Convolutional Neural Networks

L01 - Deep Learning for Life Sciences Course

Convolutional Neural Networks (CNNs)

Visual Recognition

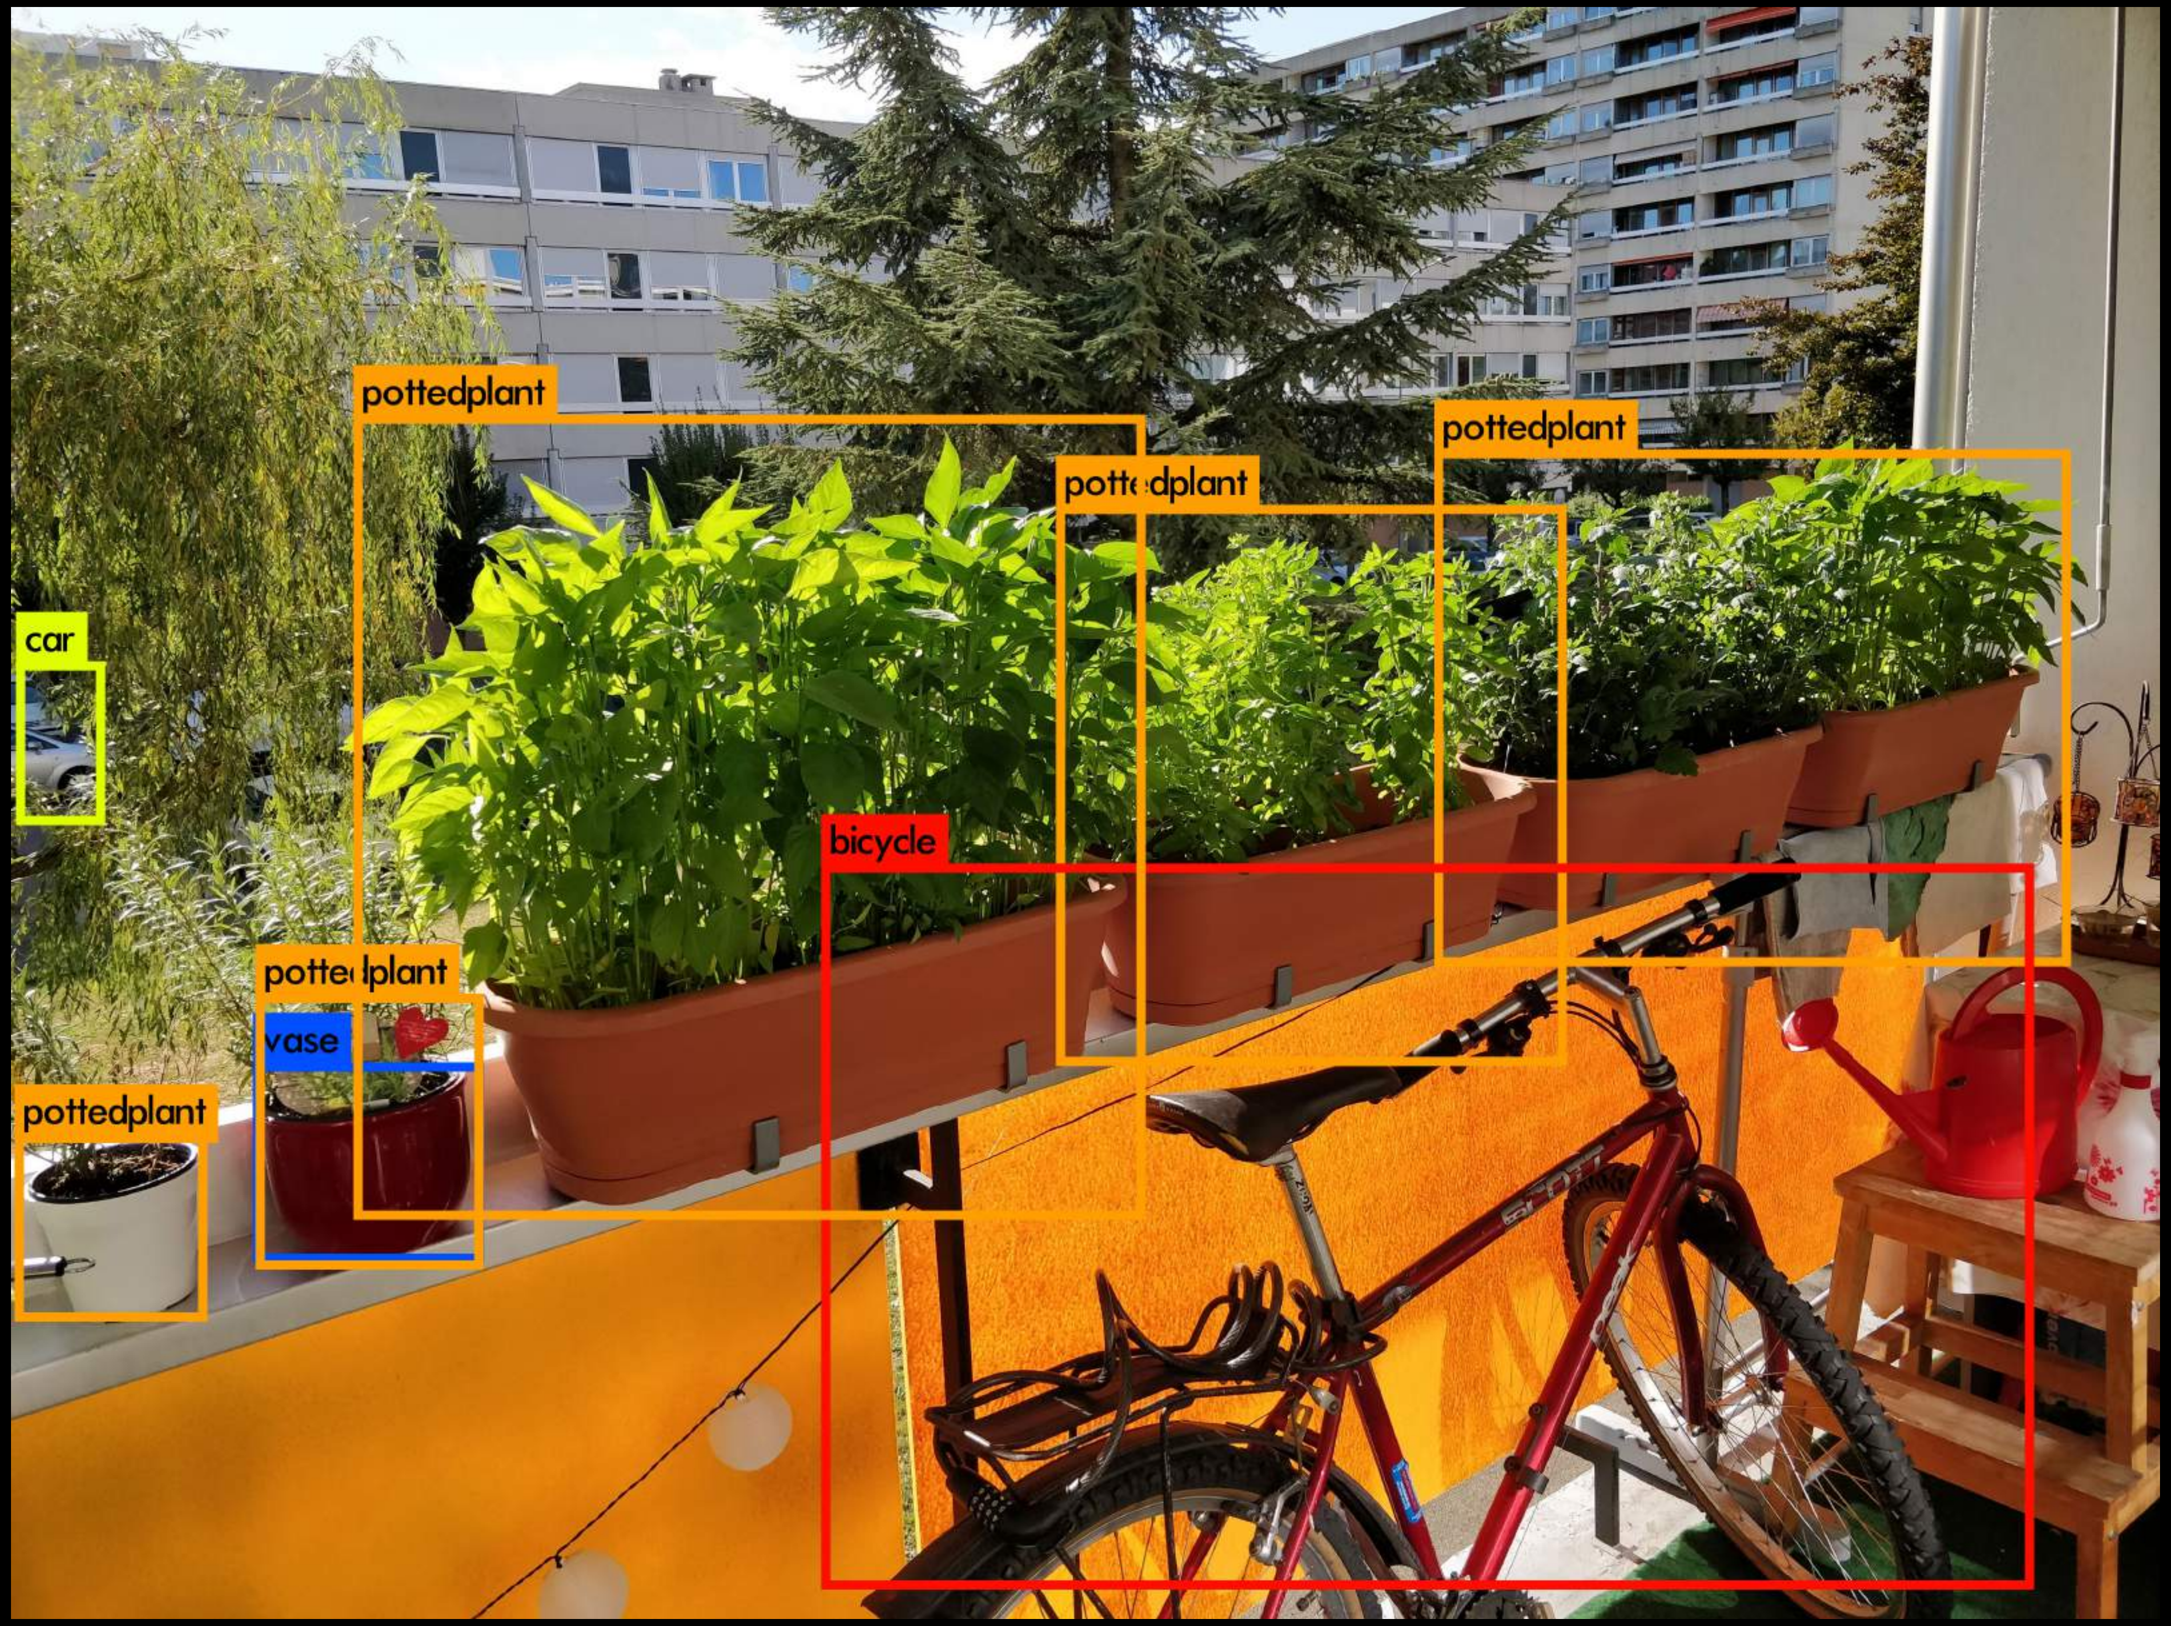

Redmon et al., 2016. *You only look once: unified, real-time object detection.*

Convolutional Neural Networks (CNNs)

Visual Recognition

- Facial Identification
- Medical Image Analysis
- Drug Design

High-Level Feature Extraction

Can you identify key features in each image category ?

High-Level Feature Extraction

Can you identify key features in each image category ?

- Eyes
- Nose
- Lips

High-Level Feature Extraction

Can you identify key features in each image category ?

- Eyes
- Nose
- Lips

- Wheels
- Windshields
- Headlights

High-Level Feature Extraction

Can you identify key features in each image category ?

- Eyes
- Nose
- Lips

- Wheels
- Windshields
- Headlights

- Doors
- Windows
- Roofs

Manual Feature Extraction

Traditional Rule-Based Methods:

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge →

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge → Define Features

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge → Define Features →

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge → Define Features → Detect Features to Classify

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge → Define Features → Detect Features to Classify

Problems ?

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge → Define Features → Detect Features to Classify

Problems ?

- Viewpoint variation
- Scale variation
- Occlusion
- Deformation

Manual Feature Extraction

Traditional Rule-Based Methods:

Domain Knowledge → Define Features → Detect Features to Classify

Problems ?

- Viewpoint variation
- Scale variation
- Occlusion
- Deformation
- Background clutter
- Illumination conditions
- Variation
- Etc

Convolutional Neural Networks (CNNs)

Solution

- Make the machine learn the features **by itself**
- Take into account the spatial proximity of features

Convolutional Neural Networks (CNNs)

Solution

- Make the machine learn the features **by itself**
- Take into account the spatial proximity of features

$$\downarrow$$

Convolutional Neural Networks (CNNs)

Solution

- Make the machine learn the features **by itself**
- Take into account the spatial proximity of features

↓

Convolutions

Sliding Window

Input

```
ada love
```

Sliding Window

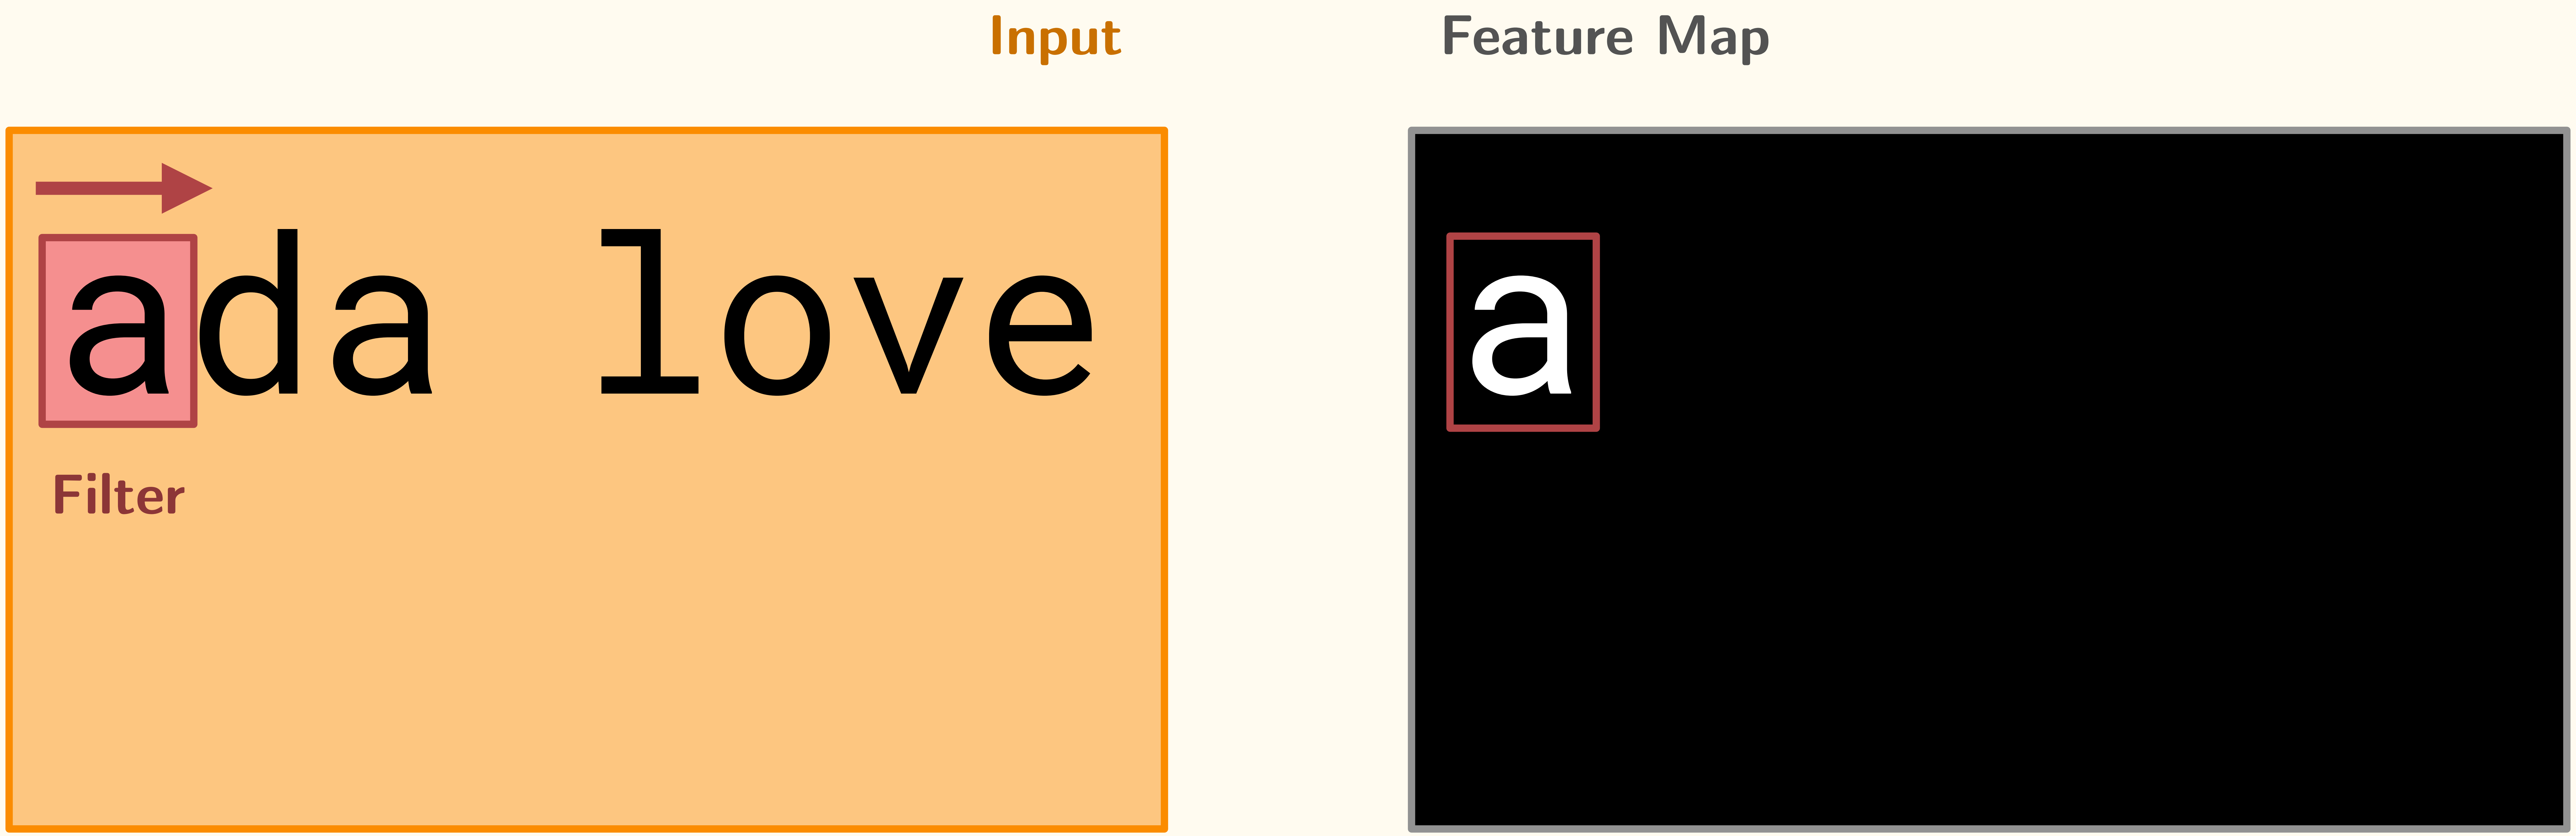

Sliding Window

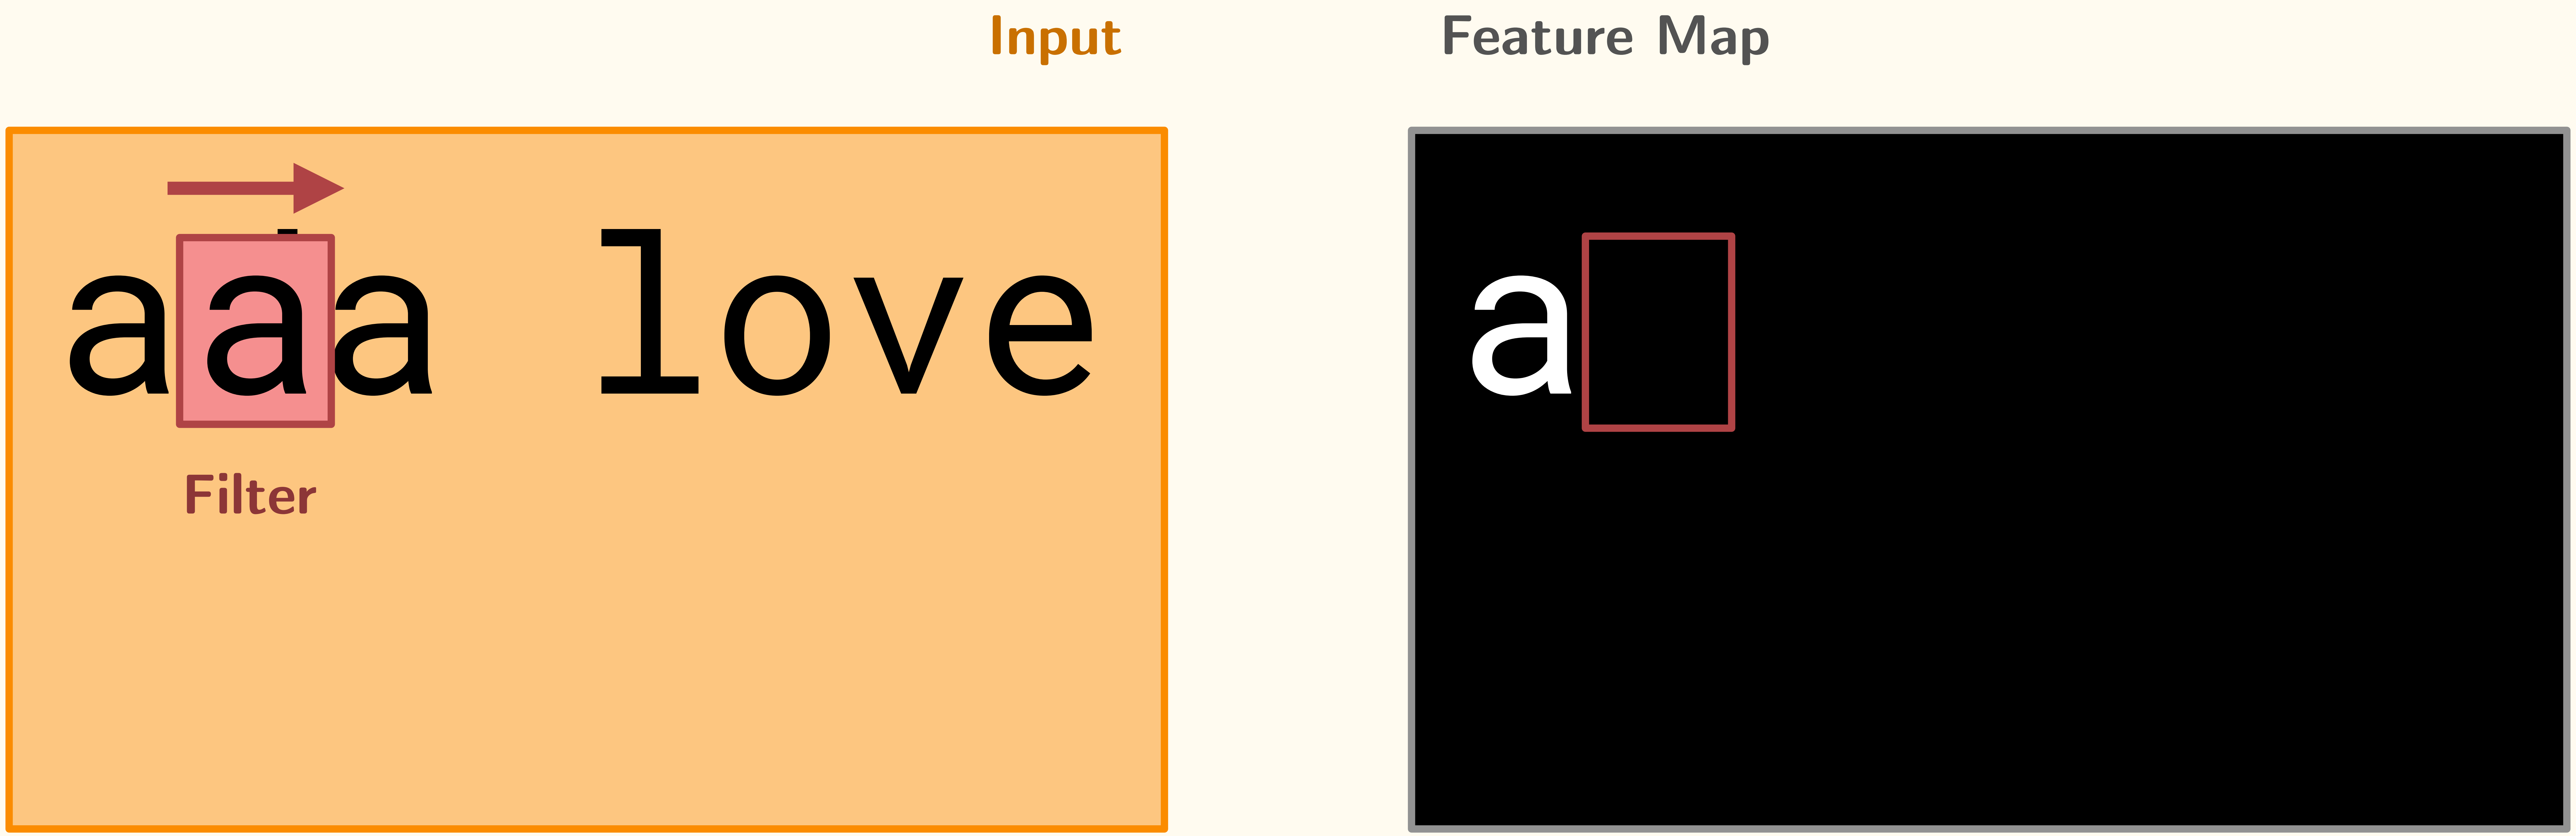

Sliding Window

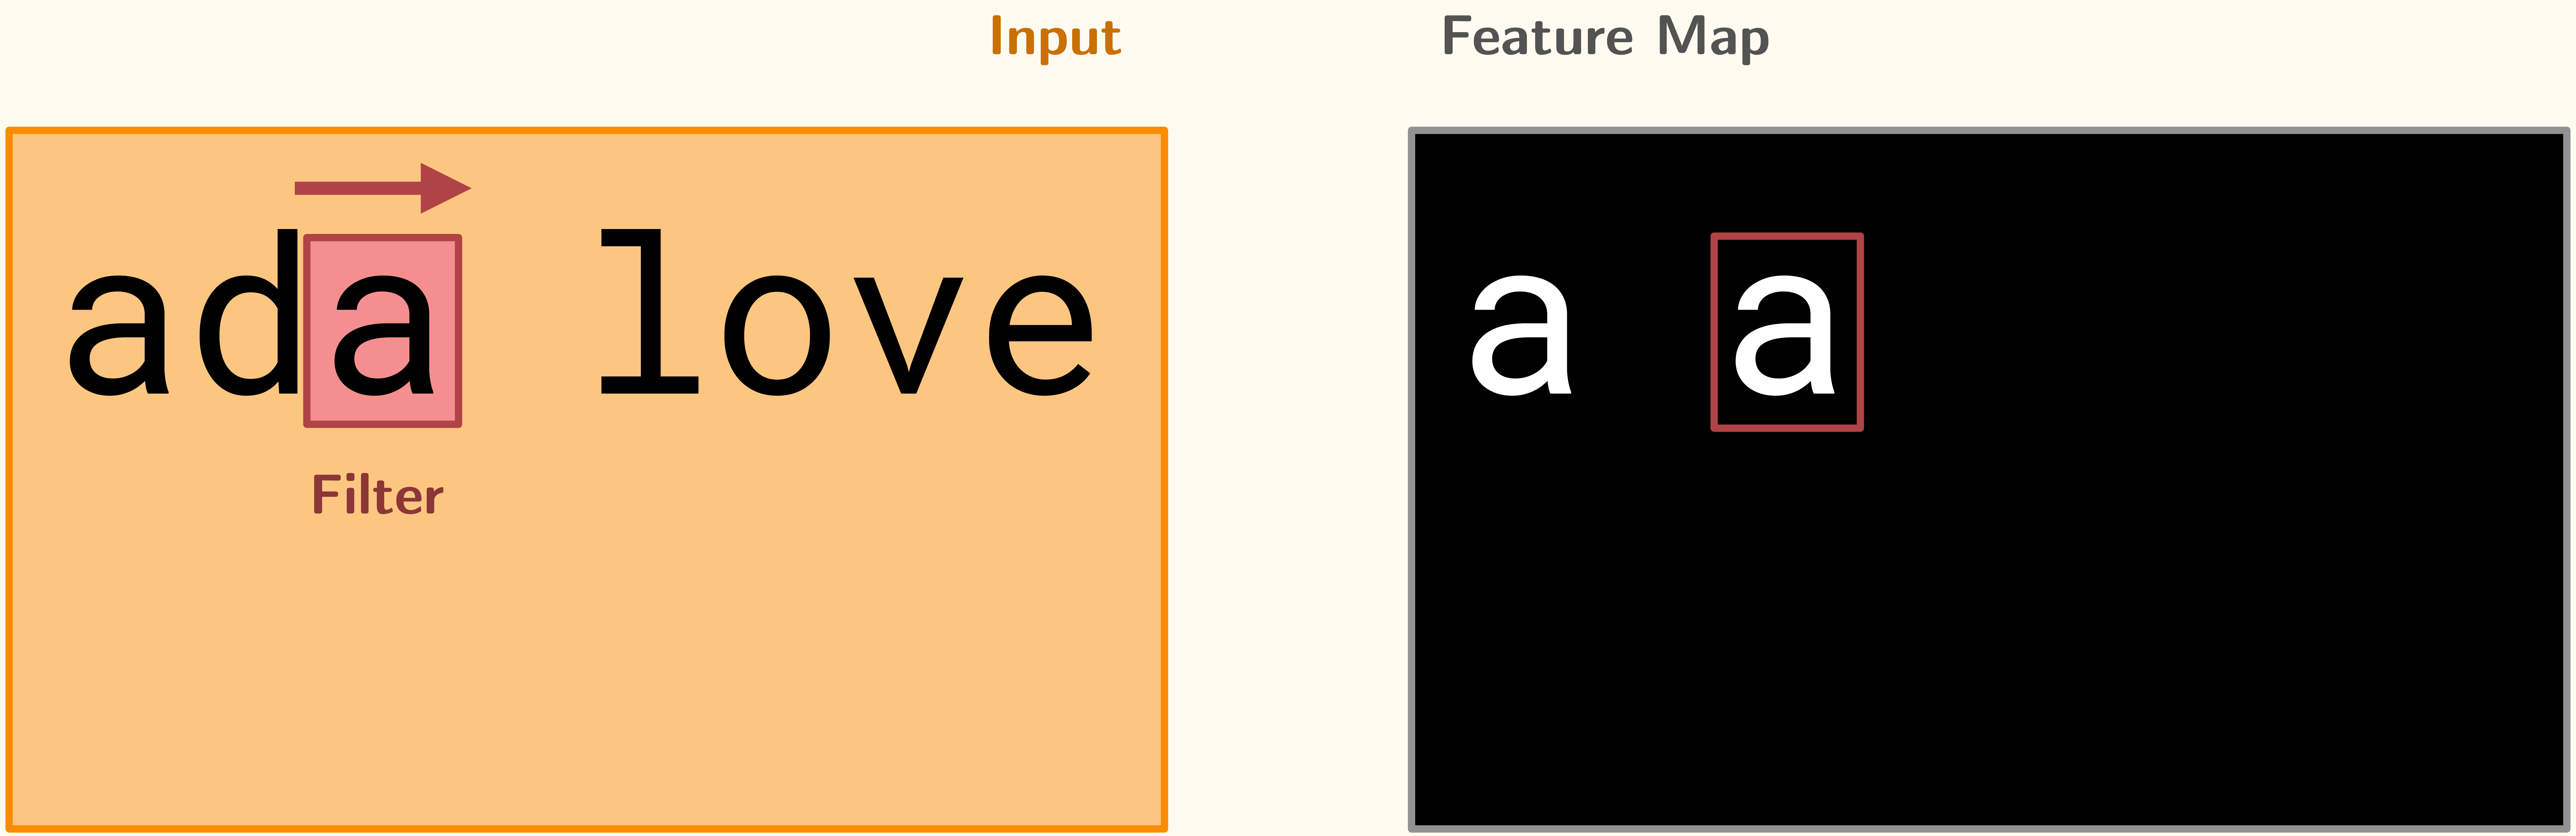

First-Layer Convolution

We use filters to extract local features.

ada lovelace was a mathematician and writer who lived in the 19th century. lovelace holds the honor of having published the very first algorithm intended to be used by a machine to perform calculations, which make lovelace the first-ever computer programmer.

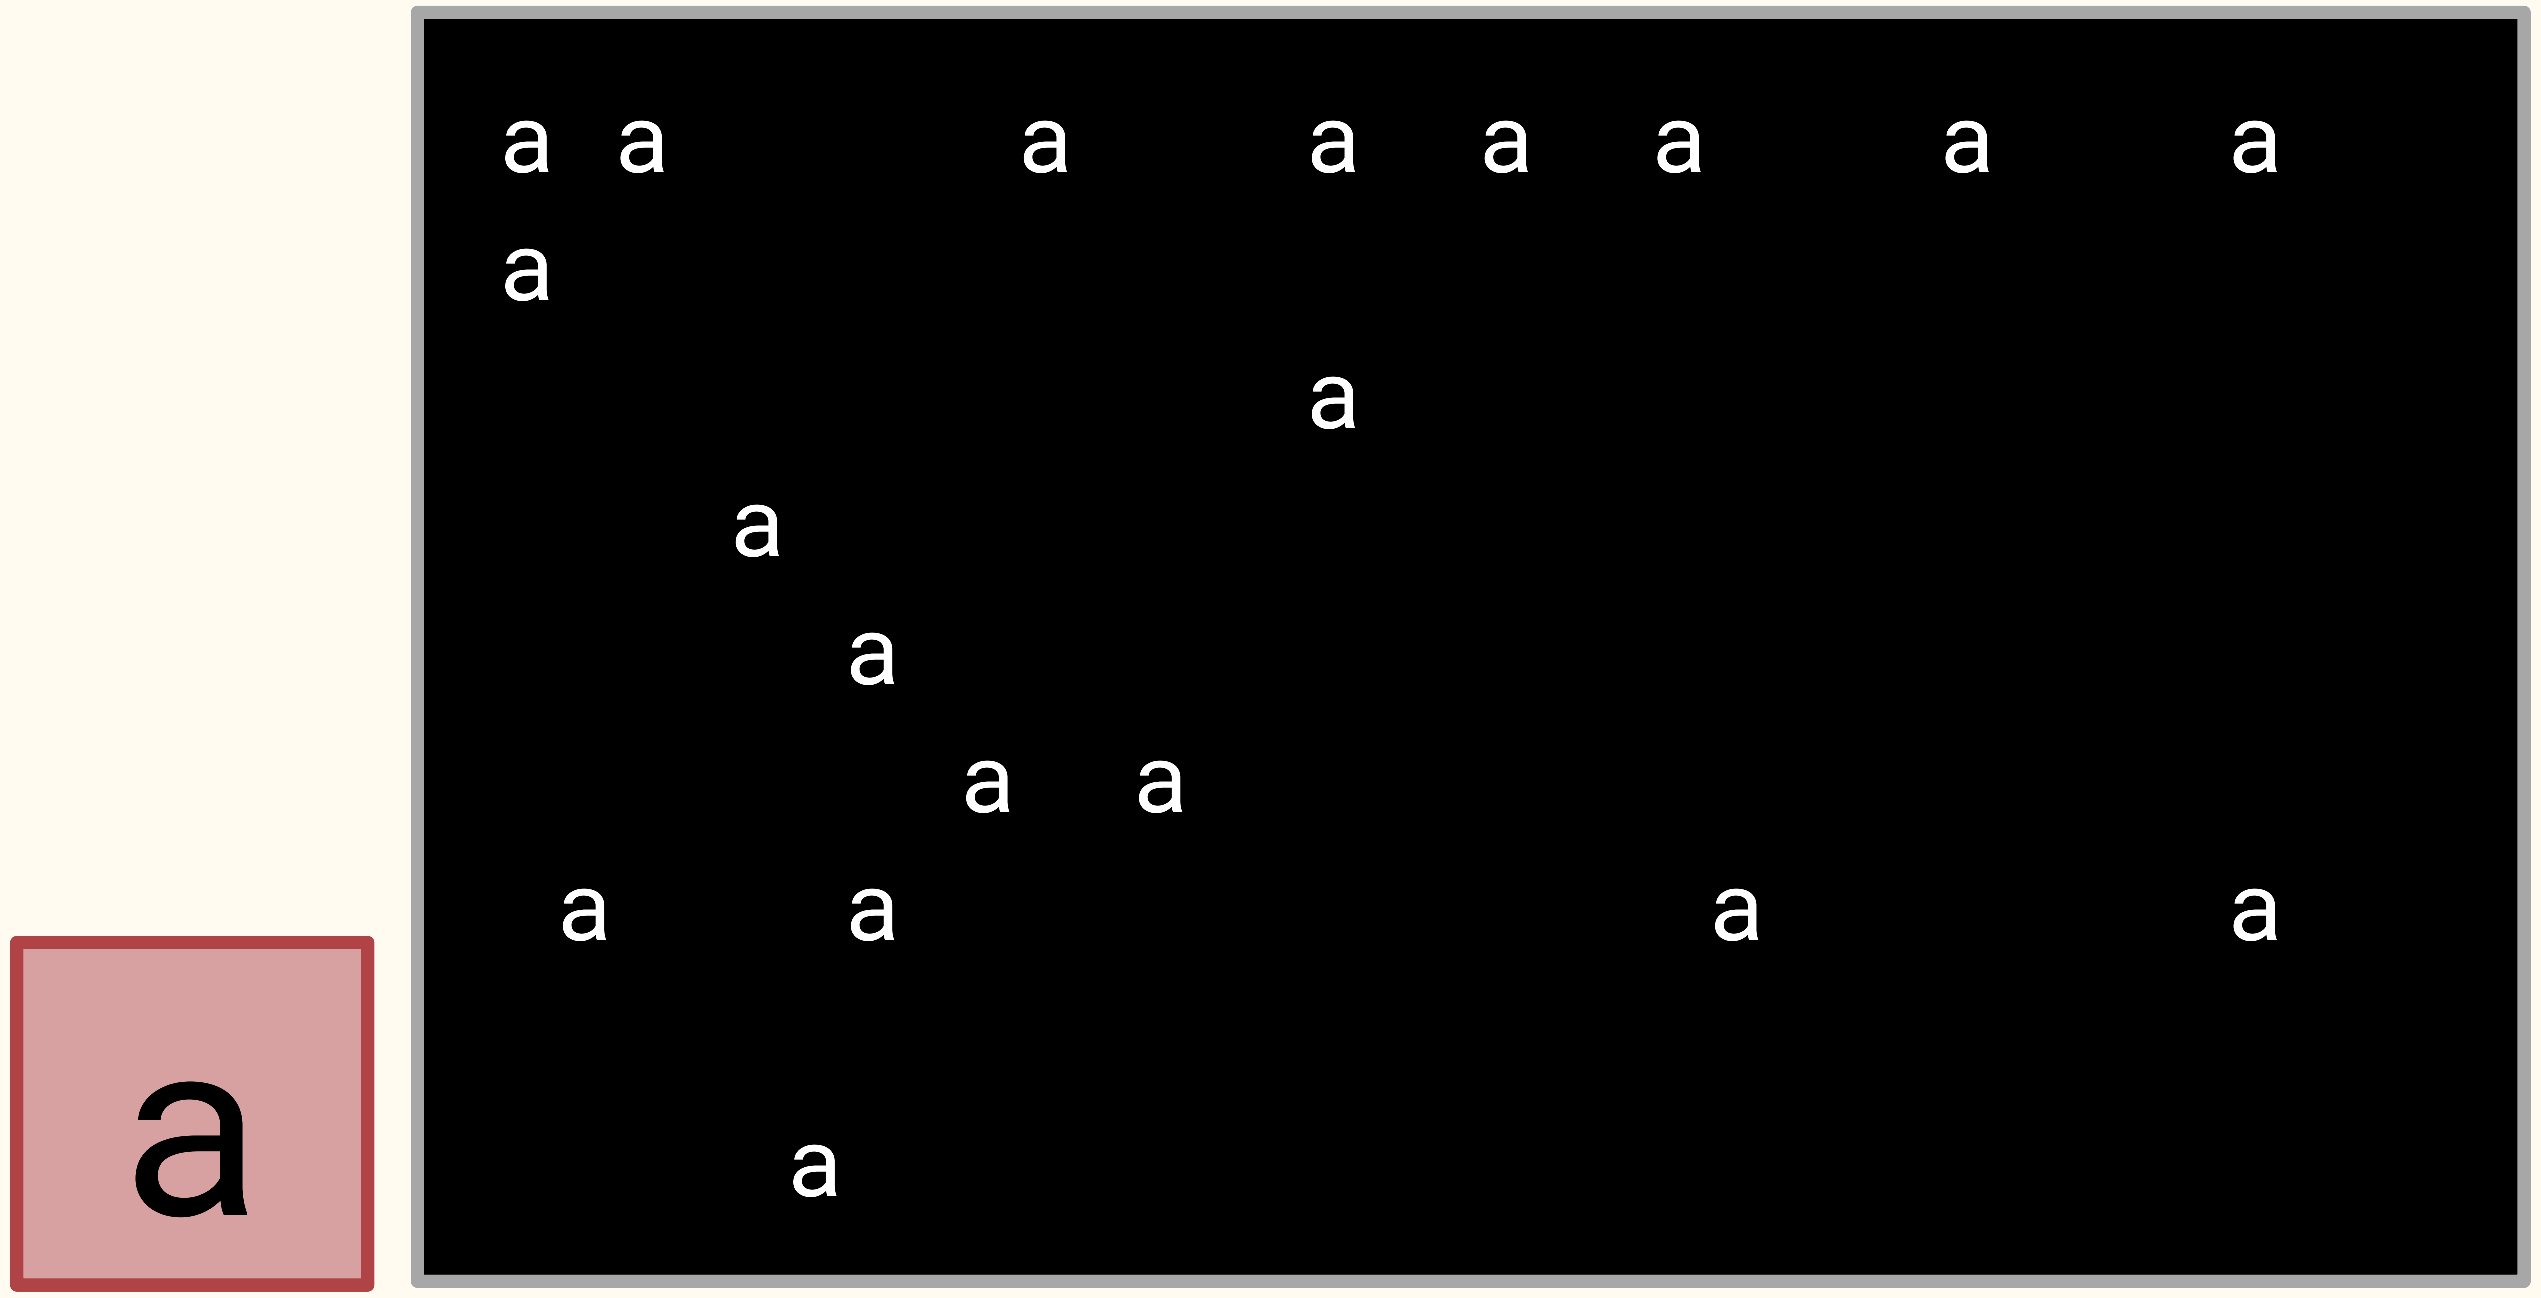

First-Layer Convolution

We use filters to extract local features.

Input

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

```
a
```

```
a a      a    a  a  a    a    a
a
              a
    a
      a
        a  a
 a    a              a        a

     a
```

First-Layer Convolution

We use filters to extract local features.

Input **Filter**

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

```
a
```

```
a a      a    a  a  a    a    a
a
              a
    a
      a
        a  a
 a    a              a        a

     a
```

First-Layer Convolution

We use filters to extract local features.

Input **Filter** **Feature Map**

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

```
a
```

```
a a      a    a  a  a    a    a
a
              a
    a
      a
        a  a
 a    a              a        a

     a
```

First-Layer Convolution

We use filters to extract local features.

Input **Filter** **Feature Map**

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

d

```
 d
  d                 d
                     d
                  d
                     d d
   d
```

Second-Layer Convolution

Second-Layer Convolution

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

Input

Second-Layer Convolution

First-Layer Filters: a d

ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.

Input

Second-Layer Convolution

First-Layer Filters: a d

ada lovelace was a mathematician and writer who lived in the 19th century. lovelace holds the honor of having published the very first algorithm intended to be used by a machine to perform calculations, which make lovelace the first-ever computer programmer.

ada

Input **Filter**

Second-Layer Convolution

First-Layer Filters: a d

Stacked First-Layer Feature Maps

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

Input Filter

Second-Layer Convolution

We use **multiple** filters to extract different features.

Stacked First-Layer Feature Maps

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

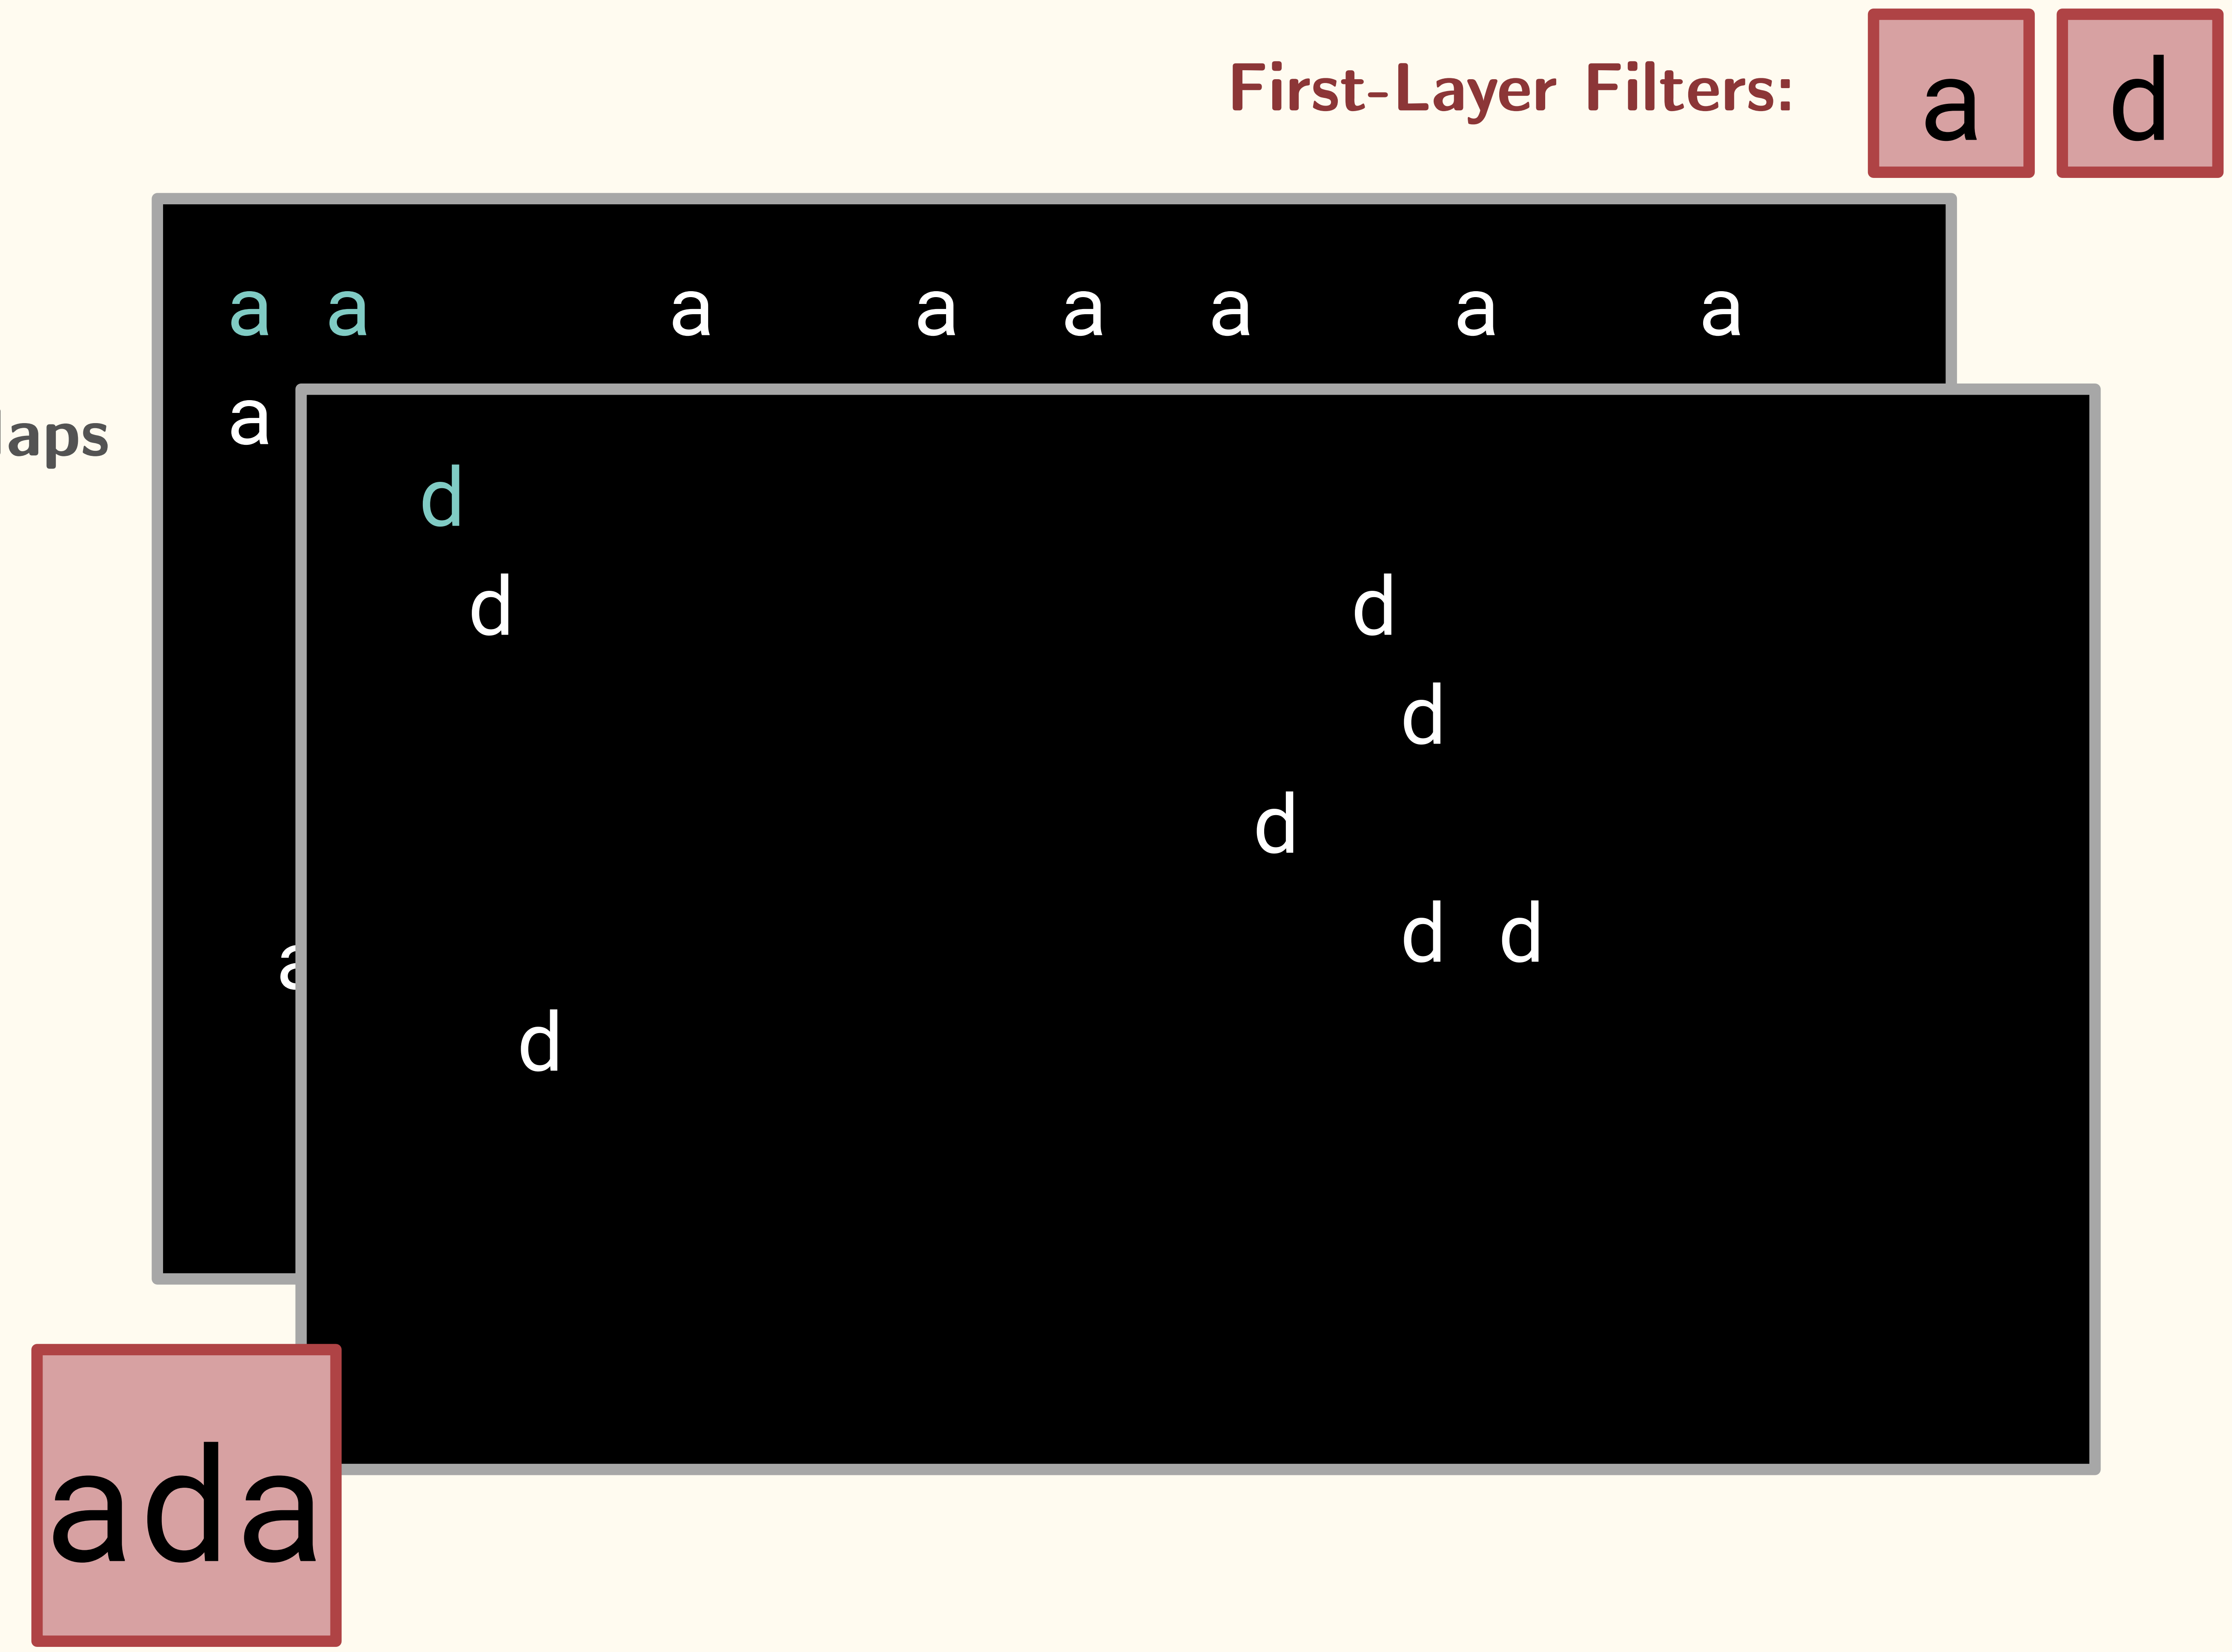

Input **Filter**

Second-Layer Convolution

We use **multiple** filters to extract different features.

Stacked First-Layer Feature Maps


```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

Input   Filter   Second-Layer Feature Map

Third-Layer Convolution

Third-Layer Convolution

ada lovelace was a mathematician and writer who lived in the 19th century. lovelace holds the honor of having published the very first algorithm intended to be used by a machine to perform calculations, which make lovelace the first-ever computer programmer.

Input

Third-Layer Convolution

First-Layer Filters: a d l o v e c

```
ada lovelace was a mathematician
and writer who lived in the 19th
century. lovelace holds the honor
of having published the very
first algorithm intended to be
used by a machine to perform
calculations, which make lovelace
the first-ever computer
programmer.
```

Input

Third-Layer Convolution

First-Layer Filters: a d l o v e c

ada lovelace was a mathematician and writer who lived in the 19th century. lovelace holds the honor of having published the very first algorithm intended to be used by a machine to perform calculations, which make lovelace the first-ever computer programmer.

ada lovelace

Input Filter

Third-Layer Convolution

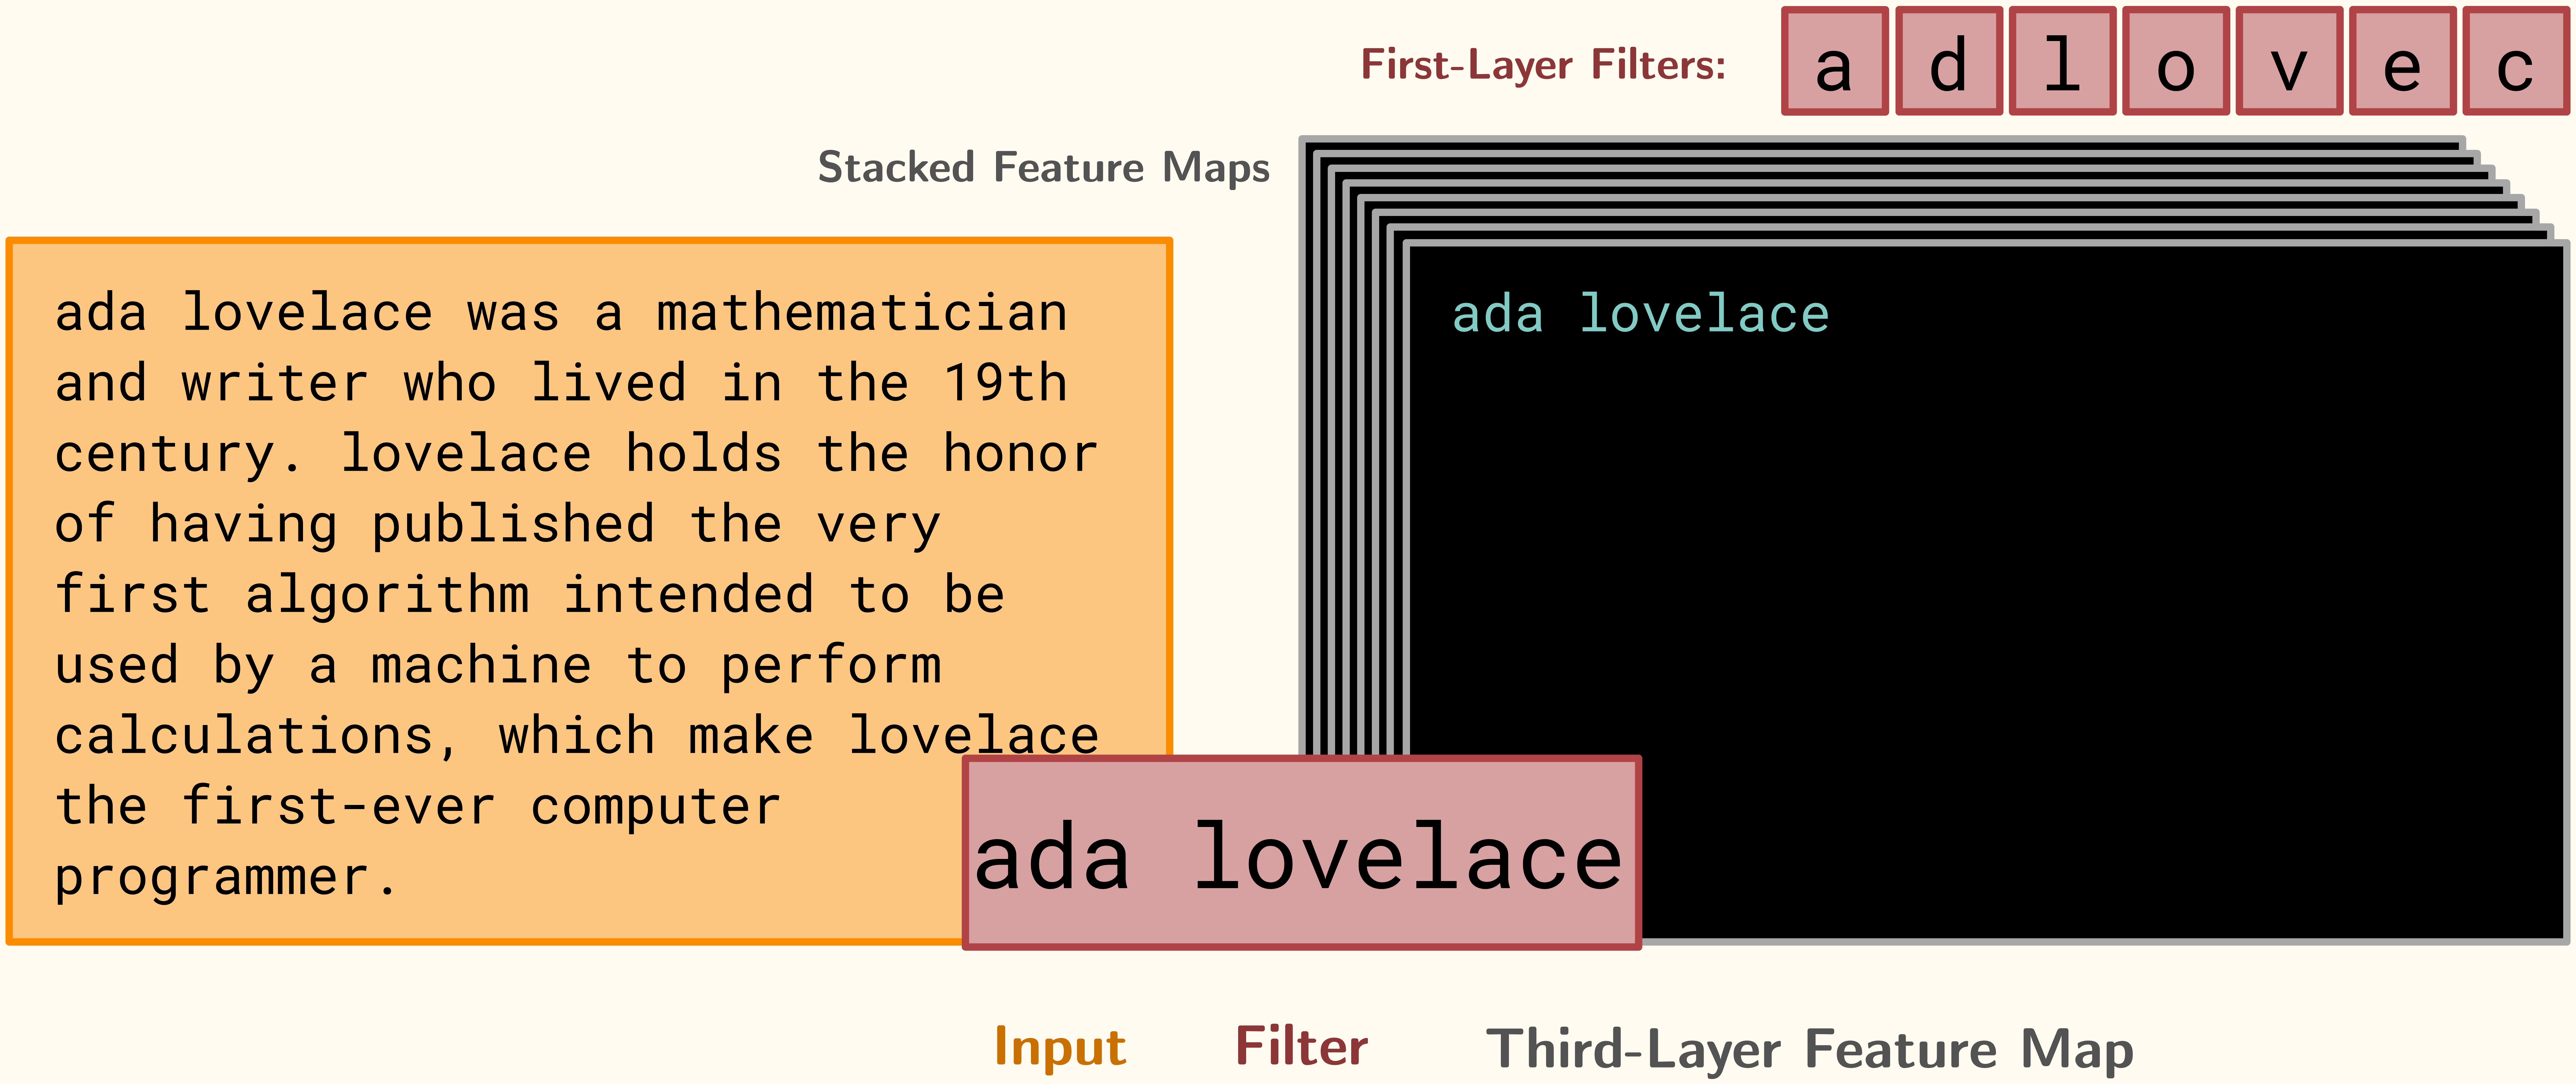

Convolutional Neural Networks (CNNs)

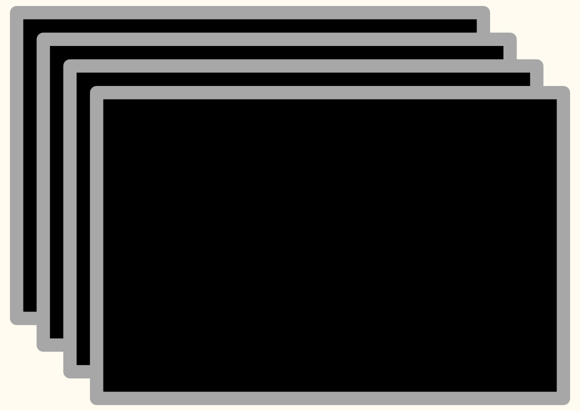

CNNs deal with greater *complexity* by having more **layers**.

Convolutional Neural Networks (CNNs)

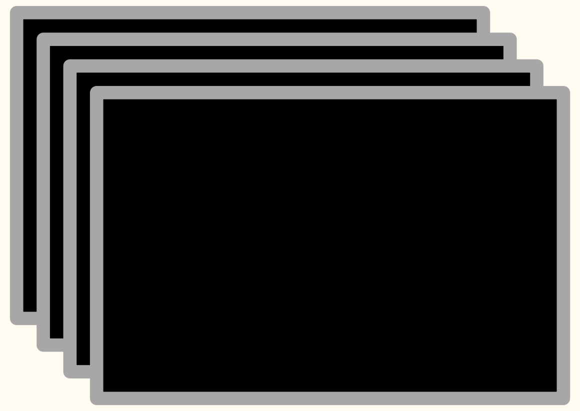

CNNs deal with greater *complexity* by having more **layers**.

Convolutional Neural Networks (CNNs)

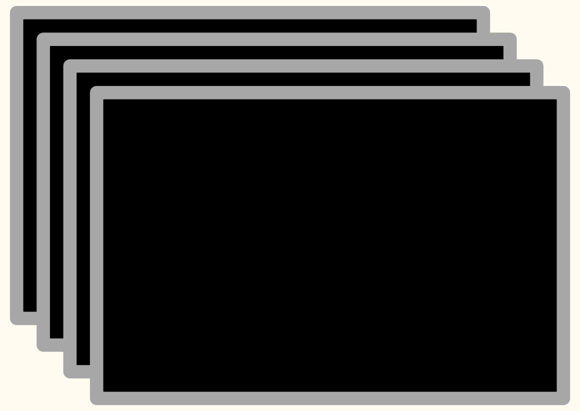

CNNs deal with greater *complexity* by having more **layers**.

Deep Learning

Convolutional Neural Networks (CNNs)

Input

Raw Image(s)

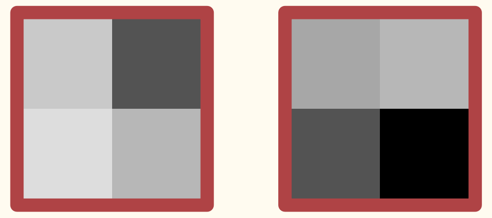

...

Filters

Feature Maps

Low-level

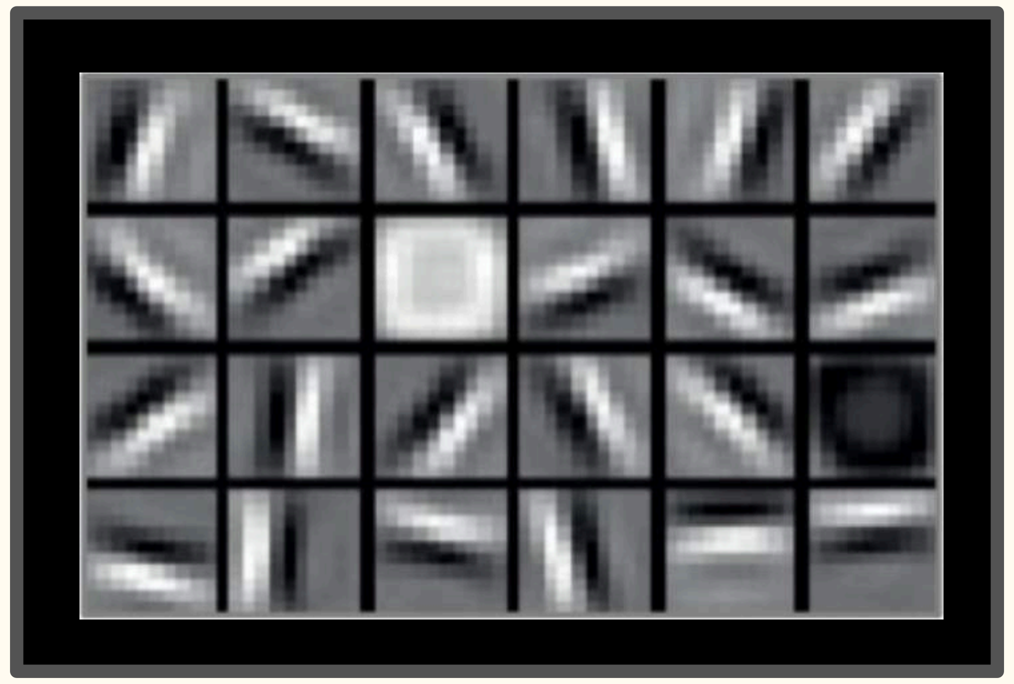

→

Mid-level

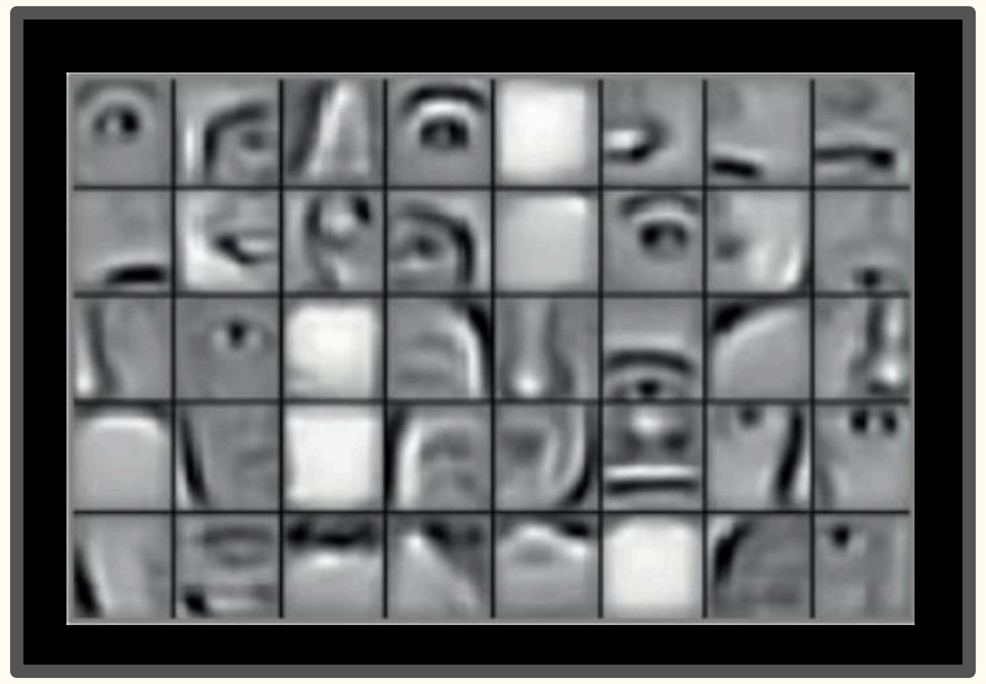

→

High-level

Convolutional Neural Networks (CNNs)

Input

Raw Image(s)

Feature Maps

Low-level

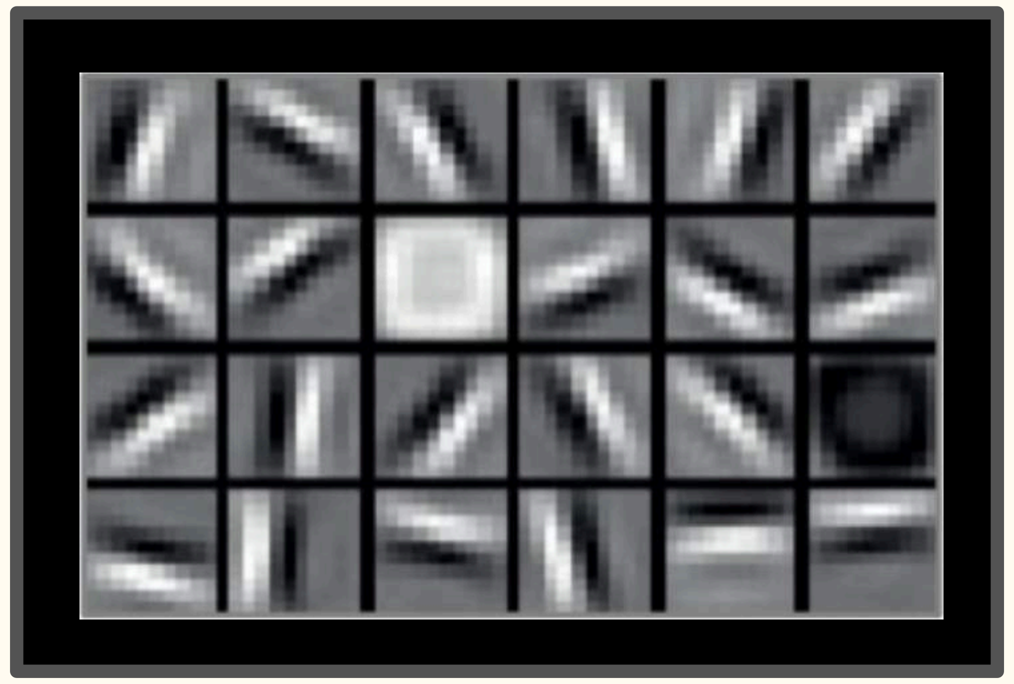

→

Mid-level

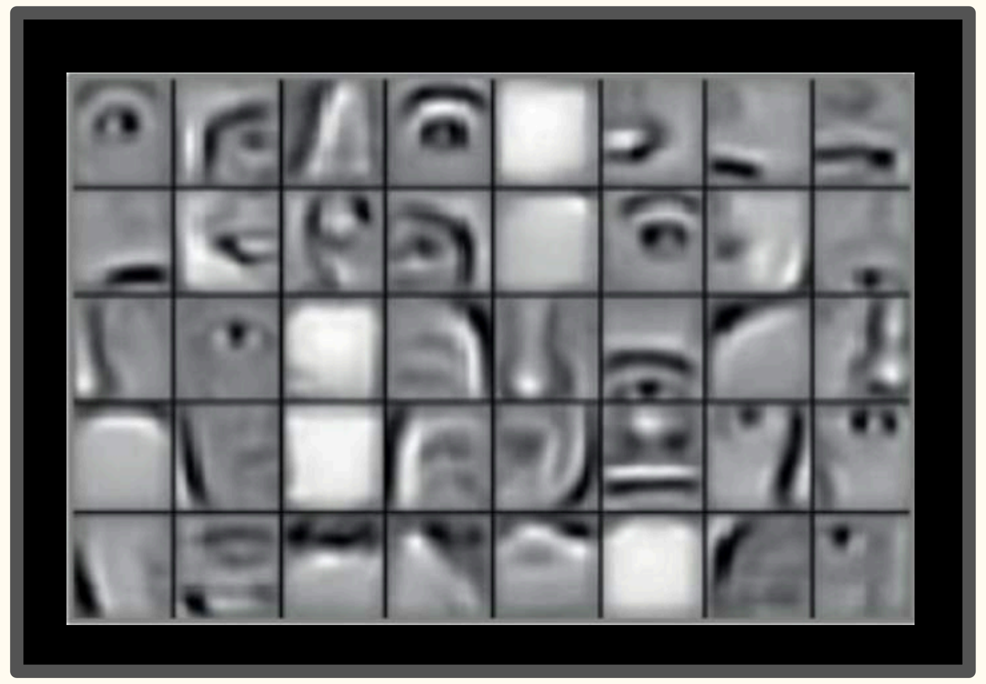

→

High-level

- Edges
- Spots

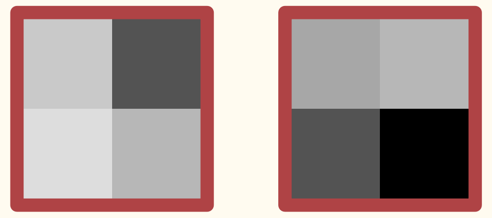

...

Filters

Convolutional Neural Networks (CNNs)

Input

Raw Image(s)

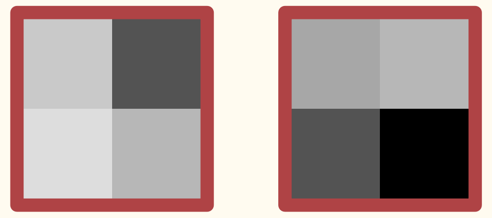

...

Filters

Feature Maps

Low-level → Mid-level → High-level

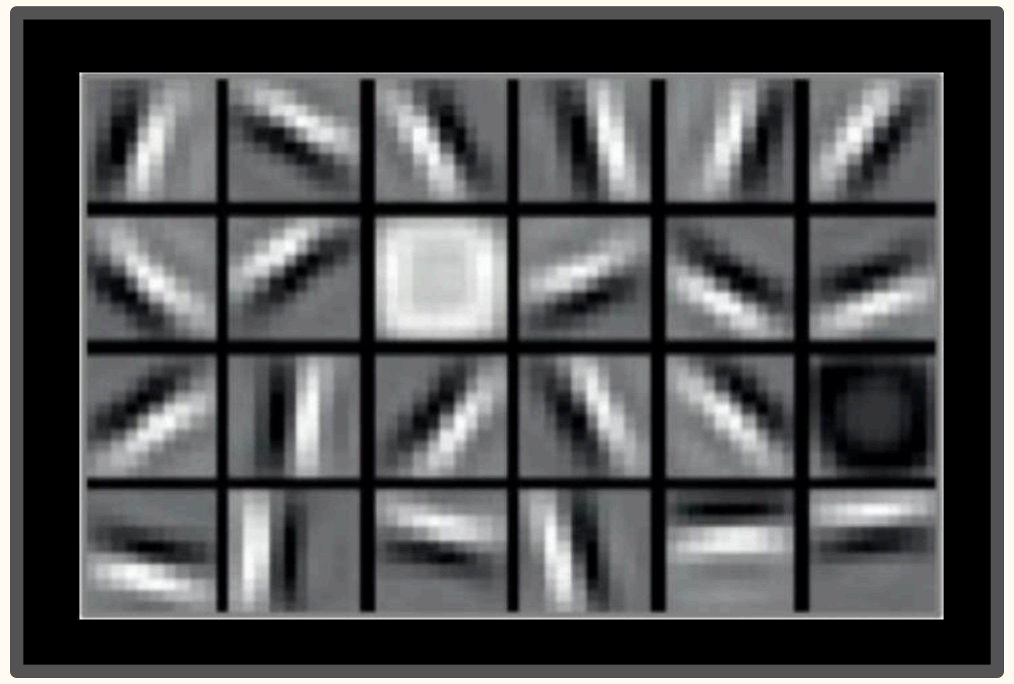

- Edges
- Spots

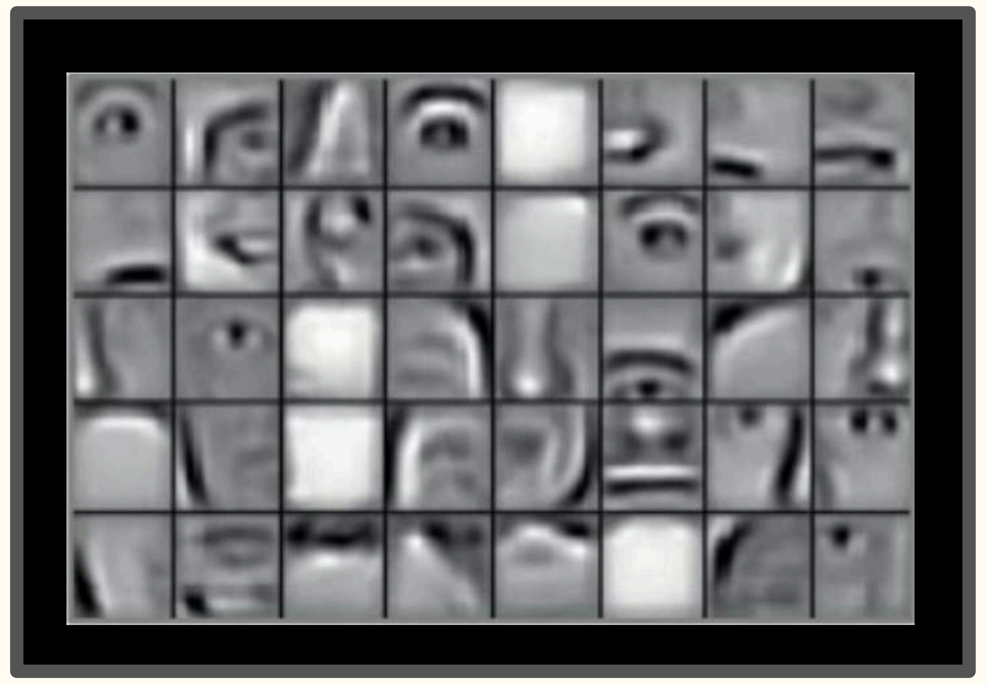

- Ears
- Eyes
- Nose

Convolutional Neural Networks (CNNs)

Input

Raw Image(s)

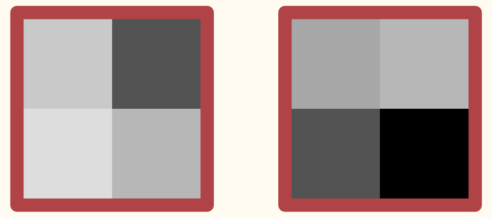

...

Filters

Feature Maps

Low-level

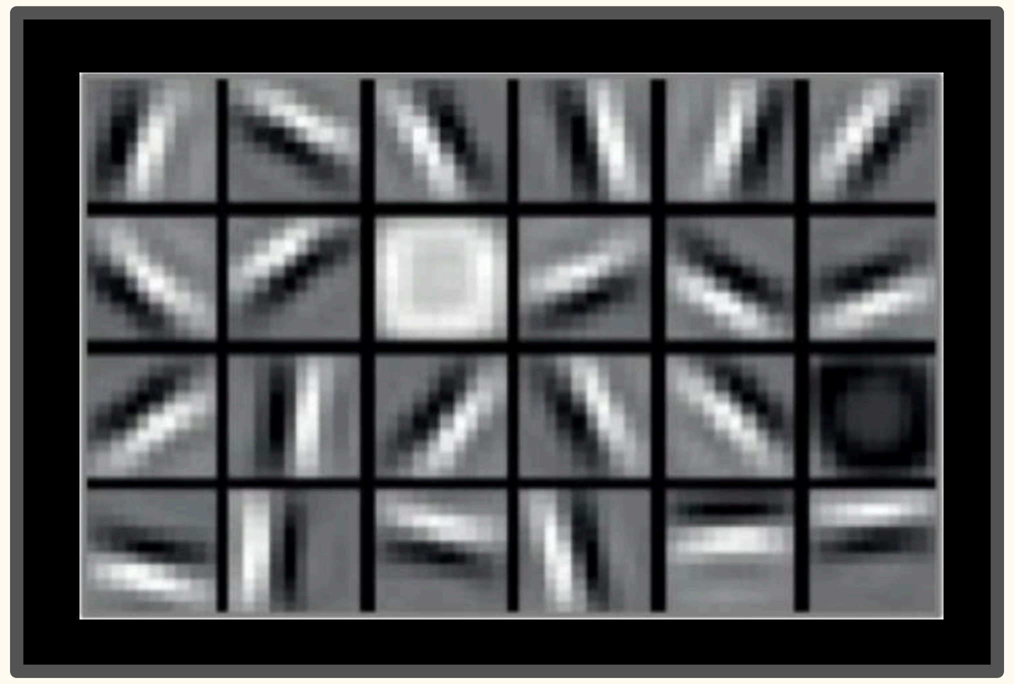

- Edges
- Spots

Mid-level

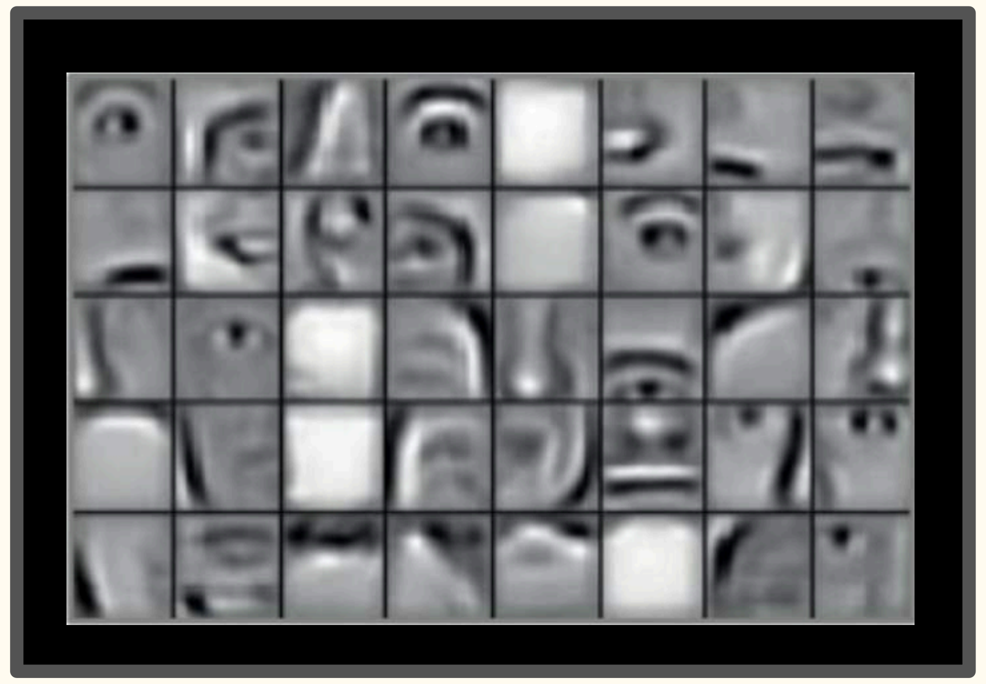

- Ears
- Eyes
- Nose

High-level

- Facial Structure

Convolutional Neural Networks (CNNs)

Input

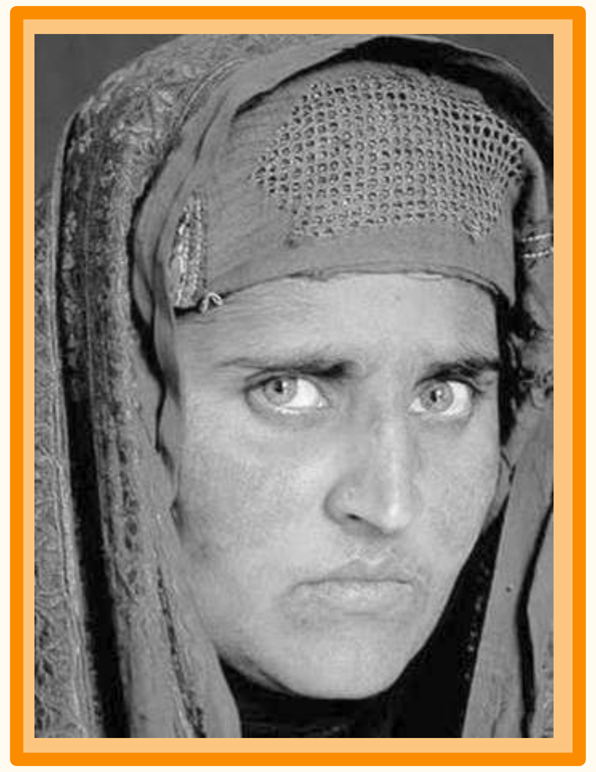

Convolutional Neural Networks (CNNs)

Input

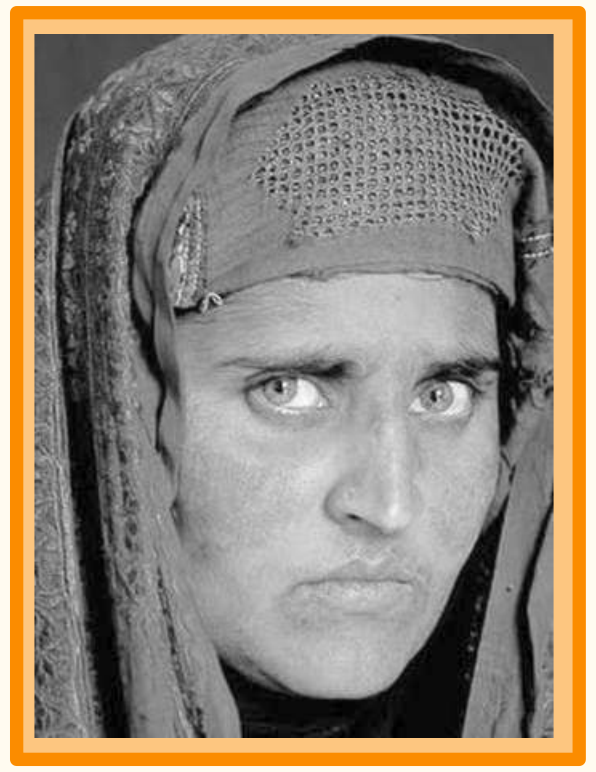

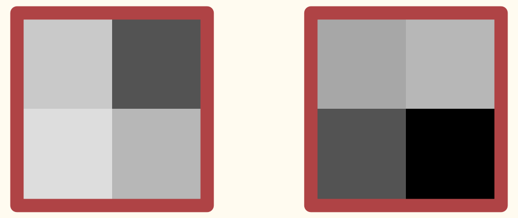

...

Filters

Convolutional Neural Networks (CNNs)

Input

Feature Maps

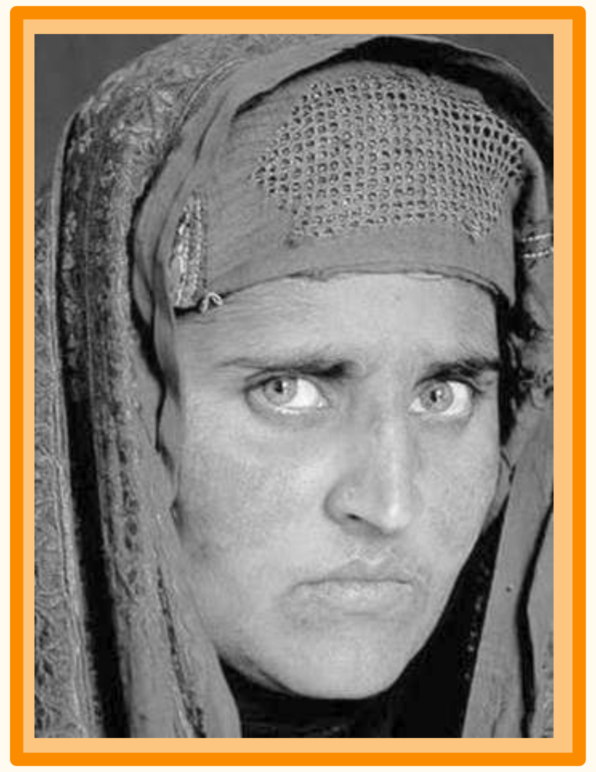

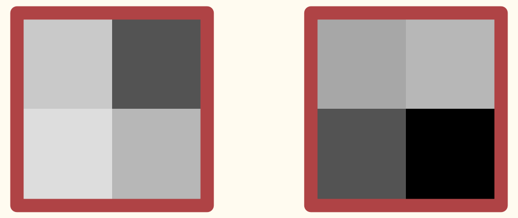

...

Filters

Convolutional Neural Networks (CNNs)

Input

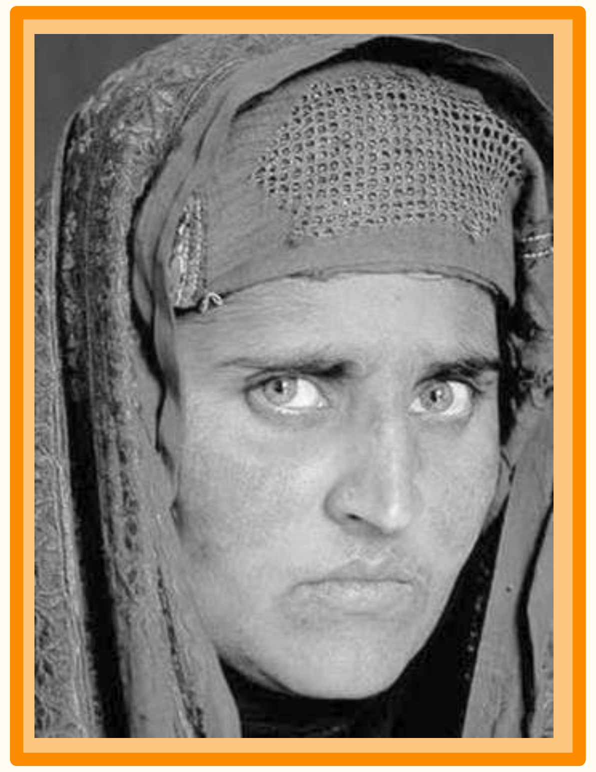

→

Feature Maps

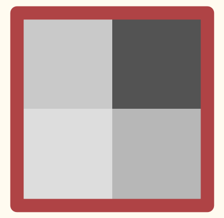

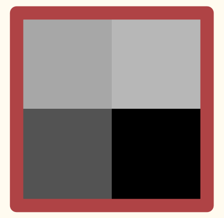

...

Filters

Convolutional Neural Networks (CNNs)

Input

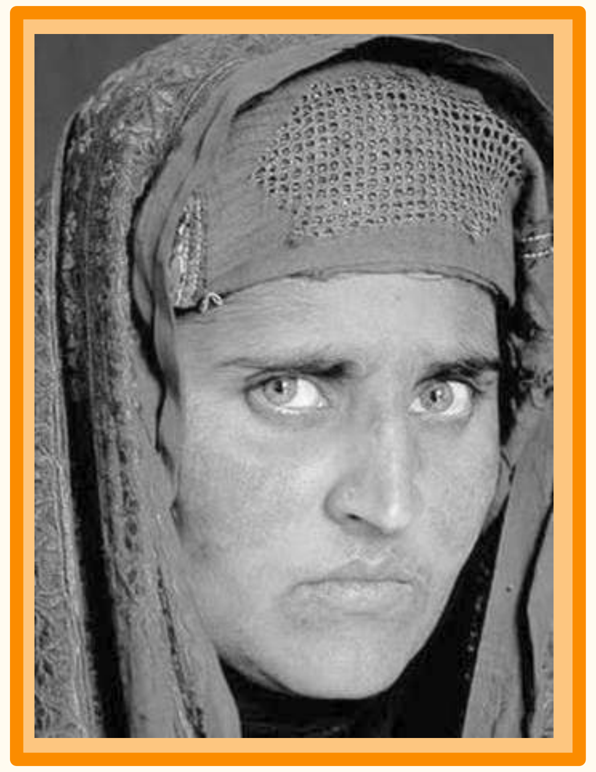

→ Low-level

Feature Maps

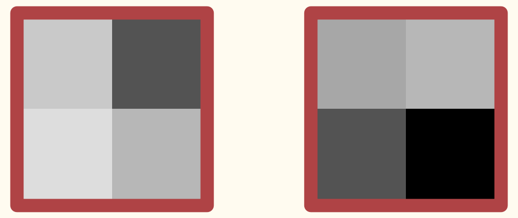

...

Filters

Convolutional Neural Networks (CNNs)

Input

Feature Maps

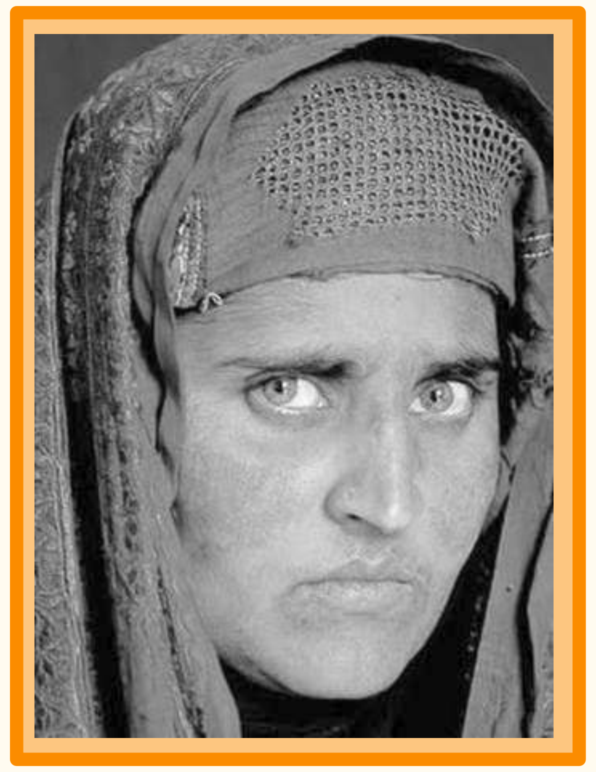

→ Low-level →

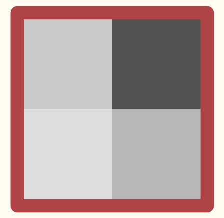

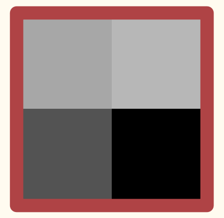

...

Filters

Convolutional Neural Networks (CNNs)

Input

Feature Maps

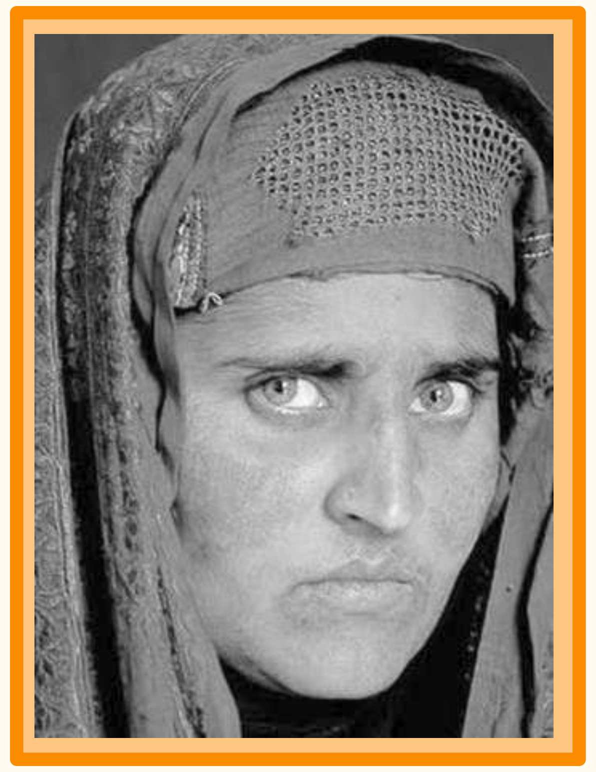

→ Low-level → Mid-level

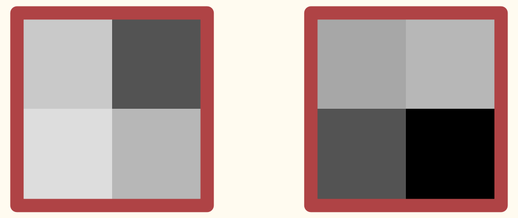

...

Filters

Convolutional Neural Networks (CNNs)

Input

Feature Maps

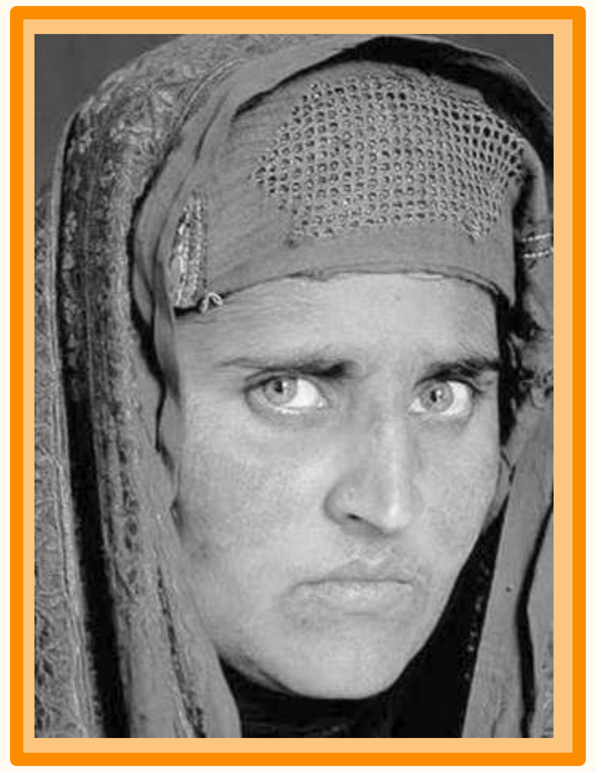

→ Low-level → Mid-level →

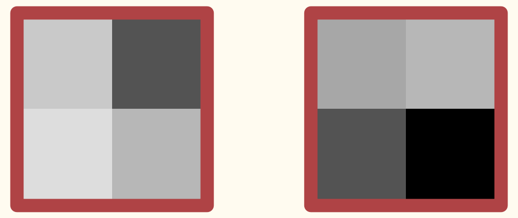

...

Filters

Convolutional Neural Networks (CNNs)

Input

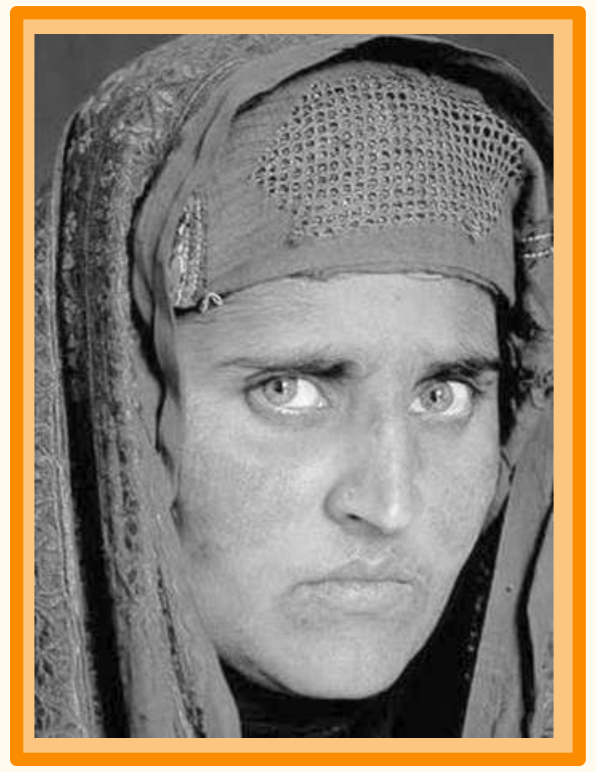

Feature Maps

→ Low-level → Mid-level → High-level

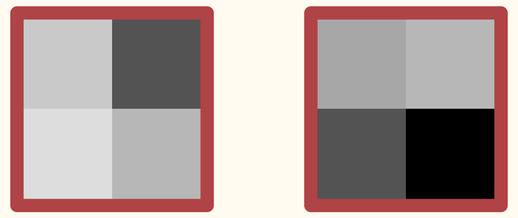

...

Filters

Convolutional Neural Networks (CNNs)

Input

Feature Maps

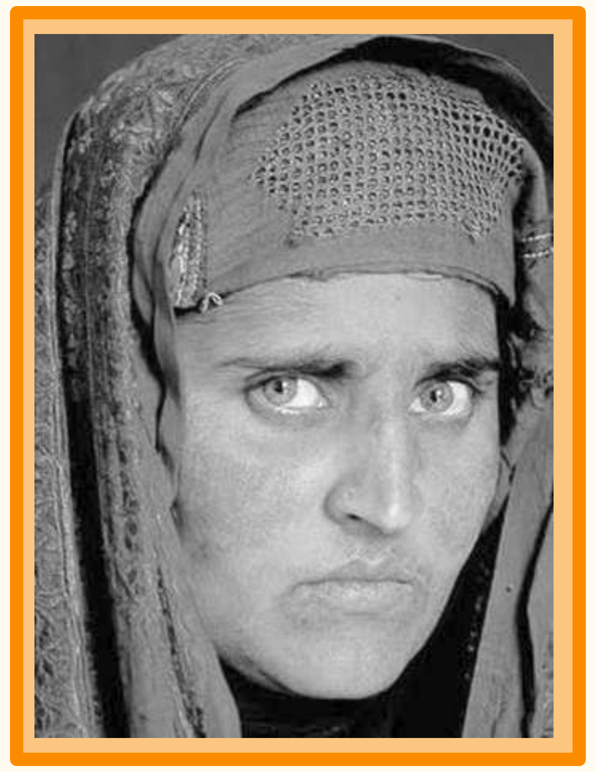

→ Low-level → Mid-level → High-level →

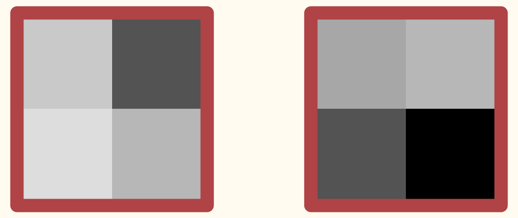

...

Filters

Convolutional Neural Networks (CNNs)

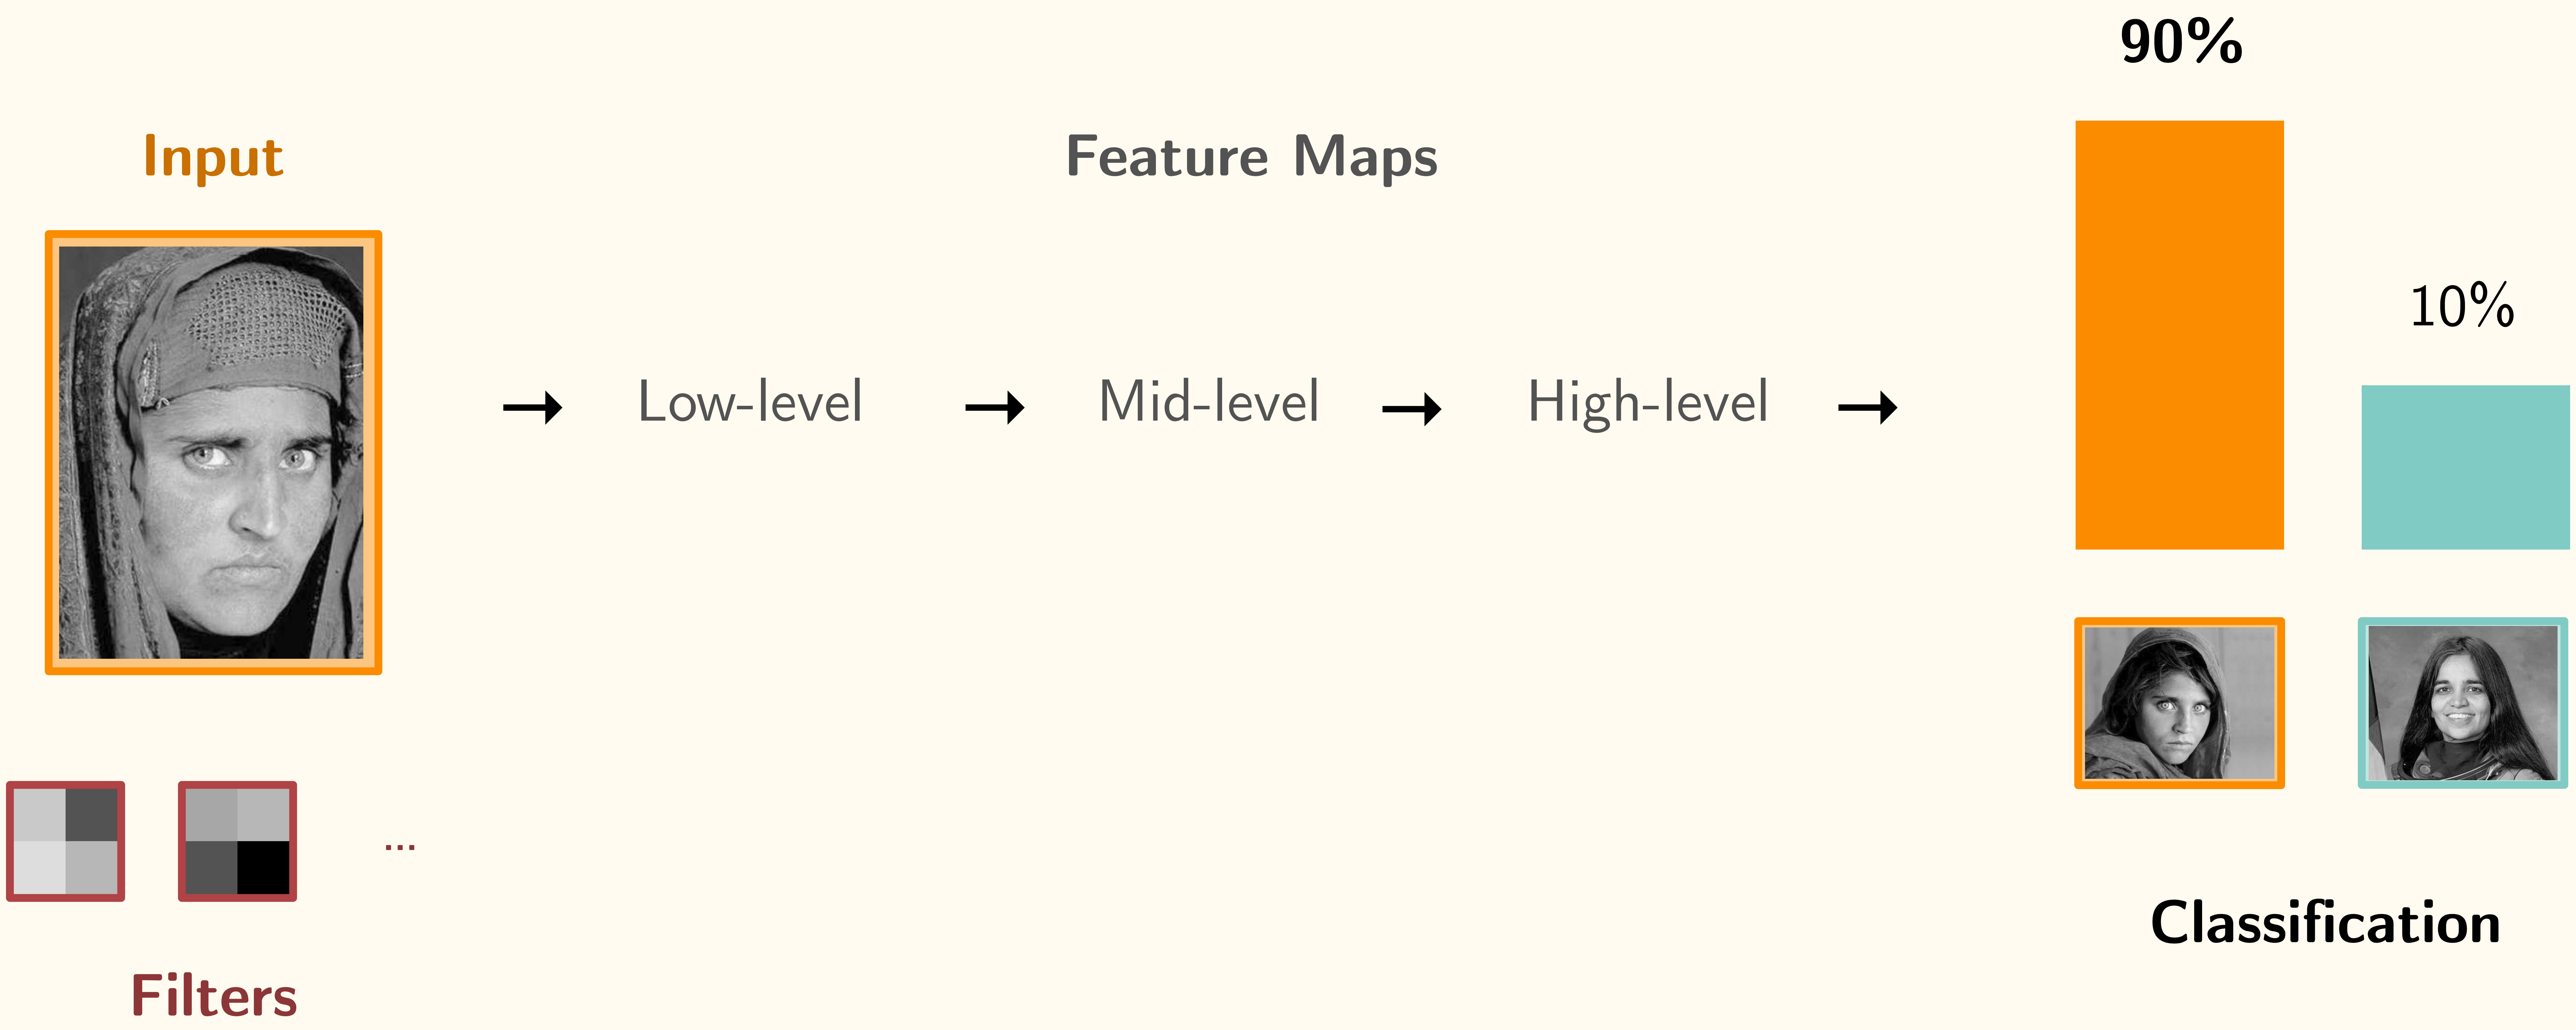

Convolutional Neural Networks (CNNs)

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier*. Medium, Towards Data Science.

Convolutional Neural Networks (CNNs)

- Works well on data with spatial relationship

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier.* Medium, Towards Data Science.

Convolutional Neural Networks (CNNs)

- Works well on data with spatial relationship
- Translation invariant

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier.* Medium, Towards Data Science.

Convolutional Neural Networks (CNNs)

- Works well on data with spatial relationship

- Translation invariant
- Scale invariant

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier.* Medium, Towards Data Science.

Convolutional Neural Networks (CNNs)

- Works well on data with spatial relationship

- Translation invariant
- Scale invariant

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier.* Medium, Towards Data Science.

Convolutional Neural Networks (CNNs)

- Works well on data with spatial relationship

- Translation invariant
- Scale invariant

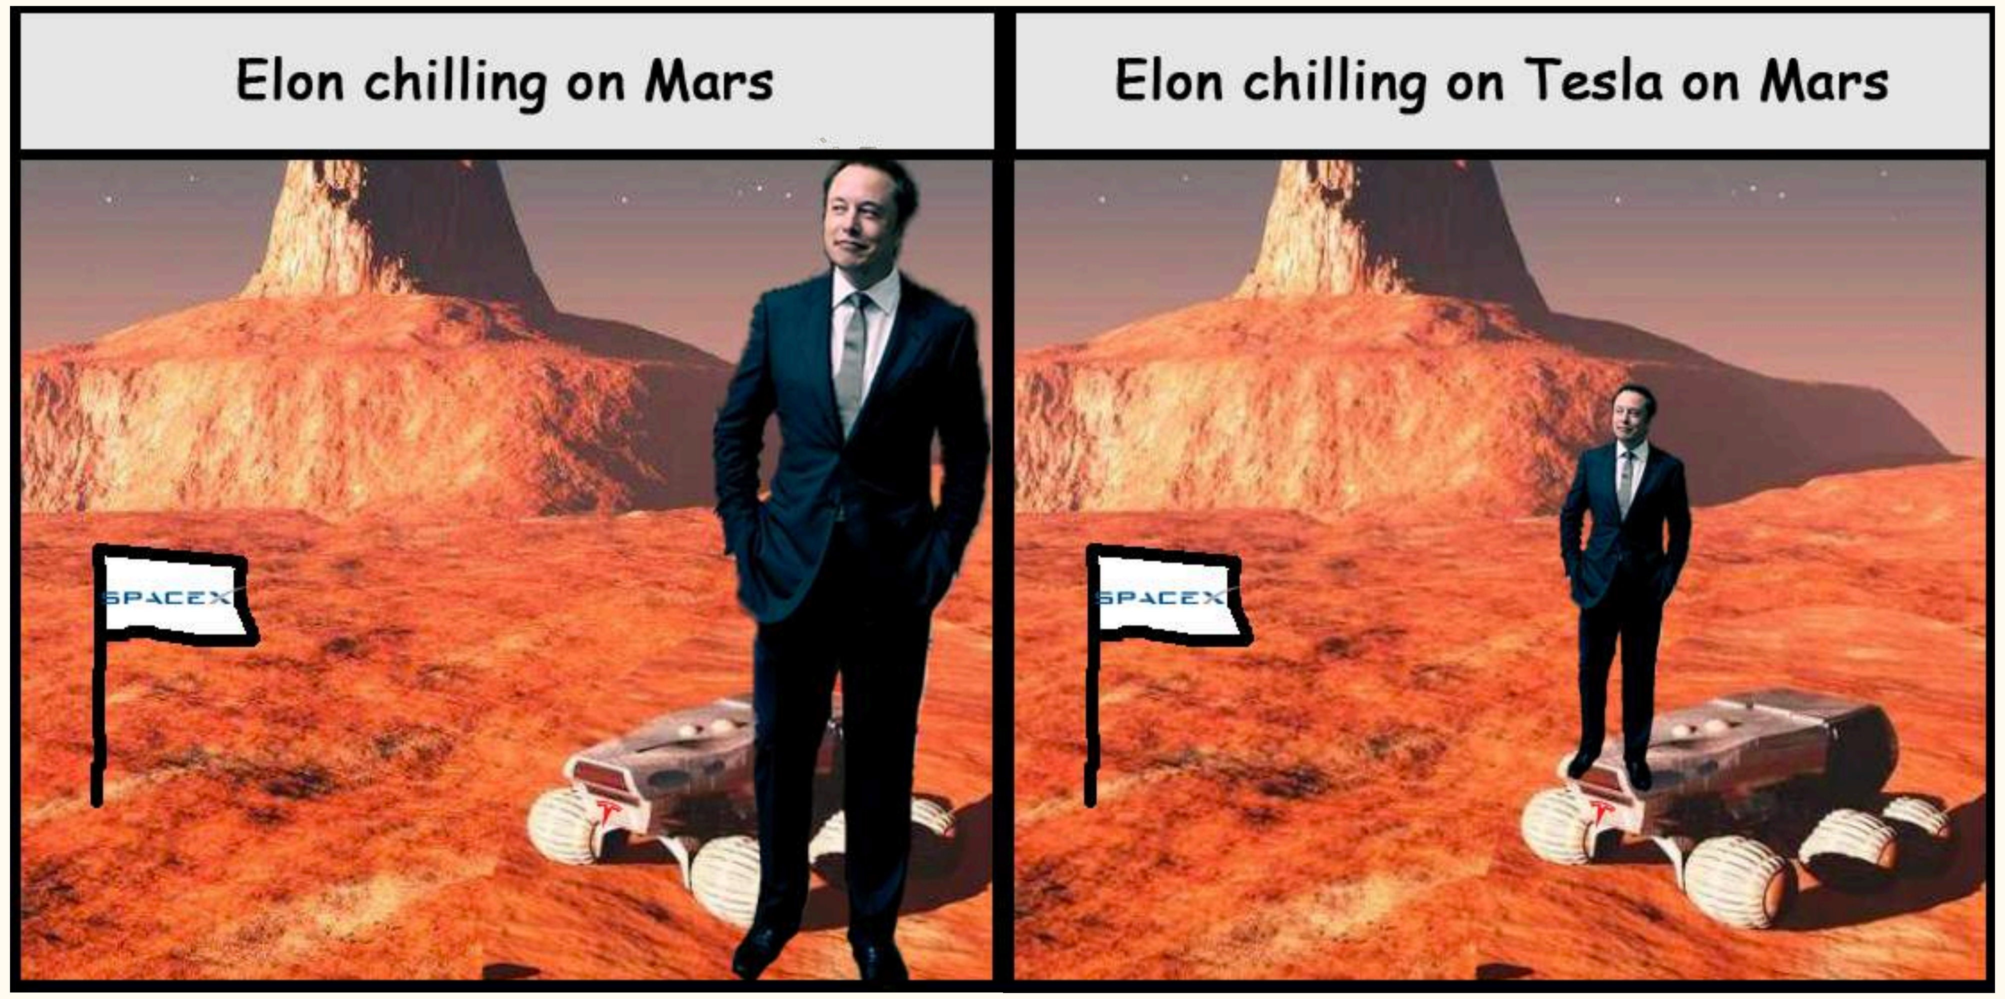

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier.* Medium, Towards Data Science.

Convolutional Neural Networks (CNNs)

- Works well on data with spatial relationship

- Translation invariant
- Scale invariant

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier.* Medium, Towards Data Science.

Convolutional Neural Networks (CNNs)

- Works well on data with spatial relationship

- Translation invariant
- Scale invariant

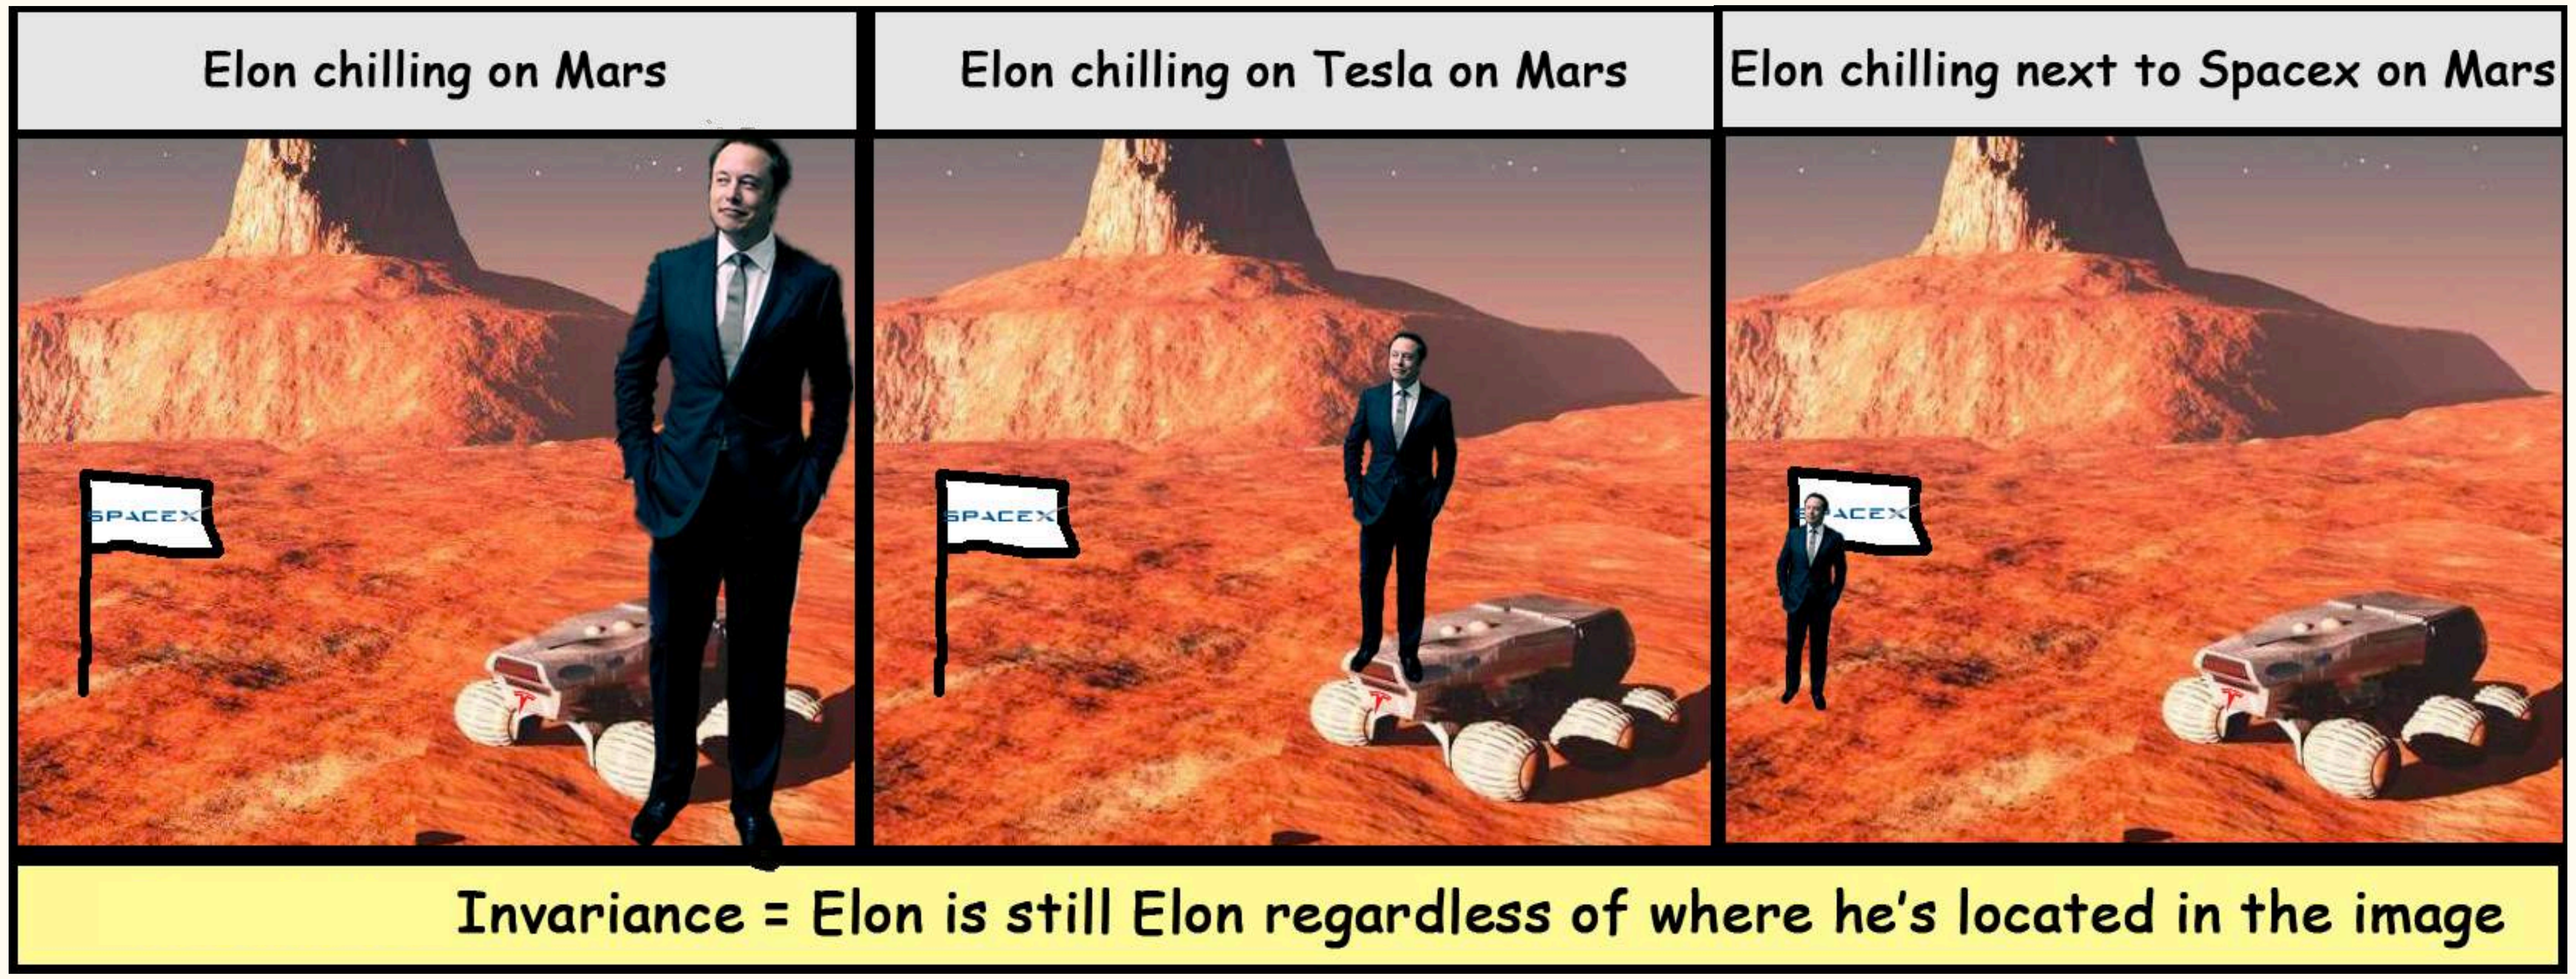

Dave Smith. 2018. *Cutting-Edge Face Recognition is Complicated. These Spreadsheets Make it Easier.* Medium, Towards Data Science.

CNNs in Technology

Wu et al., 2017, February. *Automatic alt-text: Computer-generated image descriptions for blind users on a social network service.* 2017 ACM Conference.

CNNs in Technology

Wu et al., 2017, February. *Automatic alt-text: Computer-generated image descriptions for blind users on a social network service.* 2017 ACM Conference.

CNNs in Technology

- Automatic image recognition and captioning

Wu et al., 2017, February. *Automatic alt-text: Computer-generated image descriptions for blind users on a social network service.* 2017 ACM Conference.

CNNs in Technology

- Automatic image recognition and captioning
- Used for visually impaired people

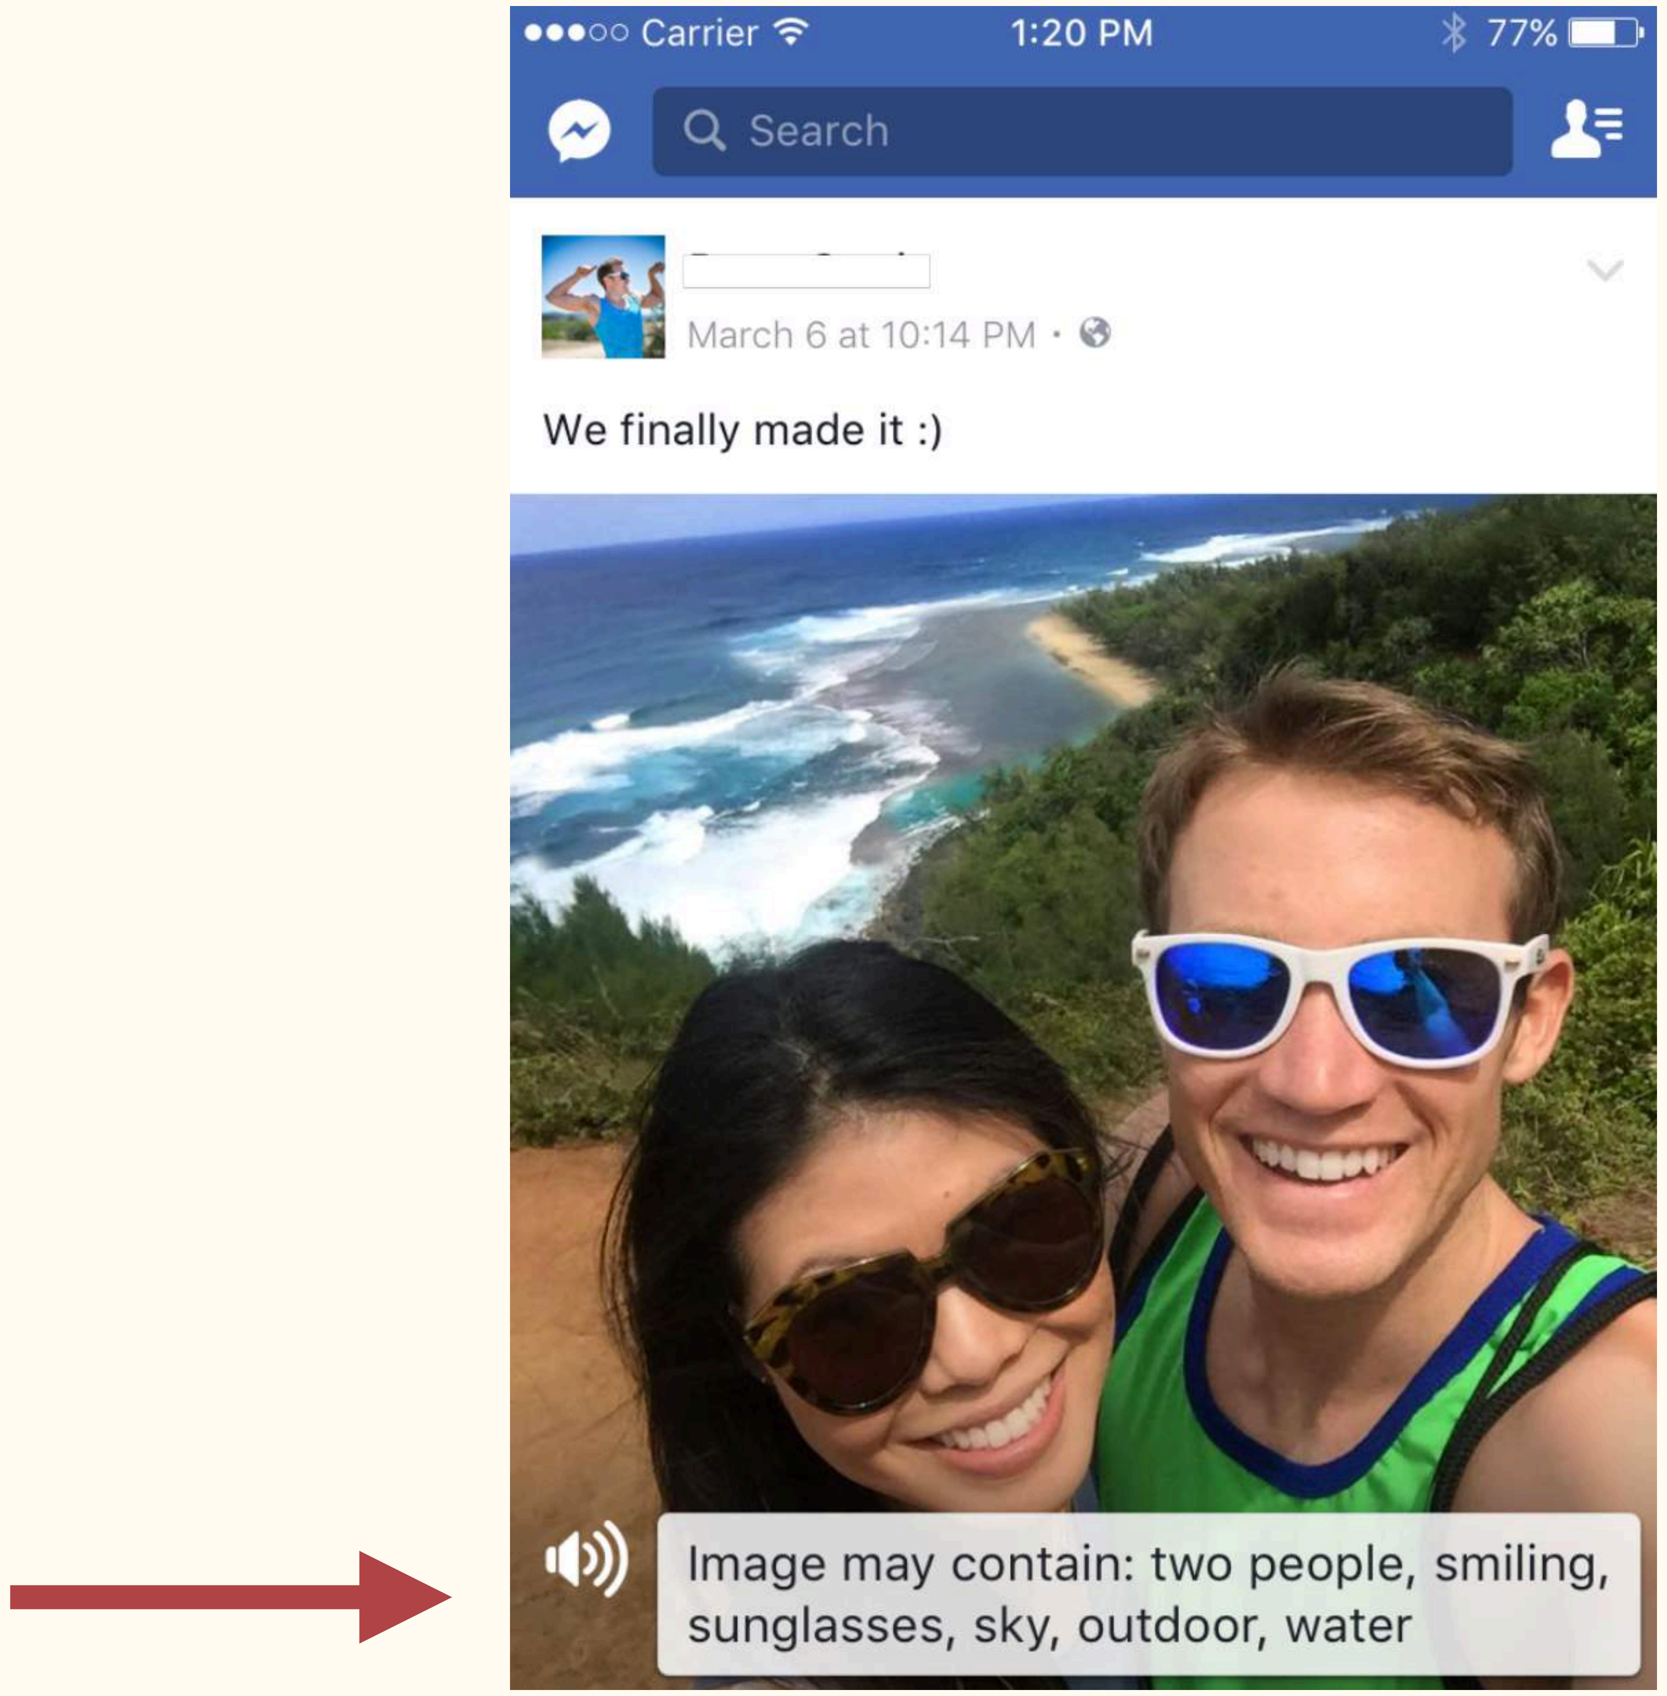

Wu et al., 2017, February. *Automatic alt-text: Computer-generated image descriptions for blind users on a social network service.* 2017 ACM Conference.

CNNs in Life Sciences

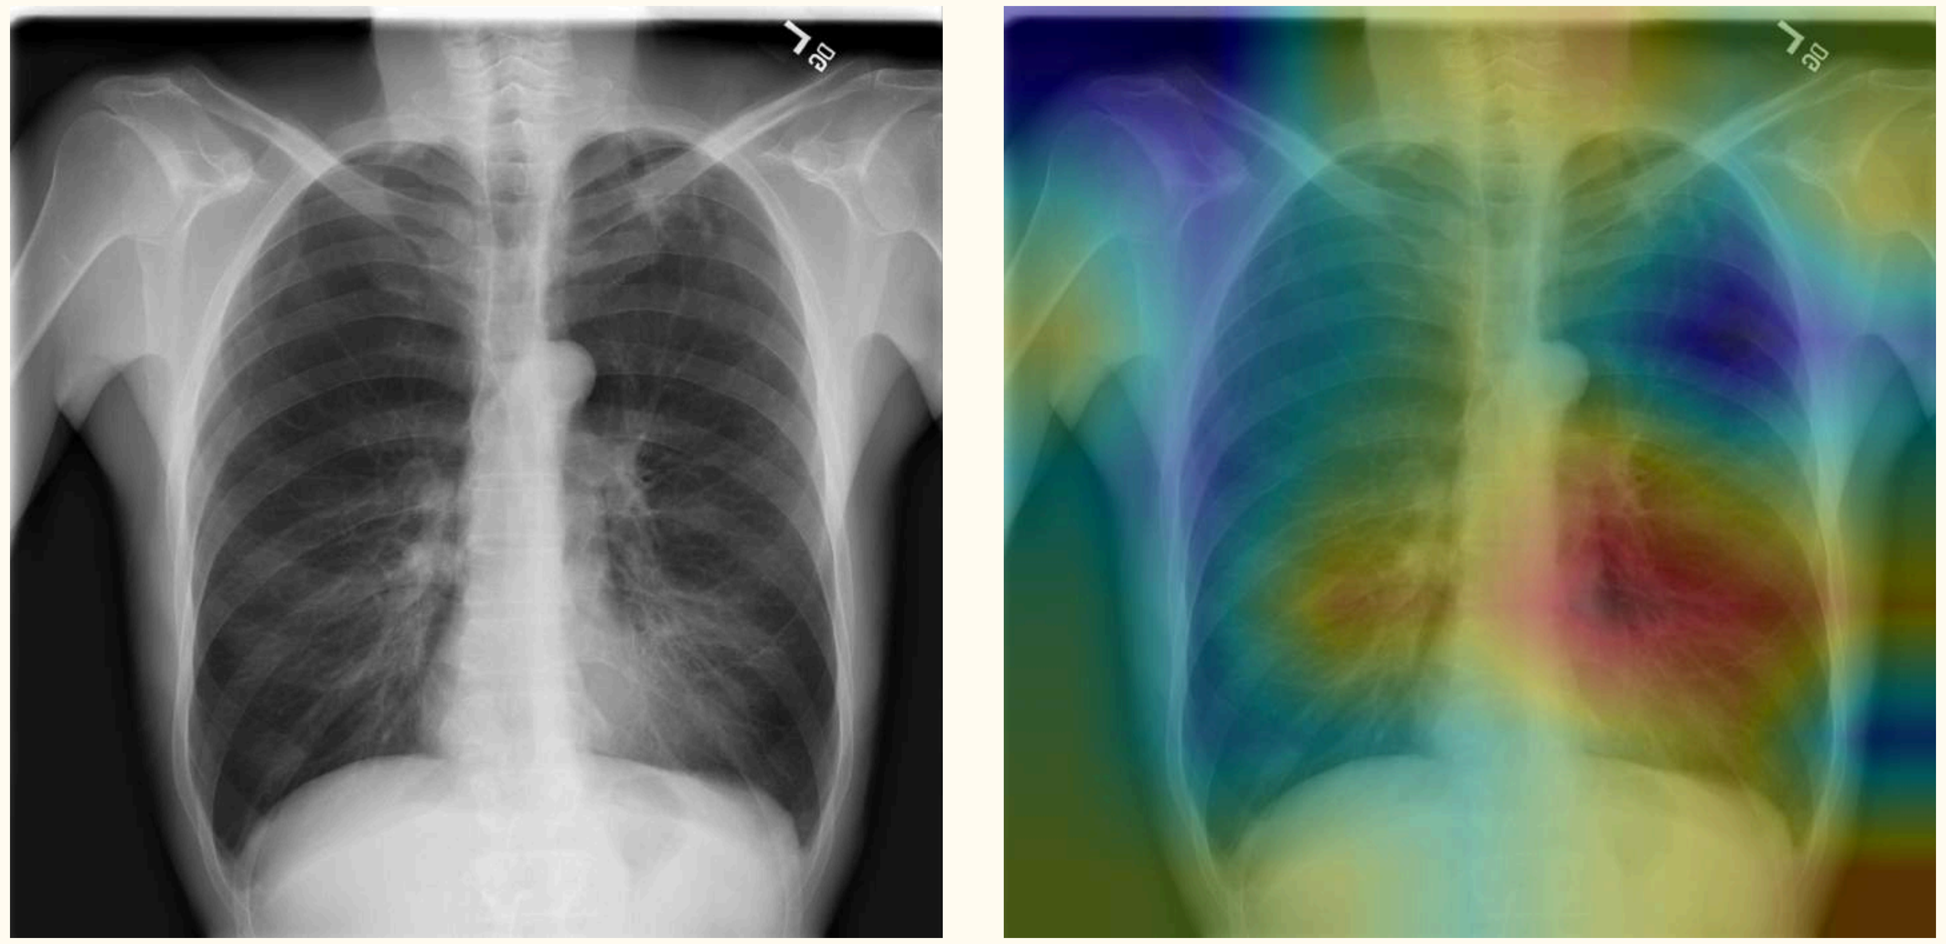

Rajpurkar et al., 2017. *Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning.*

CNNs in Life Sciences

Medical Diagnosis

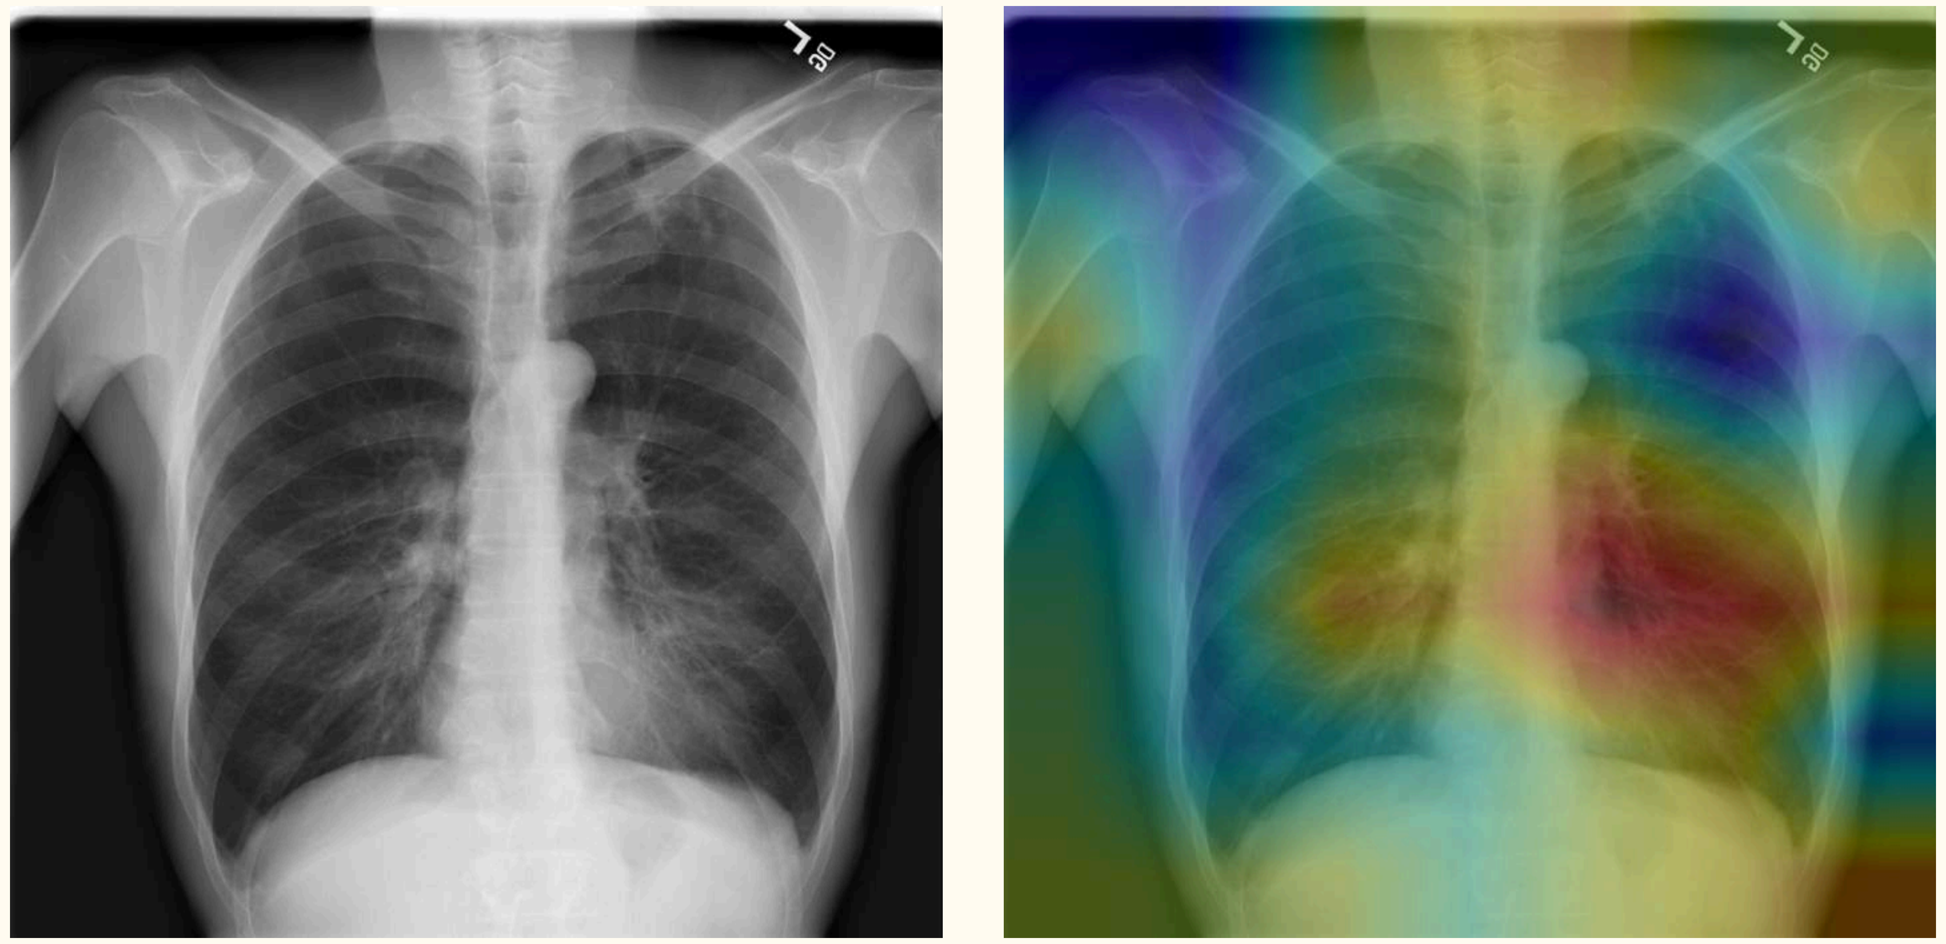

Rajpurkar et al., 2017. *Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning.*

CNNs in Life Sciences

Medical Diagnosis

- CheXNet:

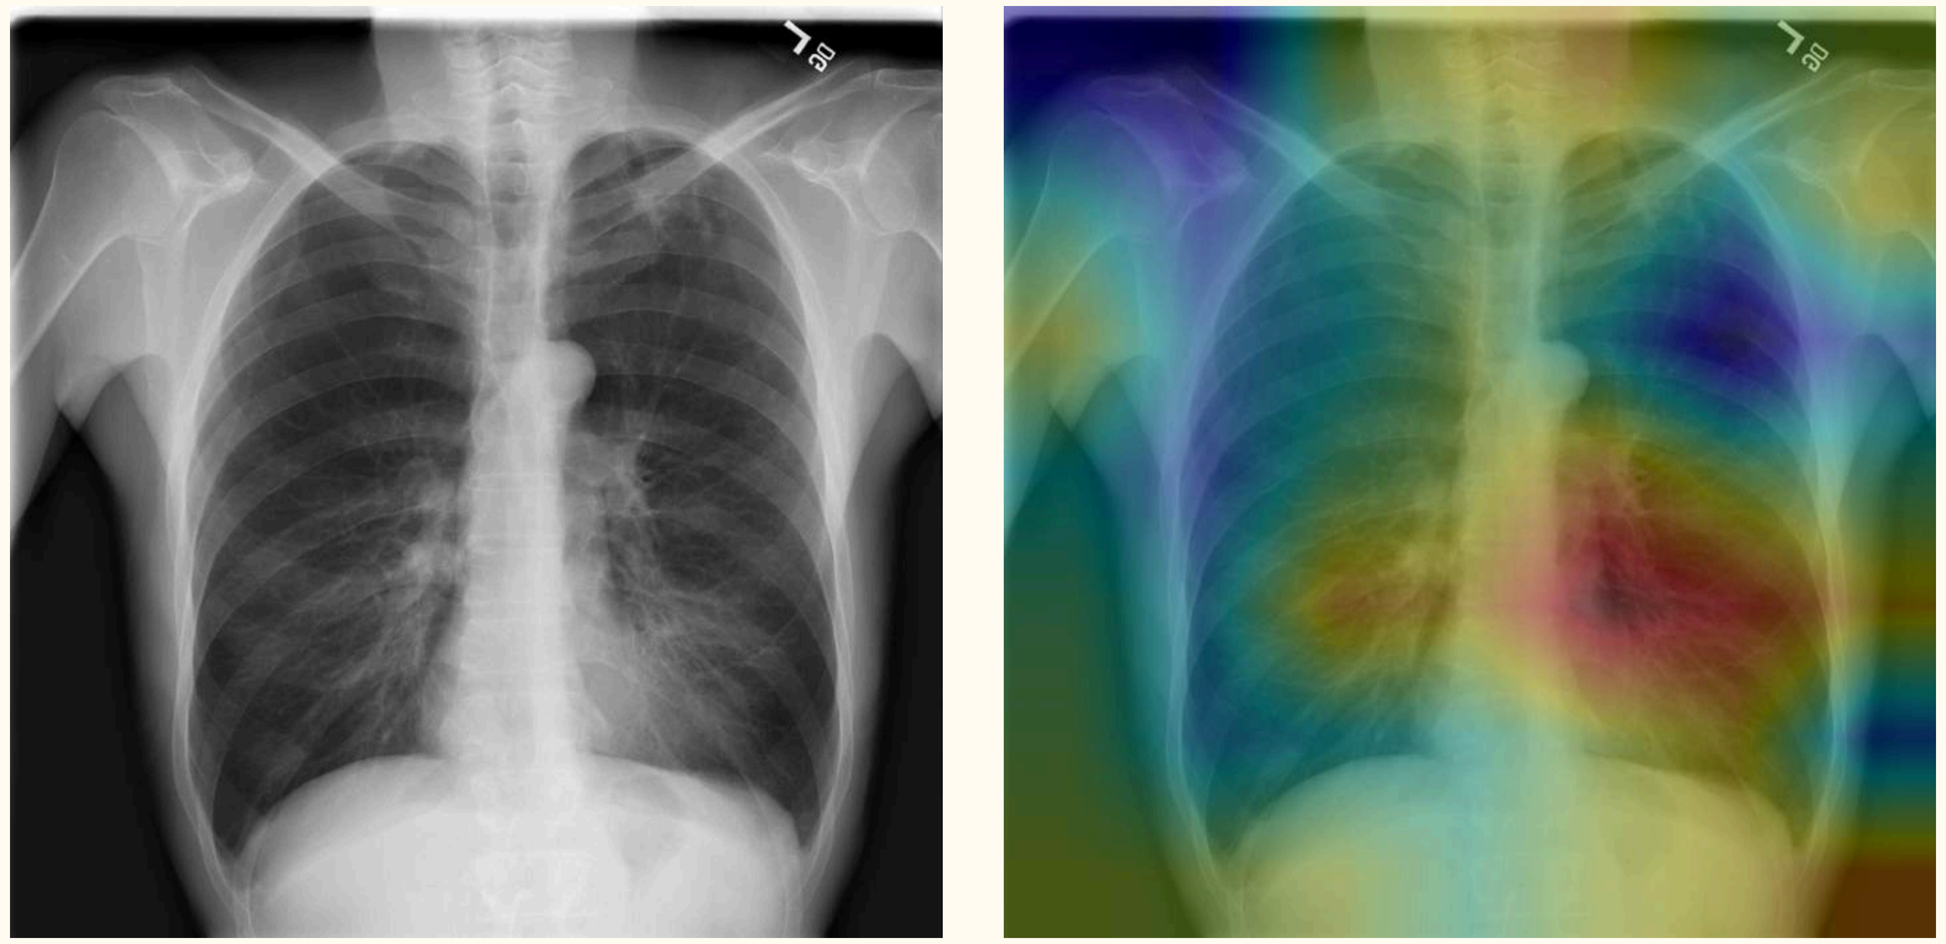

Rajpurkar et al., 2017. *Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning.*

CNNs in Life Sciences

Medical Diagnosis

- CheXNet:
  - 21-layer CNN

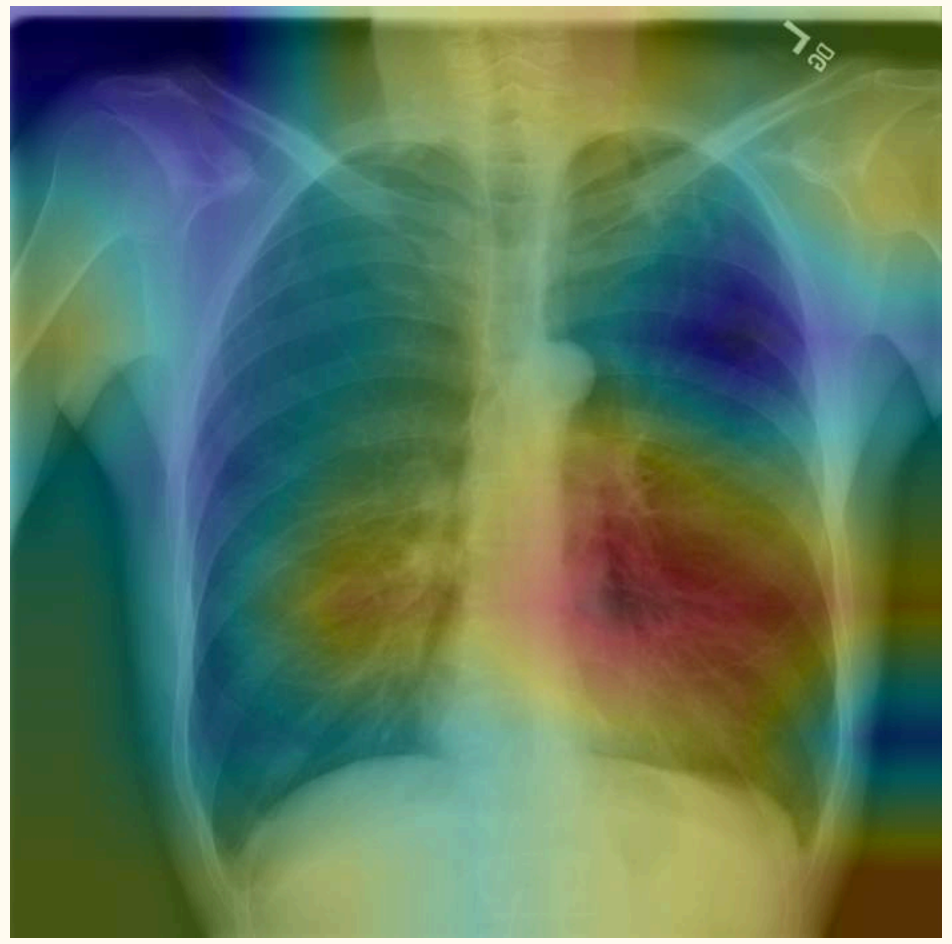

Rajpurkar et al., 2017. *Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning.*

CNNs in Life Sciences

Medical Diagnosis

- CheXNet:
  - 21-layer CNN
  - Input: chest X-ray image

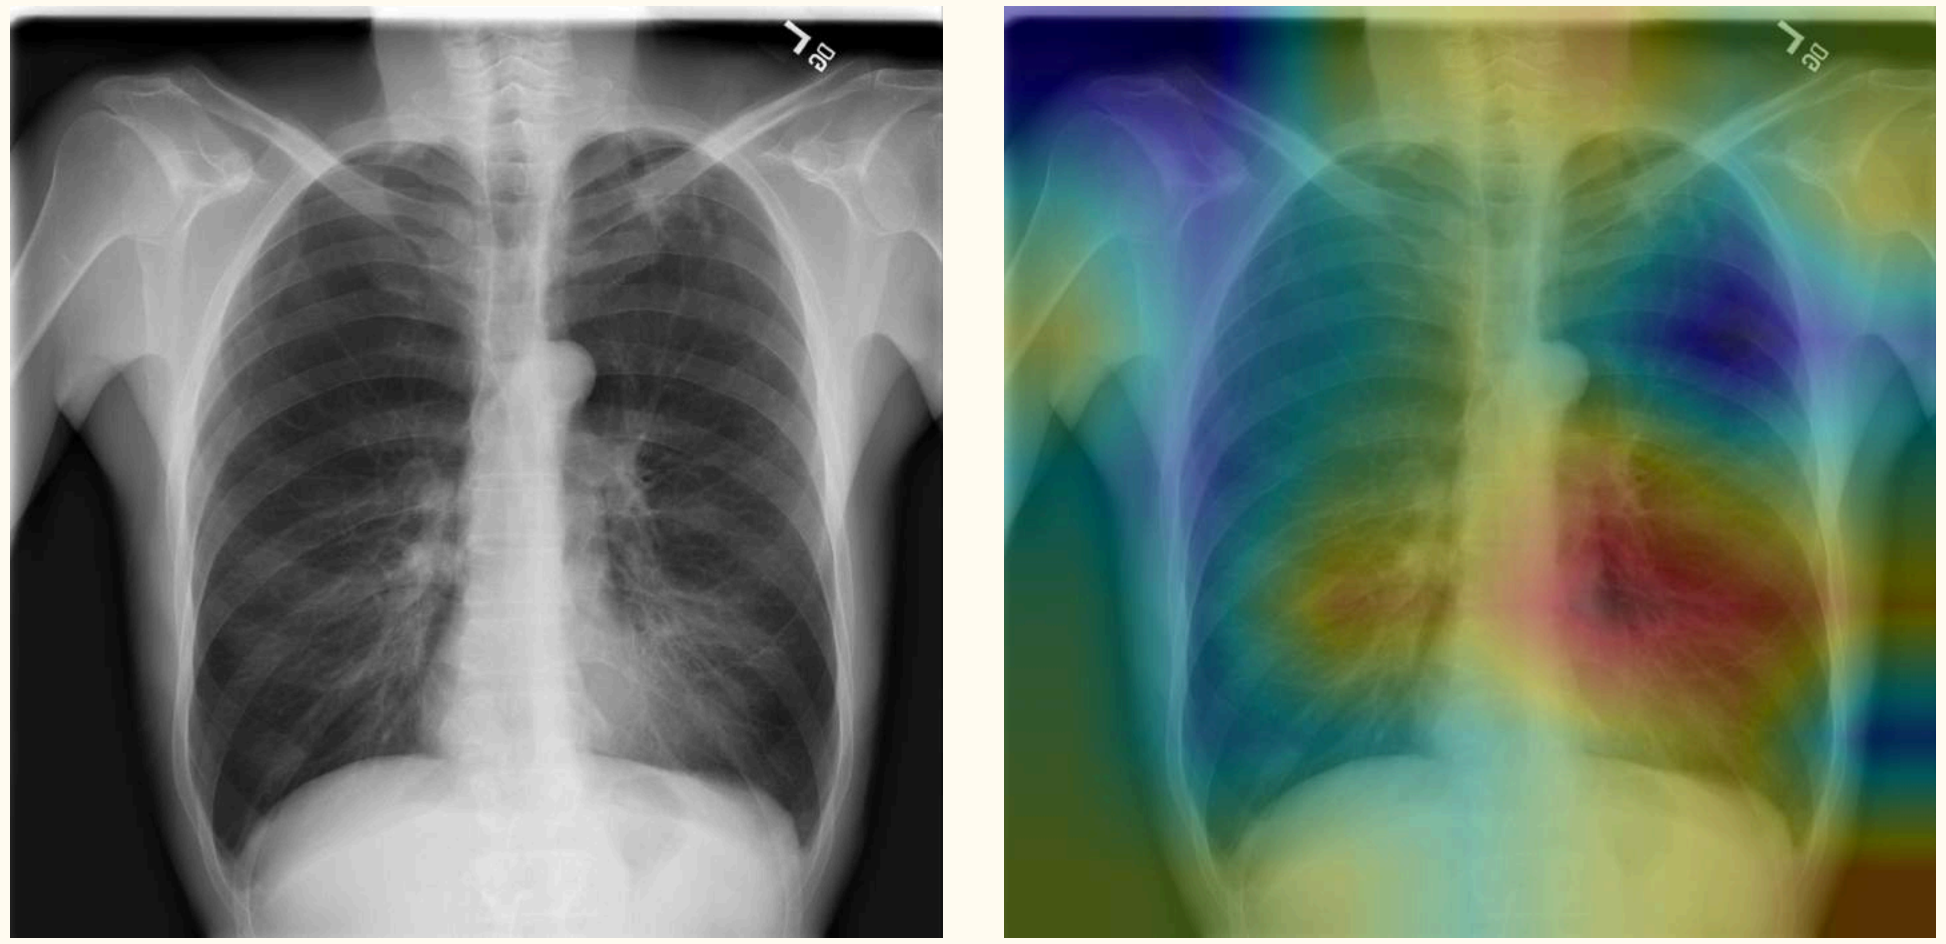

Rajpurkar et al., 2017. *Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning.*

CNNs in Life Sciences

Medical Diagnosis

- CheXNet:
  - 21-layer CNN
  - Input: chest X-ray image
  - Outputs: probability of a pathology

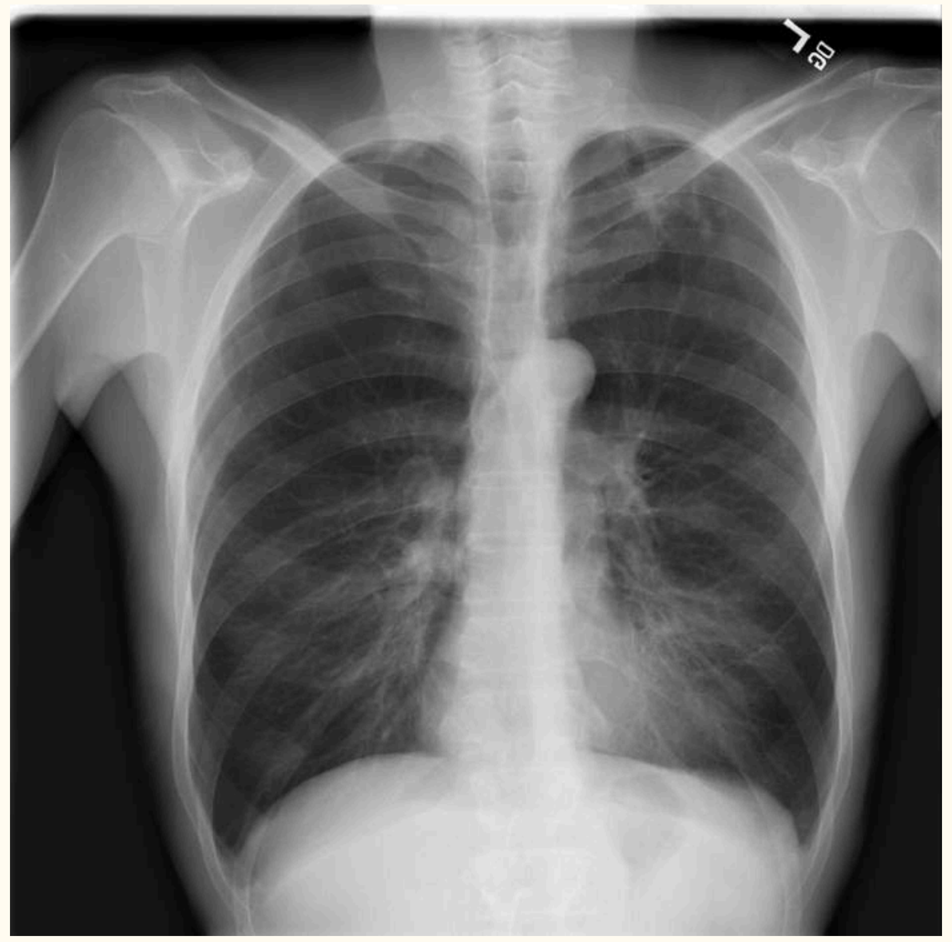
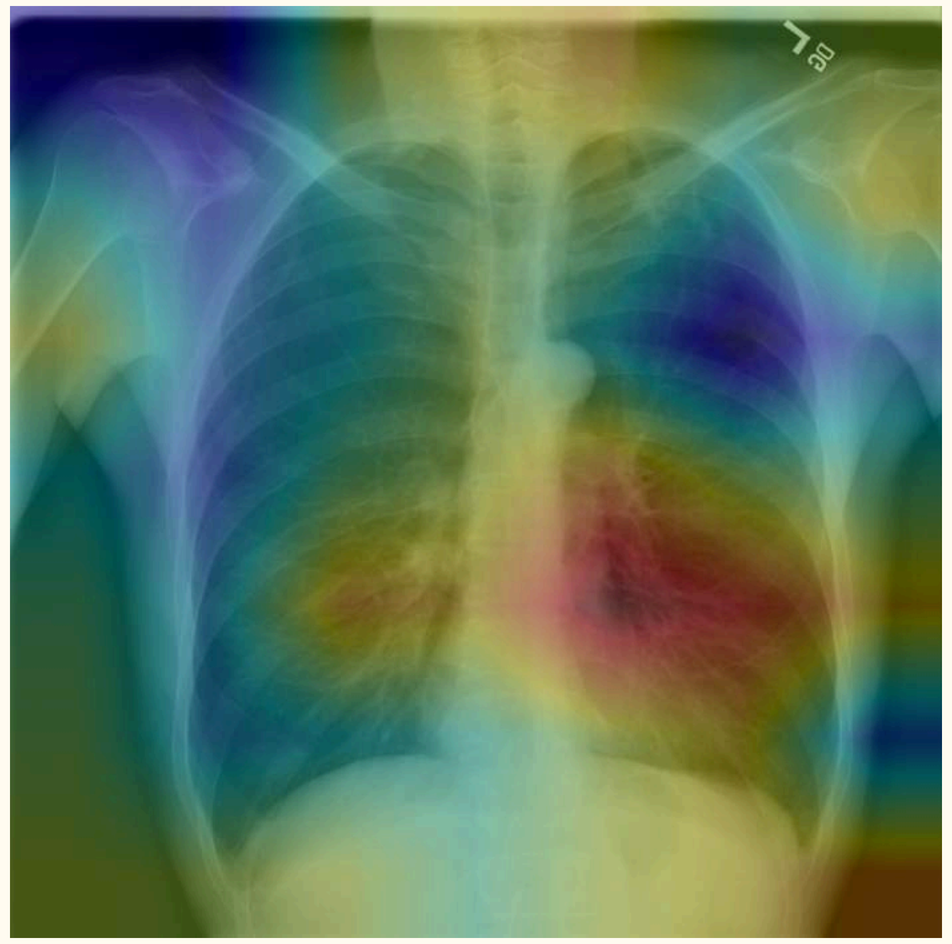

Rajpurkar et al., 2017. *Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning.*

Summary

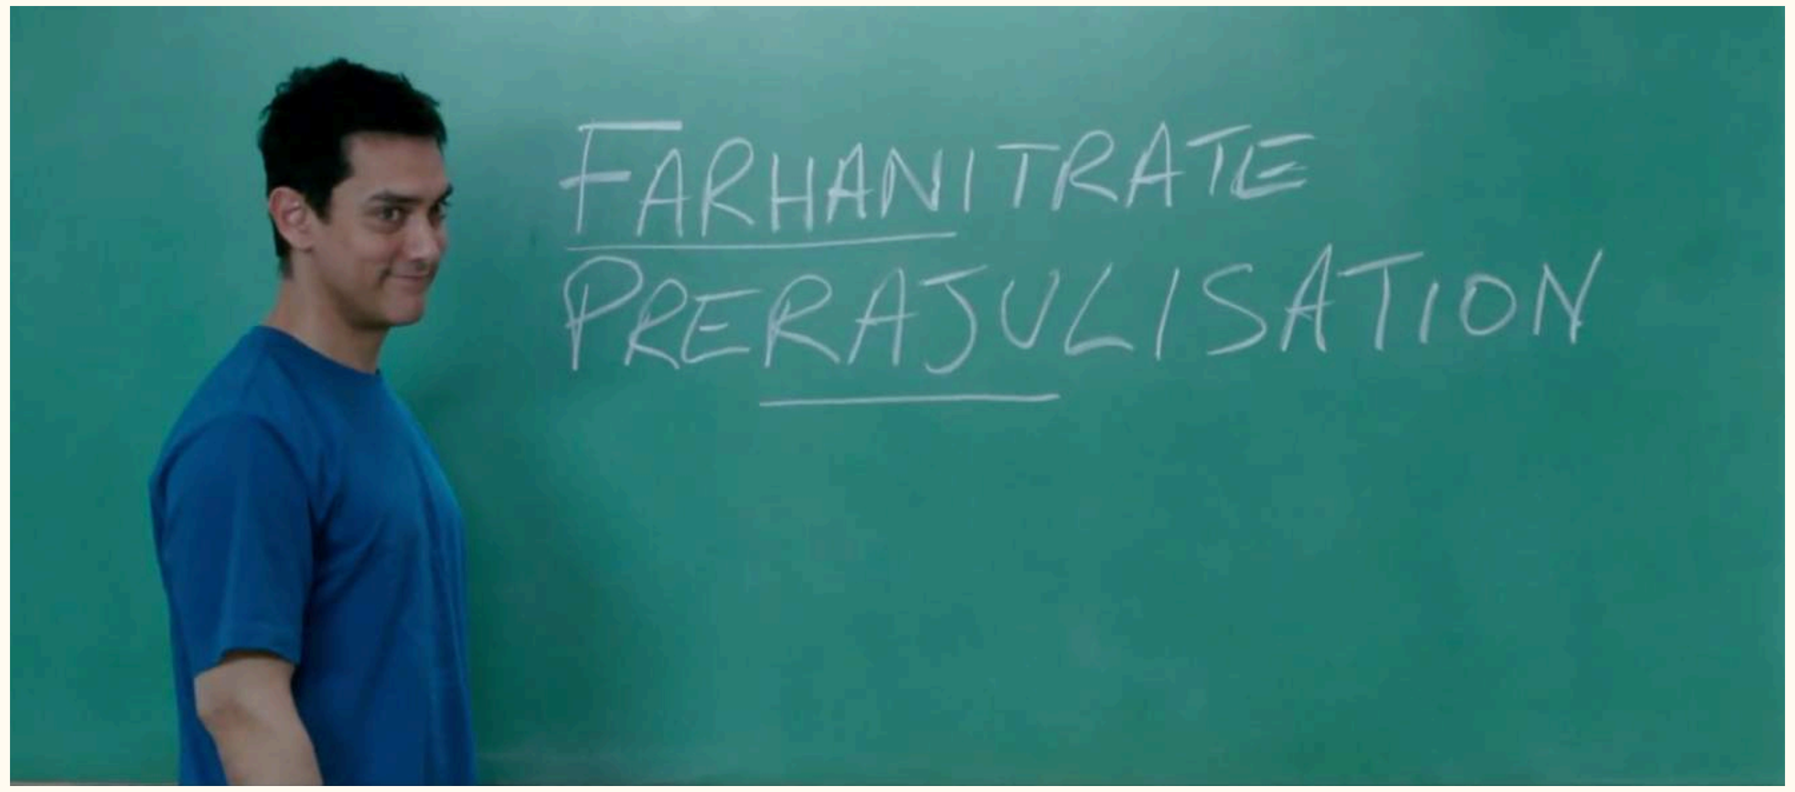

Summary

- Visual recognition

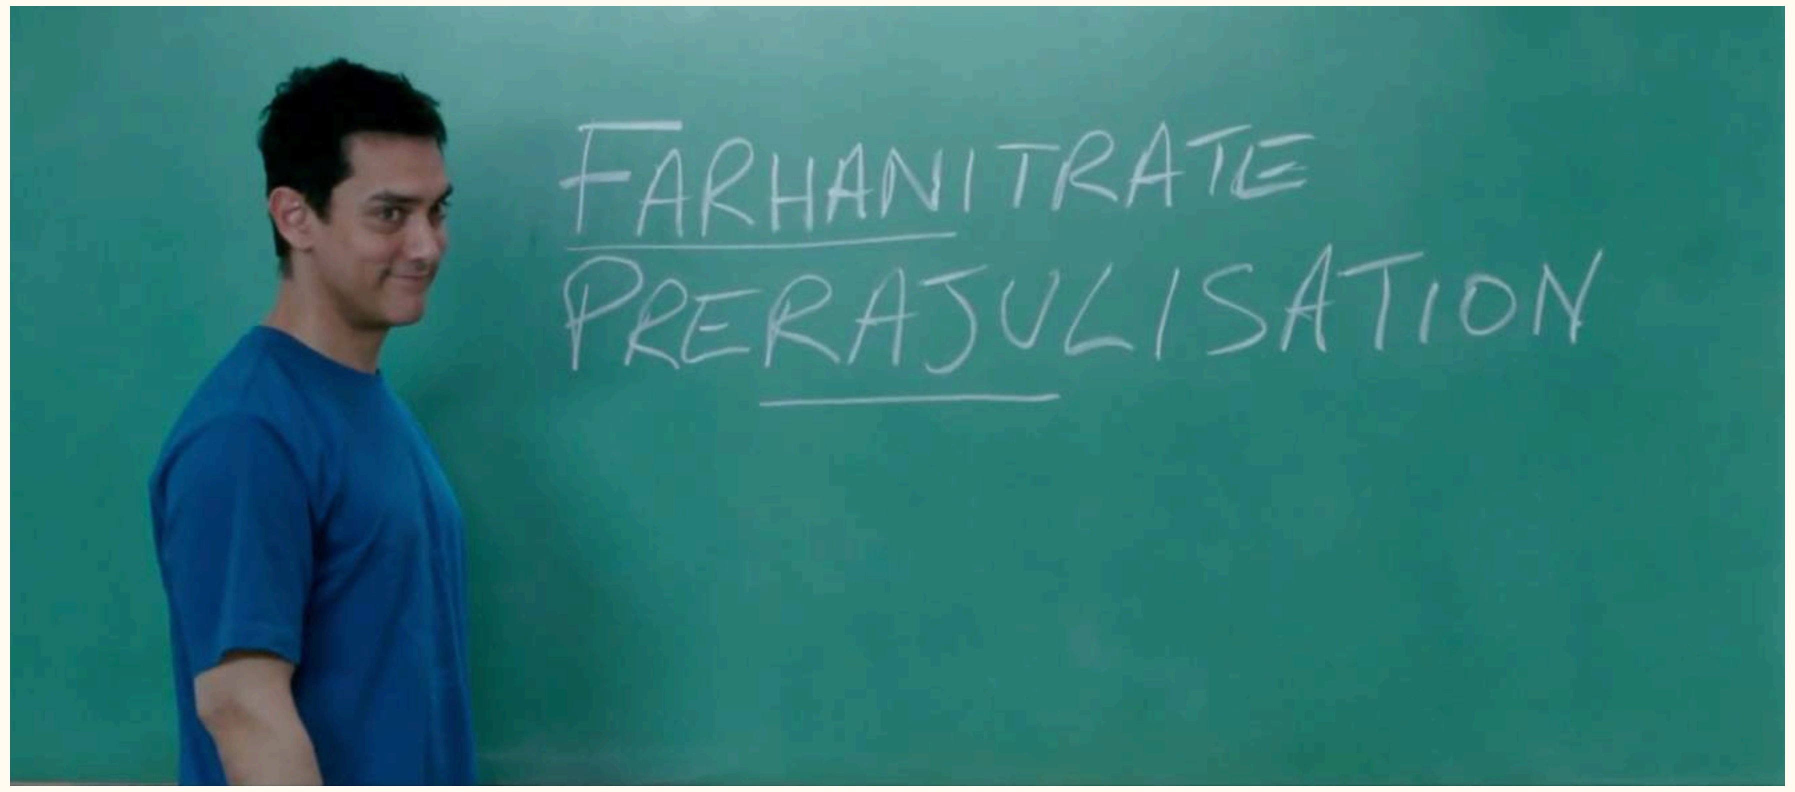

Summary

- Visual recognition
- Convolutions

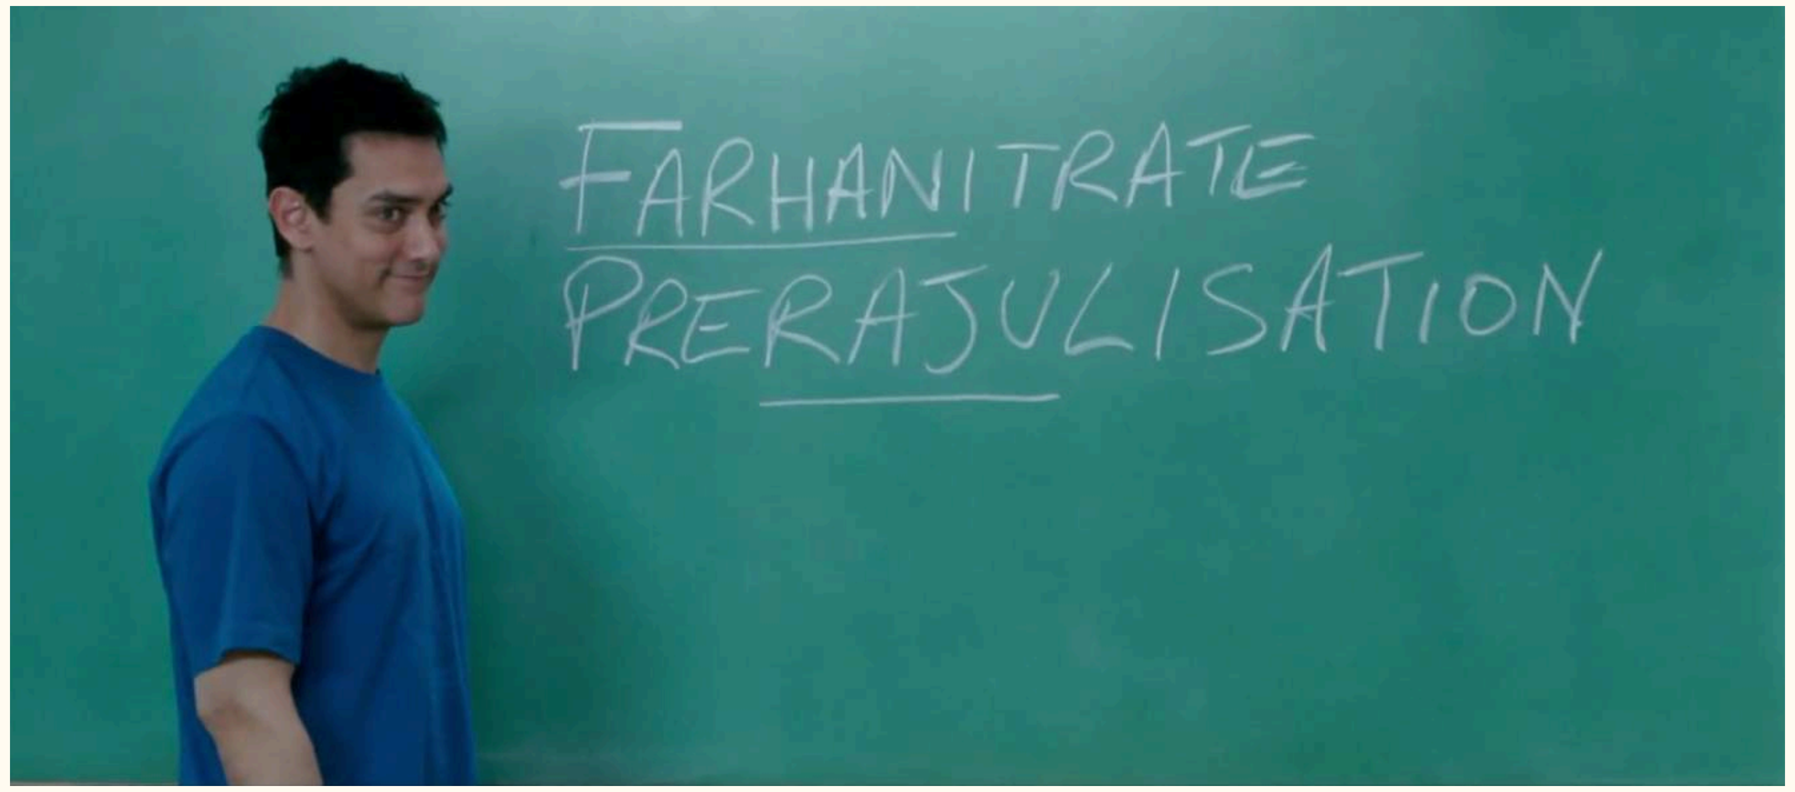

Summary

- Visual recognition
- Convolutions
  - Filters

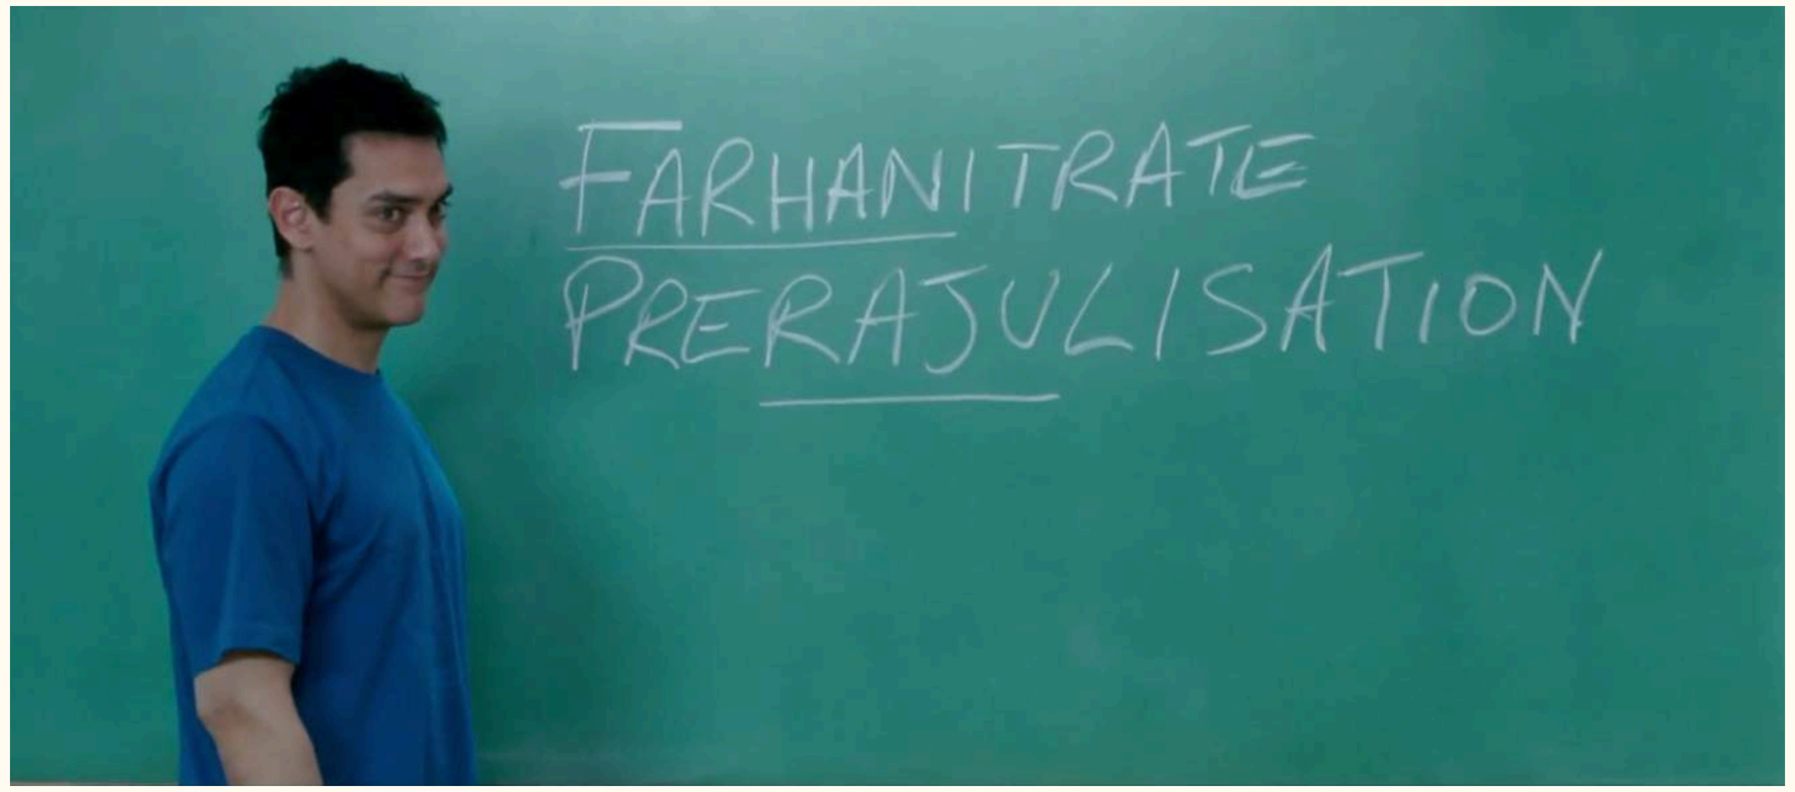

Summary

- Visual recognition
- Convolutions
  - Filters
  - Feature Maps

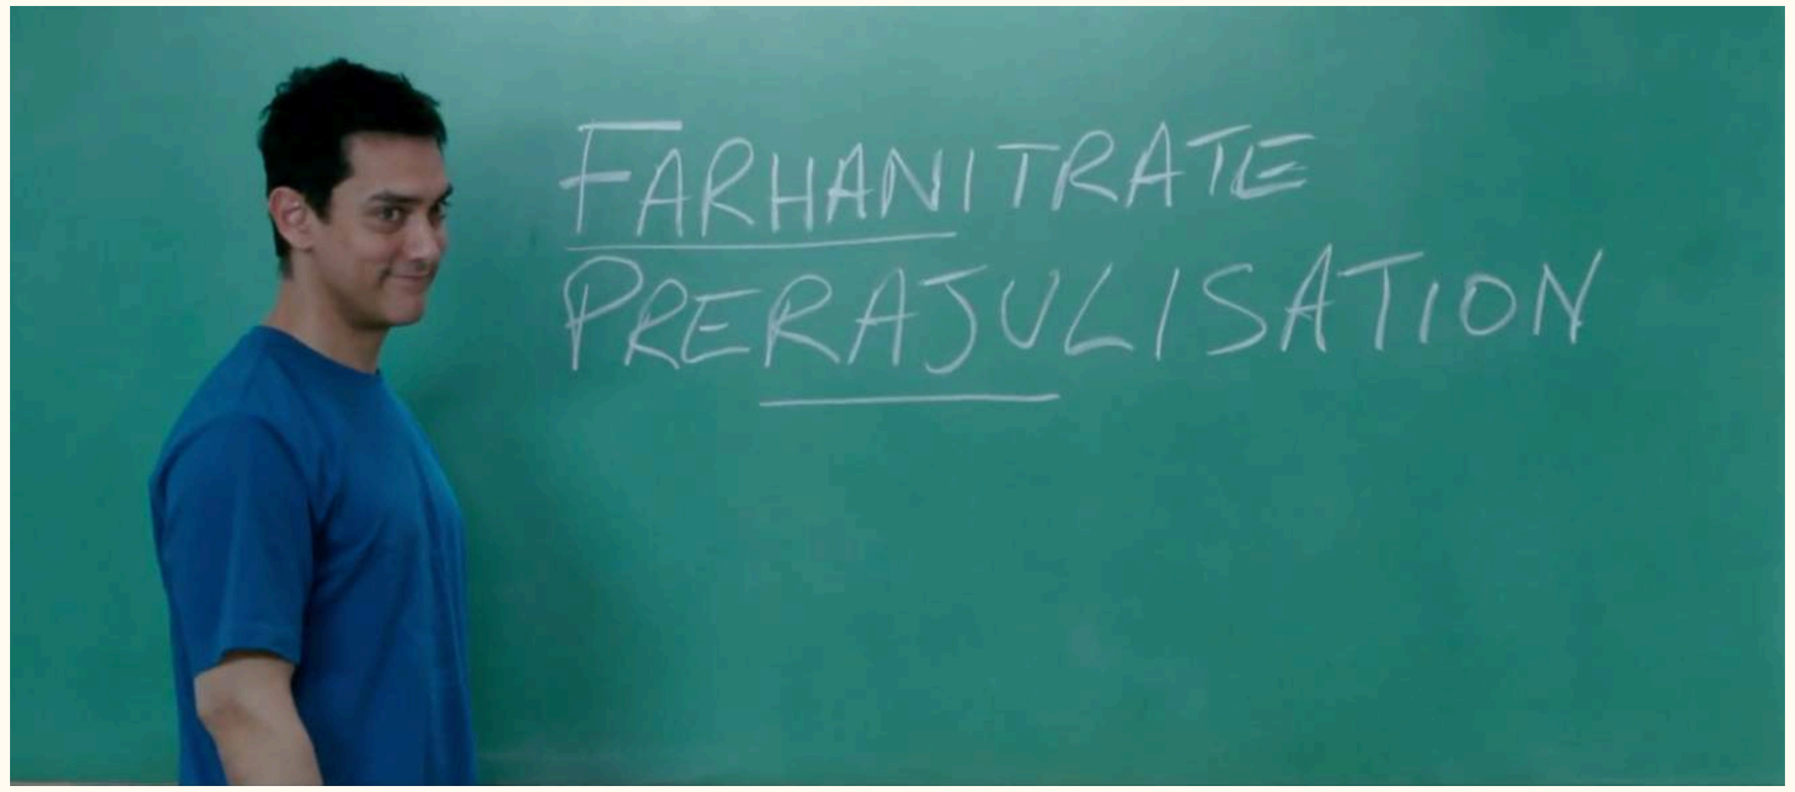

Summary

- Visual recognition
- Convolutions
  - Filters
  - Feature Maps
  - Architectures

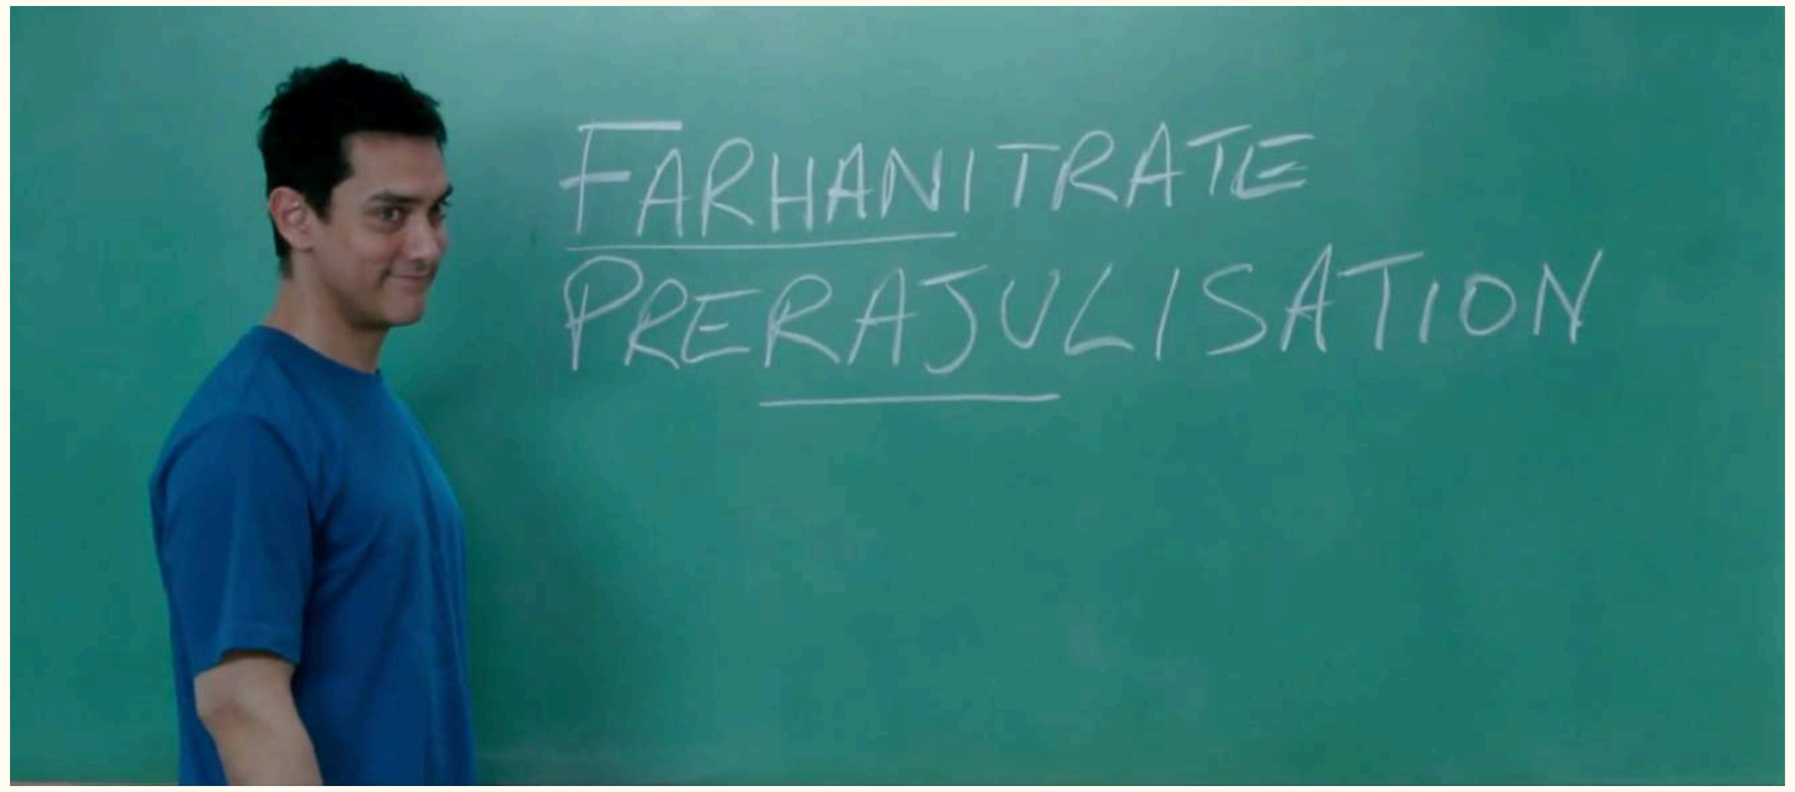

Summary

- Visual recognition
- Convolutions
    - Filters
    - Feature Maps
    - Architectures
    - Strides

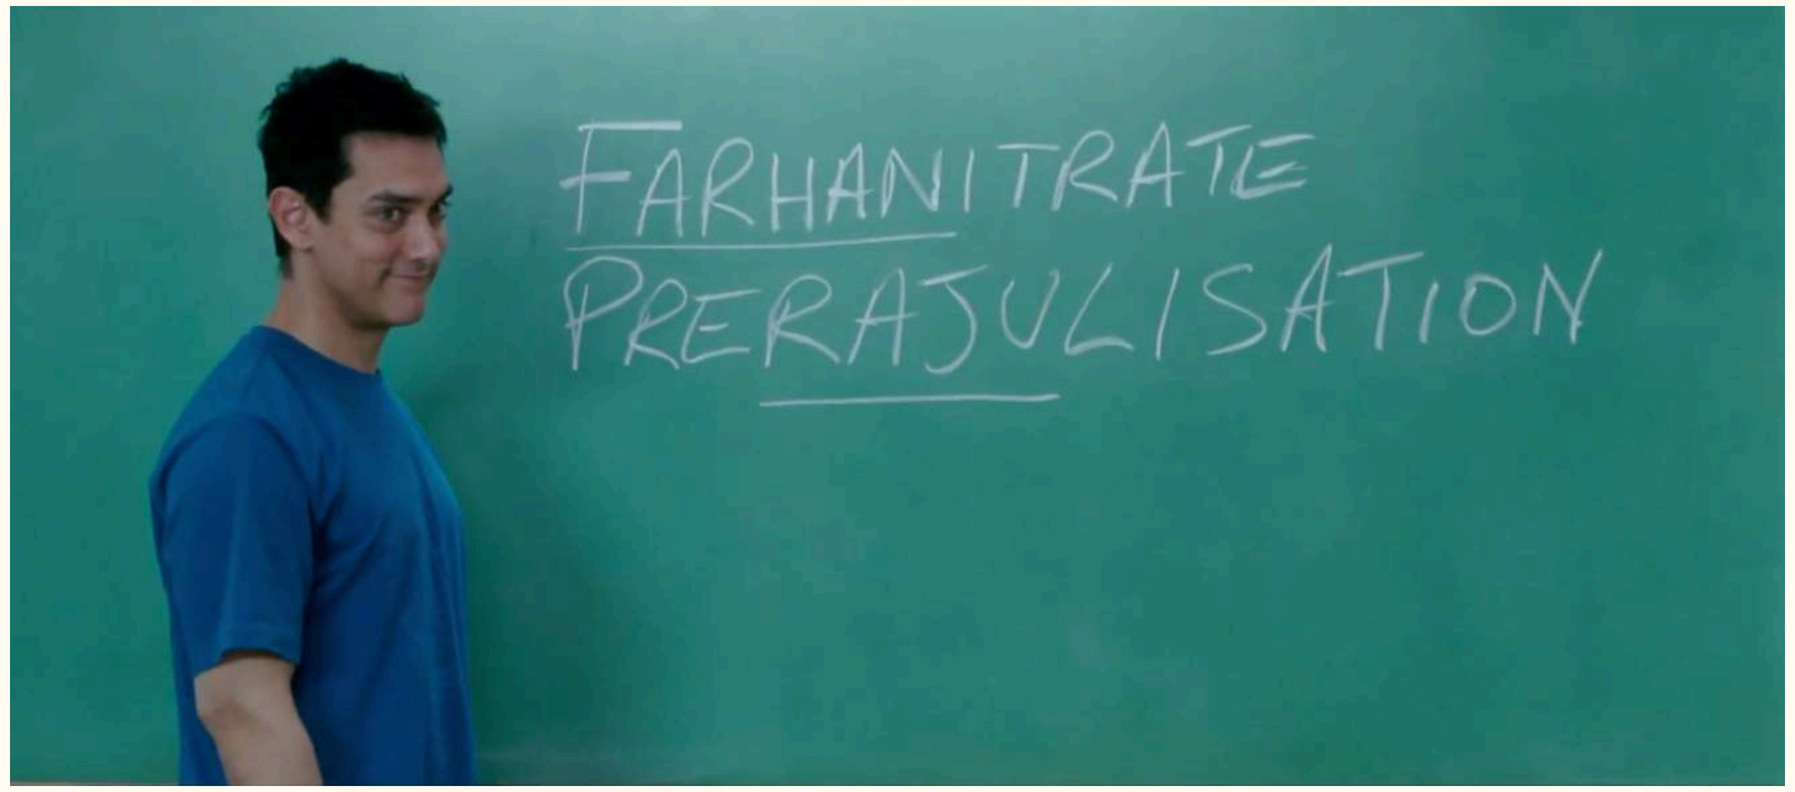

Summary

- Visual recognition
- Convolutions
  - Filters
  - Feature Maps
  - Architectures
  - Strides
  - Layers

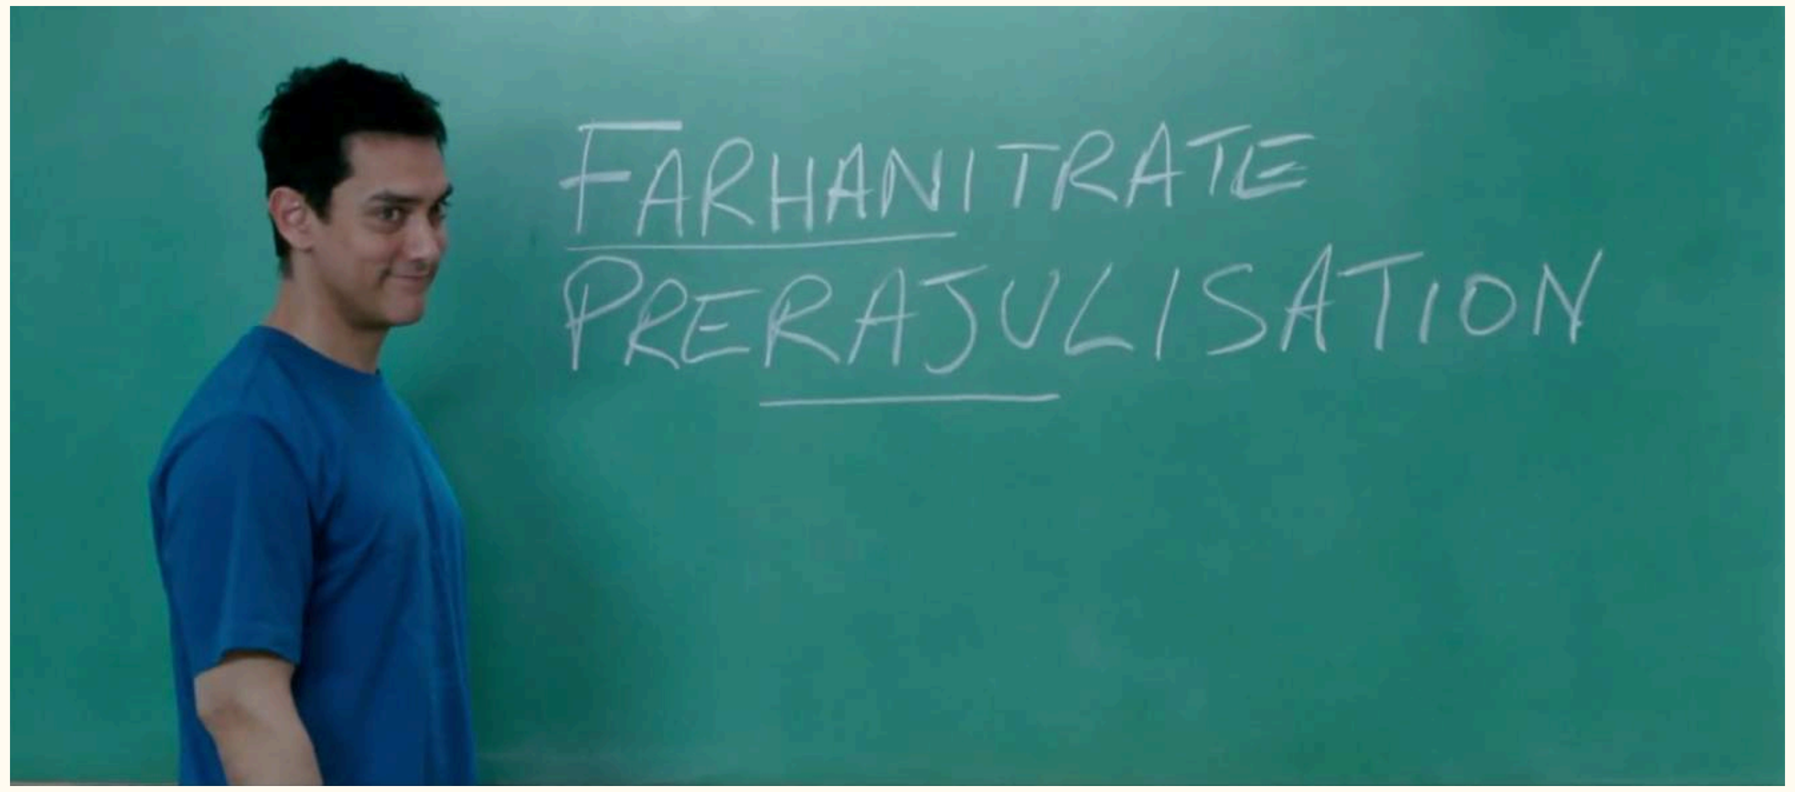

Summary

- Visual recognition
- Convolutions
  - Filters
  - Feature Maps
  - Architectures
  - Strides
  - Layers
- Handling complexity