

# Deep Learning

Introduction to TensorFlow and Keras

Prof. Dr. Jan Kirenz  
HdM Stuttgart

# ARTIFICIAL INTELLIGENCE

Early artificial intelligence stirs excitement.



# MACHINE LEARNING

Machine learning begins to flourish.



# DEEP LEARNING

Deep learning breakthroughs drive AI boom.



1950's

1960's

1970's

1980's

1990's

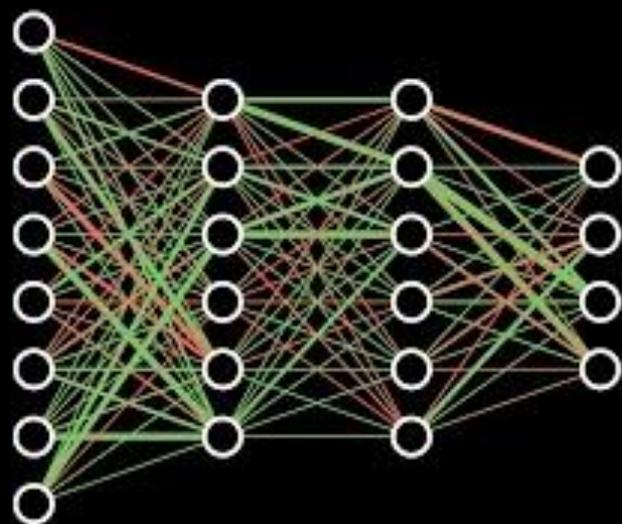
2000's

2010's

Since an early flush of optimism in the 1950s, smaller subsets of artificial intelligence – first machine learning, then deep learning, a subset of machine learning – have created ever larger disruptions.



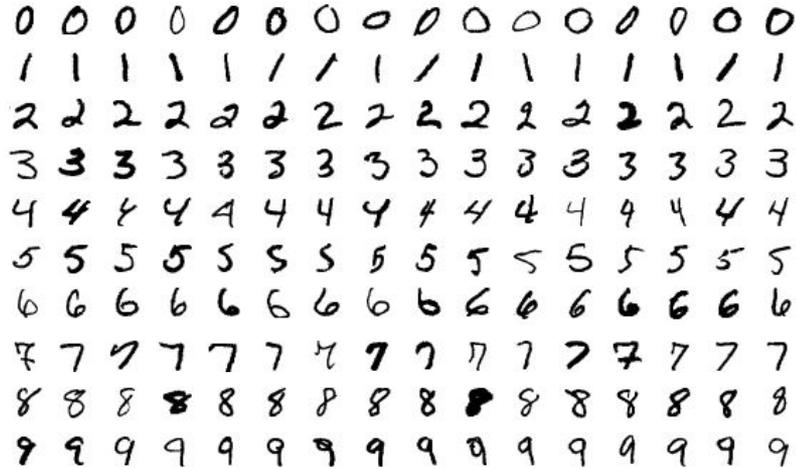
# Neural Networks



From the  
ground up

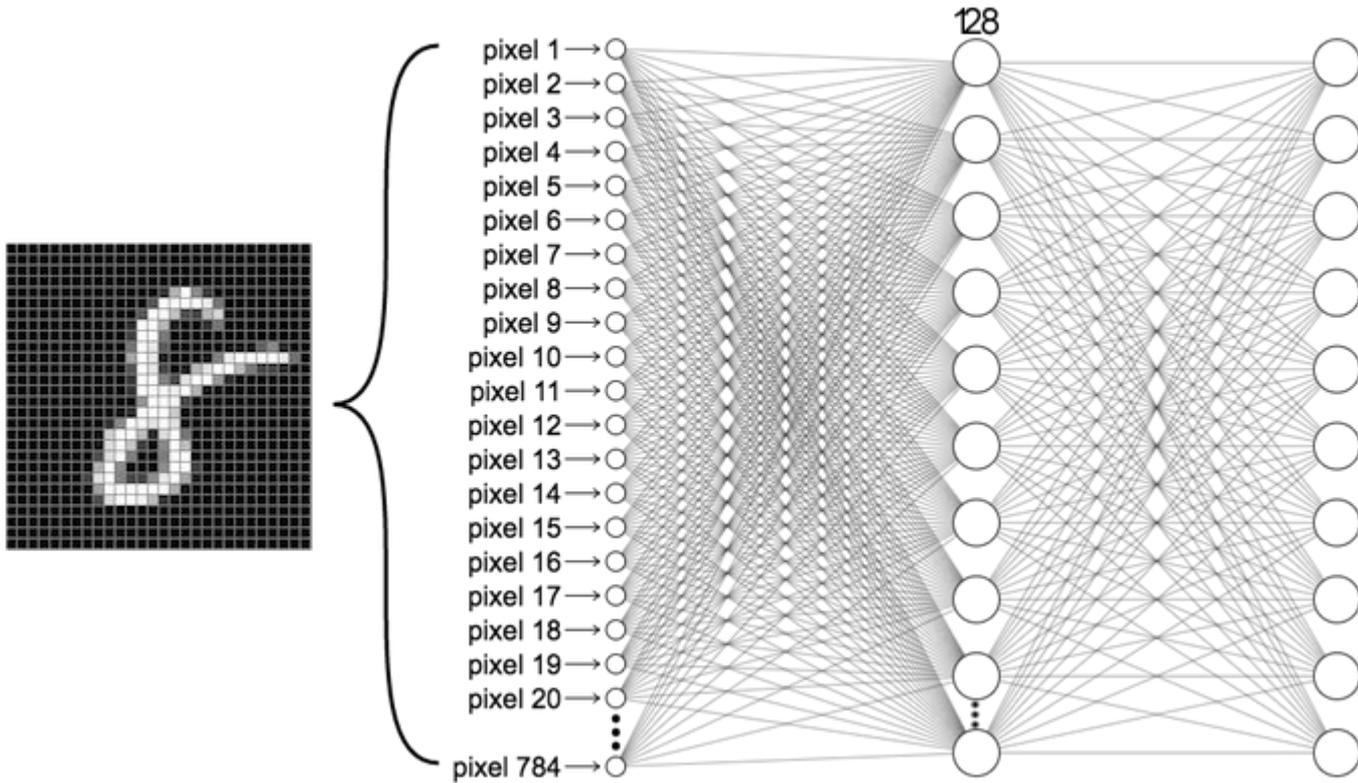


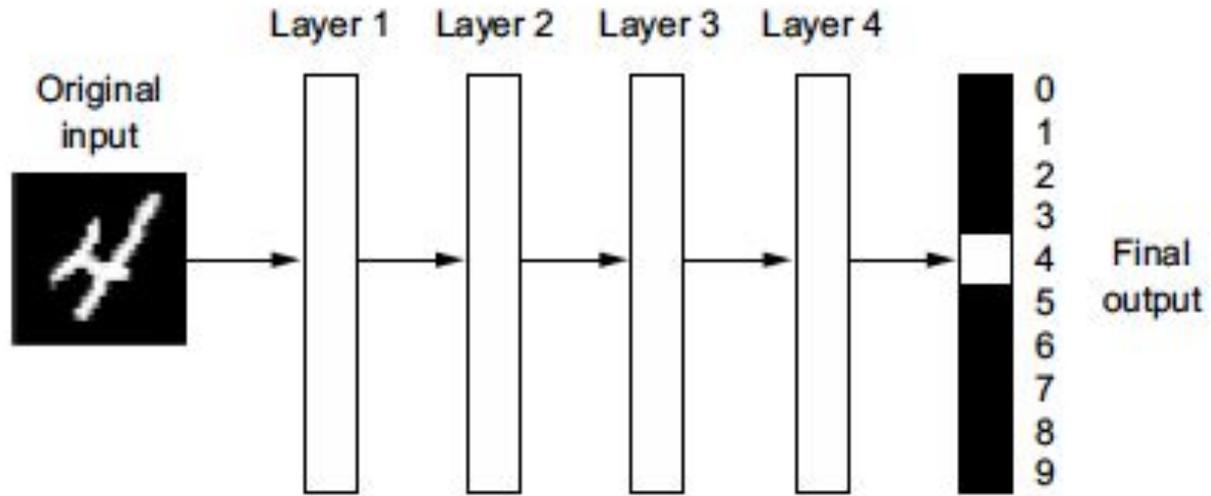
# MNIST database

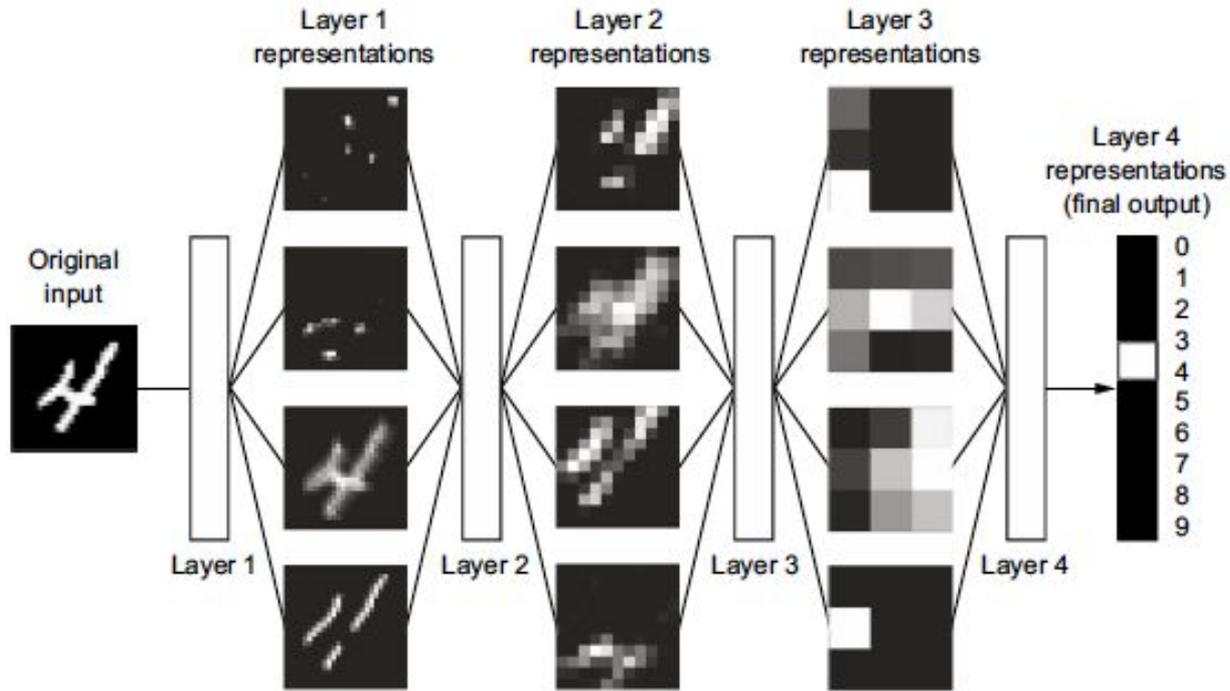


**MNIST database** (*Modified National Institute of Standards and Technology database*<sup>[1]</sup>) is a large database of handwritten digits that is commonly used for training various image processing systems









**Figure 1.6** Data representations learned by a digit-classification model

# Tinker With a **Neural Network** Right Here in Your Browser. Don't Worry, You Can't Break It. We Promise.

Epoch: 000,000 | Learning rate: 0.03 | Activation: Tanh | Regularization: None | Regularization rate: 0 | Problem type: Classification

## DATA

Which dataset do you want to use?



Ratio of training to test data: 50%



Noise: 0



Batch size: 10

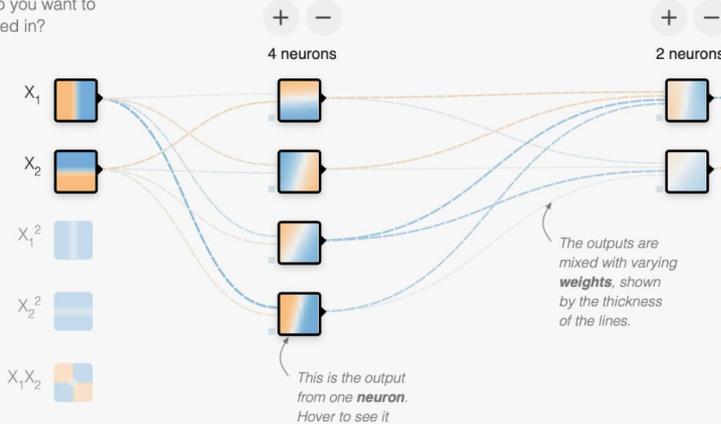


## FEATURES

Which properties do you want to feed in?

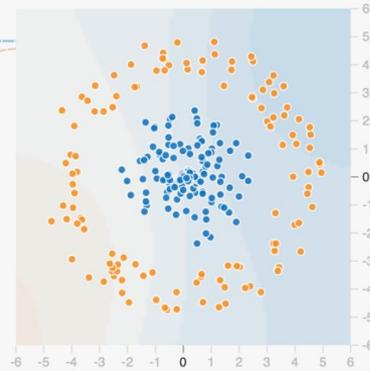
- $X_1$
- $X_2$
- $X_1^2$
- $X_2^2$
- $X_1 X_2$

## 2 HIDDEN LAYERS



## OUTPUT

Test loss 0.496  
Training loss 0.501



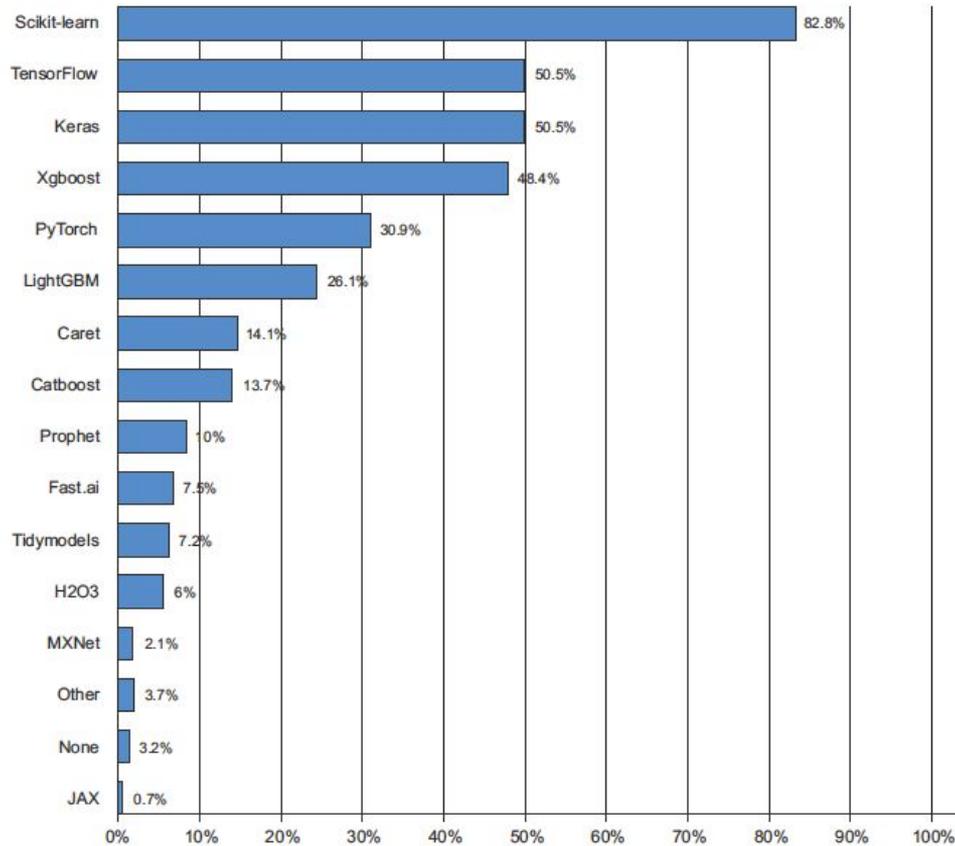


Figure 1.13 Tool usage across the machine learning and data science industry (Source: [www.kaggle.com/kaggle-survey-2020](http://www.kaggle.com/kaggle-survey-2020))

From 2016 to 2020, the entire machine learning and data science industry has been dominated by two approaches: deep learning and gradient boosted trees.



Go from zero to hero with web ML in a new online course from TensorFlow.js.

[Register now](#)

# An end-to-end open source machine learning platform

**TensorFlow**

[For JavaScript](#)

[For Mobile & IoT](#)

[For Production](#)

The core open source library to help you develop and train ML models. Get started quickly by running Colab notebooks directly in your browser.

[Get started with TensorFlow](#)



## Companies using TensorFlow

[See case studies →](#)





Simple. Flexible. Powerful.

Get started

API docs

Guides

Examples

```
from tensorflow import keras
from tensorflow.keras import layers

# Instantiate a trained vision model
vision_model = keras.applications.ResNet50()

# This is our video encoding branch using the trained vision_model
video_input = keras.Input(shape=(100, None, None, 3))
encoded_frame_sequence = layers.TimeDistributed(vision_model)(video_input)
encoded_video = layers.LSTM(256)(encoded_frame_sequence)

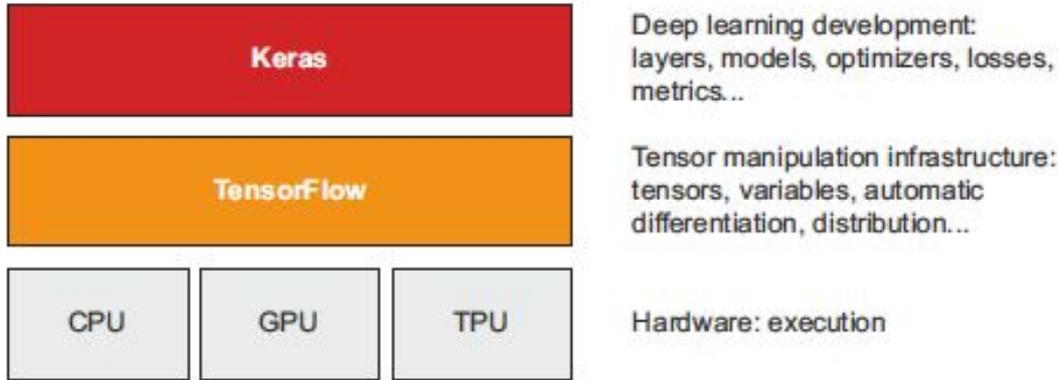
# This is our text-processing branch for the question input
question_input = keras.Input(shape=(100,), dtype='int32')
embedded_question = layers.Embedding(10000, 256)(question_input)
encoded_question = layers.LSTM(256)(embedded_question)

# And this is our video question answering model:
merged = keras.layers.concatenate([encoded_video, encoded_question])
output = keras.layers.Dense(1000, activation='softmax')(merged)
video_qa_model = keras.Model(inputs=[video_input, question_input],
                              outputs=output)
```



## A vast ecosystem.

Keras is a central part of the tightly-connected TensorFlow 2 ecosystem, covering every step of the machine learning workflow, from data management to hyperparameter training to deployment solutions.

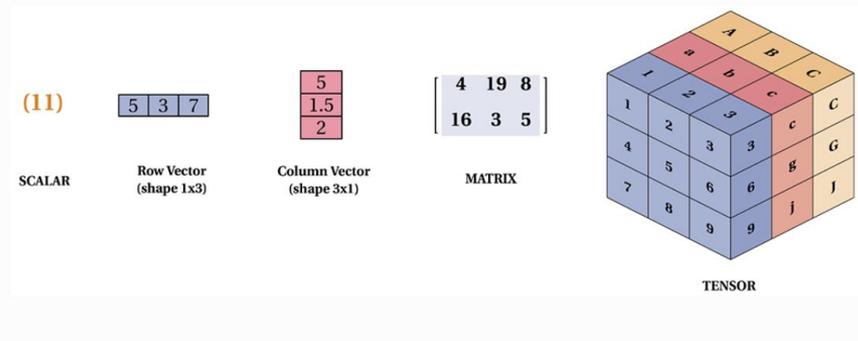


**Figure 3.1 Keras and TensorFlow: TensorFlow is a low-level tensor computing platform, and Keras is a high-level deep learning API**

# What is a Tensor?

## What is Tensor?

Tensor holds a multi-dimensional array of elements of a single data type which is very similar with numpy's ndarray. When the dimension is zero, it can be called a scalar. When the dimension is 2, it can be called a matrix. When the dimension is greater than 2, it is usually called a tensor. If you are very familiar with numpy, then understanding Tensor will be quite easy.



Go from zero to hero with web ML in a new online course from TensorFlow.js.

Register now

TensorFlow > Learn > TensorFlow Core > Guide

War das hilfreich?

## Introduction to Tensors

Auf dieser Seite

- Basics
- About shapes
- Indexing
  - Single-axis indexing
  - Multi-axis indexing
- ...

Run in  
Google Colab

View source on  
GitHub

Download  
notebook

# A tensor is an N-dimensional array of data



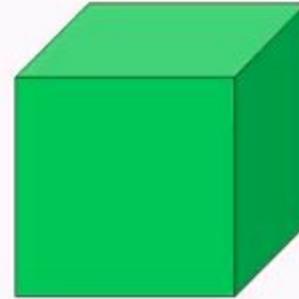
Rank 0  
Tensor  
scalar



Rank 1  
Tensor  
vector



Rank 2  
Tensor  
matrix



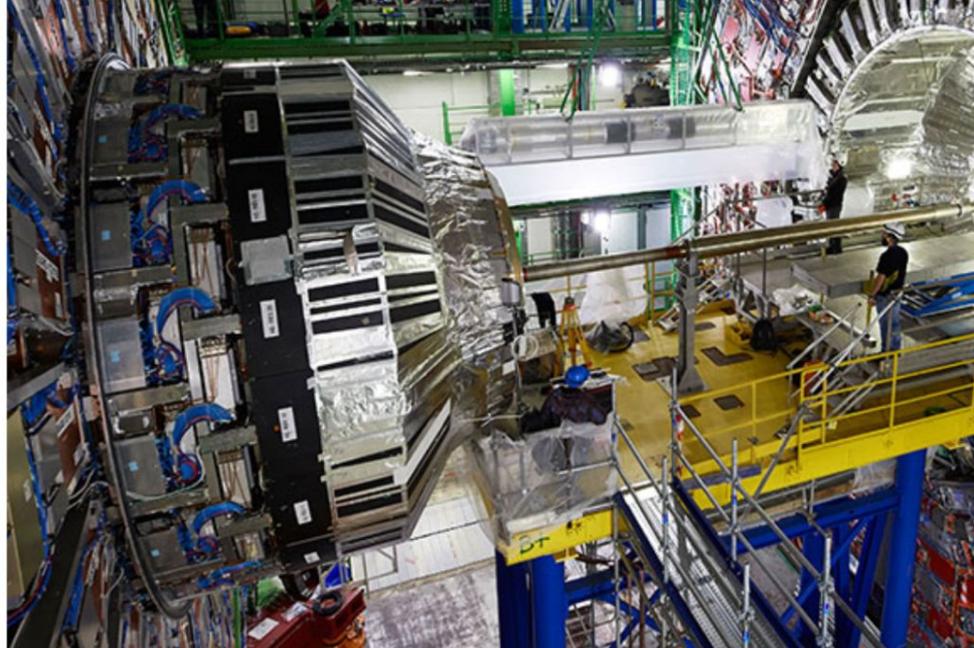
Rank 3  
Tensor



Rank 4  
Tensor

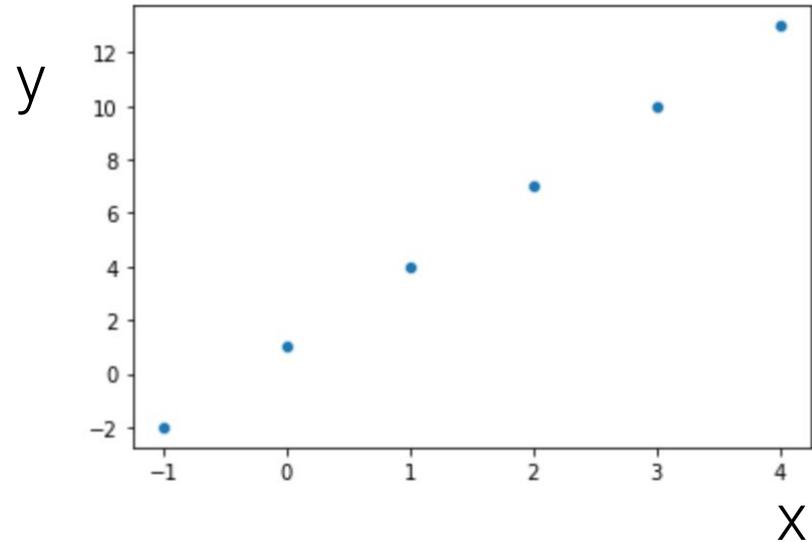
## State-of-the-art research.

Keras is used by CERN, NASA, NIH, and many more scientific organizations around the world (and yes, Keras is used at the LHC). Keras has the low-level flexibility to implement arbitrary research ideas while offering optional high-level convenience features to speed up experimentation cycles.



# Simple regression example with TensorFlow & Keras

| Variable | Value 1 | Value 2 | Value 3 | Value 4 | Value 5 | Value 6 |
|----------|---------|---------|---------|---------|---------|---------|
| x        | -1      | 0       | 1       | 2       | 3       | 4       |
| y        | -2      | 1       | 4       | 7       | 10      | 13      |



# import libraries

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

## # Data

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

## # Model definition

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

## # Model fitting

```
model.fit(x, y, epochs=50)
```

## # Model prediction

```
print(model.predict([10.0]))
```

# create data

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

## # Data

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

## # Model definition

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

## # Model fitting

```
model.fit(x, y, epochs=50)
```

## # Model prediction

```
print(model.predict([10.0]))
```

# define model architecture

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

## # Data

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

## # Model definition

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

## # Model fitting

```
model.fit(x, y, epochs=50)
```

## # Model prediction

```
print(model.predict([10.0]))
```

we use a single layer

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

```
# Data
```

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

```
# Model definition
```

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

```
# Model fitting
```

```
model.fit(x, y, epochs=50)
```

```
# Model prediction
```

```
print(model.predict([10.0]))
```

with one neuron

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

```
# Data
```

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

```
# Model definition
```

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

```
# Model fitting
```

```
model.fit(x, y, epochs=50)
```

```
# Model prediction
```

```
print(model.predict([10.0]))
```

and only one input x

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

```
# Data
```

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

```
# Model definition
```

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

```
# Model fitting
```

```
model.fit(x, y, epochs=50)
```

```
# Model prediction
```

```
print(model.predict([10.0]))
```

we compile the model

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

```
# Data
```

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

```
# Model definition
```

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

```
# Model fitting
```

```
model.fit(x, y, epochs=50)
```

```
# Model prediction
```

```
print(model.predict([10.0]))
```

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

generate a guess for y  
sgd = stochastic gradient descent

### # Data

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

### # Model definition

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

### # Model fitting

```
model.fit(x, y, epochs=50)
```

### # Model prediction

```
print(model.predict([10.0]))
```

and calculate the error

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

```
# Data
```

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

```
# Model definition
```

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

```
# Model fitting
```

```
model.fit(x, y, epochs=50)
```

```
# Model prediction
```

```
print(model.predict([10.0]))
```

we do this 50 times

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

```
# Data
```

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

```
# Model definition
```

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

```
# Model fitting
```

```
model.fit(x, y, epochs=50)
```

```
# Model prediction
```

```
print(model.predict([10.0]))
```

predict y for x=10

```
import numpy as np
import tensorflow as tf
from tensorflow import keras
```

```
# Data
```

```
x = np.array([-1.0, 0.0, 1.0, 2.0, 3.0, 4.0], dtype=float)
y = np.array([-2.0, 1.0, 4.0, 7.0, 10.0, 13.0], dtype=float)
```

```
# Model definition
```

```
model = tf.keras.Sequential([keras.layers.Dense(units=1, input_shape=[1])])
model.compile(optimizer='sgd', loss='mean_squared_error')
```

```
# Model fitting
```

```
model.fit(x, y, epochs=50)
```

```
# Model prediction
```

```
print(model.predict([10.0]))
```

# Resources

The slides are based on the excellent video tutorial “Intro to Machine Learning (ML Zero to Hero - Part 1)” by Lawrence Moroney.

