Day 04

Wan Fang
Southern University of Science and Technology

Day 4

- 8:00 – 8:50 Review of the Initial Ideation
- 9:00 – 10:30 Lecture : Convolutional Networks/RNN/GAN
  Exercise: Play the notebooks
- 10:50 – 12:10 Workshop: Data collection

- 2:00 – 5:00 Exercise: Prototyping + testing with data/model
- 5:00 – 6:00 Review of the day

AI –> Machine Learning–> Deep Learning

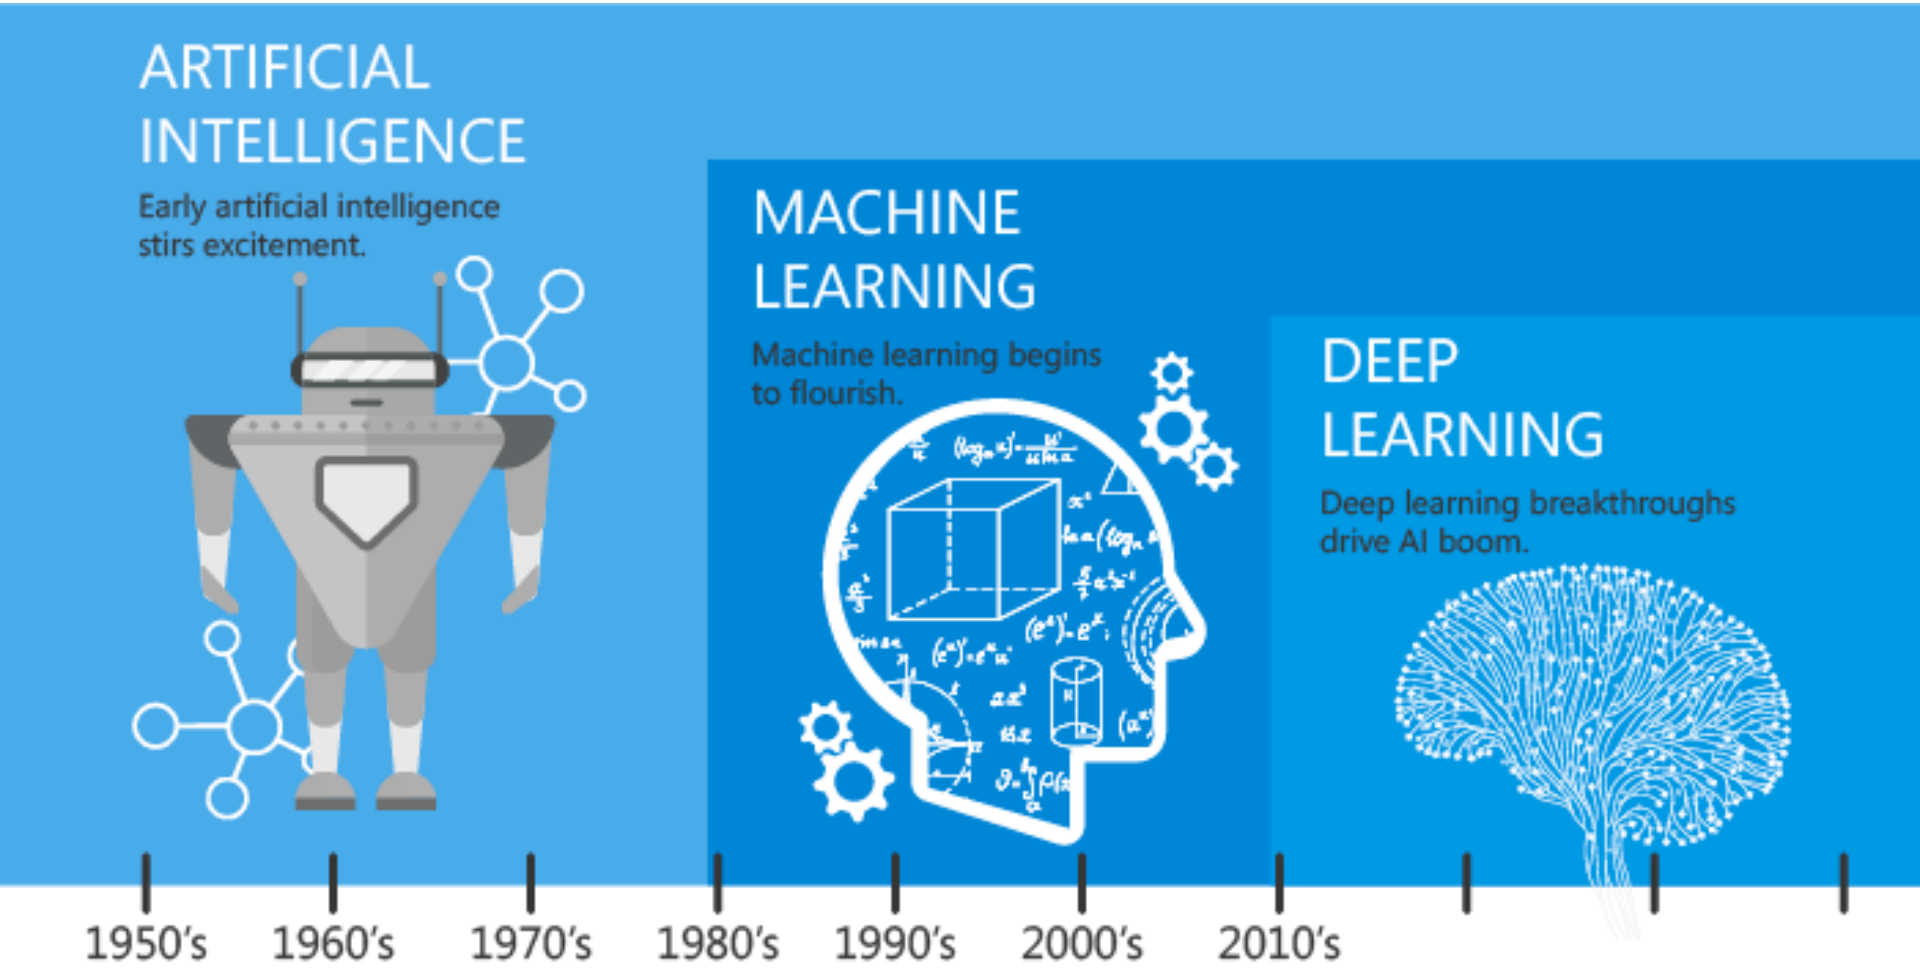

Image: Linked In | Machine Learning vs Deep learning

Convolutional Networks

Convolutional Networks Applications

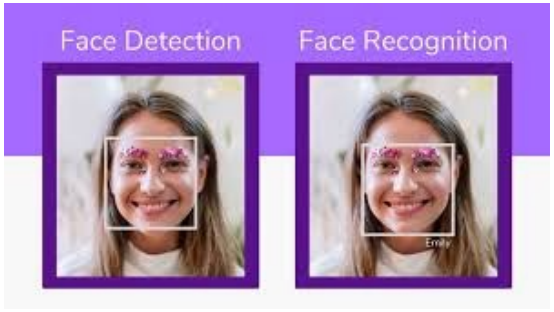

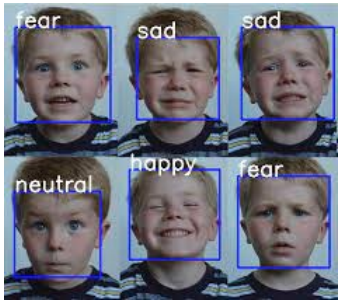

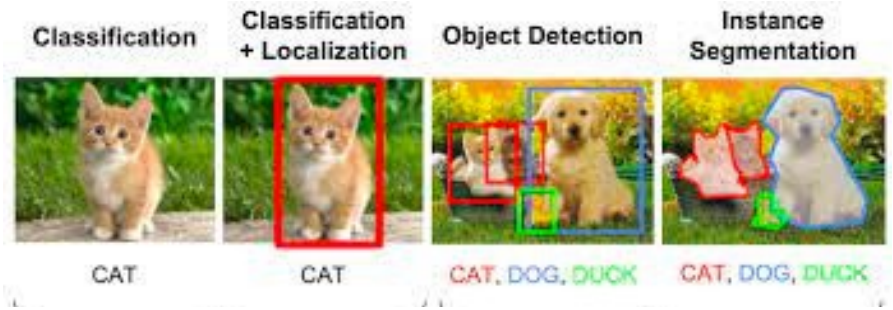

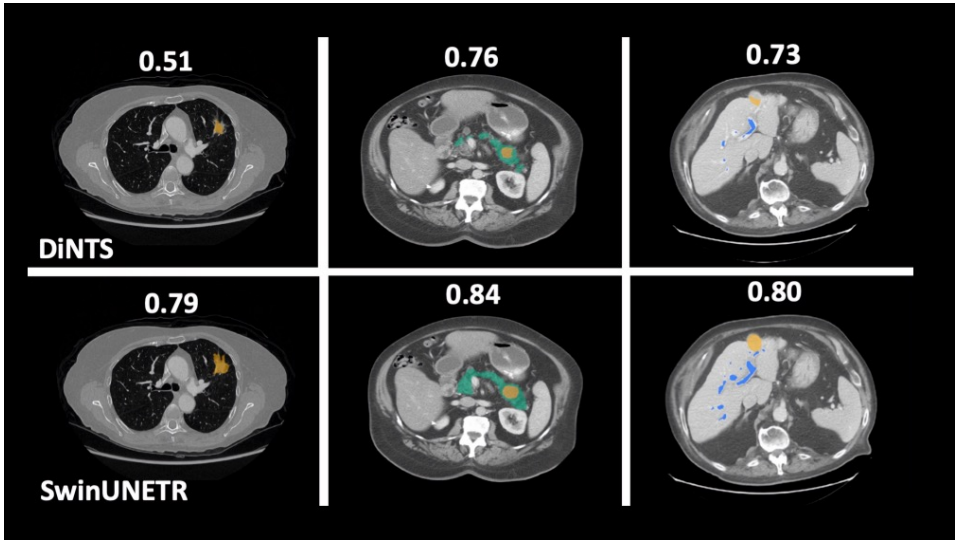

A Design Challenge with Increasing Dimensions

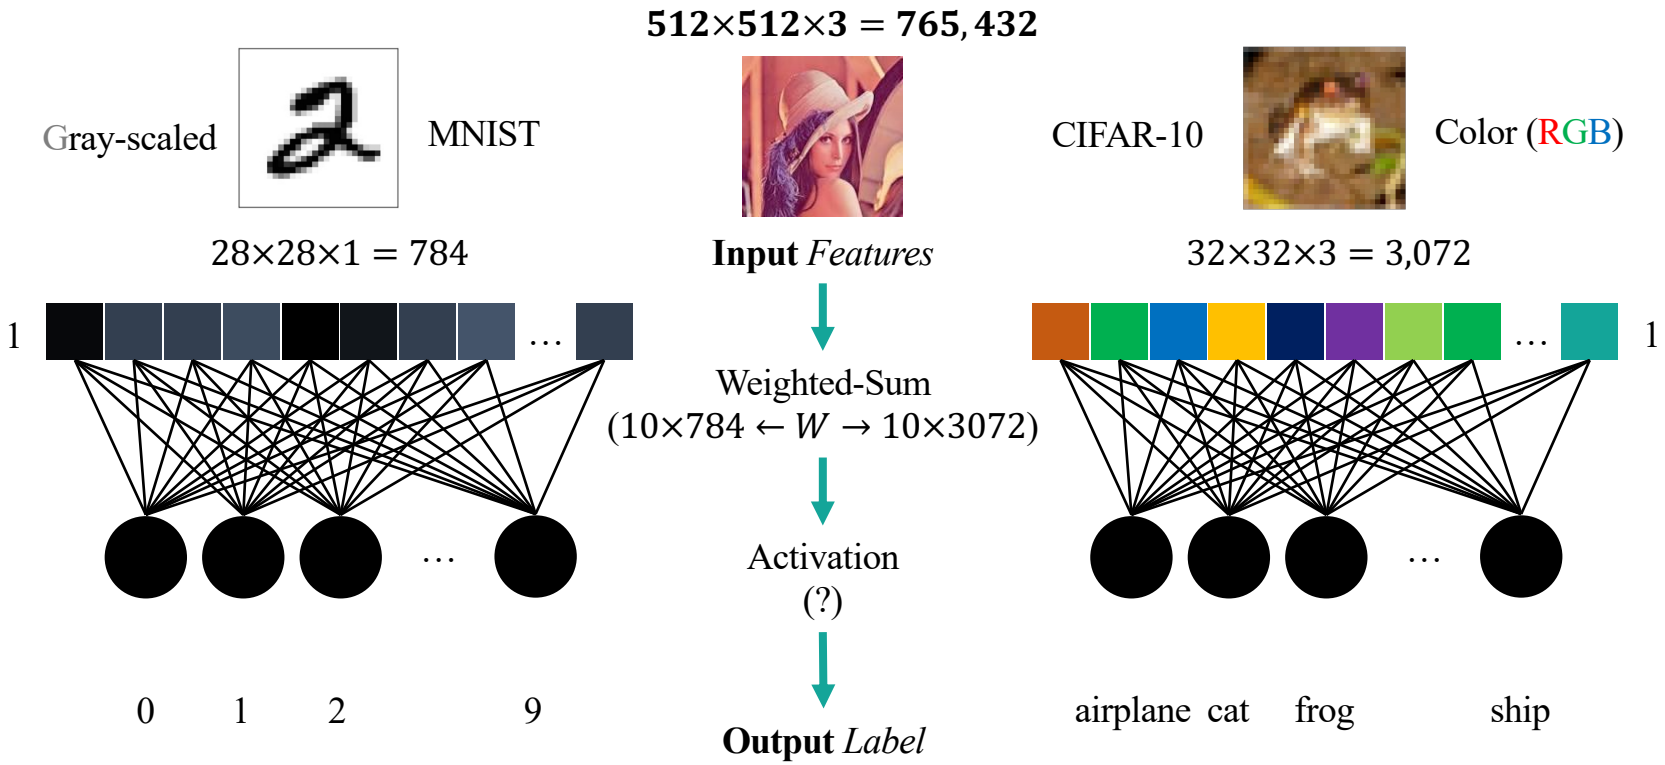

Regular Neural Nets don't scale well to full images

Convolutional Operation

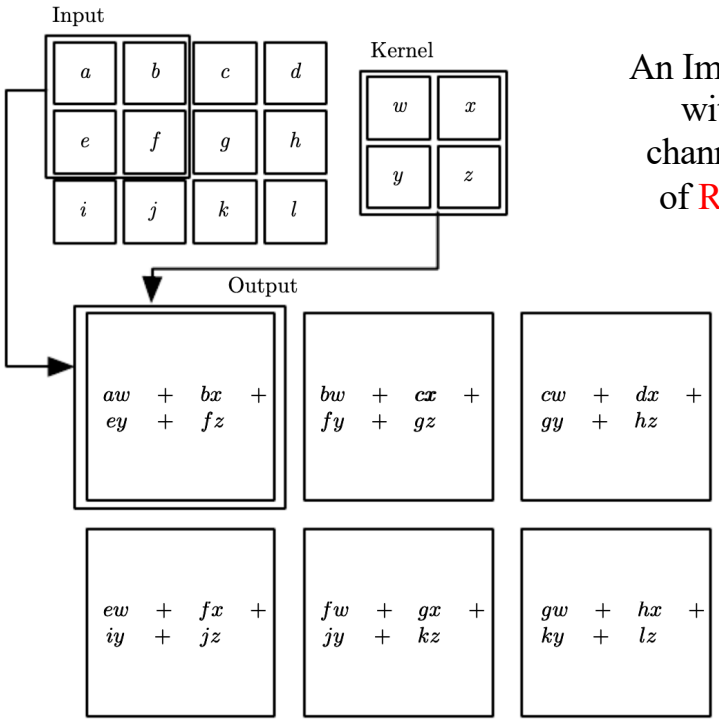

An Image with 3 channels of RGB

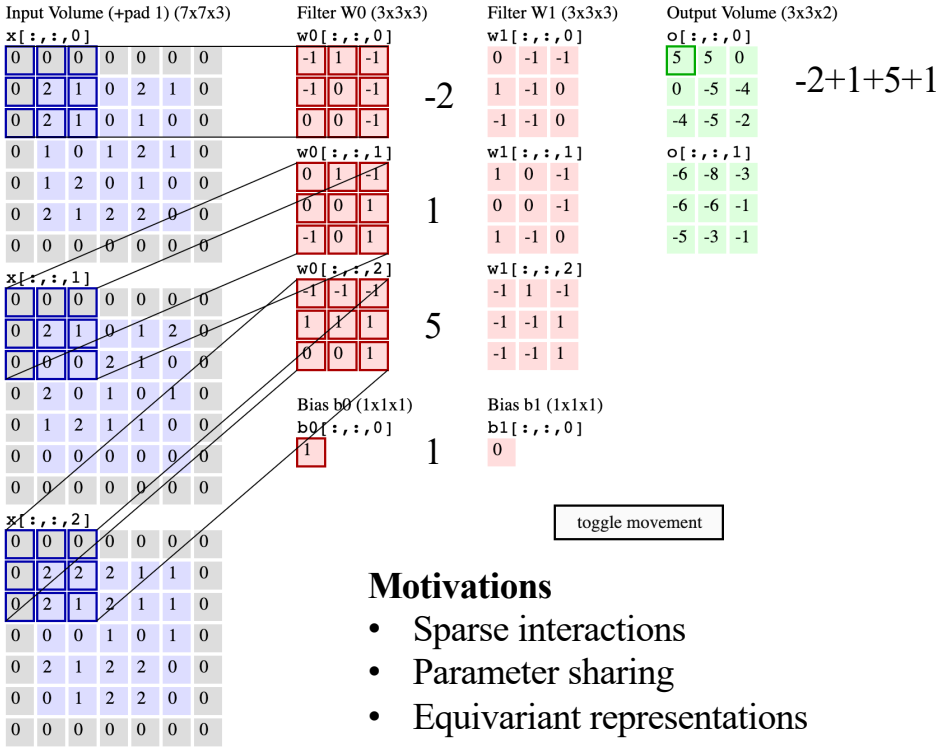

Motivations

- Sparse interactions
- Parameter sharing
- Equivariant representations

Convolution in 3D Volumes

Preserved spatial structure between the input and output volumes in width, height, number of channels

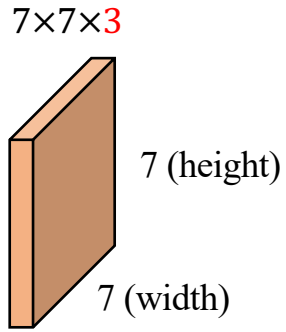

Input Volume in 3D

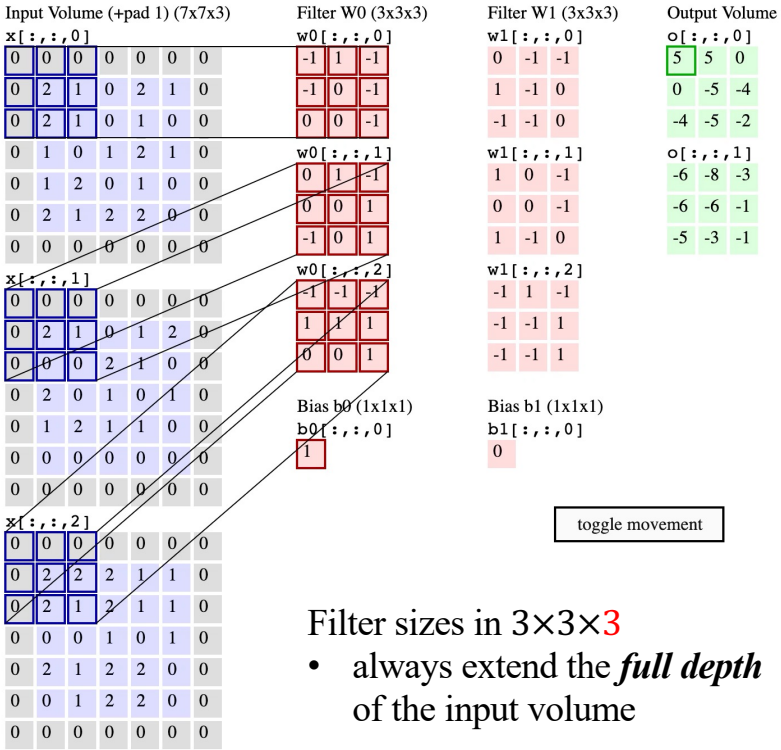

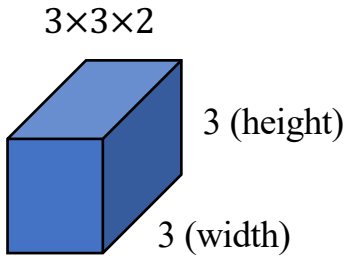

Output Volume in 3D

Layers in a ConvNet:
Transform an input 3D volume to an output 3D volume with some differentiable function that may or may not have parameters.

Filter sizes in 3×3×3
- always extend the ***full depth*** of the input volume

Convolve the filter with the image
i.e. "*slide over the image spatially, computing dot products*"

The Design of a Convolutional Layer

A convolutional layer is defined by the filter (or kernel) size, the number of filters applied and the stride

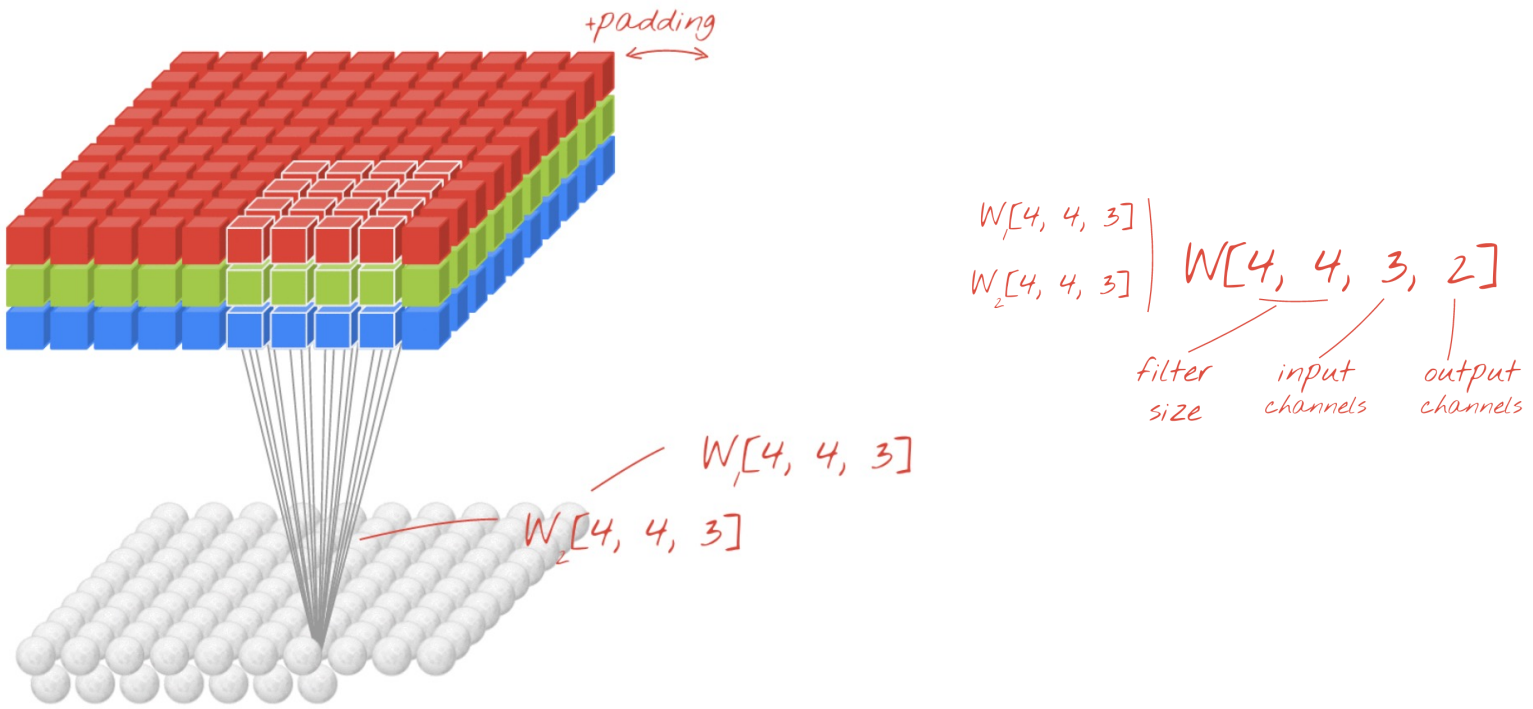

Output Volume Size

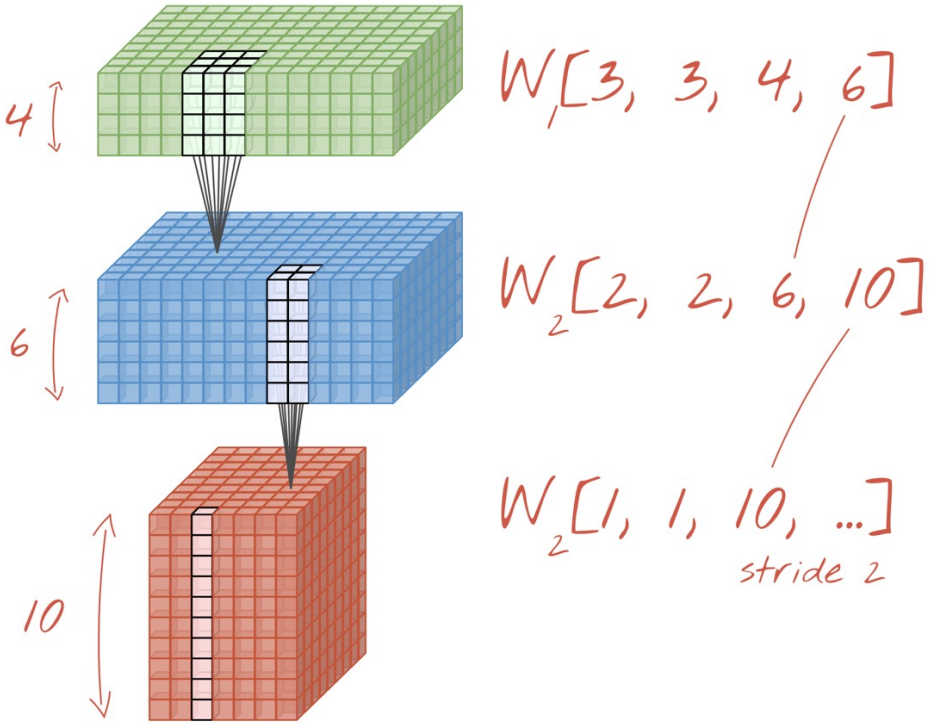

- Depth (number of channels):
  - *adjusted by using more or fewer filters*
- Width & Height:
  - *adjusted by using a stride >1*
  - *(or with a max-pooling operation)*

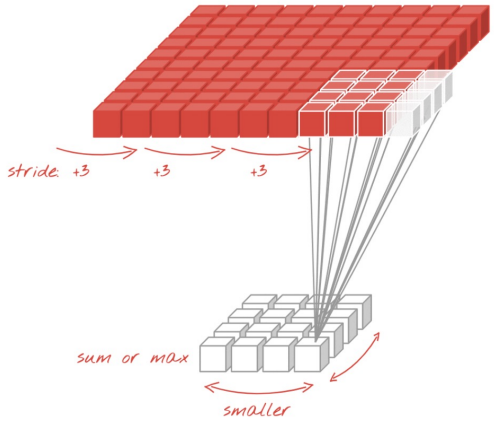

Defined by the filter (or kernel) size, the number of filters applied and the stride

The Last Layer

From a Cubic Volume in 3D to predicted labels

Similar like a normal neural network

Expensive in #weights

But preseves the location data (*x*, *y*)

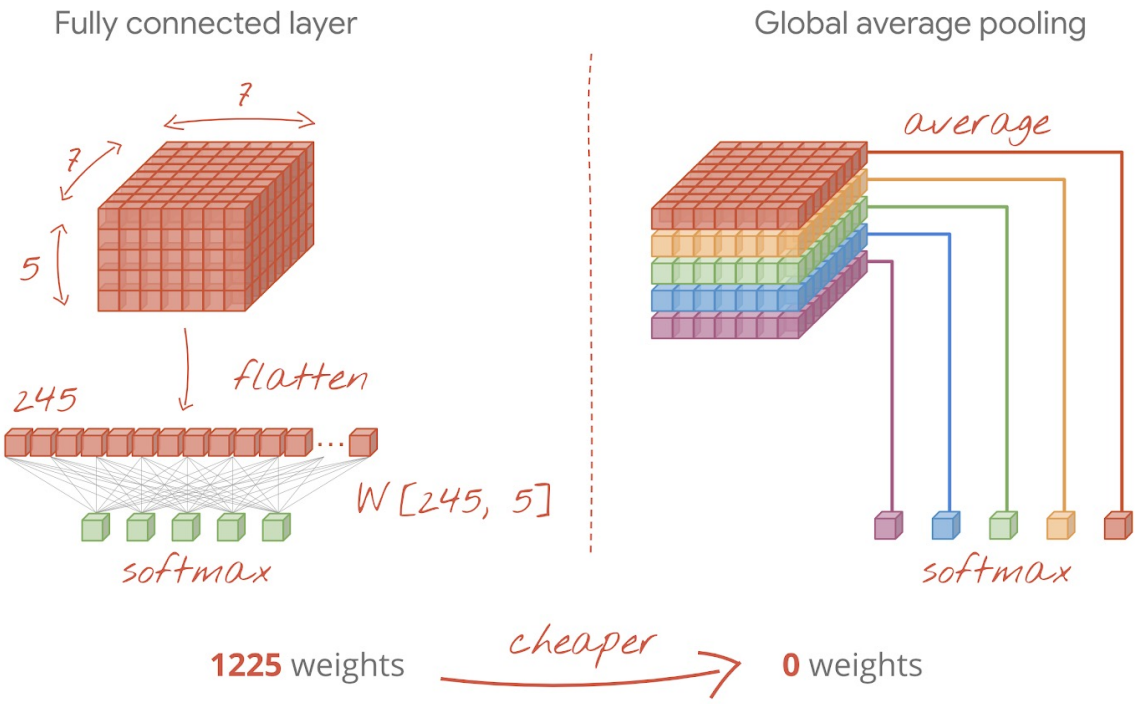

Much lighter in calculation

The average pooling explicitly discards all location data

Stacking Up a ConvNet

Layer-by-layer

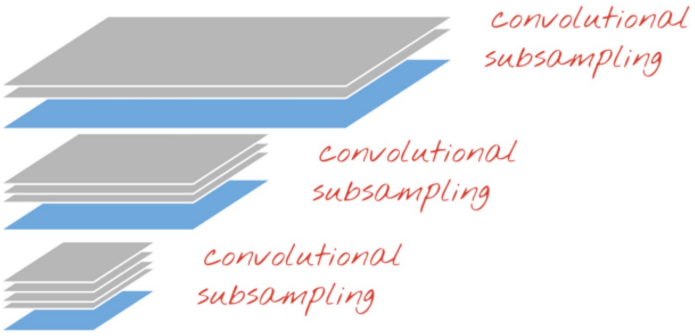

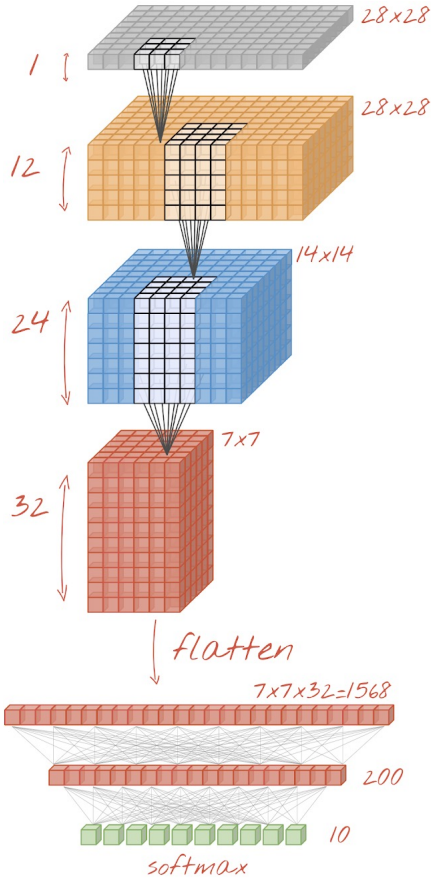

Convolutional 3x3 filters=12
$W_1[3, 3, 1, 12]$

Convolutional 6x6 filters=24
$W_2[6, 6, 12, 24]$ stride 2

Convolutional 6x6 filters=32
$W_2[6, 6, 24, 32]$ stride 2

Dense layer
$W_4[1568, 200]$

Softmax dense layer
$W_5[200, 10]$

Recurrent Neural Networks

Supervised Learning with Sequences

- Many supervised learning problems deals with ordered sequences
  - Financial time series
    - Input: ordered sequence of past series values
    - Output: ordered sequence of future series values

Supervised Learning with Sequences

- Many supervised learning problems deals with ordered sequences
  - Natural Languge Processing
    - Input: ordered sequence of words or characters
    - Output: ordered sequence of characters

Supervised Learning with Sequences

- Many supervised learning problems deals with ordered sequences
  - Machine Translation
    - Input: ordered sequence of words (language X)
    - Input: ordered sequence of words (language Y)

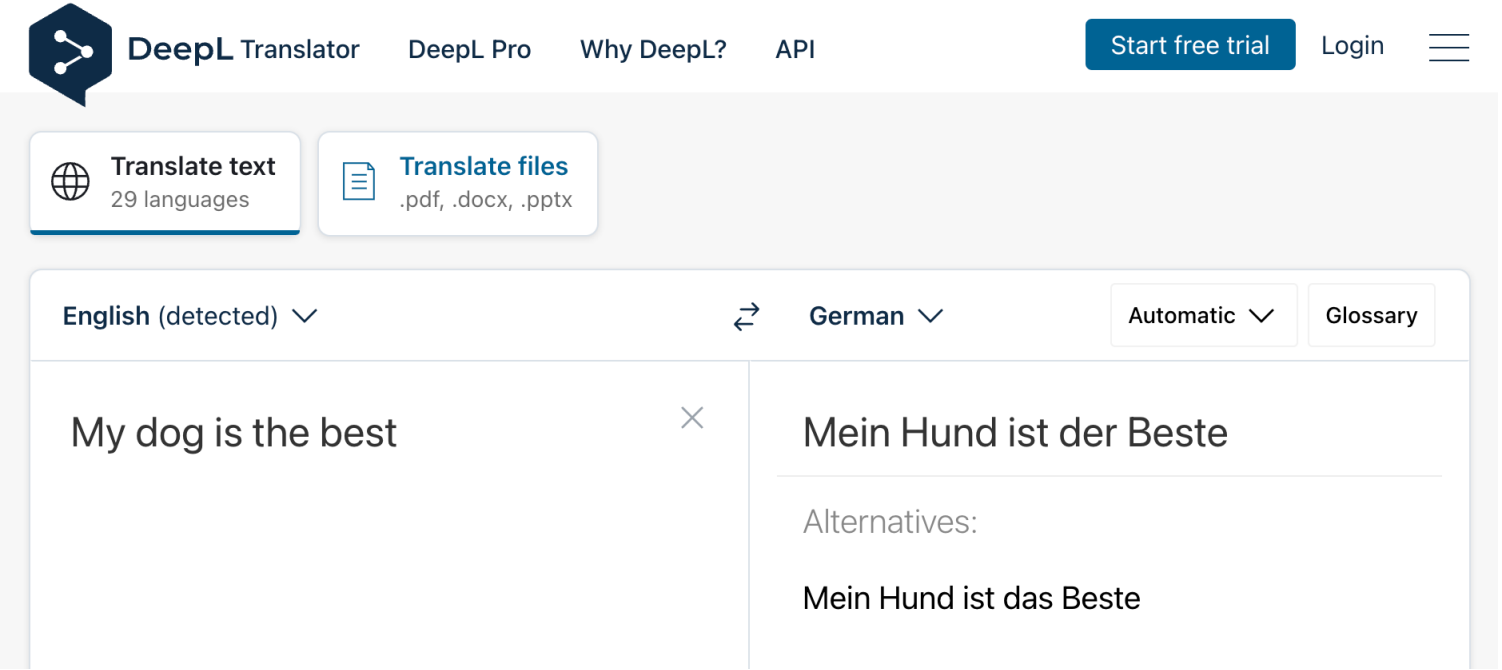

Supervised Learning with Sequences

- Many supervised learning problems deals with ordered sequences
  - Speech Recognition
    - Input: ordered sequence of audio signal
    - output: ordered sequence of words

Why Recurrent Neural Networks?

- RNN were created because there were a few issues in the feed-forward neural network:
  - Cannot handle sequential data
  - Considers only the current input
  - Cannot memorize previous inputs

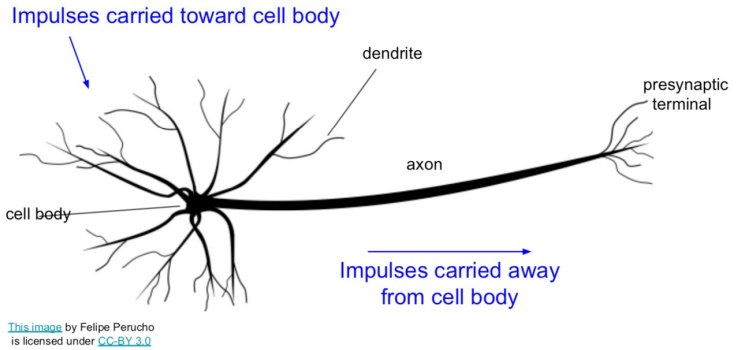

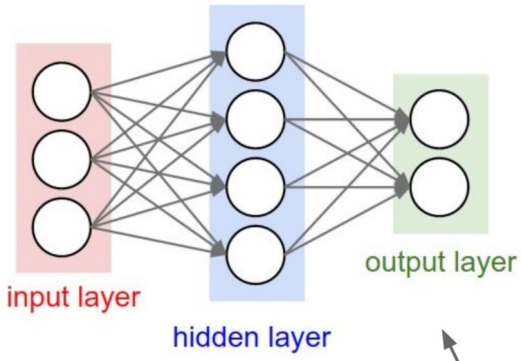

"2-layer Neural Net", or
"1-hidden-layer Neural Net"

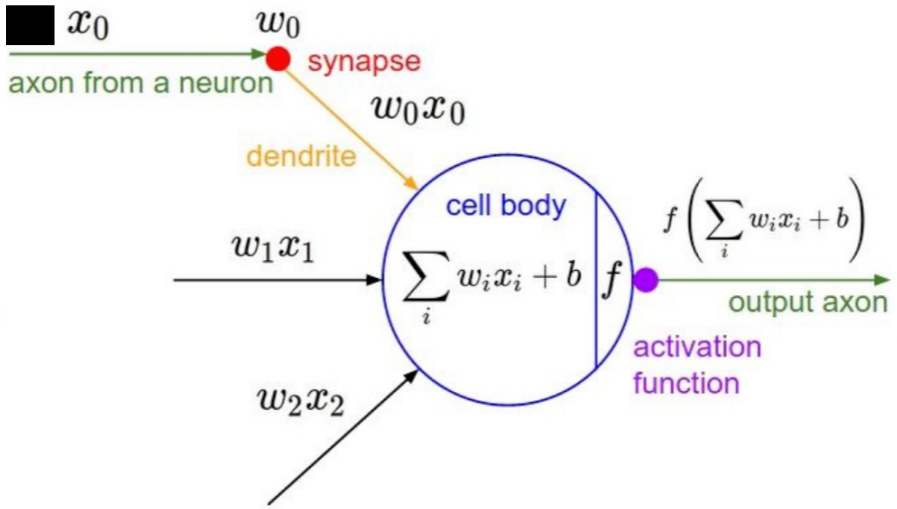

"Fully-connected" layers

What Is a Recurrent Neural Network (RNN)?

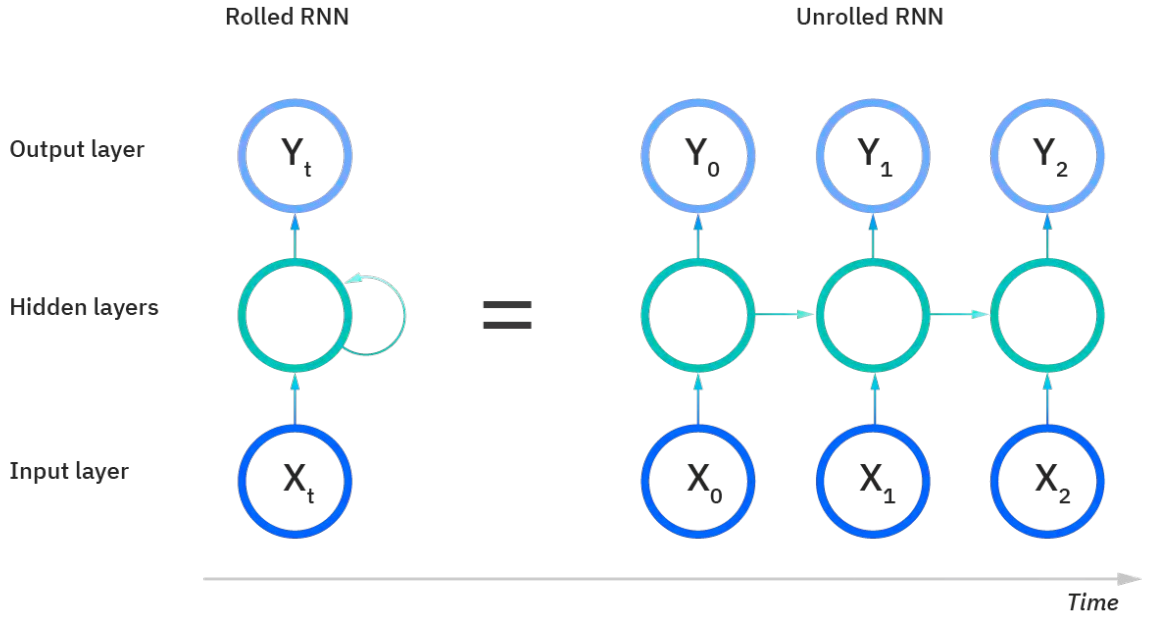

- RNN are distinguished by their "memory" as they take information from prior inputs to influence the current input and output.
- RNN share parameters across each layer of the network

https://www.ibm.com/cloud/learn/recurrent-neural-networks

Types of Recurrent Neural Network (RNN)?

One-to-one:

One-to-many:

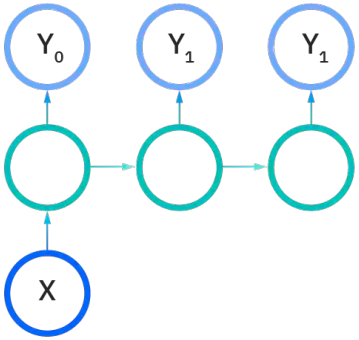

Many-to-one:

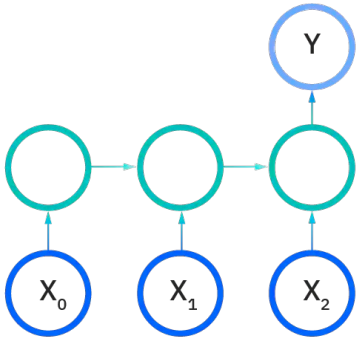

Many-to-many:

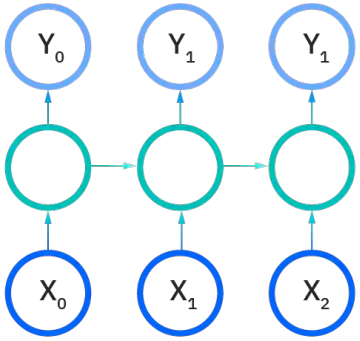

Many-to-many:

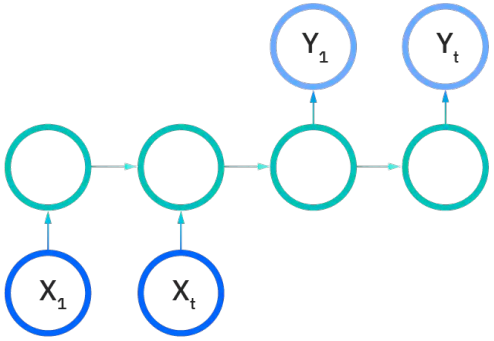

https://www.ibm.com/cloud/learn/recurrent-neural-networks

Variant RNN architectures

- Long Short-Term Memory Networks
  - LSTMs are a special kind of RNN — capable of learning long-term dependencies by remembering information for long periods is the default behavior.

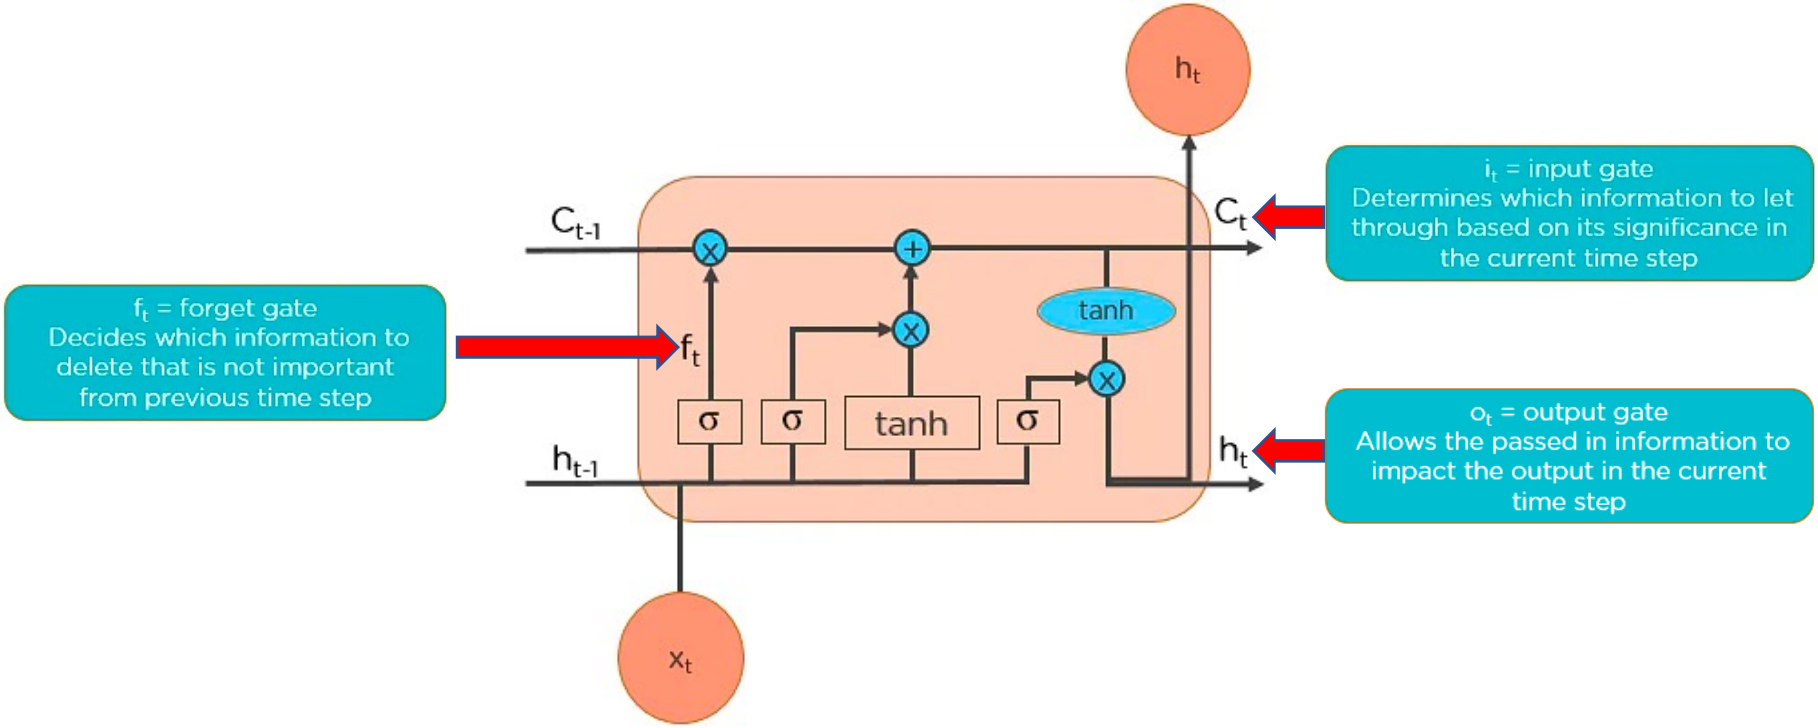

Generative Adversarial Networks

Generative Adversarial Networks

- GANs can be trained on the images of
  - humans to generate realistic faces.
  - cartoon characters for generating faces of anime characters as well as Pokemon characters.

Text to Image

What are Generative Adversarial Networks?

- Generative Adversarial Networks (GANs) were introduced in 2014 by Ian J. Goodfellow
- GANs perform unsupervised learning tasks in machine learning.
- It consists of 2 models that automatically discover and learn the patterns in input data.

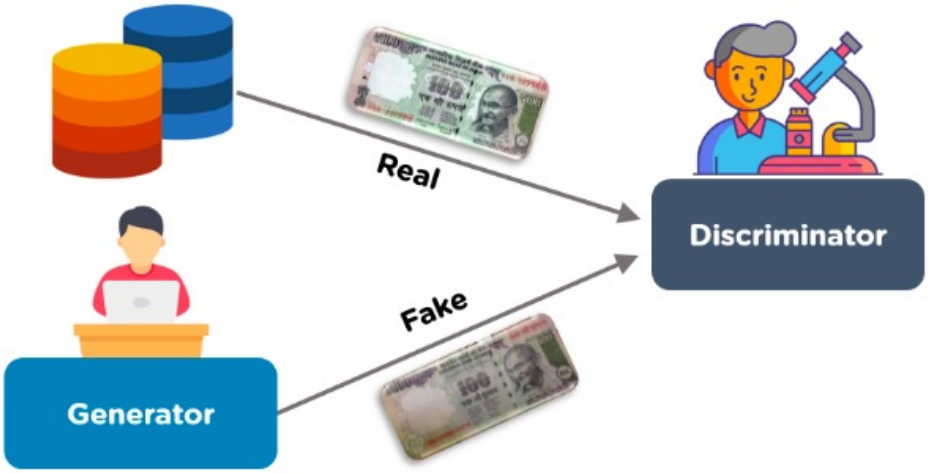

What are Generative Adversarial Networks?

- A Generator in GANs is a neural network that creates fake data to be trained on the discriminator. It learns to generate plausible data.

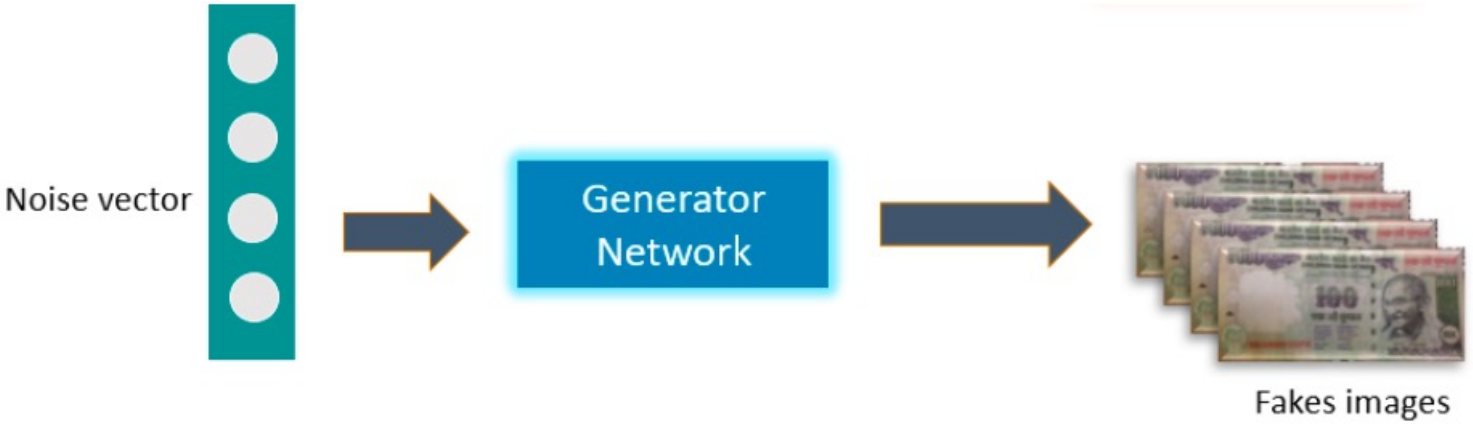

What are Generative Adversarial Networks?

- The Discriminator is a neural network that identifies real data from the fake data created by the Generator. The discriminator's training data comes from different two sources

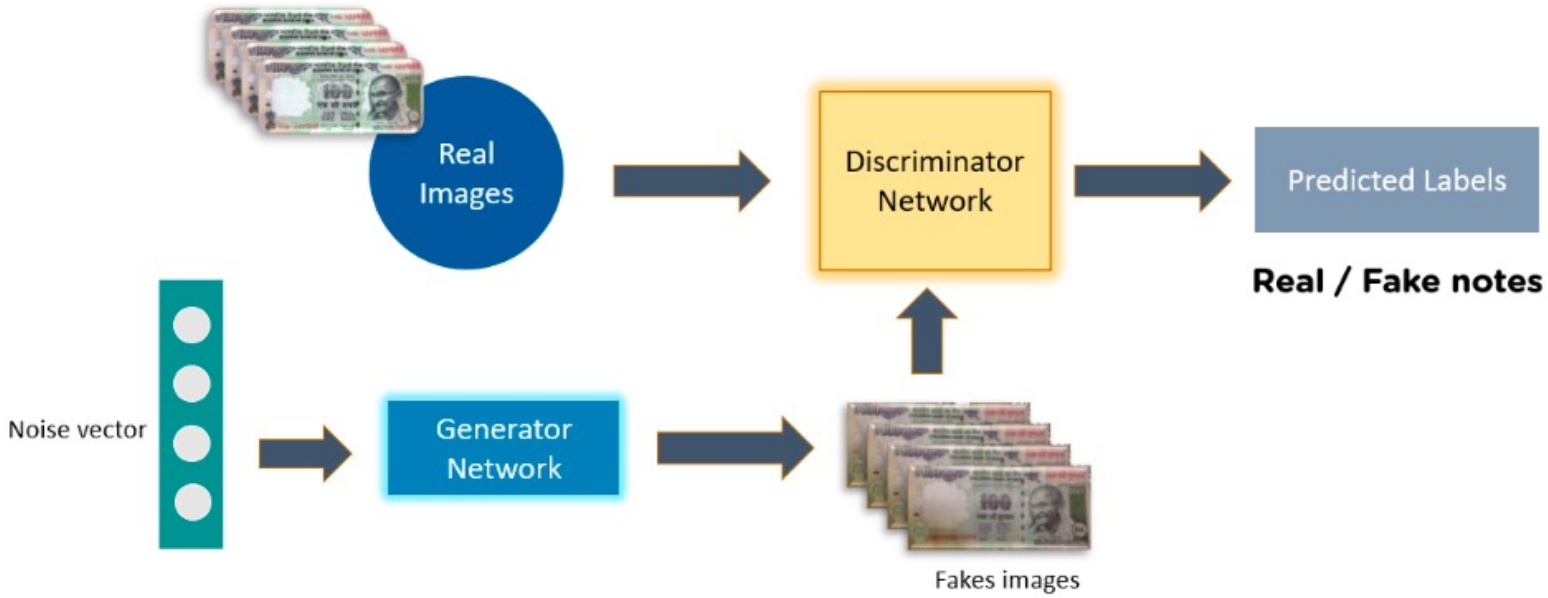

Steps for Training GAN

1. Define the problem
2. Choose the architecture of GAN
3. Train discriminator on real data
4. Generate fake inputs for the generator
5. Train discriminator on fake data
6. Train generator with the output of the discriminator

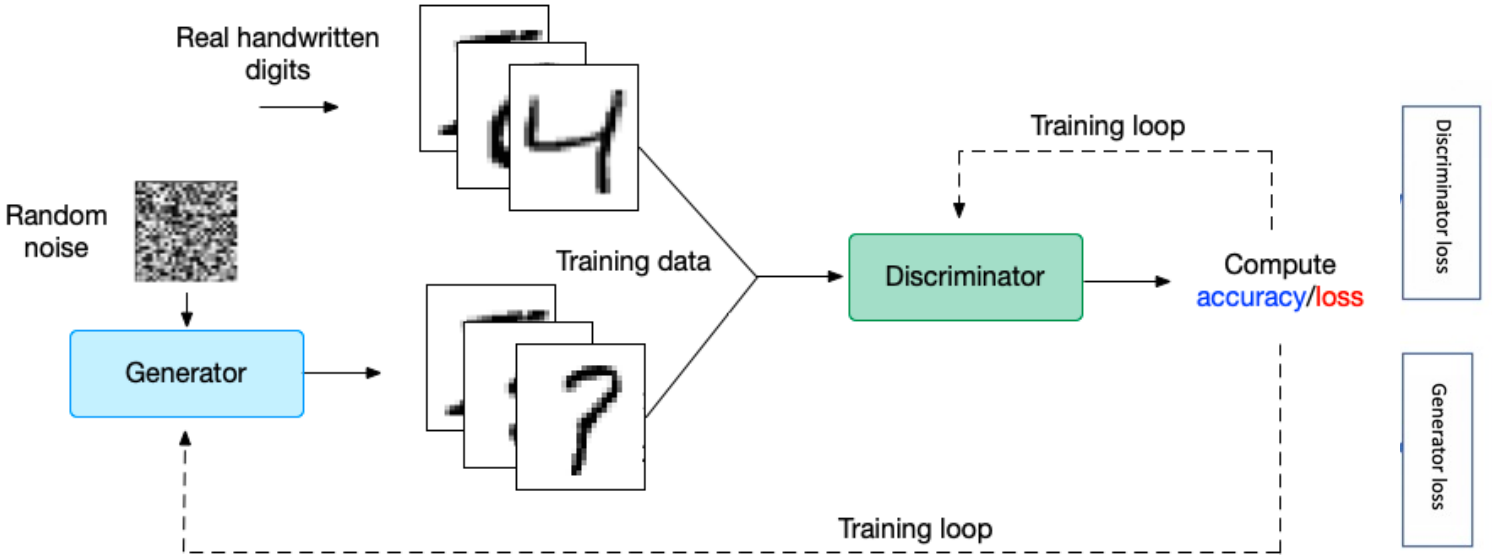

Exercise with Julia

- CNN: Handwritten digits classification
- RNN: AI Generates Shakespeare-like text
- Deep Convolutional GANs (DCGANs): Generate images from noise

Workshop: Data collection

结构化通用表征非结构环境的人机交互系统

应用场景

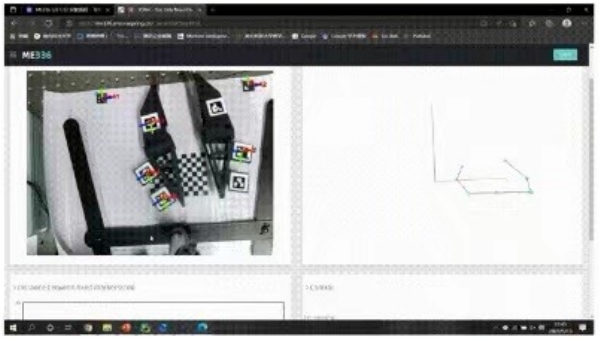

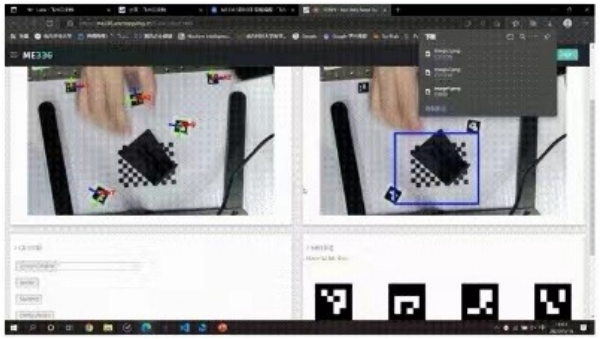

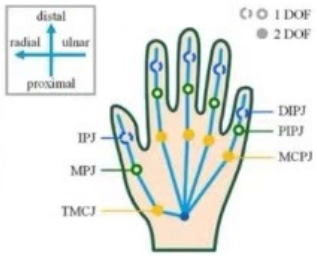

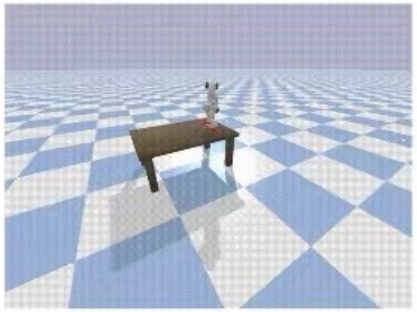

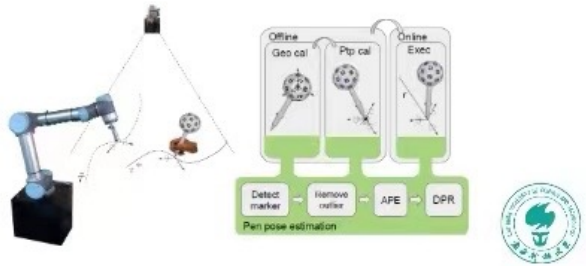

User Mannual

- 识别区域

实时显示摄像头画面并且检测二维码

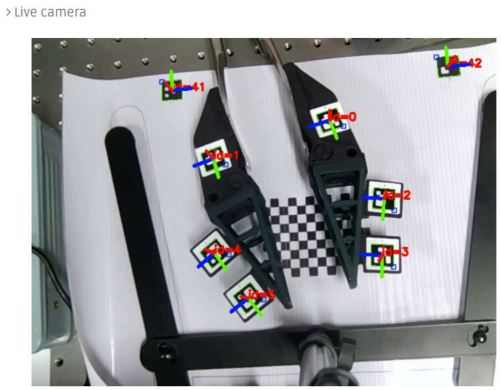

- 数据流绘图区域

该区域实时显示桌垫上两二维码的距离，请在该距离基本稳定时进行实验，如出现剧烈、频繁的波动，则存在光照等影响识别的因素，请适当调整位置使其保持稳定。

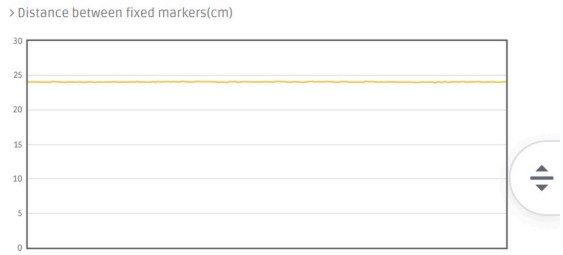

- Webgl辅助区域

对定义完后的物理意义进行可视化，直观的感受区别

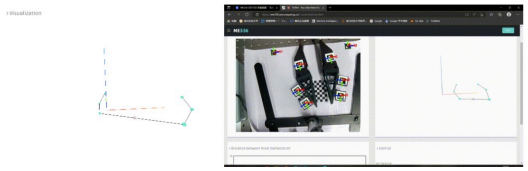

- 控制区域

该区域为控制的按钮

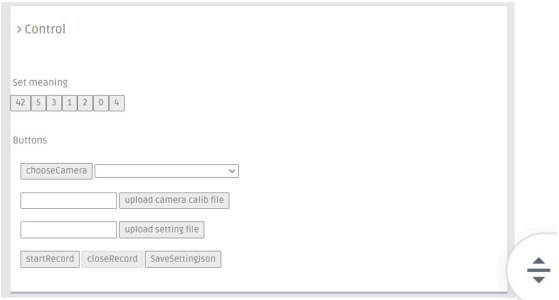

当视野中识别到二维码时，Set meaning区域会出现相对应ID的按钮，点击可以进行物理意义的赋予

• DS323: AI in Design

- 意义赋予流程

1. 该id二维码所对应的物理意义

输入物理意义后，可视化模块会出现该id对应的青色方块

2. 与该二维码id有关联的第二层级id

输入整数 代表另一个相关联的二维码id

3. 这两个id二维码第二层级的物理意义

输入物理意义

4. 这两个二维码第二层级的可视化意义

输入“line”则在可视化界面中会对两个二维码进行连线

输入“midPoint”则在可视化界面中会对两个二维码的中点绘制 灰色物块表征中点

5. Level3 层级的设置本实验不涉及，不输入怎意义赋予流程在此结束

Exercise

- Connect the USB camera to your computer and visit the following website and you will see a web-based GUI for data collection.
  - https://me336.asyst.design/CustomSetting.html

Tool:
Assessing feasibility for idea selection

Use this flowchart to quickly assess how feasible / viable your AI idea is

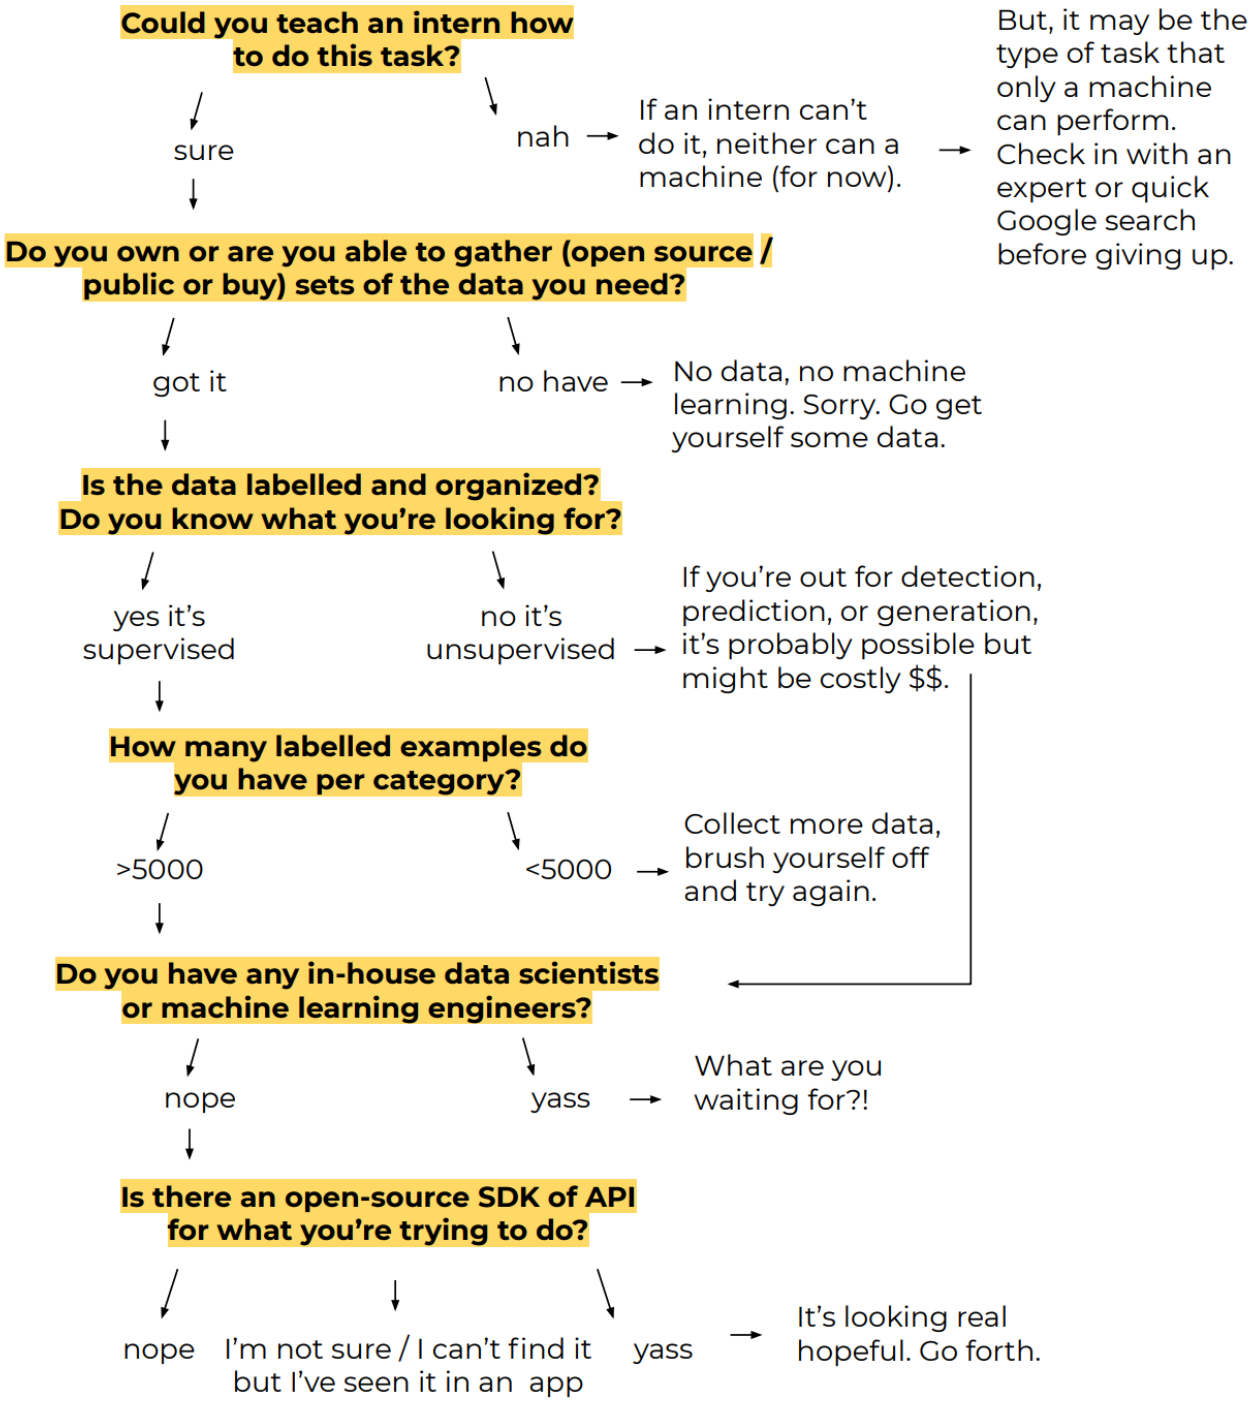

Activity:
Framing your task for concept development

Google HCML team speak from experience when they say:
"Find experts who can be the best possible teachers for your machine learner—people with domain expertise relevant to whatever predictions you're trying to make. We recommend that you actually hire a handful of them, or as a fallback, transform someone on your team into the role. We call these folks "content specialists" on our team."

The strength of machine learning is that we don't have to program the rules explicitly. At this stage of the process, it is helpful to think about them and try construct a logic based on how we humans perform the task.

1. Start with the classic exercise: describe the way a human expert would perform the task or answer the question.

   If you were to ask 10 people, would they agree on the method (for the most part)? If some do it better or differently - what can we learn from their approach?

   Especially if what you're predicting is (highly) subjective, spend extra time on this step.

2. Imagine you're onboarding a new person for this job. What do they need to understand? What assumptions would you want them to make? How would you respond so they improve over time?

3. What's the nature of the task? Can you box it as an clustering, classification, or regression problem? Refer back to the crash course in the beginning of this toolkit to find the vocabulary. Knowing this will help you understand the task as well as communicate with your tech team.

In the example of Spotify's Discover Weekly, **the human expert** would be a music lover on the hunt for new music.

Do you have data of past well-executed and completed tasks? This could be used as an initial training data set.

Tip:
Draw a diagram of the current workflow including IFTT statements and data required to make decisions.

Activity:
Plotting your model for concept development

By plotting a simple flowchart, we can begin forming a rough idea of the inputs, outputs, and logic required for our model to create value. We're also surfacing our assumptions and unknowns in the process.

1. **Objective** - What is the question we're trying to answer and asking the machine?
   Output - How is the machine's answer presented and interpreted?

2. **Features** - What data points do you need or are important factors in answering the question?
   Input - Which data sets does that data reside in? What data will the model be trained on? What data does the user input?
   + Draw connections between the assumed features and data sets they reside in.

3. **User experience** - How does the outcome get presented to and help the user?
   Business value - How does the solution return value to the organization?

AI answers (mostly) in probabilities with a confidence level. Formulate your output as a **probability**.

Do you know which features go into the answer? Think about the variables and patterns humans look at when performing this task or answering this question.

Do you have this **data to input**? If not, how do you acquire it?

Activity:
Plotting your model for concept development

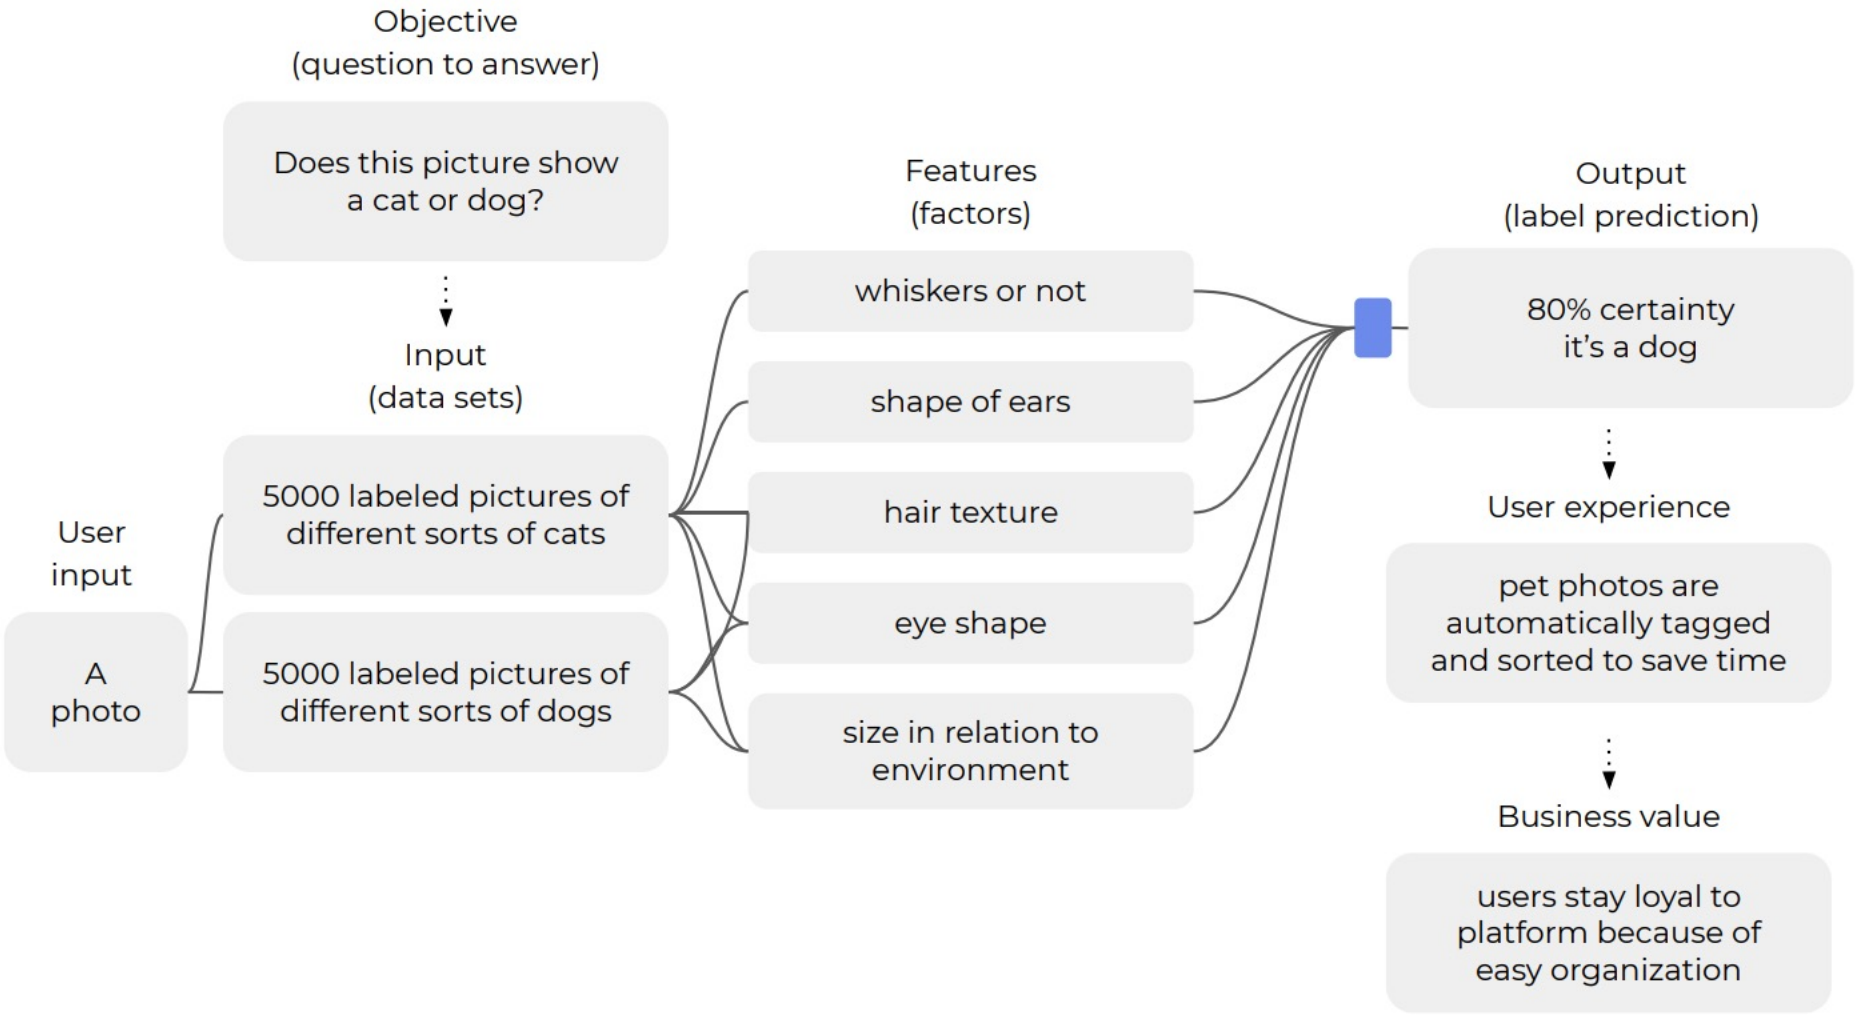

Worksheet:
Plotting your model

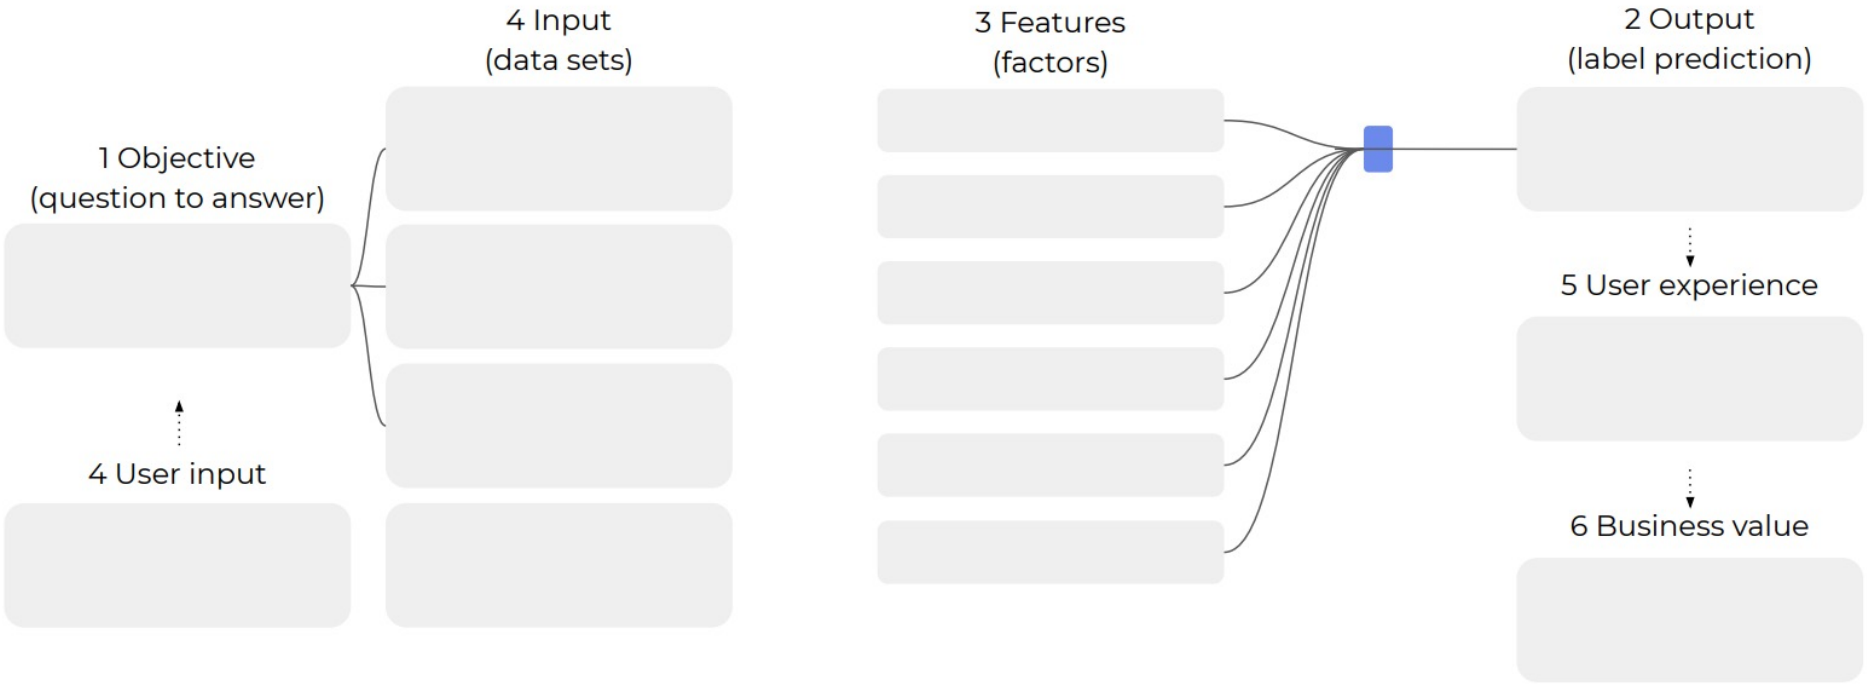

1 Objective
What is the question we're trying to answer and asking the machine?

2 Output
How is the machine's answer presented and interpreted?

Formulate your output as a probability.

3 Features
What data points do you need or are important factors in answering the question?

Do you know which features go into the answer? Think about the variables and patterns humans look at when performing this task or answering this question.

4 Input
Which data sets does that data reside in? What data will the model be trained on? What data does the user input?

Do you have this data to input? If not, how do you acquire it?

+ Connect
Draw connections between the assumed features and data sets they reside in.

5 User experience
How does the outcome get presented to and help the user?

6 Business value
How does the solution return value to the organization?

Prototyping + testing

You're with a handful of ideas and it's time to get more in-depth with your user research. Through prototyping and testing, you (in)validate your AI ideas and their design and implementation specs.

Do users want and need your solution? Are they open to adoption? Are they willing to share data and invest themselves into training the model (if necessary)? How can we test rather than just ask? How can we prototype the experience of adaptive intelligent systems?

In this chapter you will find:

User research & feedback
to know what to inquire about in addition to the usual

Prototyping & testing
to explore how to prototype and test AI applications

Activity:
User research & feedback for assessing desirability

1

Assuming you did initial user research to inform your concepts so far, now it's time to go out and (in)validate your value proposition in more detail. First assess your need as you do for any problem, asking:

- What problem does it solve or opportunity does it tap into?
- Who benefits and in what scenario?
- How pressing is the problem? For how many?
- What do they gain from the new solution? How and how much better is it than the current solution? What other advantages do they see?

Activity:
User research & feedback for assessing desirability

Iterate on your value proposition statement based on your learnings and get ready to prototype for deeper insights.

2

Once you've validated that this is indeed a problem worth solving, gather insights about your users' perspective on the AI aspects of your concept(s).

Mental models
What are their notions about having an intelligent, adaptive system work for them? Are they willing to adopt it? How important is transparency? Depending on how visible your AI elements are, this might be more or less important.

Defining success and failure
How accurate must the model be to offer user value? How high are the costs of mistakes? What would best vs worst behavior look like?

Machine teaching
What does the user need to invest to get value out of the system? Are they willing to share the data your model needs? Are they willing to provide the necessary feedback and teach the model?

Ethical & experiential concerns
What concerns do they have? Do major ethical concerns arise? Unintended consequences, edge cases, and extreme users?

Activity:
Prototyping & testing
for assessing desirability

1

Prototype

To test desirability, opt to simulate the experience without building the model and observing the responses.

Testing the concept offering can be done with product / service posters or app marketplace.

Common prototyping techniques for AI are:
Role playing
Wizard of Oz
Personalized wireframes.

Where possible, gather and use real-life personal data in your prototypes rather than placeholder content.

Provotypes (prototypes that provoke) can also be a great way to build an understanding of your users' needs.

"Fake it till you make it. If forced to choose, it's leaps-and-bounds more useful to prototype your UX with a user's real content than it is to test with real ML models - as it affords you genuine insights into the way people will derive value and utility from your (theoretical) product."

by Google Clips' team on UX of AI

Activity:
Prototyping & testing
for assessing desirability

2

Testing

Do user testing as usual and observe users' behavior. Ask them to think out loud as they're interacting with your artefact.

Keep in mind that while testing is important to understand your user, working with adaptive systems requires the designer to sacrifice a certain level of control over the final user experience exactly because it will adapt to each user and over time.

3

Analysis & selection

Analyse and synthesize your findings. Based on all your findings, decide which idea(s) (if any) to move forward with.

It can help to revisit some of the activities in idea selection phase and reconsider feasibility, viability, desirability, and responsibility.

Day 02
AI Meets Design II

Thank you~

Wan Fang
Southern University of Science and Technology