

Cryptography for **Privacy-Preserving Machine Learning**

Morten Dahl

The tf-encrypted Project and Dropout Labs

Applied ML Days, AI & Trust track, EPFL, January 2019

Why?

Machine Learning Process



data set



training



prediction service

Machine Learning Process

IMAGENET



data set



training



prediction service



Machine Learning Process

IMAGENET




data set



training



prediction service



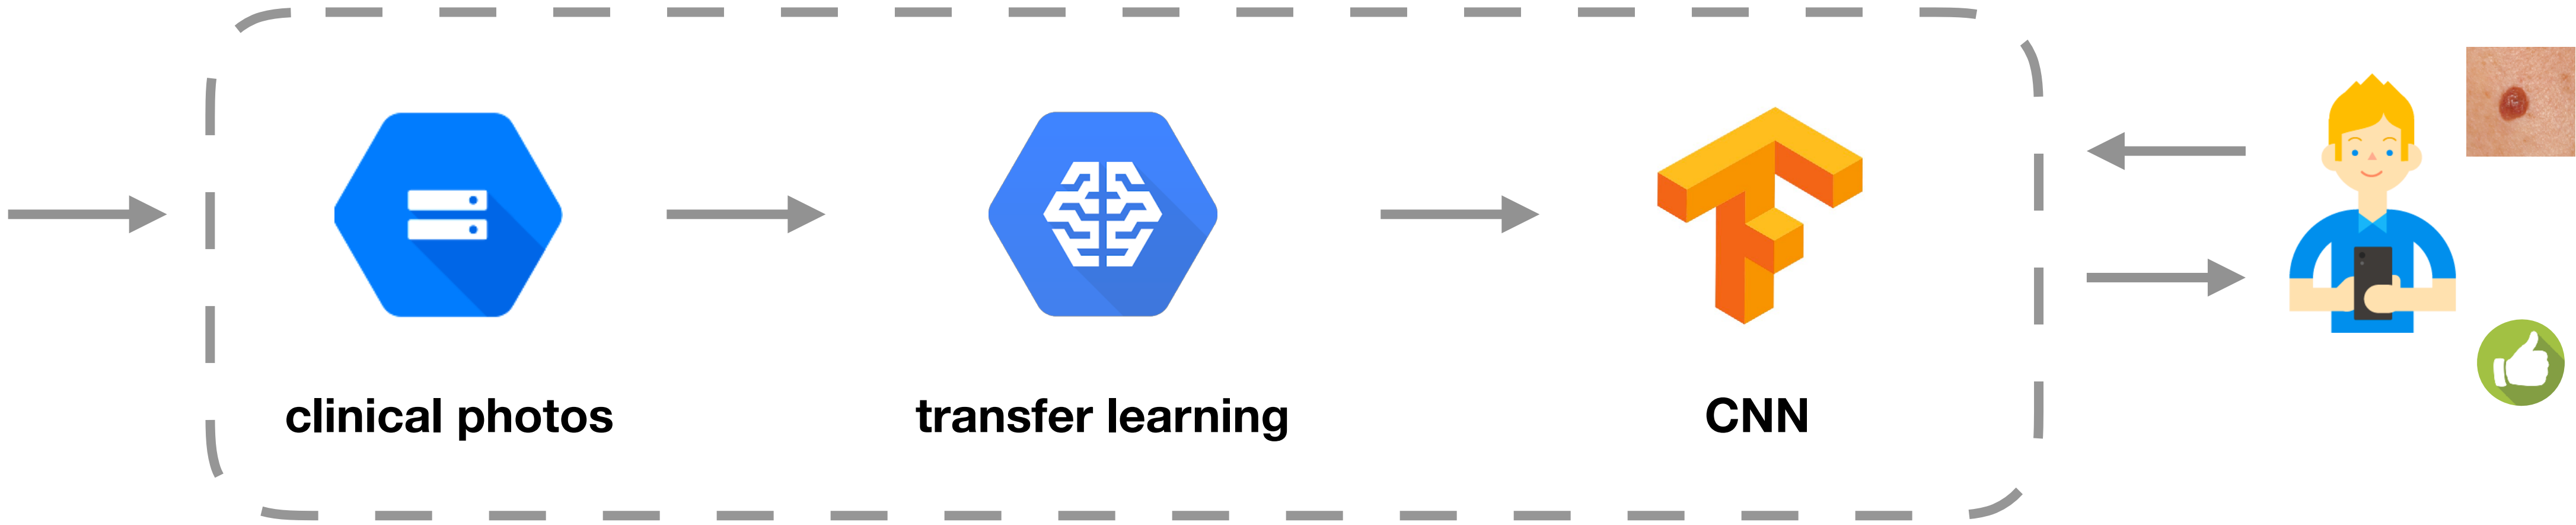



Skin Cancer Image Classification

Brett Kuprel

12:30-12:40pm

Join Brett Kuprel, and see how TensorFlow was used by the artificial intelligence lab and medical school of Stanford to classify skin cancer images. He'll describe the project steps: from acquiring a dataset, training a deep network, and evaluating of the results. To wrap up, Brett will give his take on the future of skin cancer image classification.




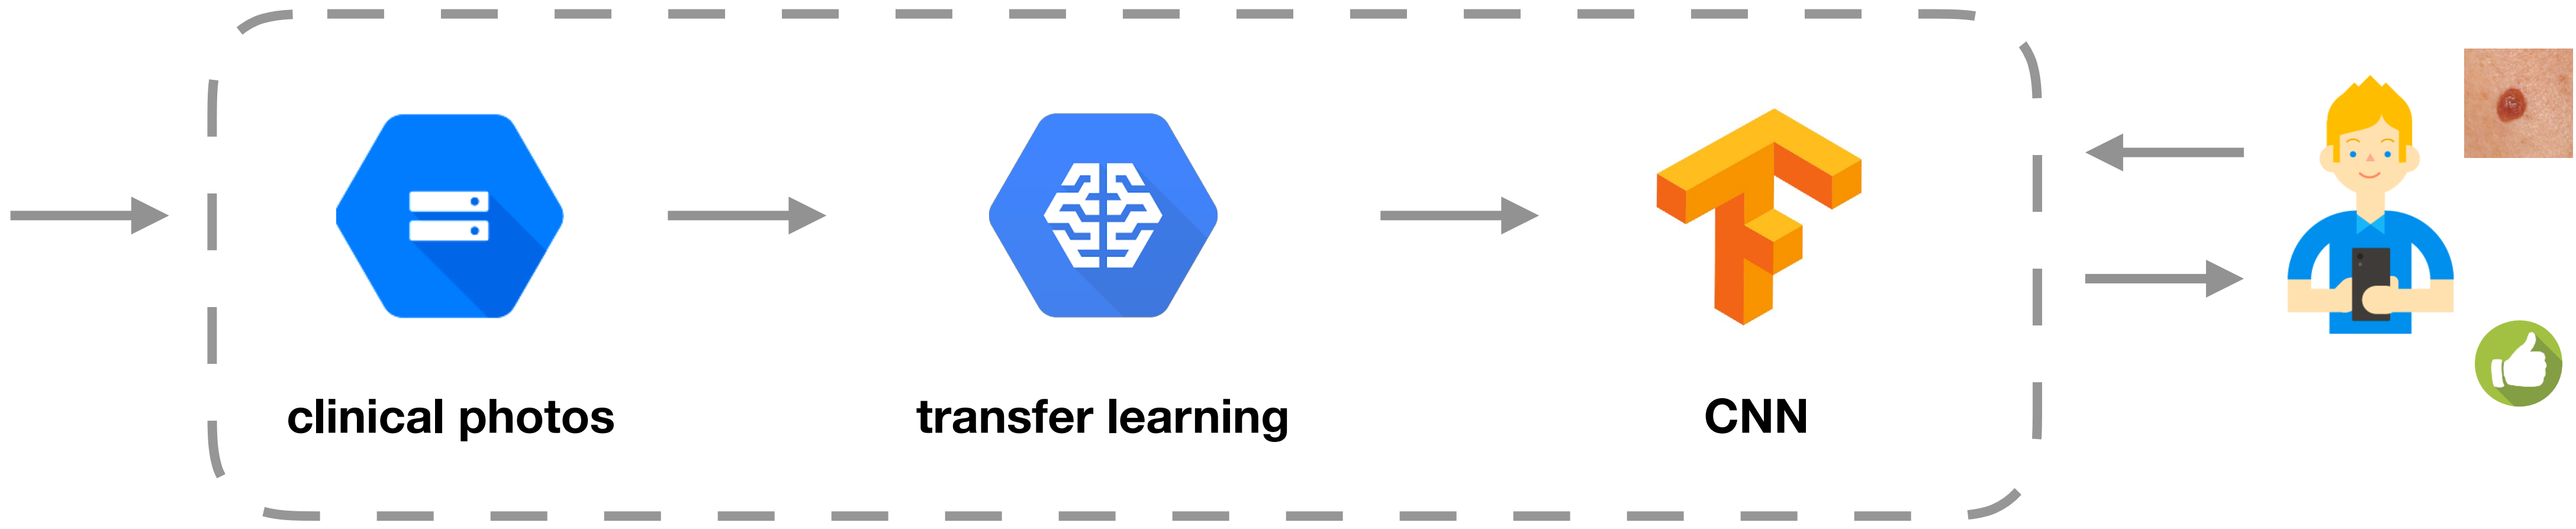


Skin Cancer Image Classification

Brett Kuprel

12:30-12:40pm

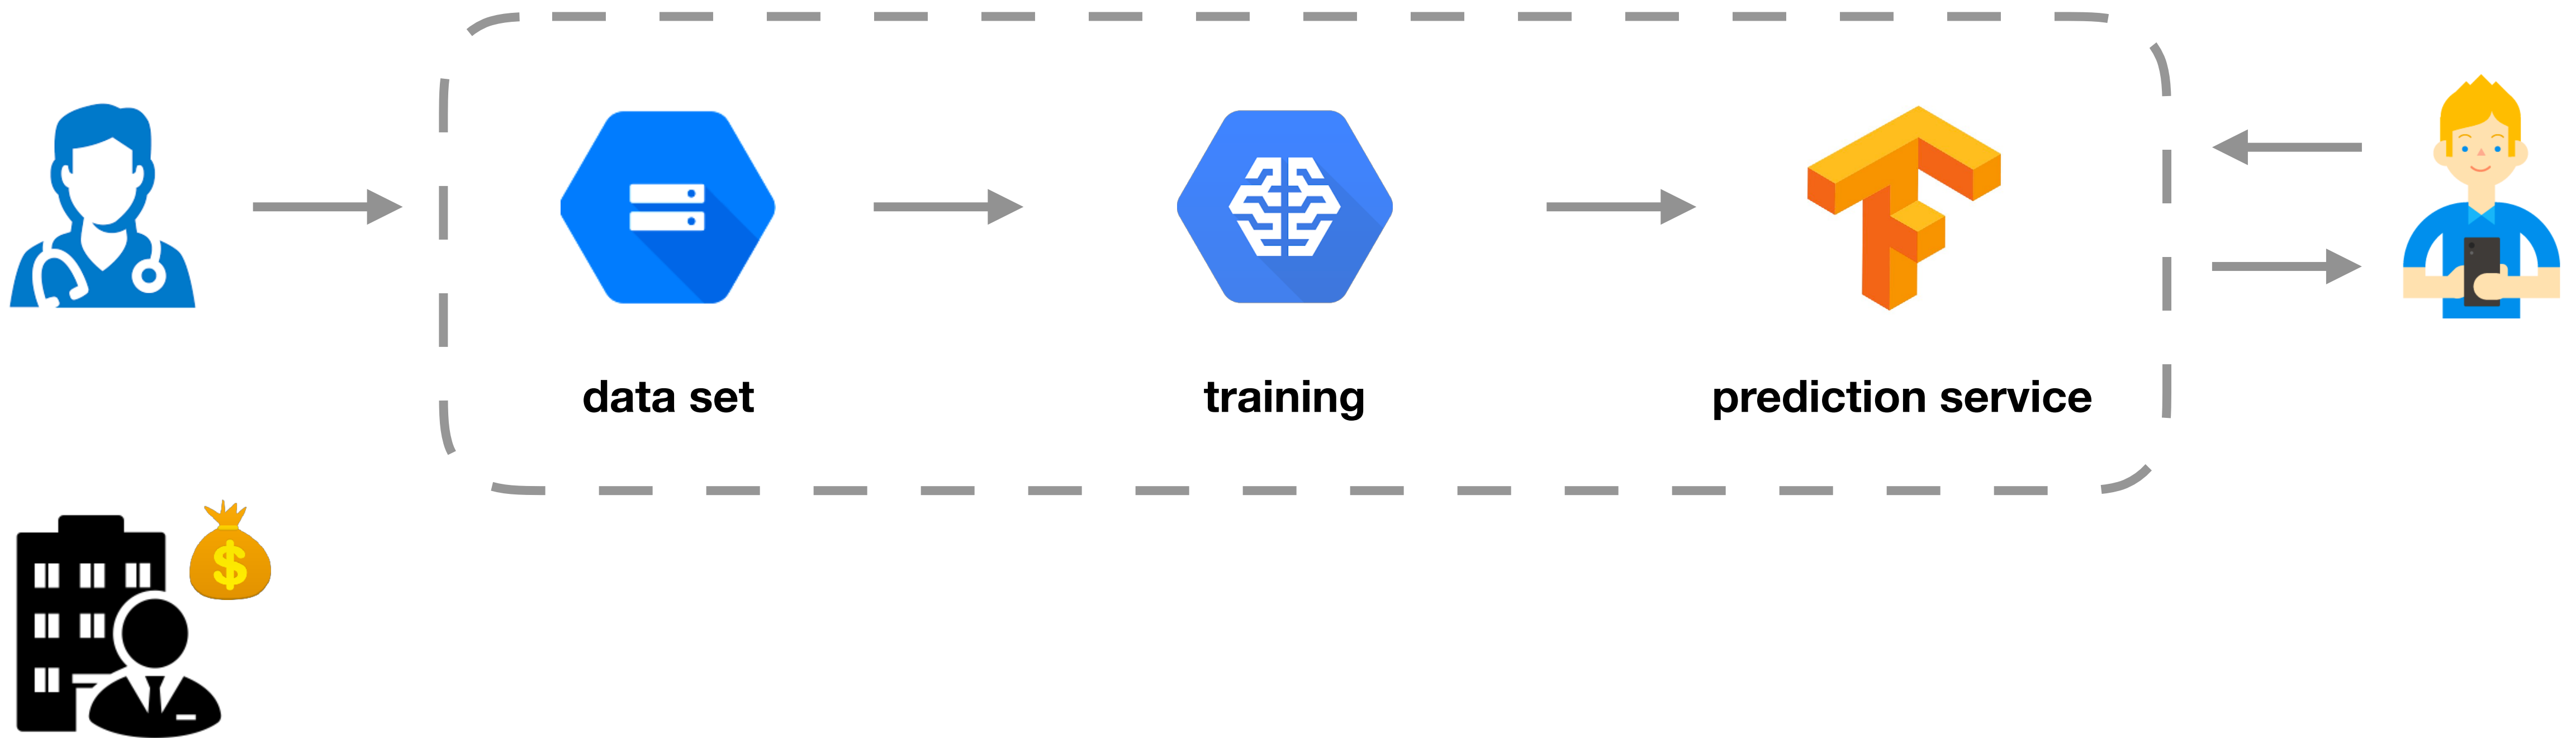
Join Brett Kuprel, and see how TensorFlow was used by the artificial intelligence lab and medical school of Stanford to classify skin cancer images. He'll describe the project steps: from acquiring a dataset, training a deep network, and evaluating of the results. To wrap up, Brett will give his take on the future of skin cancer image classification.



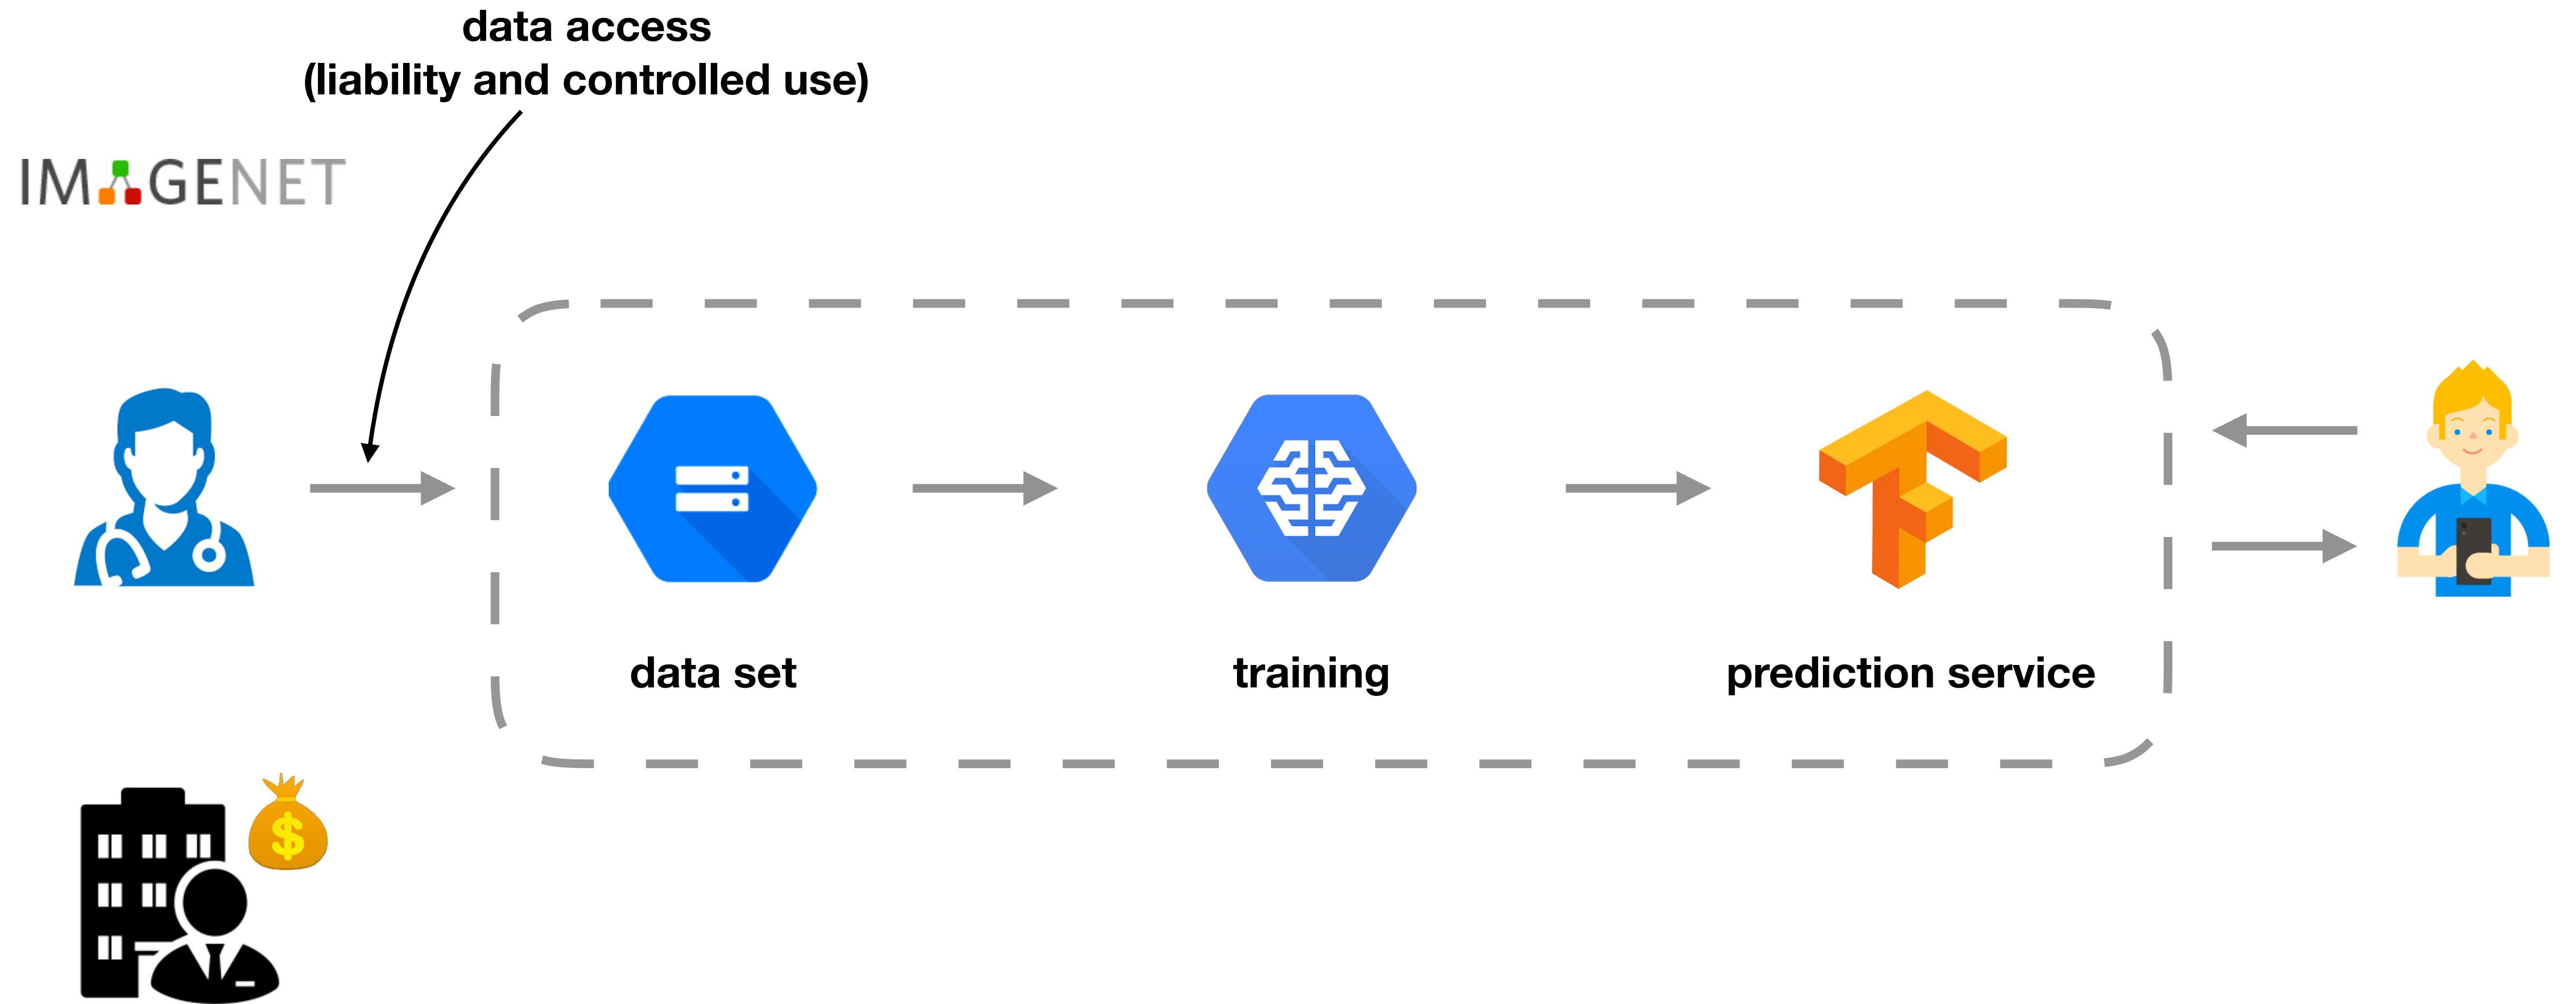
machine learning positioned to have huge impact on health care

Potential Bottlenecks

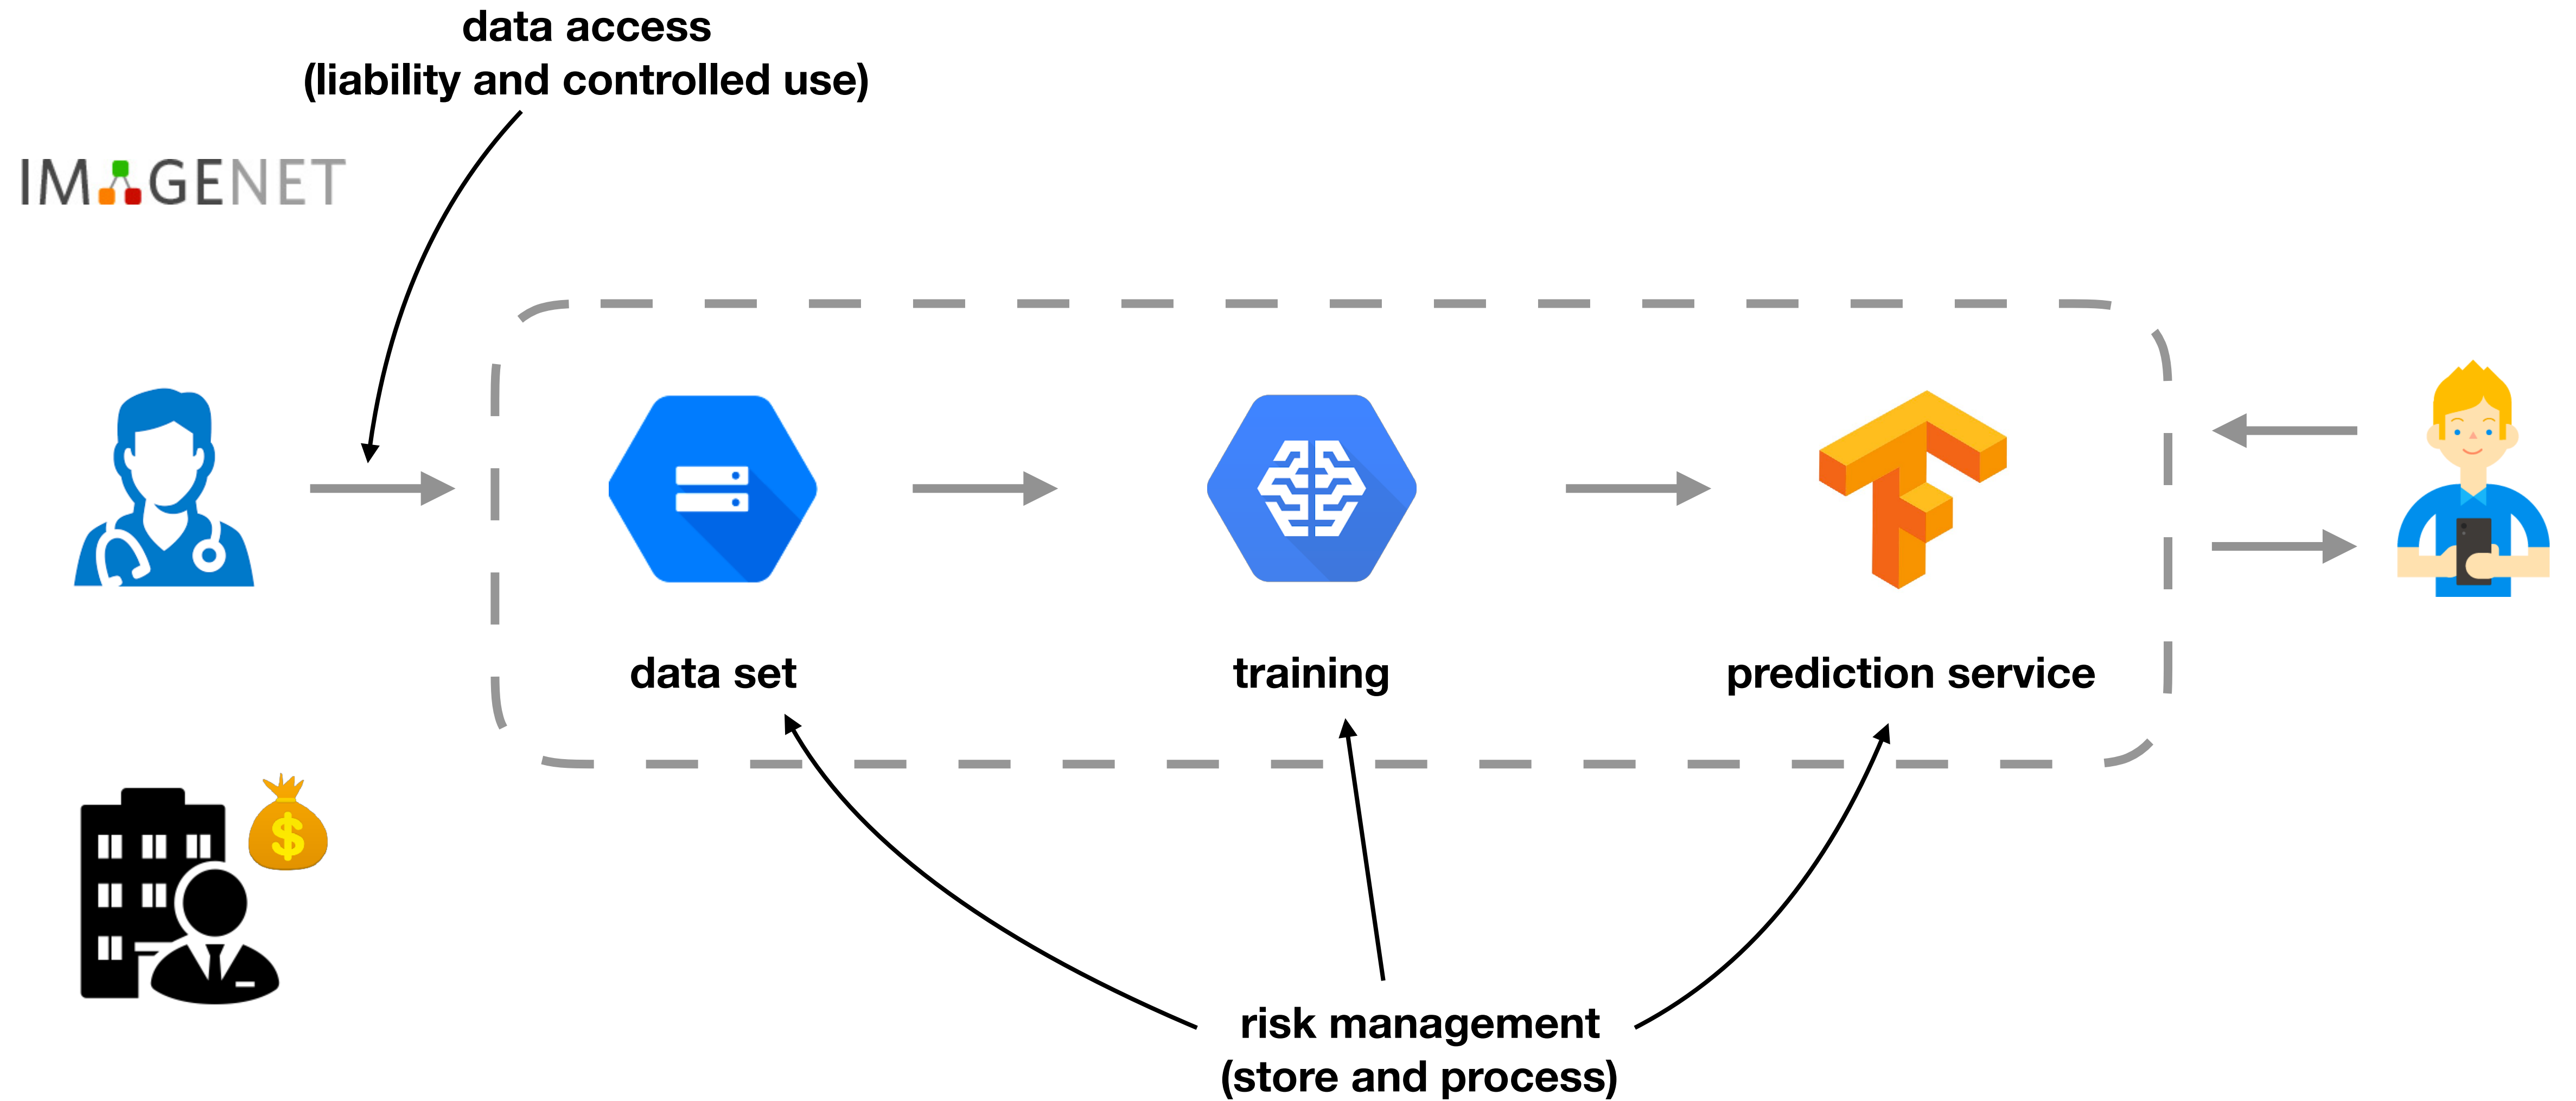
IMAGENET



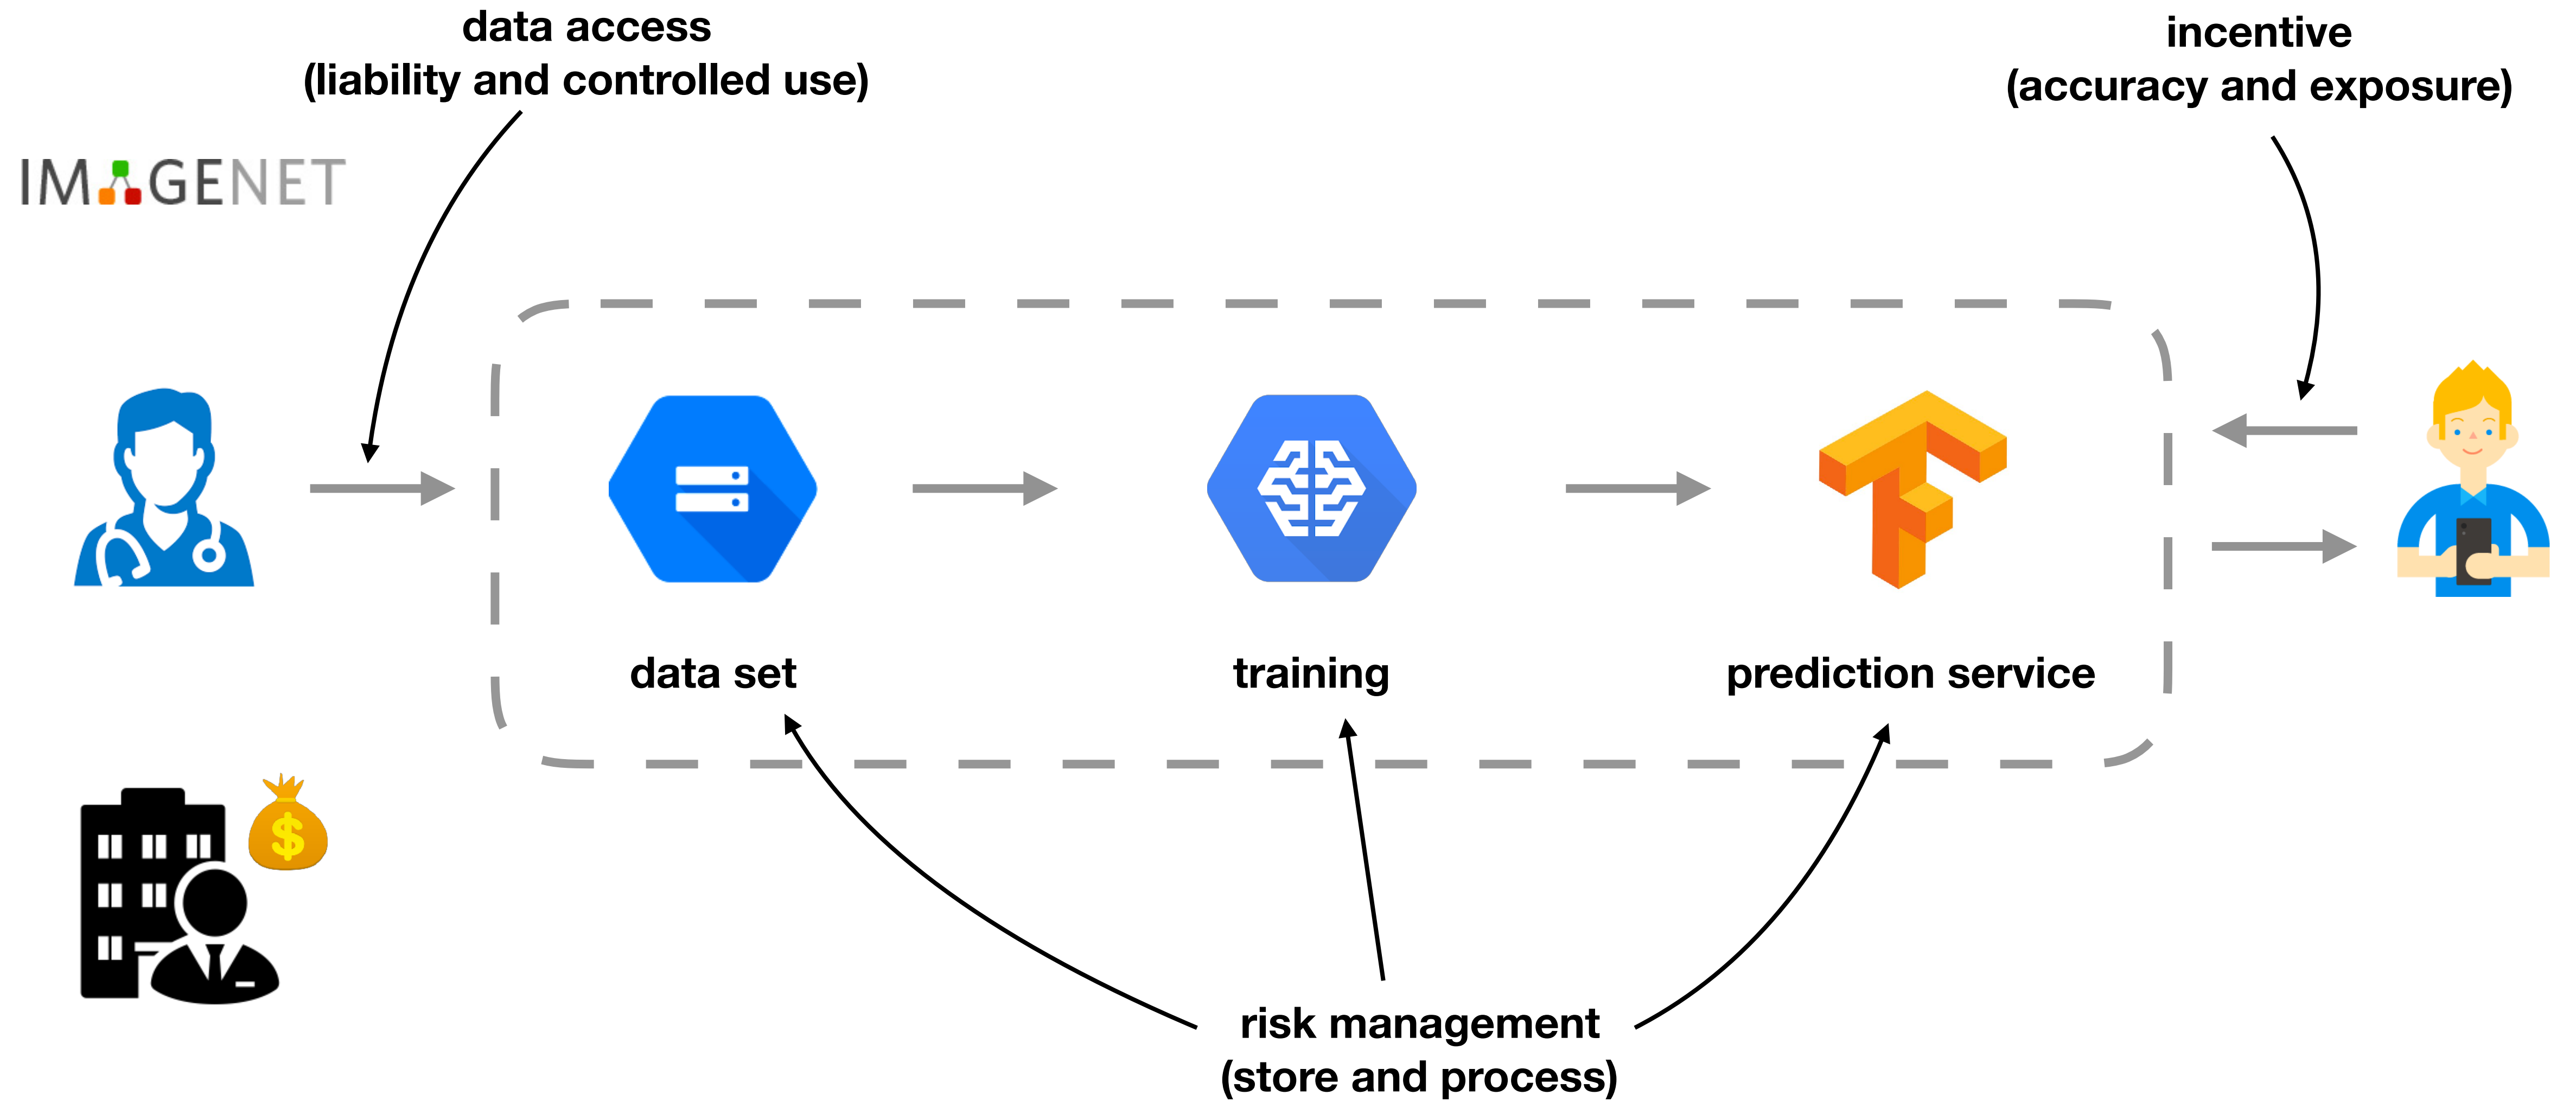
Potential Bottlenecks




Potential Bottlenecks



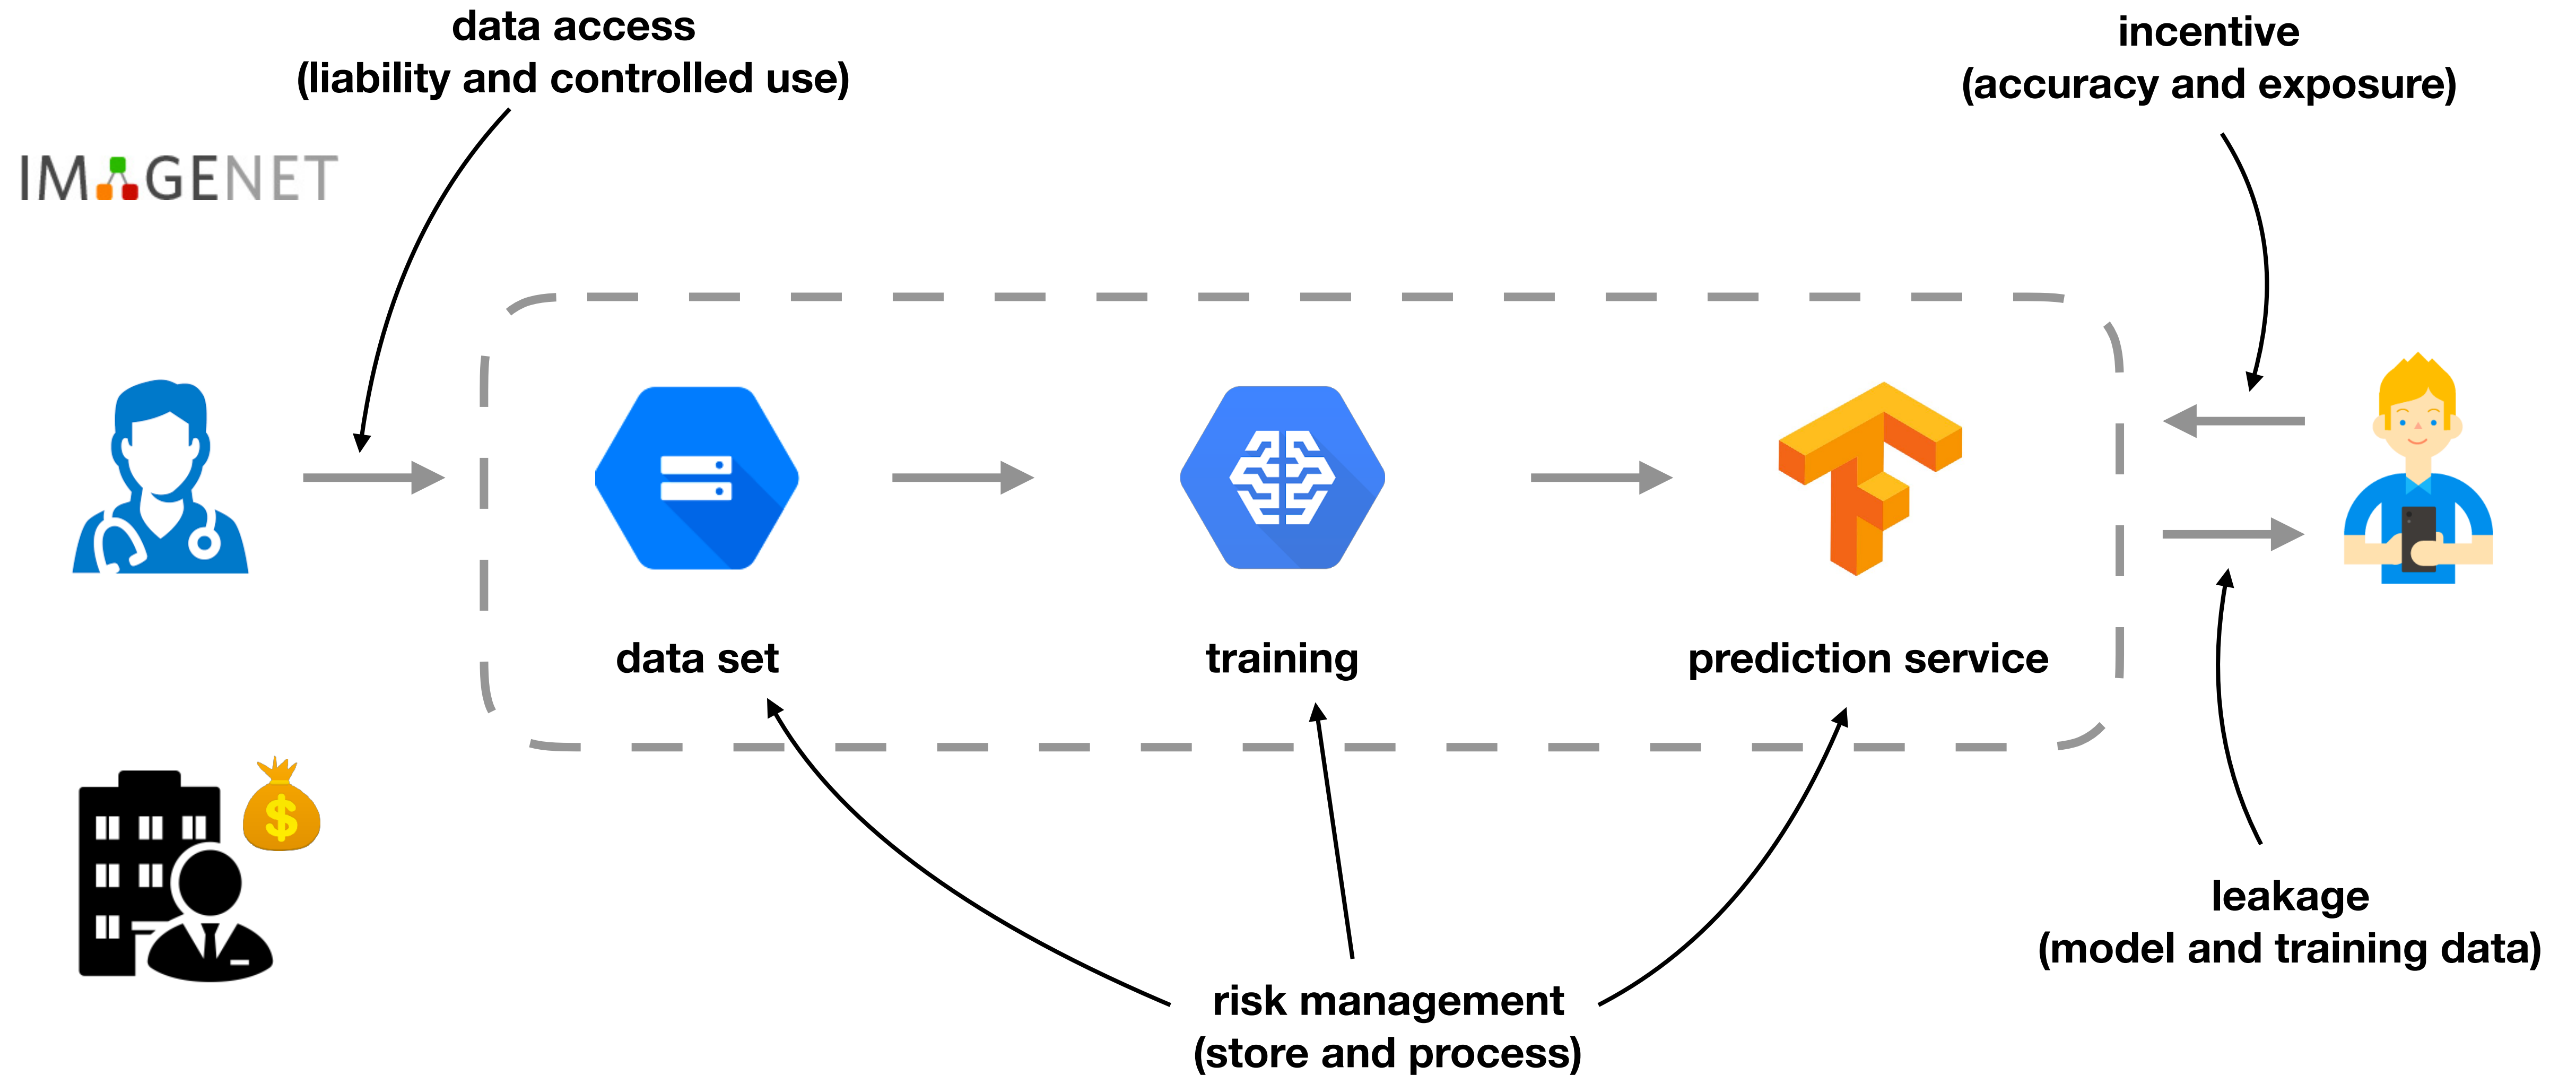
Potential Bottlenecks



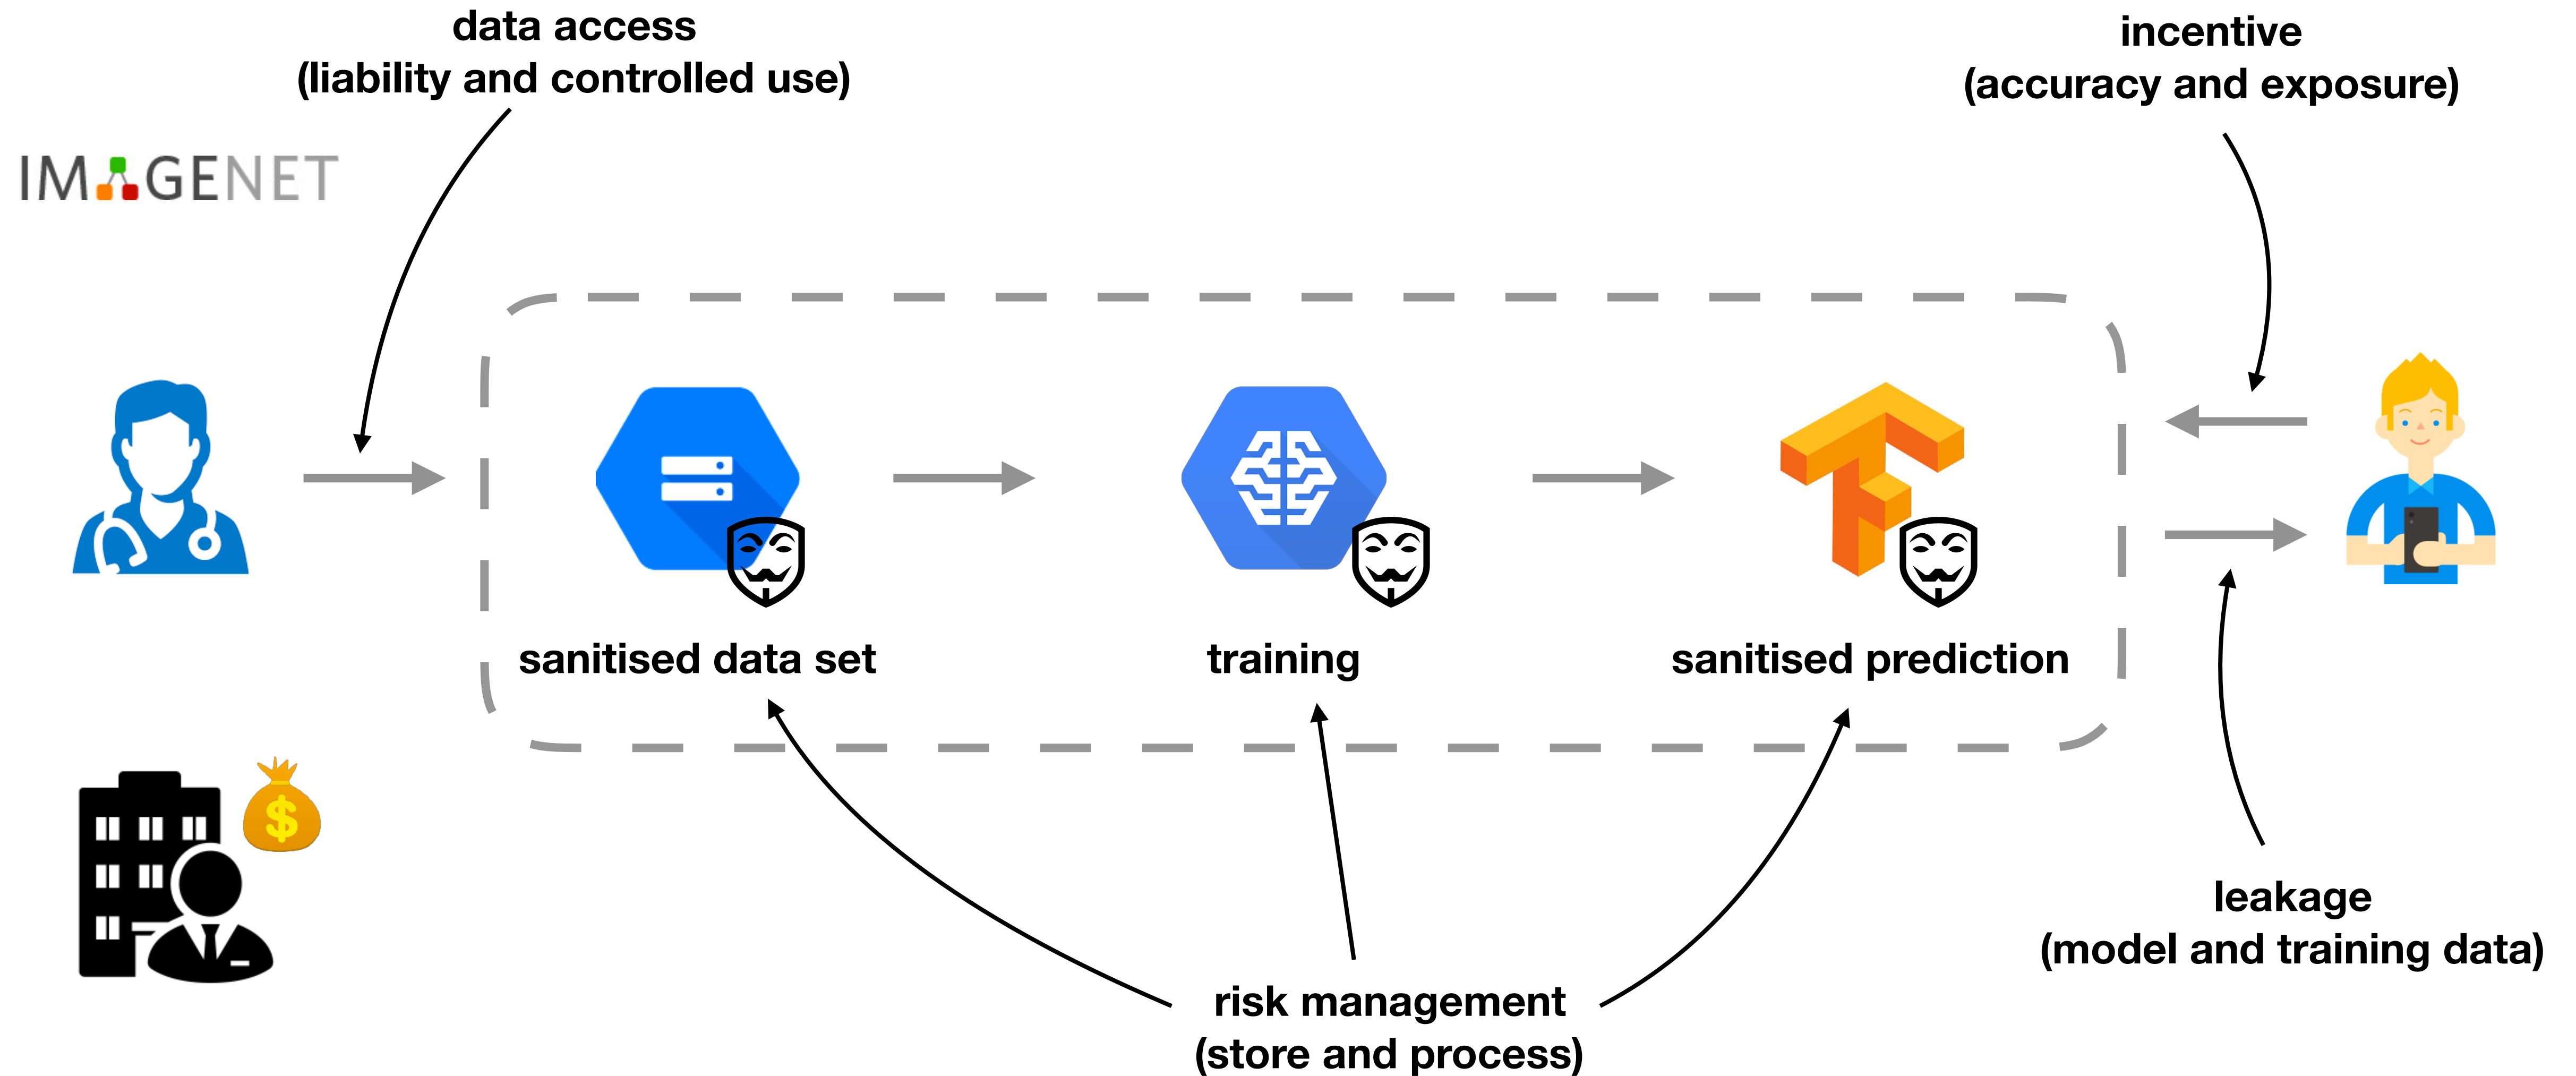
Potential Bottlenecks



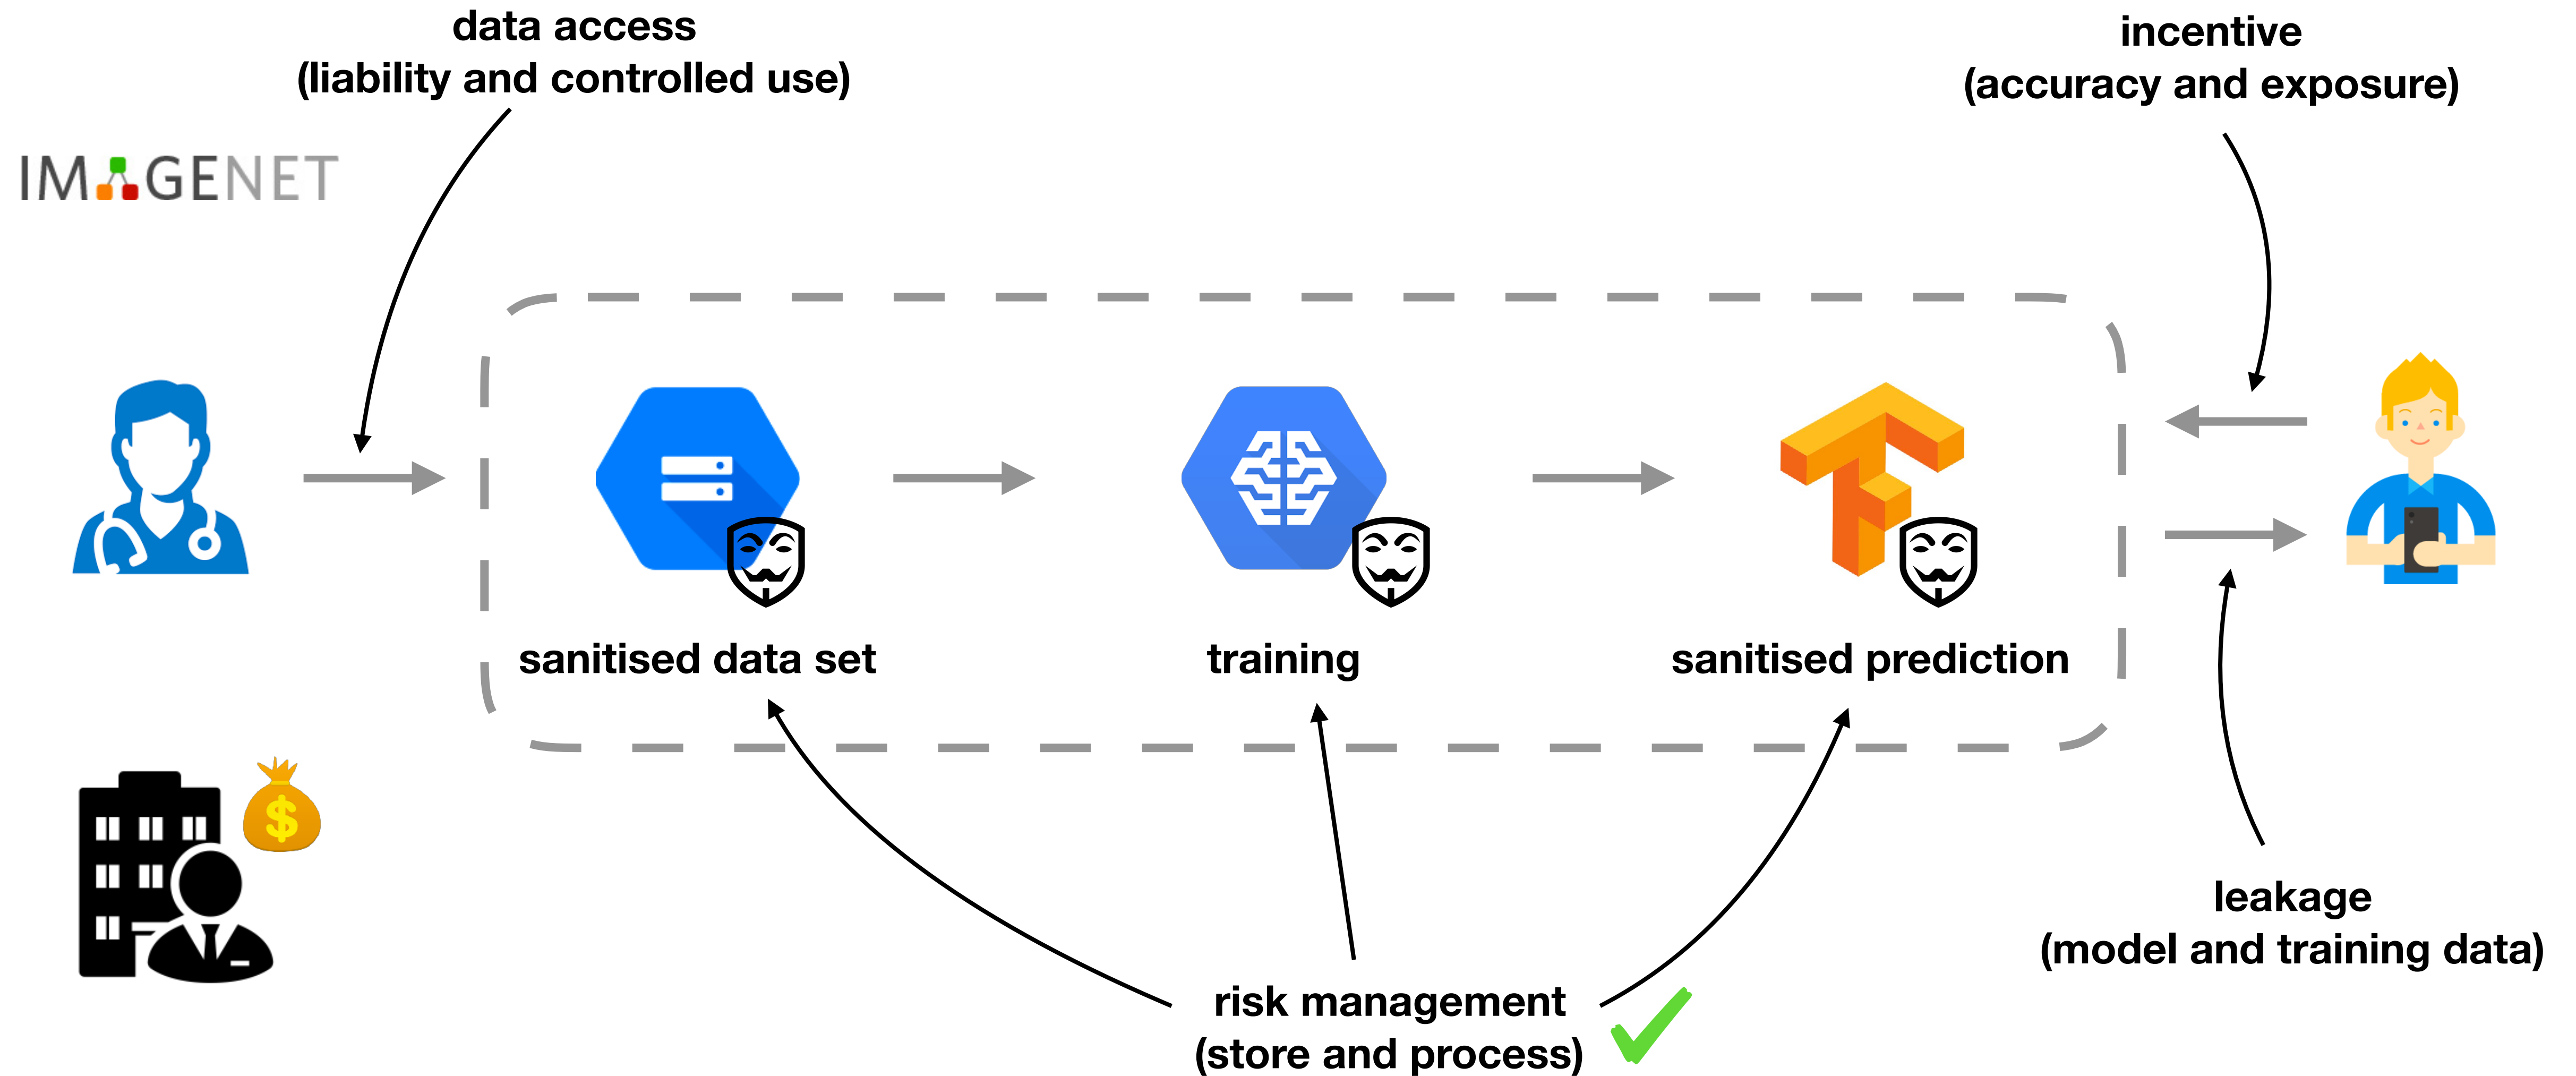
Differential Privacy



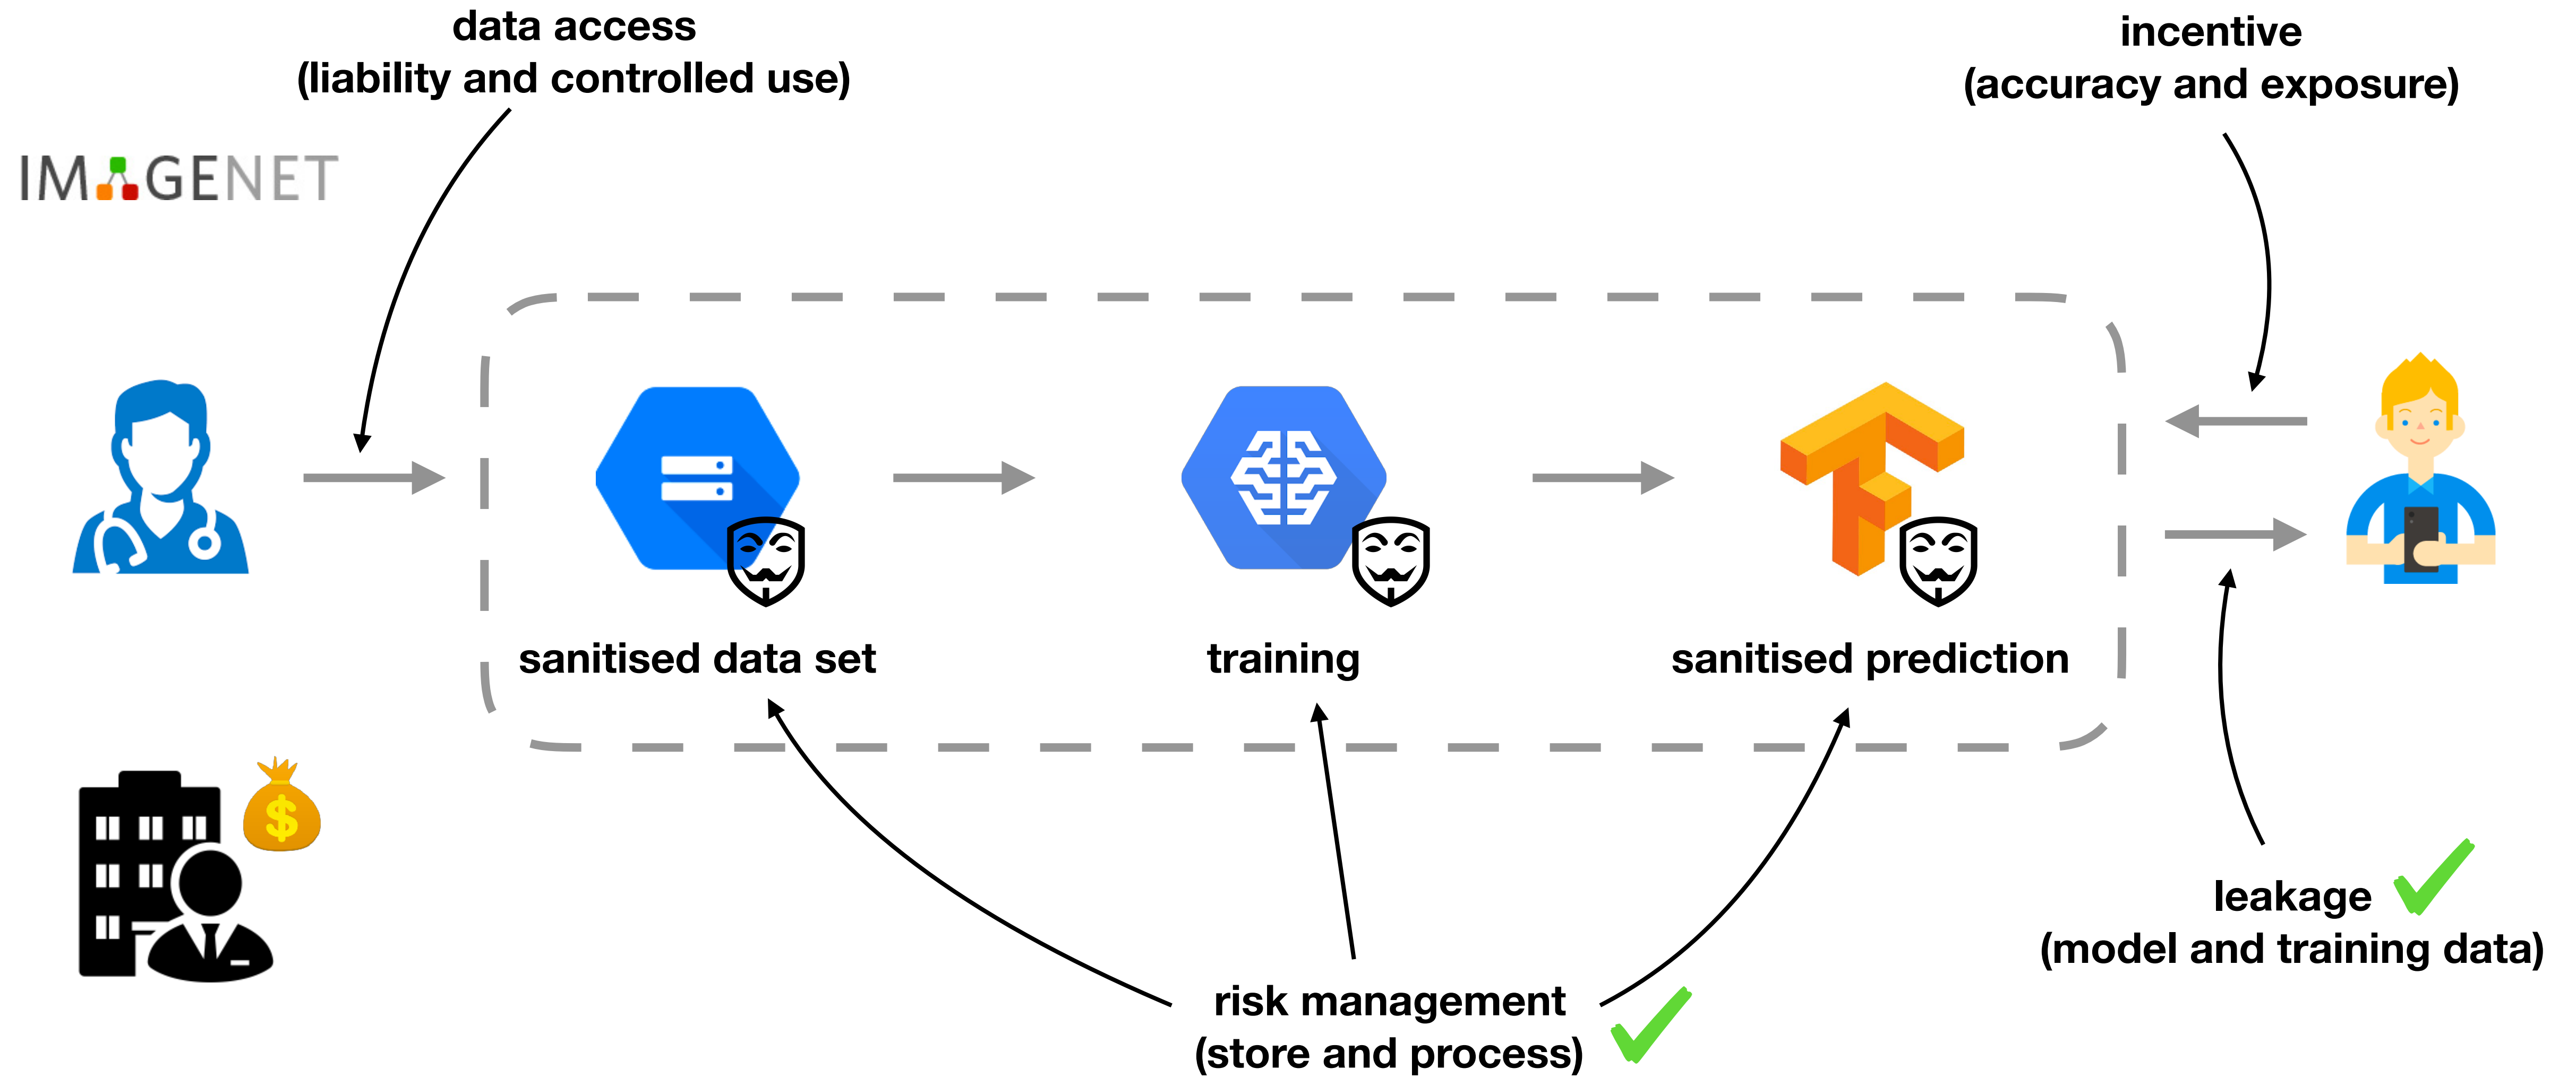
Differential Privacy



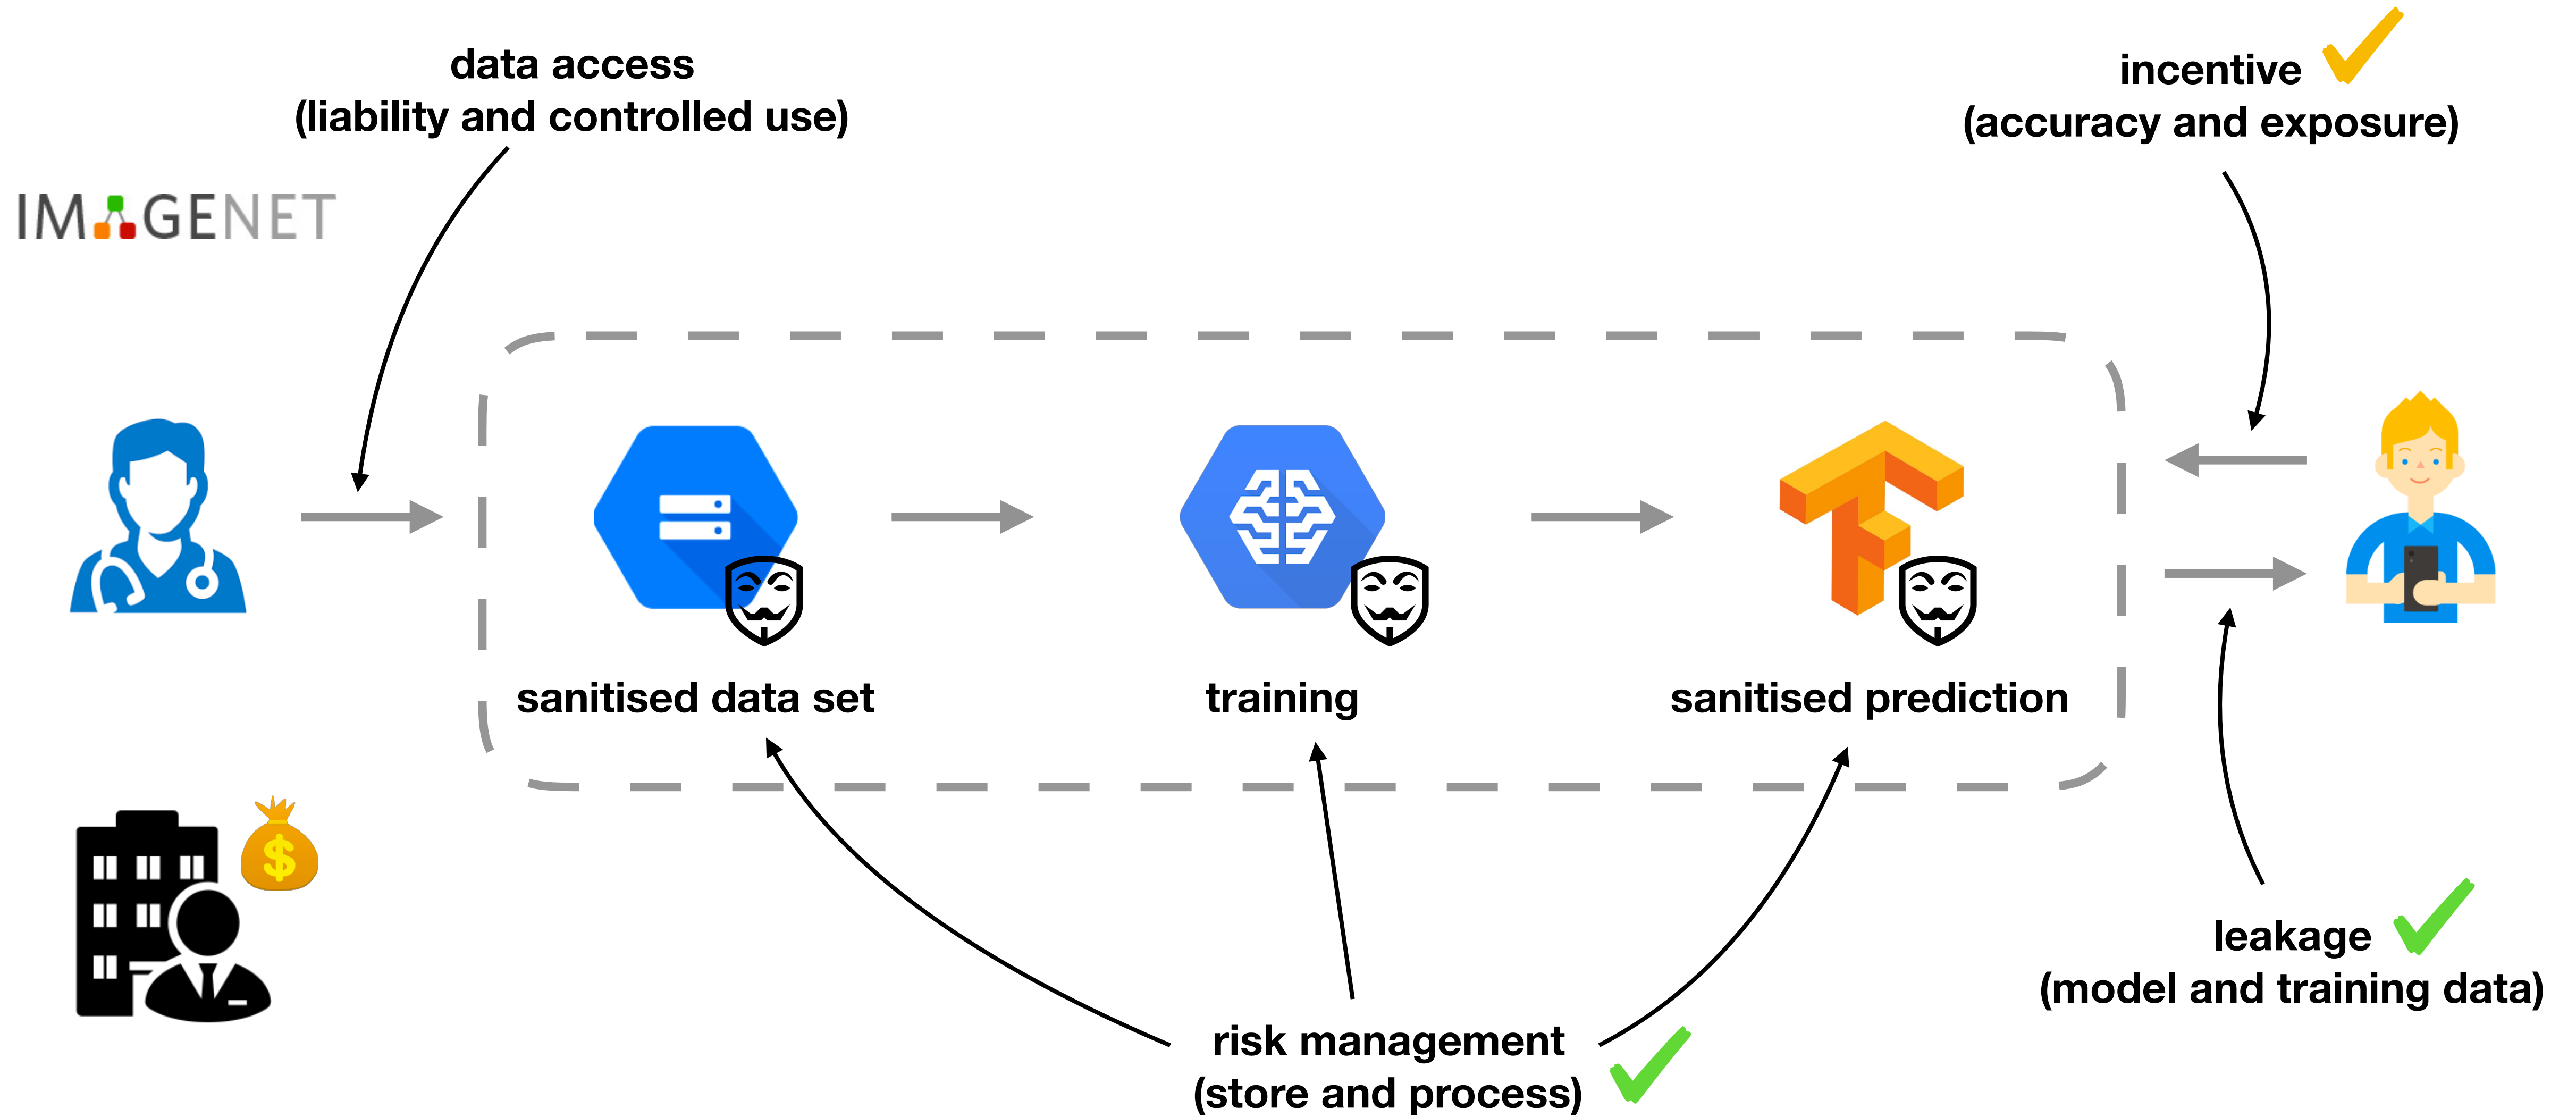
Differential Privacy



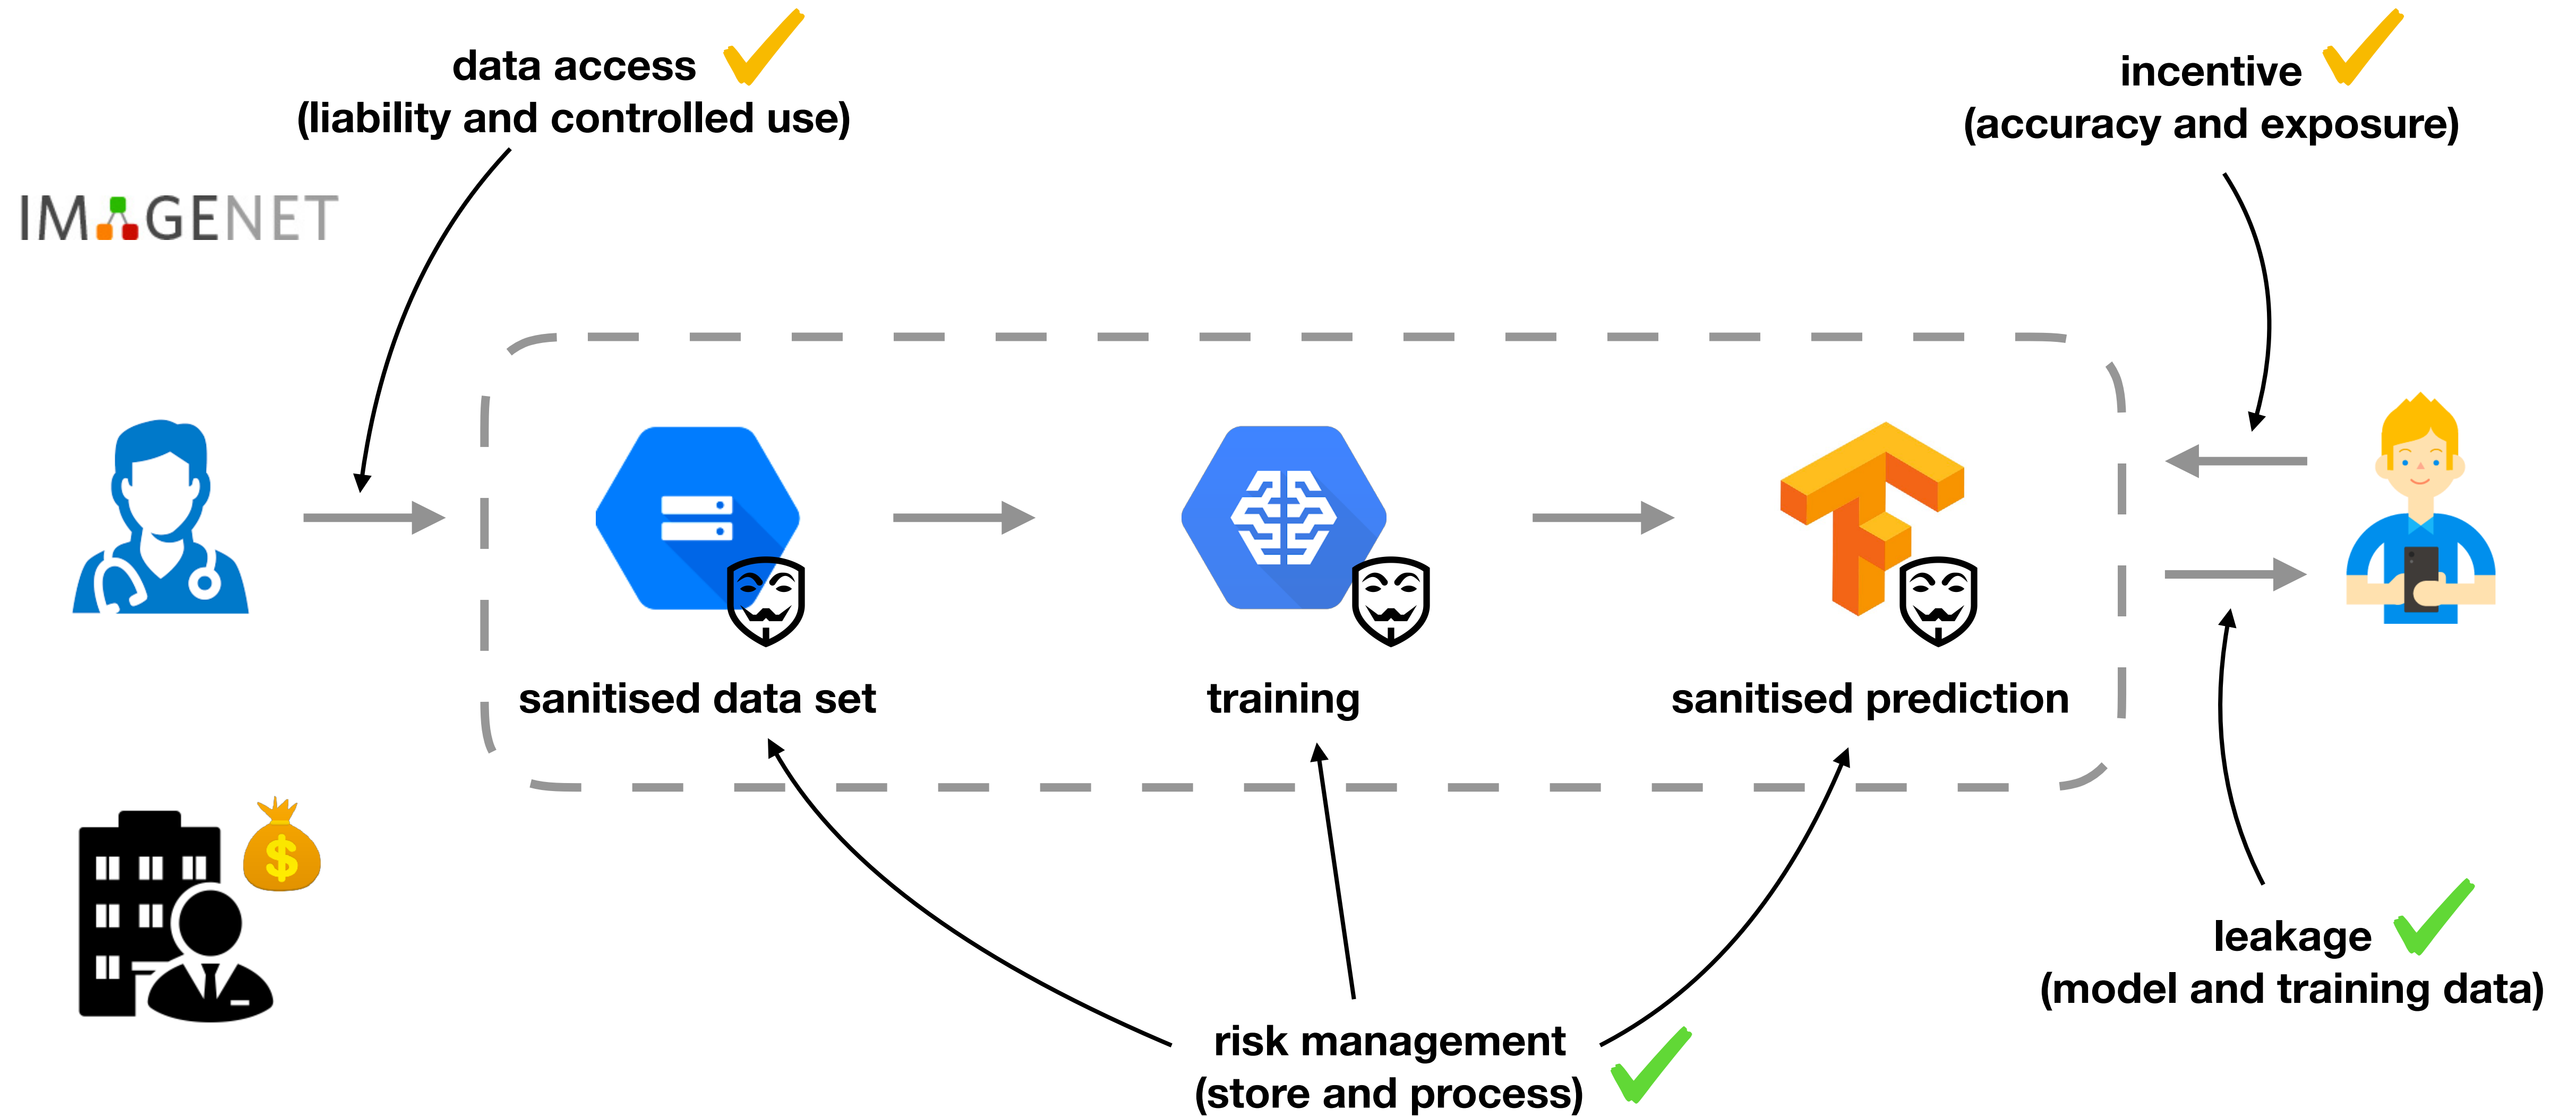
Differential Privacy



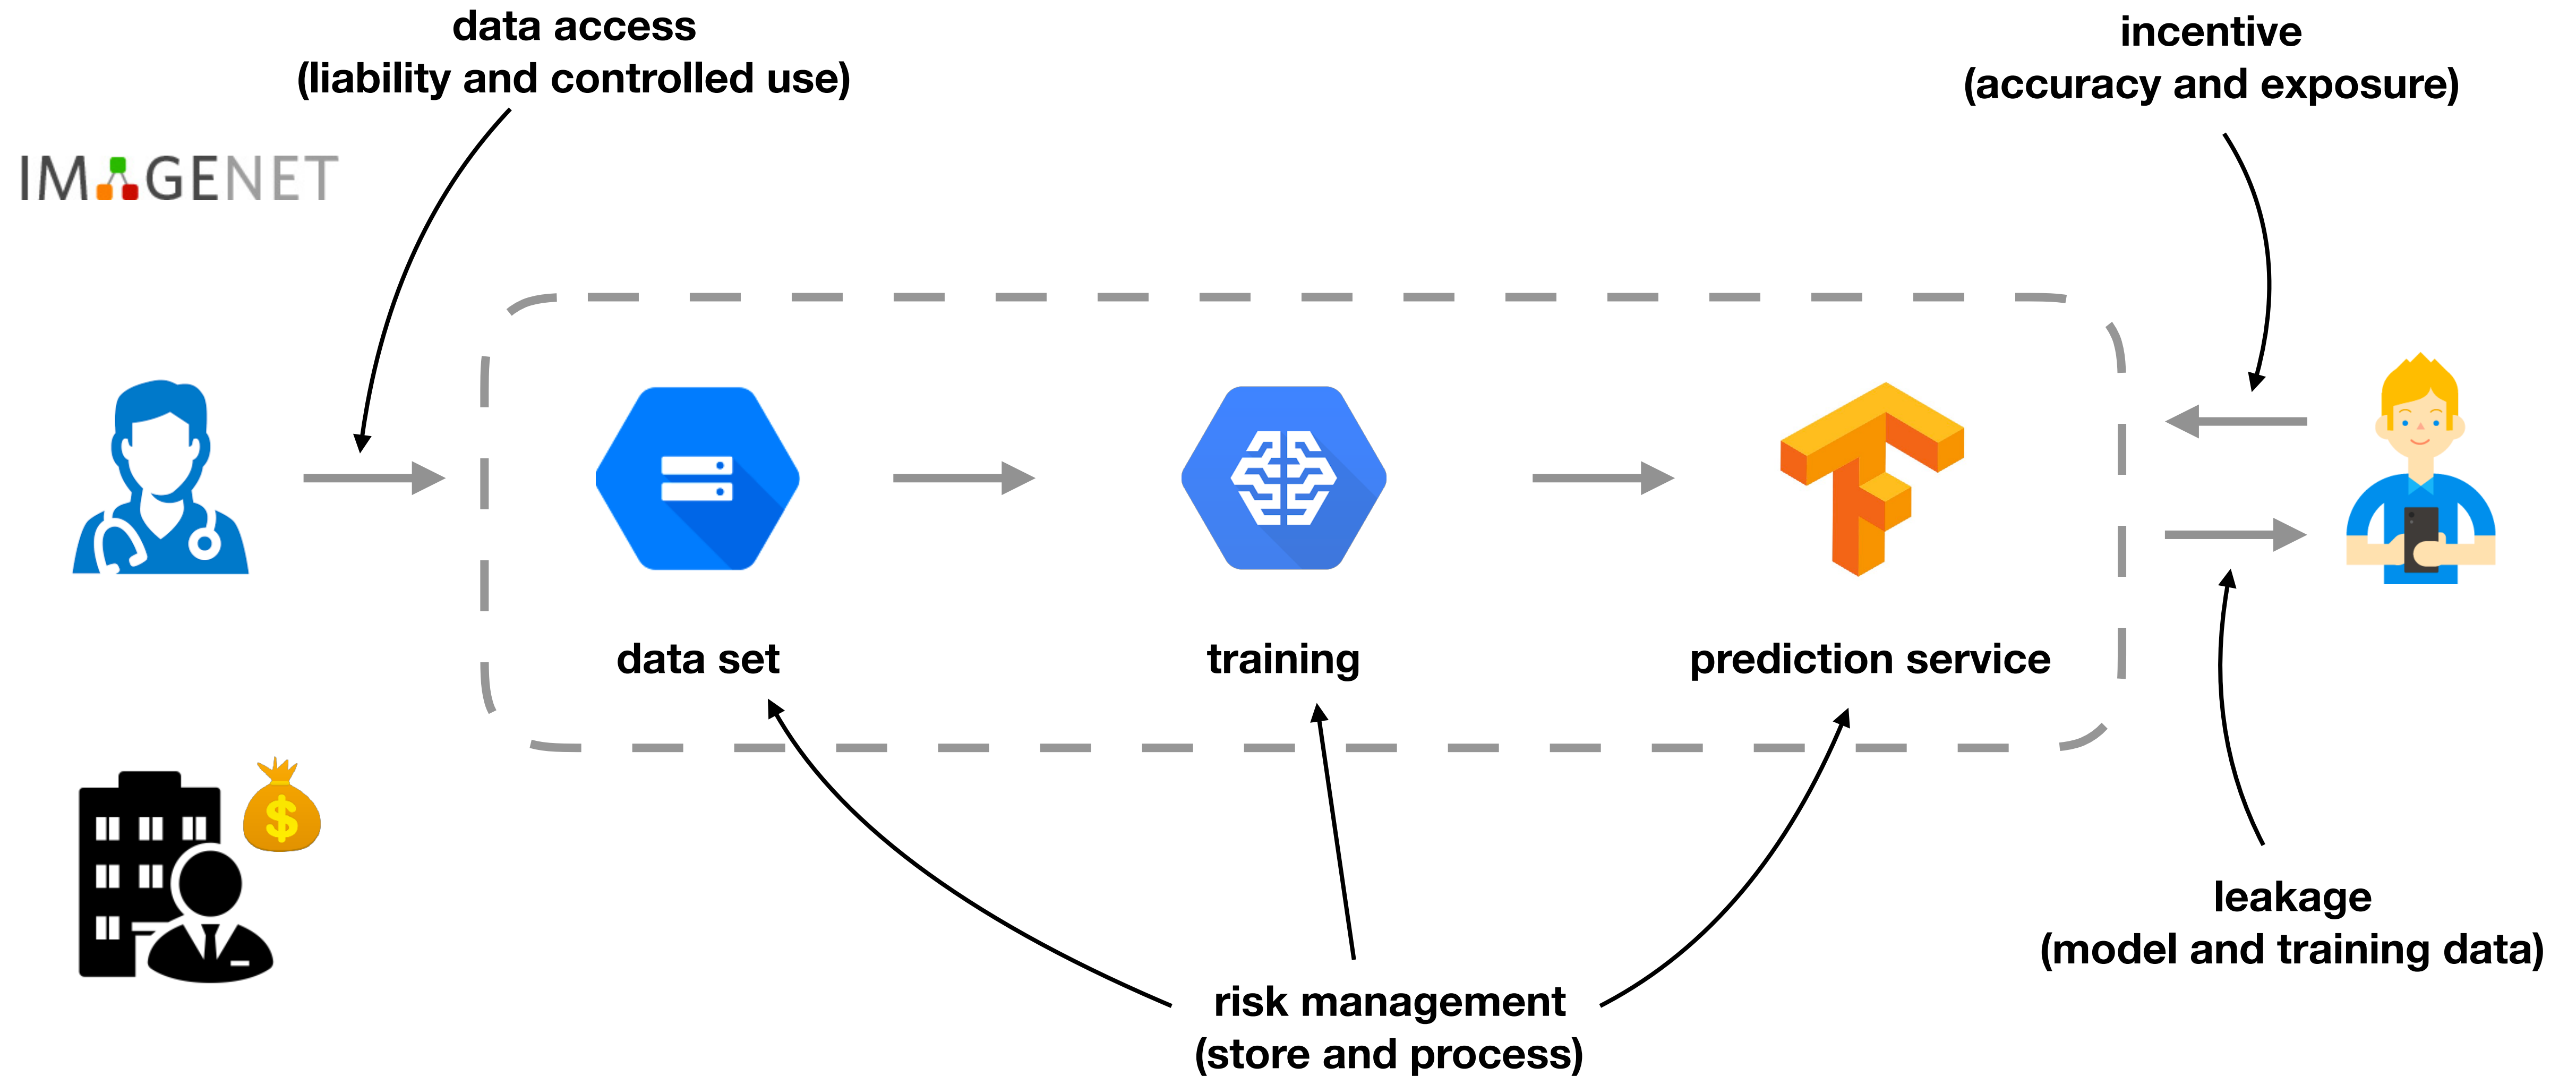
Differential Privacy



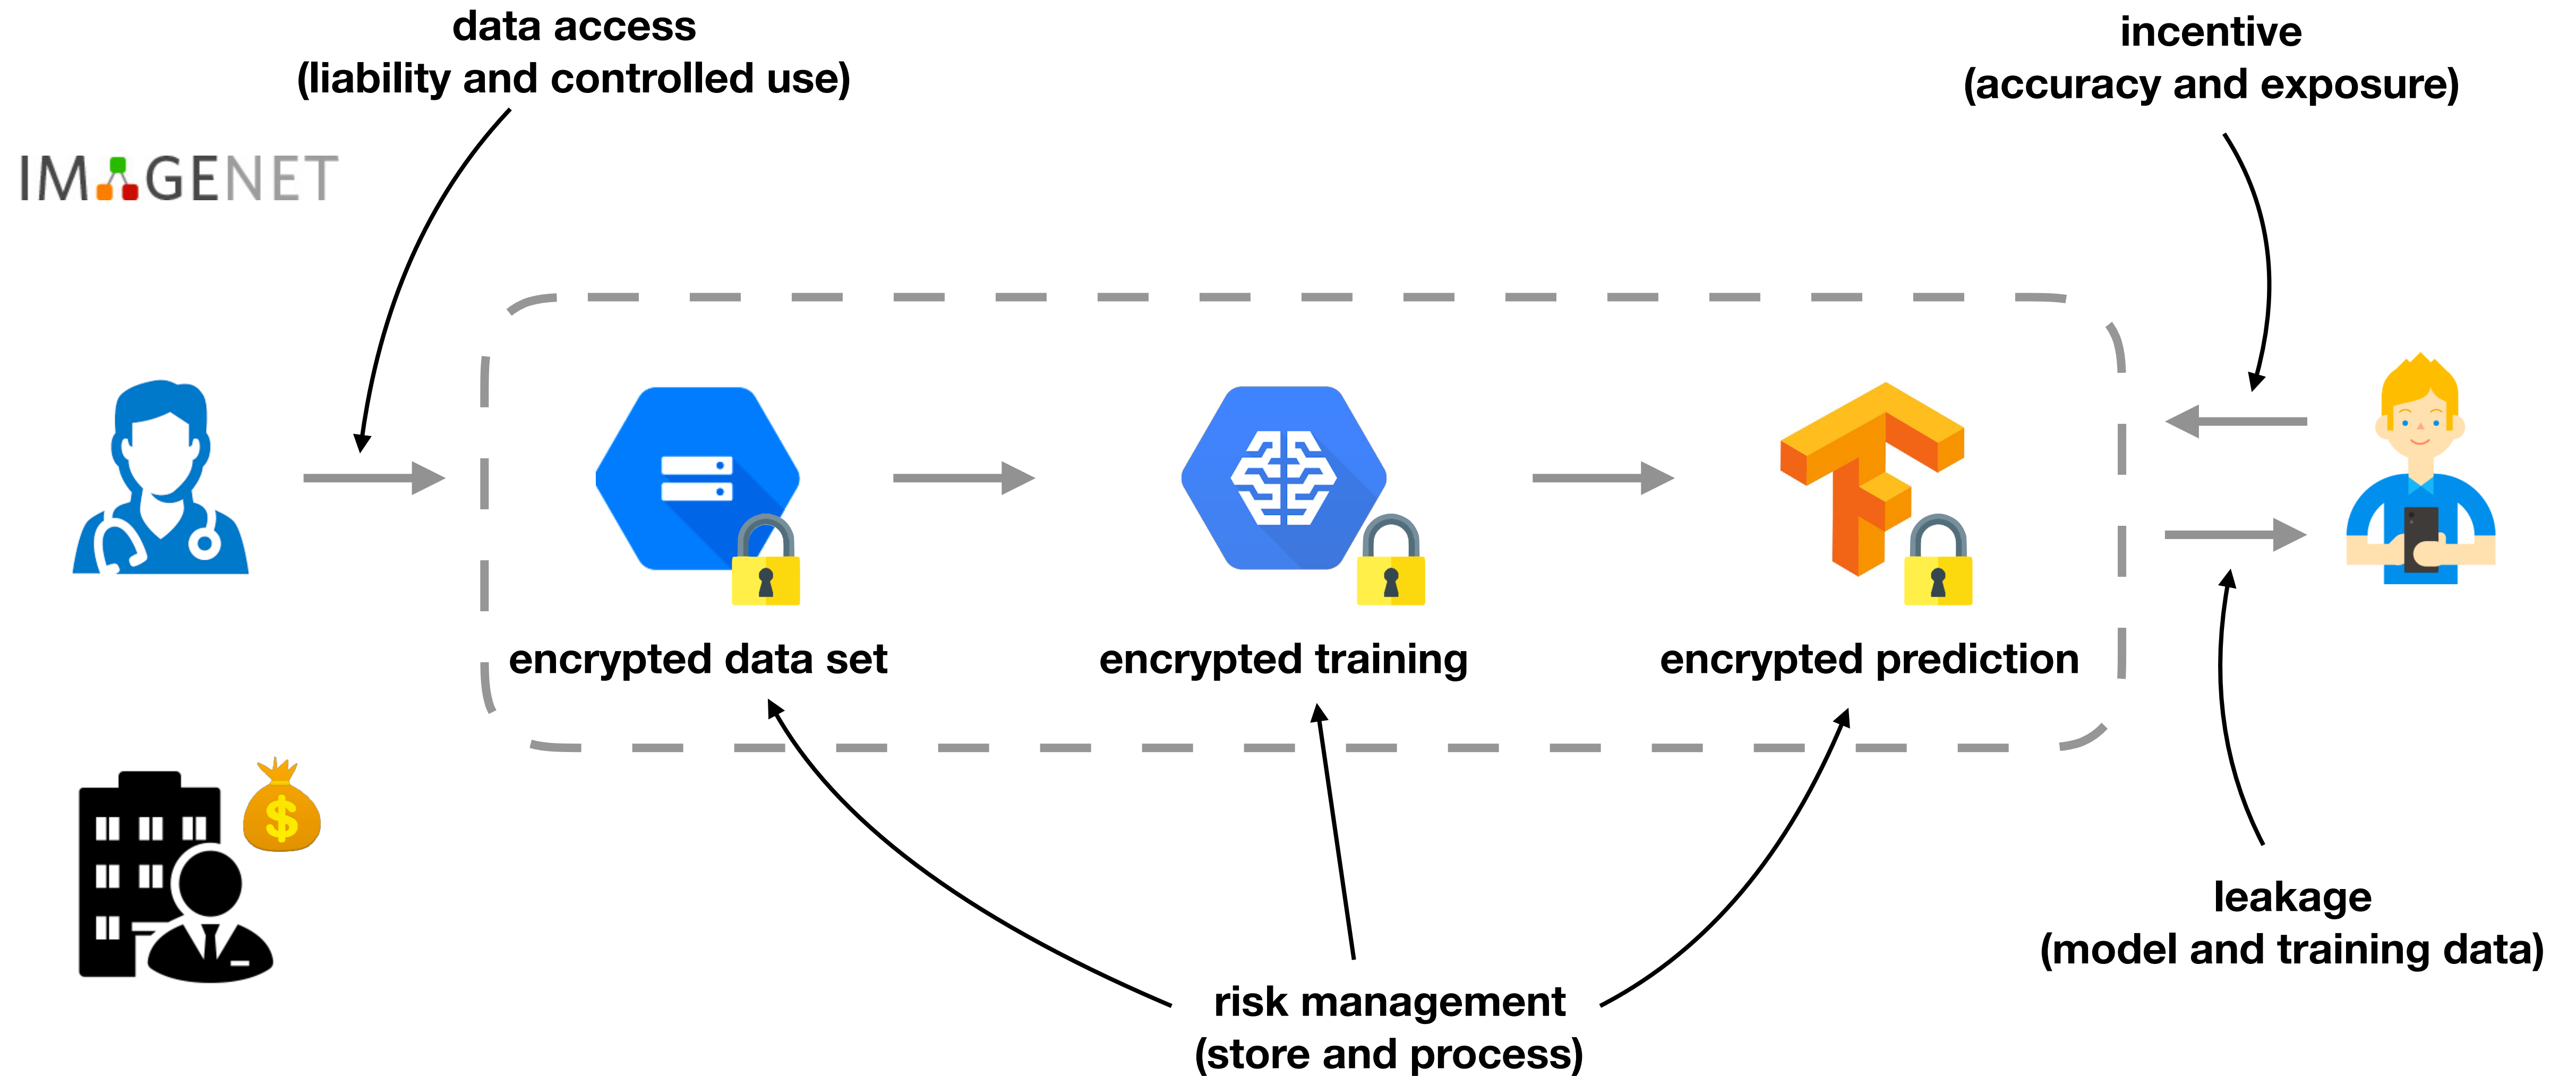
Differential Privacy



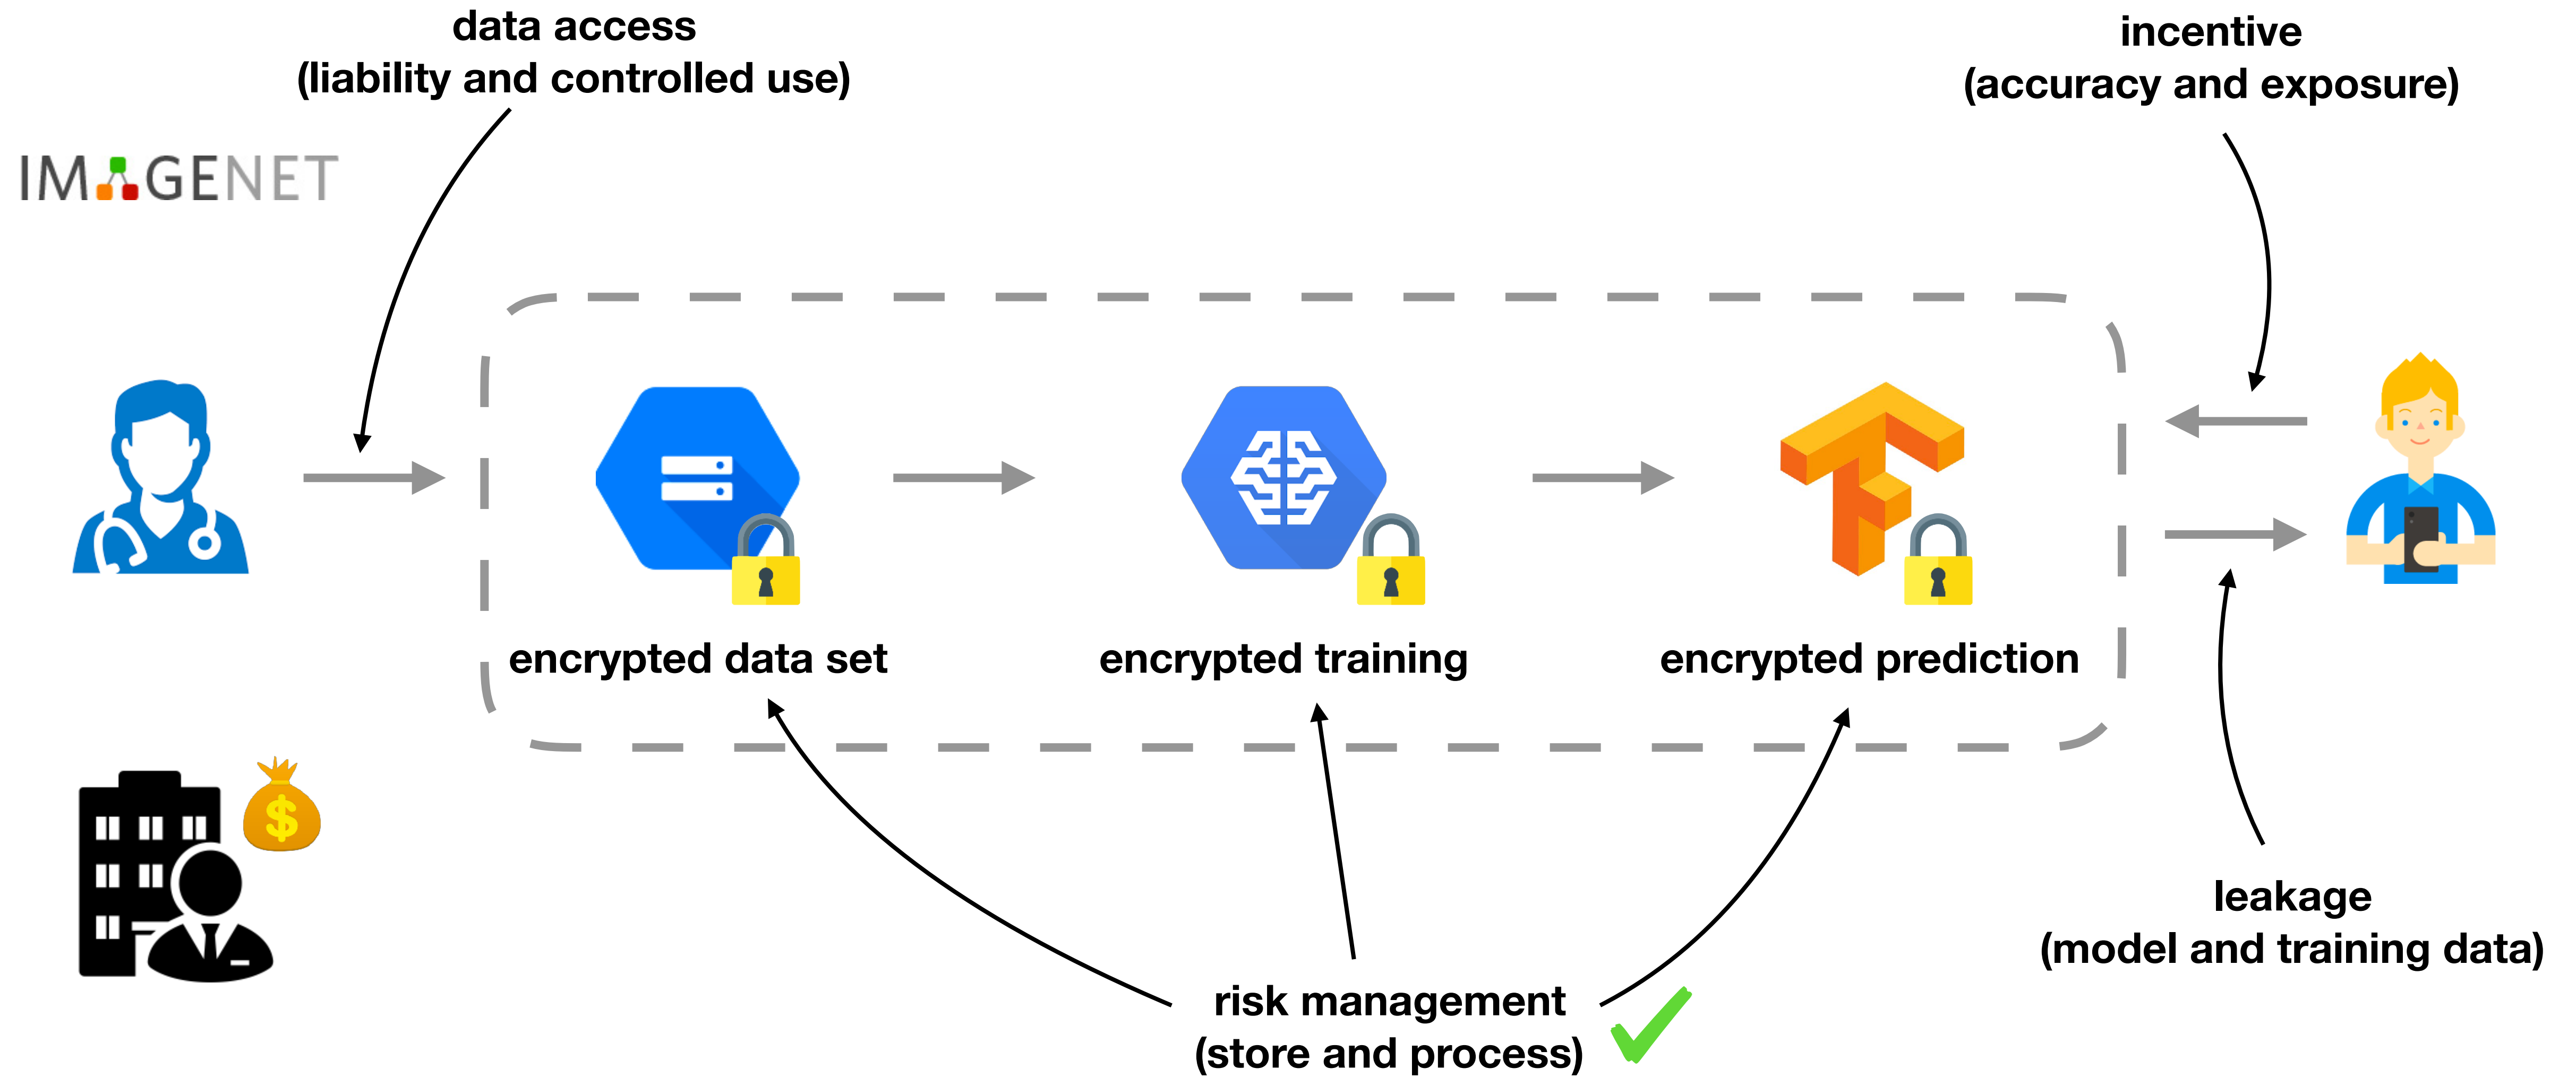
Secure Computation



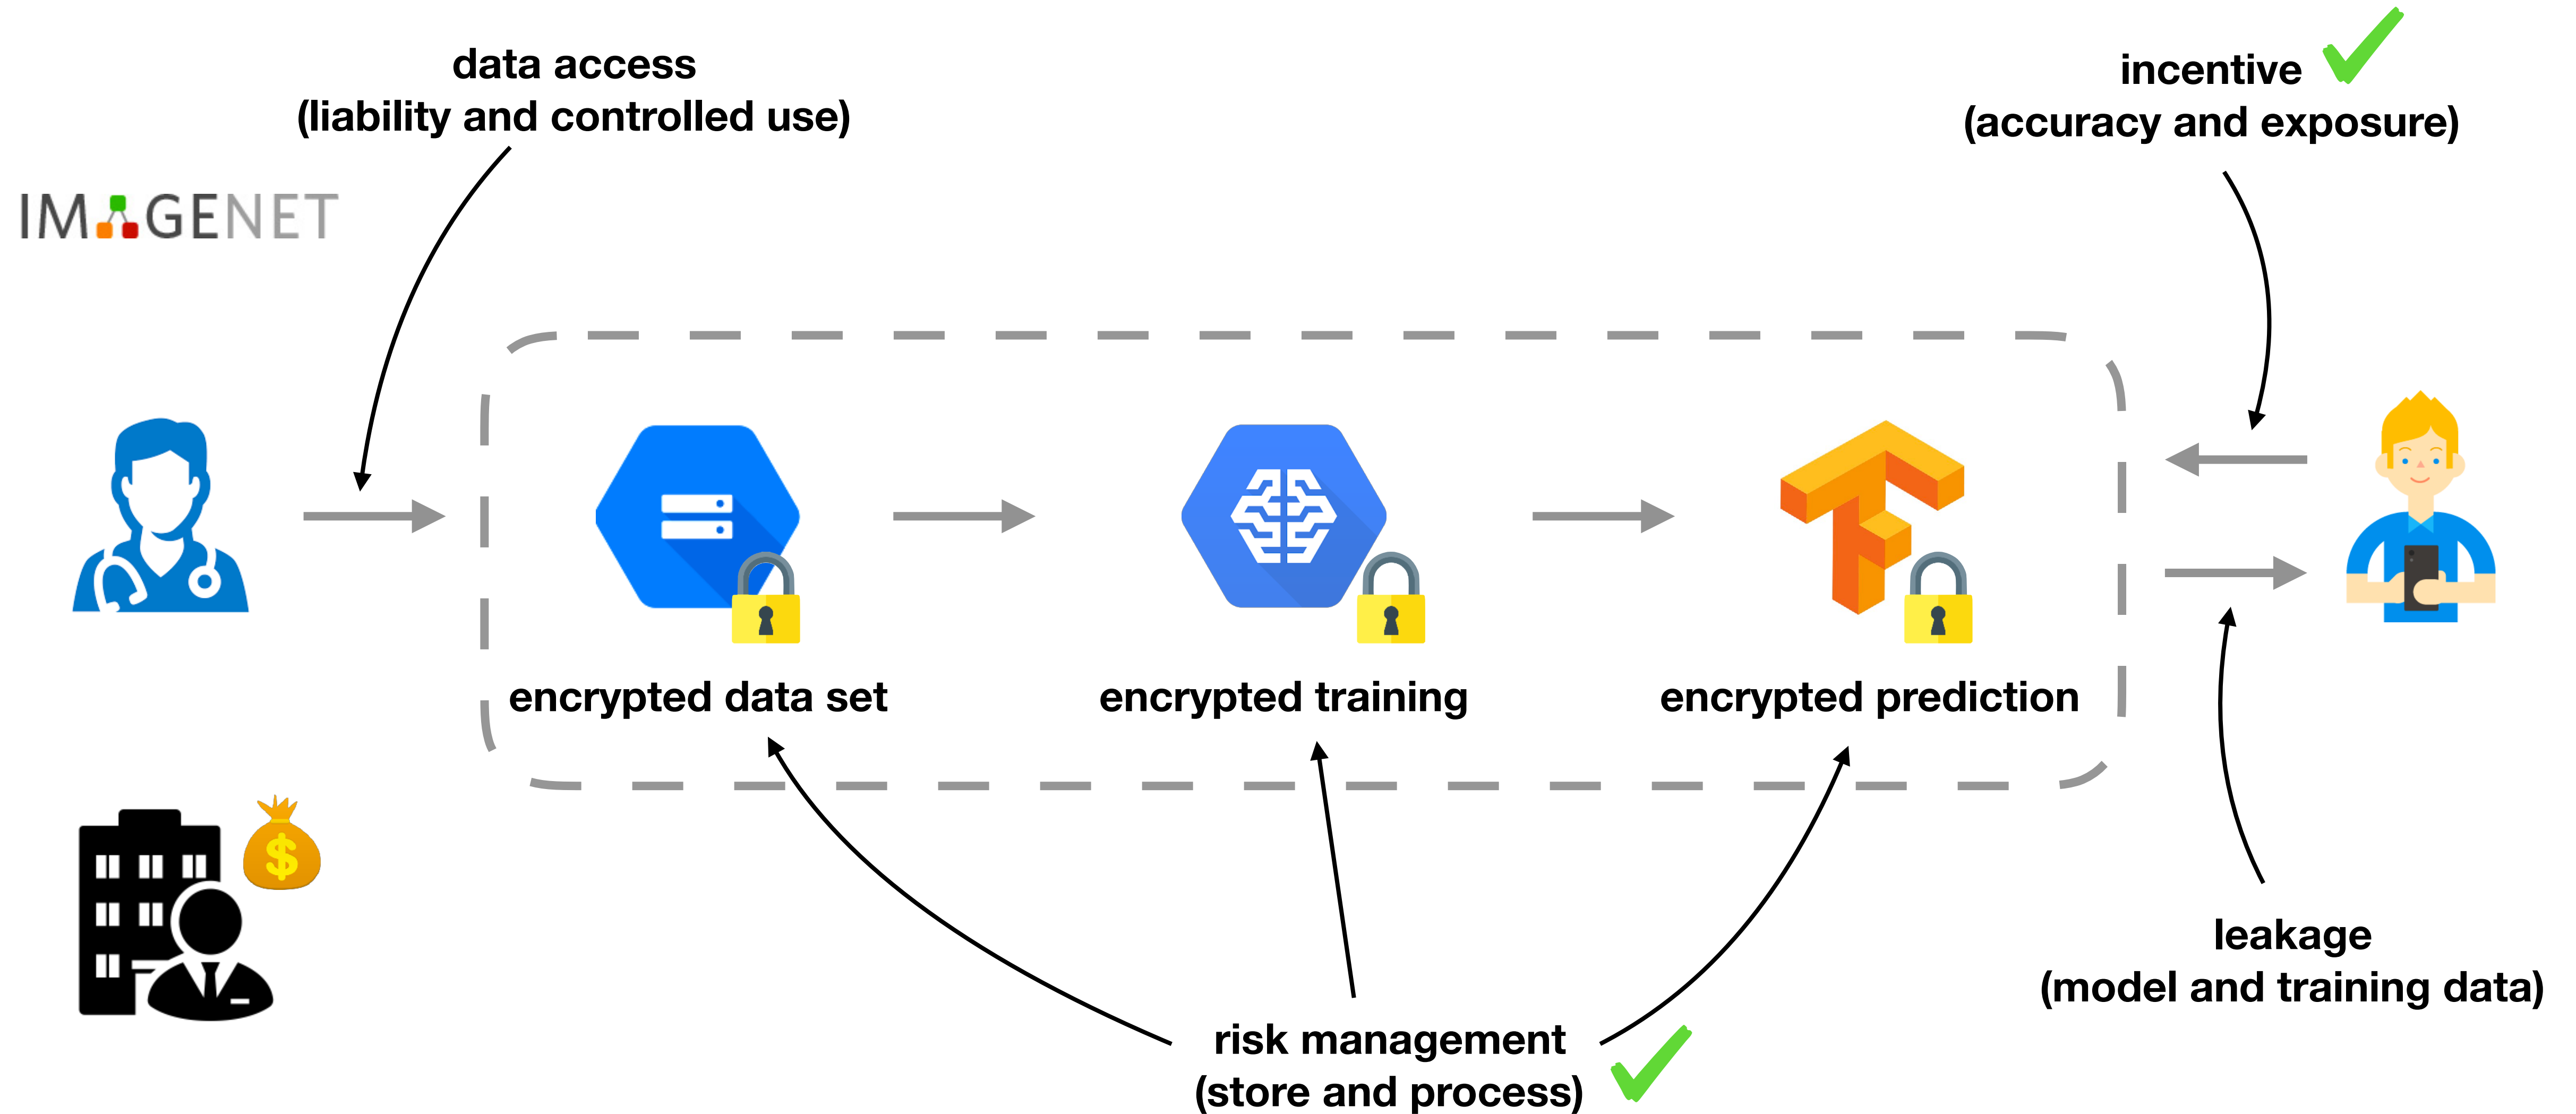
Secure Computation



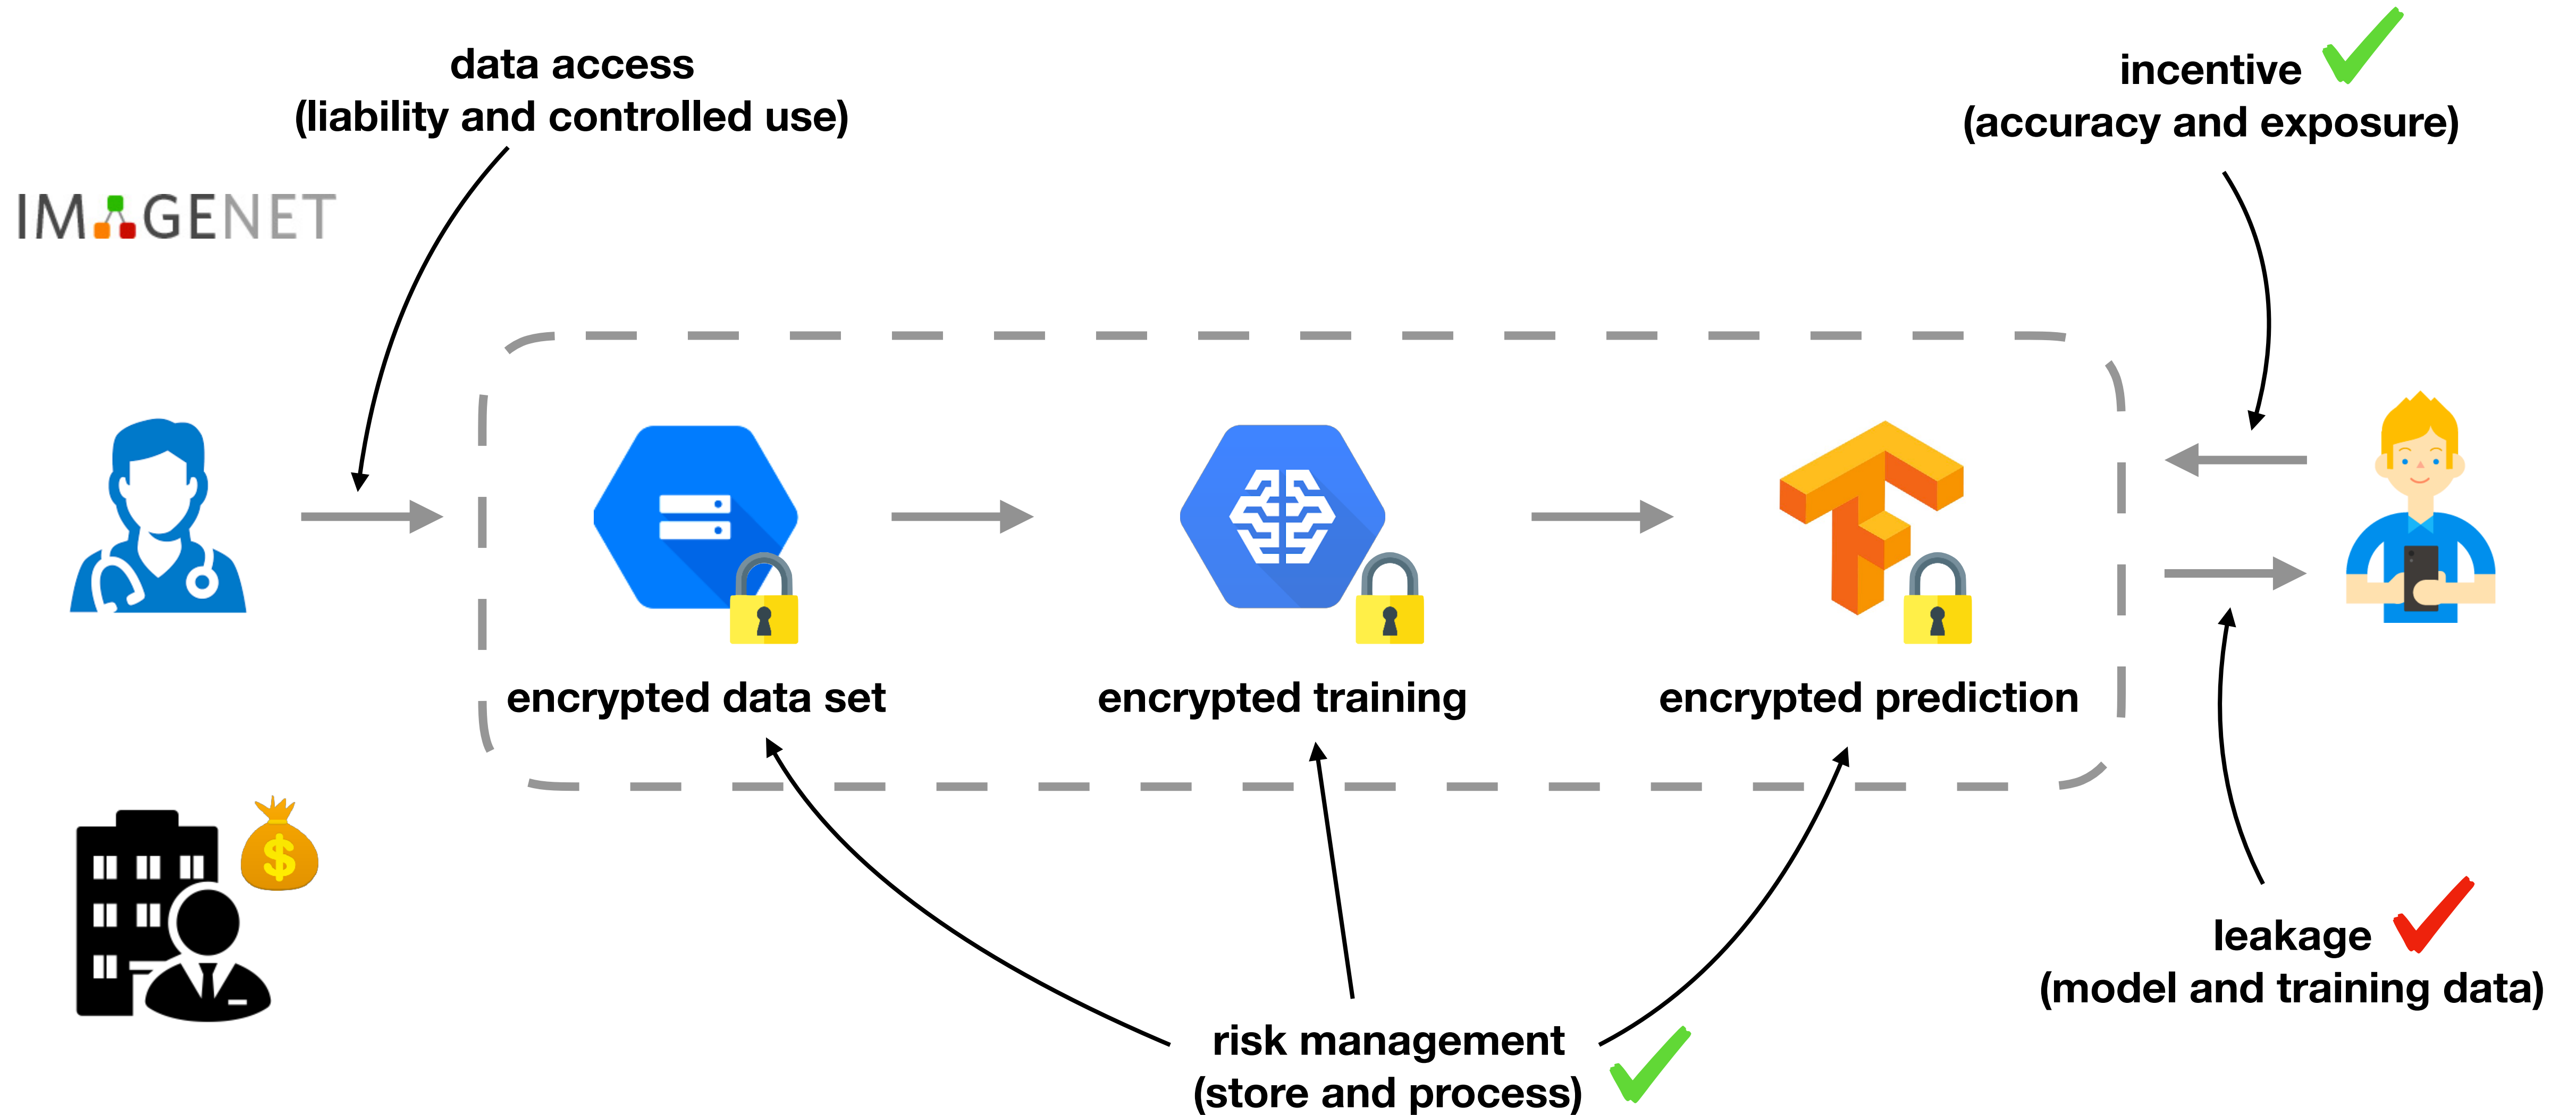
Secure Computation



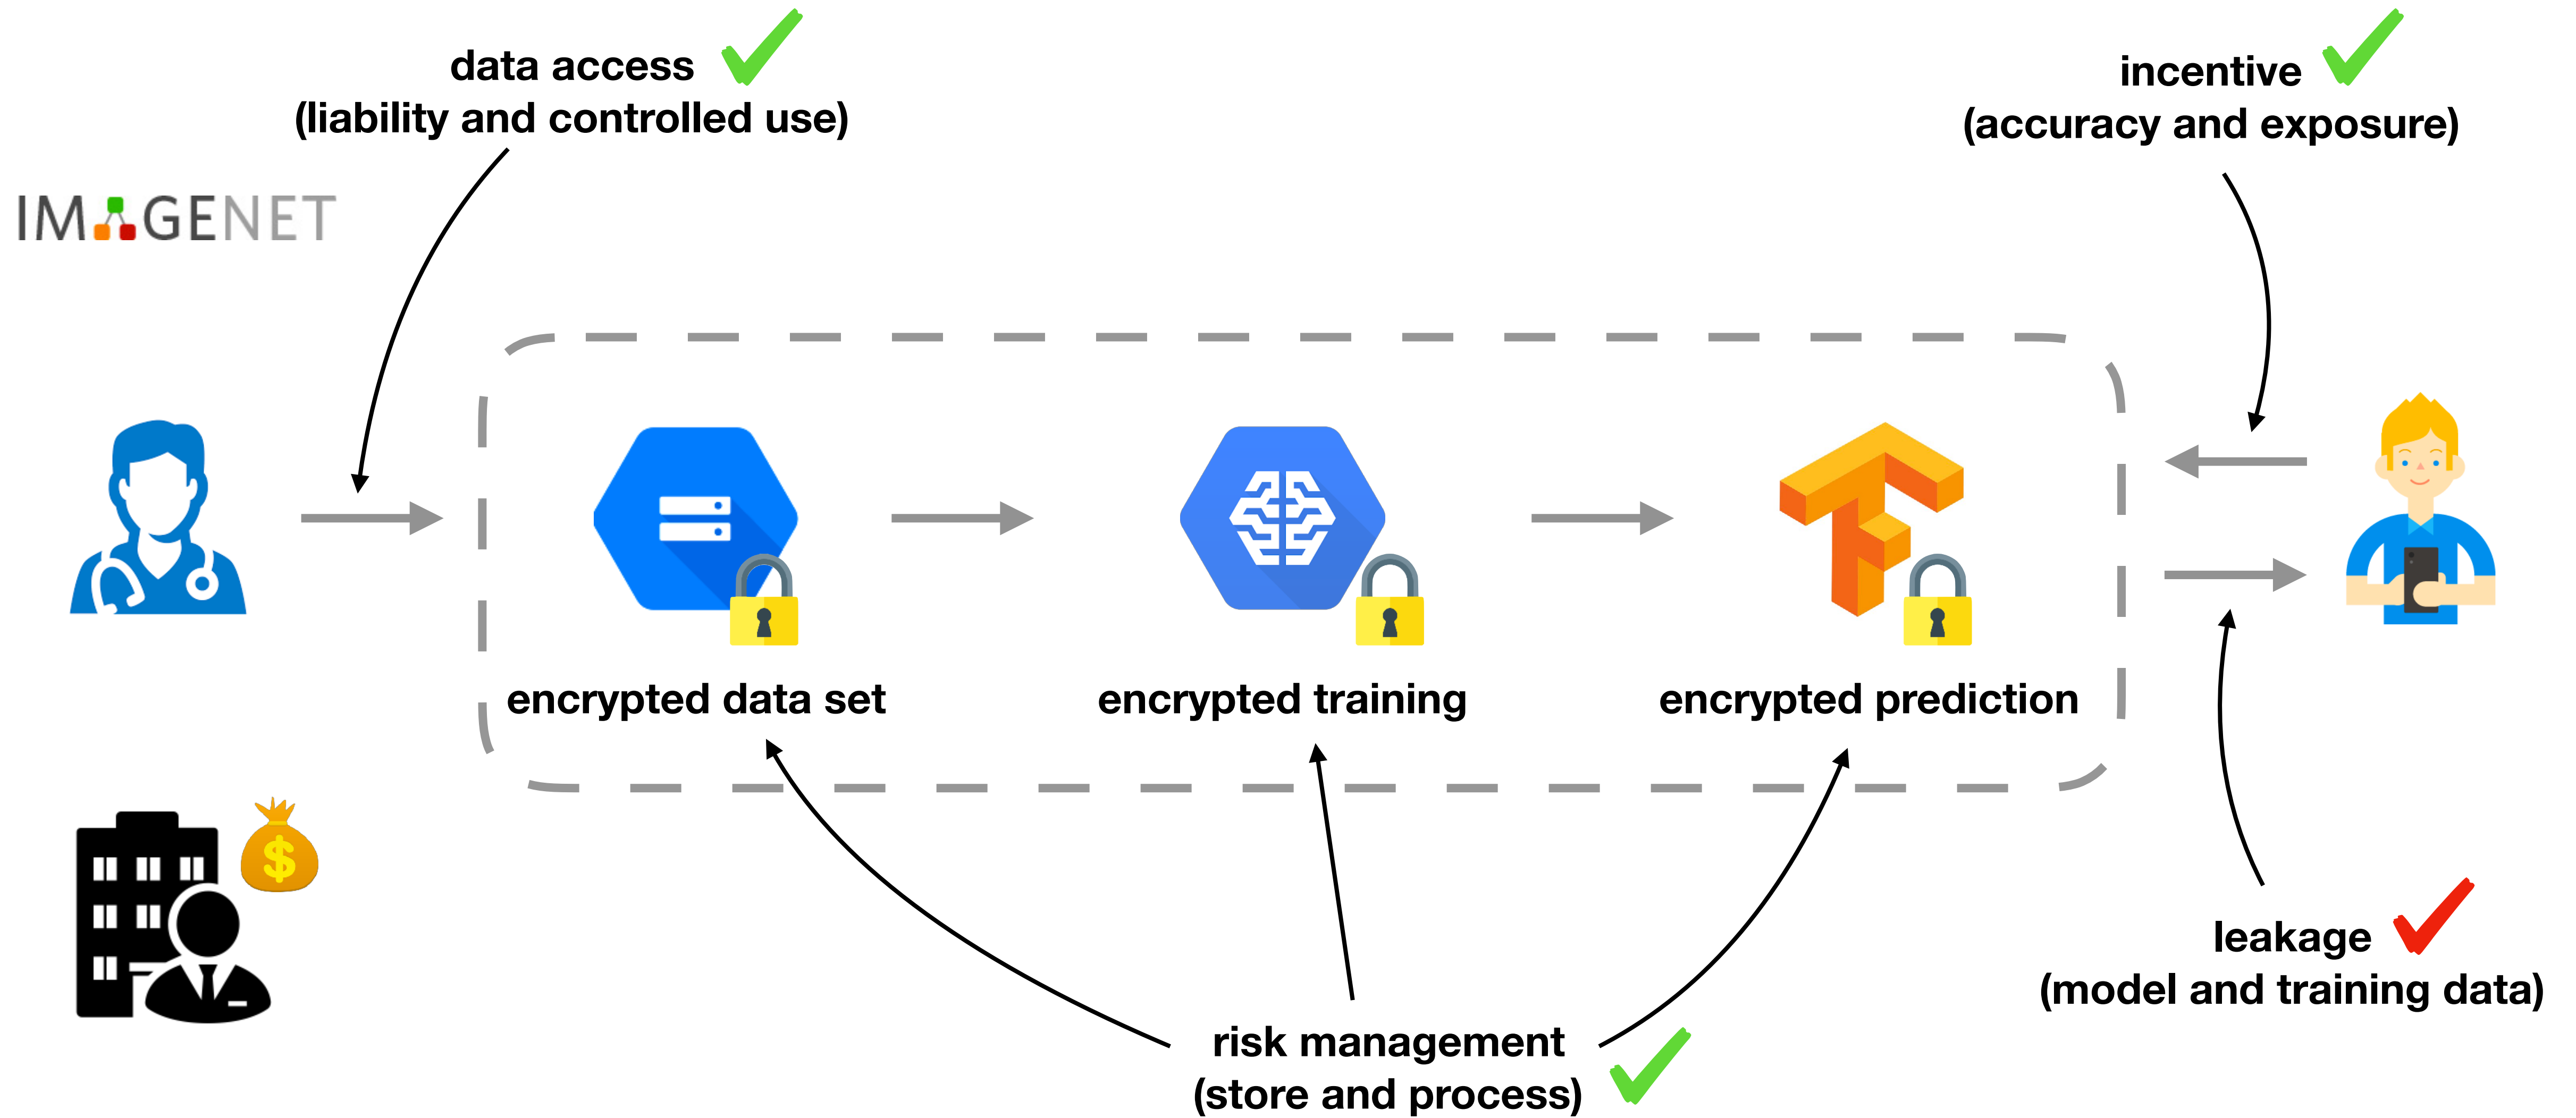
Secure Computation



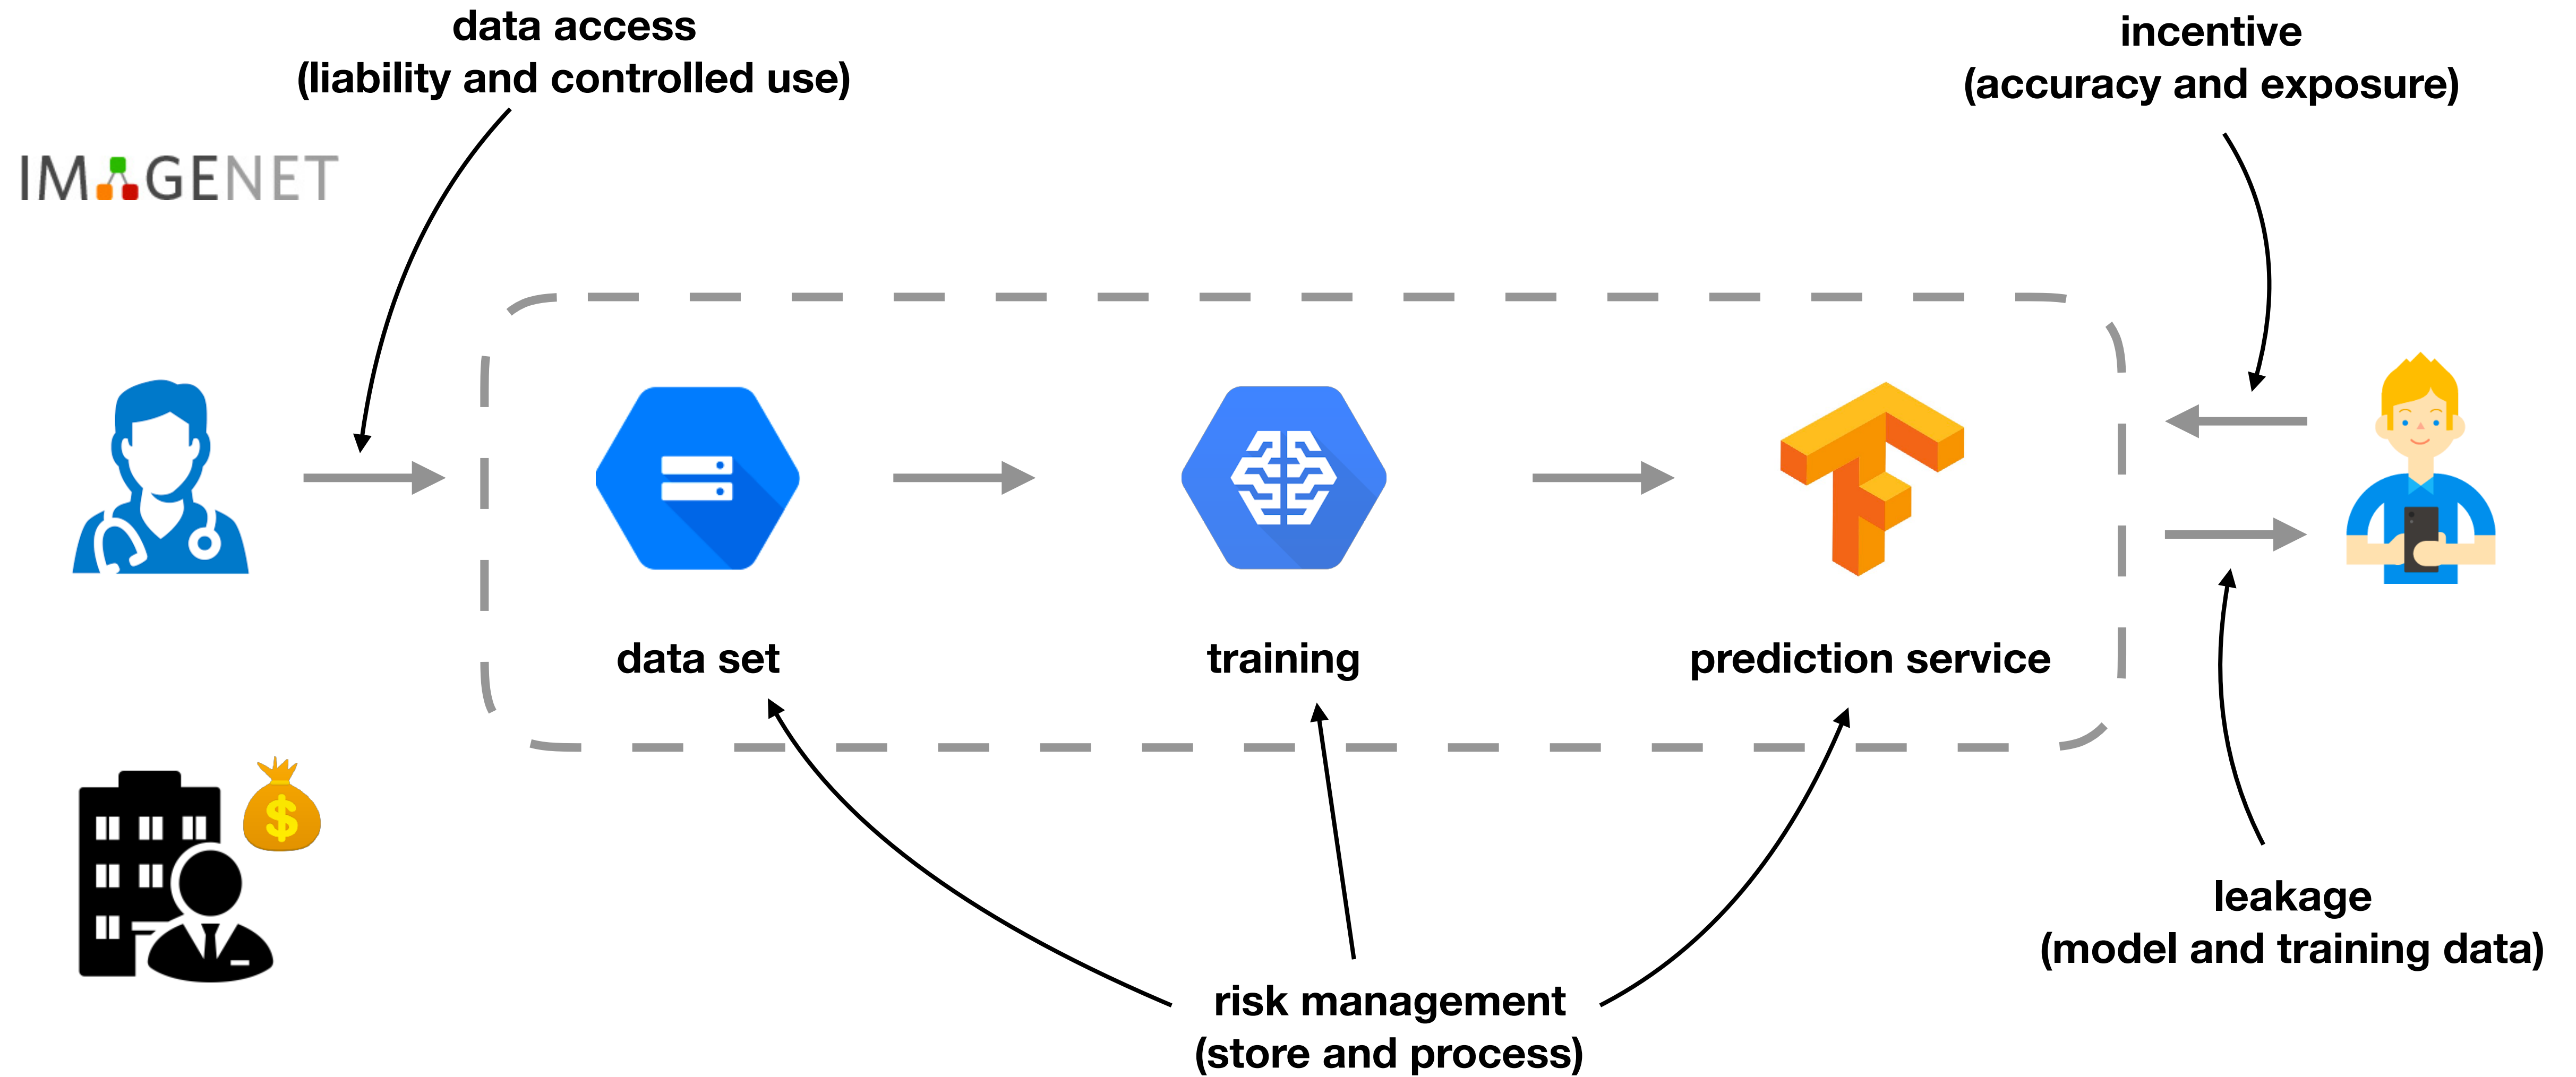
Secure Computation



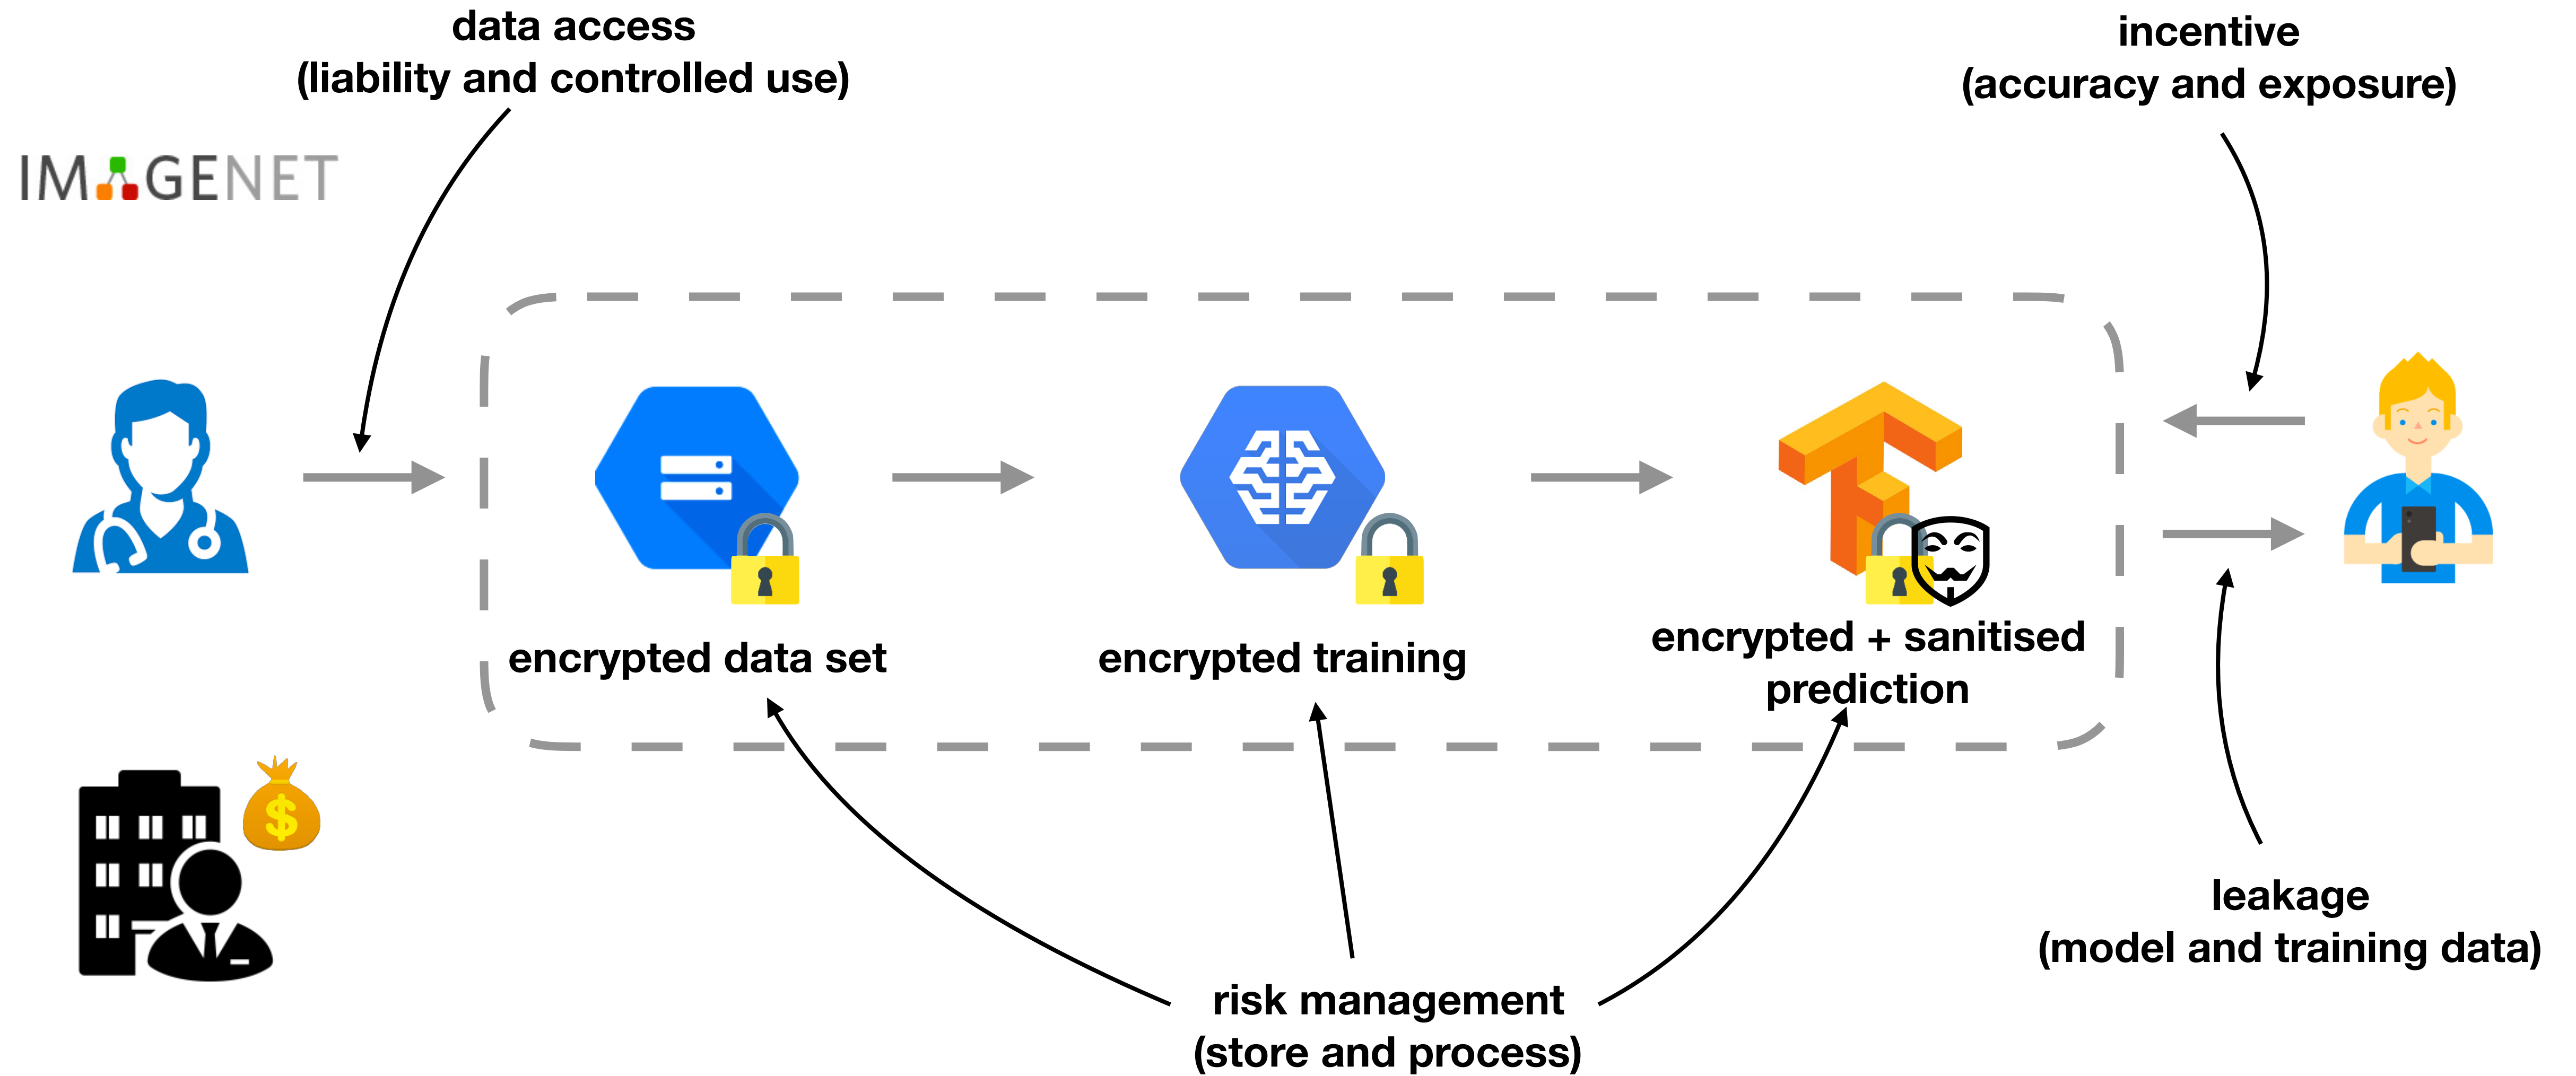
Secure Computation



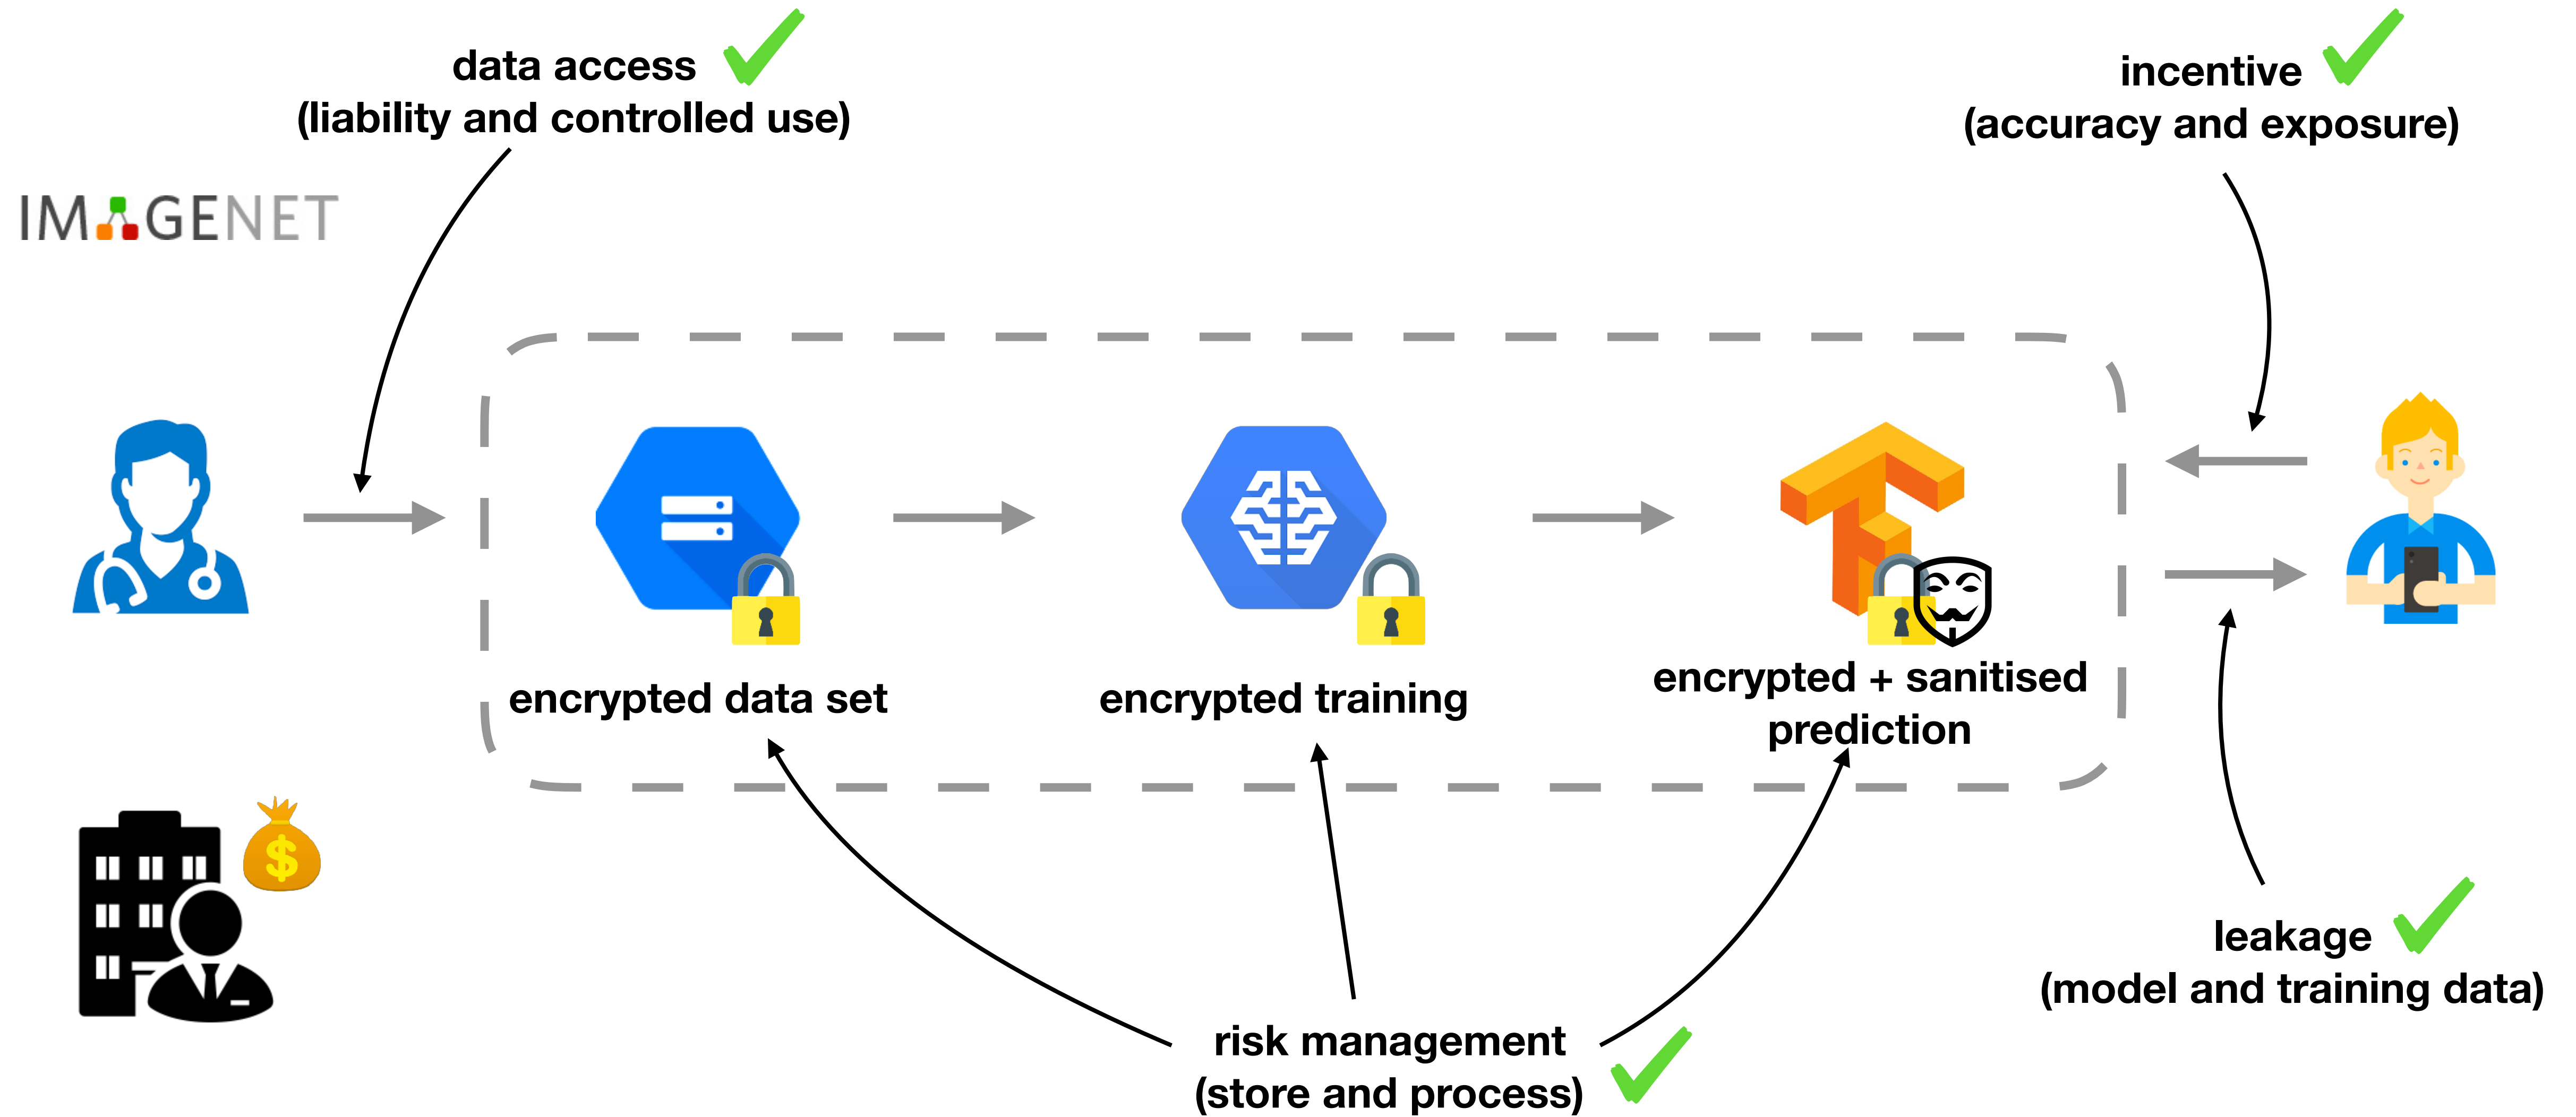
Hybrid



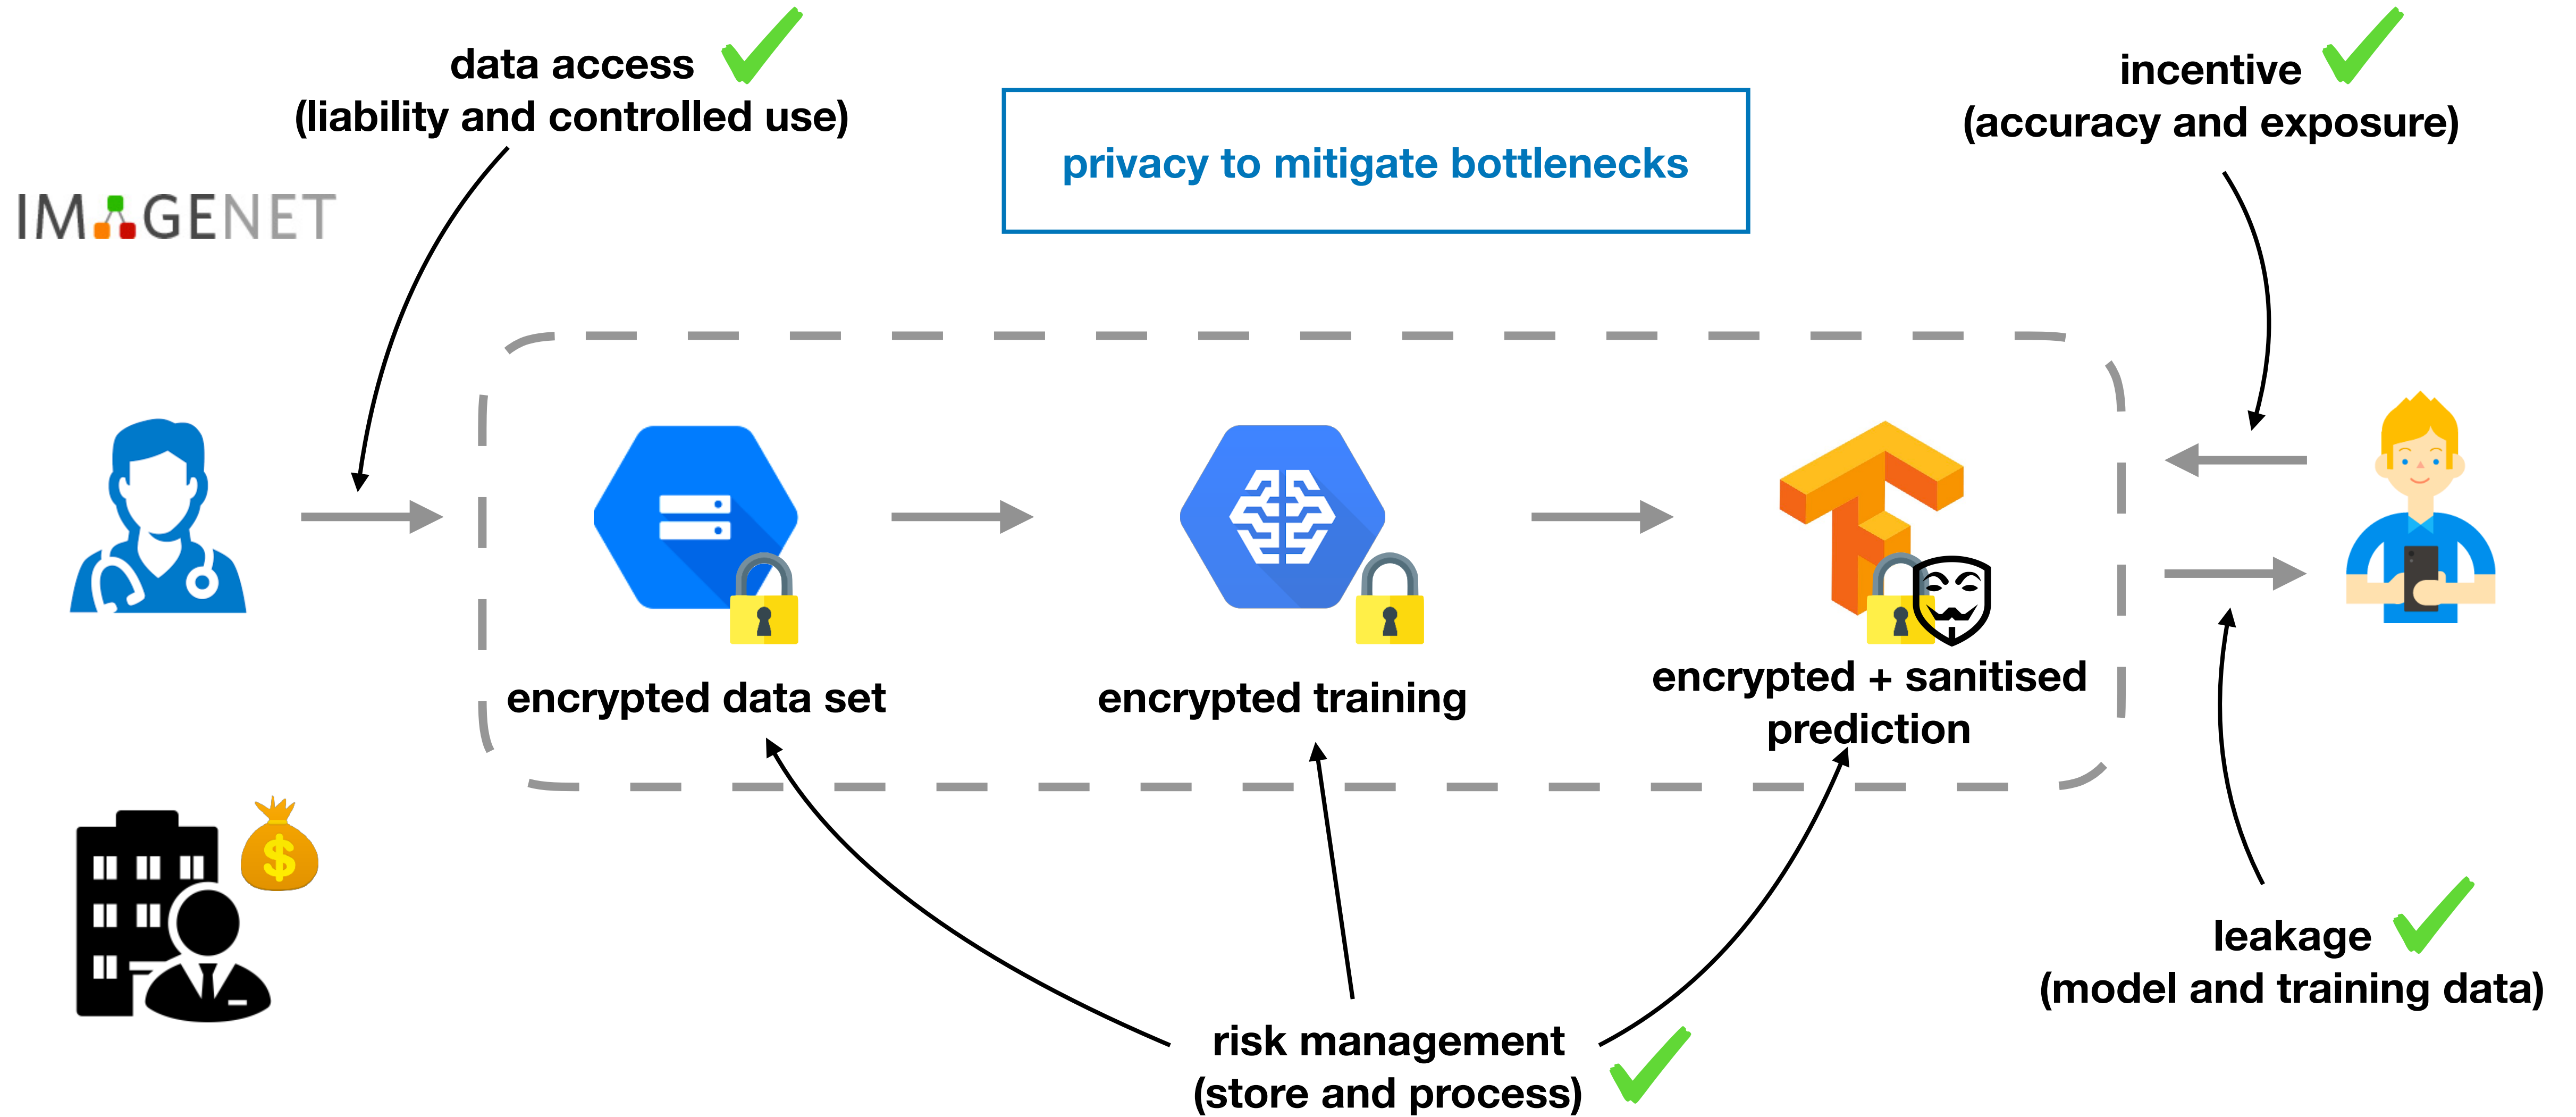
Hybrid



Hybrid



Hybrid



Prediction

Prediction with Linear Model


w




x



Prediction with Linear Model



Prediction with Linear Model



$$\text{dot}(\begin{matrix} x1 & x2 & x3 \end{matrix}, \begin{matrix} w1 \\ w2 \\ w3 \end{matrix}) = x1 * w1 + x2 * w2 + x3 * w3$$

... using Homomorphic Encryption

W



X

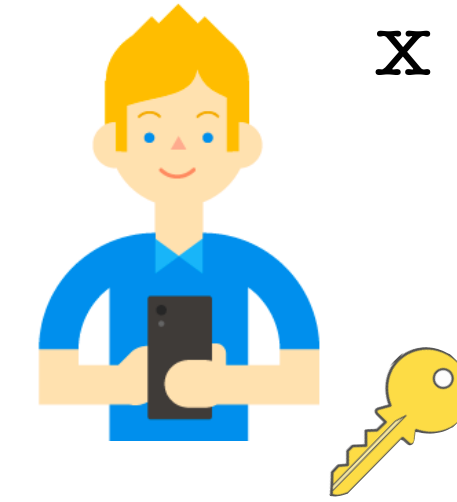


... using Homomorphic Encryption


W



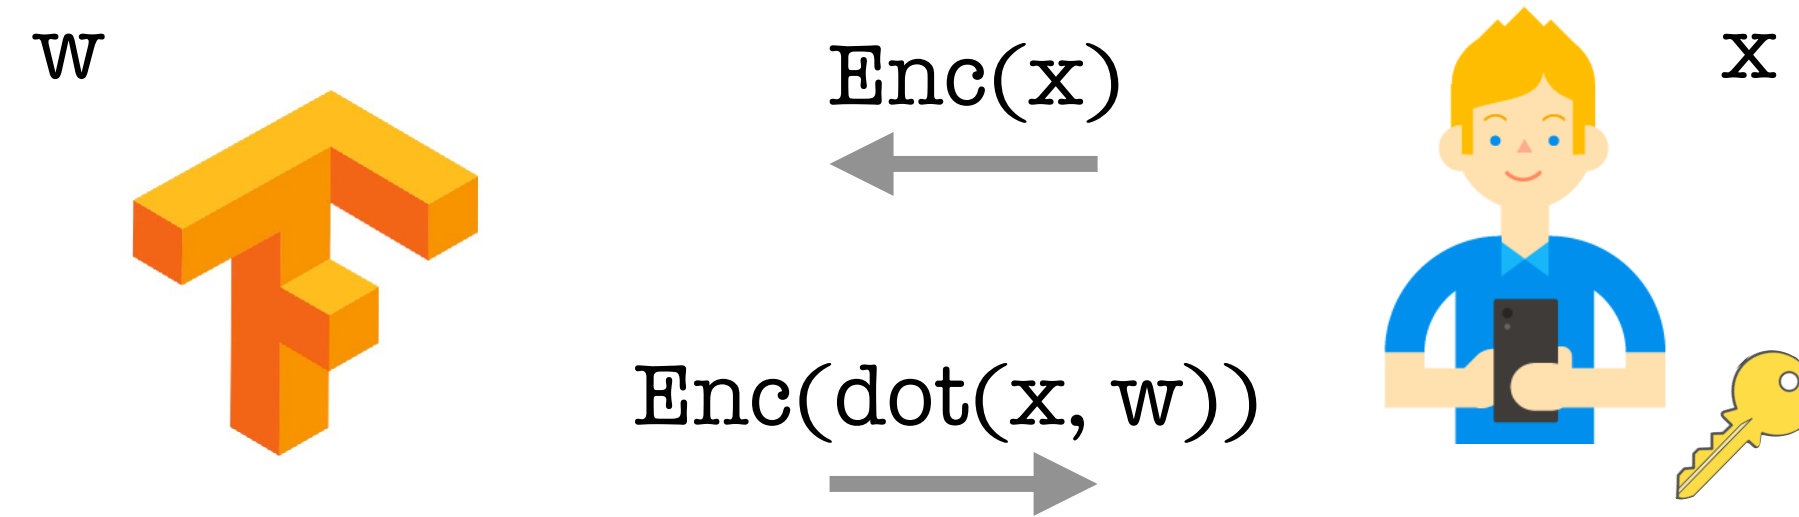
X



... using Homomorphic Encryption




... using Homomorphic Encryption



$$\text{dot}(\text{Enc}(x_1) \quad \text{Enc}(x_2) \quad \text{Enc}(x_3), \begin{matrix} w_1 \\ w_2 \\ w_3 \end{matrix}) = \text{Enc}(x_1) * w_1 + \text{Enc}(x_2) * w_2 + \text{Enc}(x_3) * w_3$$

... using Homomorphic Encryption




public multiplication

$$\text{dot}(\text{Enc}(x_1) \quad \text{Enc}(x_2) \quad \text{Enc}(x_3), \begin{matrix} w_1 \\ w_2 \\ w_3 \end{matrix}) = \text{Enc}(x_1) * w_1 + \text{Enc}(x_2) * w_2 + \text{Enc}(x_3) * w_3$$
$$= \text{Enc}(x_1 * w_1) + \text{Enc}(x_2 * w_2) + \text{Enc}(x_3 * w_3)$$

An arrow points from the text "public multiplication" to the first equation above.

... using Homomorphic Encryption



$$\begin{aligned} \text{dot}(\text{Enc}(x_1) \quad \text{Enc}(x_2) \quad \text{Enc}(x_3), \begin{matrix} w_1 \\ w_2 \\ w_3 \end{matrix}) &= \text{Enc}(x_1) * w_1 + \text{Enc}(x_2) * w_2 + \text{Enc}(x_3) * w_3 \\ &= \text{Enc}(x_1 * w_1) + \text{Enc}(x_2 * w_2) + \text{Enc}(x_3 * w_3) \\ &= \text{Enc}(x_1 * w_1 + x_2 * w_2 + x_3 * w_3) \end{aligned}$$

The diagram includes two arrows pointing to the equations: "public multiplication" points to the first equation, and "private addition" points to the second equation.


... using Homomorphic Encryption



Enc(x)



x



Enc(dot(x, w))



homomorphic encryption scheme

public multiplication

private addition

$$\begin{aligned}
 \text{dot}(\text{Enc}(x_1) \quad \text{Enc}(x_2) \quad \text{Enc}(x_3) , \begin{matrix} w_1 \\ w_2 \\ w_3 \end{matrix}) &= \text{Enc}(x_1) * w_1 + \text{Enc}(x_2) * w_2 + \text{Enc}(x_3) * w_3 \\
 &= \text{Enc}(x_1 * w_1) + \text{Enc}(x_2 * w_2) + \text{Enc}(x_3 * w_3) \\
 &= \text{Enc}(x_1 * w_1 + x_2 * w_2 + x_3 * w_3)
 \end{aligned}$$

... using Secret Sharing

w



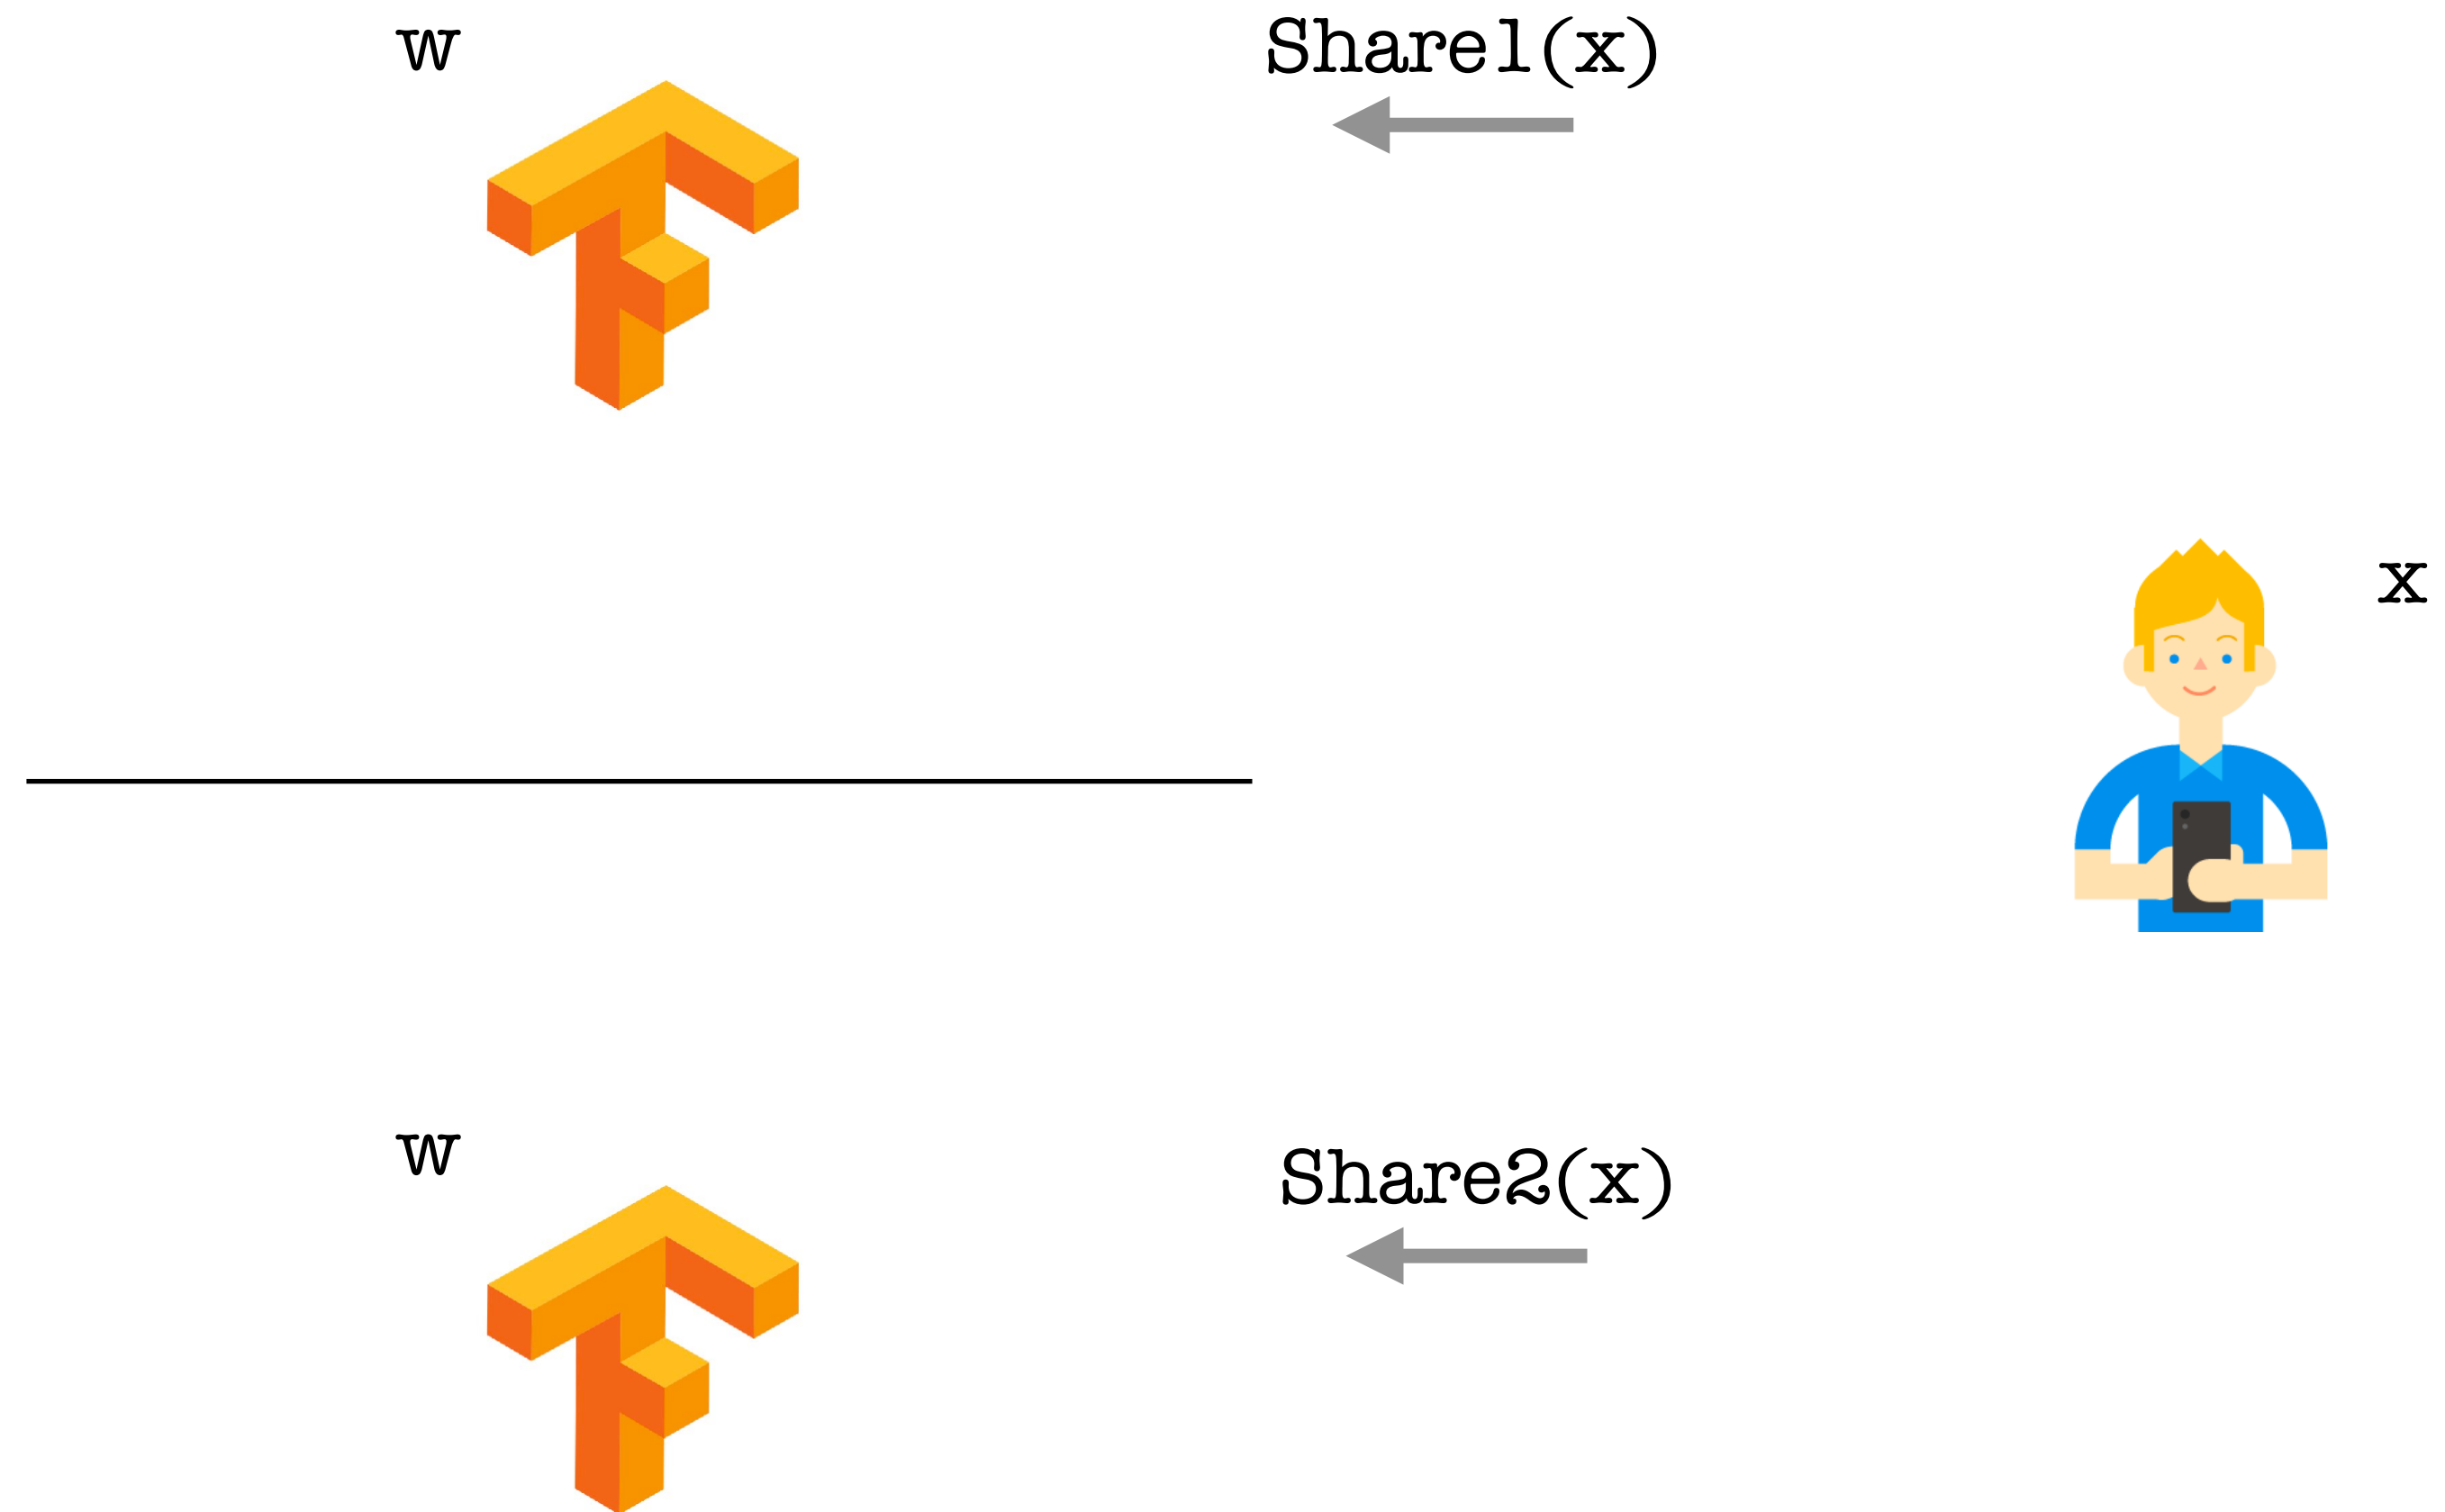
x



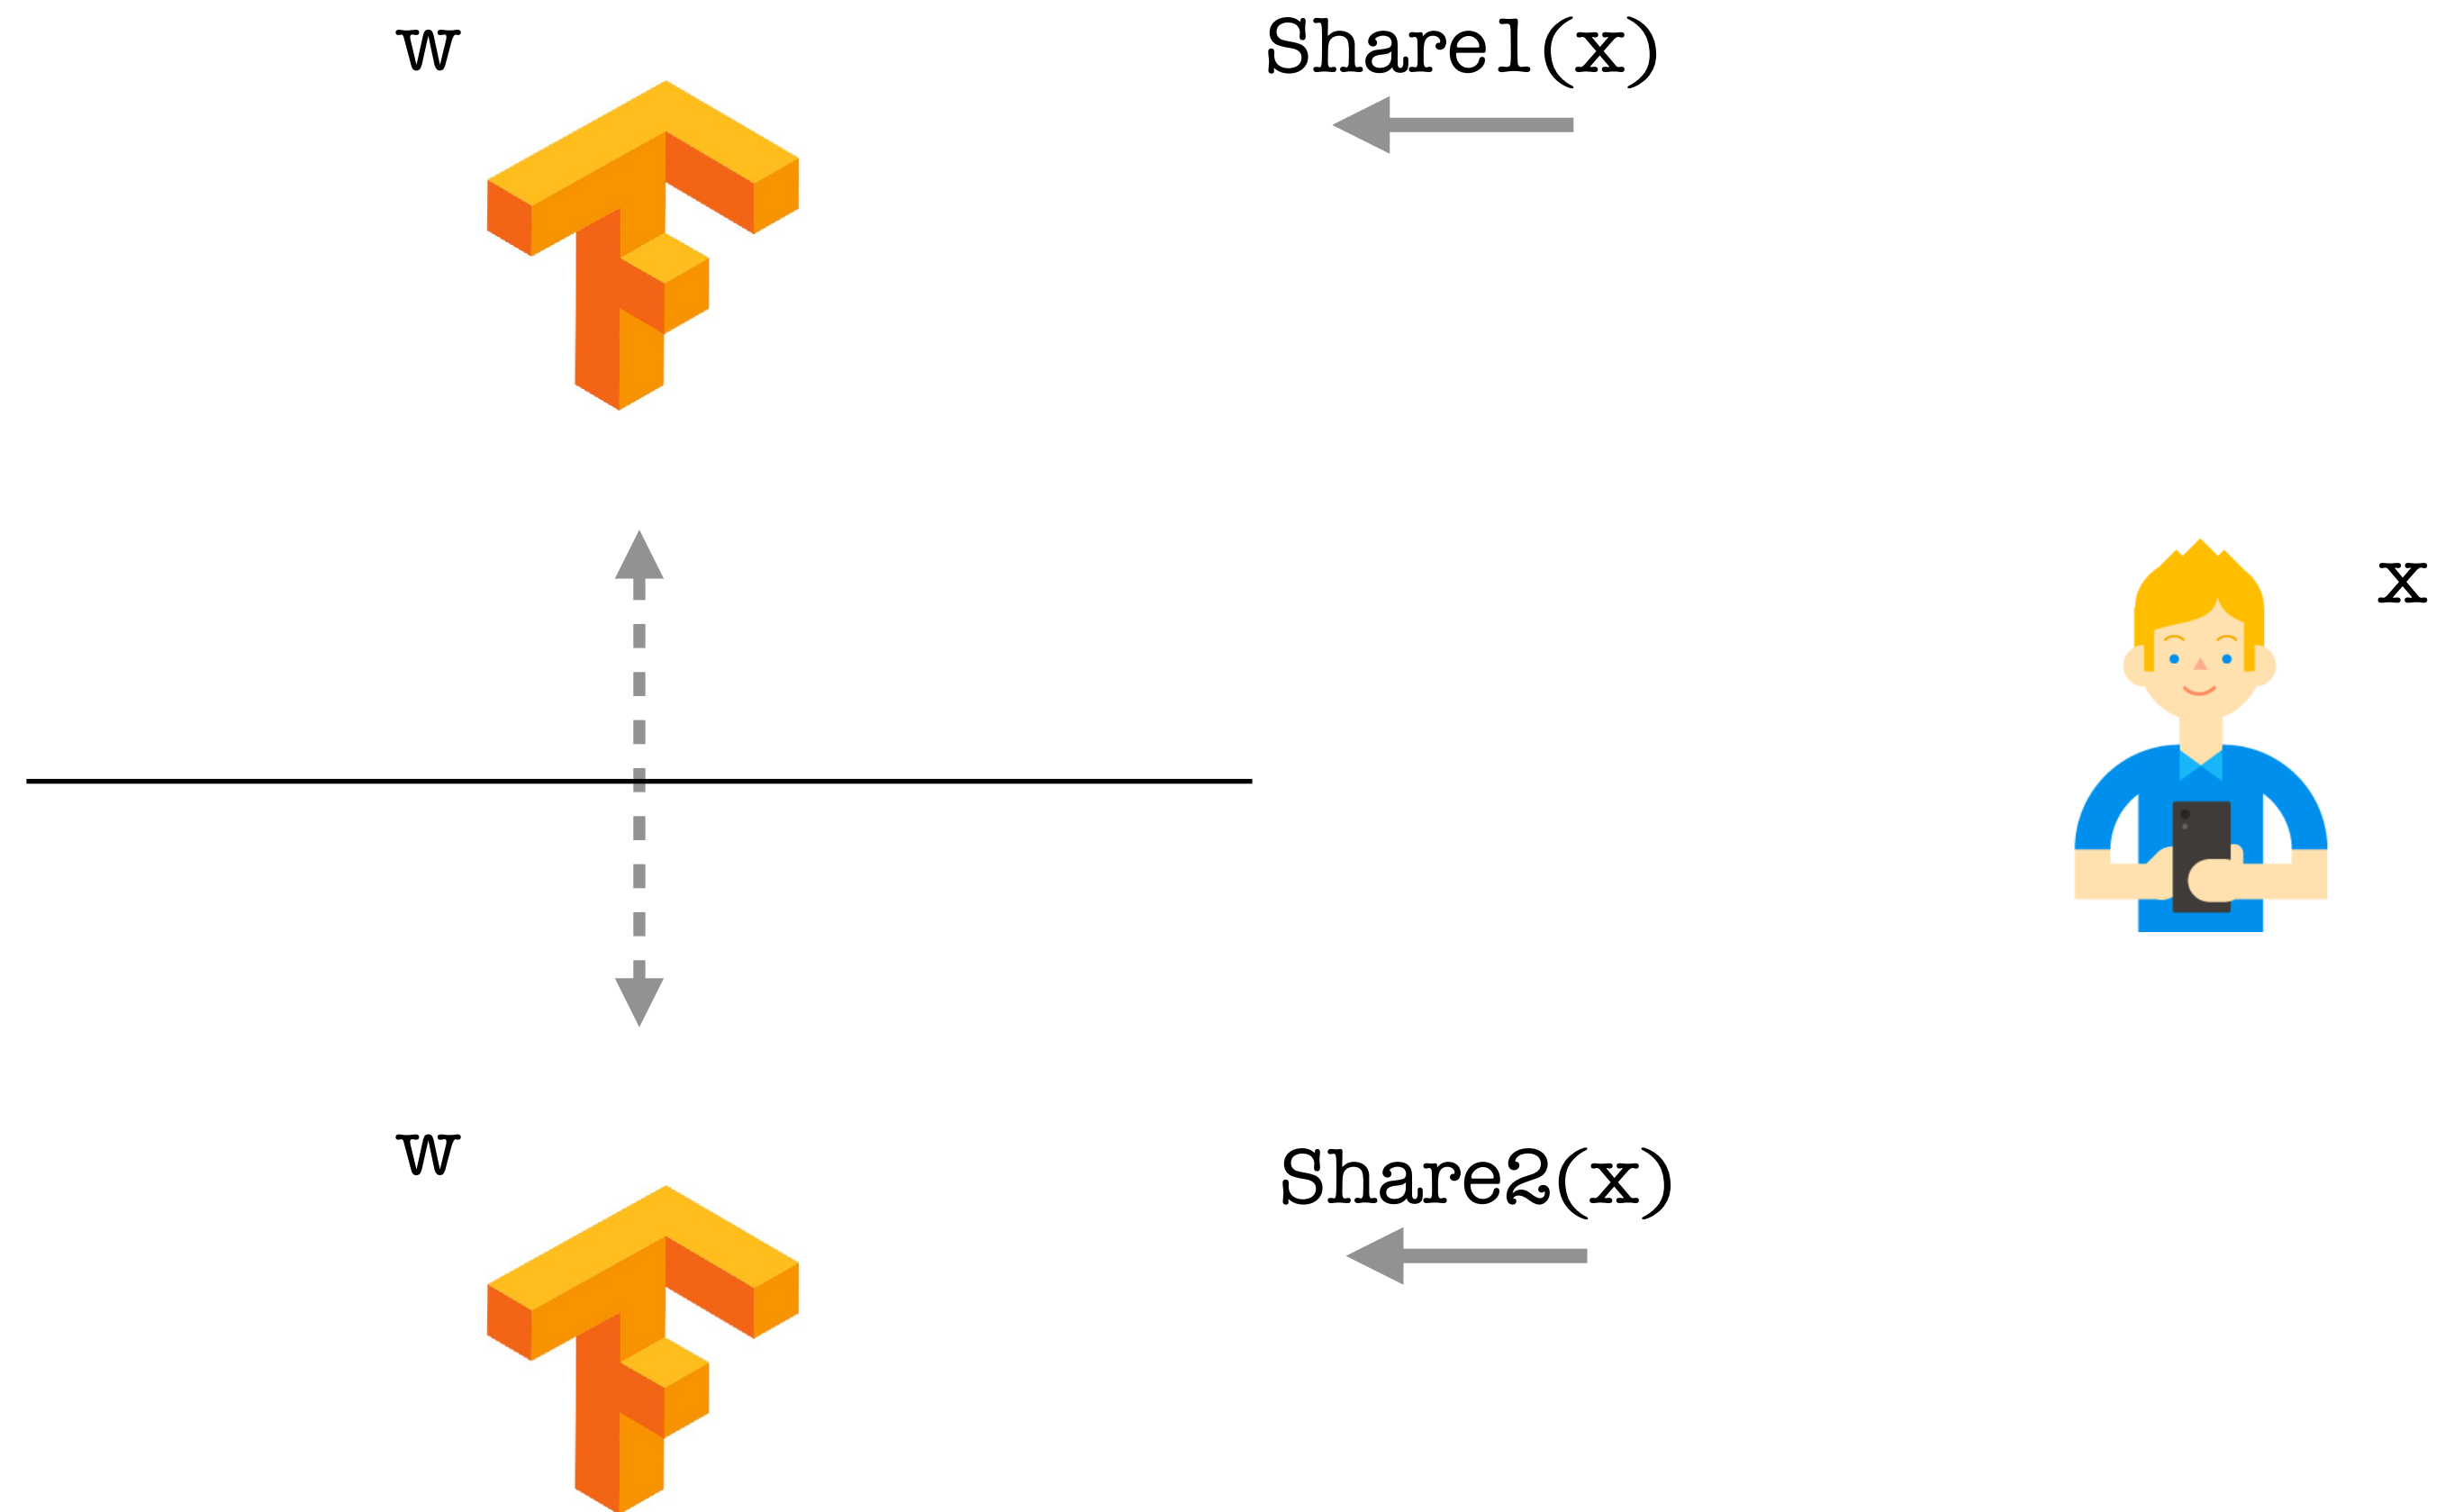
w




... using Secret Sharing



... using Secret Sharing



... using Secret Sharing



... using Secret Sharing

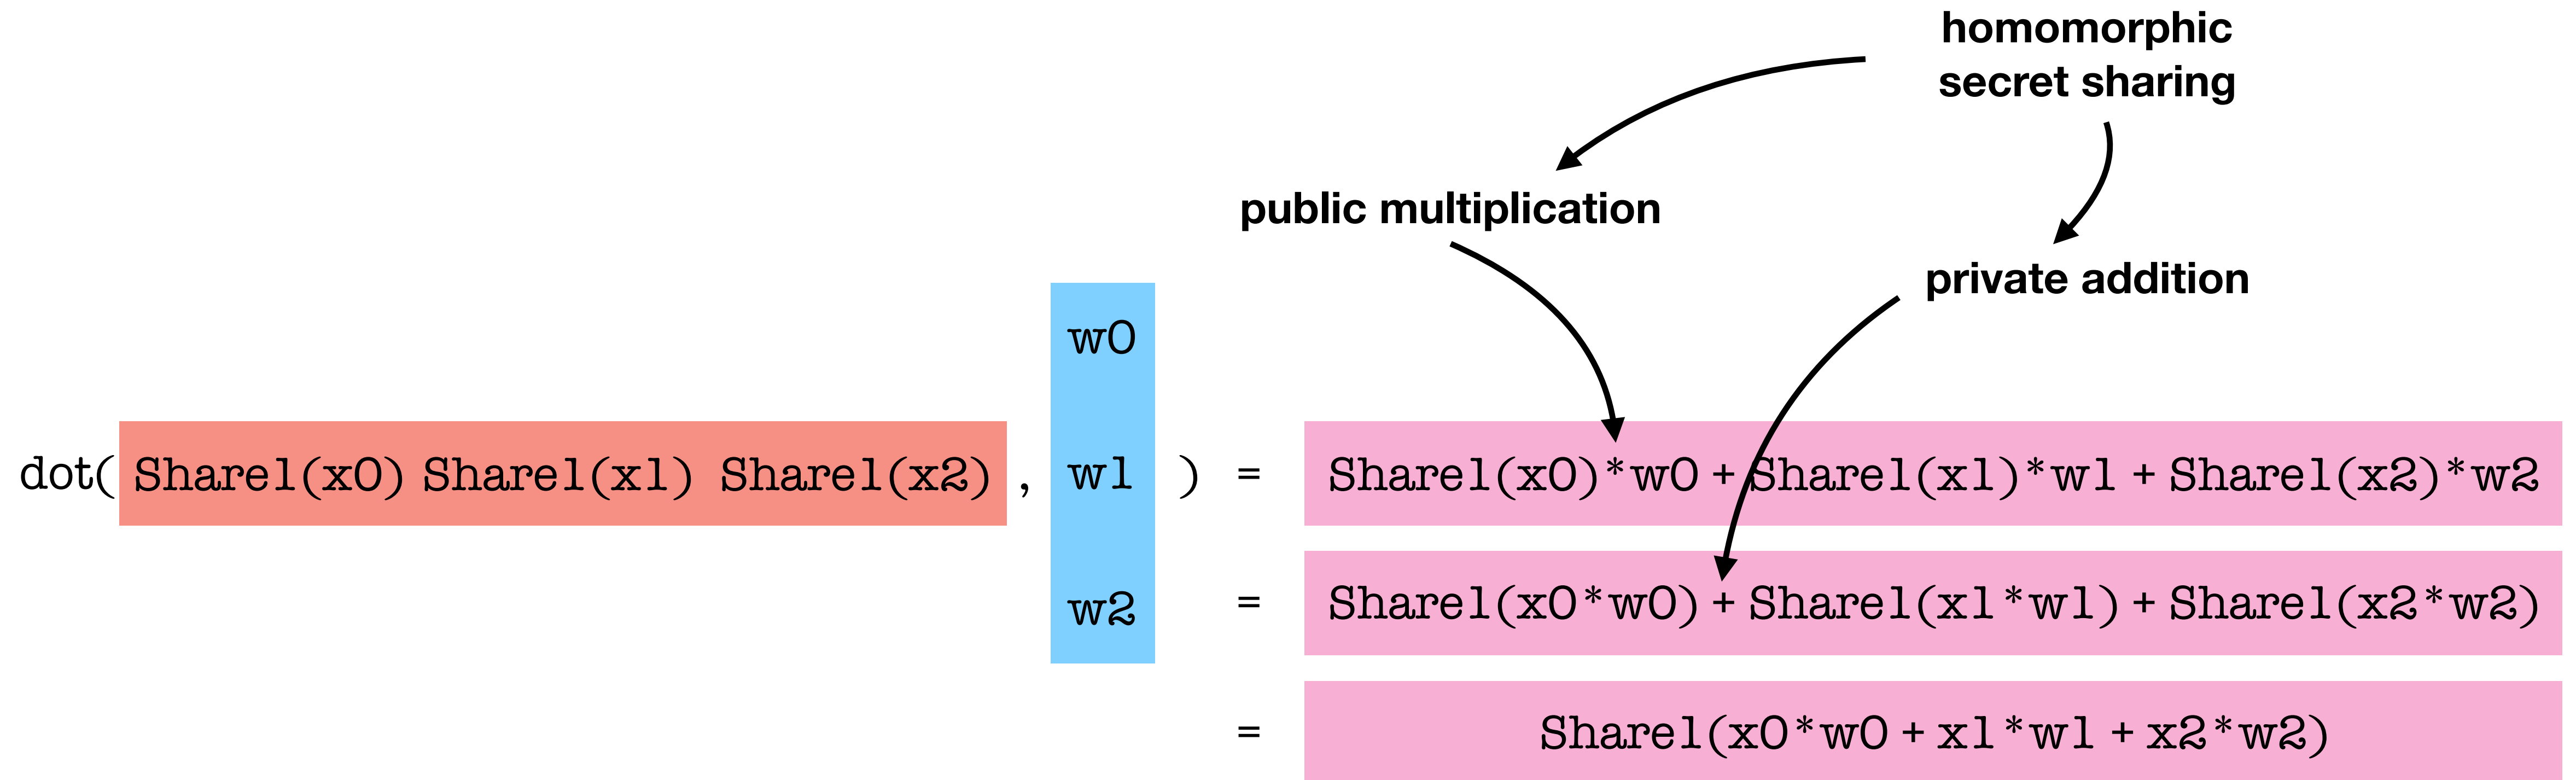
w



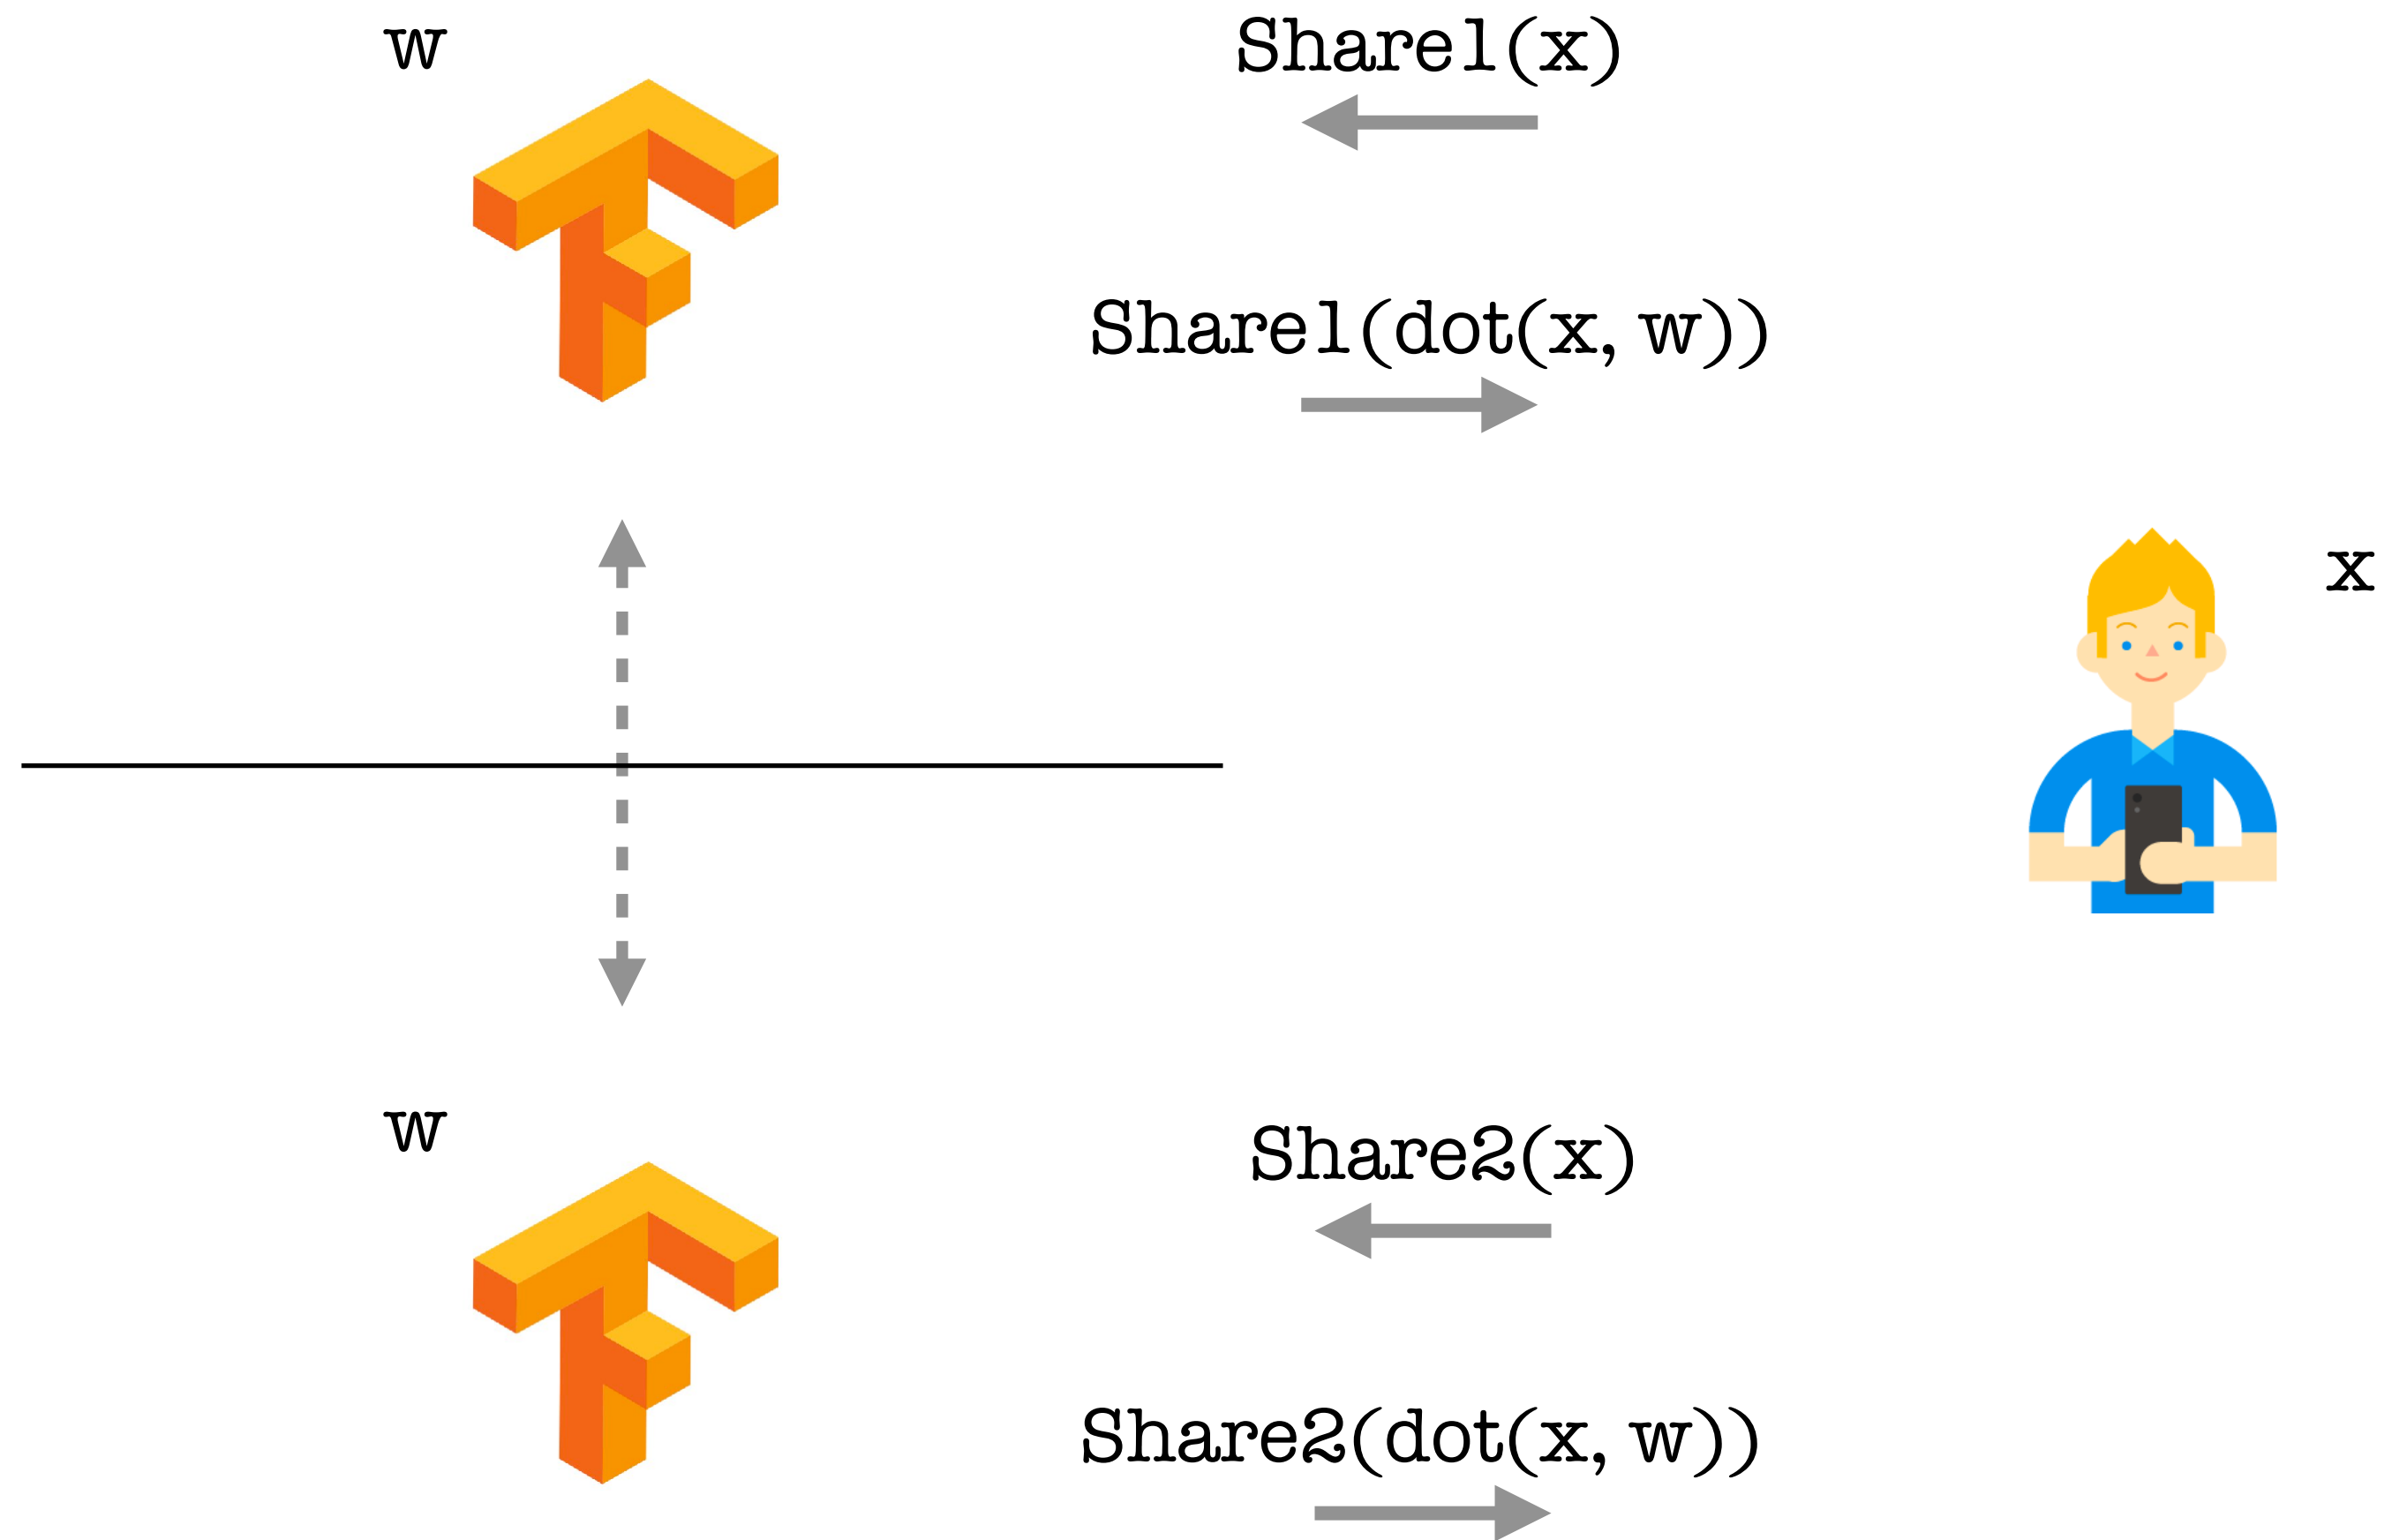
Share1(x)



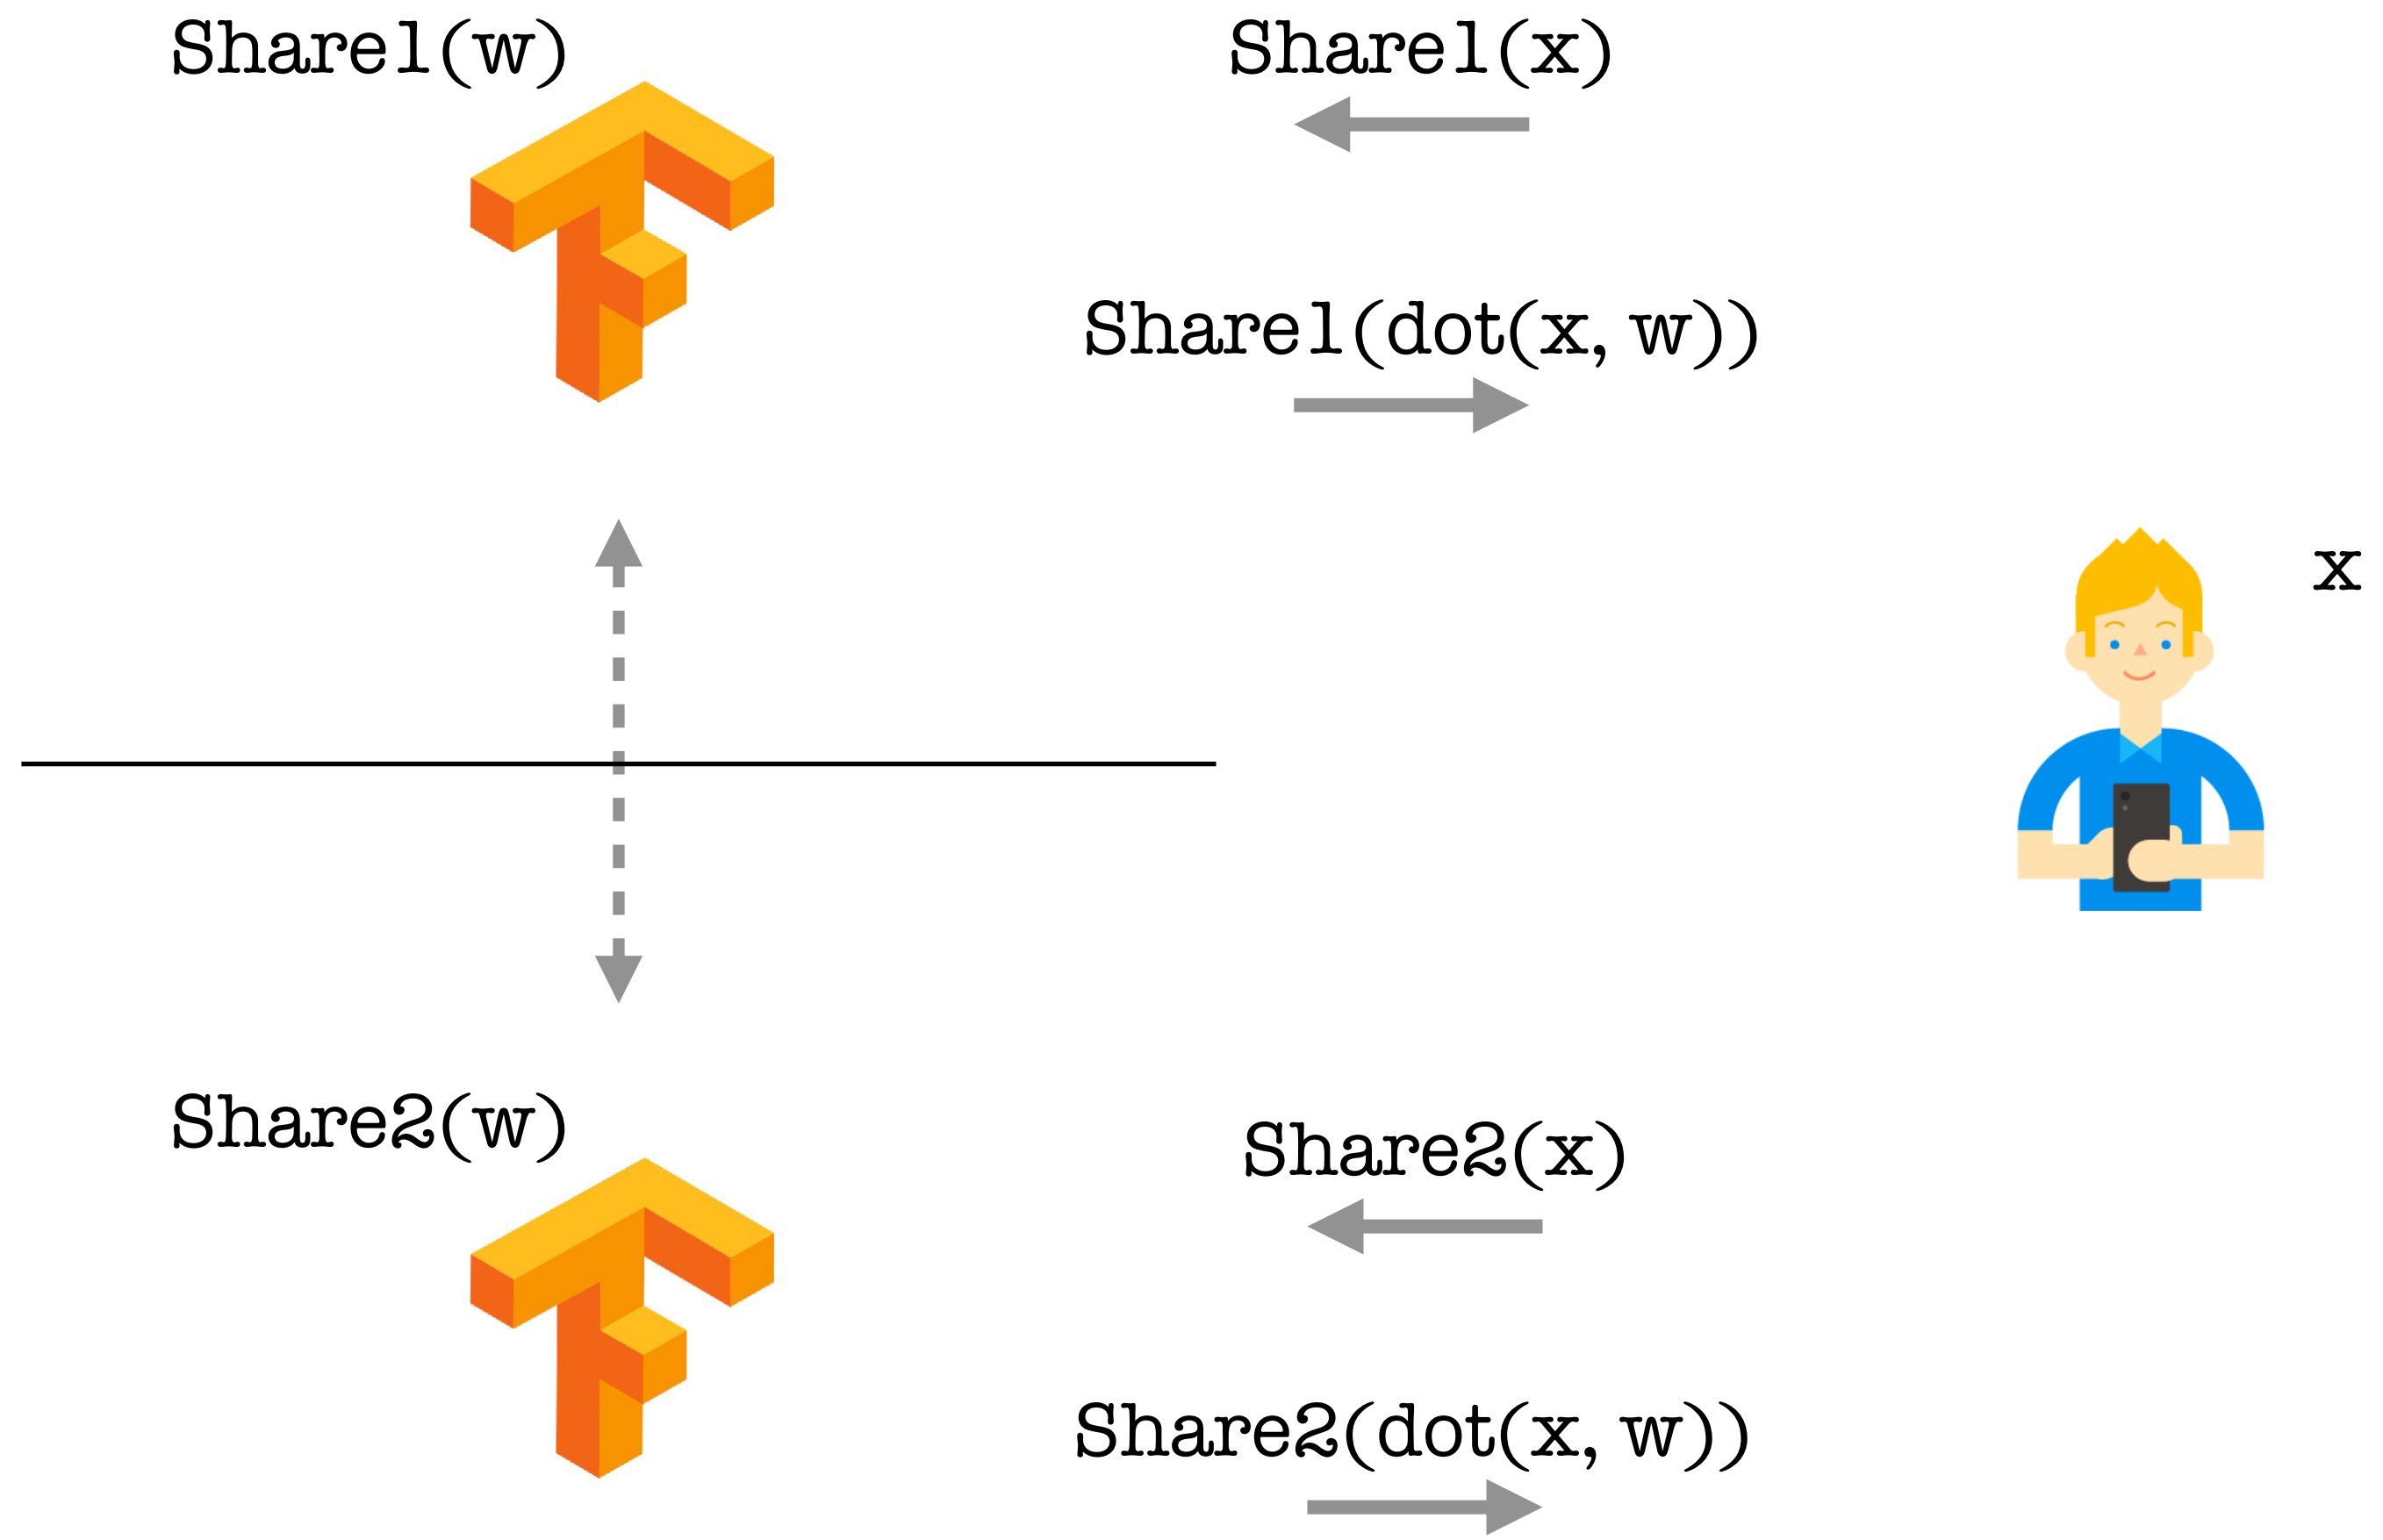
Share1(dot(x, w))



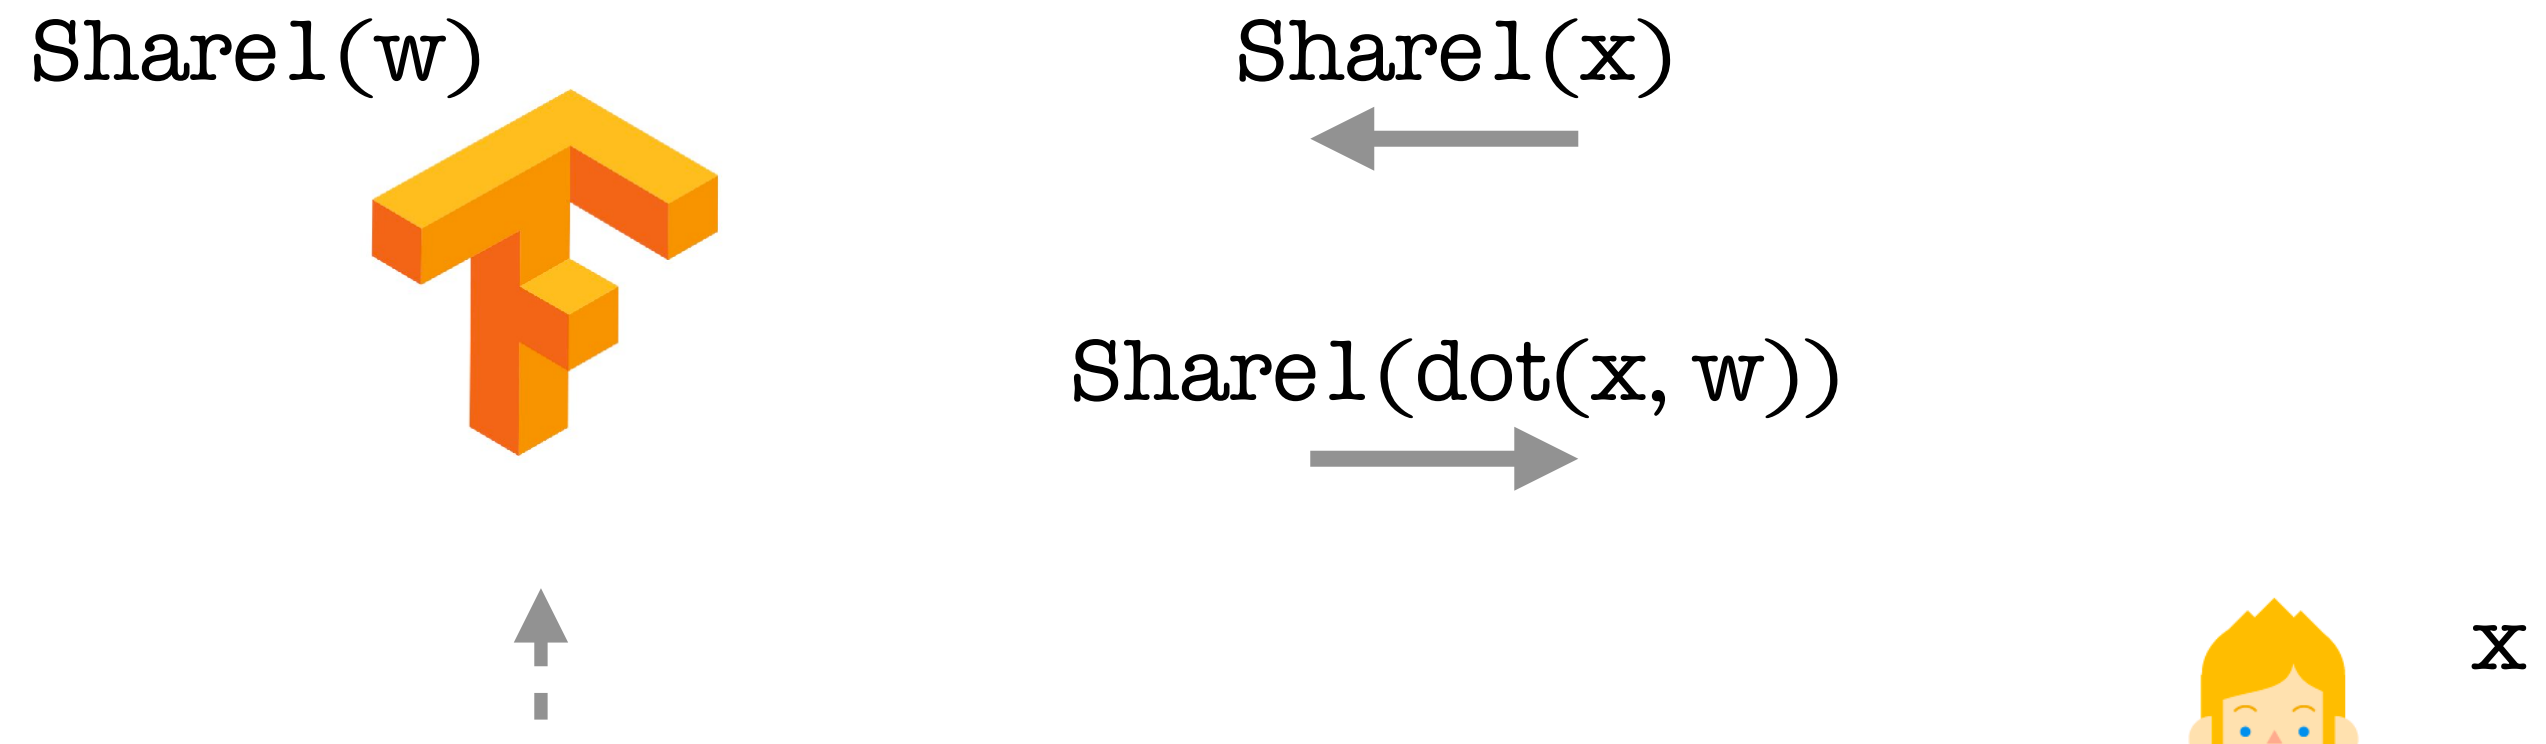
... using Secret Sharing



... using Secret Sharing



... using Secret Sharing




private multiplication

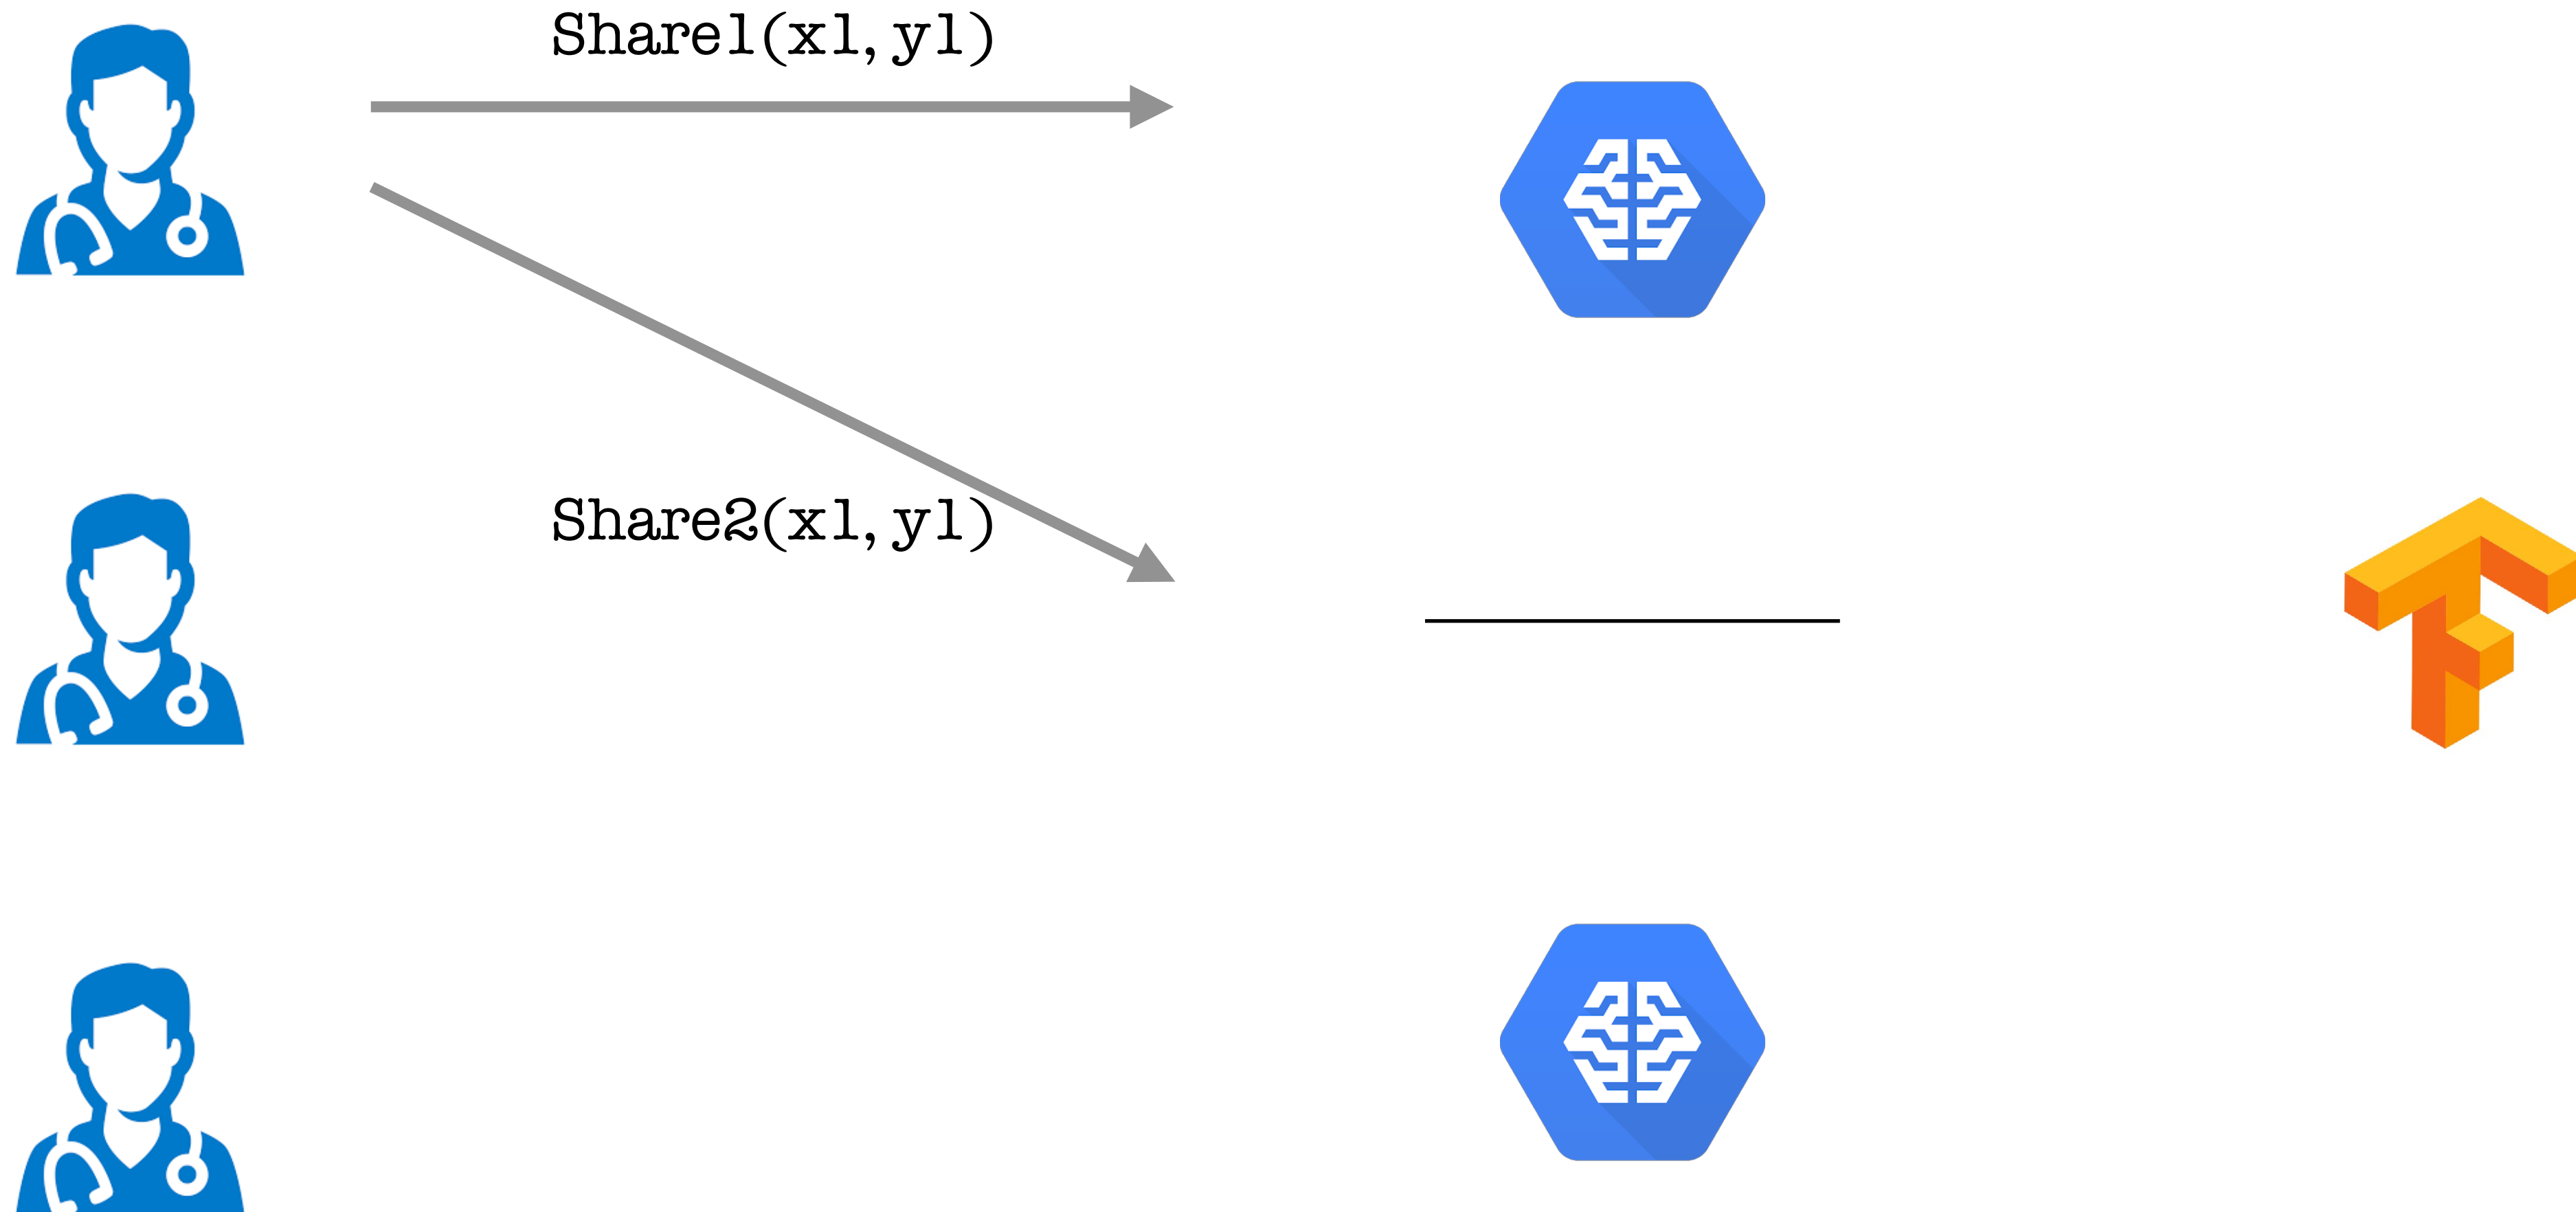
$$\text{dot}(\text{Share1}(x_0) \text{ Share1}(x_1) \text{ Share1}(x_2), \begin{matrix} \text{Share}(w_0) \\ \text{Share}(w_1) \\ \text{Share}(w_2) \end{matrix}) = \text{Share1}(x_0) * \text{Share}(w_0) + \dots$$
$$= \text{Share1}(x_0 * w_0) + \dots$$
$$= \text{Share1}(x_0 * w_0 + x_1 * w_1 + x_2 * w_2)$$

Training

Server-Aided



Server-Aided




Server-Aided




Share1(x2, y2)



Share2(x2, y2)




Server-Aided



Share1(x3, y3)

Share2(x3, y3)


Server-Aided




Server-Aided




Share1(w)



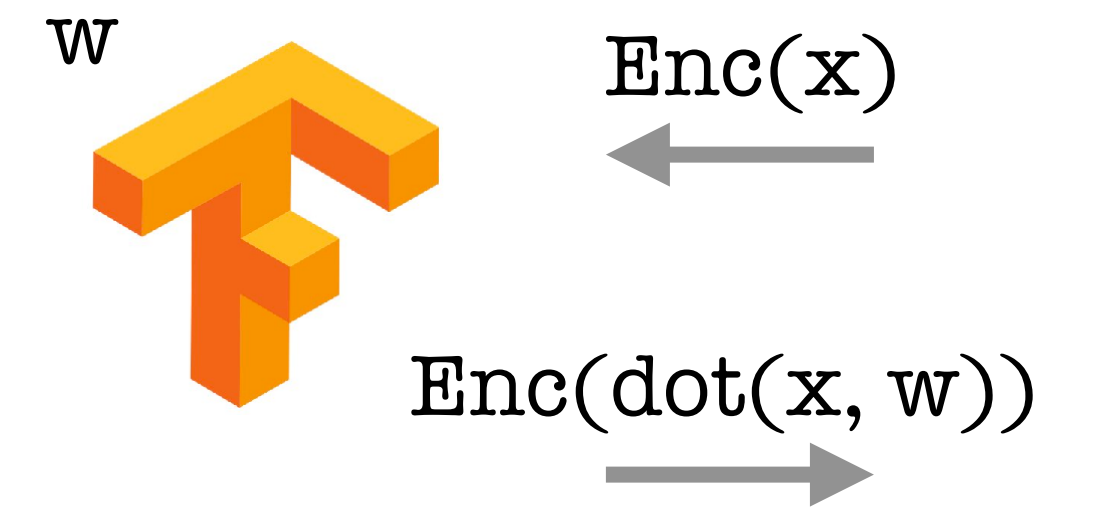
Share2(w)





Server-Aided



Server-Aided




Server-Aided




Server-Aided



Share1(w)



Share2(w)



Server-Aided




Share1(w)




Share1(x)



Share1(pred)



Share2(w)



Share2(x)




Share2(pred)




Multi-Party




Multi-Party




Multi-Party




Multi-Party



Share1(w)



Share2(w)



Making It Accessible

Projects and Literature

Projects and Literature

Recent research papers using secure computation

CryptoNets: *Applying Neural Networks to Encrypted Data with High Throughput and Accuracy*, Dowlin et al.

SecureML: *A System for Scalable Privacy-Preserving Machine Learning*, Mohassel and Zhang

DeepSecure: *Scalable Provably-Secure Deep Learning*, Rouhani et al.

Gazelle: *A Low Latency Framework for Secure Neural Network Inference*, Juvekar et al.

ABY3: *A Mixed Protocol Framework for Machine Learning*, Mohassel and Rindal

SecureNN: *Efficient and Private Neural Network Training*, Wagh et al.

Blind Justice: *Fairness with Encrypted Sensitive Attributes*, Kilbertus et al.

(also great summary in <https://eprint.iacr.org/2017/1190>)

Projects and Literature

Recent research papers using secure computation

CryptoNets: *Applying Neural Networks to Encrypted Data with High Throughput and Accuracy*, Dowlin et al.

SecureML: *A System for Scalable Privacy-Preserving Machine Learning*, Mohassel and Zhang

DeepSecure: *Scalable Provably-Secure Deep Learning*, Rouhani et al.

Gazelle: *A Low Latency Framework for Secure Neural Network Inference*, Juvekar et al.

ABY3: *A Mixed Protocol Framework for Machine Learning*, Mohassel and Rindal

SecureNN: *Efficient and Private Neural Network Training*, Wagh et al.

Blind Justice: *Fairness with Encrypted Sensitive Attributes*, Kilbertus et al.

(also great summary in <https://eprint.iacr.org/2017/1190>)

Secure computation frameworks

SCALE-MAMBA (<https://homes.esat.kuleuven.be/~nsmart/SCALE/>)

MP-SPDZ (<https://github.com/n1analytics/MP-SPDZ>)

ABY (<https://github.com/encryptogroup/ABY>)

OblivC (<http://oblivc.org/>)

(much more at <https://github.com/rdragos/awesome-mpc>)

Projects and Literature

Recent research papers using secure computation

CryptoNets: *Applying Neural Networks to Encrypted Data with High Throughput and Accuracy*, Dowlin et al.

SecureML: *A System for Scalable Privacy-Preserving Machine Learning*, Mohassel and Zhang

DeepSecure: *Scalable Provably-Secure Deep Learning*, Rouhani et al.

Gazelle: *A Low Latency Framework for Secure Neural Network Inference*, Juvekar et al.

ABY3: *A Mixed Protocol Framework for Machine Learning*, Mohassel and Rindal

SecureNN: *Efficient and Private Neural Network Training*, Wagh et al.

Blind Justice: *Fairness with Encrypted Sensitive Attributes*, Kilbertus et al.

(also great summary in <https://eprint.iacr.org/2017/1190>)

Specialised projects

tf-encrypted (<https://github.com/mortendahl/tf-encrypted>)

PySyft (<https://github.com/OpenMined/PySyft>)

Secure computation frameworks

SCALE-MAMBA (<https://homes.esat.kuleuven.be/~nsmart/SCALE/>)

MP-SPDZ (<https://github.com/n1analytics/MP-SPDZ>)

ABY (<https://github.com/encryptogroup/ABY>)

OblivC (<http://oblivc.org/>)

(much more at <https://github.com/rdragos/awesome-mpc>)

Multidisciplinary Challenge

Multidisciplinary Challenge

Cryptography
(techniques, protocols, trust)

Multidisciplinary Challenge

Cryptography
(techniques, protocols, trust)

Machine learning
(models, approx, precision)

Multidisciplinary Challenge

Cryptography
(techniques, protocols, trust)

Machine learning
(models, approx, precision)

Engineering
(distributed, multi-core, readability)

Multidisciplinary Challenge

Data science
(use-cases, workflow, monitoring)

Cryptography
(techniques, protocols, trust)

Machine learning
(models, approx, precision)

Engineering
(distributed, multi-core, readability)

Multidisciplinary Challenge

Data science
(use-cases, workflow, monitoring)

Cryptography
(techniques, protocols, trust)

Machine learning
(models, approx, precision)

Engineering
(distributed, multi-core, readability)

need for common language

tf-encrypted

tf-encrypted

open source project for exploring and experimenting with
privacy-preserving machine learning in TensorFlow

tf-encrypted

open source project for exploring and experimenting with
privacy-preserving machine learning in TensorFlow

separate concerns, take expertise out of equation,
and provide tight integration with ecosystem

Private Prediction with tf-encrypted

Private Prediction with tf-encrypted

```
1 import tensorflow as tf
2
3 def provide_weights(): """ Load from disk """
4 def provide_input(): """ Pre-process """
5 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7 # get model weights
8 w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer1, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```

Private Prediction with tf-encrypted

```
1 import tensorflow as tf
2
3 def provide_weights(): """ Load from disk """
4 def provide_input(): """ Pre-process """
5 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7 # get model weights
8 w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer1, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```


Private Prediction with tf-encrypted

```
1 import tensorflow as tf
2
3 def provide_weights(): """ Load from disk """
4 def provide_input(): """ Pre-process """
5 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7 # get model weights
8 w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer2, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```

Private Prediction with tf-encrypted

```
1  import tensorflow as tf
2
3  def provide_weights(): """ Load from disk """
4  def provide_input(): """ Pre-process """
5  def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7  # get model weights
8  w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer1, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```

Private Prediction with tf-encrypted

```
1 import tensorflow as tf
2
3 def provide_weights(): """ Load from disk """
4 def provide_input(): """ Pre-process """
5 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7 # get model weights
8 w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer1, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```

```
1 import tensorflow as tf
2 import tf_encrypted as tfe
3
4 def provide_weights(): """ Load from disk """
5 def provide_input(): """ Pre-process """
6 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
7
8 # get model weights as private tensors from owner
9 w0, b0, w1, b1, w2, b2 = tfe.define_private_input("model-owner", provide_weights)
10
11 # get prediction input as private tensors from client
12 x = tfe.define_private_input("prediction-client", provide_input)
13
14 # compute private prediction on servers
15 layer0 = tfe.relu((tfe.matmul(x, w0) + b0))
16 layer1 = tfe.relu((tfe.matmul(layer0, w1) + b1))
17 logits = tfe.matmul(layer1, w2) + b2
18
19 # process result of prediction on client
20 prediction_op = tfe.define_output("prediction-client", logits, receive_output)
21
22 # run secure graph execution in a tf.Session
23 with tfe.Session() as sess:
24     sess.run(tf.global_variables_initializer())
25     sess.run(prediction_op)
```

Private Prediction with tf-encrypted

```
1 import tensorflow as tf
2
3 def provide_weights():""" Load from disk """
4 def provide_input(): """ Pre-process """
5 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7 # get model weights
8 w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer1, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```

```
1 import tensorflow as tf
2 import tf_encrypted as tfe
3
4 def provide_weights():""" Load from disk """
5 def provide_input(): """ Pre-process """
6 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
7
8 # get model weights as private tensors from owner
9 w0, b0, w1, b1, w2, b2 = tfe.define_private_input("model-owner", provide_weights)
10
11 # get prediction input as private tensors from client
12 x = tfe.define_private_input("prediction-client", provide_input)
13
14 # compute private prediction on servers
15 layer0 = tfe.relu((tfe.matmul(x, w0) + b0))
16 layer1 = tfe.relu((tfe.matmul(layer0, w1) + b1))
17 logits = tfe.matmul(layer1, w2) + b2
18
19 # process result of prediction on client
20 prediction_op = tfe.define_output("prediction-client", logits, receive_output)
21
22 # run secure graph execution in a tf.Session
23 with tfe.Session() as sess:
24     sess.run(tf.global_variables_initializer())
25     sess.run(prediction_op)
```


Private Prediction with tf-encrypted

```
1 import tensorflow as tf
2
3 def provide_weights():""" Load from disk """
4 def provide_input(): """ Pre-process """
5 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7 # get model weights
8 w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer1, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```

```
1 import tensorflow as tf
2 import tf_encrypted as tfe
3
4 def provide_weights():""" Load from disk """
5 def provide_input(): """ Pre-process """
6 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
7
8 # get model weights as private tensors from owner
9 w0, b0, w1, b1, w2, b2 = tfe.define_private_input("model-owner", provide_weights)
10
11 # get prediction input as private tensors from client
12 x = tfe.define_private_input("prediction-client", provide_input)
13
14 # compute private prediction on servers
15 layer0 = tfe.relu((tfe.matmul(x, w0) + b0))
16 layer1 = tfe.relu((tfe.matmul(layer0, w1) + b1))
17 logits = tfe.matmul(layer1, w2) + b2
18
19 # process result of prediction on client
20 prediction_op = tfe.define_output("prediction-client", logits, receive_output)
21
22 # run secure graph execution in a tf.Session
23 with tfe.Session() as sess:
24     sess.run(tf.global_variables_initializer())
25     sess.run(prediction_op)
```

Private Prediction with tf-encrypted

```
1 import tensorflow as tf
2
3 def provide_weights(): """ Load from disk """
4 def provide_input(): """ Pre-process """
5 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
6
7 # get model weights
8 w0, b0, w1, b1, w2, b2 = provide_weights()
9
10 # get prediction input
11 x = provide_input()
12
13 # compute prediction
14 layer0 = tf.nn.relu((tf.matmul(x, w0) + b0))
15 layer1 = tf.nn.relu((tf.matmul(layer0, w1) + b1))
16 logits = tf.matmul(layer1, w2) + b2
17
18 # process result of prediction
19 prediction_op = receive_output(logits)
20
21 # run graph execution in a tf.Session
22 with tf.Session() as sess:
23     sess.run(tf.global_variables_initializer())
24     sess.run(prediction_op)
```

```
1 import tensorflow as tf
2 import tf_encrypted as tfe
3
4 def provide_weights(): """ Load from disk """
5 def provide_input(): """ Pre-process """
6 def receive_output(logits): return tf.Print([], [tf.argmax(logits)])
7
8 # get model weights as private tensors from owner
9 w0, b0, w1, b1, w2, b2 = tfe.define_private_input("model-owner", provide_weights)
10
11 # get prediction input as private tensors from client
12 x = tfe.define_private_input("prediction-client", provide_input)
13
14 # compute private prediction on servers
15 layer0 = tfe.relu((tfe.matmul(x, w0) + b0))
16 layer1 = tfe.relu((tfe.matmul(layer0, w1) + b1))
17 logits = tfe.matmul(layer1, w2) + b2
18
19 # process result of prediction on client
20 prediction_op = tfe.define_output("prediction-client", logits, receive_output)
21
22 # run secure graph execution in a tf.Session
23 with tfe.Session() as sess:
24     sess.run(tf.global_variables_initializer())
25     sess.run(prediction_op)
```

Wrap-Up

Wrap-Up

Privacy-preserving ML mitigate **bottlenecks** and
enable access to sensitive information

Wrap-Up

Privacy-preserving ML mitigate **bottlenecks** and **enable access** to sensitive information

Secure computation **distributes trust and control**, and is complementary to e.g. differential privacy

Wrap-Up

Privacy-preserving ML mitigate **bottlenecks** and **enable access** to sensitive information

Secure computation **distributes trust and control**, and is complementary to e.g. differential privacy

Privacy-preserving ML is a multidisciplinary field benefitting from **adaptations** on both sides

Wrap-Up

Privacy-preserving ML mitigate **bottlenecks** and **enable access** to sensitive information

Secure computation **distributes trust and control**, and is complementary to e.g. differential privacy

Privacy-preserving ML is a multidisciplinary field benefitting from **adaptations** on both sides

Focus on **usability** and **integration**

Wrap-Up

Privacy-preserving ML mitigate **bottlenecks** and **enable access** to sensitive information

Secure computation **distributes trust and control**, and is complementary to e.g. differential privacy

Privacy-preserving ML is a multidisciplinary field benefitting from **adaptations** on both sides

Focus on **usability** and **integration**

Thank you!