Fraunhofer
IAPT

Fraunhofer Research Institution
for Additive Manufacturing
Technologies IAPT

Artificial Intelligence in additive production: applications and opportunities

Alliance Deep Dive | 2023

exclusive

ADDITIVE
ALLIANCE

Contact

Fraunhofer Research Institution for
Additive Manufacturing Technologies IAPT

Am Schleusengraben 14
21029 Hamburg-Bergedorf
Germany

+49 40 48 40 10-500

marketing@iapt.fraunhofer.de

www.iapt.fraunhofer.de

www.linkedin.com/company/fraunhofer-iapt

www.youtube.com/FraunhoferIAPT

1. Abstract

Motivation

- Artificial Intelligence is evident in daily life and industries, proving valuable in encountering unstructured and changing data.
- AI has proven valuable while encountering unstructured and changing data when desiring structured outcomes, boosting sales and enhancing customer satisfaction.

Approach

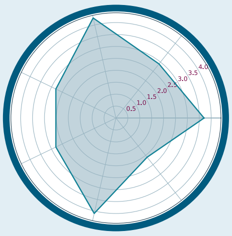

- We explore and offer an overview of artificial intelligence and its history.
- We identify current challenges in Additive Manufacturing and propose viable solutions for these challenges.
- We evaluate the solutions.

Results

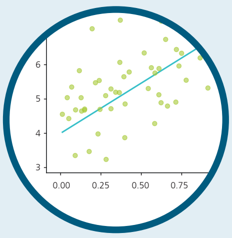

- Artificial Intelligences highlights its potential as a key enabler in optimizing and advancing Additive Manufacturing.
- Despite existing hurdles, mainly in data availability and preparation, the timely and correct deployment of AI can greatly benefit the industry.
- AI can potentially serve as a co-pilot to designers, engineers and developers, contributing to overcoming the grand challenges of Additive Manufacturing.

2. Content

3. Acknowledgement

Deep Dives are reports intended exclusively for the members of the Additive Alliance®. They provide the latest insights into the science behind Additive Manufacturing (AM). We extend our profound thanks to the members listed below. This Deep Dive could not have been created without their financial support.

Mesago Messe Frankfurt, the organizer of Formnext, has been the official sponsor and cooperation partner of the Additive Alliance® and Fraunhofer IAPT since 2020.

4. About the authors

Maximilian Vogt M.Sc. has been with the Fraunhofer IAPT and therefore excellence in research regarding AM for over six years now. Now as the head of the research group for production control systems, but also leading the IAPTs endeavors on AI and Digitalization, focusing on industrialization of AM through intelligent methods of planning AM processes and interfacing with additive machines.

Joining in is Johannes Helmholz M.A., B.Eng. who has been working on the business end of process implementation, oversight, and execution for over a decade as an officer in the German armed forces, now taking a closer look in the Production Control Systems Team on how to make industrial AM control systems smarter.

Their goal is to get you informed and involved with Artificial Intelligence (AI) in AM. Making the challenges of AM apparent and devising concrete use cases and their respective value proposition for the AM industry.

Maximilian Vogt, M. Sc.
Head of Control Systems Team

Johannes Helmholz, M. A.
Control Systems Team

Contact

Maximilian Vogt, M. Sc.
Head of Control Systems Team
Telephone +49 40 48 40 10-749
maximilian.vogt@iapt.fraunhofer.de
www.iapt.fraunhofer.de

Special thanks go our guests, friends and AI-experts, Arian Musa, Rutger Boels and Souaybou Bagayoko (Eraneos), Thomas Rost (SKAD) and Robert Friedrich (Deloitte) for their help, insights and input. We are looking forward to working with you again!

5. Motivation

There can be no doubt that the recent development and the astounding pace at which the development of Artificial Intelligence is progressing is impressive and might even scare or shock us from time to time. The pace of progression and the unbelievable success AI has had in its recent iterations must be reason enough to take a deeper look into »what is the deal with this AI thing«. And there are very good reasons to get involved.

As we face a digital revolution and the world gets reshaped by digitalization and the rise of AI, our daily lives and interactions are undergoing a transformation, as do industries on a large scale. As you are reading this introductory text, let's say it's not the first thing you did after waking up, there is a chance that is far greater than zero that you have already encountered AI in one form or another. If you have already checked your mobile device, listened to music, or browsed the web, read through your emails, there surely was AI involved. That does not necessarily mean that you were observed or that something learned from your behavior, but at least, you met unstructured and partly unknown data that otherwise would have eluded the grasp of classical algorithms.

These interactions not only reach us in our personal lives, but they also extend to other sectors as well, including manufacturing. And for that matter, Additive Manufacturing. Although the number of direct interactions per day paints a different image, that percentage is on a steep slope and may rise to 100% [Tho22; Sto22]. But there is still a lot of awareness to be built on the topic [Ken23]. Which partly relates to the understanding of AI as human-like or human-imitating behavior. On one hand, the impact and range of AI implementation are commonly understated by the general populace. On the other hand, Large Language Models (LLM) sparked fear in some, overstating and fearing the level of advancement, AI has already made. We can clearly conclude that AI is everywhere and that there is buzz around it, for one reason or another. But why implement AI in the first place? There are clear reasons: There is a competitive edge to be gained or advancements to be made (In other words, there is money to be made and insights to be gathered). And between those two, all the shades and combinations are possible. Most of the time, they do align, sometimes they benefit the customer and the people, and sometimes only the entity deploying them gains any benefit at all. For most of us, a competitive lead, better products, or better services would be the edge that keeps employees employed, shareholders and stakeholders engaged, and the company alive, which in the end generates a monetary benefit.

Technology adaptation and development lead to a competitive edge – which, in turn, enables the above. AI has proven to be valuable.
Coming back to our initial example, of you encountering AI today. If you decided to start your day with coffee and the digital news, those are probably curated by AI. Funneling your browsing behavior, personal interests, engagement, and other multidimensional variables which must not be inherently observable. If you're more on the page of paper, you might listen to music while reading, turning on your favorite playlist, which at some point will have all songs played, and shortly after that, you get handed personalized music curated by AI. We could delve into more examples. As a rule of thumb, you'll encounter AI where there is a lot of unstructured and changing data, but you desire a structured outcome. Over the last few years, it was clearly shown that AI, to say the least, boosts sales and improves customer satisfaction. [Tri20] Although interesting, those use cases show day-to-day consumer interaction. Recent examples show that AI, and Machine Learning (ML) for that matter, may also be used to compute otherwise incomprehensibly large models of real-world objects, with results matching or even surpassing our most sophisticated numerical simulations [Kar21].

AM faces several challenges. These are cost per part, as a general challenge, quality assurance, developing new materials, automating AM processes, enabling the workforce, sustainability, and optimizing the supply chain. Solving all of these ensures competitive, smart, and sustainable AM.

In essence, our goal for this Deep Dive is straightforward: Can AI, with its all-encompassing influence, be the missing piece of the puzzle in tackling the as-yet unsolved challenges of AM? Our belief is grounded in the idea that we have all the prerequisites in place for a successful deployment of AI. There is a competitive edge to be carved out, a need for structured results in an environment of unstructured or fluctuating data, and a myriad of unknown variables waiting to be unearthed for a better comprehension of the underlying mechanics. Add to that, we now have in our hands a fresh, swiftly evolving, and potent toolbox in the form of AI, which might just give AM the push it needs to be the go-to choose for complex product challenges. Through this exploration, we aim to put forth potential AI solutions to navigate the towering challenges of AM. The question is not whether AI can contribute to the realm of AM but how we can best leverage its strengths to drive this industry forward.

»AI might just give AM the push it needs to be the go-to option for complex product challenges.«

6. Approach of the Deep Dive

Both fields we are trying to tackle are their own disciplines with their very own subfields and specialists in each of those fields. Every subfield brings with it its very own and specific set of problems and solution strategies. Concerning AI, we want to give you an overview of the state of the art of AI in industry applications. Surely, we cannot cover every aspect and intricate mathematical detail, thus, we try to make it as applicable as possible, to give you an overview, an idea, and maybe a guide to your very own use case.

We will cover the basic ideas of ML and AI in general, and we have a staff picked list of concepts from the realm of AI that might be of interest for you, your production or research and development efforts. For those of you wanting to go way deeper, there is a list of references to online courses and good reads at the end of each paragraph. We won't let you figure out your use cases all by yourself though. Naturally, we covered the AM process chain as well. While looking into different AM-processes and the challenges each of the different processes is facing one might get drowned in a flood of information.

Our approach is not to identify each single process or problem specific challenge and trying to match AI on top of the problem, but to give you an idea of what the overarching, the grand challenges of AM might be. We reviewed the process chain bottom up, from single step singly purpose processes to the combined manufacturing efforts in AM.
At the end, we identified the challenges the industry is facing right now, and derived possible solutions for those challenges. We are more than willing to assist in implementing these solutions, ensuring a smoother path forward.

7. State of the art

In recent years, AI has emerged as a transformative technology, revolutionizing numerous fields, and impacting nearly every aspect of our lives. AI encompasses a wide range of approaches and techniques that enable machines to mimic or surpass human intelligence in various tasks.
The state of the art in AI spans multiple domains, including ML, natural language processing, computer vision, robotics, and more. ML has witnessed remarkable progress with the advent of DL models, which leverage artificial Neural Networks (NN) to automatically extract complex patterns and make predictions. These models have achieved groundbreaking results in areas such as image recognition, speech synthesis, and language translation.

In this section, we will explore the state of the art in each of these domains. By understanding the current landscape, we can appreciate the potential of AI, identify areas for improvement, and envision the future possibilities that lie ahead.

7.1 Brief history and introduction to Artificial Intelligence

Although, unorthodox at first glance, most of our technological innovations where first and foremost incubated by the means of science fiction literature. The golden age of science fiction in the early 20th century is one of those examples, where the human mind wandered, and though of unimaginable futures for humanity will live in the next decades and centuries.

The notion of AI, that mimics human intelligence, is also derived from science fiction and dates to the late 19th century. The well received quote from Isaac Asimov on the very nature of science fiction is accompanied by the notion that a science fiction way of thinking is necessary for the making of the right decisions.

»No sensible decision can be made any longer without considering not only the world as it is, but the world as it will be (…) but the core of science fiction, its essence, the concept about which it revolves, has become crucial to our salvation, if we are to be saved at all.« [Hol78]

This quote illustrates how thinking and forward thinking in humanity might evolve. Fantastic science fiction and fantasy stories constitute the mindset and way of thinking of generations of humans. The same goes for the story of AI. Science fiction laid the very foundation for mental models of a whole generation of scientists starting in the 1950s and 1960s. The most prominent of them, regarding computer science, is no other than Alan Turing [Tur50]. Turing asked the question »can machines think« and in the end argued, that we might hope, that one day, they can. Clearly, research and thinking about future possibilities went on. Sparking the first Ideas, thinking machines became a reality. First of its kind was the Logic Theorist [McC04] a program solely designed to prove mathematical theorems, which it sometimes did better than humans. This was achieved by a branching tree search, logic symbolic programming and heuristics. Which by the way; the now used word heuristic was coined by the inventors of the Logic Theorist Allen Newell, Cliff Shaw, and Herbert Simon themselves. The concepts of managing branching explosions with heuristics is still in broad use today in the field of AI.

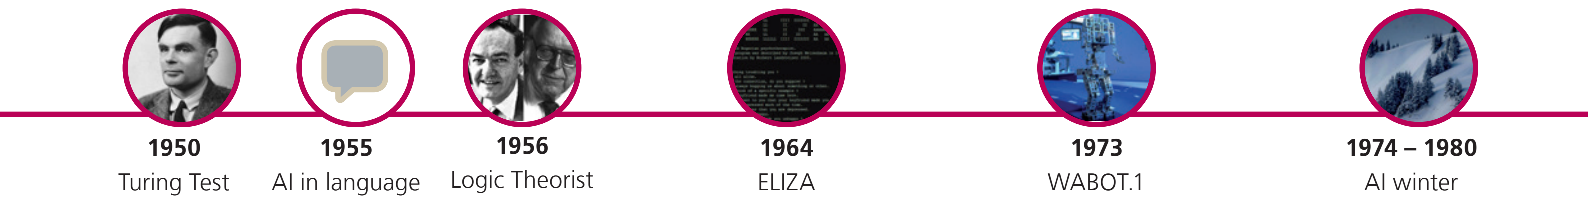

Further reading

A good start for Deep Learning and Machine Learning. ISBN: 9780262537551 and ISBN: 9780262542524
Try ELIZA and scan the QR code:

The Logic Theorist was first introduced at the Darthmouth summer research project on AI, a conference which catalyzed the years of AI research to come. Still asking the question, what human logic and reasoning is about, scholars quickly realized and have already realized decades before that humans have a distinct set of symbols, which can be combined to form language. Natural Language Processing (NLP) was first successfully tackled by Joseph Weizenbaum in 1966 with ELIZA, a »psychotherapist». Arguably ELIZA was not that good about »thinking« per se but was very good at pattern matching and giving its users the illusion of a deeper language understanding. It was most famously known for its tranquilizing ability to bind users for hours on end to the screen. Which even came to the surprise of Weizenbaum himself, as users often anthropomorphized the chatbot. [Wei66] The later years yielded not much groundbreaking success in the field of AI.

Interest in AI plummeted in the early seventies due to the well-known hype cycle decline and remained in a deep slumber. The expectations for AI were high, as was the funding up to this point. But mid-century computers were just too weak.

Popular culture got inspired by the early successes of AI but painted a more sinister picture. Movies such as »2001 a Space Odyssey« (1968) or the novel series »Dune« (1965) may have inspired a later generation of researchers and later depictions of AI in popular culture such as in »Terminator« (1984) or »The Matrix« (1999) [Cla74; Her05; Cam84; Wac99].

By the mid-80s funding rose and interest in AI was at a high level again. This was also the time DL emerged from research. Although the nature of biological information processing and storage was discussed in the 50s [Ros58] the emergent computational properties of NNs where first discussed and described by John Hopfield in the 80s. Those mathematical descriptions laid the basis for DL as we know it today [Hop82]. At the same time, expert systems, as described by Edward Feigenbaum, were heavily financed [Fei81].

Due to Moore's law still holding up, predicting the doubling of transistor count roughly every two years, formulated by Gordon Moore in 1965, [Sch97] the early 90s until the mid-2010s saw another rise in development of AI, due to finally having the raw computing power available to tackle complex calculations on a broader scale. The slow and steady rise was accompanied by some media impactful successes for the field, as some researchers endeavored in the journey of beating human made games such as chess, Jeopardy, AlphaGo and later even DOTA 2.
With a world champion in chess beaten by Deep Blue in 1997, IBM's Watson winning at jeopardy in 2011, AlphaGo beating a Grandmaster in Go in 2015 and a world class team in DOTA2 beaten by OpenAI in 2017 [Ope19]. As a side note, current research shows that the way an AI »thinks« is still very narrow, and that no system understands what some of the founding ideas of those games really mean [Wan22b]. At the same time, NLP and other fields got better month by month and integrated tightly into our devices. Speech recognition and text to speech or speech to text nowadays is an integral part of interaction with our devices. One first broadly used tool was Apple's Siri, introduced in 2011.

Today, researchers, companies and government agencies are pushing hard for AI development. There are daily developments and news articles pointing towards advancements. Those advancements are fueled by the very data-intensive and connected nature of our society. Nourishment for AI in the form of data, be it structured or unstructured, is everywhere. Computational prowess is at a point, where we can mimic some of the functions of a biological way of thinking. Advancements are everywhere, be it in computer linguistics, physics, biology, image processing and generation in every field or clustering of data. There is no doubt about large language models (LLM) impressing us the most. Language constructs reality, interaction, and emotion. In our next paragraph, we will elaborate further on the landscape of today's AI and methods used.

7.2 Artificial Intelligence overview and map

Before we go deep into the realm of AI, allow us a little remark and a definition. For each of the following subsections exist more than a few books, papers, and thousands of pages in form of mathematical proofs, description, application, design, and best practices, coming not only from engineering, but physics, biology, psychology, medicine, and every other research discipline you might think of. Arguably, AI is too vast to boil it down to a few pages, but we did try anyway.

Subfields of Artificial Intelligence

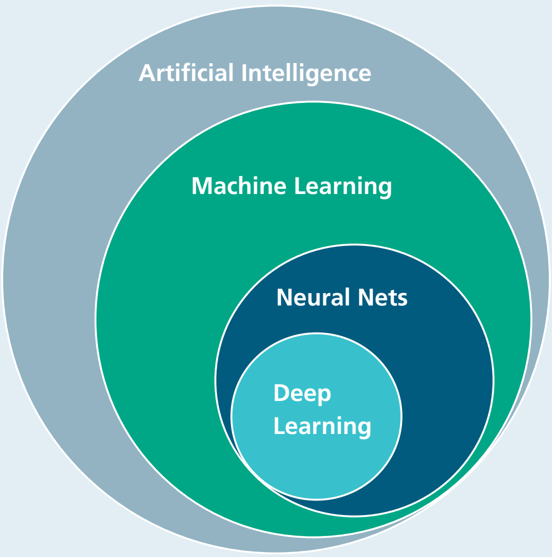

Figure 1: Subfields of Artificial Intelligence

The map on this Page depicts the different subfields of AI, in general ML, NNs and DL are subfields of AI. As a very basic definition, AI and all subfields differ from basic algorithms in that an algorithm is just a recipe that takes an input, performs a task, and yields a result.

You could compare that to a cooking recipe, that requires an input – ingredients, performs a task – the cooking itself, and has an output – a finished meal. The nature of AI differs from that in a very meaningful sense. Compared to the cooking recipe, you desire a specific meal. But the ingredients in your fridge do not meet the standard recipe and you are missing a few cooking utensils. Cooking with AI would now be able to yield a new recipe that fulfills your desire for a meal.

The analogy aims at the very core of what AI should be doing: it will flexibly learn to fulfill its directive like a human would. It is much like a chef will change the course if he misses the poultry but has only beef available. But do not expect fish if you only have pork, that's something AI won't be able to do for you.

As a rule of thumb, AI must be able to change its algorithms, based on data availability and desired output and mimics human intelligence and behavior. We will elaborate on that further in the following subsections. ML as a subset of AI describes this concept in more detail, describing the process of machines extracting information from data and learning autonomous from it. NNs and therefore DL describe a subset of ML where you mimic the neuronal structure of the human brain. In Neural Nets and DL, we are leveraging the emergence of intelligence that happens when you wire neurons in a net. We picked some of the concepts out of ML and took a better part from Neural Nets and DL and going to explain some of the concepts in the next section.

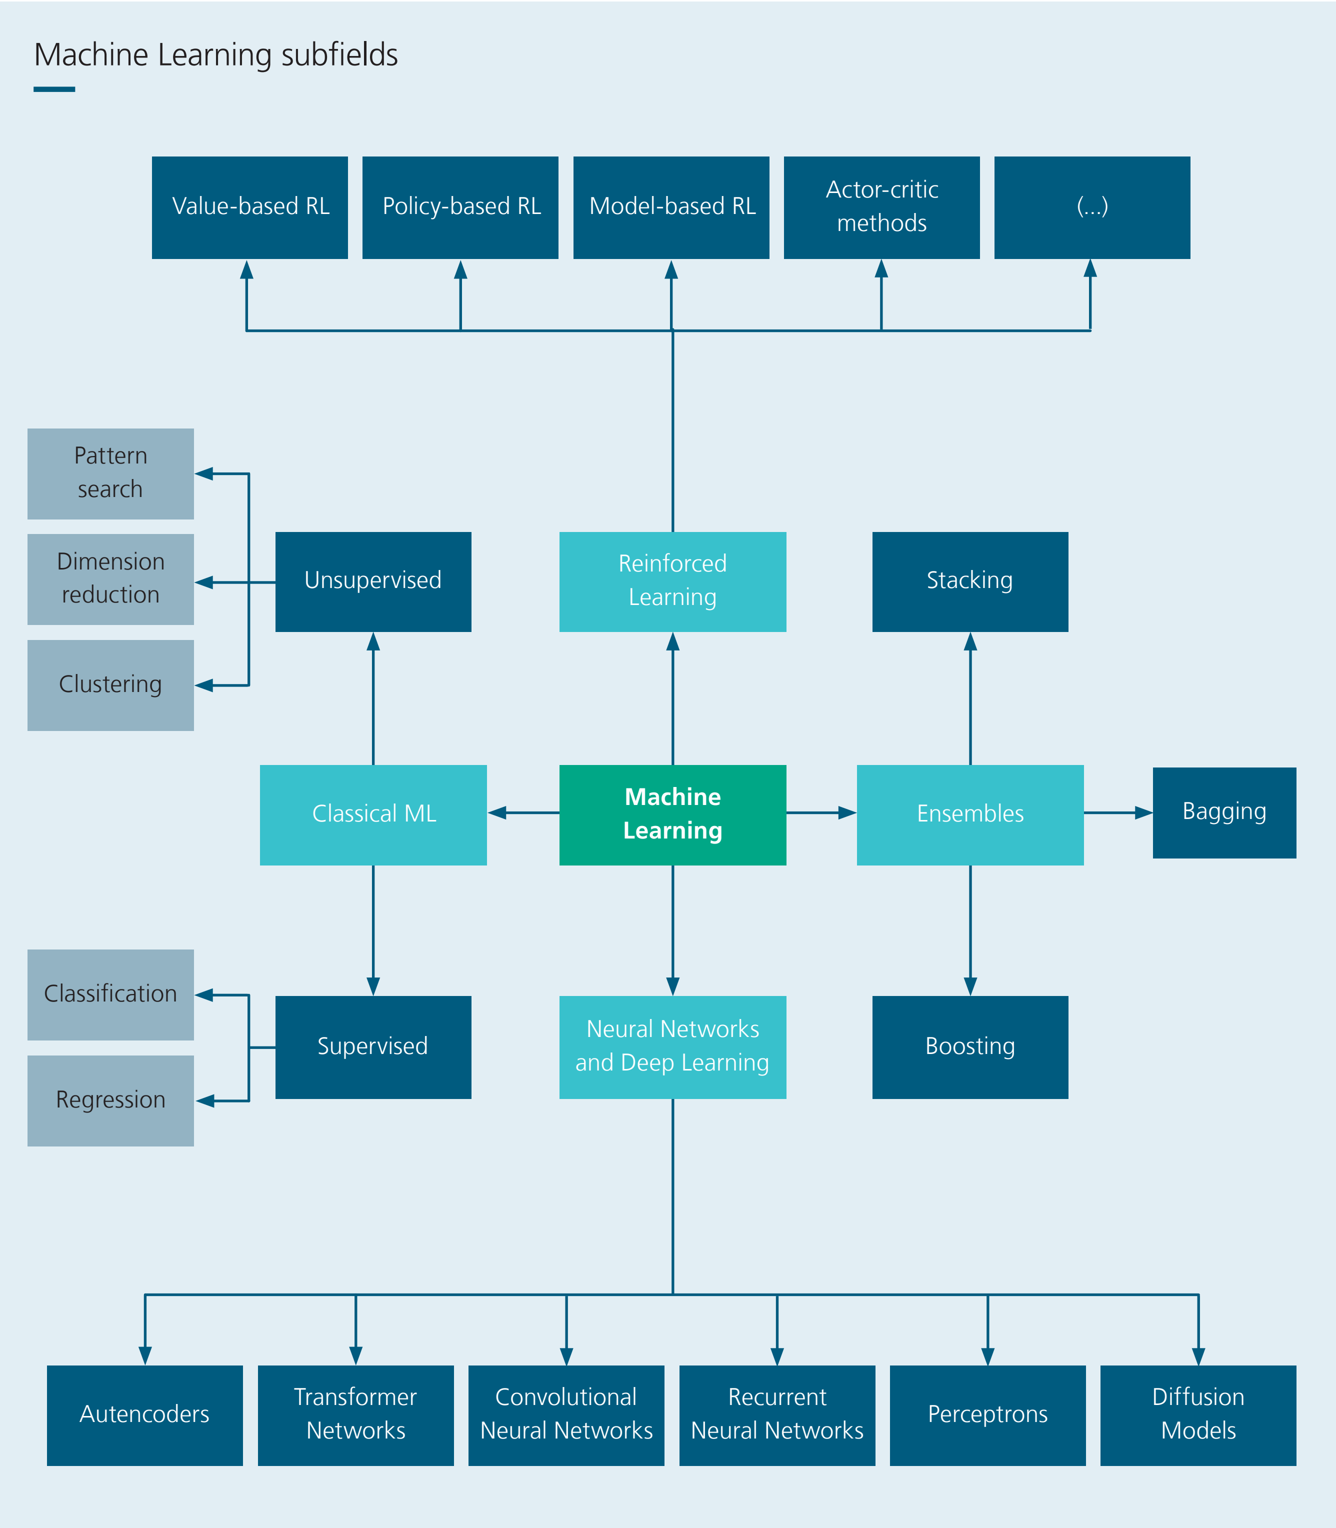

Figure 2: An overview of Machine Learning [Gol16; Ert21]

7.3 Introduction to Machine Learning

As already mentioned, in ML, the machine learns. The elephant in the room clearly is, how does it do that? There are different methods inside ML which are appropriate for different kinds of data and data availability. The following graphic shows some of the broader approaches. There is classical ML, ensemble techniques, Reinforced Learning (RL) and NNs and DL. Each of those can again be subdivided into other subfields. For example, in NNs, there are Convolutional NNs, Adversarial NNs or DL NNs and more. But to get into the single definitions and descriptions, we need to first clarify terminology and concepts.
But first we would like to start with the importance of data. After that we will discuss the difference between supervised and unsupervised learning.

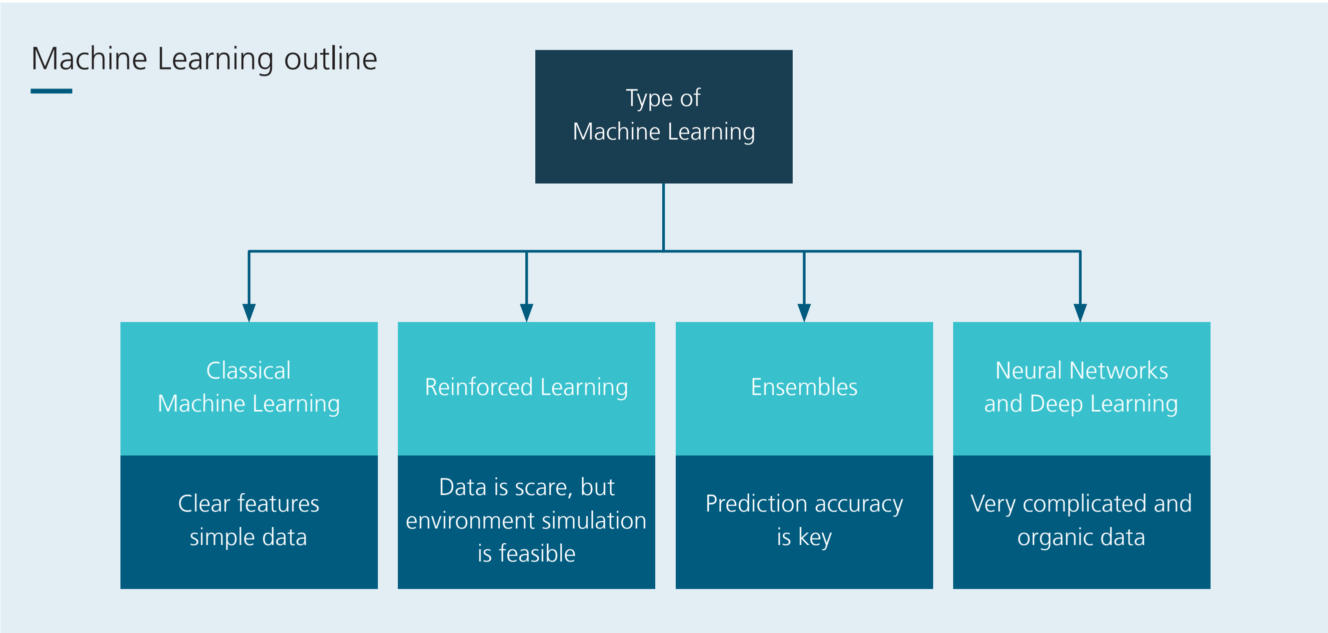

Figure 3: General overview of Machine Learning Types [Gol16]

»With data to insight, we elevate Additive Manufacturing with Artificial Intelligence.«

7.4 The importance of data

Garbage in, garbage out. One of the most repeated mantras in programming and other computer science related fields of expertise. To get involved with AI we need to take a step back and consider, what does that even require? Data comes in a variety of forms and therefore also in a variety of quality and level of readiness to be deployed in statistical analysis and ML. If you are familiar with statistical analysis, you might know, that different tests for p-values, or the level of correlation between two or more variables is dependent on certain prerequisites.
Tests often require the data to be invariant or a normal distribution, ranked data differs in tests from metric data and so forth. Outliers can have a huge influence on how certain tests perform and sometimes you are testing the wrong correlations altogether, or even trying to find causalities when there are none or misinterpreting correlation for causality even through a lack of knowledge and interpreting the results of the statistical tests, which do require a lot of background knowledge in the field they were conducted for. A significant portion of statistical analysis is dedicated to acquiring the data in the first place and then cleaning them or developing them into new features, or even collecting more data if the first set did not suffice or hinted at another interesting correlation. All that is true for engineering, business, social sciences, physics, medicine and many more. And that still holds for AI [Gol16].
Nowadays we've given fields involved with the acquisition and cleaning and interpreting the data names, data engineering and data science, those fields are no longer tools in the toolbox of other fields of expertise, but have grown so large and powerful, that not only universities dedicate own fields of studies to them, but also whole businesses employ persons just for fulfilling the tasks. Both are equally important and should be considered in your endeavors for AI.

Data engineering on one hand consists primarily of the acquisition and maintenance of data pipelines, data science on the other hand then tries to do something with that data using diverse modeling and visualization methods. If we apply this to AM, you will need a robust pipeline, that can collect relevant data from processes, machines and other systems partaking in your production efforts, this could consist of in-process data, test data and other broader key performance indicators. This data then needs to be stored in databases and processed on machines, all with scalability in mind, ensuring availability and integrity of data for other branches and stakeholders. The workflow is rather straightforward and less explorative.

In contrast to that, the data science aspect focuses on exploration of data, what data could be relevant to extract useful information. Applied to AM, this could be to iteratively find out what in process data is useful for a given task, how to refine the data, what kind of features need to build from the data. The second job is making assumptions and predictions from the data and visualize that data for stakeholders. Both branches need to be covered in some way or the other. The least you would need is a robust pipeline, storage and preparation strategy and a robust toolset in statistics, analytics, and visualization. [Bol18]

7.5 Supervised learning and unsupervised learning

Supervised and unsupervised learning are both concepts, which are always heard or mentioned when talking about different problems tackled by ML. The biggest difference between the two is, that in supervised learning, there is a ground truth – the dataset is labeled. The goal is, to find a function or representation of how input and output correspond.

As expected, except in numerical solvable relationships, there will be no function to describe every data point with every label, so a loss function is introduced. One such example is the mean square error often used in regression. But we are looking for a something close to visualized, it might look something like the following graphic. We generated some data with noise, with an almost linear relationship. We are looking for the weights in the regression problem: $\hat{y} = wx + b$. The best fit, as depicted in the graphs, is for $w = 3$.

Unsupervised learning on the other hand has no ground truth, there is no labeled data, and it is widely used to find patterns and therefore clustered data. The following Figure 5 shows the K-means clustering, where, with an initialization of cluster

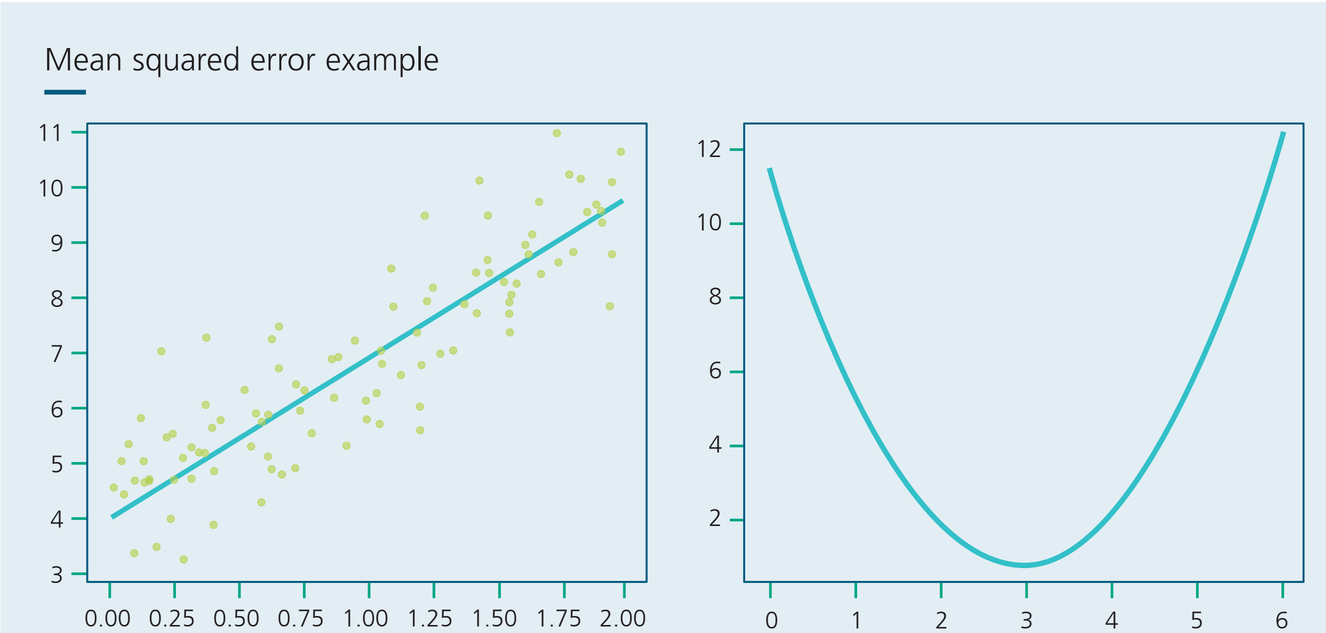

Figure 4: Linear regression and best fit example

centers, the distance between the data points and the cluster centers is expressed in a loss function.
Supervised and unsupervised learning differ greatly in where they are applicable. As a rule of thumb, supervised learning is conditions. Regression, classification, time series forecasting and simple form of speech to text would be fields, where supervised learning is applicable. If you do not know the outcome of your conditions and want to explore a dataset, you may use unsupervised learning for clustering, learning association rules or detect anomalies. The term learning in both supervised and unsupervised learning refers to the internal optimization of the algorithm to better fit the objective function. Simple data structures, clear label and difference between data points are requirements for classical approaches. But if you have those prerequisites fulfilled, there is typically no need to throw any other approach on that problem. Freely according to the motto: tolerances as exact as necessary – or models as complex as necessary.

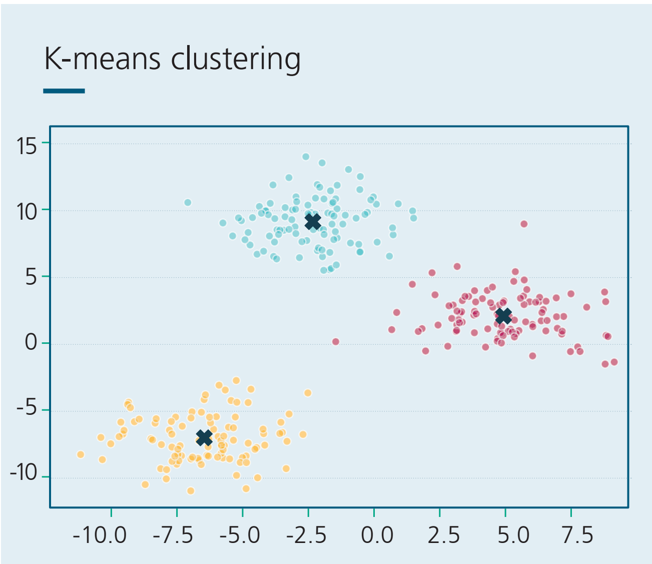

Figure 5: Example of K-means clustering

7.6 Reinforcement Learning

The first contrasting contender to classical learning approaches is Reinforcement Learning. The term RL is often mentioned in the same sentences NNs and DL are mentioned. This does apply for newer methods, but there are also methods that only use the relationships and basic ideas of RL and do not utilize any form of artificial neurons. The very core of describes a relationship between an agent and an environment. Whereas the environment has certain states, the agent can undertake certain actions and is either rewarded or penalized by undertaking its actions. There is also a certain level of risk involved for the agent, which is expressed by the connection between exploitation and exploration.

Exploration meaning an incentive to explore new states, even if old methods yield good enough results. Exploitation meaning, how much of old results should be used. Let's make this more approachable. Let's say you have a dog and have a reward function, namely a treat for said animal. It gets the treat whenever it fulfills a task for you or follows a specific command. If it does something it's not supposed to do, there is either no reward or a punishment. You may also find analogies for exploration and exploitation in this. In our analogy the dog is the Agent, the environment is where the dog undertakes its actions. The actions are expected actions as reactions to a command. Following a sit command, the dog is expected to sit, but may, especially in the beginning try out other actions, for example barking or standing on its hind legs. This is our analogy for exploration. The further the training gets, the more a dog may exploit already working behaviors, for example to sit even if the command was not to sit. The dog knows at this point, that sitting yielded a reward. With more training the dog gets better at following diverse commands and reaching its goal faster. Let's consider another rather simple example, utilizing a simple simulation.

Imagine an environment, a 100x100 grid containing a marker at a fixed position. In this environment is a dot placed at a fixed position, this is our Agent. The agent has certain actions at its disposal, it may move up, down, left, right and diagonally. For each simulation step, it may move one field, thus having eight degrees of freedom. The learning can for example be achieved by Q-learning. The actions the ball can take are represented in the Q-table. It now can either explore and choose the next action at random, or it may exploit and choose the next action based on the highest Q-value currently assigned. Moving gives the dot a reward. Getting closer to the marker gives a +1, getting farther a -1, moving out of bounds a -10 and finding the marker gets +10. The Q-value is updated and used in further moves. At some point the state space and action space of simple algorithms gets too big for classical approaches, for those, deep reinforced NNs where introduced. For problems with an extreme number of states, the Q-table at some point gets very large. If a simple approach works for the problem at Hand, use the simple approach. Before we delve deeper into the rabbit hole, there again is some interesting further reading and watching attached on the right hand site [Sut98].

Agent environment relationship

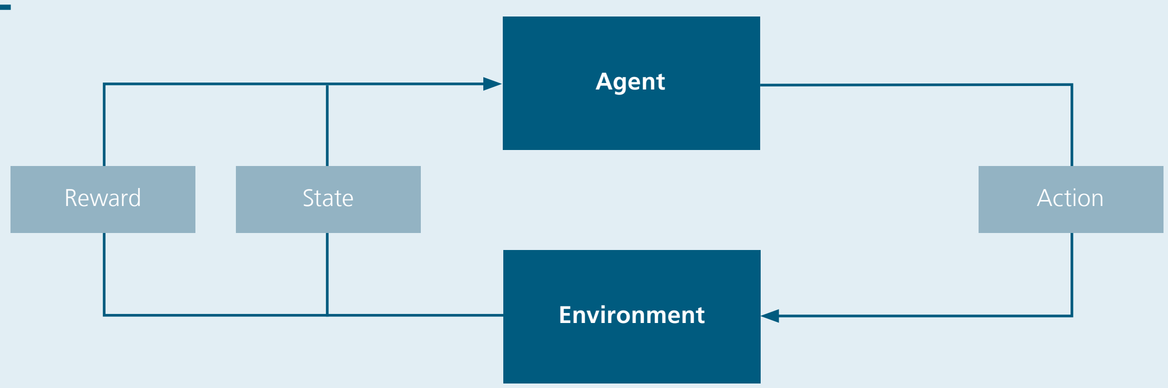

Figure 6: Reinforced Learning agent/environment

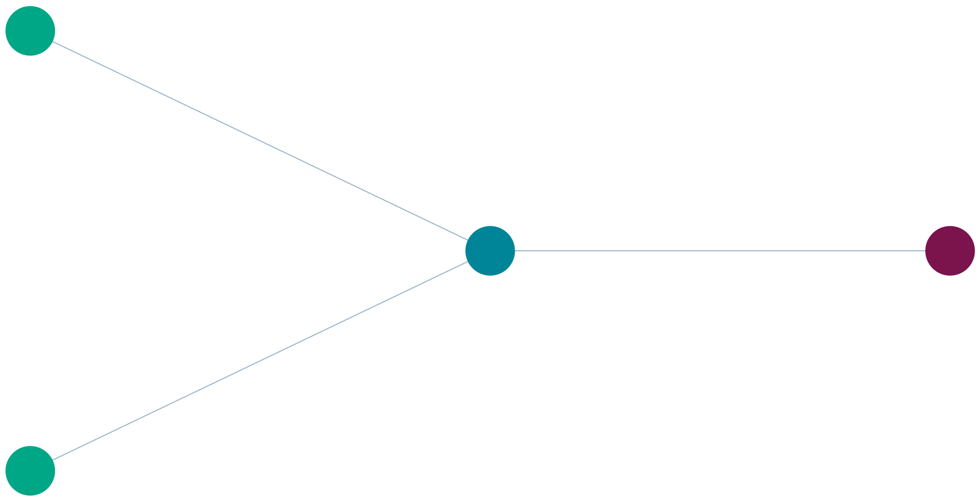

7.7 Single layer perceptron

Let's start by elaborating on the first concept in the realm of NNs, which is the very basic perceptron [Kan03; Ros58; Ros57]. To understand how those networks work, we need to at least develop a very basic understanding of the neurons in our brain. These interactions where reason to transfer biology and neuroscience ideas to computing. On a very basic level, the neurons in our brain activate if their input is reaching a chemically defined threshold, called firing threshold. Single neurons are connected via axons and can be connected in a various arrangements and be connected to different neurons.

A single neuron activates if the sum of its input is higher than the firing threshold and it will consequently transfer the chemoelectrical impulse. The graphic above shows a perceptron structure which is modeled to mimic the logical AND. The blue neuron will only activate above a certain threshold which in turn outputs a signal to the wine-red output. The green dots are inputs, and both can take a binary value, so either 0 or 1. Let's assume the activation function of the blue neuron is a step function, that activates when both inputs are 1, you would have input weights of 1 for both inputs and a bias of -1. In practice, those weights and biases would be achieved through a training process [Duk18; Ban23]. The logical AND can be represented by a dataset linearly separable. The sum of weights and biases is:

$$w_{green1} + w_{green2} + bias; \text{ or } 1+1-1, \text{ should equal or exceed } 1$$

Back to the neurons in our brain. As simple binary representation might not suffice. Our neurons and other cells as well have more intricate activation functions, which are in fact, not fully understood and even single cell behavior, which is much more complicated, is still undergoing active research. [Alt20] By changing the activation function of the neuron you can change the output behavior. Inputs can be a floating range and you get partially active states that activate the neuron. Other logical functions may require additional layers.

Green 1	Green 2	Red
0	0	0
0	1	0
1	0	0
1	1	1

Table 1: Truth table of logical AND

Further reading

Stanford class on reinforcement learning:

AI learns to escape:

How to navigate a maze with RL:

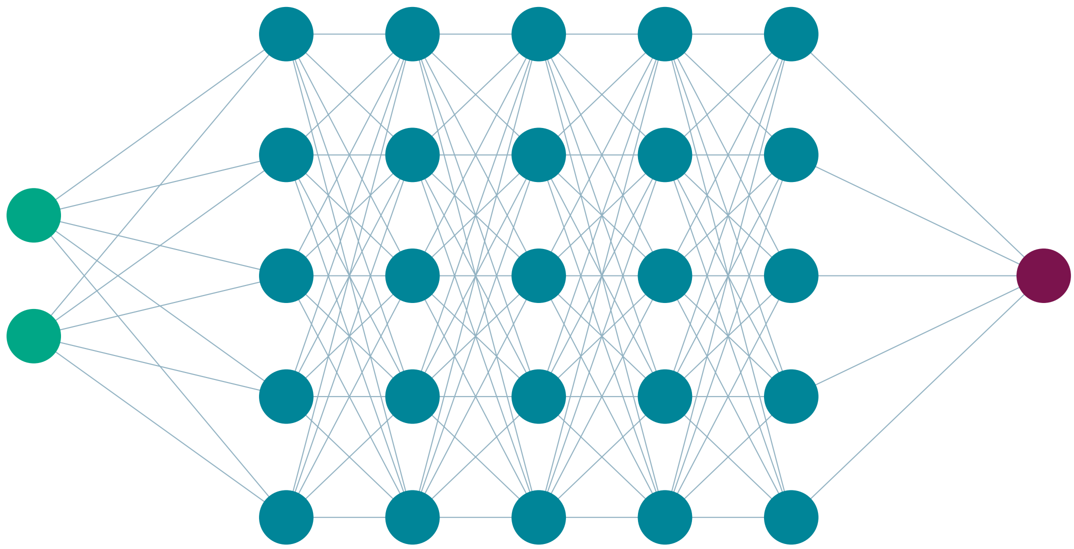

7.8 Deep (Learning) Neural Networks

The next step is understanding how densely connected networks work. Exemplary in the graphic above you see a Deep Neural Network (DNN). This network consists of two inputs, one output and five densely connected hidden layers with neurons each. This network could for example take two values in and output another, where the output is the label of the two inputs. The inputs and label are somehow intertwined by a not yet seen connection. The task of the network would be to discover that connection and apply it to other values.

On the technical side of things, those could be processing parameters like temperature and speed and labeled by the hardness of test specimen. The beauty of those networks is, that you can add, stack them, and expand them, so multiple inputs and multiple outputs can be chosen and thus better applied to multivariate technical processes. Training a network in its very basic configuration requires multiple steps. Biases and weights are often initialized by small random values. We then use these initialized values to feed data through the network for one epoch, meaning one full, or predetermined length of the dataset. For this we could again use the mean squared error. We get a number which uses our calculated values for the given output and compares the error to the actual output. The goal is to minimize this error. This is commonly achieved by back-propagation or gradient descent. On a very high level, back-propagation is used to calculate how much every weight in the network contributes to the overall losses. The back-propagation tells us what gradient each neuron has [Rum86].

To adjust the weights we can use gradient descent, which tells us in which direction we should adjust the weights to reduce the loss. This works because a NN is basically a chain of functions. Back-propagation and gradient descent both leverage that chained functions can be differentiated by the chain rule of calculus. Weights and biases can be both adjusted using this technique. There are numerous other techniques to train networks, back-propagation and gradient descent are often utilized for regression tasks. After the first iterations, the model might fit the training data exceptionally too well, thus predicting training data well, but performing poorly on other data, for this regularization might be necessary.
If those steps are completed, evaluation and fine-tuning may begin. If all has gone well, your model will produce accurate predictions for unknown data. Underlining the whole process are mathematically ideas derived at some point from observing nature. The learning process can be better described if we replace the labeled dataset with actions undertaken by a human for completing specific tasks. At some point you most likely learned how to drive a car or ride a bicycle.
At the beginning you were unsure of what controls do what, how much do I need to steer to go around a corner, or to park, what gear is appropriate for what speed and inclination?

Further reading

Backpropagation in more detail:
Learning representations by back-propagating errors, D. Rumelhart, Geoffrey E. Hinton, Ronald J. Williams: DOI:10.1038/323533a0

You are in a raw uncertain state, just as the initial state of the NN. As you are getting instructed your knowledge changes, you will get better in regards of how much steering, breaking and acceleration will change the behavior of your car. This is a phase of training, you refine all the weights of the individual parts involved in driving. You are being told what you did wrong and how much it influences the handling of the car in different conditions.

You want to minimize the mistakes, which is minimizing the loss overall and make better predictions on how the car reacts do different actions and to make predictions on what actions to take to get desired reactions out of the car. This is like the training of a NN. When you feel confident with your driving, you might change cars and learn to drive new vehicles rather quickly. Which is equivalent to generalization or retraining. This is a very generalized picture, as what eventually makes us good drivers is more complex. We have multiple direct feedback loops from our environment, we generalize actions, so two actions form one further down the line; shifting, steering, buttons and pedals becomes an act of driving. The biggest difference tough lies in our brains. For comparison, GPT-3 has 175 billion parameters, roughly equivalent to connections [Bro20]. The human brain is estimated to have around 100 trillion connections, which is only the brain, the other parts of the sensory and nervous system not even counting [Cat23]. Also, nodes and connections are malleable, and structures are not given by design but change all the time [Drag06; Kol98].

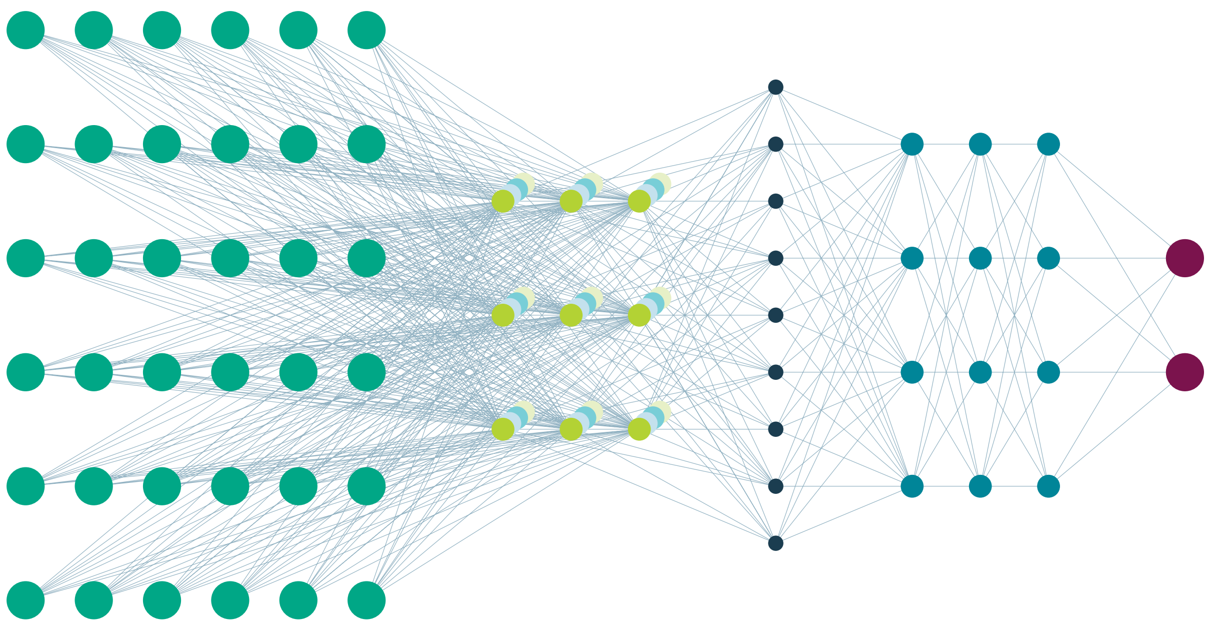

7.9 Convolutional Neural Networks

For Convolutional Neural Networks (CNN) the basic NN principles still hold up, but for image recognition tasks there are a few more layers involved. The image above shows a simple and stylized representation of a convolutional NN. This stylized representation contains of hidden densely connected layers, as discussed before and additionally there is an input in green, a representation of convolution and dimensional reduction or pooling represented by the lime-green nodes. Dark blue represents a first densely connected layer and light blue the other densely connected layers. Wine-red again is the output.

The CNN is comparable with looking through a tiny window and trying to grasp a vast landscape. You would need to look at different parts of the landscape from different angles behind the window and memorize features, like straight lines and where you found them. This is in a sense what the convolutional operation does, but it also involves introducing non-linearity by an activation function. After the initial identification the dimension of the input gets reduced and further, more complicated features extracted, for example corners or more intricate shapes in a second convolutional layer.

Stylized filters or kernels

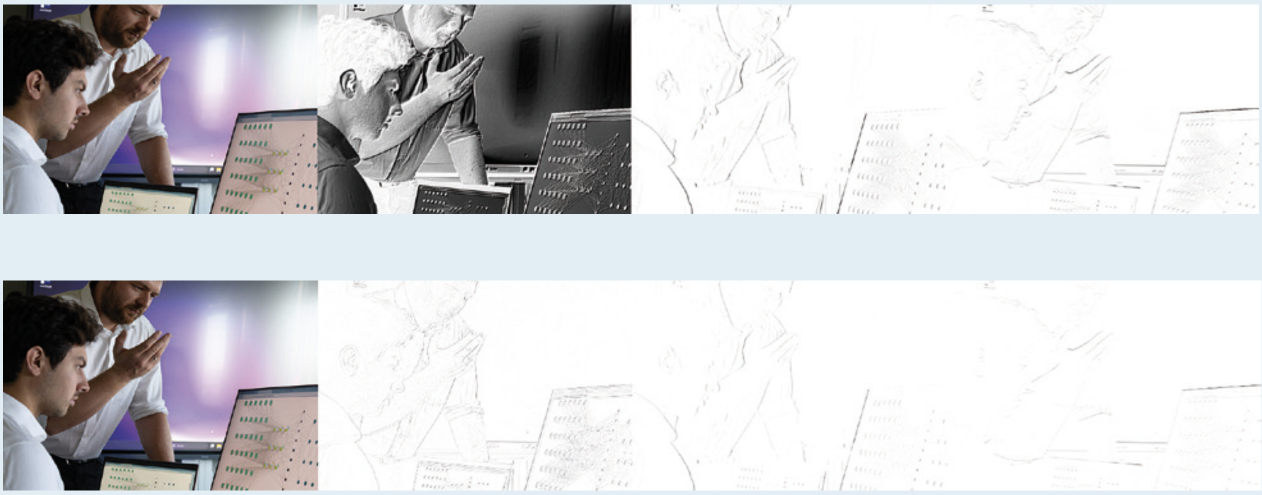

Figure 7: Kernel representations from left original image to right last manipulation

Different filters could look like this, see Figure 7. Left is the original image, followed by an embossing, followed by vertical lines, followed by horizontal lines in a first pass and a second pass where corners and longer edges are detected.

Generally, the architecture would be a cascade of convolutional layers and pooling layers. Kernel values are learned through a training process like traditional NNs without a convolutional layer. The last output of those operations would then be fed into either a flattening operation or directly into a densely connected layer. Flattening meaning the conversion of the multi-dimensional tensor into a one-dimensional vector. The CNN is then trained, quite similarly to the already described NNs. The training is more involved as other parameters as stride (how much will the filter move) or padding (how much of an image will be cut on its corners) are necessary [Alz21; Ert21].
CNN could be described by how a child learns to classify different things. Imagine teaching a child what a dog is, the child points at the animal and you say: »this is a dog«, which would be a first data point and label for a multisensory input. Upon seeing different breeds of dogs, the child updates and refines what features define and how much each feature partakes in definition of what a dog is. Now the child sees a cat and proudly announces to have seen a dog. You now label this new creature with »no, this is a cat«, which prompts the child to update what it has learned and adjust the internal weights of features making an animal a dog or a cat. The defining feature of the dog might not be that it has four legs and is furry, but a pronounced snout and the movement and posture of said animal might be better suited to classify the dog, this would be analogous to the filters which get adjusted in the network.

CNNs in combination with other techniques are not only suited for image recognition tasks with labeled data, for example, deployed as autoencoders, they might be utilized to reduce the spatial dimension of images and apply traditional clustering algorithms to the output [Alz21].

Further reading

This is a well-made CNN demonstration:

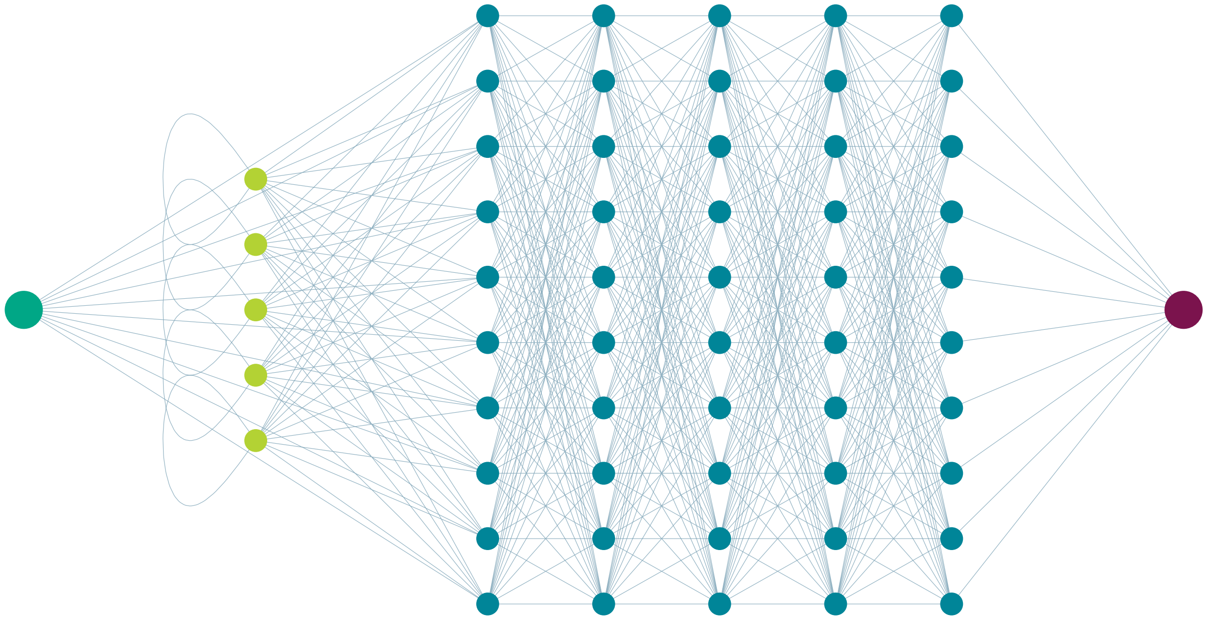

7.10 Recurrent Neural Networks

As we discussed the workings of other networks, we often found analogies to the human brain and learning quite helpful. Recurrent Neural Networks (RNN) try to mimic another important factor of how humans learn. We do know that our environment is not constant, and learning does only happen if some time passes. We are able to make decisions based on past states and how we want future states to be. This only works because we can remember things. The initially described NN has no idea of its past states, this is what changes in RNN.
The illustration above shows a schematic of a RNN, where the green nodes represent inputs, the lime-green recurrent layers, the blue one's dense layers and the wine-red the output. Recurrent networks loop previous state information back on themselves, so they can record the hidden state they were in in a previous step in time [Rum85].

Besides input information they also require information about time, or a series of values which occur after each other. Consider this example. If you have data points over time, the oldest is fed into the network and the output of the recurrent layer looped back and fed with the second value, then the last and so forth. The oldest value in the series is also contributing to the overall prediction. Herein lies the problem. The weight of the loopback is hard to adjust in such networks, because it is hard to find the gradient on back-propagation, which is either a vanishing or exploding gradient problem. A vanishing gradient means, that given enough loops the weights add up so, that the information, which is the gradient vanishes, much like an echo vanishes in a long tunnel. On the other hand, high weight values add up quickly, much like an avalanche picking up more and more snow forcing its way down a hill.

A solution to this problem is a Long Short-Term Memory Network (LSTM), which we are not going into great detail about [Sch19]. LSTMs can remember long-term memory data and short-term memory data. The crucial difference lies in the long-term memory, which is not as much and directly influenced by adjusting weights, which mitigates the problems of training RNNs. LSTMs utilize a set of gates to manage memory effectively. These include a forget gate, an input gate, and an output gate. The forget gate determines how much long-term memory is preserved, the input gate adds to the long-term memory comprised of the short-term memory and the input. The output gate determines the next short-term memory from long-term memory the current short-term memory and the input.

Further reading

A fun demonstration of how one could interface a robot and do human like handwriting:

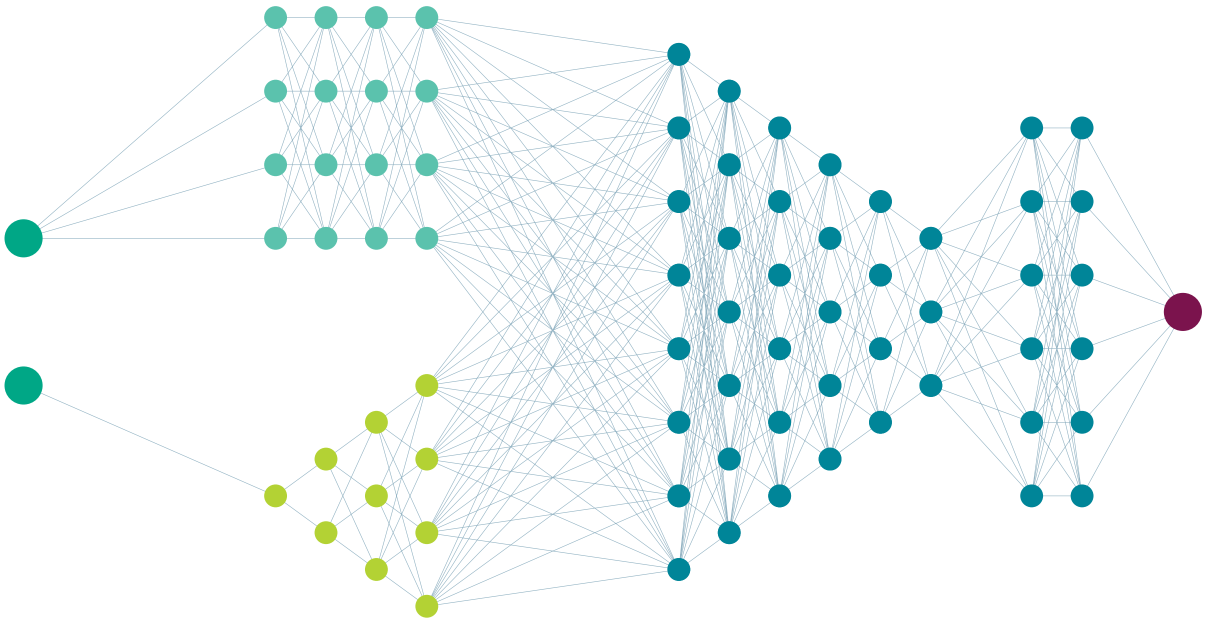

7.11 Generative Adversarial Networks

The last concept we like to briefly explain are Generative Adversarial Networks (GAN). As the name suggests, the Idea behind GANs is, it is in contest with itself. The building blocks of the network consist of a generator network, a network to process real data and a discriminator, in blue. The generator, in lime-green, produces fake date from a random input, it then produces a sample from that input. On the other hand, there is also real data fed into the network, for example images [Goo14].

The discriminator then tries to differentiate real data from fake data. Subsequently the loss of the discriminator is calculated and by back-propagating through the network the weights are adjusted.

Further reading

Try out StyleGAN – Human on Huggingface:

As the training of the network is commencing, the discriminator gets better in distinguishing fake from real data, as also the generator gets better at producing data, that looks like real data to the discriminator. At some point the generator might produce data that is indistinguishable from the real data by the discriminator. In short terms: The generator wants the discriminator to classify false data as real and the discriminator wants to correctly classify false from real data. A common analogy to the inner workings of GANs would be that of art forgers and experts.

The goal of the forger is to make a look alike that is indistinguishable from a real art piece of an artist. The experts on the other hand want to tell if a piece is genuine or not. They find clues of why the piece is not genuine and the forger will improve its technique, this goes on until either of them can't get any better, either at forging or telling fake from genuine.

7.12 Other types of Neural Networks

In the previous parts we discussed very basic principles of how NNs and different approaches for different workloads. As you might have guessed there are many more types of networks and mechanisms on how they provide the results and how their predictions are made. One of the more prominent examples of a different kind of architecture are Transformer Networks, such as Googles Bard, Metas Llama 2 or OpenAI's ChatGPT [Vas17]. Those networks often are a mix and match of different models combined into one, where each of the models is needed to do different tasks, such as image recognition and natural language processing.

Furthermore, in the realm of image generation, there are also so-called latent diffusion models, which basically generate noise until it resembles an image prompt, those models are exceptionally well suited for text to image. Networks and techniques can be combined with each other. If you want to combine RL with methods explained before, we would call it deep reinforcement learning. The actions an agent wants to undertake in those scenarios are determined by the NN and no longer simpler functions, this would be useful for when there are many actions available to the agent.

The same goes for other methods, the outputs of a CNN and a DNN could be combined in an extra regression layer to better predict certain outcomes. And this cascading or parallelization of models can be made even bigger, making the output of one model the input or one of the inputs of another and so forth. If you get the base ideas of what the individual tools are capable of, you can develop a sense of what you might achieve through remixing and matching. The flexible toolboxes are already there waiting to be utilized.

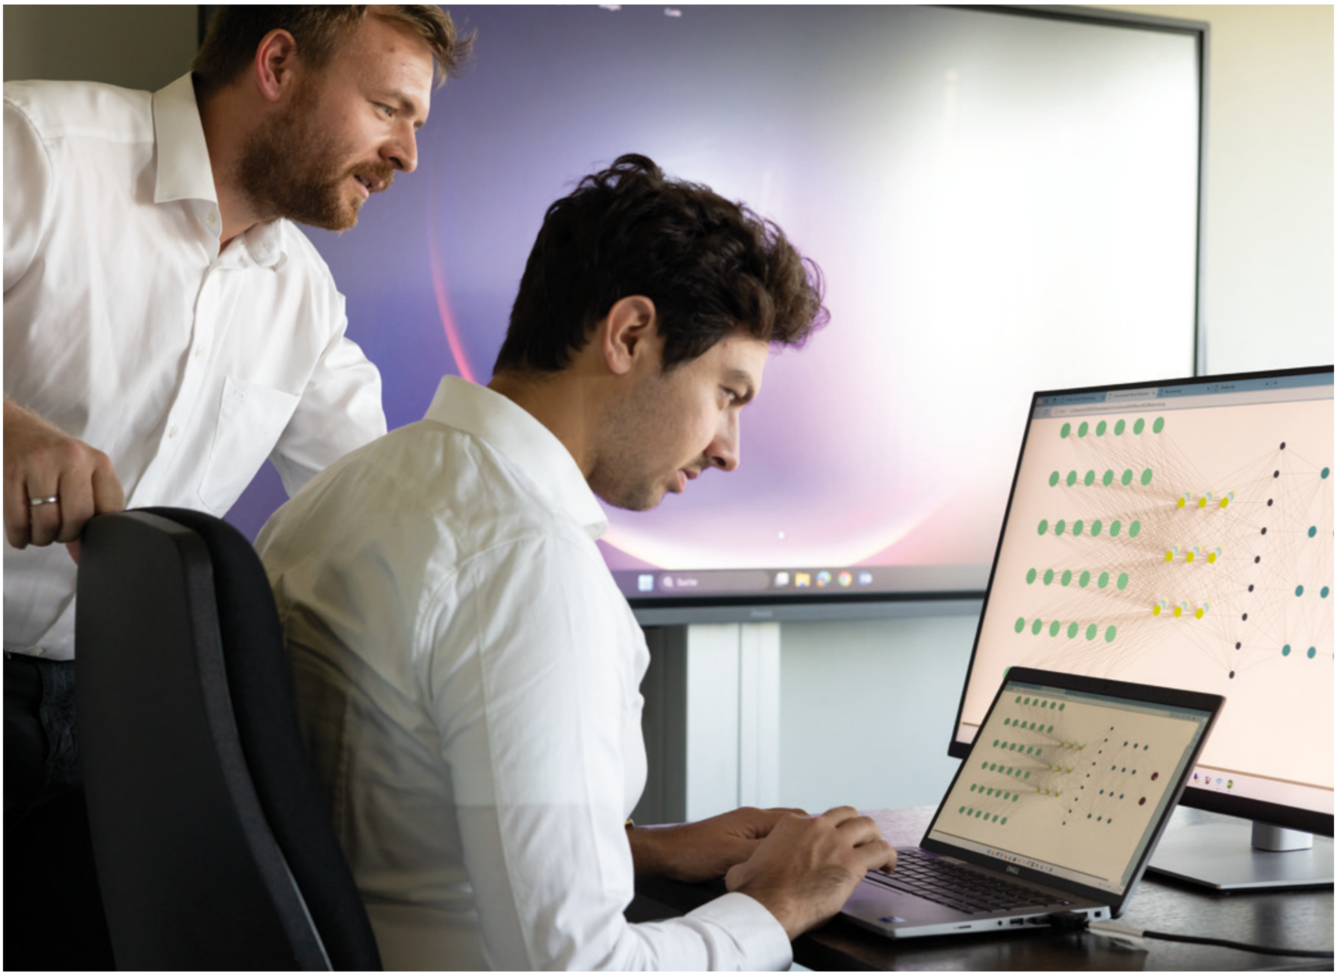

8. Methods and application areas

Having outlined the fundamentals of AI and ML, we will now discuss our approach to generating use cases, the rationale for introducing them, and specific examples built on various methods and challenges. First some groundwork is needed. AM is a broad field and consists of plethora of processes, machines, and procedures. In this Deep Dive we are approaching the use cases process wise from Laser Powder Bed Fusion (LPBF).

As on an industrial scale, LBPF is by far the most utilized and is nowadays considered the workhorse of industrialized AM. The second one is more of an assumption. We assume, that at least the data sources for the use cases are available, at some point feature engineering and data preparation is necessary, but we assume, that you do not need to build specialized sensors for a given process but have that covered already. The last one would be, that we are not looking for use cases on a petabyte data level, which would require a completely different approach regarding acquisition, pipeline, storage, processing, and evaluation. We are not developing use cases operating on »big data« amounts of data, and with that we mean the before mentioned petabytes of data. To summarize, we developed use cases for day to day AM operations either on an enterprise integrated scale, as in »the AM branch of the enterprise« or on a shop floor level.

All in all, this does not mean, that those use cases are not applicable for larger operations but might require a lot more involvement in terms of data engineering.

8.1 Design of the use cases

Current research and industry statements define challenges AM is facing right now [Liu23; Fra22; Amf19; Amp23]. Those include cost-per-part, sustainability, quality assurance, up-skilling workforce, new materials, supply chain management and automation. Each of those categories represent a challenge in further advancing AM from »niche high tech« to the »manufacturing of choice«. We worked out, what defines those challenges and discussed them with stakeholders in the AM industry and experts from different fields of the AI and data science community. For this, we choose to conduct a workshop with both groups.

This workshop consisted of a multi staged approach. We first elaborated, to get the AI and data science experts on the same page, what challenges AM is facing. This resulted in a fine granulation of sub-challenges in each of the forementioned categories. After that, we elaborated on those challenges. Firstly, we collected ideas of potential fields of improvements stemming from the challenges. Secondly, we took these ideas and asked: How are we facing the hurdles right now? What are the instruments in tools involved and what data is our foundation for an informed decision without the use digital state of the art methods? We then expanded on those questions and pondered on how to use our existing tools and methodologies and expand them rather than completely substitute them. This also includes how we do processes right now, without the use of AI. We advanced on this contextualization by expanding on selected ideas, finally conveying ideas into use cases. Involved in the contextualization and idea to use case transformation where also some thoughts about evaluation and insights into potential performance analysis of those use cases.

8.2 Evaluation design

Because you should never change a running system, we did primarily focus on the challenges described and from there we build our use cases. Thus, we defined our concrete use case benefit and what we consider added value. Secondly, we designed them in such a manner, that they ensured an otherwise taken care of problem would not be solved just differently because it was already solved before. In a nutshell: Either make a process more reliable, faster or cover a yet uncovered area and advance AM by utilizing completely new approaches. The use cases all highlight different problems and, at least on small scale scientific studies, have been rated by industry experts and yielded promising results in previous research. To achieve this, we questioned three distinct groups of experts regarding impact on the given hurdles. We also asked the experts to provide their opinion on the probable implementation difficulty and scalability of the use cases. Lastly, we sorted the challenges by their importance to the AM community. This resulted in a portfolio matrix, displaying the importance of each of the use cases for AM and the difficulty in implementation. Furthermore, we added the challenges as dimensions to the use cases expressing their respective domain. The obtained figures and dimension characteristics where then further discussed with experts of both domains alike.

»

AM is a key enabler to build future technologies. The same is true for AI.«

9. Challenges of Additive Manufacturing

9.1 Costs-per-part

There are several driving factors that factor into the cost of AM, that are symptomatic of new and emerging technologies. New technologies, that offer a significant advantage over other technologies, or previous generations, tend to be priced higher, then conventional ones. This is especially true if there are few competitors on the market. Although initial investments in AM capability is steadily decreasing and with more and more competitors on the market. This trend is presumably going to continue, material prices are still higher than conventional bar stock for example, which gets even more pronounced if you are working on simple, close to stock parts. With dropping additive material prices, more and more usage areas and geometries will economically be feasible with AM. Factoring into cost is more then material price. One of the driving factors is the complexity of getting a process to have stable, continuous, and high yields with an expected and satisfactory quality. In these chains, the complex parts engineering section directly factors in the low throughput of processes and high machine occupancy. You want to get a part the first time right, even more so if it is an expensive part of complex geometrical shape. This often means that those parts, where high yield is critical are going to be printed as a single part on a given printer, increasing overall machine time needed, powder to be recycled and machine volume utilization to be low. Which also ups the time needed to finish a sequence of jobs on a given machine.
Most of that meaning is, that AM can not play out is core strengths on complex geometries and less material waste, as material prices are too high, yield is comparably low and processes are slow leading to high costs. There are also other factors involved in the high cost-per-part. AM is still a niche high technology, suitable only when other means of production do not achieve the desired properties. Furthermore, maintenance of machines is rather costly, this is partly mitigated by AM not needing any expensive tooling if it is used as the only manufacturing technology involved in producing the part. Often that is also not the case, as most additively produced parts need subtractive post processing [Amp23].

9.2 Sustainability

Legislature, customers and manufacturers are shifting their focus to more sustainable means of production [Bun23]. Although material waste on a single process is smaller than most conventional production methods, other hurdles are in the way of truly sustainable AM. Especially the legislative focus shift to limit carbon emissions in industry is challenging.

The energy consumption of additively manufactured parts is, in some cases orders of magnitude higher, then conventional manufacturing. This even is apparent if you are looking at the raw power consumption data from conventional manufacturing technologies (e.g. mills or lathes) in comparison to LPBF systems, where only high-volume subtractive manufacturing machines come even close to the energy demand of AM machines [Yoo14]. To be clear, this differs greatly between different processes. This is equally true for raw materials. The process of powder production is energy intensive and another step that is not needed in other manufacturing processes. Although mixing raw powder with used powder is a valid approach in some cases. It is not entirely feasible for every material, especially if oxidation due to handling, sieving, and mixing is a concern. This often leads to the need to recycle powders similarly to how other metal scraps handled in the industry, negating the advantages. As decarbonization gets more important by the hour, AM and other branches need to overcome their shortcomings and play out their strengths even more, reducing resource consumption and overall carbon emissions along the entire value chain.

9.3 Quality assurance

In quality assurance defines that certain standards in production are met, they should be met consistently. In AM quality assurance is particularly difficult as a lot of defining factors come together. The first one being, that most of the processes are just now beginning to mature to a point, where you can get consistent and reliable quality while maintaining specific process parameters.

A second factor is, that processes stability and quality can differ immensely while doing new parts, even on the same machine. This is partly due to the complexity of process parameters and the underlying physical principles involved in AM but also part due to the lack of common standards and databases regarding process parameters. Material quality can change from batch to batch and from manufacturer to manufacturer. Furthermore, there are only self-imposed quality standards. This leads to a mediocre repeatability even with the same process parameters, if you are introducing even more parameters, like you would with sinter-based methods, achieving an understandable quality is a challenge. There is a wide range of sensors and methods available for monitoring ongoing prints, such as thermal imaging, cameras in general, flow sensors and many more. Using the data to comprehensively predict an outcome of a print poses challenges, as it requires deep process understanding in a non-standardized environment. Trials to get quality assurance in process working are generally expensive and time consuming and differ from subtractive manufacturing in that manner a lot, where you can home in your process while setting up the machines for relatively cheap. Running in a process to a satisfactory degree involves a lot more time and money to be spent on a particular part. Lastly, the processes are highly manual and in-depth quality assurance requires to test single parts, often destructively. Furthermore, stock manufacturers in conventional manufacturing often guarantee their quality or their stock is certified. This takes some risks away from the part manufacturers, as they are not liable for what the insides of a material look like. The work is happening on the surface. The freedom to design the inner workings and structures of additively manufactured parts adds another dimension to quality assurance, not only being responsible for the surface of a part, but its entire volume.

9.4 Upskilling workforce

Work on vocational training and anchoring AM in technicians training has just started. But as of now, there is still no formal training on additive machines. This is just one of the puzzle pieces in enabling the workforce on additive processes and made much more complex through lacking in other fields.

One of those fields is the dire need for standardization. This lack of standardization starts with the general usability of the machines themselves. In traditional manufacturing you find a uniform iconography that is identical or very similar between different machine manufacturers and machines. The second being the code that governs the machines behavior, which overall is described by g-code. This goes even so far, that experienced technicians can manipulate machine operations on a code level. Both enabling fast switching to newer machines or different machines on a shop floor altogether. Further difficulty in training the workforce arises from how different even machines from the same product line might behave in operation. Successful training is only achieved on a company level or through instructions by machine manufacturers. Besides technicians operating machines and having no formal training available to them the engineering part is also quite diffuse. Product designers and mechanical engineers alike lack the formal education and thus the necessary skills to adequately design for AM and exploit the strengths of the technology. This partly can be attributed to slow adaptation of new technologies in higher education, but also to a lack of guidelines on how to properly design for AM.

Knowledge on proper designs often is a key advantage over competitors in the tight space that is the additive industry, but also leads to a general lack of knowledge about the technology outside of the industry. Overall AM is lacking proper training, which is partly due to the emerging nature of the technology, but a big part can be assigned to a lack of open standards and guidelines for training, design, and machine interoperability.

9.5 New materials

Acquiring and testing new materials, or even materials that were used as stock before in other manufacturing methods presents itself as a highly manual and expensive process. Materials used in conventional processes might not necessarily display the same qualities after being used in LPBF. Materials react differently to temperature changes and heating and cooling cycles, might require different atmospheres to be processed and even not work at all in the process they are trialed to. New materials are also restricted by the availability of powders. Bluntly speaking, you could just get about any feedstock of material for a milling process and keep testing different tools and techniques until a satisfactory result arises. Most of the additively manufactured materials are at some point derived from traditional products because they promised to be easier worked on then with traditional methods, for example the goal was to limit the consumption of tools on especially bad to machine materials such as titanium or nickel alloys.
There is certain additive only materials development, but its not the norm. Lastly, novelty is plaguing AM again. Whereas traditional manufacturing had decades to develop materials for casting, machining or welding, the development for AM just does not have this headstart.

9.6 Supply chain management

One of the core strengths of AM is, that it can produce intricate, structurally sound, and high-quality parts with a small machine shop and comparably lightweight raw materials. This initially makes it ideal for distributed manufacturing, or on-site manufacturing. Obstacles arise from the need for accreditation of processes or parts. Industries which would profit most from on site manufacturing usually have the highest requirements regarding accreditation, such as the aerospace industry. This usually goes hand in hand with long procedures to get new processes accredited and in case of the aerospace industry, flying. Recent events have also shown that supply chains, not specific to AM, often lack the robustness and resilience to keep supply up when interrupted by external events. As logistics will inevitably get more expensive, if on road transportation with trailers remains the workhorse of the transportation industry AM could heavily benefit from distributed supply chains, which are not as dependent on large quantities of material shipping. Still, most manufacturing is heavily dependent on exact processing stages and order concerning the whole manufacturing chain and AM is no exclusion.

9.7 Automation

Although AM is often pitched as the digital manufacturing technique, it is filled with manual process steps. This can be, in some part attributed to the comparatively novelty of the whole process chain. But in automation AM is its own adversary. The number of different shapes and diversity of processes and post processing steps makes automation difficult.
A single job running on a machine can have a multitude of parts, where each part needs to be post processed differently. If every part on the substrate then has a different geometry, automatic extraction gets even more difficult without intelligent systems in place, either in robotics or in planning. Material management is also challenging. E.g. in fused filament fabrication, filaments are handled quite easily by automated hoppers and other conveyor systems, but micro particulate powders pose a different challenge altogether. Automating powder delivery and handling is challenging. Environmental concerns regarding impact on air quality needs to be considered as well as explosion risks and even how to get the material from a to b without contaminating the conveyor systems with different materials. This not only applies while getting the powder into the machine, but also for extracting powder and feeding it back to a recycling loop. Those steps are predominantly done manually these days. Other fields, with lack automation is process control, monitoring and even tasks as simple as scheduling machine operations, which may sound benign, but usually humans are good at these tasks, so often it is done manually.
In conclusion automation would also benefit heavily from standardization, common interfaces and a simpler data handling.

10. Application and evaluation of AI in the industry

As of now, there is a lot of research underway concerning AM and AI. In the Scopus literature database, articles and journals concerning themselves with a combination of both topics are on a steep rise. Looking at each of those topics alone, doesn't hinder the trend, the growth is even more substantial [Sco23]. The main research focus is process related, starting from in-situ defect detection to classification and research on new materials and materials research. As a field that relies heavily on simulation-based approaches, a new field of research emerging in AM and AI are Physics Informed Neural Networks (PINNs) [Has21; Pra22; Sil19]. The second field undergoing a new wave of research is on how to manage production and how to further advance our production environments with AI, where AM might not be the central process, but could benefit a lot from current research [Sch22; Zha21].
We will guide you through the first use case in more depth, as this was the use case, which was in favor of the AM community. The first one being in-situ defect detection and quality prediction, this use case concerns itself with monitoring printing processes and make predictions based on what can be extrapolated from the information collected.

Papers published on AI and AM

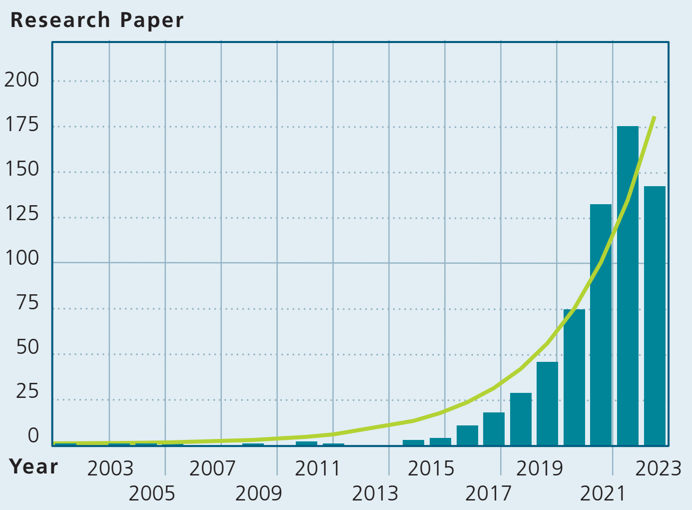

Figure 8: Trend AI in AM; date of evaluation: September 2023

The second use case we came up with deals with the energy consumption of additive machines and additive production in general, this may come to fruition, as more and more emphasis is being put on carbon imprint of industrial processes. The third one is how we could make production planning easier, especially in a job shop environment, where traditional approaches often struggle due to the immense compute power needed. This approach could free up human resources to tackle more human adequate tasks. The fourth case concerns parameter prediction, at its core delivering methods and tools for a first-time right set of parameters for a given machine and material. The fifth usage of AI puts emphasis on the scalability and routability of additive processes and tries to deliver on a method to find the adequate manufacturing routes for a given part not based on what material might have been ordered but based on the usage and the desired properties of that part. And lastly, a use case, that leverages the novel developments regarding PINNs to supplement classical fluid dynamics and finite elements simulation for nesting and build job preparation. This use case is really at the forefront of what might be possible with an multi staged, multi discipline approach.

We have structured our use cases based on their primary purpose to begin with, we defined probable results and success criteria. Furthermore, we defined scaling opportunities and implementation stages, the first one making statements about what a fully implemented system might be useful for in the future, the latter describing how we could implement the approach piece by piece to get the desired result.

Another element being the actions and methods involved with the use case and certain prerequisites, risks, and limitations of each approach. Lastly, we assessed the suitability for LPBF and if not suitable, what specific AM process might be worth considering. This also involves the respective domain knowledge to get the tool underway, also concerning users and stakeholders. After we have discussed the use cases, we like to give you an overview about expert opinions regarding the use cases. This involves sorting them by their respective domain and potential influence on the hurdles AM is facing right now and to give you a brief analysis on the potential and implementation difficulty of those use cases.

» The flexible AI toolboxes are already there waiting to be utilized.«

10.1 Concrete use cases of AI in Additive Manufacturing

10.1.1 In-situ defect detection and quality prediction

On our first use case, we want to be more detailed then on the others, graphically and content wise. Workshop and questionnaire results showed, that one of the most pressing challenges is process stability and thus, monitoring and making predictions in-situ. This directly translates into what we are expecting to be the purpose of this use case, namely, to make valid predictions on the quality and thus potential defects of parts, while the printing process is still ongoing. It might sound like a noble but potentially very hard endevor, but there are several advantages of the methods we like to present. But first, we need to define a scope, this translates to the probable implementation stages of the use case. The interactions in each build job are highly dependent on the process parameters, but also on the parts interacting with each other thermally. Interactions inside the build job also depend on the material and obviously on the process involved. A first stage could therefore encompass simple geometries, like test specimen of a given material in LPBF. In first stages, the goal is to make predictions on a well understood process and material. This could then be broadened to different materials, geometries and packing strategies, if the first stages are well understood and working within set quality parameters. Those quality parameters are defined as use case success criteria, with the first being, the monitoring system should reliably detect in process failures, and later act accordingly. What defines failures and what differentiates a good from a bad print must be assigned as the target of the given system and is dependent on e.g. the application area, material or part certifications.

Careful consideration of quality criteria is essential. One example might be the porosity must not be lower then 99.5%. Another success criterion would be the exclusion of false positives, which are a function of the sensitivity and reliability. A given system should only react when it is sure that there is a failure. This could also be measured by a threshold, which in turn would be dependent on what the level of acceptable non-detection is. Normally, you would use the current state of rejected parts to be lower then they where before. Lastly, to use all available scaling opportunities, a given detection and advisory system must be real time capable, meaning that an adequate reaction and compensation time is key. One example of differences in real time is the traffic collision and avoidance system (TCAS) found in aircraft, which must advise pilots basically instantaneous, audible and commence immediate action. A completely different control system is the temperature control of a home, which is quite sluggish and room temperature usually does not want to react instantly to changes, which is neither needed nor desired. Polling rate is thus not a huge issue and you can utilize longer time steps. Both systems are considered real time, as they are operating in an adequate time frame for their given job. Scaling wise, we could think of different opportunities, the first being, to push an advisory system to an automated correction system, handling minor corrections in process and monitoring if those corrections yielded a success. By monitoring the error occurrence and need for corrections, you could get insight in degrading machine performance, making the prediction tool a valid component of predictive maintenance. The same is applicable for raw material properties and quality, an unusual high error rate for a certain batch of powder might hint at a bad batch overall. A tool being able to achieve this requires certain prerequisites, with the first being a very good understanding of the given AM process. You need to have experience in the field and better still, a wide variety of already acquired monitoring data. This makes having or building a robust data pipeline obligatory. On a sensor level there are several approaches available. For LPBF monitoring based on sensor fusion has shown to yield potential. [Wan22a]

High resolution imaging could form a baseline for image processing, thermal imaging a second layer, this could also be further elevated by accompanying x-ray sensors. The imagebased sensors could be supplemented by sensors measuring discrete variables, e.g.gas flow, chamber temperatures or oxygen level. All data must be coupled with the time of their occurrence, as all printing processes are time dependent and errors in each layer can also develop an impact further down the process chain. [Tah23; Gai20] The raw data must undergo an engineering process; you do not want to hand over a plethora of image data and discrete data points to a NN architecture. If you are only interested in localized effects, shrinking the viewport could be a valid approach. Batch processing for labeling a time series also comes in handy. The approach we are suggesting involves a multi input long short-term memory CNN with discrete and image data inputs. The combination of multiple images has proven to be successful, for example combining multiple inputs in a CNN showed better results, then single input CNN for example in human emotion detection [Che22]. You would use thermal imaging data, the visual data and for example x-ray or other spectra as inputs to a CNN. To further improve the accuracy, using an autoencoder to reduce the features of the images to a more manageable size might be beneficial. As processes are time dependent, you would then employ an LSTM for satisfy the temporal nature of the process, inputs would be the output of the CNN and discrete data. There are certainly other approaches and techniques you could use, one strategy to get a better understanding of certain structures

Use case: In-situ defect detection

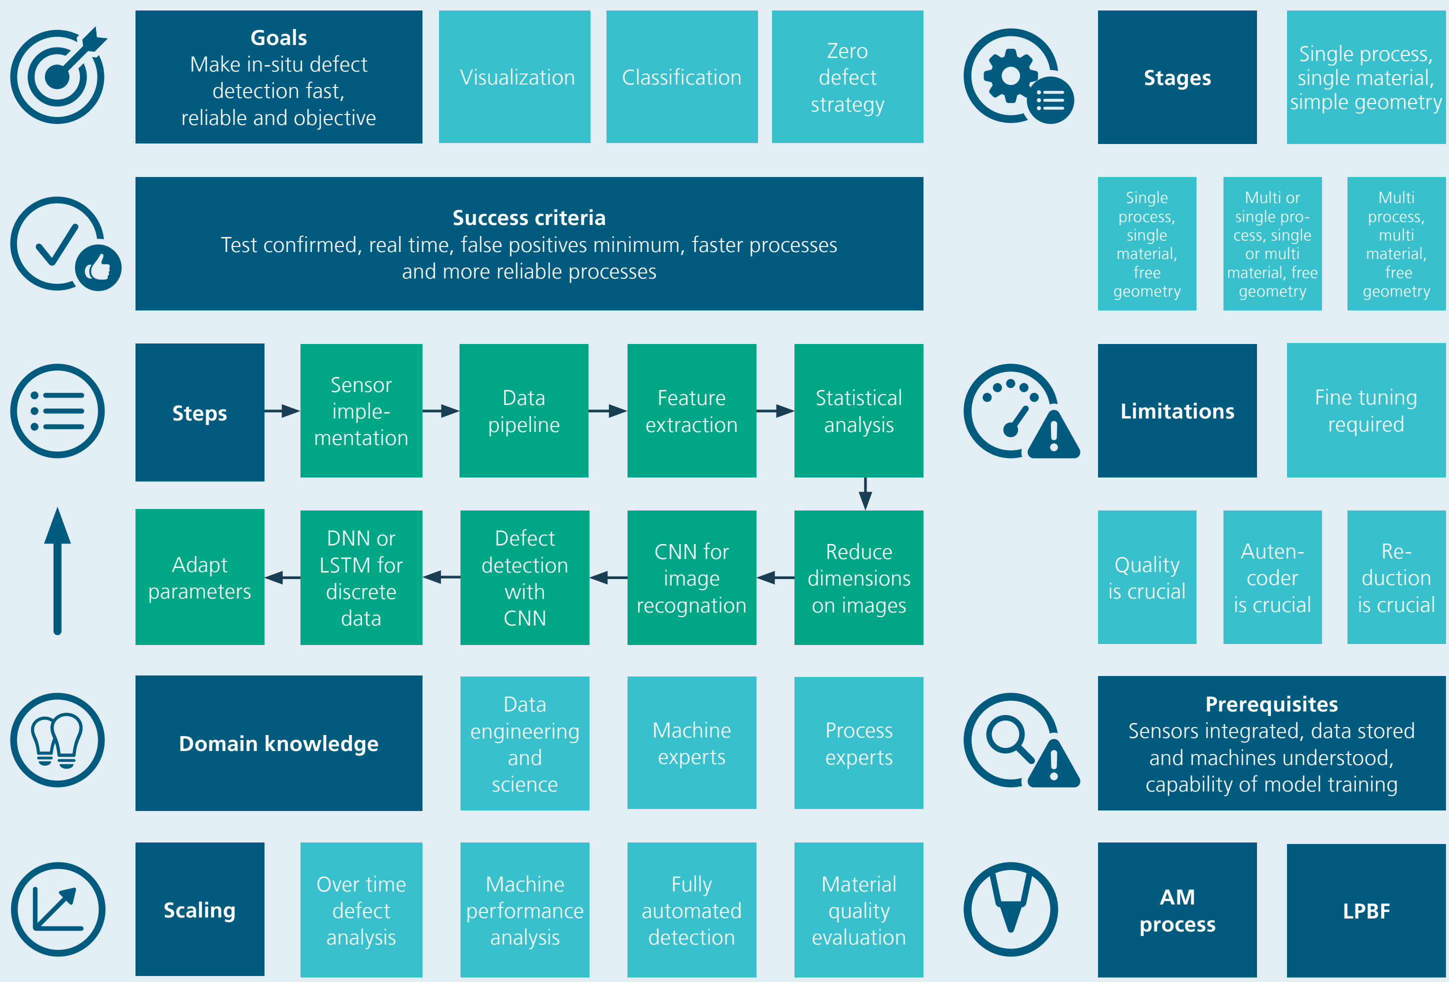

might also be to use a NN to inpaint future images [Zho24]. Acquiring the data might be challenging, as collecting the appropriate images for training requires a good amount of storage and exact timing is crucial. For this kind of work, you would need a wide variety of domain knowledge, in this case, first and foremost is a way to interface with a machine and collect the data, so either you need to deploy your own sensors or use the data the manufacturer hands over to you, which might not be sufficient. There will be several engineering challenges involved and thus engineering knowledge is needed. For data handling and consequently building a model you would need specialist's knowledge in data engineering and data science. This also includes the evaluation and, in the end, visualization of the data you acquired. This could be used to build an interface on top of the model output. Even with a sophisticated user interface and data processing, you will still need a technician to factor in other variables to complete the picture, act accordingly and supplement the tool. As for the risks, with ML tools, especially NNs, transparency is not always guaranteed, if high standards are required. For example you could deploy Shapley Additive Explanations (SHAP). SHAP tries to understand the contribution of features in each model for output importance. [Lun17] User groups and stakeholders involved in this use case are certainly manufacturers and machine producers, but also material suppliers, as the tool could gather valuable performance insights.

One of the more pressing issues is, that CNNs require a lot of fine tuning and are prone to overfitting. Furthermore, the acquisition of samples might be tedious and to some degree expensive, as you would need good and bad printing results and have measured discrete data in process. To mitigate those limitations, or when you might hit a certain ceiling with a use-case, you should subdivide the use cases. In this example, even classifying prints on visual cues is still a valid use case, for monitoring and intervening manually. The same goes for evaluating, even without AI, you could use the collected data for automatic statistical analysis and build a database for later projects and production insights. Thus, the steps to get fully automatic process correction underway include data acquisition, outlier detection, quality prediction, material properties prediction and in the end automatic process correction. Each of those would be an invaluable asset in production.

Brief use case descriptions

Energy and consumption forecasting	Accelerated production planning	Parameter prediction and development
Make production predictions for energy demand	Automated multi machine scheduling	Make material qualification and parameter identification faster
Time series, regression, NNs, random forests, DRNN	Real data observation RL in combination with DNN	Image data and discrete data. Physical material properties, parts properties and parameters: CNN and DNN

10.1.2 Energy and consumption forecasting

The second use case revolves around the energy usage and green production in AM, here, we are trying to solve several challenges currently occupying AM and manufacturing in general. For simplicity this and the other following use cases will be much shorter in detail. The approach to forecast energy consumption and therefore carbon imprint involves detailed LCAs and strict production monitoring as well as external sources. Forecasting energy consumption yields potential in giving out more detailed LCAs for your manufactured parts, which comes in handy if regulatory bodies strive to make carbon content for manufacturing a bigger part of the overall calculation. The second implication would be, that you could buy energy on demand and incorporate energy consumption in production planning and even logistics calculations for distributed manufacturing. Lastly, you could automatically buy energy on exchanges and have a much lower price and more sustainable production, as low energy prices often occur with high availability of green energy. The last step could be to automatically control machines based on energy availability and could potentially increase production resilience. The benefits in AM are obvious, machines energy consumption is much higher than subtractive manufacturing. It also leverages the possibility to manufacture where energy prices are low and make logistics less costly. If you also factor in resource consumption and make an end-to-end LCA, including product lifetime, you could get very comprehensive results regarding carbon footprint. This would need a robust production planning in place, which is our next use case. This use case could potentially benefit from LSTMs or time series forecasting.

10.1.3 Accelerated production planning

Manufacturing must cope with range of obstacles and unforeseeable circumstances that requires a high level of flexibility in an ever-changing environment. Production per se is more organic and grown as strictly planned due to the entities and number of processes involved in manufacturing goods. This especially applies to AM, where parts do not only not run on a single line but often get matched and mixed with other parts and charges. Yet another challenge is, that managing production is often done by a human, who knows all the ins and outs of the manufacturing route, thus leaving companies vulnerable if said employee would leave the company.
Solving this problem with classical computational methods is rather difficult, as production planning counts as NP-Complete, thus not deterministically solvable in polynomial-time leaving only approaches such as metaheuristics or genetic algorithms, which do yield good results – but often are difficult to expand and incorporate a multitude of objectives. Furthermore, they are not quite capable of doing large planning jobs, as the possible decisions of a given system rise with the number of states it can occupy. For production planning, we are suggesting a mix of observing real production plans and simulating new datasets. This would be a deep reinforced learning approach. Scaling could involve starting with a few machines and use inputs and outputs to scale into other domains, such as post and or preprocessing. A crucial benefit is the speed in which once trained networks act, they can not only predict but also react flexible to changing circumstances, which in turn would free up valuable human resource for more important work.

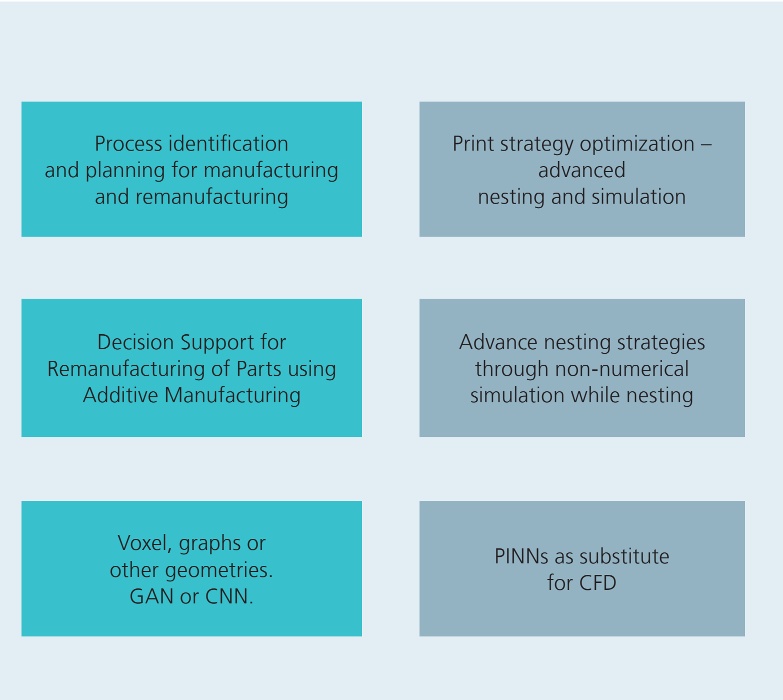

10.1.4 Parameter prediction and development

Our fourth use case concerns itself with parameter prediction and development. This ties directly into the first and leverages some of the technologies, especially in the sensor and data domain. The data involved here is more sophisticated, as you would also want to analyze the microstructure of a given sample. This is also a more static approach, which wants to map given sample image data directly to the process parameters involved in producting the specimen. Here we rely largely on image recognition, which on itself could cluster a given dataset in quality levels. This could advance parameter development. This is complemented by discrete data, such as the process parameters. For further improvement, labeled data is required. Depending on the direction, from parameter to material property or from property to parameter, you could predict one or the other. It would also be conceivable to make prediction on how a structure would look if you will further advance the process by utilizing a diffuser or a GAN. Depending on how deep you would go into material science a system that can come up with new alloys solely for AM would be the masterpiece. The feasibility of this is shown in medicine and materials science, where new drugs and protein structures are conceived by AI [Sas23; Wat23].

10.1.5 Process identification and planning for manufacturing and remanufacturing

Within this use case lies an amalgamation of production planning and materials properties prediction. Some systems allow customers to choose different materials for their on-demand prints. This is widely employed but demands a degree of knowledge from said customers. One use case might be, to implement this expert knowledge into a system and let it decide, depending on usage and load cases and geometries, what material is applicable for a given job. Albeit this would require a lot of parts and geometries already manufactured, with some sort of indication if said parts fulfilled their job and or had the right properties, it would open the way for AM into a wider audience and improve remanufacturing of parts on a large scale. One example is remanufacturing of spare parts for discontinued machines or vehicles. Materials have developed a lot in recent years and many polymers were simply not available at the time of building older machines. You could easily get the same material properties from a polymer today, that you only got from a metal 20 years ago. This opens completely new manufacturing routes. The same goes for load balancing in each production. Maybe another alloy, that is available would offer the capabilities needed but neither you nor the customer thought about using it. You could even go so far as to use a different process altogether. There is no reason to use LBPF over directed energy deposition or metals over polymers if the desired properties could be archived by either. Managing remanufacturing routes and advising customers and manufacturers alike on processes is the goal of this use case. There are numerous ways to archive that, but using CNNs to analyze the geometry and compare that to the desired properties seems the most feasible.

10.1.6 Print strategy optimization – advanced nesting and simulation

The last use case is about simulating and developing novel nesting strategies in advance. The big advantage of this use case is, that you would not really require real world data, as we have simulation tools, that yield exact results by calculating partial solutions to the governing differential equations such as the Navier–Stokes equations describing fluid dynamics. Those constraining equations could form the basis of a voxel based PINN simulation approach. Training of this model is conducted, by simulating different geometries and nesting strategies and using the data to train the model itself. First steps would only consider simple geometries, which later could be more advanced as the model gets build up. The goal here would be to shorten simulation times and proposing nesting strategies for the desired parameters. Another application would be a fast check system, that does not require long simulations that take hours or even days, to check if technicians and engineers made some process braking errors during nesting. The biggest challenge is the simulative expense for generating the data and to get the level of detail right to fit the demands of AM. Feasibility has been shown recently by simulating whole windparks with PINNs [Bho22].

10.2 Potential assessment and evaluation

The aforementioned use cases obviously need some context, first in foremost, the two largest challenges AM is facing right now are quality assurance and automation. Both are covered in the use cases, but only covering the digital side of things. Even the most sophisticated digital systems, can not fully solve challenges regarding automation, if we do not change the physical processes, such as parts handling or powder handling. Clearly there are many improvements being made, especially in terms of computer vision and robotics, but we are not tackling those here.
That said, each of the use cases has a different impact on the specific challenges, we asked AM and AI experts alike, where the use cases would make the strongest impact. All of the use cases, seem to target automation to some degree, while most of them are specific in one or the other domain. For showing how use cases would potentially impact the given challenges, we used spiderweb diagrams, where each of the web edges shows the value for the given challenge, with one being low and five being the highest.

The diagrams show, that the best overall performance lines up very good with the demand we diagnosed before. Each of those diagrams shows, that all use cases to a certain degree lower the cost per part, this goes to show, that the overarching challenge AM is facing is the inherently higher cost, compared to other manufacturing methods. Each of the use cases has a different strong point, which aligns with our preliminary definition of said use cases.

Importance of challenges

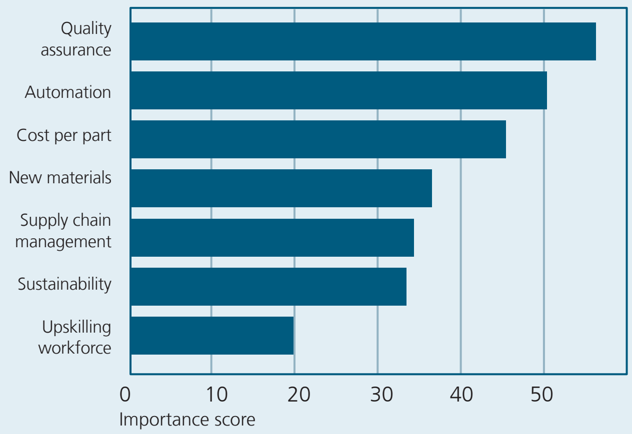

Figure 9: Perceived importance of the grand challenges

While energy forecasting focuses on sustainability, advanced nesting and simulation focuses on the automation aspect. We also scored the use cases based on the challenge ranking. This roughly is described by the area under the spiderweb, multiplied by a weighted score from our assessment of we showed in the bar graph at the beginning of the paragraph. This data is used, to devise an overall weighted score of the use cases. A second group of experts also has judged the development complexity, here we applied a weight consisting

Overall score for different categories

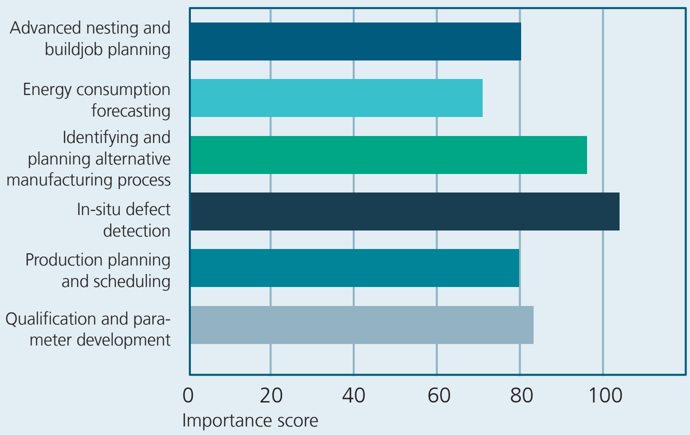

Figure 10: Overall score / Implementation height matrix

Normalized portfolio matrix

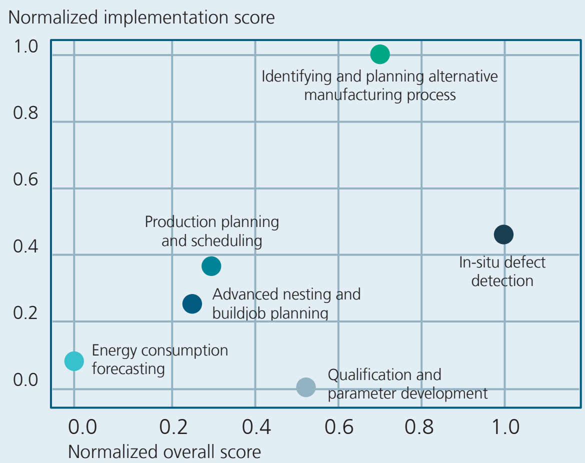

Figure 11: Overall score of use cases

Challenge relation of use cases

Figure 12: Challenge relation of use cases a) energy consumption and forecasting, b) parameter development, c) advanced nesting and simulation, d) in-situ monitoring, e) manufacturing rerouting, f) production planning

of ranked potential development and implementation cost. Both resulted in a portfolio matrix consisting of the impact, described by a normalized overall score and the implementation difficulty described also by a normalized score from the aforementioned values. The values seem to be very far apart in the graphic, this is for visualization purposes.
To give context to distance, we also included a bar graph showing the overall scores. Especially the second graphic shows, that the use cases do not differ that much in overall score but mirror the expectations we have for tackling the challenges AM is facing. The easy grab would be to automate the parameter prediction, it has a good overall score, and the implementation height is not too high. Energy forecasting is of now not a very pressing issue, but the implementation height is low and might be a good entry point for production planning and scheduling. Which comfortable sits in the middle ground, making it an ideal candidate to automate AM and get pipelines underway and acquaint personell with ML workflows.

11. Summary & Conclusion

Although AM is not as new as it used to be, it is widely still considered an emergent technology and herein lies big hopes of automation and digitization. The history of AM is layered with the same hurdles that also AI had to take. More in the realm of developing the machinery, but also managing and storing 3D files, which has also only been around since a few decades. In this sense, both technologies are somewhat similar. This is clearly shown by research in both technologies and by the connections in research in both fields. Both considered as enabler and key technologies in industry 4.0. Those connections become more apparent if you look at bibliographic analysis and the clusters both are involved in. This analysis is the parting point of this Deep Dive. So, allow us a few conclusions, starting with the challenges and concluding on how significant AI is and might be in the future of AM technology.

First and foremost, AM is facing a wide variety of challenges, with the most pressing being quality assurance and process and manufacturing automation. This leads to the conclusion, that other challenges might be important, but not as important as general automated and over a wide variety of parts and conditions stable production that can deal with new materials and emphasize and leverage the unique characteristics of AM. All those challenges influence the overall cost of additively manufactured parts, which obviously needs to be lower than it is now. We gave you a rough outline of the technologies regarding ML and AI that are out there, with a lot of them already in productive use, or in the case of medicine, helping patients. They all still require some degree of oversight from human operators in critical applications. Non-critical applications though are another field, where text to speech and speech to text applications regarding natural language processing are the bread and butter of modern device interactions. Advancements in computational power made a lot of the processing available locally and secure, where the only computationally intensive part is training the respective models. This further shows the realtime readiness of most applications.
AI has potential in all our daily, but also professional lives. We just scratched the surface of applications and use cases that where not possible and are just now becoming possible and doable. The mentality of how we use AI is different from how we normally would proceed in engineering. Most of the tools to implement AI are out there and available for you to test, try and implement. You don't need to mangle with in depth mathematical model building or trying out completely novel approaches. The AM industry is proficient looking beyond the box and as the technology becomes cheaper from a machine perspective, it is time to make it more industrialized.

Our exploration into the use cases AI can play in additive manufacturing demonstrates the diversity and depth of possible applications. From in situ defect detection to advanced production planning and process identification, AI most likely stands as a key enabler in optimizing and advancing AM. But even now, there are hurdles to overcome. Those are mostly regarding the basis of our data, if you implement a new technology, or machine, or process you should ask yourself, how can you gather data and thus insights from this process. This is valid, even if it is just for visualization purposes. This is especially true, as we ran into the big wall that is data availability and preparation. If an industry wants to stay competitive today, no matter the market, it must optimize and gather insights into its processes.

So how are we proceeding with that?
AM is a key enabler to build future technologies not feasible by subtracting material and letting it go to waste, its here to stay. The same is true for AI and unless we are not heading for some sort of Frank Herbert and Dune kind of ban of thinking machines, AI can be a great benefit, if deployed timely, ethically, and correctly. Future possibilities might involve AI being the designers, engineers or developers' copilot or right hand, freeing up valuable creative resources to make the tedious and not automated work better and human centric. Making processes involving repetitive tasks obsolete. If we combine this with the unique geometric freedom, and the very resource friendly nature of AM, we might even overcome the dreaded skilled labor shortage and be a bit more competitive than throwing valuable material away.

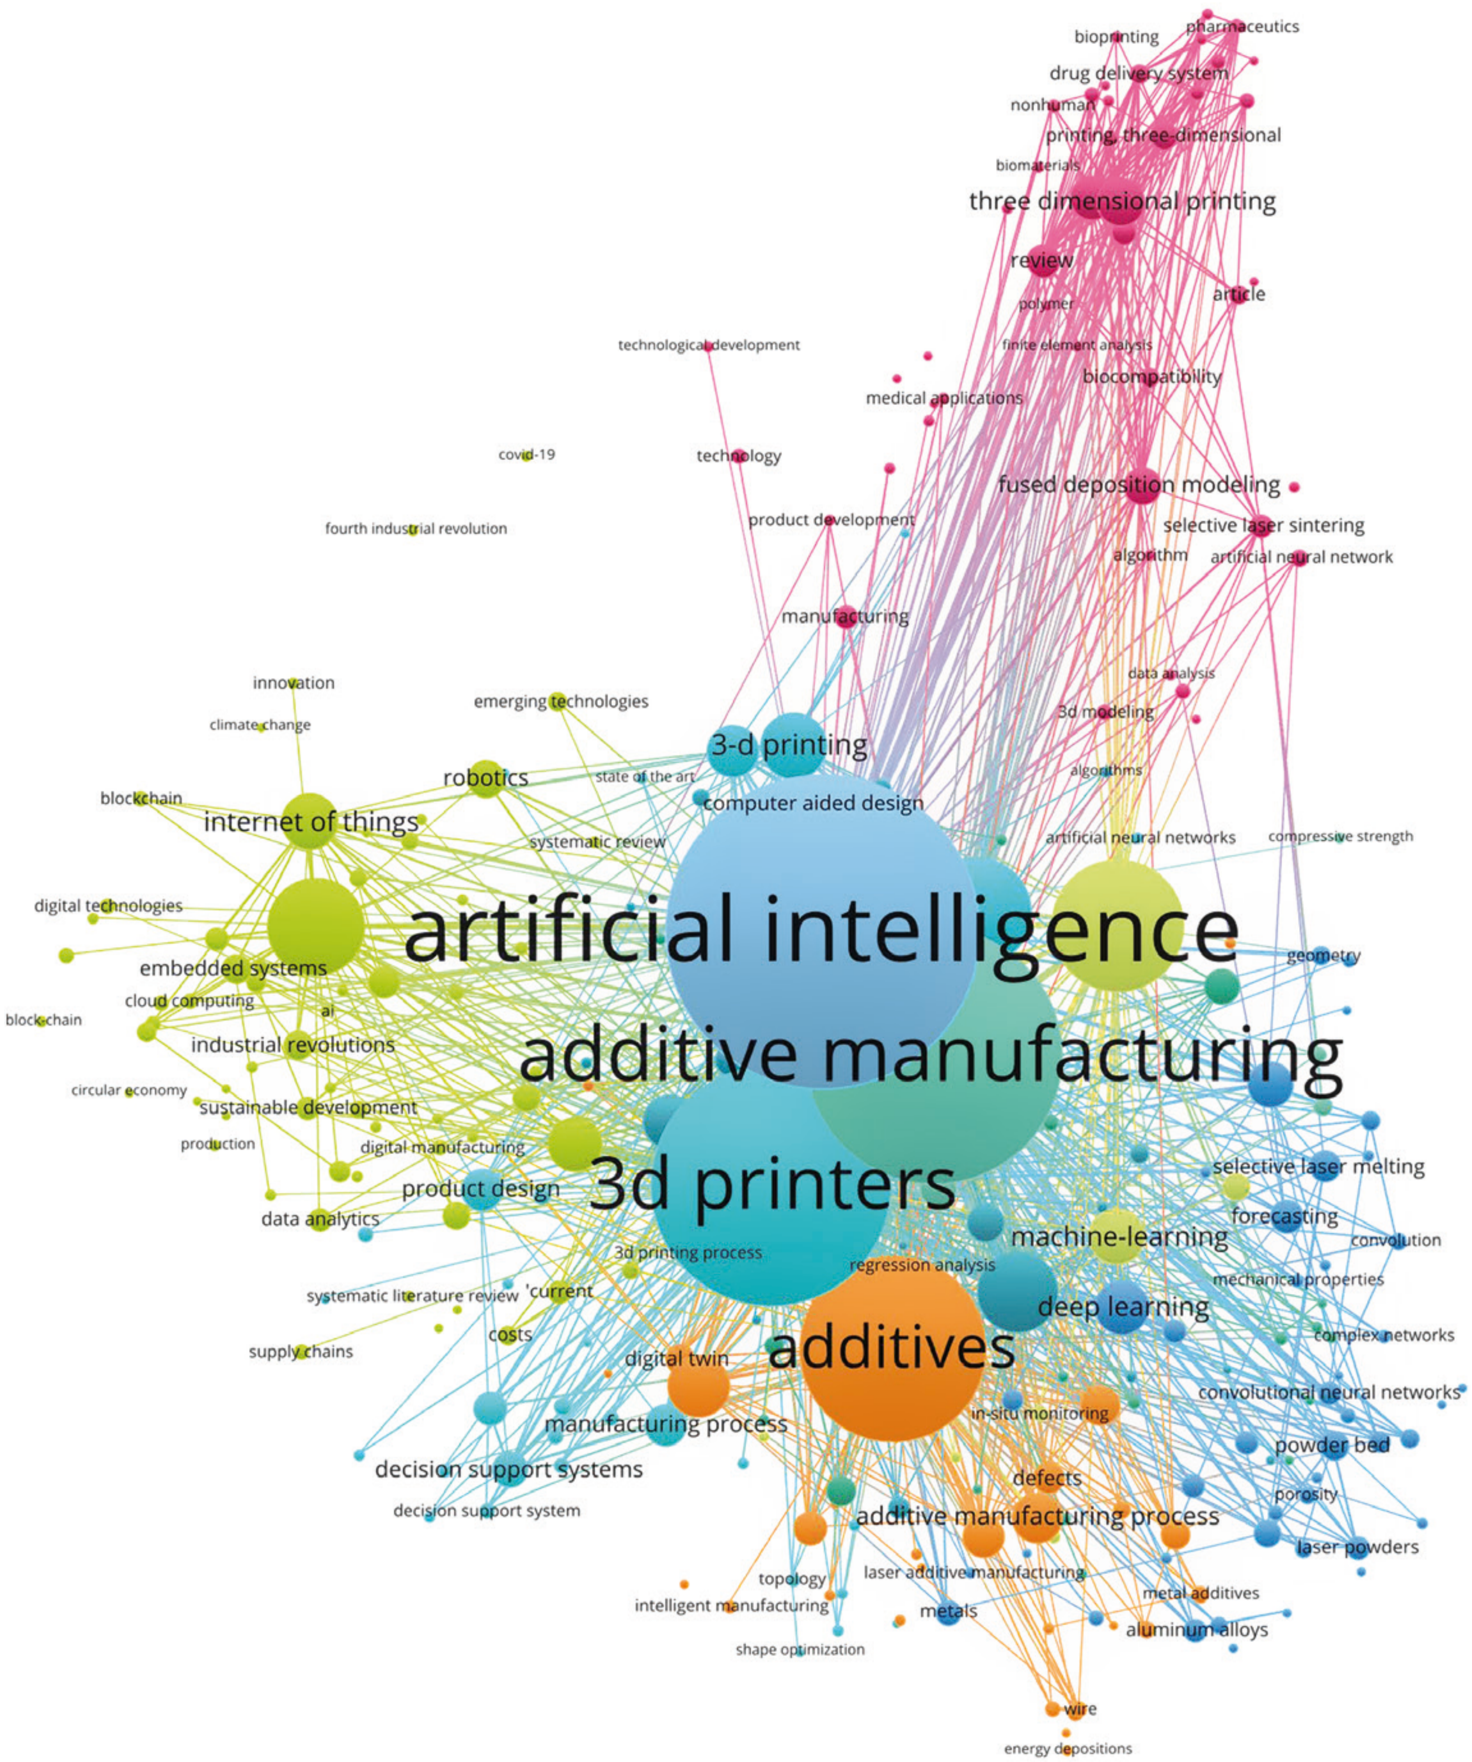

Figure 13: Scopus data knowledge graph and bibliographic analysis

12. References

[Alt20] ALTIMUS C. M., et al., "The Next 50 Years of Neuroscience", The Journal of neuroscience : the official journal of the Society for Neuroscience, vol. 40 no. 1, 2020, p. 101-106.

[Alz21] ALZUBAIDI L., ZHANG J., HUMAIDI A. J., AL-DUJAILI A., DUAN Y., AL-SHAMMA O., SANTAMARÍA J., FADHEL M. A., AL-AMIDIE M., FARHAN L., "Review of deep learning: concepts, CNN architectures, challenges, applications, future directions", Journal of big data, vol. 8 no. 1, 2021, p. 53.

[Amf19] AMFG, amfg.ai/2019/10/08/10-of-the-biggest-challenges-in-scaling-additive-manufacturing-for-production-expert-roundup/, 2019.

[Amp23] AMPower, Additive Manufacturing Market Report, 2023.

[Ban23] BANOULA M., www.simplilearn.com/tutorials/deep-learning-tutorial/perceptron, 2023.

[Bho22] BHOOMI G., blogs.nvidia.com/blog/2022/03/22/siemens-gamesa-wind-farms-digital-twins/, 2022.

[Bol18] BOLARD C., towardsdatascience.com/data-engineer-vs-data-scientist-bc8dab5ac124, 2018.

[Bro20] BROWN T. B., et al., Language Models are Few-Shot Learners, arxiv.org/pdf/2005.14165v4, 28.05.2020.

[Bun23] Bundesministerium für Wirtschaft und Klimaschutz, Kabinett beschließt Energieeffizienzgesetz, www.bmwk.de/Redaktion/DE/Pressemitteilungen/2023/04/20230419-kabinett-beschliesst-energieeffizienzgesetz.html.

[Cam84] CAMERON J., Terminator, USA, 1984.

[Cat23] CARUSO C., hms.harvard.edu/news/new-field-neuroscience-aims-map-connections-brain, 2023.

[Che22] CHEN P., ZOU B., BELKACEM A. N., LYU X., ZHAO X., YI W., HUANG Z., LIANG J., CHEN C., "An improved multi-input deep convolutional neural network for automatic emotion recognition", Frontiers in neuroscience, vol. 16, 2022, p. 965871.

[Cla74] CLARKE A. C., 2001: A space odyssey; a novel. London: Hutchinson, 1974.

[Drag06] DRAGANSKI B., GASER C., KEMPERMANN G., KUHN H. G., WINKLER J., BÜCHEL C., MAY A., "Temporal and spatial dynamics of brain structure changes during extensive learning", The Journal of neuroscience: the official journal of the Society for Neuroscience, vol. 26 no. 23, 2006, p. 6314-6317.

[Duk18] DUKOR O. S., medium.com/@stanleydukor/neural-representation-of-and-or-not-xor-and-xnor-logic-gates-perceptron-algorithm-b0275375fea1, 2018.

[Ert21] ERTEL W., Grundkurs Künstliche Intelligenz: Eine praxisorientierte Einführung. Wiesbaden, Heidelberg: Springer Vieweg, 2021.

[Fei81] FEIGENBAUM E. A., "Expert systems in the 1980s", State of the art report on machine intelligence. Maidenhead: Pergamon-Infotech, vol. 23, 1981.

[Fra22] Fraunhofer IGCV, ETH Zürich, An Additive Manufacturing Breakthrough: A How-to Guide for Scaling and Overcoming Key Challenges, www3.weforum.org/docs/WEF_Additive_Manufacturing_Breakthrough_2022.pdf, 2022.

[Gai20] GAIKWAD A., GIERA B., GUSS G. M., FORIEN J.-B., MATTHEWS M. J., RAO P., "Heterogeneous sensing and scientific machine learning for quality assurance in laser powder bed fusion – A single-track study", Additive Manufacturing, vol. 36, 2020, p. 101659.

[Gol16] GOLLAPUDI S., Practical machine learning: Tackle the real-world complexities of modern machine learning with innovative and cutting-edge techniques. Birmingham, Mumbai: Packt Publishing, January 2016.

[Goo14] GOODFELLOW I. J., POUGET-ABADIE J., MIRZA M., XU B., WARDE-FARLEY D., OZAIR S., COURVILLE A., BENGIO Y., Generative Adversarial Networks, arxiv.org/pdf/1406.2661v1, 10.06.2014.

[Has21] HASSANIN H., ZWEIRI Y., FINET L., ESSA K., QIU C., ATTALLAH M., “Laser Powder Bed Fusion of Ti-6Al-2Sn-4Zr-6Mo Alloy and Properties Prediction Using Deep Learning Approaches”, Materials (Basel, Switzerland), vol. 14 no. 8, 2021.

[Her05] HERBERT F., Dune. London: Hodder & Stoughton, 2005.

[Hol78] HOLDSTOCK R., ed., Encyclopedia of science fiction. London: Octopus Books, 1978.

[Hop82] HOPFIELD J. J., “Neural networks and physical systems with emergent collective computational abilities”, Proceedings of the National Academy of Sciences of the United States of America, vol. 79 no. 8, 1982, p. 2554-2558.

[Kan03] KANAL L. N., “Perceptron”, in:, Encyclopedia of Computer Science, 2003, p. 1383-1385.

[Kar21] KARNIADAKIS G. E., KEVREKIDIS I. G., LU L., PERDIKARIS P., WANG S., YANG L., “Physics-informed machine learning”, Nature Reviews Physics, vol. 3 no. 6, 2021, p. 422-440.

[Ken23] KENNEDY B., TYSON A., SAKS E., www.pewresearch.org/science/2023/02/15/public-awareness-of-artificial-intelligence-in-everyday-activities/, 2023.

[Kol98] KOLB B., Whishaw I. Q., “Brain plasticity and behavior”, Annual review of psychology, vol. 49, 1998, p. 43-64.

[Liu23] LIU W., LIU X., LIU Y., WANG J., EVANS S., YANG M., “Unpacking Additive Manufacturing Challenges and Opportunities in Moving towards Sustainability: An Exploratory Study”, Sustainability, vol. 15 no. 4, 2023, p. 3827.

[Lun17] LUNDBERG S. M., LEE S.-I., “A Unified Approach to Interpreting Model Predictions”, Curran Associates, Inc.

[McC04] MCCORDUCK P., Machines who think: A personal inquiry into the history and prospects of artificial intelligence. Natick, Mass.: A.K. Peters, 2004.

[Ope19] OpenAI, openai.com/research/openai-five-defeats-dota-2-world-champions, 2019.

[Pra22] PRATAP A., SARDANA N., UTOMO S., AYEELYAN J., KARTHIKEYAN P., HSIUNG P.-A., “A Synergic Approach of Deep Learning towards Digital Additive Manufacturing: A Review”, Algorithms, vol. 15 no. 12, 2022, p. 466.

[Ros57] ROSENBLATT F., The Perceptron: A Perceiving and Recognizing Automaton, 1957.

[Ros58] ROSENBLATT F., “The perceptron: a probabilistic model for information storage and organization in the brain”, Psychological review, vol. 65 no. 6, 1958, p. 386-408.

[Rum85] RUMELHART D. E., Learning Internal Representations by Error Propagation, 1985.

[Rum86] RUMELHART D. E., HINTON G. E., WILLIAMS R. J., ”Learning representations by back-propagating errors” Nature, vol. 323 no. 6088, 1986, p. 533-536.

[Sas23] SASIDHAR K. N., SIBONI N. H., MIANROODI J. R., ROHWERDER M., NEUGEBAUER J., RAABE D., “Enhancing corrosion-resistant alloy design through natural language processing and deep learning”, Science advances, vol. 9 no. 32, 2023, eadg7992.

[Sch19] SCHMIDT R. M., Recurrent Neural Networks (RNNs): A gentle Introduction and Overview, arxiv.org/pdf/1912.05911v1, 23.11.2019.

[Sch22] SCHULZ D., FAY A., MATIASKE W., SCHULZ M., eds., dtec.bw-Beiträge der Helmut-Schmidt-Universität / Universität der Bundeswehr Hamburg: Forschungsaktivitäten im Zentrum für Digitalisierungs- und Technologieforschung der Bundeswehr dtec.bw. Hamburg: Helmut-Schmidt-Universität Bibliothek, Dezember 2022.

[Sch97] SCHALLER R. R., ”Moore’s law: past, present and future”, IEEE Spectrum, vol. 34 no. 6, 1997, p. 52-59.

[Sco23] Scopus, Data on AI in AM, scopus.com, 2023.

[Sil19] SILBERNAGEL C., AREMU A., ASHCROFT I., "Using machine learning to aid in the parameter optimisation process for metal-based additive manufacturing", Rapid Prototyping Journal, vol. 26 no. 4, 2019, p. 625-637.

[Sto22] STONE P., BROOKS R., BRYNJOLFSSON E., CALO R., ETZIONI O., HAGER G., HIRSCHBERG J., KALYANAKRISHNAN S., KAMAR E., KRAUS S., LEYTON-BROWN K., PARKES D., PRESS W., SAXENIAN A., SHAH J., TAMBE M., TELLER A., Artificial Intelligence and Life in 2030: The One Hundred Year Study on Artificial Intelligence, arxiv.org/pdf/2211.06318v1, 31.10.2022.

[Sut98] SUTTON R. S., BARTO A., Reinforcement learning: An introduction. Cambridge, Massachusetts, London: The MIT Press, 1998.

[Tah23] TAHERKHANI K., ERO O., LIRAVI F., TOORANDAZ S., TOYSERKANI E., "On the application of in-situ monitoring systems and machine learning algorithms for developing quality assurance platforms in laser powder bed fusion: A review", Journal of Manufacturing Processes, vol. 99, 2023, p. 848-897.

[Tho22] THORMUNDSSON B., www.statista.com/statistics/1155960/customers-frequency-ai-enabled-interactions-organizations/, 2022.

[Tri20] TRIVEDI S., PATEL N., "The Role of Automation and Artificial Intelligence in Increasing the Sales Volume: Evidence from M, S, and, MM Regressions", SSRN Electronic Journal, 2020.

[Tur50] TURING A. M., "I. – Computing machinery and intelligence", Mind, LIX no. 236, 1950, p. 433-460.

[Vas17] VASWANI A., SHAZEER N., PARMAR N., USZKOREIT J., JONES L., GOMEZ A. N., KAISER L., POLOSUKHIN I., Attention Is All You Need, arxiv.org/pdf/1706.03762v7, 12.06.2017.

[Wac99] WACHOWSKI L., WACHOWSKI L., The Matrix, USA, 1999

[Wan22a] WANG R., GARCIA D., KAMATH R. R., DOU C., MA X., SHEN B., CHOO H., FEZZAA K., YU H. Z., KONG Z. J., "In situ melt pool measurements for laser powder bed fusion using multi sensing and correlation analysis", Scientific reports, vol. 12 no. 1, 2022, p. 13716.

[Wan22b] WANG T. T., GLEAVE A., TSENG T., PELRINE K., BELROSE N., MILLER J., DENNIS M. D., DUAN Y., POGREBNIAK V., LEVINE S., RUSSELL S., Adversarial Policies Beat Superhuman Go AIs, arxiv.org/pdf/2211.00241v4, 01.11.2022.

[Wat23] WATSON J. L., et al., "De novo design of protein structure and function with RFdiffusion", Nature, vol. 620 no. 7976, 2023, p. 1089-1100.

[Wei66] WEIZENBAUM J., "ELIZA – a computer program for the study of natural language communication between man and machine", Communications of the ACM, vol. 9 no. 1, 1966, p. 36-45.

[Yoo14] YOON H.-S., LEE J.-Y., KIM H.-S., KIM M.-S., KIM E.-S., SHIN Y.-J., CHU W.-S., AHN S.-H., "A comparison of energy consumption in bulk forming, subtractive, and additive processes: Review and case study", International Journal of Precision Engineering and Manufacturing-Green Technology, vol. 1 no. 3, 2014, p. 261-279.

[Zha21] ZHAO Y., ZHANG H., "Application of Machine Learning and Rule Scheduling in a Job-Shop Production Control System", International Journal of Simulation Modelling, vol. 20 no. 2, 2021, p. 410-421.

[Zho24] ZHOU H. A., ZHANG S., KEMMERLING M., LÜTTICKE D., SCHLEIFENBAUM J. H., SCHMITT R. H., "Fabrication Forecasting of LPBF Processes Through Image Inpainting with In-Situ Monitoring Data", in: Klahn C., Meboldt M., Ferchow J. (eds.), Industrializing Additive Manufacturing, Cham, Springer International Publishing, 2024, p. 147-158.

Imprint

Fraunhofer Research Institution for
Additive Manufacturing Technologies IAPT
Am Schleusengraben 14
21029 Hamburg-Bergedorf
Germany

Telephone +49 40 48 40 10-500
www.iapt.fraunhofer.de
marketing@iapt.fraunhofer.de

A legally dependent entity of
Fraunhofer-Gesellschaft zur Förderung
der angewandten Forschung e.V.
Hansastrasse 27c
80686 Munich
Germany
www.fraunhofer.de
info@zv.fraunhofer.de

Status of the Deep Dive, October 2023

Contact

Christian Böhm, M. Sc.
Head of Additive Alliance®
Telephone +49 40 48 40 10-636
eMail christian.boehm@iapt.fraunhofer.de

The Fraunhofer IAPT does not warrant that the information in this Deep Dive is correct.
No statement in this Deep Dive shall be considered as an assured prediction. The reader should not make any decision based exclusively on the information presented.

Photo credits:
all image rights are held by Fraunhofer IAPT.