


[320] Pre-Machine Learning: Intro

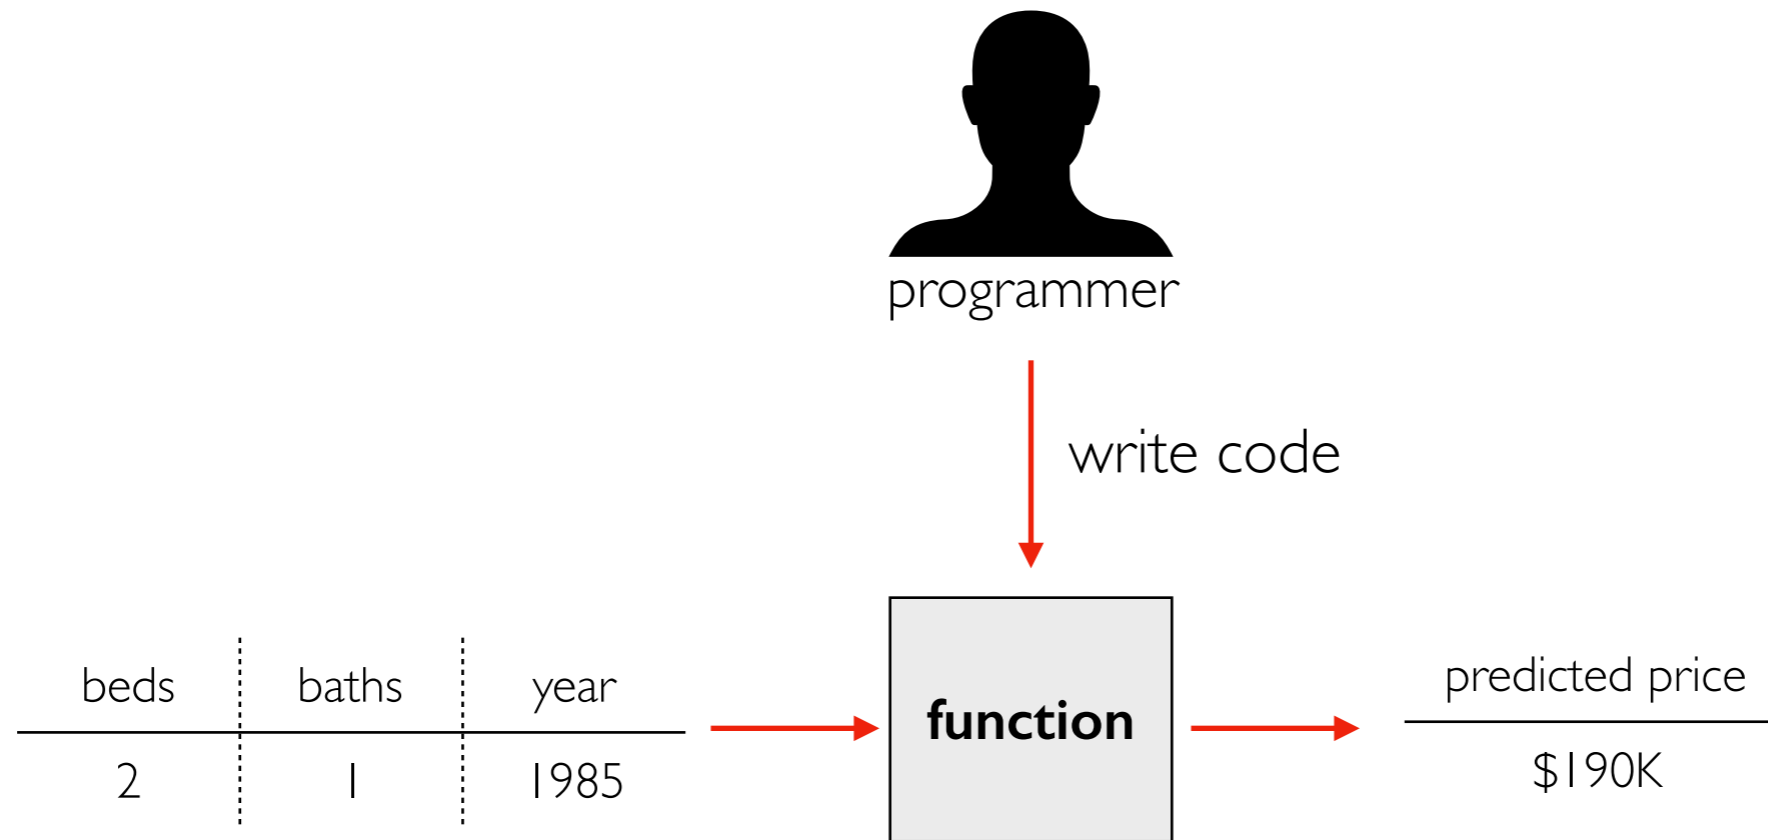
Tyler Caraza-Harter

Functions/Models

How do we make functions?

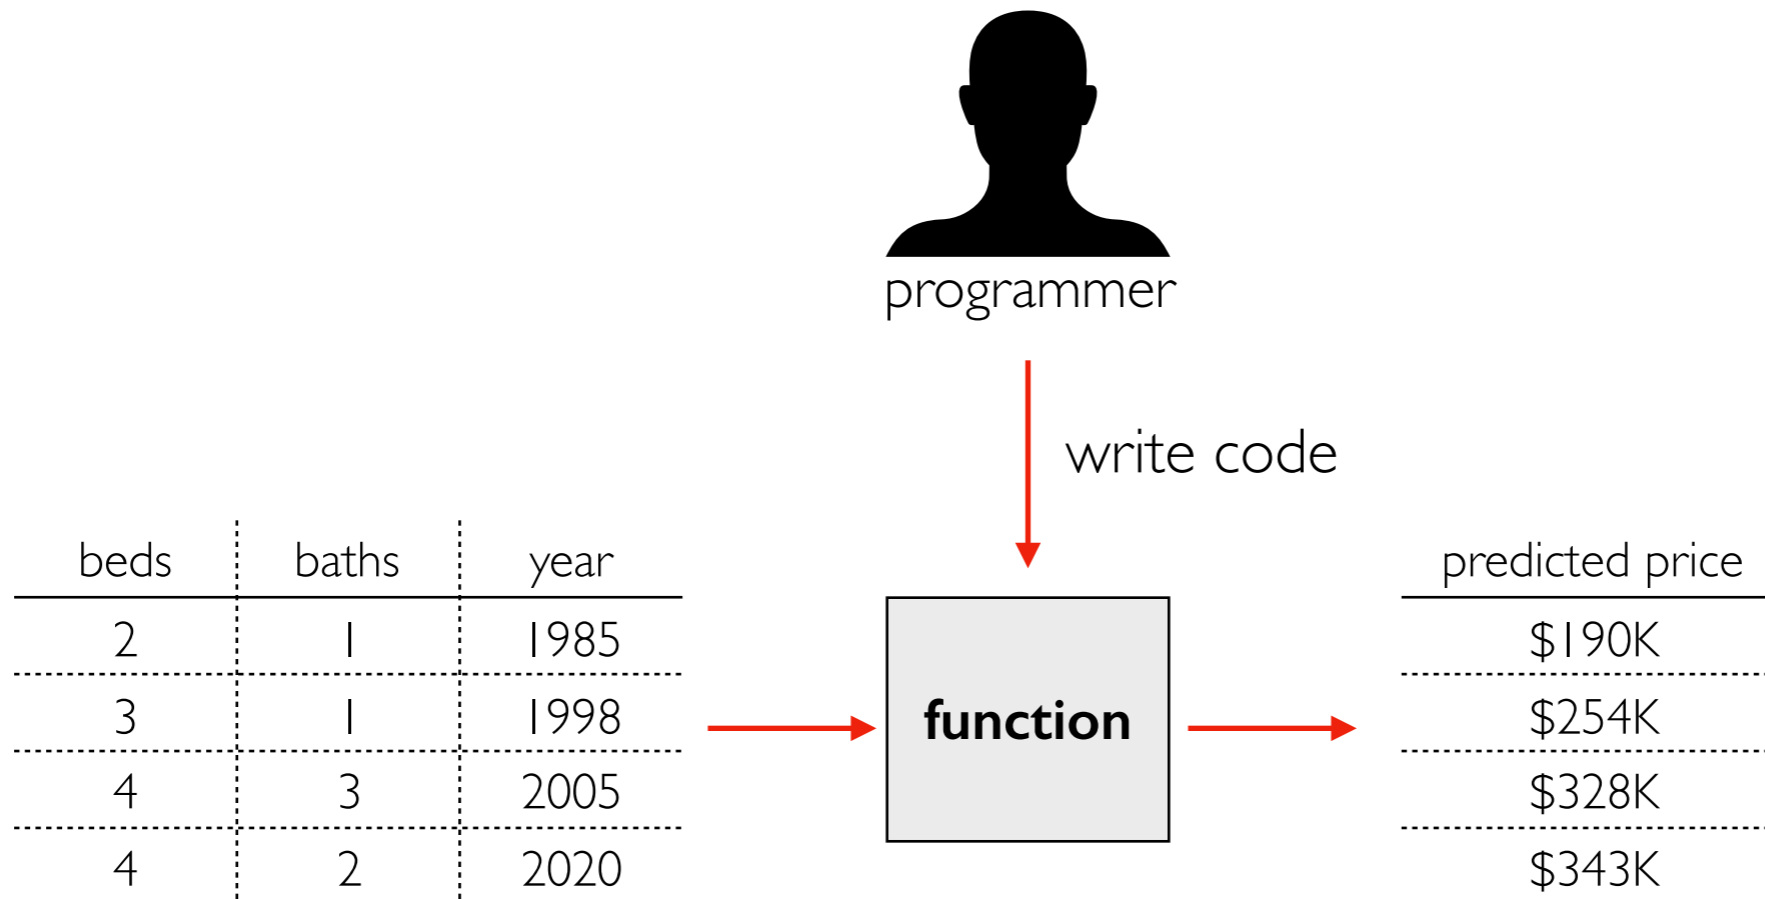


How do we make functions?




many functions are **models** that can be used to predict

How do we make functions?

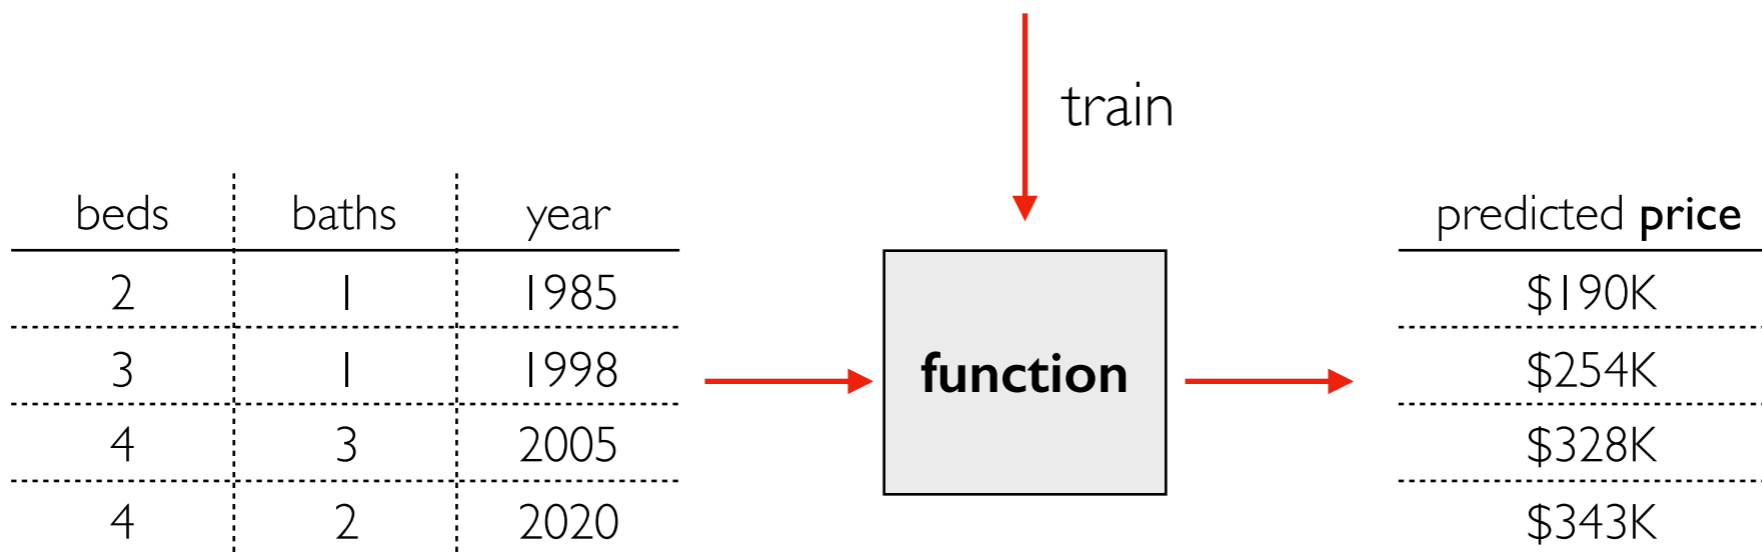


many functions are **models** that can be used to predict

How do we make functions?



Machine Learning Algorithm




many functions are **models** that can be used to predict

How do we make functions?

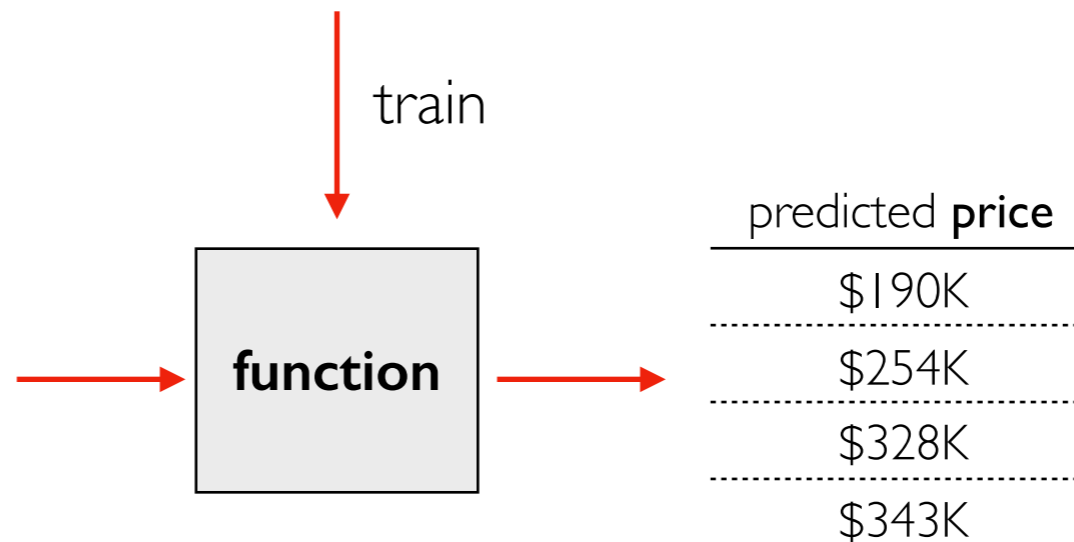
training data

beds	baths	year	price
1	1	1980	\$140K
3	1	1990	\$240K
3	4	2004	\$295K
4	3	2018	\$350K



live data

beds	baths	year
2	1	1985
3	1	1998
4	3	2005
4	2	2020



many functions are **models** that can be used to predict

Kinds of Machine Learning

Main Categories of Machine Learning

learning from data

1

Supervised Machine Learning

data is **labeled**, we know what we want to predict

2


Unsupervised Machine Learning

data is **unlabeled**, we're just looking for patterns

3

Reinforcement Learning

not covered in CS 320



Main Categories of Machine Learning

learning from data

1

Supervised Machine Learning

data is **labeled**, we know what we want to predict

2

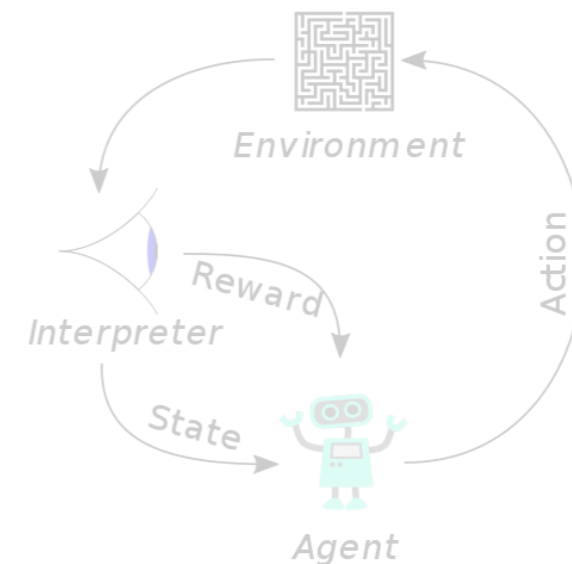
Unsupervised Machine Learning

data is **unlabeled**, we're just looking for patterns

3

Reinforcement Learning

not covered in CS 320




Supervised Learning

supervised because data is labeled,
we know what we want to predict

training data

beds	baths	year	price
1	1	1980	\$140K
3	1	1990	\$240K
3	4	2004	\$295K
4	3	2018	\$350K



Machine Learning Algorithm

live data

beds	baths	year
2	1	1985
3	1	1998
4	3	2005
4	2	2020

train

function

predicted price

\$190K

\$254K


\$328K

\$343K

Supervised Learning: Regression


training data

beds	baths	year	price
1	1	1980	\$140K
3	1	1990	\$240K
3	4	2004	\$295K
4	3	2018	\$350K



live data

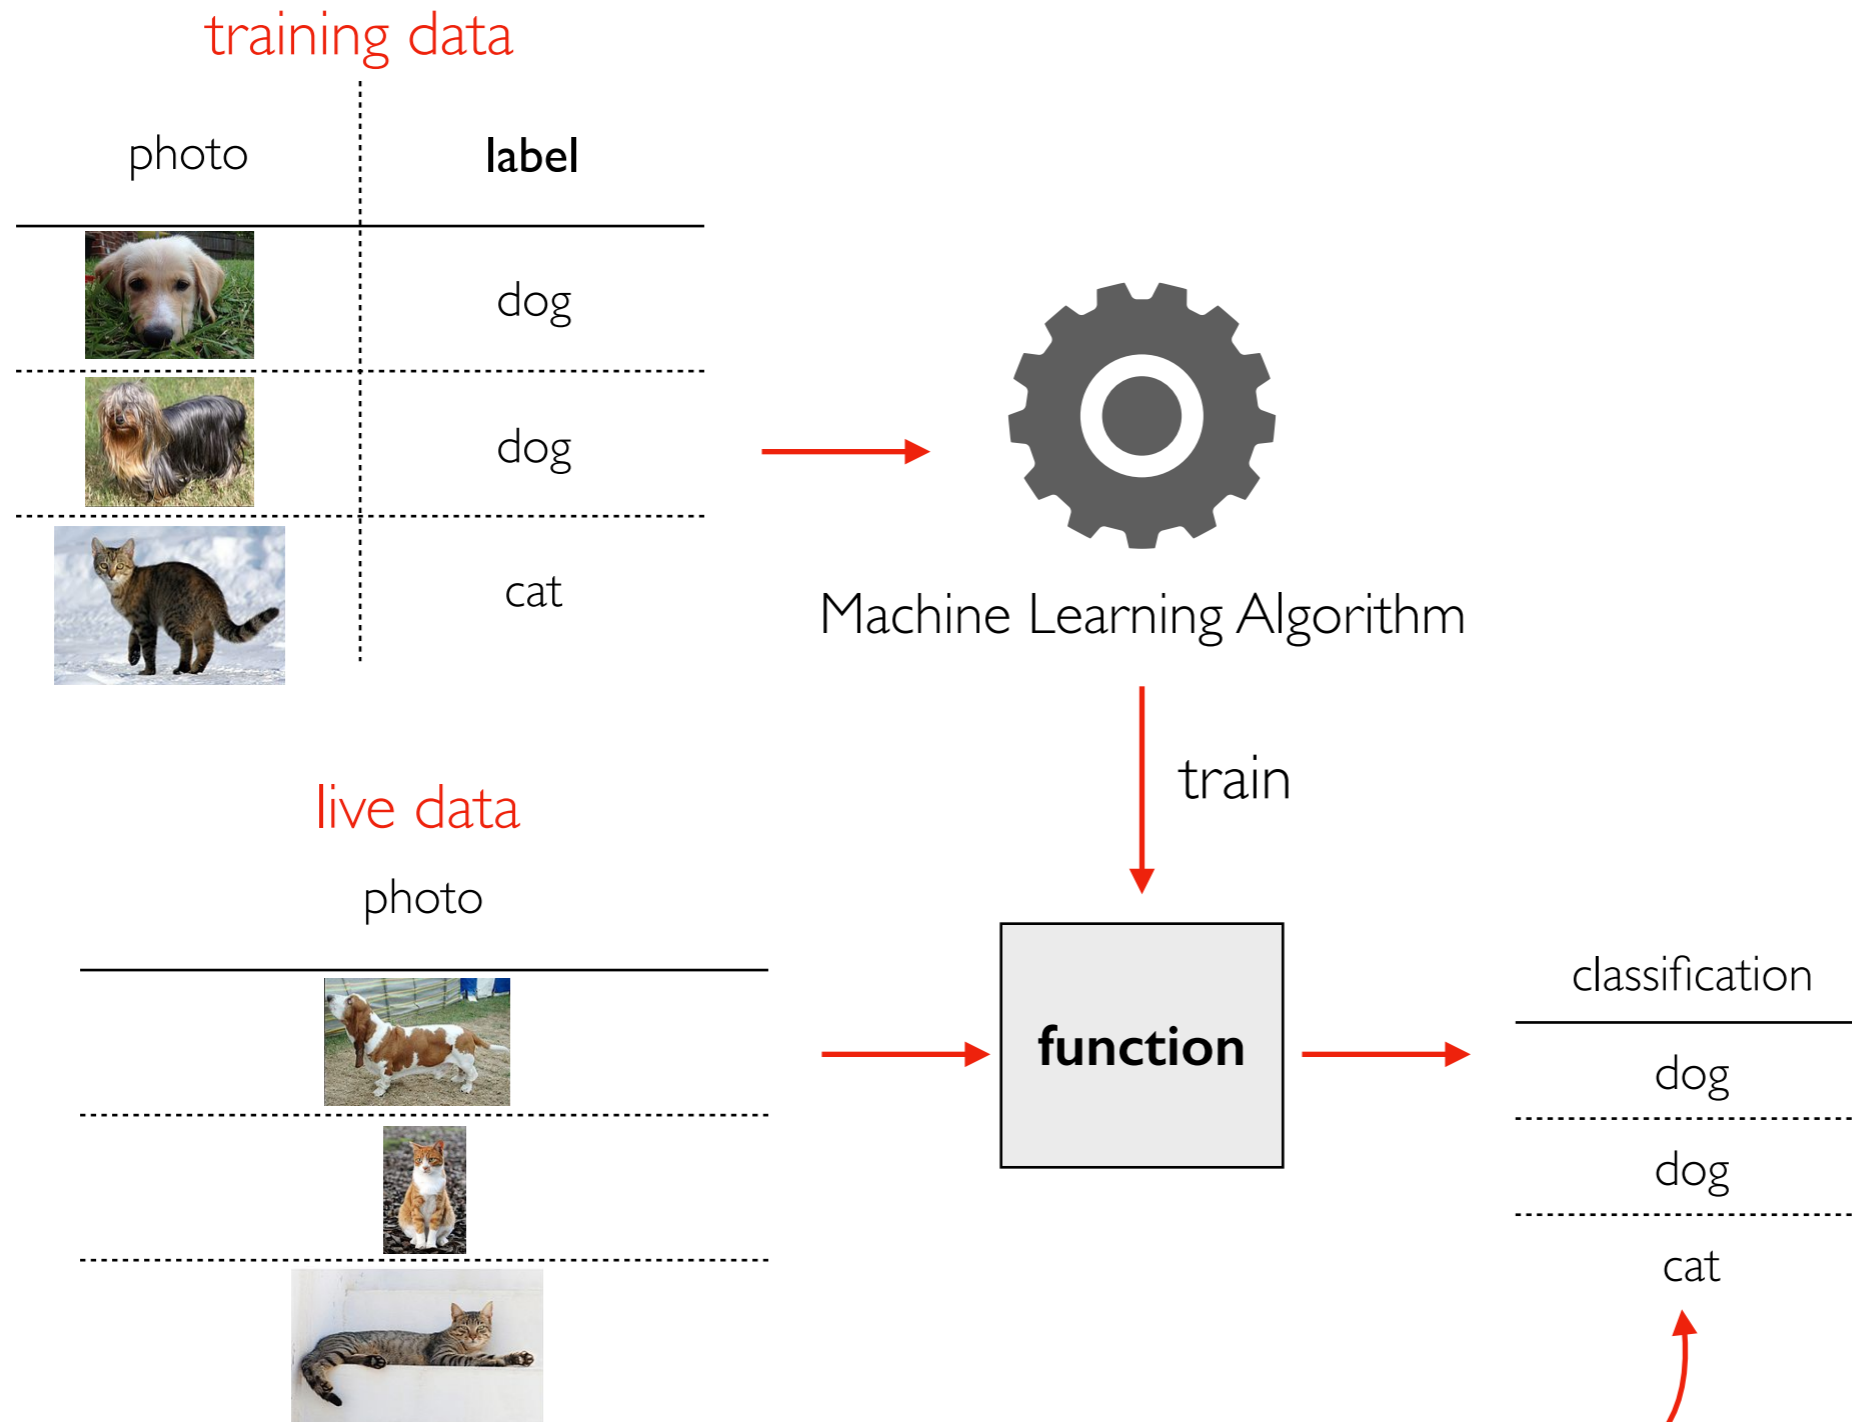
beds	baths	year
2	1	1985
3	1	1998
4	3	2005
4	2	2020



predicted price
\$190K
\$254K
\$328K
\$343K

our function is a **regressor** because it's outputting continuous data

Supervised Learning: Classification



our function is a **classifier** because it's outputting discrete data

Main Categories of Machine Learning

learning from data

1

Supervised Machine Learning

data is **labeled**, we know what we want to predict

2


Unsupervised Machine Learning

data is **unlabeled**, we're just looking for patterns

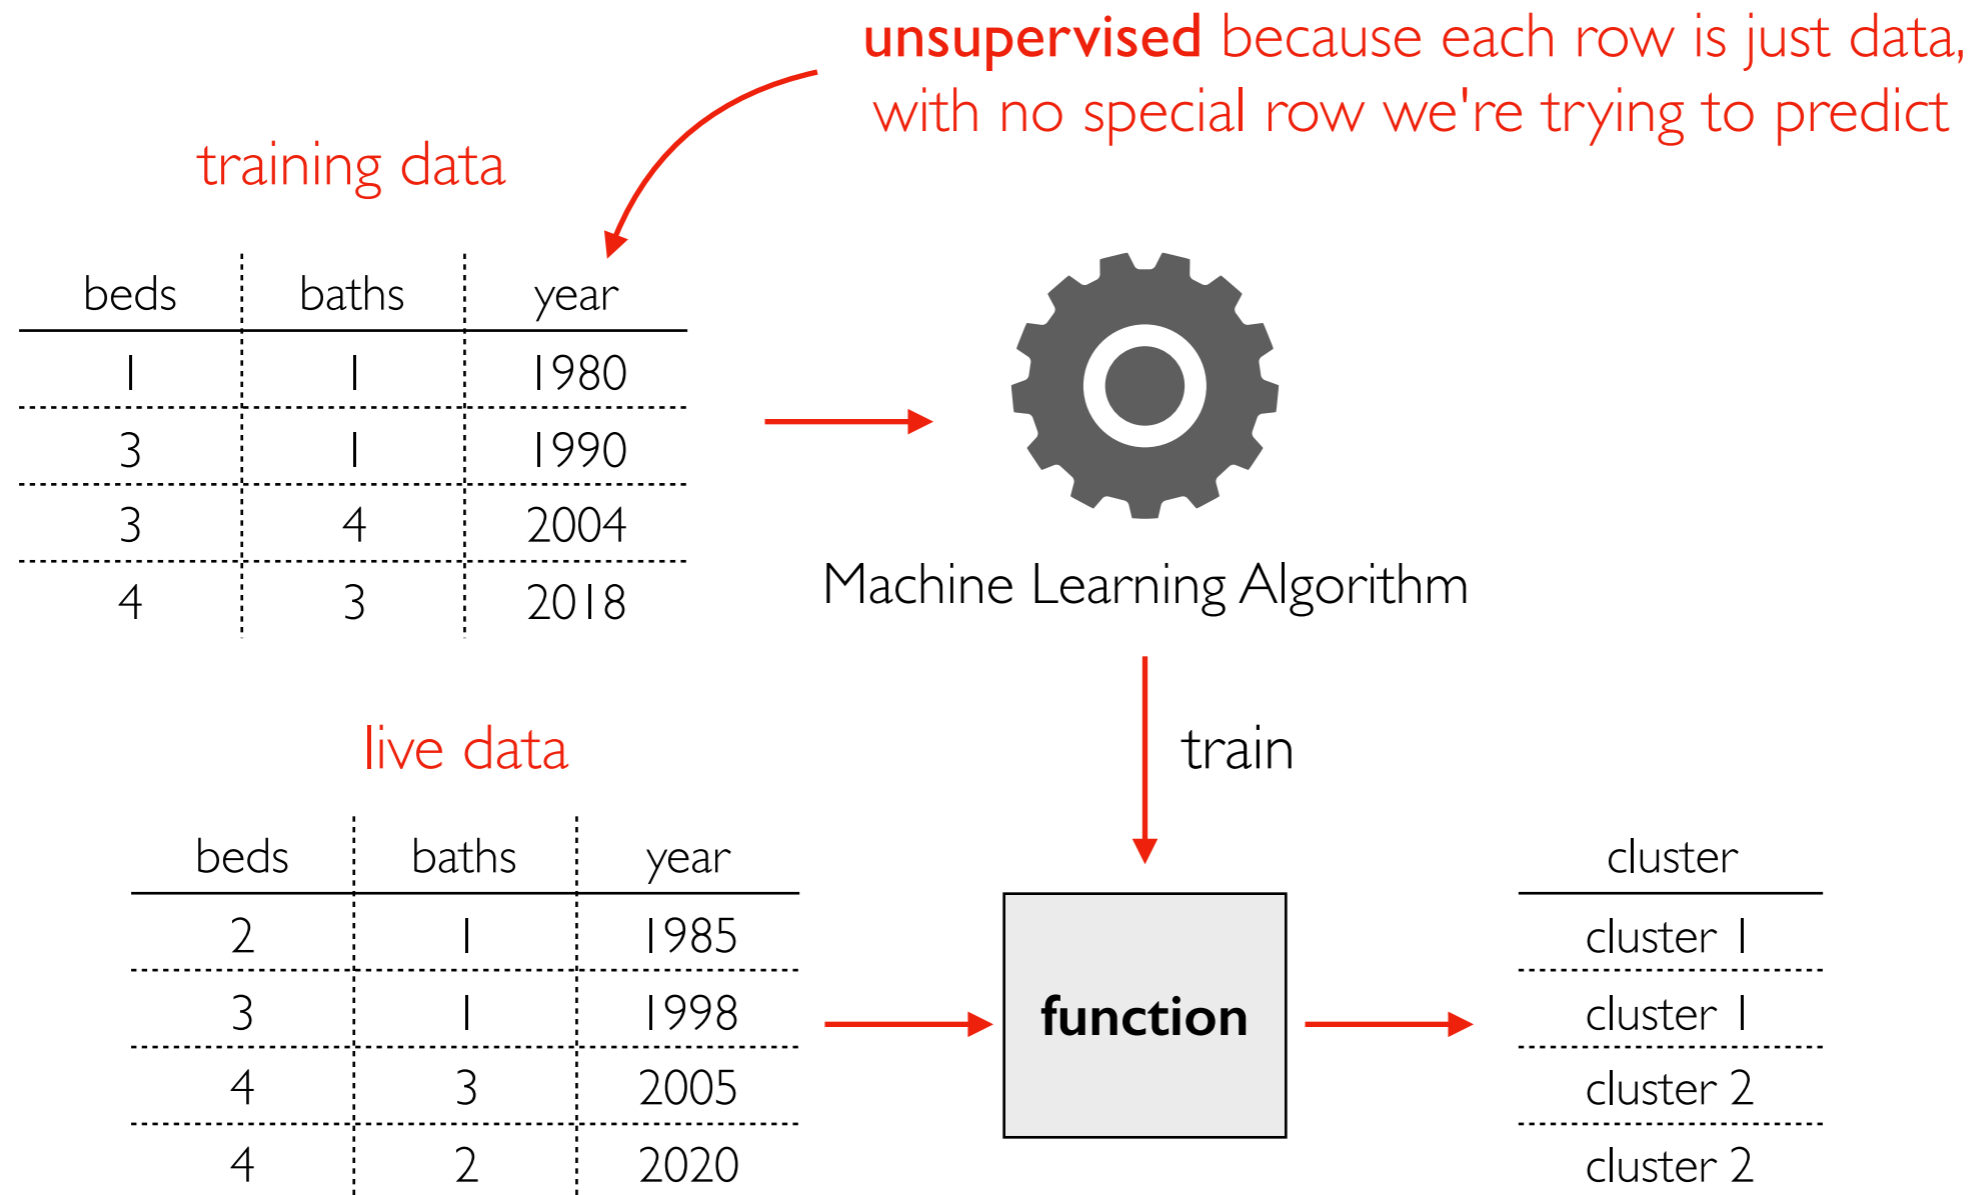
3

Reinforcement Learning

not covered in CS 320



Unsupervised Learning



unsupervised clustering algorithms try to identify groups of similar data. The algorithm decides the groups.

Sometimes (but often not) they'll correspond to things we describe. E.g., cluster 1: old houses with few bathrooms; cluster 2: new houses with many bathrooms

Foundations

Important Packages

We'll be learning the following to do ML and related calculations efficiently:

1

numpy

2


pytorch

3


scikit-learn

```
pip3 install numpy torch torchvision scikit-learn
```


Linear Algebra




with matrices...



Linear Algebra




with matrices...




$$\begin{matrix} A \\ \begin{bmatrix} 2 & 1 & 1985 \\ 3 & 1 & 1998 \\ 4 & 3 & 2005 \\ 4 & 2 & 2020 \end{bmatrix} \end{matrix} \cdot \begin{matrix} x \\ \begin{bmatrix} 41.46 \\ 10.36 \\ 1.70 \end{bmatrix} \end{matrix} + \begin{matrix} b \\ -3292.31 \end{matrix}$$

dot product

Linear Algebra



with matrices...



$$\begin{matrix} A \\ \begin{bmatrix} 2 & 1 & 1985 \\ 3 & 1 & 1998 \\ 4 & 3 & 2005 \\ 4 & 2 & 2020 \end{bmatrix} \end{matrix} \cdot \begin{matrix} x \\ \begin{bmatrix} 41.46 \\ 10.36 \\ 1.70 \end{bmatrix} \end{matrix} + \begin{matrix} b \\ -3292.31 \end{matrix}$$

dot product

```
import numpy as np
A = df.values
y = np.dot(A, x) + b
```

Linear Algebra

$$y = x^{**2} \quad \text{not linear}$$

$$y = 3*x_0 + -2*x_1 + 0.5*x_2 + \dots + 10*x_{49} \quad \text{linear}$$

with matrices...

$$\begin{bmatrix} 2 & 1 & 1985 \\ 3 & 1 & 1998 \\ 4 & 3 & 2005 \\ 4 & 2 & 2020 \end{bmatrix} \longrightarrow \boxed{y = Ax + b} \longrightarrow \begin{bmatrix} 190000 \\ 254000 \\ 328000 \\ 343000 \end{bmatrix}$$

$$\begin{matrix} A \\ \begin{bmatrix} 2 & 1 & 1985 \\ 3 & 1 & 1998 \\ 4 & 3 & 2005 \\ 4 & 2 & 2020 \end{bmatrix} \end{matrix} \cdot \begin{matrix} x \\ \begin{bmatrix} 41.46 \\ 10.36 \\ 1.70 \end{bmatrix} \end{matrix} + \begin{matrix} b \\ -3292.31 \end{matrix}$$

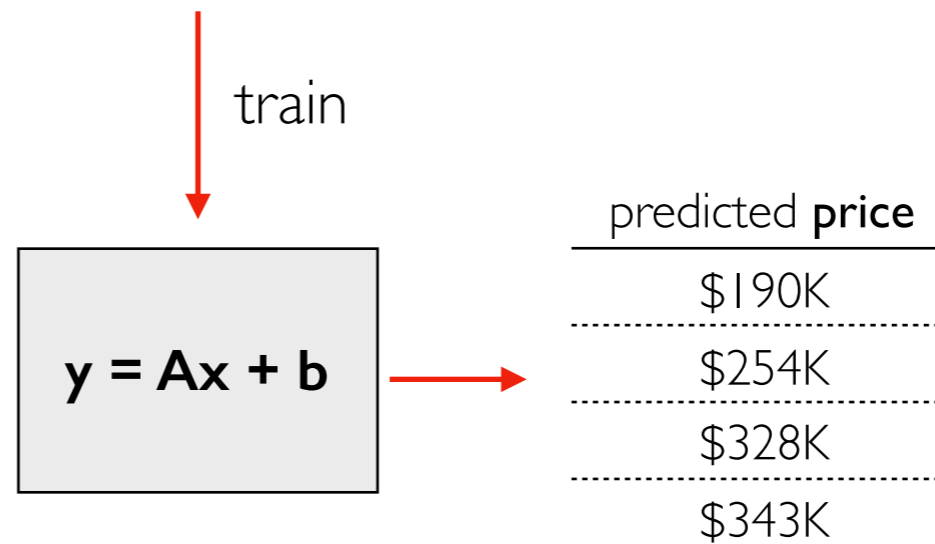

dot product

```
import numpy as np
A = df.values
y = np.dot(A, x) + b
```

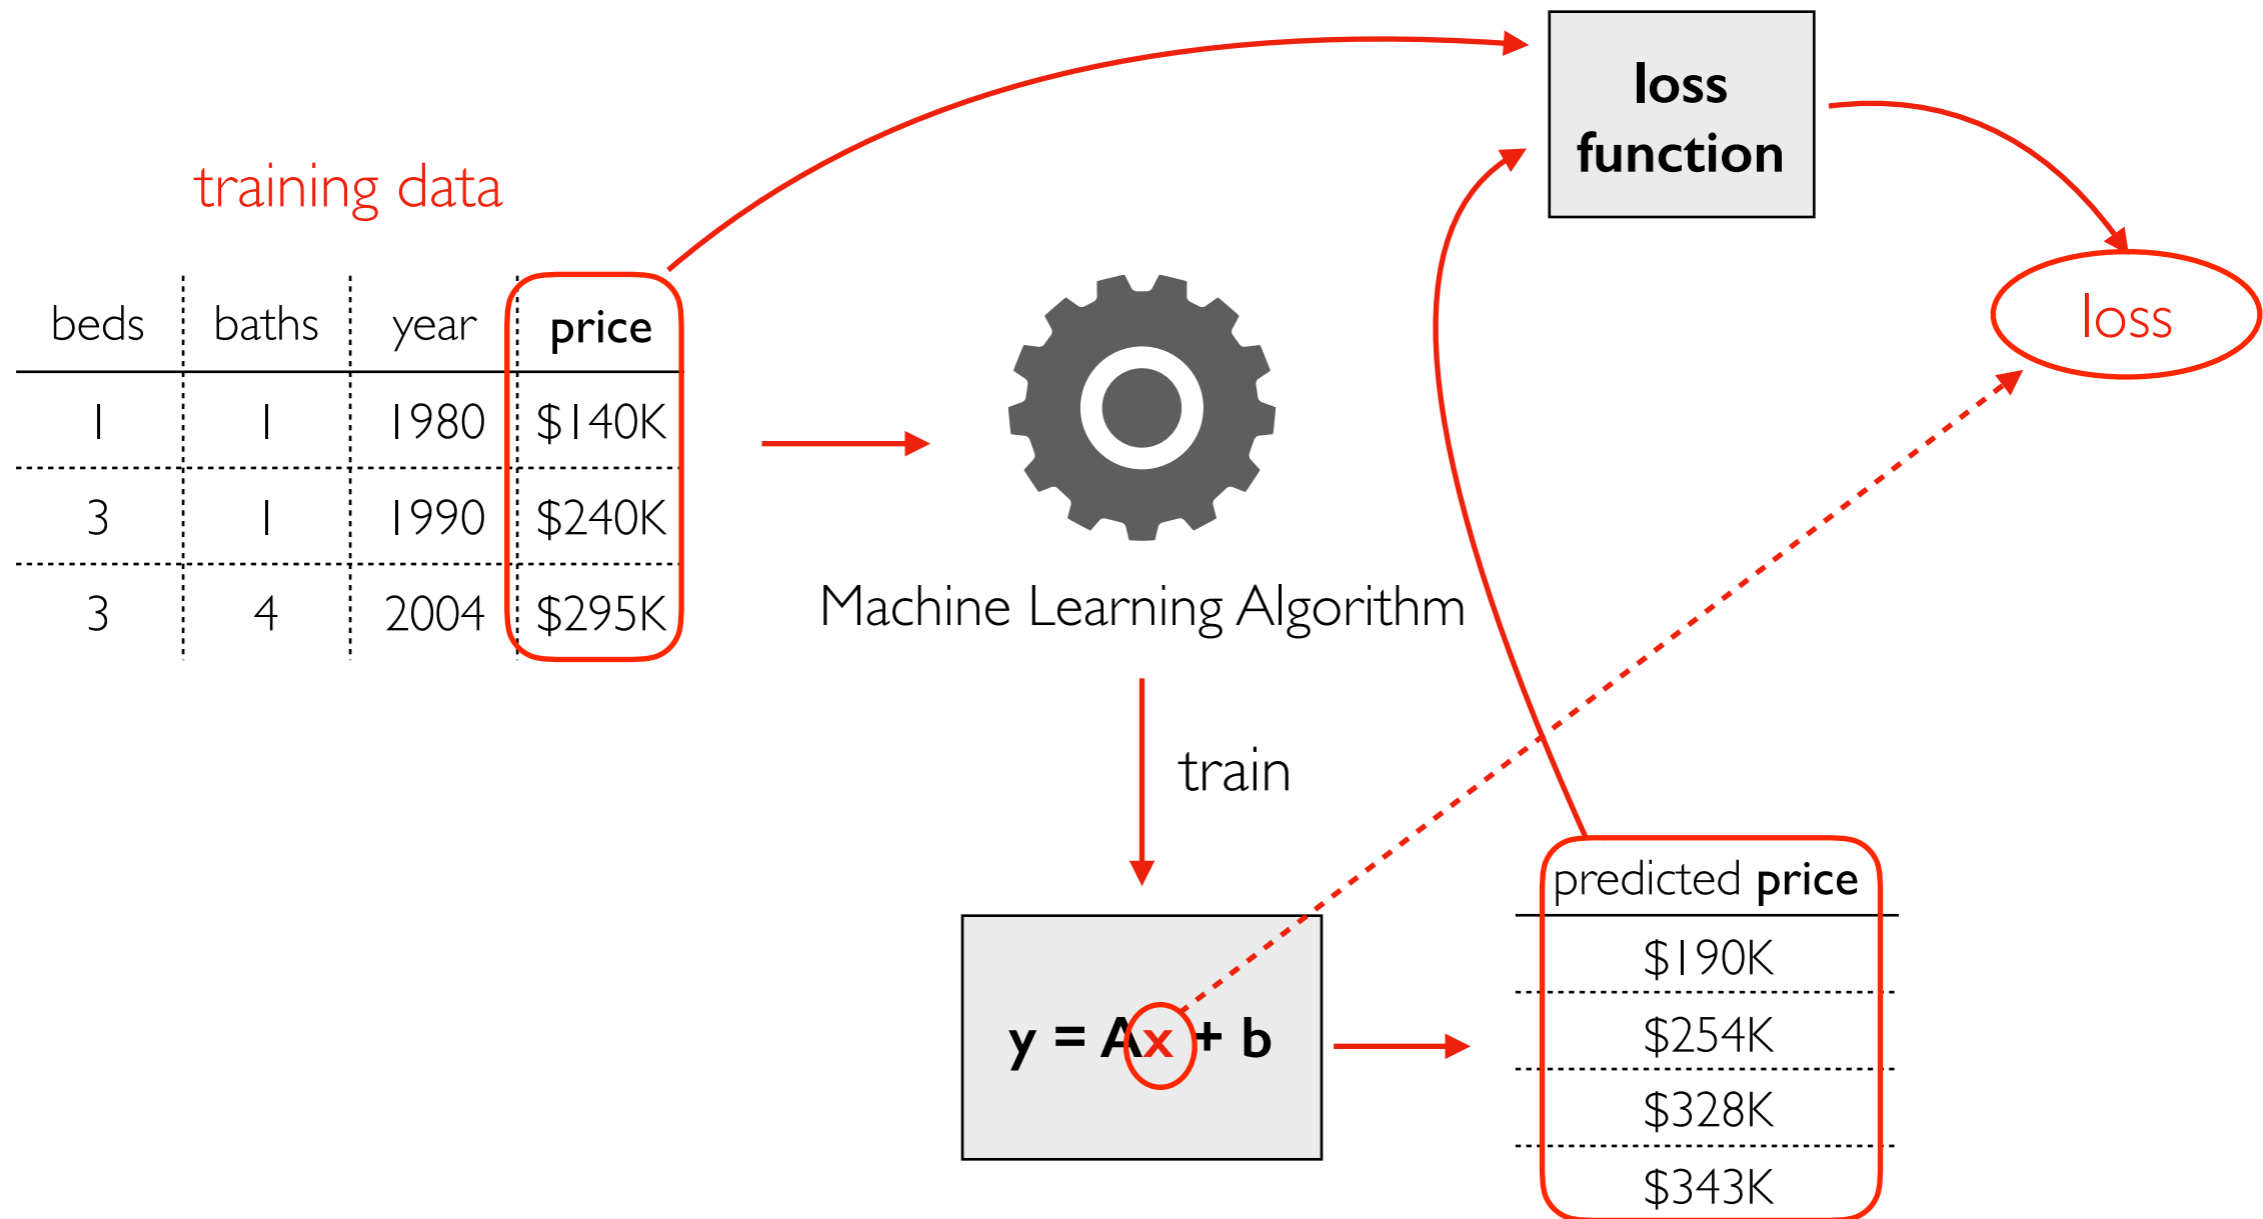
Calculus

training data

beds	baths	year	price
1	1	1980	\$140K
3	1	1990	\$240K
3	4	2004	\$295K



Calculus



how we we optimize x to minimize **loss**?
Important concepts: derivative, gradient

Parallelism

Parallelism

- doing multiple things at the same time
- requires multiple cores

GPUs (graphics processing units)

- graphics involves many of the same operation
- better to have many weaker cores working at once than fewer faster cores
- modern GPUs may have 1000s of cores (in contrast to 10s for CPUs)

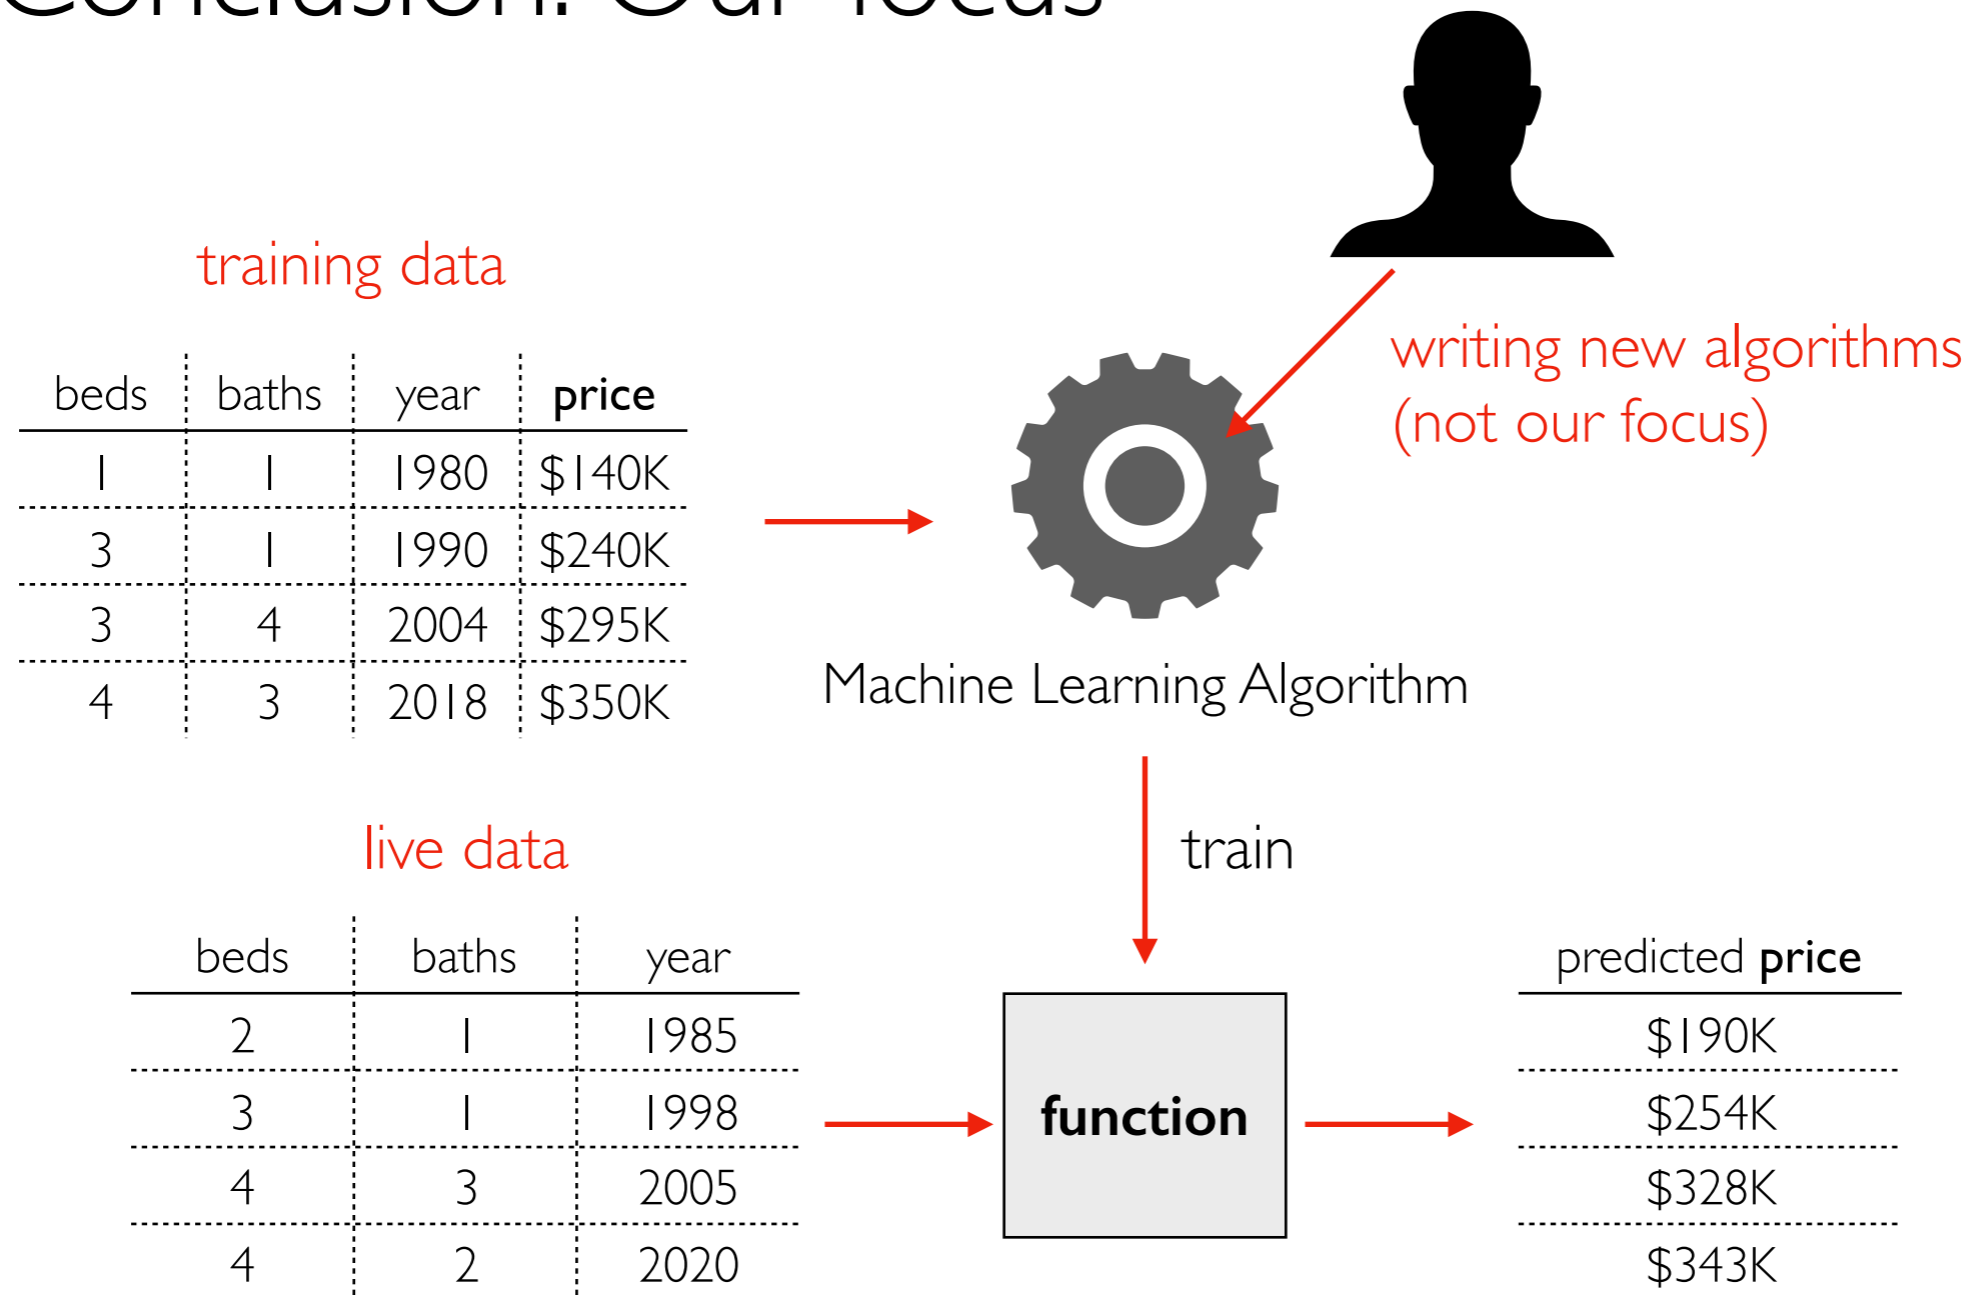
Scientific Computing

- GPUs can greatly speed up key ML operations
 - multiplying matrices
 - computing gradients
- We'll learn `pytorch` for this...



Practioners

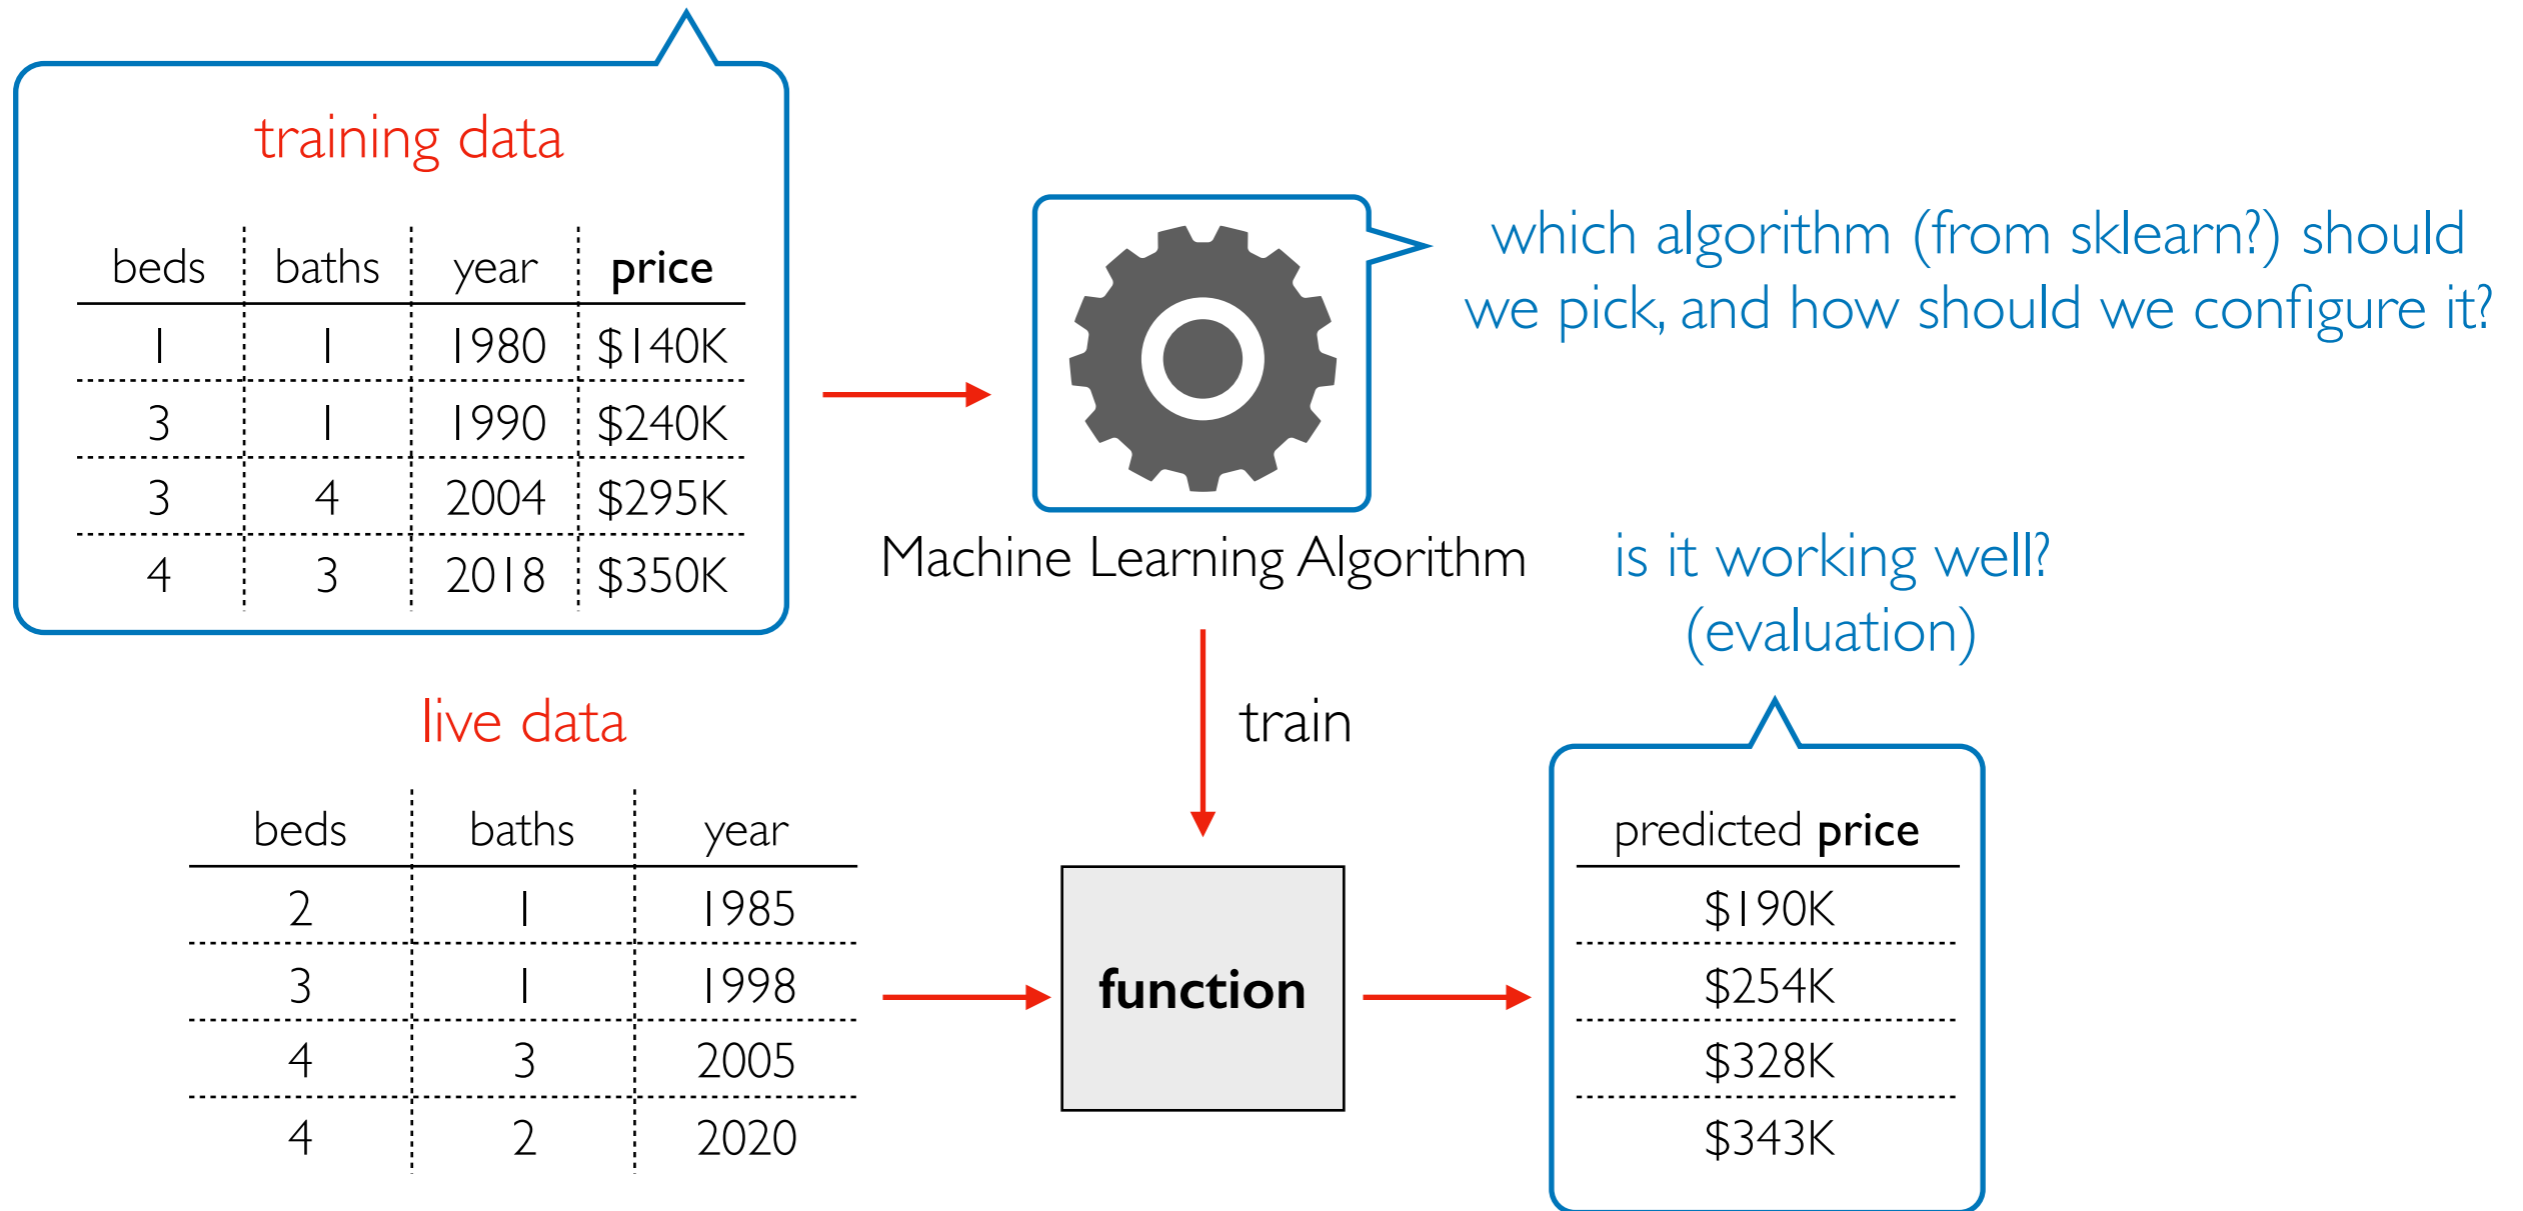
Conclusion: Our focus



many functions are **models** that can be used to predict

Conclusion: Our focus

how can we clean this up?



many functions are **models** that can be used to predict