

Brief Overview of Artificial Intelligence

John E. Laird
University of Michigan
April 16, 2019

Artificial
Intelligence

Machine
Learning

Artificial
Neural
Networks

Deep
Learning

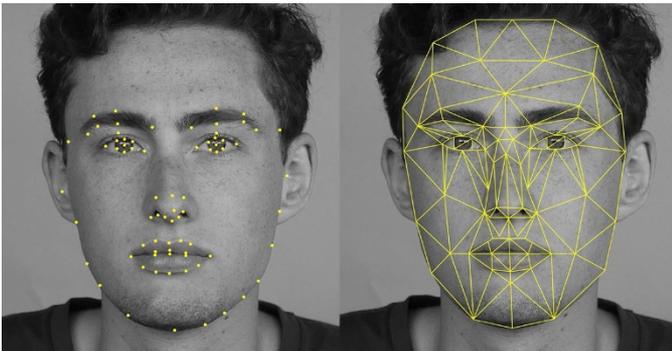
Brief History of AI: 1950-2011

- Emphasis on internal reasoning through **programming**
 - Problem solving, search, planning, scheduling, knowledge representation, natural language understanding, ...
 - Symbolic and probabilistic processing
- Expert systems, Deep Blue chess player, Jeopardy (Watson), planning and scheduling, ...



- *Tasks that take humans seconds to minutes (System 2).*

- Emphasis on perception and classification through **learning**
 - Practical use of large simulated neural networks
 - Vision processing, speech understanding, language translation, object classification, ...
- Face recognition, machine translation, health care, self-driving, Go, ...



- Learn from *extensive* training on *massive* data sets
- *Tasks that take humans < 1 second (System 1)*

Game Playing

- **Go, Chess:** Alpha Zero achieved superhuman ability in Chess, Shogi, and Go; learning by self-play.
- **Poker:** CMU program defeated 4 top humans in a tournament of 120,000 games of two-player, heads up, no-limit Texas Hold'em.
- **Strategy Game Playing:** AlphaStar beat human experts in the strategy game StarCraft.

Health Care

- **Skin Cancer Classification:** An AI system trained on a data set of 129,450 clinical images of 2,032 different was capable of classifying skin cancer at a level of competence comparable to the dermatologists.
- **Prostate Cancer Grading:** Google developed a deep learning system that can achieve an overall accuracy of 70% when grading prostate cancer in prostatectomy specimens. The average accuracy of achieved by US board-certified general pathologists in study was 61%.

Chinese - English Translation

- A Microsoft machine translation system achieved human-level quality and accuracy when translating news stories from Chinese to English.

Self-Driving Cars

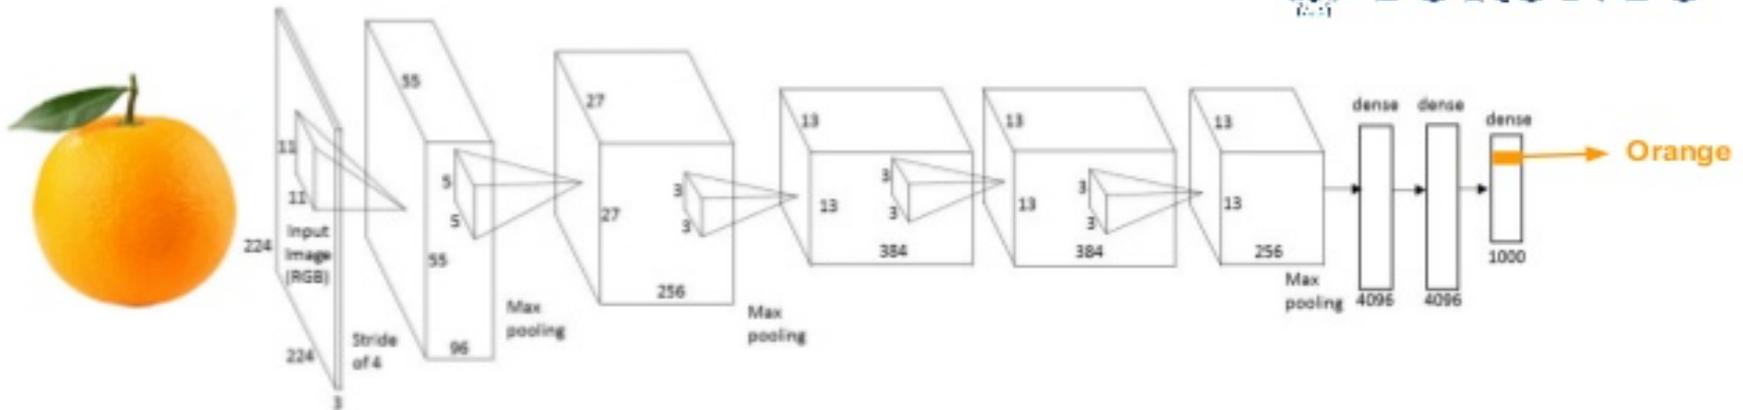
- Audi announces the A8 will be first Level 3 production vehicle.

What Changed in 2011?

1. Advances in machine learning & neural network algorithms.
2. Availability of large bodies of data for training.
3. Development of very fast, low-cost parallel computers.
4. Advent of free tools (Tensorflow, PyTorch) that made it easy for almost anyone with moderate skill to build neural network models.

2012: ImageNet Large Scale Visual Recognition Challenge (ILSVRC)

- 1,200,000 labeled images that contain 1000 object classes.
- Previous best model achieved 26.2% error.
- AlexNet (7 layer neural network) achieved 15.3% error.



A Krizhevsky, I Sutskever, GE Hinton "[Imagenet classification with deep convolutional neural networks](#)" Part of: [Advances in Neural Information Processing Systems 25 \(NIPS 2012\)](#)

Slide credit: Junting Pan, "[Visual Saliency Prediction using Deep Learning Techniques](#)" (ETSETB-UPC 2015)

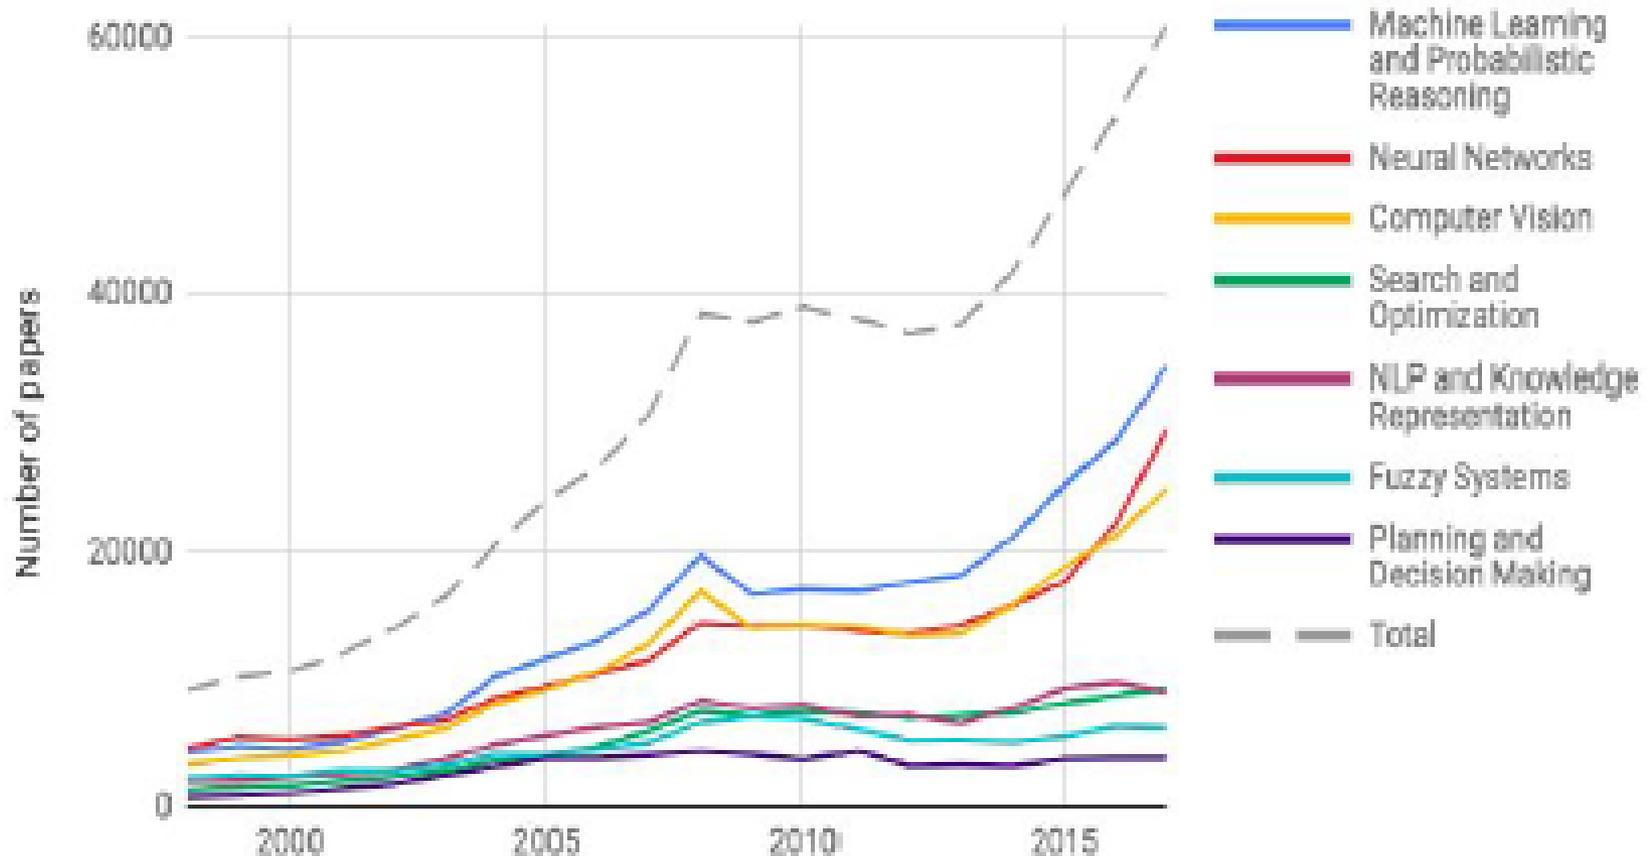
- Error rate is now <5%.

Explosion in AI Publishing

AI papers on Scopus have increased 8x since 1996. CS papers increased 6x during the same timeframe.

Number of AI papers on Scopus by subcategory (1998-2017)

Source: Elsevier



Worldwide map (Publication time range: 1998 - 2017)



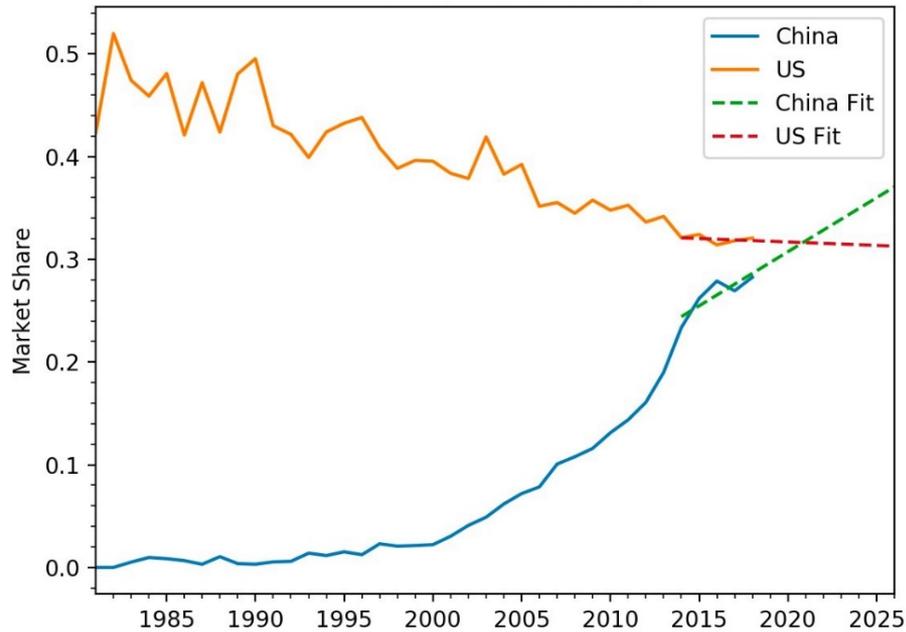
83% of 2017 AI papers originate outside the U.S.:

Europe (28%), China (25%) and the U.S. (17%).

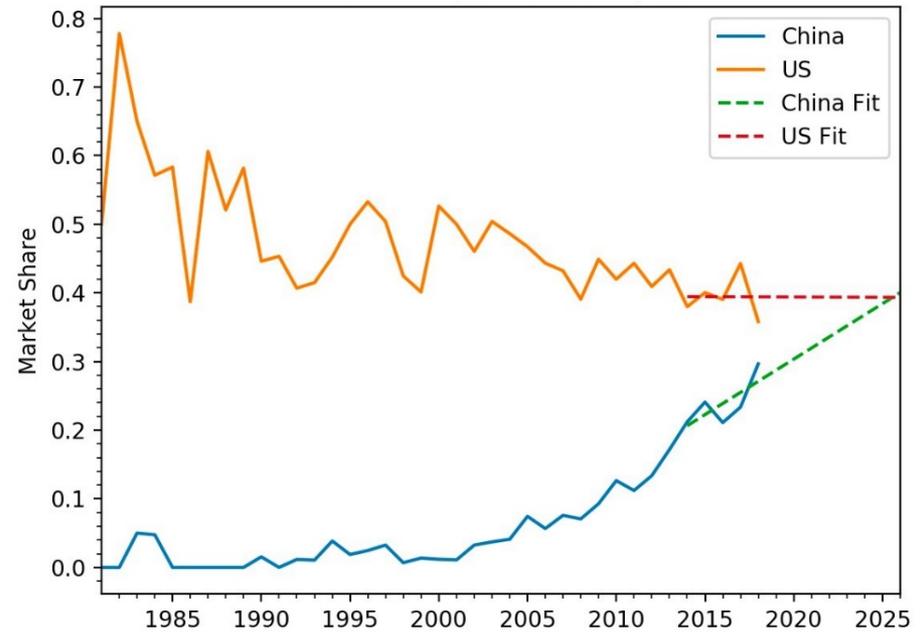
U.S. authors are cited 83% more than the global average.

Rise of China in AI Citations

Share of top 10% of AI Papers

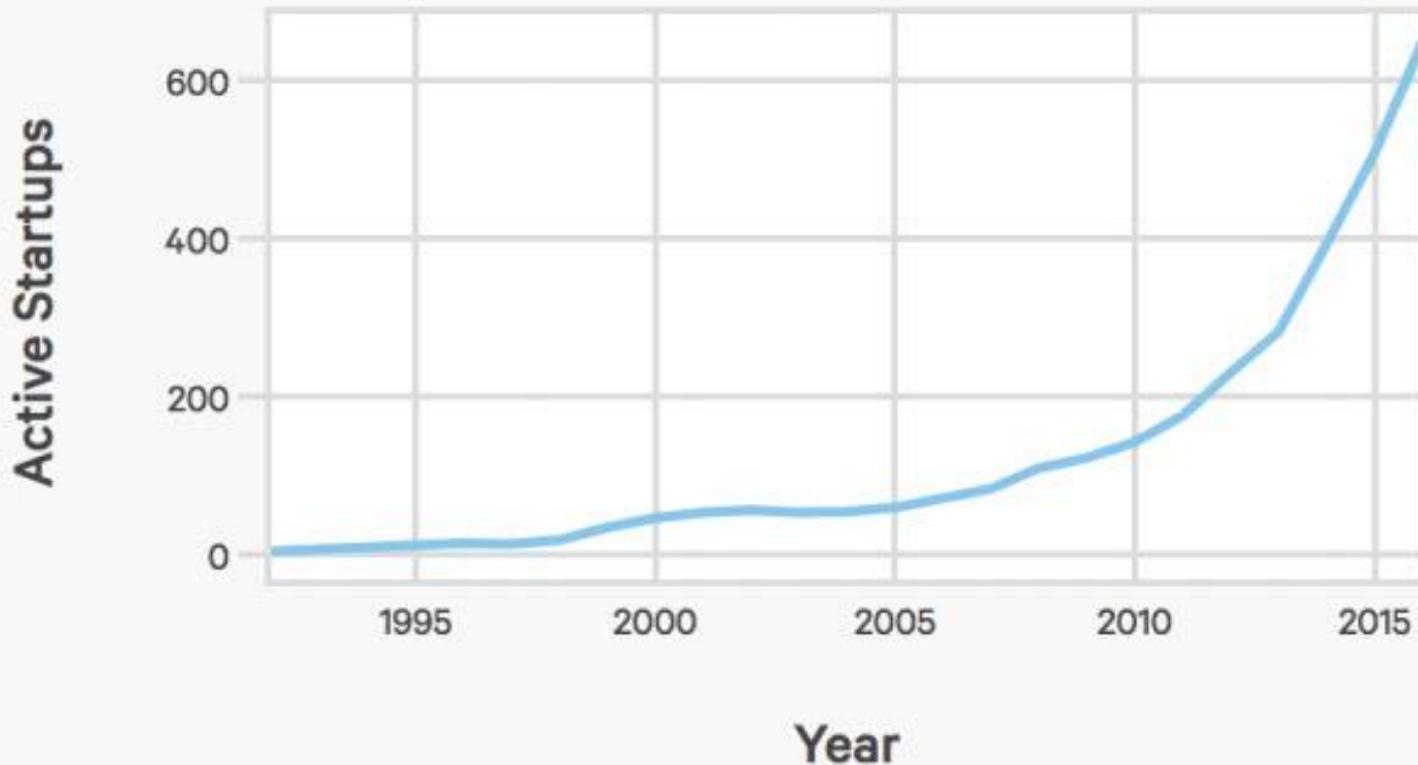


Share of top 1% of AI Papers



Explosion in AI Startups

Startups Developing AI Systems



- Large commercial investment in AI *research* groups
 - Facebook, Microsoft, Google, IBM, JP Morgan, ...
 - DeepMind (Alphabet) has >700 employees in AI research
 - Allen AI Institute (> 100 employees) \$125M over 3 years
- DARPA: \$2B over four years
- Autonomous driving - \$\$\$\$\$ in > 50 companies
 - Waymo, Ford, GM, Toyota, ...
- Chinese government plans to spend \$150B on AI by 2030
 - Chinese don't have HIPAA: they have lots of health data...

Which of these are Real People?

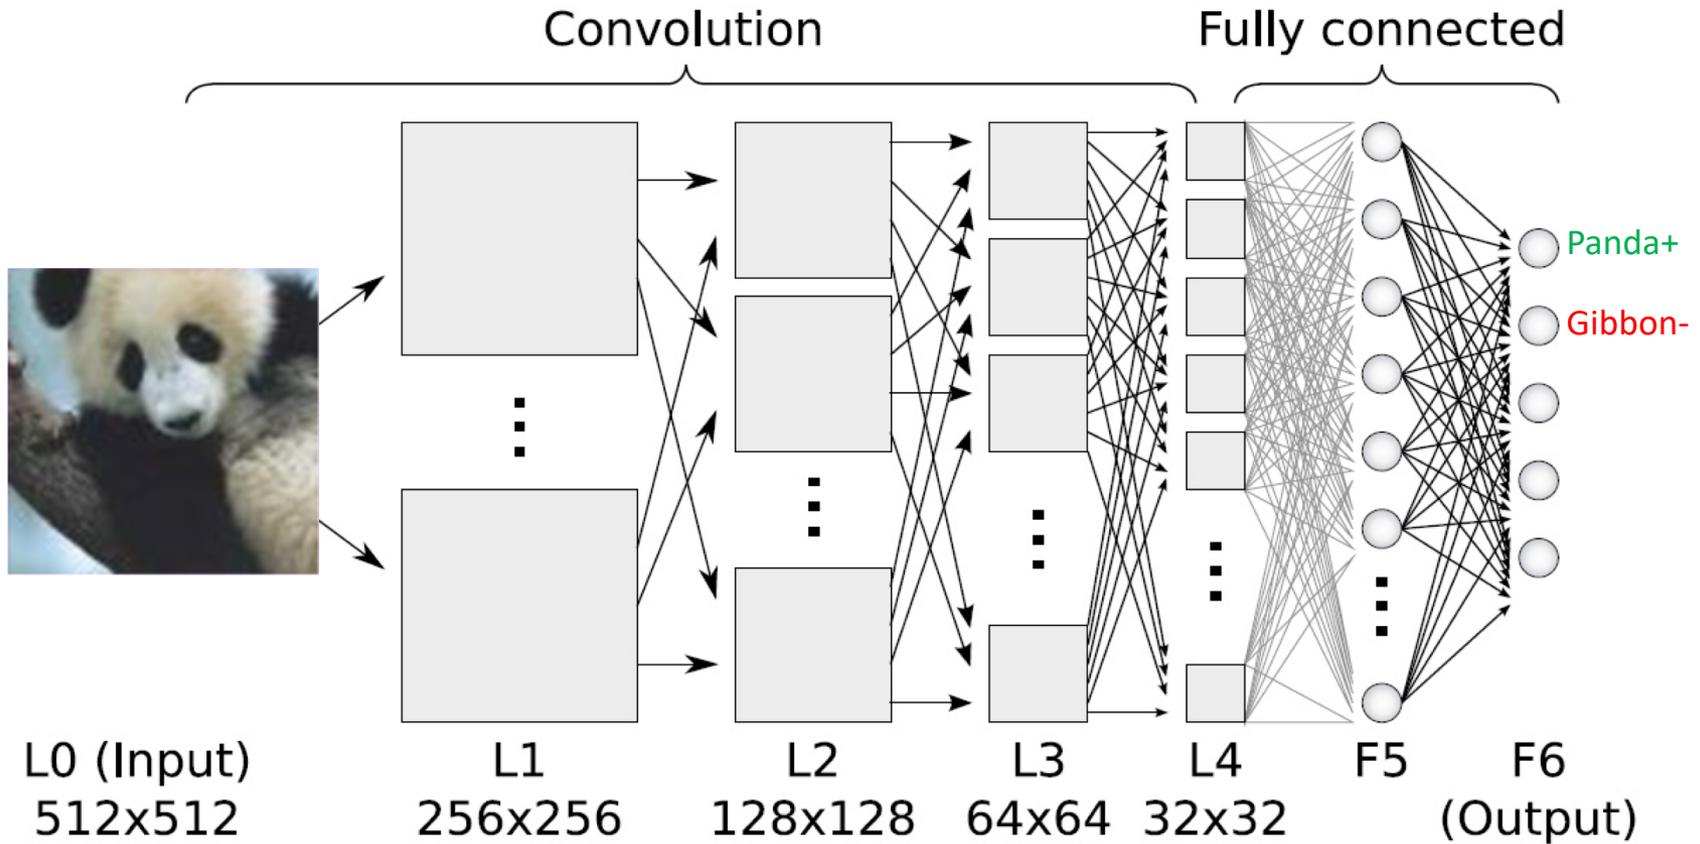


Tero Karras, Timo Aila, Samuli Laine, Jaakko Lehtinen , “Progressive Growing of GANs for Improved Quality, Stability, and Variation”, 2017

Panda vs Gibbon



Deep Learning: Panda Classifier



57.7% sure this is a panda



Overlay a carefully selected “adversarial pattern”



99.3% sure this is a gibbon



Stop sign vs. 45 MPH sign



- Pros

- Incredibly good at finding patterns in complex data
- Effective for “perception heavy” tasks
 - Lots of tools to make it easy to use
- Scales to large problems and datasets

- Cons

- Needs **LOTS** of *labeled training examples* and *computation*
- Difficult to understand how it works
 - Hard to know which features classifications are based on
 - How will it work in new situations?

- AI research, development, and applications are growing incredibly fast.
- Worldwide, not just in U.S.
- How long will it continue?

- Many unanswered questions about impact on society...
 - Privacy, jobs, security, ...