SIEMENS
Ingenuity for life

Technologien

Münchner Kreis: Das „Data Analytics“ Zeitalter

Siemens Corporate Technology, Volker Tresp

SIEMENS
Ingenuity for life

1 Digitalization of Industries – a Paradigm Shift

2 From Individual Solutions to Analytics Platforms

3 Building Advanced Analytics Framework

4 Towards Artificial Intelligence

Digitalization changes everything

The world is becoming more digital
... also in industrial environments

SIEMENS
Ingenuity for life

Manual machine configuration		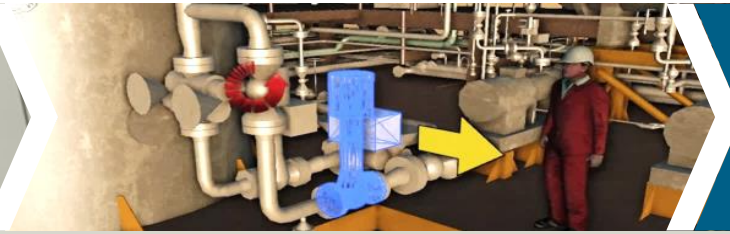	Virtual commissioning
Large power plants		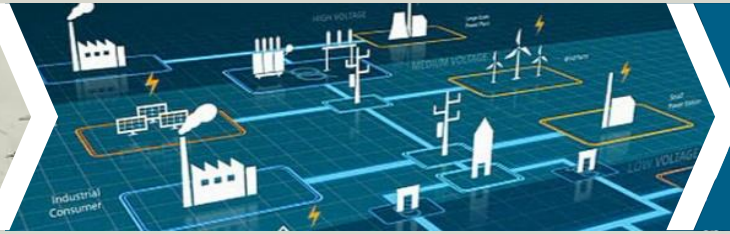	Virtual power plants
X-ray photography		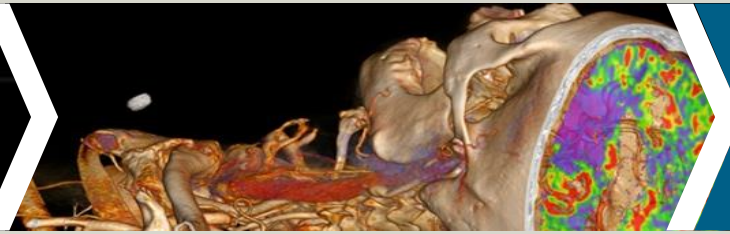	Digital imaging and analysis
Fixed maintenance intervals			Predictive maintenance

Massive pervasion of technologies driven by exponential growth of *computational power* are enablers for digitalization

SIEMENS
Ingenuity for life

Communication
and
Connectivity

Modeling
and
Simulation

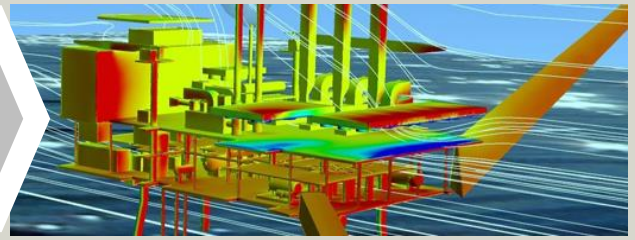

Autonomy
and
Intelligence

Big Data
and
Data Analytics

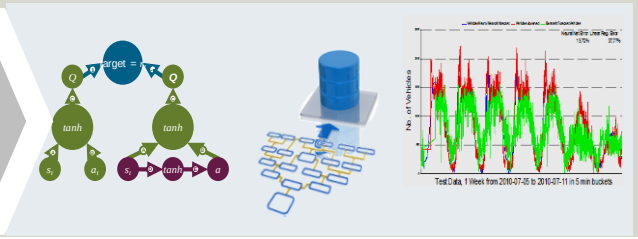

The availability of data and the possibility to get value out of the data is increasing rapidly

> 40 Zettabyte of data
in 2020 expected …

≈ including 20 Zettabyte **machine generated data**

Source: Oracle 2012, IDC CEO Summit 2015, misc. internet

Industry 4.0: Increasing complexity leads to new value systems
… and productivity, speed, flexibility remain the biggest challenges

SIEMENS
Ingenuity for life

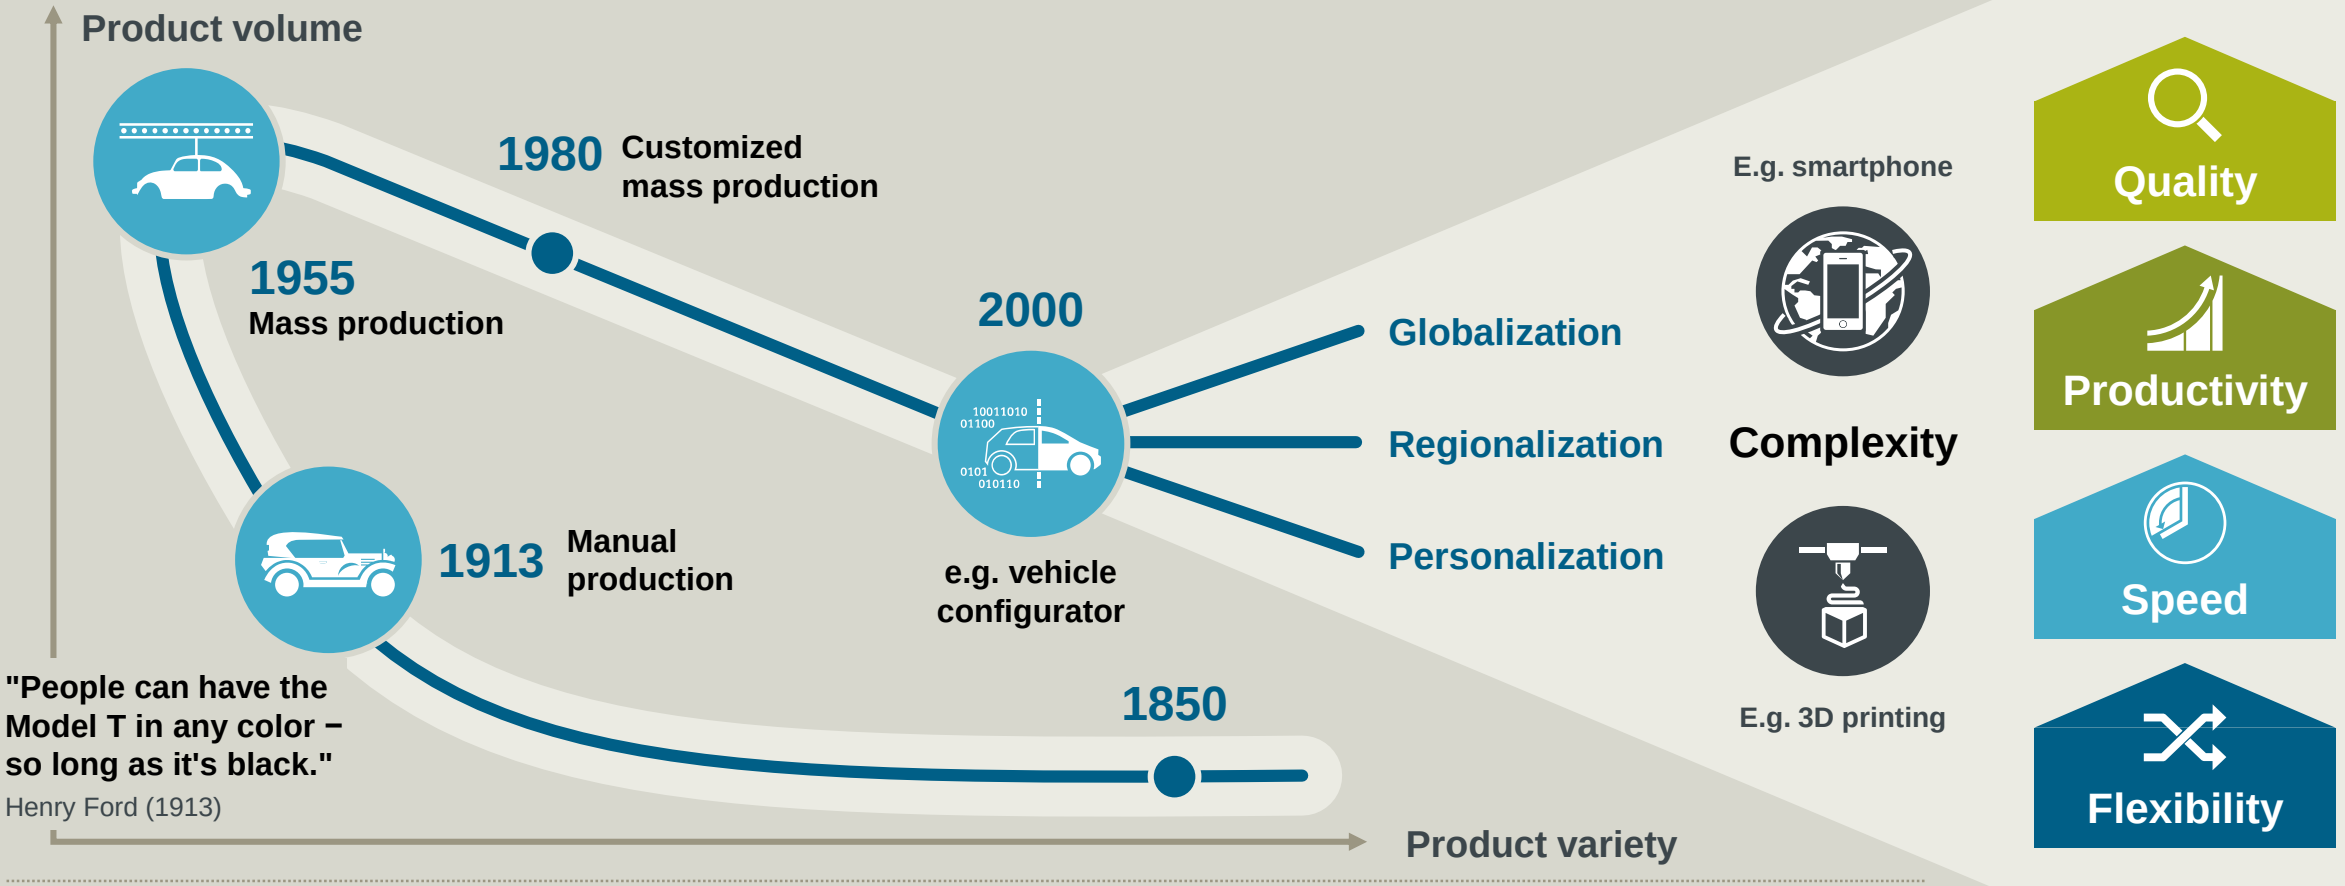

Based on: The Global Manufacturing Revolution; sources: Ford, beetleworld.net, bmw.de, dw.de.

Digitalization is disrupting entire value chains

SIEMENS
Ingenuity for life

- Data analytics
- Artificial Intelligence
- Simulation tools
- Cloud & platform technology
- Secure connectivity
- Cyber-Security

Enabling the next level of ...

... productivity and time-to-market ...

... flexibility and resilience ...

... availability and efficiency ...

Design and engineering → **Automation and operation** → **Maintenance and services**

… smart products and business models

Flexibility: Siemens AG, Gas Turbine Factory – 3D printing of gas turbine components

"The innovative selective **laser sintering** process **shortens production times** compared to conventional production processes, so much so that **function-critical components can** already **be tested during the product development stage**."

Sebastian Piegert, Siemens AG.

Source: Platform Industrie 4.0, Siemens AG, Berlin Gas Turbine Factory.

Productivity: Data Analytics supports machine optimization e.g. improved efficiency of wind parks

SIEMENS
Ingenuity for life

- More than 200 GB of sensor data sensors from ≈ 7.800 wind parks
- Early detection of divergent behavior
- Autonomous learning with Neural Networks

1-3% increase of annual energy

Availability: Data analytic supports availability of systems e.g. health check for CERN's Large Hadron Collider

SIEMENS
Ingenuity for life

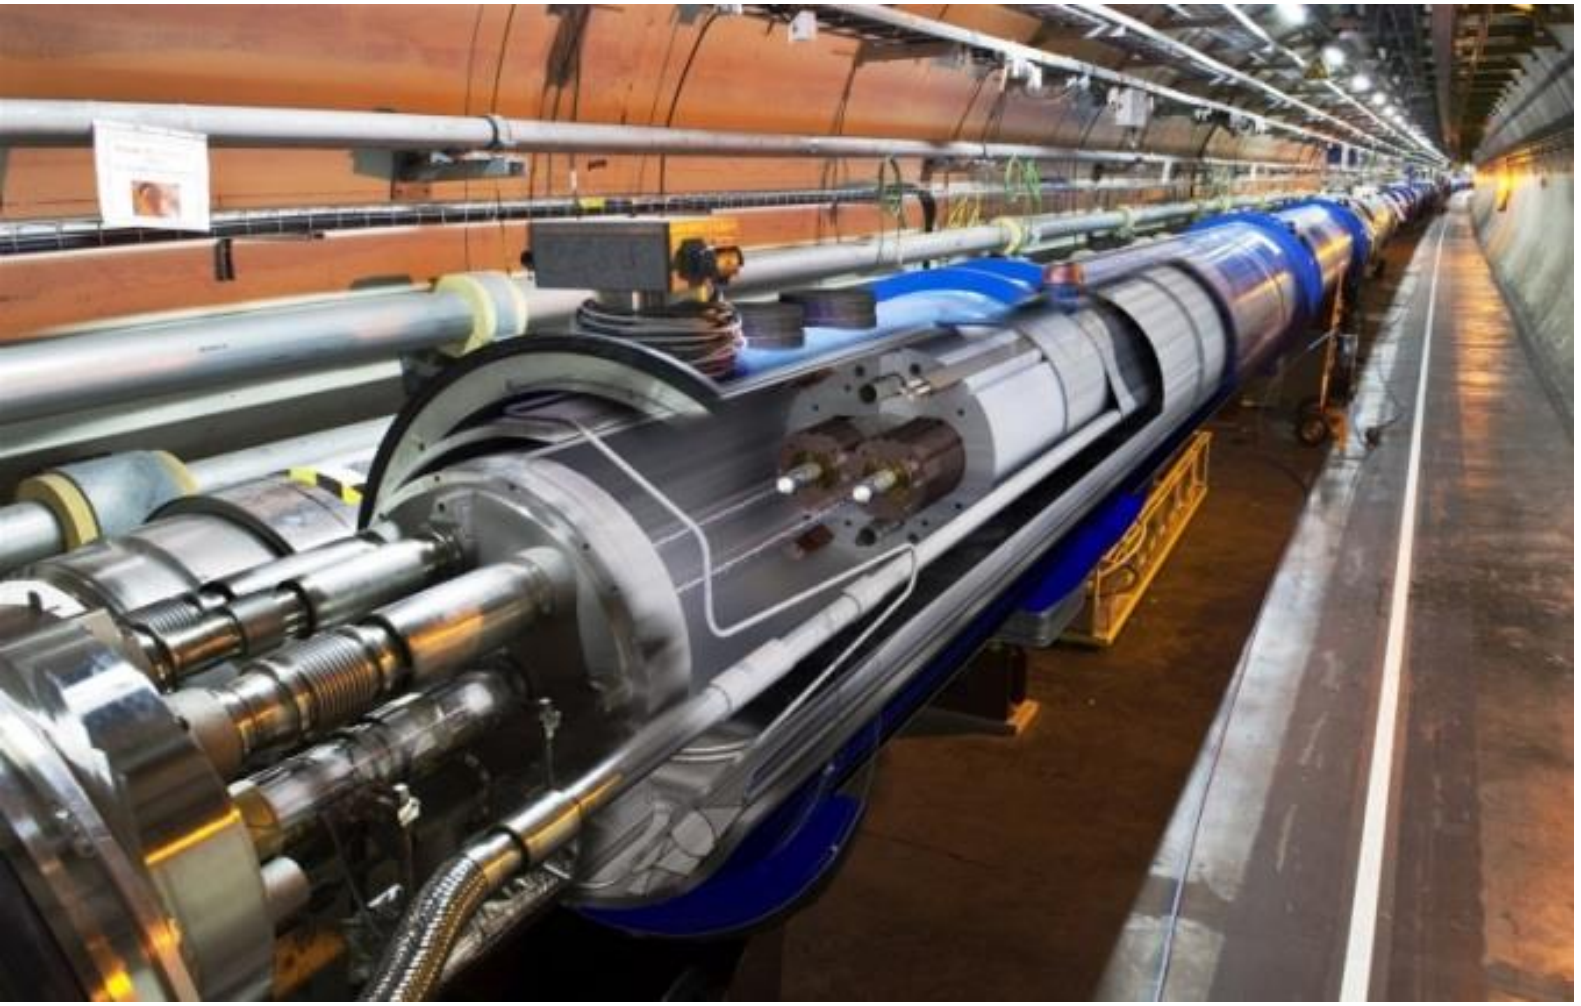

99.9999991% the speed of light

The biggest detectors ever …

… 600 million collisions per sec

Huge supervisory system and hundreds SIMATIC systems controlling the production

With rule and pattern mining methods increase operating hours

Value Chain: On the way to Industrie 4.0 – Driving the Digital Enterprise – Digital Enterprise Software Suite

SIEMENS
Ingenuity for life

- Italian passion – Virtually developed – Efficiently built
- Maserati increasing its competitiveness through digitalization of production

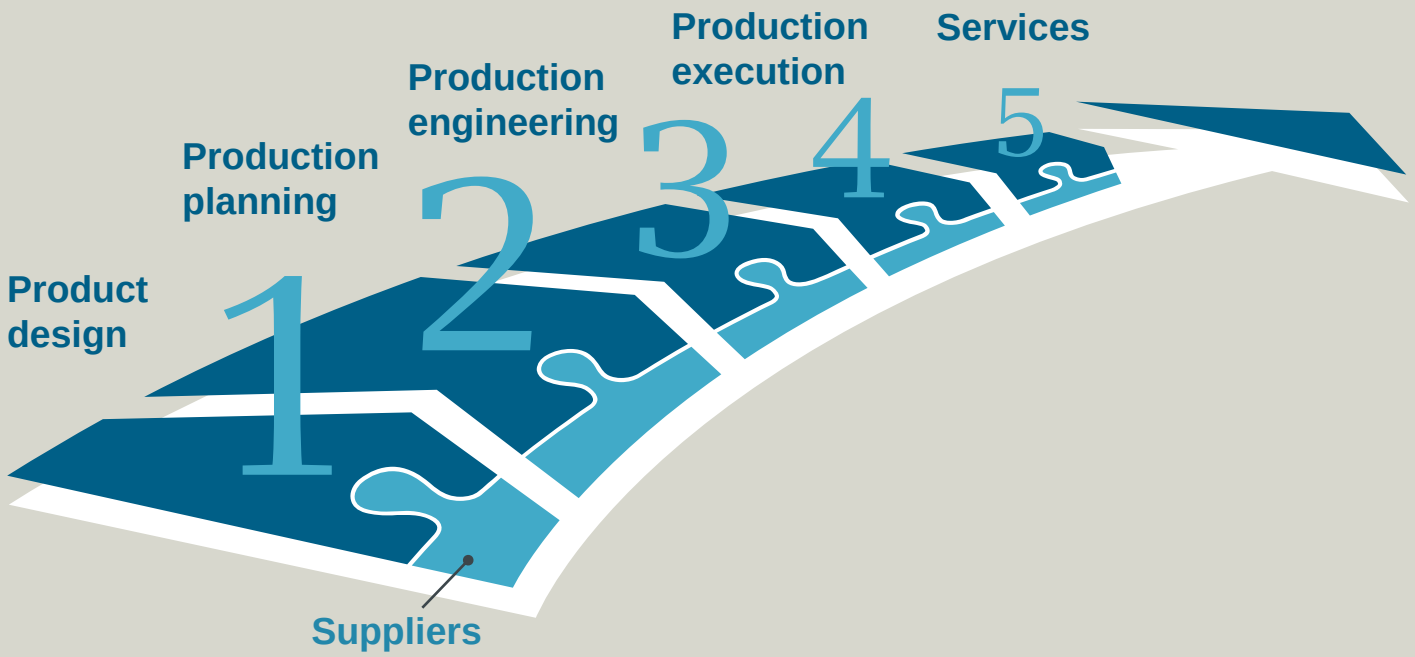

SIEMENS
Ingenuity for life

1 Digitalization of Industries – a Paradigm Shift

2 From Individual Solutions to Analytics Platforms

3 Building Advanced Analytics Framework

4 Towards Artificial Intelligence

MindSphere – The cloud-based, open Industrial Internet-of-Things (IIoT) operating system

SIEMENS
Ingenuity for life

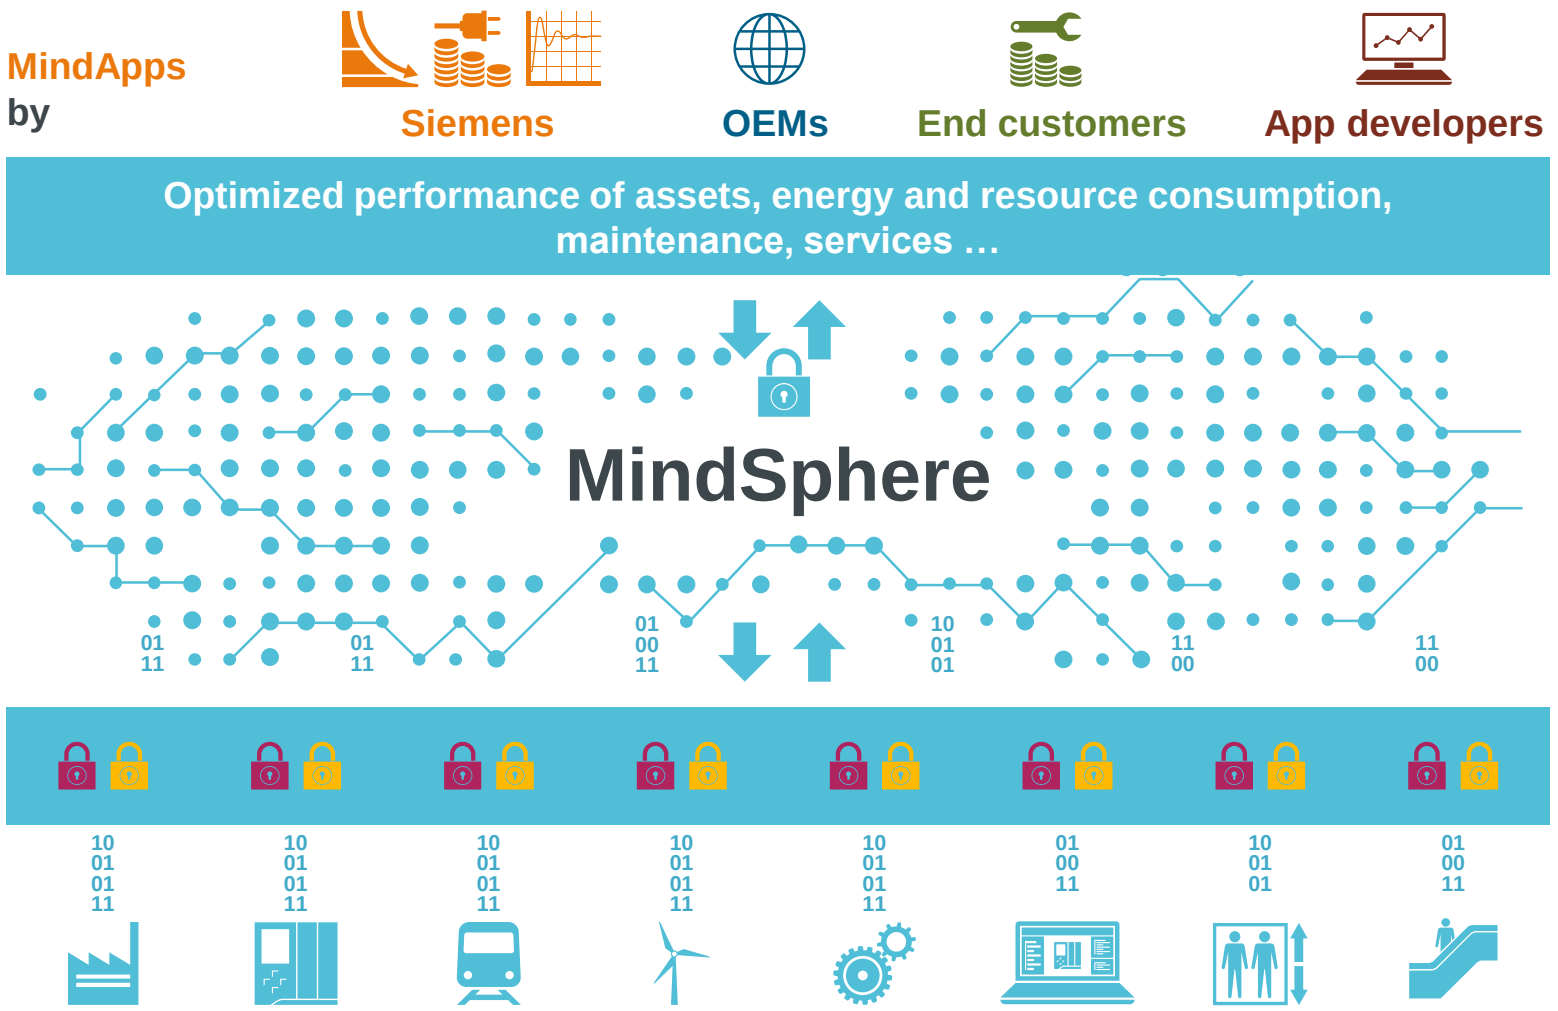

MindApps

- Developed by Siemens, OEMs, end customers and App developers
- Asset transparency and analytical insights, for example, predictive maintenance
- Subscription-based pricing model
- Fleet management

MindSphere

- Open interface for developing customer specific apps (MindApps)
- Various cloud infrastructures: public, private, or on-site

MindConnect

- Open standards (for example, OPC UA) for connectivity (also to third-party products)
- Secure plug and play connection of Siemens products

4 Technology Pillars for Analytics @MindSphere

SIEMENS
Ingenuity for life

Specific Analytics Building Blocks

Signal
API for time-series analysis of sensors

Anomaly
API for anomaly detection in process

KPI
API for ISO KPI calculations

Performance
API for performance analysis

Trend Prediction
API for predictive trending

Failure Modes
API for various standard failure modes

General Analytics Building Blocks

Statistics
API for statistical functionalities

Clustering
API for Clustering (e.g. K-Means)

Regr. Analysis
API for Regression Analysis (e.g. Linear)

Rules
API for building rules

Event Analytics
API for Sequence Pattern Mining

Classification
API for Classification (e.g. NN or SVM)

→ 4 Technology Pillars

- **Specific Industrial Analytics Algorithms**
- **Performance**
- **Analytics Design Patterns / Architecture**
- **UI / User Experience**

Siemens internal digitalization projects experience + { **Analytics KnowHow** | **Performance Expertise** | **Specific SW Development** | **UI / User Experience** } = **Basic Analytics @MindSphere**

SIEMENS
Ingenuity for life

1 Digitalization of Industries – a Paradigm Shift

2 Step 1 – From Individual Solutions to Analytics Platforms

3 Step 2 – Building Advanced Analytics Framework

4 Step 3 – Towards Artificial Intelligence

Data Analytics Solution Packages help to reduce use case development costs

Solution Packages for Smart Data Analytics

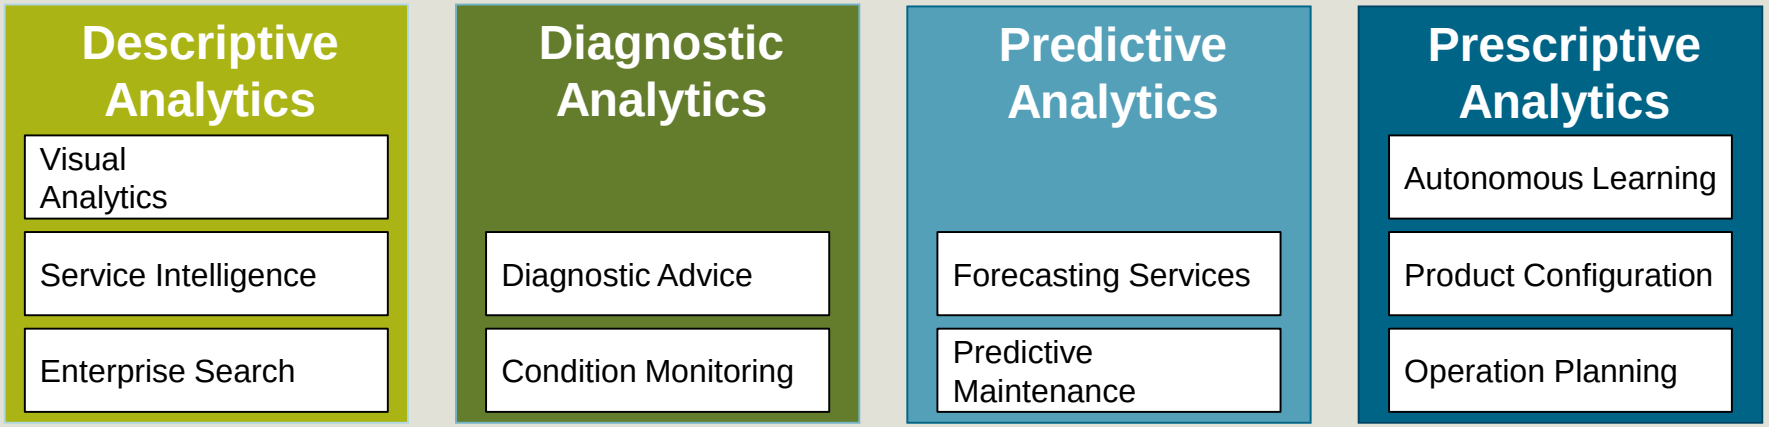

Purpose:

- Reusable data analytics elements that can be easily combined and applied for different branches and divisions

Screenshot of Smart Grid Price Forecasting

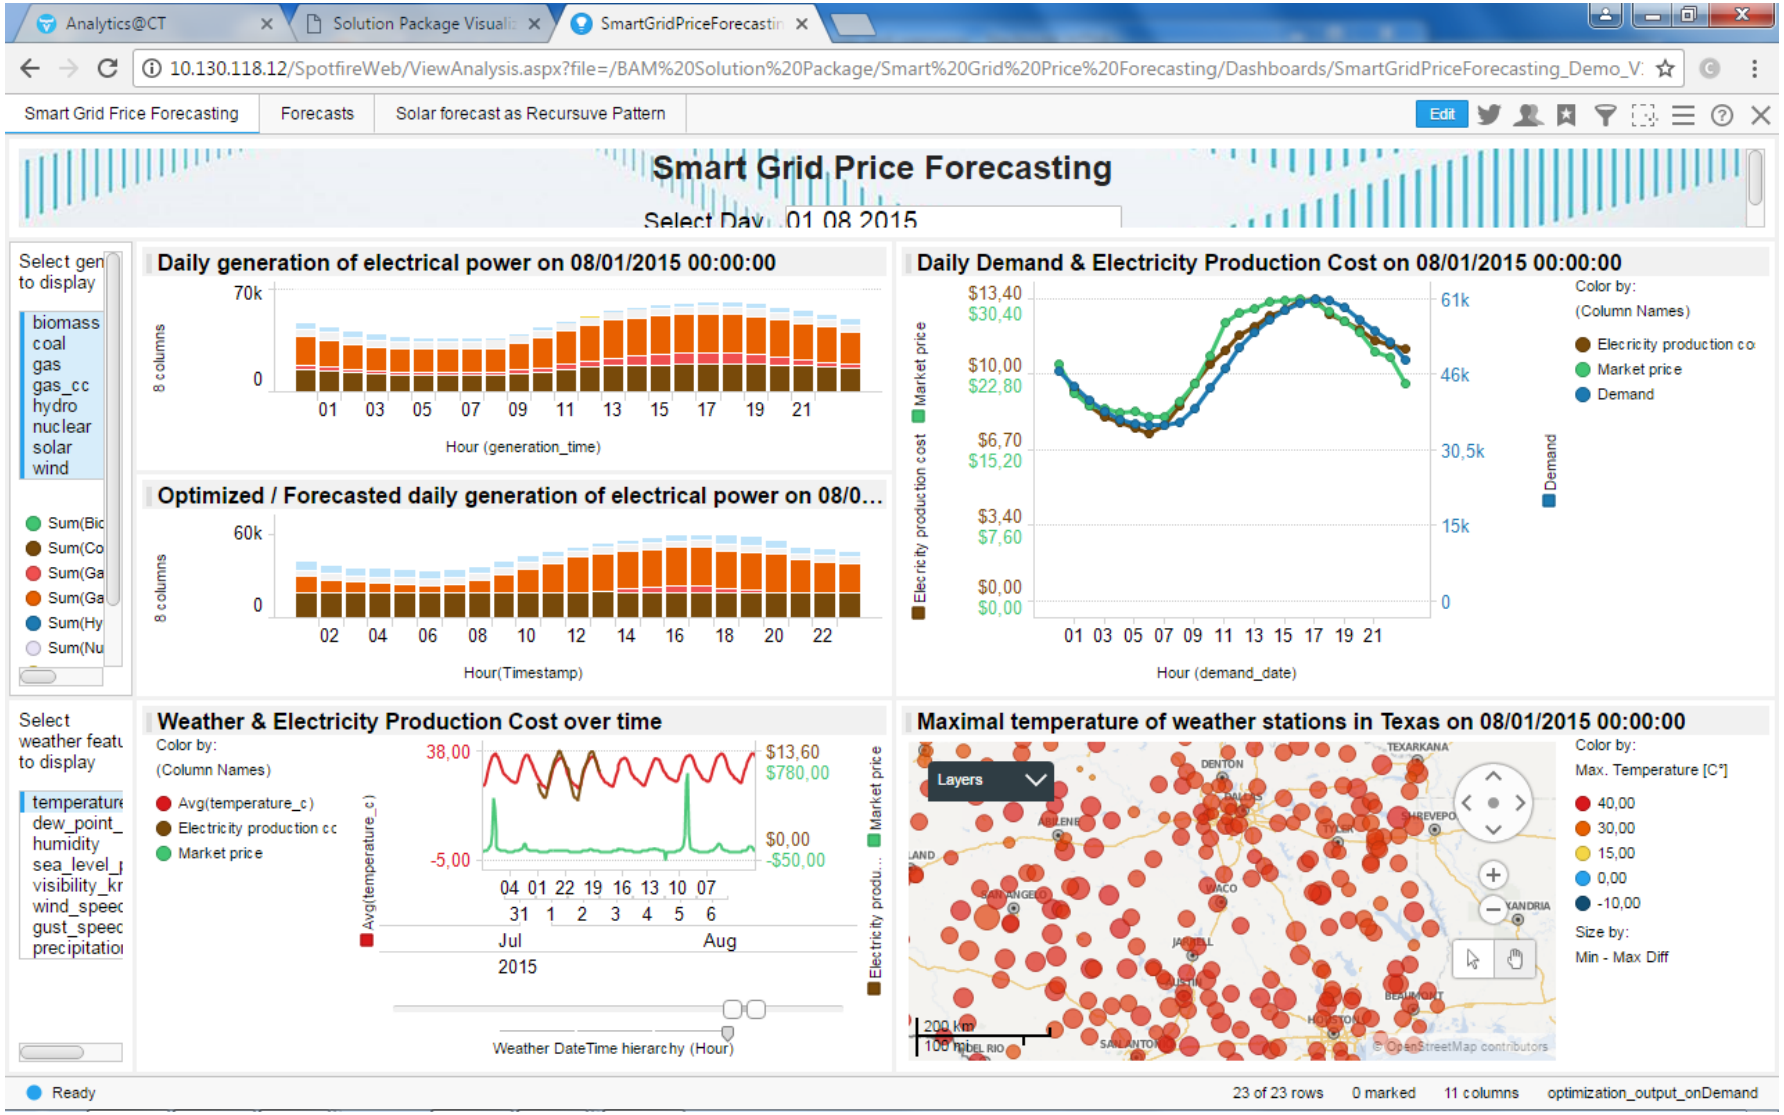

SIEMENS
Ingenuity for life

1 Digitalization of Industries – a Paradigm Shift

2 From Individual Solutions to Analytics Platforms

3 Building Advanced Analytics Framework

4 Towards Artificial Intelligence

AI changes everything

Artificial Intelligence

- *Creating machines that perform functions that require intelligence when performed by people (Kurzweil, 1990)*

Games

Translation

Q&A

Face Recognition

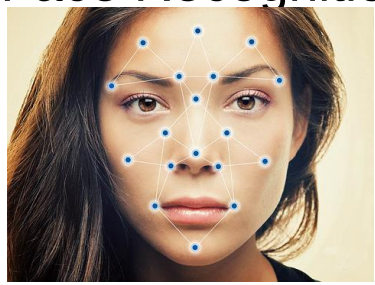

Auton. Driving

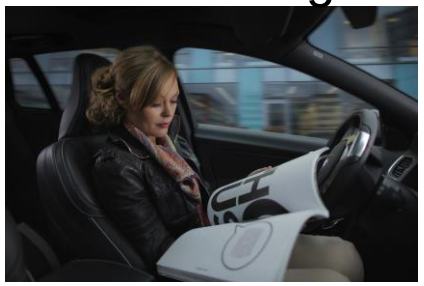

Speech Recognition

Drones, Robots

?

Many major breakthroughs in AI have occurred since 2011

SIEMENS
Ingenuity for life

Before 2011

1946
Zuse's *Z3*, first programmable electronic computer

1977
IBM Deep Blue defeats world's chess champion Kasparov

2005
Honda's humanoid robot "Asimo" comes to life

2011 – 2016

2011
Watson wins "Jeopardy!" against most successful contestants

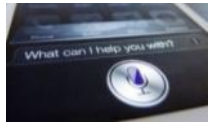

2011
Apple's intelligent assistant debuts on the iPhone

2016
AlphaGo beats Lee Sedol in a Go match

Expected by 2030+

202X
All-over virtual personal assis-tants as interface for consumers

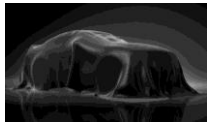

~202X
Fully autonomously driving cars become market-ready

20XX
Robots may build robot "children" on their own

Recent drivers

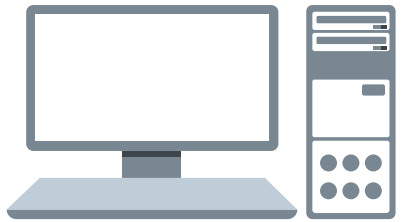

Increasing computer power

Increasing dataset sizes

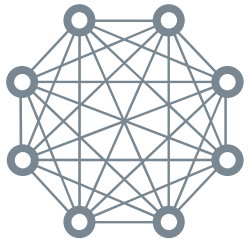

Improved efficiency of algorithm

Deep X Technologies behind Artificial Intelligence

- **Deep Learning**; Machine Learning; Data Mining; Statistics
  - More (Labeled) Data
  - Deeper Models
  - New Algorithms
  - End-to-End Training
  - Computational Power
  - Community
- **Deep Knowledge**; QA; information retrieval; search
  - Rapid IE
  - **Knowledge Graphs**
  - Web Content, Wikipedia
  - Maps and GPS
  - Ubiquitous Big Data
  - Profiles (genetics, …)

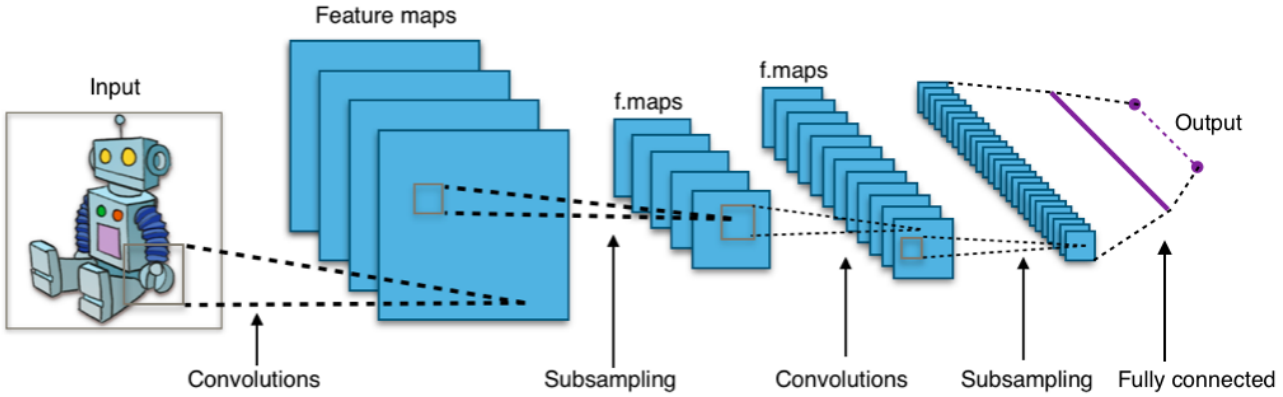

Deep Learning

- Deep Learning is the reason for the emerging huge interest in AI
  - Convolutional DL
  - Recurrent DL
  - Reinforcement DL

Face Recognition

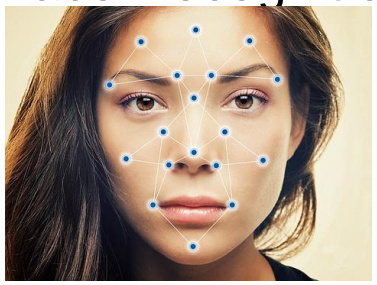

Translation

Speech Recognition

Games

Student Magic: Visual Q&A

Image	Policy Gradient		Tempered Policy Gradient	
	Is it in left?	No	Is it a person?	No
	Is it in front?	No	Is it a **vehicle**?	Yes
	Is it in right?	Yes	Is it a **truck**?	Yes
	Is it in middle?	Yes	Is it in front of photo?	No
	Is it person?	No	In the left half?	No
	Is it ball?	No	In the middle of photo?	Yes
	Is it bat?	No	Is it to the right photo?	Yes
	Is it **car**?	Yes	Is it in the middle of photo?	Yes
	Status:	Failure	Status:	Success
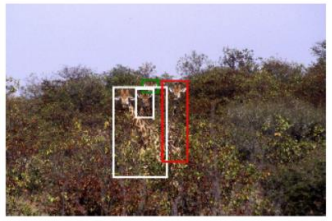	Is it in **left**?	No	Is it a giraffe?	Yes
	Is it in front?	Yes	In front of photo?	Yes
	Is it in right?	No	In the **left half**?	Yes
	Is it in middle?	Yes	Is it in the middle of photo?	Yes
	Is it person?	No	Is it **to the left** of photo?	Yes
	Is it giraffe?	Yes	Is it to the right photo?	No
	Is in middle?	Yes	**In the left** in photo?	No
	Is in middle?	Yes	In the middle of photo?	Yes
	Status:	Failure	Status:	Success

Convolutional DL
+ Recurrent DL
+ Reinforcement DL

Talents, Talents Talents!

The Oracle Model

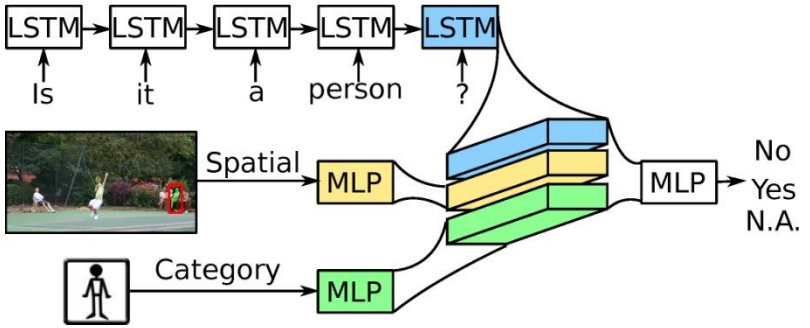

The Question-Generator Model

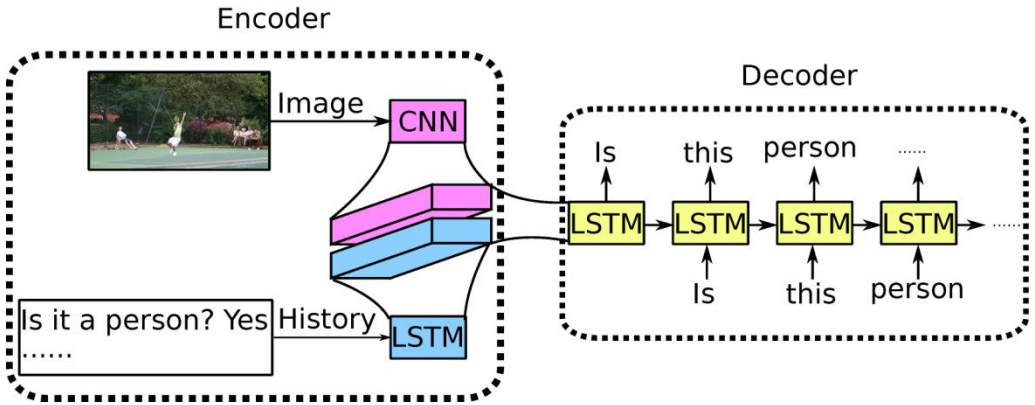

The Guesser Model

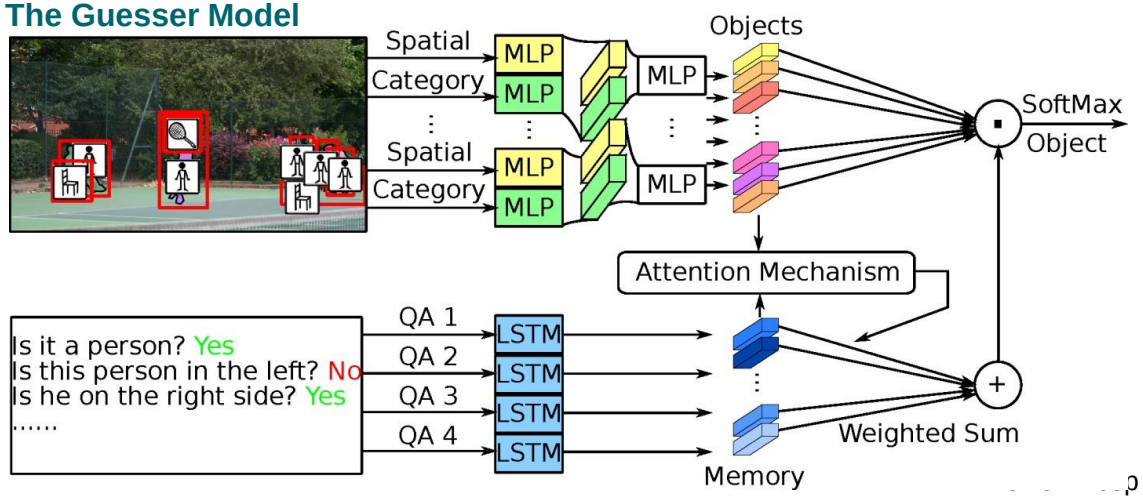

A Tradition of Deep Learning at Siemens:
Online Learning Deep Learning deployed in dozens of Steel Plants

A Tradition of Deep Learning at Siemens: Self-Optimizing Gas Turbines learn to reduce NoX Emissions

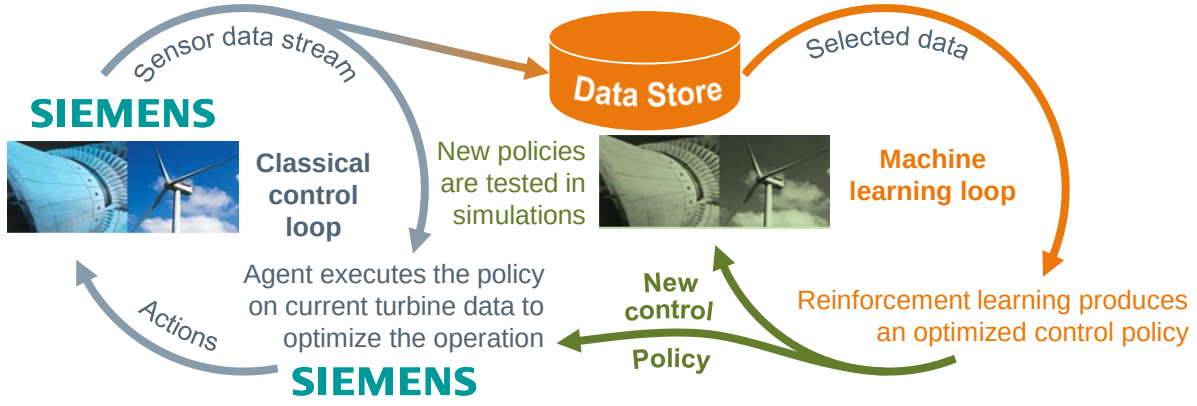

Product / use case details

- **Learn a improved control policy** for a gas turbine from data, e.g. for emission reduction or better yield
- **Simulate the behavior of gas turbine using neural networks** based on historical data
- Optimize the policy using **reinforcement learning**
- **Deploy** improved policy on the turbine

Customer value-add

- **Emission reduction** of gas turbines by 10-15%
- Performance improvement can be achieved **without need for new hardware**

Status and next steps

- In **productive use** at customer sites
- Deployment for further classes of gas turbines planned
- Extension to **wind turbines** (field test ongoing)
- **Many further opportunities in other Siemens' domains!**

Deep Knowledge: Knowledge Graphs

- The Google Knowledge Graph is a major break through in the field of Knowledge Representation
  - Scalability: > 100B fact
  - Reliability: >99% fidelity
  - Maintainability
  - Usefulness: Search, Q&A, text understanding
- The basis for the Google Knowledge Graph are facts

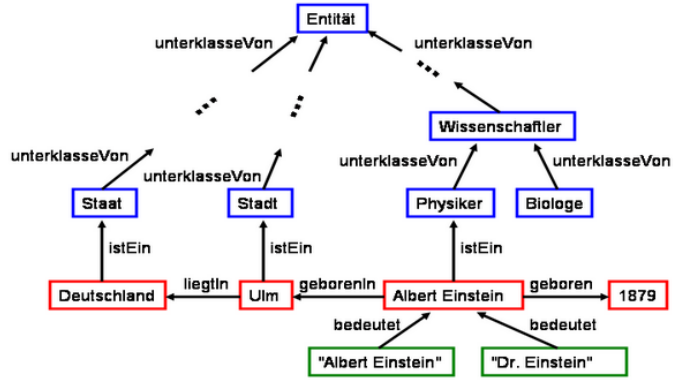

>100B Facts

Industrial Knowledge Graphs for capturing Siemens Domain Knowledge

SIEMENS
Ingenuity for life

1
Isolated Data Silos with hand-crafted expert systems

2
Domain-specific Knowledge Graphs generated from DBs

3
Connected Knowledge Graph via automated structure and link discovery

4
Learning Memories extract expert knowledge from observations

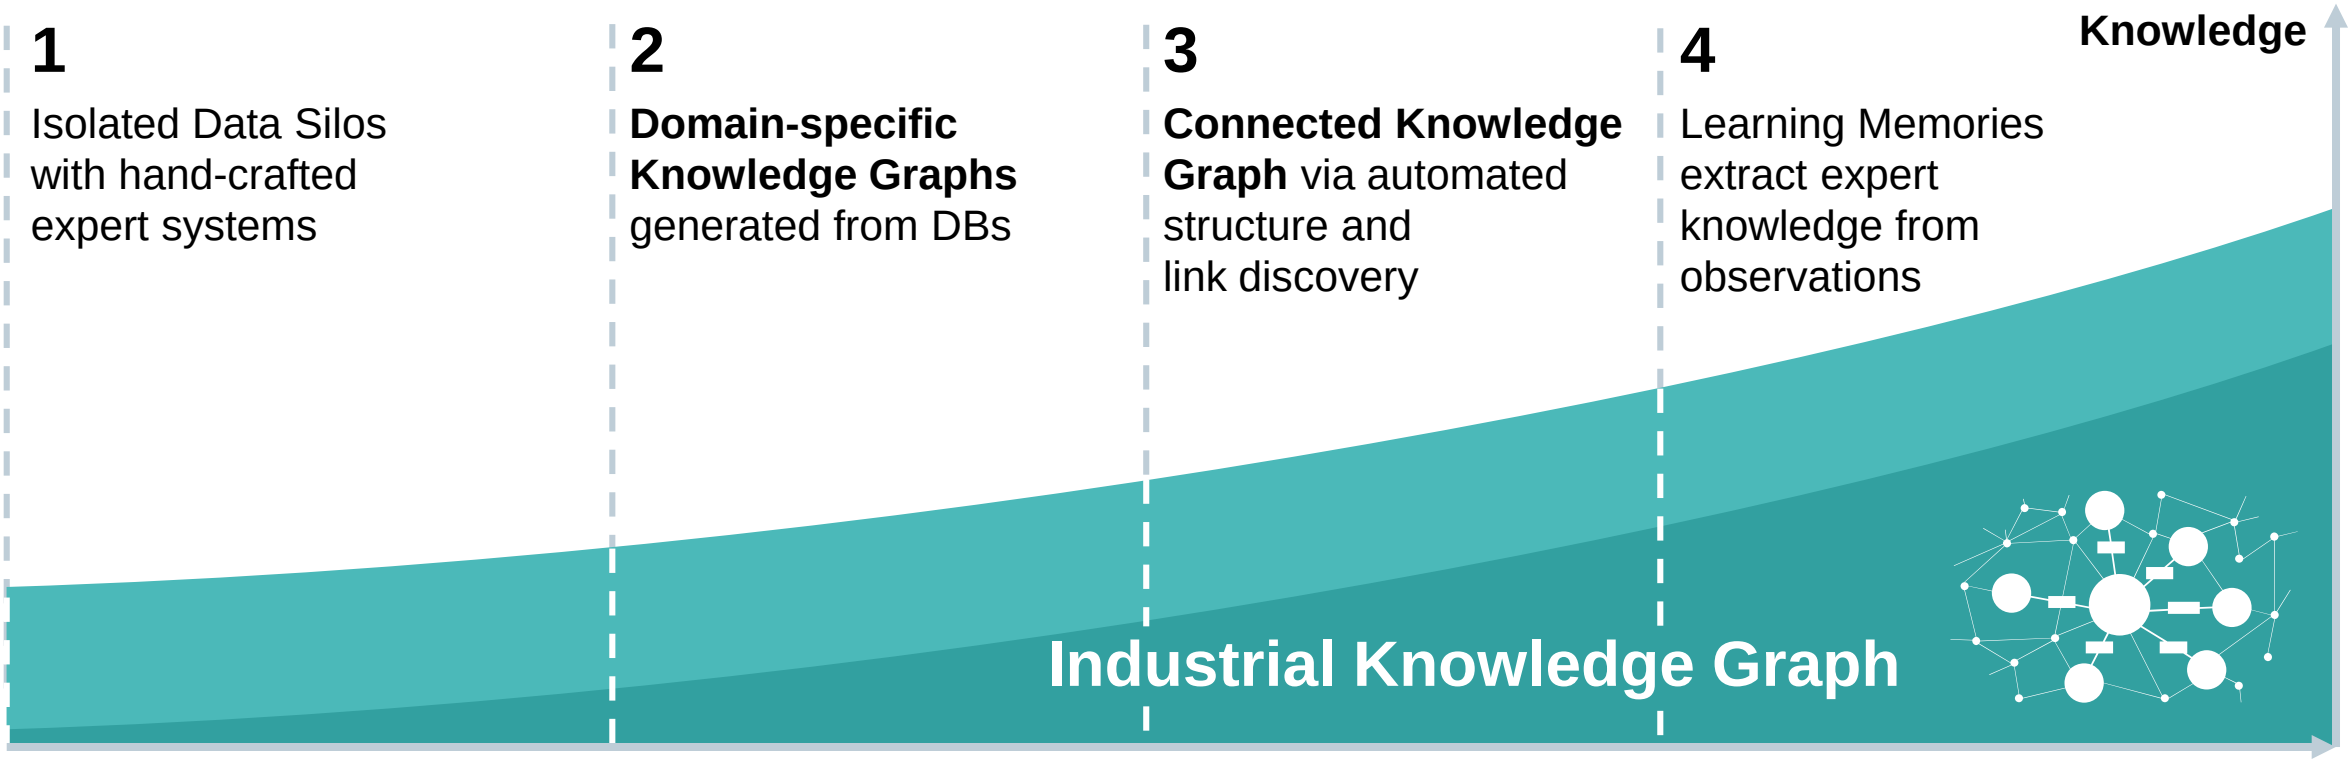

Degree of automated knowledge digitalization

From isolated data silos to learning memories

Collected data Digitalized Knowledge (via reasoning and learning)

One day in the life of an artificial assistant @ Siemens

SIEMENS
Ingenuity for life

Industrial Knowledge Graph

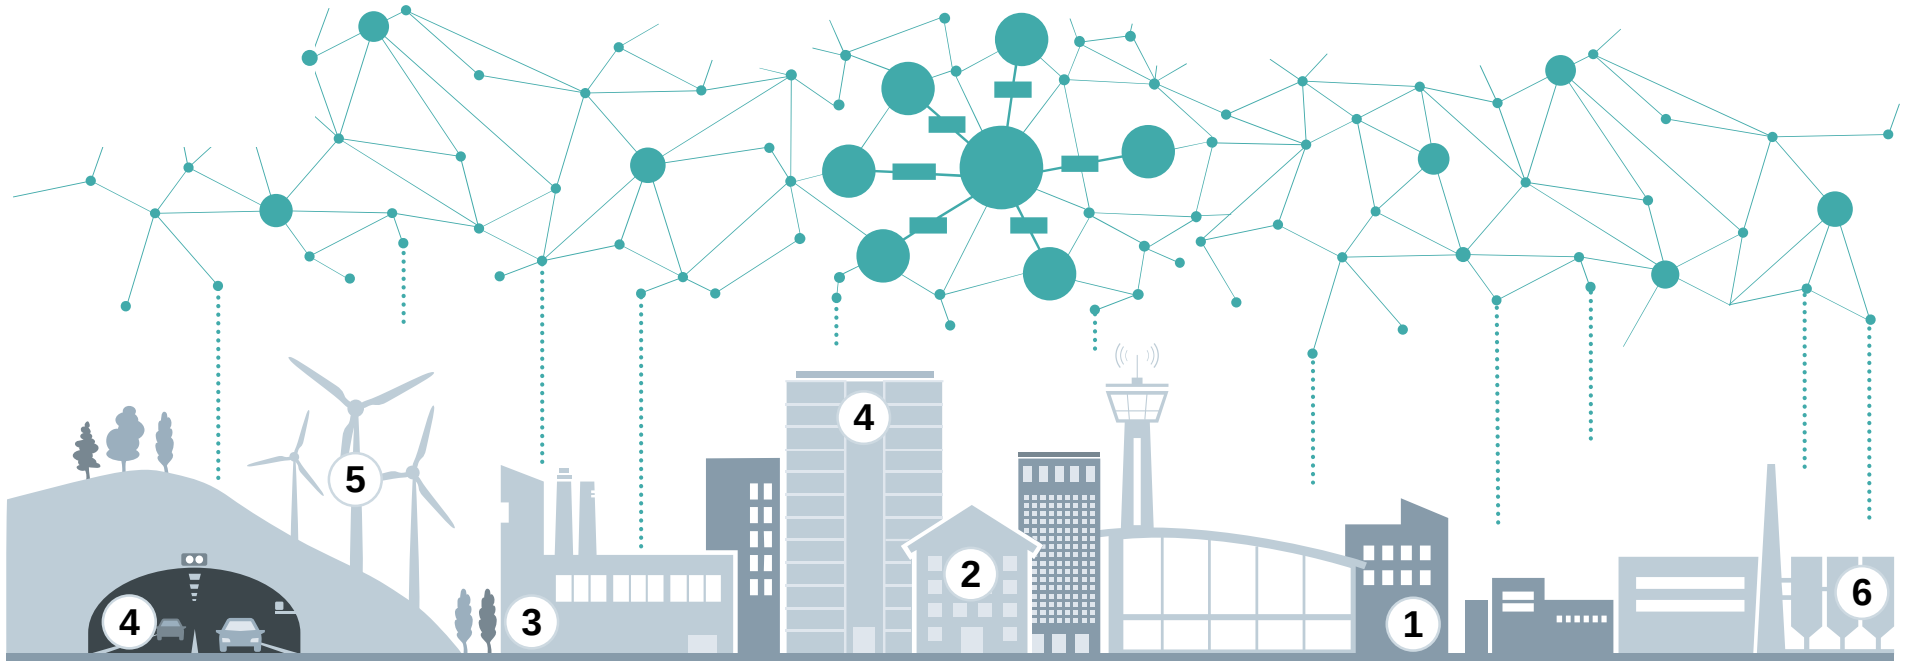

1. **09:00 – Understand**
   Tender Search
2. **11:00 – Offer**
   Configure Solution
3. **12:00 – Produce**
   Flexible Manufacturing
4. **13:00 – Supply**
   Supply Chain Mgmt.
5. **15:00 – Maintain**
   Maintenance planning
6. **18:00 - Repair**
   Advanced Diagnostics

AI Algorithms

Industrial Knowledge Graph @ Work – Augmentation of tender analysis process

SIEMENS
Ingenuity for life

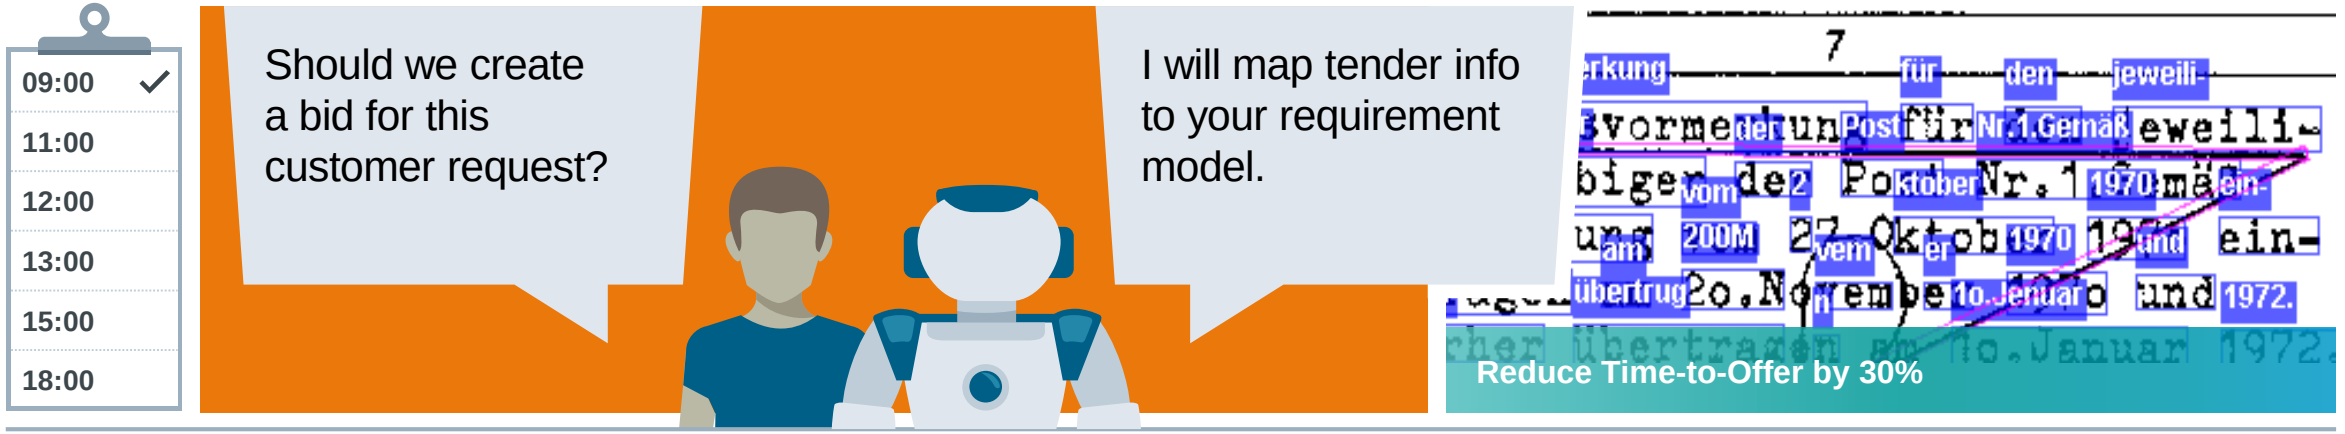

Knowledge Graph

- **Requirements model** formally capturing the scope of each tender document
- Related **domain models** to understand the context
- **Risk model** to describe critically of requirements and tenders

+

AI Algorithms

- **NLP** to extract and formally present requirements information
- **Case-based reasoning** to identify similarity in past tenders
- **Semantic mapping** to formally similarity of requirements
- **Machine learning** for automatic requirement domain assignment

=

Augmented Intelligence

- **Tender Comparison:** Automatically compare tenders tracking (critical) requirements from historic tenders
- **Automatic Bid Creation Process**: Automatic extraction of key configuration values as basis for automated bid creation

Industrial Knowledge Graph @ Work – Flexible manufacturing for Industry 4.0

SIEMENS
Ingenuity for life

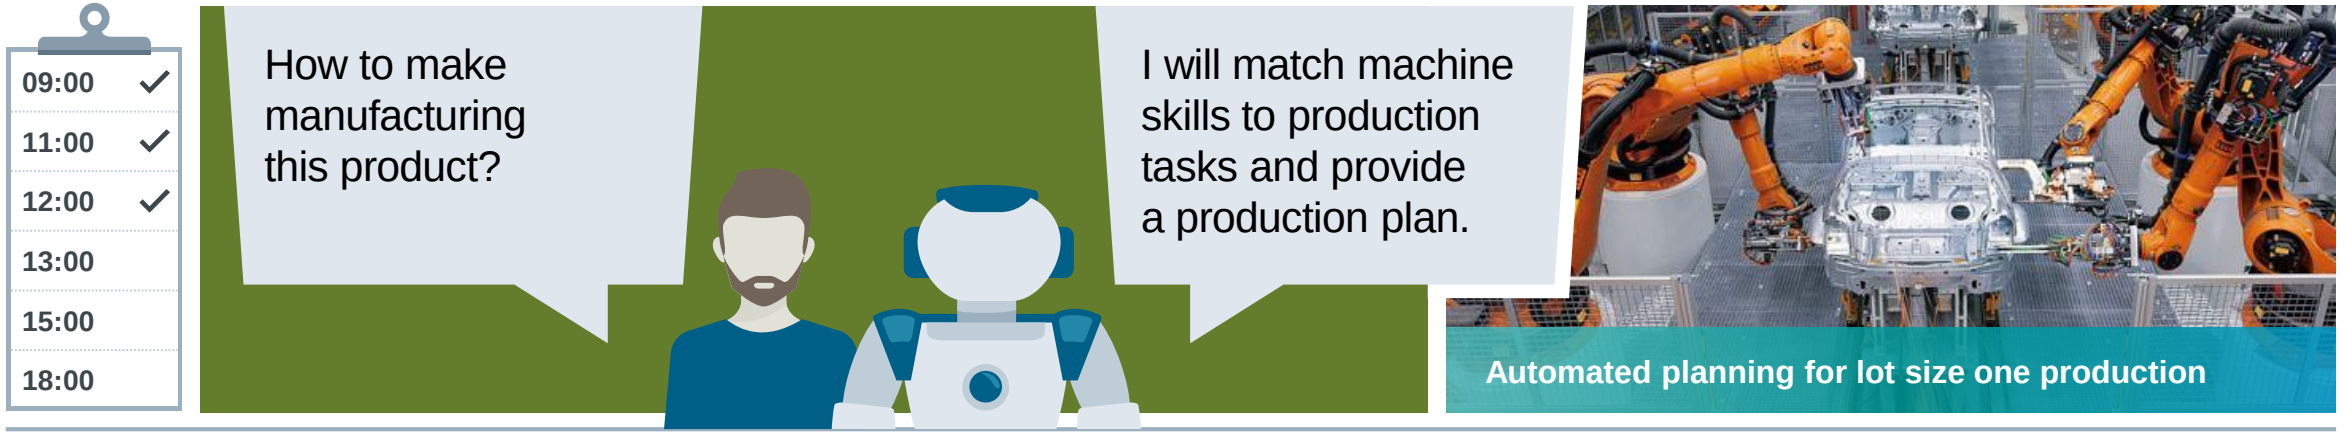

Knowledge Graph

- **Skills of a machine**, e.g. equipment, products, material, processes
- **Tasks** required to manufacture a **product**

+

AI Algorithms

- **Reasoning** about operation taxonomies, material or spatial constraints, etc.
- **Semantic Matching** of machine capabilities against required production steps
- **Constraint / Goal-driven planning** of production workflows based on machine knowledge

=

Augmented Intelligence

- **Smart Factories:** Enable flexible production scenarios
- **Smart Machines:** Act autonomously based on self-awareness
- **Smart Products:** Steer their own production utilizing machine skills

Industrial Knowledge Graph @ Work – Market intelligence for supply chain management

SIEMENS
Ingenuity for life

Knowledge Graph

- External and internal information about **suppliers** and their **products**
- Facts extracted from **social media, news, company websites** and **industrial forums**

+

AI Algorithms

- **NLP** to extract domain knowledge from reports and other sources
- **Machine Learning** for extracting and integrating facts

=

Augmented Intelligence

- **Decision support to** select vendors
- **Risk monitoring** of supplier portfolio and sourcing strategies
- **Optimization** of sourcing and supply chain processes
- **Merge internal and external information** into one coherent view

Industrial Knowledge Graph @ Work – Advanced diagnostic systems for power plants

SIEMENS
Ingenuity for life

Knowledge Graph

- **Failure** catalogues
- Historical **maintenance** information
- **Expertise of service engineers** based on solved issues
- Power **plant design** and workflow environment

+

AI Algorithms

- **Case-based** reasoning to retrieve most similar case(s)
- **Natural Language processing** to extract features from technical tickets
- **Machine learning** to cluster similar cases

=

Augmented Intelligence

- **Diagnosis Support** combining experience-based information and machine data
- **Recommendation** based on knowledge gained through past successful failure handling
- Improved workflow integration and **knowledge management**

Machine Learning with Knowledge Graphs

- Our RESCAL model is based on an approximation of the Knowledge Graph adjacency tensor
- It was the basis for further research in our group, but also in many other groups

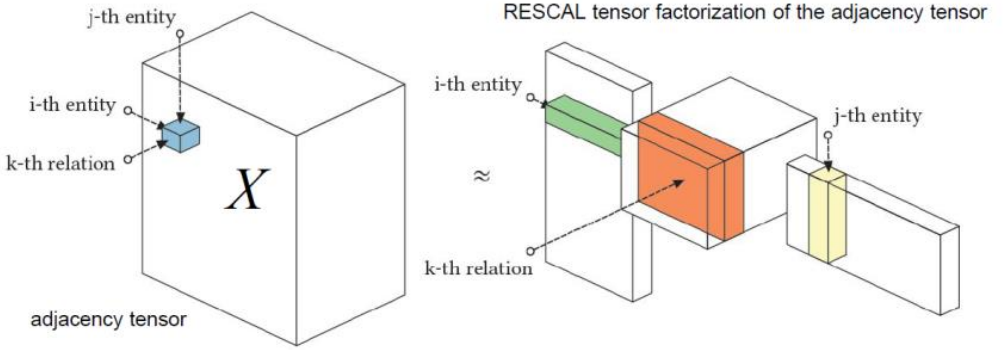

Training Data:

$x_{s,p,o} = 1$ If *(s,p,o)* is known to be true

$x_{s,p,o} = 0$ otherwise

After factorization (RESCAL2: constr. Tucker2):

$$P((s,p,o)) = \mathrm{sig}(\theta_{s,p,o})$$

$$\theta_{s,p,o} = \sum_{r_1} \sum_{r_3} a_{e_s,r_1} a_{e_o,r_3} g(r_1,p,r_3)$$

$$\Theta = G \times_1 A \times_2 A$$

- Inferential queries
  - What disease does Jack likely have?
- Automatic filling of KG
  - Knowledge Vault projects
- KG priors to understand text and images
- Detection of KG errors
- Learning Database
- Use as background information (compressed as latent factors) that can be used in other applications (predictions, decision support)

Smart Perception: Integrating Knowledge Graphs with Deep Learning

- Understanding the world means knowing the world

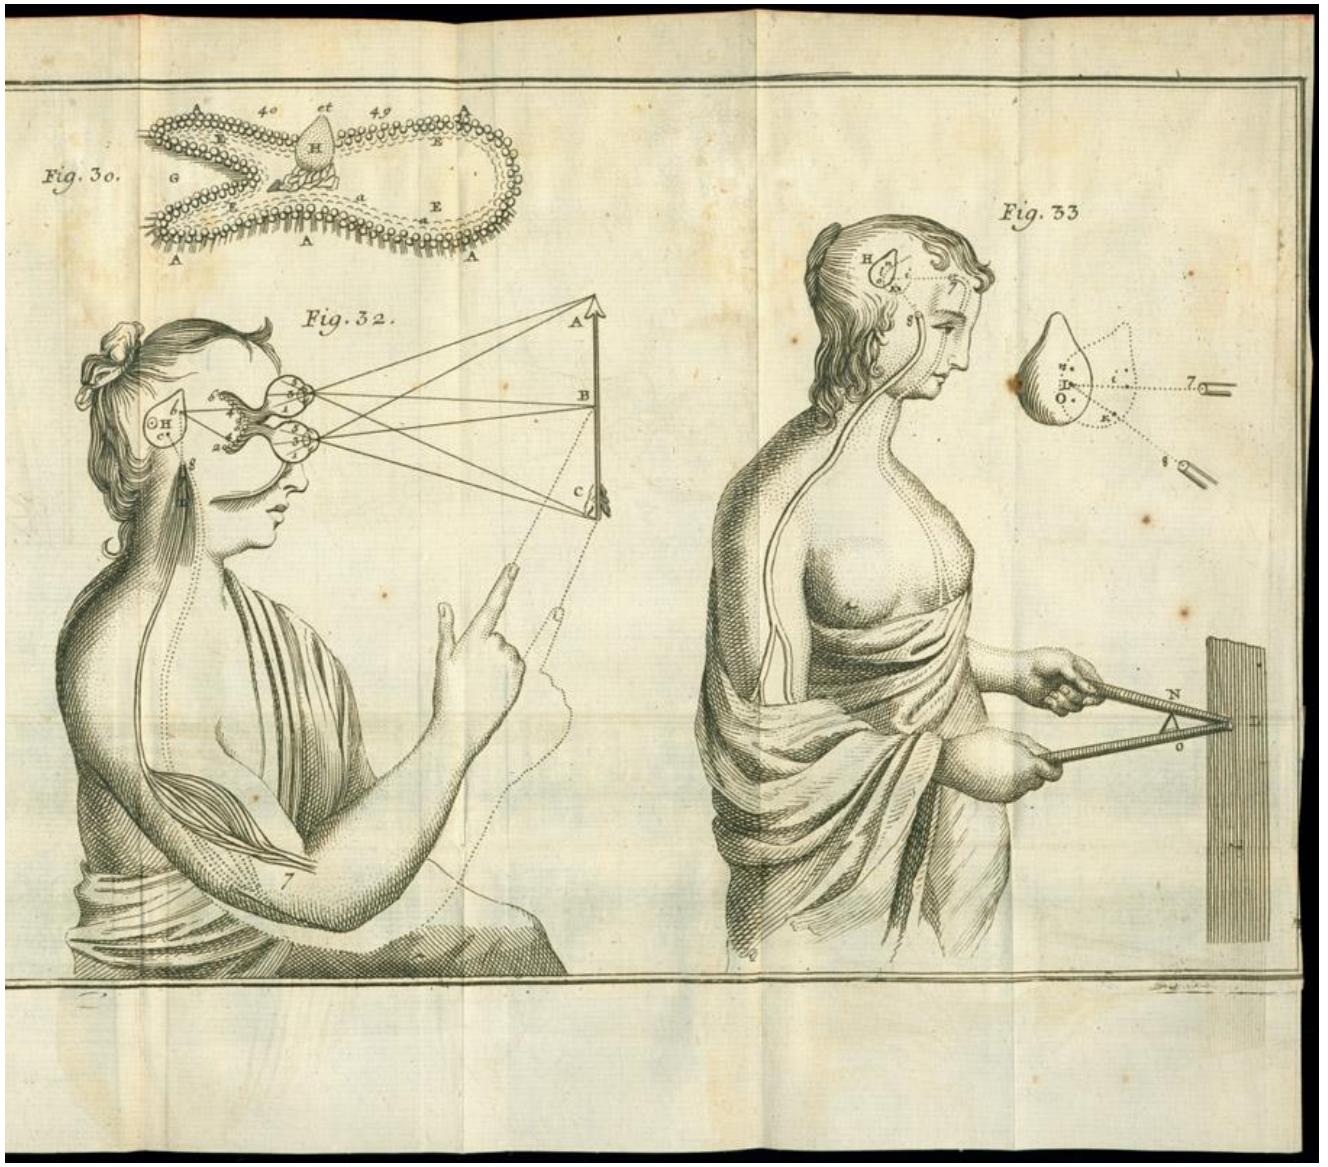

Smart Perception: Integrating Knowledge Graphs with Deep Learning

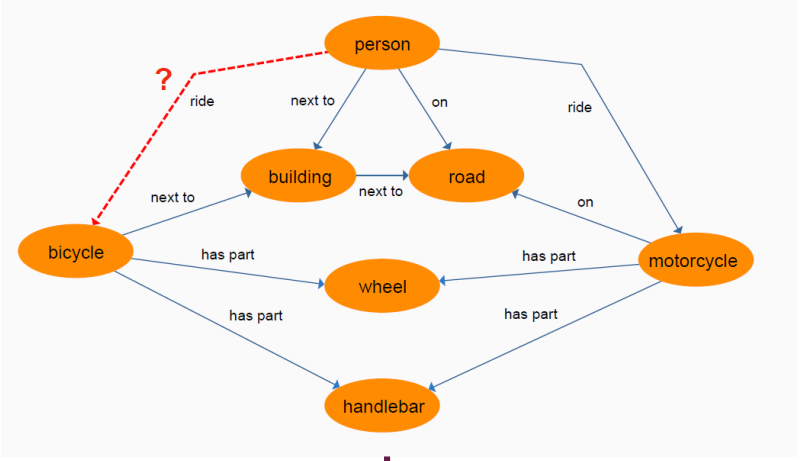

Learning Knowledge Graph

Deep Learning

RCNN

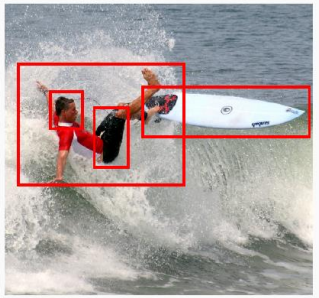

Deep Learning

CNN_e → $P(s|i_s)$

CNN_r → $P(o|i_o)$

$P(p|i_p)$

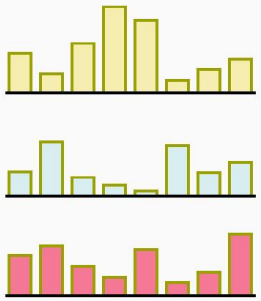

S = Person
P = nextTo
O= Surfboard

- By using a KG prior, we obtained better results than the Stanford group: *Lu, Krishna, Bernstein, Fei-Fei, 2016*
- Best student paper at the ISWC 2017
- *MLwin - Maschinelles Lernen mit Wissensgraphen, Tresp, Schütze, Weikum, Cremers, et al.*

Baier, Ma, Tresp. Improving Visual Relationship Detection using Semantic Modeling of Scene Descriptions, ISWC 2017

Learning Decision Support with Static and Temporal Knowledge Graphs; Integrating Deep Learning

- Decision Support and automated decision making are keys to success in many application domains

* The future of healthcare 2nd October 2017, The Economist (http://thefutureishere.economist.com/thefutureofheathcare-infographic.html)

Static and Temporal Knowledge Graphs for Decision Modeling

Smart Data

- Decision Modeling
- Endpoint Prediction
- sKG, eKG, LSTM, ..

BMWi Smart Data Project:
“Clinical Data Intelligence”

Semantic Knowledge Graph	Episodic Knowledge Graph	Unstructured Data
▪What do I know about the patient?	▪What happened to the patient?	▪Medical Reports ▪Radiology ▪Genetics

Decision Support

▪Best treatment? (Tumor Board)

Smart Data

By disagreement: Physicians find our proposed decisions better than the ones of their peers

	Acceptable alternative	Don‘t agree	Don‘t agree at all
Re-T-Board	11%	64%	23%
ML	33%	58%	8%

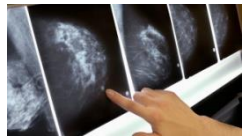

Universitätsklinikum Erlangen

CHARITÉ UNIVERSITÄTSMEDIZIN BERLIN

LMU KLINIKUM DER UNIVERSITÄT MÜNCHEN

age-related macular degeneration

Cognition: a Perspective for AI

- Learning from human cognition

Perception and Memory

Sensor Processing
- Fast, skillful reaction
- Human declarative capabilities
  - Deep understanding of sensory inputs; **declarative decoding**; with a link to language

Episodic memory (“events we remember”)
- Recall a sensory impression of past events
- Human **declarative memory**

Semantic memory (“facts we know”)
- **“**Obama is ex-president of the United States”; “Munich is in Bavaria”
- Human **declarative memory**

More: decisions; prediction; reasoning; action; learning from episodes, ….

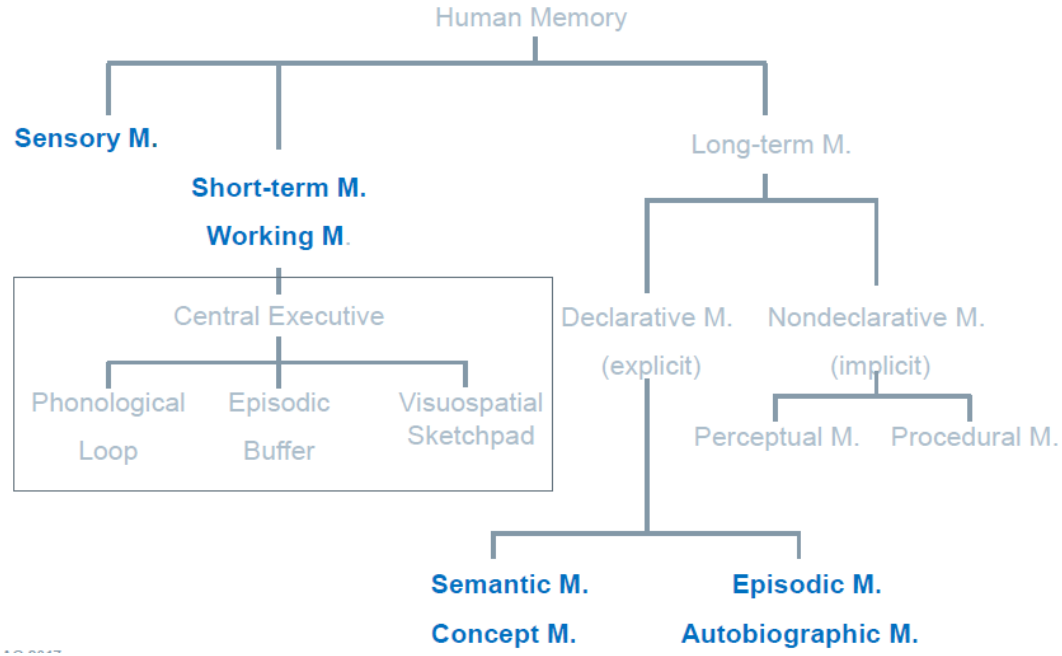

Declarative!

AI is Many Things

Deep Learning with a Cognitive Perspective

Coordination of Thousands of Modules

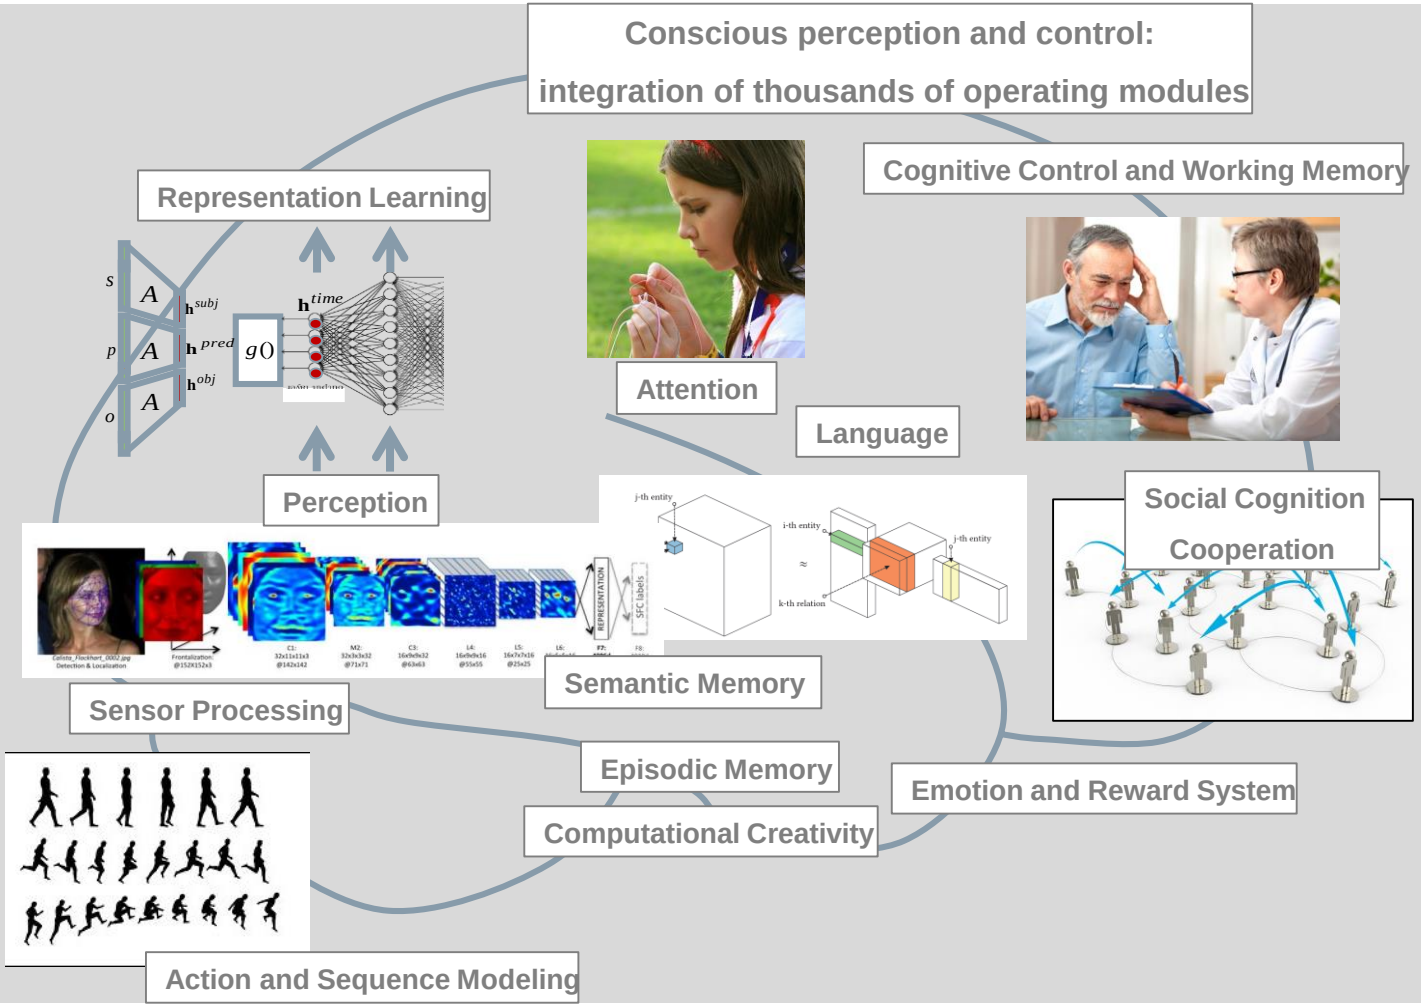

Integrated Intelligence

Integrated Intelligence

Perception

(understand)

Integrated Intelligence

Episodic Memory
(remember)

Perception
(understand)

Integrated Intelligence

Semantic Memory
(know)

Episodic Memory
(remember)

Perception
(understand)

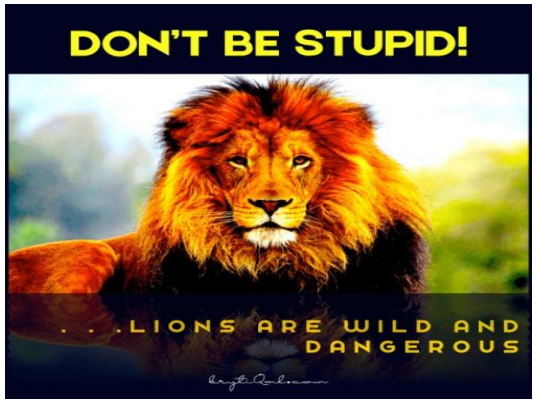

Integrated Intelligence

Semantic Memory
(know)

Episodic Memory
(remember)

Perception
(understand)

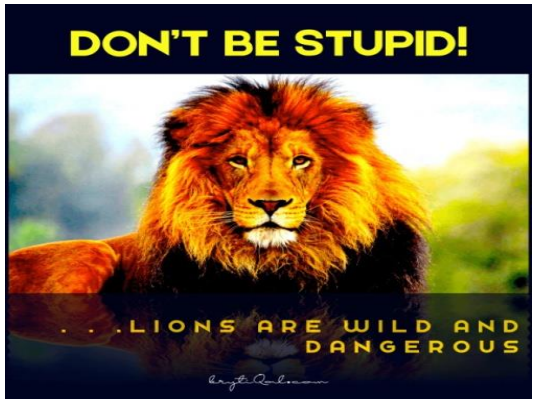

Integrated Intelligence

Working Memory

- Prediction
- Decision Making, Reinforcement Learning
- Planning

Semantic Memory

(know)

Episodic Memory

(remember)

Perception

(understand)

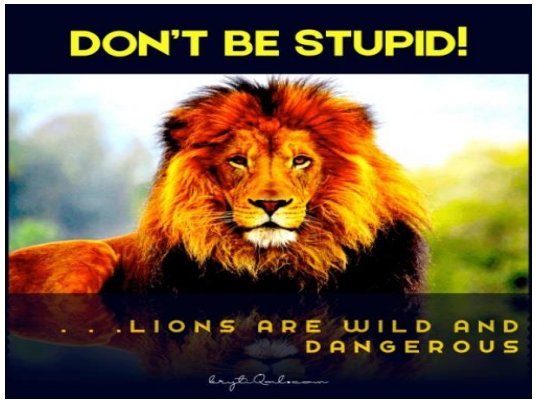

Integrated Intelligence

Working Memory

- Prediction
- Decision Making, Reinforcement Learning
- Planning

Semantic Memory

(know)

Episodic Memory

(remember)

Perception

(understand)

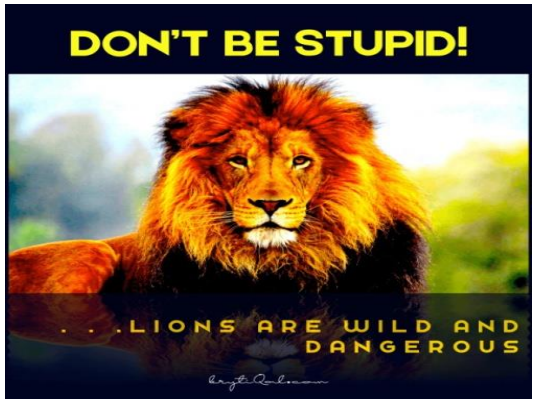

The Future: Quantum Machine Learning?

- The quantum world computes

- …It always bothers me that, according to the laws as we understand them today, it takes **a computing machine an infinite number of logical operations** to figure out what goes on in **no matter how tiny a region of space, and no matter how tiny a region of time**. How can all that be going on in that tiny space?

  – Richard Feynman

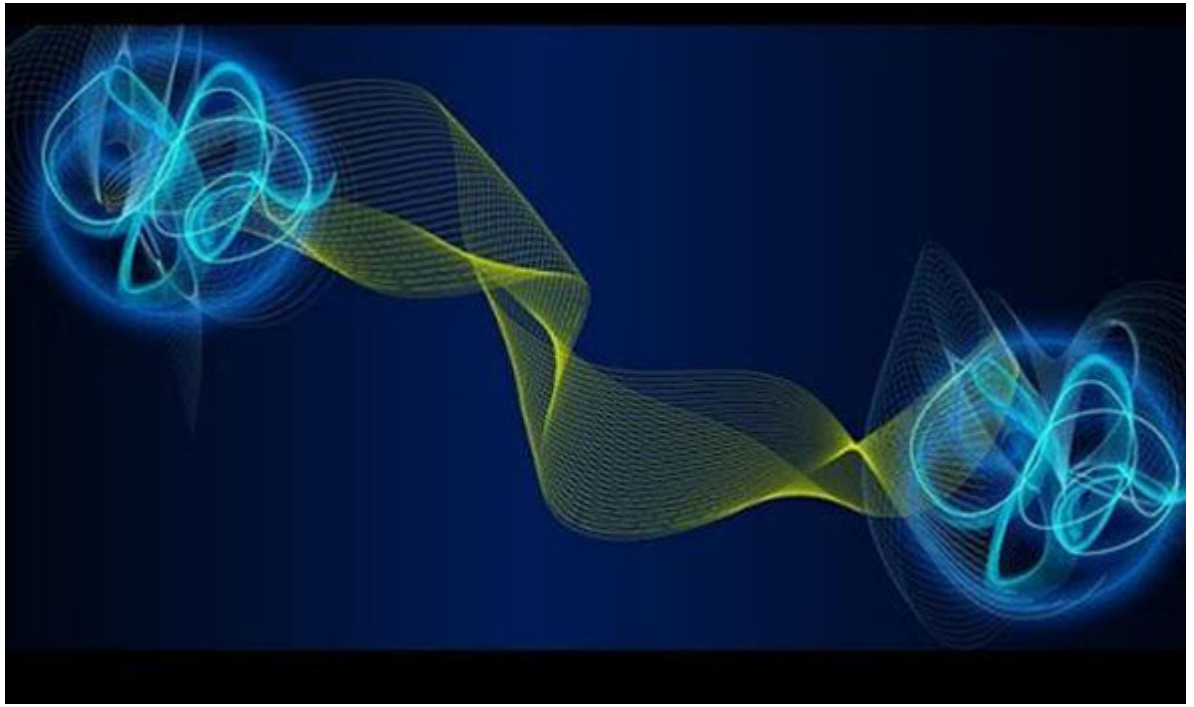

Key Quantum Effects

Superposition

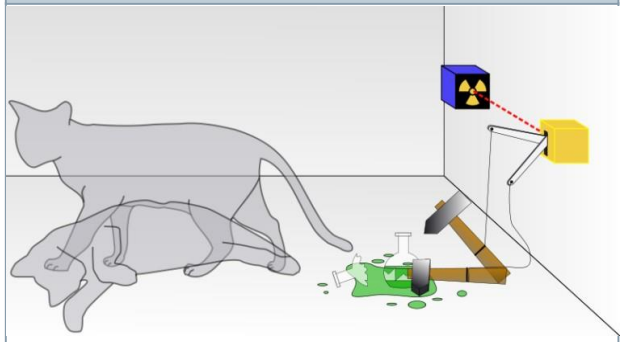

- Schrödinger's cat
- Fundamental principle of quantum mechanics
- Quantum state as sum of other distinct states
- Solution to the Schrödinger equation

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + |1\rangle\right)$$

Entanglement

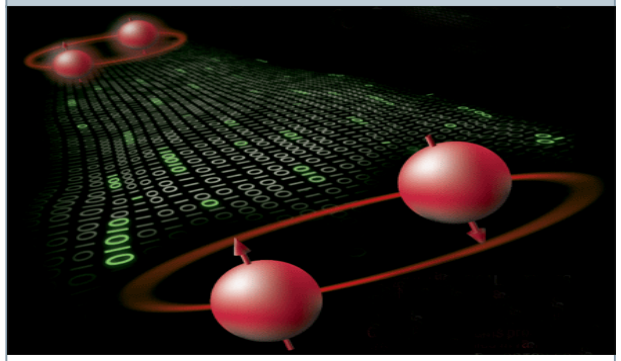

- Fundamentally non-classical phenomenon
- Entangled system as an inseparable whole
- The state of entangled system is in superposition
- Equivalent to quantum non-locality
- “Spooky action at a distance” - Albert Einstein

Quantum Tunneling

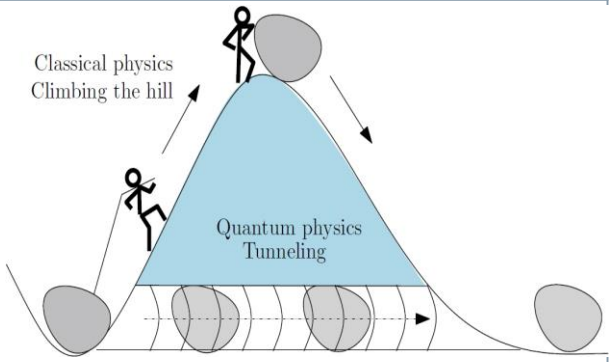

- Particle tunnels through a energy barrier, classically impossible
- Matter in quantum mechanics has properties of waves and particles
- Heisenberg uncertainty principle
- Wave function described by Schrödinger equation

Quantum Machine Learning

I would predict that in 10 years there's nothing but quantum machine learning—you don't do the conventional way anymore.
- *Hartmut Neven, Google*

Accelerate Machine Learning

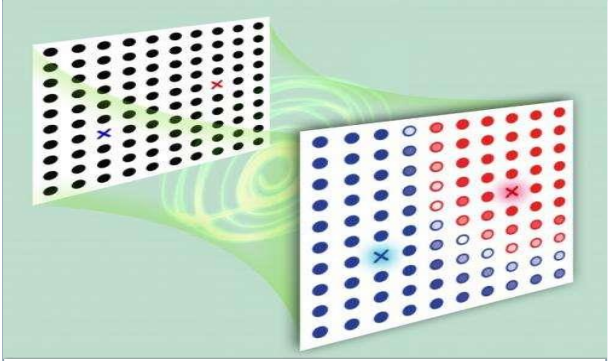

- Exploit entanglement to classify high-dimensional vectors for quantum accelerated ML
- Supervised and unsupervised machine learning
- **Exponential speed-up in solving systems of linear equations (HHL)**
- Represent vectors with quantum states, entangle them before calculating the distance
- Four-qubits entangled state to classify eight-dimensional vectors

Quantum approach to Big Data

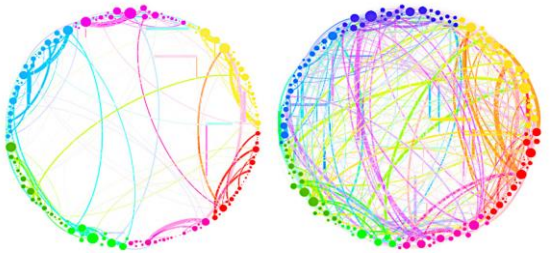

- Quantum analysis on enormous, complex sets of data
- Shown above is quantum analysis involving billions of brain cells
- Useful for analyzing connections in complex networks, e.g. wiring of brain, or power grid
- **Q-based topological analysis provides exponential speed-up**
- E.g. a data set with 300 points, conventional approach need 2^300 units, QC: only 300 Qbits

Quantum effect improves AI

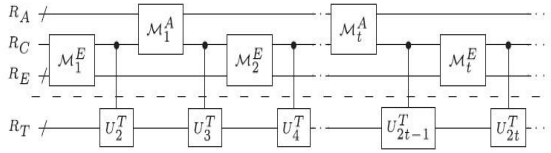

- Quantum effects improve agent-environment interaction
- Learning through interaction (e.g. **reinforcement learning**) amenable to quantum enhancements
- With quantum superposition, agent performs many update steps simultaneously in interactive learning task
- Quantum effects improve agents in automated cars, or smart factories

Thank you!

Prof. Dr. Volker Tresp,
Principal at Siemens Corporate Technology and
Professor at the Ludwig Maximilian University of Munich
volker.tresp@siemens.com