

SSP 100 - 10

How Computers See


10 / 01 / 2026

Instructor: Michael Eckmann


Today's Topics

- Questions? / Comments?
- Continuation of an incomplete and concise history of computer vision from early days to present


History of computer vision



(a) Traditional vision pipeline



(b) Classic machine learning pipeline



(c) Deep learning pipeline


History of computer vision

- Some takeaways from discussion so far
 - geometry
 - linear algebra
 - Physics of light and interaction of light and surfaces
 - Human vision influences (for example stereo because we have 2 eyes)


History of computer vision

- Skip ahead now to utilizing hand crafted features as inputs to machine learning


History of computer vision



(a) Traditional vision pipeline



(b) Classic machine learning pipeline



(c) Deep learning pipeline

History of computer vision

- The computer vision researcher would do feature engineering (deciding what to compute and how) possibly including existing feature descriptors as well as new ones
- Then precompute these features of the images and feed the features into a machine learning algorithm to figure out how to process those features for something like say classification

History of computer vision

- Some examples:
 - Dalal & Triggs HOG for Human Detection, 2005
 - Used Histogram of Oriented Gradients (hand crafted feature) fed into an SVM (Support Vector Machine)
 - Viola & Jones Robust Real-time Face Detection, 2004
 - Used some simple features fed into AdaBoost

Definition of a machine learning algorithm

- A machine learning algorithm is an algorithm that is able to learn from data.
- Mitchell (1997) “A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E .”

Performance Measure

- Performance = a quantitative measure of performance
- One example: accuracy rate (or conversely error rate)
- Note: accuracy rate is the number of accurately predicted test set examples / number of examples tested
- Whereas the error rate is the number of incorrectly predicted test set examples / number of examples tested
- Top 5 accuracy --- that is in datasets with large number of classes (e.g. 1000 classes in ImageNet) model outputs what it thinks are the 5 most likely correct classes for a test image and if the correct class is among the 5, it is considered correct
- Conversely there is a top 5 error rate


Experience

- Experience is the dataset that the learning algorithm is provided.
- For supervised learning this includes the correct label for each.


History of computer vision

- Currently (since 2012) researchers focus on deep learning where the input is the raw pixel values of an image and the machine learning algorithm learns what features are important for the problem at hand.


History of computer vision



(a) Traditional vision pipeline



(b) Classic machine learning pipeline



(c) Deep learning pipeline

History of computer vision

Image classification

- ImageNet dataset
- Had yearly competitions
- In 2010 and 2011 classical machine learning techniques won
- In 2012 AlexNet won, the first time a Deep CNN won
- Sparked much innovation into deep learning for image classification with major advances in the following years
- Let me show a table of winners and accuracy rates

History of computer vision

AlexNet

- Let's look at the features it learned
- Let's also see how many citations that paper currently has

History of computer vision

Let's discuss the TED talk by Fei-Fei Li that you watched.

Did anyone note what year the talk was given and why we might care?