


Supervised Learning Reliance on Data



Overfitting




Overfitting

Characteristics

- The model performs very well **during training**



Dog ✓
Cat ✓ **99% Accuracy**



Overfitting

Characteristics

- The model performs very well **during training**



Dog ✓
Cat ✓ **99% Accuracy**

- The model performs very poorly on **different data**



Cat ✗
Dog ✗ **30% Accuracy**






Overfitting

Possible Reasons


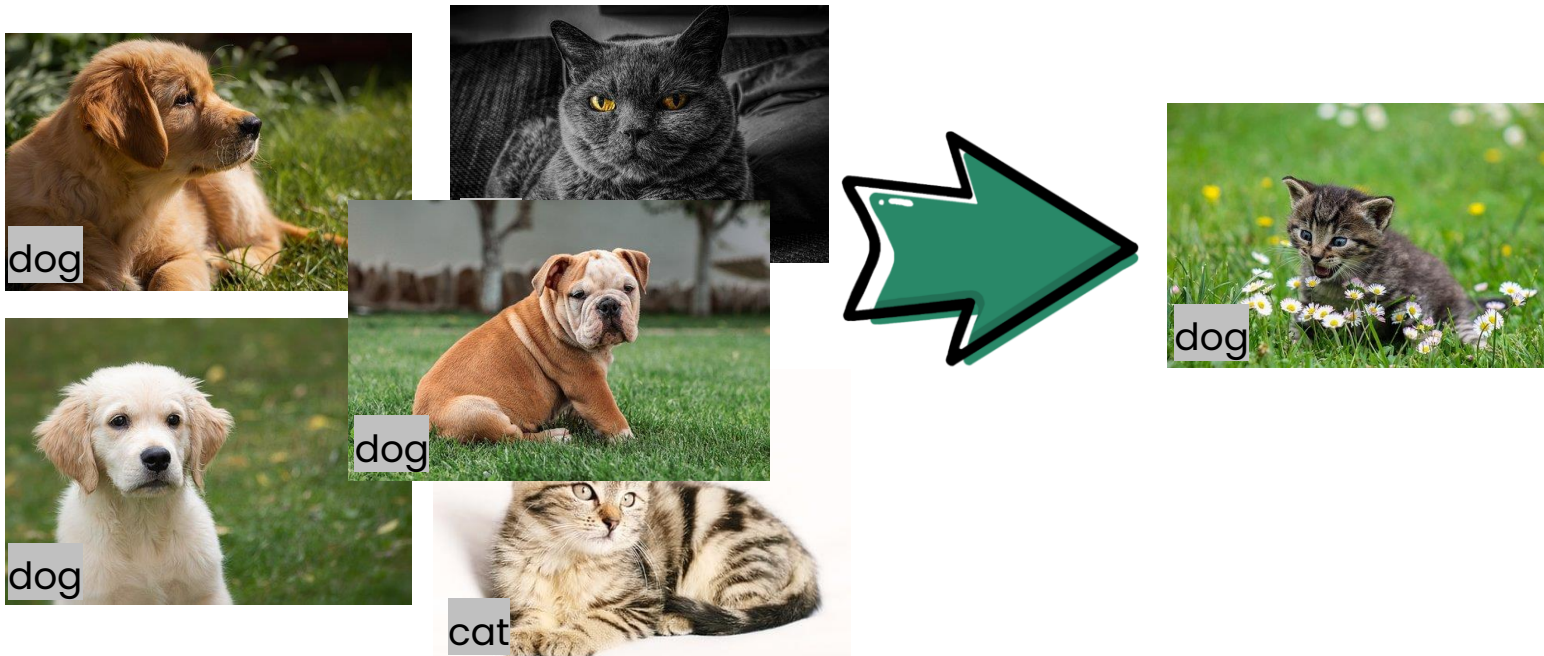
- The model did **not** learn **general enough** features
- The model focuses on the **wrong features**



Overfitting

Example 1


Most images of **dogs** in the training set contain **grass**, therefore every image with **grass** will be classified as a **dog**.



Overfitting

Example 2


All images are **closeups** of cats, therefore a image of a cat in a field will not be reconized.





Overfitting


What examples can **you** think of?



Overfitting

What examples can **you** think of?

- Focus on hand **holding** object
- Focus on person or object in the **background**
- Focus on body **posture** instead of facial expression
- Focus on **hair** and **jewelry** instead of face
- Focus on hand **position** instead of gestures
- All training images where shot in **daylight**
- ...




Overfitting

Possible solutions

- **More** data
- Better **generalized** data
- **Processing** data to reduce insignificant differences
- Increase **variety** by:



scaling



blurring



...

desaturating

flipping



rotating



Underfitting




Underfitting

Characteristics

- the model just does **not** provide **usable results**



Cat ✗
Cat ✓ **50% Accuracy**



Underfitting

Possible Reasons

- **Too few** samples
- Training **data** is **too general**
- (Training parameters are way off)



cat




dog

CAT

cat




dog



cat




dog



Underfitting

Possible Solutions

- Use **more/better** samples for trainings-data
- Adjust training parameters (more iterations, ...)



Biases and Fairness



Biases and Fairness

Imagine you create a **chatbot** that should act like a normal human. Therefore you use **live-chats** from **streaming platforms** like Twitch or YouTube for training.



Biases and Fairness

Imagine you create a **chatbot** that should act like a normal human. Therefore you use **live-chats** from **streaming platforms** like Twitch or YouTube for training.

- How well do you think this bot would perform?
- Do you see any **problems** that could occur?




Biases and Fairness

- Network **replicates** what it has **seen/learned**
 - racism, sexism, vulgarity, ...
 - bias towards streamers with a loud audience
 - bias towards names and norms that are relevant at the time of training
 - ...



Biases and Fairness

- Network **replicates** what it has **seen/learned**
 - racism, sexism, vulgarity, ...
 - bias towards streamers with a loud audience
 - bias towards names and norms that are relevant at the time of training
 - ...
- **Hard** to create **unbiased algorithms**
- Ethical question on **where to use** such algorithms
 - Where should machines be allowed to make decisions?
 - How can we guarantee an fair and unbiased result?




Training Time and Transfer Learning




Training Time

- Training a **good model** from scratch can take **weeks**
 - If you include finding the right parameters it can even take months or years...



Transfer Learning

- Uses an **already trained model** as a starting point
- Requires much **less trainings-data**
 - usable results often with a few dozens of images
- Requires much **less time**
 - usually seconds to minutes
- Many **pre-trained models** are **freely available**
 - especially for image recognition





Up next

Training a **real world** image classification **model** and use it to **control** a game of **Snake!**

