN) has ~500MB parameters

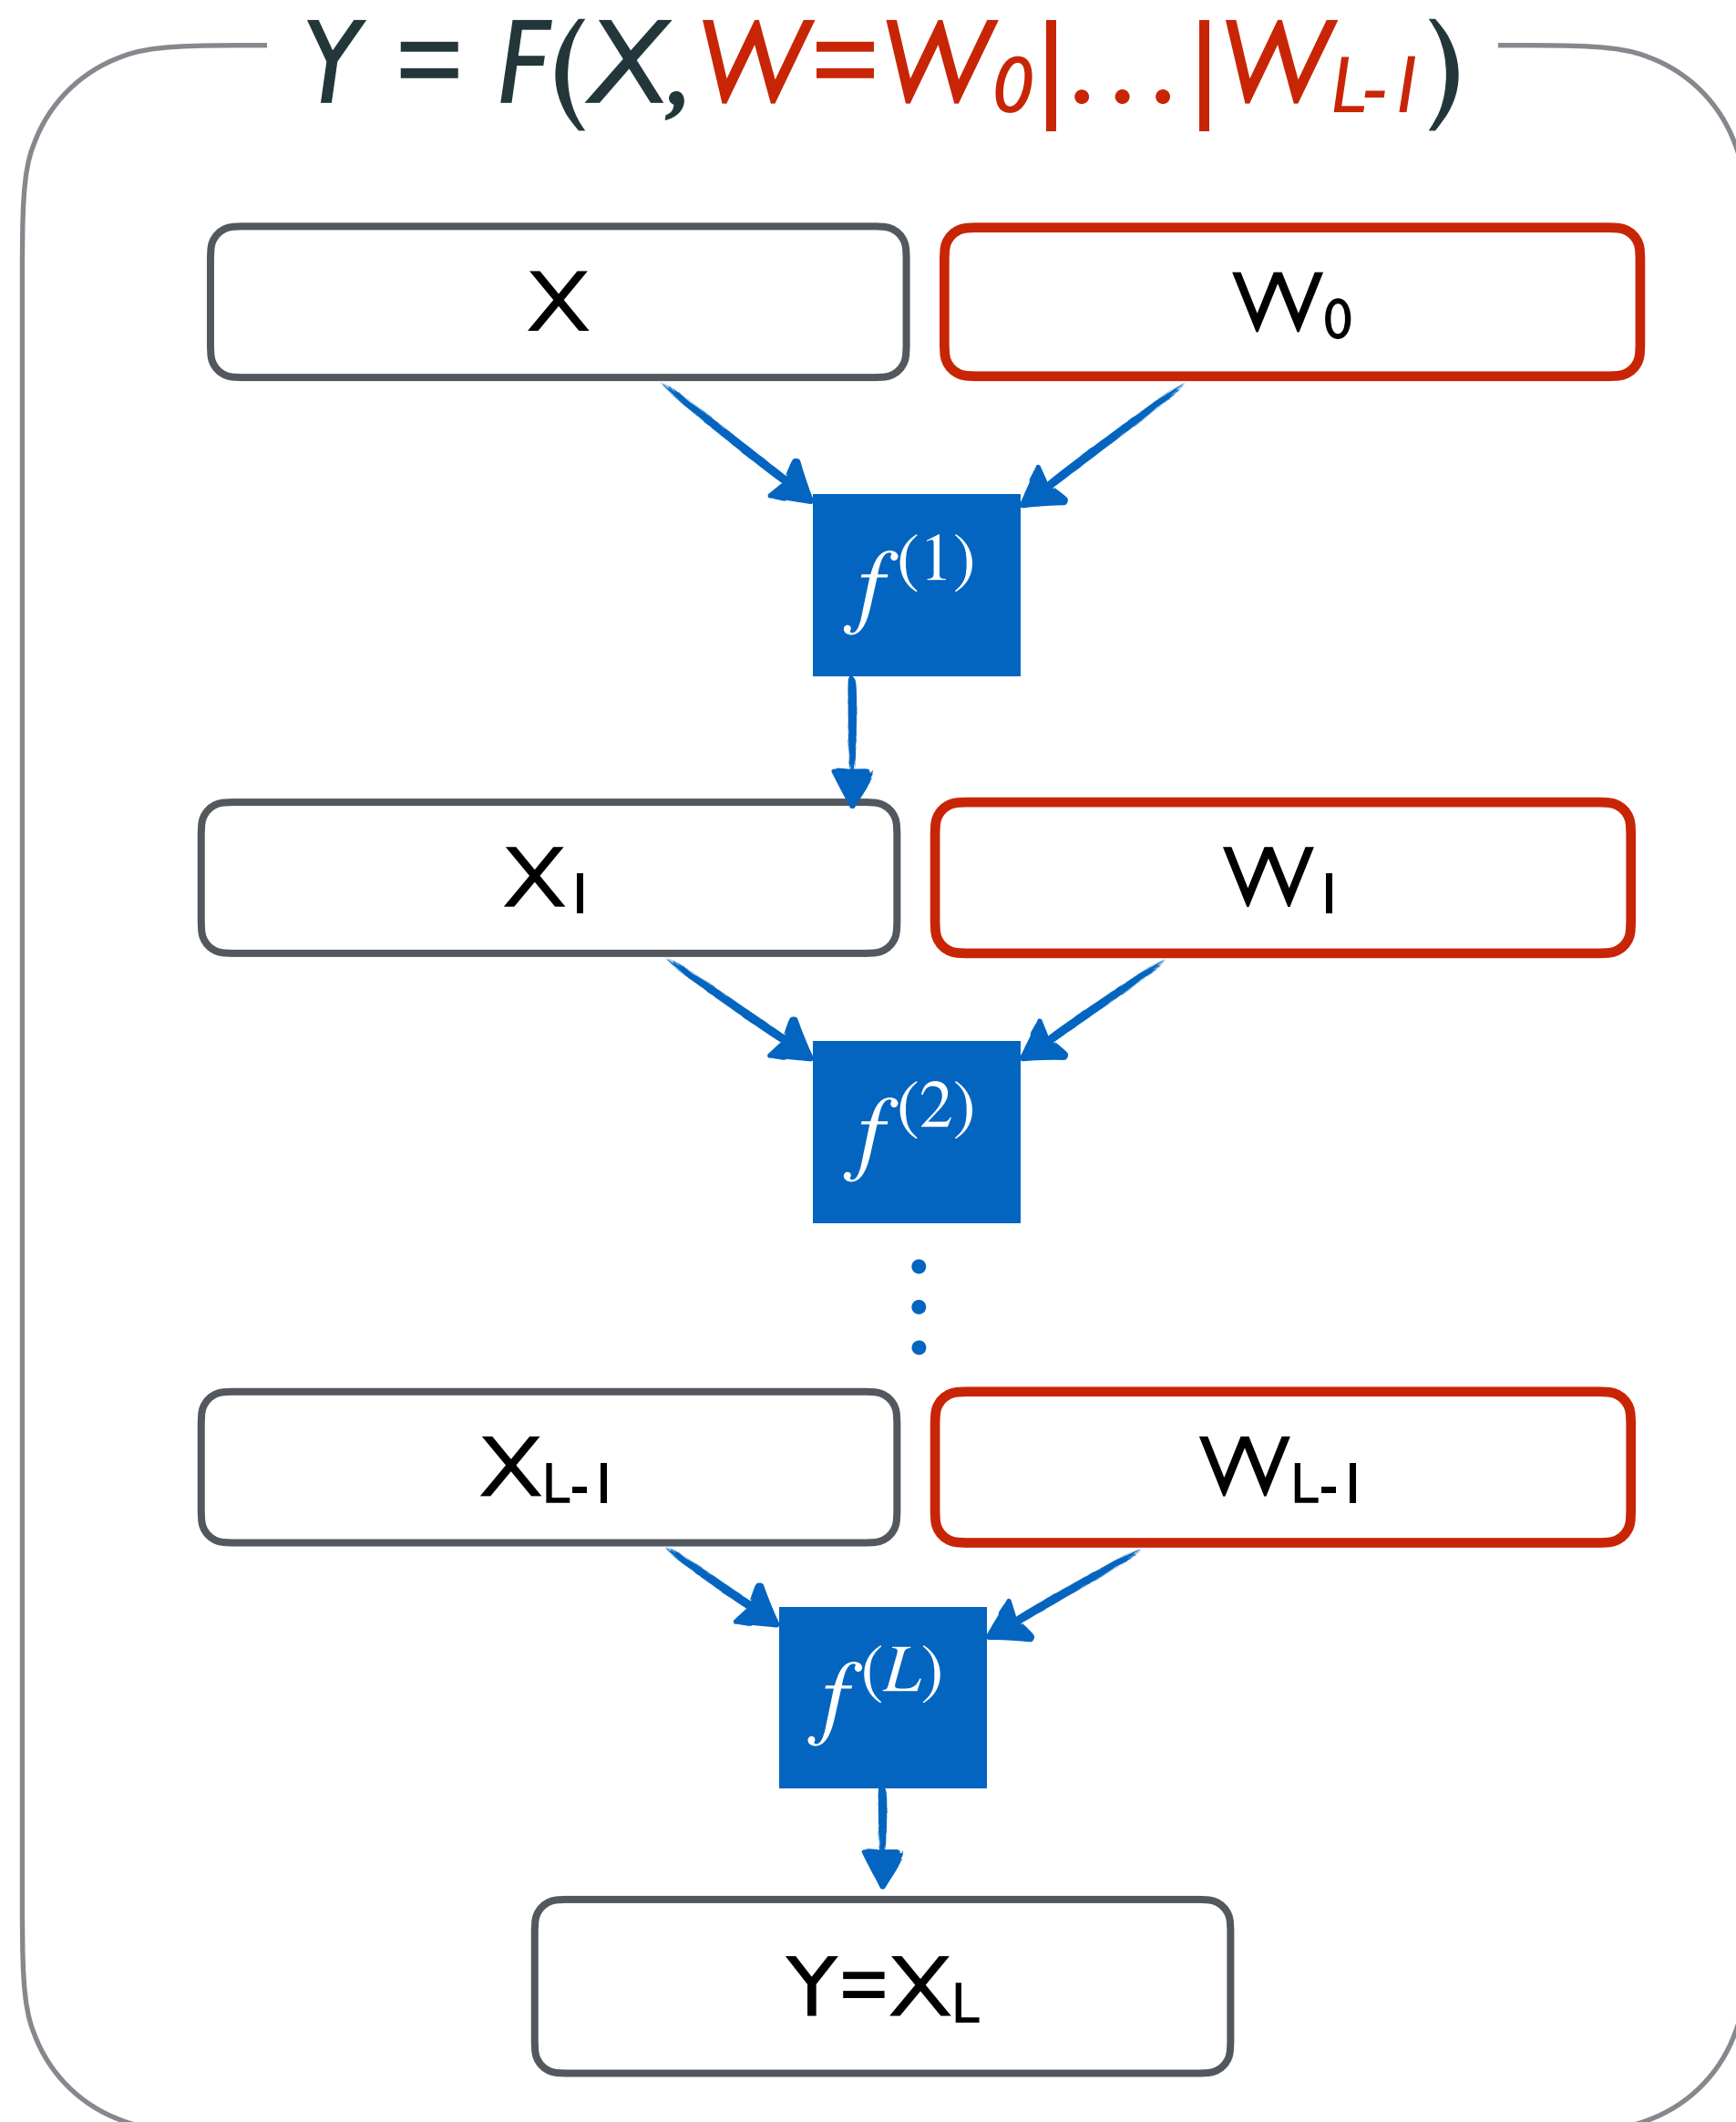
Dense matrix operations, non-linear layers (NN), threshold&comparison gates (DT)

Proof verification — can leverage SNARKs

Proof generation — the most challenging, we'd like proving time $O(|F|)$ and concretely close to $|F|$

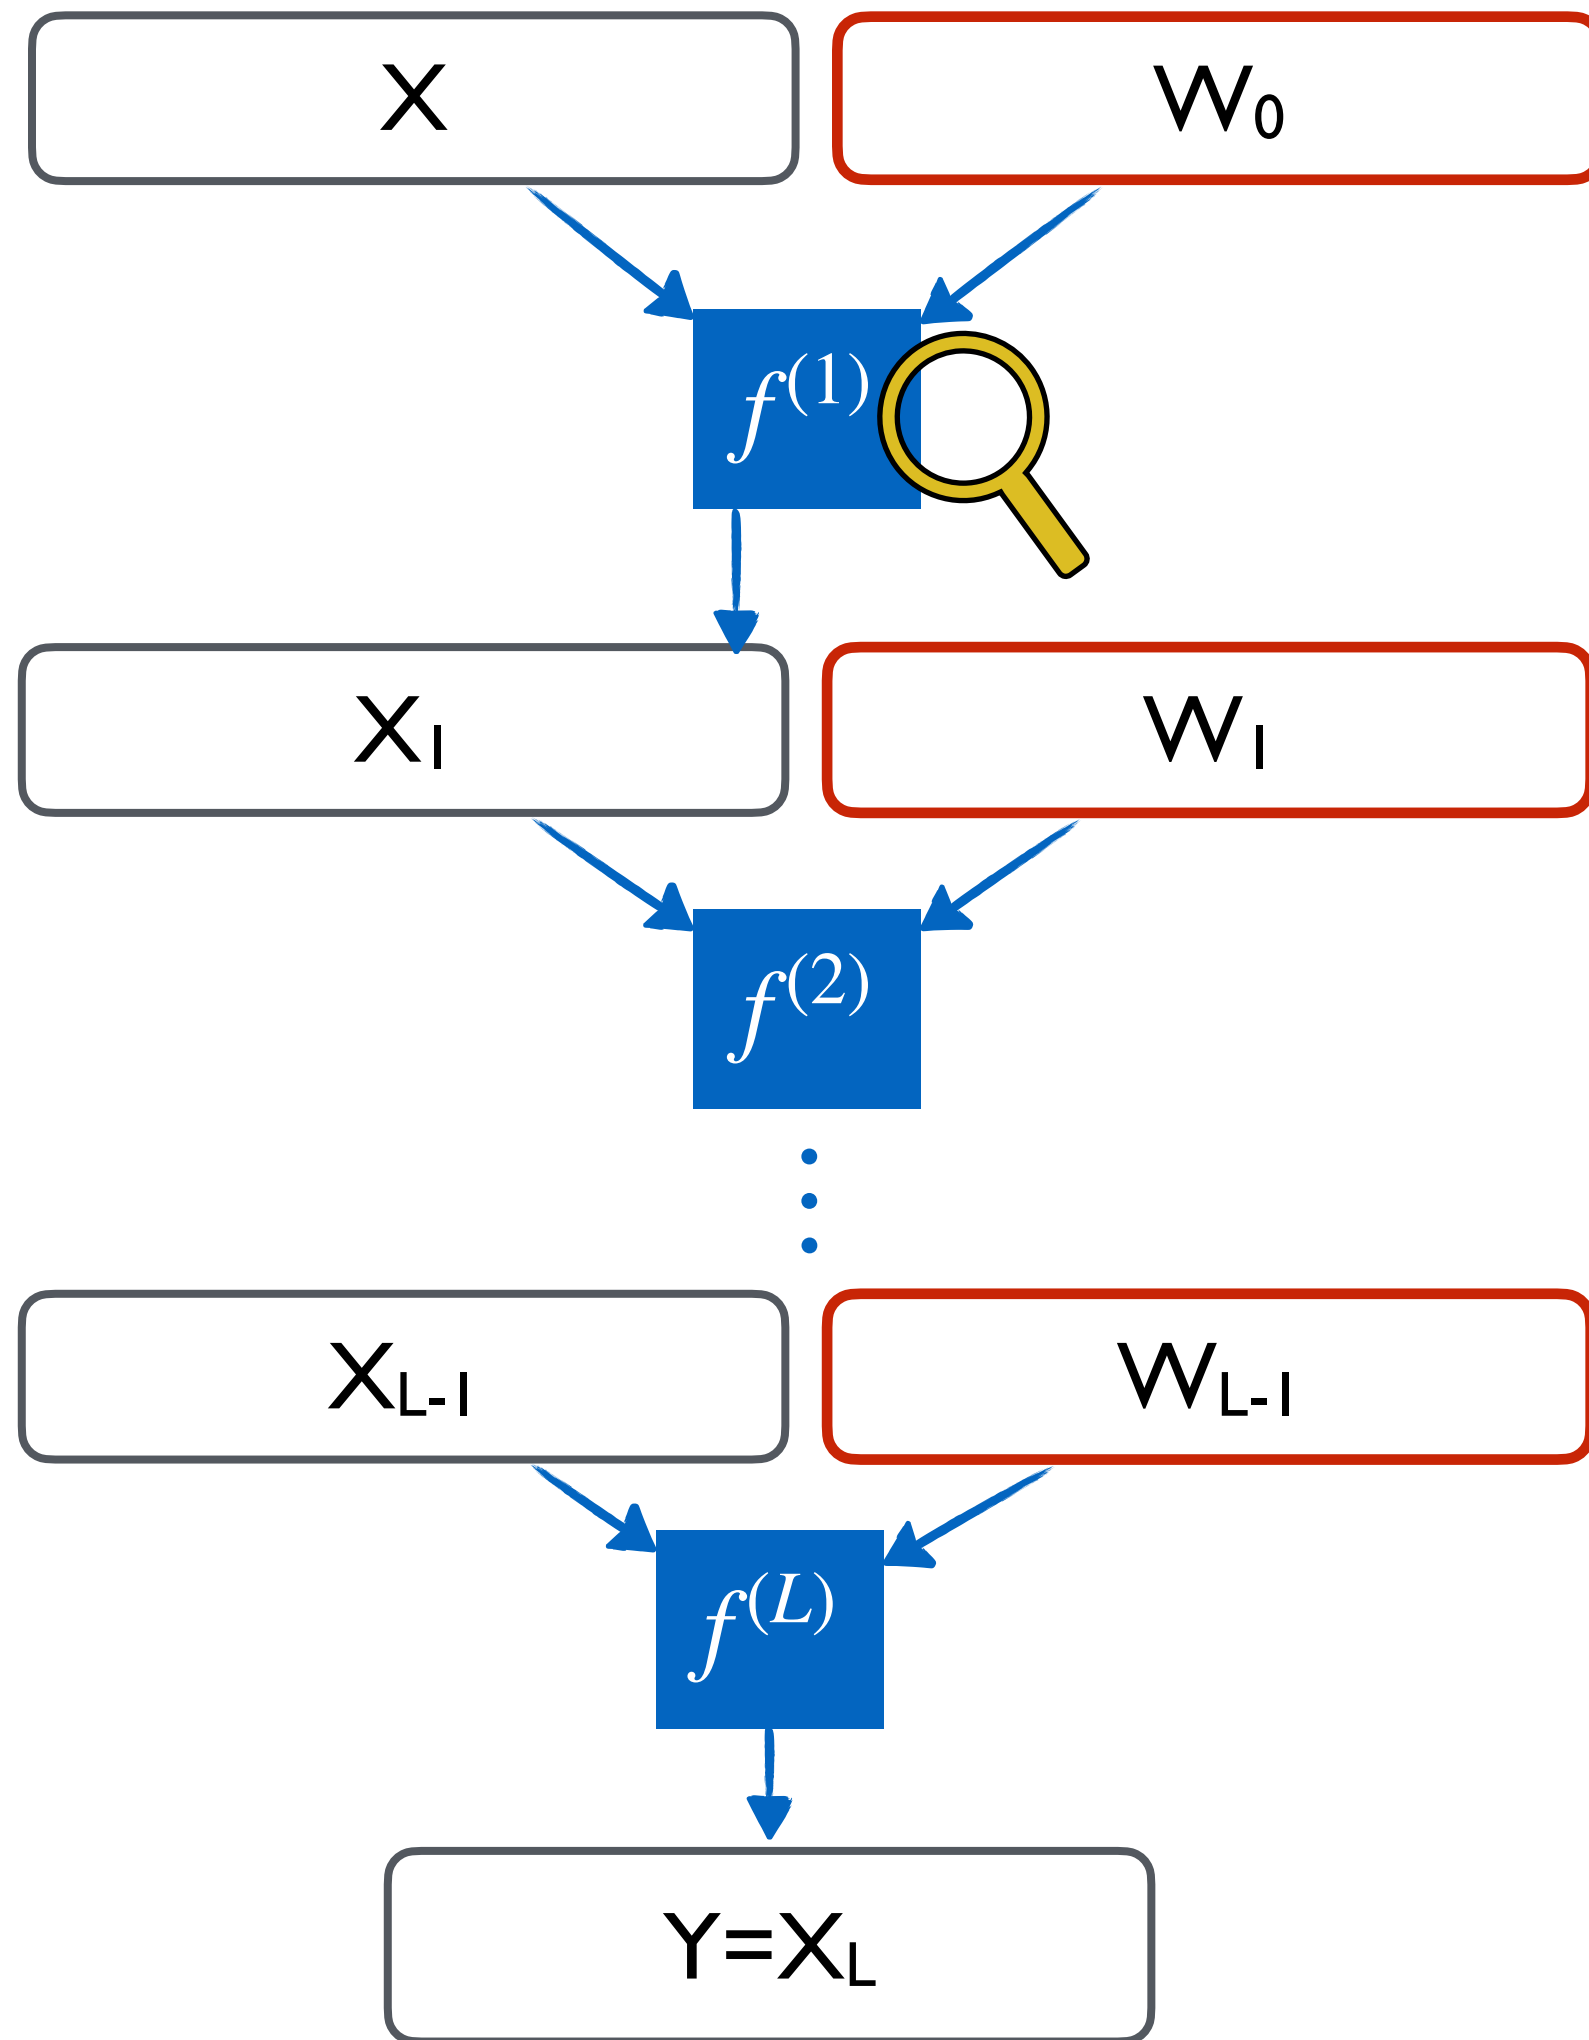
Solution: special-purpose ZKPs!

Convolutional Neural Networks



Convolutional Neural Networks

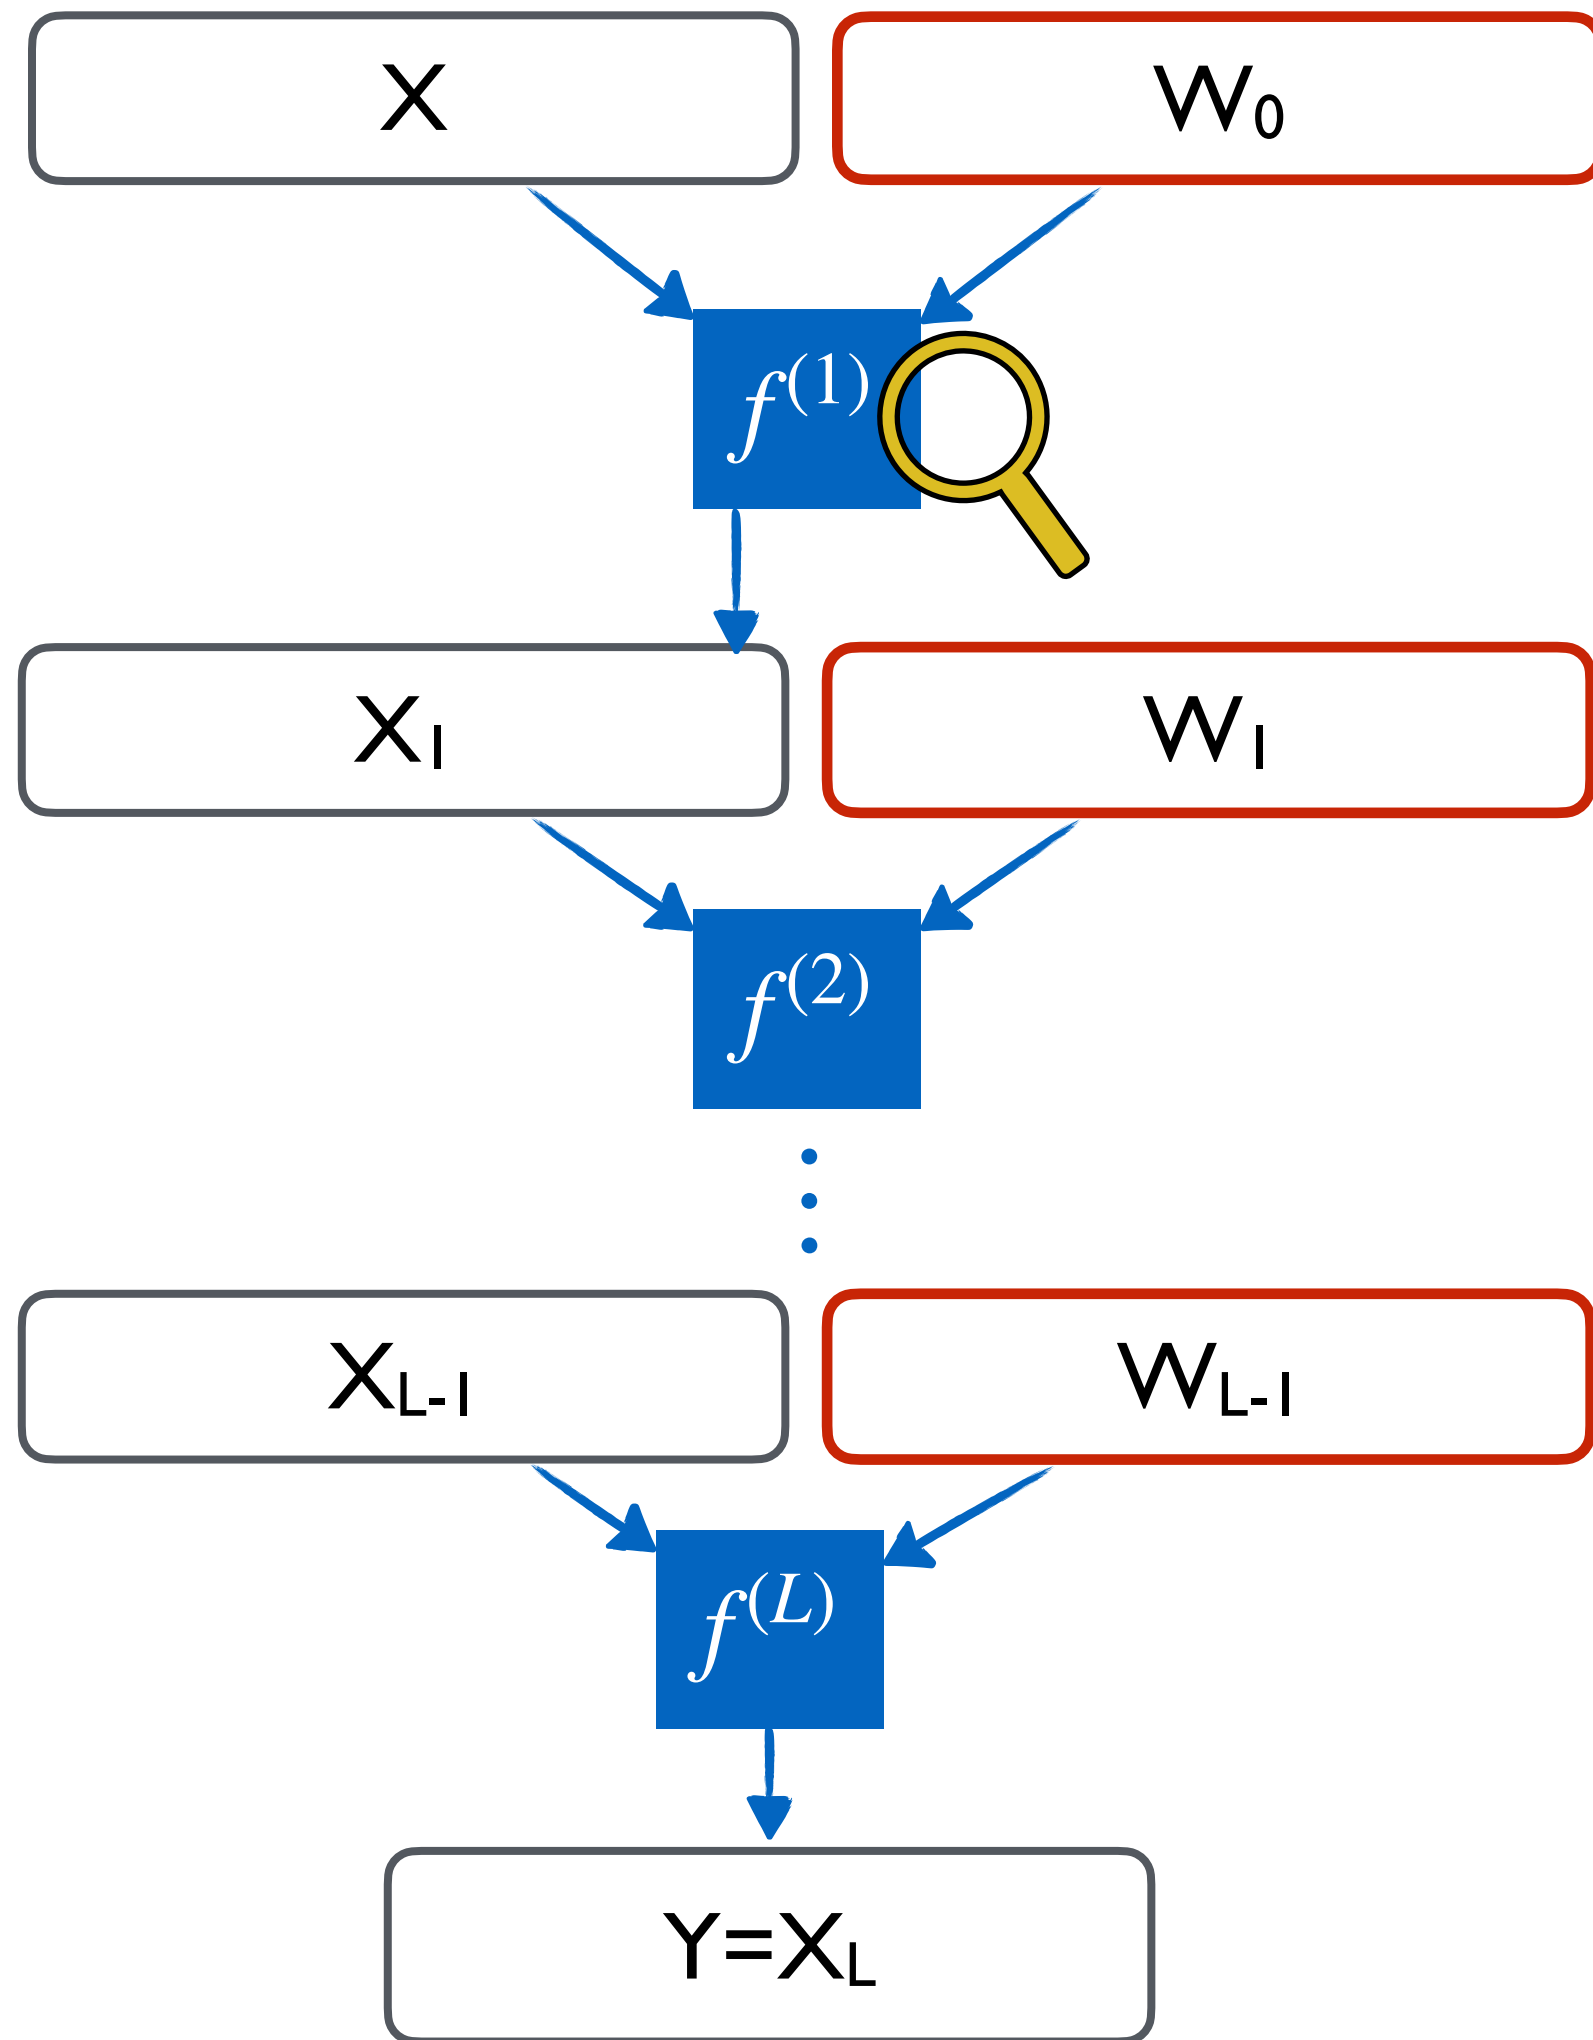
$$Y = F(X, W=W_0 | \dots | W_{L-1})$$



$$X_{k+1} = f^{(k)}(X_k, W_k)$$

Convolutional Neural Networks

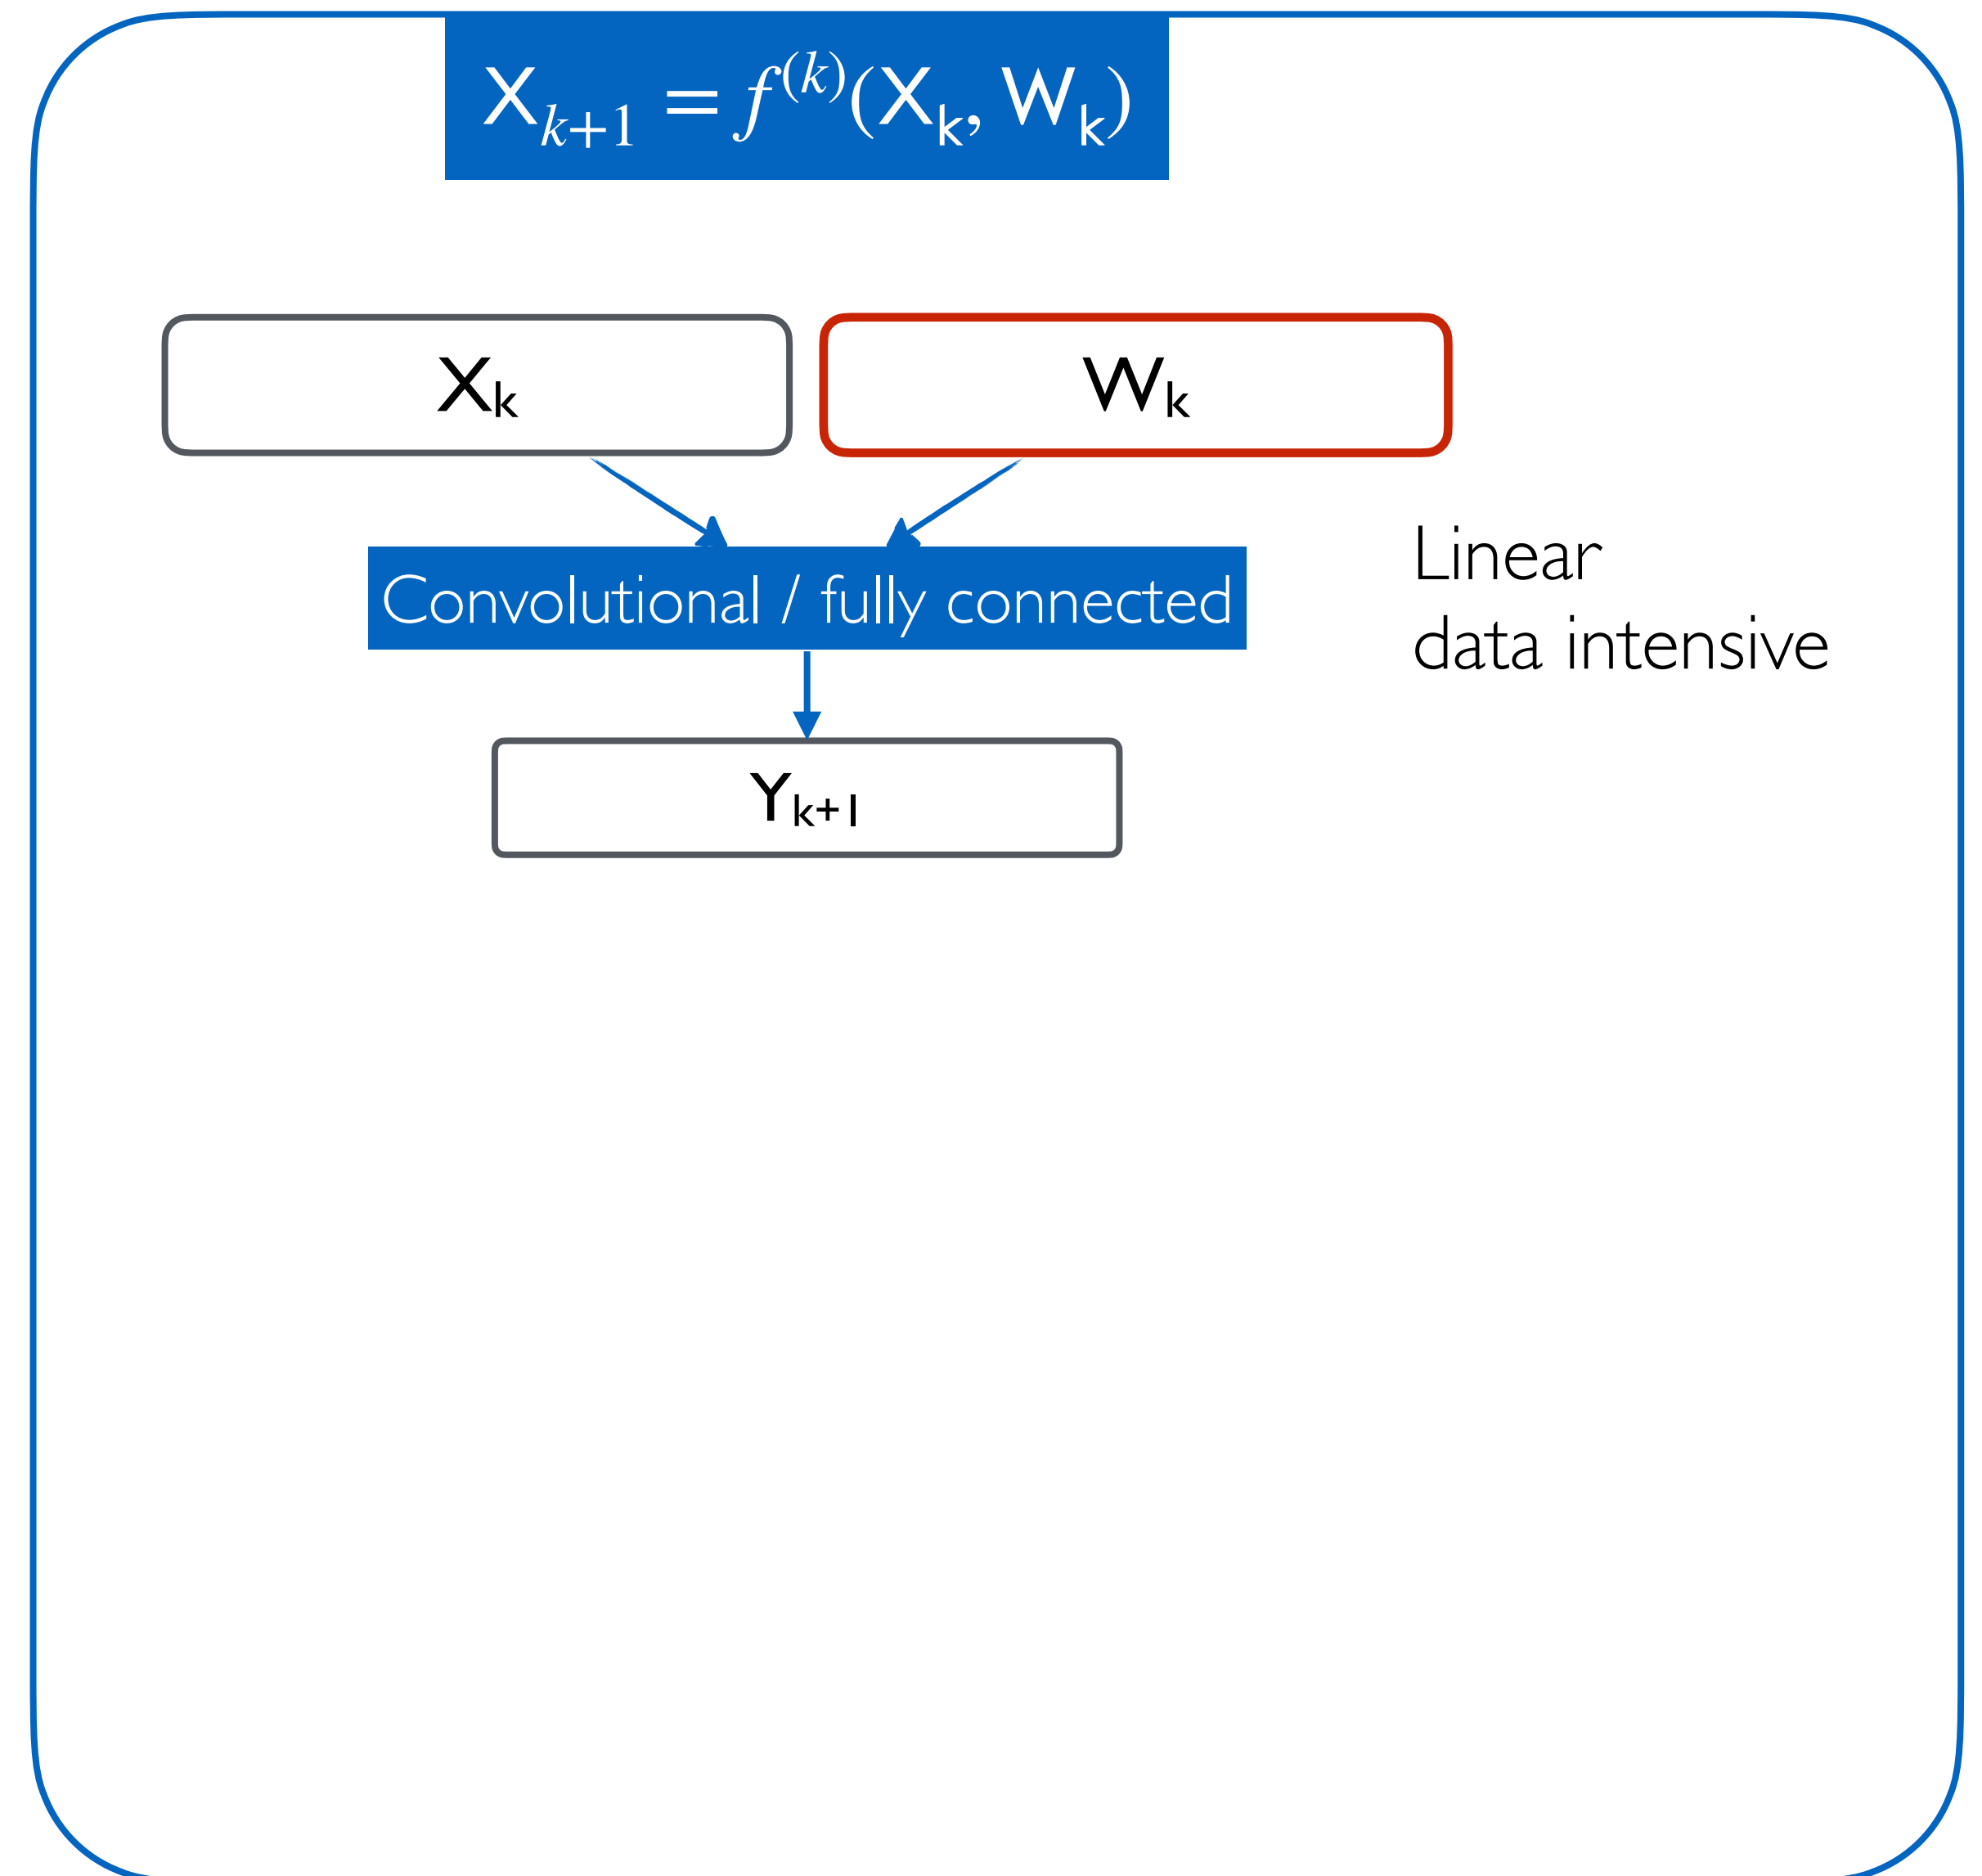
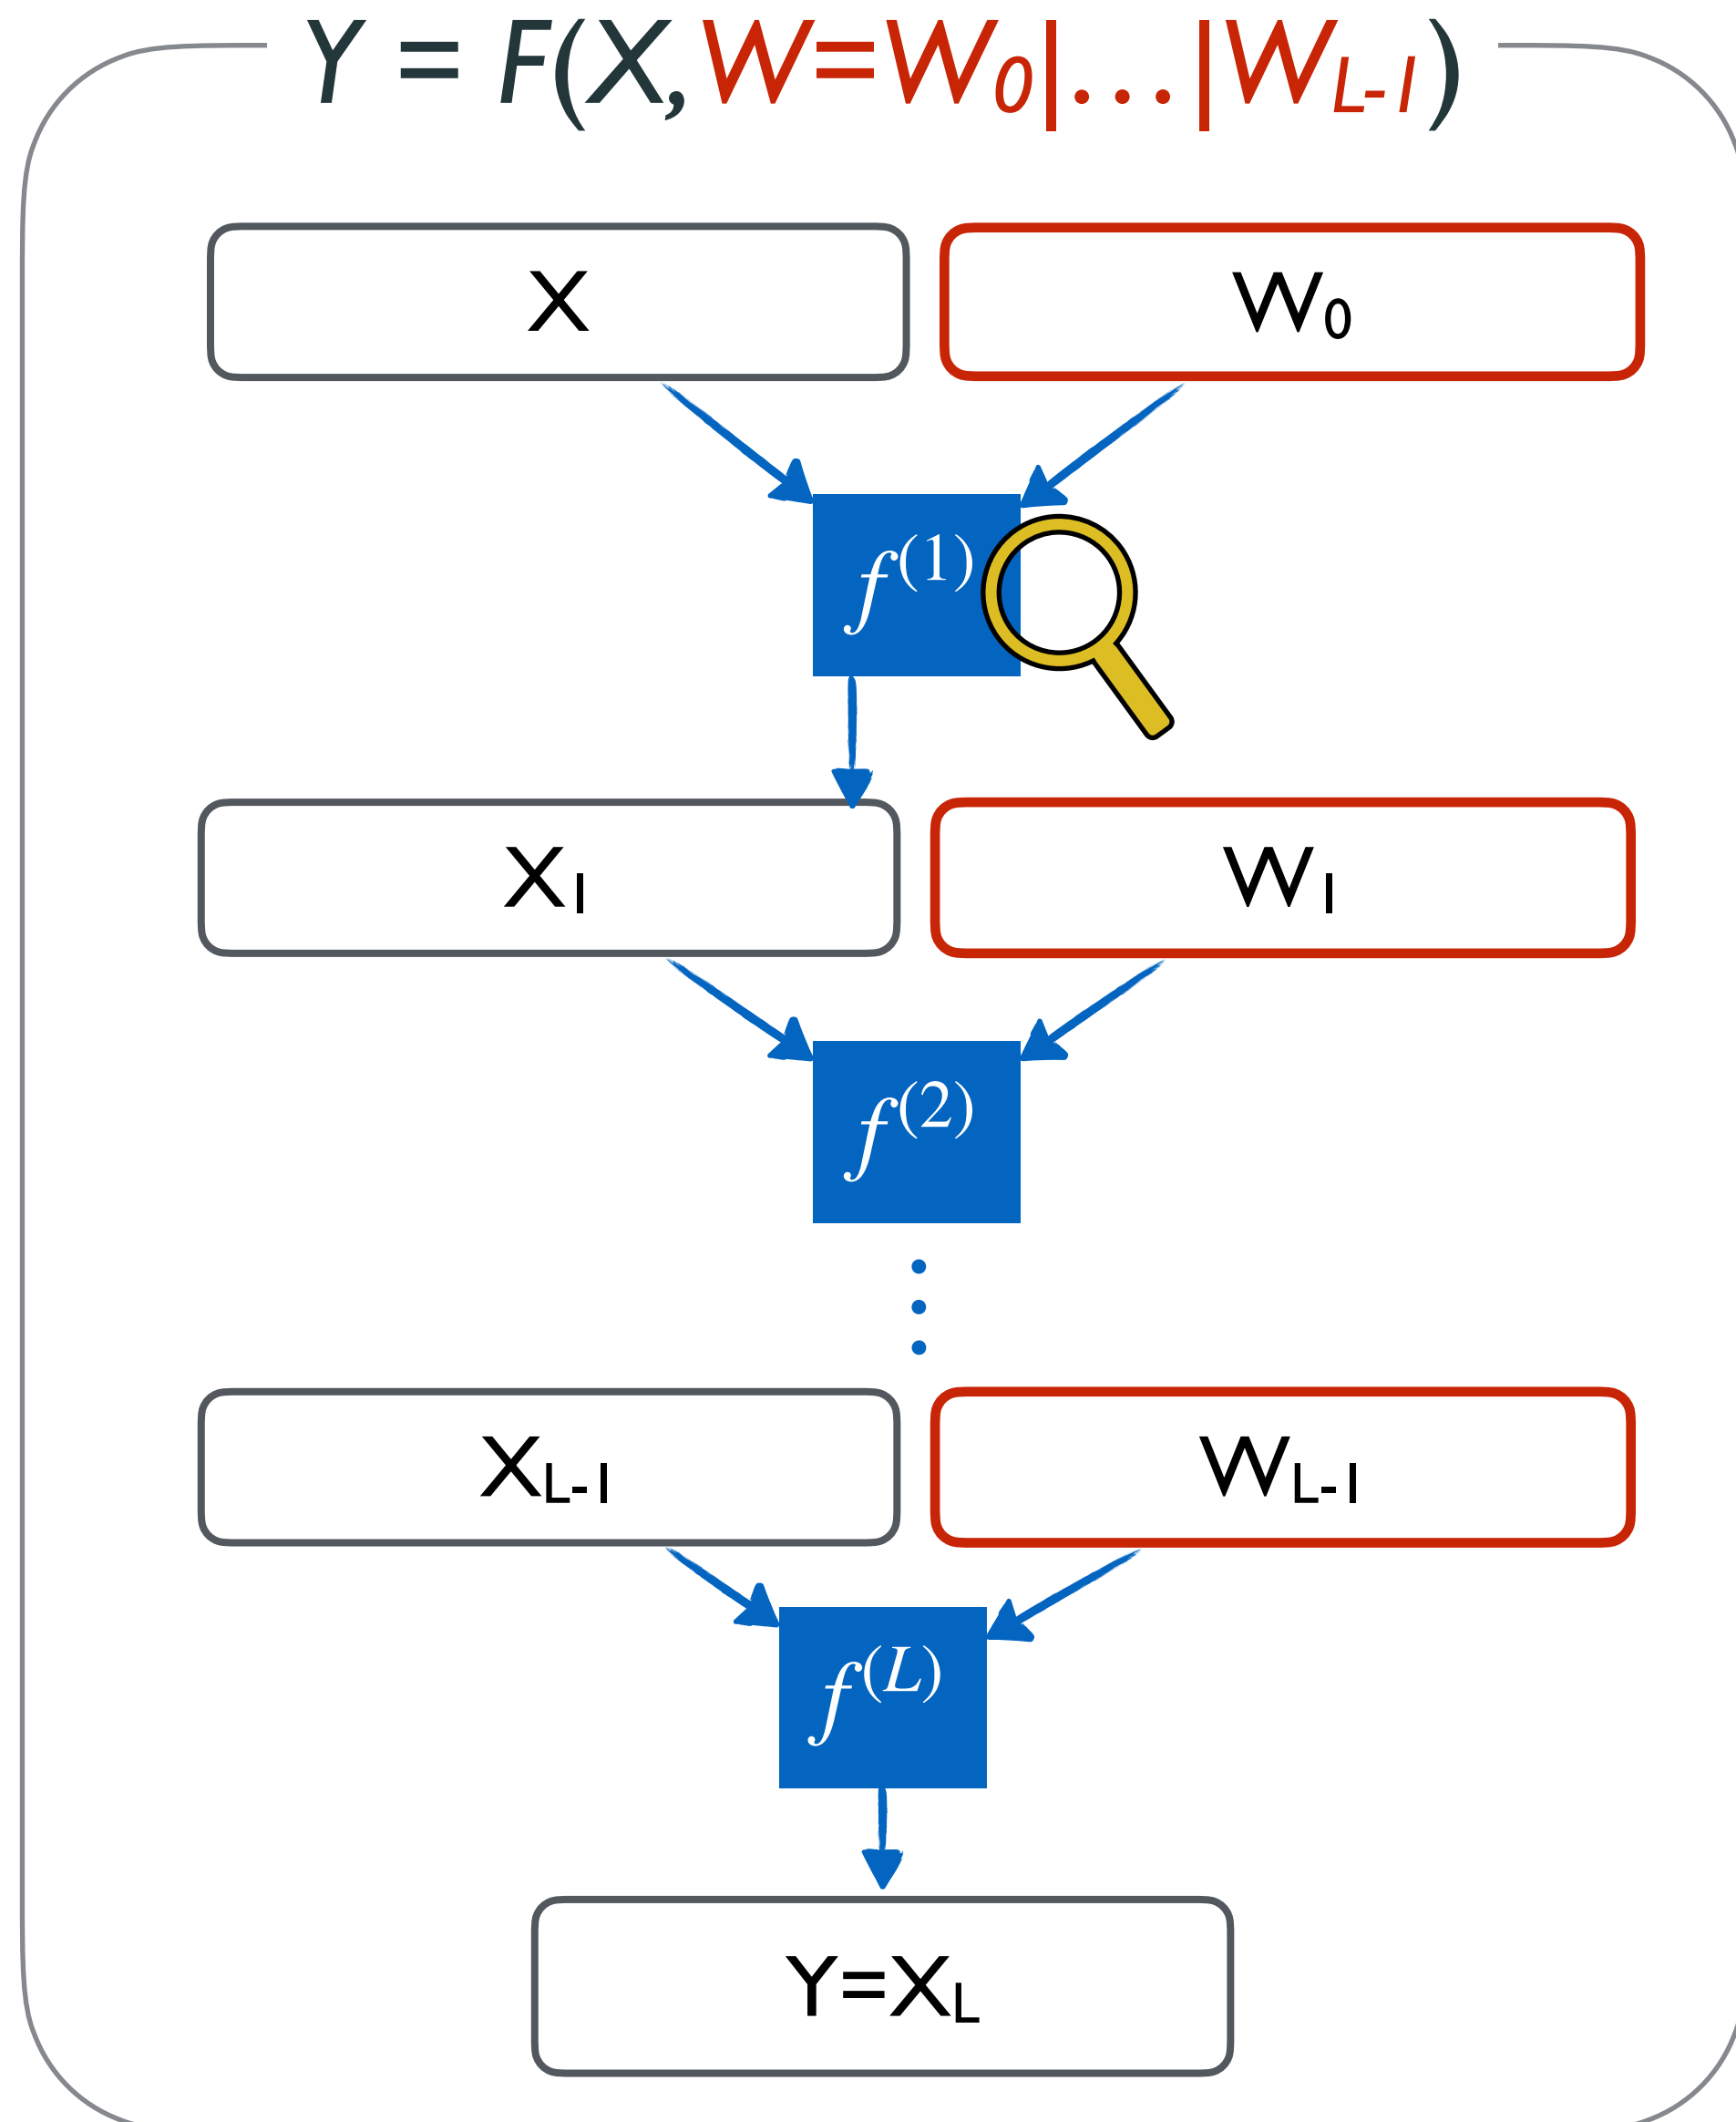
$$Y = F(X, W=W_0 | \dots | W_{L-1})$$



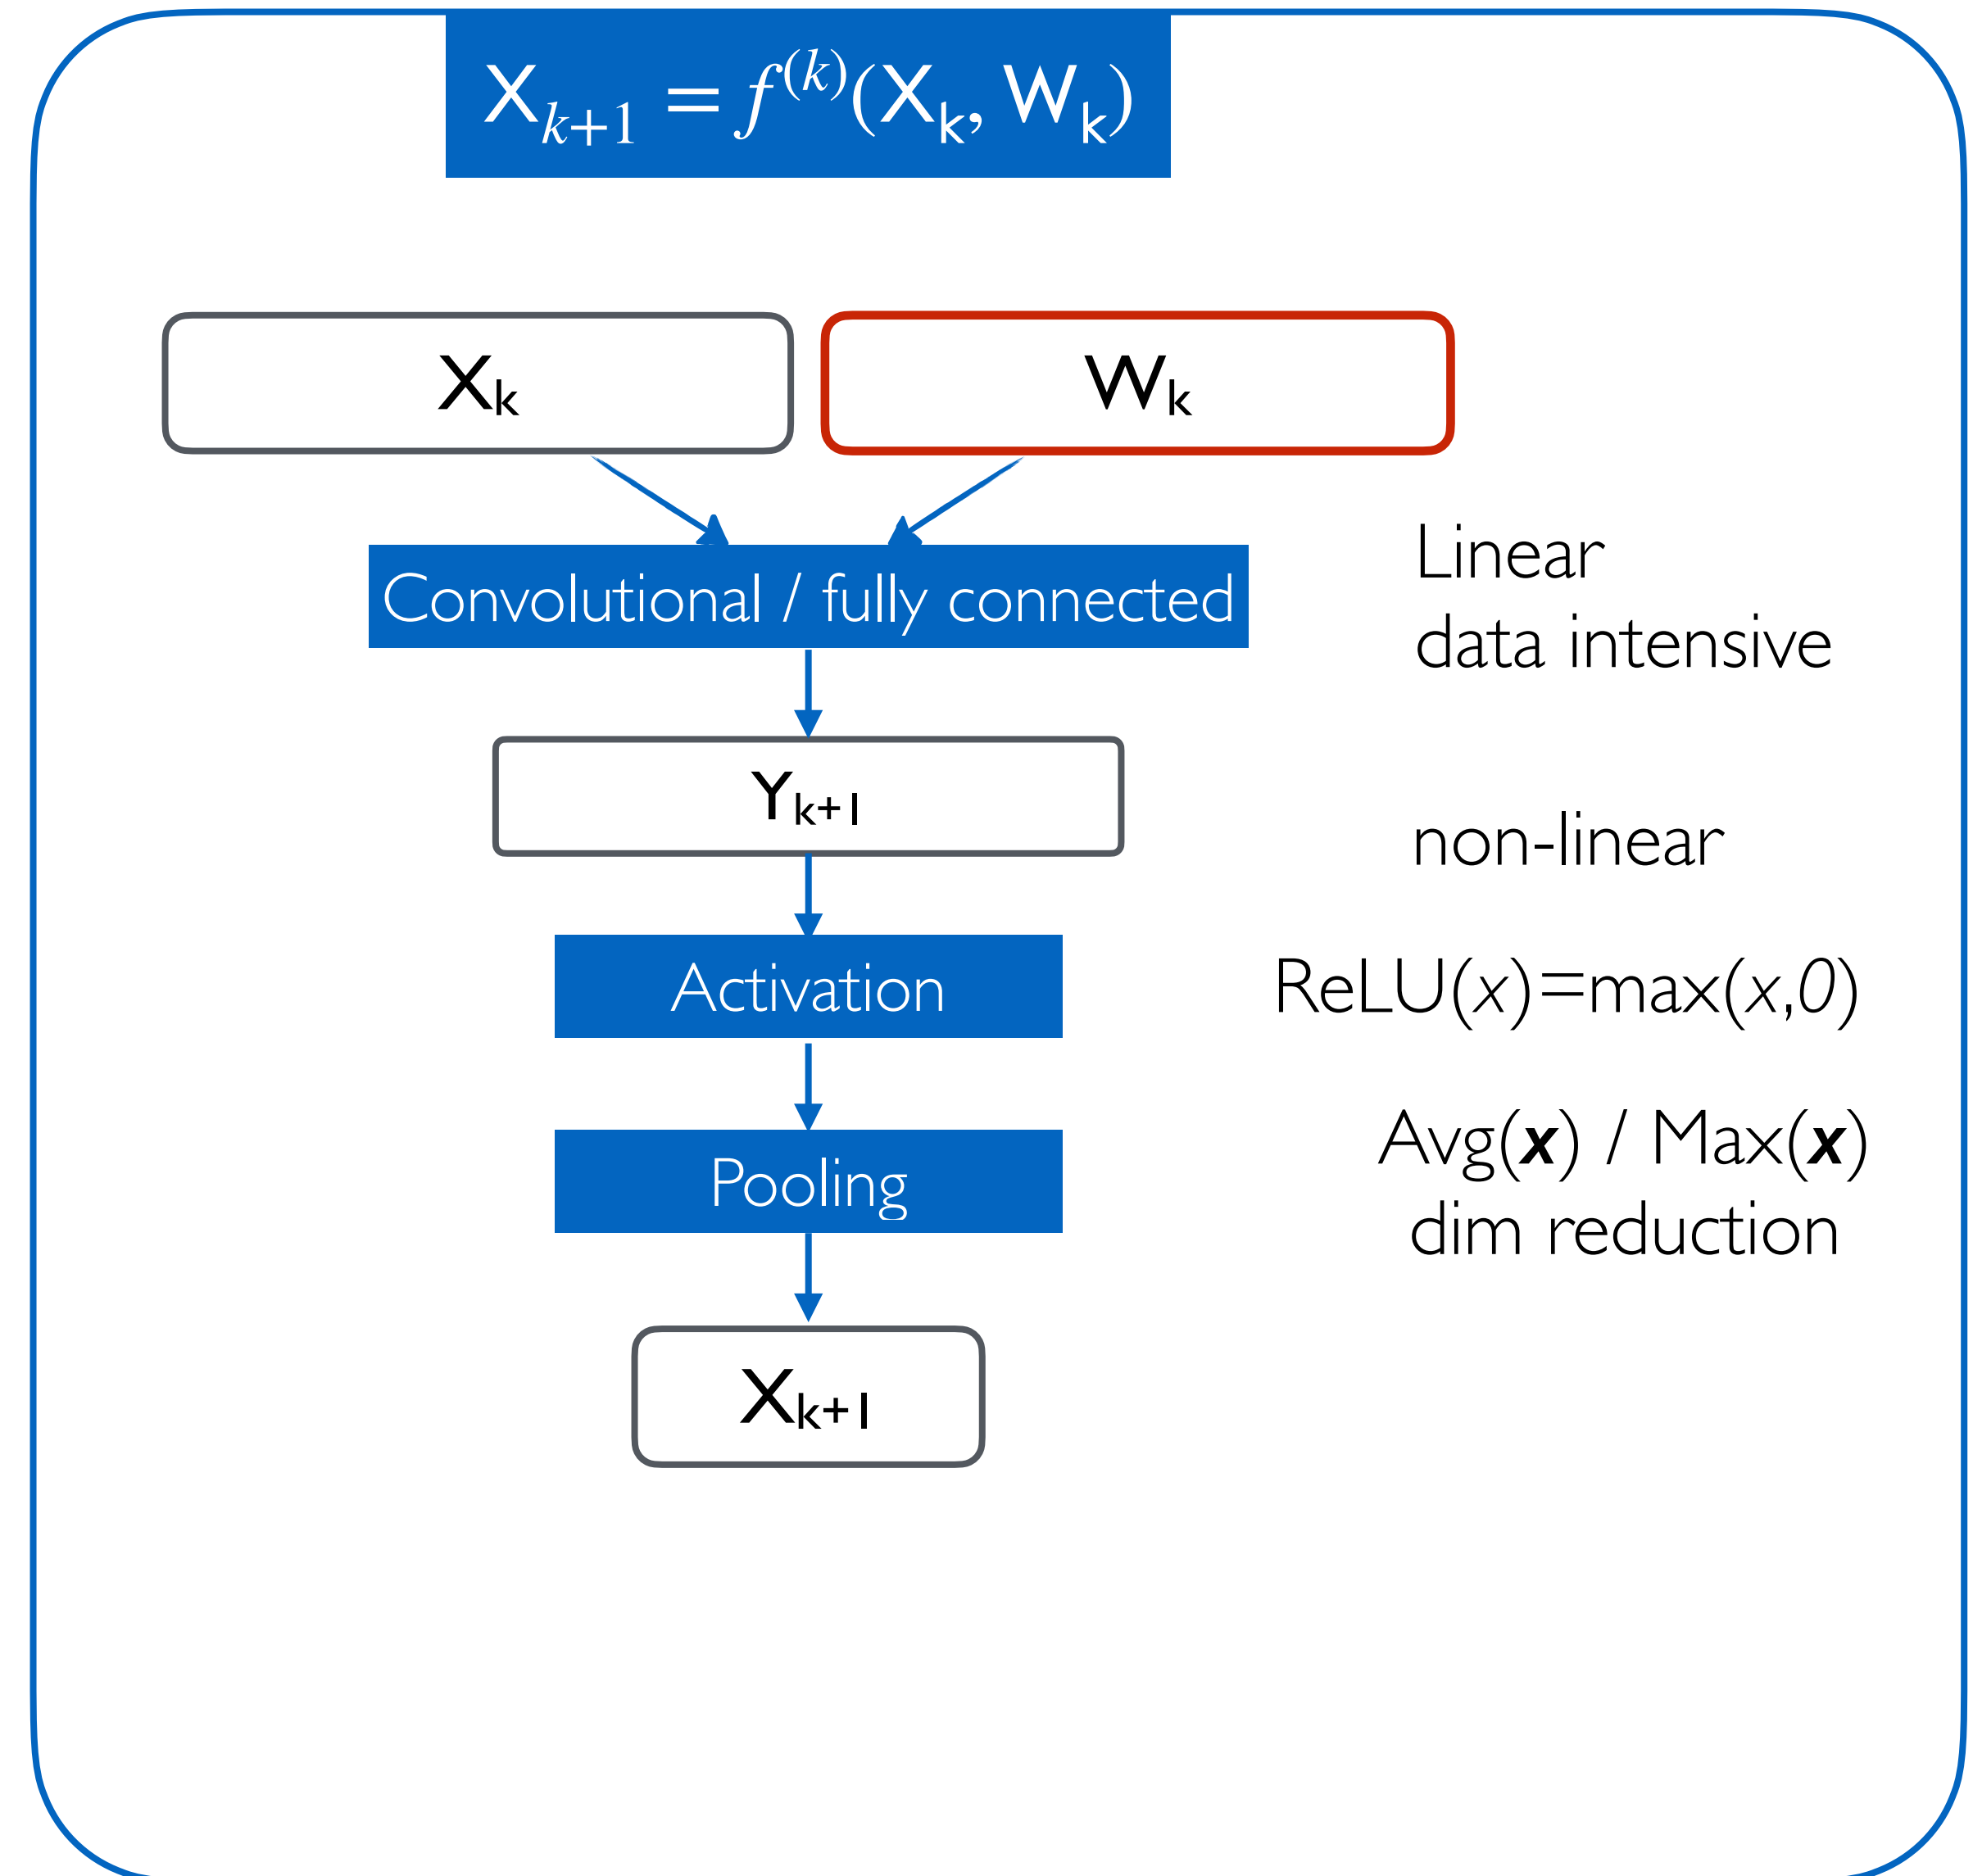
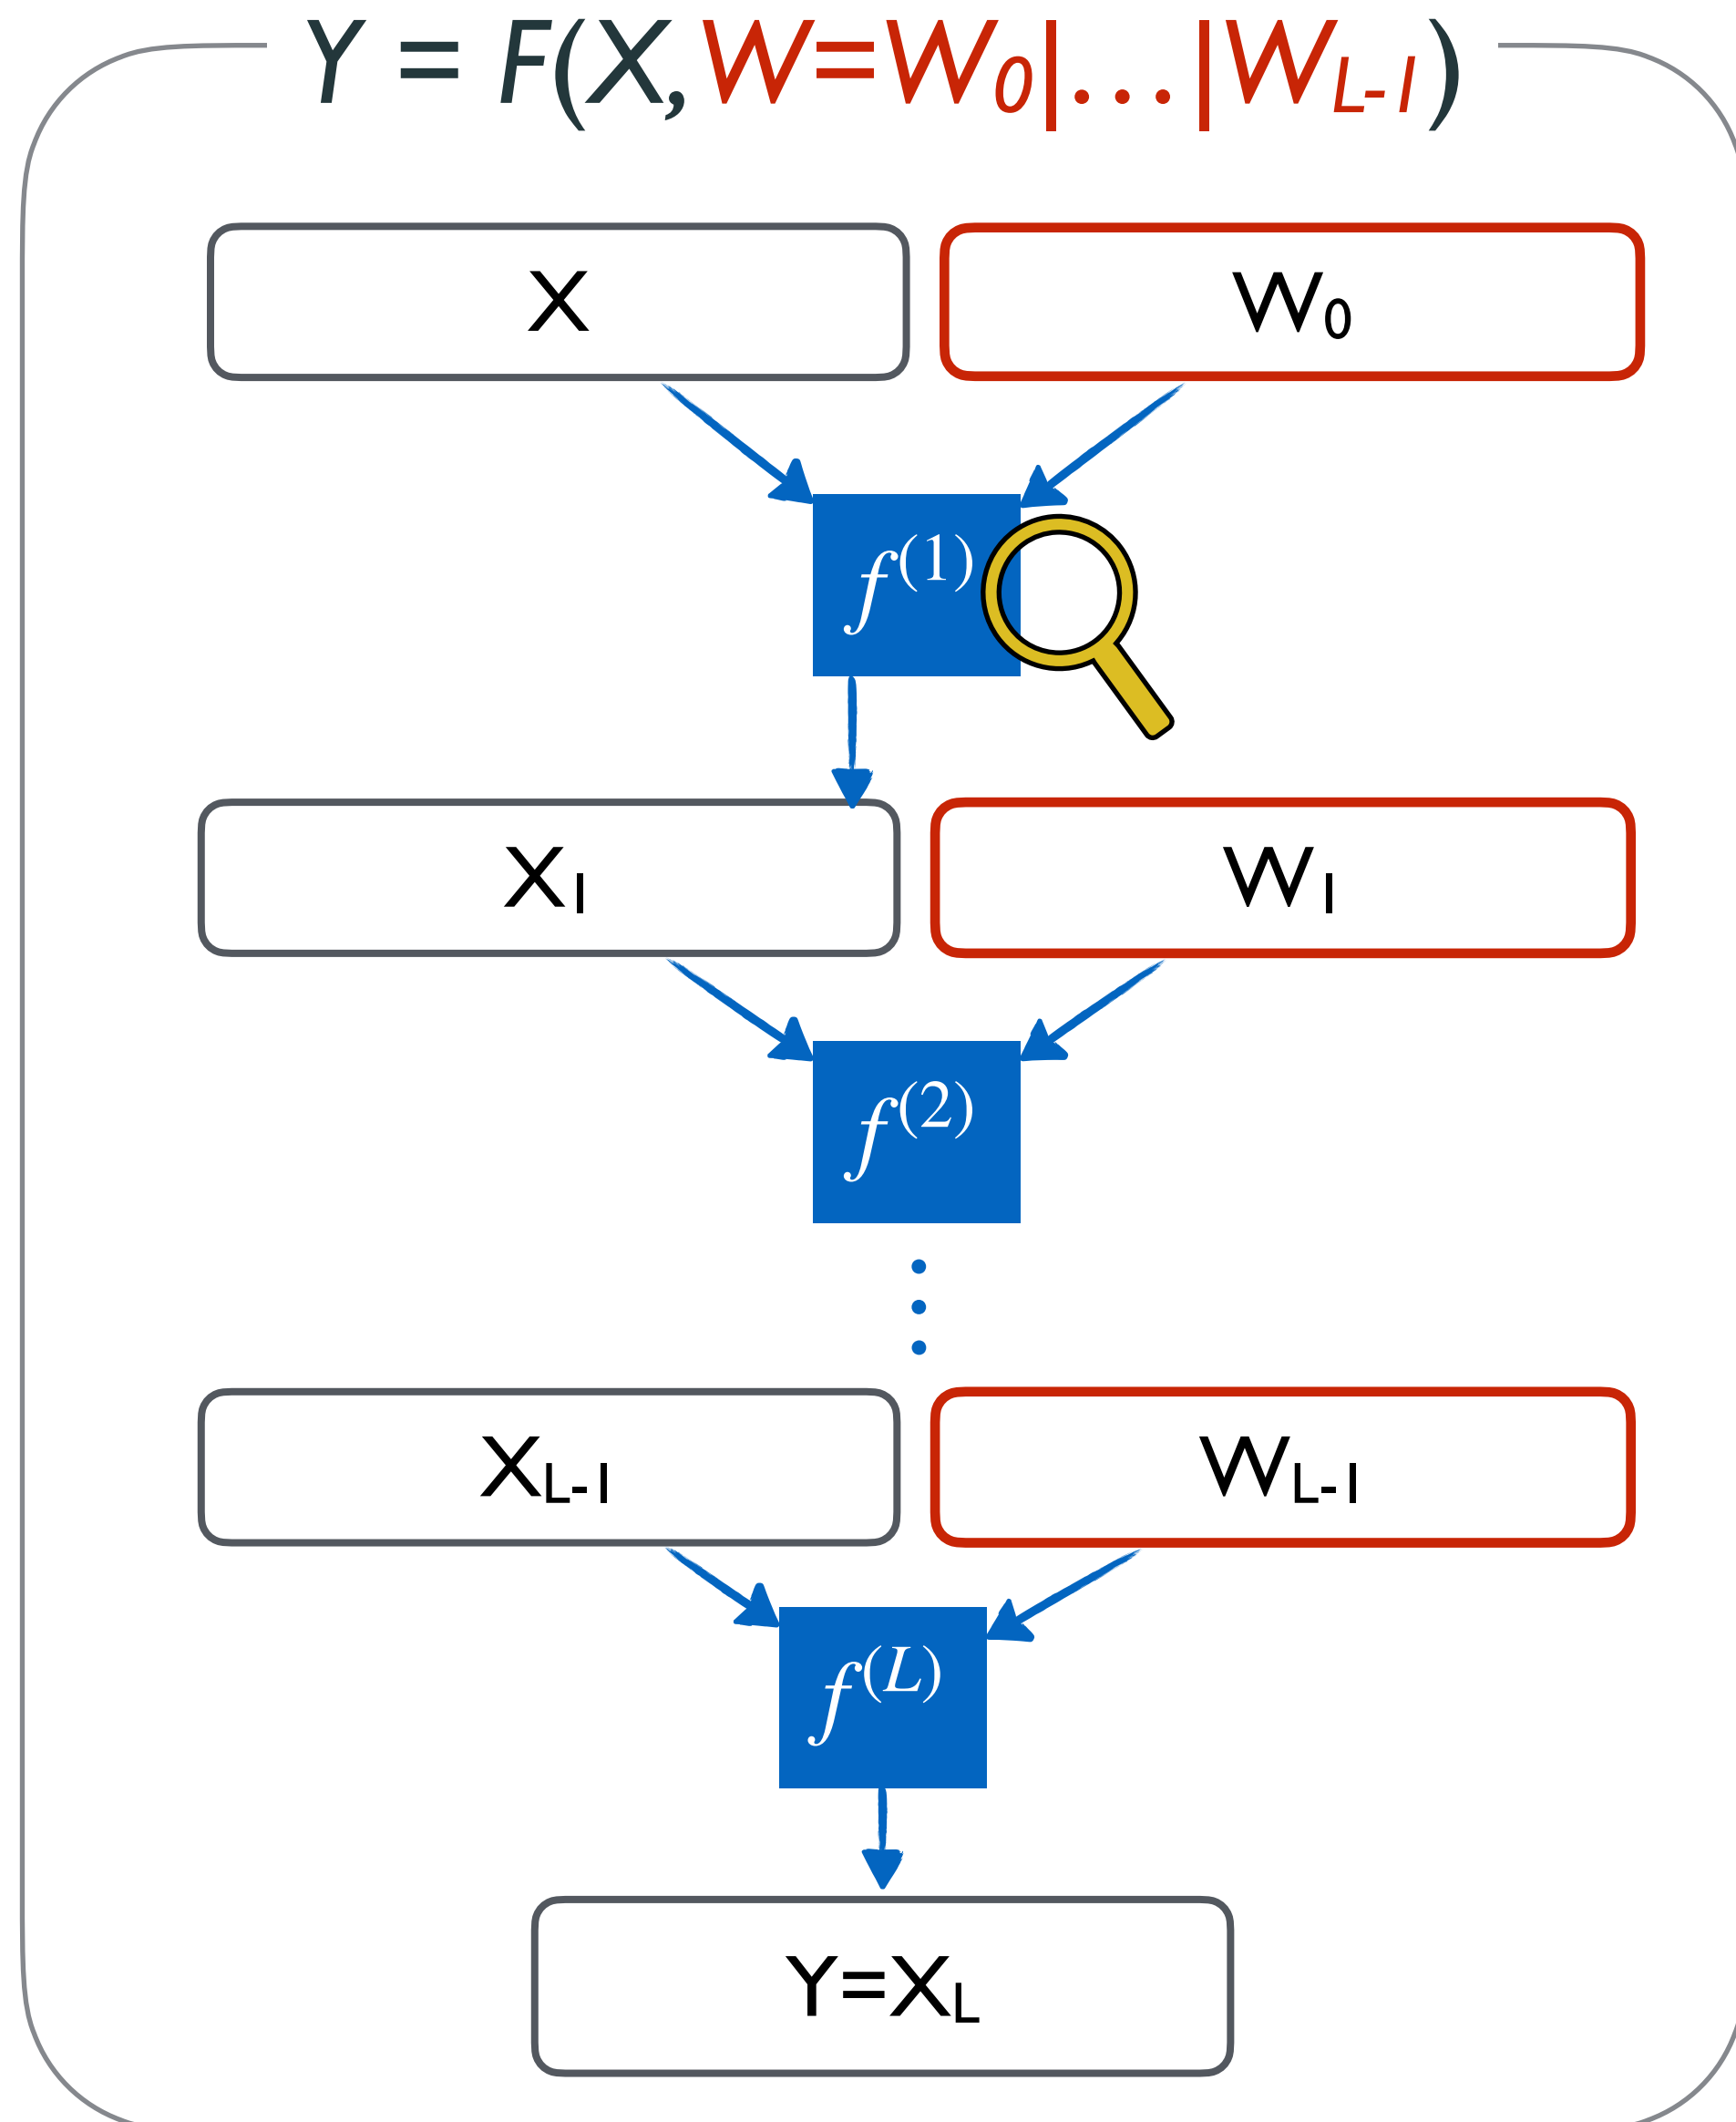
$$X_{k+1} = f^{(k)}(X_k, W_k)$$



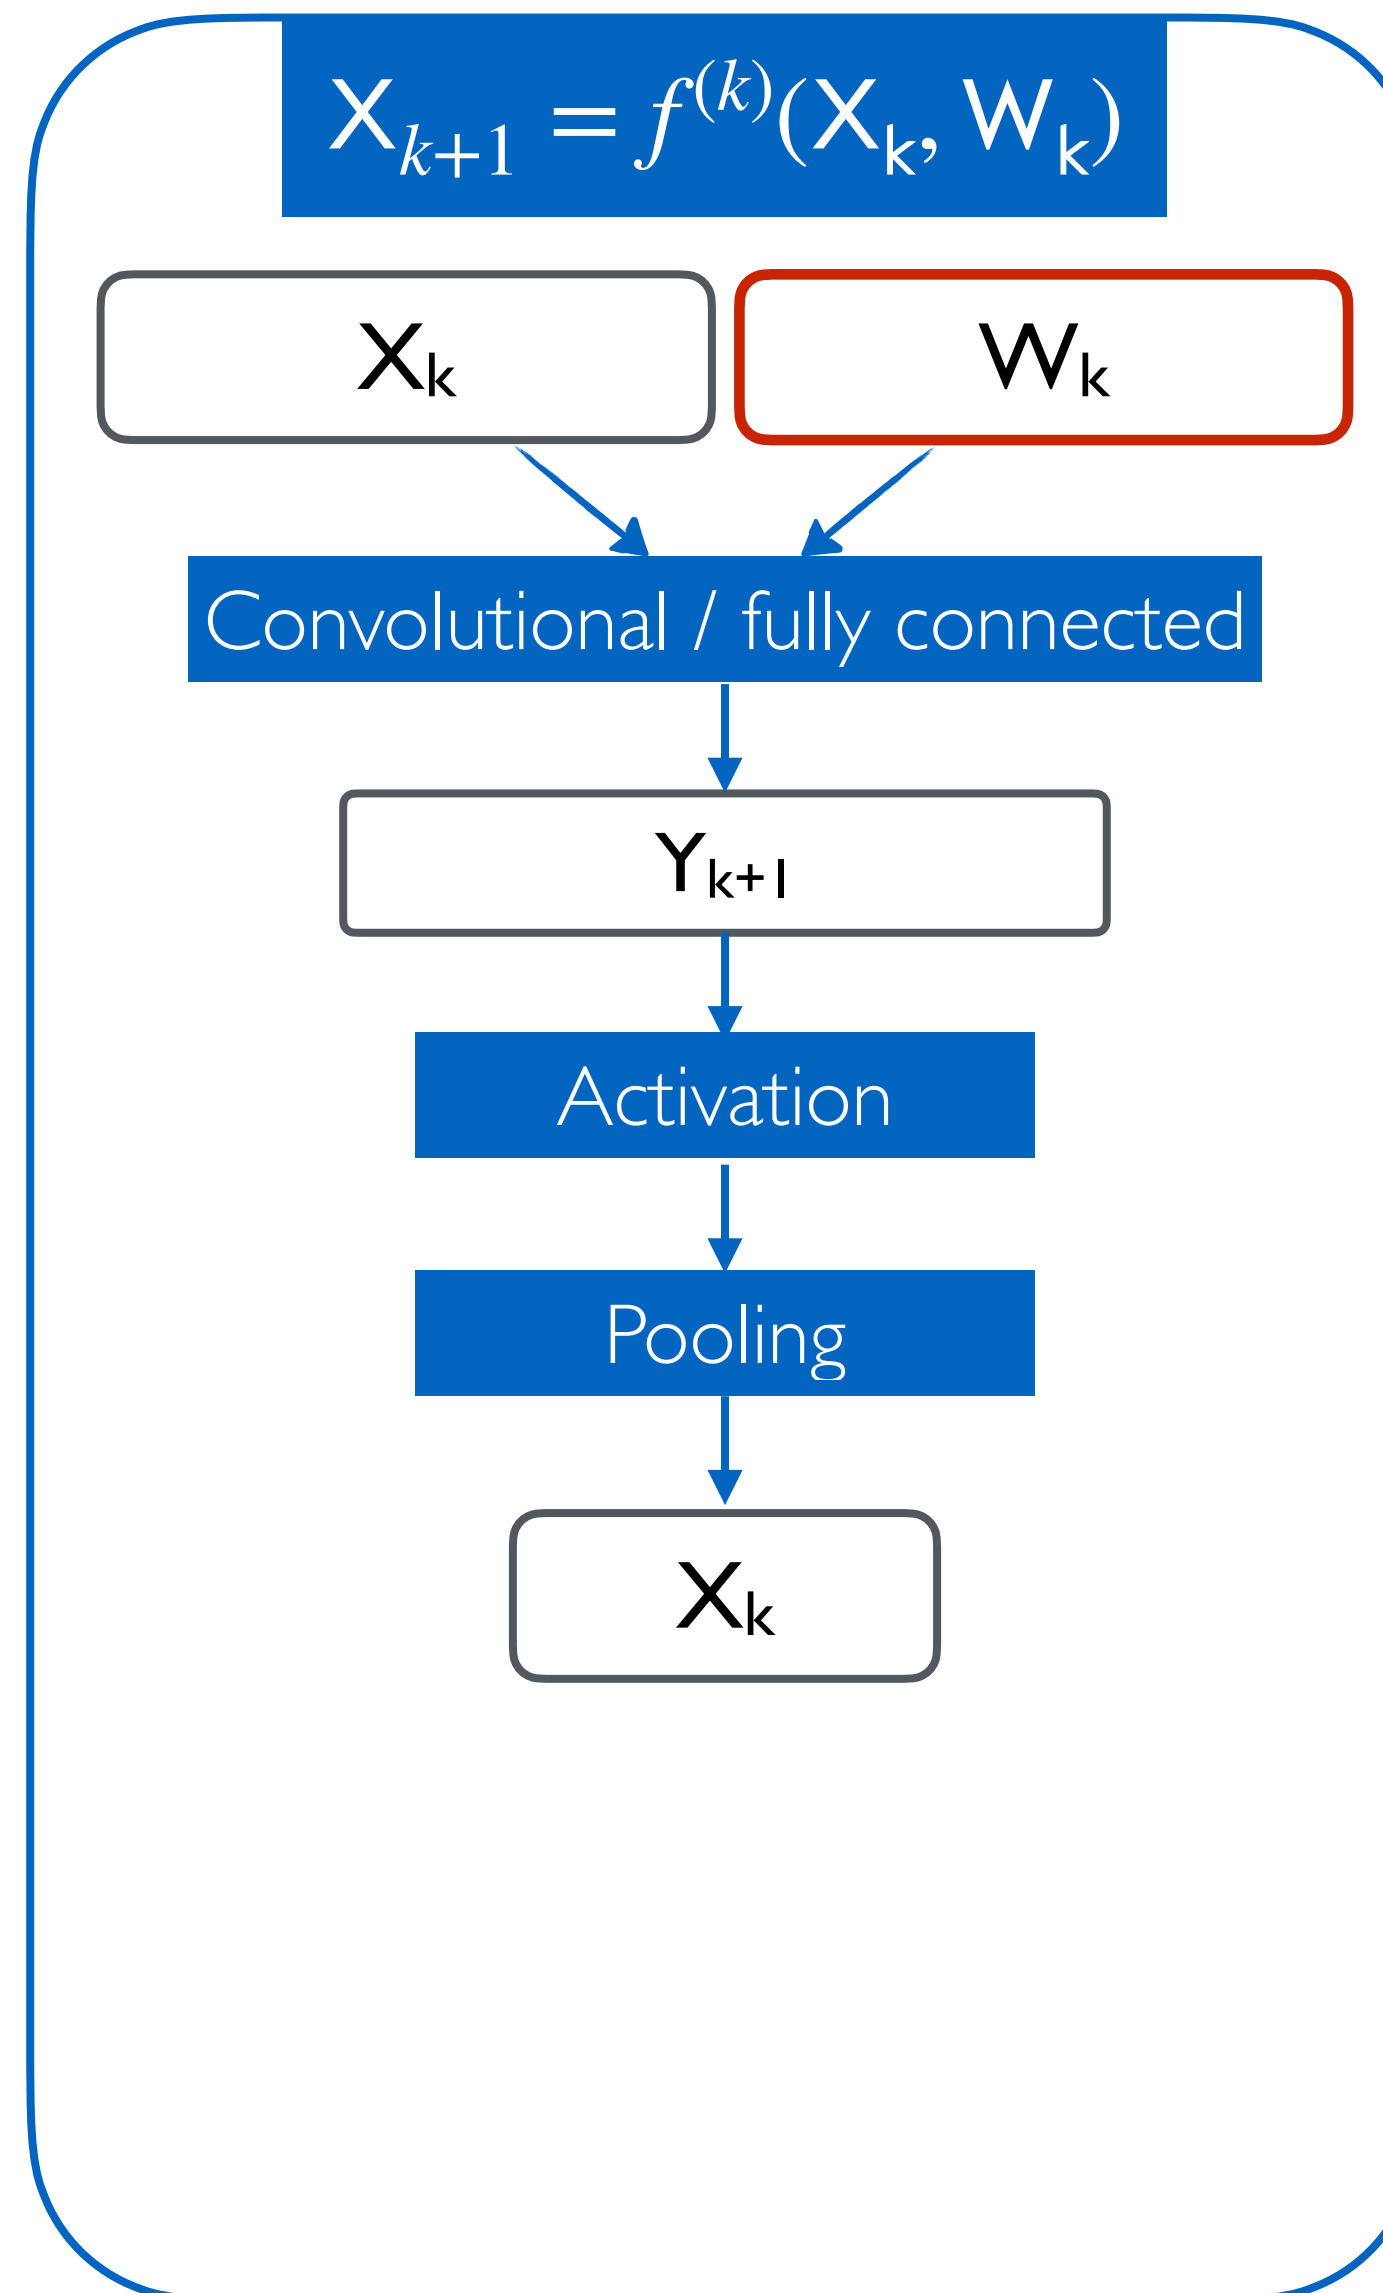
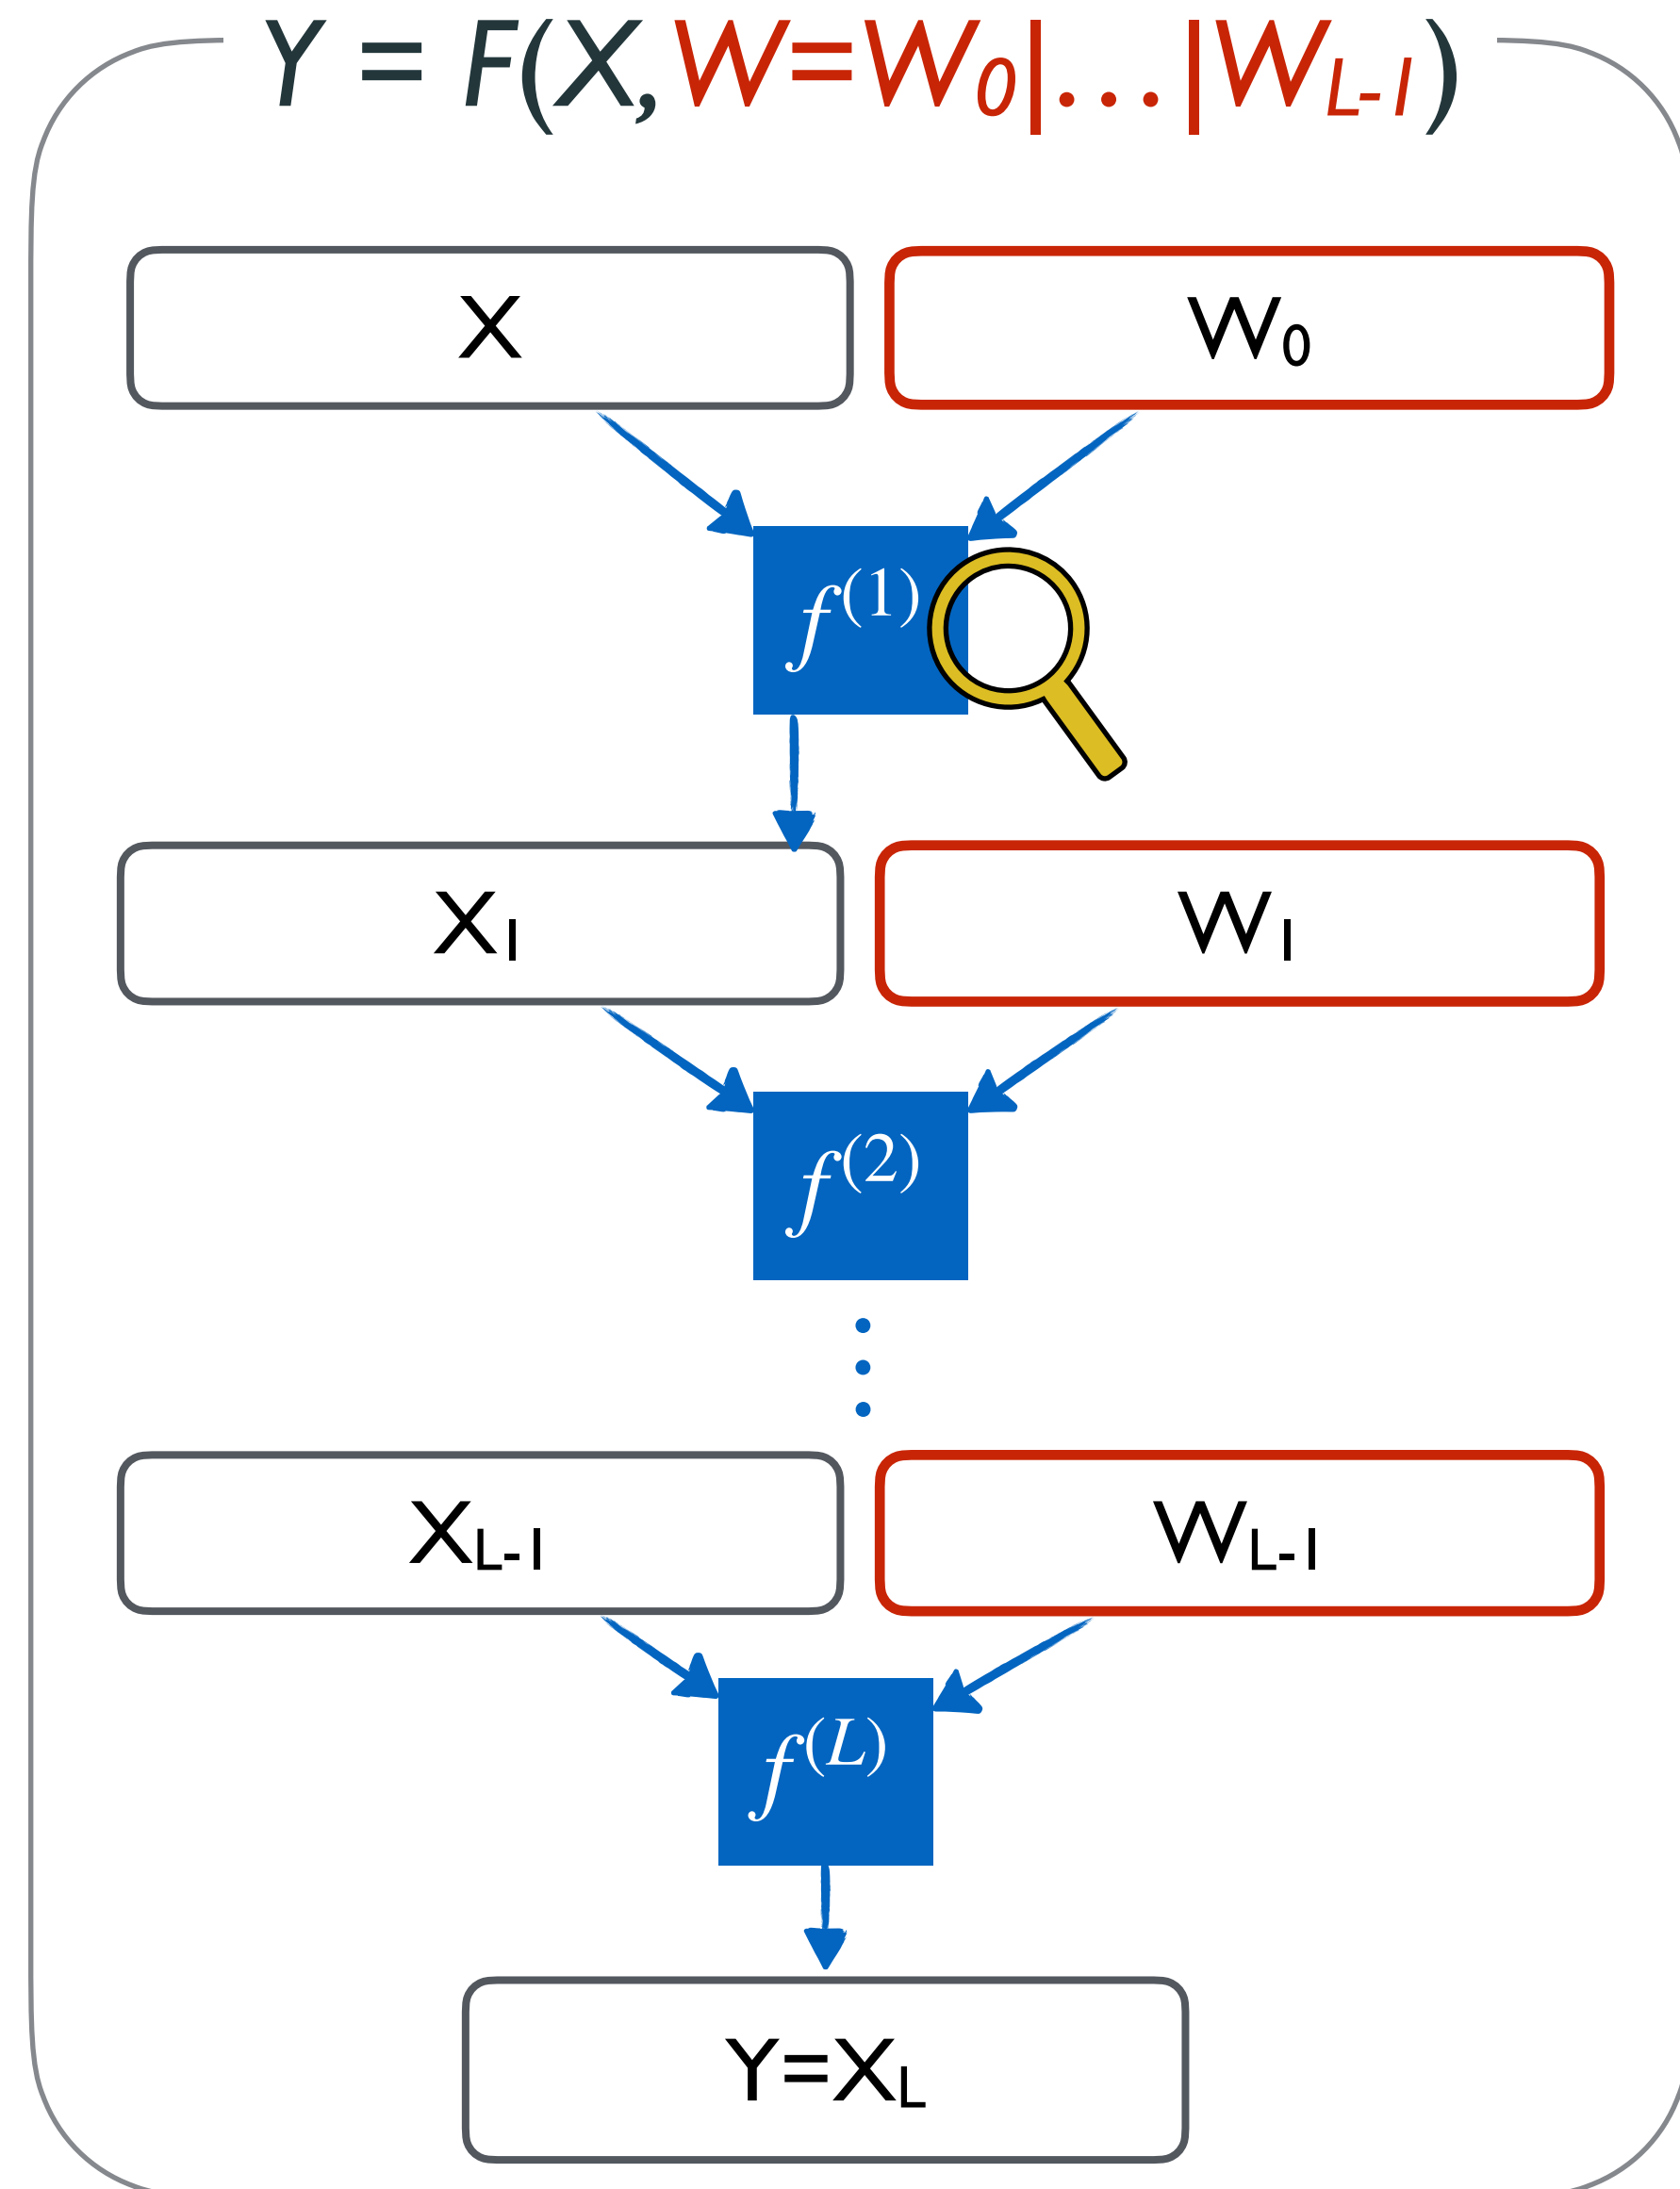
Convolutional Neural Networks



Convolutional Neural Networks



ZKPs for CNNs

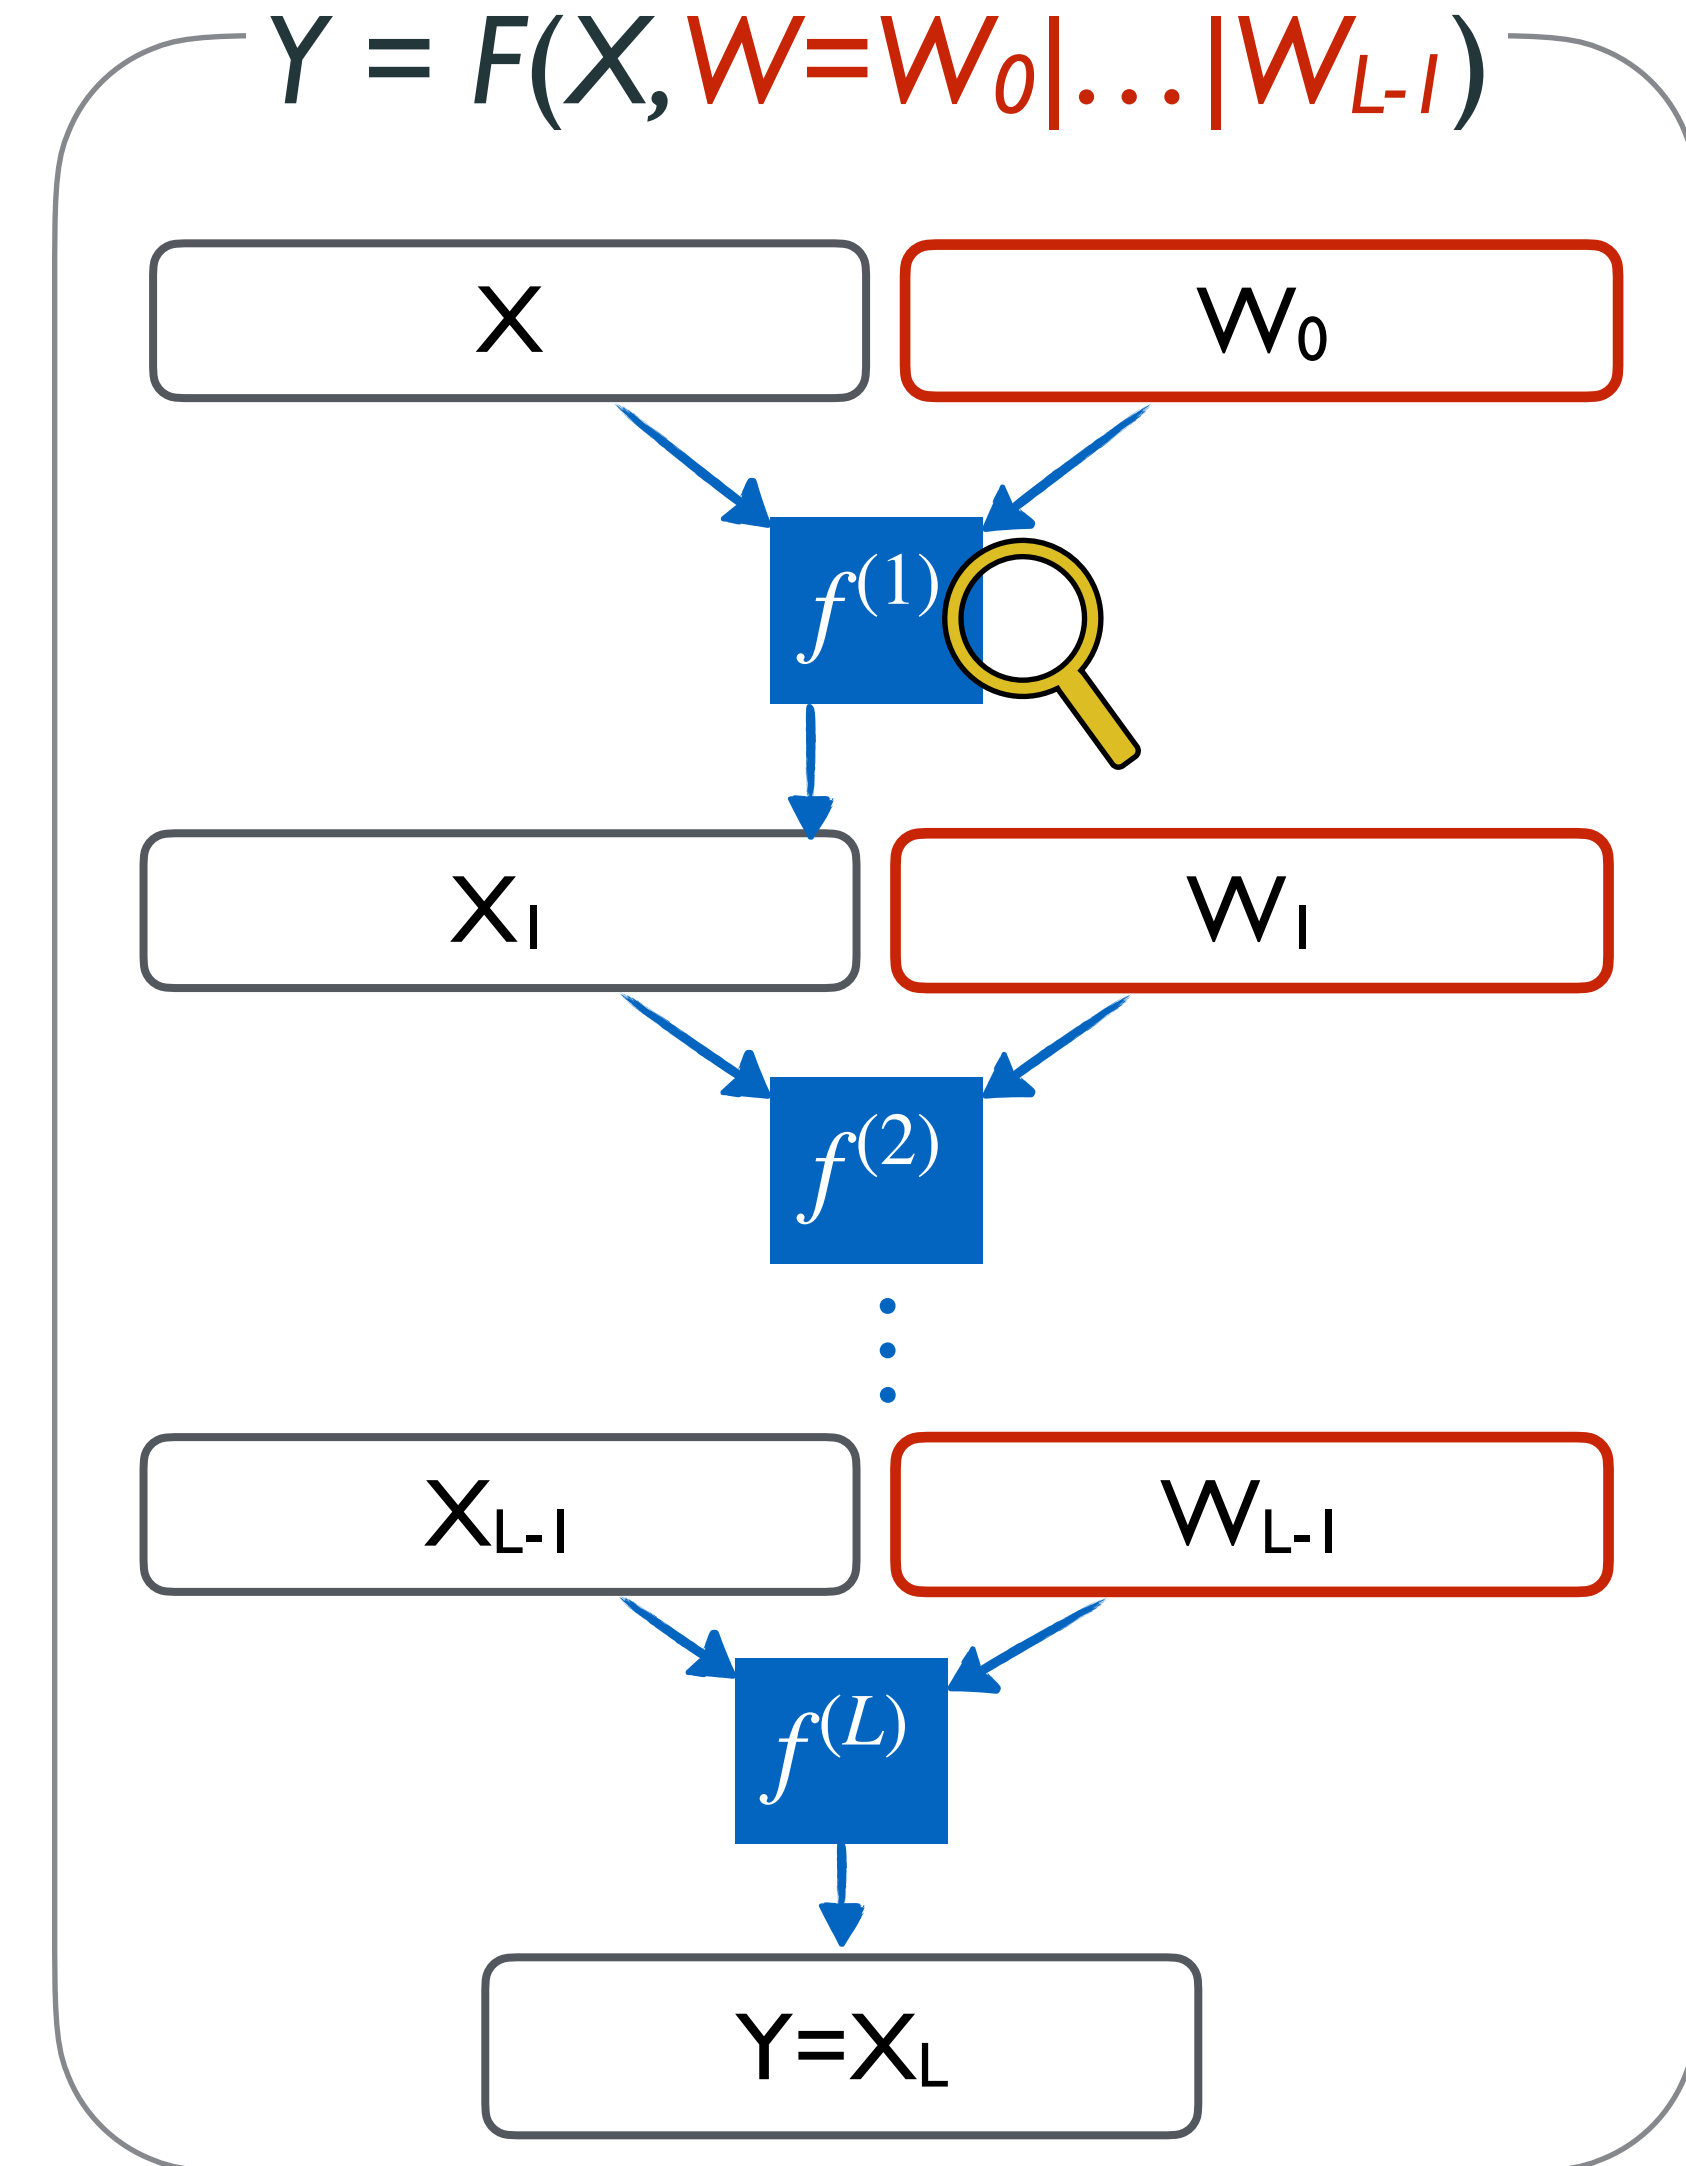
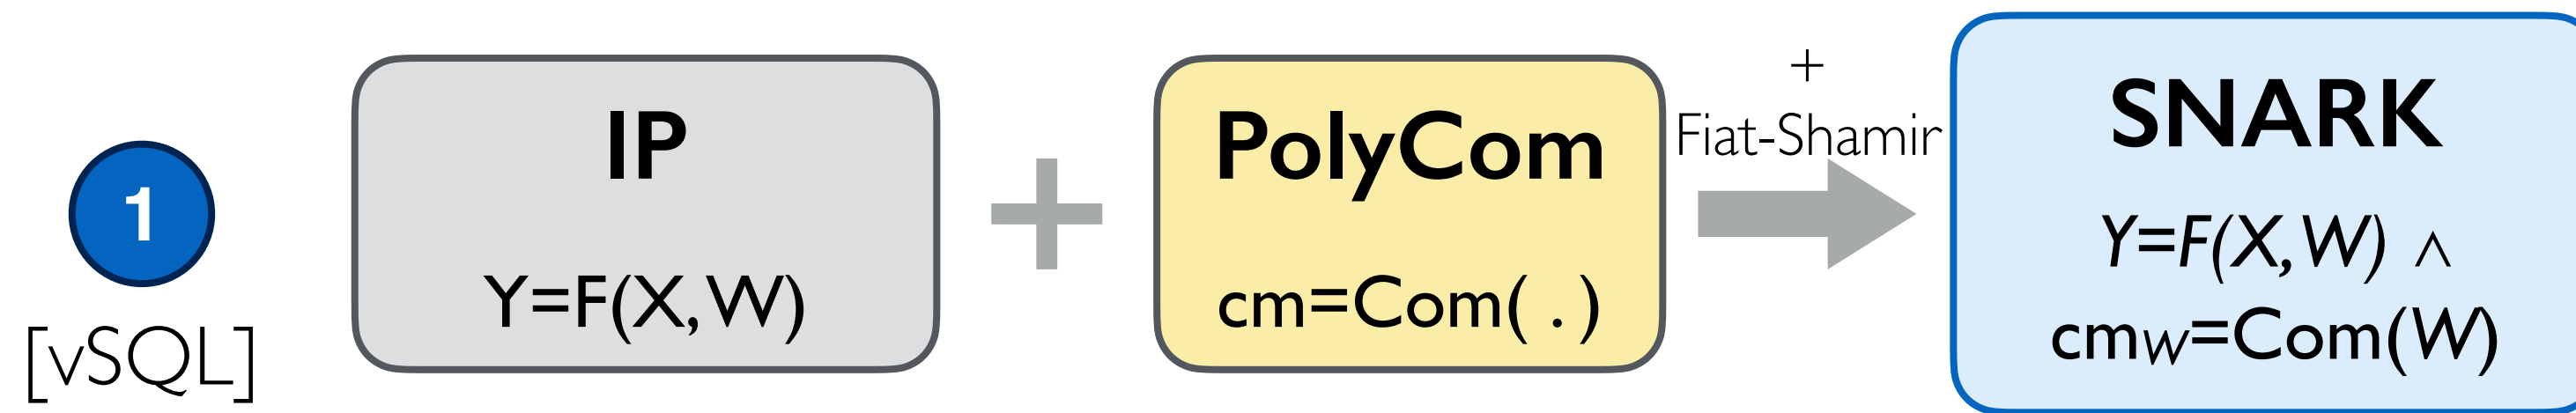


- Sumcheck-based proofs [LXZ21, BFGRS23]
- Suitable for layered computation
- Proof generation mostly information-theoretic. Cryptographic work is only $O(|X| + |W|)$

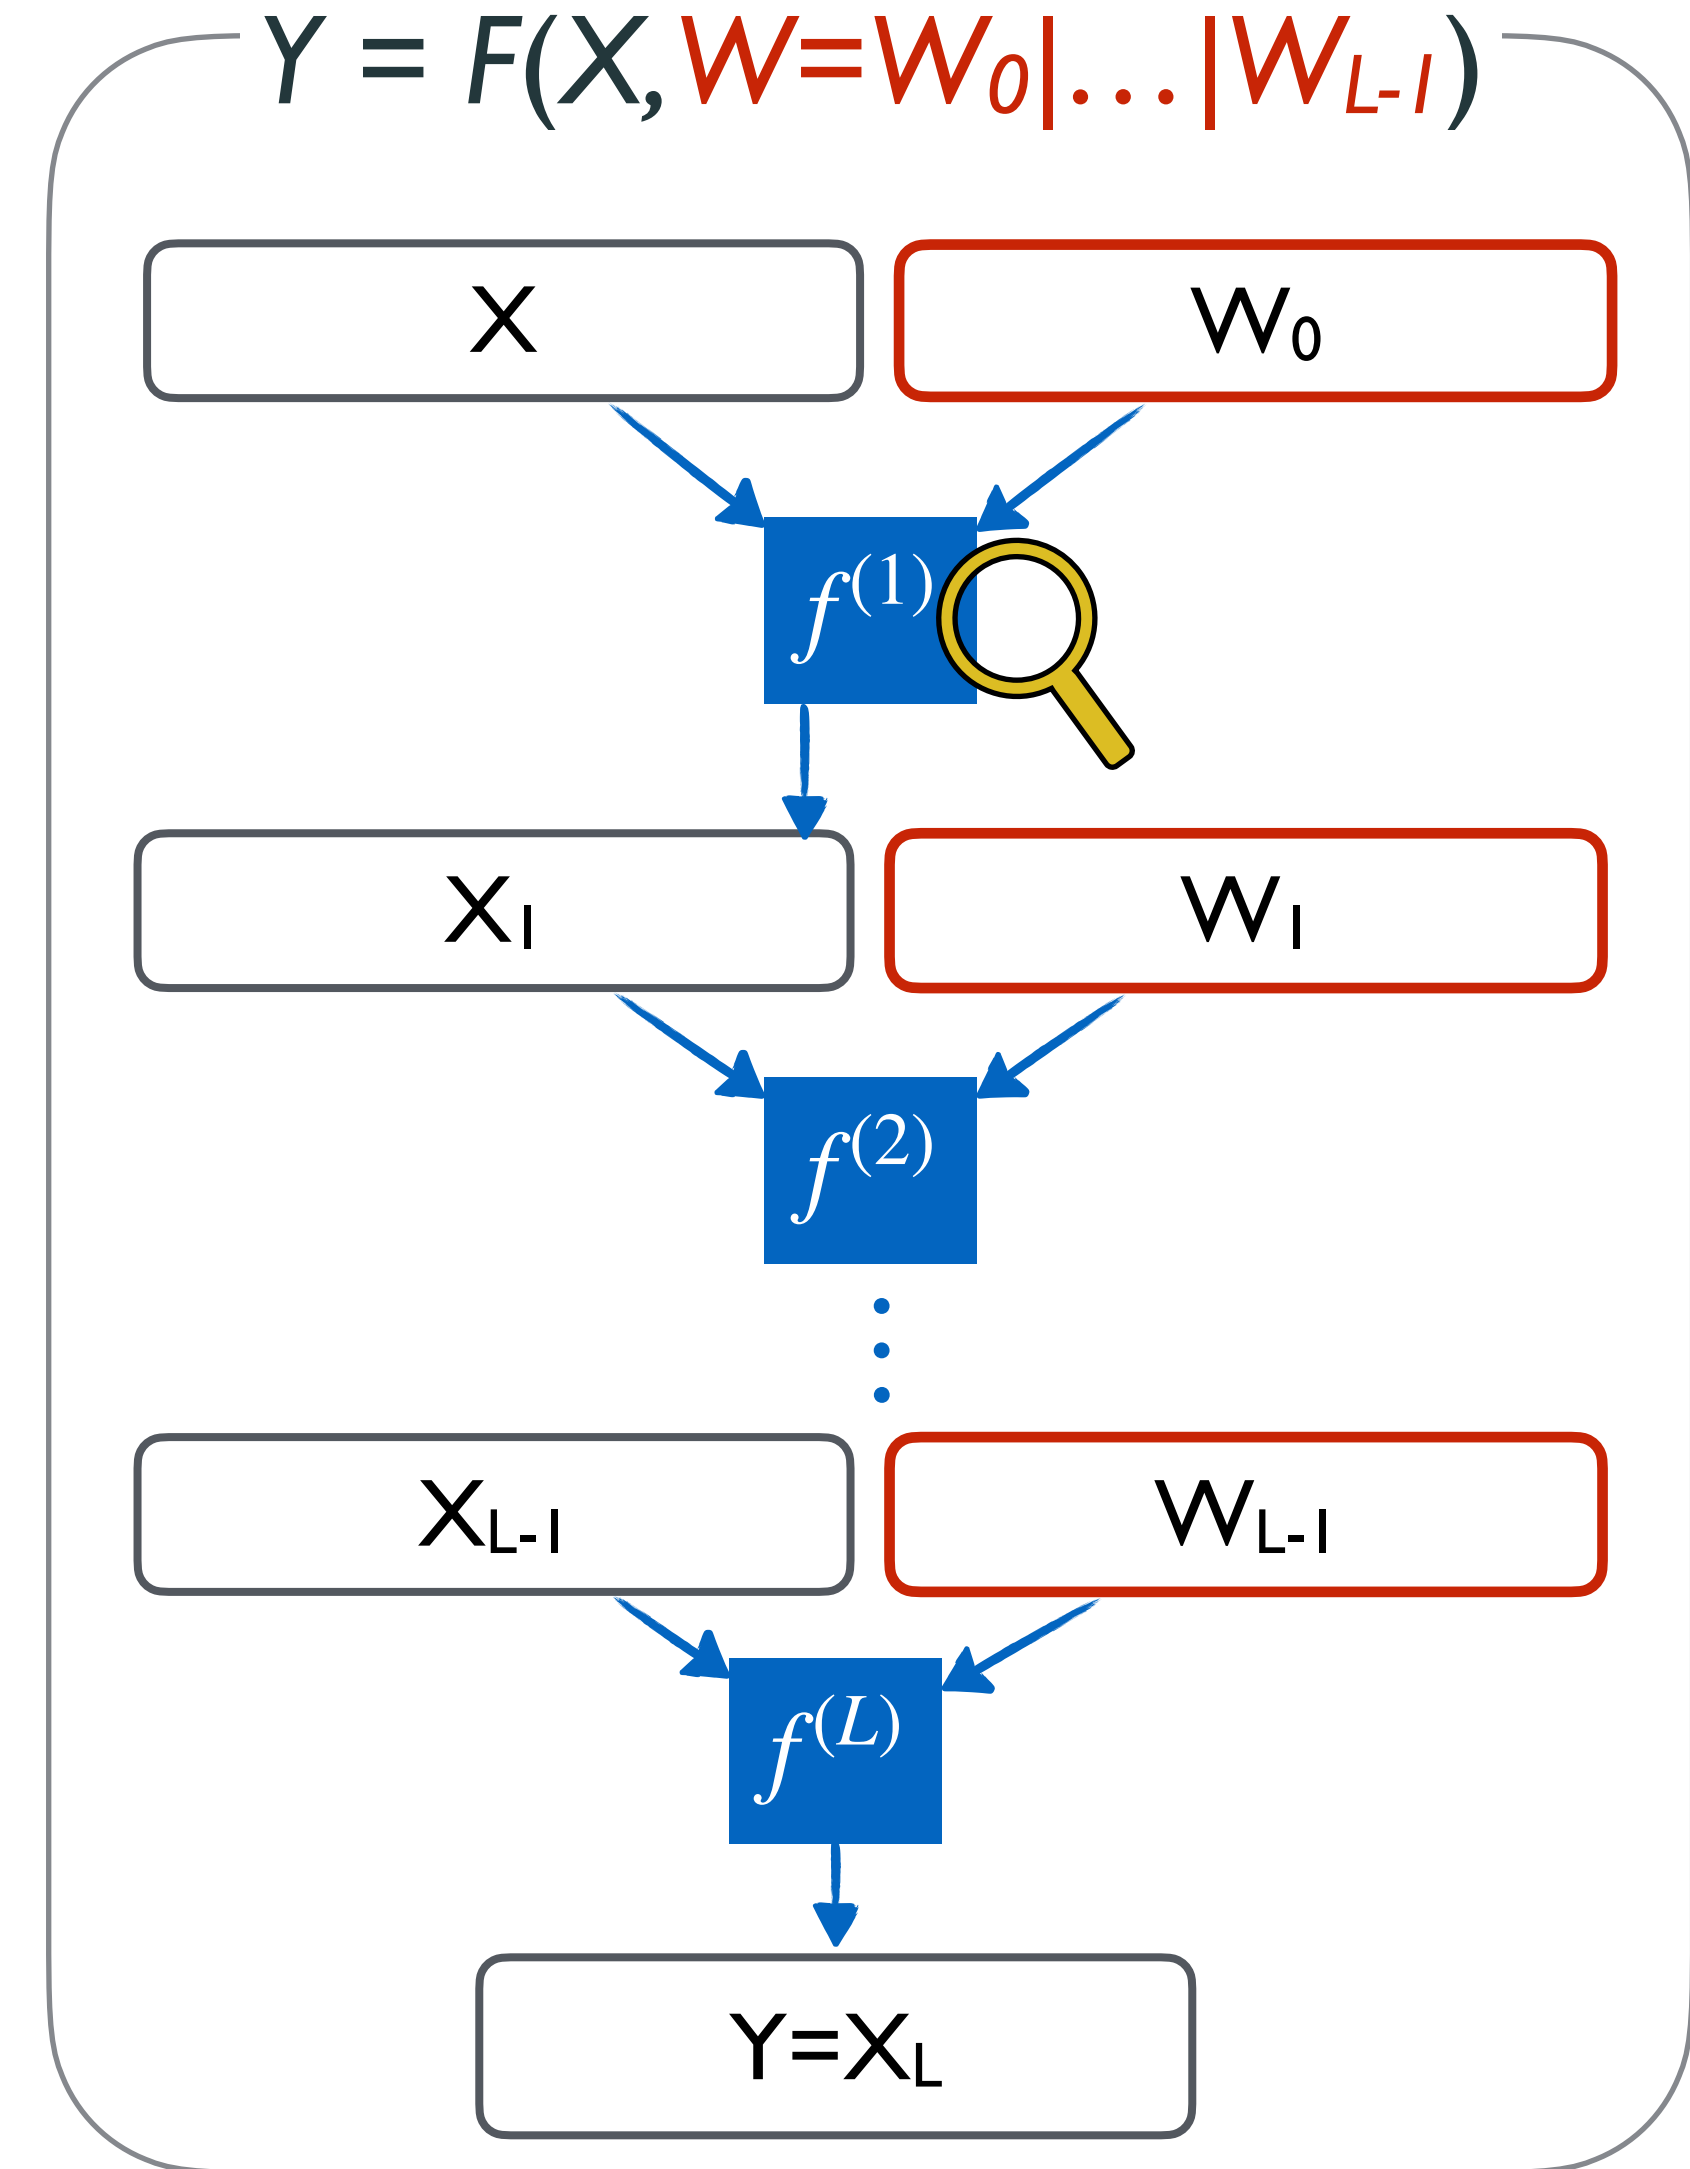
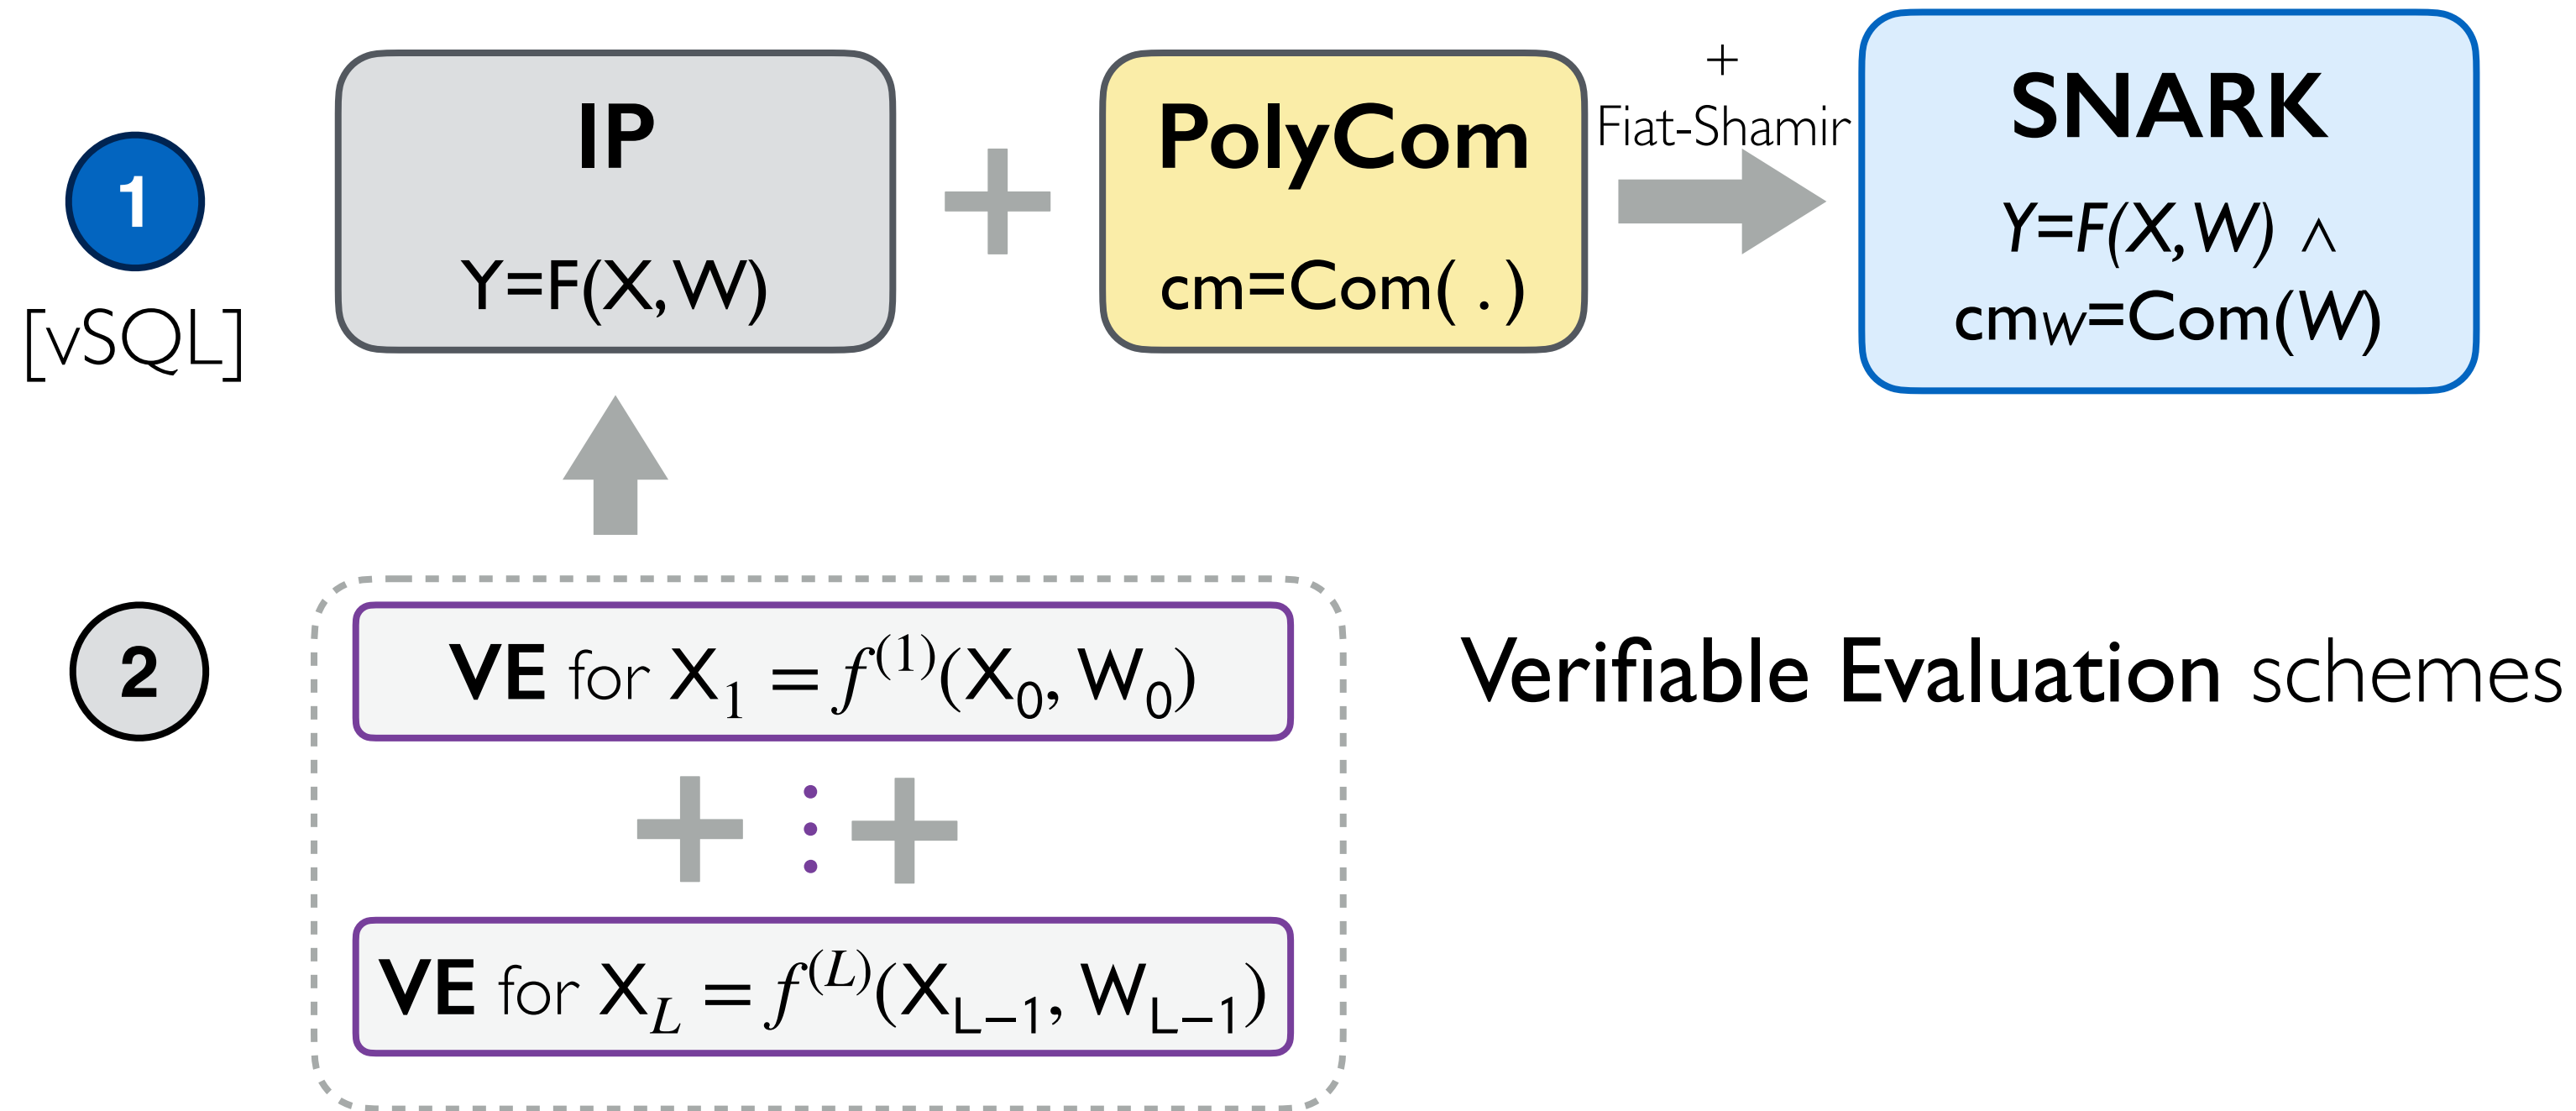
[LXZ21] Liu, Xi, Zhang. *zkCNN: Zero Knowledge Proofs for Convolutional Neural Network Predictions and Accuracy*. CCS 2021

[BFGRS23] Balbás, Fiore, Gonzalez-Vasco, Robissout, Soriente. *Modular Sumcheck Proofs with Applications to Machine Learning and Image Processing*. CCS 2023

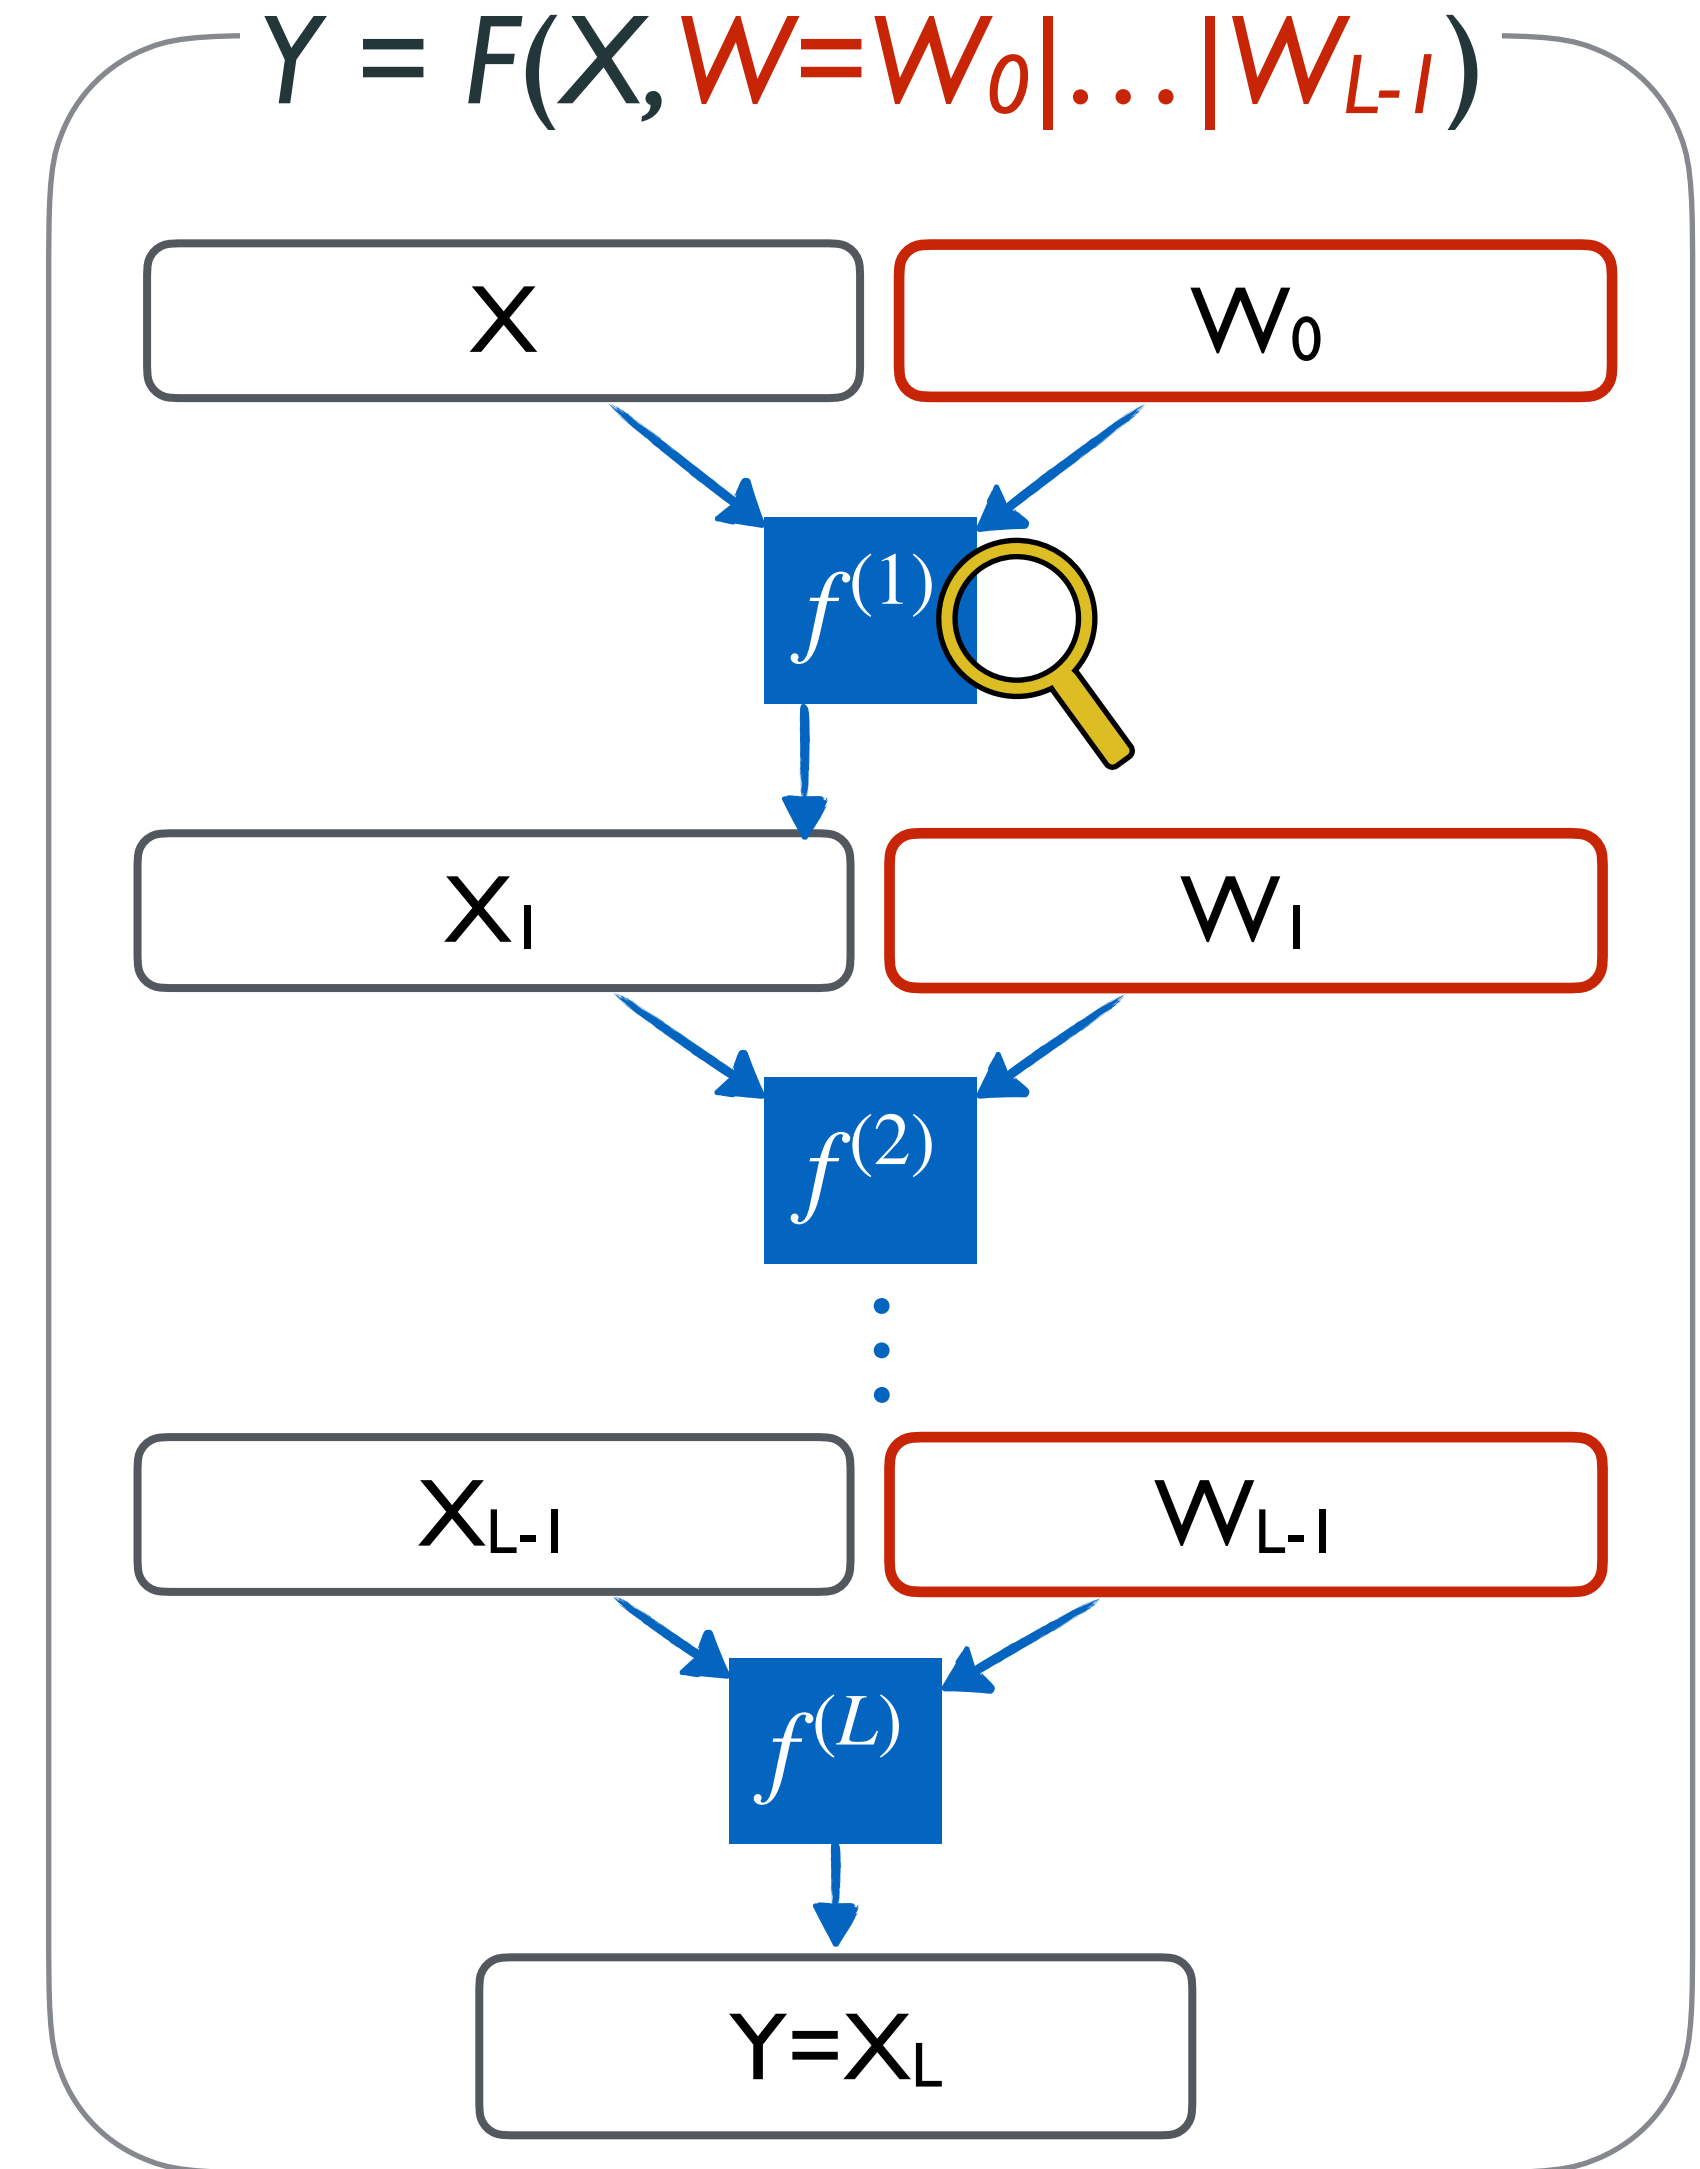
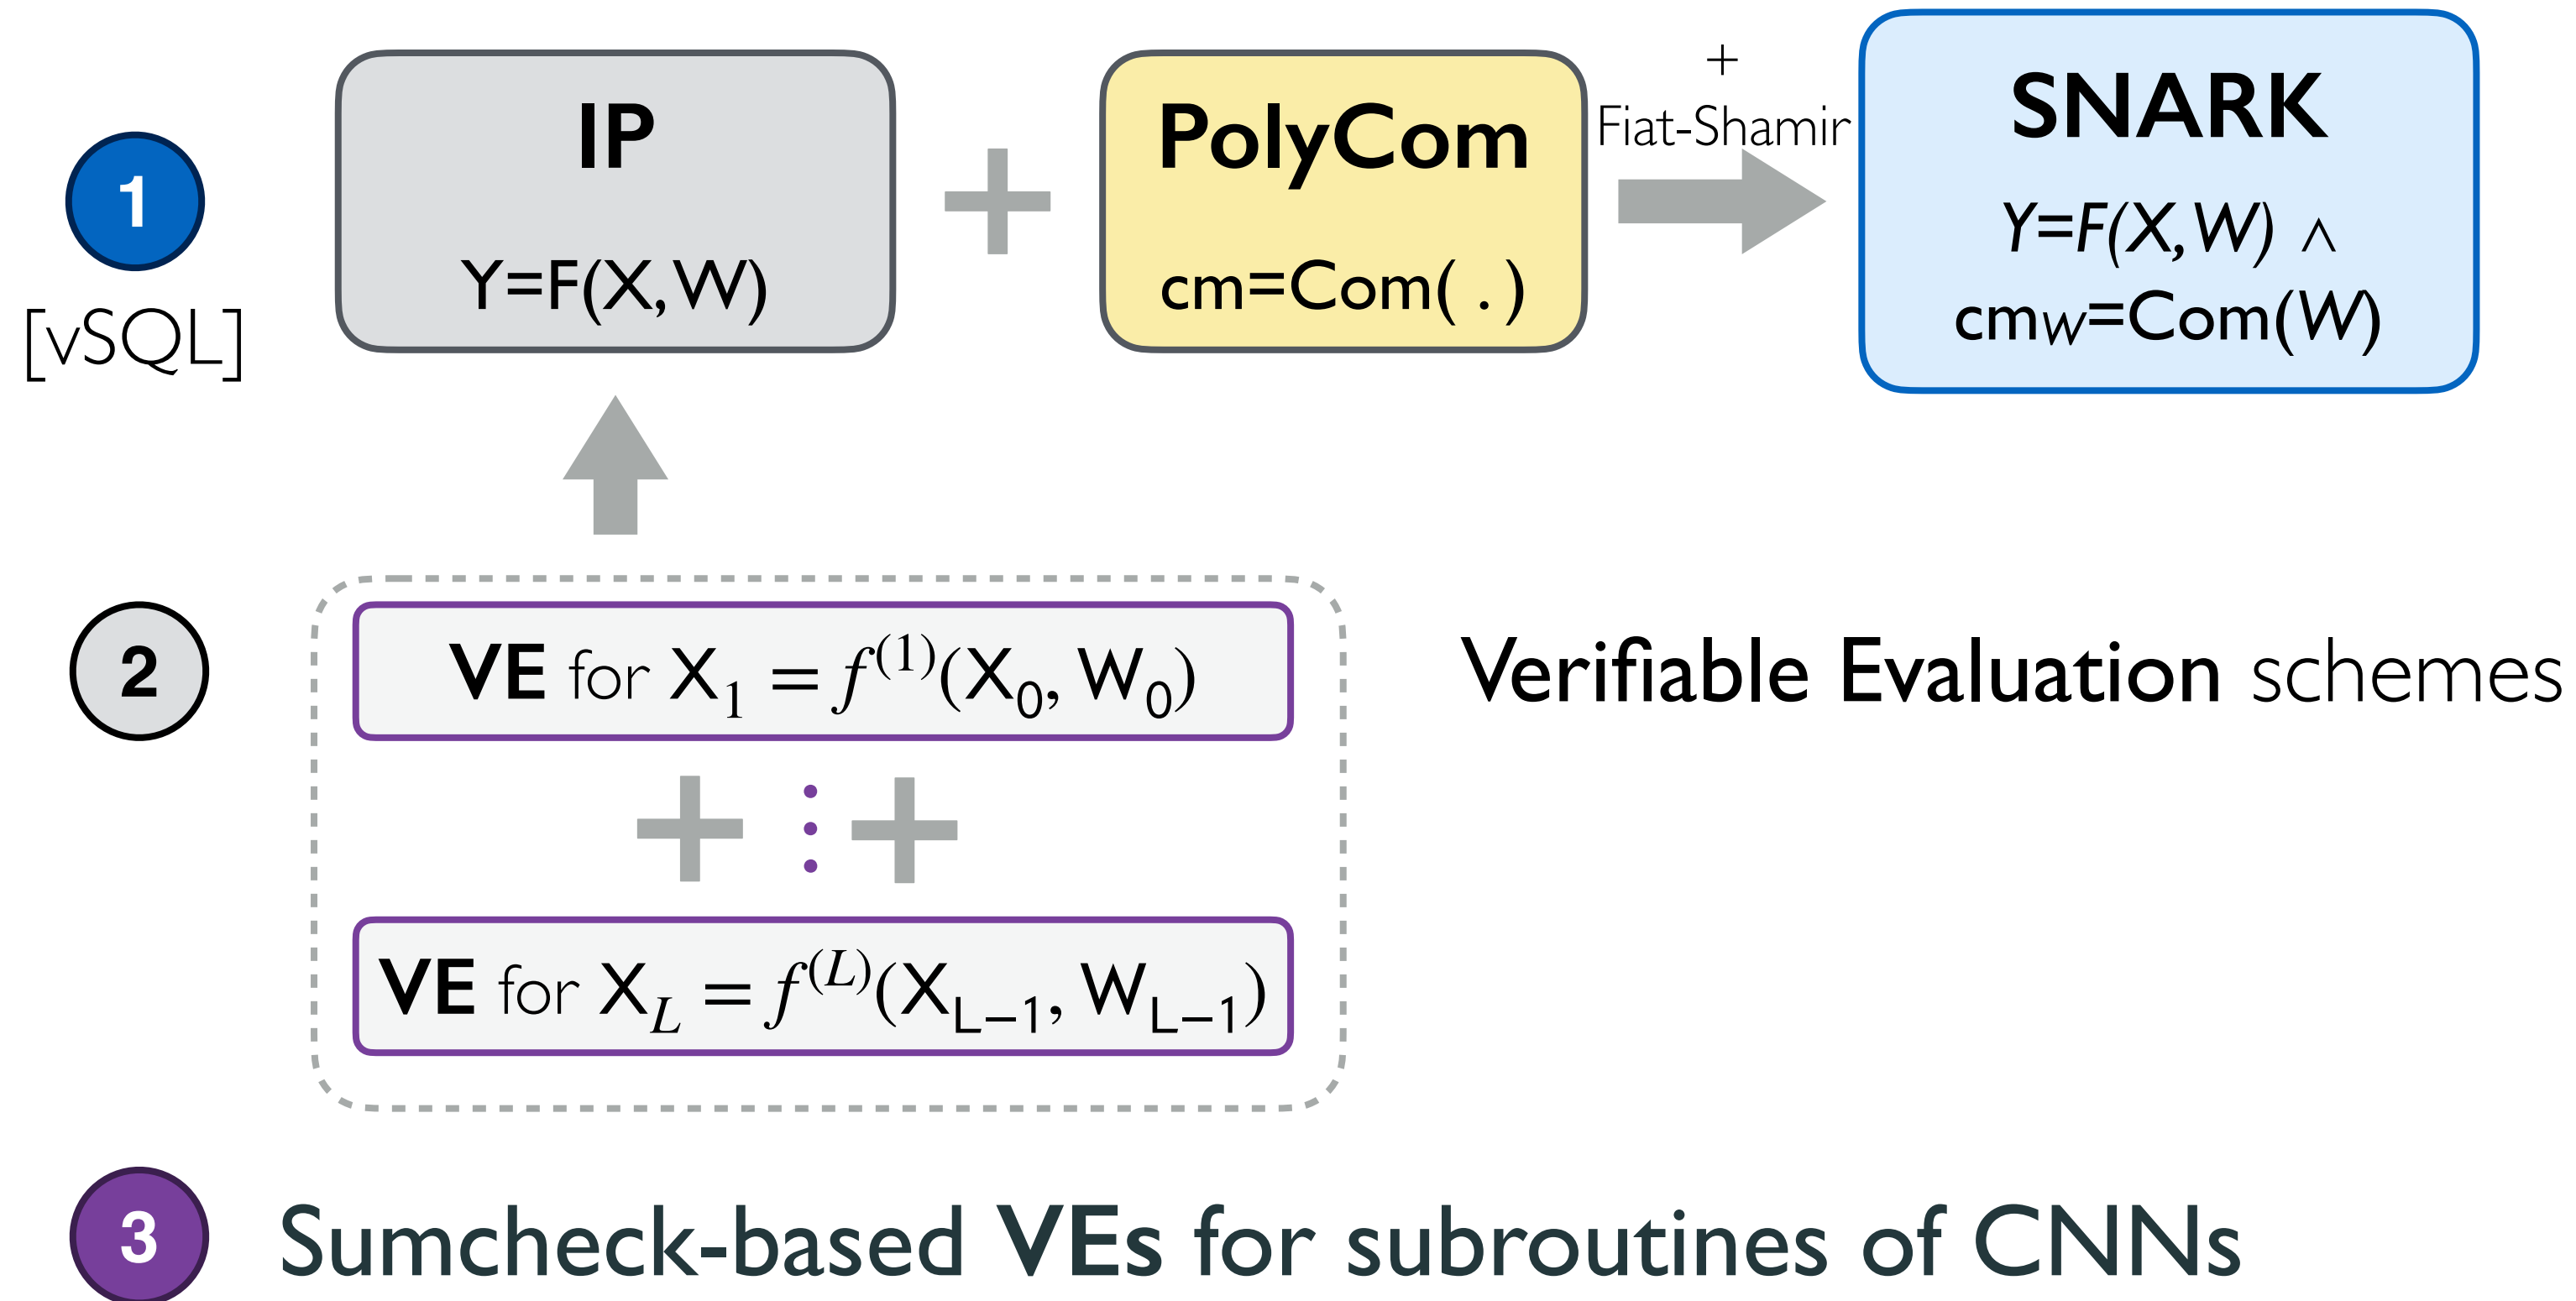
Modular approach for CNNs [BFGRS23]



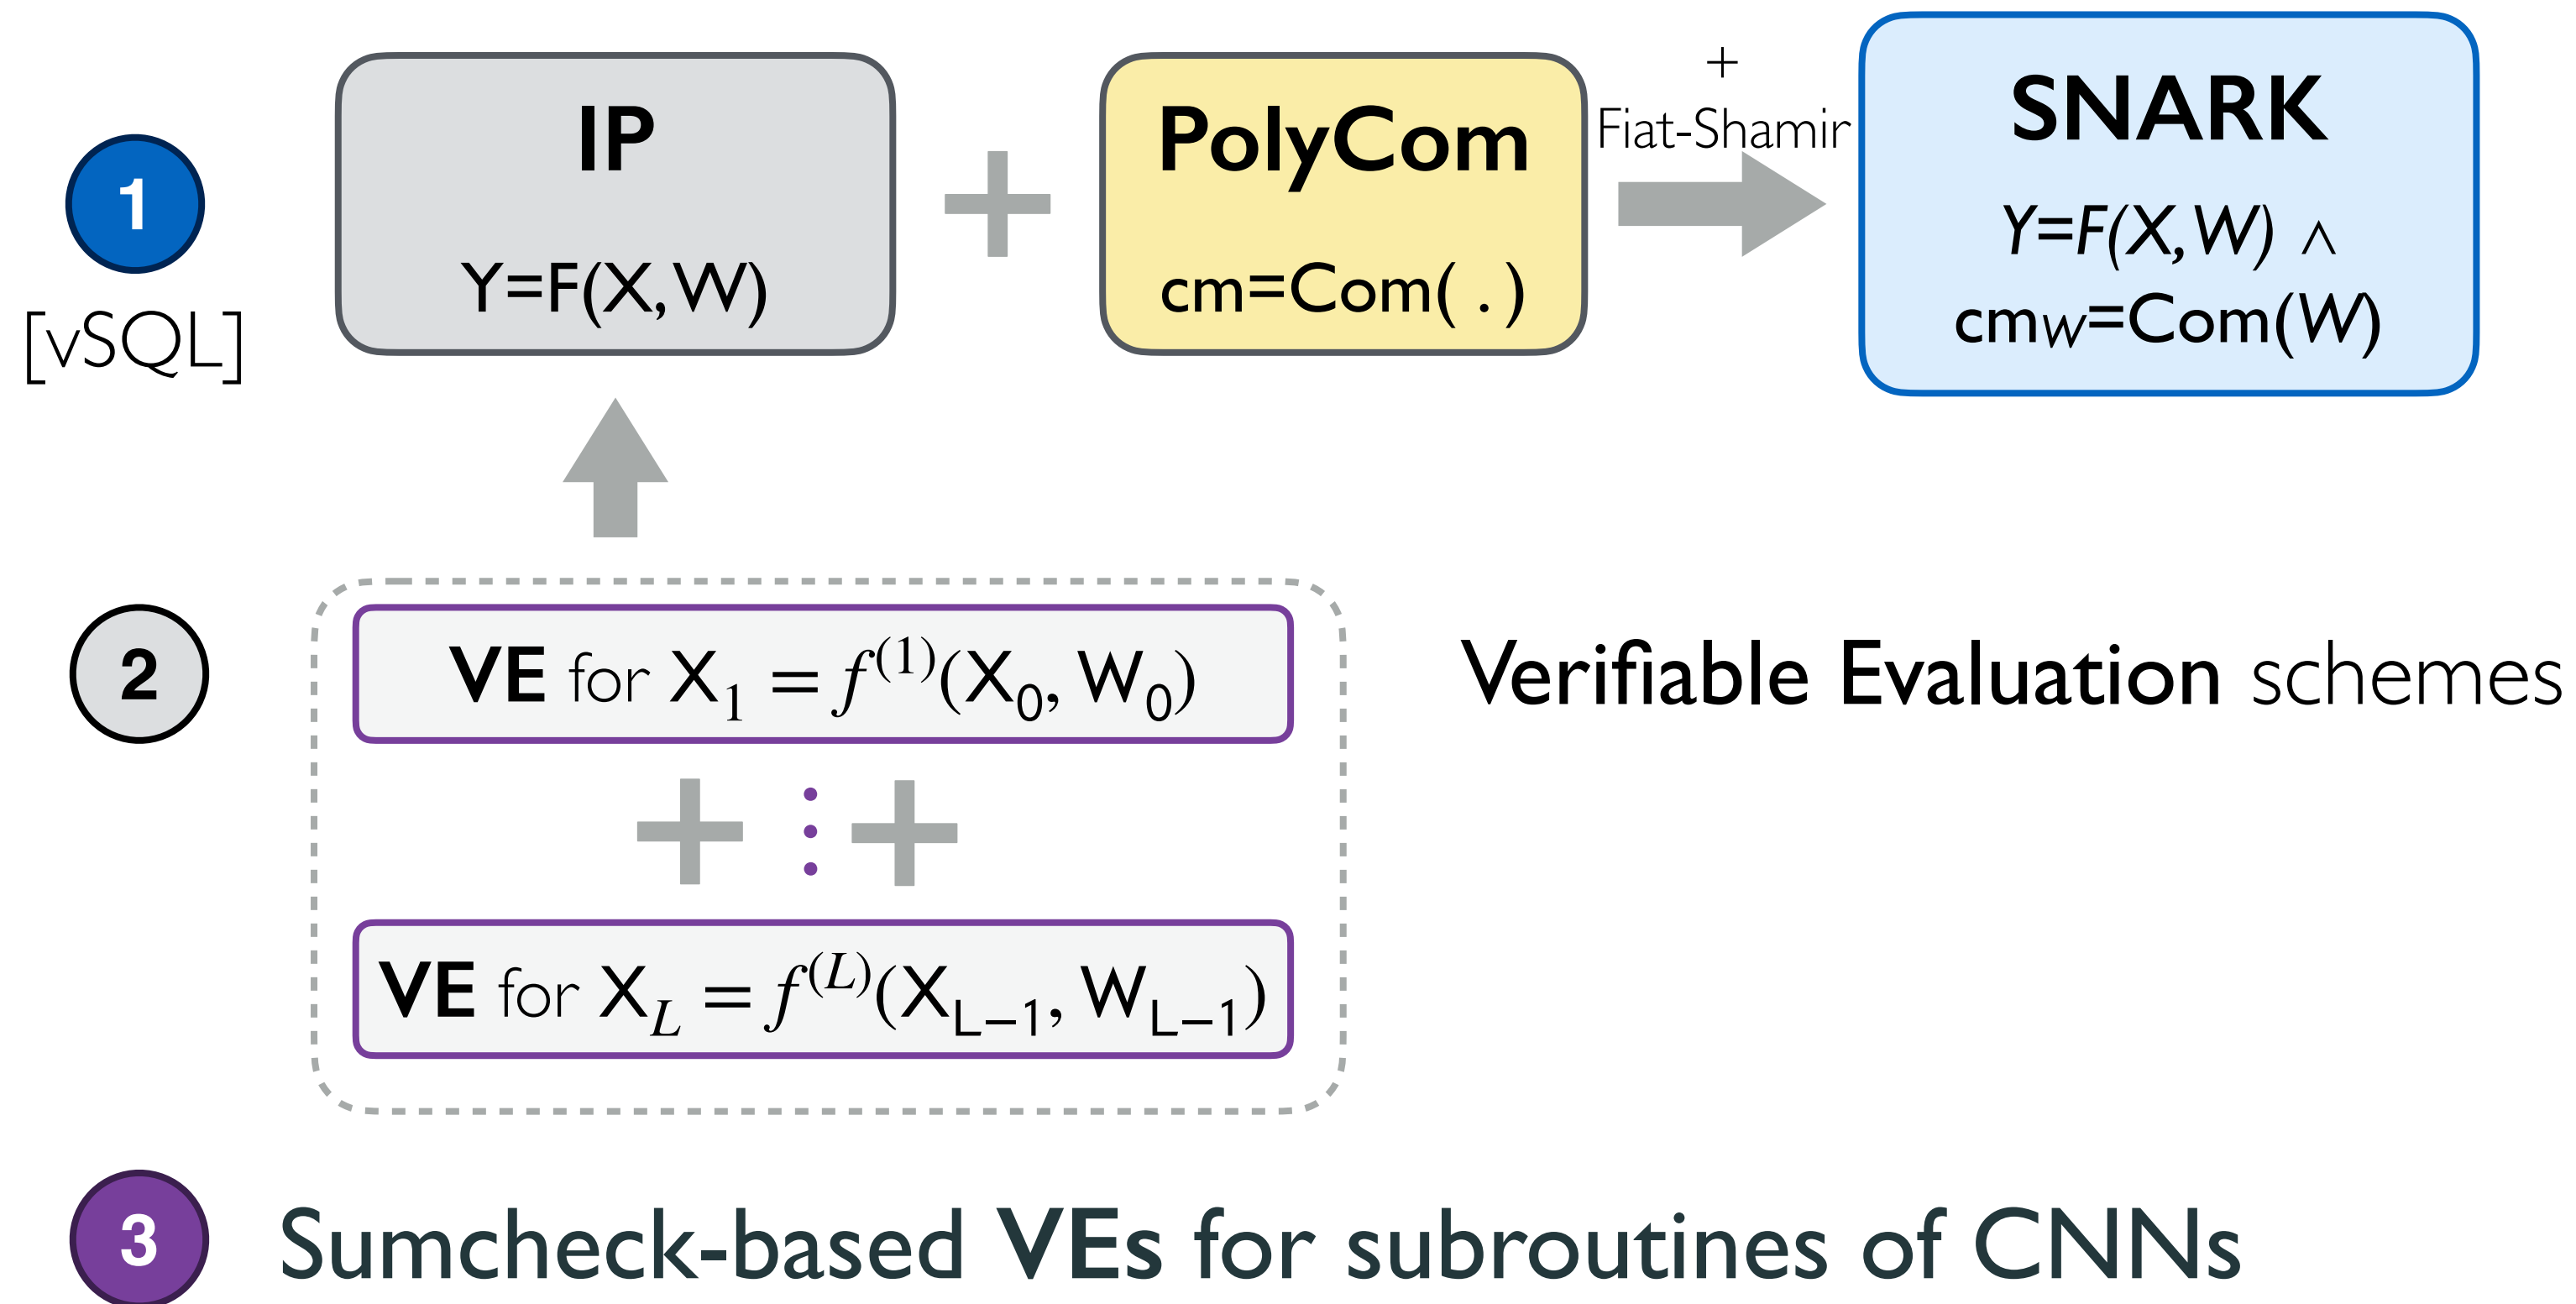
Modular approach for CNNs [BFGRS23]



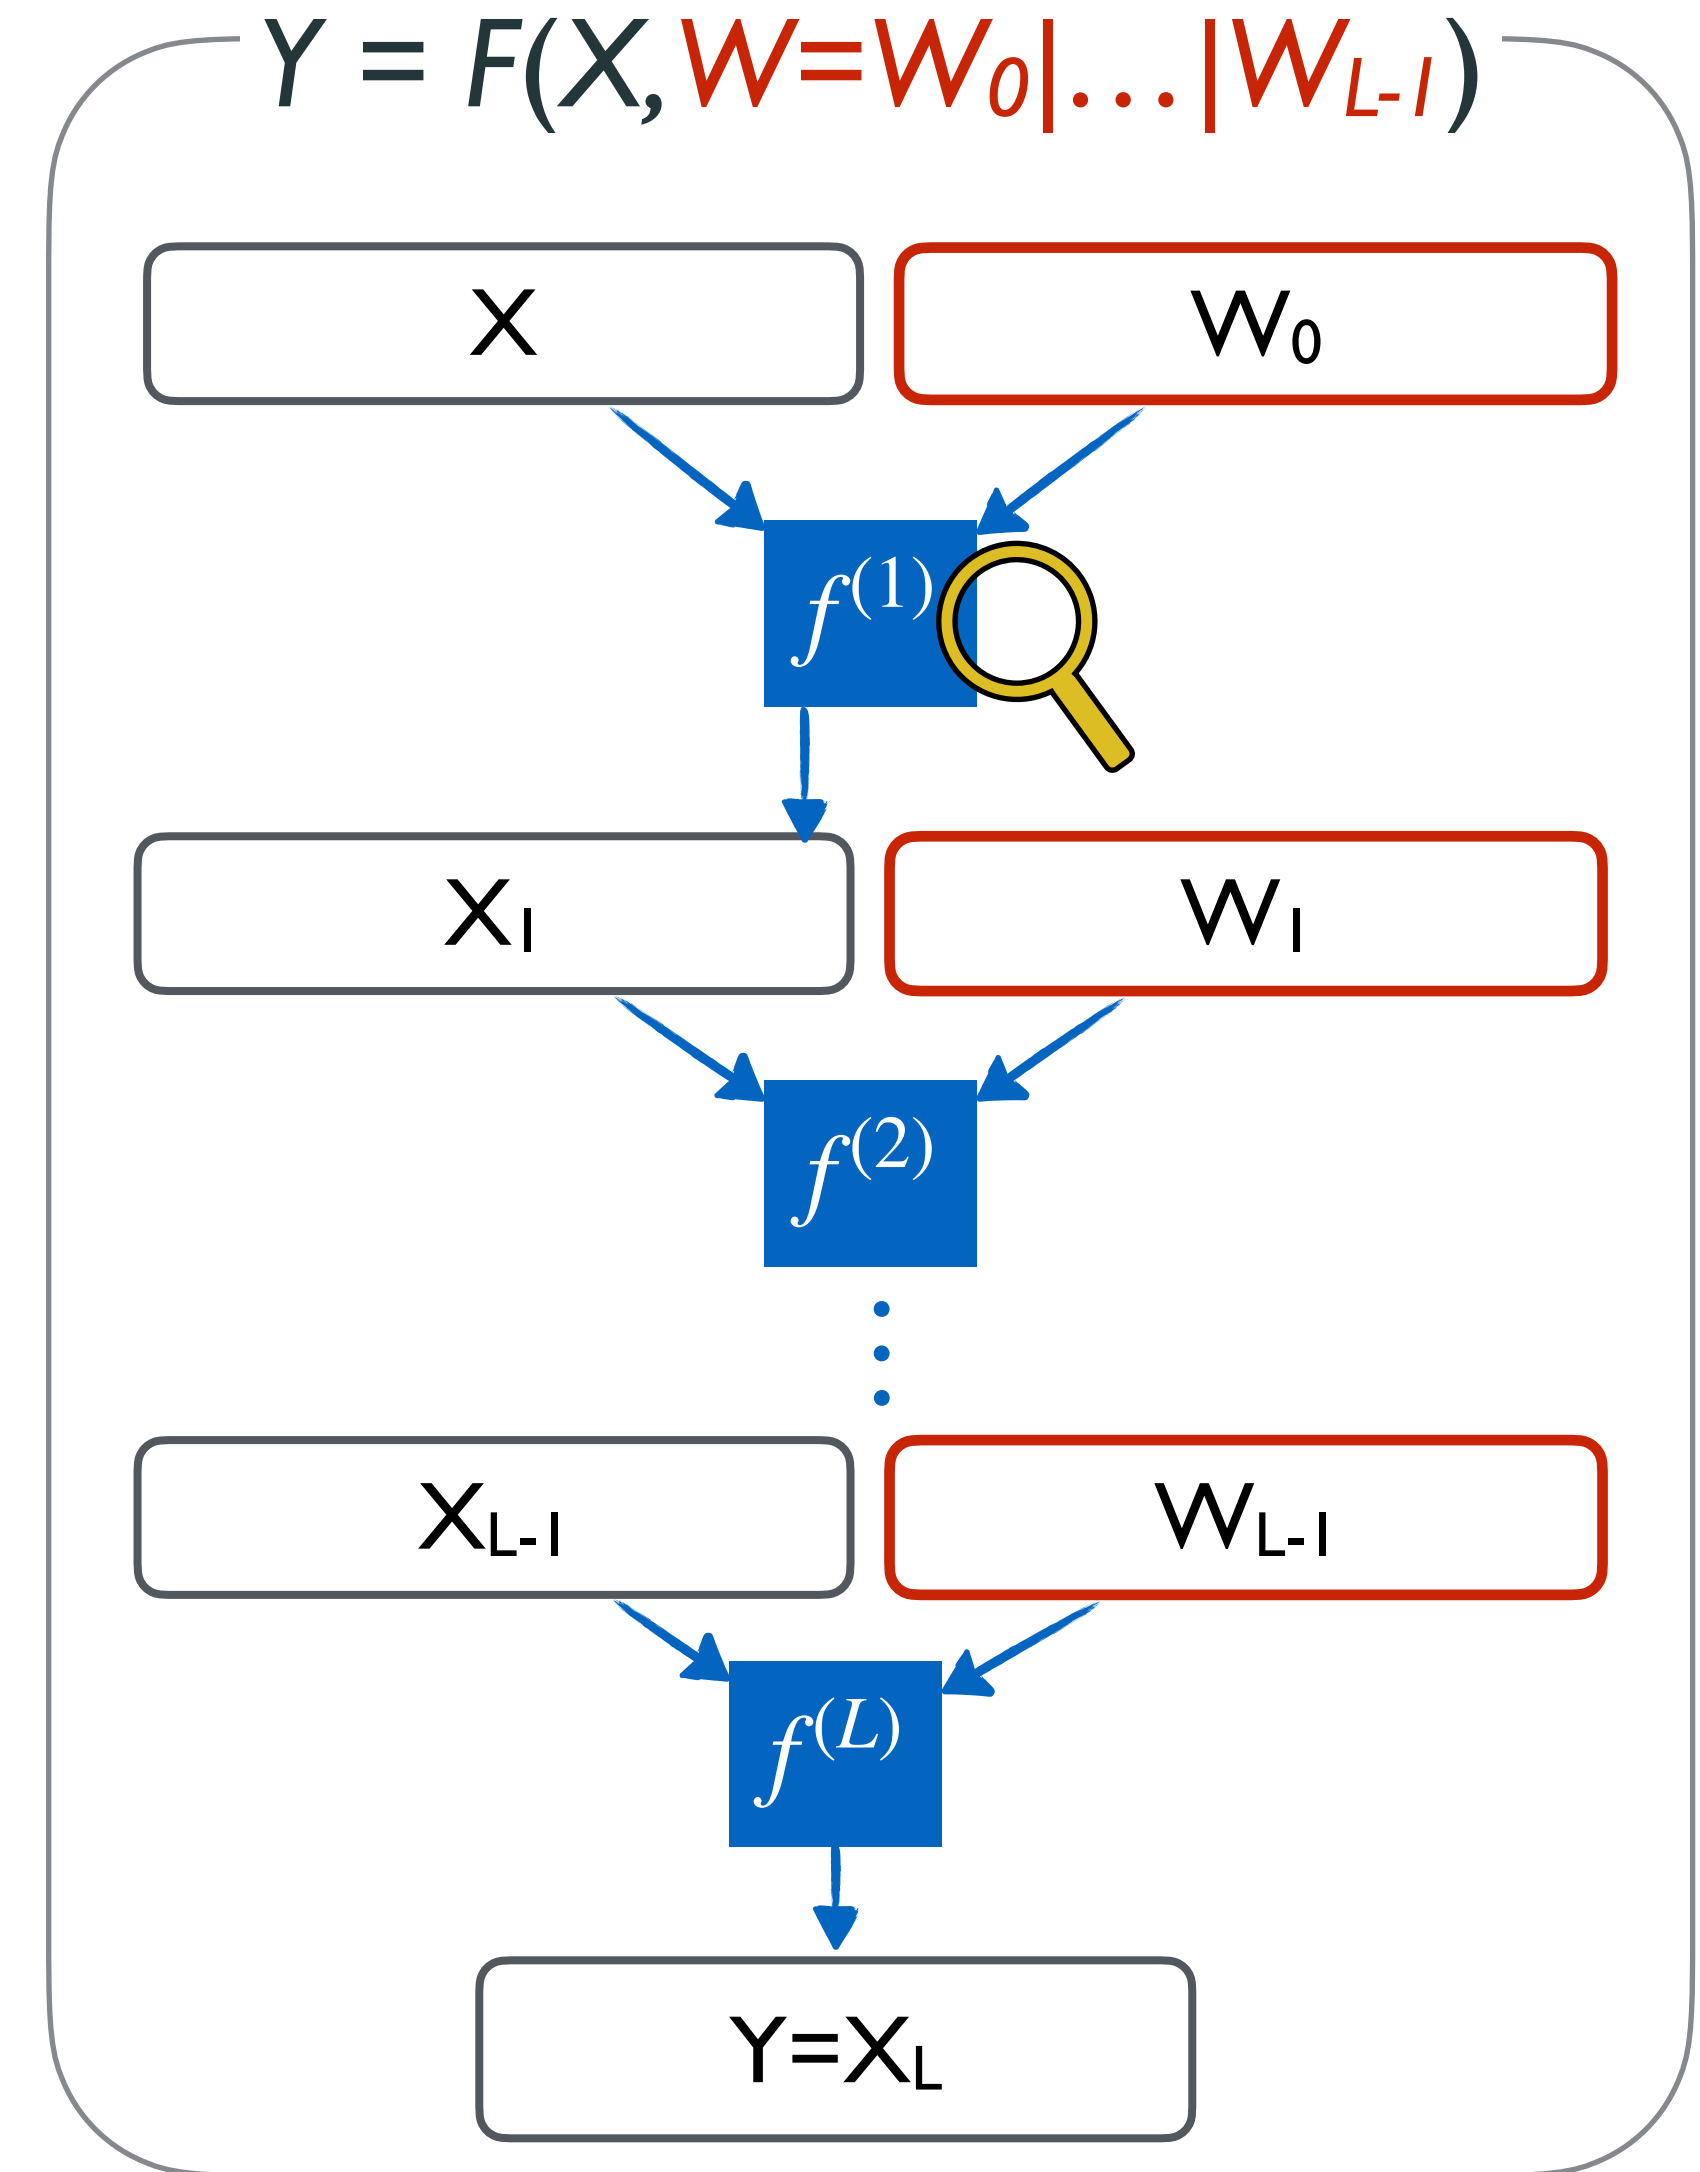
Modular approach for CNNs [BFGRS23]



Modular approach for CNNs [BFGRS23]



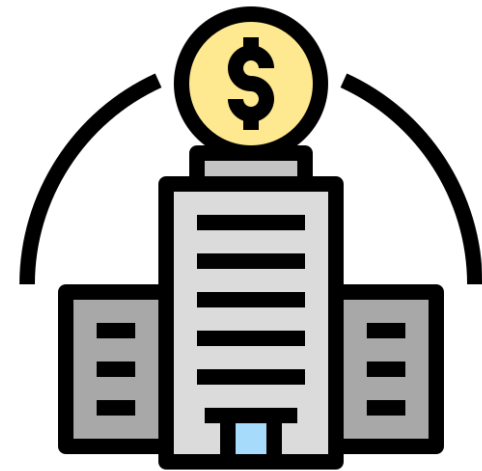
- Modular design&composition generalization of GKR-style IPs
- Easier to focus on designing efficient specialized VEs



Interactive Proofs and Fingerprints

$\langle \mathcal{P}, \mathcal{V} \rangle(f, x, y) \rightarrow b : \mathbf{IP}$ for language $\mathcal{L}_F = \{(f, x, y) : f(x) = y\}$ complete and sound

$\mathcal{P}(f, x, y)$

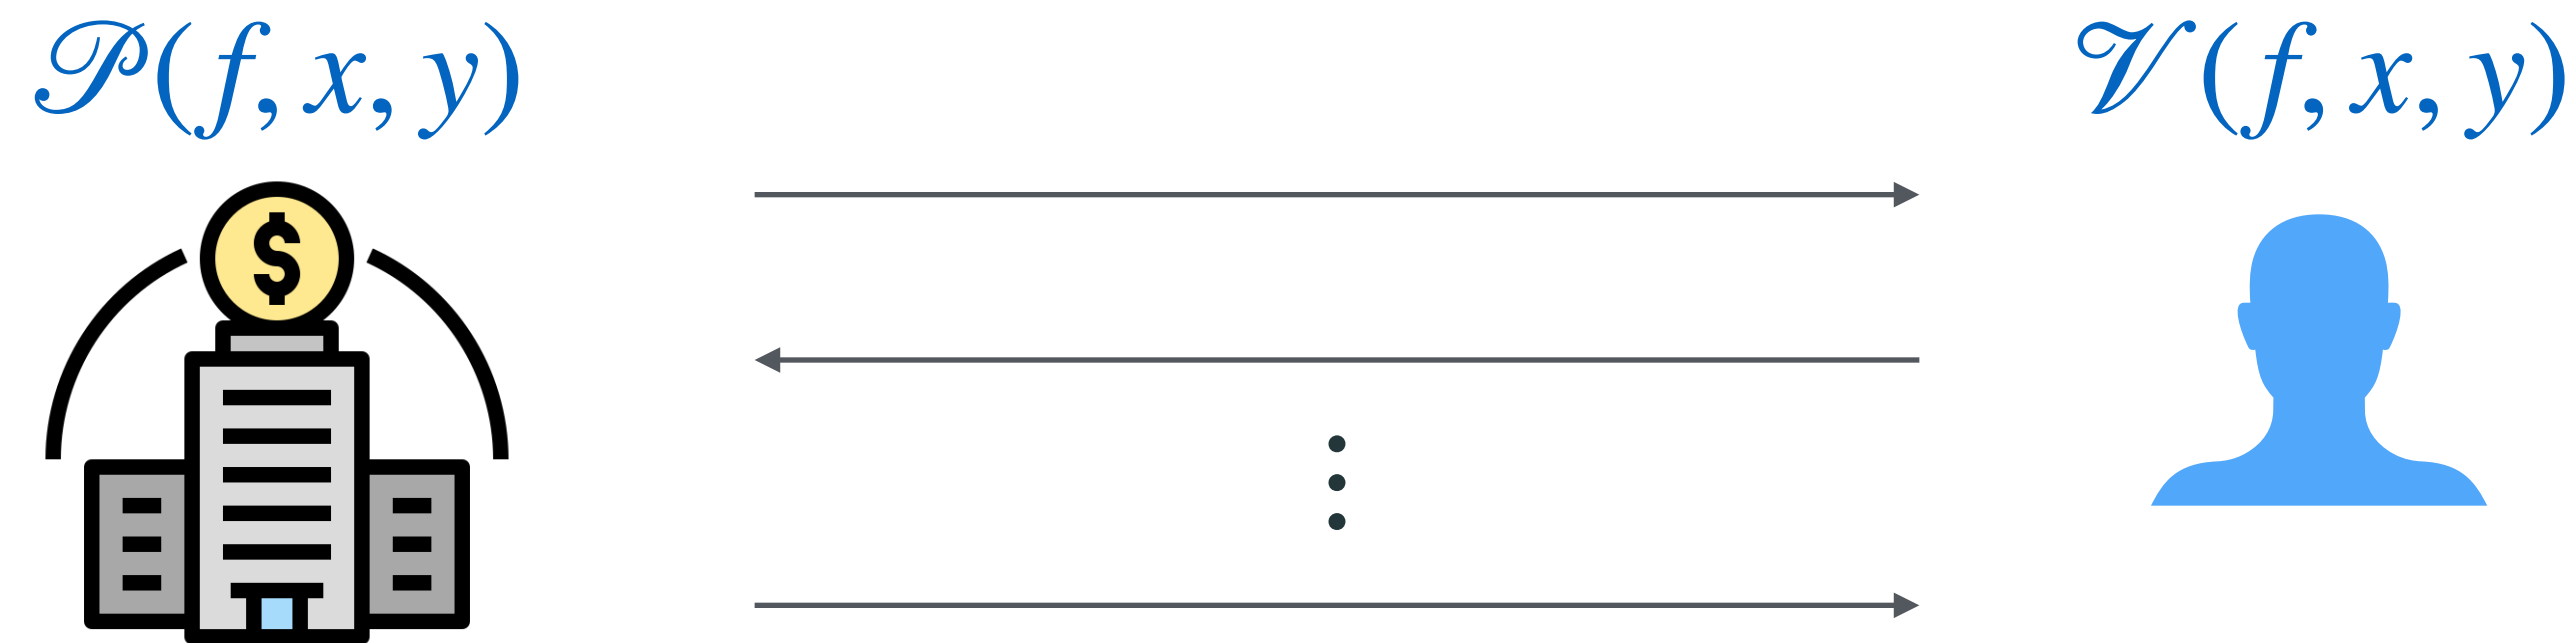


$\mathcal{V}(f, x, y)$



Interactive Proofs and Fingerprints

$\langle \mathcal{P}, \mathcal{V} \rangle(f, x, y) \rightarrow b : \mathbf{IP}$ for language $\mathcal{L}_F = \{(f, x, y) : f(x) = y\}$ complete and sound

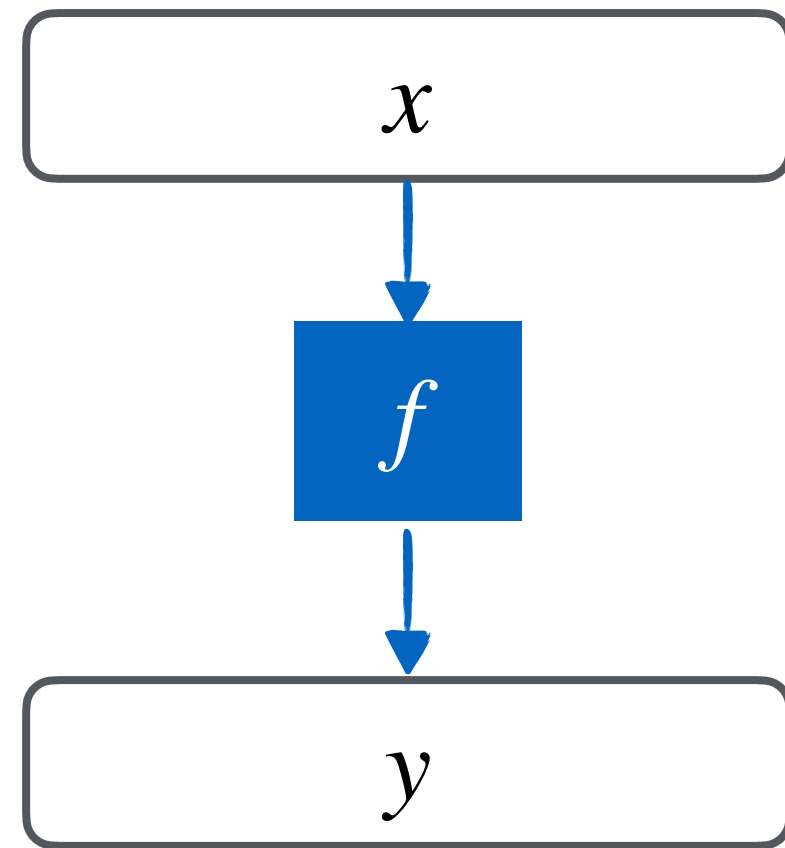


Fingerprint of x on $r : c_x \leftarrow H(x, r)$

- Compressing: $|c_x| \ll |x|$
- Statistically binding: $\forall x \neq x^*, \Pr_r[H(x, r) = H(x^*, r)] = \text{negl}(\lambda)$

Example: for $\vec{x} \in \mathbb{F}^n, \vec{r} \in \mathbb{F}^{\log n}$, H is the MLE evaluation $H(\vec{x}, \vec{r}) = \tilde{x}(\vec{r})$

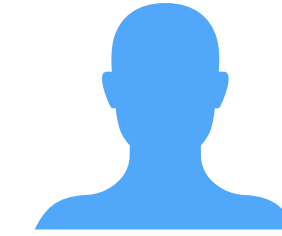
Structure of common IPs



$\mathcal{P}(f, x, y)$

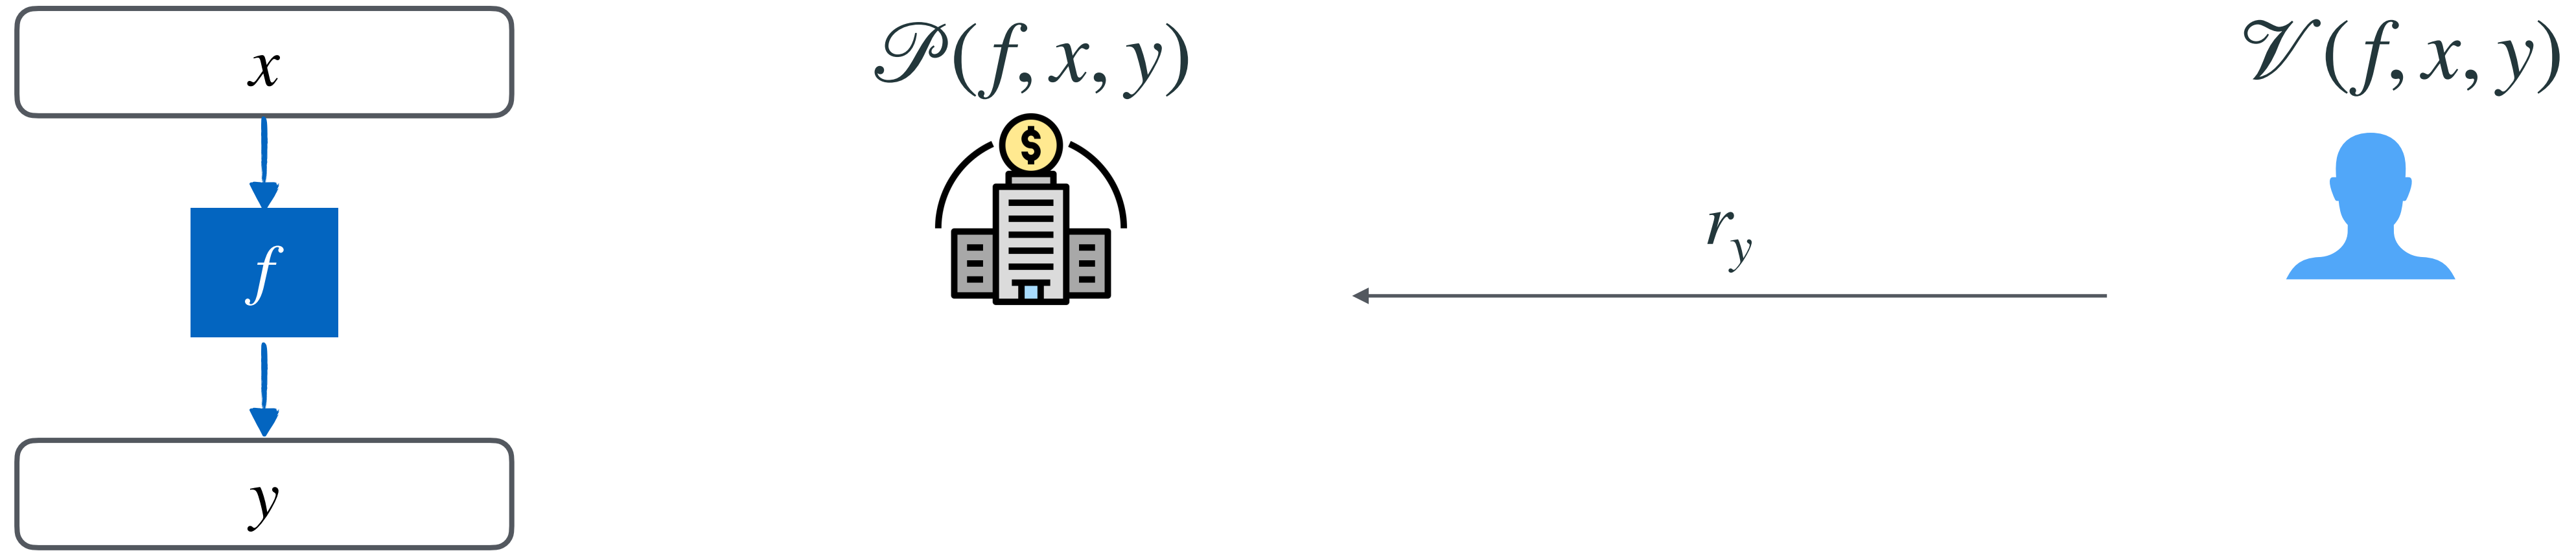


$\mathcal{V}(f, x, y)$



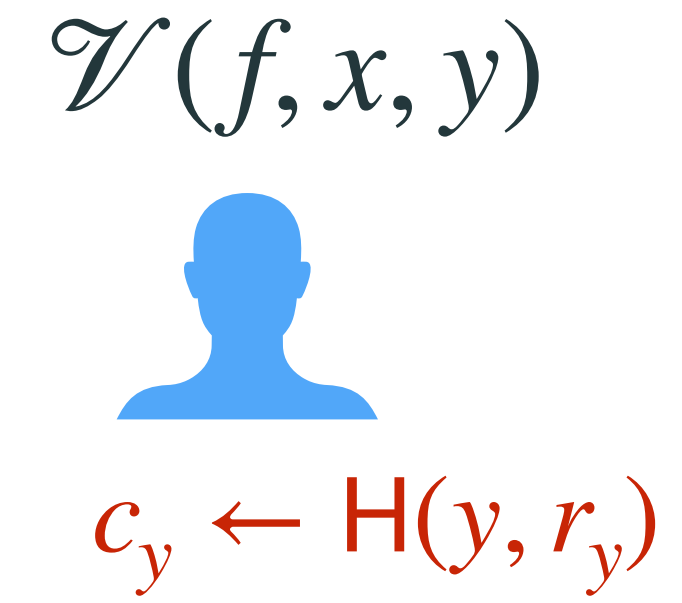
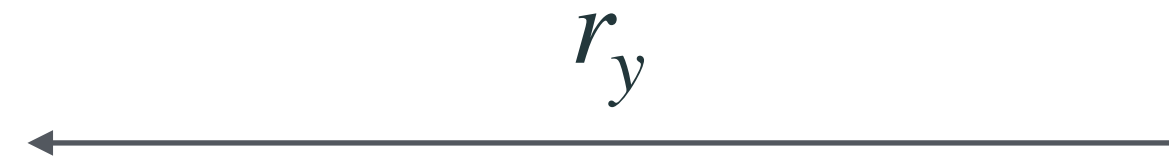
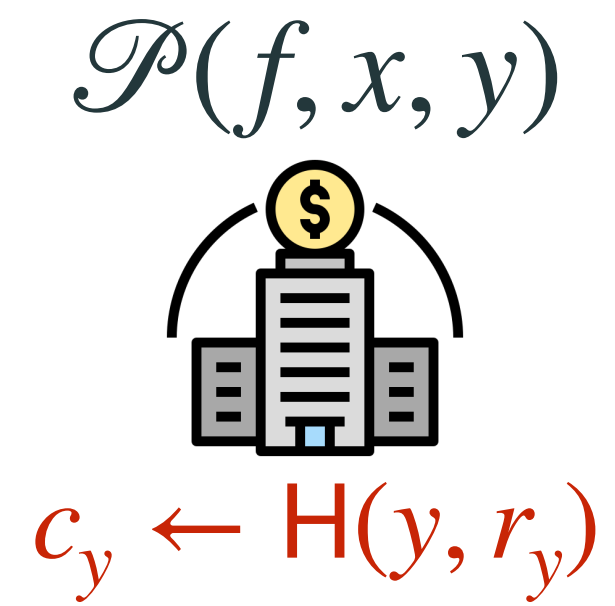
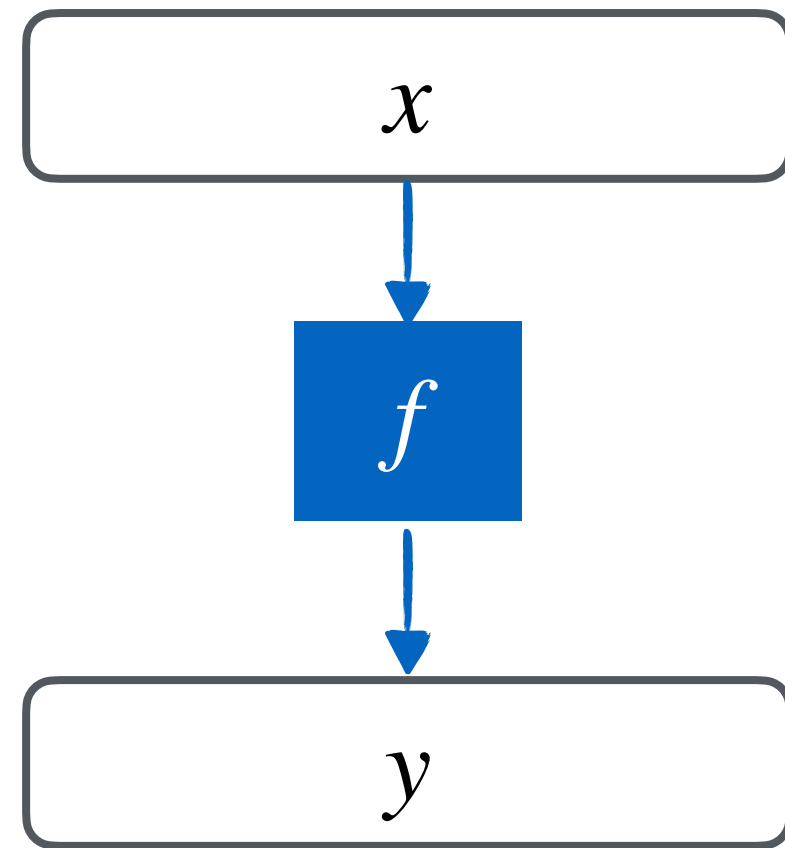
- Public coin verifier

Structure of common IPs



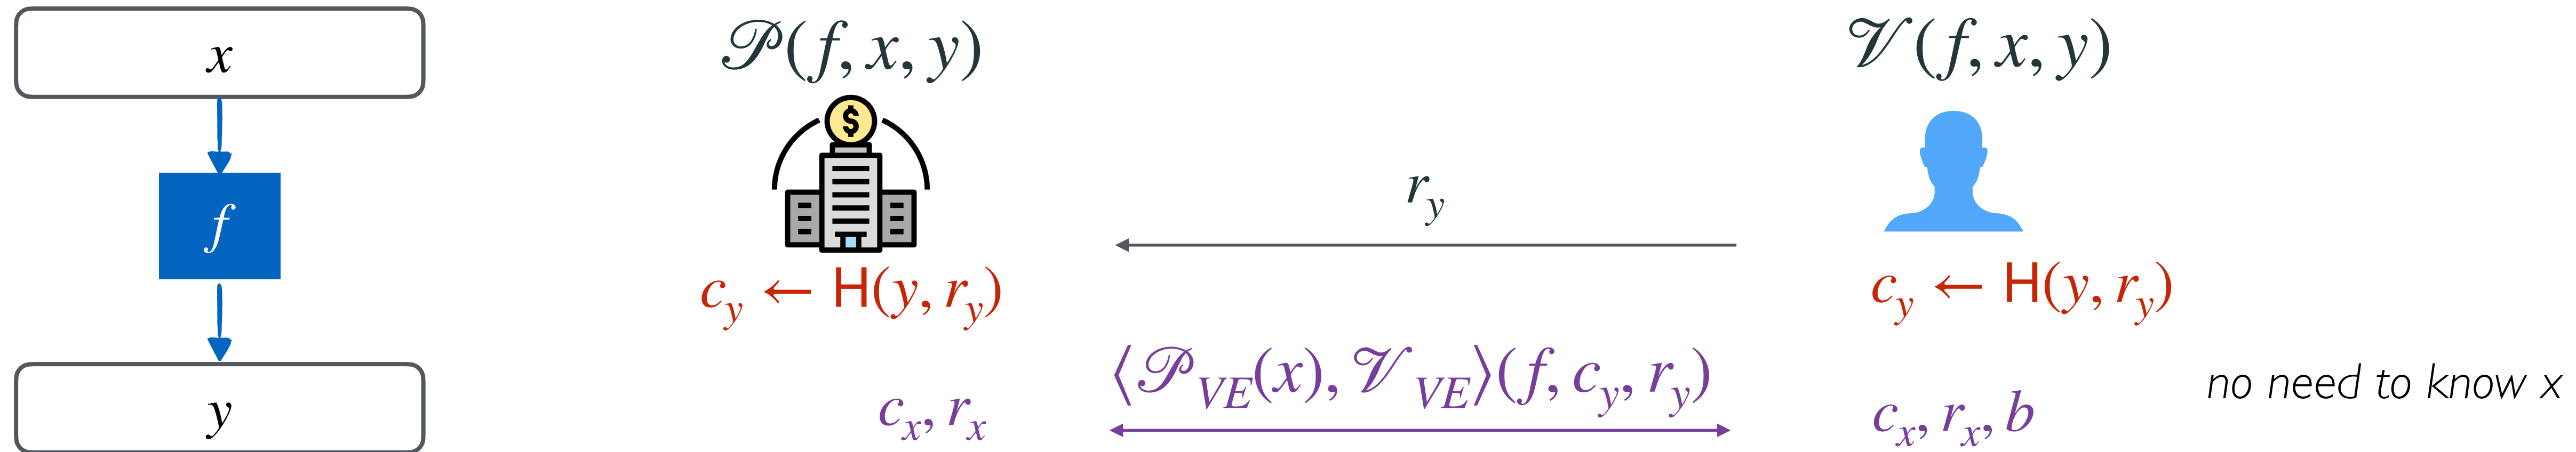
- Public coin verifier

Structure of common IPs



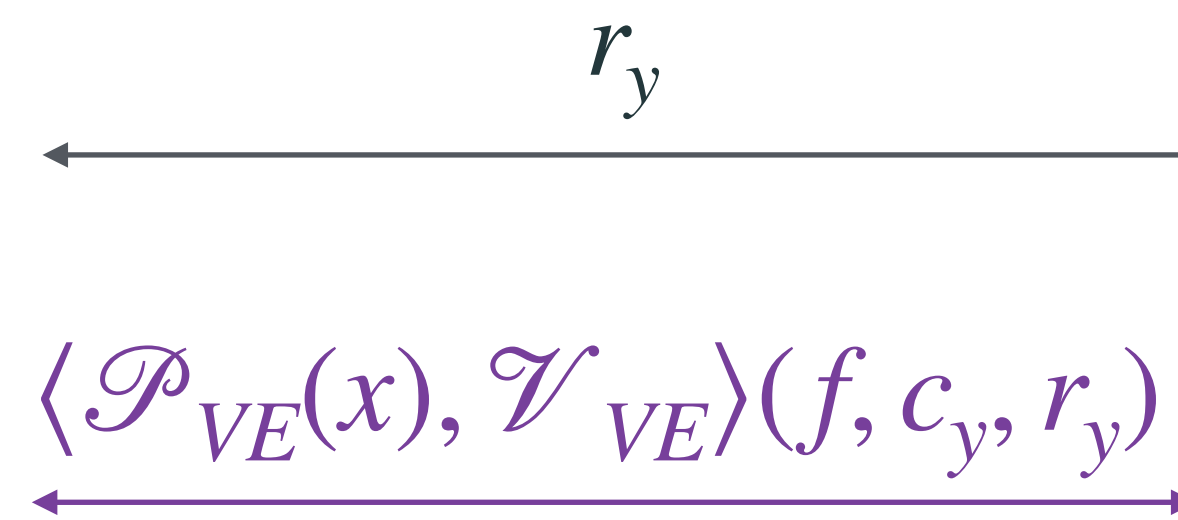
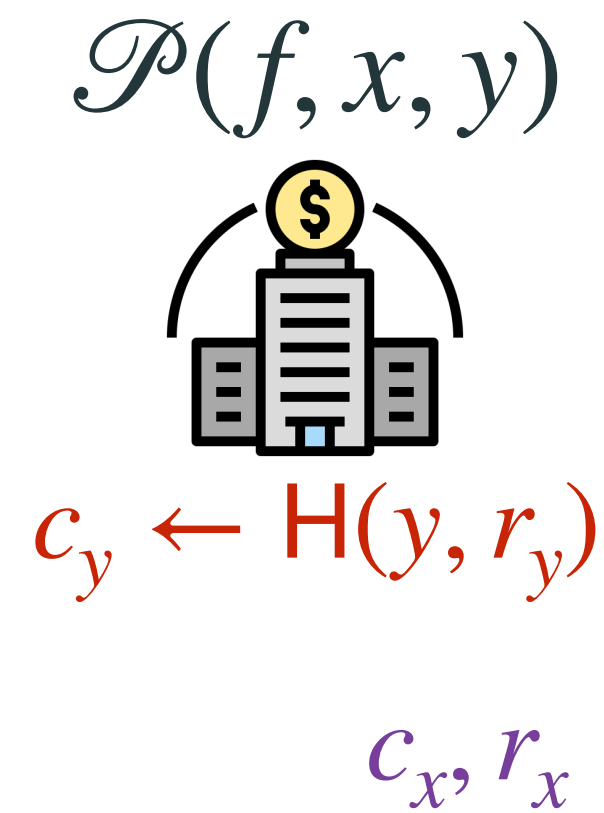
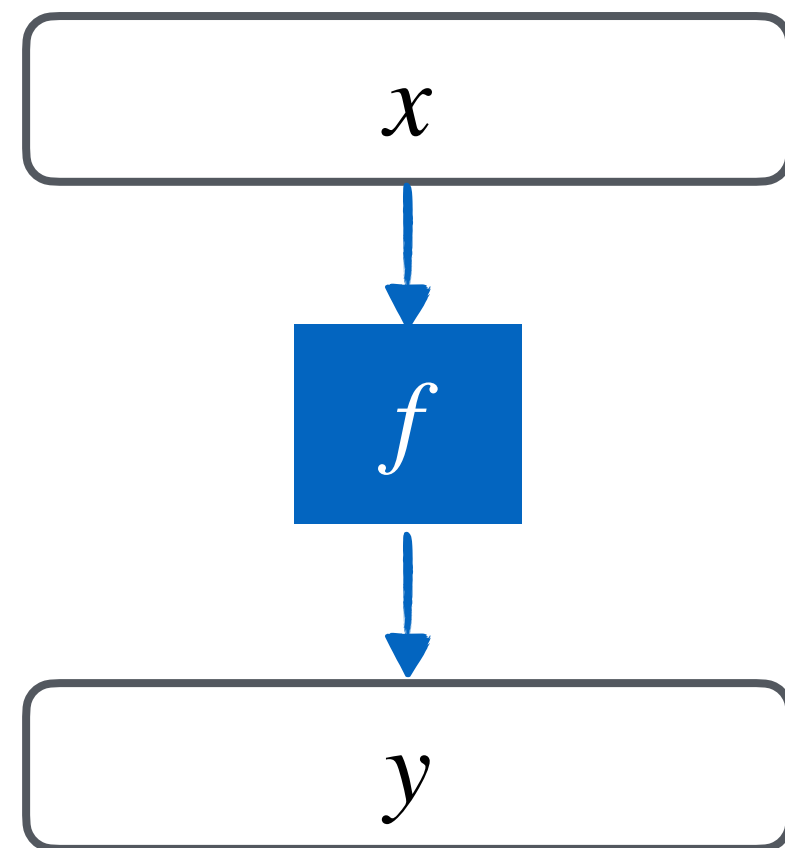
- Public coin verifier
- Output fingerprint

Structure of common IPs



- Public coin verifier
- Output fingerprint
- Subroutine: **verifiable evaluation (VE)** schemes on **fingerprinted** data (\mathcal{V}_{VE} runs w/o x, y)

Structure of common IPs



$\mathcal{V}(f, x, y)$



$c_y \leftarrow H(y, r_y)$

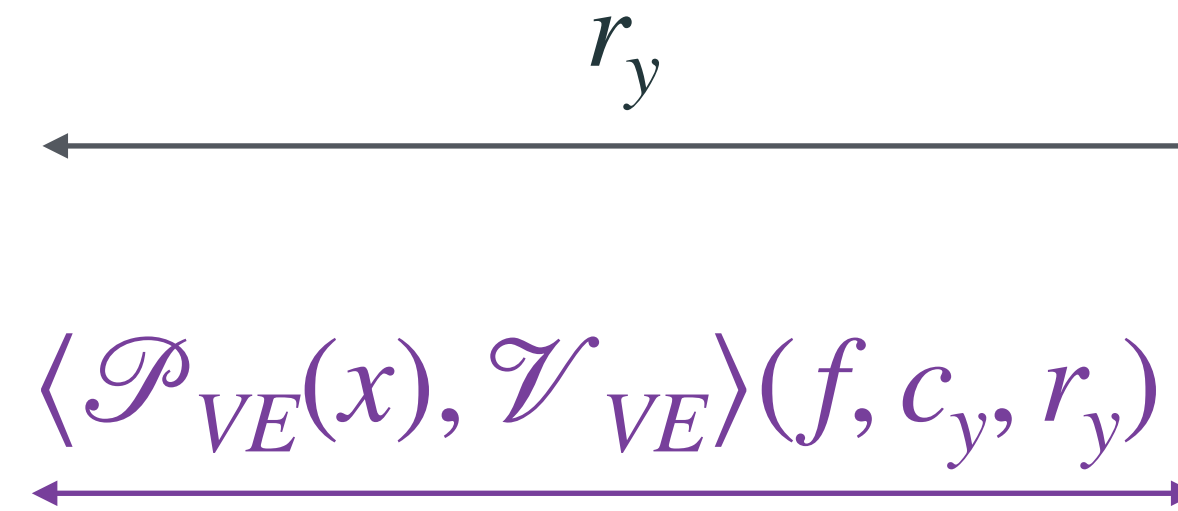
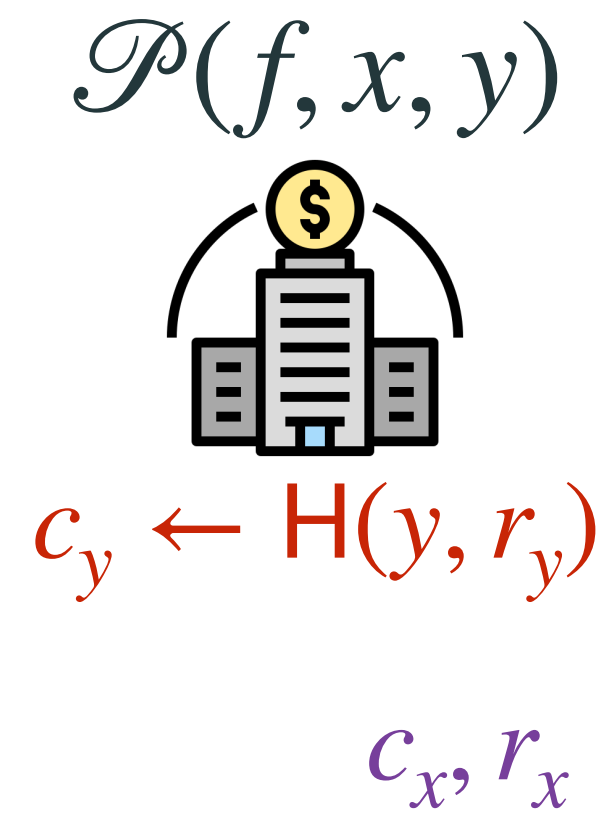
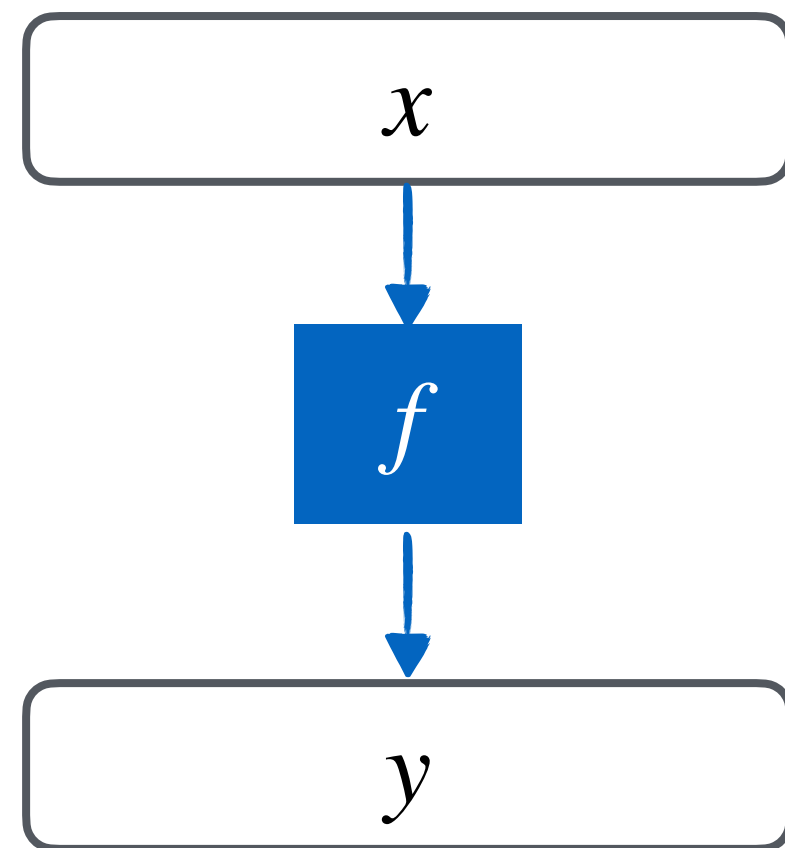
c_x, r_x, b

no need to know x

Return $b \wedge c_x \stackrel{?}{=} H(x, r_x)$

- Public coin verifier
- Output fingerprint
- **Subroutine: verifiable evaluation (VE)** schemes on **fingerprinted** data (\mathcal{V}_{VE} runs w/o x, y)
- Input fingerprint

Structure of common IPs



$\mathcal{V}(f, x, y)$



$c_y \leftarrow H(y, r_y)$

c_x, r_x, b

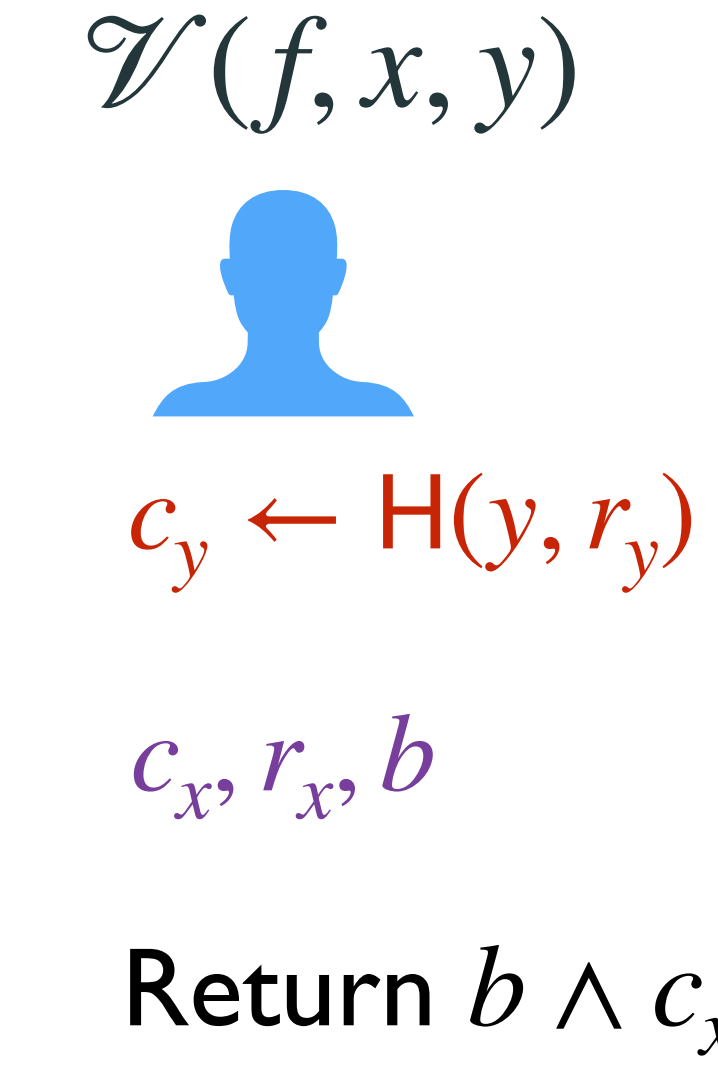
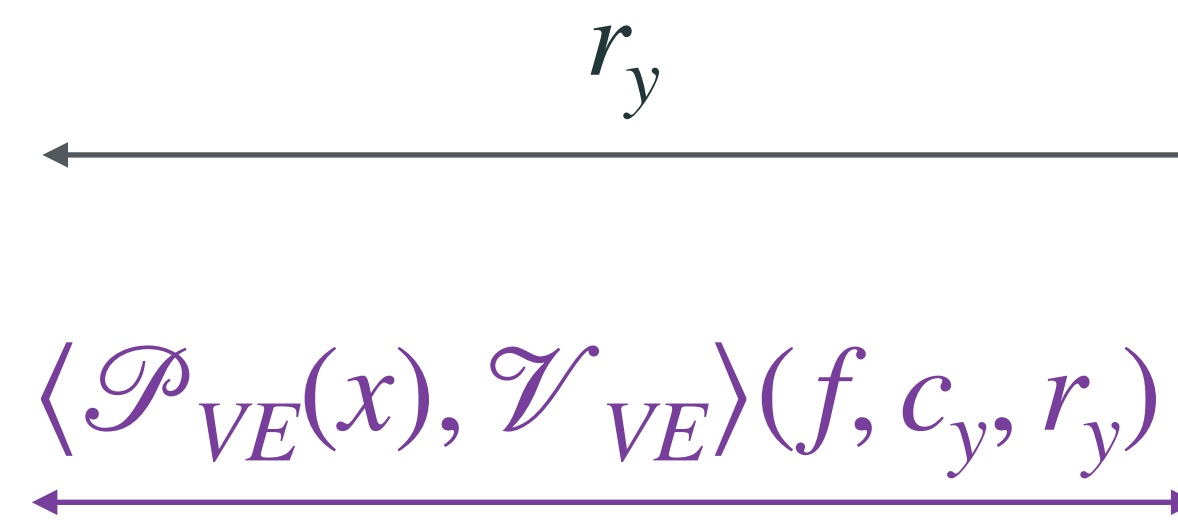
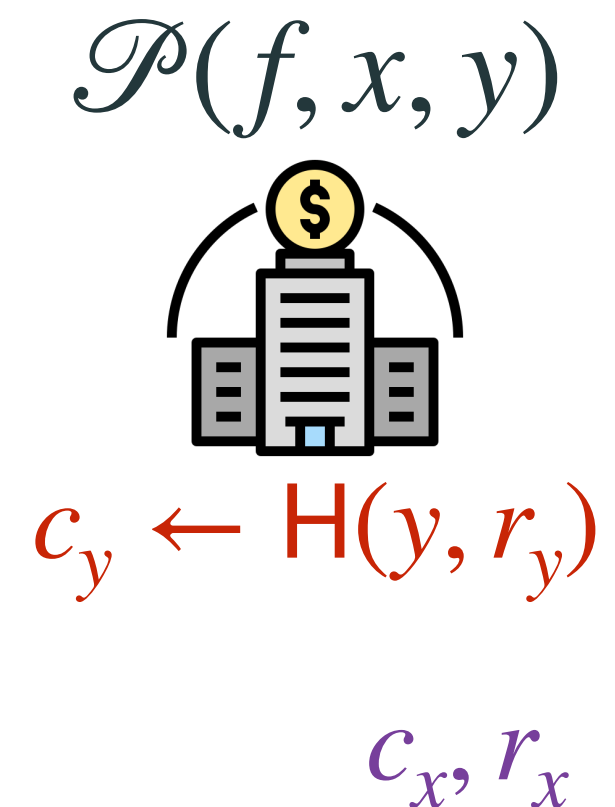
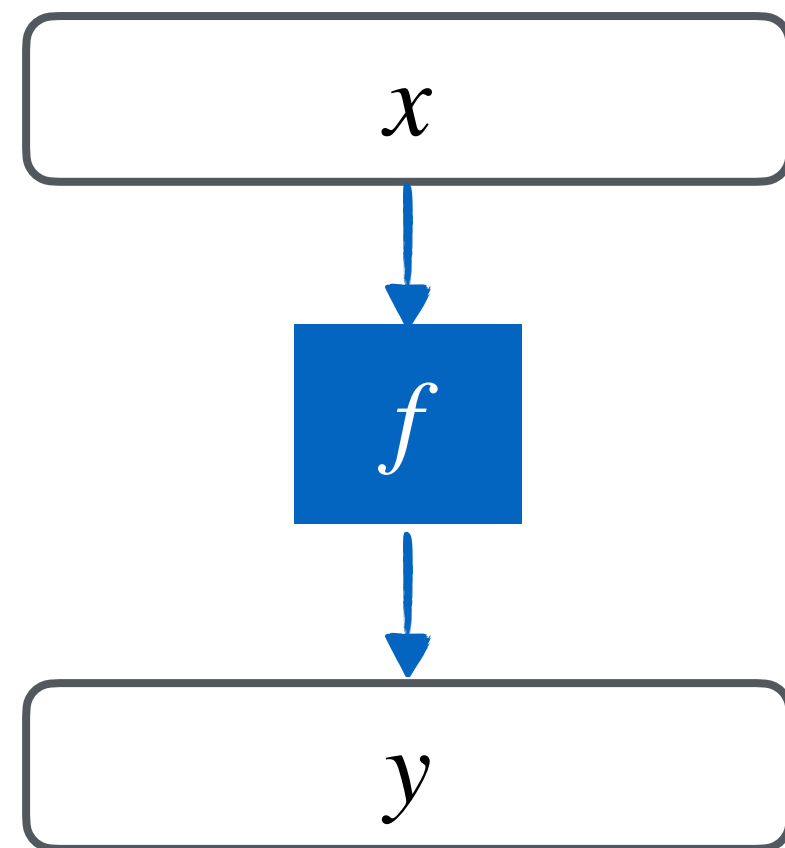
no need to know x

Return $b \wedge c_x \stackrel{?}{=} H(x, r_x)$

- Public coin verifier
- Output fingerprint
- **Subroutine: verifiable evaluation (VE)** schemes on **fingerprinted** data (\mathcal{V}_{VE} runs w/o x, y)
- Input fingerprint

Examples of VE-based IPs: sumcheck protocol, GKR,

VE Soundness

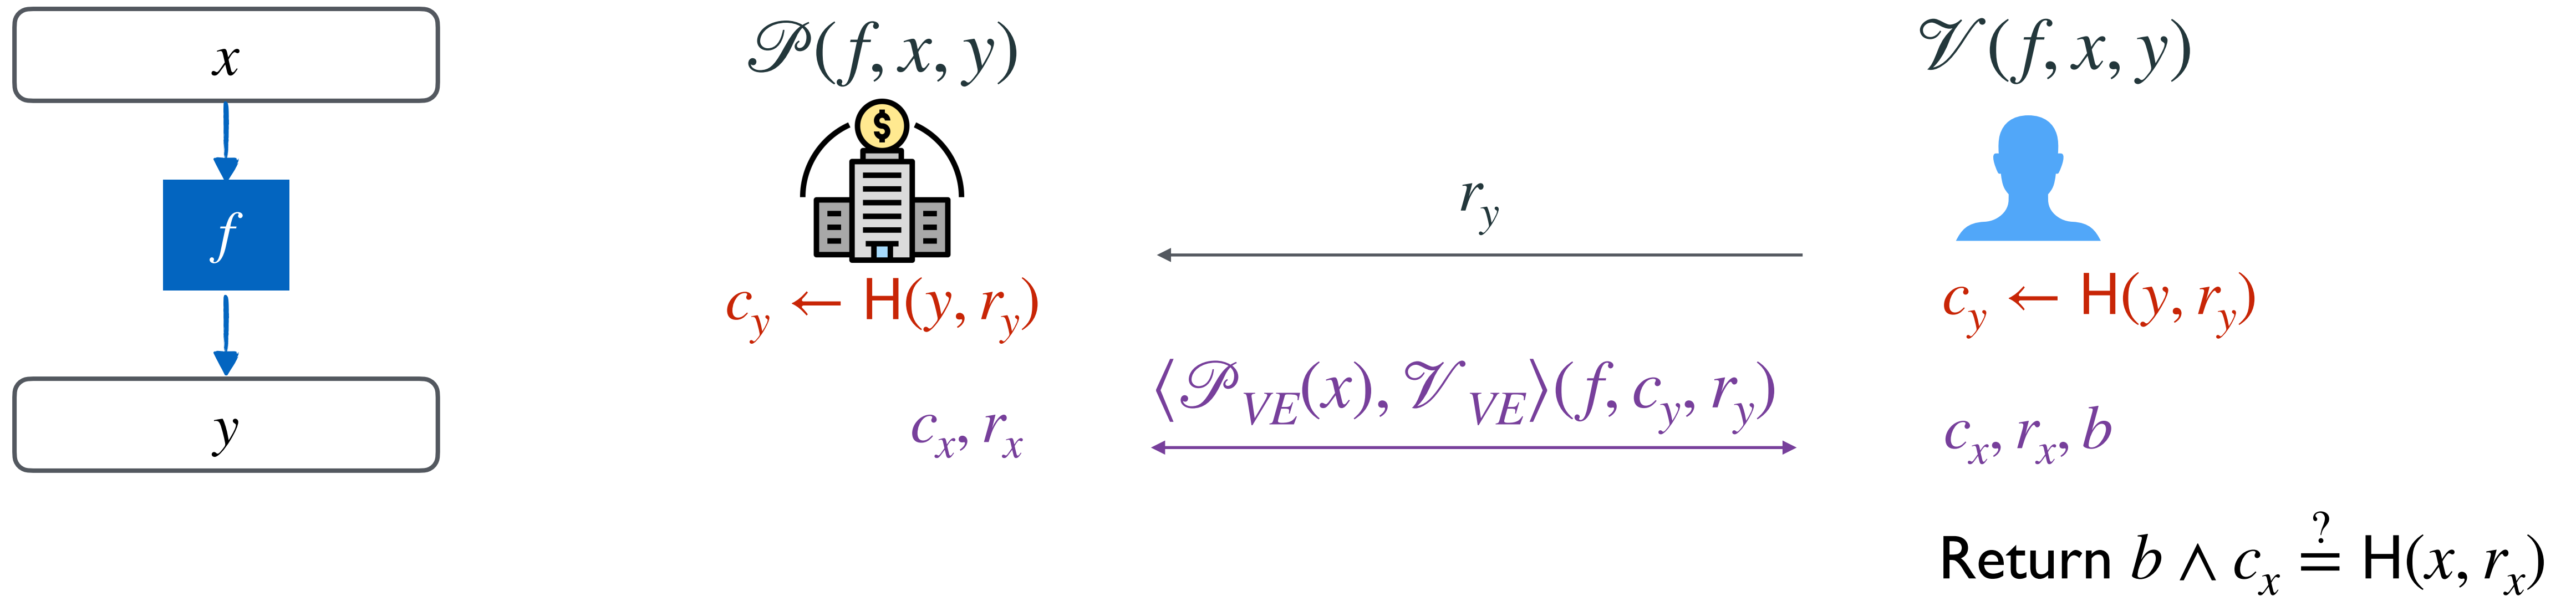


Soundness of VE subroutine: for random r_x, r_y and any unbounded \mathcal{A}

$$\Pr \left[\begin{array}{l} c_y^* \neq H(f(x), r_y) \\ \wedge b \end{array} \mid \begin{array}{l} (f, c_y^*, x) \leftarrow \mathcal{A}(r_y) \\ (c_x^*; r_x; b) \leftarrow \langle \mathcal{A}(x), \mathcal{V}_{VE}(r_x) \rangle(f, c_y^*, r_y) \\ c_x^* = H(x, r_x) \end{array} \right] = \text{negl}$$

Correct input fingerprint
 $\Rightarrow c_y^*$ fingerprints correct output

VE Soundness



Soundness of VE subroutine: for random r_x, r_y and any unbounded \mathcal{A}

$$\Pr \left[\begin{array}{l} c_y^* \neq H(f(x), r_y) \\ \wedge b \end{array} \mid \begin{array}{l} (f, c_y^*, x) \leftarrow \mathcal{A}(r_y) \\ (c_x^*; r_x; b) \leftarrow \langle \mathcal{A}(x), \mathcal{V}_{VE}(r_x) \rangle(f, c_y^*, r_y) \\ c_x^* = H(x, r_x) \end{array} \right] = \text{negl}$$

Correct input fingerprint
 $\Rightarrow c_y^*$ fingerprints correct output

Sound VE + Binding Fingerprint \Rightarrow Sound IP

Sumcheck protocol as VE-based IP [LFKN92]

$x \in \mathbb{F}[T_1, \dots, T_\ell]$ multilinear polynomial. **Goal:** prove $y = f(x) = \sum_{t_1, \dots, t_\ell \in \{0,1\}} x(t_1, \dots, t_\ell)$

$\mathcal{P}(x, y)$

$\mathcal{V}(y)$

$$f_1(T_1) = \sum_{t_2, \dots, t_\ell \in \{0,1\}} x(T_1, t_2, \dots, t_\ell)$$



r_1



\vdots

$$f_\ell(T_\ell) = x(r_1, r_2, \dots, r_{\ell-1}, T_\ell)$$



$$\bigwedge_i [f_i(0) + f_i(1) \stackrel{?}{=} f_{i-1}(r_{i-1})]$$

$$\wedge x(r_1, \dots, r_\ell) \stackrel{?}{=} f_\ell(r_\ell)$$

Sumcheck protocol as VE-based IP [LFKN92]

$x \in \mathbb{F}[T_1, \dots, T_\ell]$ multilinear polynomial. **Goal:** prove $y = f(x) = \sum_{t_1, \dots, t_\ell \in \{0,1\}} x(t_1, \dots, t_\ell)$

$\mathcal{P}(x, y)$

$\mathcal{V}(y)$

$$f_1(T_1) = \sum_{t_2, \dots, t_\ell \in \{0,1\}} x(T_1, t_2, \dots, t_\ell)$$



r_1



\vdots

$$f_\ell(T_\ell) = x(r_1, r_2, \dots, r_{\ell-1}, T_\ell)$$



r_ℓ



$$c_x \leftarrow x(r_1, \dots, r_\ell)$$



$$\bigwedge_i [f_i(0) + f_i(1) \stackrel{?}{=} f_{i-1}(r_{i-1})]$$

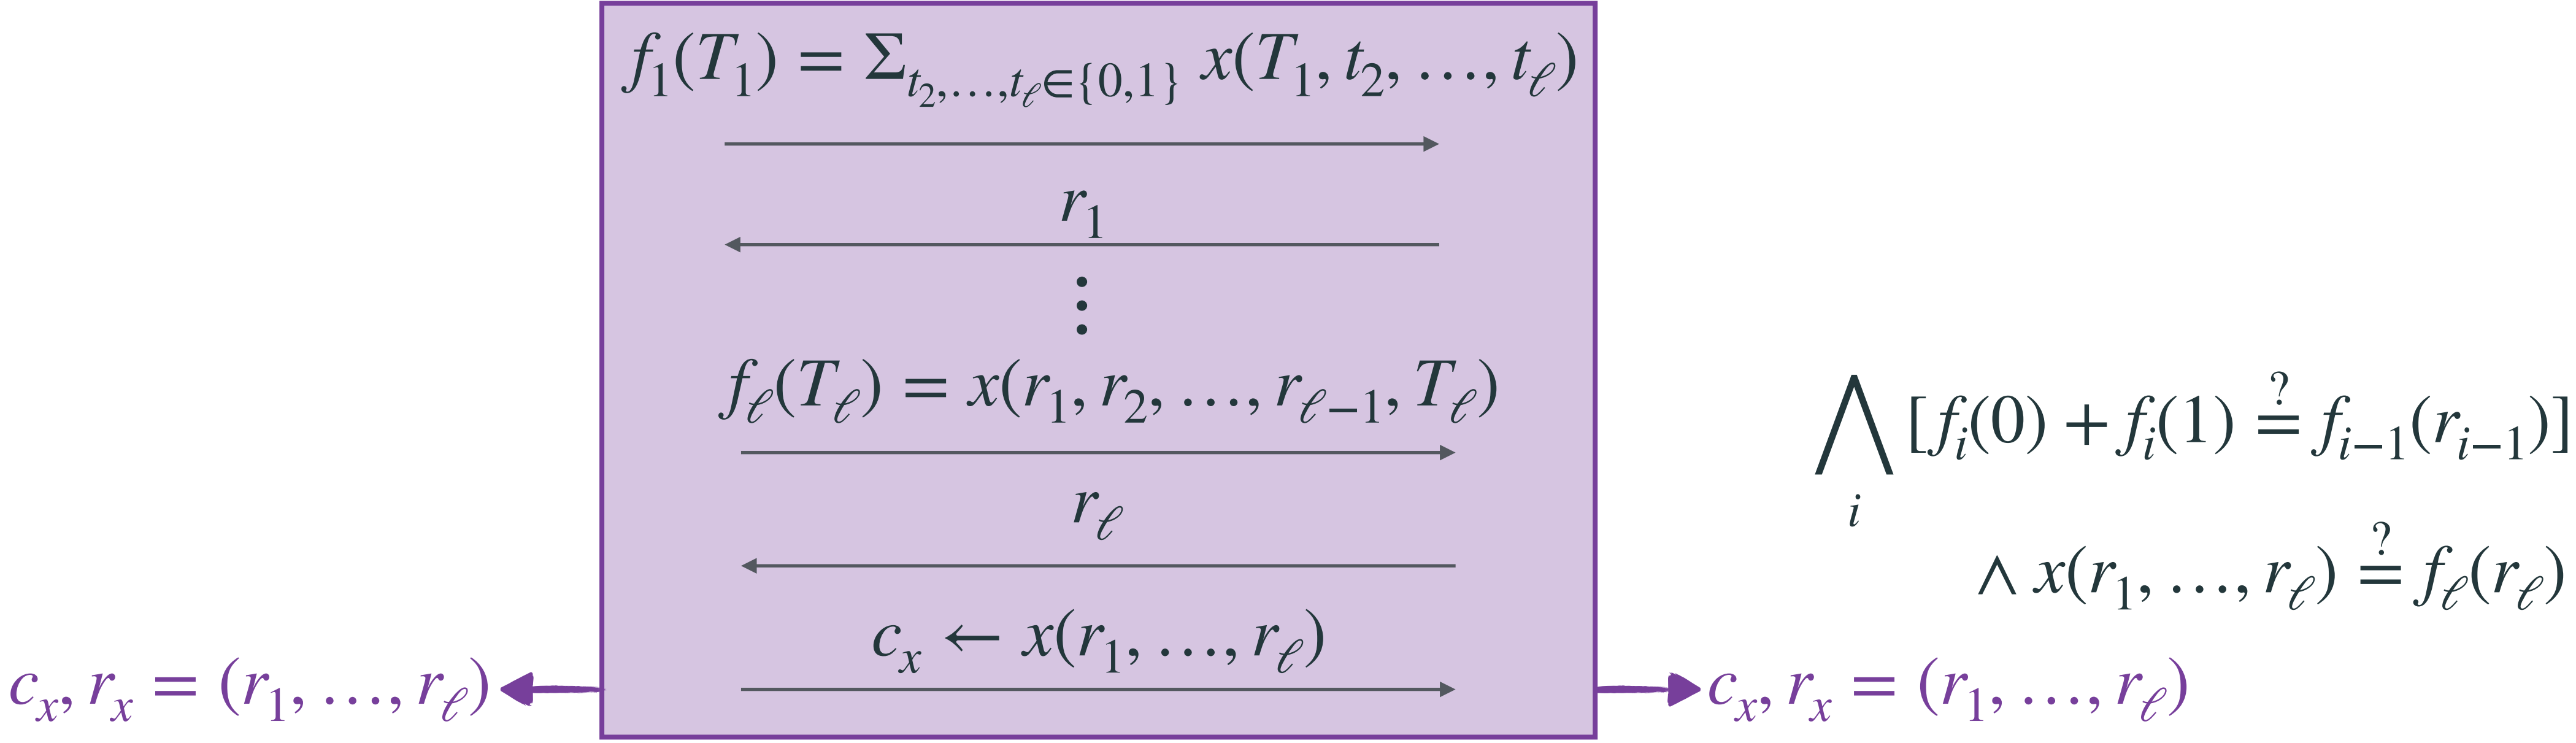
$$\wedge x(r_1, \dots, r_\ell) \stackrel{?}{=} f_\ell(r_\ell)$$

Sumcheck protocol as VE-based IP [LFKN92]

$x \in \mathbb{F}[T_1, \dots, T_\ell]$ multilinear polynomial. **Goal:** prove $y = f(x) = \sum_{t_1, \dots, t_\ell \in \{0,1\}} x(t_1, \dots, t_\ell)$

$\mathcal{P}(x, y)$

$\mathcal{V}(y)$

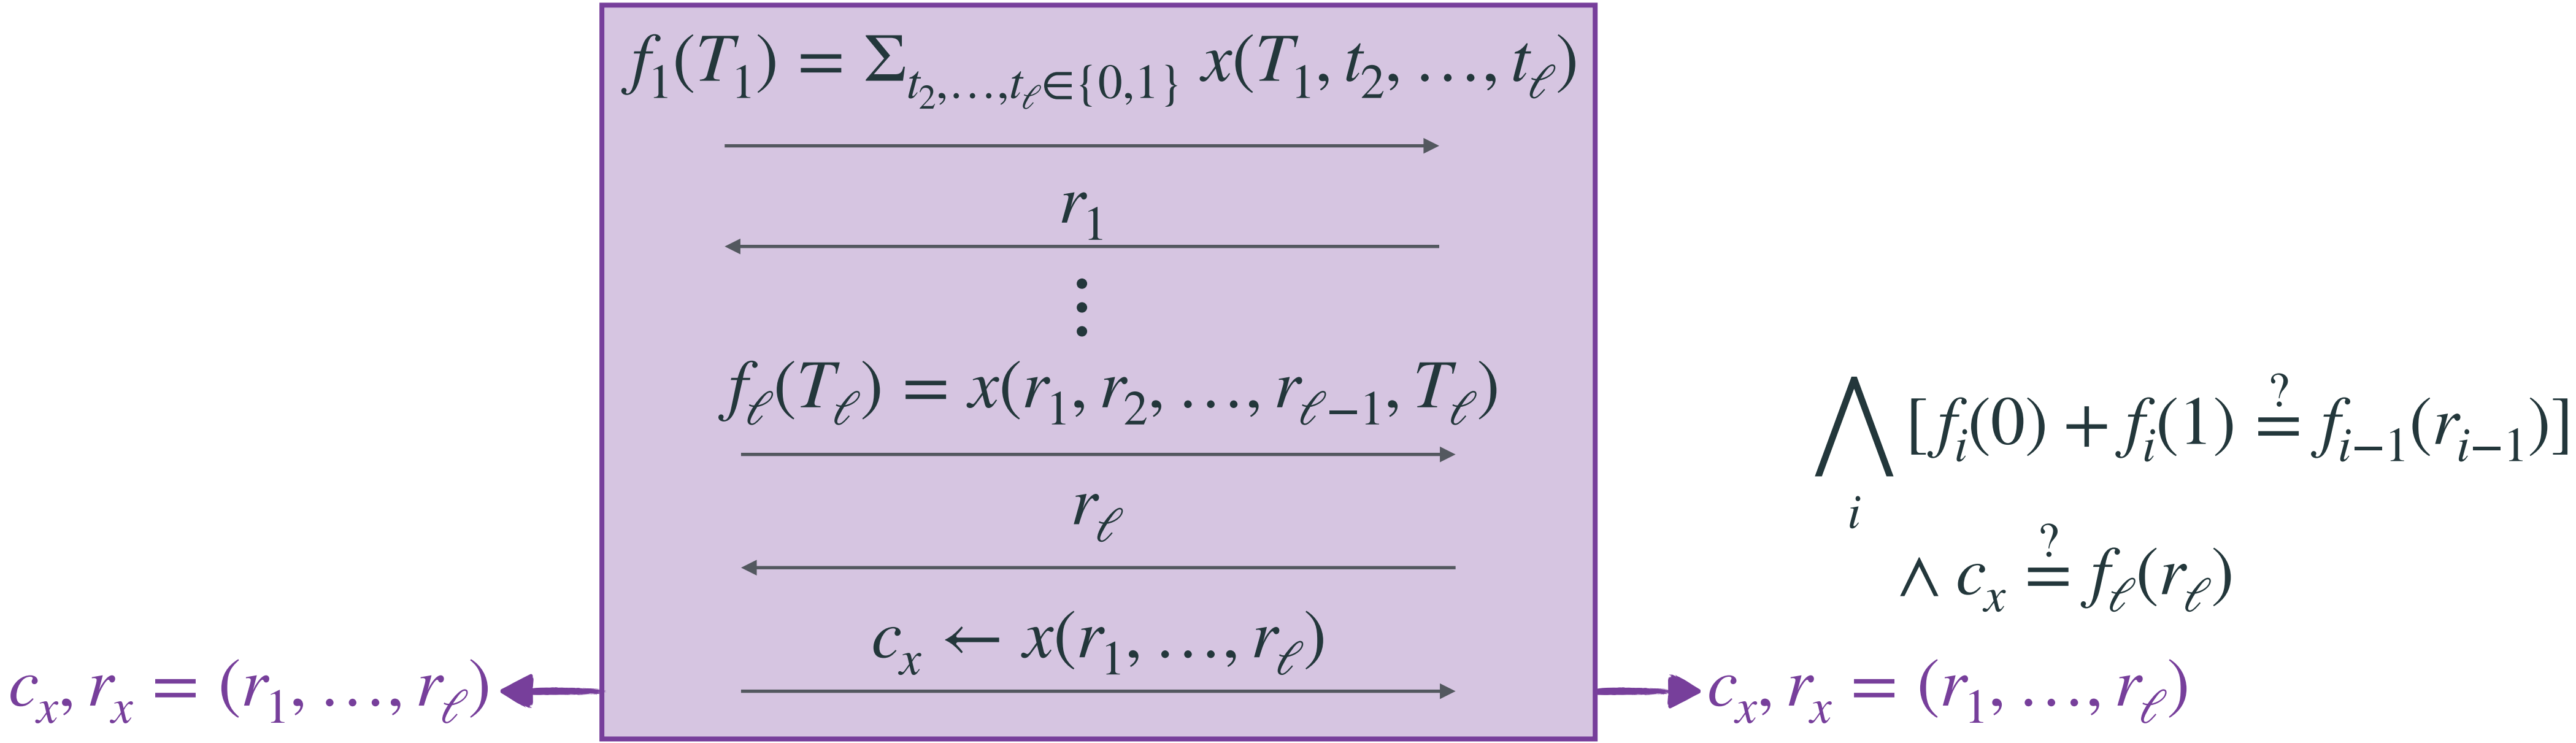


Sumcheck protocol as VE-based IP [LFKN92]

$x \in \mathbb{F}[T_1, \dots, T_\ell]$ multilinear polynomial. **Goal:** prove $y = f(x) = \sum_{t_1, \dots, t_\ell \in \{0,1\}} x(t_1, \dots, t_\ell)$

$\mathcal{P}(x, y)$

$\mathcal{V}(y)$

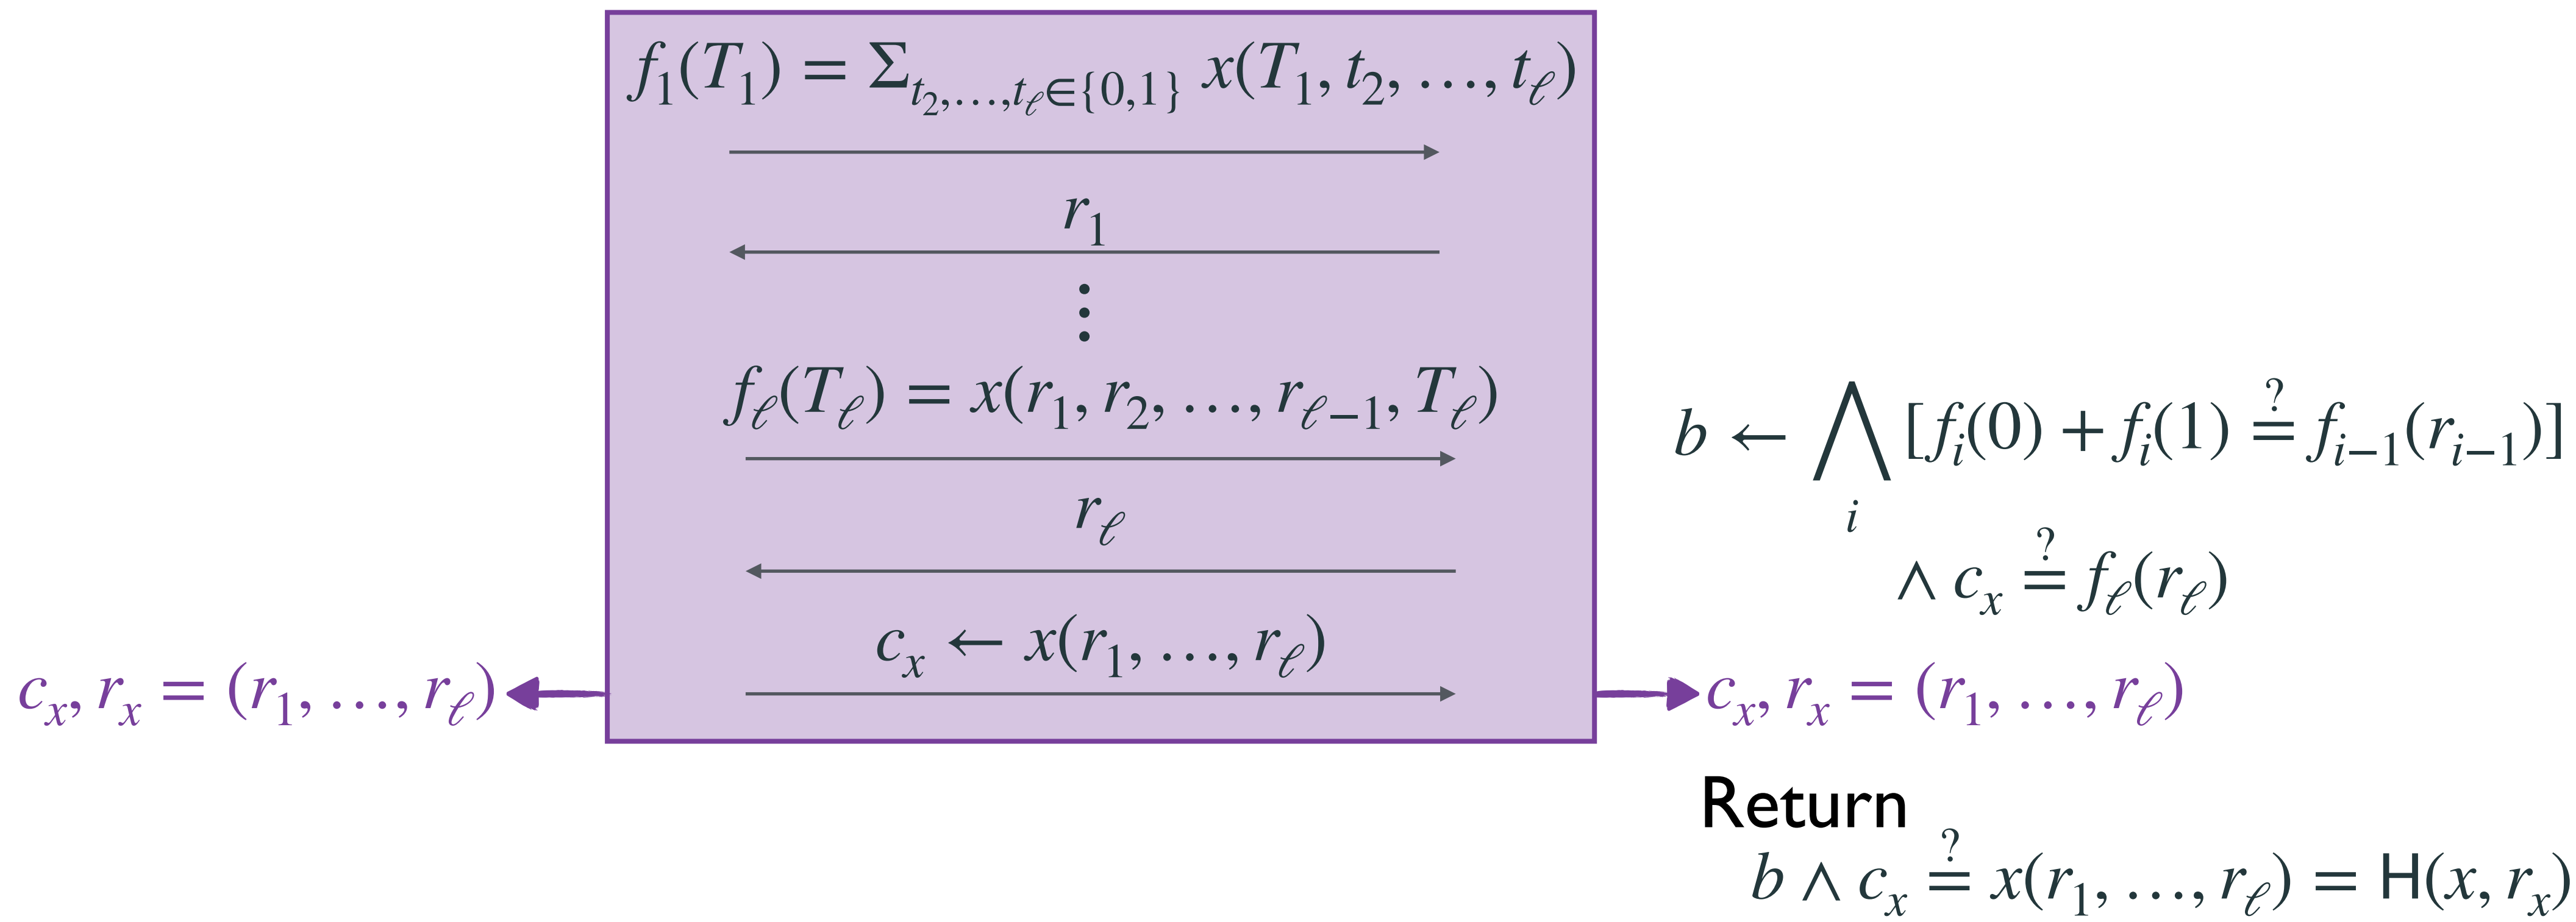


Sumcheck protocol as VE-based IP [LFKN92]

$x \in \mathbb{F}[T_1, \dots, T_\ell]$ multilinear polynomial. **Goal:** prove $y = f(x) = \sum_{t_1, \dots, t_\ell \in \{0,1\}} x(t_1, \dots, t_\ell)$

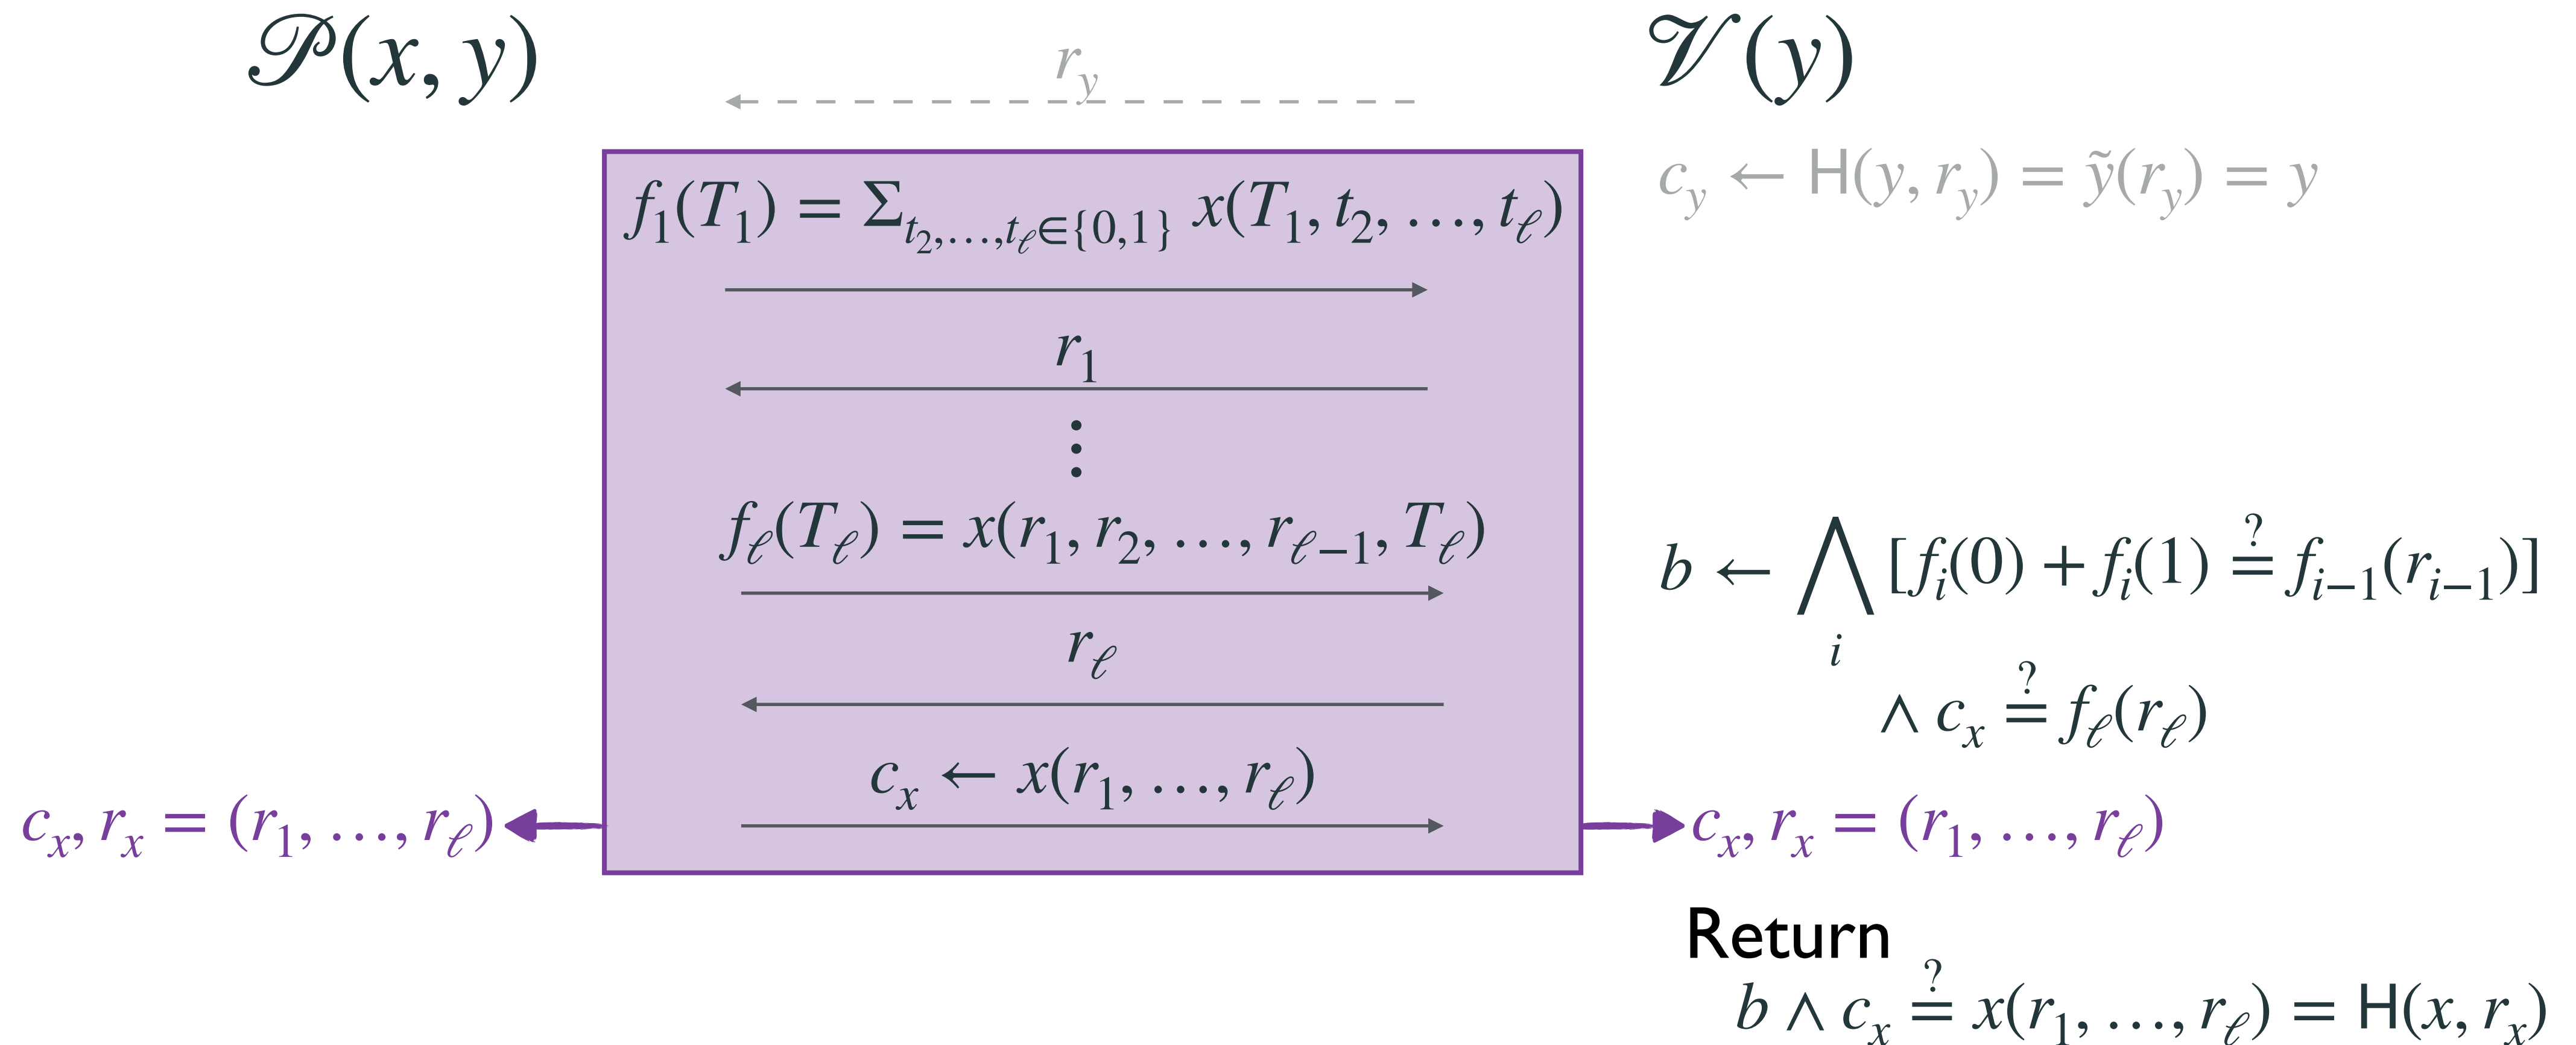
$\mathcal{P}(x, y)$

$\mathcal{V}(y)$



Sumcheck protocol as VE-based IP [LFKN92]

$x \in \mathbb{F}[T_1, \dots, T_\ell]$ multilinear polynomial. **Goal:** prove $y = f(x) = \sum_{t_1, \dots, t_\ell \in \{0,1\}} x(t_1, \dots, t_\ell)$



GKR as VE-based IP

$\mathcal{P}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

r_y

$\mathcal{V}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

GKR as VE-based IP

$\mathcal{P}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

Sumcheck claim: $c_y = \tilde{V}_d(r_y) = \sum_{b_1, b_2 \in \{0,1\}^\ell} \tilde{add}_{d-1}(r_y, b_1, b_2)(\tilde{V}_{d-1}(b_1) + \tilde{V}_{d-1}(b_2)) + \tilde{mult}_{d-1}(r_y, b_1, b_2)(\tilde{V}_{d-1}(b_1) \cdot \tilde{V}_{d-1}(b_2))$

\vdots

Sumcheck claim: $c_1 = \tilde{V}_1(r_1) = \sum_{b_1, b_2 \in \{0,1\}^\ell} \tilde{add}_1(r_1, b_1, b_2)(\tilde{x}(b_1) + \tilde{x}(b_2)) + \tilde{mult}_1(r_1, b_1, b_2)(\tilde{x}(b_1) \cdot \tilde{x}(b_2))$

$\mathcal{V}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

no need to know x

$\tilde{x}(r_x)$



GKR as VE-based IP

$\mathcal{P}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

$\mathcal{V}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

r_y

GKR

no need to know x

c_x, r_x

c_x, r_x, b

$\tilde{x}(r_x)$

GKR as VE-based IP

$\mathcal{P}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

$\mathcal{V}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y) = H(y, r_y)$

r_y



GKR

no need to know x

c_x, r_x



c_x, r_x, b

$\tilde{x}(r_x) = H(x, r_x) \stackrel{?}{=} c_x$

GKR as VE-based IP

$\mathcal{P}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y)$

$\mathcal{V}(f, x, y)$



$c_y \leftarrow \tilde{y}(r_y) = H(y, r_y)$

r_y

GKR

no need to know x

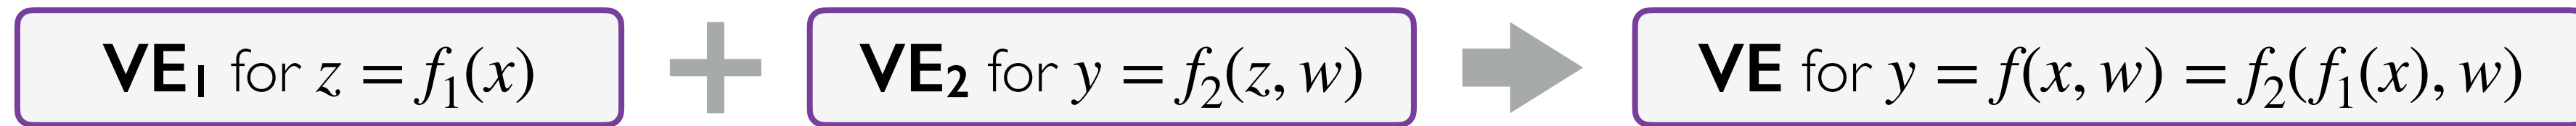
c_x, r_x

c_x, r_x, b

$\tilde{x}(r_x) = H(x, r_x) \stackrel{?}{=} c_x$

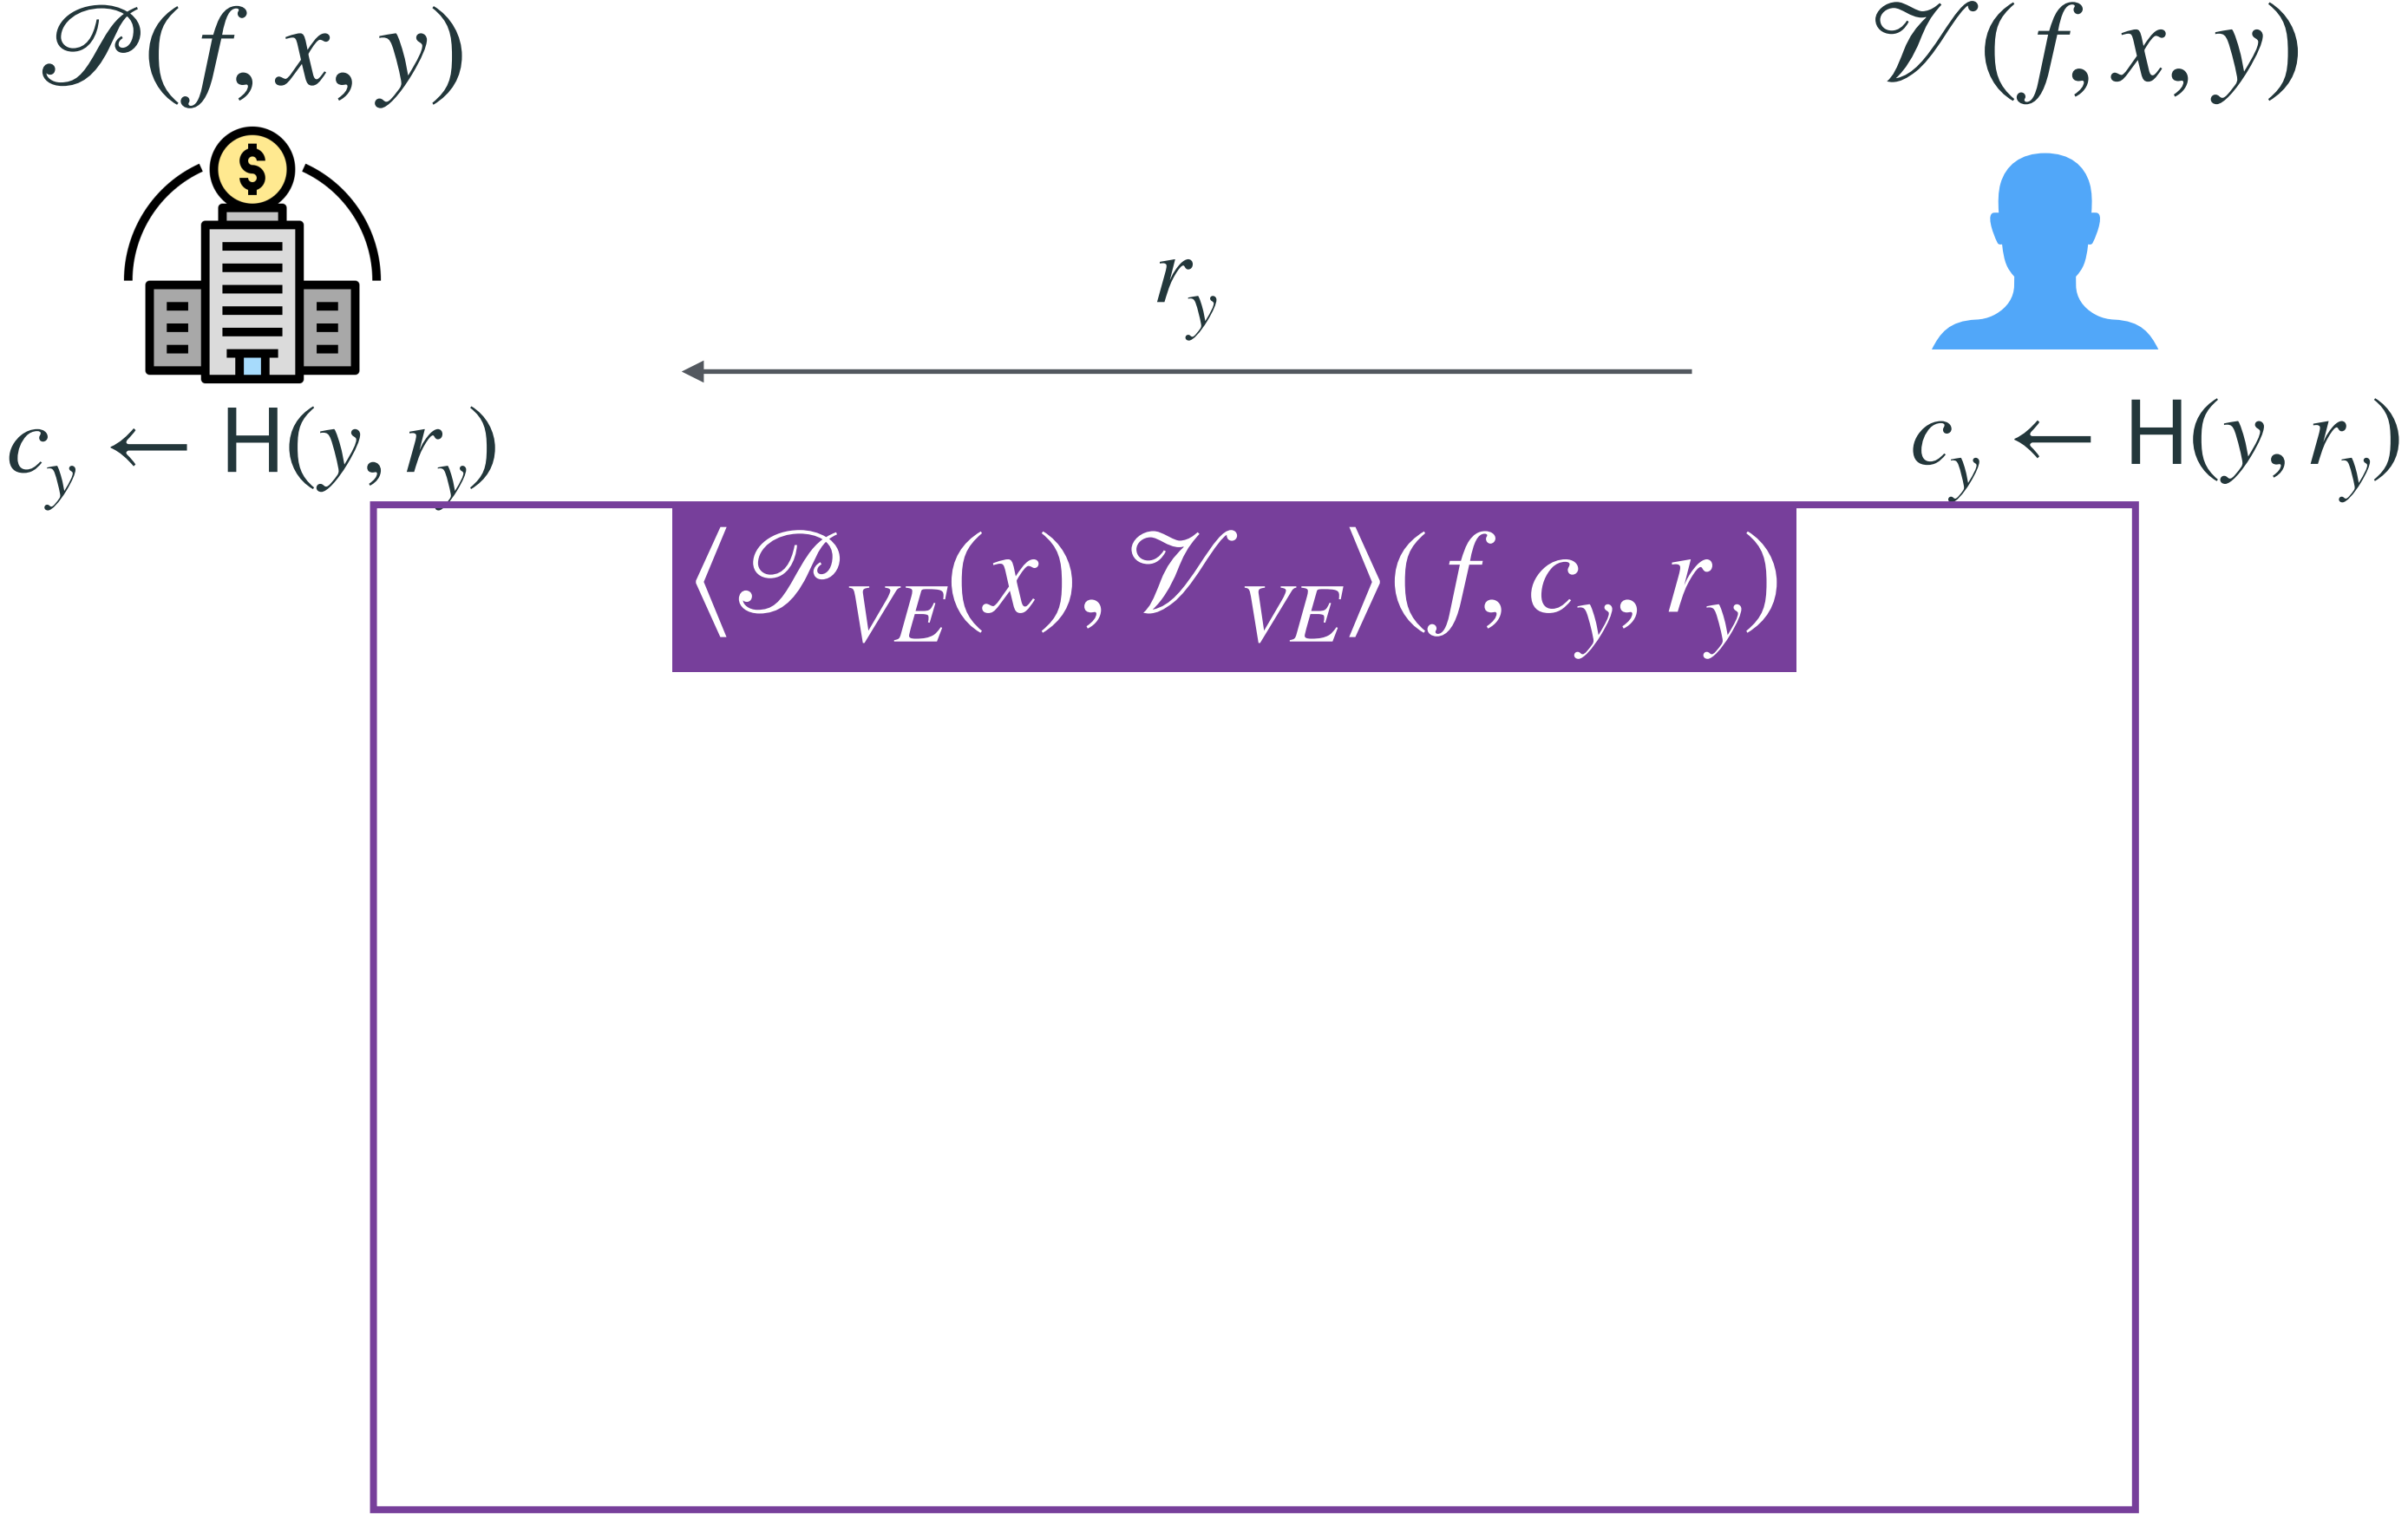
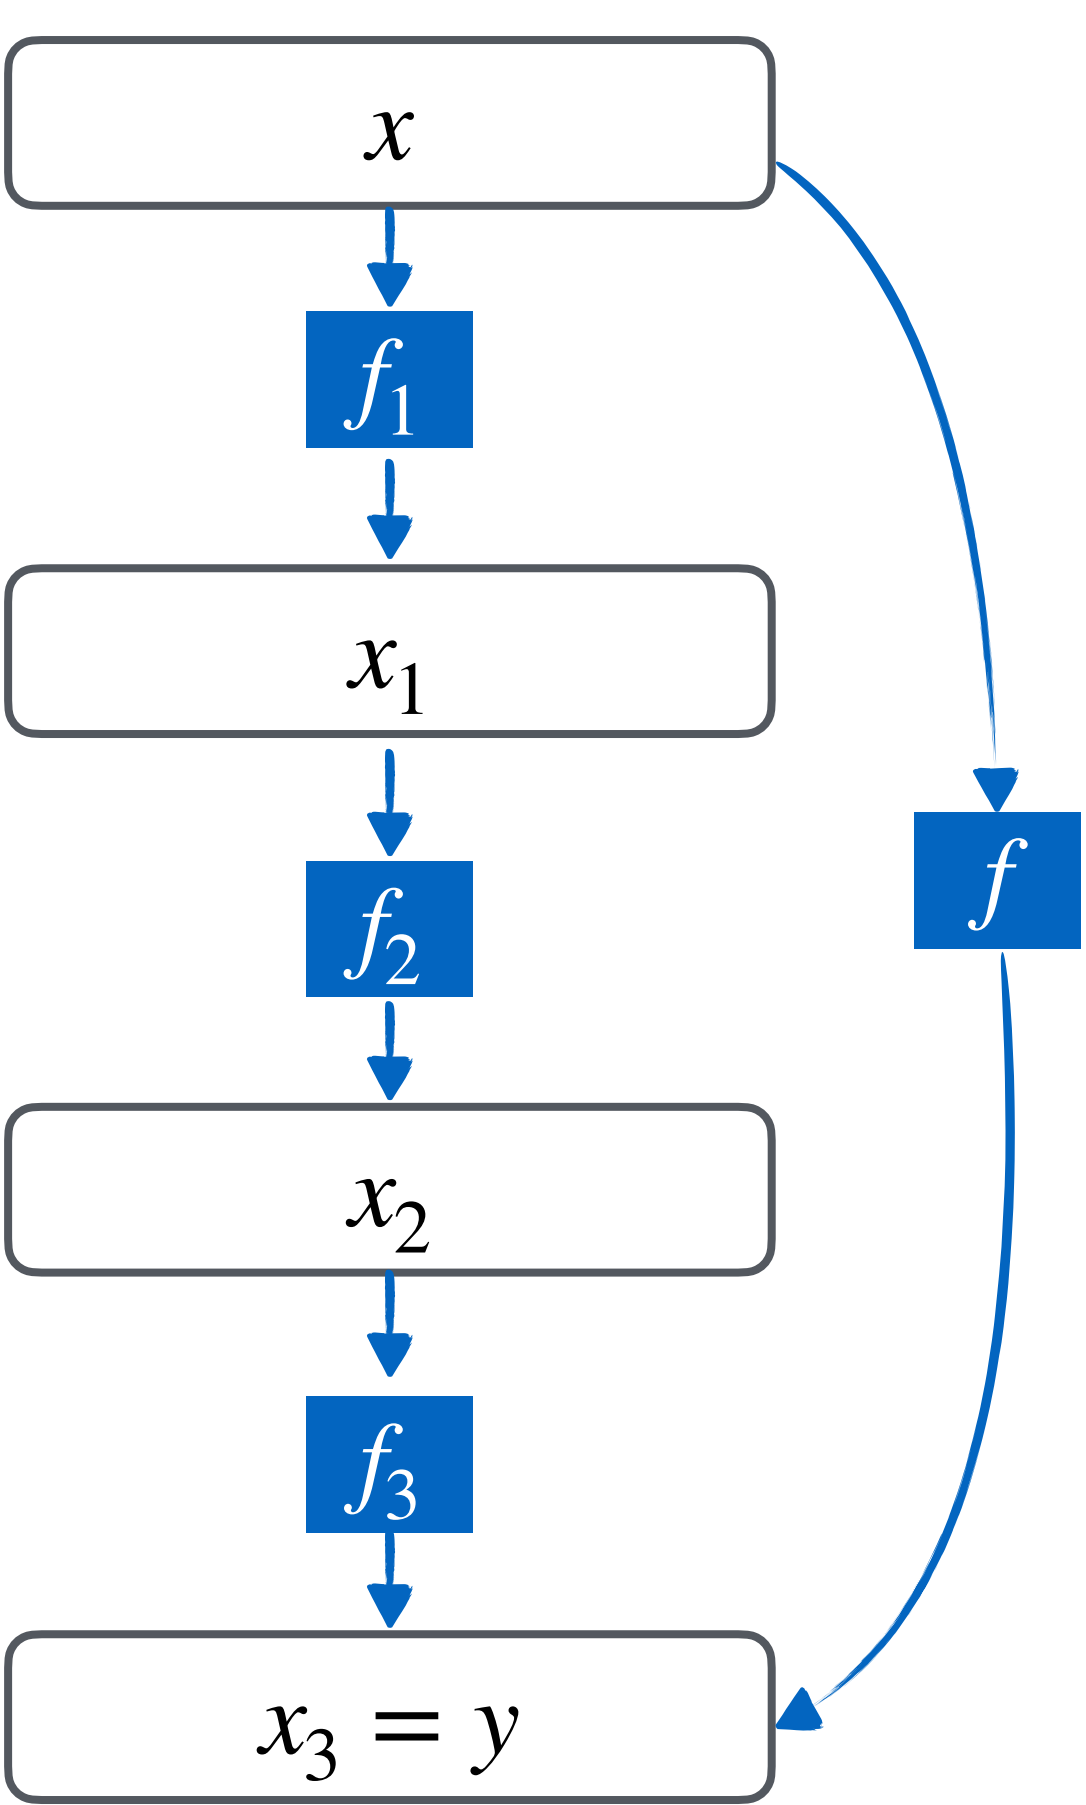
And even the layer subprotocols of GKR can be seen as VEs

Sequential composition of VEs



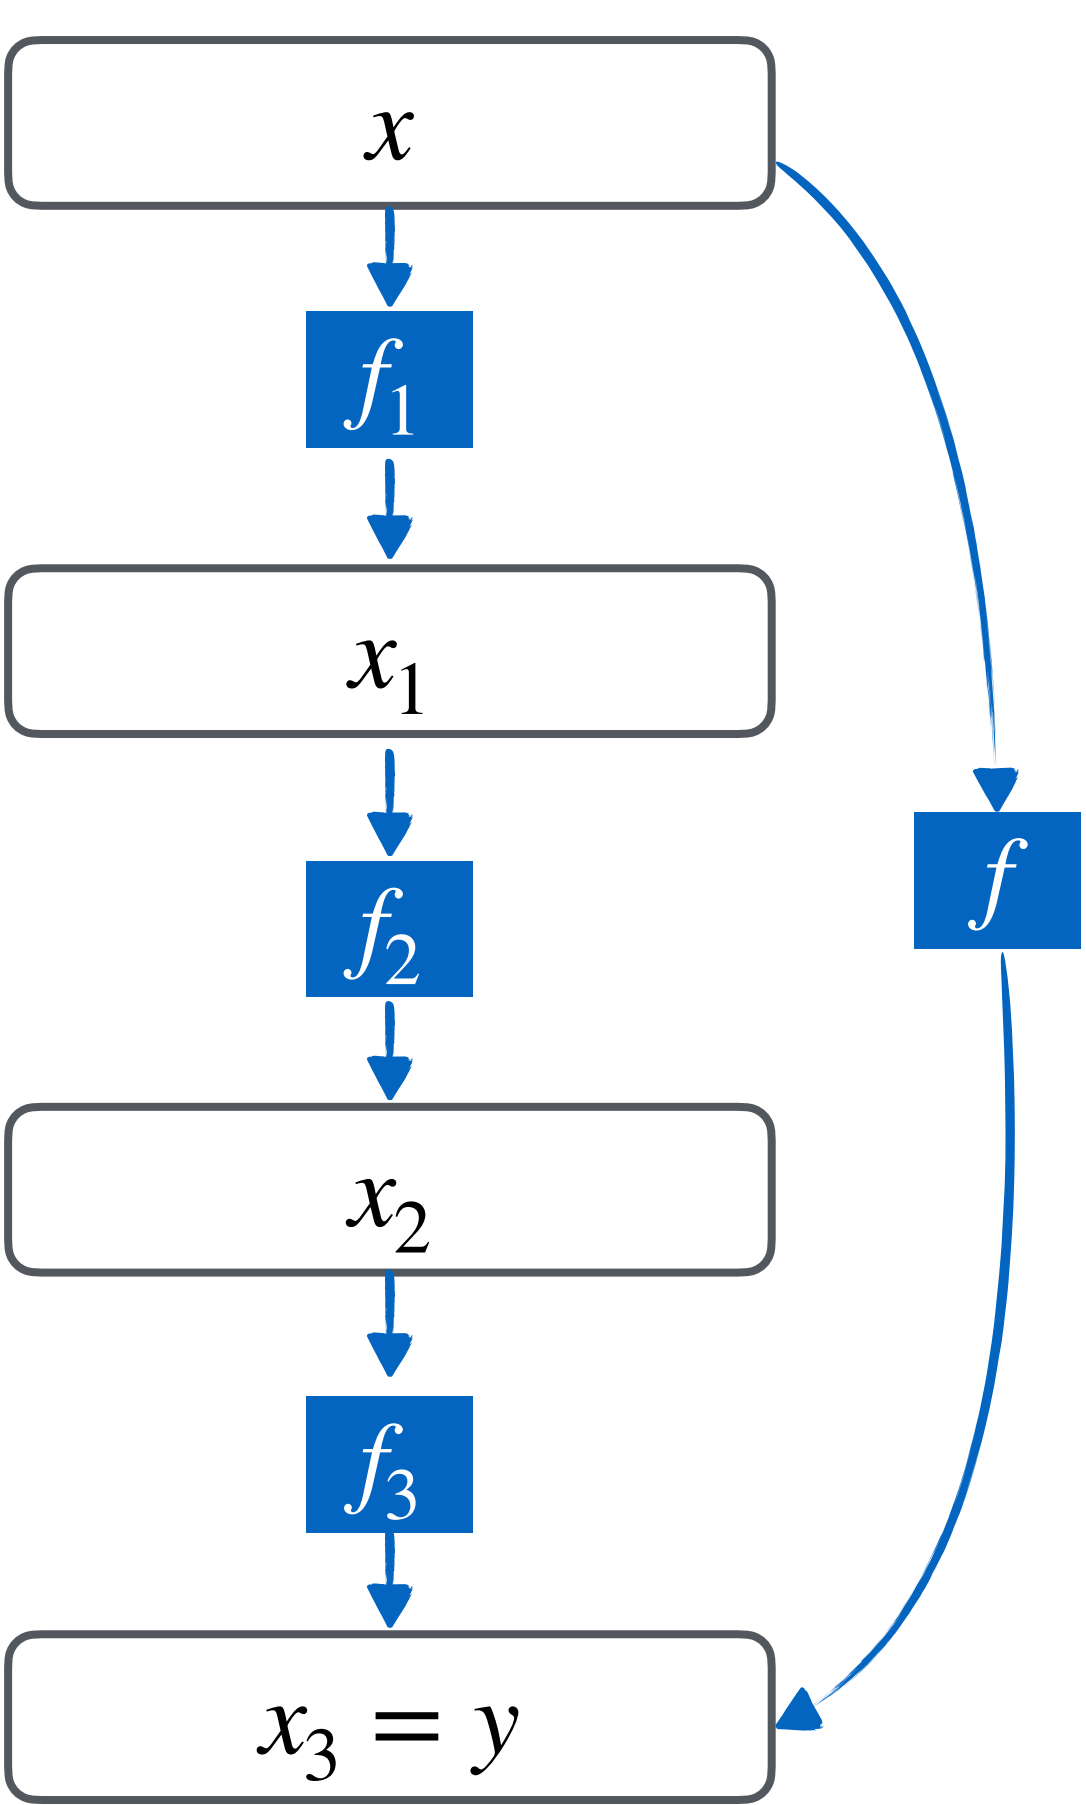
Sequential composition of VEs

$$\boxed{\text{VE}_1 \text{ for } z = f_1(x)} + \boxed{\text{VE}_2 \text{ for } y = f_2(z, w)} \rightarrow \boxed{\text{VE for } y = f(x, w) = f_2(f_1(x), w)}$$

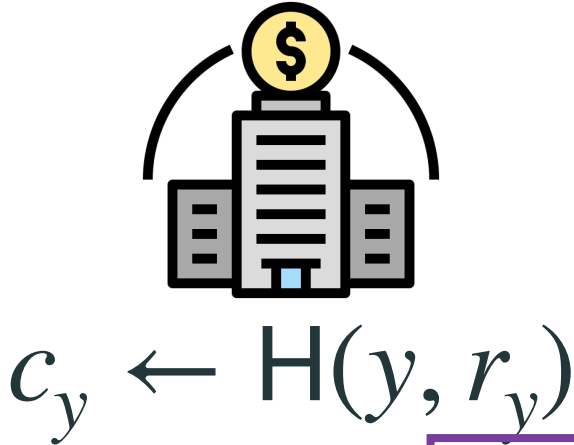


Sequential composition of VEs

$$\boxed{\text{VE}_1 \text{ for } z = f_1(x)} + \boxed{\text{VE}_2 \text{ for } y = f_2(z, w)} \rightarrow \boxed{\text{VE for } y = f(x, w) = f_2(f_1(x), w)}$$



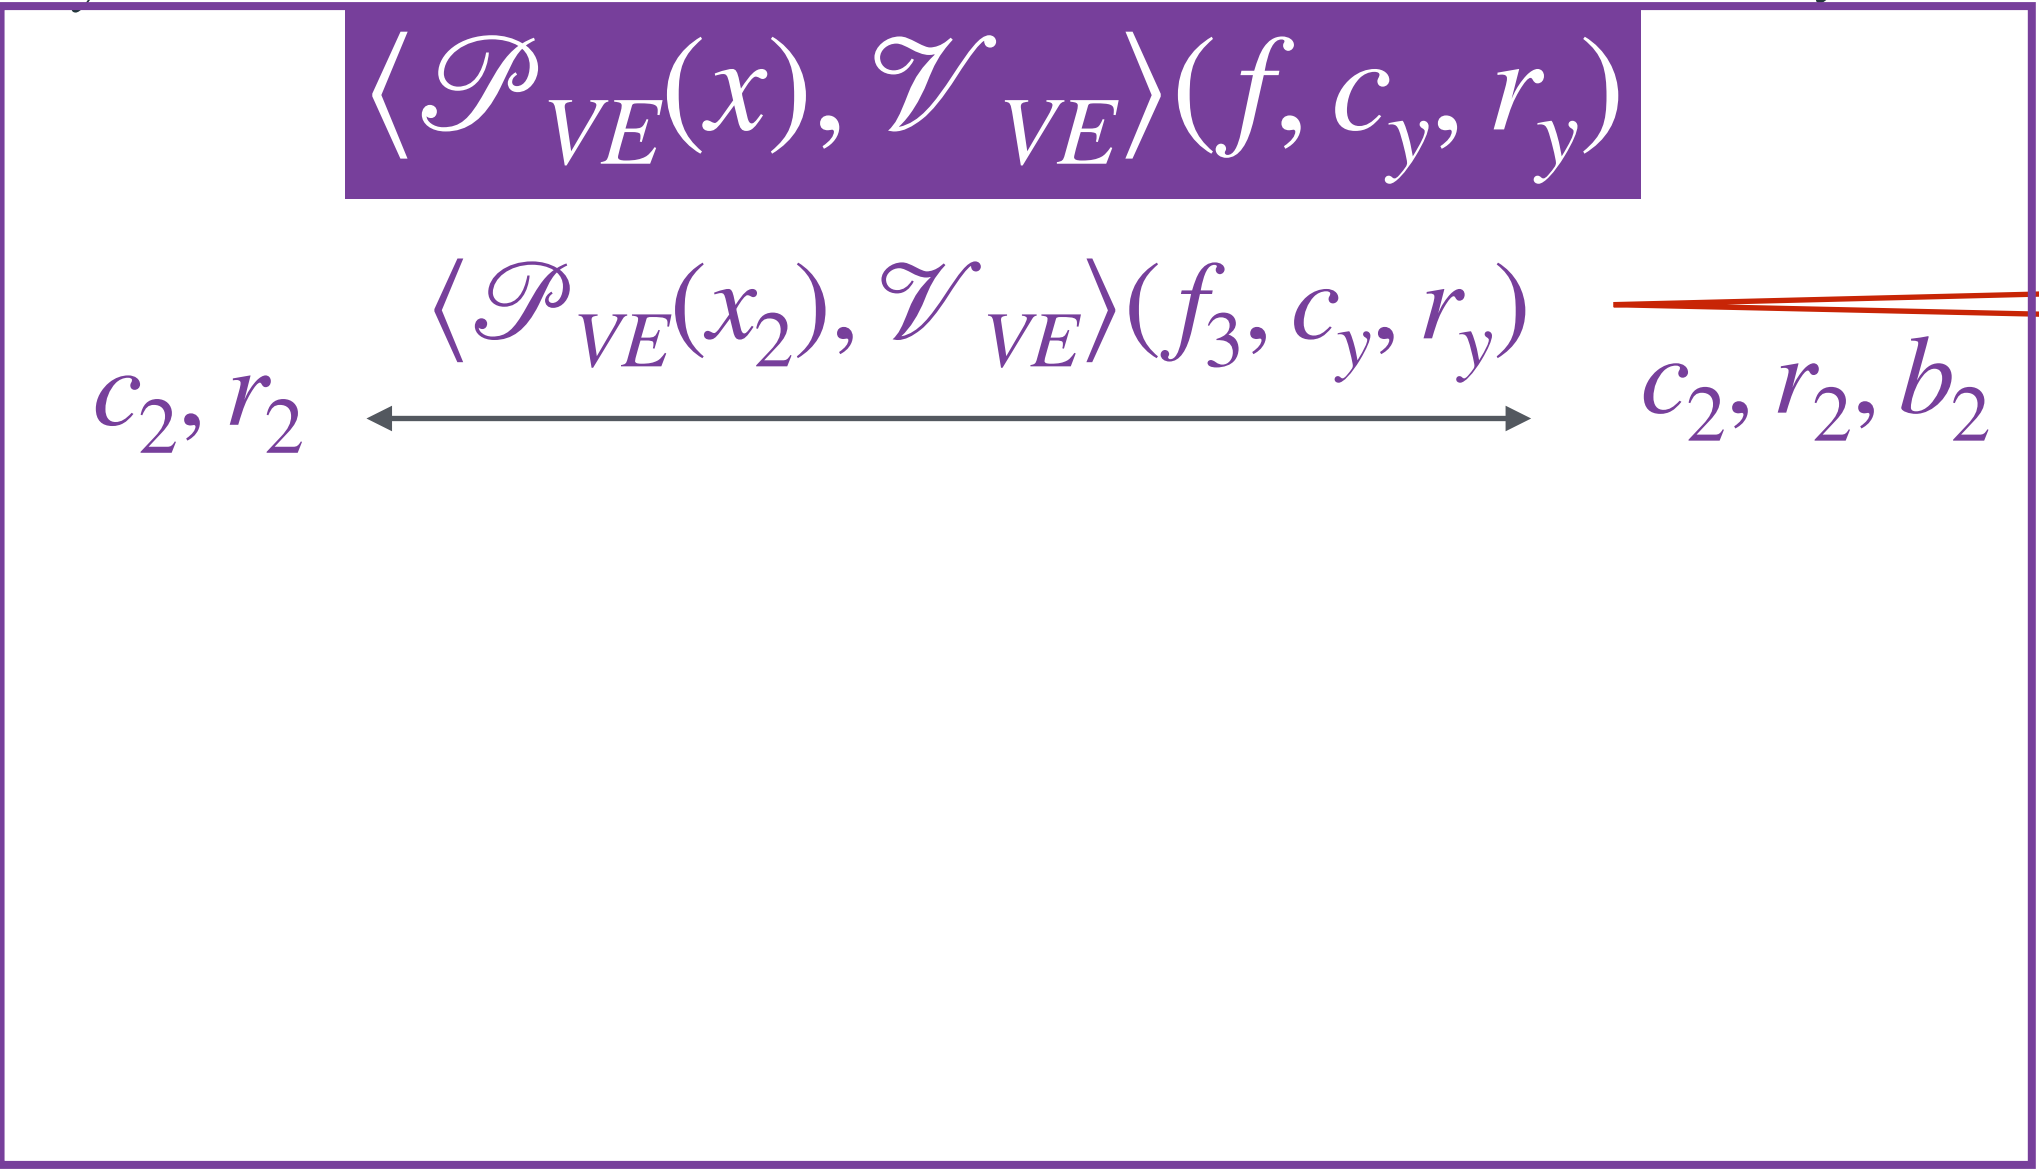
$\mathcal{P}(f, x, y)$



$\mathcal{V}(f, x, y)$



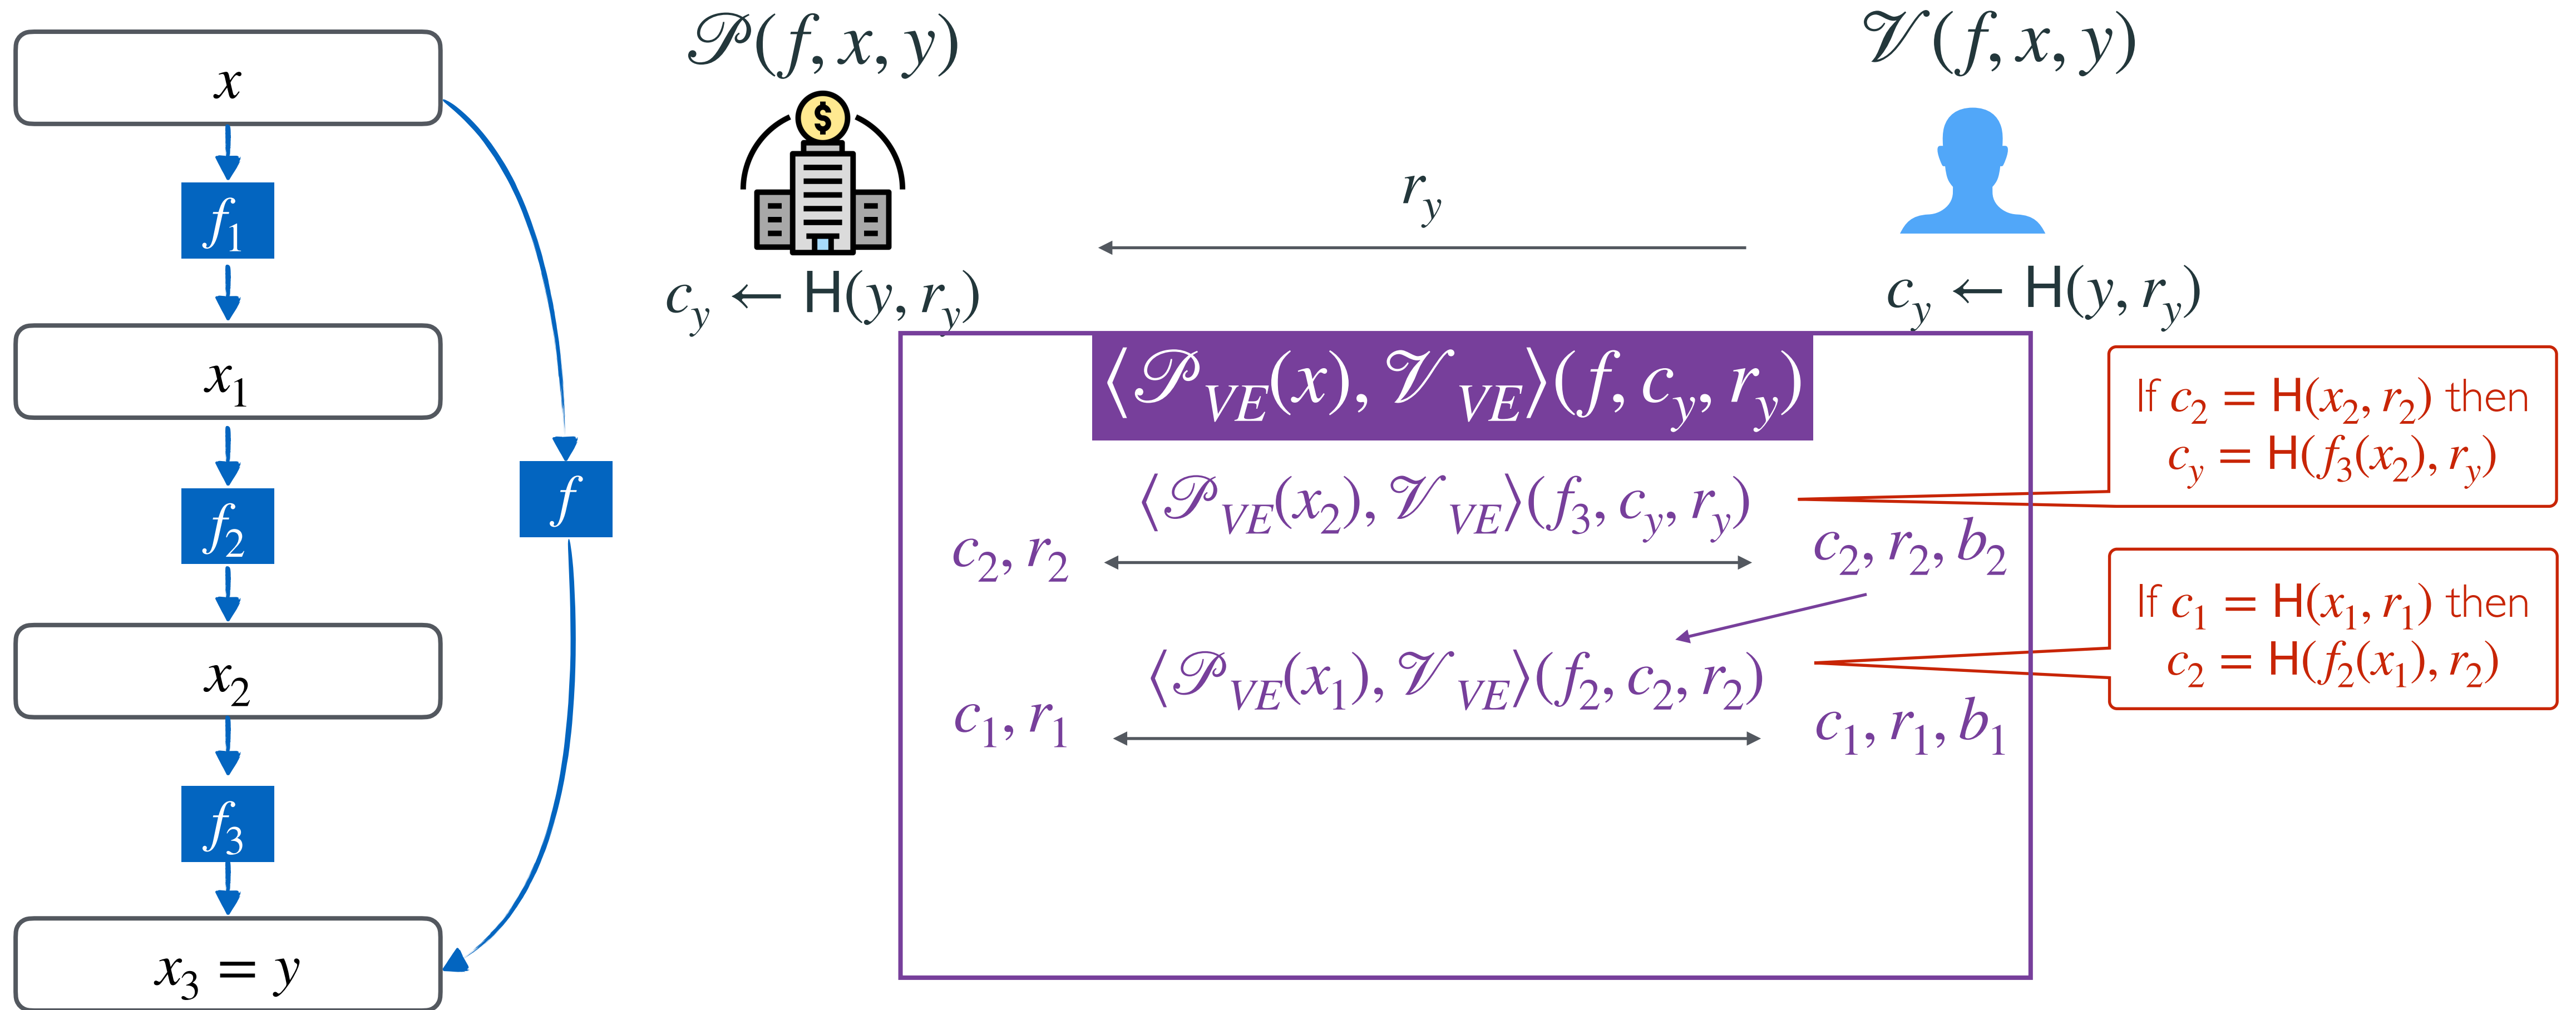
$c_y \leftarrow H(y, r_y)$



If $c_2 = H(x_2, r_2)$ then $c_y = H(f_3(x_2), r_y)$

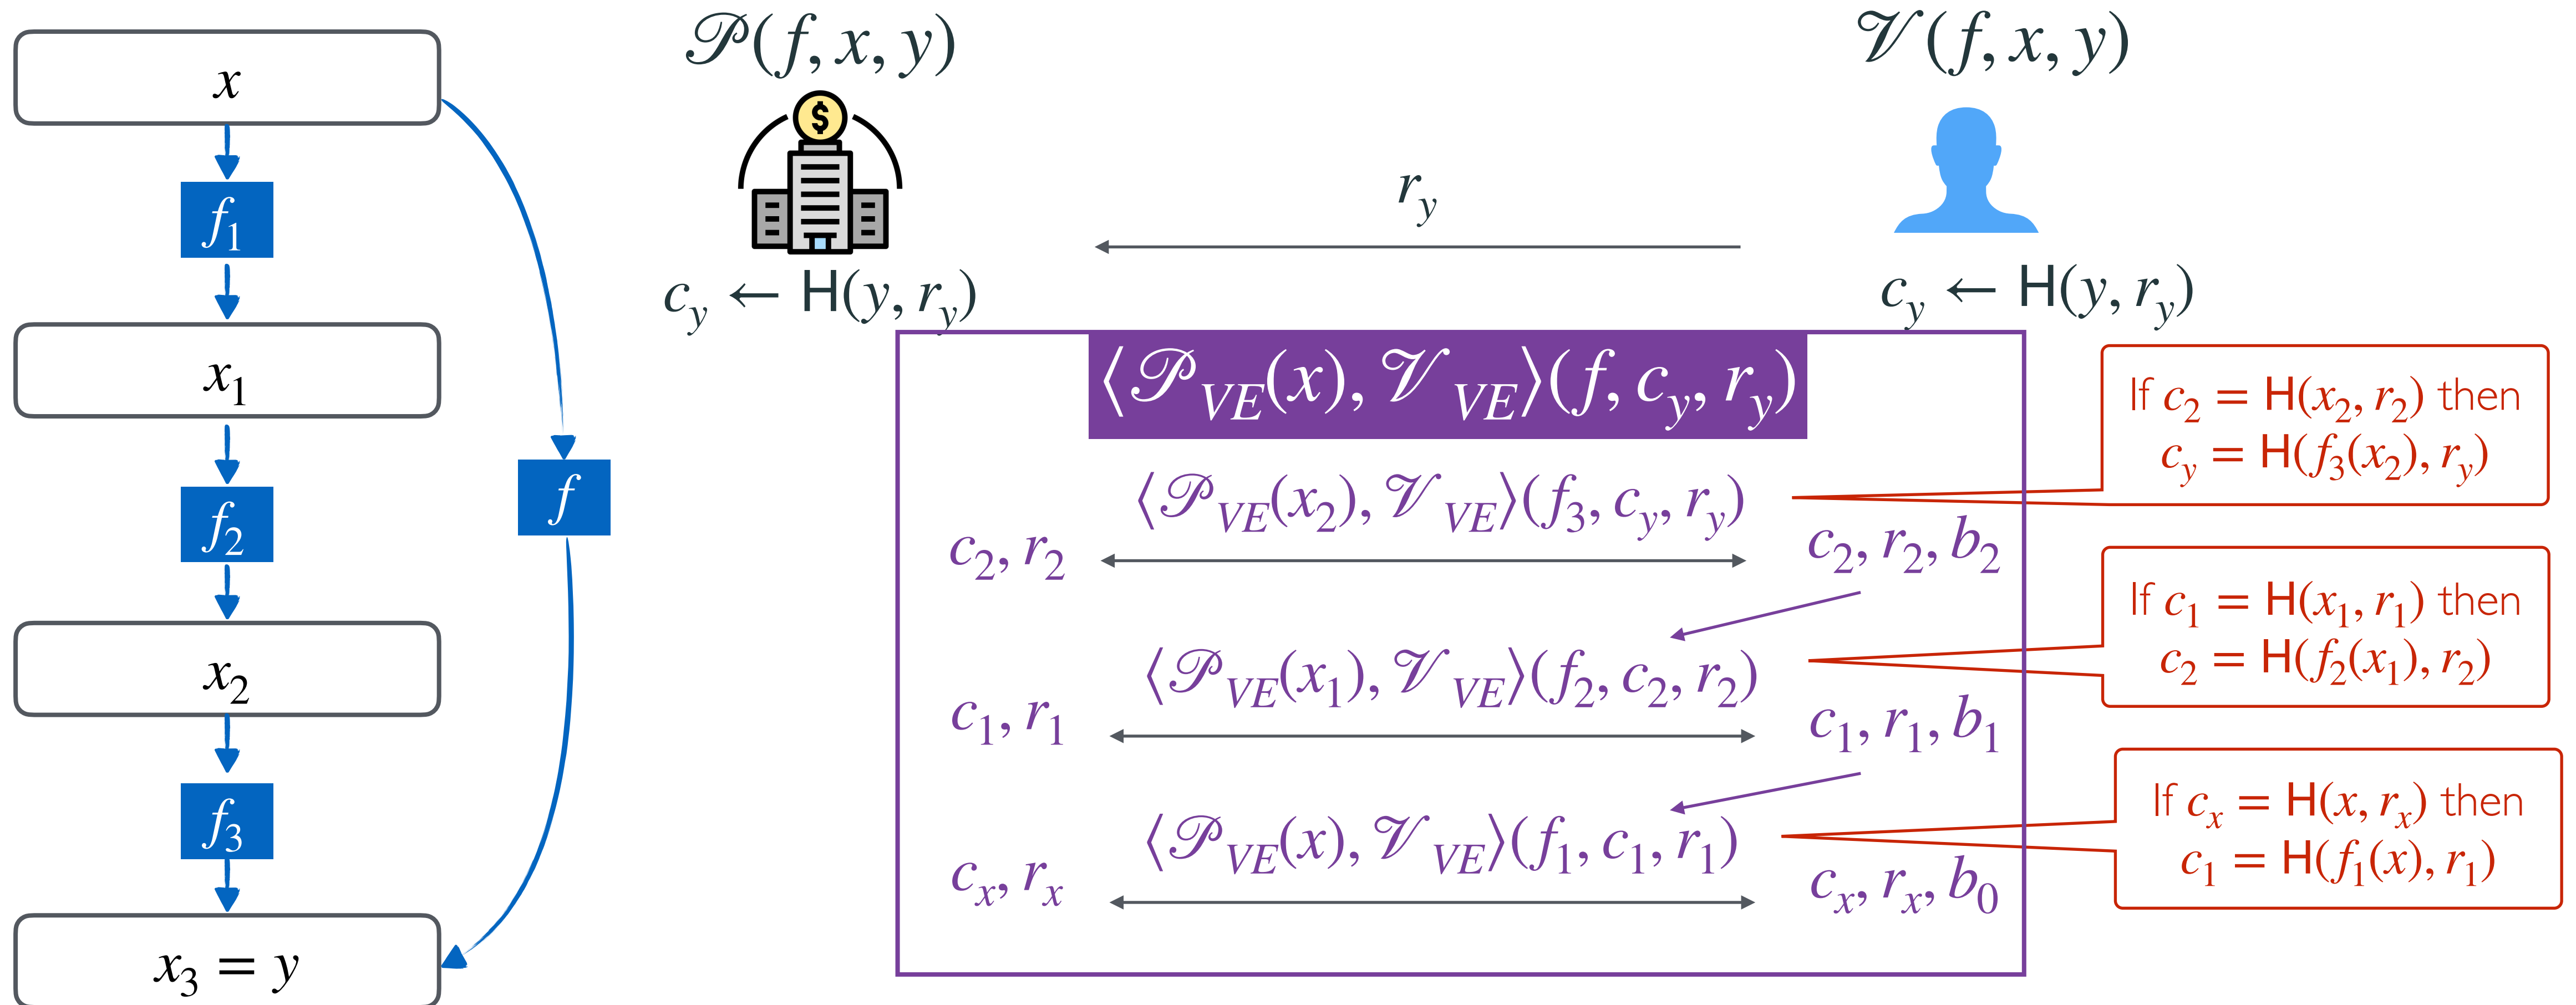
Sequential composition of VEs

$$\boxed{\text{VE}_1 \text{ for } z = f_1(x)} + \boxed{\text{VE}_2 \text{ for } y = f_2(z, w)} \rightarrow \boxed{\text{VE for } y = f(x, w) = f_2(f_1(x), w)}$$

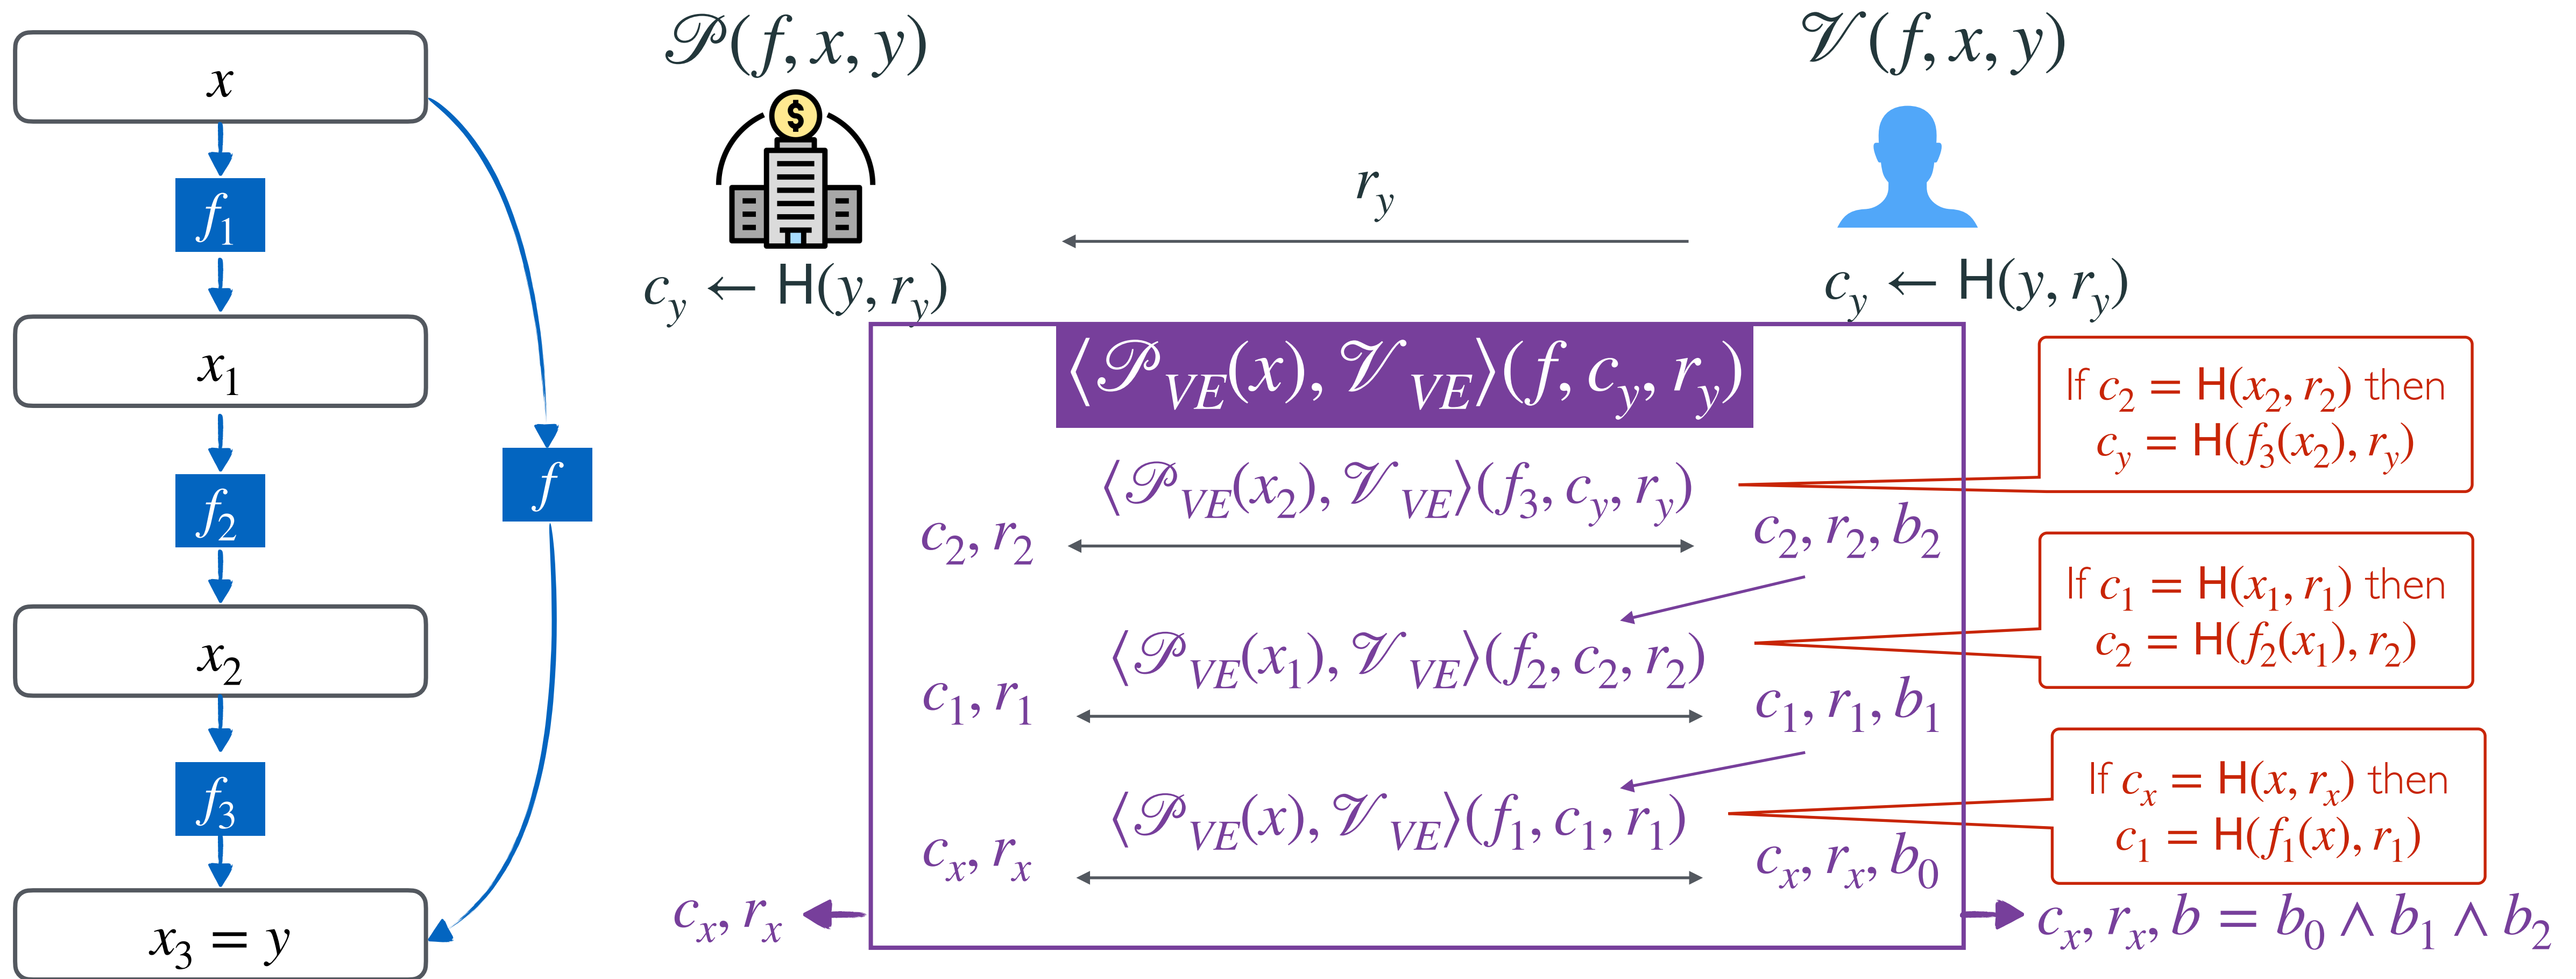
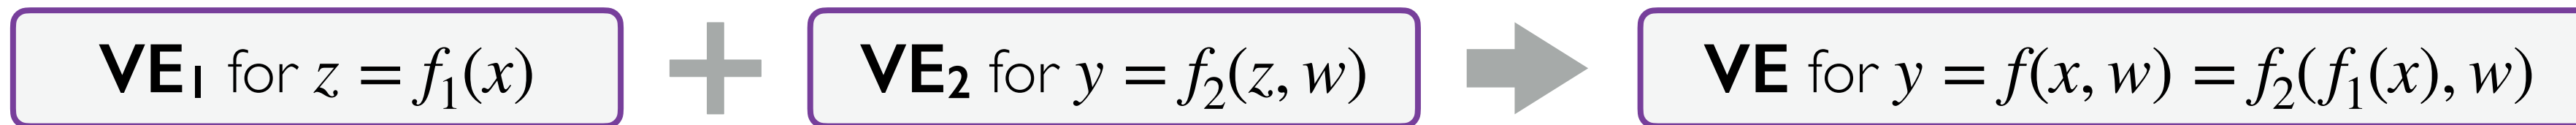


Sequential composition of VEs

$$\boxed{\text{VE}_1 \text{ for } z = f_1(x)} + \boxed{\text{VE}_2 \text{ for } y = f_2(z, w)} \rightarrow \boxed{\text{VE for } y = f(x, w) = f_2(f_1(x), w)}$$

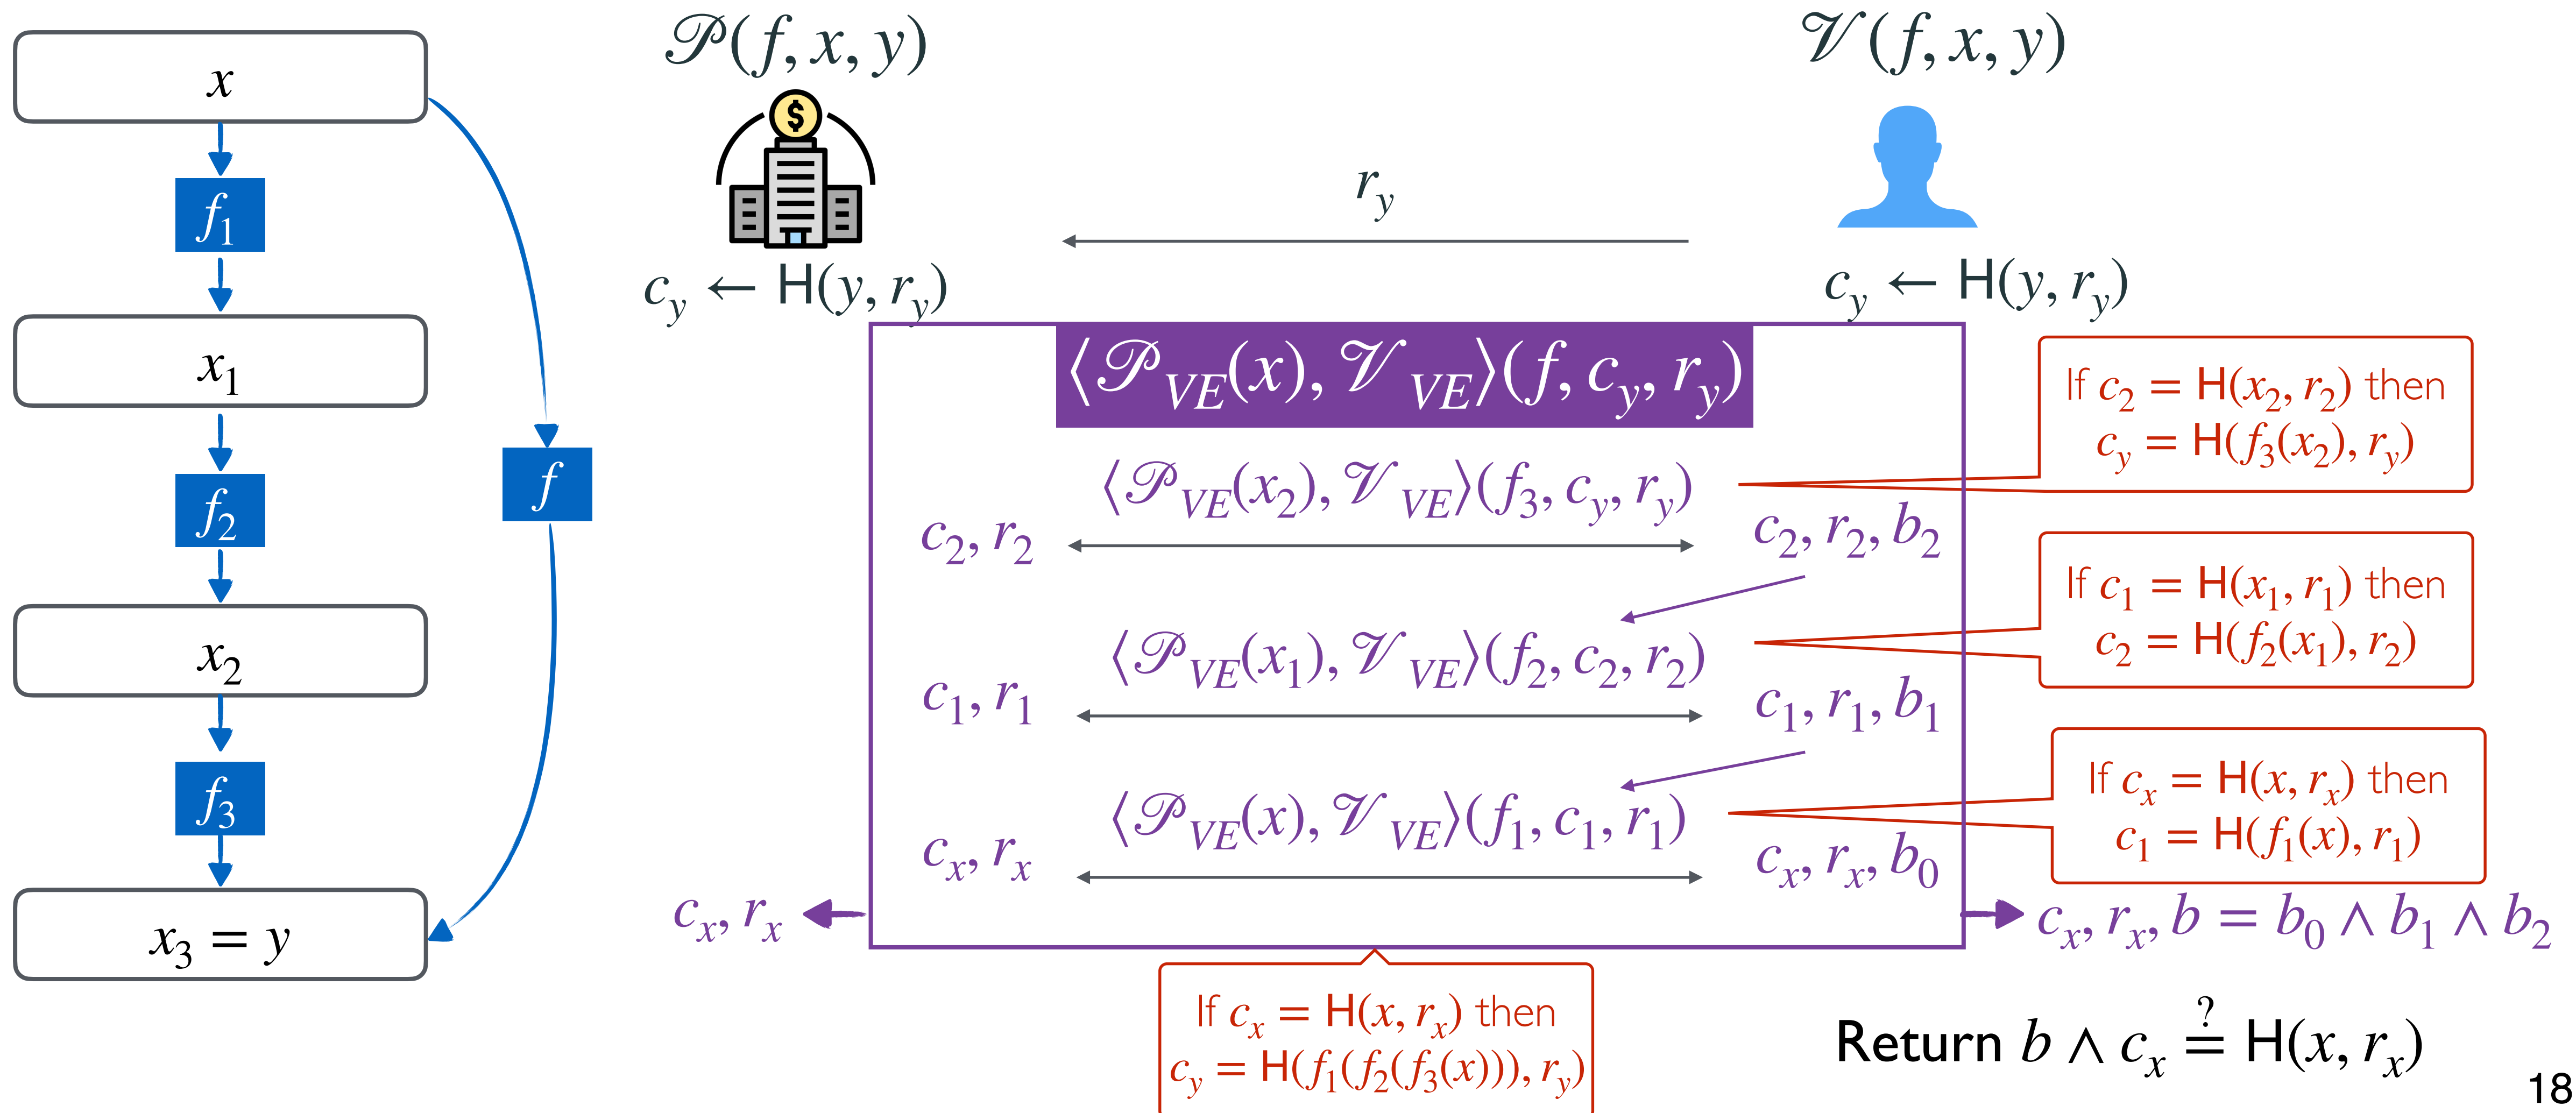
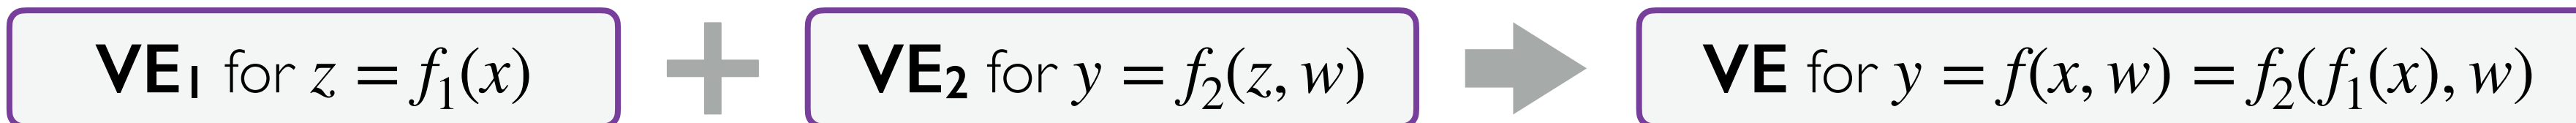


Sequential composition of VEs

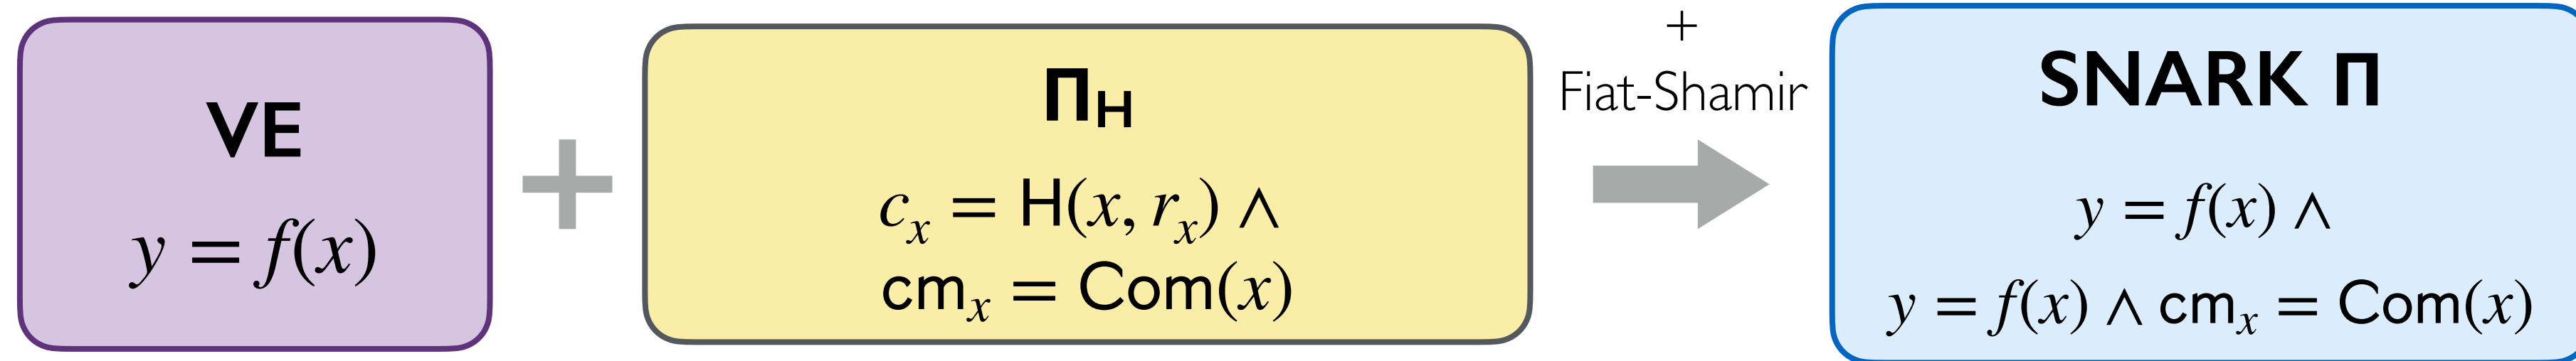


Return $b \wedge c_x \stackrel{?}{=} H(x, r_x)$

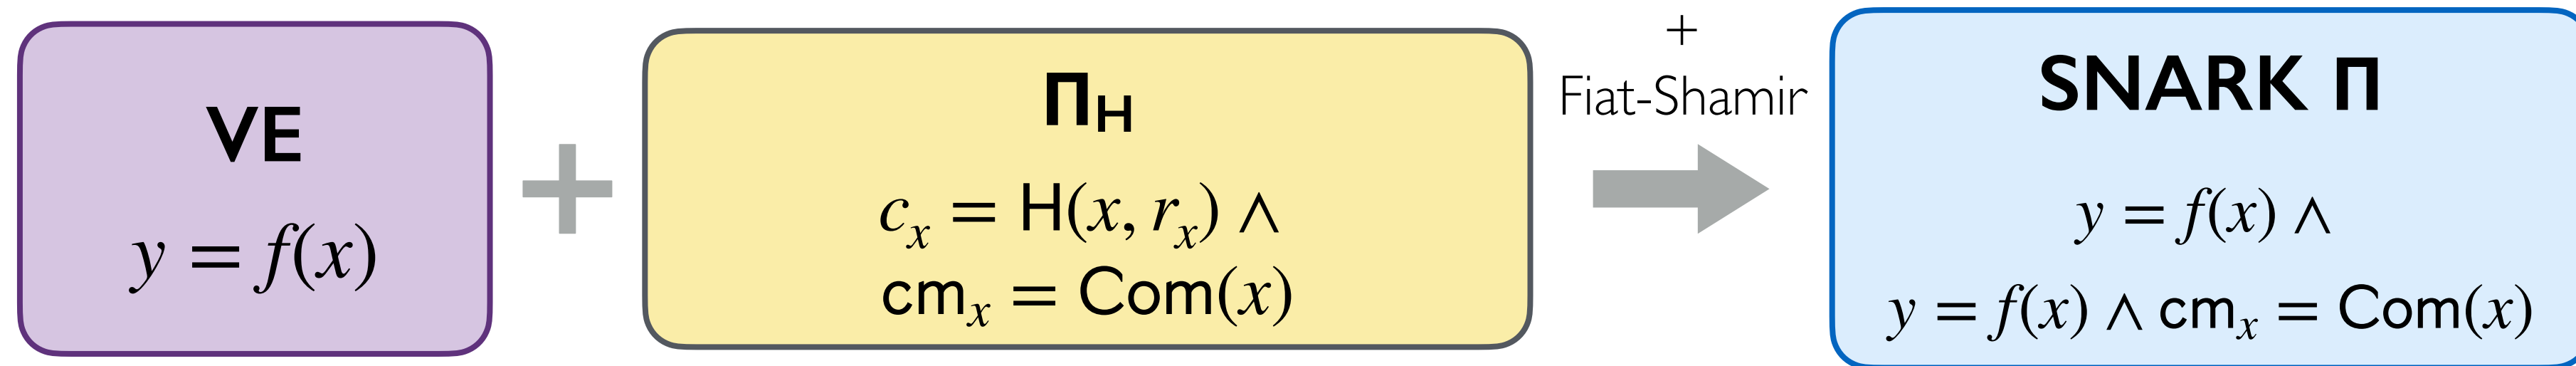
Sequential composition of VEs



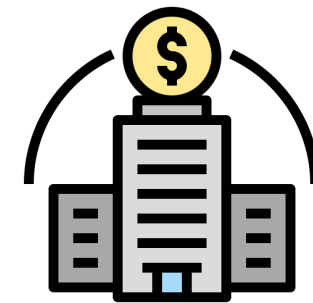
From VE-based IP to AoK [vSQL]



From VE-based IP to AoK [vSQL]



$\Pi . \mathcal{P}(f, x, y)$



$c_y \leftarrow H(y, r_y)$

c_x, r_x

$\langle \mathcal{P}_{VE}(x), \mathcal{V}_{VE} \rangle (f, c_y, r_y)$

$\Pi . \mathcal{V}(f, cm_x, y)$



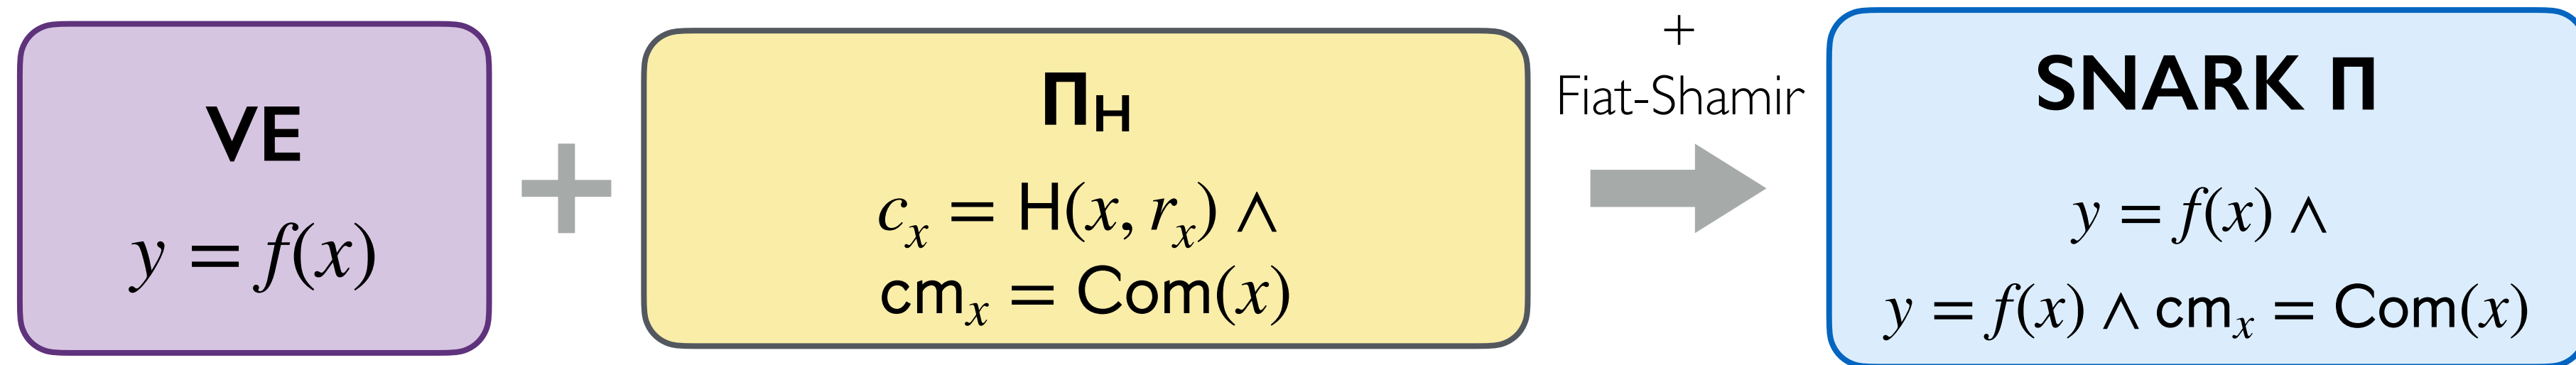
$c_y \leftarrow H(y, r_y)$

c_x, r_x, b

VE-based IP

r_y

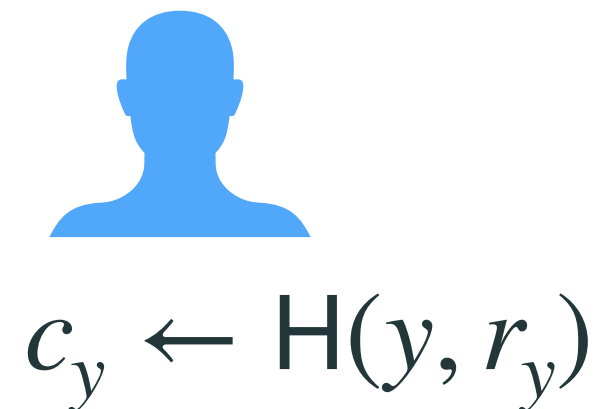
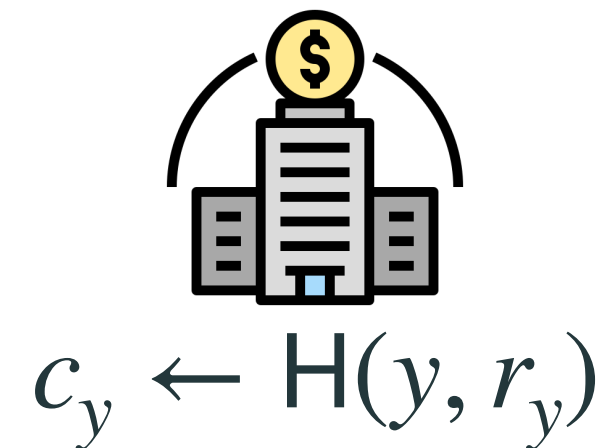
From VE-based IP to AoK [vSQL]



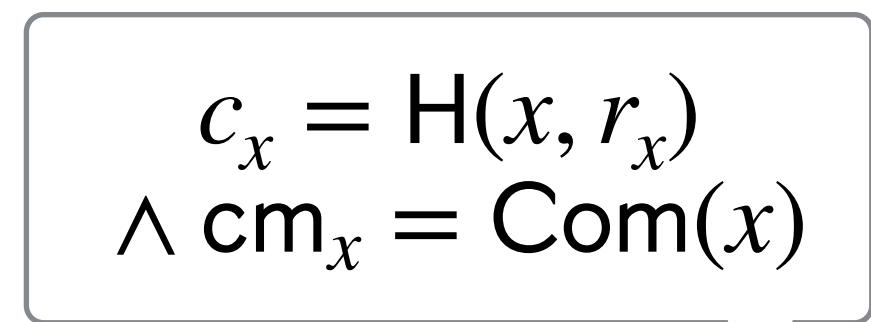
$\Pi . \mathcal{P}(f, x, y)$

$\Pi . \mathcal{V}(f, \text{cm}_x, y)$

VE-based IP



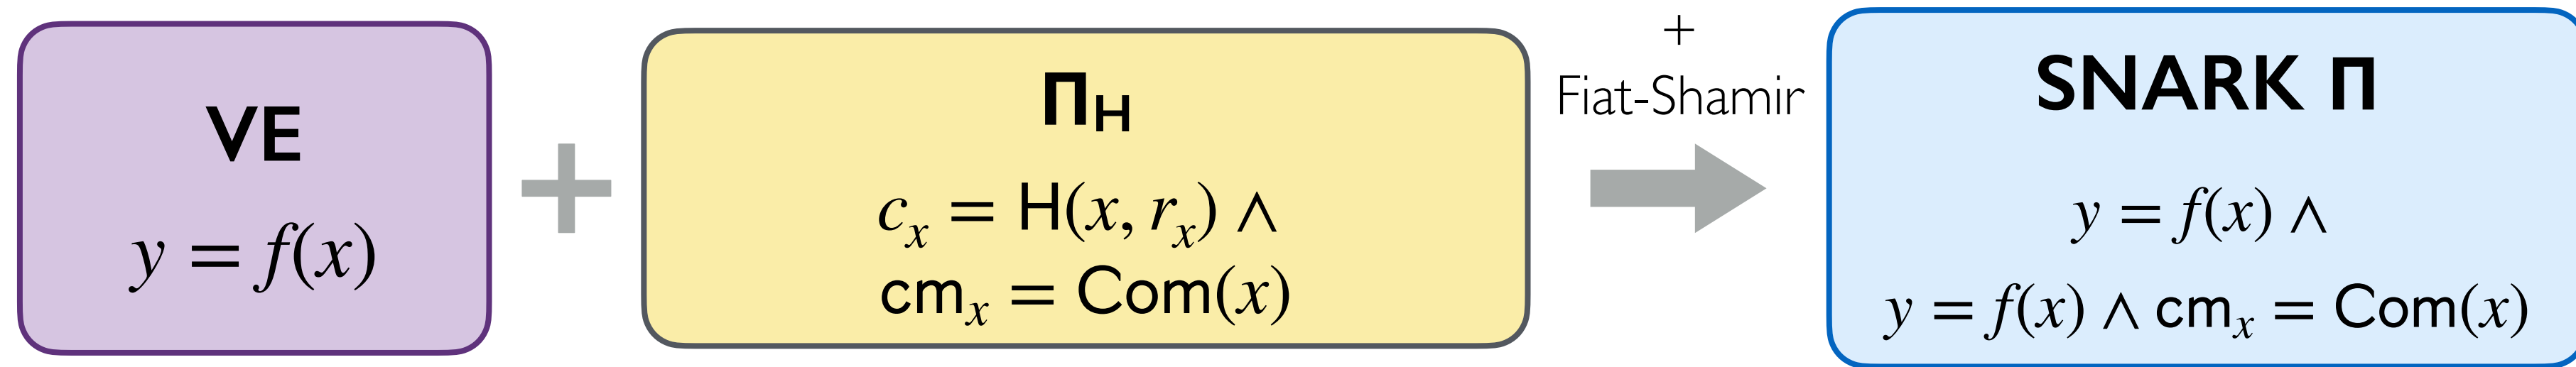
Prove that input fingerprint is correct w.r.t. cm_x



$\pi_x = \Pi_H . \mathcal{P}((\text{cm}_x, c_x, r_x), x, \pi_x)$

Return $b \wedge \Pi_H . \text{Ver}(\text{cm}_x, c_x, r_x, \pi_x)$

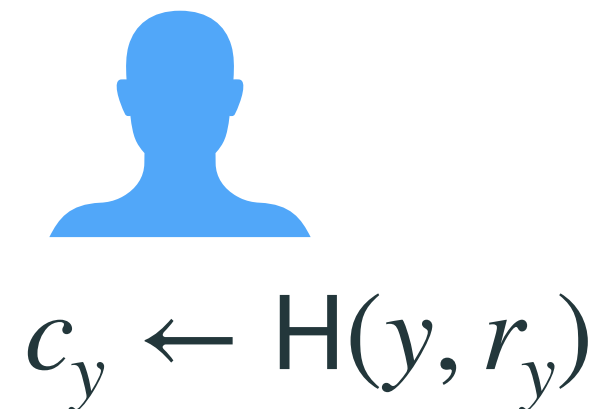
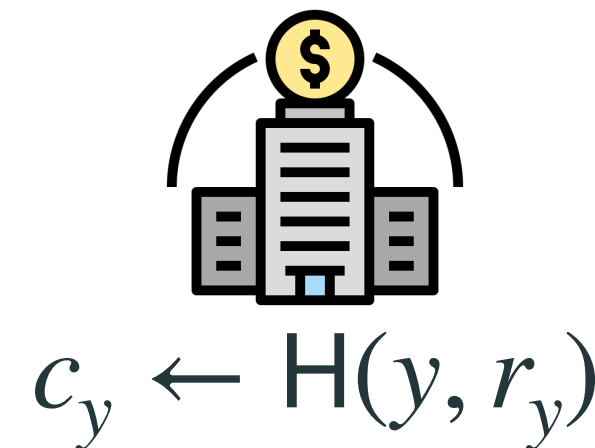
From VE-based IP to AoK [vSQL]



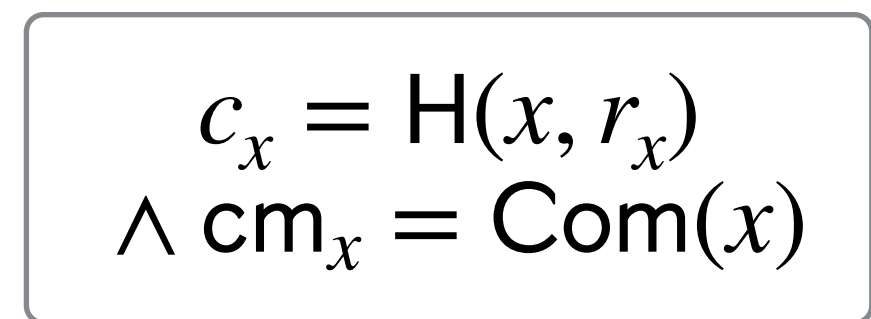
$\Pi . \mathcal{P}(f, x, y)$

$\Pi . \mathcal{V}(f, \text{cm}_x, y)$

VE-based IP



Prove that input fingerprint is correct w.r.t. cm_x

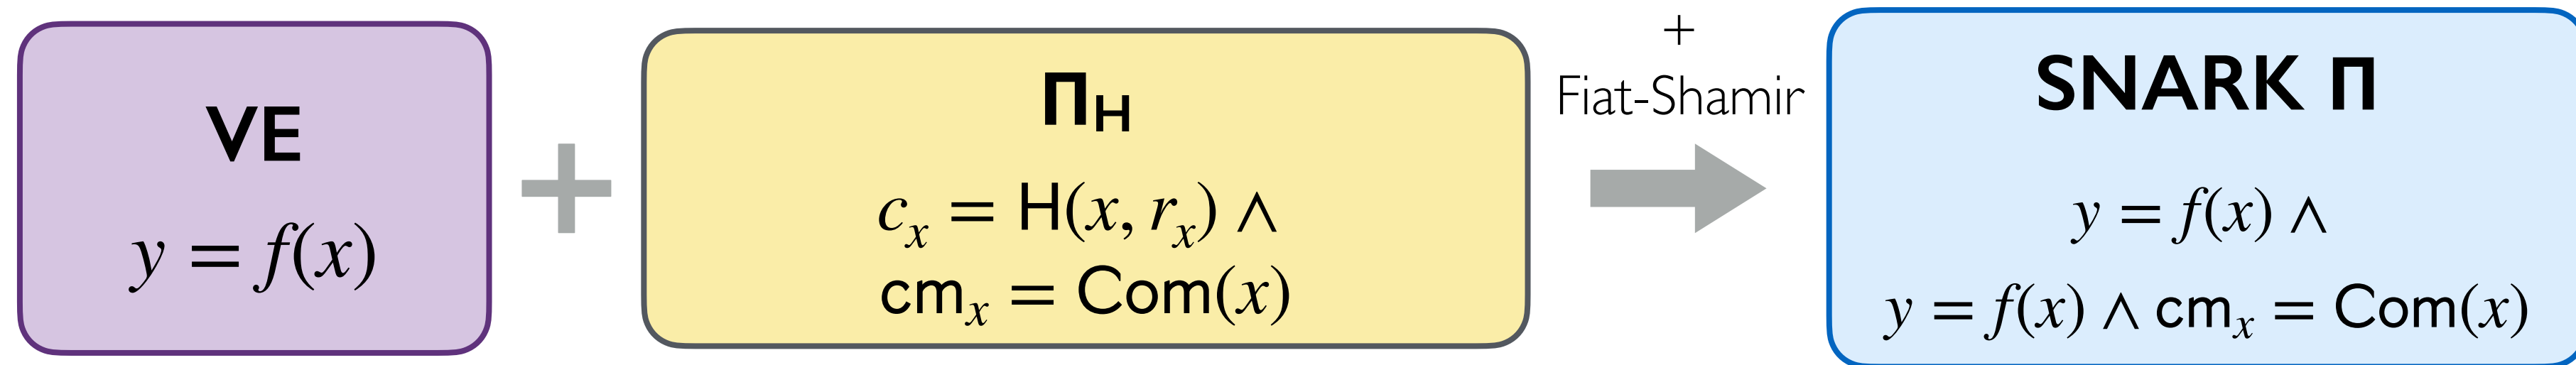


$\pi_x = \Pi_H . \mathcal{P}((\text{cm}_x, c_x, r_x), x, \pi_x)$

Return $b \wedge \Pi_H . \text{Ver}(\text{cm}_x, c_x, r_x, \pi_x)$

When $H(x, r_x) = \tilde{x}(r_x)$, Π_H can be instantiated with a multilinear polynomial commitment

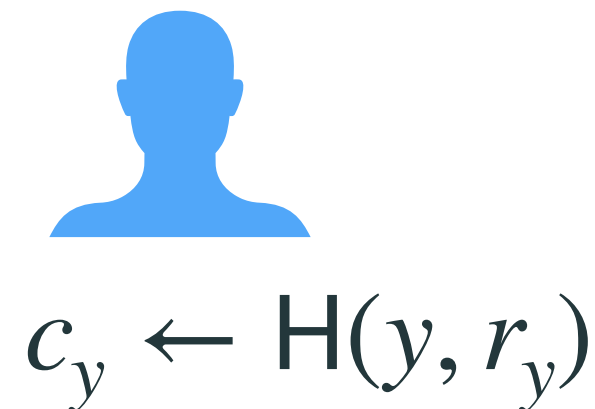
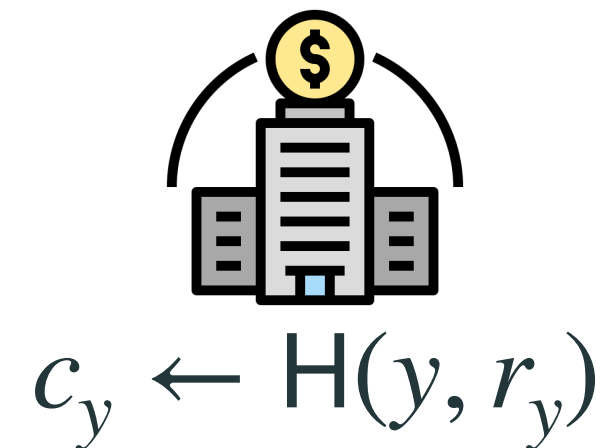
From VE-based IP to AoK [vSQL]



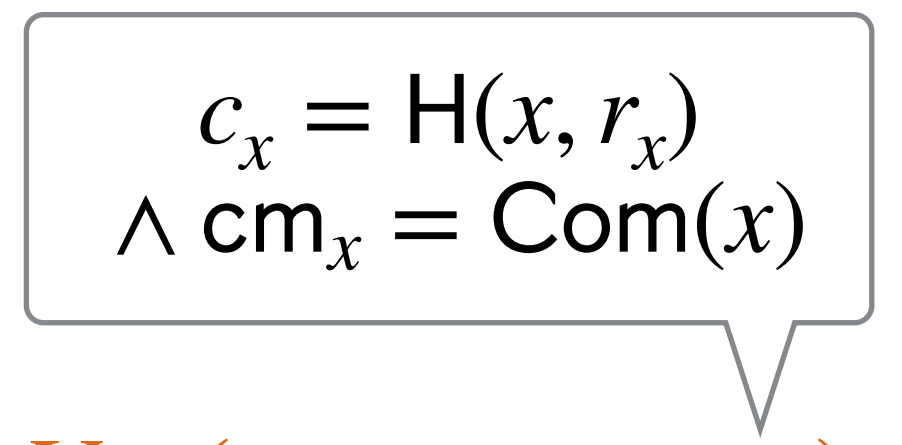
$\Pi . \mathcal{P}(f, x, y)$

$\Pi . \mathcal{V}(f, \text{cm}_x, y)$

VE-based IP



Prove that input fingerprint is correct w.r.t. cm_x



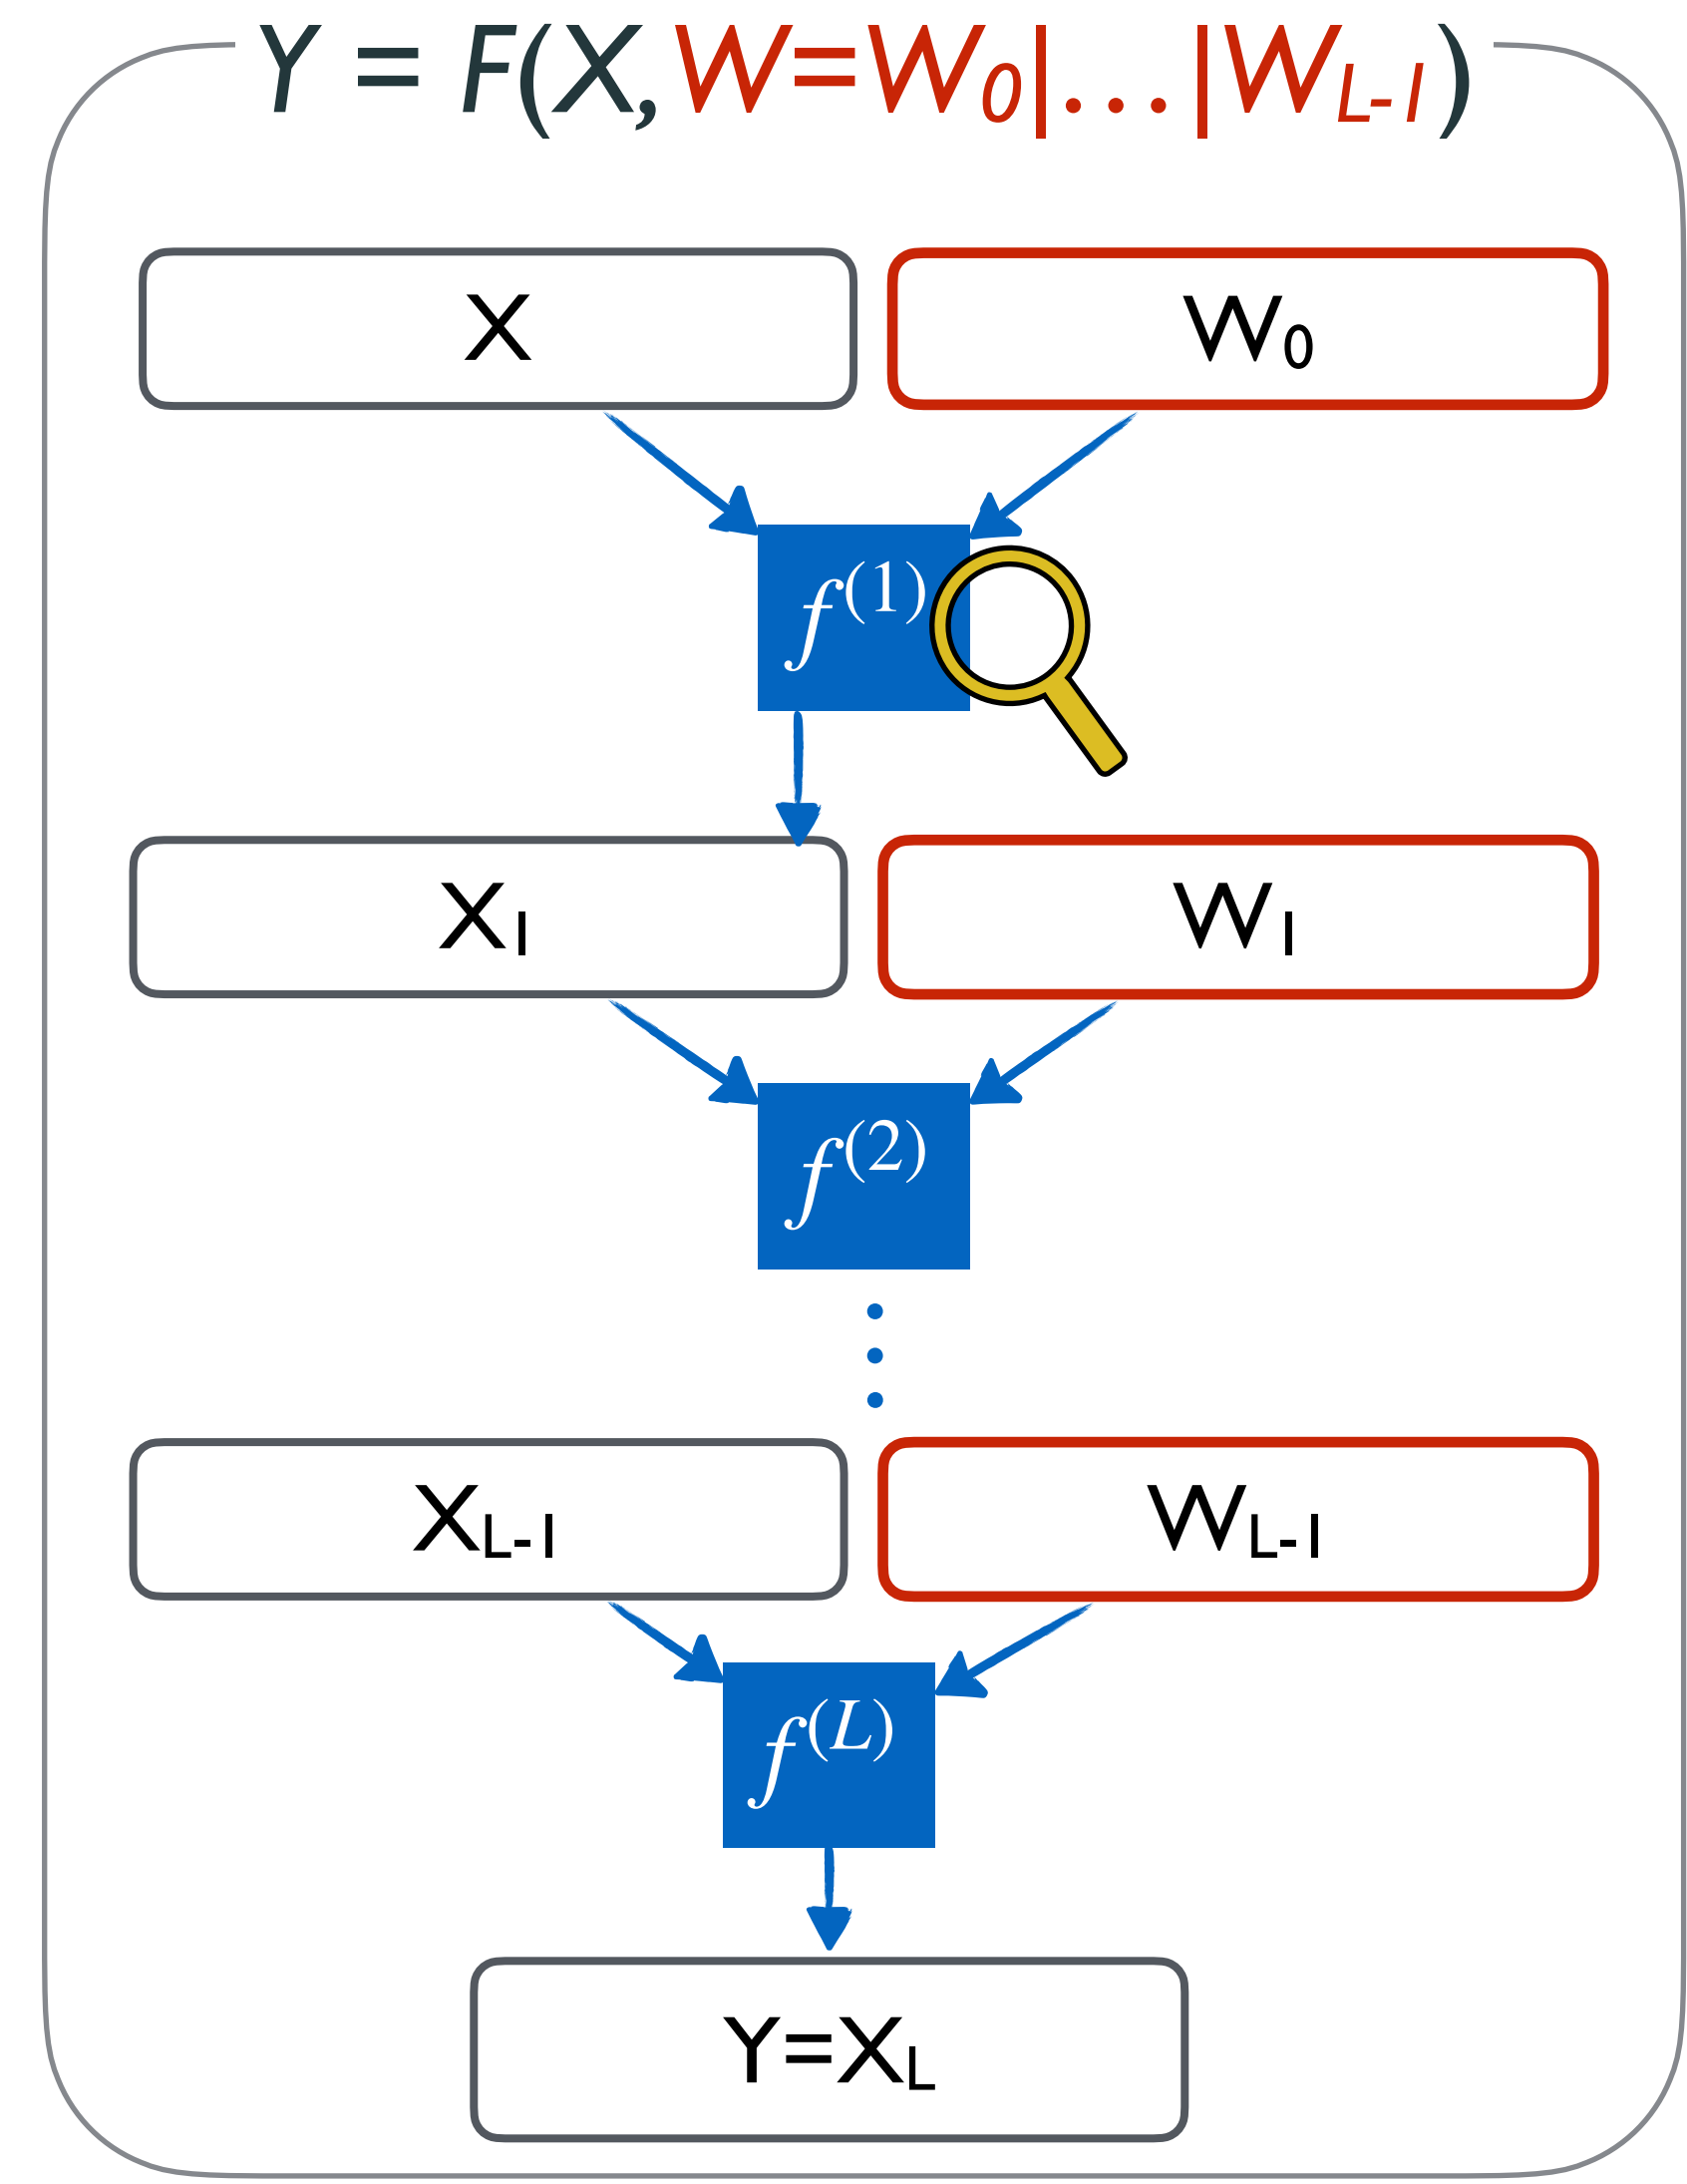
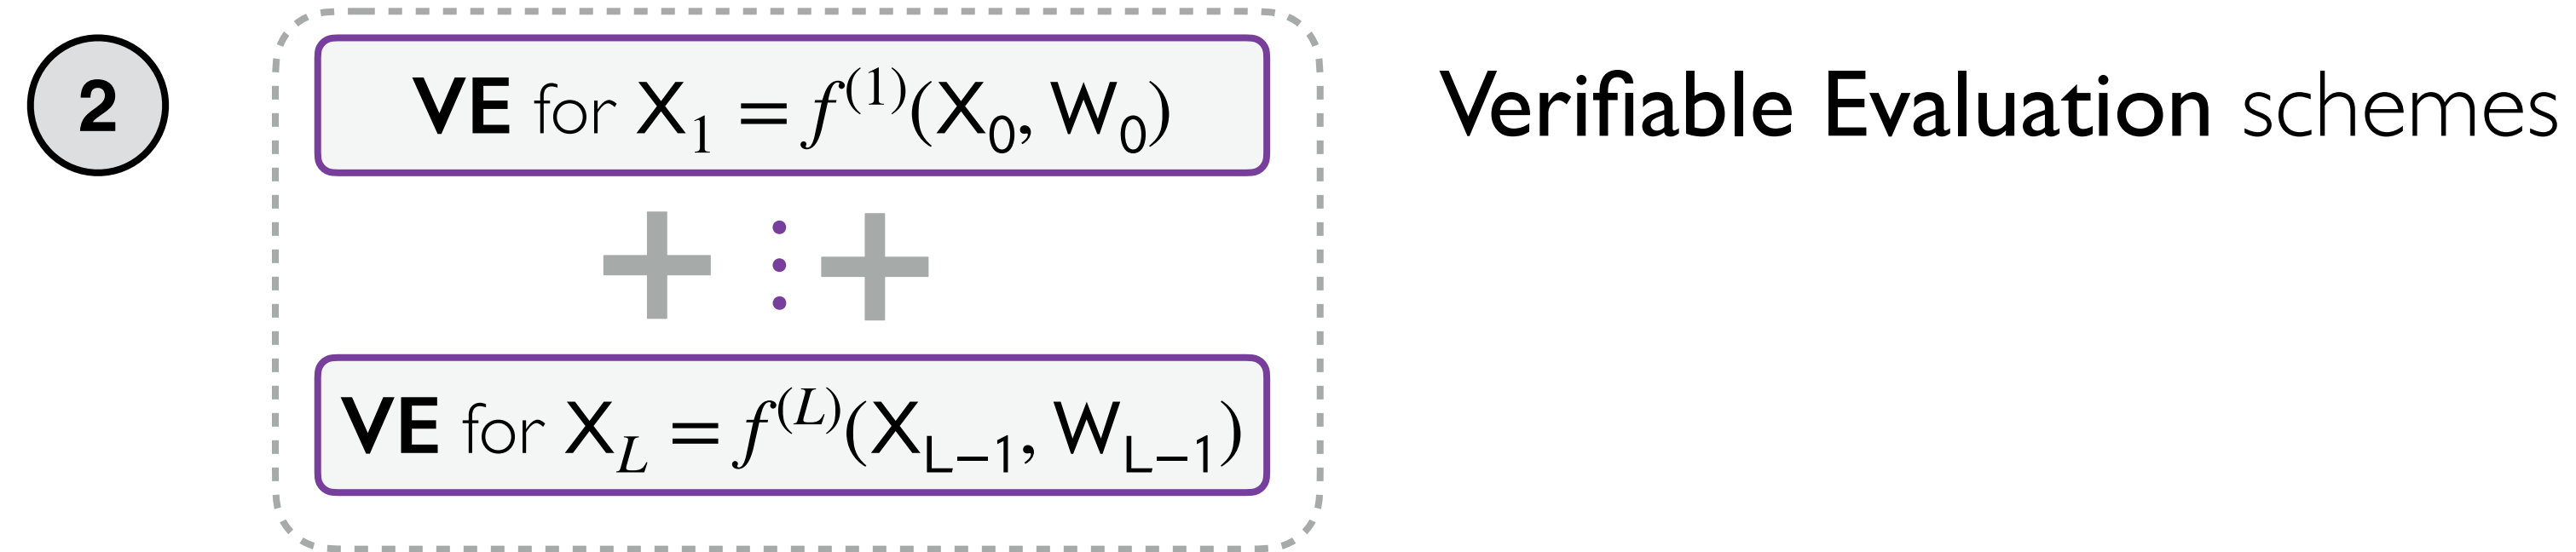
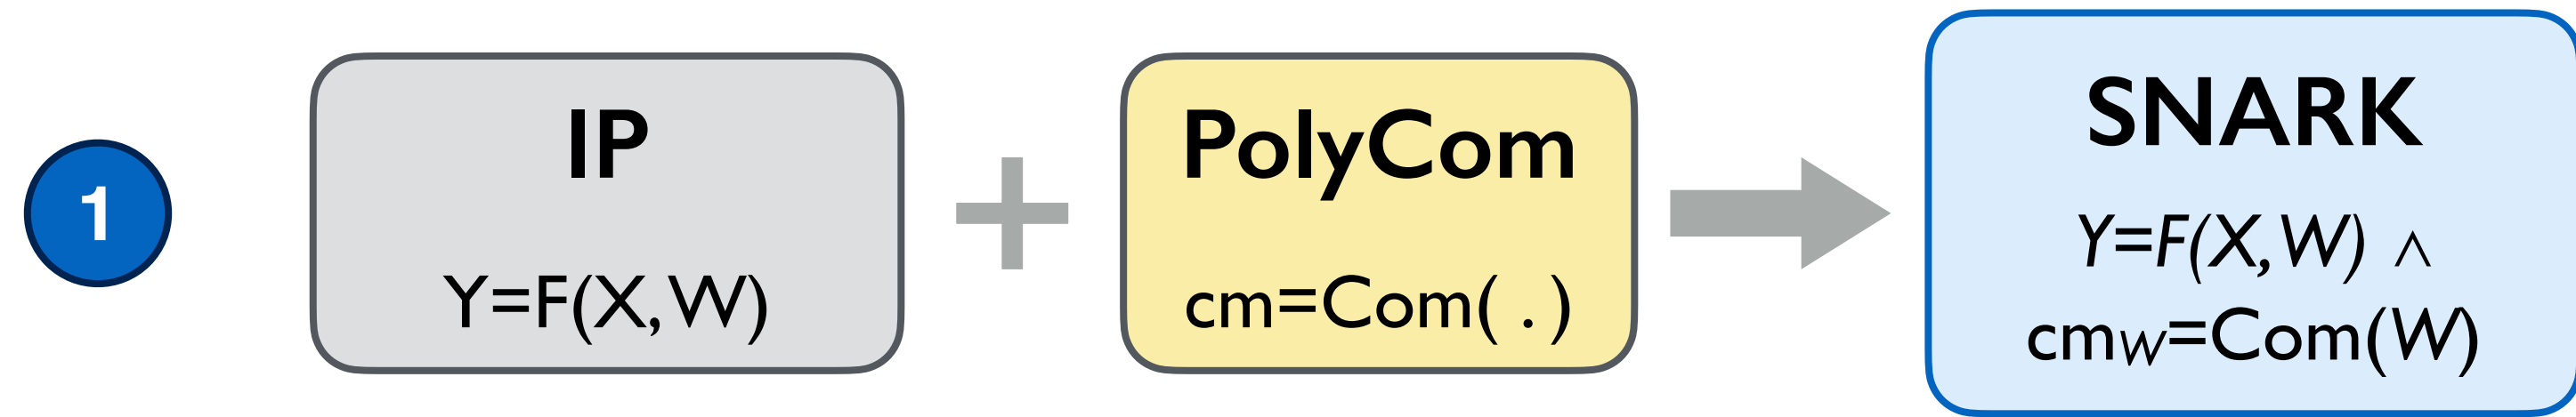
$\pi_x = \Pi_H . \mathcal{P}((\text{cm}_x, c_x, r_x), x, \pi_x)$

Return $b \wedge \Pi_H . \text{Ver}(\text{cm}_x, c_x, r_x, \pi_x)$

When $H(x, r_x) = \tilde{x}(r_x)$, Π_H can be instantiated with a multilinear polynomial commitment

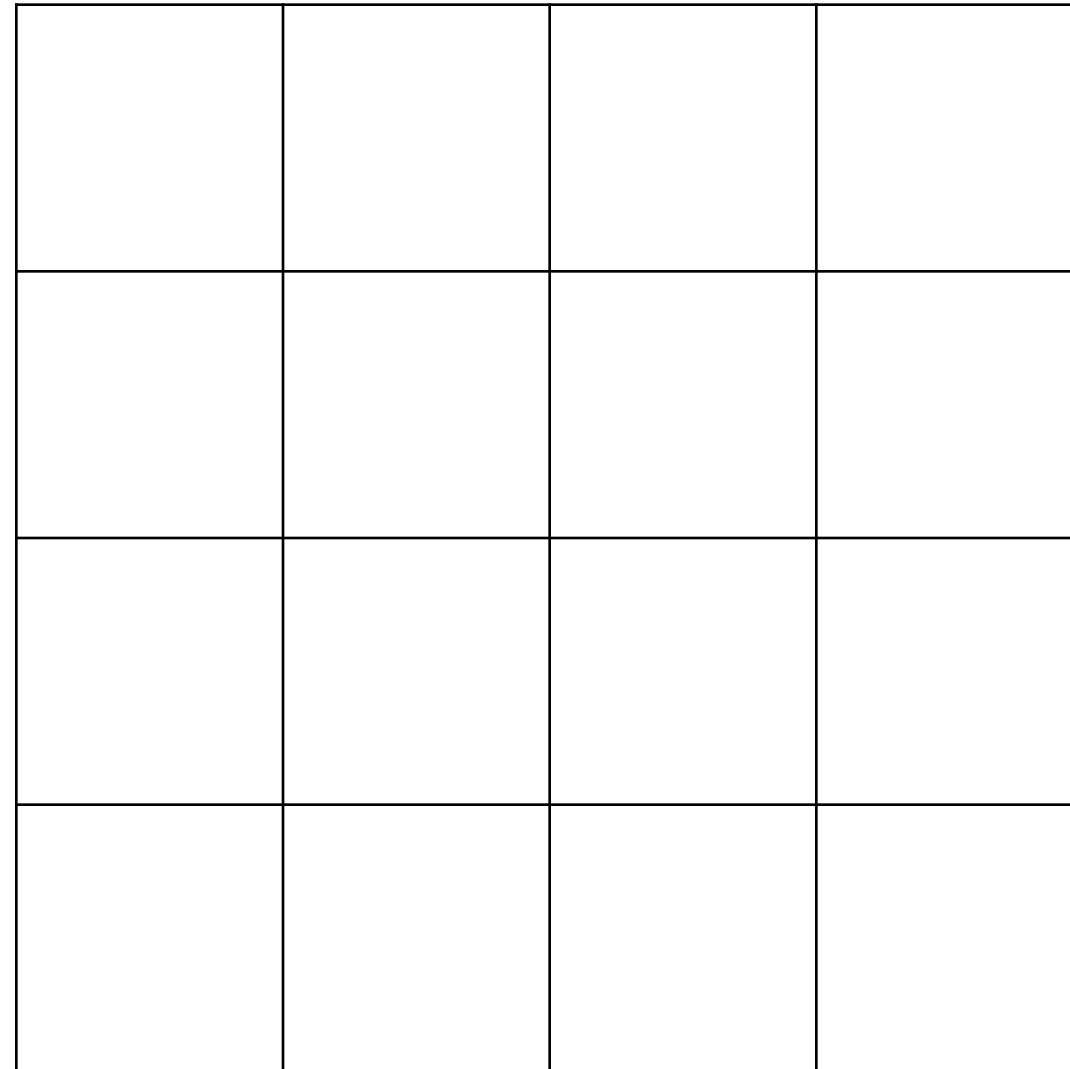
To get ZK: hiding of Com + ZK of Π_H + “ZK of the IP” [Libra] (or “committed IP” [zk-vSQL, Hyrax])

Modular approach for CNNs

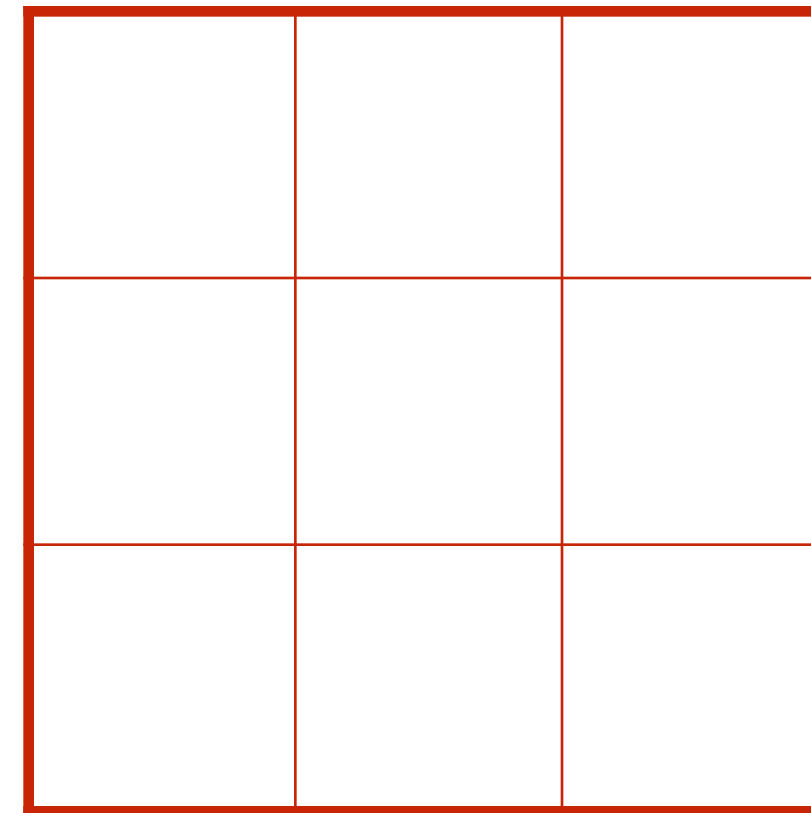


Convolution

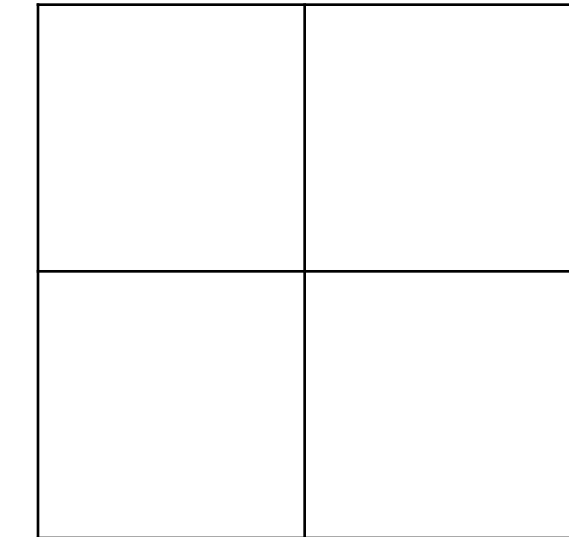
Input X
($n \times n$)



Kernel W
($m \times m$)



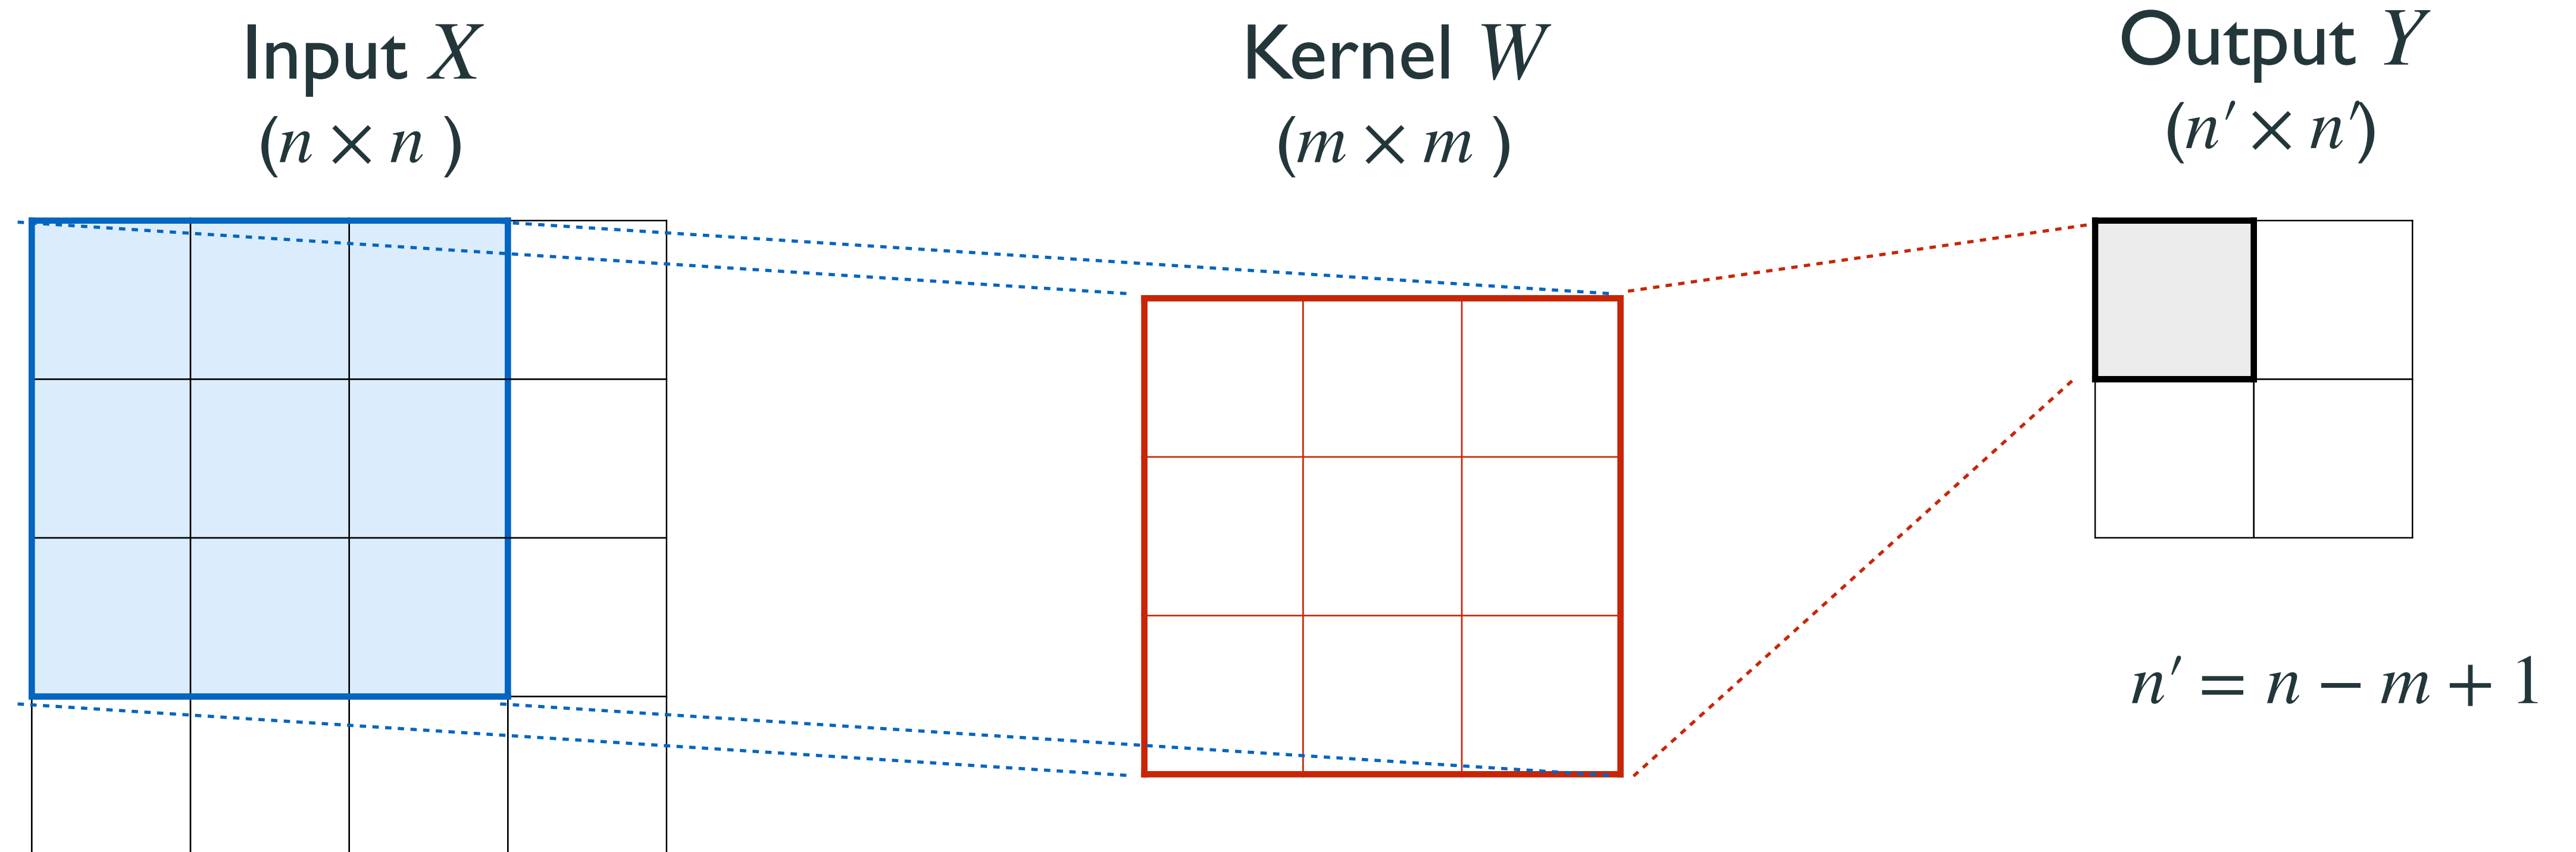
Output Y
($n' \times n'$)



$$n' = n - m + 1$$

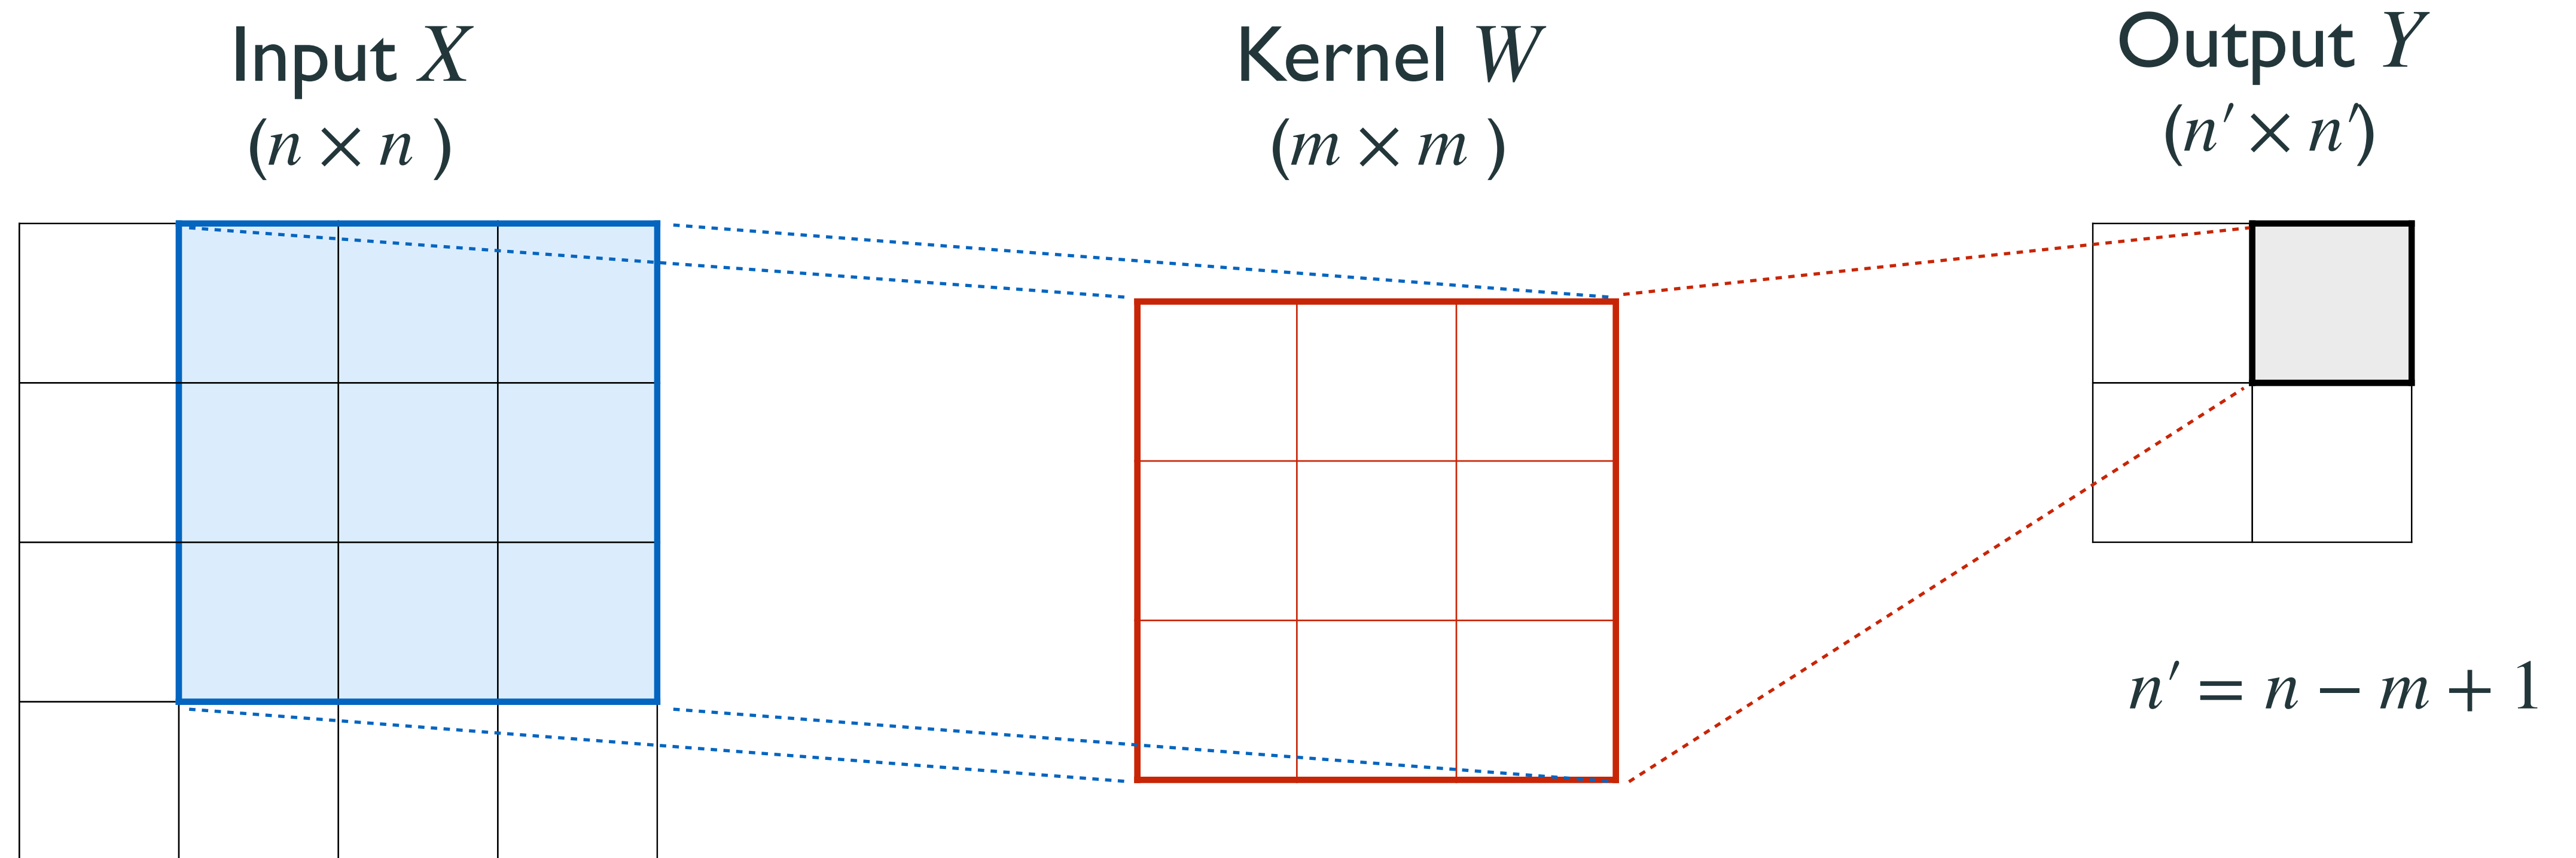
$$\text{Output } Y[u, v] = \sum_{i,j=0}^{m-1} X[u + i, v + j] \cdot W[i, j]$$

Convolution



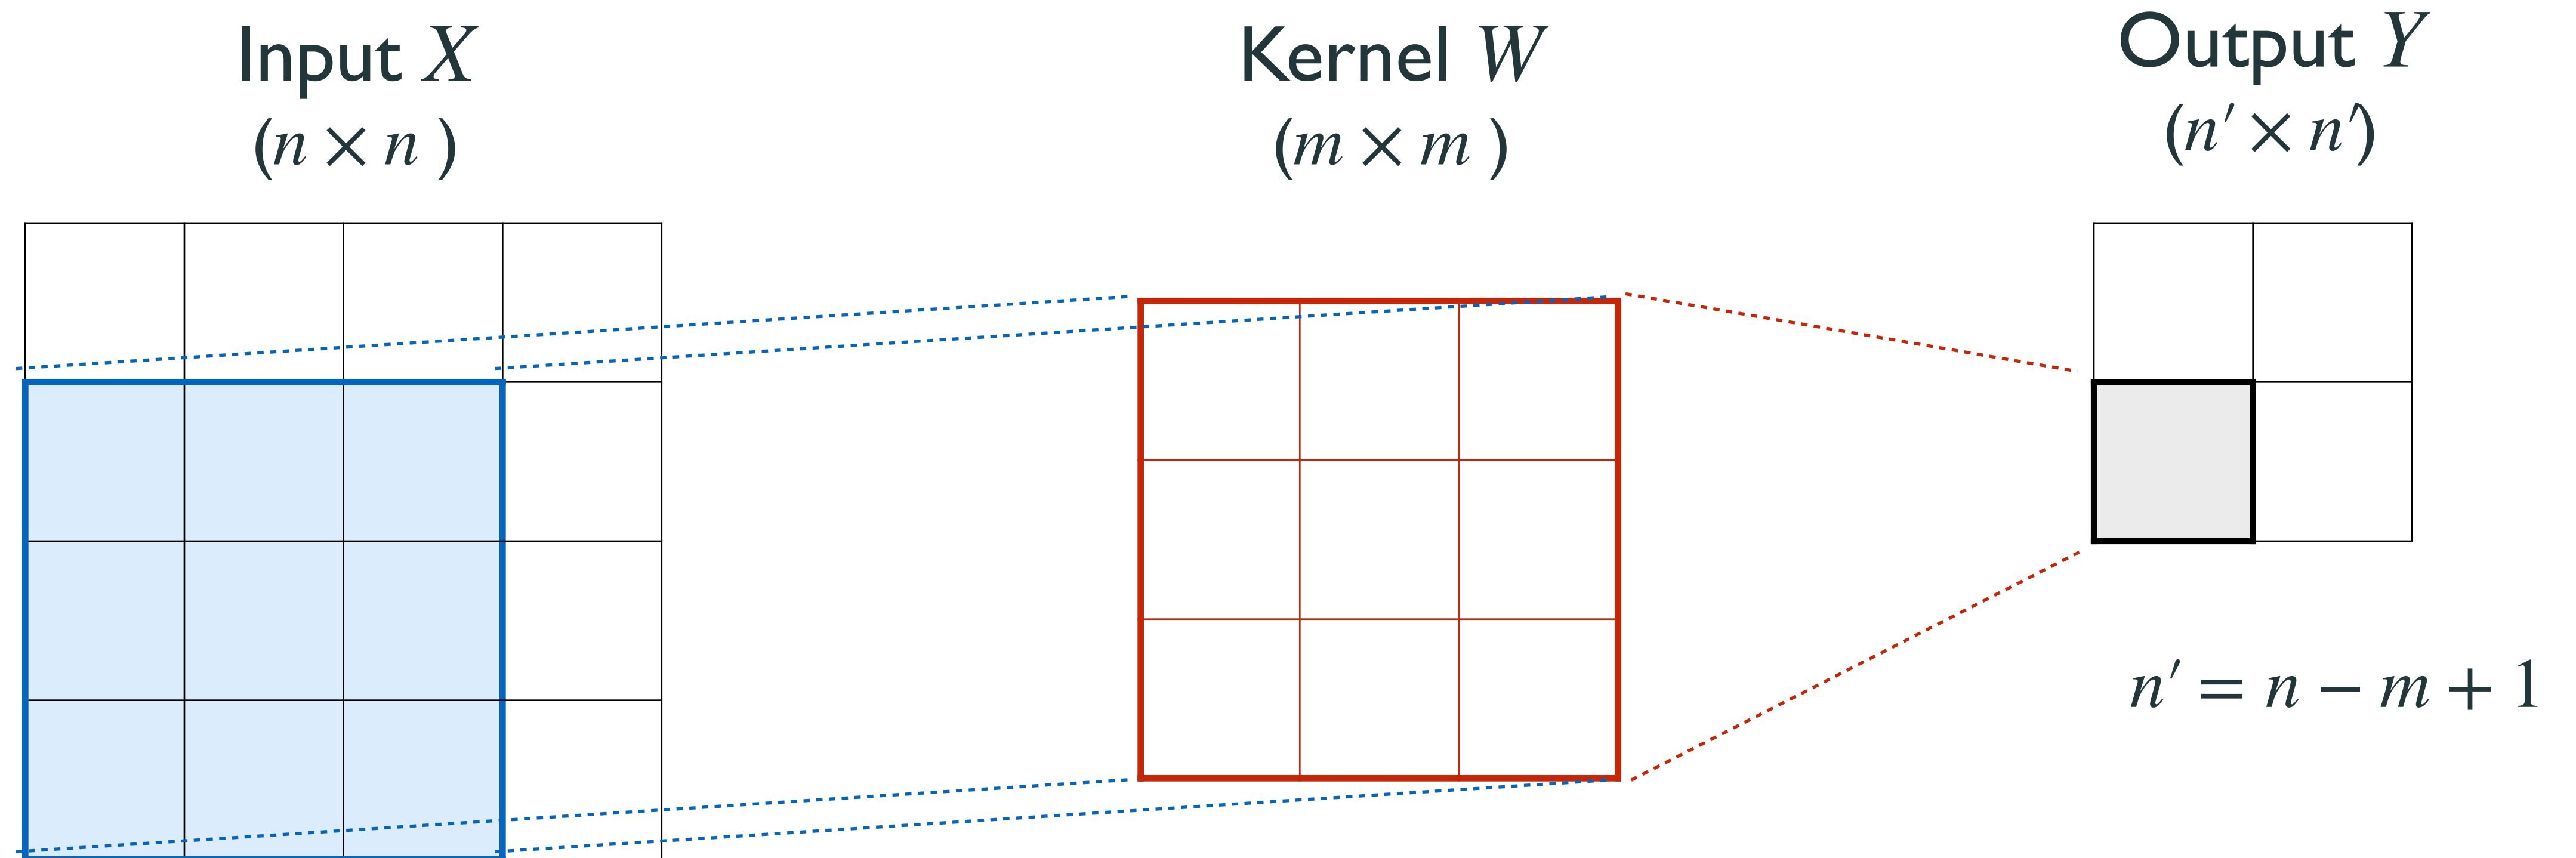
$$\text{Output } Y[u, v] = \sum_{i,j=0}^{m-1} X[u + i, v + j] \cdot W[i, j]$$

Convolution



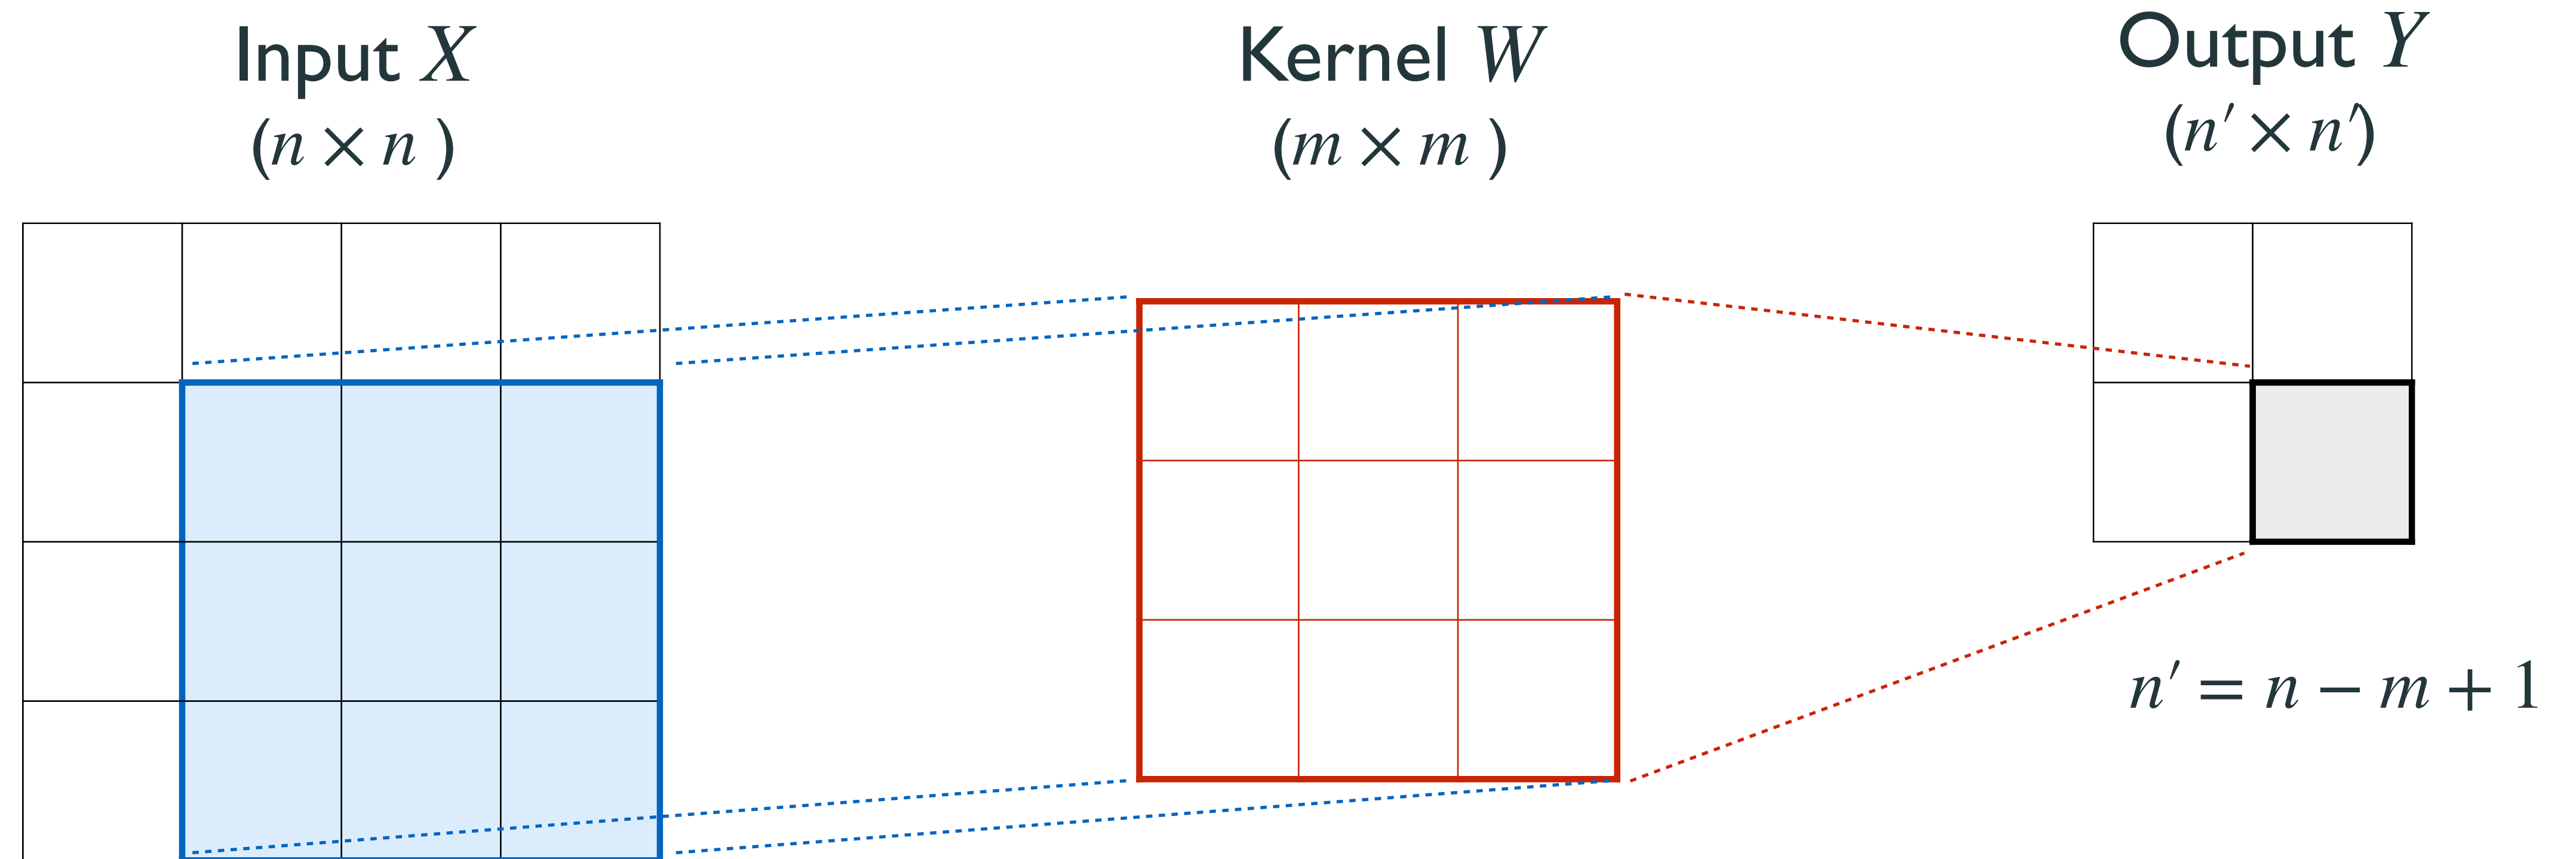
$$\text{Output } Y[u, v] = \sum_{i,j=0}^{m-1} X[u + i, v + j] \cdot W[i, j]$$

Convolution



$$\text{Output } Y[u, v] = \sum_{i,j=0}^{m-1} X[u + i, v + j] \cdot W[i, j]$$

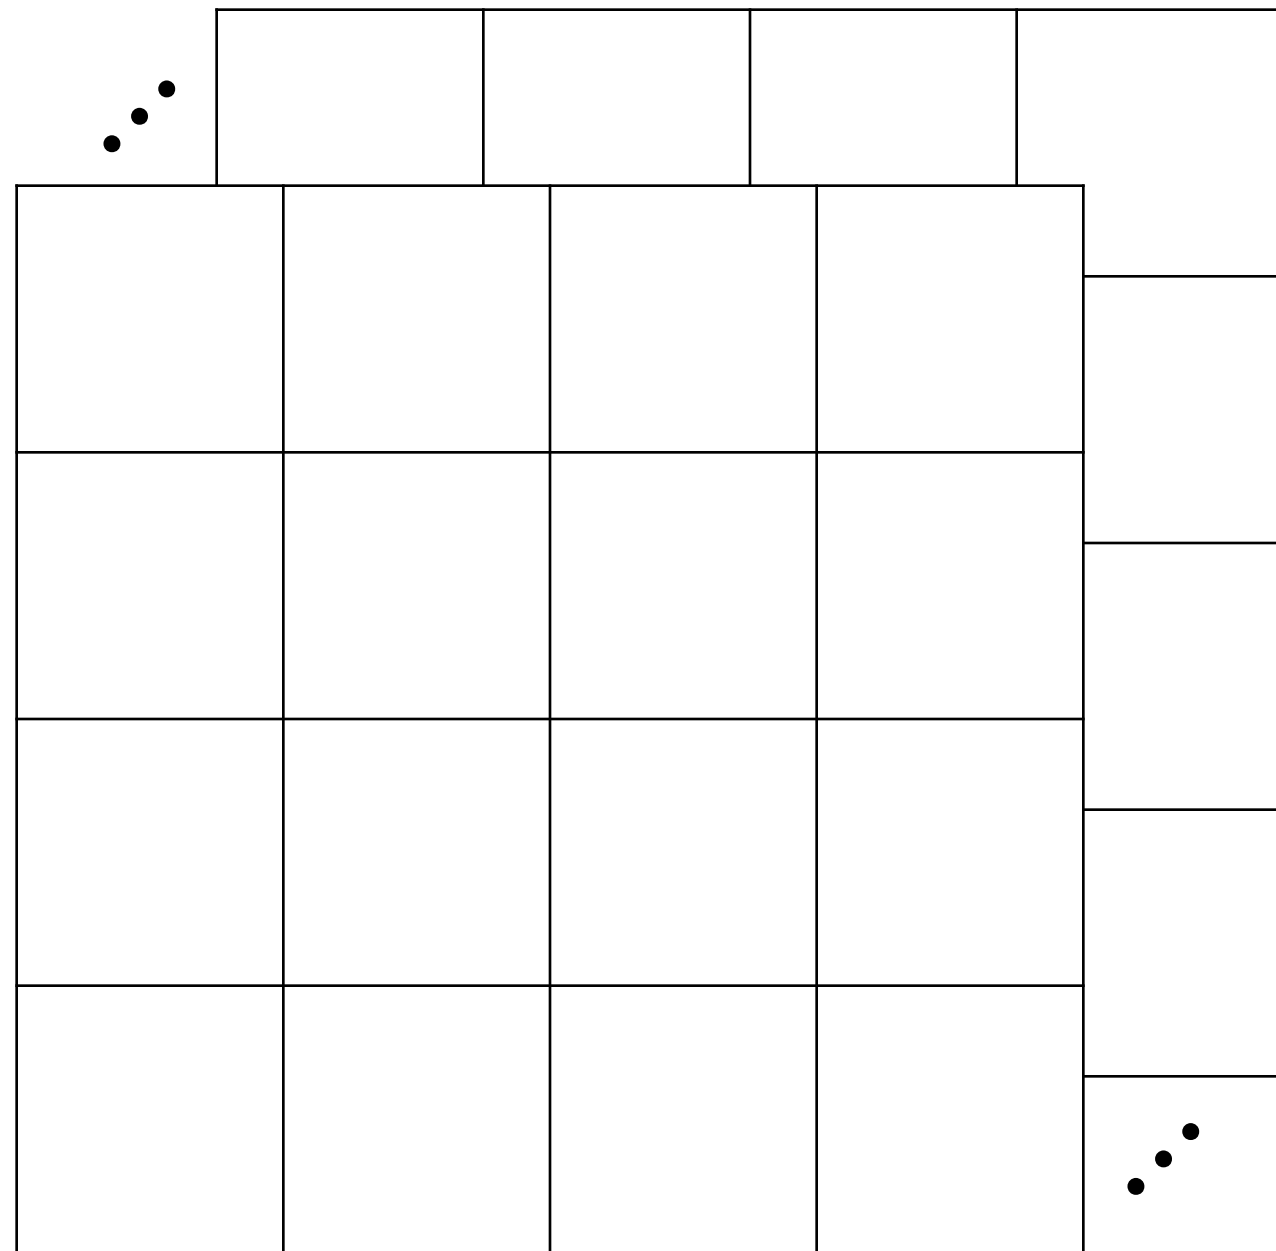
Convolution



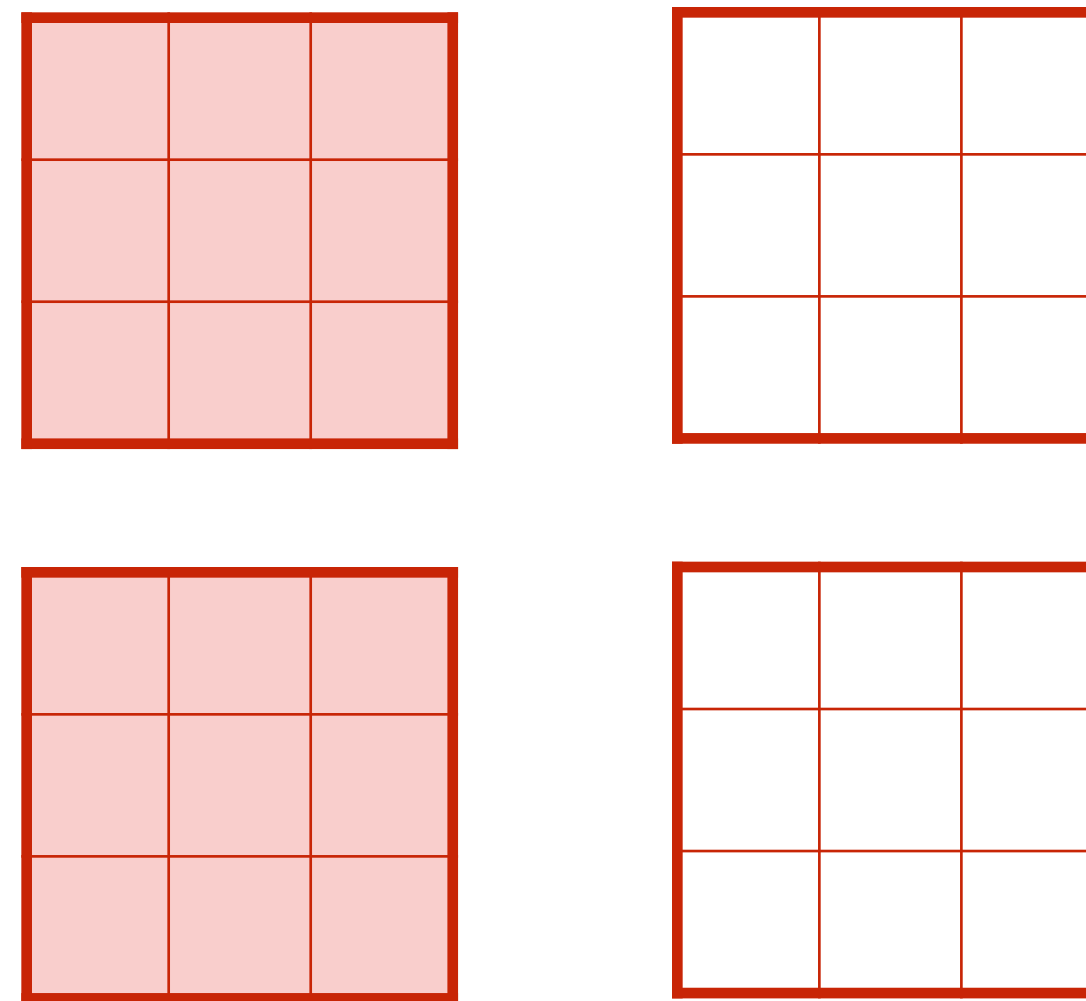
$$\text{Output } Y[u, v] = \sum_{i,j=0}^{m-1} X[u + i, v + j] \cdot W[i, j]$$

Multichannel Convolution of CNNs

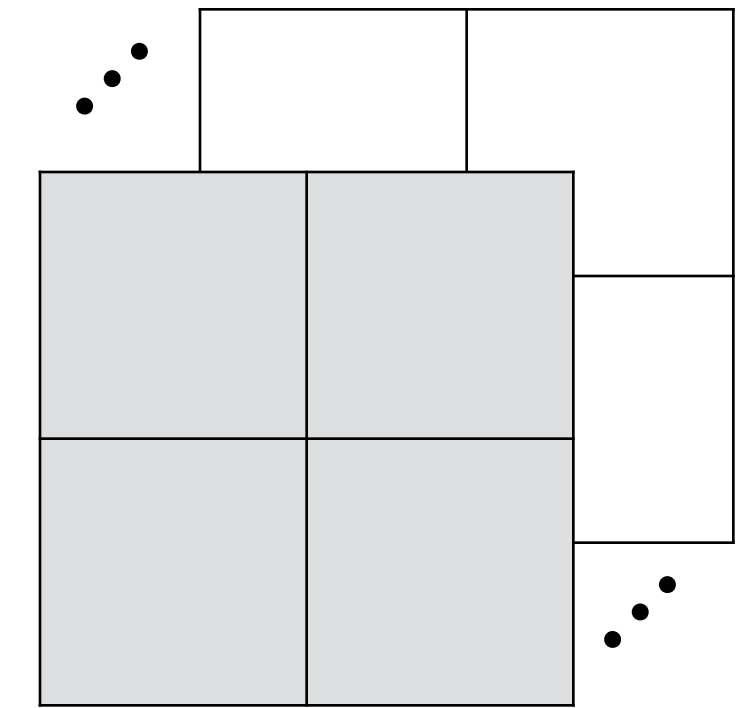
Input $\{X_{\sigma}^{(k)}\}_{\sigma=0}^{c^{(k)}-1}$
 $(n^{(k)} \times n^{(k)})$



Kernel $\{W_{\sigma,\tau}^{(k)}\}_{\sigma,\tau=0}^{c^{(k)}-1, c^{(k+1)}-1}$
 $(m^{(k)} \times m^{(k)})$



Output $\{X_{\tau}^{(k+1)}\}_{\tau=0}^{c^{(k+1)}-1}$
 $(n^{(k+1)} \times n^{(k+1)})$



$$\text{Output } X_{\tau}^{(k+1)}[u, v] = \sum_{\sigma=0}^{c^{(k)}-1} \sum_{i,j=0}^{m^{(k)}-1} X_{\sigma}^{(k)}[u+i, v+j] \cdot W_{\sigma,\tau}^{(k)}[i, j]$$

Proving convolution

$$Y_\tau[u, v] = \sum_{\sigma=0}^{c-1} \sum_{i,j=0}^{m-1} X_\sigma[u+i, v+j] \cdot W_{\sigma,\tau}[i, j] \quad \text{expensive as a circuit, } \mathcal{O}(cd |Y_\tau| \cdot |W_{\sigma,\tau}|)$$

Proving convolution

$$Y_\tau[u, v] = \sum_{\sigma=0}^{c-1} \sum_{i,j=0}^{m-1} X_\sigma[u+i, v+j] \cdot W_{\sigma,\tau}[i, j] \quad \text{expensive as a circuit, } \mathcal{O}(cd |Y_\tau| \cdot |W_{\sigma,\tau}|)$$

Special-purpose techniques: convolution \rightarrow structured matrix multiplication

Proving convolution

$$Y_\tau[u, v] = \sum_{\sigma=0}^{c-1} \sum_{i,j=0}^{m-1} X_\sigma[u+i, v+j] \cdot W_{\sigma,\tau}[i, j] \quad \text{expensive as a circuit, } \mathcal{O}(cd |Y_\tau| \cdot |W_{\sigma,\tau}|)$$

Special-purpose techniques: convolution \rightarrow structured matrix multiplication

zkCNN [LXZ21]

(2dim)Convolution \rightarrow 1 dim-convolution \rightarrow poly mult \rightarrow FFT \rightarrow matrix multiplication

Proving convolution

$$Y_\tau[u, v] = \sum_{\sigma=0}^{c-1} \sum_{i,j=0}^{m-1} X_\sigma[u+i, v+j] \cdot W_{\sigma,\tau}[i, j] \quad \text{expensive as a circuit, } \mathcal{O}(cd |Y_\tau| \cdot |W_{\sigma,\tau}|)$$

Special-purpose techniques: convolution \rightarrow structured matrix multiplication

zkCNN [LXZ21]

(2dim)Convolution \rightarrow 1 dim-convolution \rightarrow poly mult \rightarrow FFT \rightarrow matrix multiplication

[BFGRS23]

Convolution \rightarrow matrix multiplication with reshaped X, W

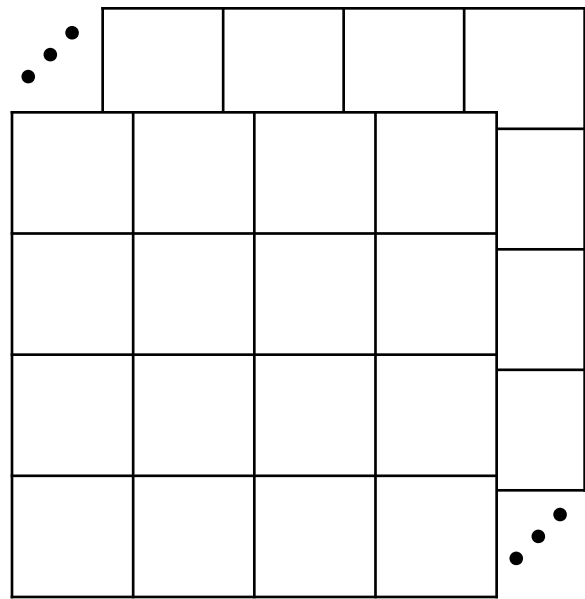
Reshaping convolution

Let $X = \begin{bmatrix} x_0 & x_1 & x_2 \\ x_3 & x_4 & x_5 \\ x_6 & x_7 & x_8 \end{bmatrix}$ and $W = \begin{bmatrix} w_0 & w_1 \\ w_2 & w_3 \end{bmatrix}$. Compact convolution as

$$\text{vec}(Y) = \begin{pmatrix} w_0x_0 + w_1x_1 + w_2x_3 + w_3x_4 \\ w_0x_1 + w_1x_2 + w_2x_4 + w_3x_5 \\ w_0x_3 + w_1x_4 + w_2x_6 + w_3x_7 \\ w_0x_4 + w_1x_5 + w_2x_7 + w_3x_8 \end{pmatrix} = \begin{bmatrix} x_0 & x_1 & x_3 & x_4 \\ x_1 & x_2 & x_4 & x_5 \\ x_3 & x_4 & x_6 & x_7 \\ x_4 & x_5 & x_7 & x_8 \end{bmatrix} \begin{bmatrix} w_0 \\ w_1 \\ w_2 \\ w_3 \end{bmatrix} = \hat{X} \cdot \hat{W}$$

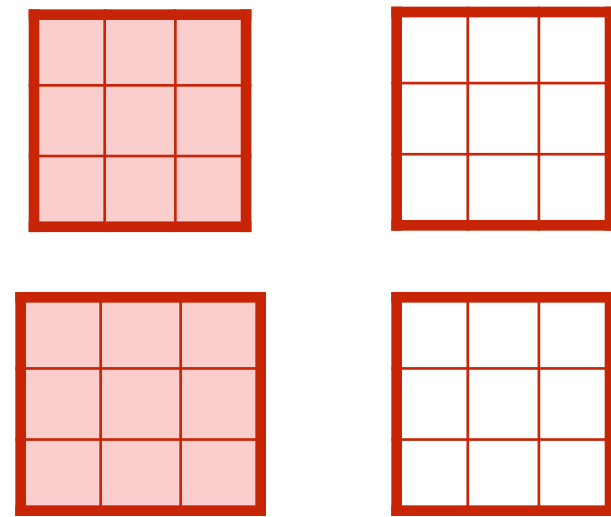
Reshaping multi-channel convolution

Inputs X_σ
($n \times n$)



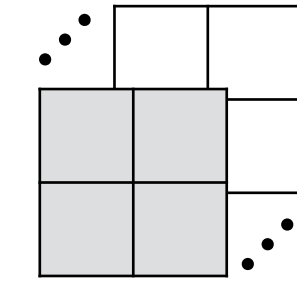
reshaped \hat{X}_σ
($(n')^2 \times m^2$)

Kernel $W_{\sigma,\tau}$
($m \times m$)



reshaped $\hat{W}_{\sigma,\tau}$
(m^2)

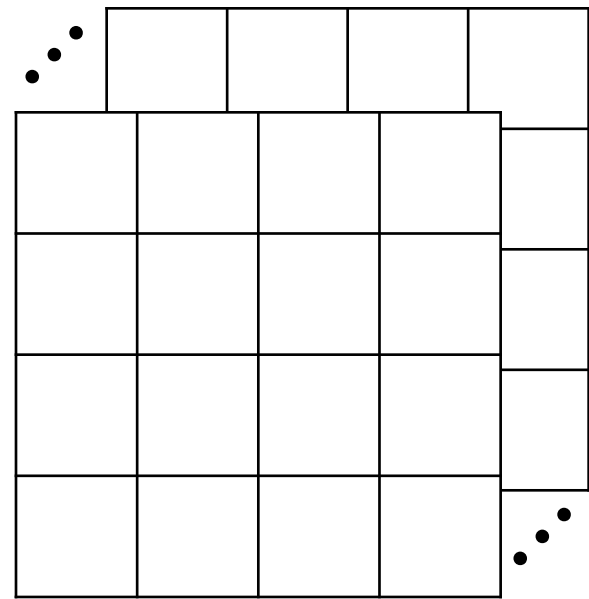
Outputs Y_τ
($n' \times n'$)



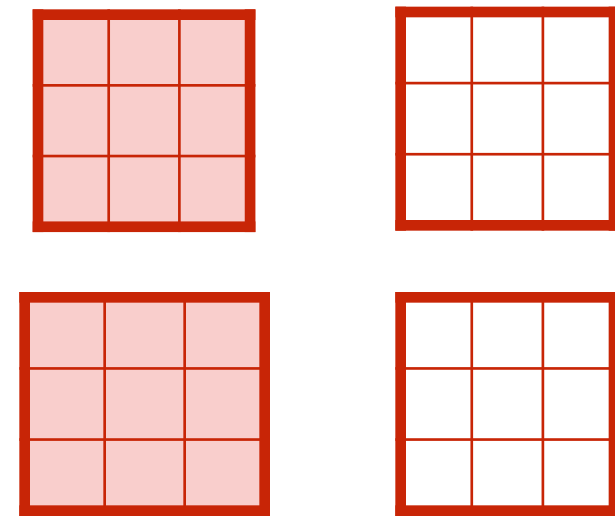
$\text{vec}(Y_\tau) = \hat{X}_\sigma \cdot \hat{W}_{\sigma,\tau}$
($(n')^2$)

Reshaping multi-channel convolution

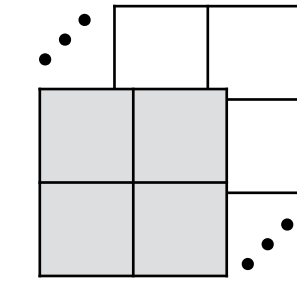
Inputs X_σ
($n \times n$)



Kernel $W_{\sigma,\tau}$
($m \times m$)



Outputs Y_τ
($n' \times n'$)



reshaped \hat{X}_σ
($((n')^2 \times m^2)$)

reshaped $\hat{W}_{\sigma,\tau}$
(m^2)

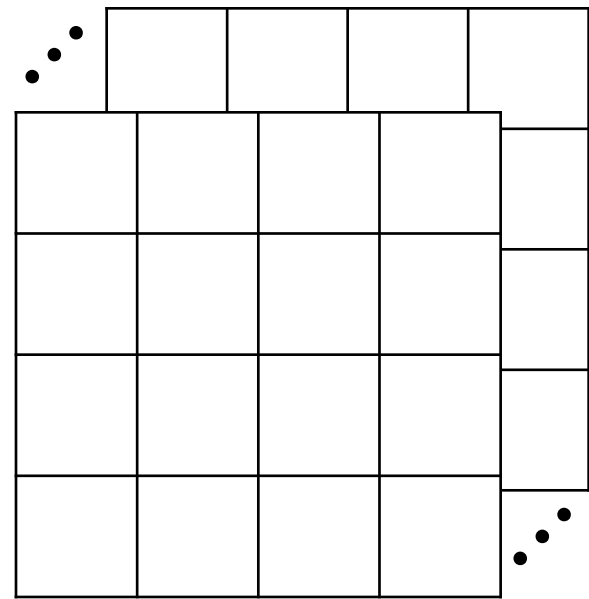
$\text{vec}(Y_\tau) = \hat{X}_\sigma \cdot \hat{W}_{\sigma,\tau}$
($((n')^2)$)

For multiple channels

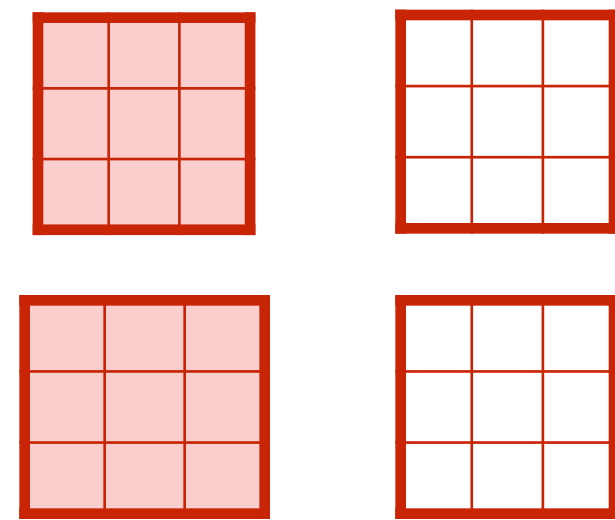
$$Y = [Y_1 | \cdots | Y_d] = \sum_{\sigma} \hat{X}_\sigma \cdot [\hat{W}_{\sigma,1} | \cdots | \hat{W}_{\sigma,d}] = \sum_{\sigma} \hat{X}_\sigma \cdot \hat{W}_\sigma \quad \text{sum of matrix multiplications}$$

Reshaping multi-channel convolution

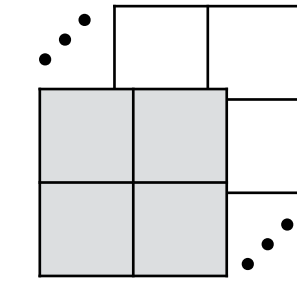
Inputs X_σ
($n \times n$)



Kernel $W_{\sigma,\tau}$
($m \times m$)



Outputs Y_τ
($n' \times n'$)



reshaped \hat{X}_σ
($((n')^2 \times m^2)$)

reshaped $\hat{W}_{\sigma,\tau}$
(m^2)

$\text{vec}(Y_\tau) = \hat{X}_\sigma \cdot \hat{W}_{\sigma,\tau}$
($((n')^2)$)

For multiple channels

$$Y = [Y_1 | \cdots | Y_d] = \sum_{\sigma} \hat{X}_\sigma \cdot [\hat{W}_{\sigma,1} | \cdots | \hat{W}_{\sigma,d}] = \sum_{\sigma} \hat{X}_\sigma \cdot \hat{W}_\sigma \quad \text{sum of matrix multiplications}$$

VE for convolution ← VE for (sum of) matrix multiplications

VE for matrix multiplication [Thaler13]

Prove $C = A \cdot B$ for $A, B, C \in \mathbb{F}^{n \times n}$. Using MLE: $\forall \vec{i}, \vec{k} \in \{0,1\}^{\log n} : \tilde{C}(\vec{i}, \vec{k}) = \sum_{\vec{j} \in \{0,1\}^{\log n}} \tilde{A}(\vec{i}, \vec{j}) \cdot \tilde{B}(\vec{j}, \vec{k})$

VE efficiency: communication & verification $O(\log n)$, prover $O(n^2)$, faster than computing $A \cdot B$ in $O(n^3)$

VE for matrix multiplication [Thaler13]

Prove $C = A \cdot B$ for $A, B, C \in \mathbb{F}^{n \times n}$. Using MLE: $\forall \vec{i}, \vec{k} \in \{0, 1\}^{\log n} : \tilde{C}(\vec{i}, \vec{k}) = \sum_{\vec{j} \in \{0, 1\}^{\log n}} \tilde{A}(\vec{i}, \vec{j}) \cdot \tilde{B}(\vec{j}, \vec{k})$

$$\mathcal{P}^{mm}(A, B, C)$$

$$\mathcal{V}^{mm}(A, B, C)$$

VE efficiency: communication & verification $O(\log n)$, prover $O(n^2)$, faster than computing $A \cdot B$ in $O(n^3)$

VE for matrix multiplication [Thaler13]

Prove $C = A \cdot B$ for $A, B, C \in \mathbb{F}^{n \times n}$. Using MLE: $\forall \vec{i}, \vec{k} \in \{0, 1\}^{\log n} : \tilde{C}(\vec{i}, \vec{k}) = \sum_{\vec{j} \in \{0, 1\}^{\log n}} \tilde{A}(\vec{i}, \vec{j}) \cdot \tilde{B}(\vec{j}, \vec{k})$

$$\mathcal{P}^{mm}(A, B, C)$$

$$c_C = H(C, \vec{r}_C) = \tilde{C}(\vec{r}_1, \vec{r}_2)$$

$$\vec{r}_C = (\vec{r}_1, \vec{r}_2)$$

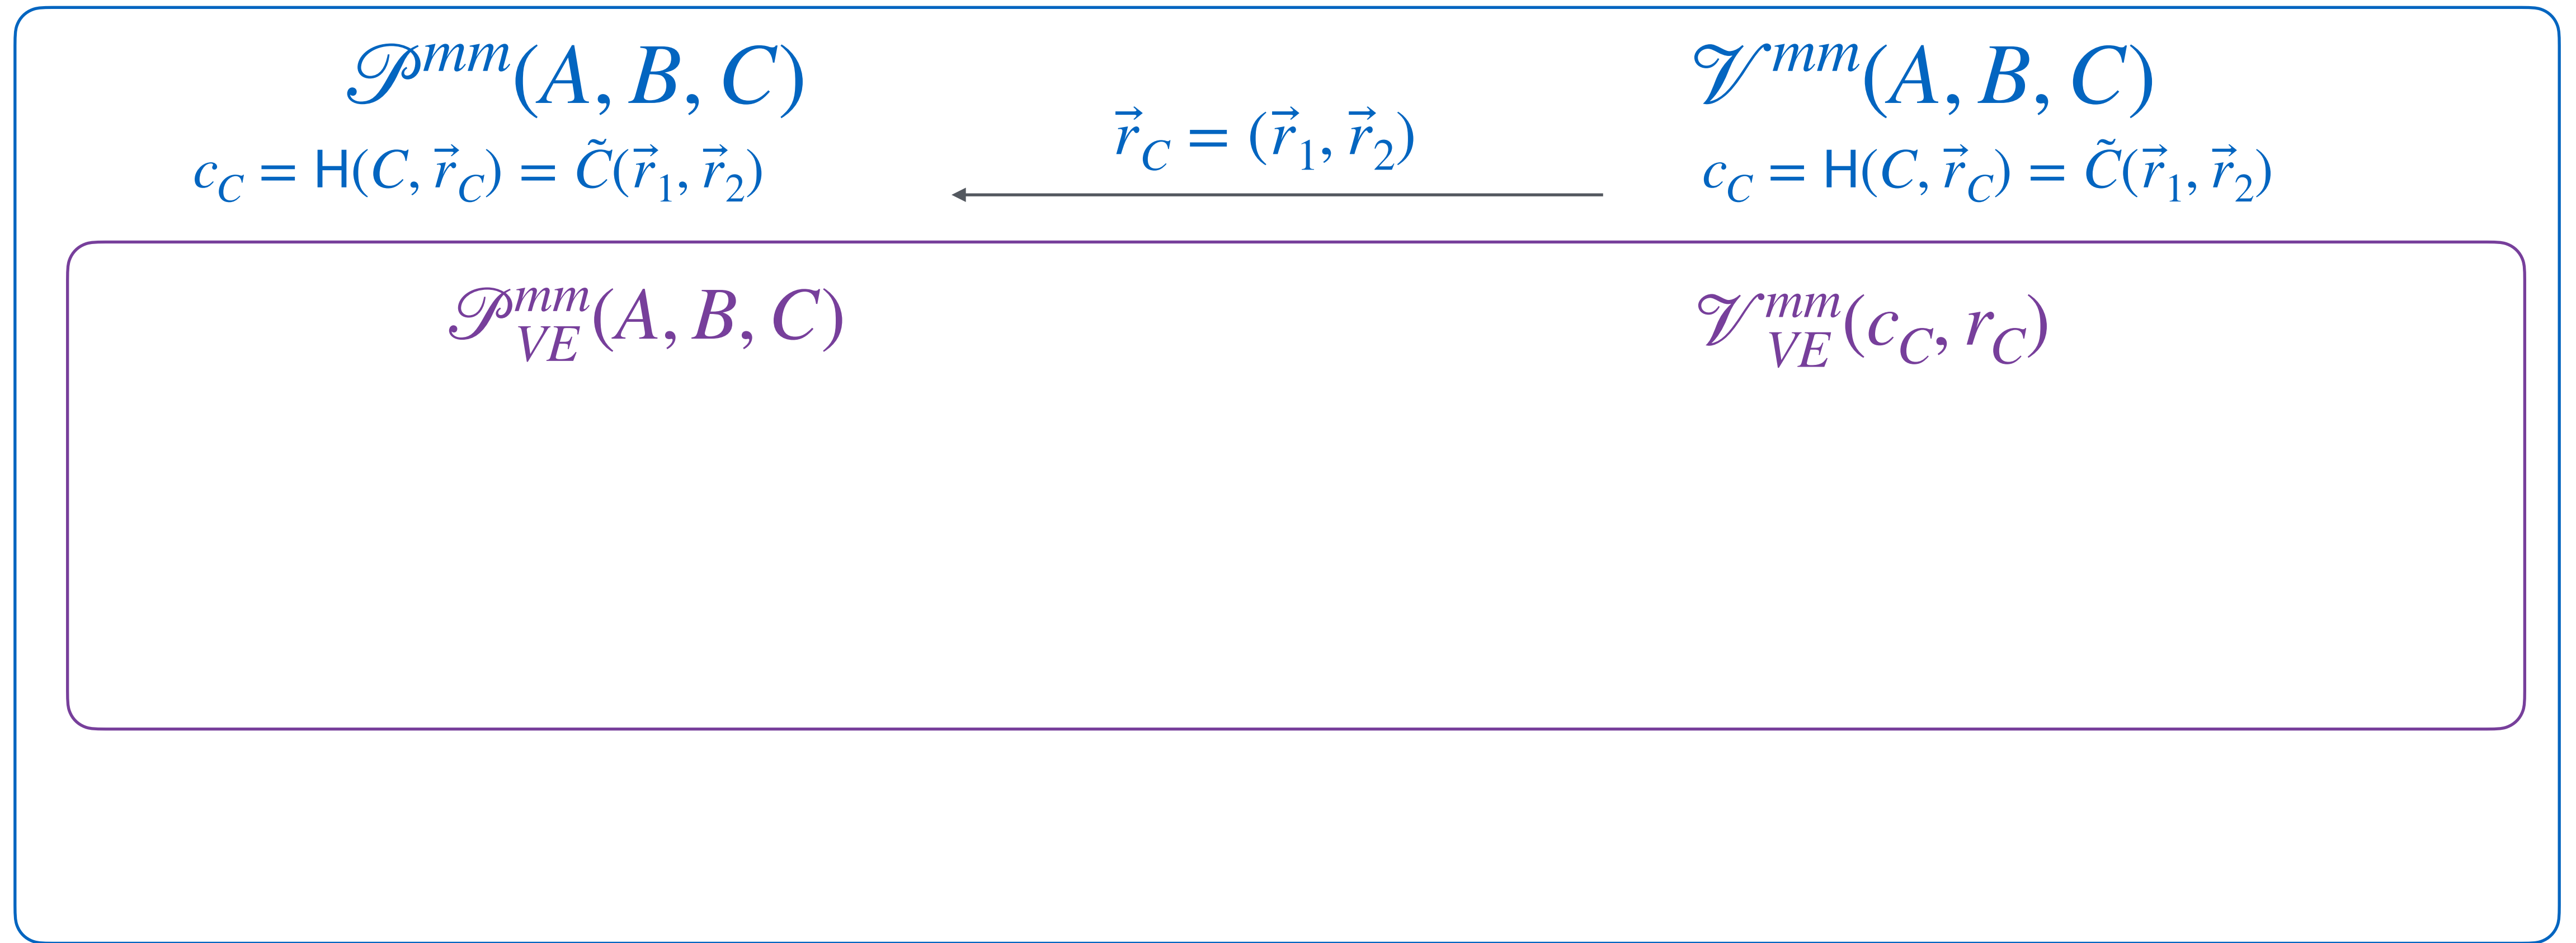
$$\mathcal{V}^{mm}(A, B, C)$$

$$c_C = H(C, \vec{r}_C) = \tilde{C}(\vec{r}_1, \vec{r}_2)$$

VE efficiency: communication & verification $O(\log n)$, prover $O(n^2)$, faster than computing $A \cdot B$ in $O(n^3)$

VE for matrix multiplication [Thaler13]

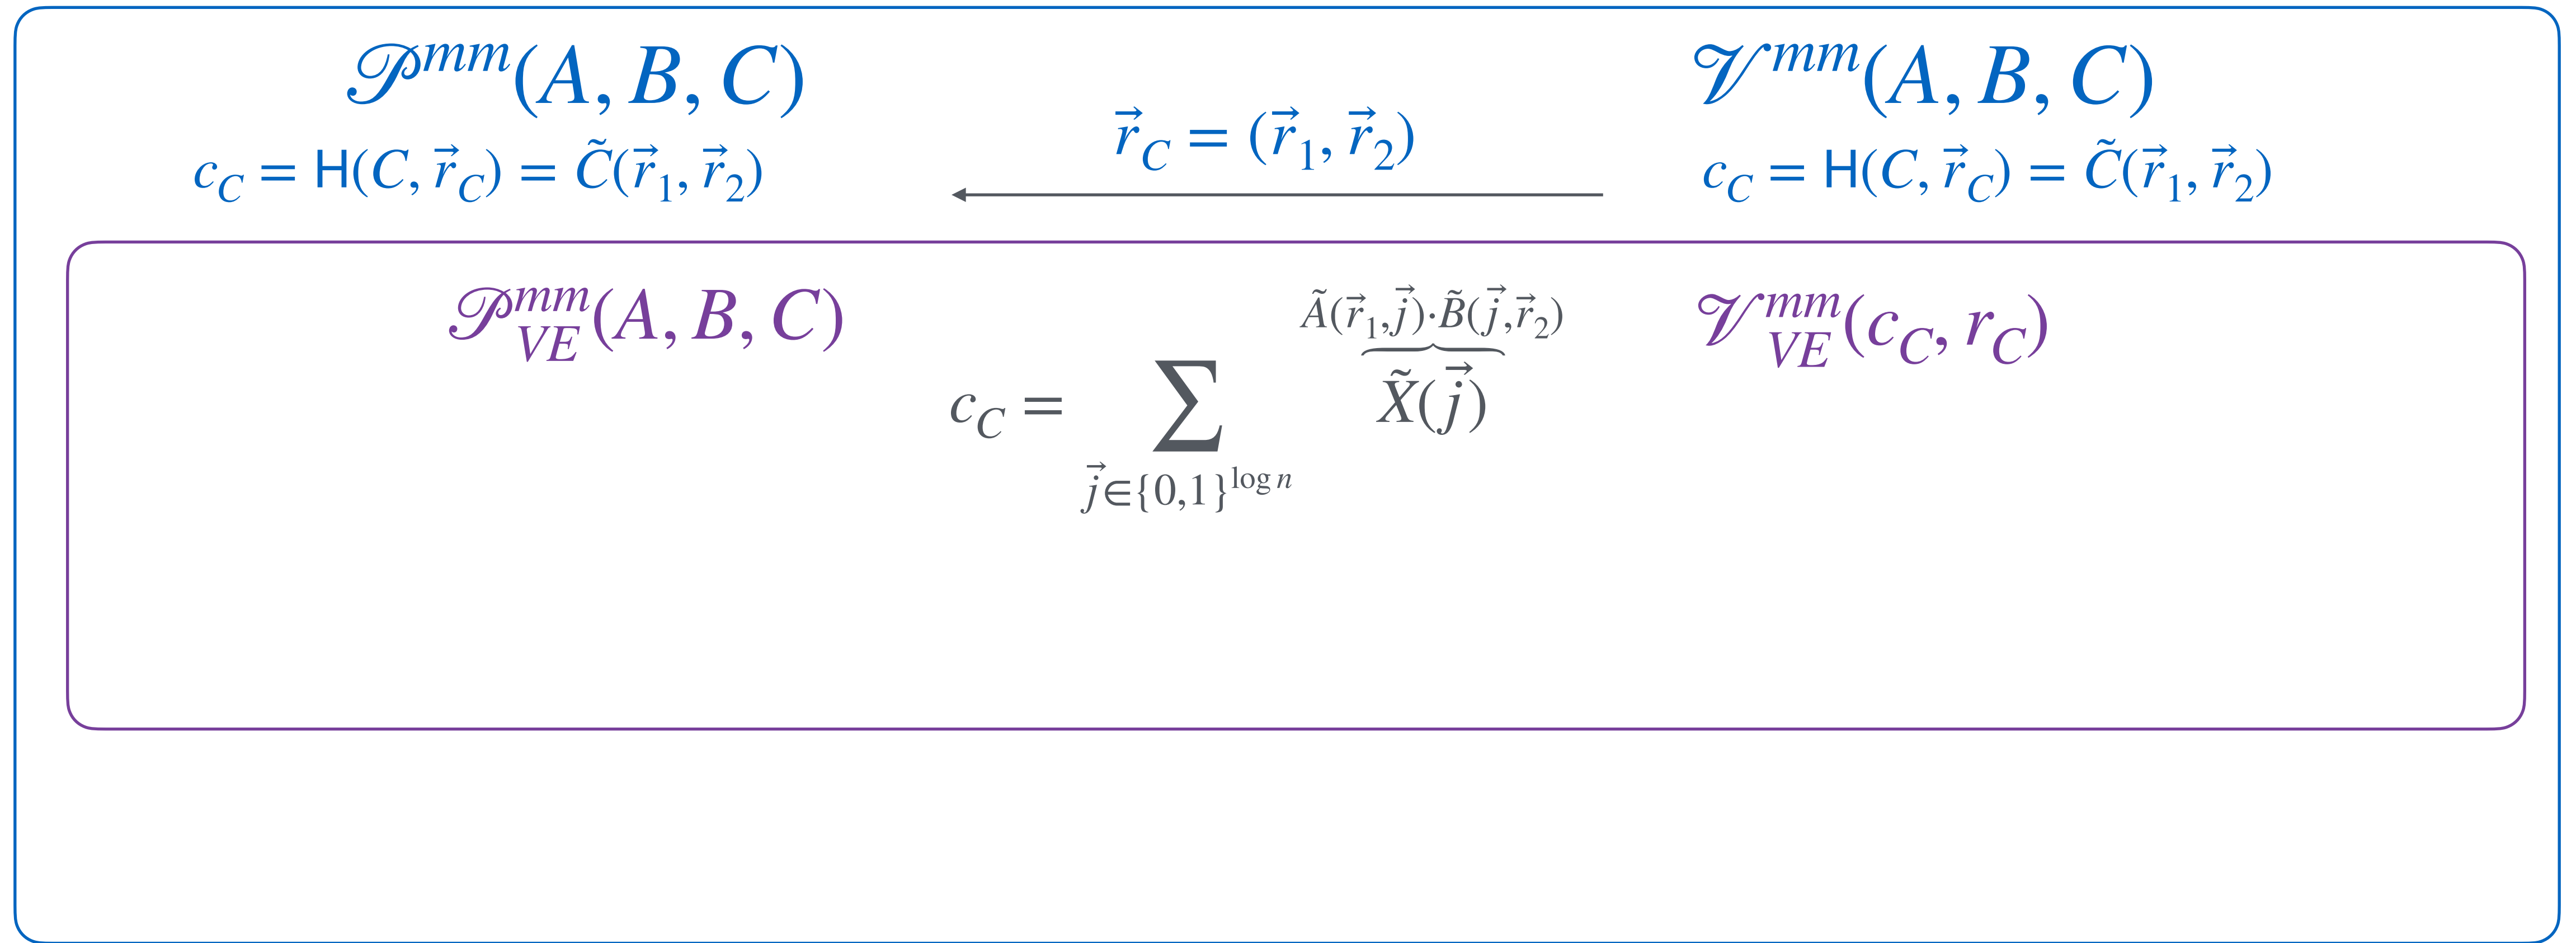
Prove $C = A \cdot B$ for $A, B, C \in \mathbb{F}^{n \times n}$. Using MLE: $\forall \vec{i}, \vec{k} \in \{0, 1\}^{\log n} : \tilde{C}(\vec{i}, \vec{k}) = \sum_{\vec{j} \in \{0, 1\}^{\log n}} \tilde{A}(\vec{i}, \vec{j}) \cdot \tilde{B}(\vec{j}, \vec{k})$



VE efficiency: communication & verification $O(\log n)$, prover $O(n^2)$, faster than computing $A \cdot B$ in $O(n^3)$

VE for matrix multiplication [Thaler13]

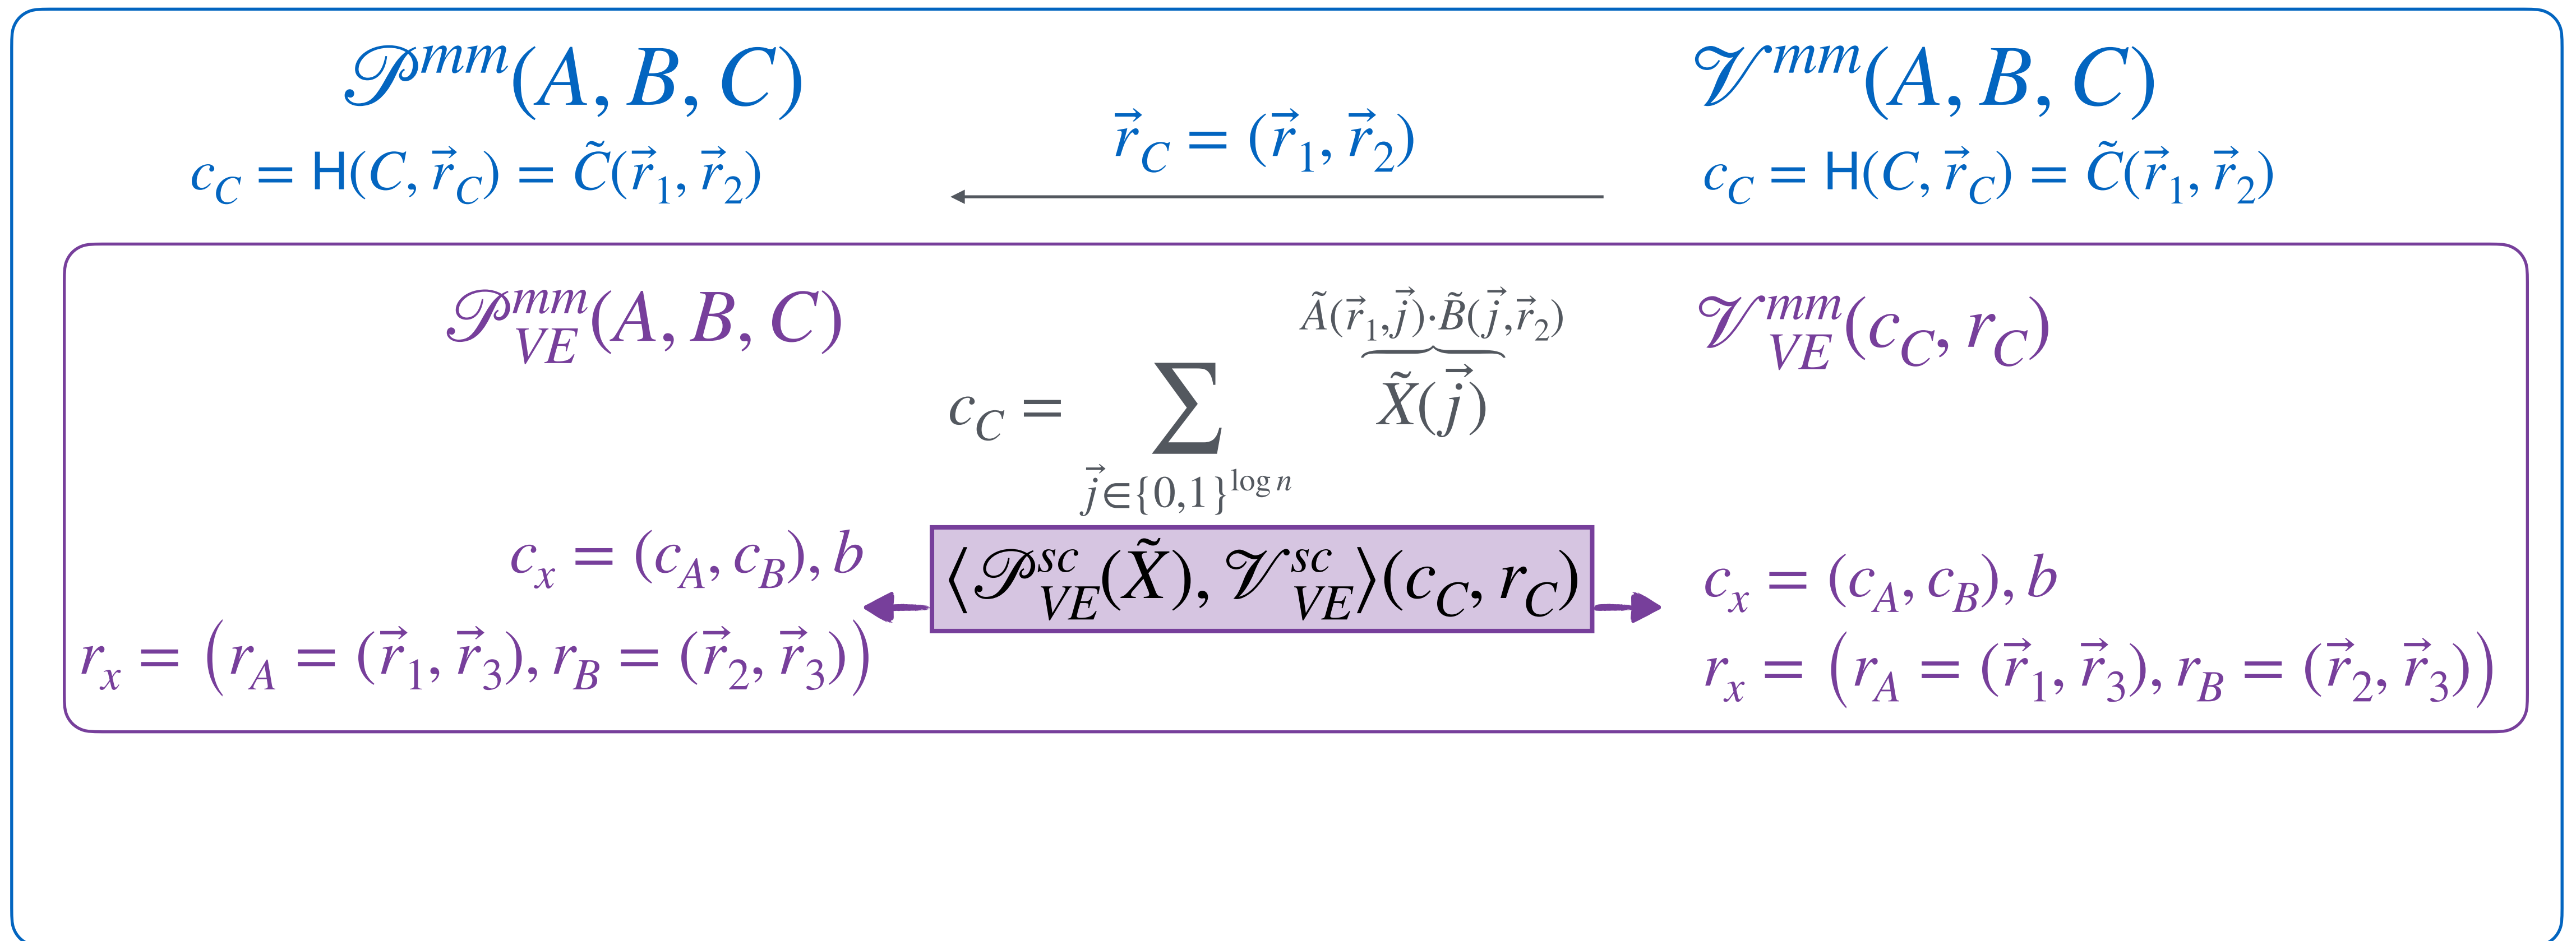
Prove $C = A \cdot B$ for $A, B, C \in \mathbb{F}^{n \times n}$. Using MLE: $\forall \vec{i}, \vec{k} \in \{0,1\}^{\log n} : \tilde{C}(\vec{i}, \vec{k}) = \sum_{\vec{j} \in \{0,1\}^{\log n}} \tilde{A}(\vec{i}, \vec{j}) \cdot \tilde{B}(\vec{j}, \vec{k})$



VE efficiency: communication & verification $O(\log n)$, prover $O(n^2)$, faster than computing $A \cdot B$ in $O(n^3)$

VE for matrix multiplication [Thaler13]

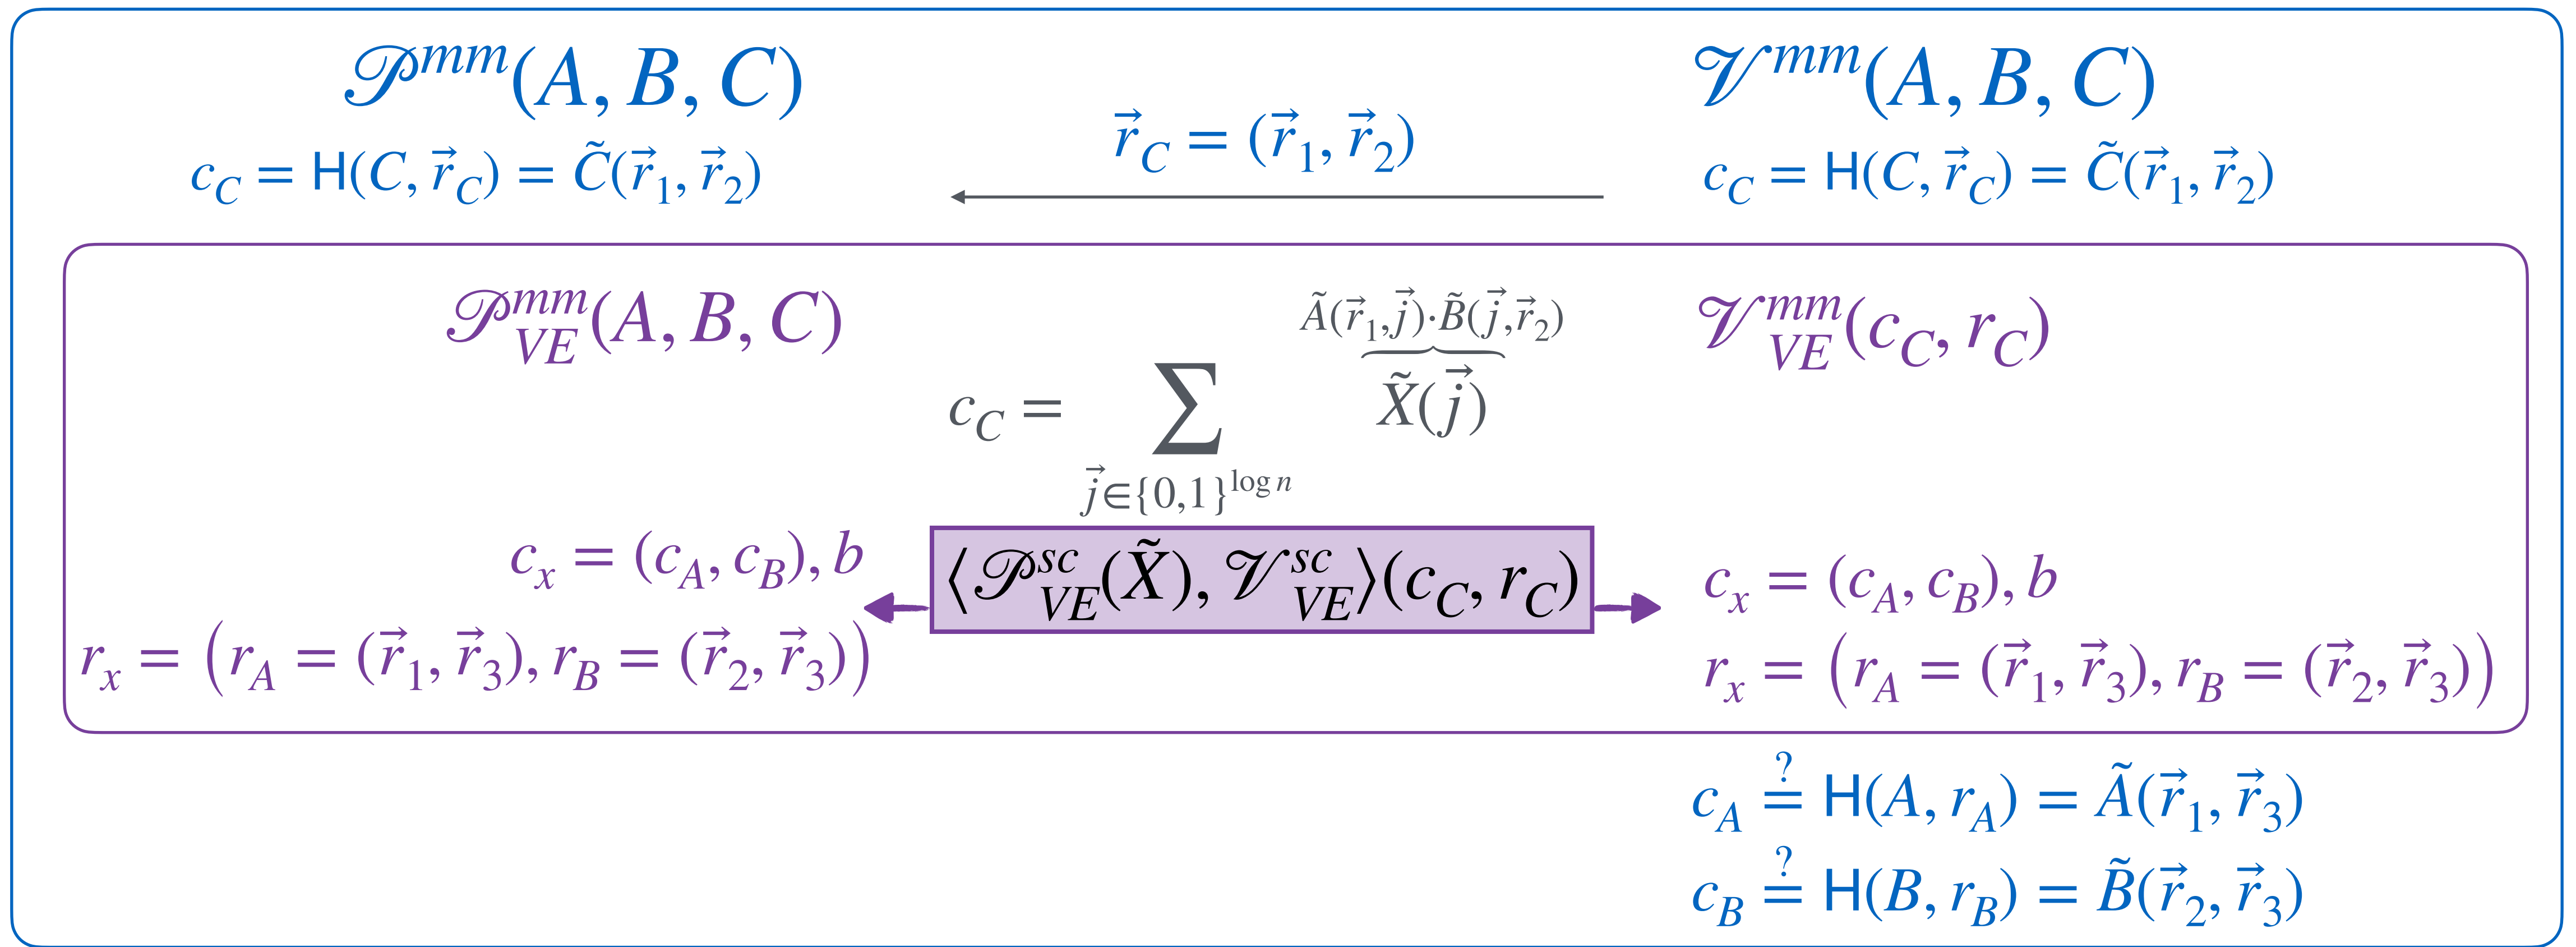
Prove $C = A \cdot B$ for $A, B, C \in \mathbb{F}^{n \times n}$. Using MLE: $\forall \vec{i}, \vec{k} \in \{0,1\}^{\log n} : \tilde{C}(\vec{i}, \vec{k}) = \sum_{\vec{j} \in \{0,1\}^{\log n}} \tilde{A}(\vec{i}, \vec{j}) \cdot \tilde{B}(\vec{j}, \vec{k})$



VE efficiency: communication & verification $O(\log n)$, prover $O(n^2)$, faster than computing $A \cdot B$ in $O(n^3)$

VE for matrix multiplication [Thaler13]

Prove $C = A \cdot B$ for $A, B, C \in \mathbb{F}^{n \times n}$. Using MLE: $\forall \vec{i}, \vec{k} \in \{0,1\}^{\log n} : \tilde{C}(\vec{i}, \vec{k}) = \sum_{\vec{j} \in \{0,1\}^{\log n}} \tilde{A}(\vec{i}, \vec{j}) \cdot \tilde{B}(\vec{j}, \vec{k})$



VE efficiency: communication & verification $O(\log n)$, prover $O(n^2)$, faster than computing $A \cdot B$ in $O(n^3)$

Proofs for convolution

zkCNN [LXZ21]

(2dim)Convolution \rightarrow 1dim-convolution \rightarrow poly mult \rightarrow FFT \rightarrow matrix multiplication

[BFGRS23]

Convolution \rightarrow matrix multiplication with reshaped X, W

Performance for c input channels, d output channels

	[BFGRS23]	zkCNN [LXZ21]
Prover	$O(c W (Y + d))$	$O(c d X)$
Verifier	$O(\log(c Y))$	$O(\log^2(c d X))$
Proof size	$O(\log(c Y))$	$O(\log^2(c d X))$

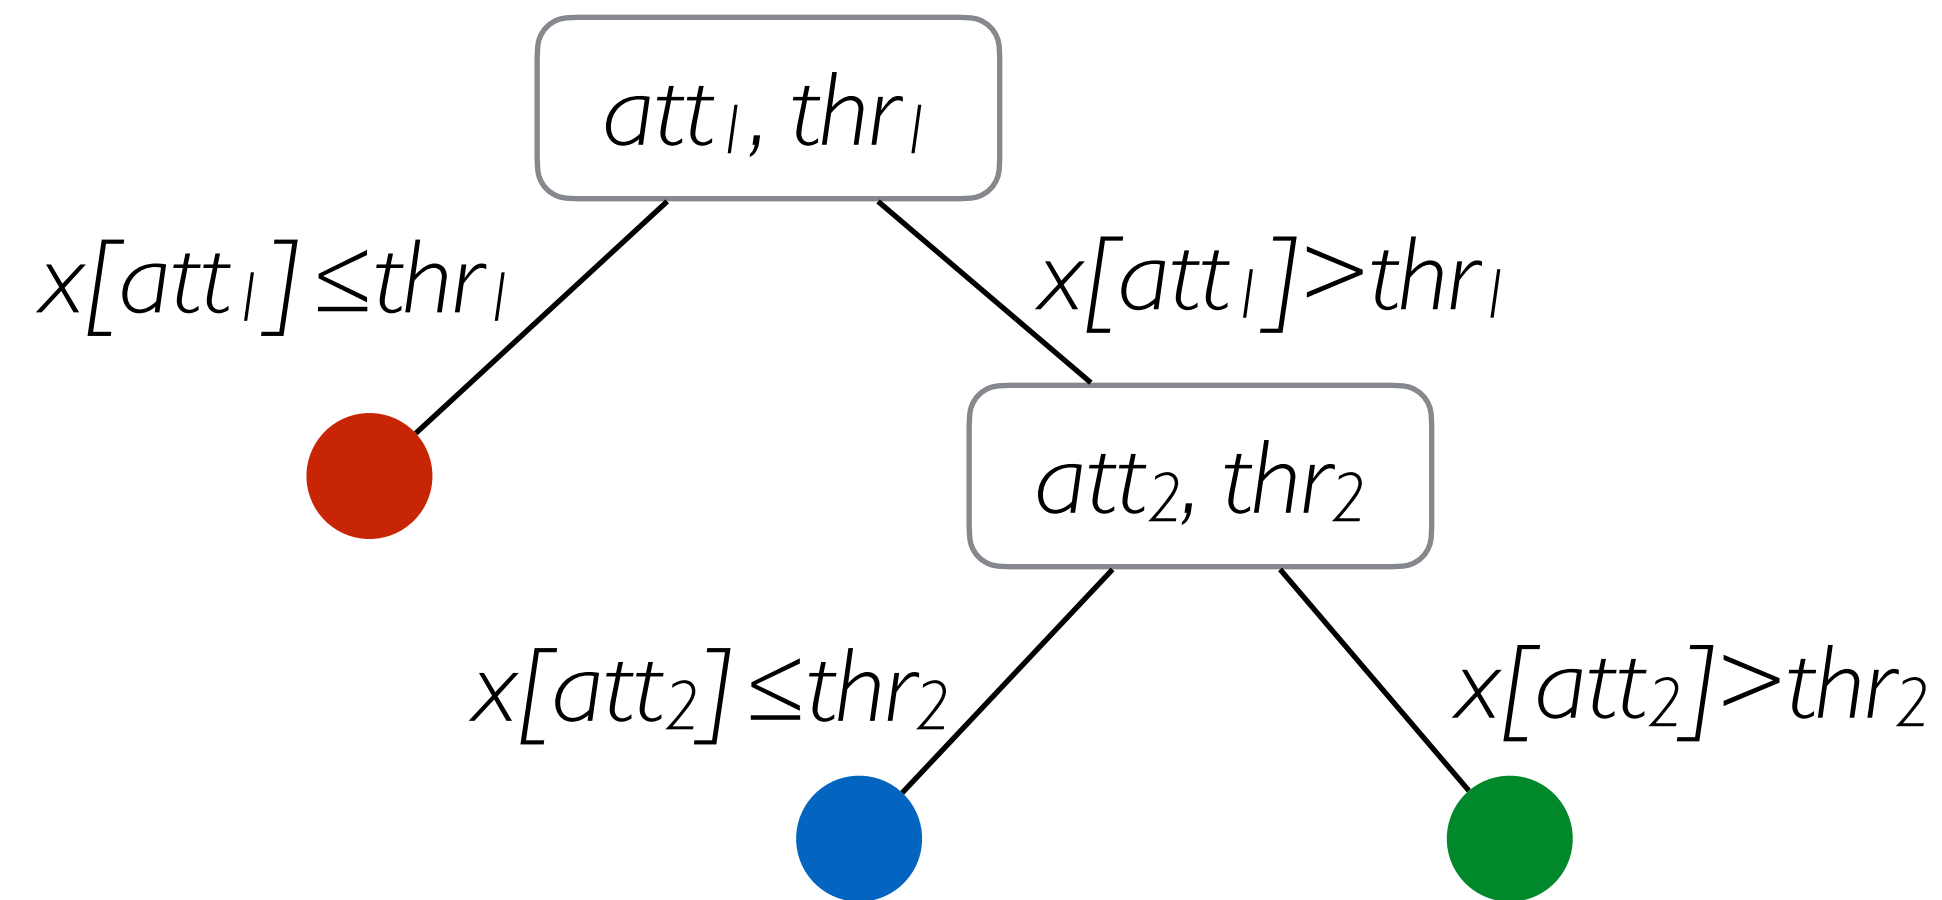
- Very efficient for small kernels $m^2 \leq d$ (VGG16 $m=3, d \rightarrow 512$)
- Confirmed experimentally (proving VGG11 $\sim 5s$)

ZKPs for Decision Trees

Tree $\mathcal{T} : \mathbb{F}^d \rightarrow [M]$

Height H , #nodes N

Classify $x = (x[att_1], \dots, x[att_d])$



[zkDT] J. Zhang, Z. Fang, Y. Zhang, D. Song. *Zero Knowledge Proofs for Decision Tree Predictions and Accuracy*. CCS 2020.



Merkle hash of T . Proving inference = proving traversal from leaf (class) to root

[CFFLL24] M. Campanelli, A. Faonio, D. Fiore, T. Li, H. Lipmaa. *Lookup Arguments: Improvements, Extensions and Applications to Zero-Knowledge Decision Trees*. PKC 2024



Matrix-encoding of T . Proving inference = proving matrix lookup (reduced to vector lookup)

ZKPs for secure ML

This talk: How to use ZKPs and Commitments to prove ML inference
Efficient solutions for CNNs and Decision Trees

ZKPs for secure ML

This talk: How to use ZKPs and Commitments to prove ML inference
Efficient solutions for CNNs and Decision Trees
Other related problems

ZKPs for secure ML

This talk: How to use ZKPs and Commitments to prove ML inference
Efficient solutions for CNNs and Decision Trees

Other related problems

Proofs of accuracy: prove that model **W** achieves a claimed accuracy level

Prove $Y_i = F(X_i, W)$ over dataset of labeled samples $\{X_i\}$ and compares results to labels

ZKPs for secure ML

This talk: How to use ZKPs and Commitments to prove ML inference
Efficient solutions for CNNs and Decision Trees

Other related problems

Proofs of accuracy: prove that model W achieves a claimed accuracy level

Prove $Y_i = F(X_i, W)$ over dataset of labeled samples $\{X_i\}$ and compares results to labels

Proofs of training: prove $W = \text{NN-Train}(\text{Data})$

Challenge: several iterations of inference-like computations

[GGJMMP23] S. Garg, A. Goel, S. Jha, S. Mahloujifar, M. Mahmoody, G. Policharla, M. Wang. *Experimenting with Zero-Knowledge Proofs of Training*. CCS 2023

[APKP24] K. Abbaszadeh, C. Pappas, J. Katz, D. Papadopoulos. *Zero-Knowledge Proofs of Training for Deep Neural Networks*. CCS 2024

Thanks!

Questions ?