Computer Vision and Al

Mathias Kölsch, Ph.D.

Harnessing AI Lecture Series



Wrist fracture?



Human alone: 80.8% sensitivity With CV: 91.5% sensitivity

Deep neural network improves fracture detection by clinicians, Lindsey et al. 2018

Computer Vision

image space

feature space

results



Computer Vision

image space

feature space

results

feature extraction

bright points, red disks edges, corners straight lines, circles shape texture and appearance change, motion, ...

Features: Decades of Craftsmanship

- SIFT Descriptor: State of the art in 2004 and for almost a decade after
- Expert-designed local appearance descriptor

feature extraction

David Lowe "Distinctive Image Features from Scale-Invariant Keypoints" IJCV 2004





Deep Neural Networks: Magic

- 100s of millions of parameters
- Extremely under-constrained optimization problem
- Non-deterministic training outcome:
 - Many re-tries, pick the best
 - Ensemble networks





Convolutional Deep Neural Networks



- Spatial correlations
- Feature sharing
- Compositional
- Invariants

Lee et al.: Convolutional Deep Belief Networks for Scalable Unsupervised Learning of Hierarchical Representations, 2009

Learned Features in a Neural Network



Erhan et al, 2009: Visualizing Higher-Layer Features of a Deep Network

Learned Features, Amplified



DeepDreamGenerator.com

CV Catalyzed Deep Neural Networks

Why?

- Local redundancies and Convolution permit training efficiencies Convolutional Neural Network – CNN
- Availability of increasingly difficult problems from OCR on digits to faces to 100 objects to 1000 to detailed bird species classification
- Training data availability: industry + Scademia
- GPU capabilities map very well to CNNs

"The Deep Neural Network will see you now"



Figure 3: Collage of some medical imaging applications in which deep learning has achieved state-of-the-art results. From top-left to bottom-right: mammographic mass classification (Kooi et al., 2016), segmentation of lesions in the brain (top ranking in BRATS, ISLES and MRBrains challenges, image from Ghafoorian et al. (2016b), leak detection in airway tree segmentation (Charbonnier et al., 2017), di-abetic retinopathy classification (Kaggle Diabetic Retinopathy challenge 2015, image from van Grinsven et al. (2016), prostate segmentation (top rank in PROMISE12 challenge), nodule classification (top ranking in LUNA16 challenge), breast cancer metastases detection in lymph nodes (top ranking and human expert performance in CAME-LYON16), human expert performance in skin lesion classification (Esteva et al., 2017), and state-of-the-art bone suppression in x-rays, image from Yang et al. (2016c).

Litjens et al, A Survey on Deep Learning in Medical Image Analysis, 2017

Self-Driving Cars



Tesla:

- 8 cameras, 12 ultrasonic sensors, radar
- Custom chip with
 2 hardware NNs,
 12 core CPU,
 GPU

Computer Vision and Al

- CV outperforms humans for more and more tasks Medical image analysis No self-driving cars without CV
- CV catalyzed deep neural networks
- Explainability, predictability, robustness

Artificial Intelligence?

- Causal reasoning, not associative or correlation alone
- "A key aspect that is often overlooked is that expert knowledge about the task to be solved can provide advantages that go beyond adding more layers to a CNN." Litjens et al, A Survey on Deep Learning in Medical Image Analysis, 2017

Thank you!

Mathias Kölsch kolsch@nps.edu

Backup

NPS Theses in Computer Vision

- AK-47 detection, other object detection
- 3D reconstruction from stereo
- Object recognition in LiDAR data
- In-flight obstacle avoidance (air to air)
- UAV onboard object detection
- Carrier landing flight path visualization from PLAT CAM video
- Star tracker
- ··· many more!

Cutting Edge: Detailed Visual Perception

- Computer vision as foundation for situational perception
 - Improved capabilities for scene understanding, context
 - Augmented reality's biggest challenge is understanding the environment

How many surveillance cameras in NYC?

180 1,800 18,000

Current Research

- Action recognition (video)
- Scene recognition (other than Eiffel tower)
- 3D scene reconstruction
- Detection of many object classes
- Fine-grained object distinction, attributes
- Higher-level reasoning

Many research programs sponsored by NSF, DARPA, DHS, ONR, IARPA

• VIRAT, VMR, PerSeas, Mind's Eye, Big Data, …