Digital Reality & AI
Realistic Rendering for Training & Validating AI with Synthetic Data

Philipp Slusallek

German Research Center for Artificial Intelligence (DFKI)

Saarland University, Computer Graphics Lab (UdS-CGL)

European High-Level Expert Group on AI (HLG-AI)

DFKI: An Overview

German Research Center for Artificial Intelligence (DFKI)

- **Overview**
  - Largest AI research center worldwide (founded in 1988)
  - Germany’s leading research center for innovative software technologies
  - 6 sites in Germany
    - Saarbrücken, Bremen, Kaiserslautern; Berlin, Osnabrück, Oldenburg
  - 20 research areas, 8 competence centers, 8 living labs
  - More than 1100 research staff & support
  - Revenues of >58 M€ in 2019 (50 M€ in 2018)
  - More than 90 spin-offs

Germany Has a Head-Start

DFKI: The World´s Largest Center for Research & Application in AI

SAP
NVIDIA
intel
SCHWARZ
LIDL
Kaufland
HARTING
ZF
JOHN DEERE
Deutsche Telekom
software AG
Microsoft Deutschland GmbH
Google
DEUTSCHE BÖRSE GROUP
accenture
EMPOLIS INFORMATION MANAGEMENT
Deutsche Messe
KIBG
Fraunhofer

CLAAS
VOLKSWAGEN AKTIENGESELLSCHAFT
DAIMLER
BMW GROUP
Rolls-Royce
BOSCH
AIRBUS GROUP
RICOH
VSE
BILFINGER
ROSEN empowered by technology
cerence

DFKI
UNIVERSITÄT DES SAARLANDES
Universität Bremen
TECHNISCHE UNIVERSITÄT KAISERSLAUTERN
Munich RE

Currently 35 Professors are Working for DFKI

15 Associated and Supernumerary Professors

20 DFKI Heads of Research Labs


Prof.
Sebastian Möller
Executive Director
Berlin


Prof.
Andreas Dengel
Executive Director
Kaiserslautern


Prof.
Frank Kirchner
Executive Director
Bremen


Prof.
Antonio Krüger
CEO


Prof.
Philipp Slusallek
Executive Director
Saarbrücken


Prof.
Joachim Hertzberg
Executive Director
Osnabrück/Oldenburg


Prof.
Rolf Drechsler


Prof.
Gesche Joost


Prof.
Jana Koehler


Prof.
Peter Loos


Prof.
Paul Lukowicz


Prof.
Wolfgang Maaß


Prof.
Volker Markl


Prof.
Niels Pinkwart


Prof.
Martin Ruskowski


Prof.
Hans Schotten

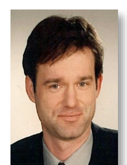
Prof.
Didier Stricker


Prof.
Oliver Thomas


Prof. Josef
van Genabith


Prof.
Oliver Zielinski


Prof.
Ziawasch Abedjan


Prof.
Klaus-Dieter Althoff


Prof.
Stephan Busemann


Prof.
Peter Fettke


Prof.
Udo Frese


Prof.
Tim E. Güneysu


Prof.
Dieter Hutter


Prof.
Jochen Kuhn


Prof.
Christoph Lüth


Prof.
Günter Neumann


Prof.
David Schlangen


Prof.
Wolfgang
Wahlster
CEA


Prof.
Hendrik Wöhrle


Prof.
Robert Wille


Prof.
Hans Uszkoreit

DFKI Covers the Complete Innovation Cycle

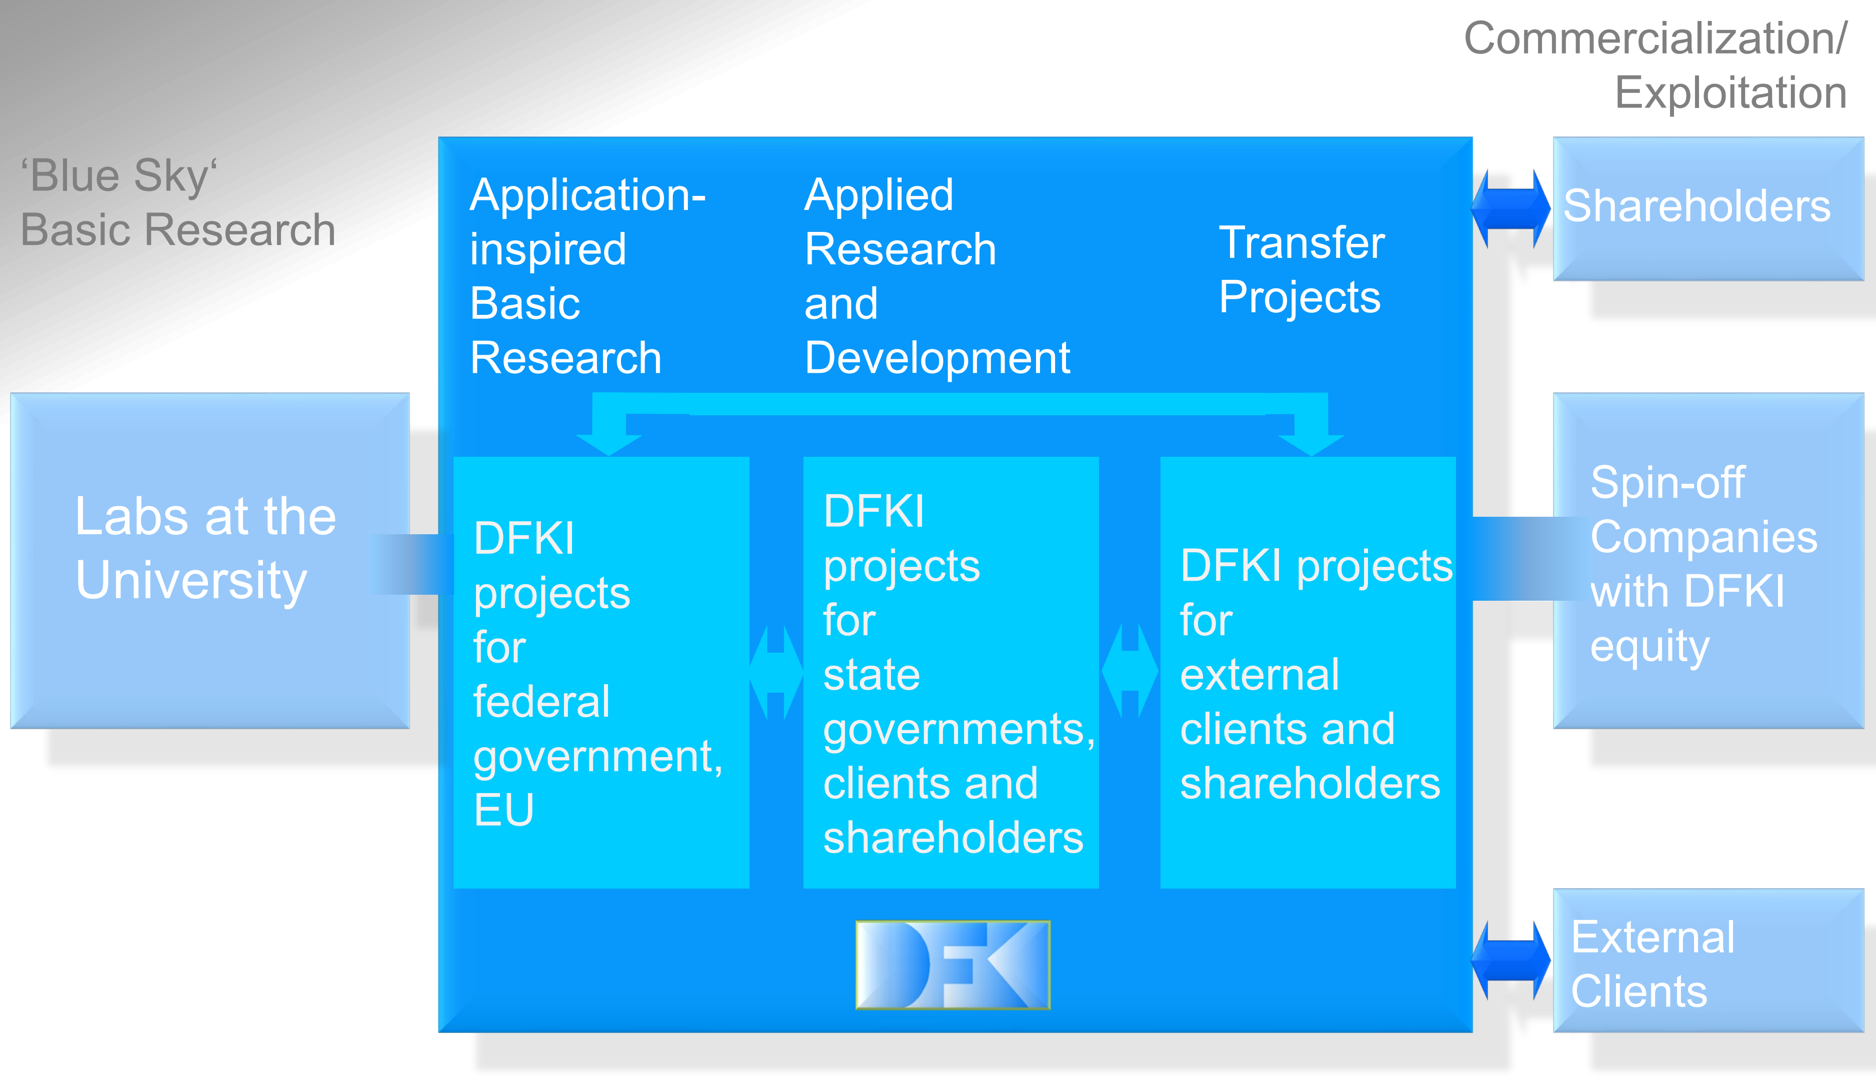

DFKI-Portfolio: Deep Expertise in AI for a Broad Innovation Spectrum

Max Planck Society	Fraunhofer	Helmholtz Society
Application-Oriented Basic Research	**Applied R&D and Transfer**	Large Test- and Demonstration Centers

The entire innovation chain in the horizontal spectrum of DFKI

The vertical specialisation of DFKI on methods and applications of Artificial Intelligence

Deep knowledge and excellence in one important section of Computer Science

DFKI Employees

Broad Methodological and Systems Competence in Artificial Intelligence

Deep Scientific Expertise in AI Technology

Deep Domain Knowledge in an Area of Application

End-to-End Demonstrations in Seven Living Labs of DFKI

Innovative Retail Lab (IRL)

Advanced Driver Assistance Systems Living Lab (ADAS)

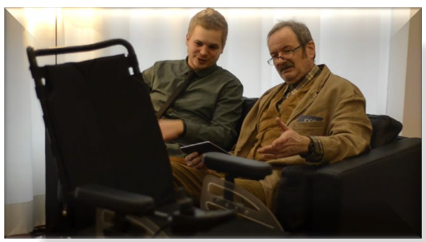

Bremen Ambient Assisted Living Lab (BAALL)

Smart Factory

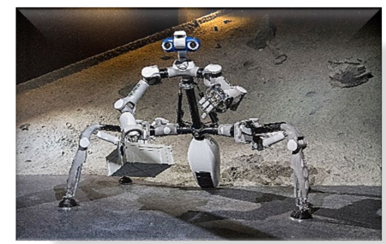

Robotics Exploration Lab (RIC)

Smart City Living Lab

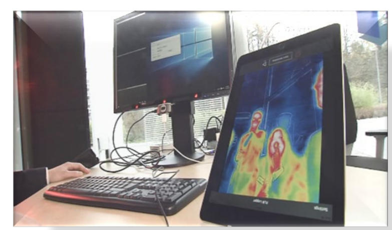

Smart Office Space

DFKI Research Area:
Agents and Simulated Reality (ASR)

Agents and Simulated Reality

AI, Graphics/Simulation, High-Performance Computing

Scientific Director
Philipp Slusallek

Christian Müller (Deputy)

Silke Balzert-Walter (Consulting)

Application Domains

Strategy Board

- **Autonomous Driving** Christian Müller
- **Computational Sciences** Tim Dahmen
- **Smart Living** Hilko Hoffmann
- **Industrie 4.0** Ingo Zinnikus
- **High-Performance Computing** Richard Membarth

Research Teams

Hybrid & Symbolic AI

- **Multi-Agent Systems** Klaus Fischer
- **Intelligent Information Systems** Matthias Klusch

ML / DL

- **Computational 3D Imaging** Tim Dahmen

AI Platform

- **Distributed & Web-Based Systems** René Schubotz
- **High-Performance Graphics & Computing** Richard Membarth

Application-Driven Teams

- **Autonomous Driving** Christian Müller
- **Smart Living** Hilko Hoffmann

Important German, European & International Cooperations:

EU High-Level Expert Group

CLAIRE

HUMANE AI

Philipp Slusallek (Co-Initiator), Silke Balzert-Walter

Lernende Systeme
DIE PLATTFORM FÜR KÜNSTLICHE INTELLIGENZ

Philipp Slusallek, Christian Müller

NPM
NATIONALE PLATTFORM ZUKUNFT DER MOBILITÄT

Philipp Slusallek

SMART LIVING
MADE IN GERMANY

Hilko Hoffmann

National Institute of Advanced Industrial Science and Technology
AIST

AIRC

Philipp Slusallek, Christian Müller

DFKI

Flexible Production Control Using Multiagent Systems at Saarstahl, Völklingen

DFKI multi-agent technology is running the steelworks, 24/7 for >12 years, 5 researchers transferred

Physically-Based Image Synthesis with Real-Time Ray Tracing

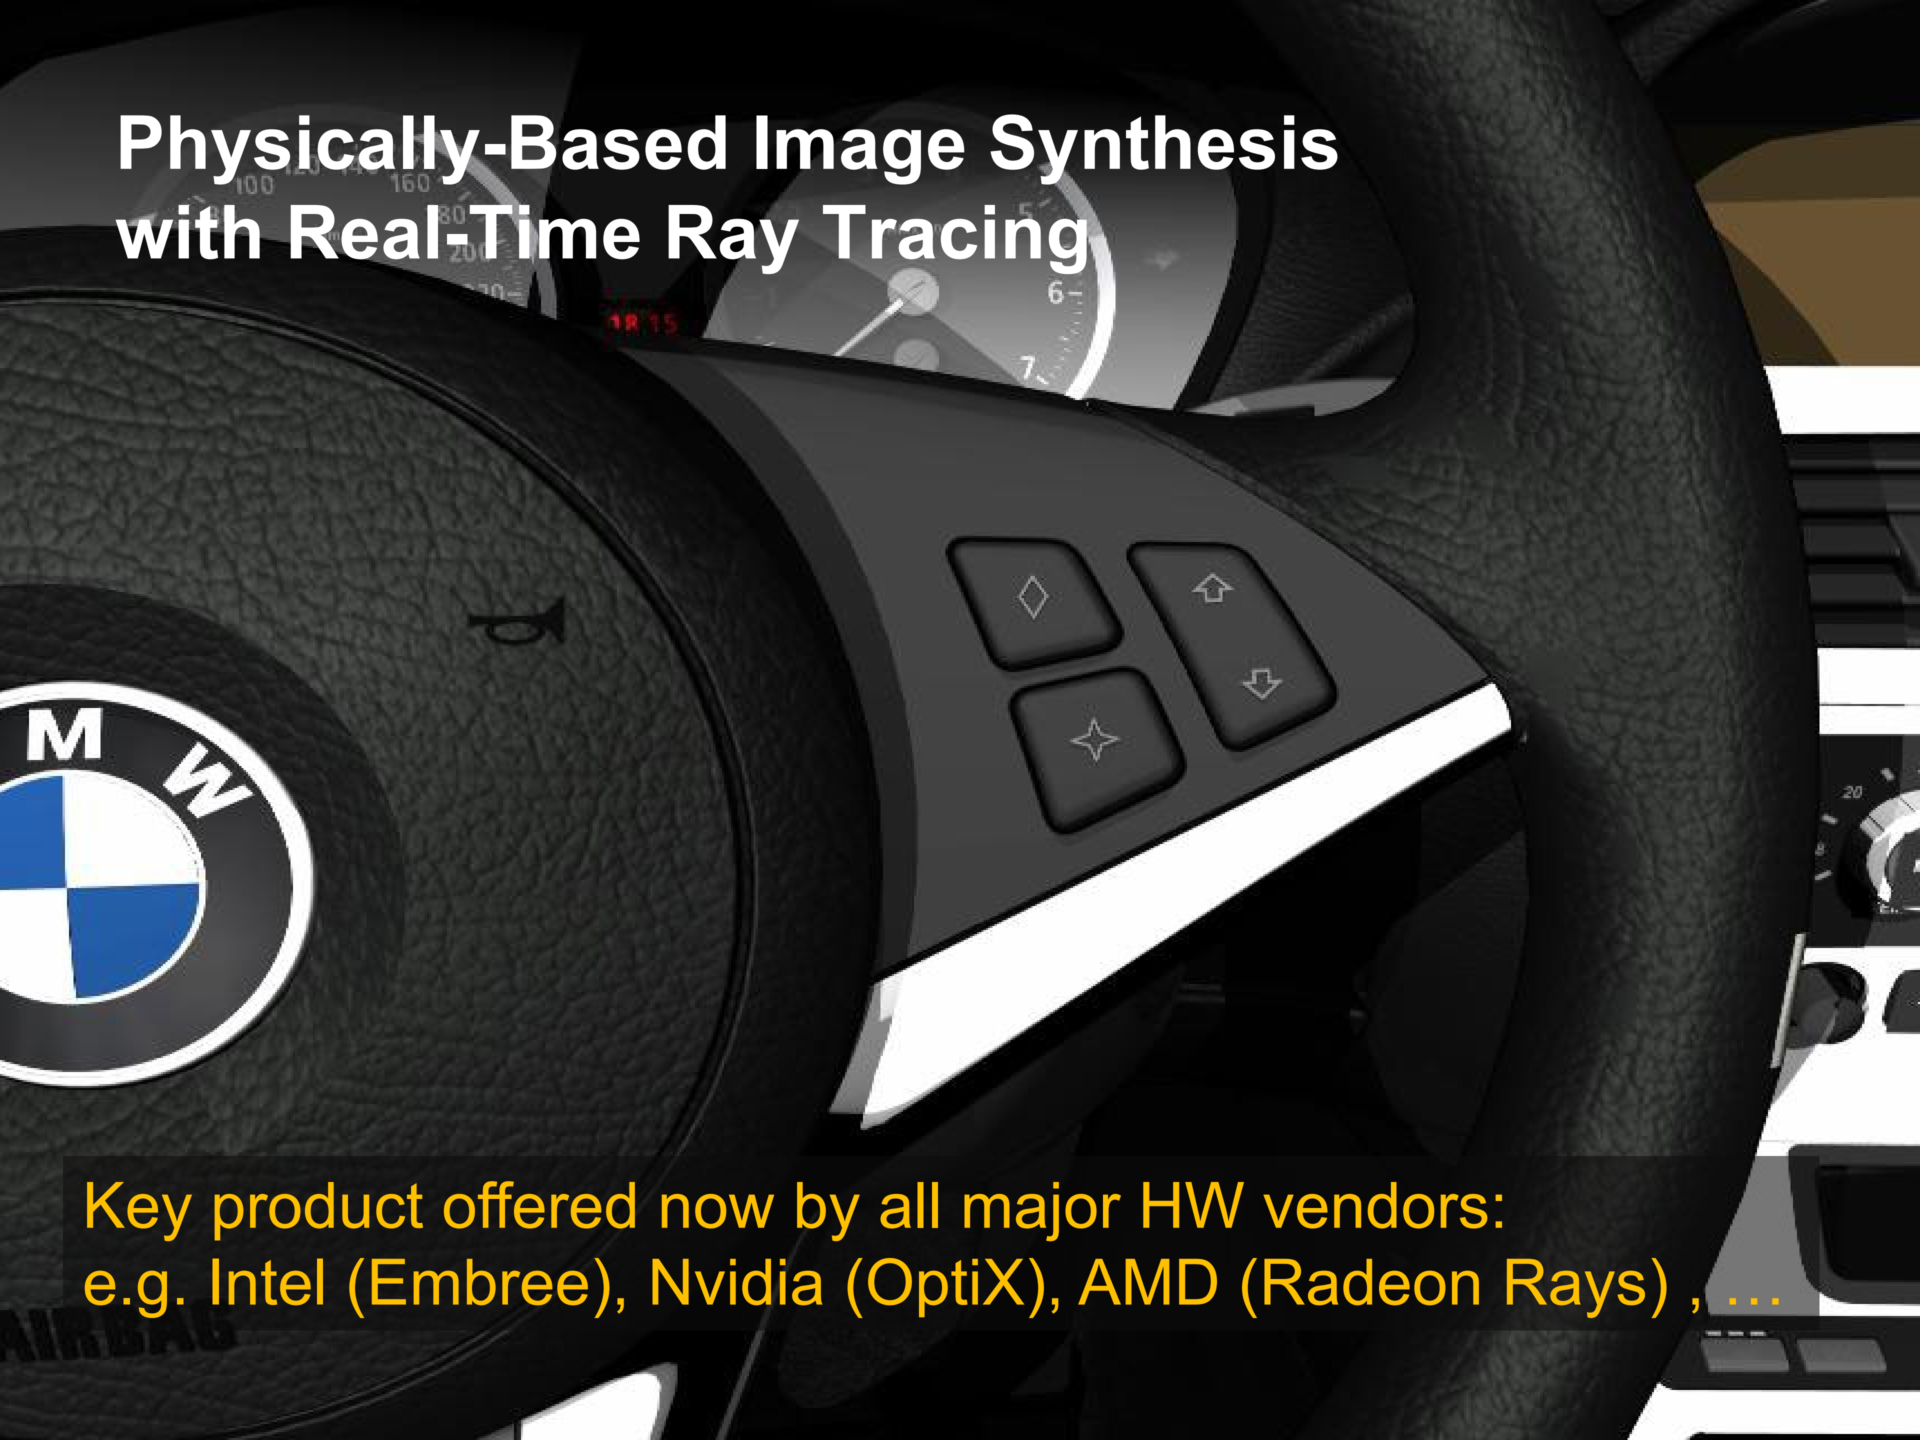

Key product offered now by all major HW vendors:
e.g. Intel (Embree), Nvidia (OptiX), AMD (Radeon Rays) , …

Efficient Simulation of Illumination: Light Propagation and Sensor Models

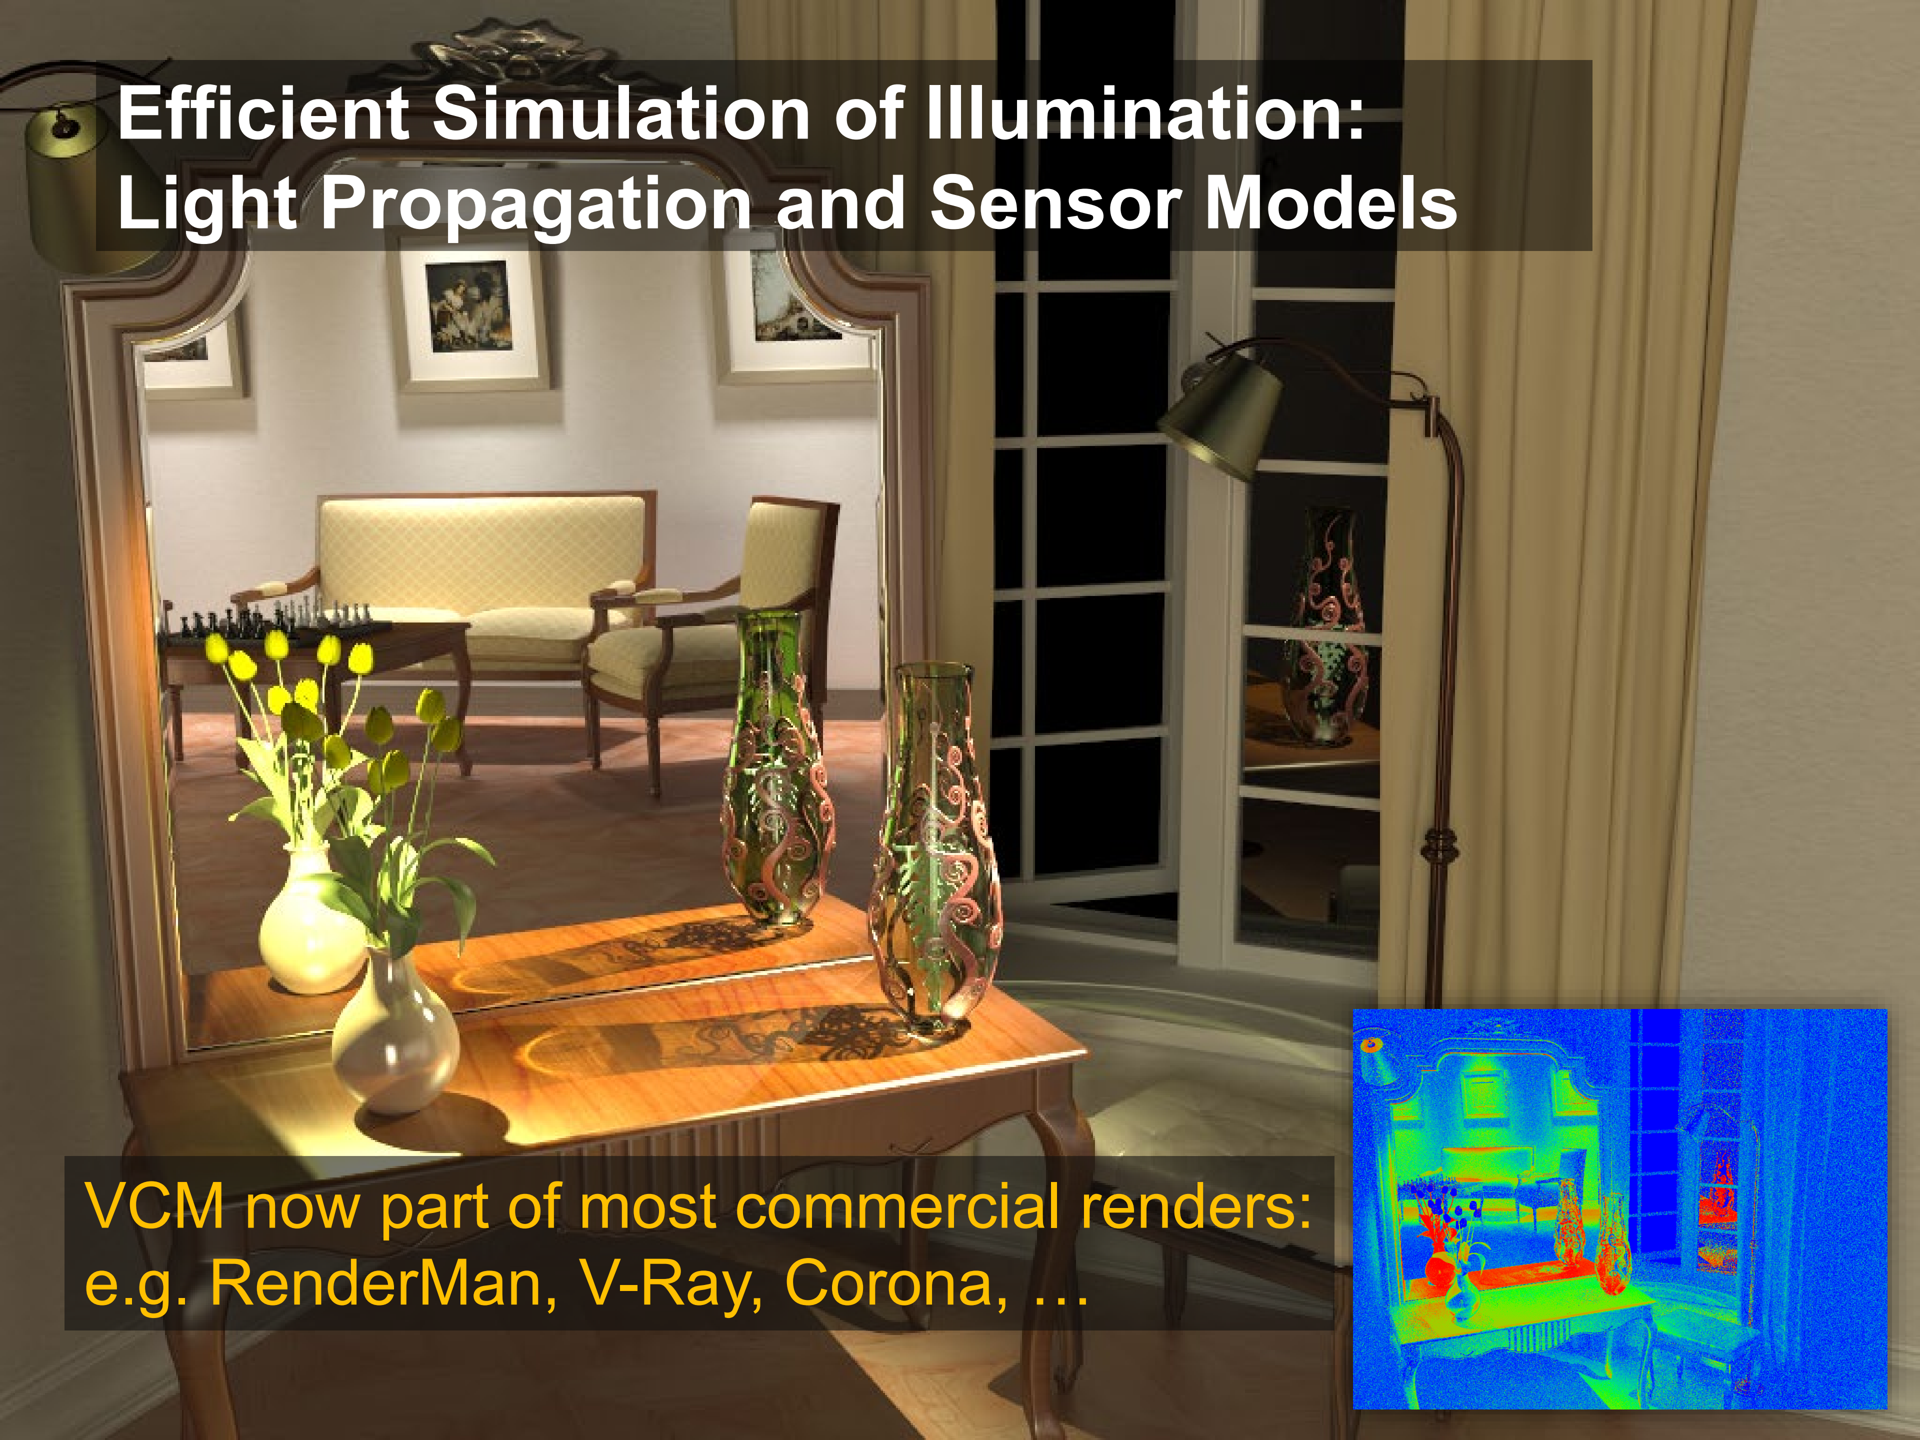

VCM now part of most commercial renders: e.g. RenderMan, V-Ray, Corona, …

Large Visualization Systems Using Ray-Tracing

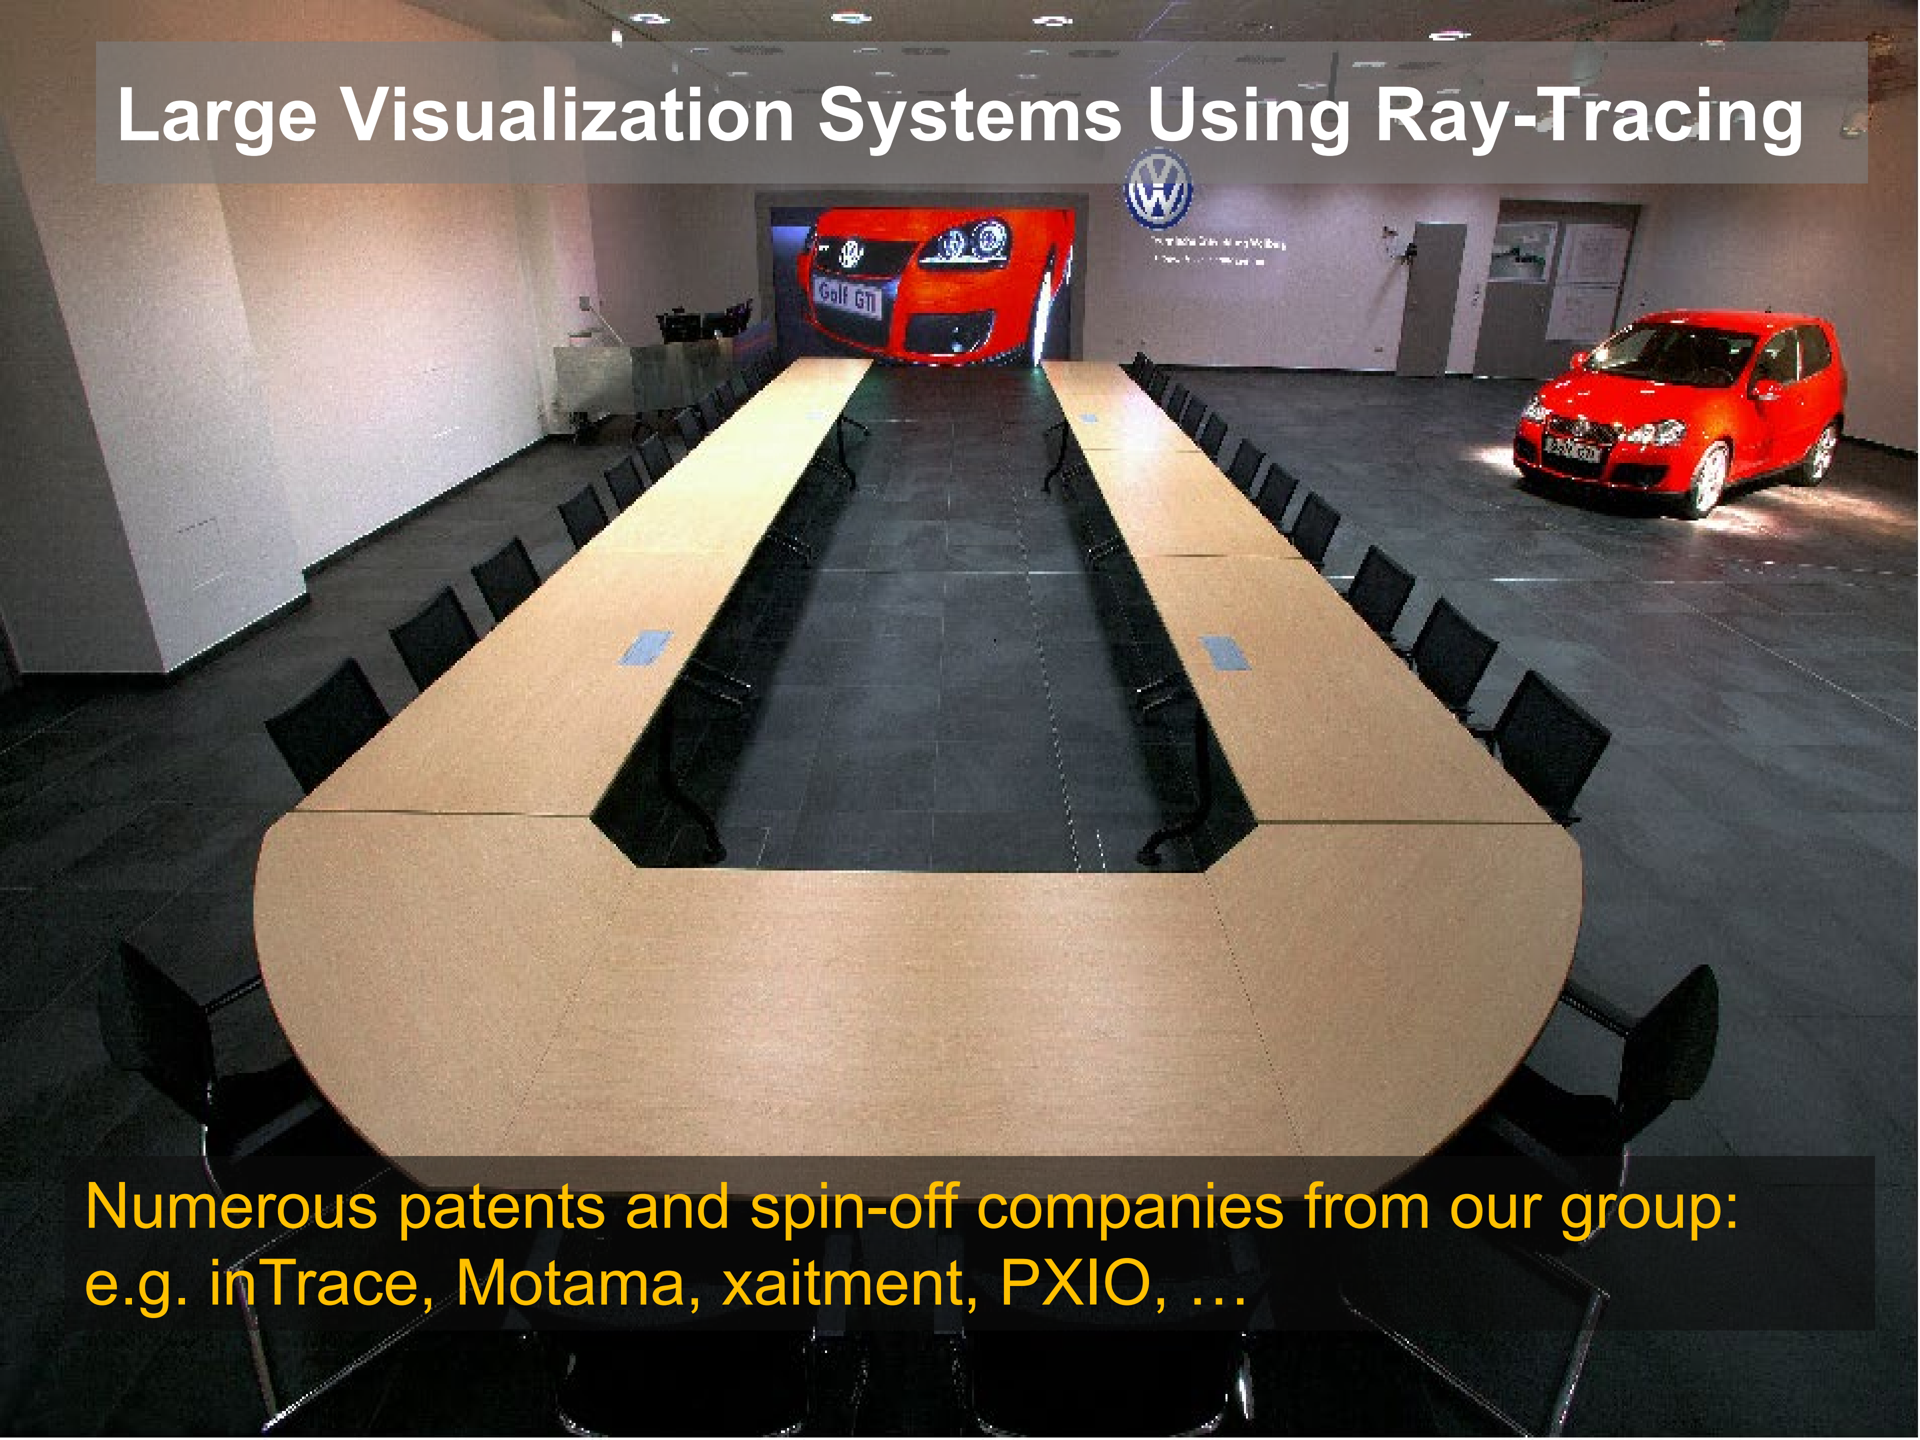

Numerous patents and spin-off companies from our group:
e.g. inTrace, Motama, xaitment, PXIO, …

Custom Ray Tracing Processor [Siggraph'05]

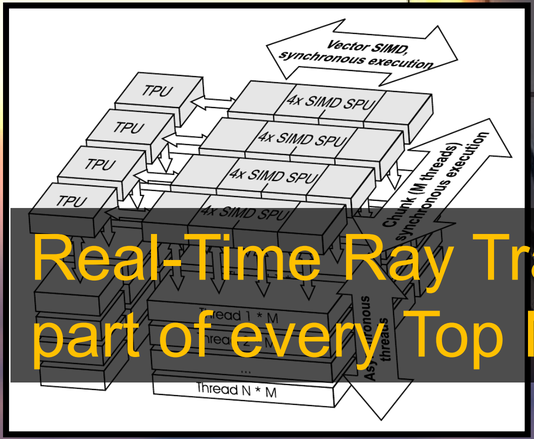

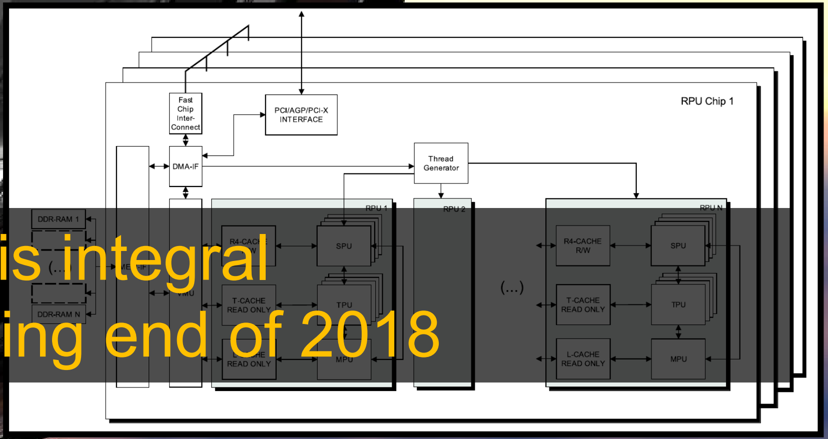

Real-Time Ray Tracing Hardware is integral part of every Top Nvidia GPU starting end of 2018

AnyDSL Compiler Framework

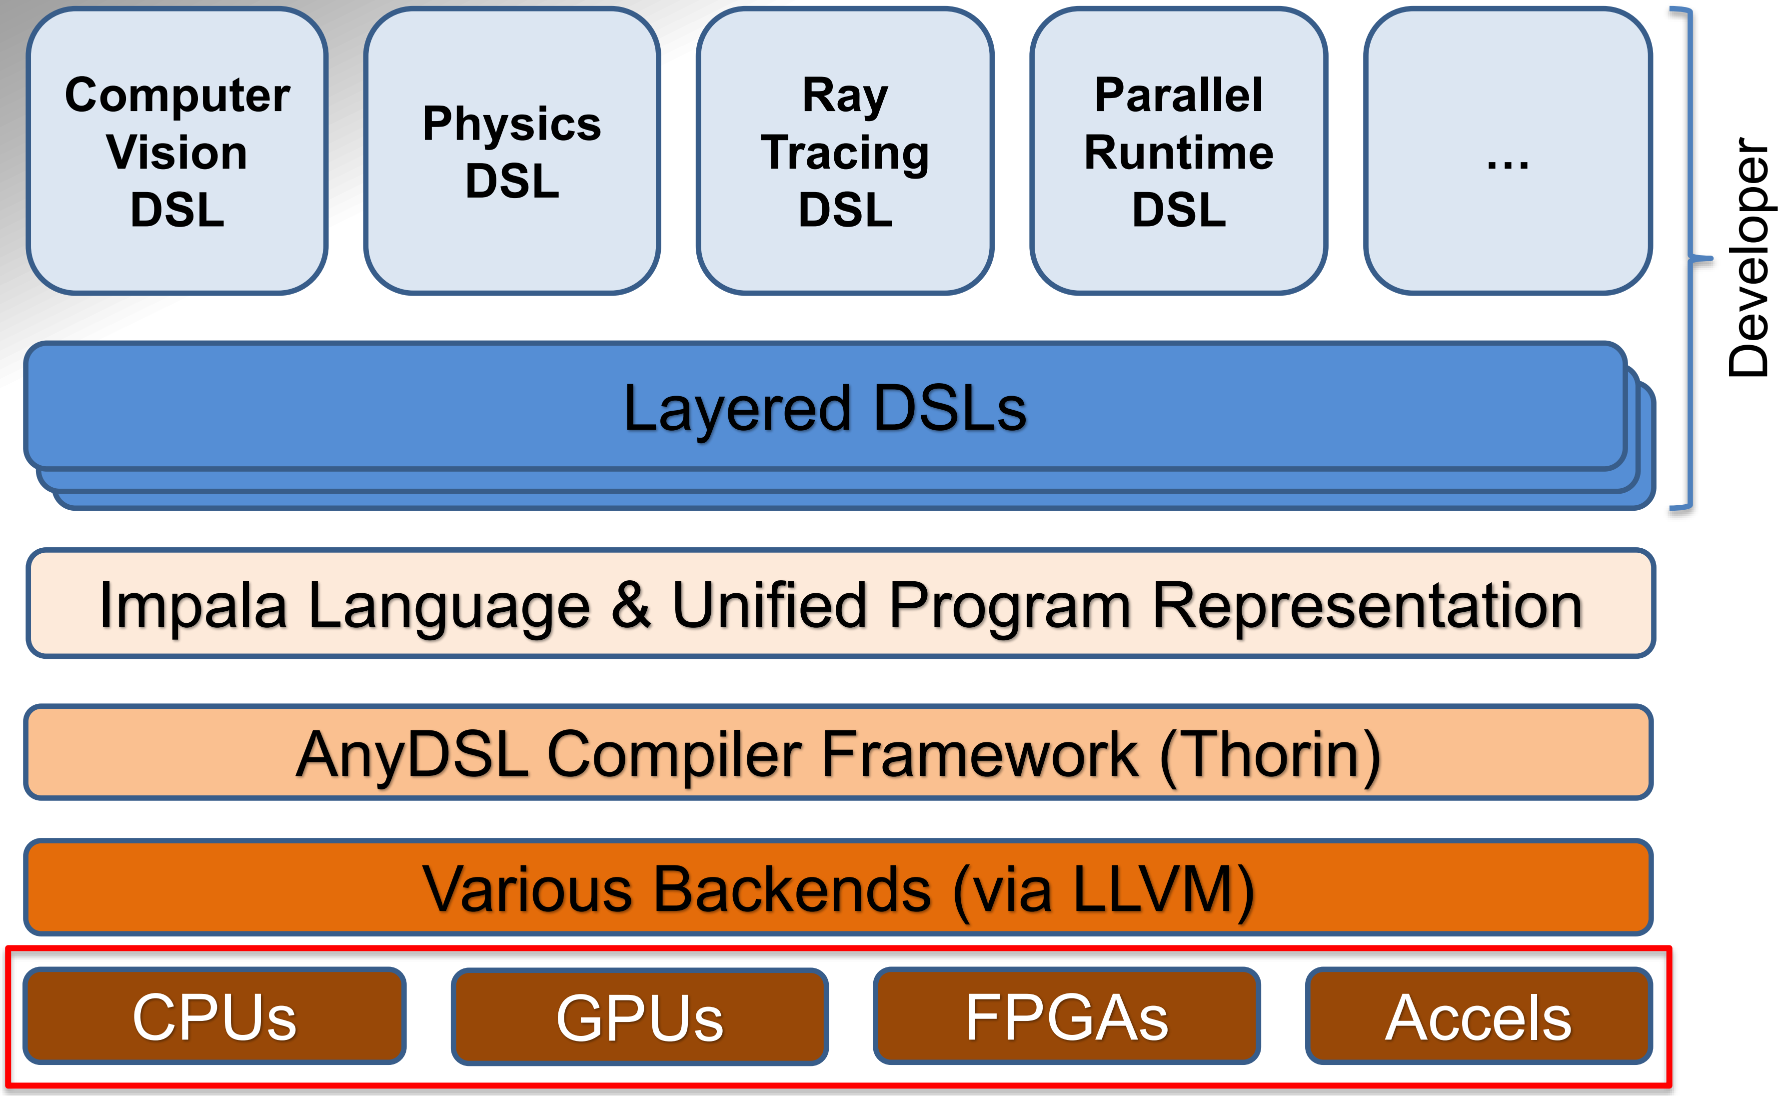

Real-Time Photorealistic Rendering on Film Sets

Material Science: Understanding & Predicting Effects of 3D Structures Across Scales

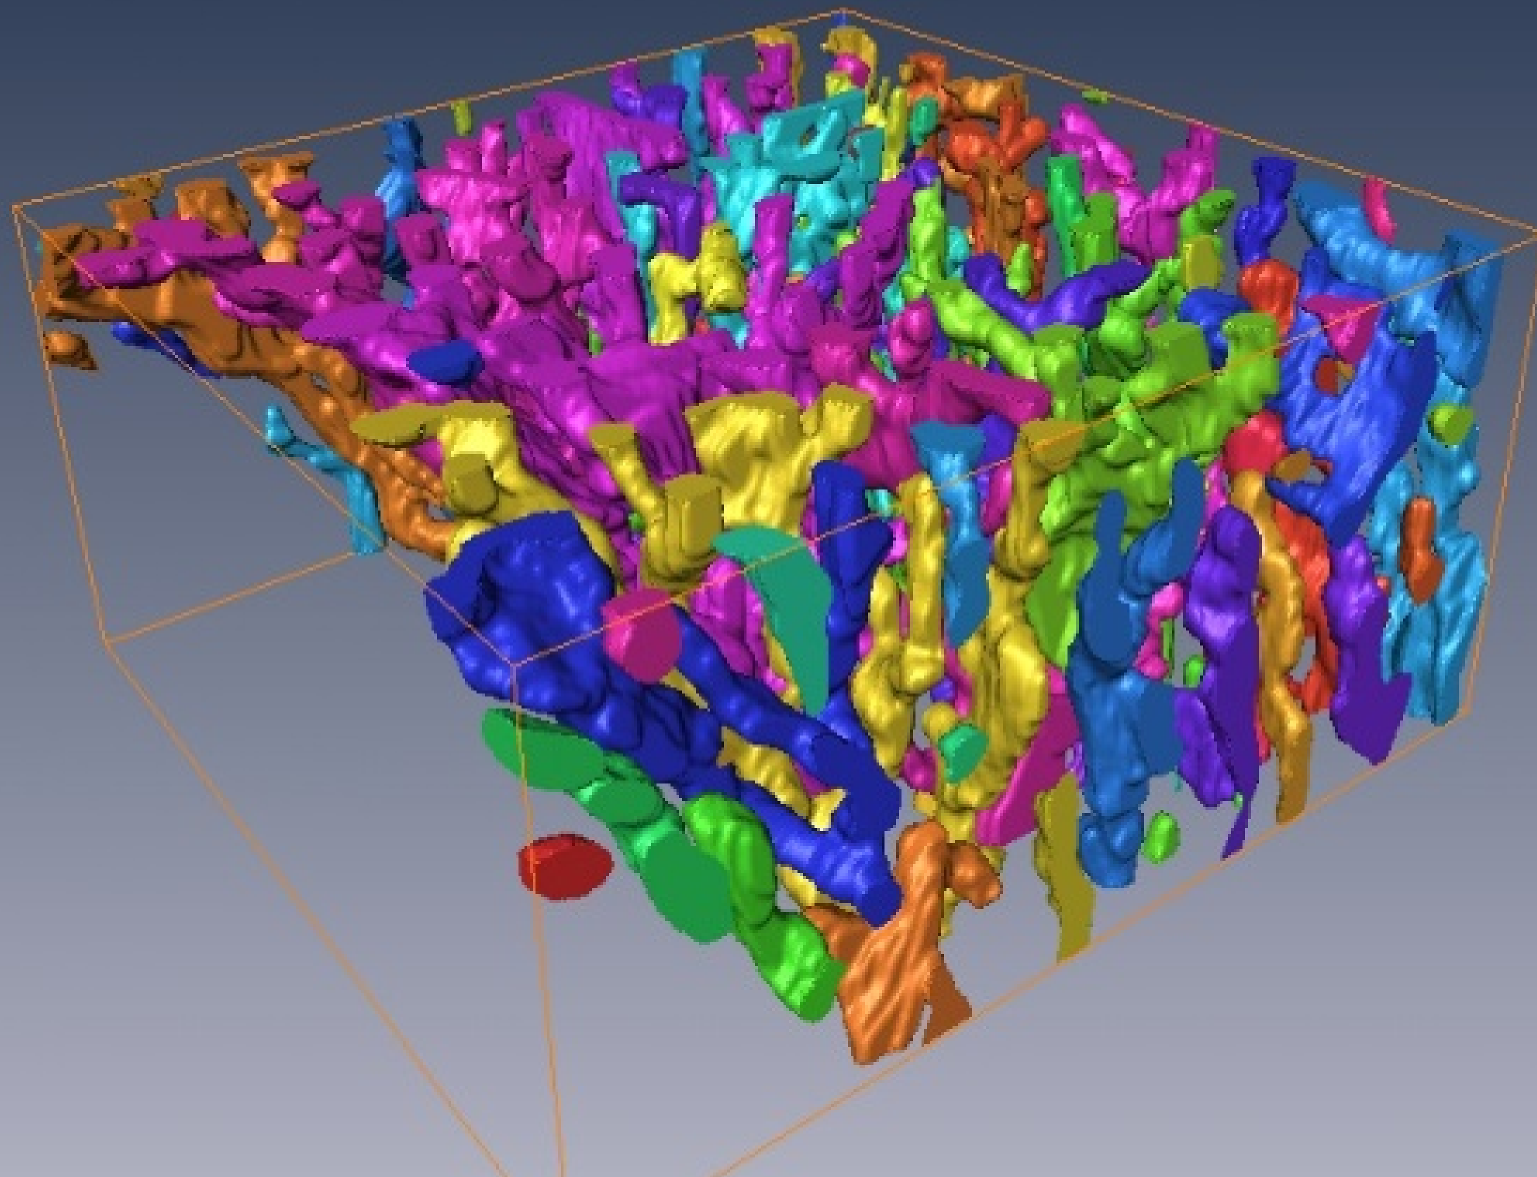

Efficient Acquisition of Imaging Data using AI (e.g. for Connectomics with EM)

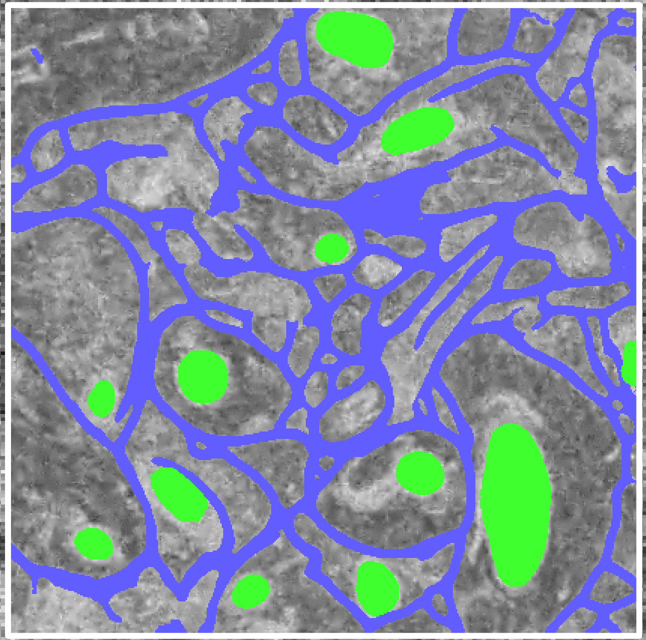

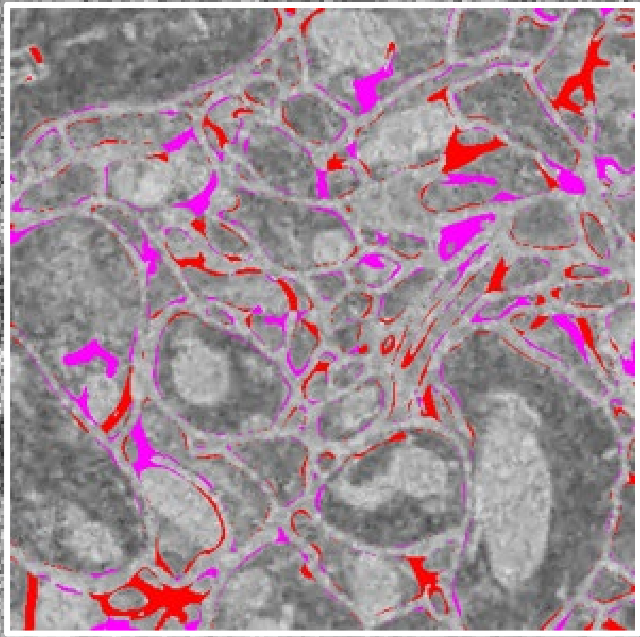

Intelligent Human Simulation, e.g. in Production Environments (Daimler, …)


Collaborative Robotics and
Simulated Reality (VW, Airbus, …)

Autonomous Driving: Training using Synthetic Sensor Data (TÜV, OEMs, Suppliers, …)

Artificial Intelligence:
What is this?

What is AI?

- **Intelligence: Hard to define**
  - “The ability or inclination to perceive or deduce information, and to retain it as knowledge to be applied towards adaptive behaviors within an environment or context”

- **Artificial Intelligence (AI)**
  - “Intelligence” exhibited by machines
  - Machines mimicking *"cognitive functions”* that humans associate with other human minds
  - Study of “Intelligent/Autonomous Agents or Systems“

cf. Russell & Norvig (AI „Bible“), Wikipedia

What is AI?

- Elements of AI Systems ("See – Think – Act" Paradigm):
  - **Perception**
    - … to find out what the world around "looks" like (incl. digital sources)
  - **Machine Learning (Deep Learning)**
    - … to automatically build/adapt models of reality
  - **Knowledge Representation**
    - … to represent, process, and communicate models of the environment
  - **Reasoning**
    - … to generalize and infer new models based on existing knowledge
  - **Planning**
    - … to plan and simulate in the virtual world, exploring possible options, evaluating their performance, and making decisions
  - **Acting**
    - … to manipulate and move around in the real/virtual world – and start over!!
  - **Supporting Functionality**
    - Validation/Verification, Security, Ethics/Social Aspects, SW/HW-Platform
  - **AI Applications**
    - AI functionality can be applied in almost every context

Data Science

What is AI?

- **AI Research: Three distinct points of view**
  - **AI as a Tool (what DFKI has done successfully for 30+ years)**
    - Enabling intelligent, supportive, and adaptive technology
    - Systems aware of their environment and reacting accordingly
    - **AI should be perceived and handled as a *tool*!**
    - **Humans decide and are responsible how these tools are used!**
  - **AI for Cognitive Sciences (we should do more of this)**
    - Understanding the human brain and cognition
    - **Helps to better understand how we humans ”work”**
  - **AI for Artificial Humans (we should do NONE of this)**
    - A non-goal for our and almost all other AI research
    - **Main cause of fear in the general public (sci-fi movies, etc.)**

AI: Structure and Terminology

Technologies

- AI
  - Reasoning
    - Rule-Based Systems
    - …
    - Statistical Reasoning
  - Learning (Machine Learning)
    - SVM
    - …
    - Deep Learning

Applied AI Communities

Vision | Language | Robotics | Interaction | Ethics | …

Application Areas

Automotive | Health | Industrie4.0 | Sciences | Public Sector | …

What is AI?

- **AI Technologies (very roughly)**
  - Logical Reasoning (rule-based systems, symbolic AI)
    - Semantics and various forms of logic/algorithms (proof theory)
    - First-Order, Higher-Order, Fuzzy, Modal, …
  - Statistical Reasoning
    - Deriving properties via probability densities (mostly designed manually)
    - Mainly Bayesian Reasoning
  - Machine Learning (data-driven approach)
    - Learning and predictions based on data
    - Decision trees, Clustering, SVMs, Reinforcement-Learning, …
  - Deep-Learning (DL, recursive pattern recognition)
    - Layered network of “neurons” with weighted connections between them
    - “Programming” via network structure and providing training data
    - Singe framework for “end-to-end” learning and optimization!
  - **Hybrid AI: Combining ML/DL with symbolic/statistical methods**

What is Deep Learning?

- **From Neural Nets to Deep Learning**
  - Neurons
    - Brain is a vast network of connected neurons each with simple characteristics
      - Human: ~86 GNeurons, <10k synapses each

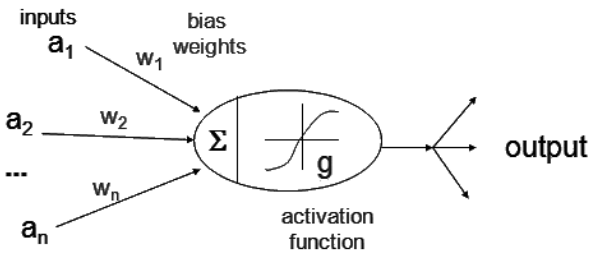

  - Early (Artificial) Neural Networks (NN)
    - A few layers of neurons (e.g. perceptron)
    - Back-propagation for learning weights (w_i)

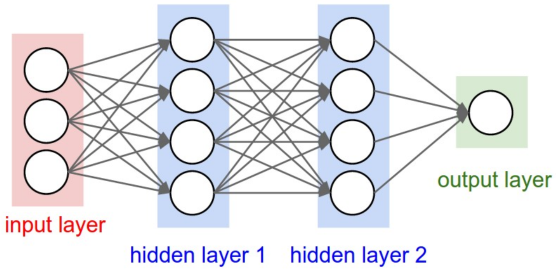

  - Deep Neural Networks (DNN)
    - Often hundreds of layers and more
    - Possible due to: Data, HW, Algorithms, …
    - Different structures (CNN, RNN, LSTM, …)

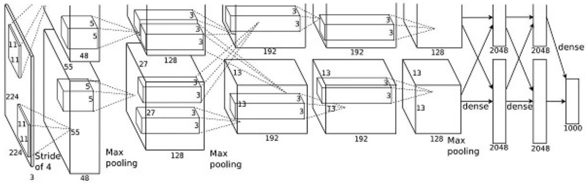

What is Deep Learning?

- **What DNN Actually Learn**

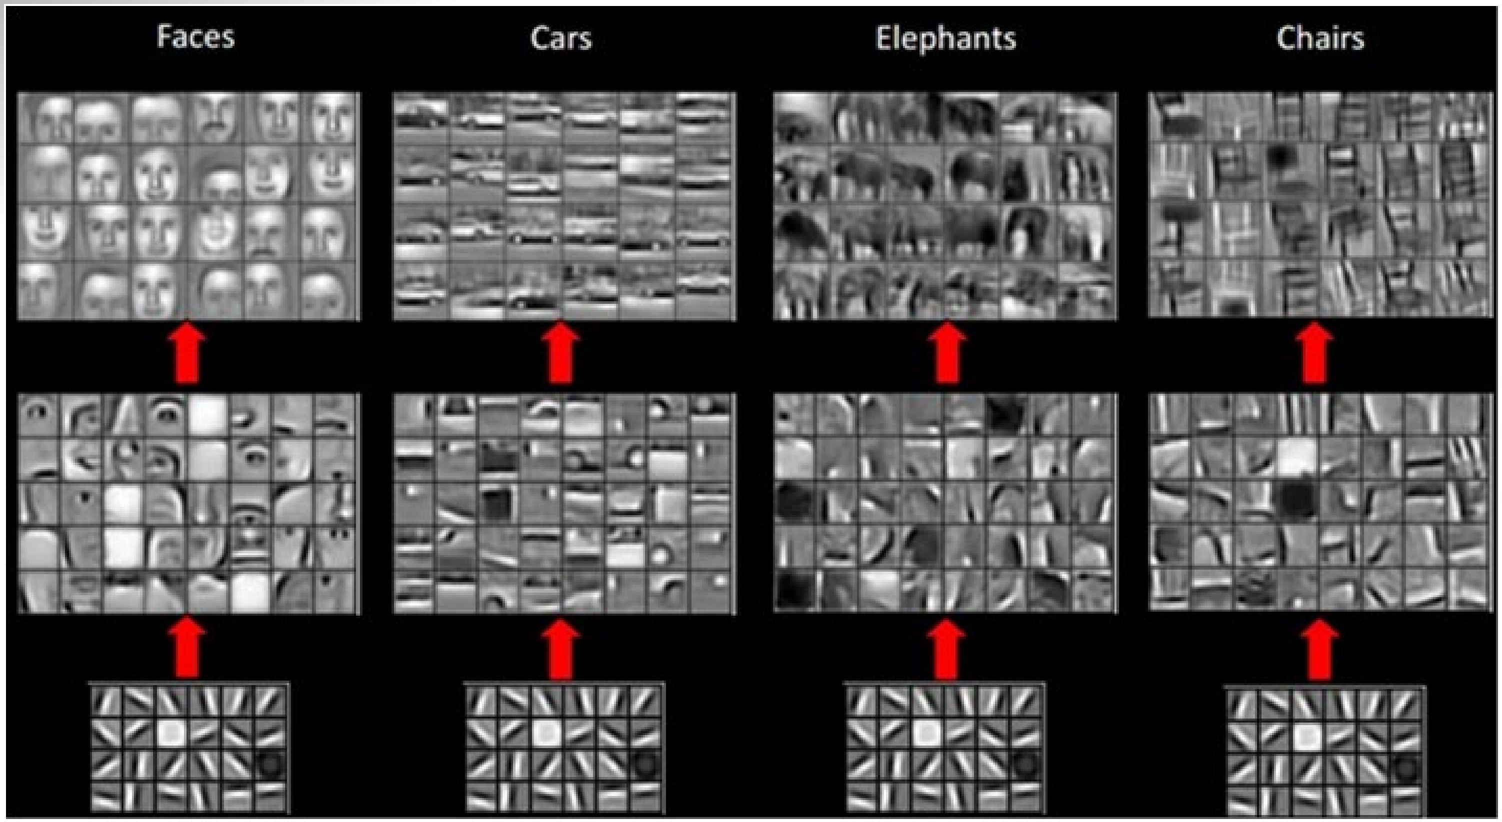

High Level

Medium Level

Lowest Level

© Honglak Lee at al.

Deep Learning: „New“ AI – With a Long Tradition

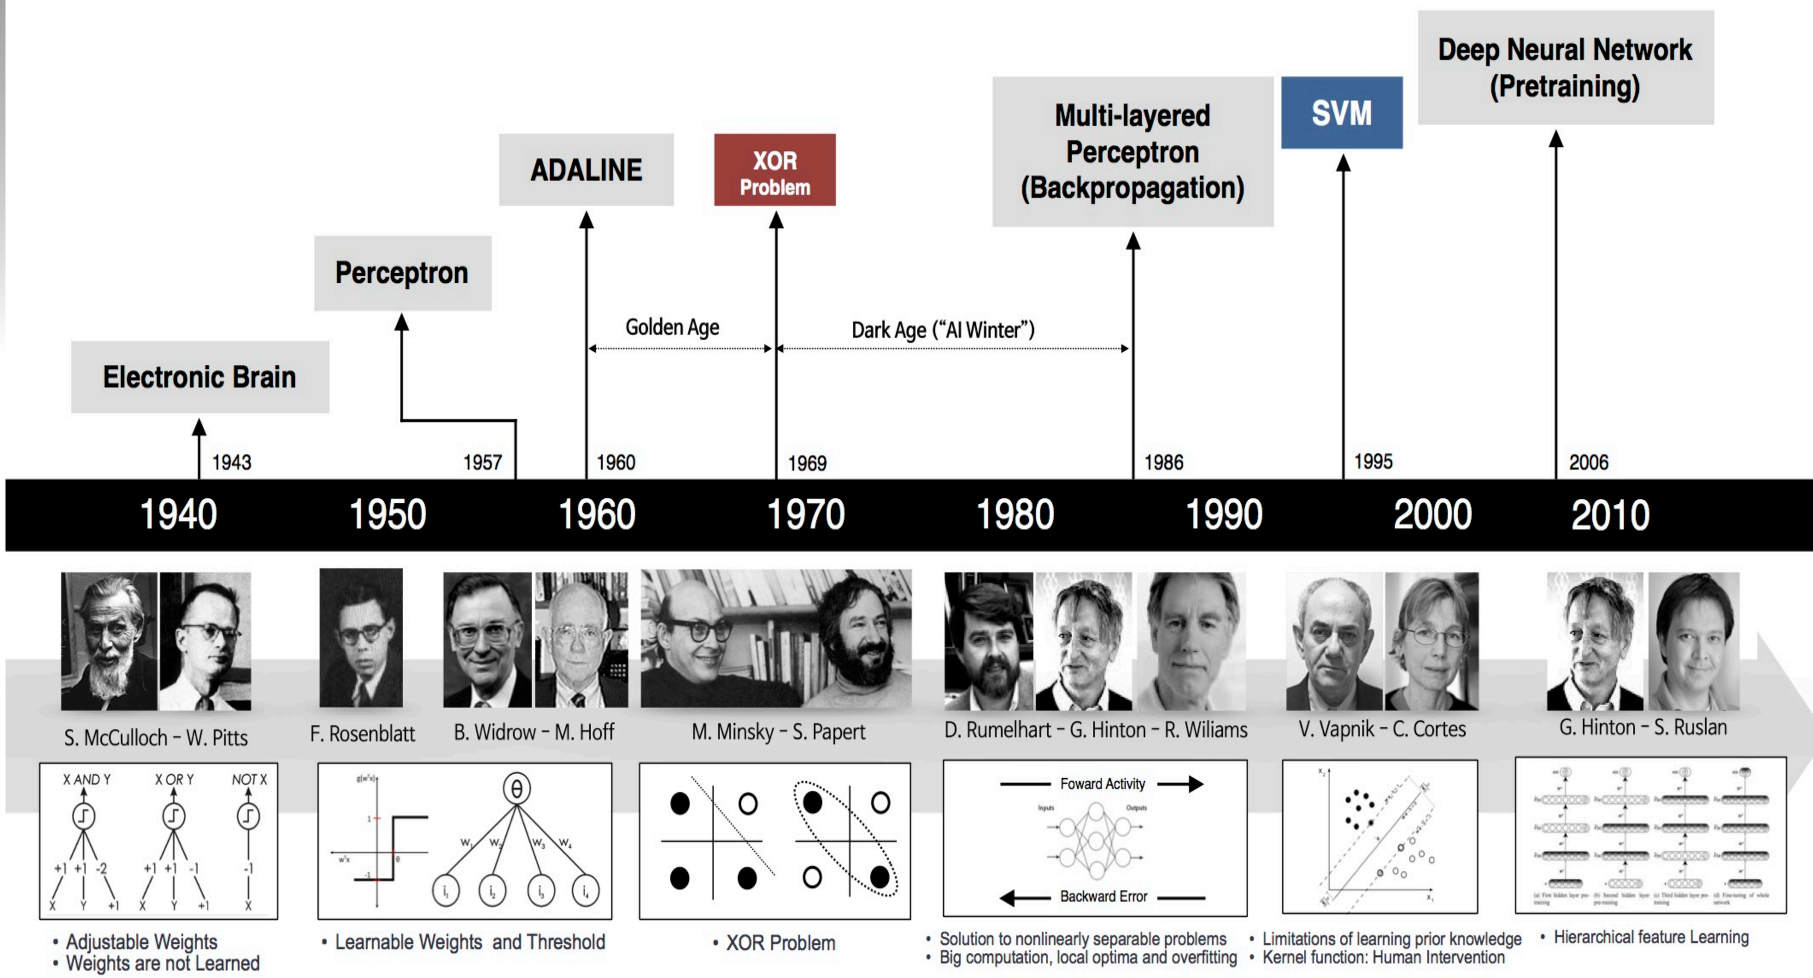

https://beamandrew.github.io/deeplearning/2017/02/23/deep_learning_101_part1.html

Other AI Success Stories:
(that hardly anyone knows about)

Pentium FDIV: The processor bug that shook the world

By Desire Athow October 30, 2014 Processors

20 years already

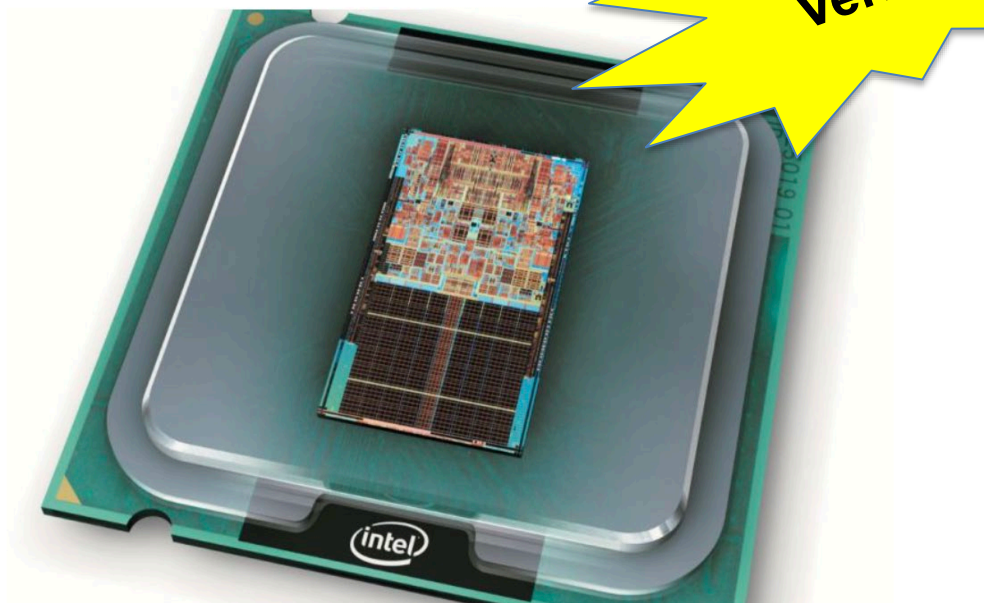

HW-Verification

1995: Intel writes off **US$ 475 Million** due to flawed microprocessor design

Other AI Success Stories: Knowledge Graphs

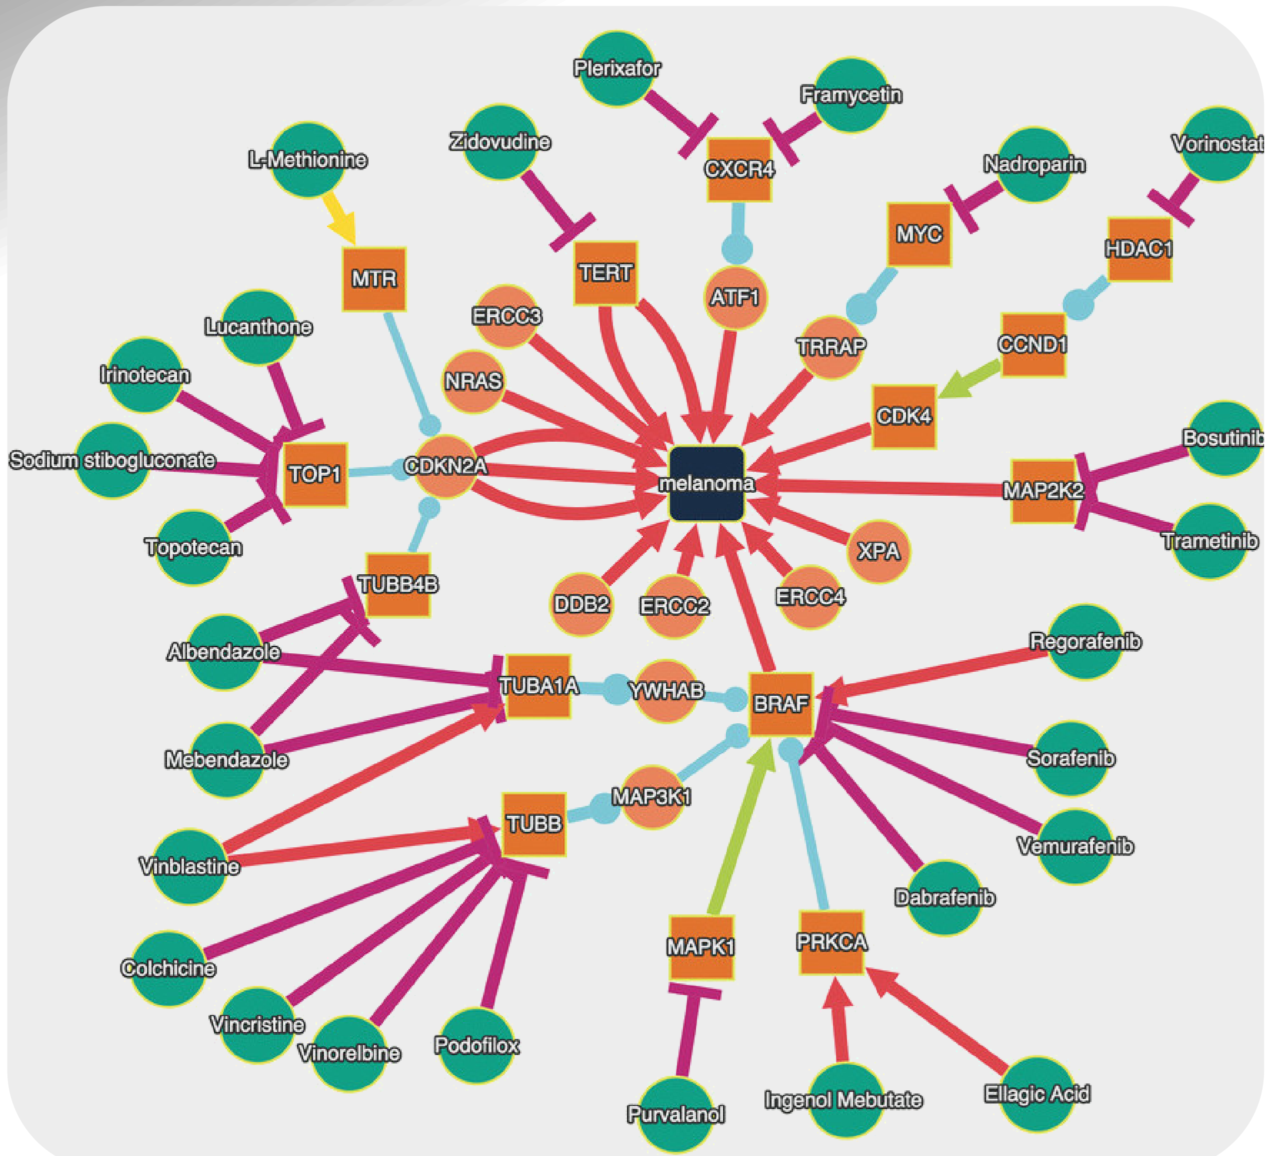

Knowledge Graphs: Up to 45-100 Billion Facts & Rules

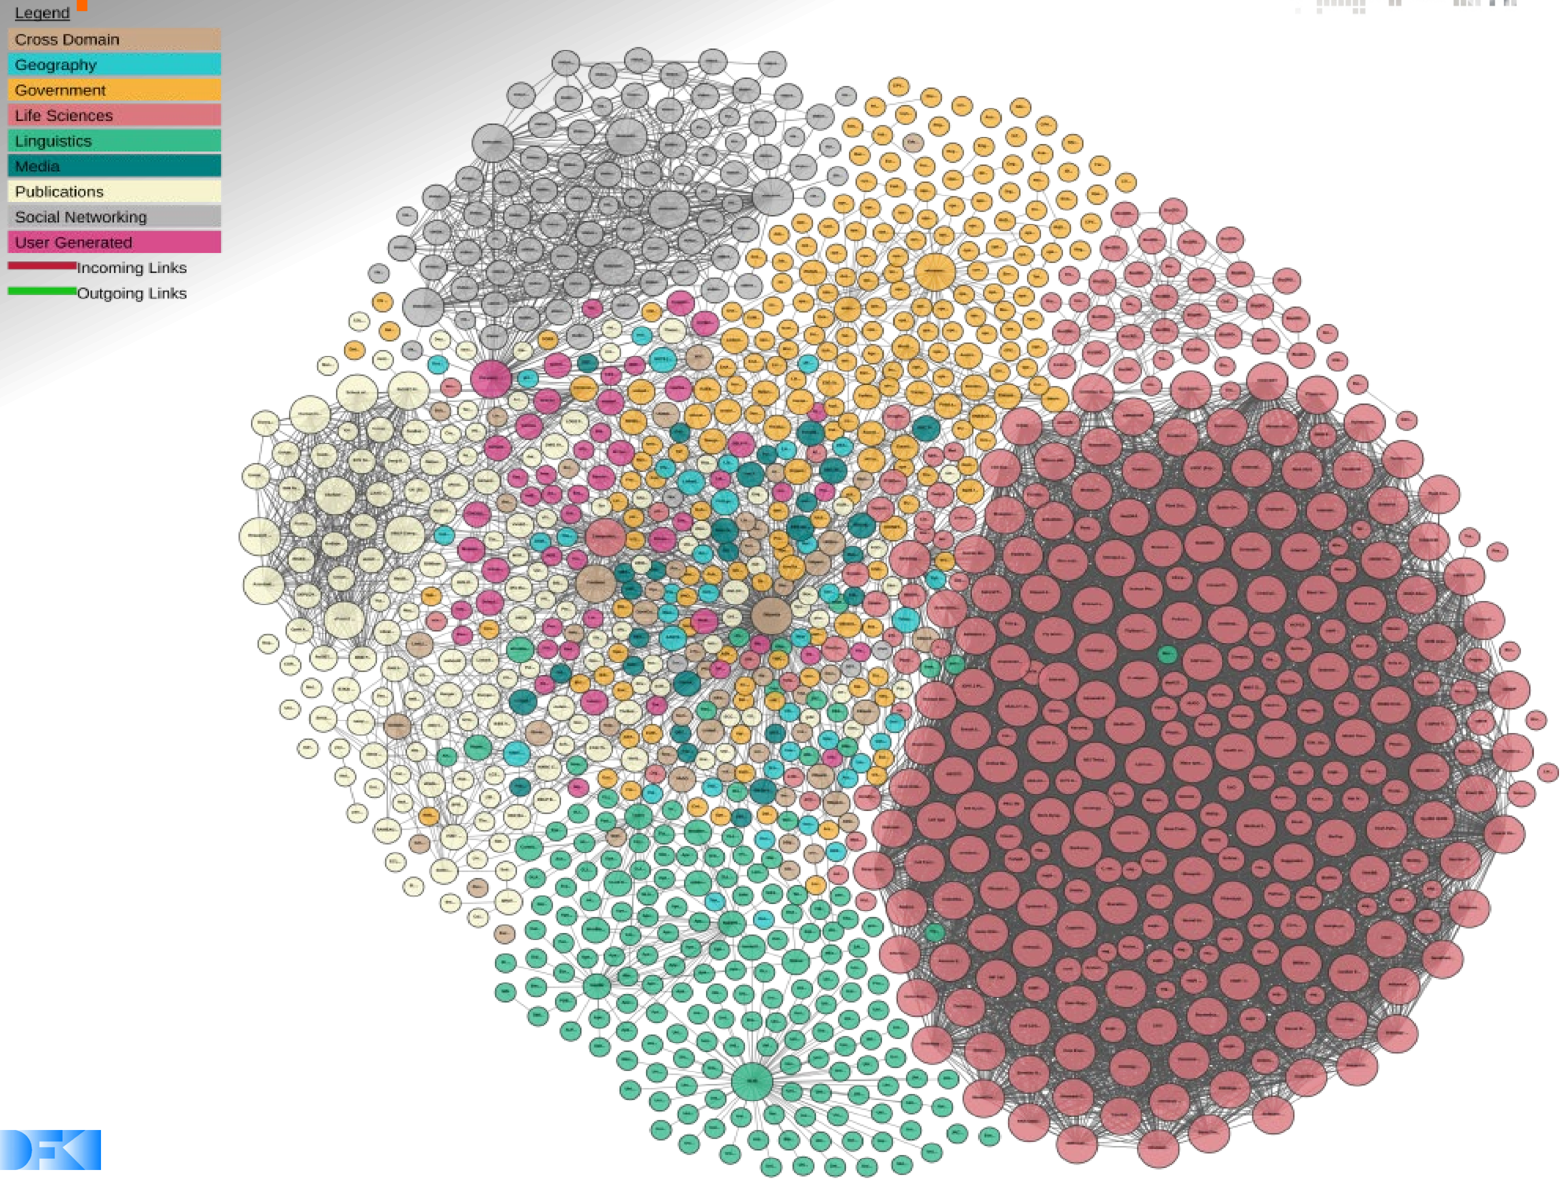

Other AI Success Stories: Search & Optimization

2008 - 2019

Size of Search Space:

Rules	Complete Search	Heuristic Search	Monte Carlo Search	Supervised Machine Learning	Reinforcement Machine Learning
10^3	10^{13}	10^{47}, 10^{20}	10^{170}		

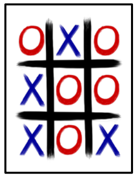
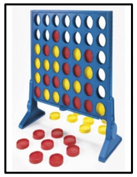



10^{80} = Number of Atoms in Universe

https://en.wikipedia.org/wiki/Game_complexity

Challenges: Functionality vs. Robustness

- AI/DL is highly capable already …

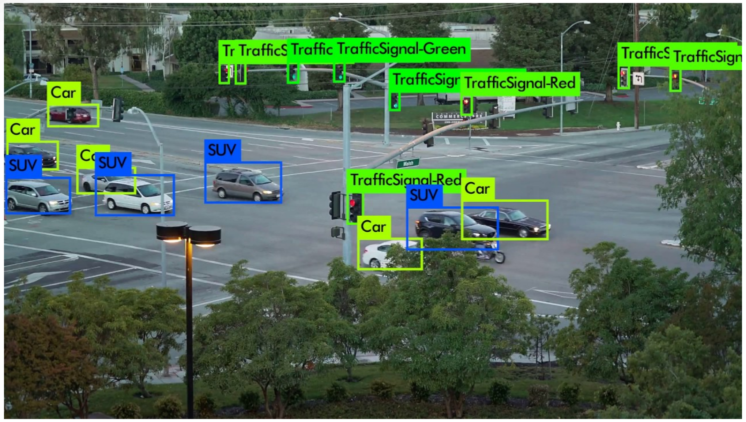

- … but we often cannot guarantee basic functionality

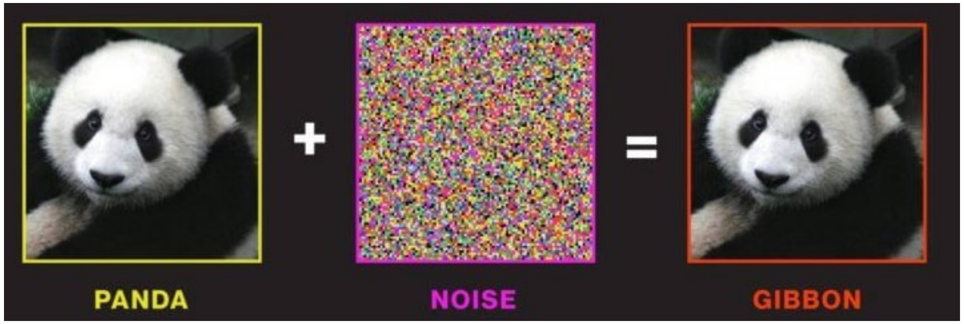

Key Challenges

- **Validation, Verification, Certification**
  - Lack of formal methods for DNN only leaves statistical ***validation*** via systematic testing (e.g. this talk)
  - To be combined with higher-level ***verification***
- **Transparency and Explanation**
  - Ability to explain behavior of AI systems to humans
  - “Black box” approach is not a bad thing (as many suggest)
- **Modularity**
  - From end-to-end learning to composable building blocks
- **Integration of Learning and Reasoning (Hybrid AI)**
  - Rich ecosystem of existing techniques, often with formal methods
- **Limits of the technology**
  - Understand limits from physics, computing, economics & theory
- **Ethical and Legal Aspects !!**

Digital Reality:
Using AI to Optimize and Certify AI
(using autonomous driving as an example)

Trustworthy AI: Using AI to Optimize & Certify AI

AI functionality is not enough –
need ability to *certify* its capabilities –
according to well-defined standards

Autonomous Systems: The Problem

- **Our World is extremely complex**
  - Geometry, Appearance, Motion, Weather, Environment, …
- **Systems must make accurate and reliable decisions**
  - Especially in ***Critical Situations***
  - Increasingly making use of (deep) machine learning
- **Learning of critical situations is essentially impossible**
  - Often little (good) data even for “normal” situations
  - Critical situations rarely happen in reality – per definition!
  - Extremely high-dimensional models

➔ **Goal: Scalable Learning from *synthetic* input data**
  - Continuous benchmarking & validation (“Virtual Crash-Test“)

Reality

- **Training and Validation in Reality**
  - E.g. driving *millions of miles* to gather data
  - Difficult, costly, and non-scalable
  - **Please: Never again use "millions of miles" argument!**
    - **No need for *more & general* data – but *better & specific* data**
    - **And better way to obtain this data (e.g. using synthetic data)**

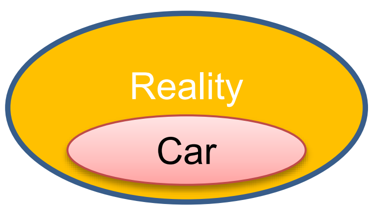

Digital Reality

- **Training and Validation in the *Digital Reality***
  - Arbitrarily scalable (given the right platform)
  - But: Where to get the models and the training data from?

Digital Reality: AI to Safeguard AI

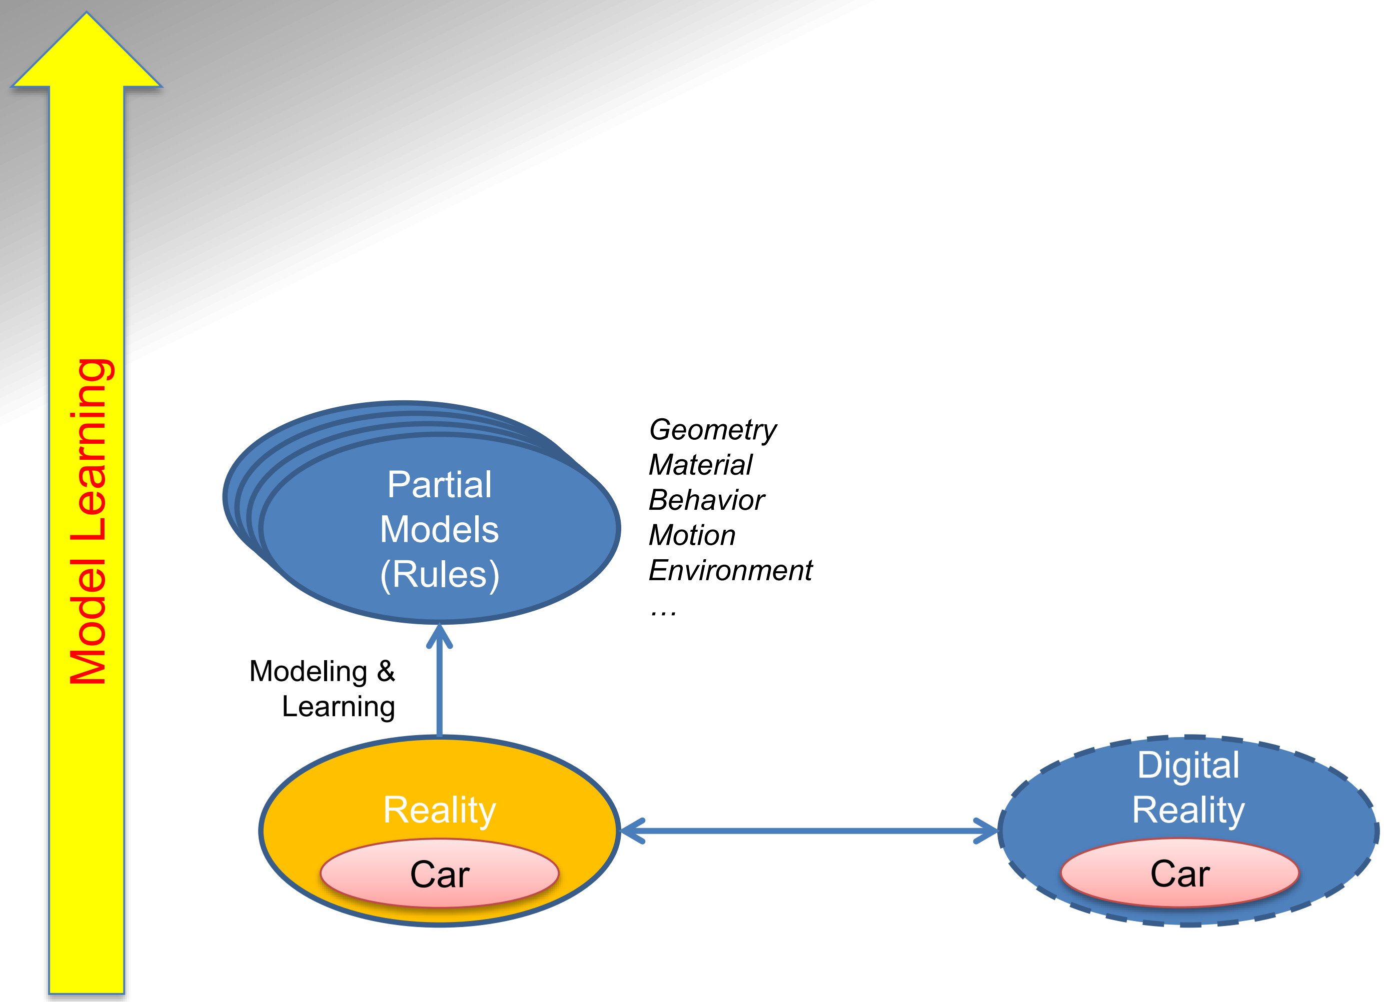

Digital Reality: AI to Safeguard AI

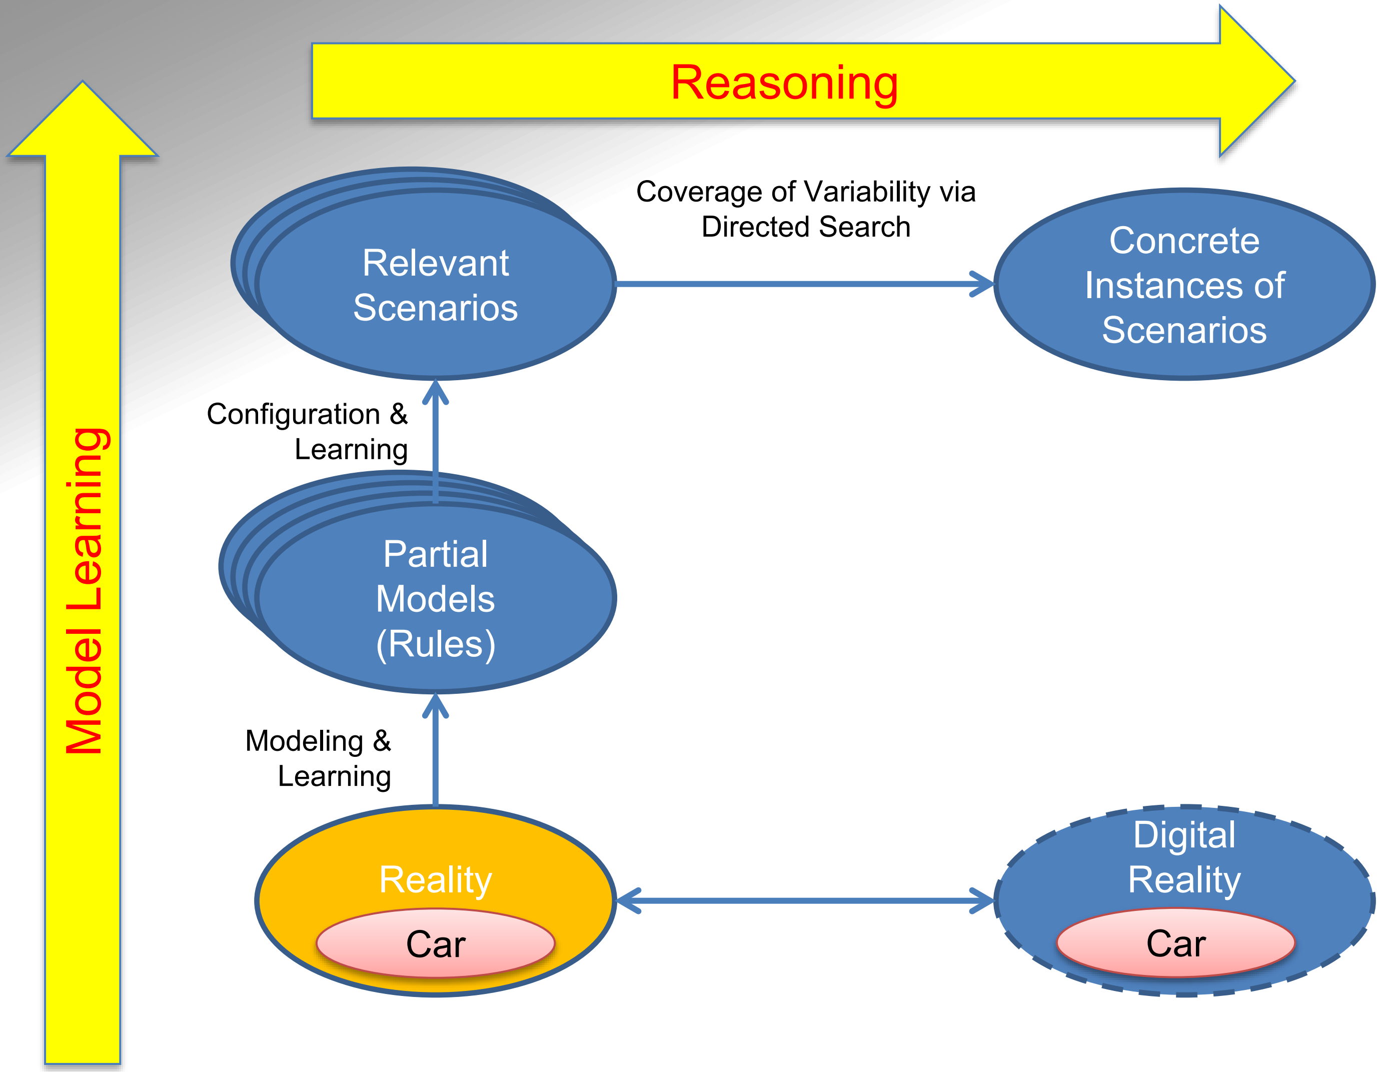

Digital Reality: AI to Safeguard AI

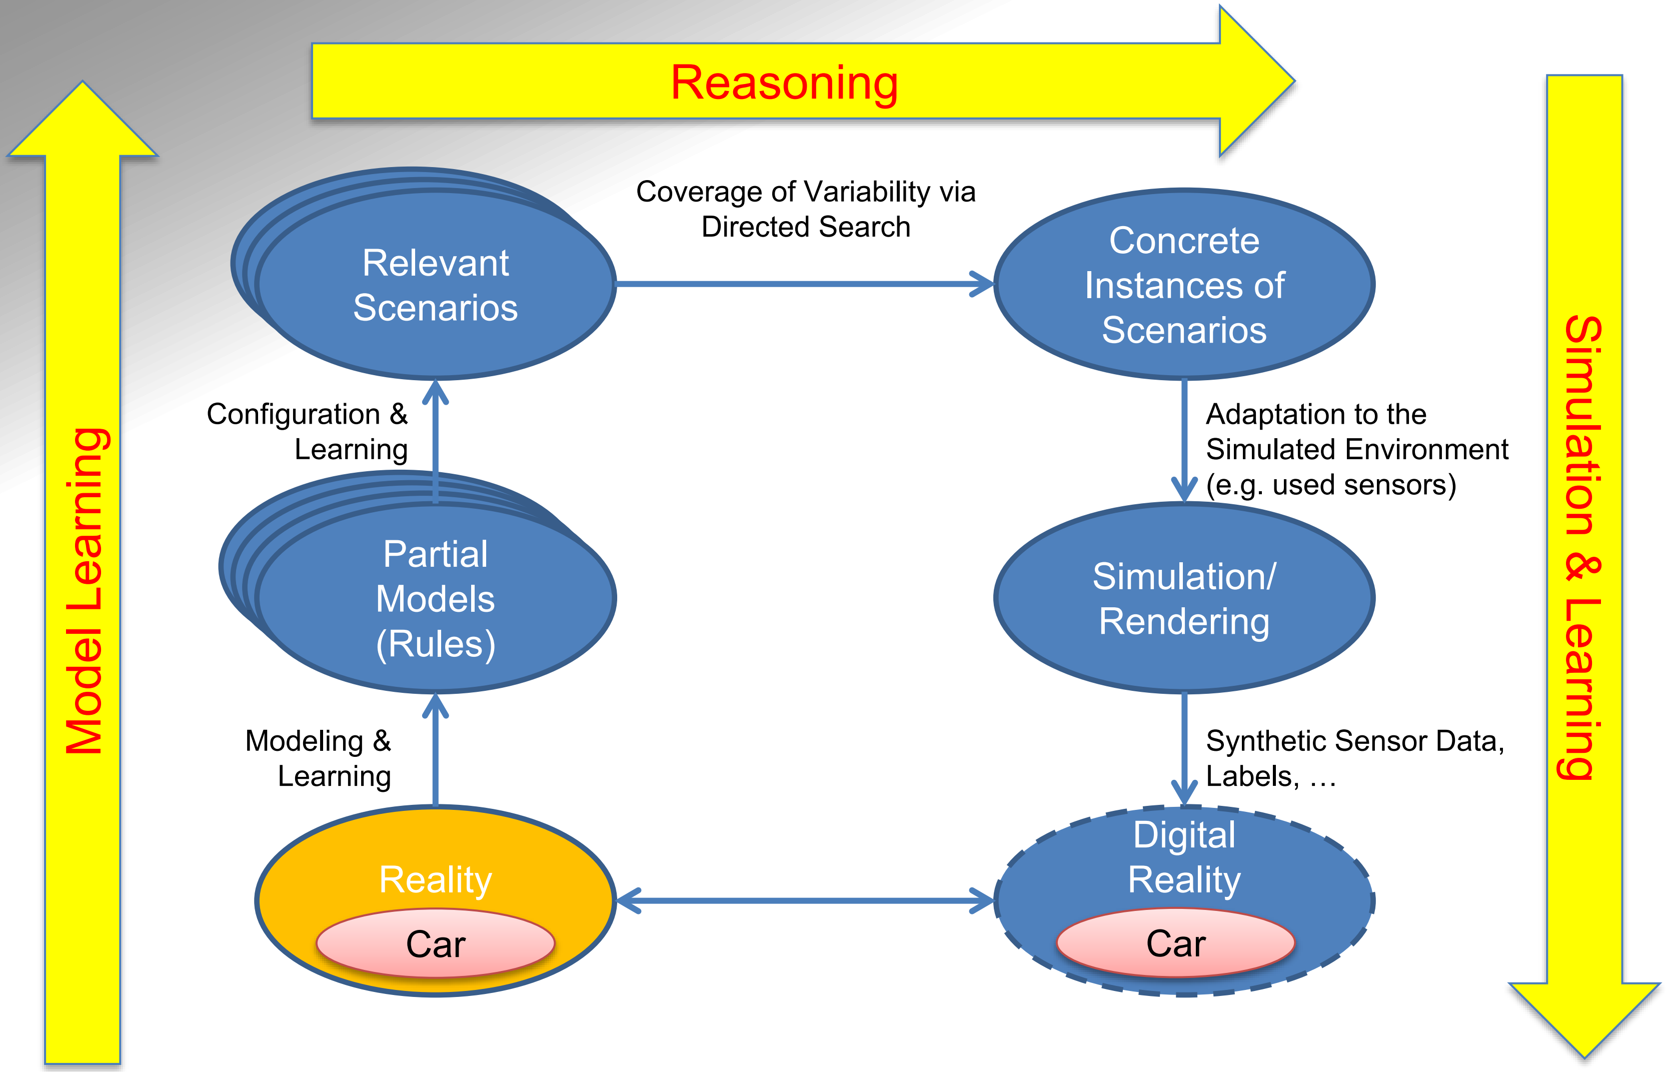

Digital Reality: AI to Safeguard AI

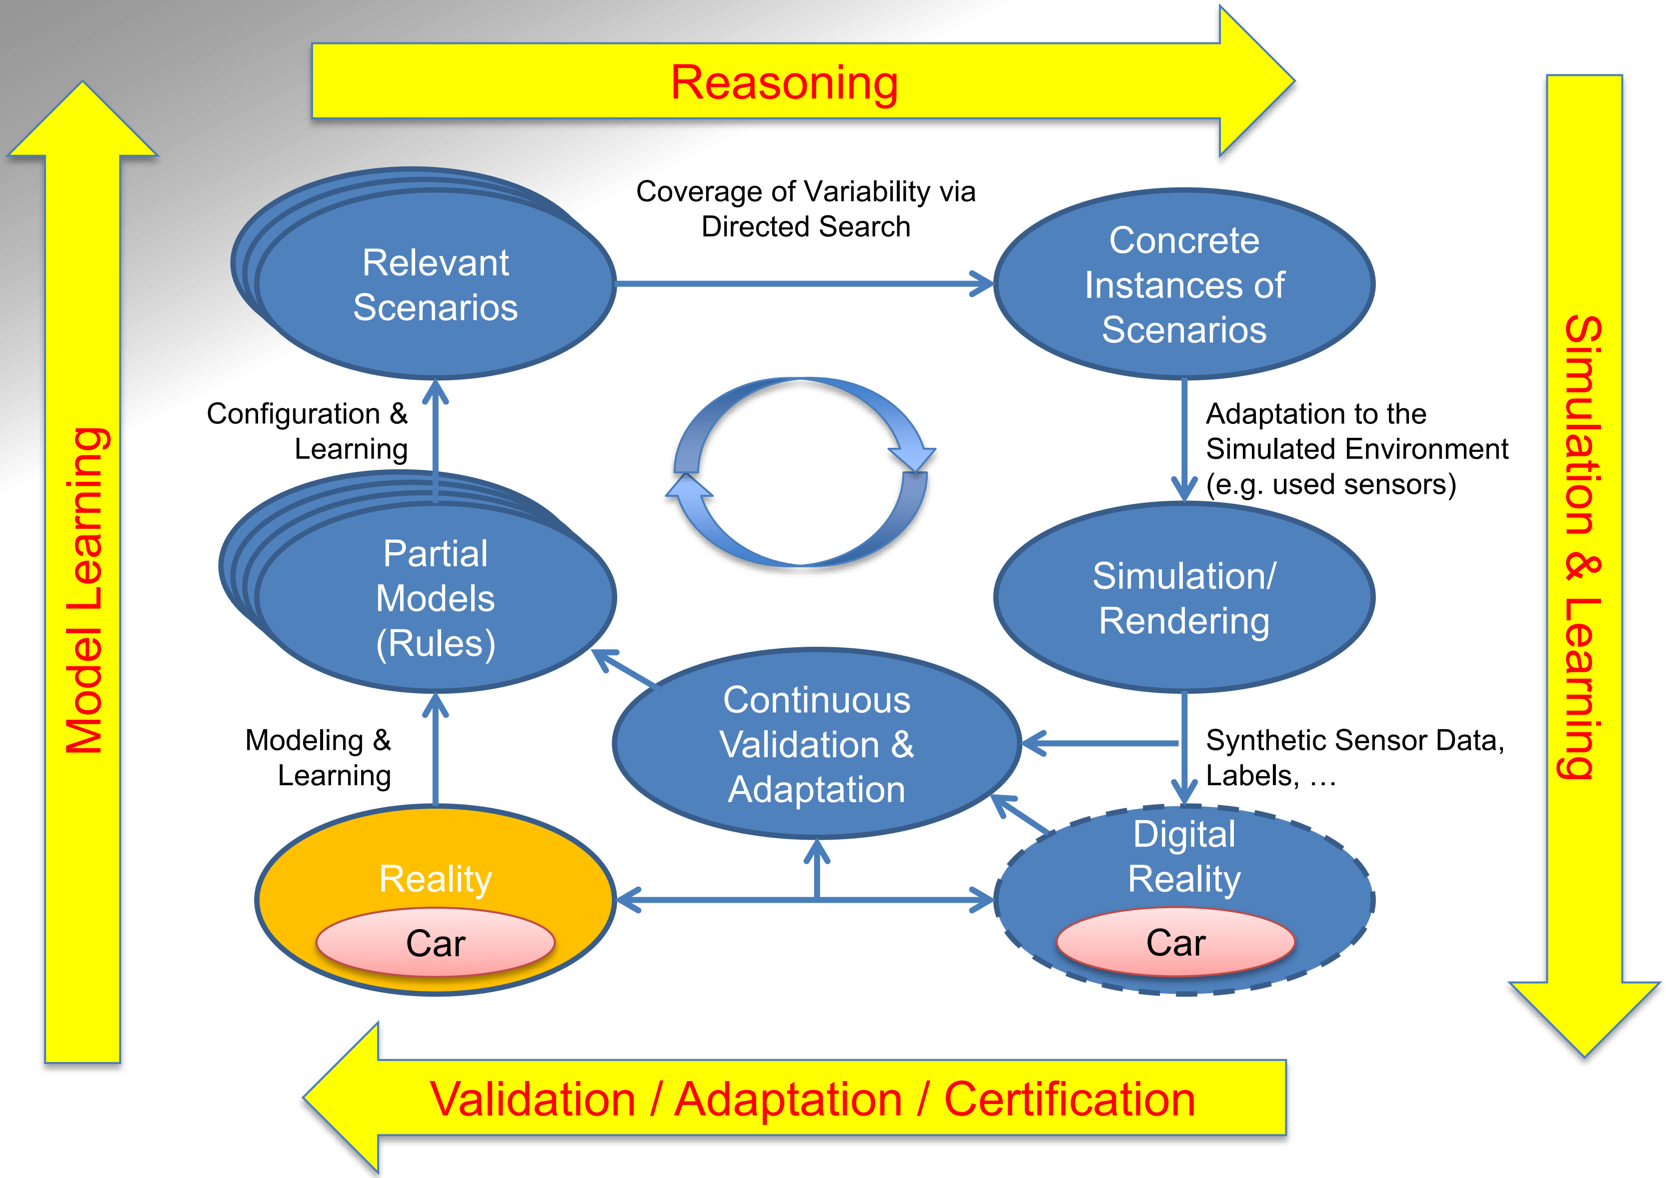

Digital Reality: AI to Safeguard AI

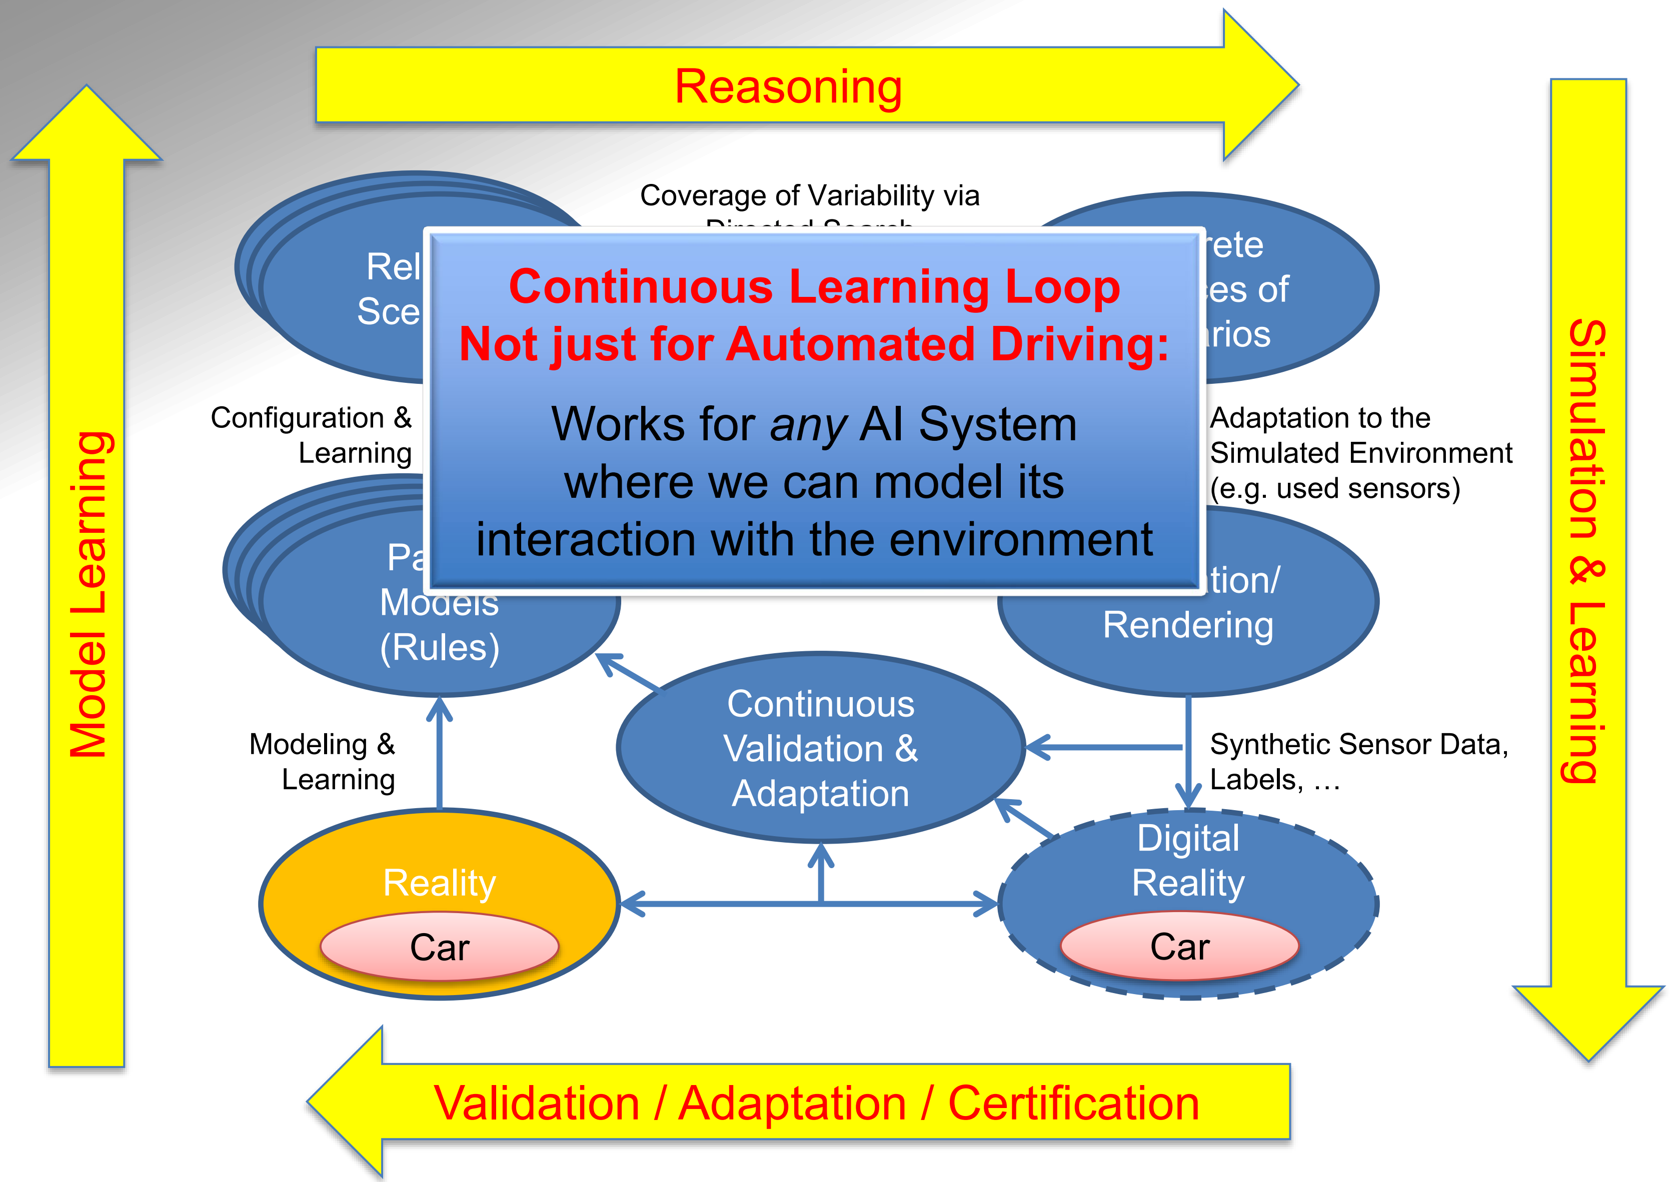

Challenge: Better Models of the World (e.g. Pedestrians)

- **Long history in motion research (>40 years)**
  - E.g. Gunnar Johansson's Point Light Walkers (1974)
  - Significant interdisciplinary research (e.g. psychology)
- **Humans can easily discriminate different styles**
  - E.g. gender, age, weight, mood, ...
  - Based on minimal information
- **Can we teach machines the same?**
  - Detect if pedestrian will cross the street
  - Parameterized motion model & style transfer
  - Predictive models & physical limits

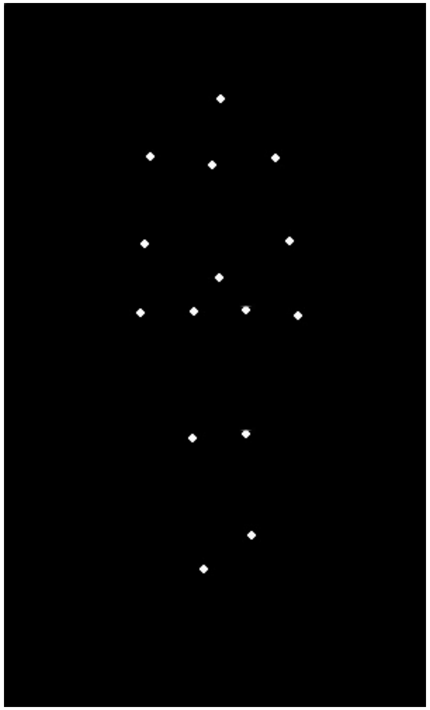

Challenge: Pedestrian Motion

- **Characterizing Pedestrian Motion**
  - Clear motion differences when crossing the street

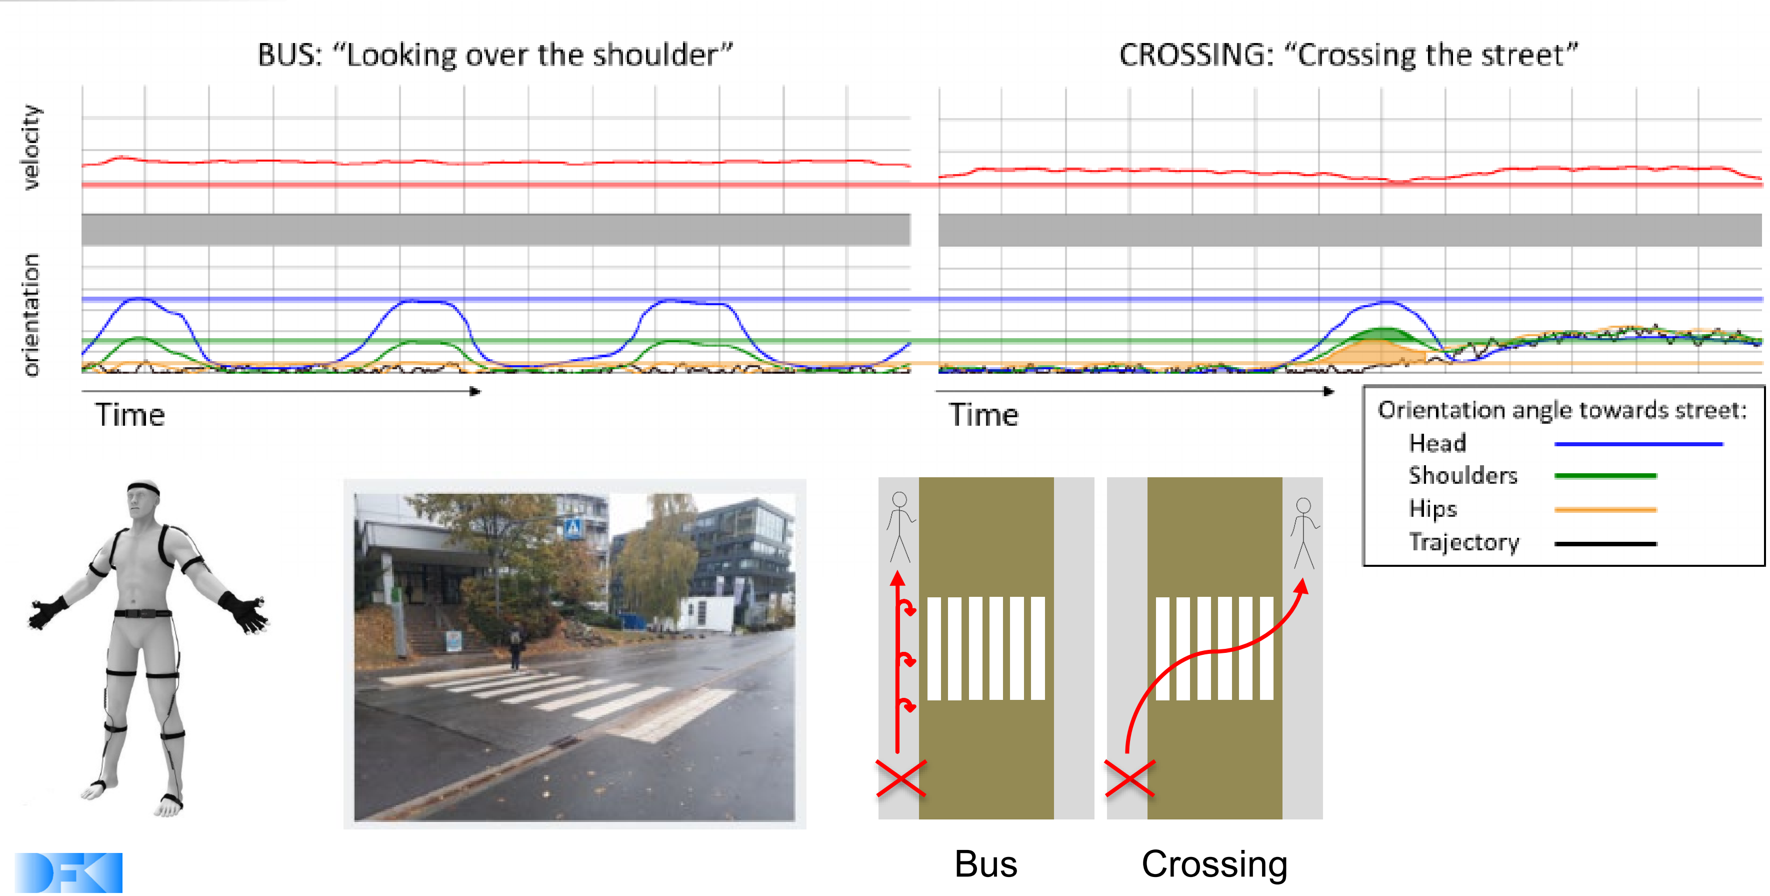

Challenge: Better Simulation (e.g. Radar Rendering)

- **Key Differences**
  - Longer wavelength (by 10^4): Geometric optics (rays) not sufficient
  - Need for *some* wave optics
    - Phase information and interference due to coherent radiation
    - Doppler effects allow for velocity information
    - Diffraction at rough surfaces and edges
    - Polarization as radar signals are strongly polarized
  - Very different goals
    - Optical: Focus on *diffuse* effects (+ some highlights, reflections, etc.)
    - Radar: Focus on *specular* transport only (i.e. caustic paths)
- **Recent Work on Caustics [Grittmann et al., EGSR'18]**
  - Identifying "useful" specular paths (using VCM)
  - Guides samples to *visible specular effects (e.g. indirect radar echos)*

Radar Simulation

- **Why Simulate Radar**
  - Widely used in mobile applications (assistance & autonomy)
    - Measures direction, distance, and relative velocity of objects
  - Quickly improving radar devices, e.g.:
    - Imaging radar with much improved direction resolution
  - Need for early fusion between modalities
    - Current devices only return preprocessed array of object data
    - Want to combine low-level data from different devices
    - Towards higher reliability and sensitivity
  - Simulation research for radar and light nearly disjoint
  - Radar could greatly benefit from CG (and maybe vice-versa)
    - CG methods are use little in radar simulation (yet).

“Geometrical Optics (GO)”

- **Essentially the same as ray tracing**
  - Typically add phase information as ray propagate
- **But cannot account for many of the wave effects**
  - E.g. diffraction

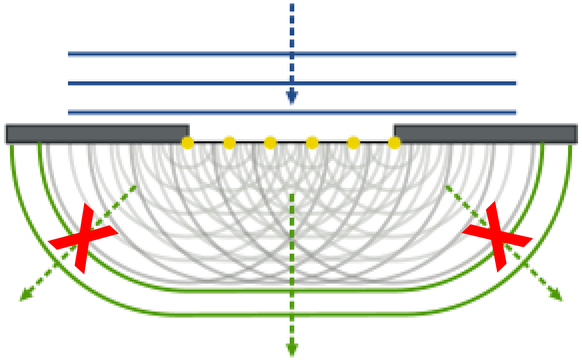

“Physical Optics” (PO)

- Popular Approach for radar simulation
- Electromagnetic wave induce surface currents
  - $\vec{J} = 2\vec{n} \times \overrightarrow{H}_{inc}$
- Treat surface as Hertzian dipoles that re-radiate
  - $\overrightarrow{H}_{ref} = -jk_0 \frac{1}{4\pi r}(\vec{r} \times \vec{J})$
- Direct use of Huygens-Fresnel Principle:

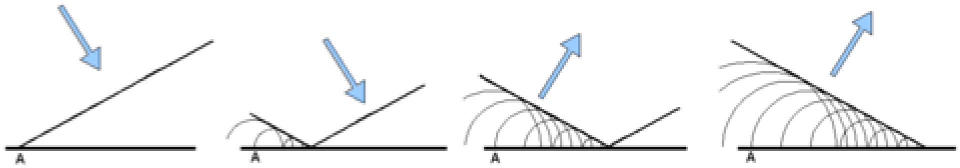

“Physical Optics” using MC

- **Sample starting point & direction on transmit antenna**
  - Equivalent to adjoint particle tracing, i.e. light tracing
- **Intersect ray with scene**
- **Perform next event estimation (NEE)**
  - Connect path to a random point on some receive antenna
- **Perform BSDF sampling using physical optics model**
  - Need to correct cosine & inverse square law
- **Stop using Russian Roulette**
- **Ideally: Use guiding to locate relevant paths**

Hybrid Approaches

- **Each bounce scaled by k_0 for physical optics**
  - For 76 GHz radar: $k_0 \approx 1{,}600$
  - Nearly all energy cancels out through interference
  - Causes problems for Monte Carlo integration
- **Mitigation: Combine both approaches (called GO/PO)**
  - Indirect bounces are handled by GO
  - Only final bounce computed using PO
- **Results become biased**
  - Do not accounts for diffraction effects along paths

Experimental Setup

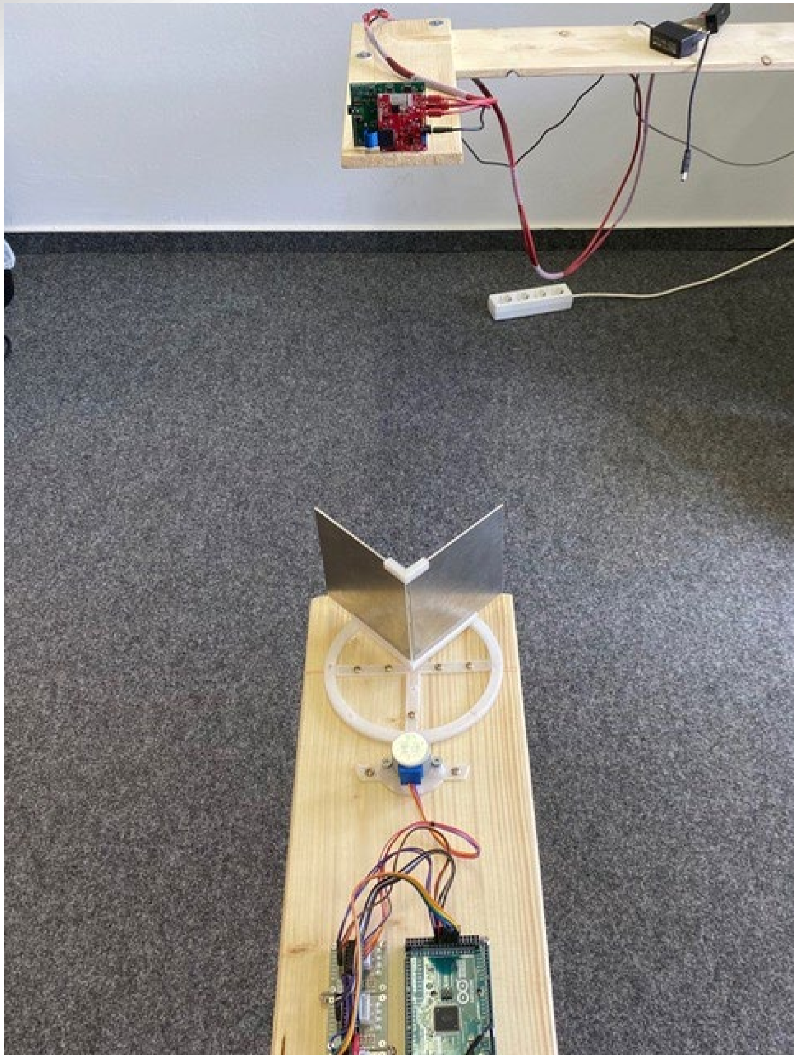

Results

- **Simulation of a concave reflector**
  - Including multiple bounces within reflector

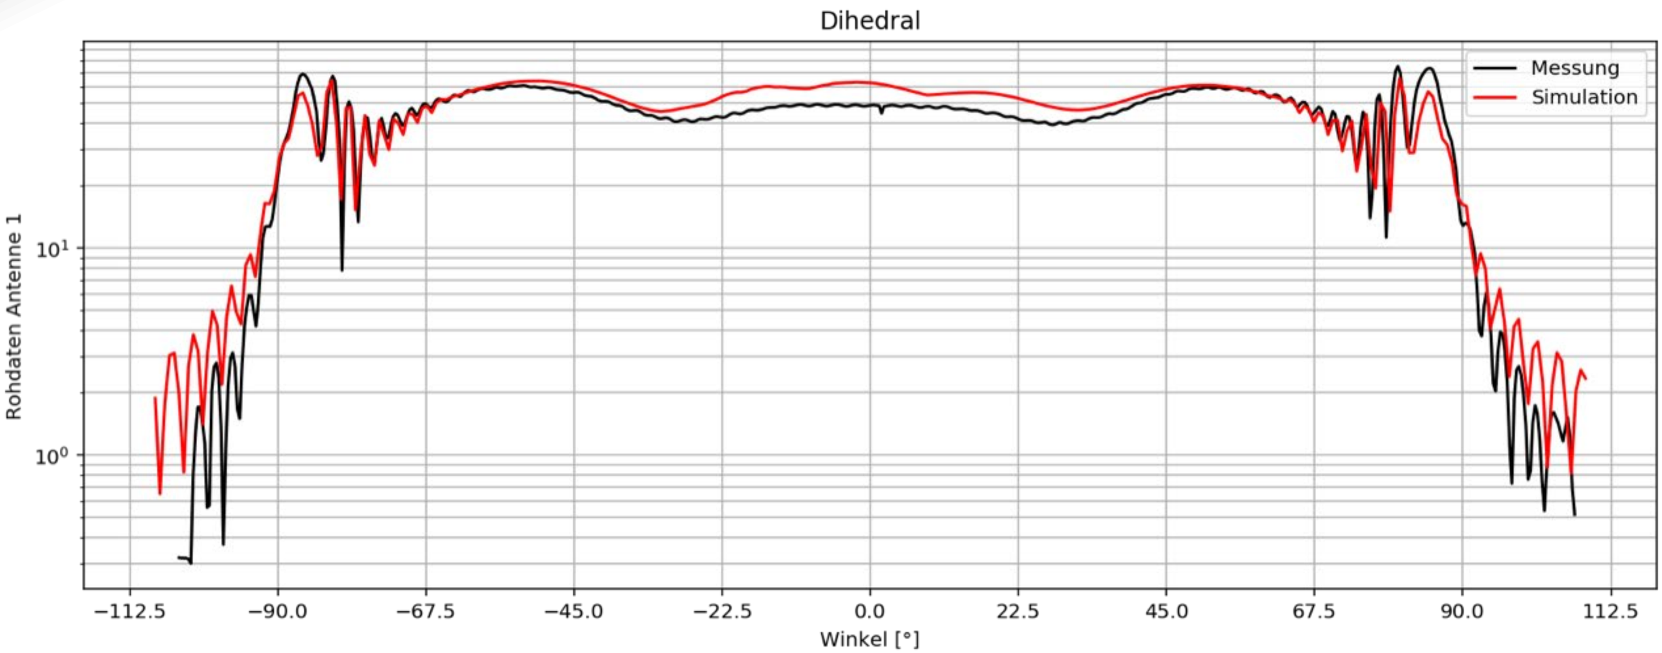

Challenge: Do we Need a Better Basis for our Simulation?

- **In the past: Two big markets, focused on nice images**
  - *Gaming*: Very nice images (at 60+ Hz)
    - Must compromise realism for frame rate
  - *Film & Marketing*: Even nicer images (at hours per image)
    - Will compromise realism for the story and artistic expression
  - Both are being used for simulations for Autonomous Driving
- **But: Strong need for *correct* images**
  - Lidar, radar, multi-spectral, polarization, measured materials, …
  - Need for "error bar per pixel" & validation
  - Existing engines unlikely to adapt to these fundamental changes
- **Towards "Predictive Rendering" engine**
  - Focused on physical accuracy ("sensor realistic") & high throughput
  - Based on latest graphics research results (and GPU-HW)

Digital Reality for Machine Vision and Image Segmentation

Digital Reality for Machine Vision & Segmentation

- **Applied *Digital Reality* in several domains**
  - Focus on difficult cases with little data
    - Optical scanning of faulty parts (cracks in chips)
    - Grain-boundary segmentation
    - One-dimensional signal classification
- **Key Results**
  - 25-33% decreased mean error rate by adding synthetic data
  - Ability to train models purely on synthetic data

➔ Digital Reality approach can have significant impact on industrial applications

Problem Description

- **Optical screening for quality control in manufacturing**
  - Detecting cracks in chip packaging
  - Training a DNN requires large amounts of data
  - Classification problem (“faulty/with crack” vs. OK)

- **Problem**
  - Only 113 real images were available (768×576)
  - Extracted 510 tiles: 255 with crack/255 OK (128x128)
    - Sets: 357 for training, 89 for validation, 64 for testing (each balanced)
- **Training Approach with Deep Neural Network**
  - ResNet with 50 layers (pre-trained on ImageNet)
    - First trained only custom “head”, then entire network (4/10 epochs)

Techniques: Capturing Material Properties

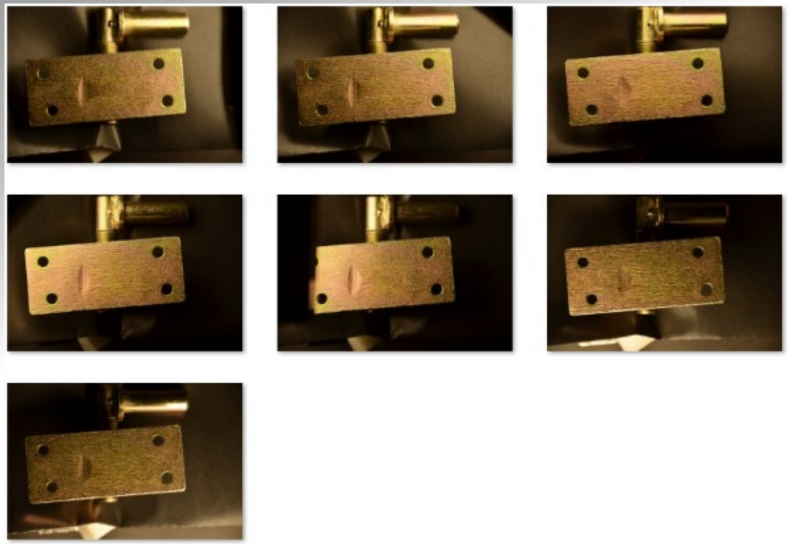

Normal Map

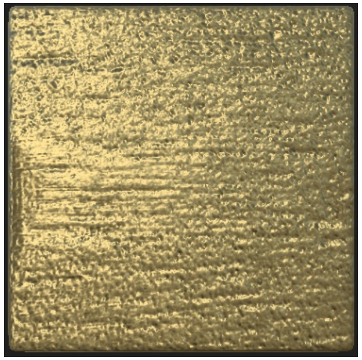

Height Map

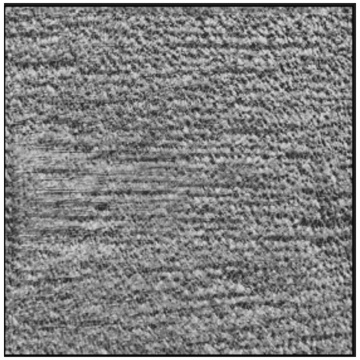

Techniques: Exemplar-Based Inpainting

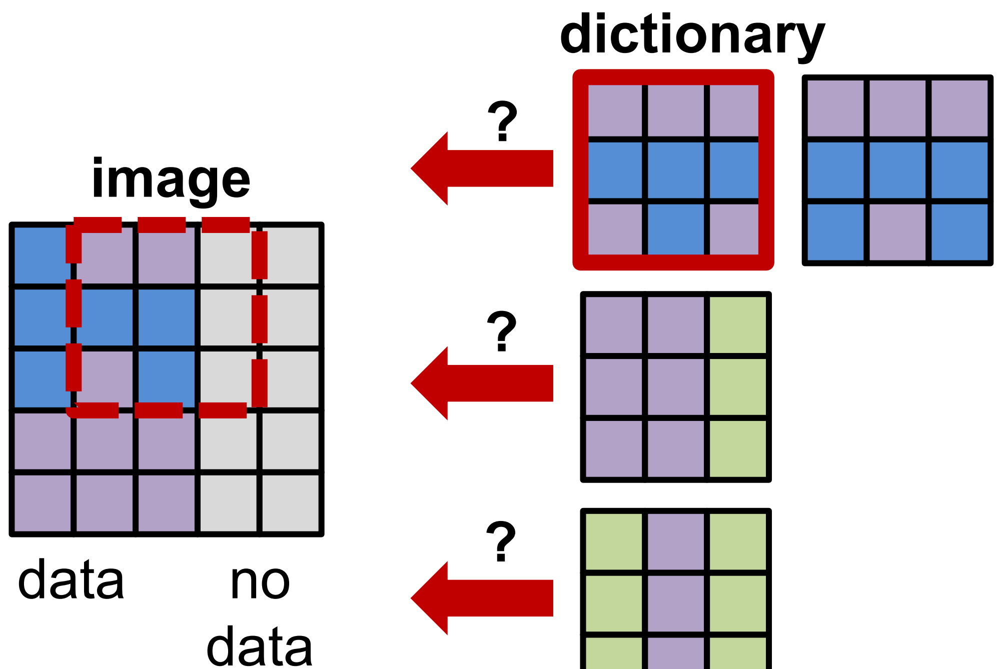

“Region Filling and Object Removal by Exemplar-Based Image Inpainting”, Criminisi et al.

Techniques: Exemplar-Based Inpainting

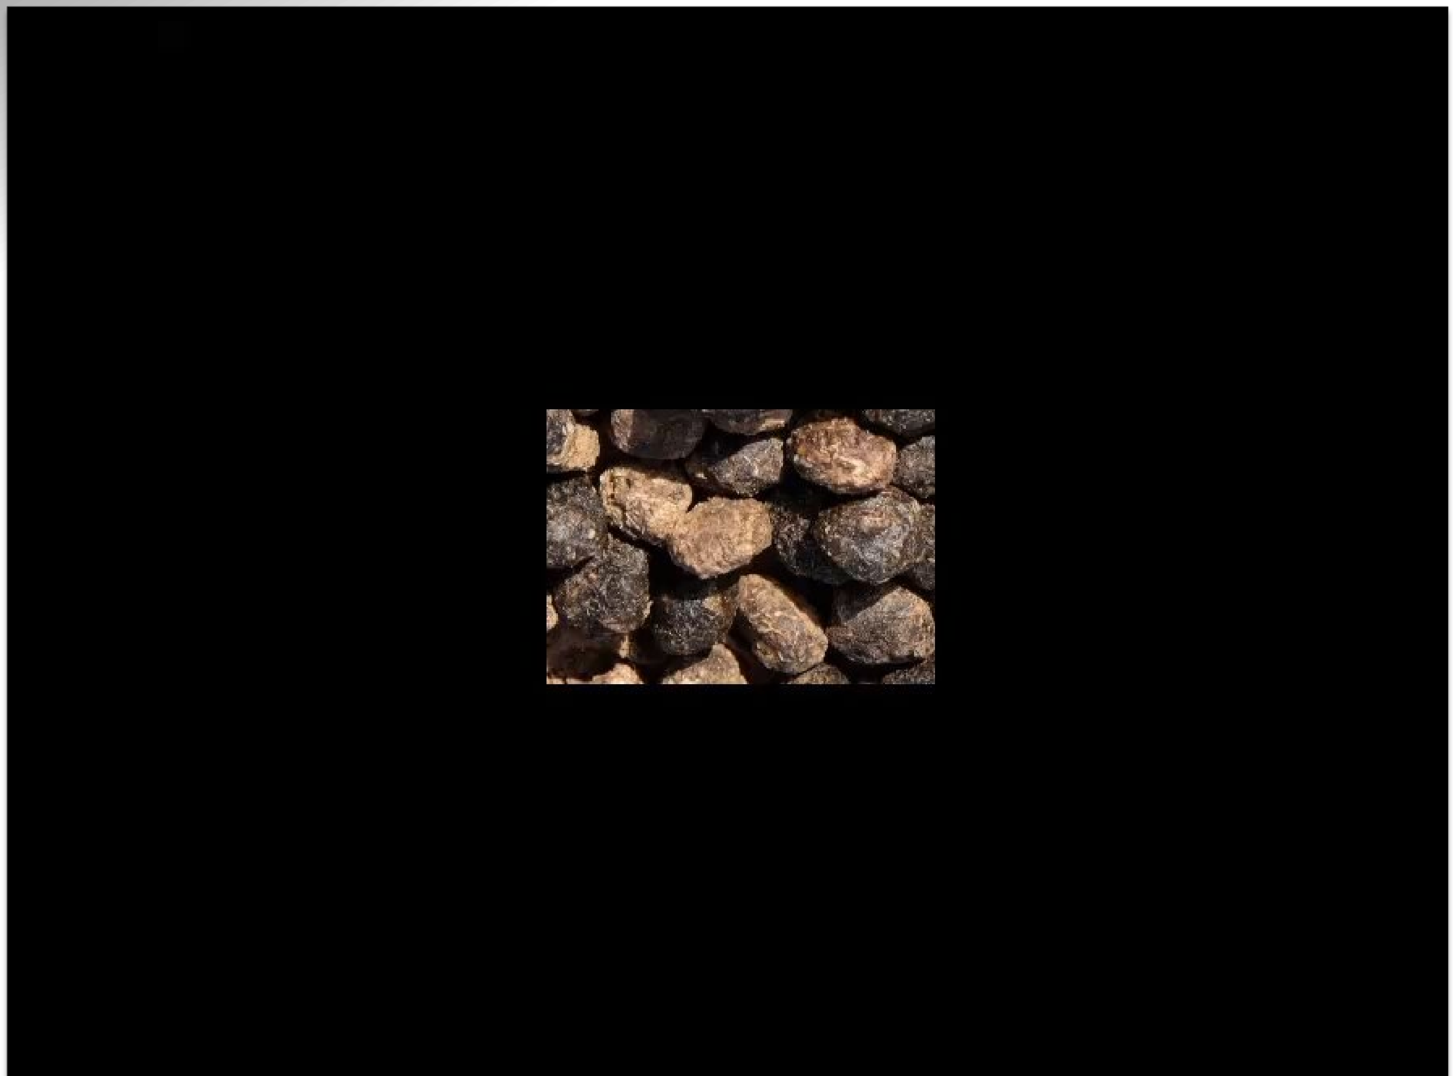

Generating Synthetic Data

- **Parametrically generating synthetic textures**
  - Extraction of clean texture patches (as dictionary)
  - Exemplar-based in-painting for generating synthetic textures
- **Parametrically generating synthetic cracks**
  - Measured statistical distribution of intensities, widths, segment lengths, angles at bends
  - Used this parametric model to generate synthetic cracks

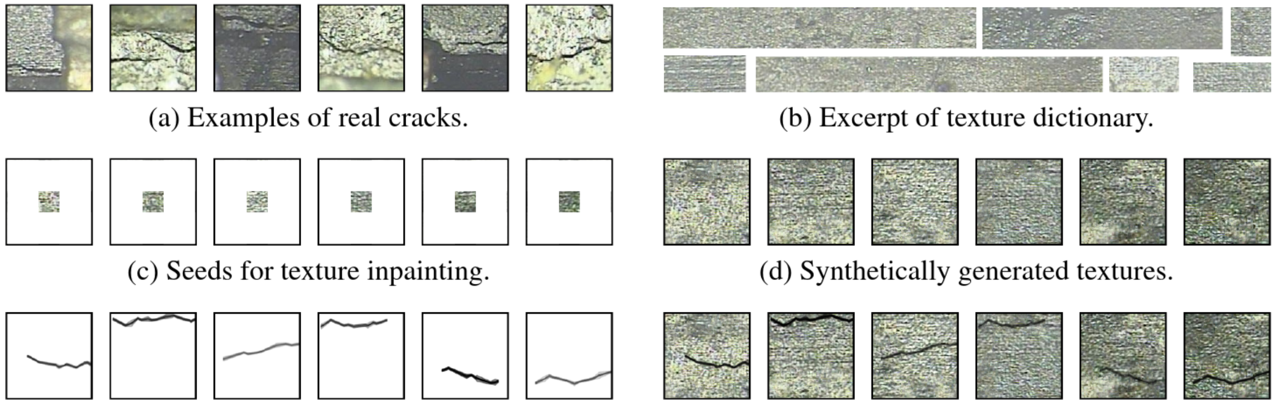

(a) Examples of real cracks. (b) Excerpt of texture dictionary. (c) Seeds for texture inpainting. (d) Synthetically generated textures. (e) Synthetically generated cracks. (f) Synthetically generated training data examples.

Details of Data Generation

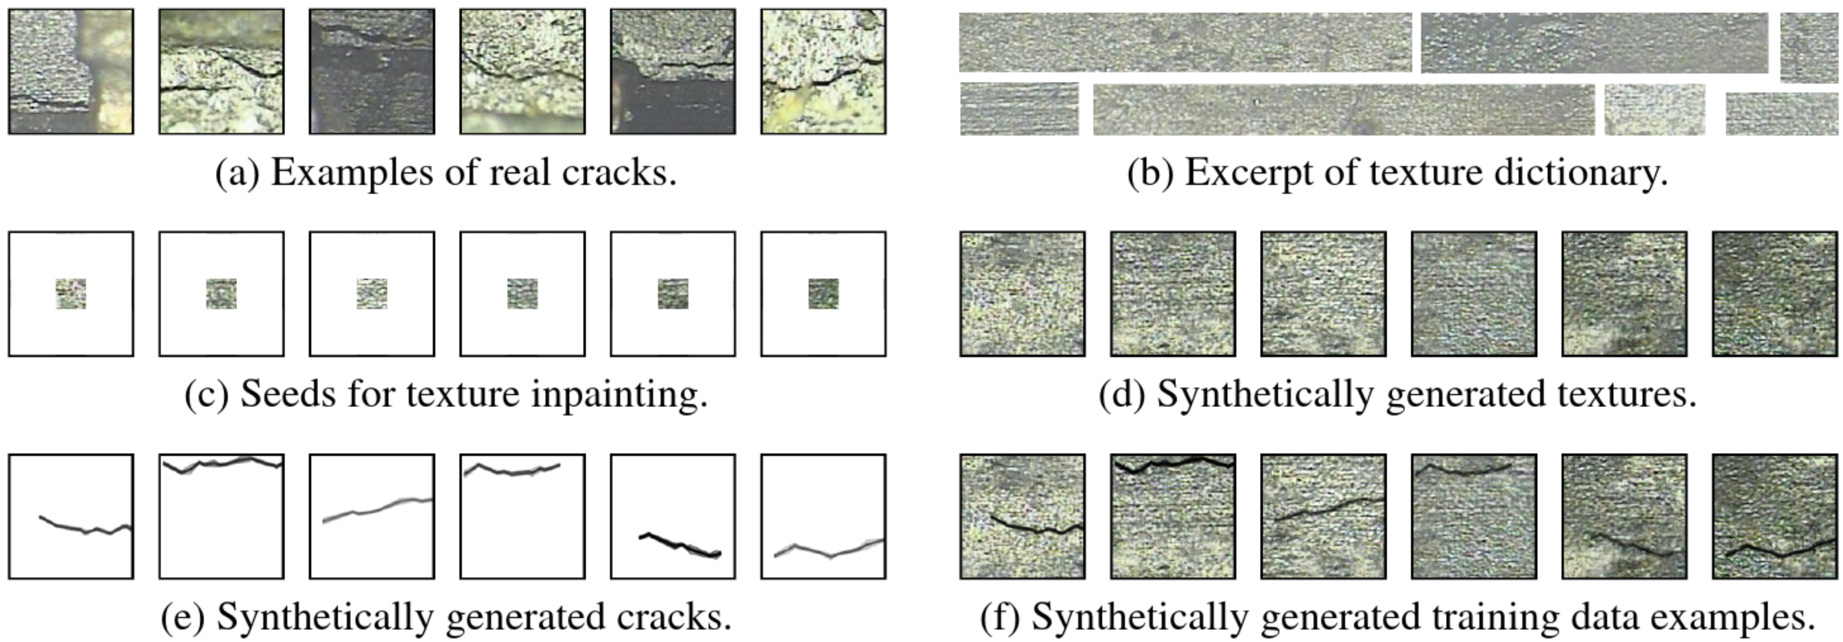

(a) Examples of real cracks.

(b) Excerpt of texture dictionary.

(c) Seeds for texture inpainting.

(d) Synthetically generated textures.

(e) Synthetically generated cracks.

(f) Synthetically generated training data examples.

Evaluation

- **Using synthetic data**
  - Generated 99,140 images (79,312 for training, 19,828 for validation)
- **Three Models**
  - Trained on real data, synthetic data only, and mixed
    - Mixed: Additional fine-tuning of synthetic model with real data
  - Statistics over 200 runs with random training/validation sets
- **Results**
  - Excellent first results on real industrial data sets
  - Mixed approach helps to overcome domain mismatch

Evaluation

- **Using synthetic data**
  - Generated 99,140 images (79,312 for training, 19,828 for validation)
- **Three Models**
  - Trained on real data, synthetic data only, and mixed
    - Mixed: Additional fine-tuning of synthetic model with real data
  - Statistics over 200 runs with random training/validation sets

		real model
test data	mean error ± standard deviation	0.273 ± 0.040
	confidence interval	$[0.267, 0.279]$
all data	mean error ± standard deviation	0.147 ± 0.039
	confidence interval	$[0.143, 0.152]$

Image Segmentation with Digital Reality

- **Images from material science (Electron Microscopy)**
  - Expensive to acquire and manually segment (12 images!)
  - Target is to do statistics on the grains (#, size, etc.)
  - Only used synthetic model to segment grains & boundaries

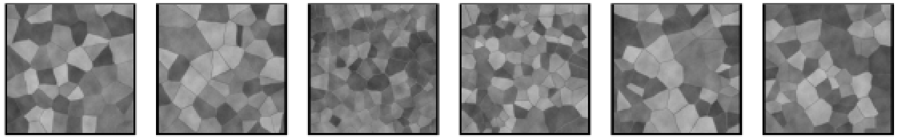

(a) Examples of synthetic images.

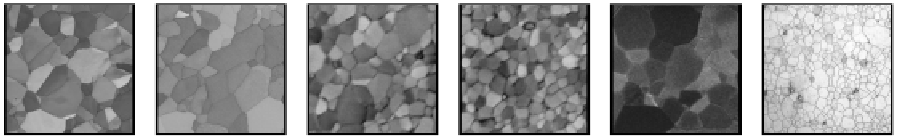

(b) Examples of real images.

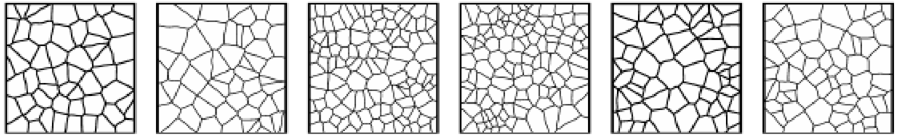

(c) Auto-generated masks for the synthetic images.

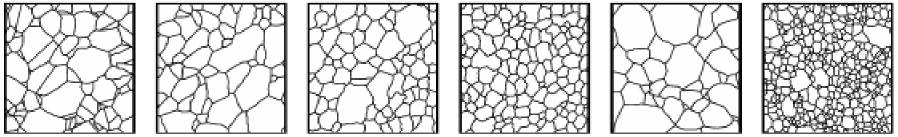

(d) Manually segmented masks for the real images.

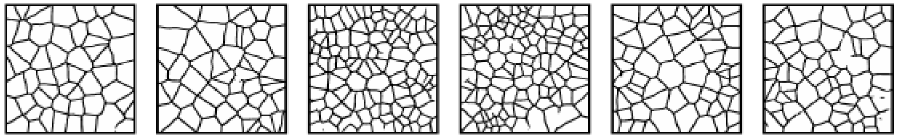

(e) Masks for the synthetic images generated with a trained model.

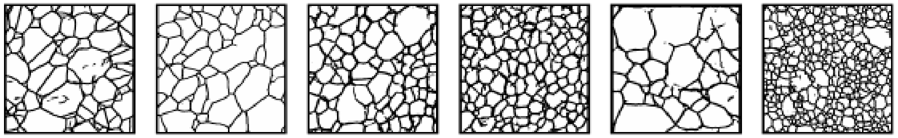

(f) Masks for the real images generated with a trained model.

Take-Aways

- **Digital Reality as a fundamental tool in AI**
  - Modeling, simulation, and leaning even in complex environments
  - **Learning and reasoning via feedback loop (e.g. RL)**
  - **Key element for future AI systems**
- **Continuous Learning Loop using Synthetic and Real Data**
  - Needed to achieve Certification and Validation of AI systems
  - **Certification & Validation required to establish trust in AI systems**
  - **Needs significant HPC resources for simulations and AI**
- **Big Challenges Ahead**
  - Many promising partial results already – but largely islands
  - **Requires closer collaboration of industry & academia**
  - **CLAIRE: Towards large-scale European initiative (“CERN for AI”)**

→ AI: A Central Component for Many Years to come


Einkaufsliste