

Heatmap-based Object Detection and Tracking with a Fully Convolutional Neural Network

Fabian Amherd
Elias Rodriguez

fabian.amherd@gmail.com
elias.rodriguez@bluewin.ch

Advisors:

Michael Bucher
Christopher Latkoczy

Kantonsschule Stadelhofen
Schanzengasse 17
8001 Zürich, Switzerland

External Advisors:

Prof. Dr. Tobias Delbrück
Yuhuang Hu
Shasha Guo

ETH Zurich / University of Zurich
Institute of Neuroinformatics (INI)
Winterthurerstrasse 190
8057 Zürich, Switzerland

Per Aspera Ad Astra

Contents

1	Introduction	1
1.1	What is Artificial Intelligence?	1
1.1.1	Definition of Artificial Intelligence (AI)	1
1.1.2	Different types of AI	1
1.1.3	A closer look at the field of AI	2
1.2	History	3
1.2.1	Earlier Epochs	3
1.2.2	AI today	4
2	Practical Part	5
2.1	Experiment	5
2.2	Materials and Methods	5
2.2.1	Data	5
2.2.2	Network Architecture	7
2.2.3	Loss Function	10
2.2.4	Training the Network	11
2.3	Analysis	15
3	Discussion	17
3.1	Applications of AI	17
3.2	Challenges regarding AI	19
3.2.1	Reliability	19
3.2.2	Interpretability	20
3.2.3	Moral Dilemma	20
3.2.4	Energy and Computing Power Demand	20
3.3	Outlook	21
4	Appendix	22
4.1	Reasons for the choice of topic	22
4.2	Acknowledgements	22

Abstract

The main topic of this paper is a brief overview of the field of Artificial Intelligence. The core of this paper is a practical implementation of an algorithm for object detection and tracking. The ability to detect and track fast-moving objects is crucial for various applications of Artificial Intelligence like autonomous driving, ball tracking in sports, robotics or object counting. As part of this paper the Fully Convolutional Neural Network "CueNet" was developed. It detects and tracks the cueball on a labyrinth game robustly and reliably. While CueNet V1 has a single input image, the approach with CueNet V2 was to take three consecutive 240 x 180-pixel images as an input and transform them into a probability heatmap for the cueball's location. The network was tested with a separate video that contained all sorts of distractions to test its robustness. When confronted with our testing data, CueNet V1 predicted the correct cueball location in 99.6% of all frames, while CueNet V2 had 99.8% accuracy. To get access to the source code of the project, scan the QR-code at the beginning of chapter 2.

1 Introduction

1.1 What is Artificial Intelligence?

1.1.1 Definition of Artificial Intelligence (AI)

Artificial Intelligence (AI) is the simulation of intelligence in a non-biological structure.

As in many fields, there does not exist one single agreed-upon definition of Artificial Intelligence. A definition often used is "Artificial intelligence (AI) would be the possession of intelligence, or the exercise of thought, by machines such as computers." [1]. So the only difference to Human Intelligence would be, that the intelligence is possessed by a machine on a semiconductor-basis, not by a human in a biological structure.

The question "Can machines think?" is very interesting to consider here. One can also have the opinion that a machine cannot possess real Intelligence, that it can

only simulate intelligent behaviour. Therefore, another definition that is around, reads as follows: "AI [can be looked at] as an engineering discipline in which researchers focus on developing useful programs and tools that perform in domains that normally require intelligence." [2]. So we can also define Artificial Intelligence as algorithms able to solve tasks, that normally require intelligence to be solved, although these algorithms do not possess "Intelligence" in a traditional sense.

1.1.2 Different types of AI

Some people may find it useful to put different algorithms or "AI systems" into categories, as a criterion for which can be used the comparison to a human. The following categorization is based on the capability of an AI. There is another popular categorization based on functionality, which won't be considered here though.

The least sophisticated type of AI is called **Artificial Narrow Intelligence (ANI)** or "weak AI". This is where we would place all of the AI systems we have encountered up to this point. As Forbes puts it: "Artificial narrow intelligence refers to AI systems that can only perform a specific task autonomously using human-like capabilities." [3]. Important to consider here is that these systems can be operating at a superhuman level of performance, but only in one specific area they were programmed for.

One step up the ladder we have **Artificial General Intelligence (AGI)** or "strong AI". These AI systems have human-level intelligence across all domains, they are as multi-functional as humans. They can understand, discuss and generalize like human beings. "AGI can think, understand, and act in a way that is indistinguishable from that of a human in any given situation." [4]

An **Artificial Superintelligence (ASI)** would surpass the capabilities of humans in every aspect. It would be able to come up with ideas that are impossible to grasp for humans. Whether it is in science, sports, understanding emotions, creative thinking, ASI will outperform humans. The potential of such a machine is tremendous and comes with unknown consequences.

It has been proposed that the emergence of ASI is almost inevitable when AGI has been achieved. To justify this, one only has to look at the advantages of a digital implementation over a biological one. In the digital implementation of a brain, the simulation can be copied

and accelerated multiple times. Thus, when we create AGI with a whole brain emulation, we can, with sufficient computational resources, create a collection of such AIs that work together at superhuman speeds. [5]

1.1.3 A closer look at the field of AI

Machine Learning (ML)

Machine Learning (ML) is the process that an Artificial Intelligence goes through when learning.

Another term for Artificial Intelligence, although less used, is "Machine Intelligence". "Machine Learning" is what leads to it.

Machine Learning is the process that an algorithm goes through when trying to search for patterns in a massive amount of data. The term refers to a subfield of AI (see figure 1), but still contains numerous other subfields (one of which is Deep Learning). The big advantage of Machine Learning is that the algorithm tries to find the patterns automatically, it learns through the experience it gets from the data that gets fed into it. Mathematically, the program tries to find the best possible function to map a certain input to a certain output. This is also why neural networks are sometimes called "universal function approximators".

There are different approaches to how a programmer can set up the environment for an algorithm to learn. One approach is called **Supervised Learning**, its principle is to present the algorithm with the input, let it output something, and then give feedback on how close the actual output was to the desired output, also called the "Loss", by calculating the difference between the actual output and the label. There are different ways to calculate this difference, which will be discussed in more detail in 2.2.3 of this paper.

A second approach often used is called **Unsupervised Learning**. The goal here is to let the algorithm automatically find the underlying structure of the data fed into it. This may help to discover connections that are not recognizable by humans alone because we can't take as much data into respect as computers can.

A third popular approach is **Reinforcement Learning**. Here the algorithm becomes an "agent" in a simulated or real environment in which it can take actions that lead to a different state of the environment. The algorithm

constantly tries to change its "policy" of taking actions to maximize the reward it gets. [6]

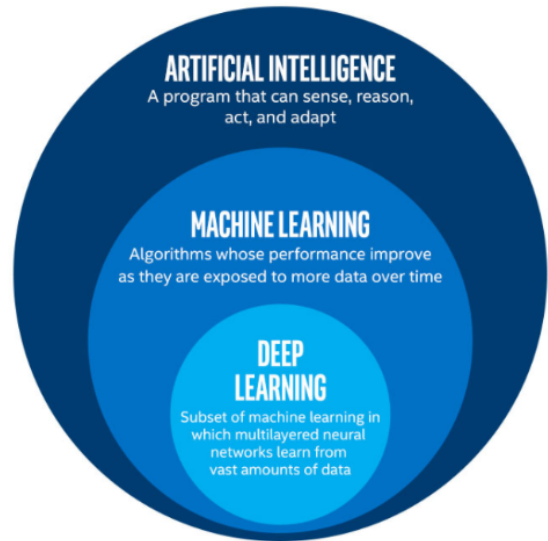


Figure 1: AI, Machine Learning and Deep Learning in a hierarchical order.

Deep Learning (DL)

If we perform Machine Learning (ML) with a Deep Neural Network (a form of an AI), it is called Deep Learning (DL).

One of the biggest subfields of Machine Learning is named "Deep Learning". This field addresses the development of **Artificial Neural Networks (ANNs)**, whose structure is strongly inspired by the human brain. ANNs will be discussed in more detail in the first part of 2.2.2. More complex tasks often require bigger networks with more subsequent layers. If an ANN consists of more than three layers, we call it a **Deep Neural Network**. [7]

There are dozens of different architectures of ANNs that excel at specific types of problems. The architectures we will have a closer look at in this paper are the most simple form, the Dense Neural Network (DNN), as well as the Convolutional Neural Network (CNN), which is being used in the practical part of this work.

Deep Learning is already being applied in all kinds of different industries all over the world, most widely known are recommendation systems as on Youtube or Amazon. Although already used so much, it is said to still have enormous potential to transform our lives completely. The

realised and potential applications of AI (and with it Deep Learning) will be discussed more closely in 3.1 of this paper.

1.2 History

The history of AI dates way back several thousand years where myths and stories of artificial beings with intelligence were being told. Amongst the old greeks there existed the myth of Talos, a guardian of the island of Crete, which was artificially forged by Hephaestus with the aid of a cyclops. His task was to throw boulders at ships of invaders and walk three times around the island's perimeter daily [8]. From 380 BC to 1900 there have been numerous mathematicians, theologians, philosophers, professors and authors who mused about calculating machines, mechanical techniques and other systems that eventually led to the concept of artificial human-like thought in non-human beings.

1.2.1 Earlier Epochs

The first real implementation of an AI was proposed by Warren McCulloch and Walter Pitts in 1943 with a mathematical model, combined with a computer model of the biological neuron. [9] At the beginning of 1950, Alan Turing and John von Neumann became the founding fathers of the technology behind AI, as they made the transition from analogue and hard-wired digital computers to computers with a stored program in a separate memory (the so-called "von Neumann" architecture). Turing and Neumann thus formalized the architecture of the computers from today. Turing raised the question of the possible intelligence of a machine in his famous article "Computing Machinery and Intelligence" from 1950. In this article, he described a "game of imitation" which later was called the "Turing test", where the machine can pass it if a human moderator can not tell if he is talking to a human or a machine. [8] The field was in its golden years, people were extremely optimistic, claiming that within ten years a digital computer will be the world's first chess champion and within ten more years a digital computer will discover and prove an important new mathematical theorem. [10]

The first AI winter 1974-1980

AI remained fascinating but the popularity of the technology fell back in the early 1970s. AI would become subject to critiques and financial setbacks. The

tremendous optimism had raised expectation impossibly high, which lead to the unfulfillment of the promises. [11] Combined with the limitations of the hardware at that time, on top of which came the cutting off from almost all funding for undirected research into AI, the technology was lead into its first winter. During the same period, Marvin Misky harshly criticised perceptrons and therefore shut down the field of neural nets almost completely for the next 10 years. [12]

Boom 1980-1987

In the 1980s, a form of AI called expert systems was adopted by corporations around the world. These expert systems are programs that can answer questions or solve problems about a specific domain of knowledge. An expert system tries to emulate the decision-making ability of a human expert, hence the name, by going through many "if-then" rules. The earliest examples were developed by Edward Feigenbaum and his students [13][14]. Expert systems only function in a specific field of knowledge and all in all were useful, which AI had not been able to achieve up to this point. [15][16]

During the same years in Japan, the government started to fund AI with its fifth-generation computer project. The Japanese Ministry set aside \$850 million for the project. The goal was to create programs and machines that could do things like translate language and reason like humans. [17] As a reaction, other countries started to fund new projects in the field of AI too. In the USA a few companies came together to form the Microelectronics and Computer Technology Corporation for funding large projects in the field of AI. The Alvey project started in the UK which cost £350 million. [18]

The second AI winter 1987-1993

In a similar pattern as the first winter, the second one occurred. Expectations were again set much too high and AI was exposed to financial setbacks. One could compare it to an economic bubble. Systems for vision and speech worked very poorly with expert systems because they contained too many edge cases, which couldn't be manually designed around. [19]

AI 1993-2011

Now some of the initial goals of AI could finally be achieved. May 11th 1997, Deep Blue, a chess-playing

computer won a game against the world chess champion Garry Kasparov [20]. AI's success started to spread throughout the technology industry. The overall success of AI was not due to some revolutionary discovery, but mostly due to the tremendous increase in the speed and capacity of computers in the 90s [21]. This dramatic increase is described by Moore's law, which states that the density of transistors on an Integrated Circuit (IC) of new computers doubles about every two years.

1.2.2 AI today

2011-present

Since the beginning of the 2010s, the computer science branch of AI has experienced a boom (again). The reason is mainly the availability of three key resources in AI: Computing power, the internet and data. The latter is available in such quantities that it has even gotten a name: "**Big Data**".

These circumstances also led to significantly increased budgets for AI R&D in the industry, the results of which we could observe in the last few years: In 2011, IBMs Watson won the TV-Show "Jeopardy", in 2016, AlphaGO beat the best human players, in 2020, Agent 57 beat the human benchmark in all 57 Atari games [22] and the list of new achievements goes on. [8]

"Who is Who" in the field of AI?

The point of this section is to give a quick overview of the most important players in the AI Industry regarding the big technology companies.

Nvidia is a company based in California, USA, that, among other things, develops Graphics Processing Units (GPUs). GPUs allow for massively parallel computation, and therefore, accelerated training of neural networks. As of Q4 in 2019, Nvidia has a GPU-shipment market share of ~70% [23]. [24] Nvidia is mostly fueled by the video game industry as heavy parallel computing power is required to run the latest video games.

Google is one of the major players in the AI industry. Google launched a separate department called "Google AI", that does R&D. Waymo, a company owned by Google, currently has one of the most advanced autonomous driving systems. In 2016, Google introduced the Tensor Processing Unit (TPU) (See chapter 2.2.1).

TPUs are integrated circuits specifically designed for machine learning with neural networks [25]. Of course, there is also the "Google Assistant", which is a virtual assistant that can take your language or keyboard input, and search the internet for answers (amongst other things). Google has even launched a machine learning platform called "Tensorflow".

At **IBM**, after breakthroughs like DeepBlue in 1996/1997 and Watson in 2011, "Project Debater" has been in development since 2012 and has already debated against the debating world-champion Harish Natarajan in 2019, though it has lost that discussion. IBM also has an R&D department called "IBM Research" that, amongst other things, does research at the cutting edge of AI. [26]

Just like Google and Amazon, **Microsoft** has its own Cloud Computing platform. It is called "Microsoft Azure" and is partly devoted to run heavy AI workloads. Microsoft is building AI into their existing products such as Office 365 to make smart recommendations. One of Microsoft's initiatives is called "AI for Good", with which it tries to have a positive impact on the present and future development of AI.

At **Amazon**, AI is not located in a single department or office. Instead, AI is distributed and applied throughout the whole company. The most apparent application is the recommendation engine on the company's online shop. Similar to Google, Amazon developed a virtual assistant called "Alexa". One major advantage for Amazon is its access to a huge data source from tracking the purchase behaviour of their clients. [27]

Like Google, **Facebook** has developed a Machine Learning library called "Pytorch", which is popular amongst researchers. Facebook also has a research department called "Facebook AI Research (FAIR)", that does research as well as publish useful tools like "Detectron 2", a platform that includes several implementations or different computer vision algorithm, including object detectors like the one of our study. [28]

2 Practical Part

This paper includes a practical project with a real-life application of an AI. This following chapter is devoted to the documentation of the experiment, as well as some theoretical explanations on the side to better illustrate the mathematics behind neural networks. The code is available under the following GitHub repository: https://github.com/FabianAmherd/CNN_ER_FA (see QR-code below).



2.1 Experiment

The challenge of the labyrinth game in figure 2 is to roll the ball to a destination without it falling into one of the holes on the board. The ball is moved by tilting the labyrinth with two knobs (one on each axis). Our given task was to develop a neural network to detect the ball on this labyrinth game for tracking it. Specifically, we developed a fully convolutional neural network. With that, the first step for a fully automated labyrinth game was developed. The data was collected and labelled by ourselves with a camera and suitable labelling software.




Figure 2: Setup: Brio labyrinth with RC servo motors, DAVIS camera and the labelling software.

2.2 Materials and Methods

The following section contains explanations and documentation to how the task described in 2.1 was solved. Also, it will cover the theoretical background of some popular Deep Learning methods or tools.

Equipment

- **Camera:** The DAVIS camera used is a dual-frame and event camera [29]. It outputs conventional frames and a stream of dynamic vision sensor (DVS) brightness change events [30]. The events can be collected to variable duration frames to drive a conventional CNN with sparse frame input, allowing tracking of fast-moving objects.
- **Camera/Labelling Software:** To record and label our data, the project has used jaAER software [31] as well as the program "After Effects".
- **Labyrinth:** The labyrinth used was a "Brio Labyrinth Game" [32]. It is fitted with a camera stand as well as two RC servo motors to control the table tilt. These servo motors can be computer-controlled from a computer using the jaAER software.
- **Machine Learning tools:** The project has used the open-source machine learning library **Pytorch** for developing, training and testing the CNN.
- **GPUs:** The project has used a CUDA-enabled NVIDIA GTX 1080 GPU to accelerate the training of the CNN.

2.2.1 Data

Machine Learning and hence Deep Learning are both strongly connected to the field of data science. That's probably why you don't hear too little about AI being all about (Big) Data.

Tensors

All the data that is being processed, passed in and out of a neural network is represented using an n -dimensional array of numbers, called a tensor.

In Machine Learning, a tensor is used as a way to represent

data that researchers and programmers use when dealing with neural networks. A tensor is the generalization of vectors and matrices and can be n-dimensional. A vector is just a one-dimensional or rank one tensor and a matrix is a two-dimensional or rank two tensor.

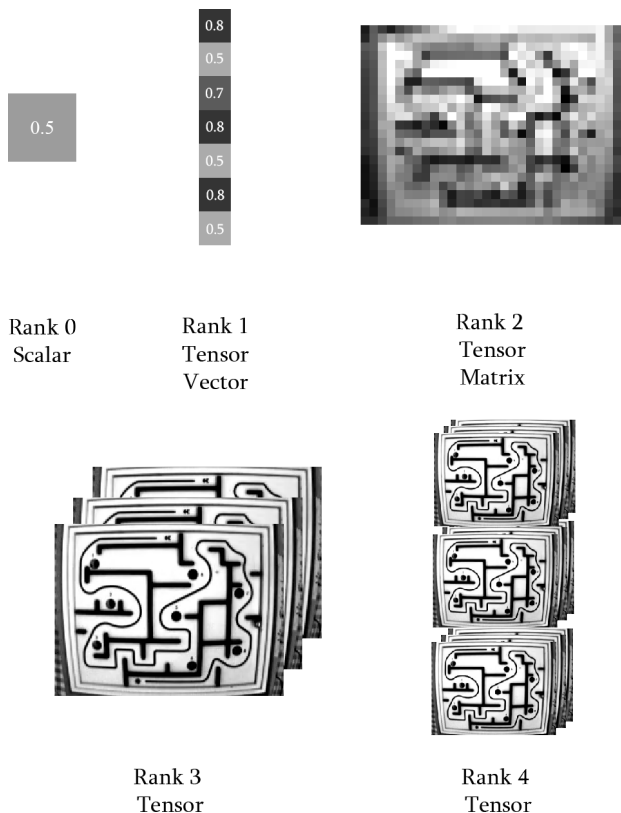


Figure 3: Graphical representation of a tensor with different ranks and dimensions.

A tensor can be represented by an n-dimensional array of numbers. Its *rank* is determined by the number of basis vectors needed to find a specific component. So if we had a two-dimensional array, we needed the row and the column of the component, so that tensor would have rank 2. The *dimension* of a tensor is in the case of our array-representation simply determined by the number of rows or columns the array has.




Figure 4: Graphical representation of a tensor with dimension three and rank two (Also called the "Cauchy Stress Tensor").

The same observation (i.e. a ball rolling in a certain direction) can be represented by different coordinate systems. Different coordinate systems rely on different *basis vectors*, so the *components* in the mathematical objects used (i.e. a vector with a magnitude and a direction to represent the velocity of the ball) also have to be different to compensate for the change of the basis vectors to still represent the same observation [33].

The power of Tensors lies in the fact that they ensure that the observation they represent (in the form of a combination of a set of basis vectors and components) stays the same for every observer using any coordinate system [34]. This means that if you rotate your coordinate system, the ball keeps rolling "in the same direction" seen from an absolute standpoint because our way to represent an observation does not affect the actual physics going on.

To conclude for our work, Tensors have served as a way to represent the data that is passed through our neural network. That is when we first input the images from the camera, as well when the images are manipulated by the neural network when processing the data. It was always a Tensor, as you can see in figure 3.

Dataset Creation

When developing an AI algorithm, a substantial part of the work goes into preprocessing the data, so the neural network can learn properly. To create the training set for the CNN, we needed to go through a couple of steps.

After recording videos of the ball rolling around on the labyrinth (table tilted by hand), we had to label the video. We could do that by importing and activating the "DvsSliceAviWriter" filter in the jAER software. Once it is activated, there is a submenu called "TargetLabeler", which lets you use your mouse to set the labels while watching the recorded video. Having done that, we used the jAER software to output a .avi video file with the recorded video, plus a .txt text file that contained the corresponding label locations/coordinates for every frame of the video that we had applied a label to using the mouse on the beforehand. To accelerate this process we made use of the program "After Effects", which already contains a tracker for different kinds of objects. With that, we could label our data faster and more accurately. In total, our dataset reached a size of about 30'000 frames.

Since the video output of the DAVIS camera had the classical frames (**APS frames**) on the red channel and the event frames (**DVS frames**) on the green channel (blue is empty), the next step was to write a python script that could both extract all the frames from the .avi-formatted video plus separate the APS and DVS frames and save them into separate folders as .png images. With that, we have prepared the input to our CNN properly.

Since our project uses a "Supervised Learning" approach (see 1.1.3), we have rate the output of the neural network automatically for the CNN to be able to continually improve its performance. This is where the labels come into play. The task now was to use the location information of the previously described .txt file to create so-called **heatmaps** which represent the ideal output we want from our neural network for each picture. By taking the coordinates of the .txt files and putting a Gaussian distribution around them, we could create the desired heatmaps. The purpose of these heatmaps is to help us tell the CNN how good or bad it is currently doing at telling us the location of the ball on the labyrinth, by providing the perfect solution. How we solved the task of rating the CNN as well as an example of a label is covered in 2.2.3 .

Once we had this collection of examples and their corresponding solutions prepared, we could start developing the architecture of the neural network itself.

Data Augmentation

Data Augmentation is a technique used in Deep Learning to artificially increase the amount of data in the dataset, as well as to increase the performance and robustness of the neural network.

The technique's trick is to artificially increase the number of learning examples in the dataset by slightly modifying existing samples. With images you could, for example, rotate it, adjust the brightness, change the virtual lighting by shading the image, add fake highlight reflections, translate it, crop it, etc. as you can see below in figure 5. The three square images below have all been augmented. They can now all be used as a separate training example apart from the original image, thus multiplying the size of the dataset! You could of course implement more different forms of image augmentation, but transformations other than rotating, as well as adjusting the brightness and contrast were sufficient for good results of our study.




Figure 5: Examples of image augmentation used in our project.

Our project used the technique of Data Augmentation with the intent to make the neural network more robust against different lighting conditions, slight rotations of the labyrinth etc. but also improve the CNNs accuracy by letting it train on a larger dataset. It also helped to increase the performance of the CNN under real-life conditions, because they have been simulated and trained on before.

2.2.2 Network Architecture

How does an ANN work?

An Artificial Neural Network (ANN) is a form of AI, whose structure is strongly inspired by the human brain.

In this chapter, we will have a closer look at how a neural network works. After that, we present the architecture of the CNN we have developed for the practical project.

The first type of a neural network we will look at is the most simple form. It is called a **”Dense Neural Network (DNN)”** or **”Fully-connected Neural Network”**. It consists of thousands to millions of nodes, that are organized into different layers. Each node is connected to all nodes in the previous layer as well as in the following layer. So there are even many more connections than there are nodes. This is why the individual layers are called **”Dense Layers”** or **”Fully Connected Layers”**, the latter of which is often abbreviated with **”FC”** [35].

Figure 6 shows a simplified DNN with fewer nodes and connections. The first layer of a neural network, in general, is called the **”input layer”** because this is where we feed in our data. The following layers are called **”hidden layers”** because we rarely access them directly and in most cases also don’t exactly know what these layers represent logically. The last layer is the **”output layer”**, this is what we receive back as a result of our neural network.




Figure 6: Dense Neural Network.

Let’s have a look at the mathematics behind it.

A single neuron has the following structure (see figure 7): It takes a set of inputs $\{x_1, \dots, x_i, \dots, x_n\}$, where n is the number of neurons it is connected to in the previous layer. In figure 7, $n = 2$. Each connection between two neurons has a **weight** w_i associated with it. So before the inputs enter the neuron, they are multiplied by the weight that’s associated with them.

$$x_i \rightarrow x_i * w_i$$

In the next step, we sum all of the weighted inputs and add a **bias** b to it. The bias is a real number that helps to control how hard it is for a neuron to get excited (i.e. to output a large number to the next layer).

$$\sum_{i=1}^n (x_i * w_i) + b$$




Figure 7: The structure of a neuron.




Figure 8: A Graph of the sigmoid activation function.

As the last step, an **activation function** is used to bring our weighted sum into a predictable form. One function that is often used is called the **”sigmoid function”**, denoted with the symbol σ . It compresses every input in the range between 0 and 1.

$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

$$\lim_{z \rightarrow +\infty} \sigma(z) = 1$$

$$\lim_{z \rightarrow -\infty} \sigma(z) = 0$$

So finally, we can write the output y of our neuron as:

$$y = \frac{1}{1 + \exp(-\sum_{i=1}^n (x_i * w_i) - b)}$$

For our network, the ReLU (Rectified Linear Unit) activation function was used, as this is a commonly used activation function for CNNs which was first inspired by observations of real cortical neurons in the laboratory of the Institute of Neuroinformatics at UZH/ETH Zurich [36]. The ReLU activation function is easy to compute and helps against the "vanishing gradient" problem, in which the gradient of the loss becomes extremely small and thus the learning progress starts to stall (more on that in chapter 2.2.4). The ReLU activation function sets every negative input to zero and leaves every positive input as it is. It is mathematically expressed as:

$$ReLU(x) = \max(0, x)$$




Figure 9: Graph of ReLU Activation function.

We now know how data gets fed through forward to produce our result. But how does a neural network learn? After outputting the result, the neural network gets feedback from a "Loss Function" which we will look at in chapter 2.2.3. Once it knows how good or bad the produced result was, it uses two techniques called "Backpropagation" and "Gradient Descent" (See chapter 2.2.4) to do better next time.

Convolutional Neural Networks

A relatively newer form of a neural network, whose structure is inspired by the visual cortex of the human brain, is called a Convolutional Neural Network (CNN) or "ConvNet". As a consequence of their structural origin, the individual "neurons" in this type of neural net correspond to certain areas of the input image, known as the "Receptive Field" when talking about the visual system of humans. One of many differences to a DNN is that a CNN can capture the spatial information contained in an image fed into it. This fact makes CNNs excel at visual tasks such as object recognition/localization, which is why a CNN has also been used in the practical part of this paper.

A typical CNN consists of two parts: First, we have a few "Convolutional Layers", whose job it is to extract the features of the input image and bring them into a form that is easier to process without losing the critical features. This information is then passed through a few "Fully Connected Layers", whose job it is to classify the received information into a category (See figure 10) [37].

A single "Convolutional Layer (ConvLayer)" consists of a set of **filters or kernels**, matrices that contain the "learnable" weights (The weights are what gets adjusted when the feature extractor learns). These filters are much smaller than the input image (See square on the image on the left side in figure 10), so they propagate over the whole image, step by step, starting at the top left, ending at the bottom right of the image. After each step, the weights are multiplied with the underlying pixel-activations in that part of the image (See figure 11). The results of each step are then put together to form a new image. This output image is then used as input to the next ConvLayer. Depending on the filter weights, different aspects of the image are enhanced. In figure 12, the patterns the individual filters of a CNN look for were visualized. In the left rectangle of figure 12, we can see that edges at different angles and certain colors are being enhanced. Moving to the right, we can see the learned filters from following ConvLayers, which are a little more complex patterns, formed out of the simple edges and colours to the left. This goes on for the rest of the layers. The patterns the filters look for will change depending on what type of images the CNN was trained on.

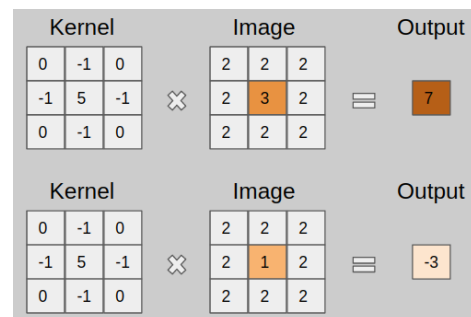


Figure 11: Example of a kernel.

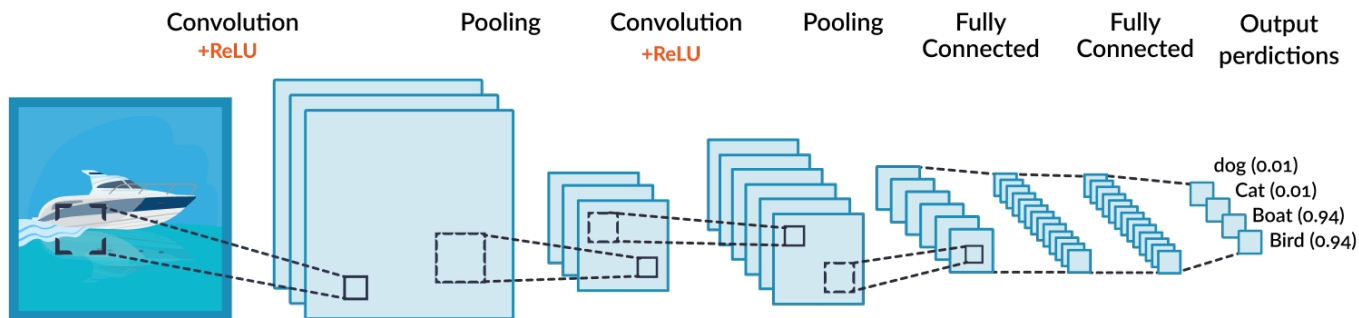


Figure 10: Example of a CNN.




Figure 12: Different filter patterns.

In figure 10, besides the ConvLayers and the Fully Connected Layers, we can also see so-called **"Pooling Layers"**. Two types of pooling operations often used are called "average-pooling" and "max-pooling". Average-pooling outputs the average pixel-activation of the current set of pixels it is applied to. The max-pooling operation outputs the maximum pixel-activation out of the patch it is applied to. These layers help to reduce the size of the input, as well as the amount of redundant information.

2.2.3 Loss Function

When training a neural network, you generally need to tell it in some way how it is doing in terms of delivering the output that you want it to. Here we can use a "Loss Function" (or also "Cost Function"). It is a function that takes two inputs: one is the output of the ANN, the other is a label that shows the desired output. It then in some way computes a "difference" between the two and outputs a number, that gets fed back into the ANN and is the origin for the learning process. The higher the number, the more different the output is from the desired output. This means that the lower this number (called "Loss" or "Cost"), the better the ANN is performing.

The heatmaps we are using to represent our labels and network output are, again, just a 2D-matrix of numbers, in which every number stands for the activation of a pixel.

In figure 14, there is a typical input image (that we got by extracting the red channel of the DAVIS video feed and converting it to black and white frames). In figure 15, you can see a typical label picked out of the training dataset. These two figures are a input/label pair, which means that at the same coordinates of the white blob on the label, the cueball is at on the input image.



Figure 14: Example of a frame used for training.




Figure 15: Example of a label used for training.

Now, at the beginning of training, the CNN has no clue what it should do and all the weights are initialised randomly. The output heatmap looks something like what you can see in figure 16. Remember, the desired output

would be something like in figure 15. Also note that the activations are very small, as you can see on the scale to the right side of the heatmap. The reason for that is the architecture of our CNN, specifically the last layer, as you can see in the last part of figure 21 . The whole heatmap can be looked at as a sort of "probability distribution" on where the CNN thinks the ball is.

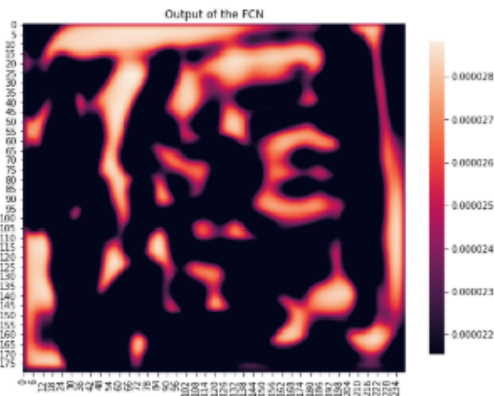


Figure 16: Output of our CNN at an early training stage.

Now we need a way to tell the network that this is not what we want. We do it by feeding back a big number to the neural network that stands for a big difference to our desired output. When dealing with images, the most simple and obvious way to calculate the difference between two pictures is to iterate over every pixel-position of the images and calculate the difference between the two pixel-values that are at the same location. You can then add up all those numbers or take the mean and you have the difference between those two pictures expressed in the form of a number.

The way just described is often used (also in this project) and it is called "L1-Loss". It can be expressed as the sum of absolute differences or mathematically as

$$Loss = \sum_{i=0}^n |a_L - a_O|$$

where the activation of a pixel in the label image is denoted by a_L , a_O for the output image, respectively.

Modified versions of L1-Loss were put to the test during the development of CueNet (the CNN developed in this project), hoping it would learn faster, but none of the approaches was superior to the classical L1-Loss.

2.2.4 Training the Network

Gradient Descent

After preprocessing the data, choosing the architecture of the network and declaring the loss function it is time to train it. In the beginning, all the weights and biases are randomly initialised. The network thus creates random outputs. With the loss function, you can measure how well the network is performing on a given input which alone doesn't get us anywhere. We have to find a way to tell the network how it should change its weights and biases so it can increase its performance and thus "learn".

To illustrate how this can be done, we will look at a simple Loss Function with just two inputs and one output expressed as:

$$Loss = L(x, y)$$

If we plot the input on the x- and y-axis and the output on the z-axis we could get the following graph:




Figure 17: 3D Plot of an example Loss Function

Our goal is to find the direction in this input space which decreases the loss most quickly, because the lower the loss, the better our network performs. The gradient of a function gives you a gradient vector whose direction points in the direction of steepest ascent. If we take the negative of the gradient vector, we get the direction of quickest descent. The gradient is calculated by packing all the partial derivative information of our loss function together.

$$\nabla L(x, y) = \begin{bmatrix} \frac{\partial f}{\partial x} \\ \frac{\partial f}{\partial y} \end{bmatrix}$$

To find a local minimum we have to compute the gradient over and over while taking small steps in the inversed

direction of the gradient vector. This process is called **Gradient Descent** and also applies to neural networks, many of which create input spaces with millions of dimensions. There are many ways to improve this process, such as "Stochastic Gradient Descent (SGD)", which we won't discuss here in detail.




Figure 18: Gradient Descent

For us to take small steps on the plane the loss function needs to have a continuous output. This is the reason why artificial neurons have continuous ranging activations, unlike biological neurons which activations behave in a binary way.

Now that we know how to descend the valley of our loss function if we know the perfect direction, we need to figure out how to find that "direction" in the multi-million-dimensional input space of our loss function each step. This is where the Backpropagation algorithm comes into play.

Backpropagation

Calculating the gradient of a neural network is a lot harder as there are many more parameters we can change and their impact on the final loss function is dependent on changes to other parameters. The core of Backpropagation is an expression for the partial derivative $\partial L / \partial w$ of the loss function L with respect to any weight w or bias b in the network. This expression tells us how fast the loss changes when we shift the weights and biases. Going further we will rely on understanding the following notation which lets us refer to weights and biases in the network in an easy way. w_{jk}^l will be used to denote the weight for the connection from the k^{th} neuron in the $(l - 1)^{th}$ layer to the j^{th} neuron in the l^{th} layer. As an example, we will look at the indicated weight w in figure 19. The weight w is the connection between the

first neuron in the second layer to the second neuron in the third layer of the network. Therefore the weight w is notated as w_{21}^3 .




Figure 19: Example network to illustrate notation.

To refer to the network's biases and activations we use a similar notation. For the bias of the j^{th} neuron in the l^{th} layer, we use b_j^l . And for the activation of the j^{th} neuron in the l^{th} layer, we use a_j^l . With that the indicated bias b and activation a in the figure below are notated as b_2^1 and a_2^3 .




Figure 20: Example network to illustrate notation.

With these notations, the activation a_j^l is related to the activations in the $(l - 1)^{th}$ layer by the following equation:

$$a_j^l = \sigma \left(\sum_k w_{jk}^l a_k^{l-1} + b_j^l \right)$$

where the sum is above all neurons k in the $(l - 1)^{th}$ layer. The sigmoid activation function is just used as an example and can be exchanged with any activation function.

We will now rewrite this expression in a matrix form to make our expression more readable and intuitively understandable. We define a weight matrix w^l for each layer l where the entries of the weight matrix are the weights connecting to the l^{th} layer of neurons, that is, the entry in the j^{th} row k^{th} column is w_{jk}^l . Analogously

the bias vector is defined for each layer l by b^l and the activation vector by a^l . The components of the bias and activation vectors are the values b_j^l or a_j^l for each neuron in the l^{th} layer. Finally, we must vectorize our activation function σ . We do this by applying the function to every element in a vector v and use the notation $\sigma(v)$ to denote this elementwise application of a function. We can now rewrite our last equation in a compact vectorized form

$$a^l = \sigma(w^l a^{l-1} + b^l).$$

Below is the written out activation vector of the 3^{rd} layer of the example network in figure 20.

$$\begin{bmatrix} a_1^3 \\ a_2^3 \end{bmatrix} = \sigma \left(\begin{bmatrix} w_{11}^3 & w_{12}^3 & w_{13}^3 \\ w_{21}^3 & w_{22}^3 & w_{23}^3 \end{bmatrix} \times \begin{bmatrix} a_1^2 \\ a_2^2 \\ a_3^2 \end{bmatrix} + \begin{bmatrix} b_1^3 \\ b_2^3 \end{bmatrix} \right)$$

When we compute a^l with our equation above we compute the intermediate quantity $z^l = w^l a^{l-1} + b^l$, which is the weighted input to the neurons in layer l , this z^l will be useful later. With this, our equation above can be rewritten as $a^l = \sigma(z^l)$.

To use Backpropagation we must make two main assumptions about the shape of our loss function. To state these assumptions we will look at the quadratic loss function as an example, which has the form

$$L = \frac{1}{2n} \sum_x \|y(x) - a^L(x)\|^2.$$

where n is the total amount of training samples. The sum is over all the samples in the training set, $y(x)$ is the corresponding desired output. L is the number of layers in the network and $a^L(x)$ is the vector of the output activations from the network for the given input x .

The first assumption we have to make about our loss function L is that the loss function can be written as an average $L = \frac{1}{n} \sum_x L_x$ over loss functions L_x for the individual training examples x . For our loss function, this is the case, as the loss for a single training example equals $L_x = \frac{1}{2} \|y - a^L\|^2$.

The reason for this is that we try to approximate the partial derivatives $\delta L/\delta w$ and $\delta L/\delta b$ by calculating the partial derivatives $\delta L_x/\delta w$ and $\delta L_x/\delta b$ for each training example x and average it. With that in mind, we will remove the subscript x and write the loss L_x as L .

The second assumption about our loss function is that it can be rewritten as a function that takes the output of our network as the input. This applies to our loss function L as it can be written for a single training example x as

$$L = \frac{1}{2} \|y - a^L\|^2 = \frac{1}{2} \sum_j (y_j - a_j^L)^2$$

and by that is a function of the output activations of our network. One may wonder why we do not take the loss L as a function of the desired output y or the training example x . This is because x and y are fixed parameters.

We will now introduce the intermediate quantity δ_j^l , which we will call the *error* in the j^{th} neuron in the l^{th} layer to ultimately compute $\delta L/\delta w_{jk}^l$ and $\delta L/\delta b_{jk}^l$. Backpropagation will enable us to compute this error δ_j^l and then will relate δ_j^l to $\delta L/\delta w_{jk}^l$ and $\delta L/\delta b_{jk}^l$. To elucidate this, imagine an error which lies in the j^{th} neuron in the layer l . When the input enters the neuron, the error adds a little change Δz_j^l to the neuron's weighted input, which changes the neuron's output $\sigma(z_j^l)$ to $\sigma(z_j^l + \Delta z_j^l)$. We must now find a Δz_j^l that minimizes the loss. If $\frac{\delta L}{\delta z_j^l}$ is large then we can lower the loss a lot by choosing Δz_j^l to have the opposite sign to $\frac{\delta L}{\delta z_j^l}$. Reciprocally if $\frac{\delta L}{\delta z_j^l}$ is small making changes to Δz_j^l won't alter the loss by much. As a consequence, we can already tell which neuron is near the desired state. We define this error δ_j^l in the j^{th} neuron and l^{th} layer by

$$\delta_j^l = \frac{\delta L}{\delta z_j^l}.$$

As before we will use δ^l as the vector of errors in layer l and equals

$$\delta_j^L = \frac{\delta L}{\delta a_j^L} \sigma'(z_j^L). \quad (1)$$

for the output layer.

This intuitively makes sense as the first term $\delta L/\delta a_j^L$ shows us how fast the loss is changing as a function of the j^{th} output activation. If L doesn't depend much on the given output by the j^{th} neuron then δ_j^L will be small. The second term $\sigma'(z_j^L)$ shows us how fast the activation function σ changes at z_j^L .

Equation 1 is easily computed as z_j^L will already be computed when passing the training data through the

network and it's only a small addition to compute $\sigma'(z_j^L)$. The form of $\delta L/\delta a_j^L$ will be dependent on the form of our loss function. For example, when using the already known quadratic loss function $L = \frac{1}{2} \sum_j (y_j - a_j^L)^2$ it would be $\delta L/\delta a_j^L = (a_j^L - y_j)$ which isn't hard to compute.

While equation 1 is a componentwise expression for δ^L we want a matrix-based form which we get by rewriting the equation as

$$\delta^L = \nabla_a L \odot \sigma'(z^L). \quad (2)$$

$\nabla_a L$ is a vector whose components are the partial derivatives $\delta L/\delta a_j^L$ and thus the rate of change of L with respect to the output neurons. We will use these equations interchangeably. Our example quadratic loss function equals $\nabla_a L = (a^L - y)$ and the fully matrix-based form can be written as

$$\delta^L = (a^L - y) \odot \sigma'(z^L).$$

The next equation we will look at will be

$$\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l). \quad (3)$$

With this equation, we can calculate the error δ^l in relation to the error of the next layer δ^{l+1} . By applying the transpose weight matrix $(w^{l+1})^T$ we are moving the error back through the network. By taking the Hadamard product $\odot \sigma'(z^l)$ we move the error backwards through the activation function, leaving behind the error δ^l in the weighted input in layer l . We are propagating backwards through the network which is the reason it is called Backpropagation.

Now with the combination of equation 2 and 3, we can compute the error δ^l for any layer in the network. We start by calculating δ^L with the 2nd equation, then use the 3rd equation to compute δ^{l-1} over and over for all remaining layers of the network.

The next equation we will introduce lets us calculate the rate of change of the loss with respect to any bias in the network as we know how to calculate δ_j^l .

$$\frac{\partial L}{\partial b_j^l} = \delta_j^l. \quad (4)$$

Equation 5 will enable us to calculate the rate of change of the loss with respect to any weight in the network.

$$\frac{\partial L}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l. \quad (5)$$

As we know how we can calculate the terms of the quantities δ^l and a^{l-1} we can now compute equation 5.

The Backpropagation algorithm will go through five steps, the first being setting the corresponding activation a^1 for the input layer. The next step is to feedforward the input activation. For each $l - L$ we compute $z^l = w^l a^{l-1} + b^l$ and $a^l = \sigma(z^l)$. The third step is to compute the error vector for the last layer L with $\delta^L = \nabla_a L \odot \sigma'(z^L)$. The fourth step is to backpropagate the error back through the network by computing $\delta^l = ((w^{l+1})^T \delta^{l+1}) \odot \sigma'(z^l)$ for all the layers. The last step is to calculate the gradient of the loss function with $\frac{\partial L}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l$, $\frac{\partial L}{\partial b_j^l} = \delta_j^l$ and change the weights and biases accordingly.

As we have already discussed we calculate the gradient for one single training example a time. This is because calculating the gradient for all the examples would take too much memory and therefore increase the computational cost. We combine Backpropagation with a learning algorithm such as Stochastic Gradient Descent, in which we compute the gradient for a chosen batch-size m and take a learning step based on the gradient of this batch, by doing this we decrease the accuracy but can learn a lot faster. [38][39]

Training phases

In our code, we define two sections: the training loop and the validation loop. In the first step, we go through the training loop where we calculate the partial derivatives for each training example and update the weights and biases according to the calculated gradient which gets multiplied by a *learning rate* α which lets us adjust the steps the network takes so that we don't jump over local minima. After we went through all the training data we jump to the validation loop in which we calculate the loss for each example in the validation set and average it. This time we won't change the weights and biases as we want to check if the network can generalize what it has learned from the training set. These two phases will be repeated multiple times in so-called *epochs*, which is a pass over all the training data. The loop won't stop until the validation loss does not decrease anymore. At this point, the network starts to overfit, which means that it has started to simply memorize the data in the training

set and doesn't generalize well. In figure 23 you can see how both the training loss and validation loss behaved when going through 20 epochs in our experiment.

2.3 Analysis

CueNet consists of an abridged version of the VGG16 convolutional neural network with a subsequent deconvolutional neural network. The 240 x 180-pixel input with a depth of 3 is passed through the VGG16 part, which reaches a depth of 256 and a resolution of 60 x 45 in the middle of the CueNet. In the DeconvNet part, upscaling layers are used to recover the dimensionality loss from the max-pooling layers. Due to the Softmax layer as the last layer of the network, the output produces a probability heatmap for the ball's location where each pixel carries a probability for the ball being there. All these probabilities sum up to 1, as the ball has to be somewhere on the board. The coordinates of the cueball equal the coordinates of the activation-peak on the heatmap. Our input has a depth of 3 channels because we want to give 3 consecutive images to CueNet, hoping that it can learn something about the behaviour of a ball, e.g. a ball cannot jump from one side of the labyrinth to the other in consecutive images. For comparison purposes, we have created an identical version of the CueNet, which can only use one image as an input. [40]

Our testing set is used to evaluate the performance of the CueNet, distinguishing between CueNet V1, which takes one image as an input and CueNet V2, which takes three consecutive images as an input. The rest of the dataset is separated into a training set and a validation set, depending on the true position of the ball (see figure 22).

Layer	Filter Size	Depth	Padding	Stride	Activation Fn
Conv1	3 x 3	64	1 x 1	1	ReLU + BN
Conv2	3 x 3	64	1 x 1	1	ReLU + BN
Pool1	2 x 2 max pooling, Stride = 2, Padding = 0				
Conv3	3 x 3	128	1 x 1	1	ReLU + BN
Conv4	3 x 3	128	1 x 1	1	ReLU + BN
Pool2	2 x 2 max pooling, Stride = 2, Padding = 0				
Conv5	3 x 3	256	1 x 1	1	ReLU + BN
Conv6	3 x 3	256	1 x 1	1	ReLU + BN
Conv7	3 x 3	256	1 x 1	1	ReLU + BN
UpS1	2 x 2 upsampling				
Conv8	3 x 3	128	1 x 1	1	ReLU + BN
Conv9	3 x 3	128	1 x 1	1	ReLU + BN
Conv10	3 x 3	128	1 x 1	1	ReLU + BN
UpS2	2 x 2 upsampling				
Conv11	3 x 3	64	1 x 1	1	ReLU + BN
Conv12	3 x 3	64	1 x 1	1	ReLU + BN
Conv13	3 x 3	64	1 x 1	1	ReLU + BN
Softmax					

Table 1: Parameters of CueNet.




Figure 22: ■ Training Set, ■ Validation Set.

This is to ensure that CueNet generalizes and does not overfit. By separating the training data and validation data by the true location of the ball we make sure that there are not any nearly identical images in the training and validation set. If you were to split your sets by randomly shuffling your data or by time, there could be frames in the training and validation that are almost identical, which leads to the network not generalizing but rather memorizing the frames. To optimize the weight and biases, the Adam optimizer was used. We set the number of epochs to 20 to ensure that we find the least validation loss. The entire training using a batch size of 1 took about 8 hours.




Figure 21: CueNet V2: ■ Conv Layer, ■ ReLU, ■ MaxPooling Layer, ■ Upsampling Layer, ■ Softmax Layer




Figure 23: Training Loss vs Validation Loss.

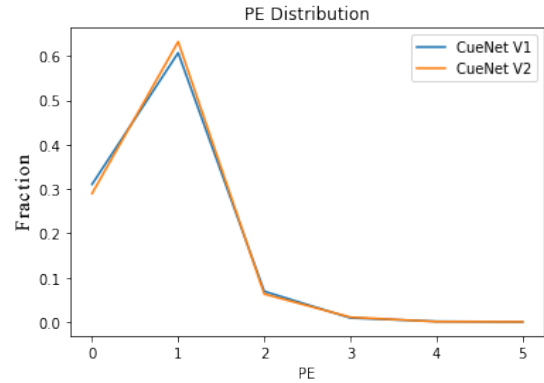


Figure 24: PE distribution of the two versions of CueNet.

The x-axis of the plot shows us the PE and the y-axis the fraction of the occurrences for this PE. This distribution shows that 99.6% of the classifications of CueNet V1 are in our set tolerance range of $PE \leq 4$ and 99.8% of the classifications of CueNet V2 respectively.

CueNet is trained to produce the "correct" heatmap output. But to allow us to quantify the accuracy of CueNet, we define a tolerance range in which the output can lie and still be classified as correct. Since the ball is about 8 pixels wide, the tolerance range is 4 pixels around the center of the real position of the ball. The peak of the heatmap from the output of CueNet has to be in this area to be counted as a correct classification. The Positioning Error (PE) is the distance between both centers in pixels. The distribution of the PE is shown in figure 25.

Apart from the standard testing set, we have created a special dataset called "Heavy Test" to push CueNet to its limits. The recordings of "Heavy Test" were made in very poor lighting conditions, as a lamp was rotated around the labyrinth at a steep angle. This resulted in very strong and long shadows on the labyrinth, which should lead to a more difficult classification task. In addition to this, the view of the board was alternately covered with hands without covering the ball, which should show that CueNet can ignore distractions.

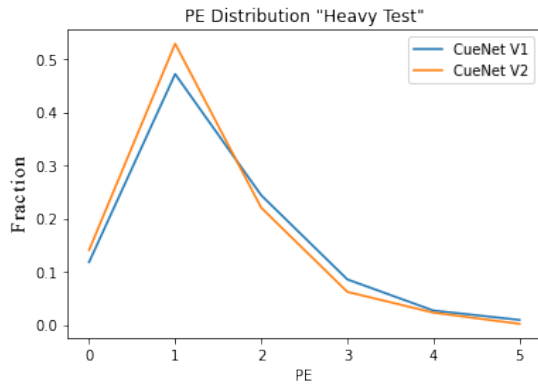


Figure 25: PE distribution of the two versions of CueNet on a more difficult testing set.

In this benchmark CueNet V1 achieved 92.9% accuracy with our given tolerance range of $PE \leq 4$ and CueNet V2 achieved an accuracy of 96.4%. When comparing the mistakes both networks made you can see that CueNet V1 often loses track of the ball completely and thus labels the ball as somewhere completely else. By contrast, CueNet V2 locates the ball just barely outside of the tolerance range. We think that CueNet V2 is more robust because it can rely on the two older images and can therefore approximate the position of the ball with the learned trajectory patterns. In addition, it can compare changes between the consecutive images, making it easier to locate the ball since the ball is the fastest moving object. The potential tradeoff for CueNetV2's higher accuracy and robustness is higher computational cost, which could be important for the robotics control.

As the last test, we combined some of the footage of our "Heavy Test" set to test the multi-object tracking capabilities of our network. Despite never having trained this, both versions of CueNet can track two objects simultaneously with sufficient result. If we had trained this we are confident that simultaneous tracking would not have been a problem since the last layer is a softmax layer and can easily split the probability distribution and thus track multiple objects.




Figure 26: ■ Output of CueNet tracking two balls simultaneously, where the heatmap output in red is overlaid over the APS image.

Conclusion

In the context of this work, we developed a Fully Convolutional Neural Network named CueNet that can effectively detect and track the cueball on the used labyrinth. We also found that CueNet V2 outperforms CueNet V1, which indicates that using consecutive images as an input to the CNN increases its accuracy overall and especially in difficult situations. Additionally, both versions of CueNet could generalize well and thus were able to track multiple objects simultaneously, despite never having been trained on tracking more than one object. By extensively testing CueNet in difficult conditions, we could confirm the robustness/reliability of the algorithm under real-life conditions.

3 Discussion

3.1 Applications of AI

In recent years, AI has found many new applications in a wide range of disciplines. In this section, we will cover some in which AI has made significant advances.

Medicine

Through the vast amount of generated and stored data in healthcare, AI today has found use in diagnosis, drug development and the creation of personalized treatments.

Deep Learning algorithms assist doctors in the evaluation of CT scans, MRI images, skin images etc. for the diagnosis of all kinds of diseases. Some of the algorithms

can achieve accuracy of up to 99% which is almost as good as a professional doctor [41]. The immense benefit of these algorithms is their low cost and speed. They can draw conclusions in seconds and the results can be shared across the world at the speed of light. This means that even people in poorer regions can receive a professional diagnosis for less money. [42]

When designing the treatment for a patient there are many variables one has to take into account. Small differences in drug doses and treatment schedules can lead to different reactions from different patients. With the help of machine learning, doctors can find correlations between the patient's characteristics and the patient's response to a particular treatment. This helps with choosing the right treatment and thus maximizing the chance of the patient's survival. [42]

Finance

The use of AI in banking dates back to 1987 where the Security Pacific National Bank in the US started using an expert system for the detection of unauthorized use of debit cards [43]. Today machine learning is used to reduce the chance of fraud and financial crimes by monitoring behavioural patterns for anomalies and flag them for human review [44].

In the world of trading, there is the never-ending competition of coming up with the best mathematical model to predict the future of the stock market. High-frequency trading represents one of the fastest-growing sectors in financial trading and of course, AI was found to be useful with this kind of problem. By 2010, 60% of all trades were executed by computers of which most work with "simple" algorithms (Algorithmic Trading). Those algorithms can make tens of thousands of trades per second. Machine Learning, on the other hand, working a lot slower, can learn sophisticated patterns from past data that could indicate the next financial crisis, a bullish/bearish trend etc. Many banks, funds and trading firms have their entire portfolios managed purely by AI systems [45]. There are ETFs (Exchange Traded Funds) whose portfolios are built with the help AI like ARK's Industrial Innovation ETF. Starting this year it gained 69.60% compared to the average yearly 10% that the S&P 500 gains, which is quite an accomplishment. [46].

Transportation

We are on the verge of a revolution in the automotive industry, there are currently over 40 companies developing self-driving cars using AI. Tesla and even tech giants like Apple and Google are all developing their own systems [47], of which Tesla's is the most advanced so far. The advantage of Tesla lies in its possession of a huge amount of real-life data. Their entire fleet collects data while driving and sends it back to Tesla, where they can train their neural network with it. Driving automation is subdivided into six levels (See figure 27), of which we have only reached level 2 on commercial products up to this day.

Tesla has already demonstrated a level 5 system, but mainstream production of anything higher than level 2 has yet to come. The reason for this is partially the lack of technological capability but also the ethical and political questions that go along with autonomous driving. Another problem of autonomous systems is that they can be exploited and thus serious accidents could occur. In chapter 3.2.1 we will talk more about the reliability and potential safety exploits of AI.

Computer Vision

Image restoration, object recognition, video tracking, scene reconstruction, motion estimation are all sub-domains of computer vision [48]. The practical part of this paper goes into object recognition and video tracking, as we are trying to track a cueball from a continuous stream of frames. Object tracking and recognition are crucial for autonomous systems because they are built upon them. For them to be able to make decisions, they first have to see. In figure 28, an algorithm tracks and recognizes different kinds of vehicles.



Figure 28: Tracking of vehicle movement using Artificial Intelligence.

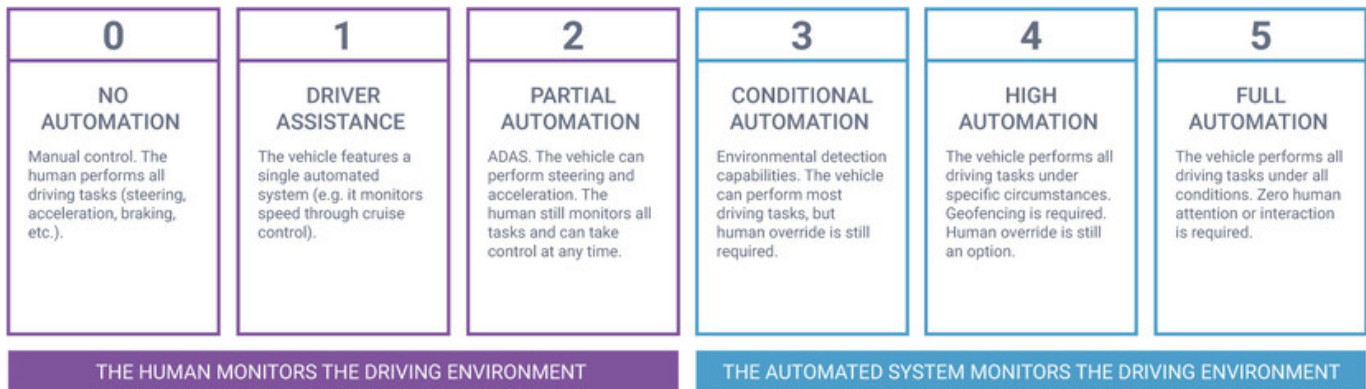


Figure 27: Levels of Vehicle Autonomy.

In image restoration, the goal is to remove unwanted noise from pictures or add missing details. This is typically done by feeding the network with millions of images until it can fill the missing content considering the other content of the image. [49]



Figure 29: Image restoration using Artificial Intelligence.

When does applying AI make sense?

Implementing an AI system makes sense when working with large amounts of data. In this large amount of data, the algorithm can approximate a generalization for a problem that can handle data which it has never seen before. Getting clean, balanced and unbiased data can sometimes be very difficult, though. The **Return on Investment (ROI)** is a value that is crucial when deciding over the implementation of an AI system. This means weighing up the cost and complexity of the implementation versus the benefit the AI can give while operating. Simpler tasks like for example regulating the heating of a house can simply be done autonomously without an AI operating in the background - the ROI of the AI solution might not be worthwhile here. Regarding more complex tasks, such as autonomous driving, the ROI

gets dramatically better. Solving this task without the help of AI would seem infeasible.

Most tasks that can be completed with a data-defined pattern can be done with AI. Depending on the difficulty of the pattern it tries to approximate, it may be more efficient to manually program in all the rules for each case. As AI continues to advance, gathering sufficient data and implementing the algorithm is going to become simpler and simpler from a business perspective. The ROI for AI solutions will continue to ascent and AI will make its entrance into more and more sectors.

3.2 Challenges regarding AI

3.2.1 Reliability

The further AI advances, the more complex and important tasks it will be able to take on. More important tasks also mean greater responsibility, especially as an increasing number of humans rely on that system. Therefore, **technical AI safety** becomes increasingly more important. While being able to empirically prove the accuracy or performance of an AI system in the lab, real-world scenarios always hold an unforeseen situation. This effect is called "data shift" and is an important problem to regard when trying to build a reliable AI system.

The technical safety of an AI can be divided into three areas: specification, robustness and assurance. Defining the right purpose of the system and trying hard to avoid ambiguities and side-effects caused by the way the system is designed is what falls into the area of specification. The robustness of a system is measured by its ability to withstand unwanted interference. This can either be done

by proper training (see practical part of this work) or by error-correction and fallback mechanisms for recovery. The last area, assurance, is mainly about monitoring (see next section) and checking the system activity while in production. [50]


3.2.2 Interpretability

To be able to use the power of AI for important decisions without blindly relying on it, we need to know based on what factors the AI decides.

Today, we mostly don't know how to interpret/explain the reasoning behind complex AI systems. That is why one of the problems we will probably have to work on is the interpretability of the algorithms. The term "**Explainable AI (XAI)**" refers to AI systems that possess some amount of "transparency", which means that humans can understand in at least some detail, why the algorithm came to the solution it did [51]. The challenge of creating XAI is to improve the "transparency" of the model to a level where it can prove useful to humans without compromising the performance of the model too much. Also, an "effective explanation" of the model can vary between different industries and users of the AI system. A developer, for example, would rather like to know possible areas of improvements, while a judge would rather want to know if the AI is making fair decisions [52].

3.2.3 Moral Dilemma

"Moral dilemmas are situations in which the decision-maker must consider two or more moral values or duties but can only honour one of them; thus, the individual will violate at least one important moral concern, regardless of the decision." [53]. In moral dilemmas, because all choices are bad, deciding is difficult and relies heavily on personal biases. The question regarding AI is the following: Should we give AI systems the responsibility of making ethical decisions? How do we determine in which way an AI system reacts to moral dilemmas? To feel how difficult deciding becomes in dilemma situations, visit <https://www.moralmachine.net/> (see QR-code below).



As you've probably seen, it can sometimes be very difficult, and as self-driving cars or other AI systems carry more responsibility, how we solve such situations becomes more important.

Another stage where moral questions have to be answered is when AI becomes sentient or even self-aware. When AI starts to be conscious, is it correct to just unplug it or delete it? Or even further, should the AI system have human or civil rights?

There are still a lot more questions regarding morality. For example, should AI be banned for war purposes like chemical and biological weapons since the Geneva Protocol? Is it morally reprehensible to further develop AI systems regarding all the future threats? We cannot and don't want to try to give any answers to these questions, but draw attention to the fact that they exist.

3.2.4 Energy and Computing Power Demand

In 2016, when AlphaGo defeated Lee Sedol in the game of Go, it used 1202 CPUs and 176 GPUs to compute its next move. To operate, it consumed about 1 megawatt, which, compared to Lee Sedol, is about 50'000 times more power [54].

Using such big amounts of energy is impractical, expensive and environmentally harmful. There is a general rule that the more complex the patterns to find or the task to solve, the more complex and deep your network should be. But this also comes at the cost of more computations and energy needed.

There have already been developed techniques to

drastically reduce the size/”weight” of neural networks while keeping up the good performance. One of the techniques is called ”Pruning”, in which the researcher iteratively removes connections between neurons and retrains the network afterwards to achieve a minimum level of redundant connections in the network [55]. Efficiency gains can also be achieved through tailored hardware, such as Google’s TPUs (as explained in 1.2.2 under paragraph ”Google”), as well as in the specialized hardware architectures developed all around the world, including at ETH Zurich and UZH.

3.3 Outlook

There are still many unresolved questions regarding the technology of AI. This is where we would like to leave the floor to the experts in the field, which is why we have conducted an interview with Prof. Dr. Martin Vechev, who leads the Secure, Reliable, and Intelligent Systems Lab (SRI) at ETH Zurich.

Q: Should AI be regulated more/less or is it just right and why? Are we on the way to losing control over AI? What are possible disadvantages of AI?

A: I believe we are far away from losing control over AI. While a scenario where an Artificial General Intelligence starts to improve itself in a runaway fashion is a popular topic for Science Fiction novels and might become a concern in the far future, we are still very far away from such a form of AI. Therefore, I think the time to regulate AI more strictly has not yet come, as we simply do not yet know what the AI systems of the future will look like and should avoid stifling innovation by overly strict regulations. When talking about individual tasks where AI is now used instead of other systems, it might make sense to generally regulate how decisions are made. I don’t see a reason to apply such regulations only to AI Systems and not the decision making system used before. AI is like any tool, in that it has to be used correctly to produce good results. A danger might be that it is harder to assess whether it is working well. The performance might vary significantly depending on the circumstances. For example, if the dataset used to train and evaluate a model is heavily biased in the same way, this will go undetected during testing but might have unforeseen consequences once it gets applied in a real-world environment. Therefore it is important to be careful and diligent when implementing AI solutions

starting with the gathering of data.

Q: In what ways are AI systems exploitable and what can we do against it? (Adversarial attacks etc.)

A: AI systems are exploitable in several ways. The training data might be biased, intentionally or not, manipulated data can be injected into the training procedure and the resulting behaviours later exploited, or the data evaluated by an AI system could be engineered to mislead the system. These maliciously manipulated data points can fool AI systems that are not trained to be robust to such attacks, while looking unchanged to human eyes. These weaknesses each require different countermeasures. Methods to analyse the vulnerability of an AI system to these adversarial examples and training algorithms to defend against them are an active research topic that we are also working on. Many methods exist to generate empirical or even deterministic robustness guarantees, although those come at the cost of increased computational cost and reduced accuracy.

Q: Who is responsible for AI systems when they fail?

A: Who will be responsible for a failing AI system is certainly an interesting question that will most likely have to be decided by regulators. In general, I believe it is unreasonable to expect that an AI system will be 100% under all (potentially unforeseen or maliciously engineered) circumstances, therefore defining what a failure is, is critical to answer this question. If you are referring to cases such as autonomous driving systems confusing a bright side panel of a trailer or car with the sky and failing to detect an obstacle, the entity providing the system is probably responsible for the failure of the AI system. However, this is not the same as responsibility for a potentially resulting accident. An observant driver who could have intervened, society in general exaggerating the ability of such a system or a regulator not defining appropriate certification requirements might also all be partially responsible.

Q: What’s the outlook on the optimization of the power consumption of AI algorithms? Is there a lot of room for optimization? AlphaGo consumed about 1 megawatt, which, compared to Lee Sedol, is about 50’000 times more power.

A: *Energy efficiency is an important question in AI research and there are significant efforts to reduce the energy and hardware requirements to enable the use of these systems for example on mobile phones. While there probably is still significant room for optimization, a factor of 50,000 seems unlikely. The AlphaGo example is also an extreme case demonstrating what is possible if efficiency constraints are ignored. Considering that the successor to AlphaGo, AlphaGo Zero, is estimated to have an Elo of over 5000, corresponding approximately to losing only 1 out of 1000 games against the best human go players, using more power might also be acceptable.*

Q: *Is the brain efficient? (in the sense that it only can be simulated by a thing bigger and more complicated than itself)*

A: *It is almost universally true that it requires a more complex system to fully simulate the behaviour of a system, this is called computational irreducibility, while there are some exceptions to this behaviour, none have yet been found for the brain. The question is, whether you want to actually simulate a human brain, or solve the same tasks. The latter seems like the more interesting problem and makes the question harder to answer. If you think about complex mathematical operations the human brain is orders of magnitudes slower and less efficient than even a simple calculator, while AI systems can not replicate other feats. Therefore, I will go with: The brain is probably efficient for some tasks but not for others.*

Q: *Try to estimate the impact of AI in the workplace in the short term (5-10 years) and in the long term (50+ years). Will humans become obsolete in the workplace?*

A: *I predict the biggest short term impact of AI on the workplace to be in the transportation field. Autonomous driving in combination with electrification has the potential to revolutionize both cargo and personal transportation. To make predictions about the far future would only be guessing. Everything from regulators clamping down on AI so heavily as to make it unattractive up to true Artificial General Intelligence is possible. While I don't believe that humans will become obsolete in the workplace, their work might shift to tasks focused more around human interaction.*

Q: *Is the creation of Artificial Intelligence the most*

important event in human history so far?

A: *Achieving Artificial General Intelligence could become the most important event in human history, but until then I don't believe so.*

Q: *How far are we from Artificial General Intelligence and a Superintelligence?*

A: *It does not look like the development of an AGI or Superintelligence is imminent, but whether it will take 20 or 200 years, if it ever happens, is difficult to say.*

4 Appendix

4.1 Reasons for the choice of topic

AI could unravel to be the most important invention of humankind of our time or even of all time. Because of this and the rapid development during this decade it was an easy decision for us to choose AI as the subject of our Matura Thesis. We believe that the future lies in AI and therefore it is worthwhile to have engaged with it. As potential computer science students we could imagine a specialization towards AI in the later course of our studies and thus a job in this field. So far only the tip of the iceberg has been scratched, the field of AI will take a big step in the direction of progress during our lifetimes and we strive to contribute to this development.

4.2 Acknowledgements

We would first like to thank our thesis advisor Prof. Dr. Delbrück Tobias of the Institute of Neuroinformatics. With any problems or questions we encountered he guided and encouraged us to do the better even when it was not easy. He also gave us access to the needed equipment for this project, which we profoundly appreciate.

Of course, we also thank the PhD students Yuhuang Hu and Shasha Guo, whom Tobias Delbrück provided for additional help. They were very responsive and we therefore could always rely on quick support for smaller questions.

We would also like to acknowledge the work of Michael Bucher and Christopher Latkoczy from Kantonsschule Stadelhofen as our school intern advisors and proof-readers of this thesis, we are gratefully indebted to their valuable suggestions on this thesis.

Finally, we must express our deepest gratitude to our family and close friends for providing us with unfailing support and continuous encouragement throughout the process of researching and writing this thesis. This accomplishment would not have been possible without them. Our sincerest gratitude to you.

Declaration of Consent

The authors of this paper hereby declare that this paper was written by them independently and that no sources other than those indicated were used. They also acknowledge that this paper can be checked with a plagiarism detection tool.

Location, Date

Signature

Location, Date

Signature

References

- [1] Artificial intelligence, internet encyclopedia of philosophy. www.iep.utm.edu/art-inte/, last accessed on 26.07.2020.
- [2] Philosophy of computing and information technology, philosophy of technology and engineering sciences. www.sciencedirect.com/topics/engineering/artificial-intelligence, last accessed on 26.07.2020.
- [3] 7 types of artificial intelligence. www.forbes.com/sites/cognitiveworld/2019/06/19/7-types-of-artificial-intelligence/, last accessed on 04.08.2020.
- [4] What are the 3 types of ai? www.codebots.com/artificial-intelligence/the-3-types-of-ai-is-the-third-even-possible, last accessed on 04.08.2020.
- [5] Murray Shanahan. The technological singularity, 2015. p. 86.
- [6] What is machine learning? www.technologyreview.com/2018/11/17/103781/what-is-machine-learning-we-drew-you-another-flowchart/, last accessed on 04.08.2020.
- [7] What is deep learning? www.machinelearningmastery.com/what-is-deep-learning/, last accessed on 04.08.2020.
- [8] History of artificial intelligence. <https://www.coe.int/en/web/artificial-intelligence/history-of-ai>, last accessed on 11.09.2020.
- [9] The history of artificial intelligence. <https://www.coe.int/en/web/artificial-intelligence/history-of-ai>, last accessed on 11.09.2020.
- [10] Daniel Crevier. Ai: The tumultuous history of the search for artificial intelligence, 1993. Simon and Newell 1958, p. 7-8.
- [11] Daniel Crevier. Ai: The tumultuous history of the search for artificial intelligence, 1993. p. 100-144.
- [12] Daniel Crevier. Ai: The tumultuous history of the search for artificial intelligence, 1993. p. 102-105.
- [13] Pamela McCorduck. Machines who think (2nd ed.), 2004. p. 327-335.
- [14] Daniel Crevier. Ai: The tumultuous history of the search for artificial intelligence, 1993. p. 148-159.
- [15] Daniel Crevier. Ai: The tumultuous history of the search for artificial intelligence, 1993. p. 158-159.
- [16] Peter Norvig. Artificial intelligence: A modern approach (2nd ed.), 2003. p. 23-24.
- [17] Pamela McCorduck. Machines who think (2nd ed.), 2004. p. 436-441.
- [18] Daniel Crevier. Ai: The tumultuous history of the search for artificial intelligence, 1993. p. 240.
- [19] Sebastian Schuchmann. History of the second ai winter. <https://towardsdatascience.com/history-of-the-second-ai-winter-406f18789d45>, last accessed on 24.09.2020.
- [20] Pamela McCorduck. The singularity is near, 2004. p. 480-483.
- [21] Ray Kurzweil. Machines who think (2nd ed.), 2005. p. 274.
- [22] Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturowski, Pablo Sprechmann, Alex Vitvitskyi, Daniel Guo, and Charles Blundell. Agent57: Outperforming the atari human benchmark, 2020.

- [23] Nvidia gpu market share. <https://wccftech.com/amd-radeon-and-nvidia-geforce-discrete-gpu-market-share-q4-2019/>, last accessed on 12.08.2020.
- [24] Nvidia. <https://en.wikipedia.org/wiki/Nvidia>, last accessed on 12.08.2020.
- [25] Cloud tensor processing units (tpus). <https://cloud.google.com/tpu/docs/tpus>, last accessed on 12.08.2020.
- [26] Project debater. <https://www.research.ibm.com/artificial-intelligence/project-debater/about/#>, last accessed on 12.08.2020.
- [27] Artificial intelligence in retail innovation, by amazon. <https://arekskuza.com/the-innovation-blog/amazon-and-artificial-intelligence-in-retail/>, last accessed on 12.08.2020.
- [28] Facebook ai research. <https://ai.facebook.com/>, last accessed on 12.08.2020.
- [29] C Brandli, R Berner, M Yang, S C Liu, and others. A 240×180 130 db $3 \mu\text{s}$ latency global shutter spatiotemporal vision sensor. *IEEE J. Solid-State Circuits*, 2014.
- [30] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128×128 120 db $15 \mu\text{s}$ latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circuits*, 43(2):566–576, 2008.
- [31] jaer software. <https://github.com/SensorsINI/jaer/>, last accessed on 06.08.2020.
- [32] Brio labyrinth game. www.brio.us/products/by-age/6-years-and-up/labyrinth, last accessed on 06.08.2020.
- [33] What the heck is a tensor?!? www.youtube.com/watch?v=bpG3gqDM80w, last accessed on 06.08.2020.
- [34] What’s a tensor? www.youtube.com/watch?v=f51liqUk0ZTw, last accessed on 06.08.2020.
- [35] Explained: Neural networks. <http://news.mit.edu/2017/explained-neural-networks-deep-learning-0414>, last accessed on 08.08.2020.
- [36] R. H. Hahnloser, R. Sarpeshkar, M. A. Mahowald, R. J. Douglas, and H. S. Seung. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, vol. 405, no. 6789, pp. 947–951.
- [37] A comprehensive guide to convolutional neural networks — the eli5 way. <https://towardsdatascience.com/a-comprehensive-guide-to-convolutional-neural-networks-the-eli5-way-3bd2b1164a53>, last accessed on 09.08.2020.
- [38] Grant Sanderson. Backpropagation calculus — deep learning, chapter 4, 2017. <https://youtu.be/tIeHLnjs5U8>, last accessed on 15.12.20.
- [39] Michael Nielsen. How the backpropagation algorithm works, 2019. <http://neuralnetworksanddeeplearning.com/chap2.html>, last accessed on 15.12.20.
- [40] Y. Huang, I. Liao, C. Chen, T. İk, and W. Peng. Tracknet: A deep learning network for tracking high-speed and tiny objects in sports applications*. In *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–8, 2019.
- [41] D Kermany, M Goldbaum, and Kang Zhang. Identifying medical diagnoses and treatable diseases by image-based deep learning. [https://www.cell.com/cell/fulltext/S0092-8674\(18\)30154-5?_returnURL=https%3A%2F%2Flinkinghub.elsevier.com%2Fretrieve%2Fpii%2FS0092867418301545%3Fshowall%3Dtrue#secsectitle0030](https://www.cell.com/cell/fulltext/S0092-8674(18)30154-5?_returnURL=https%3A%2F%2Flinkinghub.elsevier.com%2Fretrieve%2Fpii%2FS0092867418301545%3Fshowall%3Dtrue#secsectitle0030), last accessed on 08.10.2020.

- [42] Markus Schmitt. Artificial intelligence in medicine. <https://www.datarevenue.com/en-blog/artificial-intelligence-in-medicine>, last accessed on 08.10.2020.
- [43] Charles A Christy. Impact of artificial intelligence on banking, 17 January 1990. <https://www.latimes.com/archives/la-xpm-1990-01-17-fi-233-story.html>, last accessed on 09.10.2020.
- [44] Machine learning and ai for fraud prevention: A primer. <https://feedzai.com/resources/machine-learning-ai-fraud-prevention-primer/>, last accessed on 10.10.2020.
- [45] Algorithmic trading. <https://www.investopedia.com/terms/a/algorithmictrading.asp>, last accessed on 10.10.2020.
- [46] Artificial intelligence etf list. <https://etfdb.com/themes/artificial-intelligence-etfs/>, last accessed on 10.10.2020.
- [47] 40+ corporations working on autonomous vehicles. <https://www.cbinsights.com/research/autonomous-driverless-vehicles-corporations-list/>, last accessed on 10.10.2020.
- [48] Tim Morris. Computer vision and image processing, 2004.
- [49] Ziyu Wan, Bo Zhang, Dongdong Chen, Pan Zhang, Dong Chen, Jing Liao, and Fang Wen. Old photo restoration via deep latent space translation, 2015. <https://arxiv.org/pdf/2009.07047.pdf>, last accessed on 19.12.20.
- [50] Building safe artificial intelligence: specification, robustness, and assurance. <https://medium.com/@deepmindsafetyresearch/building-safe-artificial-intelligence-52f5f75058f1>, last accessed on 28.12.2020.
- [51] Explainable artificial intelligence. https://en.wikipedia.org/wiki/Explainable_artificial_intelligence, last accessed on 13.08.2020.
- [52] Xai—explainable artificial intelligence. <https://robotics.sciencemag.org/content/4/37/eaay7120>, last accessed on 21.08.2020.
- [53] Moral dilemmas. https://link.springer.com/chapter/10.1007/978-3-030-15191-1_2, last accessed on 21.08.2020.
- [54] Another look at alphago vs. lee sedol: The power angle. <https://www.ceva-dsp.com/ourblog/artificial-intelligence-leaps-forward-mastering-the-ancient-game-of-go/>, last accessed on 22.08.2020.
- [55] Lecture 15 — efficient methods and hardware for deep learning, August 2017. <https://www.youtube.com/watch?v=eZd0kDtYMoo&list=PL3FW7Lu3i5JvHM8ljYj-zLfQRF3E08sYv&index=15>, last accessed on 22.08.2020.

List of Figures

1	Source: https://towardsdatascience.com/what-is-deep-learning-and-how-does-it-work-2ce44bb692ac . . .	2
2	Source: Image created by authors of this paper.	5
3	Source: Image created by authors of this paper.	6
4	Source: https://en.wikipedia.org/wiki/Tensor	6
5	Source: Image created by authors of this paper.	7
6	Source: https://towardsdatascience.com/building-a-deep-learning-model-using-keras-1548ca149d37 . . .	8
7	Source: https://towardsdatascience.com/machine-learning-for-beginners-an-introduction-to-neural-networks-d49f22d238f9	8
8	Source: https://towardsdatascience.com/machine-learning-for-beginners-an-introduction-to-neural-networks-d49f22d238f9	8
9	Source: Image created by authors of this paper.	9
11	Source: https://towardsdatascience.com/types-of-convolution-kernels-simplified-f040cb307c37	9
10	Source: https://missinglink.ai/guides/convolutional-neural-networks/convolutional-neural-network-tutorial-basic-advanced/	10
12	Source: https://stats.stackexchange.com/questions/362988/in-cnn-do-we-have-learn-kernel-values-at-every-convolution-layer	10
14	Source: Image created by authors of this paper.	10
15	Source: Image created by authors of this paper.	10
16	Source: Image created by authors of this paper.	11
17	Source: https://mc.ai/gradient-descent-an-optimization-method-used-in-machine-learning/	11
18	Source: https://blog.paperspace.com/part-3-generic-python-implementation-of-gradient-descent-for-nn-optimization/	12
19	Source: Image created by authors of this paper.	12
20	Source: Image created by authors of this paper.	12
22	Source: Image created by authors of this paper.	15
21	Source: Image created by authors of this paper.	16
23	Source: Image created by authors of this paper.	16
24	Source: Image created by authors of this paper.	16
25	Source: Image created by authors of this paper.	17
26	Source: Image created by authors of this paper.	17
28	Source: https://www.wikiwand.com/en/Computer_vision	18
27	Source: https://www.synopsys.com/automotive/autonomous-driving-levels.html (modified)	19
29	Source: https://en.wikipedia.org/wiki/Inpainting	19