LMU

LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

FACULTY FOR MATHEMATICS, INFORMATICS, AND STATISTICS
DEPARTMENT OF MATHEMATICS
BAVARIAN AI CHAIR “MATHEMATICAL FOUNDATIONS OF ARTIFICIAL INTELLIGENCE”

Reliable and Sustainable AI: From Mathematical Foundations to Next Generation AI Computing

Gitta Kutyniok

Ludwig-Maximilians-Universität München

(also DLR – German Aerospace Center & University of Tromsø, Norway)

ML in PL Conference 2025
Warsaw, October 15 - 18, 2025

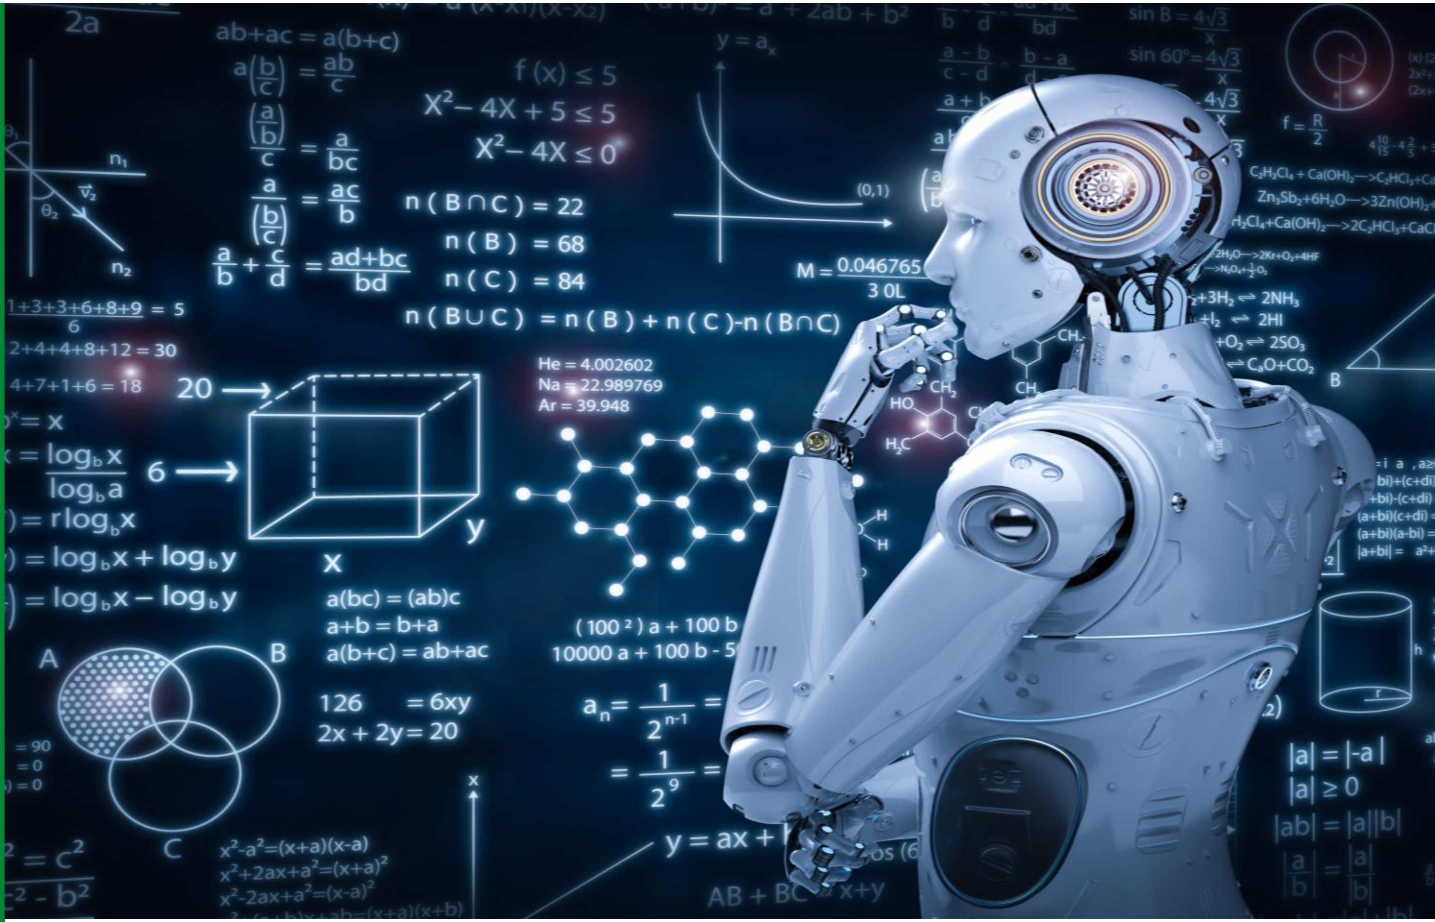

baiosphere the bavarian ai network
DFG Deutsche Forschungsgemeinschaft
Bundesministerium für Forschung, Technologie und Raumfahrt
relAI Konrad Zuse School of Excellence in Reliable AI
MCML Munich Center for Machine Learning
Bayerisches Staatsministerium für Wissenschaft und Kunst
HIGHTECH Agenda Bayern

Fourth Industrial Revolution by Artificial Intelligence

MATH 4AI

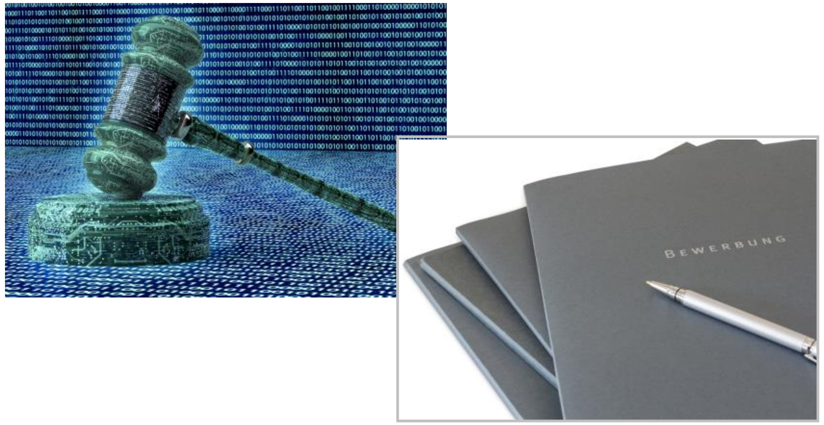

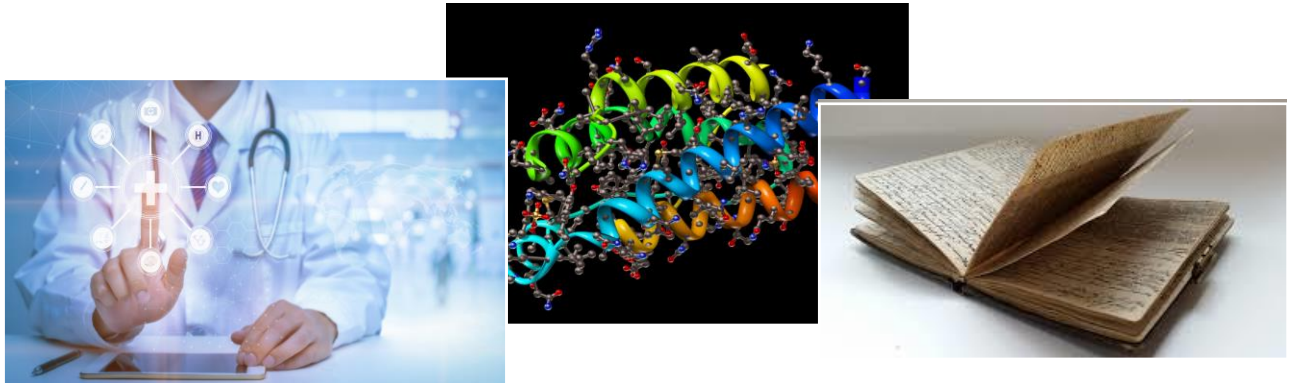

Radical Change of our Society in its Full Breadth!

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Challenges in Artificial Intelligence: Reliability

Problems with Safety

Example:
Accidents involving robots

Problems with Security

Example:
Risks of hacking into AI systems

Problems with Privacy

Example:
Privacy violations of health data

Problems with Responsibility

Example:
Black-box and biased decisions

Current major problem worldwide:

Lack of reliability of AI technology!

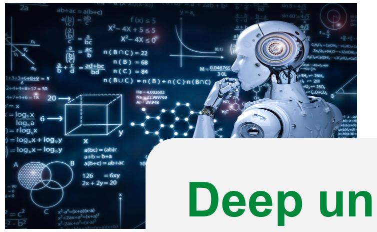

Deep understanding from a mathematical perspective!

Challenges in Artificial Intelligence: Sustainability / Energy Efficiency


Oracle will use three small nuclear reactors to power new 1-gigawatt AI data center

News By Jowi Morales published 8 hours ago

Source: Twitter/X, Sept. 2024

Challenges in Artificial Intelligence: Sustainability / Energy Efficiency

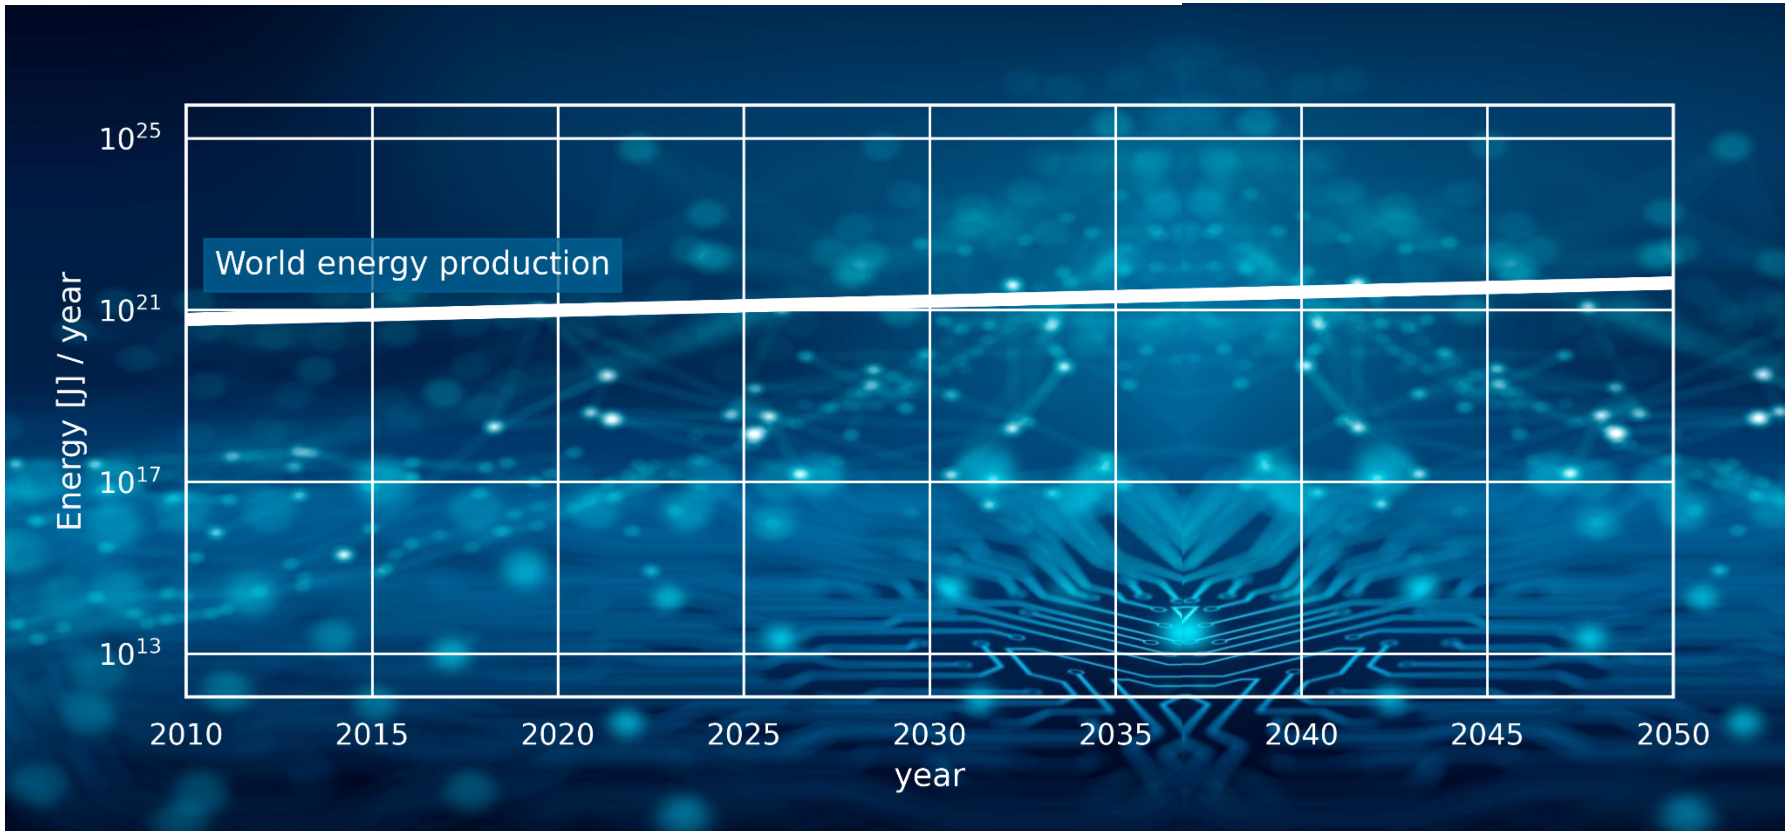

Source: Decadal Plan of the Semiconductor Research Corporation for the Biden (US) Administration, 2021

Challenges in Artificial Intelligence: Sustainability / Energy Efficiency

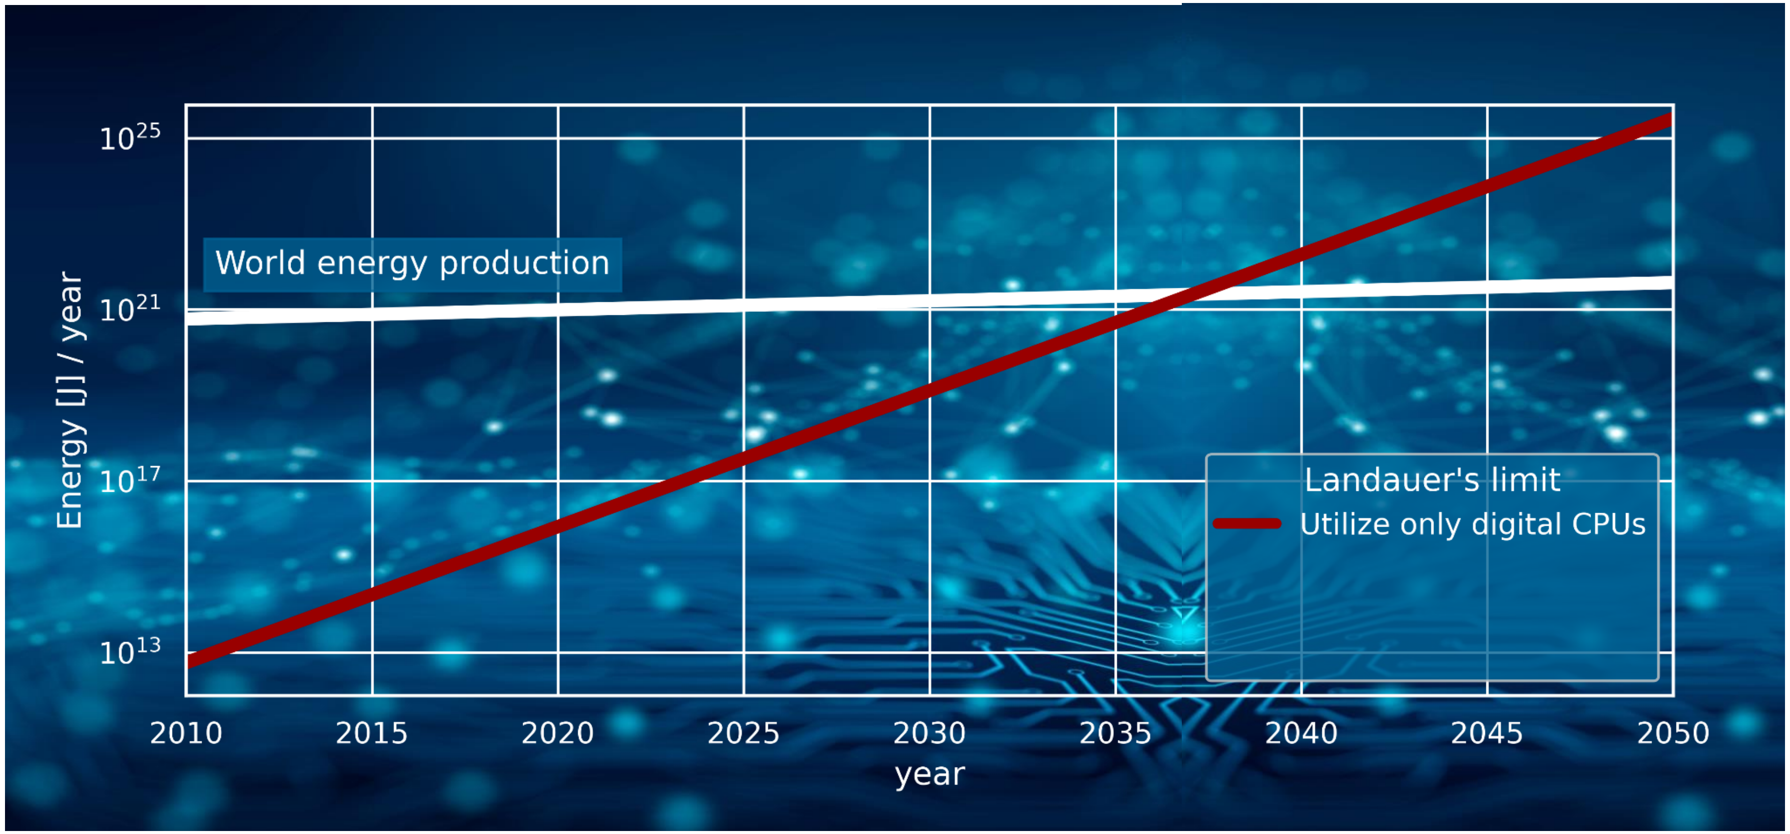

Source: Decadal Plan of the Semiconductor Research Corporation for the Biden (US) Administration, 2021

Challenges in Artificial Intelligence: Sustainability / Energy Efficiency

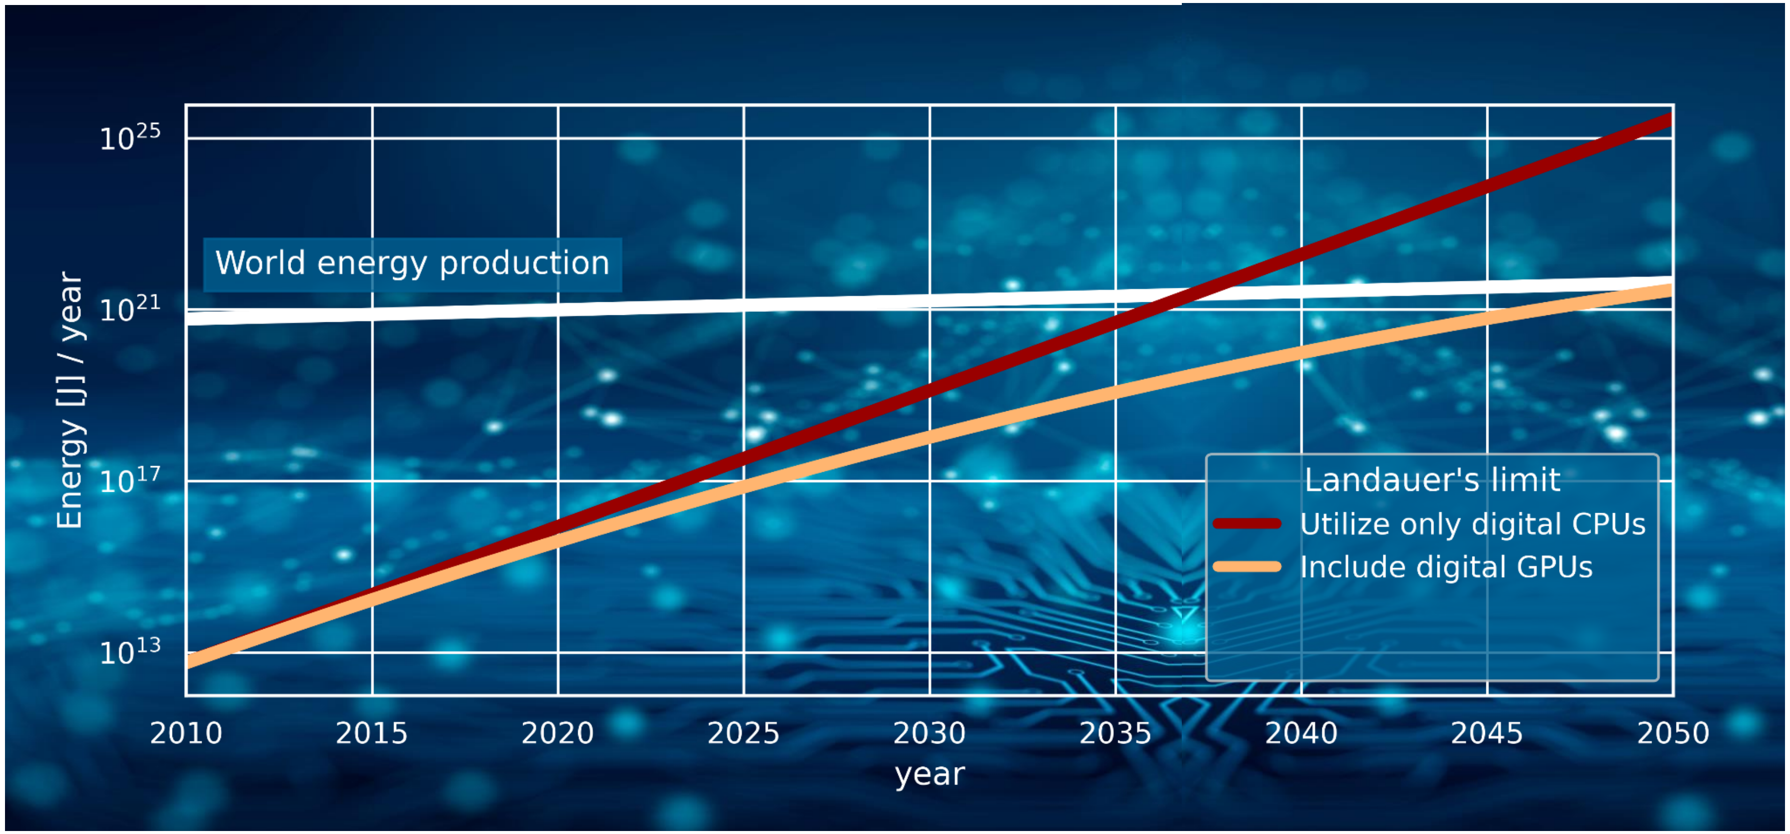

Source: Decadal Plan of the Semiconductor Research Corporation for the Biden (US) Administration, 2021

Challenges in Artificial Intelligence: Sustainability / Energy Efficiency

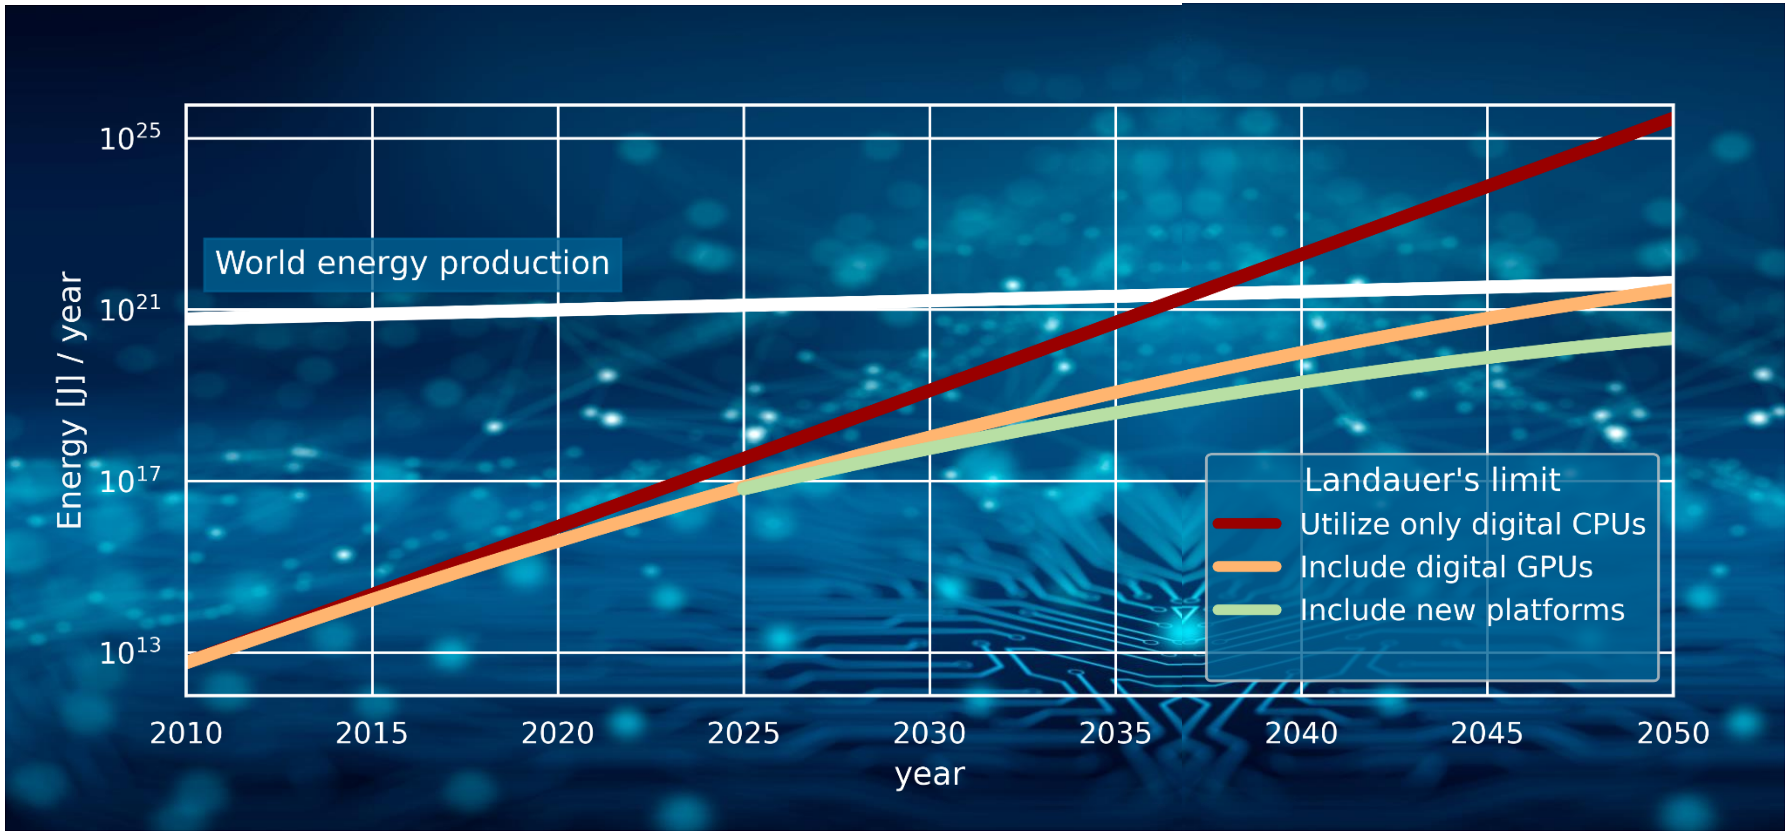

Source: Decadal Plan of the Semiconductor Research Corporation for the Biden (US) Administration, 2021

Challenges in Artificial Intelligence: Sustainability / Energy Efficiency

MATH 4AI

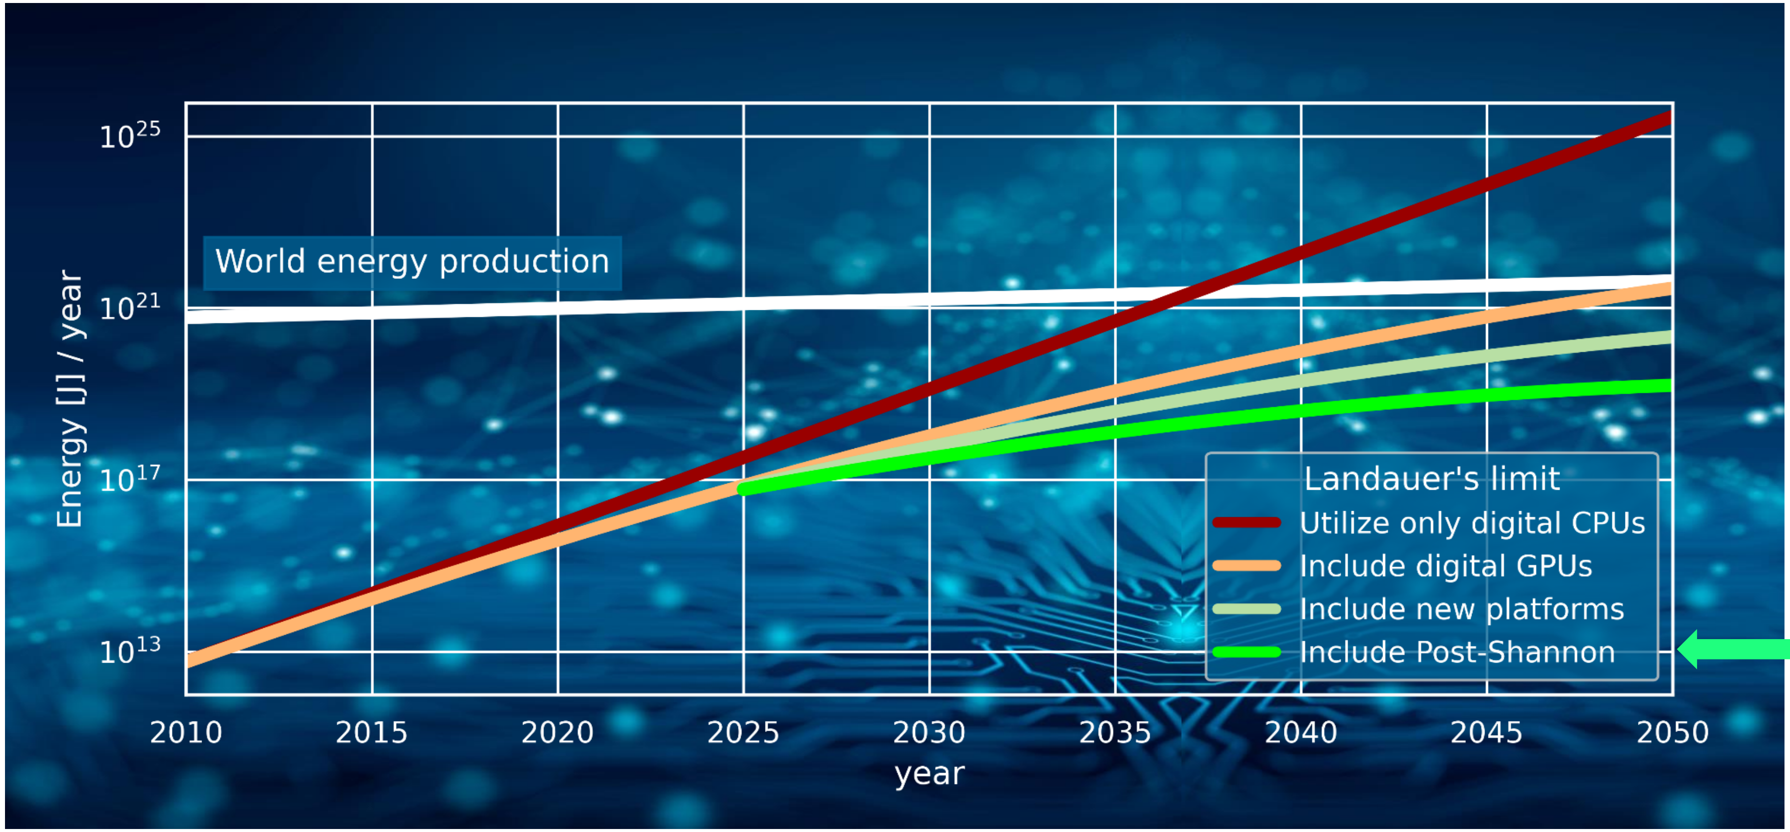

Novel Mathematical Method!

Source: Decadal Plan of the Semiconductor Research Corporation for the Biden (US) Administration, 2021

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Taking a Mathematical Perspective

Deep Neural Networks

Key Goal of McCulloch and Pitts (1943):

→ Introduce *artificial Intelligence!*

Artificial Neurons:

$$f(x_1, \ldots, x_n) = \rho\left(\sum_{i=1}^{n} x_i w_i - b\right)$$

Definition of a Neural Network:

A *deep neural network* is a function $\Phi\colon \mathbb{R}^d \to \mathbb{R}^{N_L}$ of the form

$$\Phi(x) = T_L \rho(T_{L-1}\, \rho(\ldots \rho(T_1(x)) \ldots)), \qquad x \in \mathbb{R}^d,$$

with

$T_l\colon \mathbb{R}^{N_{l-1}} \to \mathbb{R}^{N_l}$, $l = 1, \ldots, L$, where $T_l(x) = W^{(l)} x + b^{(l)}$.

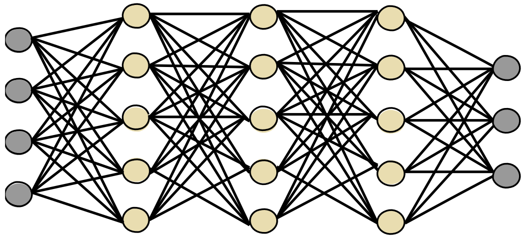

Workflow of Applying Deep Neural Networks

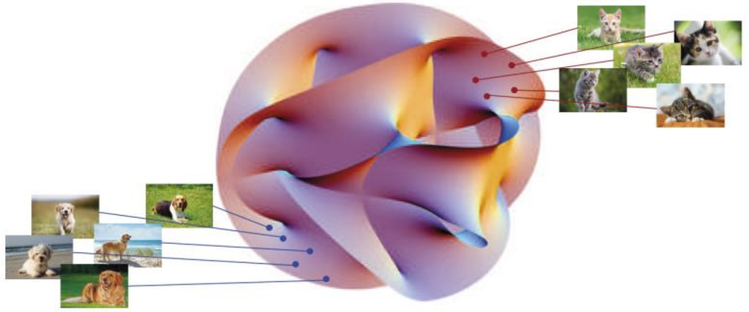

Starting Point :

→ Samples $\left(x_i, f(x_i)\right)_{i=1}^{n}$ of a function $f : \mathcal{M} \to \{1,2, \ldots K\}$.

Split into training- and test data set.

Selection of Architecture:

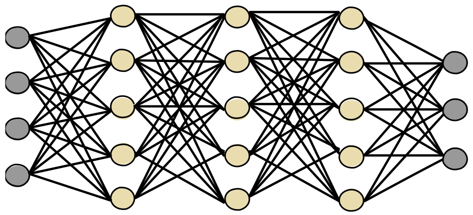

→ Choose the number of layers, the number of neurons in each layer, etc.

Training:

→ Learn the affine-linear functions $T_l(x) = W^{(l)}x + b^{(l)}, l = 1, \ldots, L$ via

$$\min_{\left(W^{(l)}, b^{(l)}\right)_l} \left(\sum_{i=1}^{m} \mathcal{L}\left(\Phi_{\left(W^{(l)}, b^{(l)}\right)_l}(x_i), f(x_i)\right) \right)$$

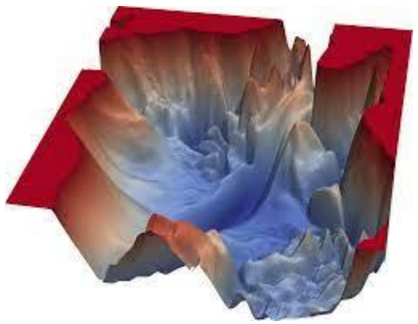

Performance Check:

→ For the test data set: $\Phi_{\left(W^{(l)}, b^{(l)}\right)_l}(x_i) \approx f(x_i)$

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Towards a Mathematical Foundation for Reliable AI

MATH ^{4}AI

Expressivity:

→ Which *aspects of a neural network architecture* affect the performance of AI-systems?

Deriving general guidelines of how to choose the network architecture!

Learning:

→ Why does *stochastic gradient descent* converge to good local minima despite the non-convexity of the problem?

Understanding how to best design the training algorithm!

Generalization:

→ Can we derive an understanding of the *performance on the test data set*?

Providing success guarantees and error bounds!

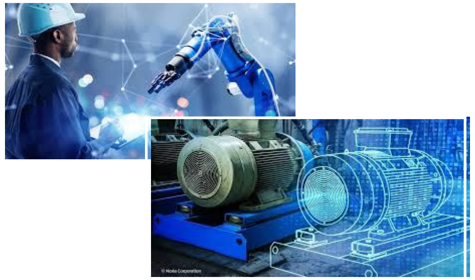

Explainability:

→ Why did a trained deep neural network *reach a certain decision*?

Ensuring trustworthiness and complying with legal regulations!

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

A Glimpse into Generalization: Mathematical Success Guarantees

Graph Neural Networks

Graph neural networks generalize classical neural networks to signals over graph domains.

Graph signal:

$$s: \text{ graph nodes } \to \mathbb{R}^c$$

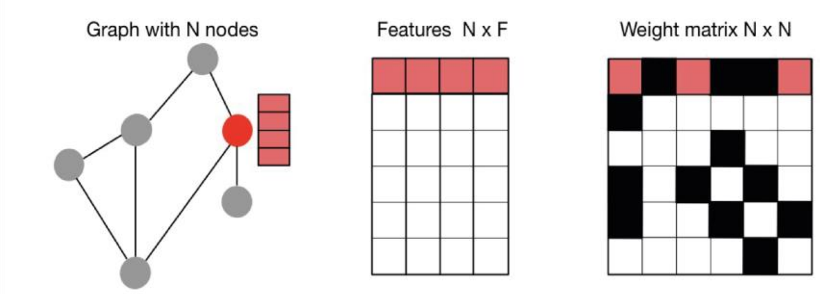

Exemplary Applications:

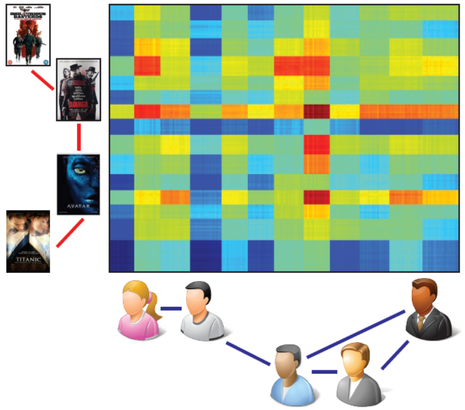

Recommender system

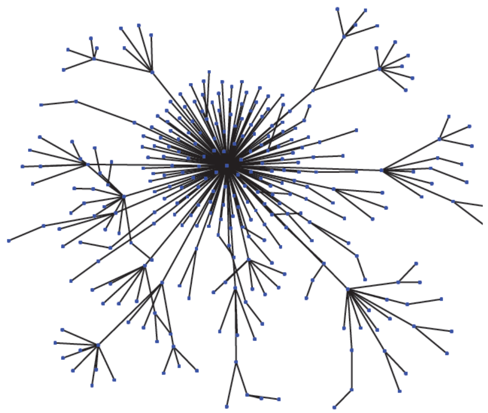

Fake news detection

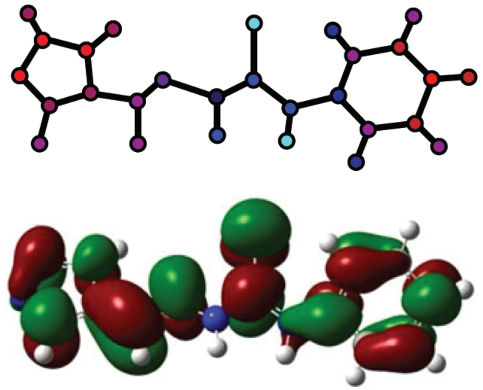

Chemistry

A Special Form of Generalization Capability

General Form of Generalization:

Graph neural networks should *generalize* to graphs and signals unseen in the training set.

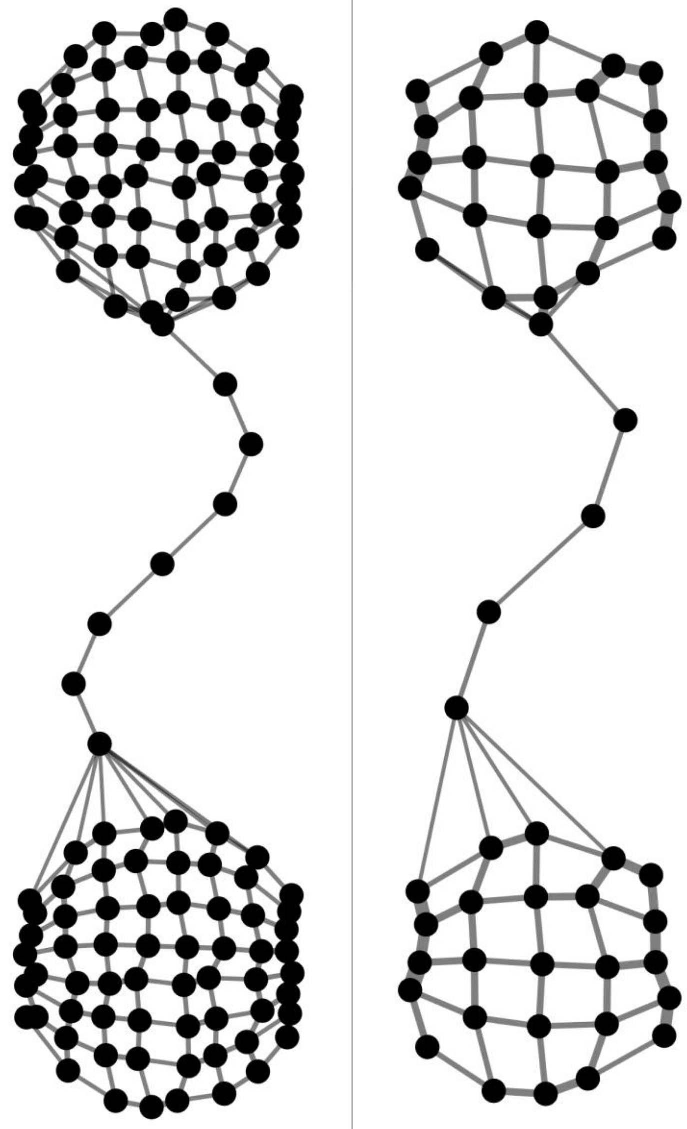

A Special Form of Generalization Capability

General Form of Generalization:

Graph neural networks should *generalize* to graphs and signals unseen in the training set.

The Concept of Transferability:

If two graphs *model the same phenomenon*, a trained graph neural network should have approximately the *same repercussion on both graphs*.

Some Common Approaches:

→ *Metric (Continuum) Space Sampling*

→ Graphon Approach

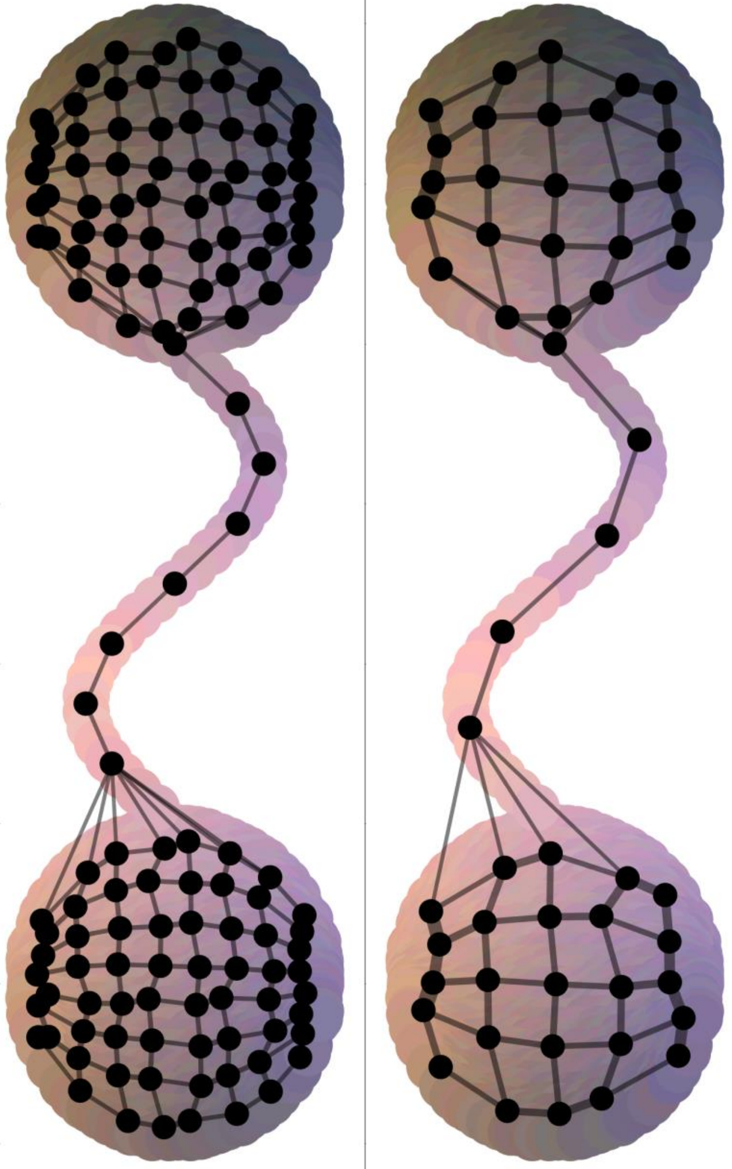

Estimate of Generalization Error

Theorem (Levie, Huang, Bucci, Bronstein, Kutyniok; 2021):

"Generalization error of graph (convolutional) neural network

$\leq$ Transferability error of graph Laplacian + Consistency error"

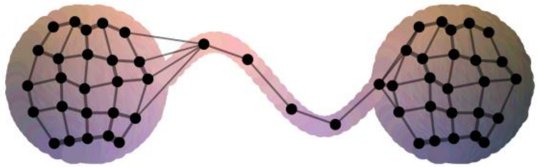

Key Idea:

→ Use graph convolutional neural networks with *specific spectral filters*; this…

- ...solves the instability problem (Levie, Isufi, Kutyniok; 2019)
- ...solves the computational problem for a large class of filters.

→ Introduce *functional analytic framework* akin the Nyquist—Shannon digital signal processing

→ Compare action of graph network on *two similar graphs via metric (continuum) space*

Further Results on Generalization Ability of GNNs

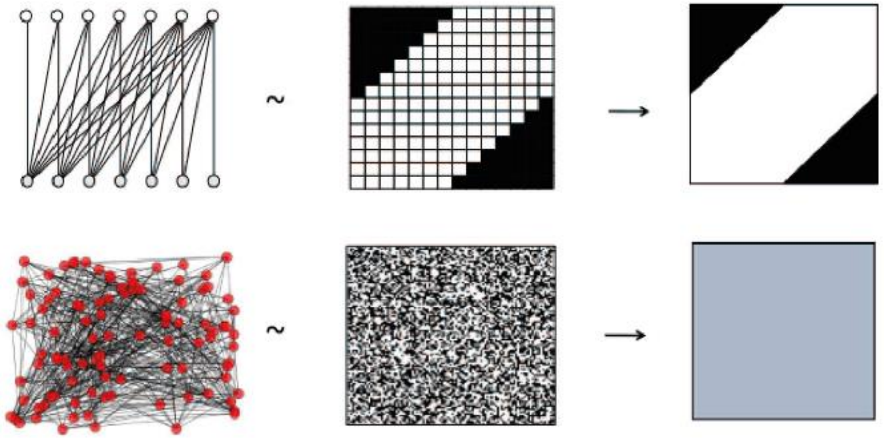

Graph Convolutional Neural Networks:

- *Similar results on transferability* for the *graphon* setting (Maskey, Levie, Kutyniok; 2022 & 2024).
- This builds on (Ruiz, Wang, Ribeiro; 2021).

Message Passing Graph Neural Networks:

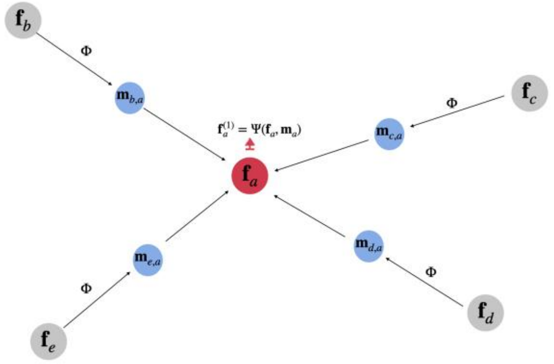

- *Non-asymptotic generalization bounds*, only depending on the regularity of the network and space (Maskey, Levie, Lee, Kutyniok; 2023).
- This builds on (Garg, Jegelka, Jaakkola; 2020), (Verma, Zhang; 2019), (Yehudai, Fetaya, Meirom, Chechik, Maron; 2022).

Towards a Mathematical Foundation for Reliable AI

Expressivity:

→ Which *aspects of a neural network architecture* affect the performance of AI-systems?

Deriving general guidelines of how to choose the network architecture!

Learning:

→ Why does *stochastic gradient descent* converge to good local minima despite the non-convexity of the problem?

Understanding how to best design the training algorithm!

Generalization:

→ Can we derive an understanding of the *performance on the test data set*?

Providing success guarantees and error bounds!

Explainability:

→ Why did a trained deep neural network *reach a certain decision*?

Ensuring trustworthiness and complying with legal regulations!

Explainability: A Mathematical Approach

Some General Thoughts about Explainability

Main Goal: We aim to *understand* decisions of ``black-box" predictors!

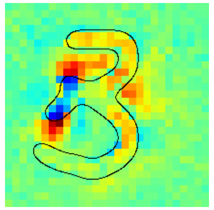

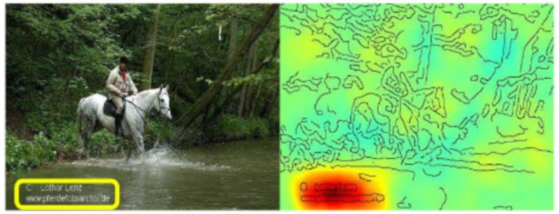

Source: Lapuschkin, Wäldchen, Binder, Montavon, Samek, Müller; 2019)

Selected Questions:

→ What *exactly* is relevance in a mathematical sense?

→ Can we develop a theory for *optimal relevance maps*?

→ Can we derive meaningful *higher level explanations*?

Vision:

Questioning the AI as a human about the reason for a decision!

The explainability approach itself needs to be reliable!

Information Theory: Rate-Distortion Viewpoint

MATH ^{4}AI

The Setting:

→ Let $\Phi : [0,1]^d \to [0,1]$ be a *neural network*.

Alice —— *Part of the Message* ——→ Bob

Message

$\Phi(x) = 0.97$

„Monkey"

Message

Original image x

Partial image S

Distorted Message

Random completion y

Distorted Message

$\Phi(y) = 0.91$

„Monkey"

Expected Distortion:

$$D(S) = D(\Phi, x, S) = \mathbb{E}\left[\frac{1}{2}\left(\Phi(x) - \Phi(y)\right)^2\right]$$

Rate-Distortion Function:

$$R(\epsilon) = \min_{S \subseteq \{1,\ldots,d\}} \{|S| : D(S) \le \epsilon\}$$

Use this viewpoint for the definition of a relevance map!

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Rate-Distortion Explanation (RDE)

Theorem (Wäldchen, Macdonald, Hauch, Kutyniok; 2021):

“Solving this problem is NP^{PP} −complete, even computing an approximation is NP −hard.“

Some Examples:

→ Planning under uncertainties

→ Finding maximum a-posteriori configurations in graphical models

→ Maximizing utility functions in Bayesian networks

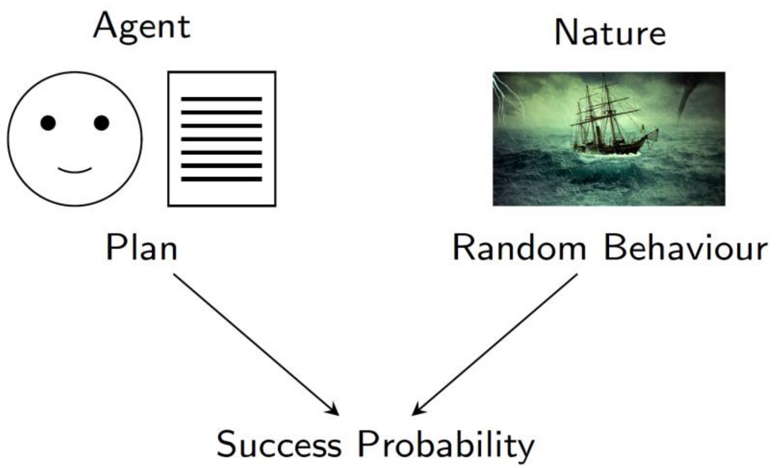

Computable Variant of RDE (Macdonald, Wäldchen, Hauch, Kutyniok, 2020):

$$\text{minimize } D(s) + \lambda\, \|s\|_1 \quad \text{subject to} \quad s \in [0,1]^d$$

…allows rigorous mathematical performance analysis!

STL-10 Experiment

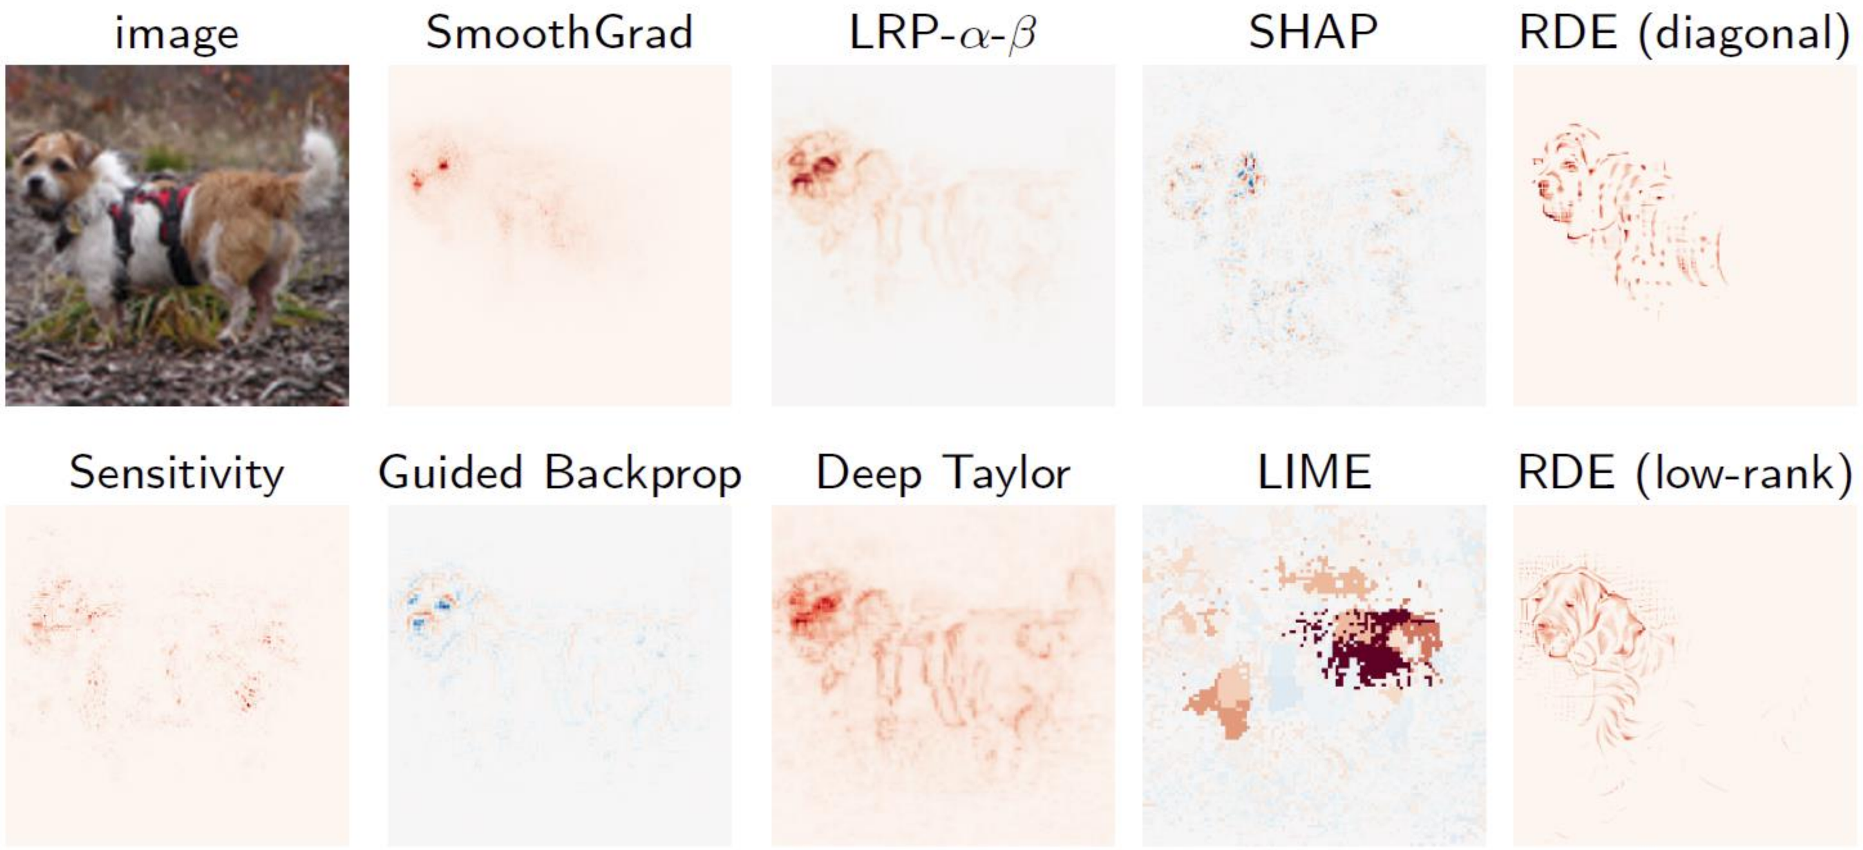

STL-10 Experiment

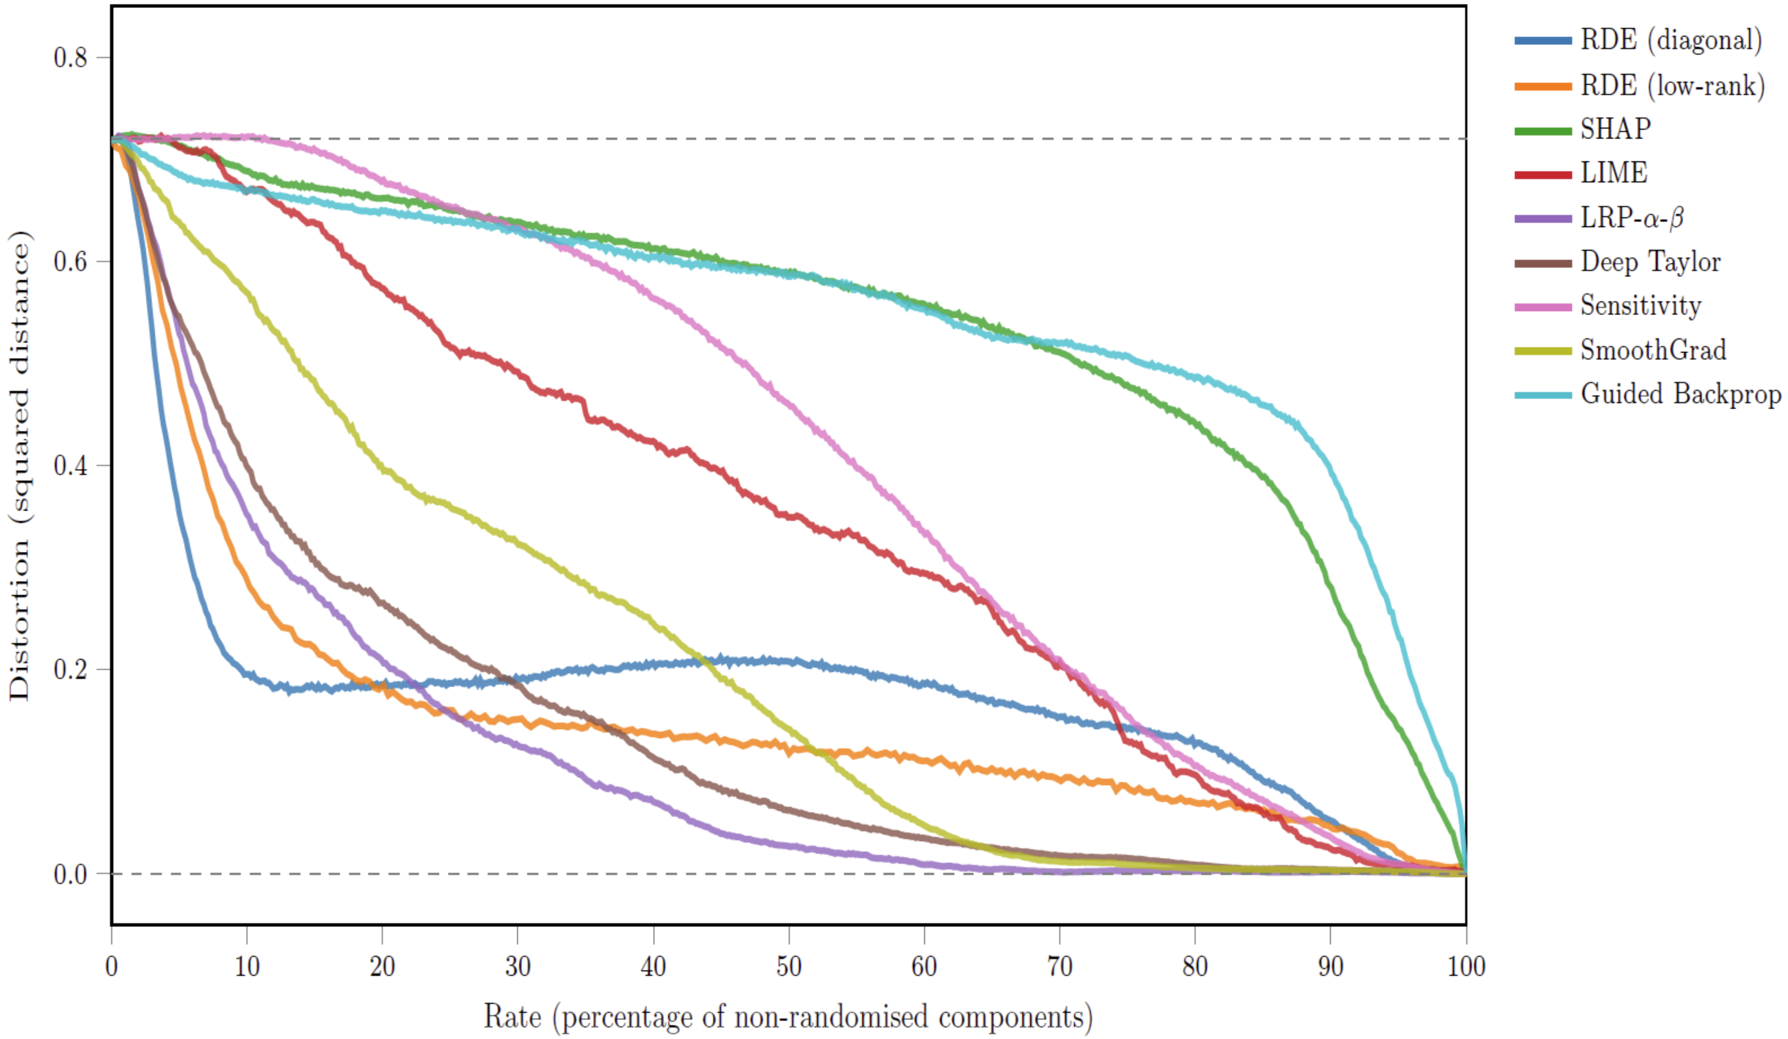

Going Beyond….

Extending to More Realistic Scenarios?

Extension 1 (Heiß, Levie, Resnick, Kutyniok, Bruna; 2020):

→ Choose the obfuscations more natural

→ Example: Apply an inpainting GAN

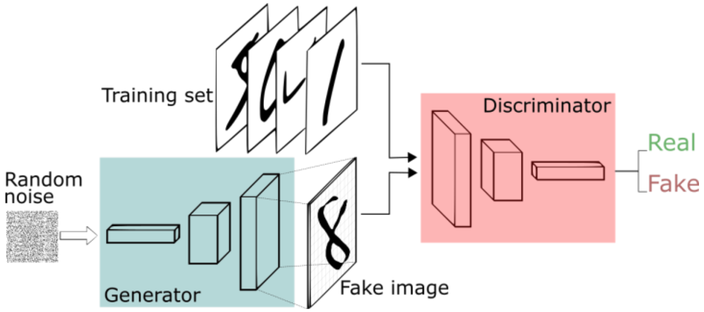

Obtaining Higher-Level Explanations?

Extension 2 (Kolek, Nguyen, Levie, Bruna, Kutyniok; 2021):

→ Apply RDE to decompositions of the data

→ Example: Take a wavelet decomposition of an image.

→ *CartoonX*

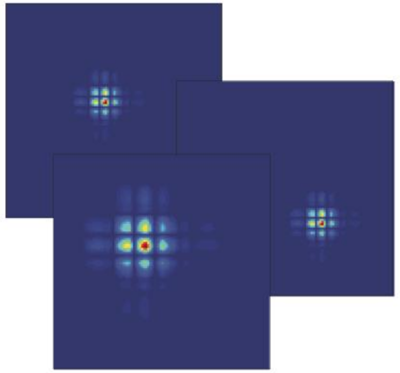

Idea of CartoonX (Kolek, Nguyen, Levie, Bruna, Kutyniok; 2022)

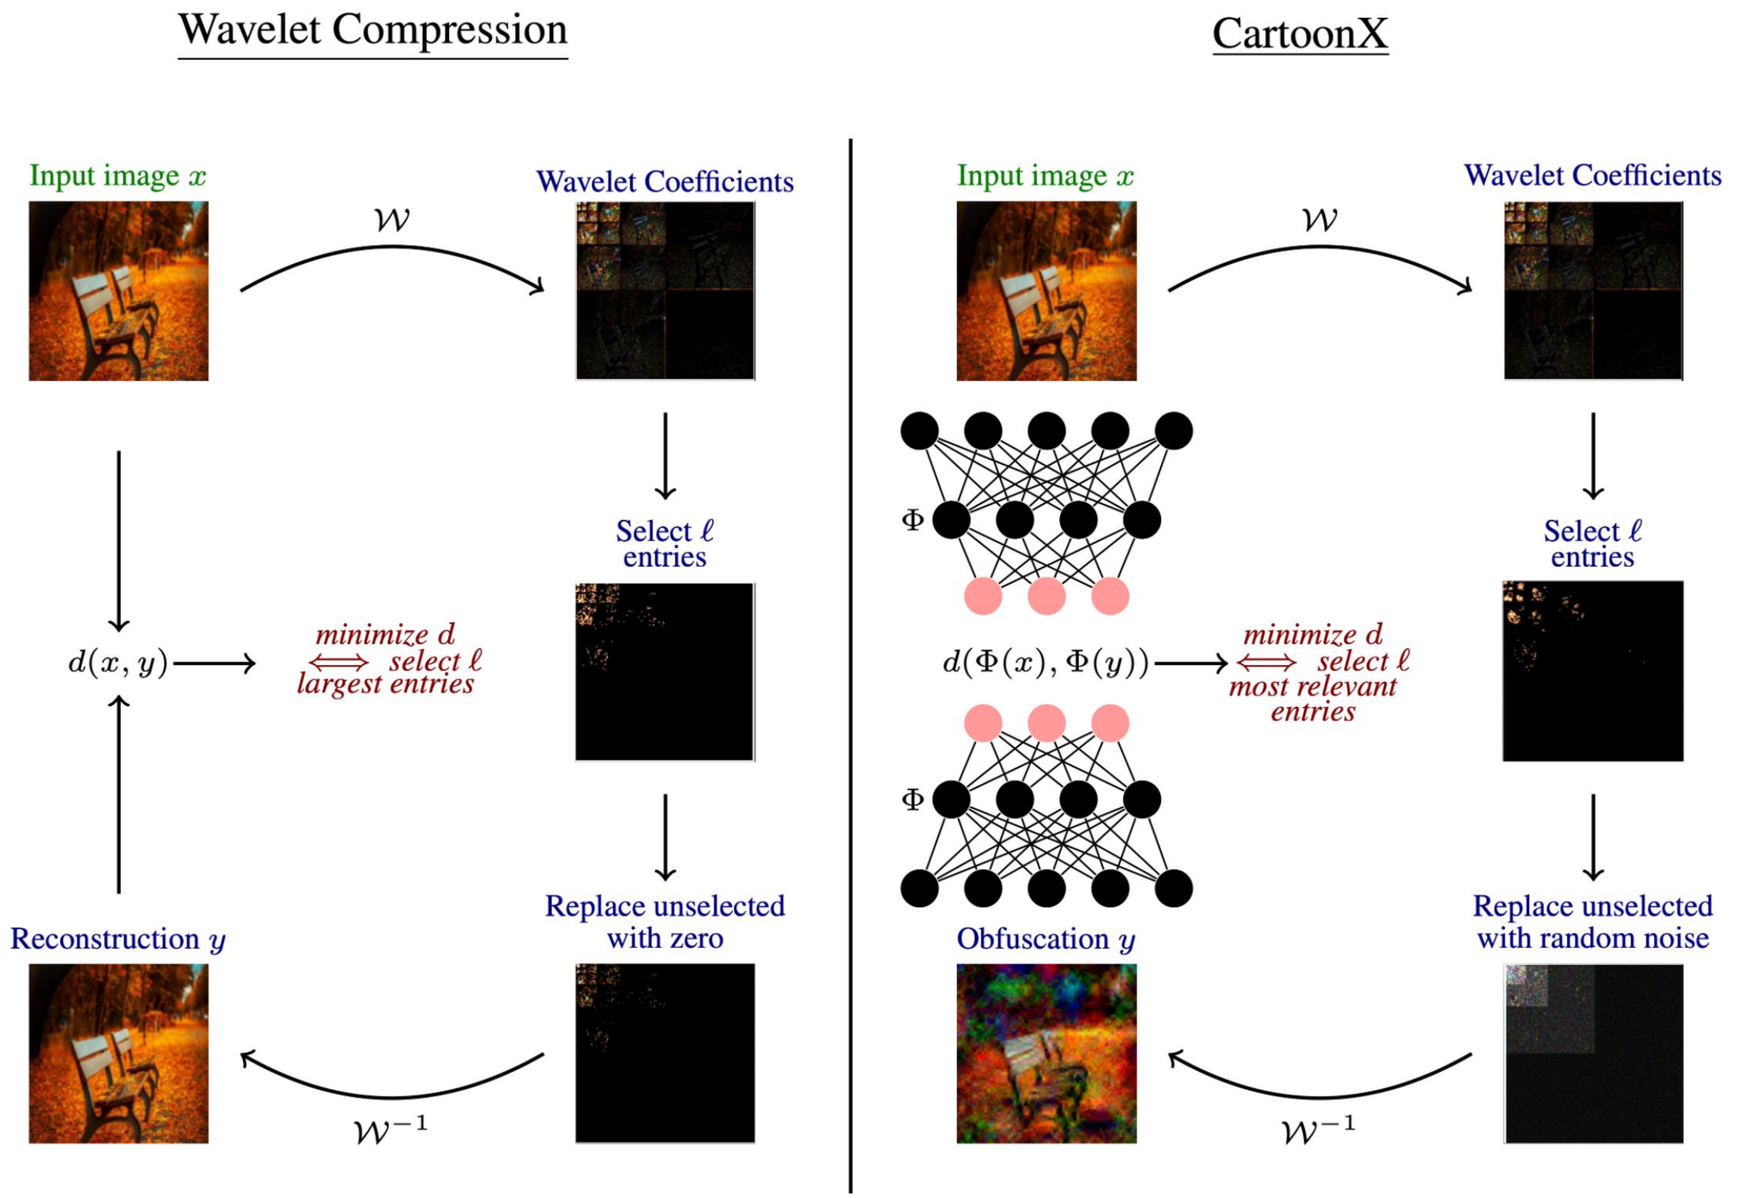

Explainability: Understanding Seemingly Wrong Decisions

Example from Telecommunication:

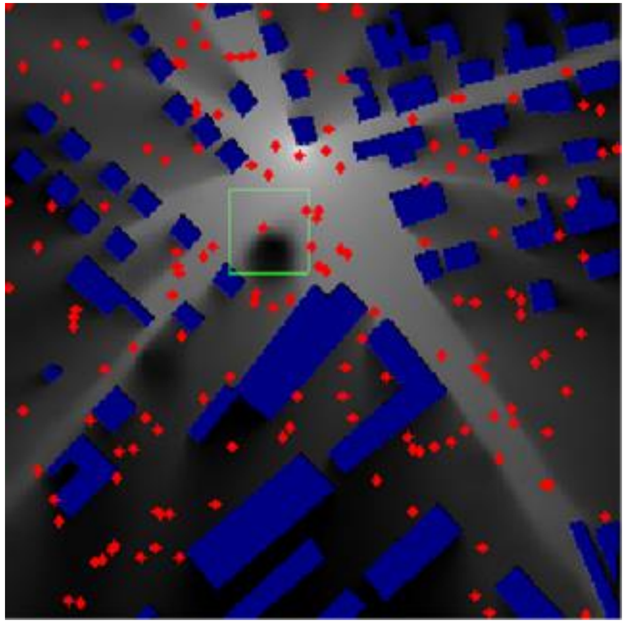

Estimated RadioMap via RadioUNet

(Levie, Cagkan, Kutyniok, Caire; 2020)

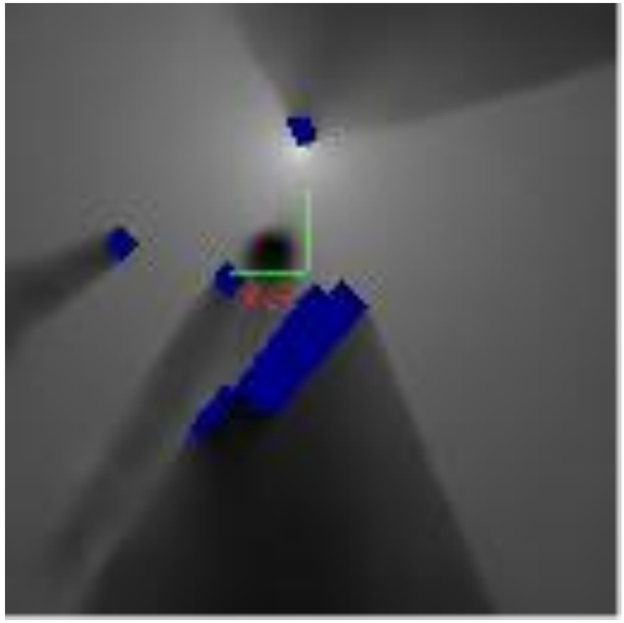

Rate-Distortion Explanation

(Heiß, Levie, Resnick, Kutyniok, Bruna; 2020):

Explainability: Understanding Wrong Decisions

Example from Imaging:

Wrong decision by AI:

Diaper

Wrong decision by AI:

Screw

Explainability: Understanding Wrong Decisions

Example from Imaging:

Explanation by CartoonX

(Kolek, Nguyen, Levie, Bruna, Kutyniok; 2021)

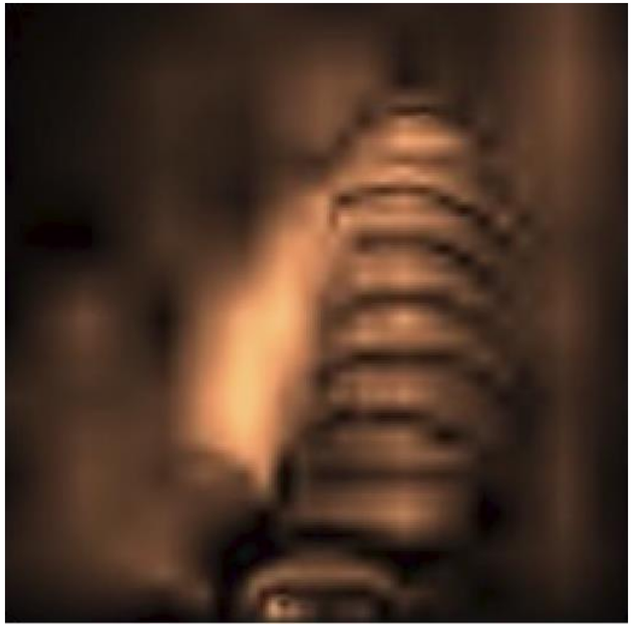

Explanation by CartoonX

Extension: ShearletX (Kolek, Windesheim, Loarca, Kutyniok, Levie; 2023)!

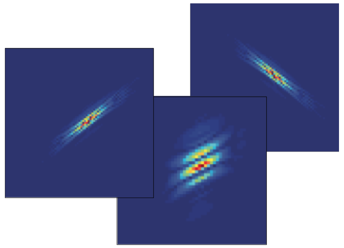

Mathematical Underpinning: Ensuring Reliability

Problem:

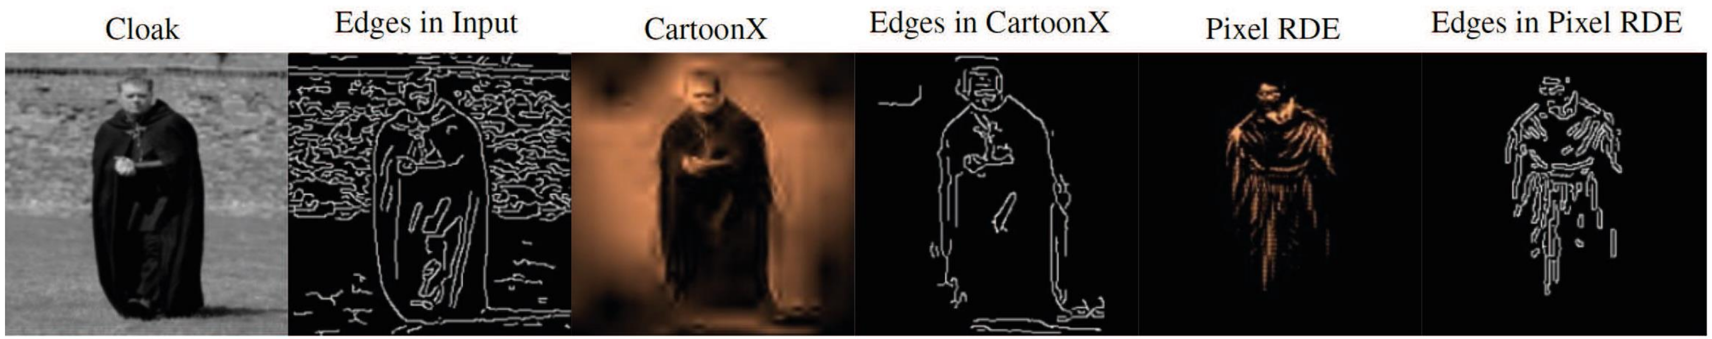

Theorem (Kolek, Windesheim, Loarca, Kutyniok, Levie; 2023):
Let $x \in L^2[0,1]^2$ be an image modeled as a L^2-function. Let m be a bounded mask on the shearlet coefficients of x and let y be the image masked in shearlet space with mask m. Then, we have

$$WF(y) \subset WF\ (x)$$

and thus masking in shearlet space does not create new edges.

Towards a Mathematical Foundation for Reliable AI

Expressivity:

→ Which *aspects of a neural network architecture* affect the performance of AI-systems?

Deriving general guidelines of how to choose the network architecture!

Learning:

→ Why does *stochastic gradient descent* converge to good local minima despite the non-convexity of the problem?

Understanding how to best design the training algorithm!

Generalization:

→ Can we derive an understanding of the *performance on the test data set*?

Providing success guarantees and error bounds!

Explainability:

→ Why did a trained deep neural network *reach a certain decision*?

Ensuring trustworthiness and complying with legal regulations!

A Glimpse into Problems of Compliance with the EU AI Act

Challenges in Artificial Intelligence: EU AI Act

Exemplary Requirements from the EU AI Act:

- Article 43: Conformity Assessment
- Article 50: Transparency Obligations for Providers and Deployers
- Article 86: *Right to Explanation* of Individual Decision-Making

Current Danger:

- *Enormous costs* for small-size companies and start-ups.
- *Uncertainty and potential disadvantage* in Europe

Differential Privacy (Formalization of "Privacy"):

The algorithm $\mathcal{A}$ is said to provide *ε-differential privacy* if, for all datasets D_1 and D_2 that differ on a single element, and all subsets S of $\operatorname{im}(\mathcal{A})$:

$$\frac{P(\mathcal{A}(D_1) \in S)}{P(\mathcal{A}(D_2) \in S)} \leq \varepsilon.$$

Algorithmic Transparency (Boche, Fono, Kutyniok; 2024):

An algorithmic implementation is *transparent* in a *given computing model* if the realization $\mathcal{A}_f$ of some function $f : \mathbb{R}^m \to \mathbb{R}^n$ by an algorithm $\mathcal{A}$ is not altered by its implementation in the computing model. We then say that f allows for a *transparent algorithmic implementation* in the given computing model.

A „Formalization" of the legal requirements of the EU AI Act would allow a fair, low-cost, and automatic verification!

Research Project of the Bavarian AI Act Accelerator

MATH ⁴AI

Bayerisches Staatsministerium für Digitales

appliedAI institute for europe

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

TUM

UTN

Exemplary Requirements from the EU AI Act:

- → Article 43: Conformity Assessment
- → Article 50: Transparency Obligations for Providers and Deployers
- → Article 86: *Right to Explanation* of Individual Decision-Making

Current Danger:

- → *Enormous costs* for small-size companies and start-ups
- → *Uncertainty and potential disadvantage* in Europe

Differential Privacy (Formalization of "Privacy"):

The algorithm $\mathcal{A}$ is said to provide *ε-differential privacy* if, for all datasets D_1 and D_2 that differ on a single element, and all subsets S of $\mathrm{im}(\mathcal{A})$:

$$\frac{P(\mathcal{A}(D_1) \in S)}{P(\mathcal{A}(D_2) \in S)} \leq \varepsilon.$$

Algorithmic Transparency (Boche, Fono, Kutyniok; 2024):

An algorithmic implementation is *transparent* in a *given computing model* if the realization $\mathcal{A}_f$ of some function $f : \mathbb{R}^m \to \mathbb{R}^n$ by an algorithm $\mathcal{A}$ is not altered by its implementation in the computing model. We then say that f allows for a *transparent algorithmic implementation* in the given computing model.

A „Formalization" of the legal requirements of the EU AI Act would allow a fair, low-cost, and automatic verification!

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

EU AI Act: The Role of the Computing Platform

Theorem (Boche, Fono, Kutyniok; 2024):

There exists an algorithm $\mathcal{A}$ with *transparent implementation in the Turing model* realizing $\mathcal{A}_f$ if and only if $f : \mathbb{R}^m \to \mathbb{R}^n$ is Borel-Turing computable.

Theorem (Boche, Fono, Kutyniok; 2024):

There exists an algorithm $\mathcal{A}$ *with transparent implementation in the analog (Blum-Shub-Smale) model* realizing $\mathcal{A}_f$ if and only if $f : \mathbb{R}^m \to \mathbb{R}^n$ is analog (BSS) computable.

Digital hardware can also cause problems of compliance with the EU AI Act!

Towards a Mathematical Foundation for Reliable AI

Expressivity:

→ Which *aspects of a neural network architecture* affect the performance of AI-systems?

Deriving general guidelines of how to choose the network architecture!

Learning:

→ Why does *stochastic gradient descent* converge to good local minima despite the non-convexity of the problem?

Understanding how to best design the training algorithm!

Generalization:

→ Can we derive an understanding of the *performance on the test data set*?

Providing success guarantees and error bounds!

Explainability:

→ Why did a trained deep neural network *reach a certain decision*?

Ensuring trustworthiness and complying with legal regulations !

…toward the core of the reliability and sustainability problem!

Computing in the 21th Century

Importance of Computing:

- Digital Transformation ⟶ *Ubiquitous Computing*
- (Generative) AI ⟶ *Large-Scale Computing*
- Virtual Reality ⟶ *Fast Computing*
- Information and Communication Technology (ICT) ⟶ *Distributed Computing*

Computing is the heart of modern technology, powering innovation, transforming industries, and shaping the future of our society!

Reliable and Sustainable AI:
The Need to Rethink Current Computing!

Are There Fundamental Limitations to Be Aware Of?

…Delving Deeper!

What can actually be *computed on digital hardware*?

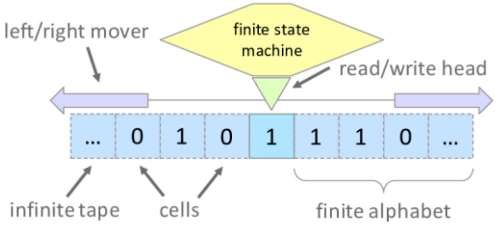

Turing-Machine

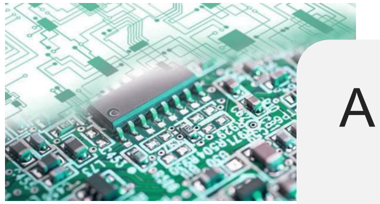

A *computable problem (function)* is one for which the input-output relation can be computed on a digital machine for any given accuracy.

What about Non-Computability?

Non-computable problems can be tackled successfully in practice, if limited precision succeeds!

But we have no guarantees of correctness, hence no reliability!

Very Disappointing News

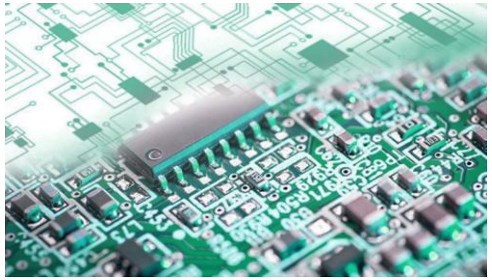

Theorem (Boche, Fono, Kutyniok; 2023):
The solution of a finite-dimensional inverse problem is *not (Turing-)computable* (by a deep neural network).

Solution Set: For $A \in \mathbb{C}^{m\times N}$ and $y \in \mathbb{C}^m$ let

$$\Psi(A,y) := \underset{x\,\in\,\mathbb{C}^N}{\text{arg min}}\, \|x\|_1 \text{ such that } \|Ax-y\|_2 \le \epsilon.$$

Theorem (Boche, Fono, Kutyniok; 2023):
Fix parameters $\epsilon \in \left(0,\frac{1}{4}\right)$, $N \ge 2$, and $m < N$. There does *not exist a (Turing-) computable function* $\widehat{\Psi} : \mathbb{C}^{m\times N} \times \mathbb{C}^m \to \mathbb{C}^N$ such that

$$\sup_{(A,y)\,\in\,\mathbb{C}^{m\times N}\times\mathbb{C}^m} \left\|\Psi(A,y) - \widehat{\Psi}(A,y)\right\|_2 < \frac{1}{4}.$$

MATH 4AI

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

More Problems with Digital Hardware

Theorem (Boche, Fono, Kutyniok; 2023):
Many classification problems are also *not (Turing) computable*!

Theorem (Boche, Fono, Kutyniok; 2023):
The Pseudo Inverse is *not (Banach-Mazur) computable*!

Theorem (Bacho, Boche, Kutyniok; 2024):
Computing the solutions to the Laplace and the diffusion equation on digital hardware causes a *complexity blowup*.

Theorem (Lee, Boche, Kutyniok; 2024):
Finding the solution of most optimization problems is *not (Turing-)computable*; it can *not even be approximated* by a Turing computable function!

What now?

Theory tells us…

Theorem (Boche, Fono, Kutyniok; 2024):
The solution of a finite-dimensional inverse problem is *computable* (by a deep neural network) on an *analog (Blum-Shub-Smale) machine*!

Reliability for certain problem settings requires novel hardware!

Exciting Future Developments:

→ Neuromorphic computing

→ Biocomputing

→ Quantum computing

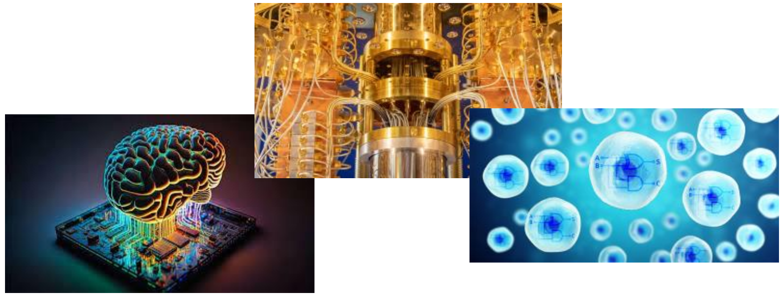

Highly energy efficient!

EcoLogic Computing

https://www.ecologic-computing.com

Reliable and Sustainable AI…by Next Generation AI Computing!

LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

Next Generation AI Computing

Neuromorphic Hardware

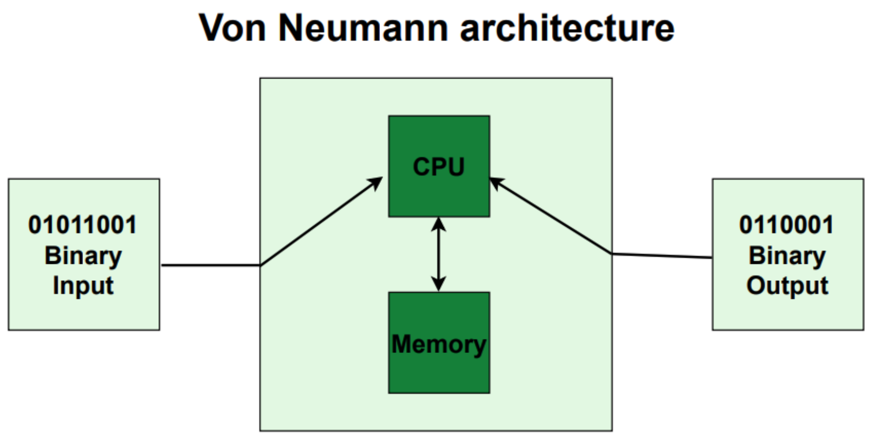

Features of Neuromorphic Hardware:

- → Closer to the human brain.
- → Energy efficiency.
- → Execution speed.
- → Robustness.
- → ….

What is the correct type of neural network?

The Framework of Spiking Neural Networks

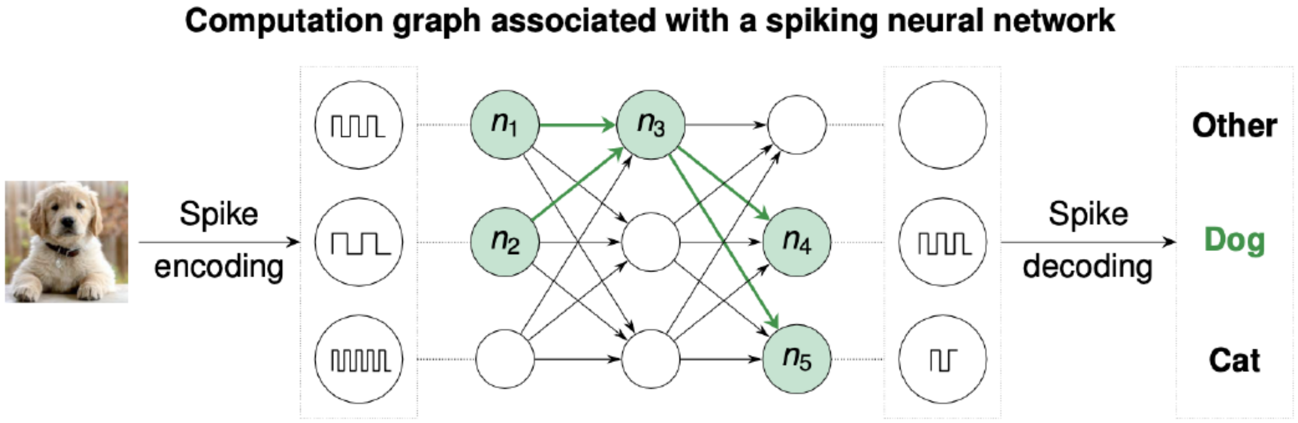

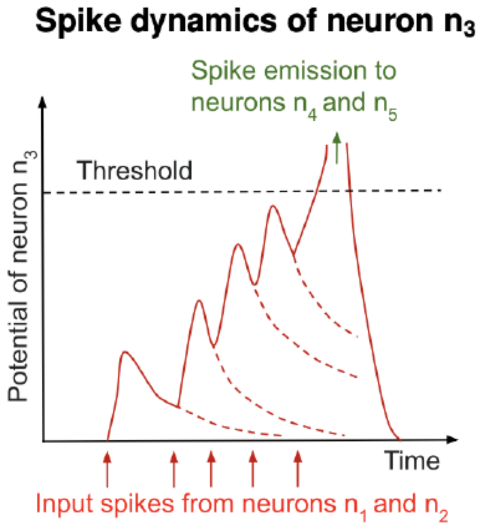

Remarks:

→ *More biologically realistic* than first and second generation artificial neurons.

→ Information is encoded in the *timing of individual spikes.*

→ Numerous models for spiking neurons exist; one of those is the *Spike Response Model.*

Time is one crucial factor in this model!

Our Focus: Expressivity

How expressive are spiking neural networks compared to classical networks?

Function approximation:

How well do the realizations of a neural network approximate a target function?

Number of linear regions:

How does the network partition the input space, affecting decision boundaries?

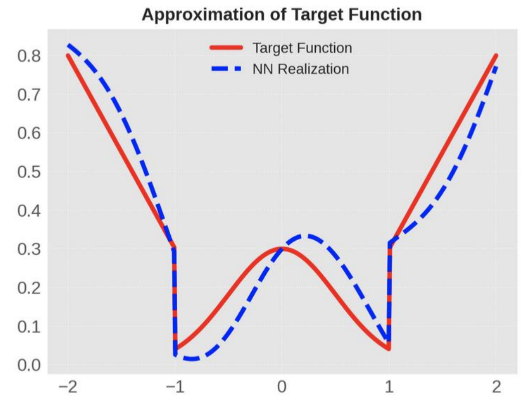

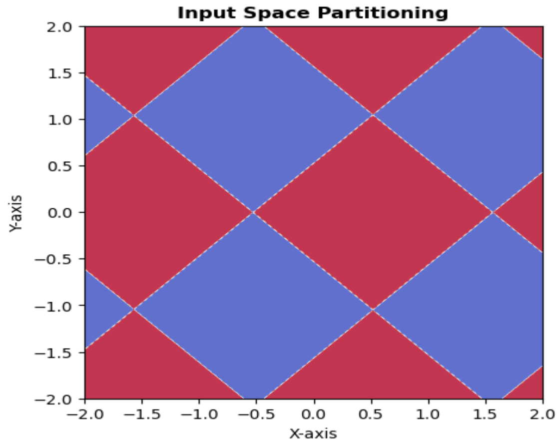

Generalization: Neuman, Dold, Petersen; 2024

The Spike Response Model (SRM)

Definition: A *SRM (spiking neural) network* Φ is a directed graph (V, E) and consists of a finite set V of spiking neurons, a subset $V_{in} \subset V$ of input neurons, and a set $E \subset V \times V$ of synapses. Each *synapse* $(u, v) \in E$ is associated with

→ a *synaptic weight* $w_{uv} \geq 0$,

→ a *synaptic delay* $d_{uv} \geq 0$,

→ and a *response function* $\epsilon_{uv} : \mathbb{R} \to \mathbb{R}$.

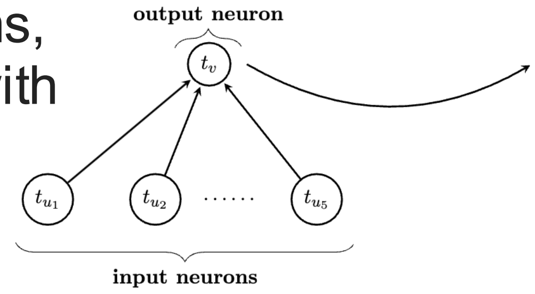

Each neuron $v \in V \setminus V_{in}$ is associated with

→ a *firing threshold* $\theta_v > 0$,

→ and a *membrane potential* $P_v : \mathbb{R} \to \mathbb{R}$,

which is given by

$$P_v(t) = \sum_{(u,v)\in E} \sum_{t_u^f \in F_u} w_{uv} \epsilon_{uv}(t - t_u^f)$$

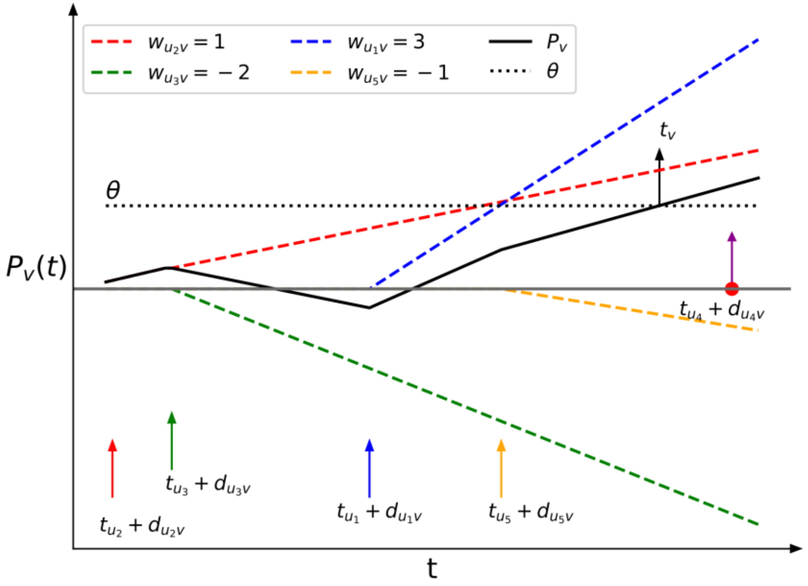

with $F_u = \{t_u^f : 1 \leq f \leq n \text{ for some } n \in \mathbb{N}\}$ being the set of *firing times* of neuron u, i.e., times t whenever $P_u(t)$ reaches θ_u.

Spike Response Model Networks: Function Approximation

Theorem (Singh, Fono, Kutyniok; 2024):

Let $L, d \in \mathbb{N}$, $[a,b]^d \subseteq \mathbb{R}$, and let Ψ be a classical ReLU-neural network of depth L and width d. Then there exists a SRM network Φ with $N(\Phi) = N(\Psi) + L(2d + 3) - (2d + 2)$ and $L(\Phi) = 3L - 2$ that realizes the output of Ψ on $[a,b]^d$.

Theorem (Singh, Fono, Kutyniok; 2024):

For $d \geq 2$, $\ell := \lceil \log_2(d+1) \rceil + 1$. For Φ being a 1-layer SRM network with one output neuron v and d input neurons $u_1, \dots u_d$ with $w_{u_i v} \in \mathbb{R}_{>0}$ for $i \in \{1, \dots, d\}$. Then

(1) t_Φ can be realized by a classical ReLU-neural network Ψ with $L(\Psi) = \ell$ and $N(\Phi) \in O\left(t \cdot 2^{2d^3+3d^2+d}\right)$.

(2) t_Φ can be realized by a classical ReLU-neural network Ψ with $L(\Psi) \in O(d)$ and $N(\Phi) \in O\left(8^d\right)$.

SRM Networks: Approximation of the Minimum Function

Theorem (Singh, Fono, Kutyniok; 2024):
For $d \geq 2$, there exists a single layer SRM network Φ, with linear response function, with one output neuron v and d input neurons, such that

$$|\Phi(x_1, \dots, x_d) - \min\{x_1, \dots, x_d\}| \leq \frac{(d-1)\theta}{2dw}, \quad \text{for all } x_1, \dots, x_d \in \mathbb{R},$$

where $\theta > 0$ is the threshold of v and $w > 0$ is the weight of each connection.

Comparison with ReLU-neural networks:

→ For any classical ReLU-neural network, irrespective of depth, to approximate *min*, *each hidden layer must have at least d neurons*.

→ Under certain assumption on the weights and data distribution, a classical ReLU-neural network of *depth 3 is necessary* to efficiently approximate *min.*

Spiking neural networks are strictly more expressive!

SRM Networks: Linear Regions

Theorem (Singh, Fono, Kutyniok; 2024):
Let Φ be a one-layer SRM network with linear response, with input dimension d and a single output neuron. Then the *maximum number of linear regions* $|\mathfrak{R}|$ satisfies the *tight* upper bound

$$|\mathfrak{R}| \leq 2^d - 1$$

Some Remarks:

→ In comparison, one ReLU neuron divides the space *only into two regions*, regardless of d.

→ A single spiking neuron divides the input space with the *same number of linear regions* as a classical two-layer neural network with d hidden neurons.

Spiking neural networks are strictly more expressive!

Vision of our Project

///gAIn Next Generation AI Computing

MATH ⁴AI


G. Kutyniok
LMU


H. Boche
TUM


S. Speidel


F. Fitzek
TECHNISCHE UNIVERSITÄT DRESDEN

Comprehensive theory-driven framework for next generation (Green) AI-systems: Optimally application-adapted hard-software combinations for maximal energy-efficiency and reliability!

Provably reliable AI-based *communication* systems which comply with the EU AI Act!

Low cost, trust-worthy *medical* AI-devices for diagnosis and therapy!

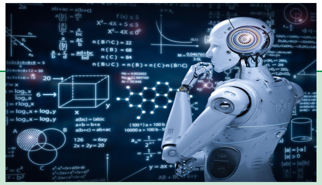

Next generation *robotics* with reliable robot brains and life-long learning capabilities!

Bayerisches Staatsministerium für Wissenschaft und Kunst

STAATSMINISTERIUM FÜR WISSENSCHAFT KULTUR UND TOURISMUS

Freistaat SACHSEN

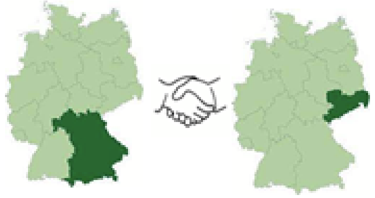

Bavaria Saxony

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Team at our Chair for gAIn Next Generation AI Computing





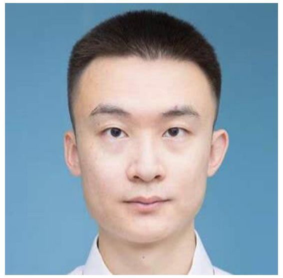

Dr. Ernesto Araya Dr. Massimiliano Datres Adalbert Fono Vit Fojtek Dr. Jianfei Li

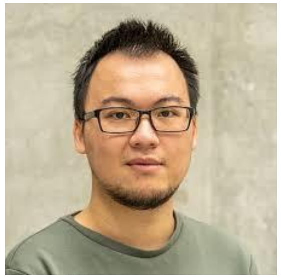

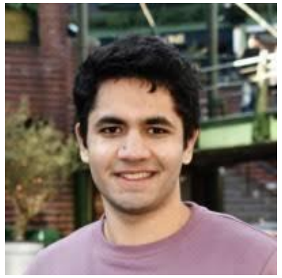

Duc Anh Nguyen Sarah Pardo Manjot Singh Dr. Juan Suarez Cardona Jonas von Berg

Conclusions

Conclusions

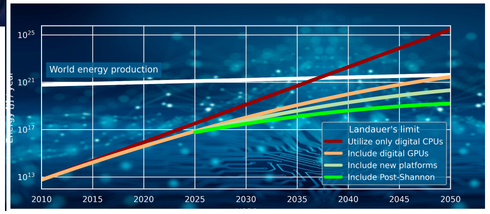

Current Problems with Reliability and Sustainability of AI!

Taking a Mathematical Perspective:

→ Analysis of *Expressivity, Training, Generalization*

→ *Explainability:* Rate-Distortion Explanation / CartoonX

Fundamental Problem with Digital Hardware!

Next Generation AI Computing:

→ *Analog hardware* such as neuromorphic computing!

→ *Analog AI systems* such as spiking neural networks!

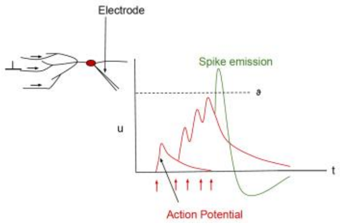

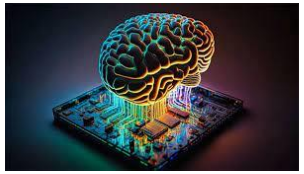

Vision: Mathematically Reliable and Sustainable AI!

Konrad Zuse School of Excellence in Reliable AI

(https://zuseschoolrelai.de)

SPONSORED BY THE

Federal Ministry of Education and Research

DAAD Zuse Schools

Konrad Zuse Schools of Excellence in Artificial Intelligence

MATH 4AI

relAI

Konrad Zuse School of Excellence in Reliable AI

Munich, Germany

Technische Universität München

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

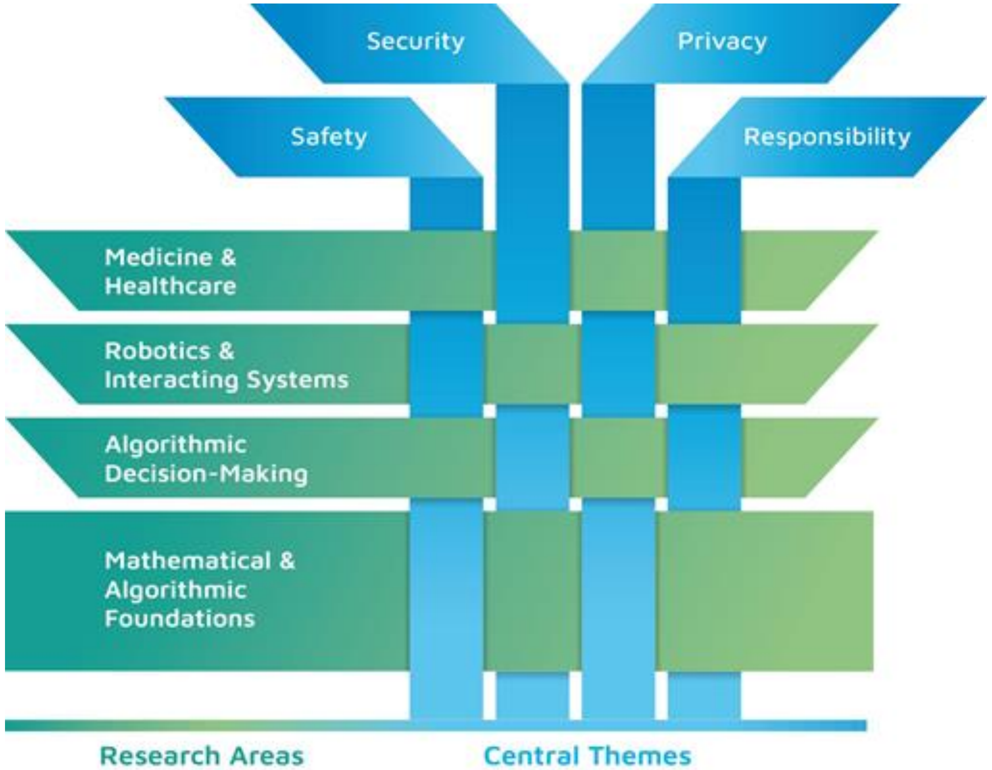

Mission: Train future generations of AI experts in Germany who combine technical brilliance with awareness of the importance of AI's reliability

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Thank you very much for your attention!

References available at:

www.ai.math.lmu.de/kutyniok

Survey Papers:

Berner, Grohs, Kutyniok, Petersen, The Modern Mathematics of Deep Learning, 2021
Fono, Singh, Araya, Petersen, Boche, Kutyniok, Sustainable AI: Mathematical Foundations of Spiking Neural Networks, 2025

Related Book:

Grohs and Kutyniok, eds., Mathematical Aspects of Deep Learning, Cambridge University Press, 2022.