ROBOTICS AND AI

Sören Schwertfeger / 师泽仁
ShanghaiTech University

https://robotics.shanghaitech.edu.cn

Outline

- Robotics at ShanghaiTech
- Robotics & AI
- Deep Learning
- Introspection
- Standards for Robotics

ShanghaiTech University

- Founded 2013 by Chinese Academy of Sciences and Shanghai Municipality
- Goal: Become a leading University in China and the world => excellent funding
  - 2700 students now; 6000 planned
- Location: Shanghai Pudong; Zhangjiang Hi-Tech Park
- Schools & Institutes:
  - School of Physical Science and Technology (SPST)
  - School of Information Science And Technology (SIST)
  - School of Life Science And Technology (SLST)

  - School of Entrepreneurship And Management (SEM) (no degrees)
  - School of Art And Creativity (in start-up / no degrees)

  - Shanghai Institute For Advanced Immunochemical Studies
  - iHuman Institute

ShanghaiTech Automation and Robotics Center STAR Center

- 7 Professors:
  - Andre Rosendo (Mechatronics & Systems)
    Hao Chen (Security & Learning)
    Laurent Kneip (Computer Vision & SLAM)
    Yajun Ha (FPGA & Smart Cars)
    Boris Houska (Model Predictive Control & Optimization)
    Jie Lu (Distributed Control & Optimization)
    Sören Schwertfeger (Systems & AI & SLAM)
- About 30 graduate students
- Working on various topics in the areas of Artificial Intelligence, Perception and Modelling, Planning, Control as well as Mechatronics.
- Applications in Safety, Security and Rescue Robotics (SSRR), Intelligent Unmanned Micro Aerial Vehicles, Service and Industry Robotics and Smart Cars.
- 1700 m^2 lab and office space

STAR Lab

- 7 graduate students (2 PhD); 1 Technician; 1 Assistant; several undergraduates
- 200 m^2 lab with advanced equipment

Robotics and AI Research at STAR Lab

- Mobile Manipulation
- Autonomous Navigation
- Smart Car
- Robot Performance Evaluation/ Map Evaluation
- Shanghai Institute of Fog Computing Technology:
  - Fog computing for robotics

RoboCup China 2017
STAR Rescue Team

https://robotics.shanghaitech.edu.cn

Research: 3D Mapping/ 3D SLAM/ Map Representation

- Efficient 3D mapping using range (LiDAR) and (stereo) vision
- 3D Simultaneous Localization and Mapping (SLAM) ”Learning Maps”
- Map representations for planning, map merging, map evaluation, …

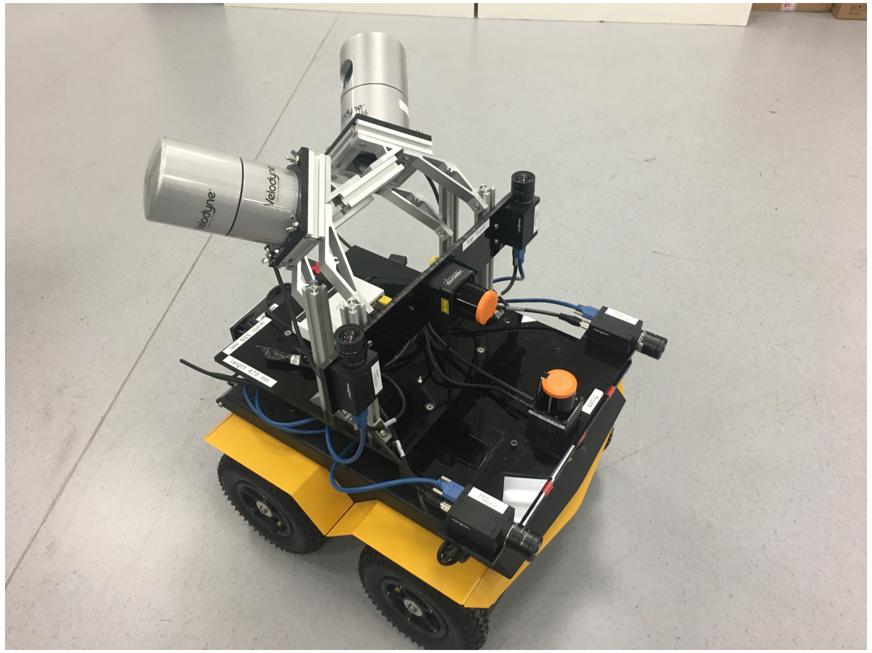

Artificial Intelligence – very big field

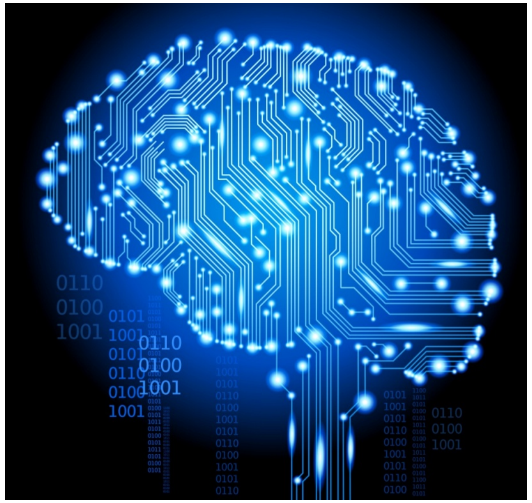

- Applications
  - Reasoning
  - Knowledge
  - Planning
  - Learning
  - Natural Language Processing
  - Perception/ Computer Vision
  - Robotics

  - Data Mining
  - Online-services (e.g. recommendation systems)

- Techniques
  - Statistical Methods
  - Symbolic AI
  - Computational Intelligence (Neural Networks)
  - Optimization
  - Search

- Strong AI vs. Weak AI
  - Strong AI: general AI as skillful and flexible as humans (and beyond)
  - Weak AI: Tuned for a specific task

Foundations of AI

- Philosophy
  - Logic, reasoning, mind as a physical system, foundations of learning, language and rationality.
- Mathematics
  - Formal representation and proof algorithms, computation, (un)decidability, (in)tractability, probability.
- Psychology
  - adaptation, phenomena of perception and motor control.
- Neuroscience
  - physical substrate for mental activities.
- Economics
  - formal theory of rational decisions, game theory.
- Linguistics
  - knowledge representation, grammar.
- Control theory
  - homeostatic systems, stability, optimal agent design.

AI and Embodyment

- Rolf Pfeifer, University of Zürich:
- ”Intelligent behavior emerges from the interplay between brain, body and world.” =>
  - Brain, body and world equally important factors for intelligent behaviors
  - NOT just outcome of internal control structure (brain, computer)
- => Only through robots can we aim to achieve Strong AI

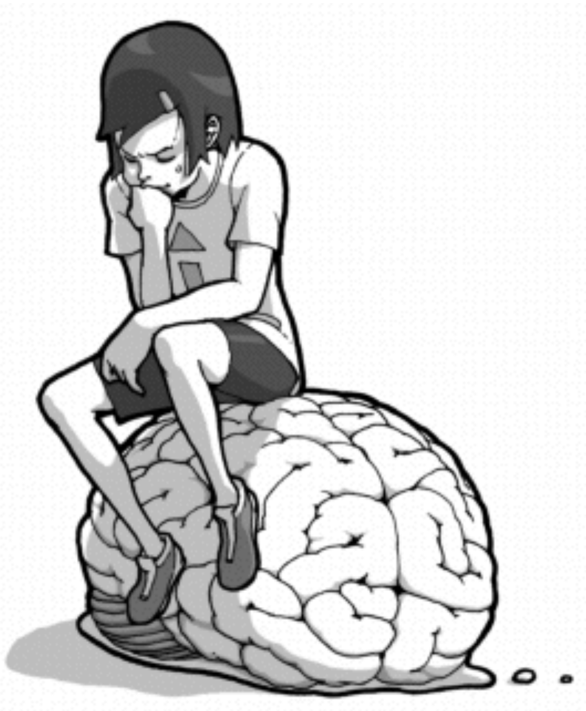

Image source: Shun Iwasawa, Rolf Pfeifer

Robotics

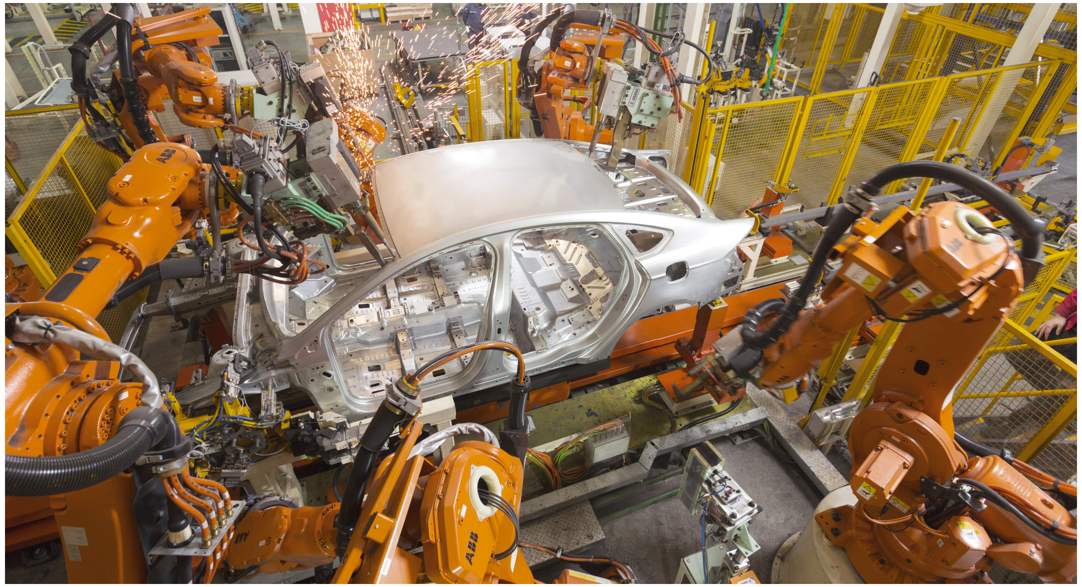

- Robotics very interdisciplinary
  - Mechatronics; Electronics; Communication; Computer Science; Biology; Psychology; Physics
- Industry Robots: Manipulators
  - 2015: 254,000 industrial robots sold
  - Over 1.6 million industrial robots installed – rising by about 9% per year
  - China biggest robot market regarding annual sales – also fasted growing market worldwide
- AI plays important role in advanced robotics:
  - Mobile robots; lots of perception; lots of uncertainty
  - No pre-planned motion

Image Sources: RobotWorx
Vanillase Wikimedia

General Control Scheme for Mobile Robot Systems

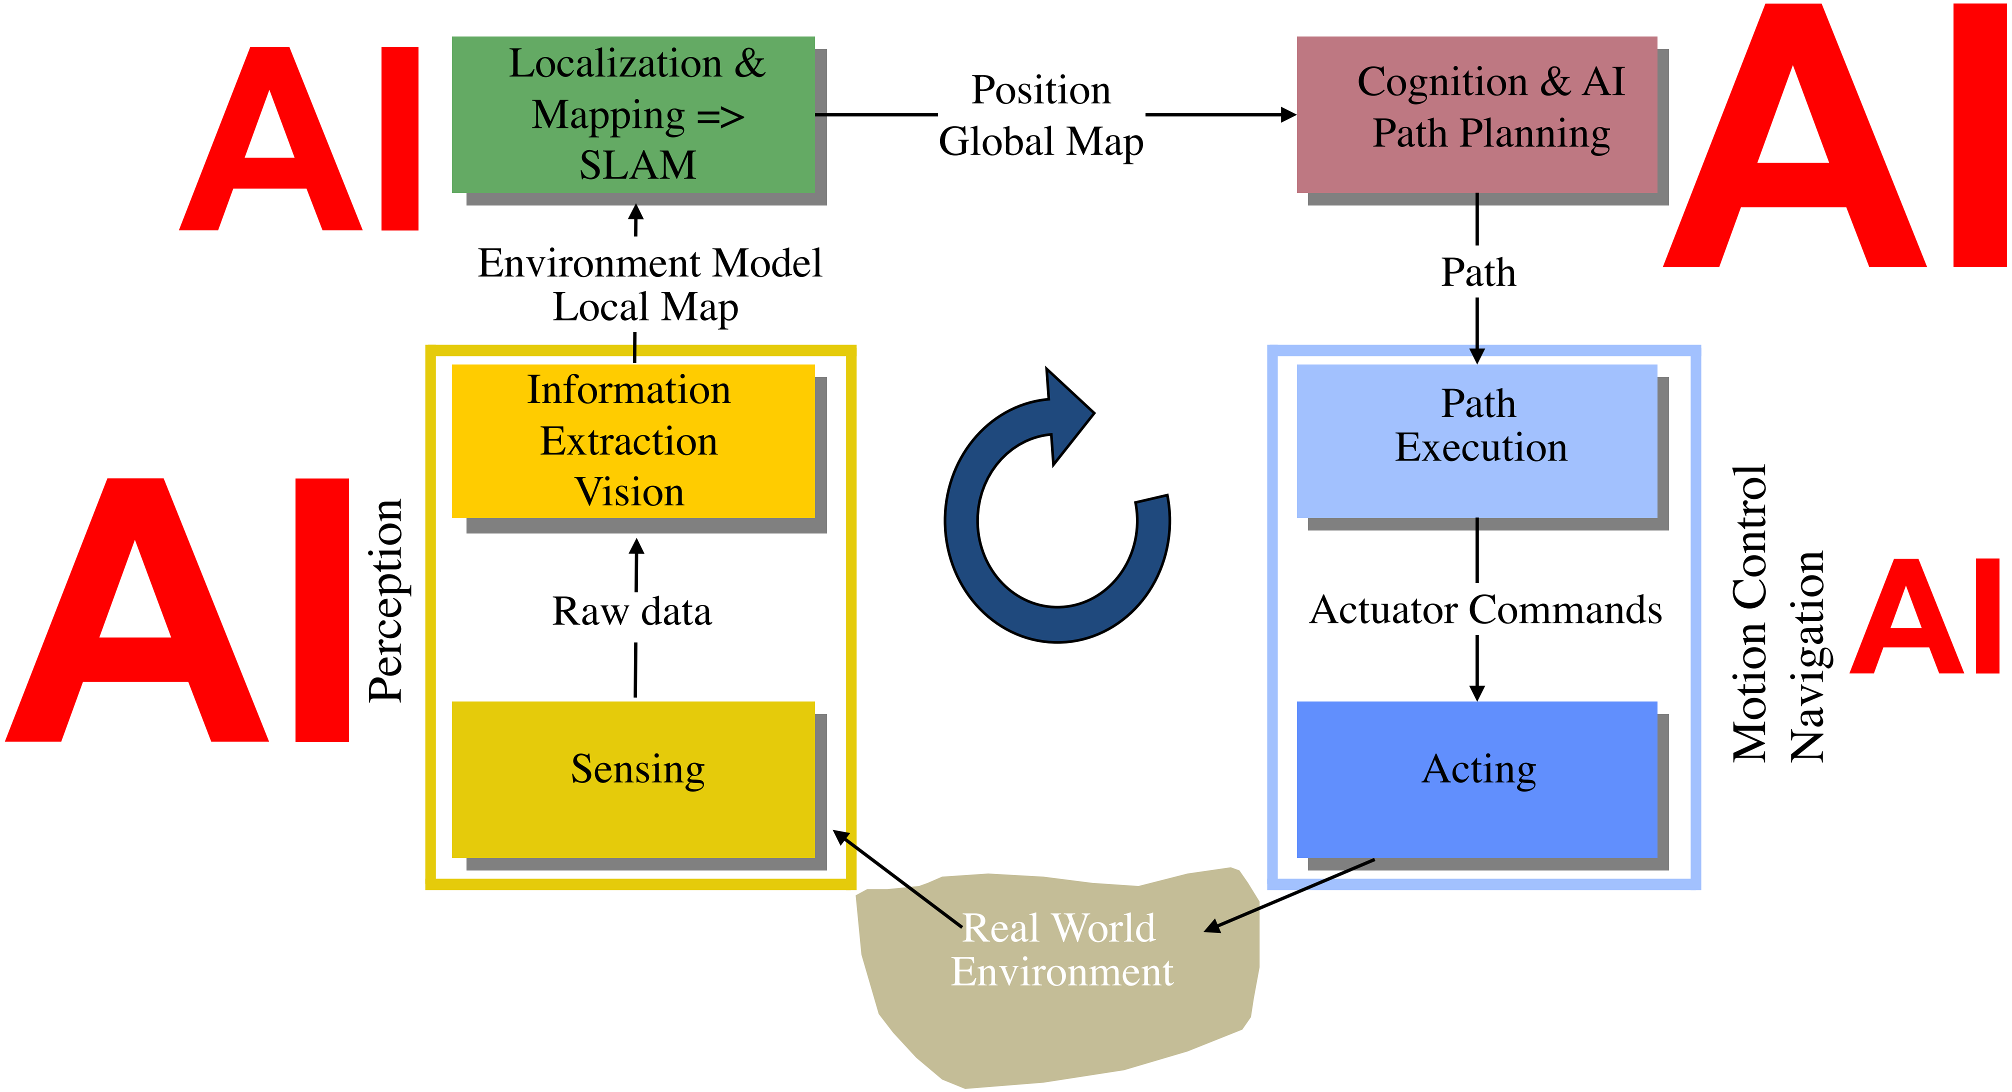

Material from Roland Siegwart and Davide Scaramuzza, ETH Zurich

Deep Learning

- Very hot trend in Machine Learning / AI
- Traditional approaches: Hand tuned features, e.g.:
  - Feature extraction (Computer Vision);
    Hidden Markov Models & Statistics (Natural Language Processing)
- Deep Learning approach:
  - Learning of parameters for Artificial Neural Network (ANN) based on training data
- Recent success based on:
  - Huge amounts of (training) data
  - Lots of computing power
  - Better DL architectures

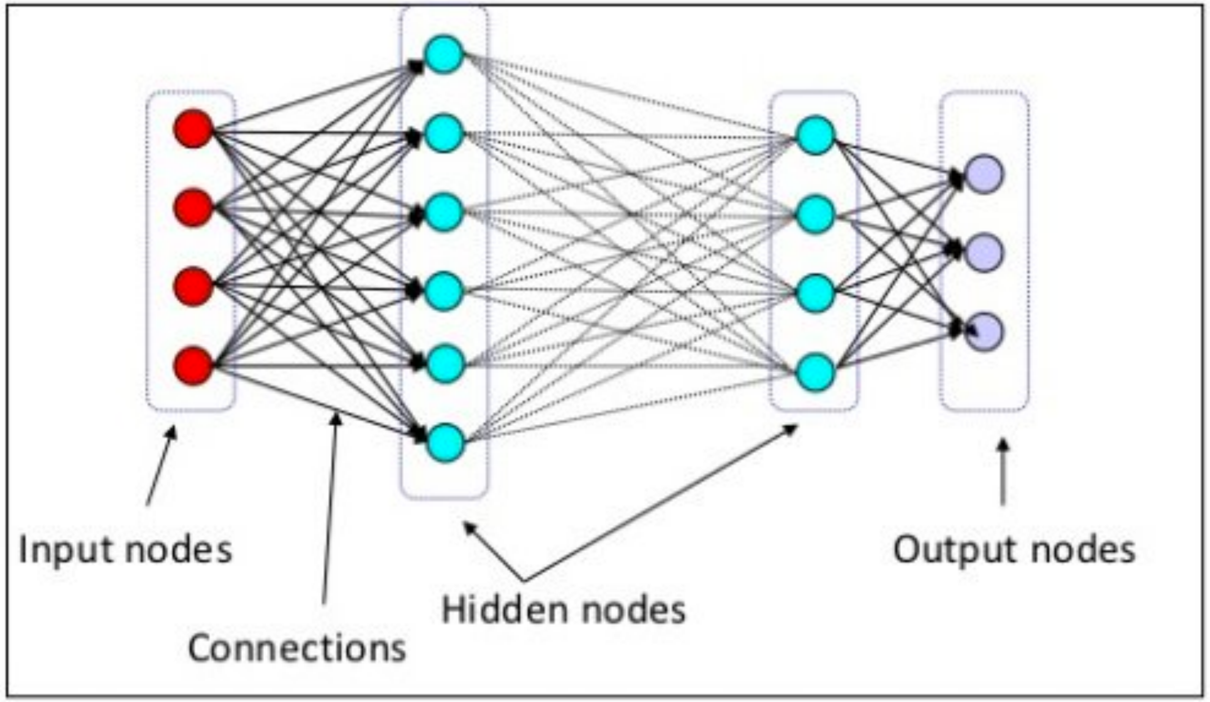

Image Source: Matthew Mayo

Traditional Pattern Recognition: Fixed/Handcrafted Feature Extractor

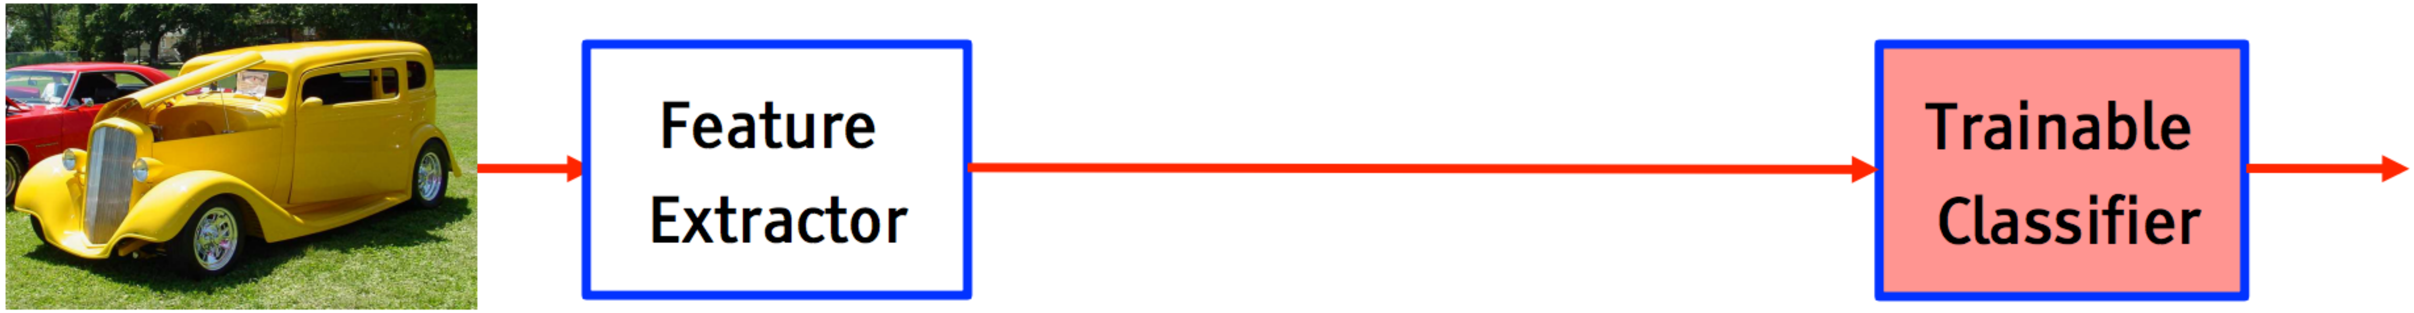

Mainstream Pattern Recognition 9until recently)

Deep Learning: Multiple stages/layers trained end to end

Source:
NIPS'2015 Tutorial
Geoff Hinton, Yoshua Bengio & Yann LeCun

Multiple Layers: Levels of Abstraction

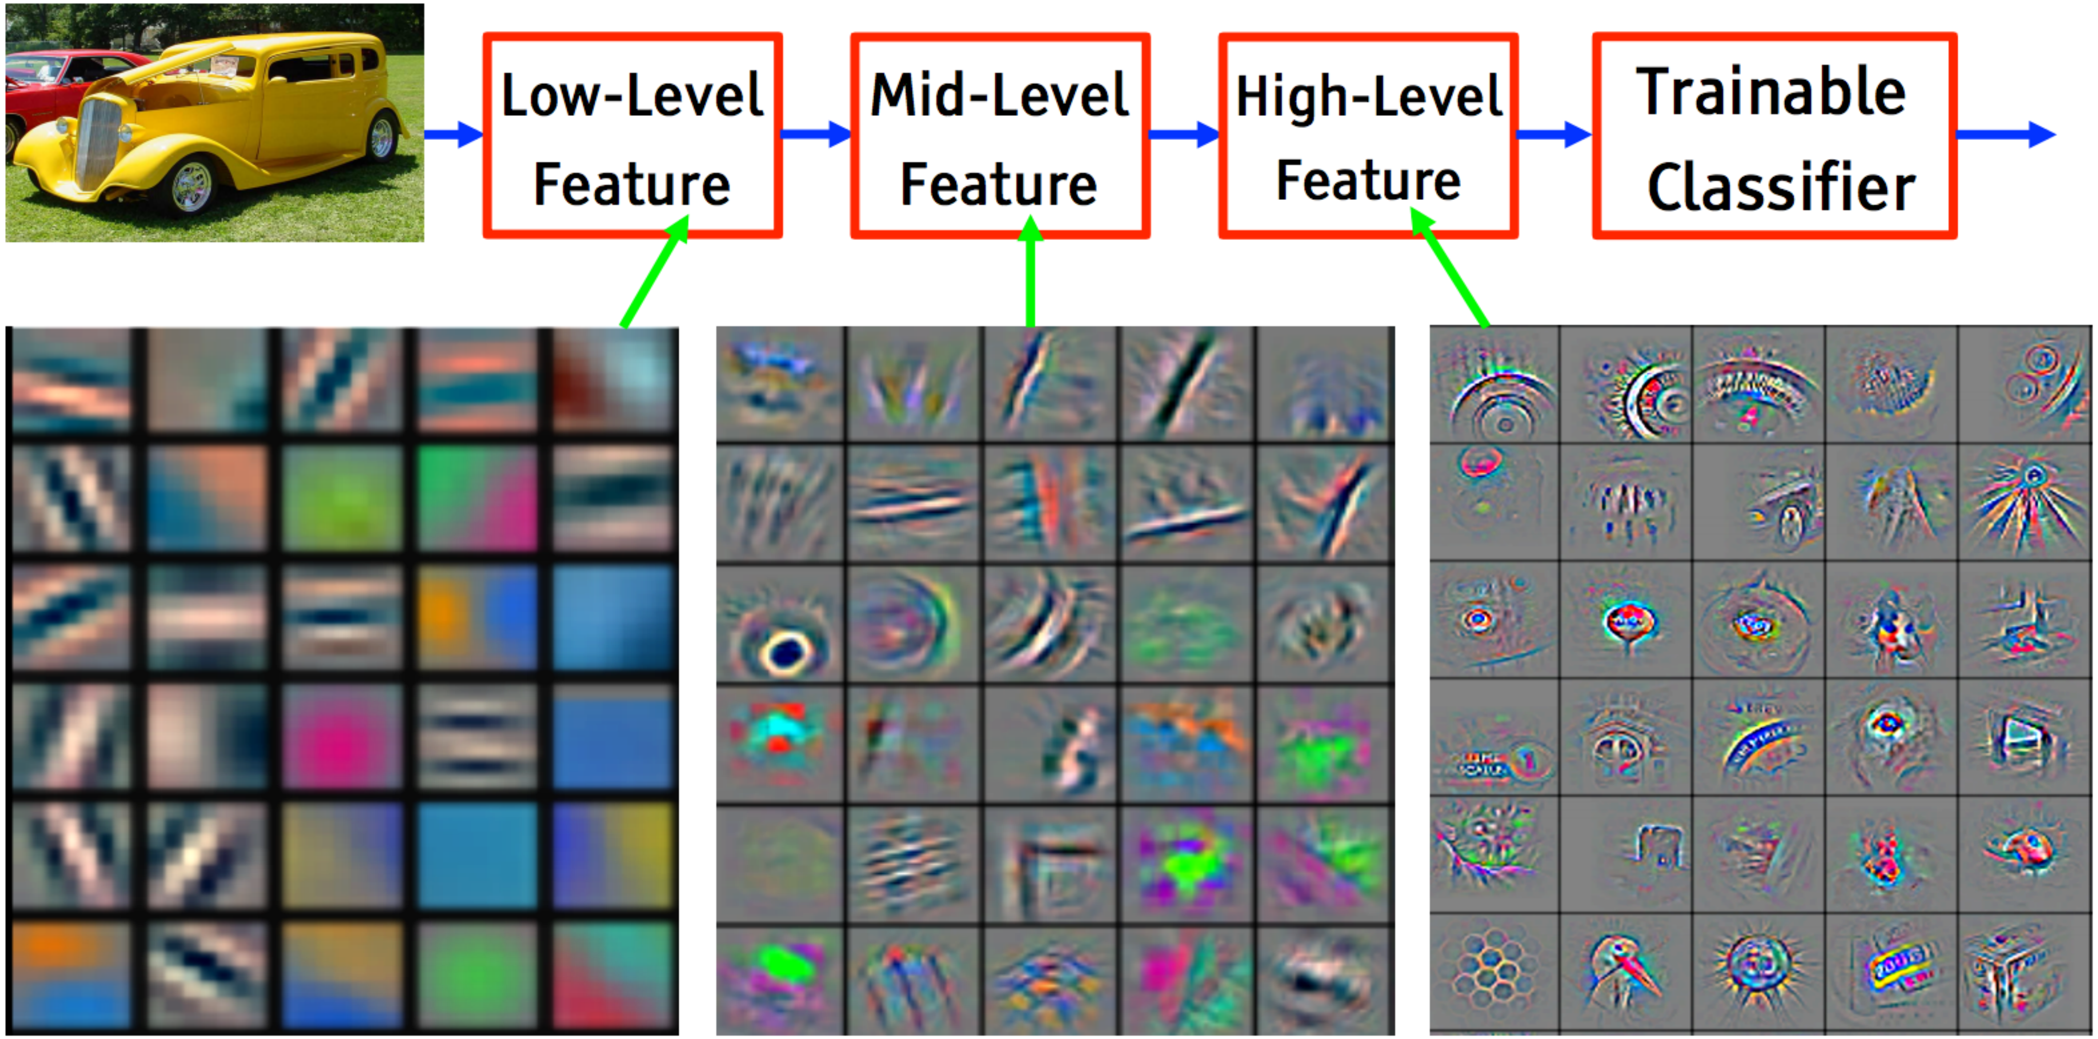

Amazon Picking Challenge

- Robot stow and take (pick) products from shelves
- Challenges in: Computer vision (find object and its pose [position & orientation]) and Manipulation (planning, grasping)
- Since 2015 – now co-located with RoboCup
- Winner 2016: Team Delft
  - Suction gripper (easy!)
  - 100 items per hour (human: 400)
  - Failure rate: 16.7 %
  - Deep learning to find objects:
    - Stereo camera for 3D
  - ROS Industrial

End-to-End Deep Reinforcement Learning

- From sensors to actuation: one layered or recurrent neural network! =>
  - NOT classical general control scheme (Perception, SLAM, Cognition & Planning, Navigation)

- Needs reward signal: sparse, noisy, delayed!
- Take time into account: input frames are related!

- Gained interest 2013 again with:
  - Deep Mind (google) playing ATARI 2600 games
  - Video: Breakout
  - Learned 7 games
  - Surpasses human expert in 3

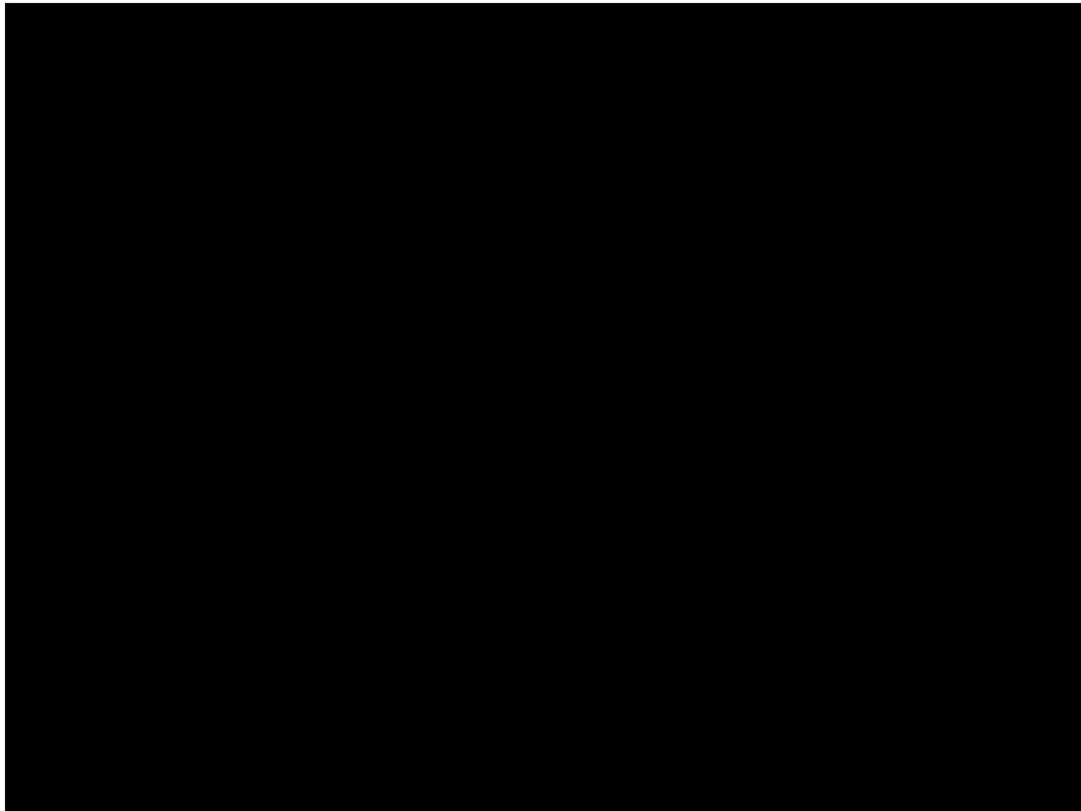

BRETT: Berkeley Robot for Learning Tedious Tasks: Deep Reinforcement Learning

- ”There are no labeled directions, no examples of how to solve the problem in advance. There are no examples of the correct solution like one would have in speech and vision recognition programs”
- Learn simple tasks in 10 minutes; learn vision and control together in 3 hours
- Pieter Abbeel of UC Berkeley;
- 2015

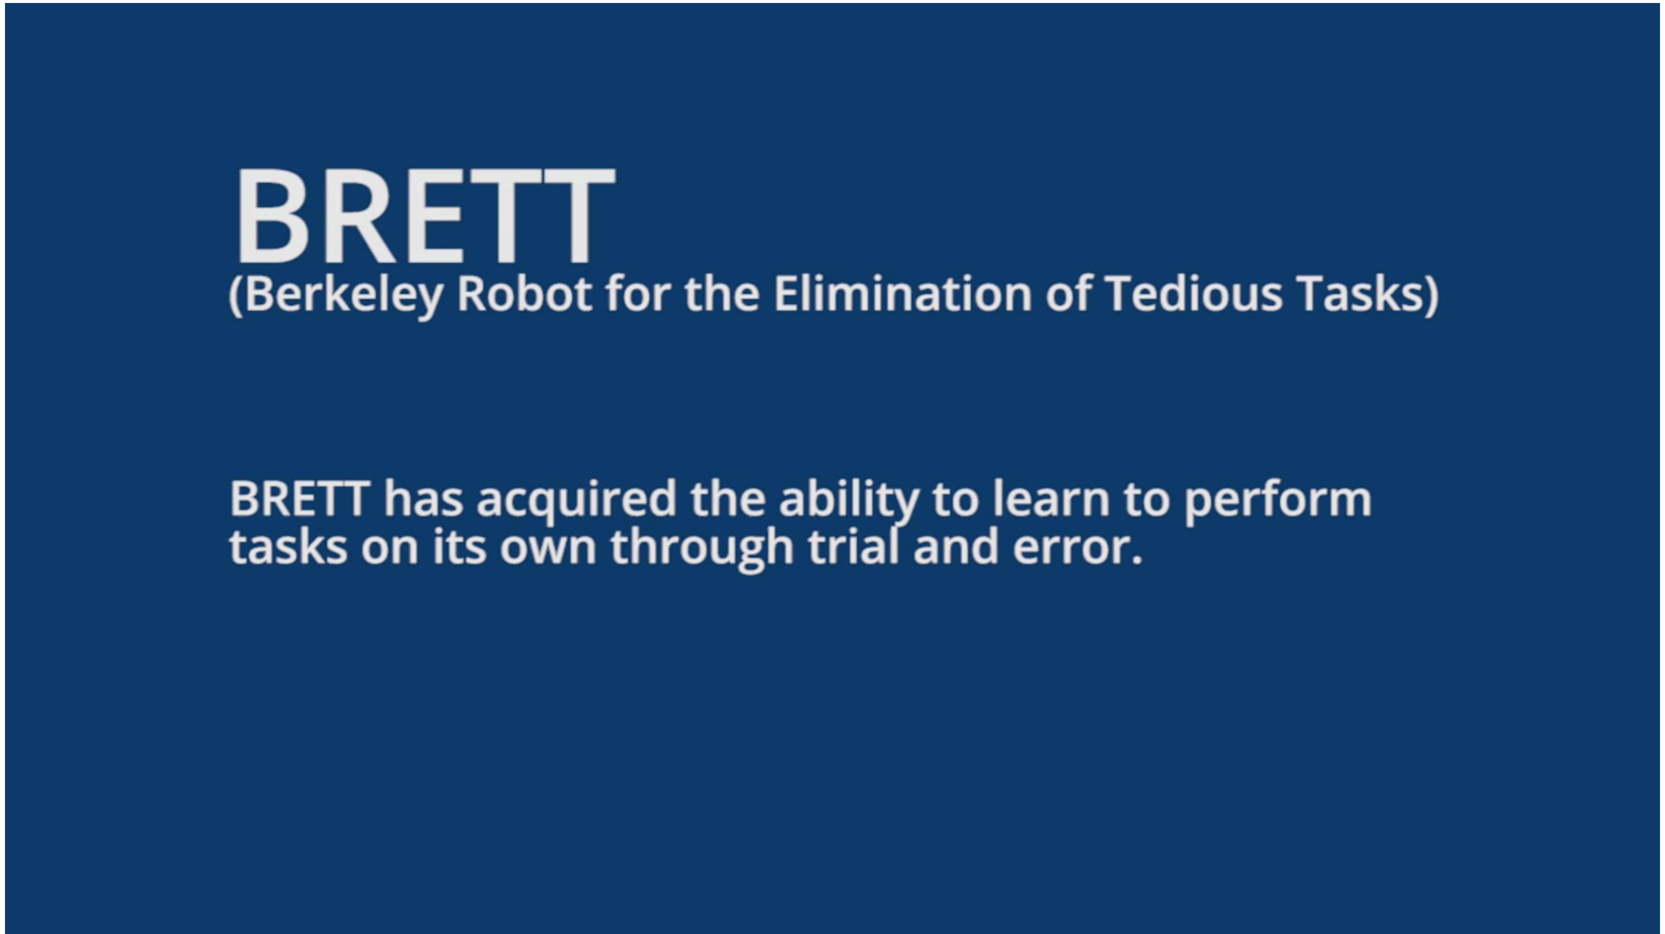

Google Door Opening Project

- Learn to open doors using Reinforcement learning
  - Learning reward: opening the door
  - Much harder than purely digital learning: very slow iterations!
  - Simulation only helps a bit:
    real world much more complex

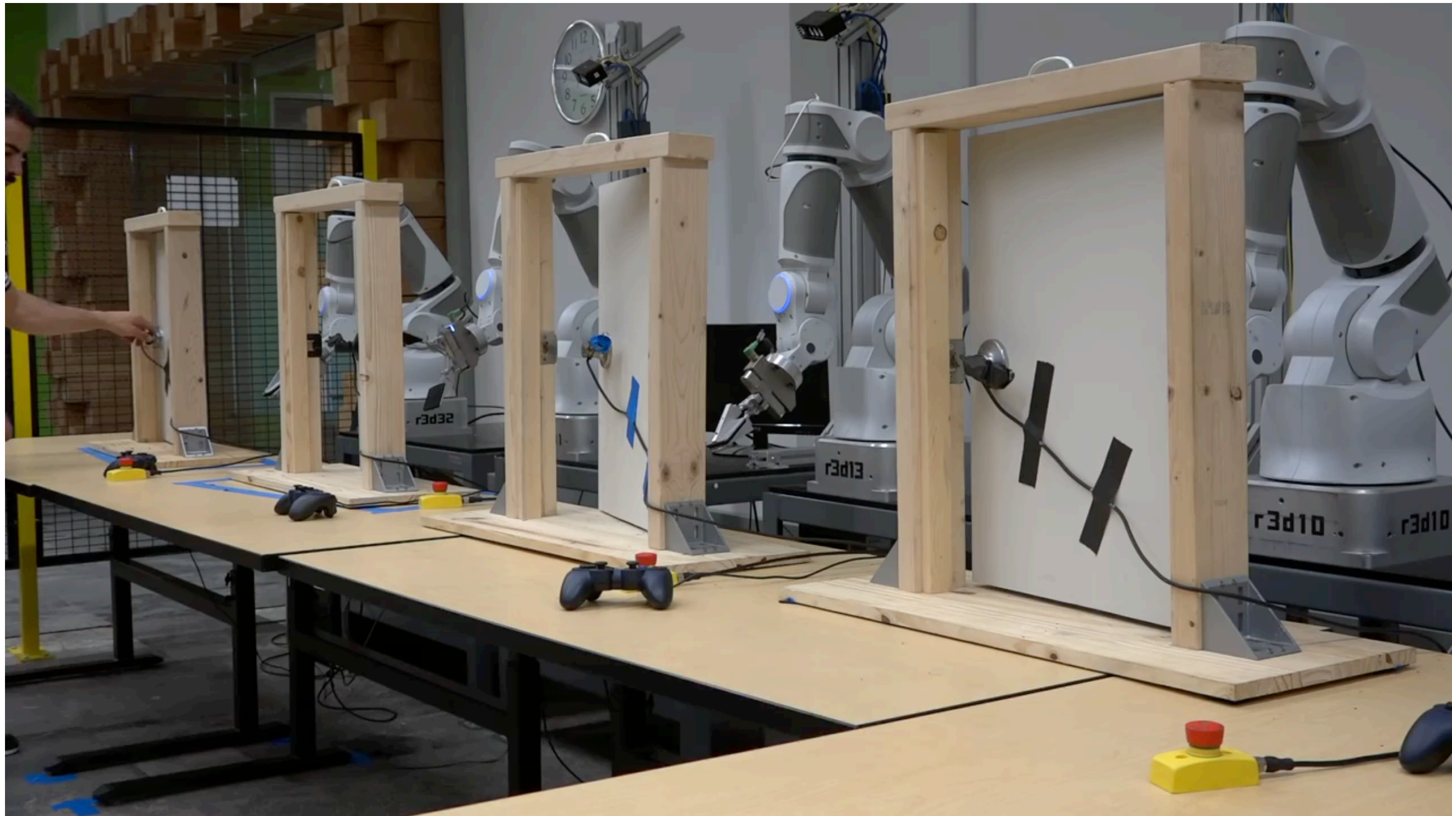

- Google and
  UC Berkeley Sergey Levine

- Google very secretive …

Nvidia end-to-end deep learning self driving car

- Raw pixel of single camera => steering command of car
  - 30 FPS; single-image control (no history)
  - Nvidia Drive PX car computer (ARM cores and Nvidia GPU)
  - Human steering angle for training
- End-to-end: NO explicit:
  - Lane detection
  - Road detection
  - Obstacle detection

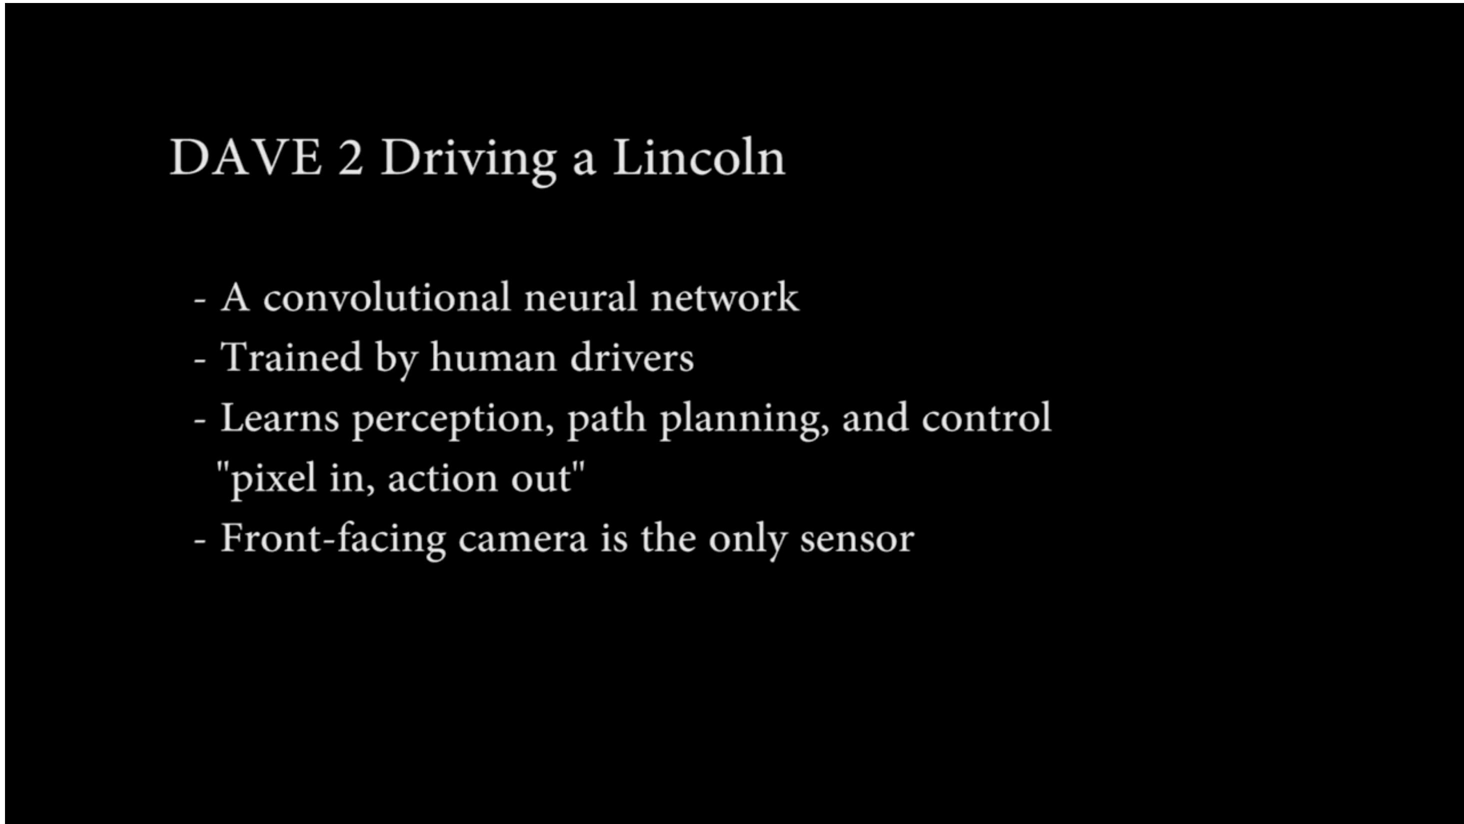

- Training:
  - Additionally use left and right camera: negative examples!
  - Highway, residential roads, unpaved roads, car parks
  - Different weather conditions

- Result: Autonomous 98% of the time => 2% driver intervention

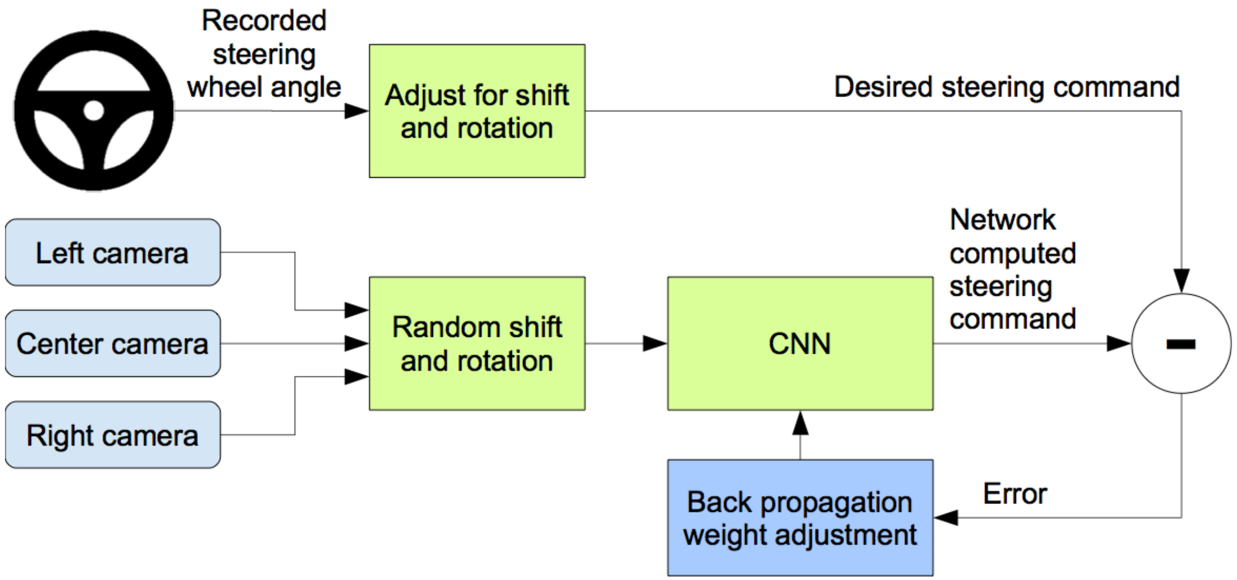

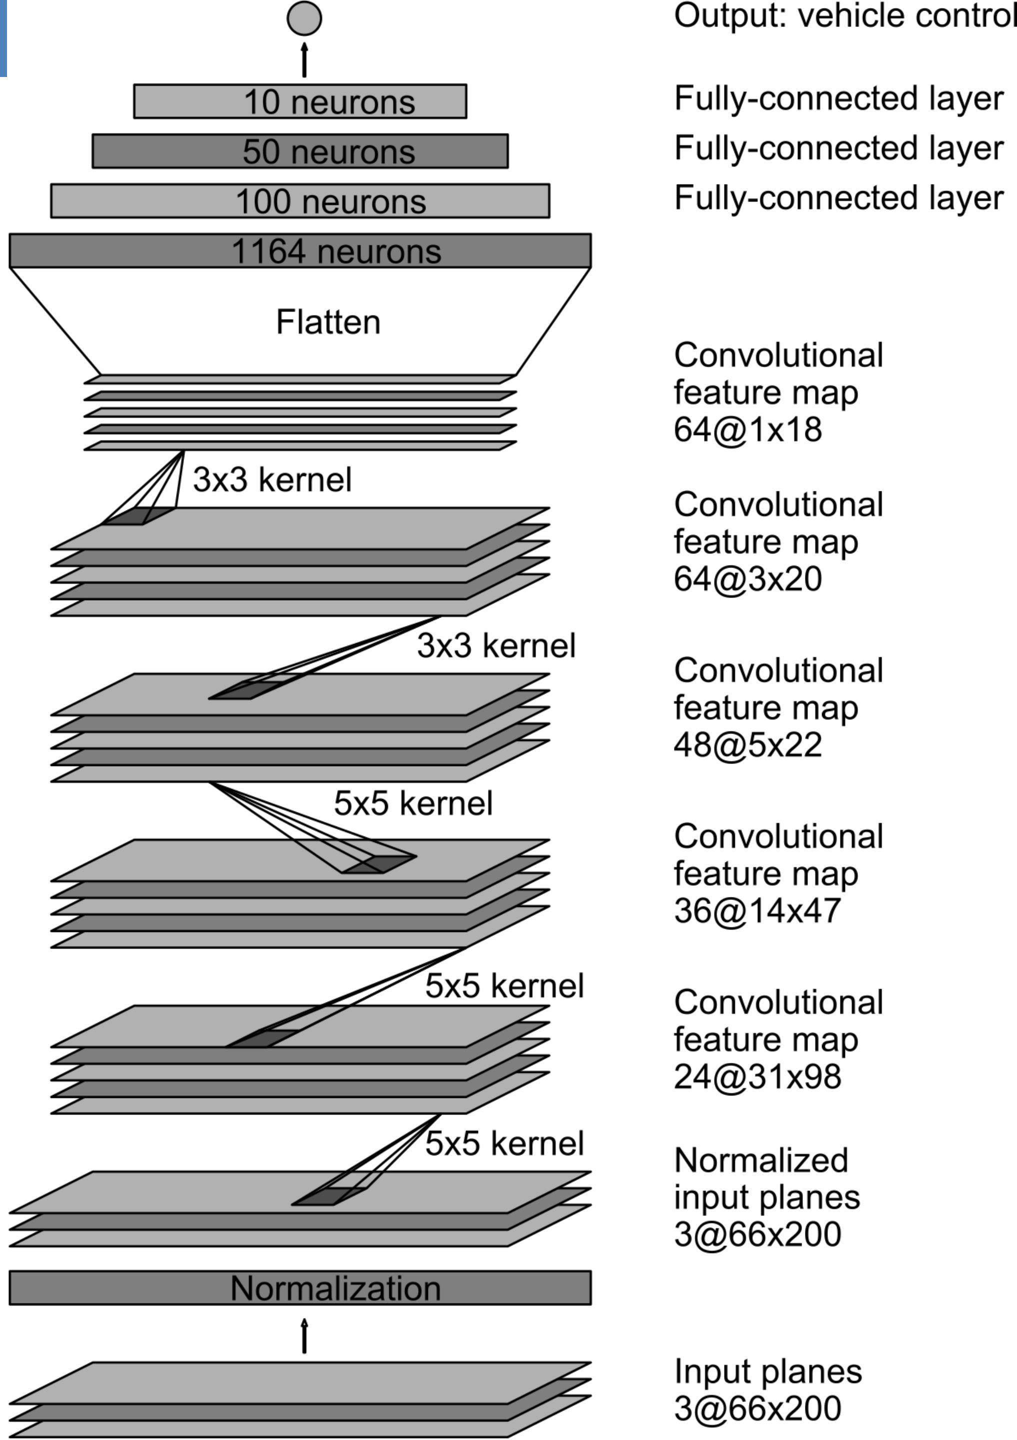

Problems with Deep Learning

- 99% success rate sounds good, but 1% failure is often unacceptable (e.g. autonomous car)
  - Failures are unavoidable =>
  - need quality estimate/ uncertainty of the result!
  - Often not available for DL ☹
- Lack of theory regarding deep learning
  - Acts like a black box…
- No introspection of how or why a DL system is behaving like it is =>
  - No safety guarantees possible
- Deep learning only part of an overall AI system
  - Hand-crafted methods can still be very powerful
  - Modelling useful (with input from DL)
  - Statistical methods
  - Reasoning
  - Planning

Introspection for Robots and AI Agents

- Internal assessment:
  - How to assess the quality of internal models, methods and sensor data, used by robots/ AI agents and how to alter their behavior upon this information?
- Analysis:
  - Failure analysis
  - Execution monitoring
- Introspection-related actions
  - Active Learning
  - Failure recovery
  - Reconfigurable robots
  - Planning with uncertainty

- IROS 2017 Workshop: Introspective Methods for Reliable Autonomy /
  IEEE Transactions on Cognitive and Developmental Systems:
  Special Issue on Introspective Methods for Reliable Autonomy

Safety for Robotics & Standards

- New generation of robots enter human workspace (industry, office, home, streets) => safe behaviors essential!
- Robots: very complex systems => difficult to give guarantees
  - Deep Learning: black box => no guarantees? Just tests?
- Safety Certification needed => Standardized Tests
  - Example: ASTM International Standards Committee on Homeland Security Applications; Response Robots (E54.09)
  - Align test: don’t drop off; needs: good perception + good control
  - developed by National Institute of Standards and Technology (NIST) US

- Many Robotics Standard Bodies (few examples):
  - ISO/TC 299 Robotics (15 Standards; 12 under development)
  - International Electrotechnical Commission (IEC) TC 62/JWG 35 and 36
  - ASTM International
  - IEEE Standard for Ontologies for Robotics and Automation (IEEE P1872)
  - Robotic Industries Association (RIA): ANSI/ RIA R15.06-2012 Robot Safety Standard
  - US: OSHA/NIOSH/RIA Alliance for robotics safety (founded Oct. 6, 2017)
  - IEEE Standards Committee on Robot 3D Map Data Representation (under development)

RoboCup Rescue

- Yearly competition for rescue robots: teleoperated or autonomous, big or small, with or without manipulation, flying, …
- ASTM Test Methods for scoring!

Safety is important ;)