

Potential enrichments in malaria diagnostics:
hyperspectral imaging and
group-equivariant neural networks
(Thesis work)

Author

C.C.J. van Engelenburg (Casper) [4237080]

Supervisors

Prof. G.V. Vdovin (Gleb)

Dr. T.E. Agbana (Tope)

Institute and direction

Master Systems and Control (S&C),
Delft Center for Systems and Control (DCSC) [1],
Delft University of Technology (TUD) [2]

December 4, 2020

1 | Abstract

Proper diagnostics are essential in the combat against severe diseases which mainly have big impacts in remote areas in poor countries. A focus direction within the NC4I group at DCSC, Delft University of Technology, is the development of new imaging modalities and the design and implementation of smarter algorithms for improved detection of parasitic diseases. The first part of my research exploits hyperspectral imagery (HI) as new potential imaging modality of thin blood smears that could highly improve on preparation time, labor intensiveness and use of materials. HI retrieves both spatial and spectral information of the observed objects simultaneously, thus providing the ability to discriminate near similar constituents within the blood smear. In doing so, it enables the possibility of label-free detection. In this thesis, the development and building of such a system is addressed and carried out. In the context of malaria, it is shown that HI is promising and lays a profound foundation for further exploration. The design and evaluation of improved generalizing neural networks characterize the essence of the second and larger part of the research. Several group-equivariant networks are evaluated and compared with conventional convolutional networks which shows that efficient and redefined integration of weights can help build smarter and more robust classifiers for the detection of parasites. In group-equivariant networks, re-interpreting the way feature maps are connected to one another manifests in the development of convolutional stages that equivary under an increased amount of transformations besides merely translations. It is shown that enlarging the heuristic of that transformation group (the extra amount of transformations the operations are equivariant under) significantly contributes to better performance without necessarily increasing the size or changing the architecture of the networks. Compared to the aforementioned baseline (conventional convolutional stages), the best network (being equivariant under 16 equidistant rotations and mirror reflections) improves approximately 2-fold on all relevant performance metrics, among which are accuracy, sensitivity, specificity, precision, and the F1-score which are common measures in the classification of malaria. The networks were tested on the Rajaraman database [3]. Furthermore, the pre-trained models are used as classifiers for a different database extracted from the microscope build by AiDx medical [4]. At least for this specific database, it is shown that the more realistic transformations the pre-trained networks equivary under, the more robust they are.

Contents

1	Abstract	3
2	The burden of malaria	13
2.1	Malaria worldwide	13
2.2	Diagnostics	15
2.3	Autonomous detection	18
2.4	Opportunities	18
2.5	Research question and proposal	20
2.5.1	Hyperspectral imaging as novel imaging modality in the diagnosis of malaria	20
2.5.2	Towards <i>smarter</i> classifiers in the diagnosis of malaria	21
I	Hyperspectral Imaging	23
3	Physical background	27
3.1	Light and optics	27
3.1.1	Geometrical optics	27
3.1.2	Physical optics: numerical aperture and spatial resolution	28
3.1.3	Bright-field microscopy and koehler illumination	30
3.2	Spectral characteristics	32
3.2.1	Light-matter interactions	33
3.2.2	Light-tissue interactions in blood	34
3.2.3	Dispersion and chromatic aberration	36
3.2.4	The electromagnetic spectrum and spectral purity	36
4	Design and realization of a hyperspectral imaging system	41
4.1	Hyperspectral imaging systems	41
4.2	Design and realization of the optical scheme	43
4.2.1	Source	43
4.2.2	Illumination path	44
4.2.3	Objective and camera	45
4.3	Evaluation and discussion	46

II	Towards Smart Generalizing Networks	51
5	Generalizing operations in neural networks	55
5.1	Artificial neural networks	55
5.2	Equivariant operations	59
5.2.1	Convolutions	60
5.2.2	Group convolutions	62
5.2.3	Steerable convolutions	69
5.2.4	Steerable filter design	73
6	Design, training and implementation of deep networks	77
6.1	Deep learning: residual networks	77
6.1.1	The residual block	77
6.1.2	Encoding the network	79
6.1.3	Decoding the network	79
6.1.4	Overall architecture	80
6.1.5	Remarks	81
6.2	Training the networks	82
6.2.1	Cost function	82
6.2.2	Activation function	83
6.2.3	Sample and batch size	84
6.2.4	Stochastic gradient descent	84
6.2.5	On the evaluation of the training process	85
7	Evaluation and results	87
7.1	Data sets and manipulations	87
7.1.1	Rajaraman data set	88
7.1.2	AiDx data set	89
7.2	Performance metrics	89
7.2.1	Accuracy, sensitivity and specificity	89
7.2.2	The F-score	90
7.3	K-fold cross-validation	91
7.4	Results	92
7.4.1	Rajamaran data set	92
7.4.2	AiDx data set	96
8	Discussion and conclusions	97

List of Figures

2.1	Annual deaths (left) and death rates (right) due to <i>P. falciparum</i> . Children under the age of 5 are most sensitive. Mortality in Africa is decreasing significantly over the last decade by 35%. Moreover, the areas outside Africa are, in terms of regulating their mortality, doing worse eventhough their absolute number of deaths have decreased as well.	14
2.2	Pathogenesis of the infiltration and ingestion of malaria parasites. There is three distinct cycles within: the <i>exo-erythrocytic</i> , <i>erythrocytic</i> and <i>sporogonic</i> cycle. The <i>exo-erythrocytic</i> cycle marks the evolution of sporozoites into merozoites and primarily takes place in the liver, the <i>erythrocytic</i> cycle resembles the infiltration of erythrocytes and formation of sexual gametocytes within the blood vessels and the <i>sporogonic</i> cycle remains for the production of fresh sporozoites which takes place in the mosquito. [5].	15
2.3	Categorization diagram for malaria diagnostic methods. On a high-level, three distinct categories exist. <i>Clinical diagnosis</i> is merely diagnosing on the basis of present symptoms of the patient, <i>parasitic detection</i> is focussed on detecting the parasite or parts of it and <i>by-product detection</i> is based on detecting either products that the parasites produce as well as products the patients themselves produce. Parasitic detection's sub-categories include morphological and nucleic acid based detection methods and by-product detection include antigen and anti-body based methods. Some specified methods are not discussed in my research such as micro-assay, serology, flow cytometry, ACC and MS.	16
2.4	Overview of conventional steps taken in light microscopy and specifically its examination which in this case is autonomous. The high-level steps are self-explanatory. Five subsequent steps should be taken in autonomous detection. <i>Data acquisition</i> constitutes the retrieval of digital images. <i>Pre-processing</i> transforms (e.g. cleans, normalizes) the raw images such that it gets denoized, becomes relevant and its redundant information thrown away. <i>Segmentation</i> partitioned the image into multiple objects, which, in the case of malaria, is the segmentation of individual erythrocytes. Arguably the hardest part is the determination of features that are used to discriminate between infected cells, determine the species and infiltration stage. This process is called <i>feature selection</i> . <i>Classification</i> is the final stage and classifies, based on the features, individual cells.	19
2.5	21

3.1	Geometrical optics: Schnell’s law and the thin lens.	28
3.2	Light-matter interactions in general. From left to right: absorption, (in-)elastic scattering, transmission and emission.	33
3.3	Light-tissue interactions in whole blood. At the top, the predominant interactions are sketched: transmission, scattering and absorption. The middle elaborates on both type of scattering that occurs (left) as well as the main component of erythrocytes, namely haemoglobin (right). The bottom shows the absorption and scattering coefficients for HbO ₂ and Hb. Here, the resolution is 2nm until 600nm and 5nm upwards the spectrum. All data is directly copied from the tabulated data in the Appendix of [6].	35
3.4	Chromatic aberration in lenses due to dispersion. On the left three types of lens configurations are sketched: the single lens (top), the achromatic doublet lens (middle), and the achromatic triplet lens (bottom). In optical terms, 656.3nm (red) is referred to as ‘C’ light, 587.6nm (yellow) as ‘d’ light, and 486.1nm (blue) as ‘F’ light. For the single lens, only ‘d’ light is perfectly focused, while achromats do so for ‘F’ and ‘C’. Apochromats do so for all three. On the left, the <i>focus shifts</i> are given for all intermediate wavelengths. In terms of overall performance in VIR-region, apochromats usually obtain the highest degrees in LCA correction.	37
3.5	Spectrum of a Halogen lamp.	38
3.6	Spectral purity as measure for the monochromaticity of a light’s spectrum. Here, the spectral purity of the frequency is given. Please note that the report uses the wavelength.	39
4.1	Categorization of HI systems [7].	42
4.2	Data retrieval in HI. In the top-left corner, the set containing all images at specific wavelengths is given. Instead of representing it as a set of features, the images get stacked in what is called a <i>3D-hypercube</i> (bottom-left). Every pixel is therefore associated with a spectral response for which some are sketched on the right. The matrix representing all individual spectral responses in its columns is referred to as the <i>voxel-set representation</i>	42
4.3	Overview of the optical scheme: staring scanning hyperspectral imaging microscope. Light travels from A to B in which it chronologically follows the route through the source, the illumination path, the sampling stage, the objective and is arriving at the camera - all highlighted in bold. Sub-stages that are together fixed in space are outlined in black: the sub-stage condenser, the sub-stage sample and the capturing sub-stage.	43

4.4	The source and spectral purity of its extension. On the left, the source is illustrated in which the outputted band is tuned by rotating the blazed grating. An optical fiber is coupled to the output of the monochromator and. The amount of light and spectral purity highly depend on the chosen configuration - size of entrance slit and fiber diameter. Spectral purity is measured for two configurations as indicated on the right. Part of the measurements are shown in the range of 600 - 620 nm. Furthermore, a Gaussian fit is used to approximate the mean and <i>full-width half-maximum</i> (FWHM) for every selected wavelength measured at. Least-squares using the Levenberg-Marquardt algorithm is used as optimizer [8].	44
4.5	Koehler aligned setup.	45
4.6	Spatial resolution of the setup for various configurations. The spatial resolution, d_λ , is theoretically computed according to the top-noted equation. Here, NA_{obj} , λ and F_{cond} are fixed (underlined in red) and the magnification of the field, M_{field} , and the diameter of the source, D_{source} , are used as configuration parameters (underlined in blue). Green dots denote the region with sufficient spatial resolution while the red-dotted legion is the non-optimal domain. . .	46
4.7	Calibrated images of Giemsa-stained thin blood smear films for several wavelengths. Crops are taken from the whole FoV which is approximately 15 times as big, thus containing about 500 erythrocytes in total. A plan-apochromat objective with 20X magnification is used.	47
4.8	Hyperspectral imaging microscope: the real setup.	48
4.9	Proposed flow-diagram for non-invas classification of malaria parasites using hyperspectral imaging. The first region concerns the acquisition and preparation of the sample. The second region shows the measurements and the third the pre-processing steps. The fourth is the protocol for segmentation and storing of individual cell data. The fifth is an automated detection protocol based on these individual cell data. It should be immediately noted that there is two parallel horizontal lines: one representing hyperspectral imaging and the other conventional imaging. This is done because we need a validation set to train and test the data on. Since there is no clue based on the hyperspectral data only how to classify every cell, we need a method on which we can trust that does so reliably. To do so we stain, after HI measuring, an equivalent sample and examine that under a conventional microscope. The same domain is evaluated and labels are given for the whole set of segmented cells. This is subsequently divided in train and test set and fed into a classification algorithm.	50
5.1	General structure of the artificial feedforward neural network.	56
5.2	Update procedure for artificial feedforward neural networks.	58
5.3	Two-dimensional discrete convolution. For every new pixel in the convolved image on the right a dot product between the input image (left) and filter (green) are taken: this scalar value is allocated at the central position of the filter. Moving around the filter and letting its center coincides with the whole domain of the image, the new altered image is retrieved.	60
5.4	Convolutional layer in a CNN.	61

5.5	Visual interpretation of the group P4 and P4M. The group P4 comprises all possible combinations of 90-degree rotating actions (thus having dimension 4), while the group P4M includes the possibility of mirror reflections on top of that (thus having dimension 8). The graphs or webs on the right provide the roadmaps on how to move from one patch to the other: a 90-degree rotation in clockwise direction is achieved by following a red arrow and a mirror reflection by following the blue pathway.	64
5.6	First G-convolution.	64
5.7	Two vital operations on group structured signals. Transformation of group structured signals (bottom) and the dot-product between two structured ones (top)	65
5.8	Full G-convolution.	66
5.9	<i>Decomposition</i> of linear feature spaces in convolutional layers. It can be seen as a <i>depth-stack</i> of two-dimensional <i>feature maps</i> (green indicates one such feature or as a <i>bundle of fibers</i> (yellow indicates one such fiber).	69
5.10	The principle of steerable filters	70
5.11	Visual clarification of steerability of the whole domain.	72
5.12	First group-equivariant convolutional stage of the SF-GCNN.	75
6.1	General form of the residual unit, or residual block (fenced by yellow rectangle). A ResNet is the consecutive follow-up of many residual blocks (left). The input parallelly flows through two independent lanes: the <i>operation branch</i> and <i>residual branch</i> . Nearly all learnable weights and thus operations are located in the operation branch while the data has negligible resistance in the other: the residual branch exclusively conveys trivial operations.	78
6.2	Overview of a residual network. The initial layer (red block) conveys an initial convolution that maps between the input field and the field over which the (G-)convolution convolves and subsequently normalizes the batches and activates the remaining signal through an pointwise activation function. The output of the initial layer is fed into the most crucial part of the network which consists of (many) consecutive residual blocks (green and yellow). After every N_{rb} residual blocks, the signal is expanded in terms of channels while simultaneously the domain is downsampled (green blocks). After exiting the last residual layer, the decoder lets the signal converge into a vector from which the predicted classes are extracted. This final element conveys two fully connected layers both directly followed by pointwise activations.	80
6.3	Visualization of the cross entropy loss: infected and uninfected red blood cells.	83
6.4	Overfitting and underfitting in machine learning.	86
7.1	Processing steps in the development of data sets that are comprised of segmented red blood cells.	87
7.2	Several samples from the Rajamaran data set (left) and AiDx data set (right).	88

7.3	<i>Diagnosis</i> confusion matrix for the logical two-class system: positive-negative. P is the amount of positives and N the amount of negatives measured. The subscript indicates whether it is correctly (t) or falsely (f) predicted. In the yellow boxes, 3 distinct performance measures are denoted - (1) <i>sensitivity</i> , the ability of the test to correctly classify if an individual is positive, (2) <i>specificity</i> , a similar measure that does it for an individual that is negative, and (3) <i>accuracy</i> , the overall correctness.	90
7.4	K-fold cross-validation.	91
7.5	Hierarchy, naming and relevant specifications of the proposed convolutional networks. The common name, as earlier specified, of the architecture used is called <i>ResNet-18</i> (comprised of 18 operational stages), which is abbreviated as RN18. In between the type and used architecture, the amount of equidistant rotations is mentioned. Furthermore, an extra "m" reflects the equivariance under mirror reflections.	92
7.6	Performance for increasing heuristic. Recall that $ p4m = c8 = 8$ and $ D8 = C16 = 16$	94
7.7	Learning curves (accuracy) for all different operations. Both training (dotted) and test (line) curves are shown. The best-scoring networks are highlighted in red (SFCNN-16m-RN18, best) and blue (GCNN-4m-RN18, 2nd best). . .	95

2 | The burden of malaria

This chapter should serve as detailed introduction to the research topic by disclosing the need for improved malaria diagnostics. I will do so by first introducing the current burden of malaria. The second part is about the history and current state-of-the-art diagnostic methods and their advantages, limitations and opportunities. Autonomous detection will be highlighted separately as one such an opportunity, solving some of the process' inherent problems simultaneously. Its main features and challenges will be highlighted. Finally, some eye-catching alternative methods including the concept of hyperspectral imagery will be introduced. After reading this chapter, my hope is that you, the reader, have noticed the relevance and importance of both the subject, thus the need for creativity and in-depth research in the development of decent diagnostics.

2.1 Malaria worldwide

Malaria is still a significant contributor of severe illness and death worldwide. Annually, between 350 and 600 thousands people have died and between 200 and 250 million cases have been diagnosed in between 2010 and 2018 according to the World Health Organization (WHO)¹. Furthermore, most of the deaths (90%) and cases (85%) are found in Africa among which the number of deaths is clearly highest for children under the age of 5. Consequently, the mortality is enormous: between 1.7 and 2.4 per thousand of diagnosed within the African area between 2010 and 2018. This shows that there is too many victims still suffering everyday. On the other hand, a vast decrease in the number of deaths and mortality in Africa brings hope and proves that we are on the right track [9, 10].

Moreover, rural areas are those hit hardest. Although explanations of this phenomenon differ among researchers, one argument is repeatedly mentioned: better health facilities in most urban settings is highly correlated with lower morbidity and mortality. Most likely, it means that in many rural areas it lacks well established local implementation of both diagnostic tools, proper treatment and prevention. [11, 12, 13, 14, 15]. Next section lays emphasis on diagnostics specifically and how this can be improved locally as well. Before we move there, I will introduce the life cycle of malaria parasites, the types of parasites that currently exist and how the it expresses itself. This knowledge can help design and build smart diagnostic tools and sophisticated treatment plans that are effective, robust, cheap, less complex and less labour intensive.

¹These are the ones caused by *P. falciparum* which is the main contributing type of the malaria parasites in both the amount of cases and deaths.

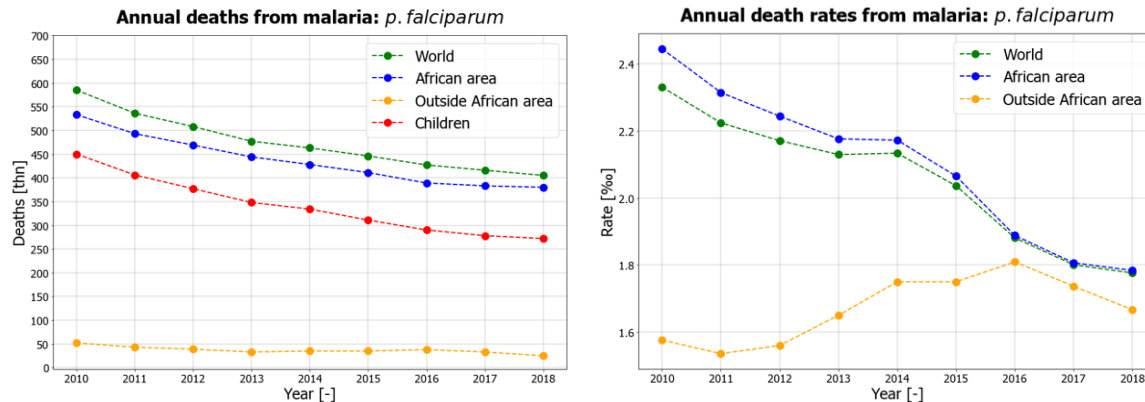


Figure 2.1: Annual deaths (left) and death rates (right) due to *P. falciparum*. Children under the age of 5 are most sensitive. Mortality in Africa is decreasing significantly over the last decade by 35%. Moreover, the areas outside Africa are, in terms of regulating their mortality, doing worse eventhough their absolute number of deaths have decreased as well.

There exist 5 known variants of malaria parasites in order of significance: *Plasmodium falciparum*, *Plasmodium vivax*, *Plasmodium ovale*, *Plasmodium malariae* and *Plasmodium knowlesi*. I will not elaborate on each type individually but one should know that for effective treatment it is beneficial to know what type is present. The molecular structure and spatial appearance enable diagnostic tools to discriminate between these types [16, 17]. The pathogenesis for these types are, at higher level, identical. For every type, the parasite's distributors are, as one should know, mosquitos. A female Anopheles mosquito inoculates the parasites into the blood vessels while performing its regular sting. The *sporozoites*, the stage of the parasites just after injection, are now present in the blood. They first enter the exo-erythrocytic cycle which occurs in the liver which boils down to an alteration of the state of the parasite. The sporozoites mature in *schizonts* that will eventually rupture and release *merozoites*. Most will abandon the liver and will enter the erythrocytic cycle in the blood. In this cycle, the victims are red bloodcells, also named *erythrocytes*. The merozoites can infiltrate the erythrocytes who provide a platform for the merozoites to undergo asexual multiplication. The infiltrated erythrocytes become fully occupied playgrounds which, depending on the occupation level, evolves from immature (also called the ring stage) to mature *trophozoites* and eventually a schizont. The cycle ends with the disruption of the schizont, consequently releasing fresh merozoites in the vessels. Some merozoites will reshape into sexual erythrocytic stages. *Gametocytes* are formed which can be either female or male. If ingested by a mosquito it enters the sporogonic cycle. Female and male gametocyte can form *ookinetes* that evolve into *oocysts*. Similar to the schizonts, oocysts disrupt and release sporozoites. The cycle is complete and the mosquito can now carry the parasite to another person [18, 19]. The cycle is visualized in Fig. 2.2.

Symptoms range between fever, chills, sweats, headaches, nausea, vomiting, body aches and malaise. These are barely specific and are highly overlapping with symptoms from different sources. Consequently, relying on clinical diagnosis only will never be enough. More severe forms of malaria can cause the failure of organs or abnormalities in the patient's blood

or metabolism. Severe malaria manifests in many forms including neurological defects, anemia, hemoglobinuria, respiratory problems, abnormalities in blood coagulation, acute kidney injury, hyperparasitemia, metabolic acidosis, hypoglycemia, nephrotic syndrome, hyperreactive malarial splenomegaly, etc. Many of these are relentless and attribute to immense pain or even death [16, 17]. Precise and reliable diagnostic tools, efficient treatment plans and rigorous prevention are key to help people and especially children in need.




Figure 2.2: Pathogenesis of the infiltration and ingestion of malaria parasites. There is three distinct cycles within: the *exo-erythrocytic*, *erythrocytic* and *sporogonic* cycle. The *exo-erythrocytic* cycle marks the evolution of sporozoites into merozoites and primarily takes place in the liver, the *erythrocytic* cycle resembles the infiltration of erythrocytes and formation of sexual gametocytes within the blood vessels and the *sporogonic* cycle remains for the production of fresh sporozoites which takes place in the mosquito. [5].

2.2 Diagnostics

Diagnostics is the discipline or practice of diagnosis, that serves for identification or characterization of the object or process searched for. In the case of malaria, proper diagnostics identify whether or not a person is infected in the first place, what type of malaria parasite it is infected with and what the stage of that infection is. Throughout the last 150 years, many of such tools have been developed, maintained and revised, all having their advantages and limitations. The type of tool used in healthcare facilities depends specifically on the availability of educated personnel, accessibility of equipment, costs and regulations [20, 21, 22]. Consequently, the type of diagnostics is often depending on the location it is used. Especially, as mentioned earlier, providing decent diagnostic tools in rural areas is essential. In my opinion, the following list resembles the most important criteria in developing proper diagnostics: accuracy of test, ability of species discrimination, quantification of parasitaemia, expenditures, labour intensiveness, complexity of test, local implementation possibilities, ro-

bustness and test duration. In the next section, I will take some time to evaluate the most prominent and promising methods based on those criteria. Some explicit quality measures should be explained that can help the assessment. *Sensitivity* is the probability of a positive sample to test positive, *specificity* the probability of a negative sample to test negative and the *limit of detection*, or LoD, is the detectable parasite threshold. Consequently, high sensitivity and specificity provide low numbers of misdiagnosis and a low LoD enables early stage detection. Therefore, diagnostics with a high sensitivity, high specificity and low LoD are wanted. Let us focus on the diagnostic methods which can be categorized as shown in Fig. 2.3.

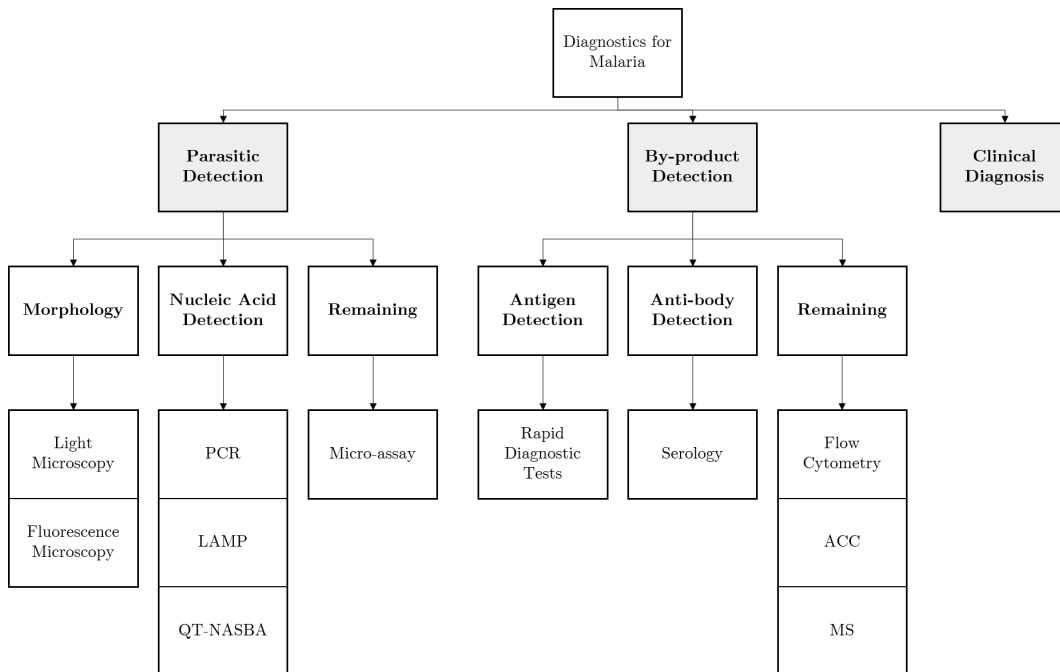


Figure 2.3: Categorization diagram for malaria diagnostic methods. On a high-level, three distinct categories exist. *Clinical diagnosis* is merely diagnosing on the basis of present symptoms of the patient, *parasitic detection* is focussed on detecting the parasite or parts of it and *by-product detection* is based on detecting either products that the parasites produce as well as products the patients themselves produce. Parasitic detection’s sub-categories include morphological and nucleic acid based detection methods and by-product detection include antigen and anti-body based methods. Some specified methods are not discussed in my research such as micro-assay, serology, flow cytometry, ACC and MS.

Morphology

Light microscopy, LM, has been the most widely accepted diagnostic tool worldwide. The diagnosis relies on the examination of blood smears under a microscope, magnifying and thus increasing the morphological detail we can see. Often, staining is used, allowing for visual distinction, consequently making the parasites traceable and their structure observable. The technique is established and versatile, enables for detection of different species, additively provides quantification of parasitaemia, has high sensitivity and specificity and has relatively

low expenditures per test. On the other hand, well educated personnel is needed, tests are subjective, starting costs are high, sample preparation is time consuming and the method is highly complex. Although the method is among the most reliable and robust, disadvantages often obstruct the usability in remote areas. [23, 24]. In laboratory settings, *fluorescence microscopy*, or FM, has shown to be an alternative morphological method that is both rapid and sensitive. It was developed to simplify and enhance conventional LM and partly succeeded. The technique uses fluorescent dyes that bind to the parasite's DNA which, under an epi-fluorescent microscope, creates sufficient contrast, allowing for easy identification of parasites. Eventhough some argue it should, in endemic areas, be the preferred diagnostic method, there is some disadvantages. It is unable to reliably discriminate between species, it requires specialized instrumentation and educated personnel, the method is complex and the costs are high. In addition, quality control is nearly impossible since the stains fade over time [25, 26, 27, 28, 29, 30].

Nucleic Acid Detection

Since the 90s, molecular diagnostic techniques were developed that amplify genetic material of the parasites. Nowadays, *polymer chain reaction*, or PCR, is the reference standard as a tool for confirmation in well-developed healthcare facilities. Highest scores for sensitivity and specificity at low levels of parasitaemia and the ability to discriminate between all species makes the technique exclusive. On the contrary, it is not routinely implemented in developing countries since it is lacking standardization, in need of highly educated personnel, time consuming, costly and is in need of a constant power supply [31, 32, 33, 34]. There are extensions of PCR that achieved similar results while being less technically complex, significantly reducing the test duration and are more cost-effective. *Loop mediated isothermal amplification*, or LAMP, does amplify DNA with high specificity, efficiency and rapidity under isothermal conditions. This method is simpler and easier to perform [35, 36, 37, 38]. *Quantitative nucleic acid sequence-based amplification*, QT-NASBA, shows similar results as LAMP but has the added feature of distinguishing sexual and asexual stages of the parasite, therefore contributing to the knowledge in which phase the patient is in [39, 40].

Rapid Diagnostics Tests

Above-described techniques are time-consuming and generally not locally implementable. As argued, this is a problem for people living in the periphery. The problem of accessibility of diagnosis for everyone has been addressed with the emerging of *rapid diagnostics tests*, RDTs, that are all based on the same principle, the detection of malaria antigen. These methods are simple, have no starting costs, are easy to perform, are in no need of accomplished doctors and are easily transported everywhere. The drawbacks are being unreliable in species discrimination, not being able to quantify the parasitaemia and a high LoD. Furthermore, researchers disagree about the sensitivity and specificity ranging between 60% and 99% [41, 42, 43].

Evidently, I showed that they exist, on the one hand, diagnostic methods such as LM or PCR, with high performance thus being able to discriminate between species, quantify the parasitaemia and are both sensitive and specific for low LoDs while, on the other hand, RDTs that show a tremendous decrease in duration of test and ability for implementation in low-resource settings. It looks that if both 'sides' lack each others benefits. The concept

of *smart diagnostics* tries to engage that challenge, that is developing diagnostics that have high performance and being able to distribute and implement it at any location by being *smart*. In LM, one of these breakthroughs is the concept of autonomous detection in which an algorithm determines the results which has the perspective of solving some of its limitations by decreasing complexity, labour intensiveness and duration of test while maintaining or even increasing performance.

2.3 Autonomous detection

In this section, I will state the importance of interplay between the human brain and the digital brain, also known as the computer, in developing smart diagnostics. As mentioned in [20], accurately testing drug-resistance, measuring drug-effectiveness and identifying the disease's stage are essential in addition to merely the diagnosis of malaria. Therefore, it becomes clear-cut that non-subjective and standardized methods that accurately count parasites are called upon. As shown, the outcome and applicability of LM depends on the experience and competence of the personnel. In the detection stage, a computer's objectiveness can help improve this method examination: a computer that, on the basis of algorithms, determines the properties interested in. This is what we call *autonomous detection*. Over the last decade, literature reveals its ability to flourish in the near future by many convincing and promising results. Its main advantages could be its reduction of workload and costs, and its ability to provide more reliable and standardized interpretation of blood films [44, 45]. Although it seems very promising, no clear indicator can (still) be subtracted that determines the current state-of-the-art. This can be attributed to small data sets and incomplete evaluations. Please have a look at the process of conventional LM in which the digital *doctor's assistant* comes into play (Fig. 2.4).

The second part of the report dives deeper in making the artificial detection machines smarter and more robust. Autonomous detection will in future times be able to solve some of the present problems at hand but there is, elsewhere in the process, room for improvement as well. Next section elaborates on some interesting alternative practices among which is the concept of hyperspectral imaging.

2.4 Opportunities

A genuinely brilliant idea was the adoption of *smartphones* in the process. They are equipped with excellent optical and sensory hardware compared to their price. Moreover, even in remote areas in Africa, people have up-to-date smartphones that could subsequently, besides normal use, be used as part of the diagnostics device. This idea has been exploited over the last years and seems auspicious [46, 47, 48, 49]. It allows for small embodiments, ease of use and reduction in starting expenditure, thus providing advantages of the use in rural areas over other devices. Its major limitation is the spatial resolution and the observable field but pixel sizes are getting smaller and sensor surfaces bigger as we speak.

Another interesting alternative technique is *mass spectrometry*. The diagnosis grants detection of *heme* molecules that, during the growth of parasites within the red blood cells, were

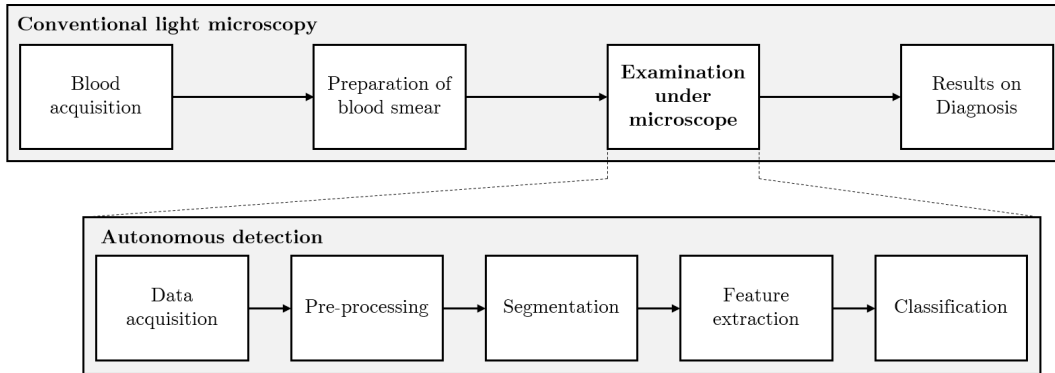


Figure 2.4: Overview of conventional steps taken in light microscopy and specifically its examination which in this case is autonomous. The high-level steps are self-explanatory. Five subsequent steps should be taken in autonomous detection. *Data acquisition* constitutes the retrieval of digital images. *Pre-processing* transforms (e.g. cleans, normalizes) the raw images such that it gets denoised, becomes relevant and its redundant information thrown away. *Segmentation* partitioned the image into multiple objects, which, in the case of malaria, is the segmentation of individual erythrocytes. Arguably the hardest part is the determination of features that are used to discriminate between infected cells, determine the species and infiltration stage. This process is called *feature selection*. *Classification* is the final stage and classifies, based on the features, individual cells.

sequestered into molecular crystals (hemozoin). During the process of infiltration, the heme molecules detach from the hemoglobin, consequently creating a volume of highly concentrated and purified biomarker. Hemozoin (also called *malaria pigment*) strongly absorbs ultraviolet light contributing to the vaporization and ionization of individual heme molecules. Laser desorption mass spectrometry, LSDM, detects these isolated heme molecules. The method provides high throughput, is among the most rapid, is automated and has a low LoD. On the other hand, it is expensive and highly complex [45, 50].

LSDM shows tremendous decrease in sample preparation, which could become crucial in the design of smart diagnostics. High performance methods have prolonged sample preparation, e.g. because amplification (for PCR) and the process of staining (for LM) are time consuming as the inherent chemical reactions are time-consuming. The imaging modality that is introduced next, which is the one exploited in this research, takes advantage of the fact that parasitic material and other constituents in the blood are inherently different by nature. This fact should enable discrimination by only looking at it if sufficient detail is provided. Providing sufficient detail is hard though. For example, as humans we only take a combination of three spectral receptors that determine the final color². Computer sensors observe in the same fashion denoting the individual spectral sensors as R, G and B being most sensitive, for as we see it, the colors red, green and blue respectively. Even though this provides proper detail in the colors we can discriminate, it can be tremendously boosted by having more of these sensors. *Hyperspectral imaging* can provide many (more) of these sensors (think of 100 instead of 3). It is our hope that providing enough spectral detail

²It is obviously more complicated. A detailed explanation is described in 3.2

one could enable accurate diagnostics that are similarly used as LM without the need for staining. Theoretically, the concept is extraordinary and could improve much on duration of test. Furthermore, chemical waste and labour intensiveness could be reduced significantly.

2.5 Research question and proposal

2.5.1 Hyperspectral imaging as novel imaging modality in the diagnosis of malaria

Hyperspectral imaging (HI) is an emerging modality in many booming domains, among which are surveillance, food quality control, agriculture and forensics [51, 52]. Moreover, it starts to mushroom in many medical applications, of which mostly developments lie within improved diagnostics [53]. Medical HI (MHI) offers a great potential for non-invasive investigation of tissue as it retrieves inherently more precise spectral responses. It is (mostly) assumed that during the disease’s progression, absorption, scattering and fluorescence characteristics change. These characteristics are possibly faster detected and/or discovered if MHI is used instead of conventional (RGB or monochrome) vision.

For above-mentioned reasons, MHI could potentially be deployed as label-free detection technique in the diagnosis of malaria. Albeit theoretically a profound direction, there are many technicalities that need to be looked at and solved, among which are the machine’s cost, acquisition (process and time), aberration-corrective capacity, spectral purity, calibration and many more. In the context of malaria, the setup’s cost is arguably the most important as standardized HI systems are very expensive. The costs for a complete microscope-integrated HI system (standardized) usually range between 45.000€ and several hundred thousands [54]. In order to deploy this modality for improved malaria diagnostics, these numbers have to drastically decrease. In doing so, many trade-offs and smart element choices must be done in order to address the other mentioned technicalities as well. The first part of my research is about the exploration of this technique. Consequently, it tries to answer the below-mentioned research question:

Could an **optimized off-the-shelf, low-cost microscope-integrated hyperspectral imaging system** be used as novel **non-invasive diagnostics tool** for improved detection of malaria parasites?

Part I addresses this question and lays emphasis on its needed physical background in Ch. 3 and the design, building and evaluation in Ch. 4. Unfortunately, time and data constraints did not allow to create proper hyperspectral imaging databases of many unlabeled samples. The idea was to use these databases as training and test sets for enhanced and generalized computer algorithms that could learn themselves how to discriminate parasitized from non-infected red blood cells. Although these databases were not created, the design and building of smart neural networks capture the essence of the second part of the research.

2.5.2 Towards *smarter* classifiers in the diagnosis of malaria

In poor countries, classification of blood films is usually done manually. This could be significantly enriched by the use of smart computer algorithms that do the job digitally. These computer algorithms could objectify and fasten the process. Furthermore, it could reduce cost as less-educated personnel could do the job accordingly. The second part of my research is about the exploitation of the operations in deep neural networks and how, if smartly designed, they could benefit overall performance metrics, such as accuracy, sensitivity and specificity. It moreover lays emphasis on the robustness of these models compared to more conventional operations, in this case, convolutions. It is significant as types of retrieving data and blood films themselves will differ depending on the type of setup, its configuration, the person's blood and the way the blood films are prepared. Part II of the research focuses on answering the following research question:

How must the **operational stages in deep convolutional networks be manipulated** such that **smaller-sized and more robust models** are found that improve on the state-of-the-art in malaria classification?

Ch. 5 provides the necessary background on how to create smarter operations in convolutional neural networks. The essence of Ch. 6 is to make the reader accustomed with the design, training and implementation of deep learning networks. Evaluation of several equally-sized networks is done in Ch. 7 in which the models are trained on a well-known online available malaria database. Robustness of the pre-trained models is checked on a more specific database. Ch. 8 finalizes with an in-depth discussion and conclusion.

Fig. 2.5 gives an overview on how to place each part and its chapters in the process of diagnosis.

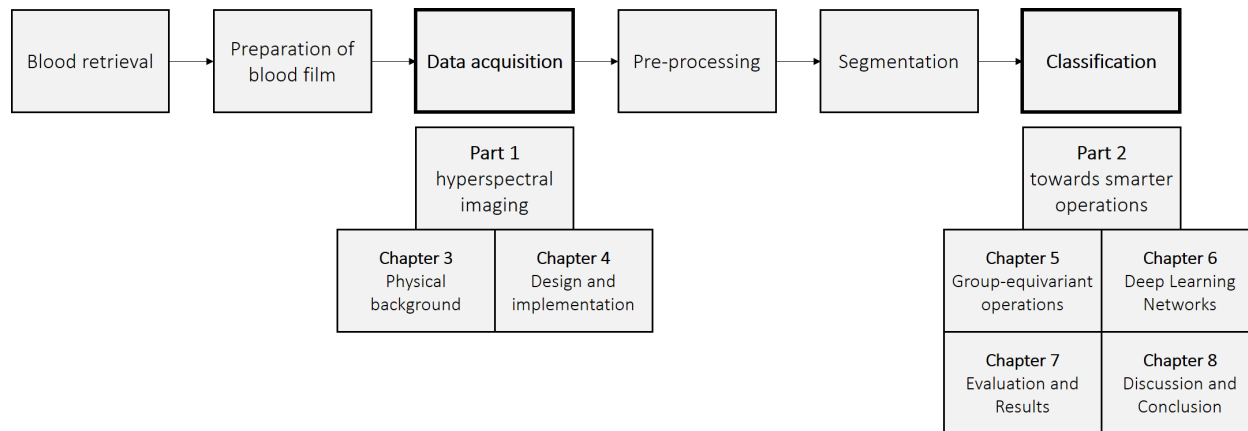


Figure 2.5

Part I

Hyperspectral Imaging

This part is dedicated to the concept of *hyperspectral imaging* (HI) and my implementation of a specific one. HI is an imaging modality that additionally retrieves spectral information on top of what a conventional microscopy retrieves, namely the 2D morphology. For every spatial location (pixel), HI tries conducts the individual spectral responses. Where standard 2D imagery provides information on either one (monochrome) or maximally three spectral bands (RGB), HI provides information among many spectral bands. In doing so, HI could enhance the search for specific constituents, identification of materials and help uncover processes. Although processing and analyzing the tremendously large data sets is computationally excessive, HIs versatility, availability of contemporary fast processors, and the rise of data science, high-resolution (both spatial and spectral) HI becomes more popular. It plays a vital role in monitoring of farmland, food processing, surveillance, mineralogy etc. Especially the fact that HI enables accurate distinction between closely related materials by means of their spectral response, made this concept interesting for label-free detection of malaria. The part is comprised of 2 chapter in which the first (Ch. 3) provides the necessary theoretical background on light and optics. Ch. 4 closes with the design, building and evaluation of an off-the-shelf, low-cost hyperspectral imaging system.

3 | Physical background

3.1 Light and optics

Optics is the study of light and how it interacts with other materials. This branch of physics has contributed to many applications, including the developments in imaging. Among the first breakthroughs in that domain, microscopes and telescopes enabled the observation of detail at scales far from perceptible with the human eye alone. This section focuses on the theory of light and specializes in Koehler aligned microscopes.

3.1.1 Geometrical optics

In geometrical optics, light is assumed a *particle*. It travels in straight lines and changes direction if it hits upon a surface (*reflection*) or is entering a new type of medium (*refraction*). The incident angle of the light path, the angle upon which it interacts with the surface or medium, determines how it reflects or refracts and is, for reflection, given by

$$\theta_1 = \theta_2 \tag{3.1}$$

and is, for refraction, known as *Snell's law*, given by

$$\frac{\sin \theta_1}{\sin \theta_2} = \frac{n_2}{n_1} \tag{3.2}$$

in which θ_1 and θ_2 are the incident and outgoing angles of the light path, both measured w.r.t. the normal of the surface at which it enters (see Vol 1, Ch. 26 in [55] and see Fig. 3.1). This knowledge was used in the development of many optical elements, such as the lens. The material a lens is build of (mostly some type of glass) has a high refractive index ($n = 1.50$, air - $n = 1$). Consequently, upon entering the medium, a light path changes direction according to Snell's law. The lens' shape is developed in such a way that it, approximately, obeys the following three pillars: (1) all particles that travel in parallel to the optical axis are refracted such that they cross the *back focal point*, (2) all particle that cross the lens at the optical axis have unchanged direction, and (3) particles that have equivalent incident angles meet at the *back focal plane* (for clarification of the pillars and terms see Fig. 3.1). The discovery of the lens was astonishing because it facilitated the ability to create *images* by seizing the opportunity to redirect light coming from an object such that it gets focused elsewhere. In *paraxial ray approximation*, in which is assumed that rays are only travelling under small angles ([56]), the distance between the image and lens, s_i , is retrieved as

$$\frac{1}{f} = \frac{1}{s_o} + \frac{1}{s_i} \text{ (thin lens equation)} \longrightarrow s_i = \frac{s_o f}{s_o - f} \quad (3.3)$$

in which f is the focal length of the lens and s_o the distance between lens and object (see Vol 1, Ch. 26 in [55]). The planes that comprise object and image are referred to as the object and image plane respectively.

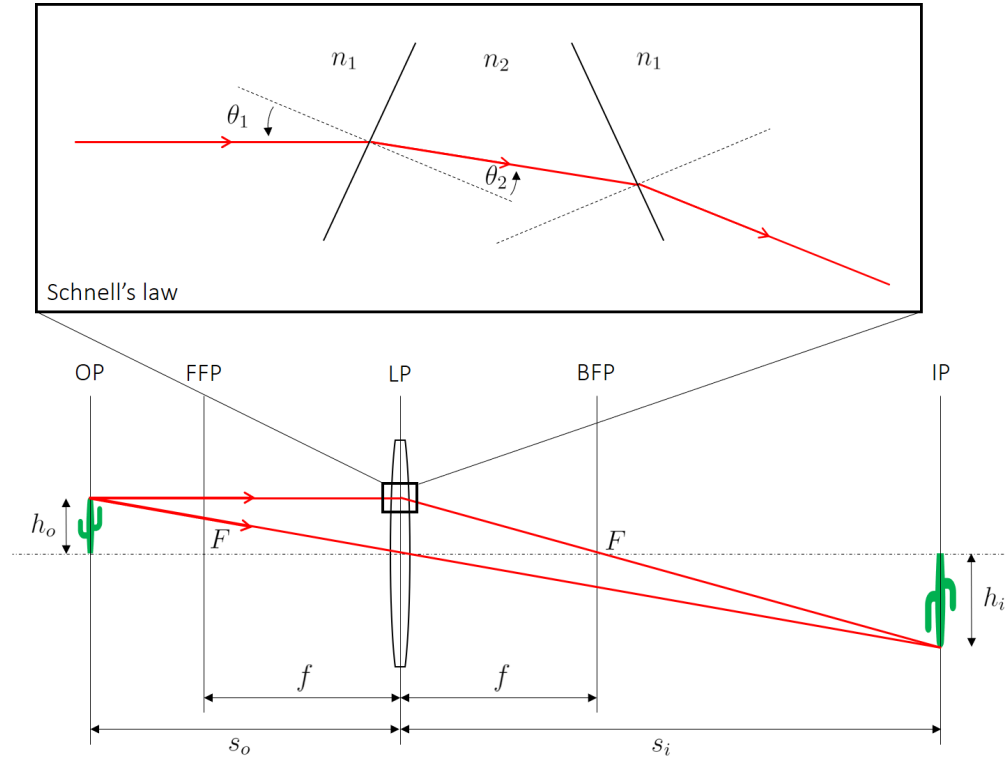


Figure 3.1: Geometrical optics: Schnell's law and the thin lens.

Geometrical optics allows us as well to grasp the concept of magnification (see Vol 1, Ch. 26 in [55]): the transversal magnification, M_T , - the ratio between the height of the object, h_o , and image, h_i - is equal to the ratio of both distances, that is

$$M_T := \frac{h_i}{h_o} = \frac{s_i}{s_o}. \quad (3.4)$$

Magnification is the main contributor to observe besides what the naked eye is possible to perceive, and therefore a key property in all microscopes.

3.1.2 Physical optics: numerical aperture and spatial resolution

Geometrical optics is not able to explain all phenomena in optics, e.g. interference, the way sunglasses work (polarization), resolution, etc. Many researchers in the domain have to be accustomed with the theory, or part of it, that explains these phenomena: *physical optics*. Physical optics treats light as waves, in which the electric and magnetic field develop in the direction of propagation (same as the arrows in Fig. 3.1), moving away from the source, but

having their forces pointing orthogonally to that direction (and orthogonally to each other as well). In time and space, both forces oscillate in sinusoidal fashion around the origin (zero). As a consequence of this ever changing field, depending on the location and time the light's intensity is measured¹, the output might differ. In physical optics, the field as function of time and space determines the intensity measured. In theory, it could be that a person, being encircled with point sources that emit light in all directions, does not perceive or observe any of the light as the sum of all fields is exactly nothing at the location he or she is at. This is in a nutshell the concept of interference in which the resultant field, the superposition of two or more waves, is either cancelled (destructive interference; the example), reinforced (constructive interference) or somewhere in between.

Another phenomenon in optical physics, which highly depends on interference or, as Richard Feynman finds, is exactly the same, is that of *diffraction* (see Vol 1, Ch. 30 in [55]). It occurs when light waves hit upon an opening or bend around a corner. The diffracting object, which, in the case of a lens, is its aperture, effectively becomes a secondary source. Due to the interference of waves, a collimated flat wave entering an ideal lens will not be focused exactly in the back focal plane, but its intensity instead be spread out around that focus point. The intensity map, or the out-blurred point, depends on the the size and shape of the aperture, which for a circular aperture is referred to as the *Airy Disk*. The variation in intensity for a given angle θ of the Airy Disk is computed as

$$I(\theta) = I_0 \left(\frac{2J_1(kD \sin \theta)}{kD \sin \theta} \right)^2 \quad (3.5)$$

in which $k = 2\pi/\lambda$, D the radius of the circular aperture and J_1 a bessel function [57]. The equation tells us that the smaller the size of the aperture, the larger the spot size. Let it sink in that diffraction happens for every point that is focused, creating possible overlap between Airy Disks arising from different points. If the Airy Disks of two near-separated points in the image plane have "too" much overlap, they become indistinguishable. The distance between two just-resolvable points is the optimal spatial *resolution*. The Rayleigh criterion specifies resolution as the point at which the first-order minimum of the first point coincides with the maximum (which is the focus point) of the second point. For a circular aperture, this translates to

$$d_\lambda = 1.22 \frac{\lambda}{n \sin \theta_{out}} \quad (3.6)$$

in which θ_{out} is the half angle of the cone that emanates the lens (see Vol 1, Ch. 30 in [55]). The denominator in the given fraction is a wide-used term in optics, called the *numerical aperture*, or NA [58]. Rewritten in terms of the NA, the spatial resolution for a single lens is given by

$$d_\lambda = 1.22 \frac{\lambda}{NA}. \quad (3.7)$$

1

Consequently, optical elements that have relatively large NAs achieve higher resolution. Practically, such elements, e.g. objectives or condenser lenses, having high NAs are harder to manufacture as their complexities generally increase, hence making them more expensive and often larger [58].

3.1.3 Bright-field microscopy and koehler illumination

Microscopy is the science of the investigation of micro-scaled objects and processes. A *microscope* is the apparatus that provides us an enlarged view of these tiny constituents by a smart interplay of (mostly many) optical elements. Microscopes are either designed to create a magnified image on our eye’s retina or on the sensor of a camera. There are many types of microscopes that, for example, differ in the way they illuminate their samples e.g. bright-field, dark-field, and phase contrast microscopy (for more details see *Light and Video Microscopy* by *Wayne* [59]). This section focuses on bright-field microscopy that is aligned in an optimal manner (mostly paraphrased from [59]).

One of the most important and fundamental techniques in light microscopy was published in 1893 by August Koehler [60]. His technique for improved illumination of the specimen, later named as Koehler illumination, is still implemented in most of the conventional light microscopes. There are three main features of a proper Koehler aligned microscope that are beneficial in optimal imaging: improved resolution, uniform concentrated sample illumination and easy interplay between contrast and resolution.

The optimal resolution d_λ , using the Rayleigh criterion, for a light microscope is computed as

$$d_\lambda = \frac{1.22\lambda}{\text{NA}_{obj} + \text{NA}_{cond}} \quad (3.8)$$

in which NA_{obj} and NA_{cond} are the NA of the objective and the condenser side respectively [61]. As mentioned earlier, improved overall resolution is realized with increased NA, which in this case holds for both the condenser and the objective side. As we will see, Koehler illumination significantly increases the NA_{cond} in most cases, therefor increasing the resolution. In the following, its general optical scheme is explained on the basis of which the main features will be proved.

The illumination path consists of several optical elements: light source, collector lens, field diaphragm (iris), field lens, condenser diaphragm (aperture) and condenser lens. Fig. elaborates on the exact placement of these elements. Taking a look at the source’s light rays residing from 3 distinct points within the filament, one observes that the points are all focused in the front focal plane of the condenser lens at different locations, resulting in differently angled collimated beams (beam pencils). Furthermore, the intersection of the beam pencils is in the back focal plane of the condenser lens, which by design is the location of the sample. This will hold for every point in the source. This has two direct implications: the sample is uniformly illuminated and all the light is condensed, being concentrated in a small

area centered around the optical axis ². The size of the beam pencils depend on the ratio between the field and condenser lens, which we call the condenser magnification, M_{cond} ,

$$M_{cond} = \frac{F_{cond}}{F_{field}} \quad (3.9)$$

in which F_{cond} and F_{field} are the focal lengths of the condenser and field lens respectively. Usually, the aim is to have a large demagnification ($M_{cond} \ll 1 \rightarrow F_{cond} \ll F_{field}$), to focus all the light on a tiny spot for which its diameter should ideally align with the FOV.

As mentioned, increasing the NA at the condenser side results in better overall resolution. Koehler illumination does so by smartly designing the ratio between the focal lengths of the field and collector lens and the use of a powerful condenser lens (F_{cond} small). This ratio between field and collector lens is called the field magnification, M_{field} ,

$$M_{field} = \frac{F_{field}}{F_{col}} \quad (3.10)$$

in which F_{col} is the focal length of the collector lens. In an optimal configuration, the height and width of the image of the source in the FFP of the condenser should equal the size of the condenser:

$$\hat{D}_{source}^{(i)} = \hat{M}_{field} D_{source} = D_{cond} \rightarrow \hat{M}_{field} = \frac{D_{cond}}{D_{source}} \quad (3.11)$$

in which D_{cond} is the diameter of the condenser lens, D_{source} is the diameter of the source and $\hat{D}_{source}^{(i)}$ is the optimal diameter of the source's image. The configuration is optimal in the sense that the condenser's limited NA is reached without wasting optical power. Moreover, it should be known that the NA is limited as

$$NA_{cond} \leq 2 \arctan \left(\frac{D_{cond}}{2F_{cond}} \right). \quad (3.12)$$

Practically, it often means in a microscopic setting that high-NA condenser lenses are generally small to ensure the possibility of a very small focal length that lives outside the interior of that lens. In a non-ideal setting, in which the image of the source in the FFP of the condenser is (too) small, the NA is computed as,

$$NA_{cond} = 2 \arctan \left(\frac{M_{field} D_{source}}{2F_{cond}} \right). \quad (3.13)$$

Since the conjugated focal planes of the sample and field iris reside at the same locations along the optical axis, the field iris' image can be seen in the sample plane. It should be noted that closing and opening the field iris does alter uniformity of illumination. The diameter of the iris directly determines the (circular) region that is illuminated on the sample. Often, this region ideally aligns with the FOV at the sample side. Given the diameter of the FOV at the sample side, $D_{FOV,sample}$, the optimal diameter of the field iris is computed as,

²In ?? it will be seen that depending on the FOV at the sample side and the diameter of the source, one has to put constraints on the minimal diameter of the field lens to assure "perfect" uniform illumination

$$\hat{D}_{field} = \frac{D_{FOV,sample}}{M_{field}}. \quad (3.14)$$

Lastly, changing the condenser aperture directly changes resolution and contrast: resolution is enhanced with opening and contrast is enhanced with closing the aperture. This easy interplay between resolution and contrast is highly beneficial in many microscopic applications and the exact settings depend on what is evaluated. In terms of the field magnification, the diameter of the source and the numerical aperture of the objective side, the resolution is computed as,

$$d_\lambda = \frac{1.22\lambda}{NA_{obj} + 2 \arctan \left(M_{field} D_{source} (2F_{cond})^{-1} \right)} \quad (3.15)$$

3.2 Spectral characteristics

By looking, feeling, smelling, and tasting, we, as humans, can discriminate between many different objects, materials, and processes. We use our eyes to distinguish lions from cheetahs, use our hands to feel different textures and temperatures, use our noses to discriminate soy from fish sauce and use our taste buds to distinguish between different types of apples. In this research, I am interested in that first mentioned sensor: the human eye and the capability to see. *How is it possible that we can discriminate between millions of different objects?* Let us have a closer look at the eye: its main components are the pupil and the retina. The pupil is a lens that focuses light on your retina which captures that light and the retina is a layer of sensory neurons which consist mainly of photo-receptors that respond to light. These receptors generate electrical pulses that via the optic nerve connect between your eye and brain where these signals are being processed. There is 4 type of photo-receptive cells and they are placed everywhere on the retina: 3 type of cone cells and one type rod cell. Different cone cells react differently to *colors* and rod cells are sensitive to the *amount of light*. More scientifically, cone cells are wavelength sensitive and thus allow the perception of color and rod cells are intensity sensitive thus allowing the perception of brightness. Moreover, since the combination of brightness and color are measured at different locations on the retina corresponding to different locations in space, our brain is able to reconstruct space as it does. It is that specific combination, the ability of perceiving spectral responses for many locations simultaneously, that enables classification of millions of objects, observing different environments and discriminate human postures. It has enabled humans to discriminate between ripe and rotten fruits, between shallow and deep waters, horses and zebras, and calm and vibrant places. Although the capacity of *seeing* is extraordinary, it clearly has its limits. These limits are best quantified in terms of the spatial and spectral resolution, which are measures of detail. Spatial resolution is defined as the closest distance at which two points can be distinguished and spectral resolution as the closest distance between two distinguishable pure colors. While spectral resolution can be enhanced by having more cone cells, spatial resolution could, theoretically though, be enhanced by increasing the area density of the cells. While the latter does not substantially differ among people, the former however is. By far most people have three types of cone cells (*trichromacy*), but some have

a lower number: 2 for *dichromacy* or "color blindness" (4 - 6% of the population), 1 for *monochromacy* (1 in 25.000 - 35.000 people) and 0 for total blindness (0.5 - 0.6%) [62]. On the other hand, there is a significant part of humanity blessed with having an increased number of cells, namely 4 (*tetrachromacy*). According to *Jay Neitz*, PhD, an ophthalmology professor at the University of Washington, trichromats are able to discriminate between 1 million colors, while tetrachromats a stunning 100 million [63]. Further increasing the number of cells or, in digital imaging, the number of channels, allows for improved vision of color. Hyperspectral imaging is exactly using that property. The following subsections focus on the interactions between light and matter which form the foundation of why there is the ability to distinguish non-self-emitting objects in the first place.

3.2.1 Light-matter interactions

Most of this subsection is paraphrased from *QED: The Strange Theory of Light and Matter* by *Richard Feynman* [64]. If a particle obstructs the light path, light interacts with it. Light-matter interactions are categorized in 4 major levels: transmission, absorption, scattering and emission. Transmission is the pass-through of light, as if it does not see the particle, absorption is the process in which light is engrossed in the particle, scattering is the return of light in a different direction, and finally, emission is the spontaneous release of some of the particle's energy in the form of radiation. Transmission, absorption and reflection only occur if incident light, hitting the matter, is present while emission solely occurs based on the particle's interior status. Fig. 3.2 depicts all four.




Figure 3.2: Light-matter interactions in general. From left to right: absorption, (in-)elastic scattering, transmission and emission.

At the atomic level, electrons "live" at discrete quantum mechanical *energy levels*. An atom is able to interact with a photon that has the energy that either equals or surpasses the gap between such energy levels. If it does, the atom is said to "absorb" the photon while moving up the electron to a higher energy level. On the other hand, atoms can as well emit photons while simultaneously letting a or a multiple of electrons fall down to a lower energy level. These two mechanisms that take place at the atomic level are the main building bricks to all mentioned light-matter interactions.

At the larger, molecular, scale, absorption manifests in the conversion of photon energy

into thermal motion or molecular vibrations. For example, absorption can occur if the frequency of the wave equals the energy of the free vibrations of an elastic molecular structure. Degree of absorption deeply correlates with the electronic structure of the molecules (and atoms), wavelength of the incident light, tissue thickness and temperature. On the contrary, scattering occurs at frequencies that are clearly different than those of the natural vibrations of the molecules. There is two distinct forms of scattering, one in which the energy of the scattered photon equals that of the incident photon (elastic), and one which bounces back a photon having a different frequency (inelastic). Elastic scattering happens more often: it is, in most cases, a million times more likely to occur.

What interactions, or mechanisms, occur in the applications one is investigating can be beneficial. At least having the knowledge of what types of interactions are dominant can be of great use in understanding and explaining certain observations. In the context of microscopy, it can help in the design of a proper optical scheme. In the next, I will discuss those interaction that are predominant in biological tissues, specifically thin blood smears.

3.2.2 Light-tissue interactions in blood

In bright-field microscopy, the sample, sometimes referred to as the field, is usually illuminated from beneath. Only that portion of light that propagates, in one way or the other, through the sample can be observed at the image plane. Consequently, photons that undergo specular reflection at the boundaries of the glass slide(s) are not present at the sensor. Furthermore, in the case of malaria infected thin blood smears, spontaneous emission is very weak and thus negligible as well [6].

Whole blood consists of many components, such as red blood cells, white blood cells, platelets, and blood plasma. The haematocrit, the volume percentage of red blood cells in whole blood, in healthy adult blood specifically is between 40 and 45 % depending on gender (40 in case of women and 45 in case of men) (see Background section in [6]). Including the specific structure they obey, and the components they are build up from make them prime targets for light to interact with in especially the visible and near-infrared range. Practically, all those interactions are due to haemoglobin. Haemoglobin, the protein responsible for the transport of oxygen and carbon monoxide, alternates between different appearances depending on what it carries. Oxy-haemoglobin (bound to oxygen), HbO_2 , and de-oxy-haemoglobin (unbound to oxygen), Hb , are most abundant among all types [6]. Oxygen saturation, $S(\text{O}_2)$, is the volume percentage of oxy-haemoglobin compared to the total amount of haemoglobin, thus computed as

$$S(\text{O}_2) = \frac{[\text{HbO}_2]}{[\text{HbO}_2] + [\text{Hb}]} \quad (3.16)$$

On the bottom of Fig. 3.3, HbO_2 and Hb the absorption and (elastic) scattering coefficients are plotted as those interactions most predominant. It is observed that, throughout the applied spectral range, one is not significantly inferior or dominant over the other for both. A natural question arises: "How do these interactions attribute to what is observed, thus seen and captures by the camera? Moreover, how should they be interpreted?". First of all,

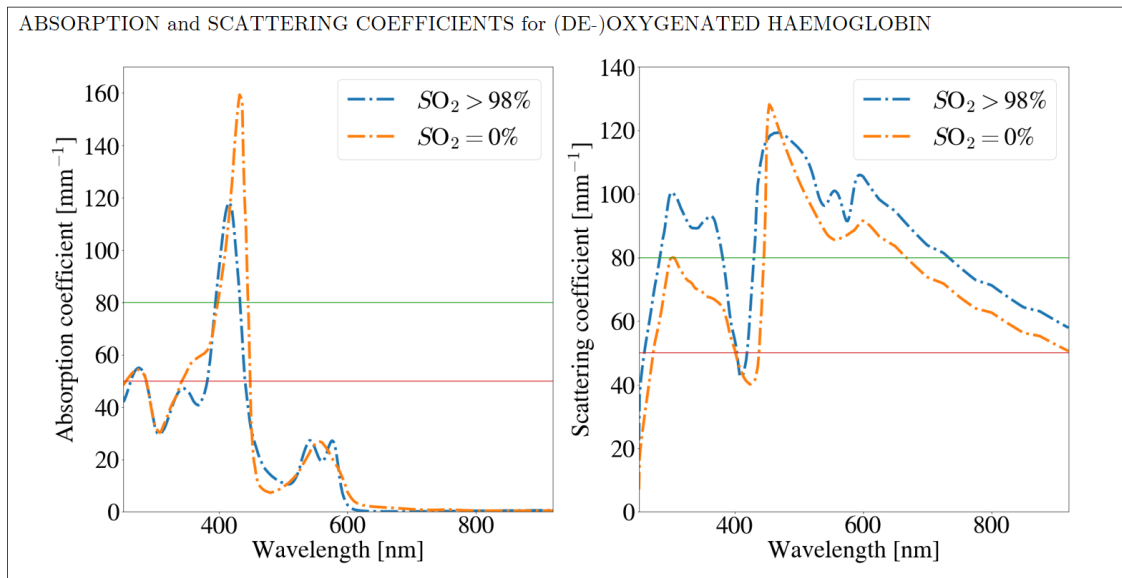



Figure 3.3: Light-tissue interactions in whole blood. At the top, the predominant interactions are sketched: transmission, scattering and absorption. The middle elaborates on both type of scattering that occurs (left) as well as the main component of erythrocytes, namely haemoglobin (right). The bottom shows the absorption and scattering coefficients for HbO₂ and Hb. Here, the resolution is 2nm until 600nm and 5nm upwards the spectrum. All data is directly copied from the tabulated data in the Appendix of [6].

it should be mentioned and noted that it are those interactions in the first place that make it possible to (visually) observe the specimen. In a proper microscopic setup, the sample is homogeneously illuminated with light coming from many angles. This means that if the sample stage is uninhabited we will observe a uniform bright field. If on the other hand tissue is present, some light is absorbed or scattered in which the amount depends on the exact substances, thus molecular structures, it comes across. Logically, this depends on the summation of all (molar) absorption and scattering coefficients of those individual molecules. In bright-field (transmission) microscopy, one consequently observes simultaneous differences in absorption and scattering of the whole scene. More specifically, it is the absorption that eventually provides the differences in the morphological and spectral landscapes. Scattering slightly diffuses the propagation of light since, even for thin blood smears, it usually scatters more than once (multiple scattering).

3.2.3 Dispersion and chromatic aberration

Strictly speaking, the refractive index for most materials is wavelength dependent. Different colors therefore refract differently upon entering a new medium (see Schnell's law, Eq. 3.2, and add $n = n(\lambda)$). Technically, this phenomenon is known as *dispersion*, as it basically disperses light in several directions (see the *single lens* in Fig. 3.4). For most transparent materials the following inequality constraints hold,

$$1 < n(\lambda_{\text{red}}) < n(\lambda_{\text{green}}) < n(\lambda_{\text{blue}}). \quad (3.17)$$

For a thin lens (see description in Eq. 3.3), light originating from the object plane will consequently be focused at slightly different spots depending on the wavelength (can be both longitudinal as transversal). All lenses that make use of refractive media (in which the light travels through transparent media different than air), are prone to this phenomenon, which is known as *chromatic aberration*. For many optical setups, it is crucial to correct therefor. Longitudinal chromatic aberration, or LCA, occurs if light is focused at different depths along the optical axis. Correction is mostly done using achromatic doublet (*achromats*) or triplet (*apochromats*) lenses. Achromats perform what is called *primary LCA correction*: two different wavelengths are focused on the same spot on the optical axis. They are mostly comprised of a positive and negative lens with different refractive index. Apochromats deploy *secondary LCA correction* - three different wavelengths are focused on the same longitudinal spot - which mostly manifests in the use of three separate lenses. Fig. 3.4 illustrates all differences.

3.2.4 The electromagnetic spectrum and spectral purity

The electromagnetic spectrum is the range of frequencies (the spectrum) of electromagnetic radiation and their respective wavelengths and photon energies. It ranges from photons having extremely high frequencies and small wavelengths (gamma rays: hundreds of Eta Hertz at several pico-meters) to extremely low frequencies and very long wavelengths (long radio waves: several Hertz at tens of thousands of meters). The classes of light of which optics is mostly concerned with are the near-ultraviolet (100nm - 400nm), the visible (400nm

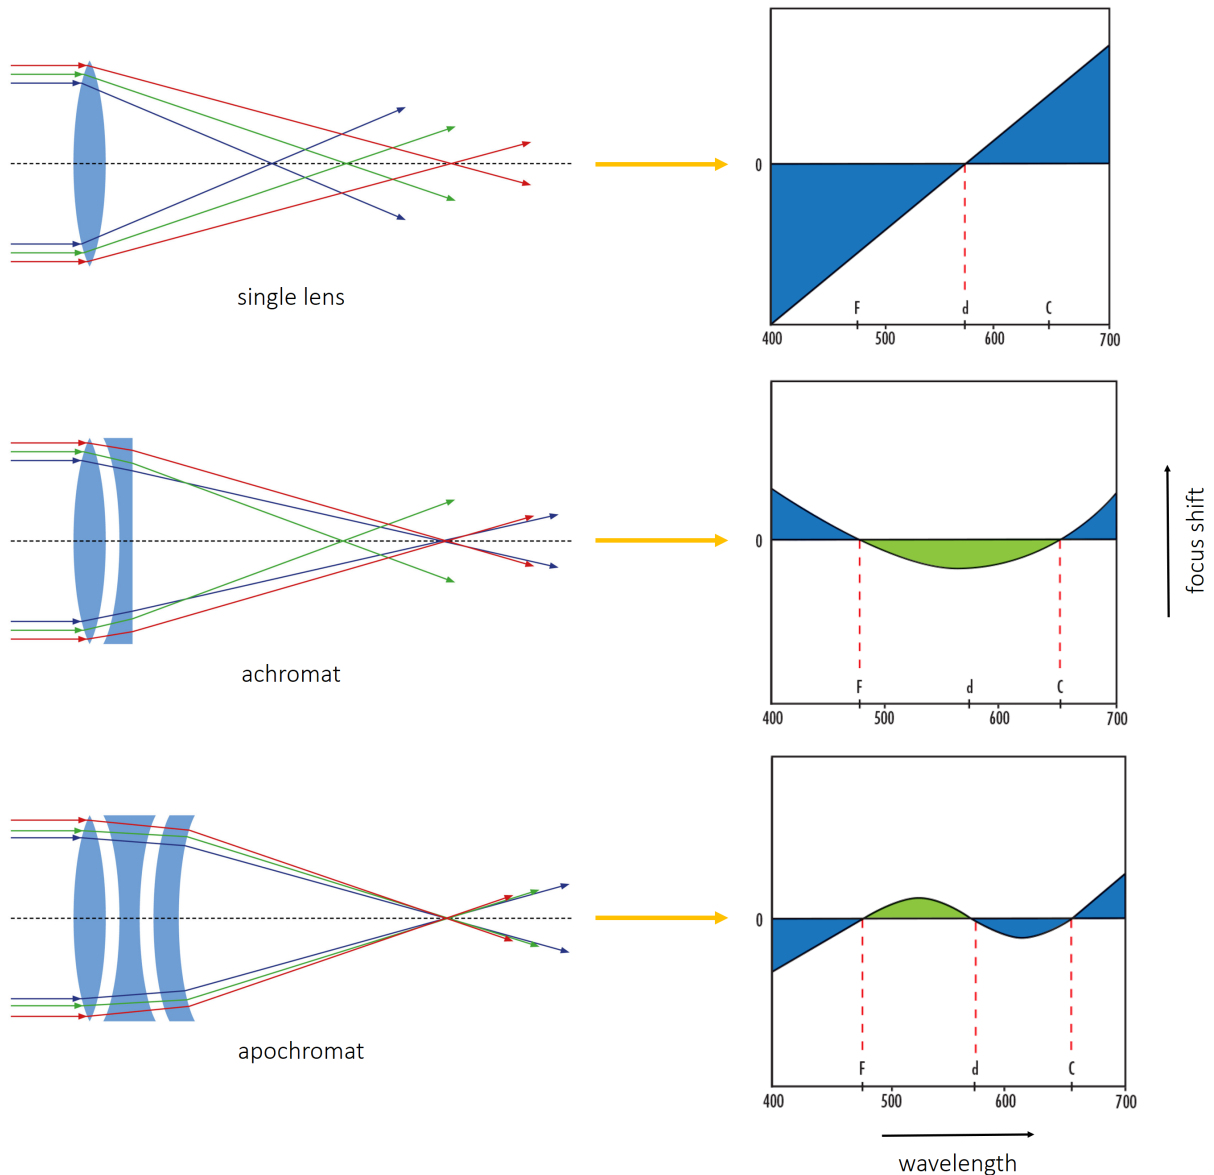


Figure 3.4: Chromatic aberration in lenses due to dispersion. On the left three types of lens configurations are sketched: the single lens (top), the achromatic doublet lens (middle), and the achromatic triplet lens (bottom). In optical terms, 656.3nm (red) is referred to as 'C' light, 587.6nm (yellow) as 'd' light, and 486.1nm (blue) as 'F' light. For the single lens, only 'd' light is perfectly focused, while achromats do so for 'F' and 'C'. Apochromats do so for all three. On the left, the *focus shifts* are given for all intermediate wavelengths. In terms of overall performance in VIR-region, apochromats usually obtain the highest degrees in LCA correction.

- 700nm) and infrared (700nm - 3000nm) regions. In microscopy, the visible region, or VIR, is usually most interesting as we, humans, are able to perceive most of its signals. In order to capture as much information of the samples as possible, the source of a microscope should optimally emit light from all over the VIR region. Halogen lamps are among the most reliable

and broad spectra possible. Like all incandescent light bulbs, a halogen lamp produces a continuous spectrum of light, from near ultraviolet to deep into the infrared. Furthermore, its lifetime is relatively long. The spectrum is depicted in Fig. 3.5.




Figure 3.5: Spectrum of a Halogen lamp.

In laser optics as well as hyperspectral imaging, it is often important to have a *pure* spectrum instead. It is a quantification of the monochromaticity of a given light spectrum - the "narrowness" of that particular spectrum - and is referred to as the *spectral purity*. Spectra coming from a laser source or emitted from a monochromator (see Sec. 4.2.1) are mostly Gaussian-like distributed for which the *full-width half-maximum* is a common metric to measure the spectral purity of the signal - the distances between both half-maxima. In multi-modal and hyperspectral imaging spectral purity ranges mostly between one and several tens of nano-meters. Spectral purity is further illustrated in Fig. 4.4.

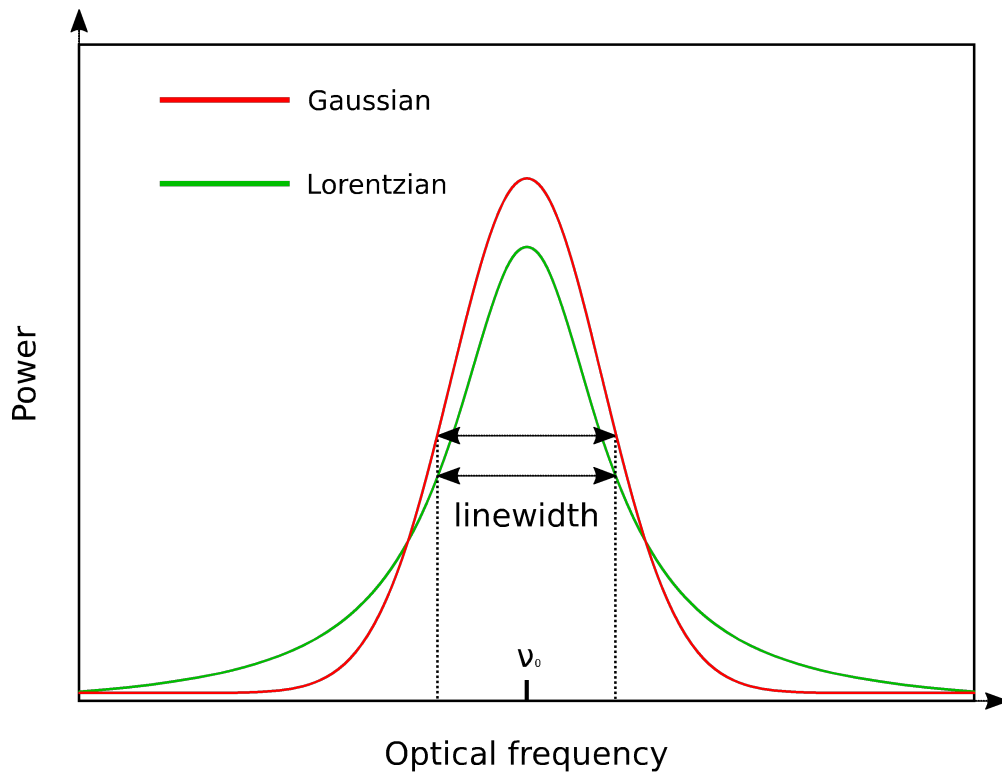


Figure 3.6: Spectral purity as measure for the monochromaticity of a light's spectrum. Here, the spectral purity of the frequency is given. Please note that the report uses the wavelength.

4 | Design and realization of a hyperspectral imaging system

This chapter focuses on the development of a hyperspectral imaging (HI) system in a transmission-microscopic context as novel investigative tool for thin micro-scaled biological samples. The first section focuses on the concept of HI systems and its categorizations, the second on the development of our own setup after which it is evaluated and discussed in the final section.

4.1 Hyperspectral imaging systems

Depending on the application, HI systems are different in how they retrieve the data. Three distinct acquisition modes can be identified: the spectral scanning mode, here referred to as staring mode, the spatial scanning mode and a snapshot mode. Systems that scan the 2D morphological landscape along different wavelengths are called *staring methods*. Usually, these methods have a spectral filter somewhere in their scheme that enables the visualization of the spatial structure for a specific narrow spectral band while filtering all other light out. In many cases, electro-optical filters such as liquid crystal tunable filters (LCTF) or acousto-optical tunable filters (AOTF) are used for this which are either placed within the illumination path or in the light path emanating the specimen. In case of the *spatial scanning mode*, the most used is *Pushbroom* that scans along a line in the spatial grid and retrieves these pixel's spectra simultaneously. It uses a dispersive device that is placed in the path emanating the sample which focuses that single line on a 2D camera in which the wavelengths are spread along lines of pixels on the camera's sensor that are orthogonal to the direction of the scanned line, therefore filling the space of the sensor. While moving along all adjacent lines, it creates a growing pack of 2D images. *Whiskbroom* is similar in that it spatially scans as well. Its difference resides in the fact that it scans, instead of lines, along individual points, therefore scanning along two directions. A snapshot retrieves *all* data in one instance. Fig. 4.1 provides an overview of all four methods and elaborates on the retrieved data in more detail. For microscopical-integrated HI setups it is beneficial to implement a staring mode because scanning along the spatial domain is magnified as well - as the precision of scanning decreases linearly with magnification - and the slits or holes need to be smaller, making them harder to be manufactured. The way data is retrieved and organized is explained in Fig. 4.2.

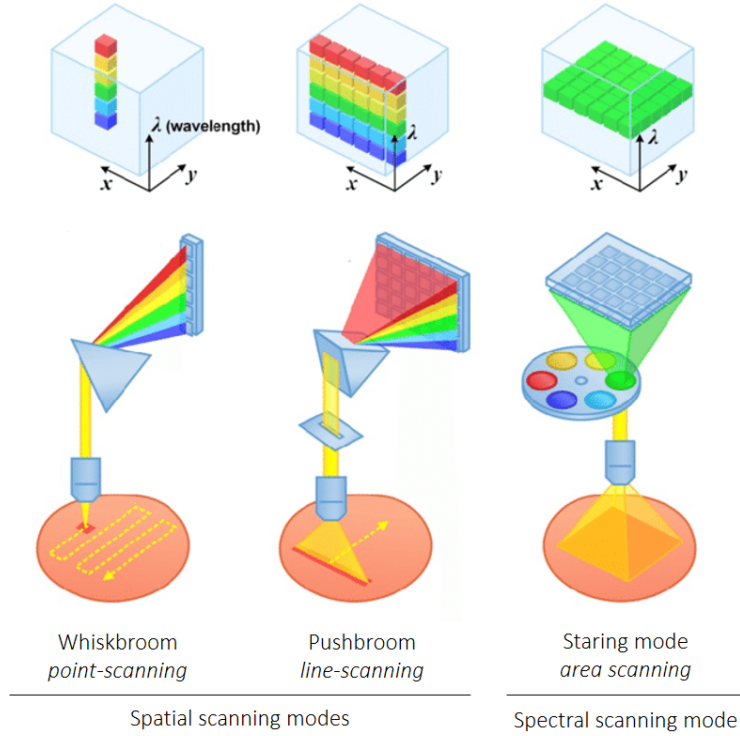


Figure 4.1: Categorization of HI systems [7].




Figure 4.2: Data retrieval in HI. In the top-left corner, the set containing all images at specific wavelengths is given. Instead of representing it as a set of features, the images get stacked in what is called a *3D-hypercube* (bottom-left). Every pixel is therefore associated with a spectral response for which some are sketched on the right. The matrix representing all individual spectral responses in its columns is referred to as the *voxel-set representation*.

4.2 Design and realization of the optical scheme

The microscopic setup that is developed is illustrated in Fig. 4.3. It comprises five distinct parts, or *stages*: (1) source, (2) illumination path, (3) sample stage, (4) objective, and (5) camera. Following subsections will discuss and explain all stages separately. The setup's discussion and conclusions are done in the final section of the chapter.




Figure 4.3: Overview of the optical scheme: staring scanning hyperspectral imaging microscope. Light travels from A to B in which it chronologically follows the route through the source, the illumination path, the sampling stage, the objective and is arriving at the camera - all highlighted in bold. Sub-stages that are together fixed in space are outlined in black: the sub-stage condenser, the sub-stage sample and the capturing sub-stage.

4.2.1 Source

The source consists of three components: a halogen lamp, a monochromator and an optical fiber. The halogen lamp is used as primal light source. It emits a broad spectrum of light, is consistent among many wavelengths, is reliable and has a relatively long life expectancy. The halogen lamp is coupled to a monochromator that enables the selection of specific wavelengths the setup is working at. It uses dispersion as its working horse and tunes by rotating a blazed grating that is in the light path. The angle at which the first order light rays emanate that element depend on the wavelength. If a small hole is placed at the end of the light path, it is, as a consequence, possible to capture a narrow spectral range of the primal source's spectrum. Visual assistance is given in Fig. 4.4 (on the left). The peak's

location is the wavelength the monochromator is tuned at. As exit slit, an optical fiber is used such that it is easily mountable and exchangeable with the other parts of the setup. It furthermore has the benefit of emitting a two-dimensional symmetrical cone of light, for which, most likely, the brightness is Gaussian-distributed. Obviously, by changing the size of the entrance slit or diameter of the fiber, the spectral purity of the signal changes - the peak becomes relatively narrower/wider and lower/higher. The difference in spectral purity for two specific slit settings is measured and given in Fig. 4.4 (on the right). Increasing the spectral purity comes at an inevitable cost, namely the brightness of the emanated signal. The second configuration is chosen, providing a high enough spectral resolution. Proper sample illumination is essential to make optimal use of the light.

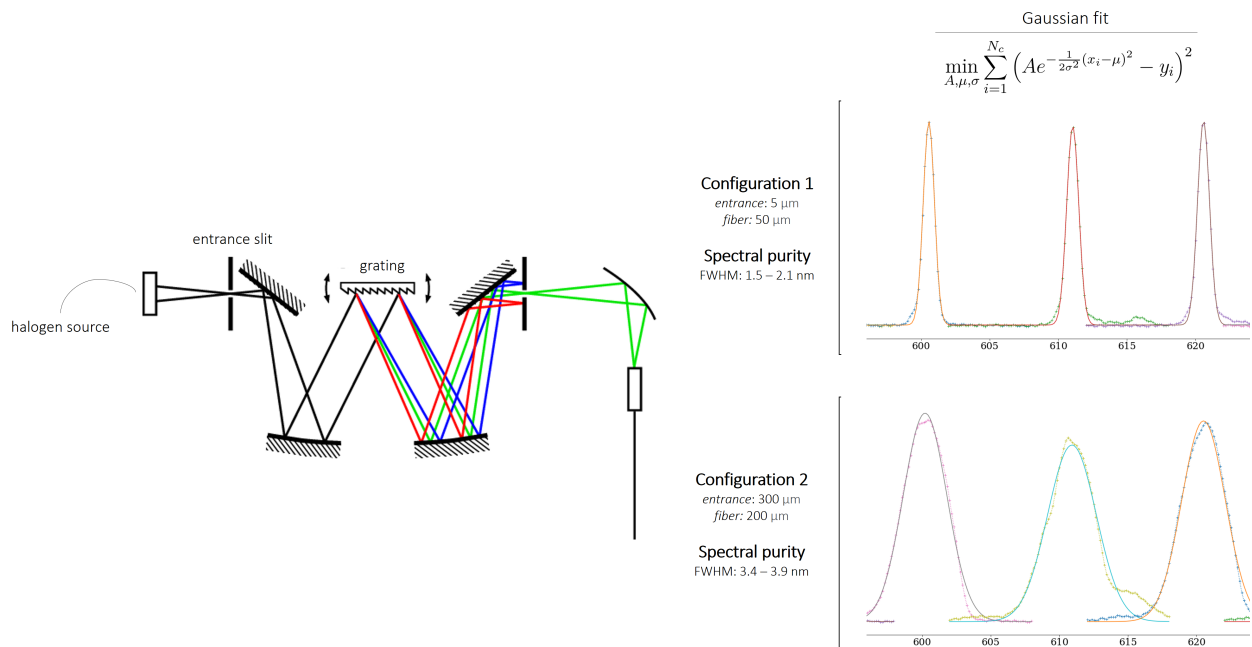


Figure 4.4: The source and spectral purity of its extension. On the left, the source is illustrated in which the outputted band is tuned by rotating the blazed grating. An optical fiber is coupled to the output of the monochromator and. The amount of light and spectral purity highly depend on the chosen configuration - size of entrance slit and fiber diameter. Spectral purity is measured for two configurations as indicated on the right. Part of the measurements are shown in the range of 600 - 620 nm. Furthermore, a Gaussian fit is used to approximate the mean and *full-width half-maximum* (FWHM) for every selected wavelength measured at. Least-squares using the Levenberg-Marquardt algorithm is used as optimizer [8].

4.2.2 Illumination path

Having an extended source that produces brightness-depressed cones of light, it becomes essential that the illumination part is optimized in the sense that all the light is focused on the samples field of view. Furthermore, it should be optimized such that the sample is homogeneously illuminated, preferably with light directed from different angles. In achieving

this, the path is aligned using Koehler alignment (see Sec. 3.1.3 and Fig. 4.5). If correctly aligned, the sample is homogeneously illuminated by a converging set of plane waves that each arise from different points of the source, thereby improving and squeezing out maximal lateral and axial resolution. The Koehler aligned subsystem consists of a multiple of lenses and diaphragms: the collector lens that collects all the light from the source, a field lens that creates an, in our case enlarged, image of the source in the front focal plane of the condenser lens. The condenser then focuses differently angled collimated beams on a small area around the optical axis. The source's image should ideally be as large as the opening aperture of the condenser lens such that its numerical aperture limit is reached without wasting optical power. By aligning the field iris with the front focal plane of the field lens, it is conjugated with the sample. Changing its diameter is immediately visible in the sample plane and helps aligning the setup. The condenser aperture on the other hand, if placed in the front focal plane of the condenser, can be used to tune between spatial resolution and contrast: opening the diaphragm results in a higher working NA and thus allows for improved resolution, while decreasing the diameter allows for better contrast. The condenser sub-stage - the condenser lens and aperture - can be moved manually across all axes, allowing easiness and handiness of use. The height (z direction) is manipulated with a circular knob and the others, x and y directions, are tuned using an precise XY-stage (cage-system fit).

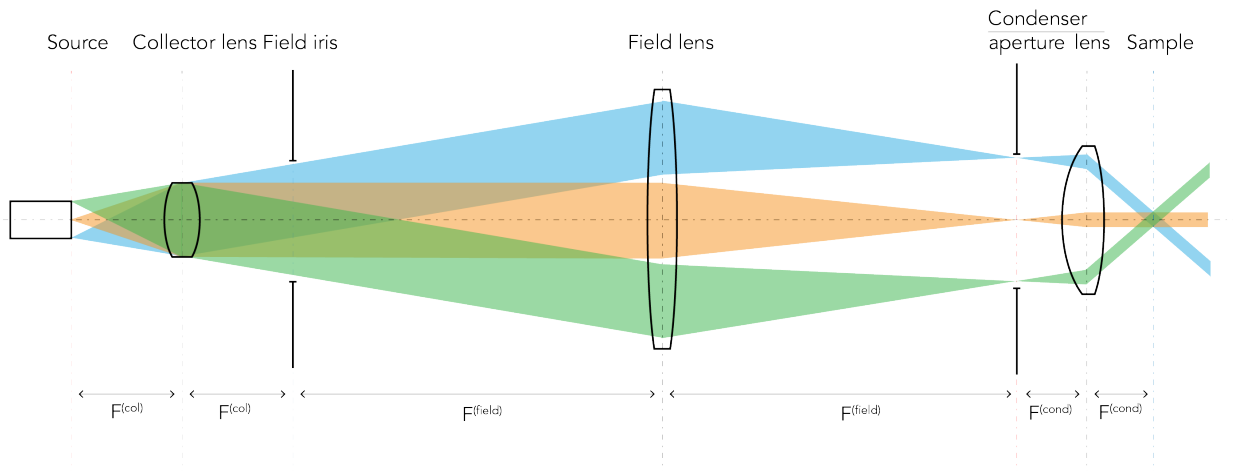


Figure 4.5: Koehler aligned setup.

4.2.3 Objective and camera

The *capturing sub-stage* of the system consists of an objective, a tube lens, and a camera. The objectives used are plan- apochromats that highly correct for chromatic aberration along the region we are interested in, highly correct for spherical aberration and have high numerical aperture ranging from 0.75 to 1.25, depending on its magnification and type of immersion. Specific magnifications used are 20, 40 and 63. A monochrome *IDS* CCD camera is adopted in the setup. The pixel size is 1.25 micron, which is among the smallest industrially accessible, has 4912 and 3684 pixels in width and height respectively (18MP) and has a pixel depth of 12-bit. Camera read-out was done using Python.

4.3 Evaluation and discussion

On the resolution The imposed resolution of the system is computed as in which both the numerical aperture on the condenser side as well as that of the objective side are vital (please see Eq. 3.15). In order to achieve a high enough resolution in the spectral range of interest (400 to 700 nm), elements in the illumination part should be chosen smartly such that that the condenser sides NA is high enough as well. The resolution for different system settings given the smallest possible objectives NA (=0.75) are plotted in Fig. 4.6. On the left, the ratio between the field and collector lens, the field magnification, is given and on the bottom the size of the source. We aim for an optimal resolution that is smaller than 0.6 micron. Therefore, it is chosen to have a source that is approximately 2 mm in diameter which is attained by placing a diffuser behind the exit of optical fiber. The magnification of the field is chosen to be 200/30 (see huge red marker in the figure).




Figure 4.6: Spatial resolution of the setup for various configurations. The spatial resolution, d_λ , is theoretically computed according to the top-noted equation. Here, NA_{obj} , λ and F_{cond} are fixed (underlined in red) and the magnification of the field, M_{field} , and the diameter of the source, D_{source} , are used as configuration parameters (underlined in blue). Green dots denote the region with sufficient spatial resolution while the red-dotted legion is the non-optimal domain.

Furthermore, it is observed that the field lens has a relatively long focal length. Consequently, the setup is developed with a horizontal optical axis that contains most of the elements comprised in the illumination path and a vertical one that contains the other sub-stages of the system in order to make the setup handable. Practically, it is found that the system works in the spectral range between 400 and 700 nm in which 4nm is used as spectral resolution, thus providing 76 channels. Some calibrated¹ examples of Giemsa-stained blood films are provided in Fig 4.7 for 2 wavelengths. Although spatial resolution is decreased if compared with other high-end microscopes, it remains sufficient for the clear-cut distinction of the parasites on the basis of their morphological features, e.g. trophozoites, the ring stage of the *P. Falciparum* malaria parasite, is easily detected (especially observe the huge ring in the center patch).




Figure 4.7: Calibrated images of Giemsa-stained thin blood smear films for several wavelengths. Crops are taken from the whole FoV which is approximately 15 times as big, thus containing about 500 erythrocytes in total. A plan-apochromat objective with 20X magnification is used.

Versatility, interchangeability and easiness of use The setup can be used for all transmission-microscope-made samples and is tweaked and optimized in a fast and user-friendly manner. The most crucial movements, the displacement of sub-stage condenser and capturing stage, are done with well-workable precision stages that are tuned with proper knobs. The system is either used as conventional microscope or as HI microscope in which interchanging between both can be done efficiently as coupling and decoupling of both sources takes a minimal amount of time. Depending on the application, the examiner can without

¹calibration is done using $\frac{I_{raw} - I_{dark}}{I_{ref} - I_{dark}}$, in which I_{raw} is the measured signal, I_{dark} the signal measured with closed-off camera and I_{ref} a reference image without sample.

much effort choose the brightness, select the wavelength, trade between resolution and contrast, alternate between specific cameras, and control the region of the sample that is focused. The setup is photographed and given in Fig. 4.8.




Figure 4.8: Hyperspectral imaging microscope: the real setup.

On the cost Compared to standardized setups, the cost of the setup is significantly lowered: approximately between 5.000€ and 13.000€ depending mostly on the camera and objective. The plan-apochromat objective is the main contributor (for Leica, ranging between 3000€ and 9000€). It is though found that plan-apochromats are near-essential as

their highly corrective property for chromatic aberration and spherical aberration makes the process more reliable without the need of refocusing every wavelength iteration.

Remarks and potential improvements The main disadvantage, although optimally illumination-aligned, is low-brightness, obstructing high-end performance in spatial resolution, spectral resolution and the *perceptible spectrum* - the spectrum of light the system is able to measure at. The spectrum between 300 and 400 is especially interesting in many biological processes and is expected to be informative for malaria as well. Inevitably, the primal source's (halogen lamp) spectrum peak is usually centered somewhere around 600 - 800 nm and is much depressed in the aforementioned region. A source which is less *dimmed* in those regions could be beneficial. Having a more powerful source on the other hand is profitable as well, as increasing the brightness along the whole spectrum improves the spatial resolution along the whole domain as well. Furthermore, it could potentially increase the spectral resolution because modifying the source settings, that is decreasing the sizes of entrance slit and fiber diameter, can be done without a disastrous fallback in spatial resolution.

Adapting to a different configuration (field magnification and extended source diameter) could squeeze out the last boost in spatial resolution, but at a cost. In moving the red circle in Fig. 4.6 deeper into the interior of the green-dotted territory, one either has to increase the diameter of the source or the magnification of the field. Increasing the diameter of the source reduces some brightness and enlarging the magnification of the field is tantamount to a more powerful condenser lens or a less powerful field lens. Both options are cumbersome as both lay harder practical constraints: (1) sample gets very close to the condenser or (2) the setup's vertical part will become substantially longer.

Retrieving data for many spectral bands is time-consuming since the wavelength should be manually tuned every single instance. This process could be enhanced by automating the capturing process in which independent motors are coupled to both the monochromator knob and fine-tune focusing knob. A natural feedback system that, on the basis of the current state of the camera, has the ability to refocus in combination with a stepper-motor that autonomously increases the wavelength could enhance and, more important, speed up the process significantly.

A serious disadvantage of plan-apochromats and lens-comprised objectives is that they never perfectly correct for chromatic differences. Especially if larger perceptible spectra are required, chromatic aberration becomes substantial in the spectrum's peripheries. A reflective objective, only using parabolic mirrors to change the beams of light, could benefit the setup. Drawbacks are serious flaws in brightness, increased spherical aberrations and reduction in NA.

Sequel A natural question arises: *Is the developed HI system good enough for non-invasive detection of parasites in thin blood-smears?* Definitely supplementary steps are to be taken in order to answer that. A proposal of steps (including those already taken) is given in Fig. 4.9.

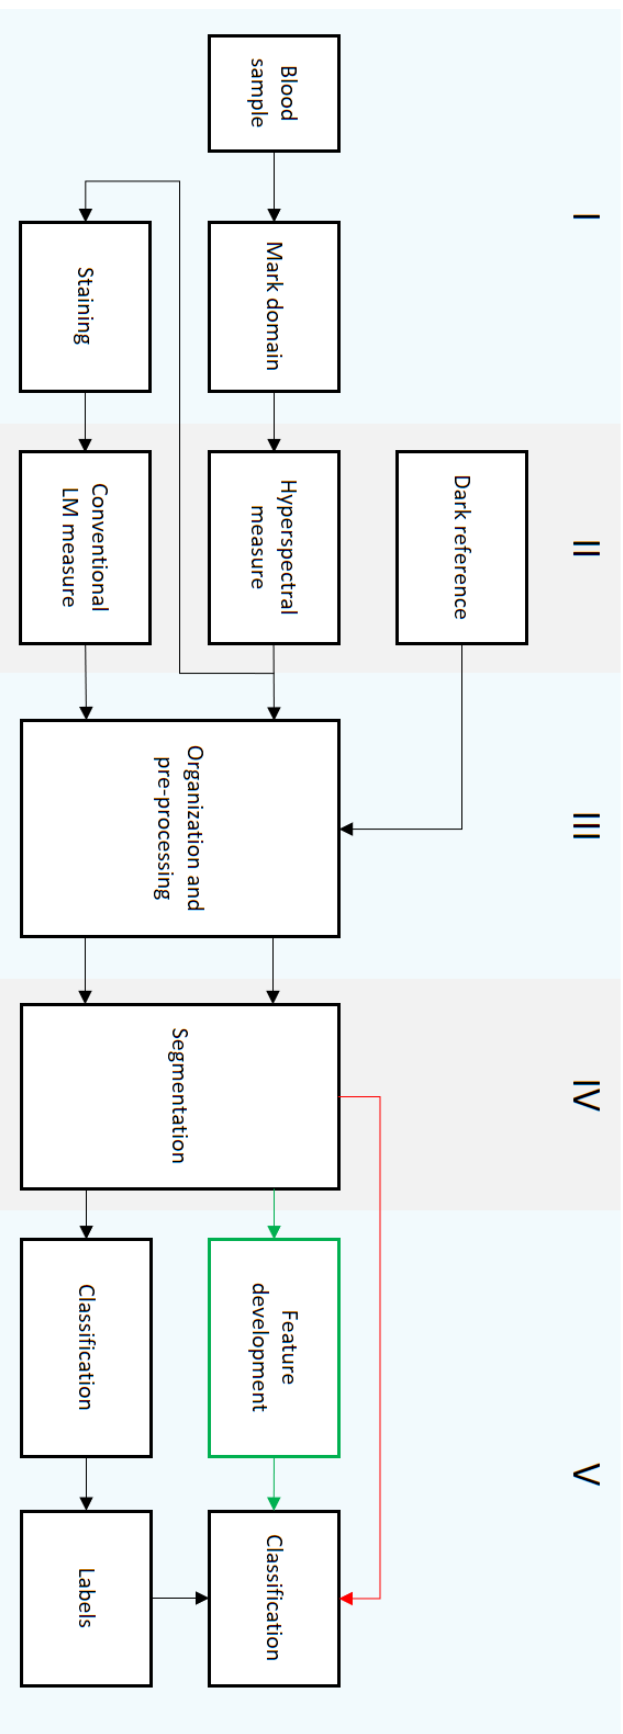


Figure 4.9: Proposed flow-diagram for non-invasive classification of malaria parasites using hyperspectral imaging. The first region concerns the acquisition and preparation of the sample. The second region shows the measurements and the third the pre-processing steps. The fourth is the protocol for segmentation and storing of individual cell data. The fifth is an automated detection protocol based on these individual cell data. It should be immediately noted that there is two parallel horizontal lines: one representing hyperspectral imaging and the other conventional imaging. This is done because we need a validation set to train and test the data on. Since there is no clue based on the hyperspectral data only how to classify every cell, we need a method on which we can trust that does so reliably. To do so we stain, after HI measuring, an equivalent sample and examine that under a conventional microscope. The same domain is evaluated and labels are given for the whole set of segmented cells. This is subsequently divided in train and test set and fed into a classification algorithm.

Part II

Towards Smart Generalizing Networks

The topic of Part I was the design and evaluation of a new potential imaging modality for the retrieval of data. As mentioned in the introduction, time and the lack of proper unstained samples made it impossible to generate large enough databases for being pushed further in the diagnosis chain (for the chain, see Fig. 2.5). This part is dedicated to the ending of that chain: *classification*. The intermediate steps (pre-processing and segmentation) will be shortly mentioned but not exploited (see sec. 7.1). For several years, deep learning algorithms have propelled in many medical domains in which the segmentation and classification of malaria has become one of lately [65, 3]. Those networks are often very large in size and no clear indication of their robustness towards other specific malaria databases is sketched yet. Most of the networks deploy relatively simple operations, such as convolutions or fully-connected layers, which could be one of the reasons why they have to be so deep. An idea is to increase the complexity of the operations such that a-priori knowledge is induced in these operations instead. The field of *geometric deep learning* is such an example, in which realistic geometric properties of the objects of interest are being mold into the operations (think of physical transformations the objects can freely undertake without changing what they are, such as rotations, reflections, scaling, etc.). This part is about the design and evaluation of several of those operations and how they contribute to overall performance on the well-known Rajaraman database [3]. Furthermore, robustness of the pre-trained models is checked on a database extracted from the microscope developed by AiDx Medical [4]. Unfortunately, the models have not been tested on hyperspectral data for above-mentioned reasons. Ch. 5 focuses on all the necessary theory to build the operational stages that are being implemented in deep models described in Ch. 6. The used databases, performance metrics and results are described in Ch. 7. At last, Ch. 8 discusses and concludes the findings.

5 | Generalizing operations in neural networks

This chapter captures the essential theory on *smart* neural network's layers that will be used to build and train deep networks in later chapters of this part. Most of the readers must be familiar with the idea of neural networks or at least the impact they have had over the last two decades. Extreme improvements have changed and benefited tasks ranging from language-processing to those associated with the monitoring of farmlands, and surveillance to those needed in the medical sector. In those tasks it becomes more and more precious to additionally develop networks that are both *generalizing* as well as *small*, or as I will sometimes refer to in combination: *smart*. In my opinion, a well-generalizing network is best described as it being equipped with the *power to capture relevant information*. Within those networks it therefore *learns important aspects* while it *discards non-relevant features* that manifests itself in (mostly) the *condition* this important information is in. Generalizing networks will be fundamental in the design and implementation of robust systems and in systems that could be used for a variety of tasks simultaneously. Small on the other hand is two-fold: small in terms of *allocated memory* and small in terms of *complexity*. More specifically, this chapter will be about the design of generalizing operations for the use in deep networks (see Ch. 6). It asks from the reader to be equipped with a proper mathematical background as it touches upon several complex domains such as *group theory*. It is divided in 2 sections. First the basic artificial neural network is described. Section 2 touches upon the core in which several equivariant operations will be explained in detail which will be the building bricks for the networks designed in the next chapter (6).

5.1 Artificial neural networks

Neural in artificial neural network, or ANN, originates from the association with the human brain. Although ANNs discard, in my opinion, many complexities necessary for the tasks *real* networks of neurons are able to collaboratively perform, it is understandable why it has been given its name. An ANN is in essence a collection of interconnected nodes in which those nodes and interconnections biologically mimic neurons and synapses respectively. Here, nodes are associated with a *value* and connections with a *function*. Interconnected nodes are thus able to exchange data determined solely by the functions that prescribe the connections. As you will find out, the ability of ANNs to adapt the function's parameters and thus the whole communication protocol between all nodes creates the possibility to learn.

The common *architecture* of ANNs consists of an input layer, several hidden layers and an output layer (see 5.1). The amount of hidden layers and nodes per individual layer are architectural *hyperparameters* and will be discussed in more detail in the later chapters. As ANNs have to learn (from data) they are restricted with the ability to input data. This is done at the input layer in which those nodes exactly represent the data. Furthermore, it needs to have the ability to evaluate the current state of the network. This is done at the output layer in which the combination of values of nodes is somehow checked with what the network ideally should have produced, also known as the *ground truth*. The network is completed by the interpositioning of hidden layers between in- and output layers. It should be noted that *all* connections are uni-directional as information should always stream from in- to output¹. This has the direct consequence that nodes in the l^{th} -layer are exclusively determined by those in the previous, $(l - 1)^{\text{th}}$ -, layer (and the way they are connected).

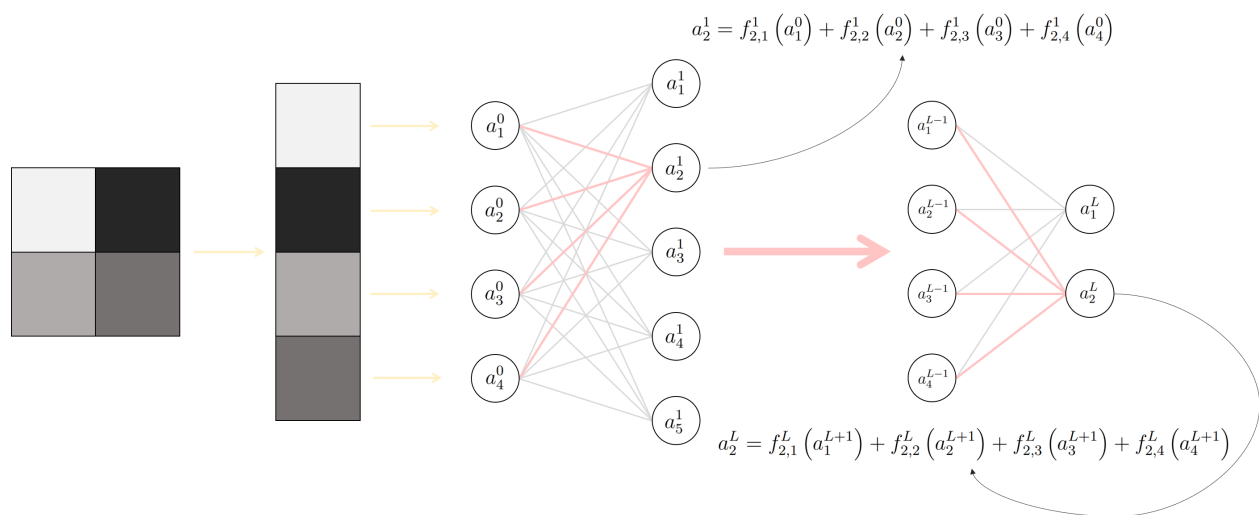


Figure 5.1: General structure of the artificial feedforward neural network.

The flowthrough of data from input to output layer is usually referred to as the *forward pass* or *feedforward*. In most general form, the values of the nodes residing in the hidden and output layers² are computed as

$$a_u^l = \sum_{v \in \mathbb{N}^{l-1}} f_{u,v}^l (a_v^{l-1}) \quad (5.1)$$

in which the subscripts u and v represent the node numbers in the (superscript) l^{th} -layer or $(l - 1)^{\text{th}}$ -layer respectively. Furthermore, \mathbb{N} describes the set of all nodes present in the indicated layer. At last, the function $f_{u,v}^l$ determines how a specific node from the $(l - 1)^{\text{th}}$ -layer passes its information to the node in the l^{th} -layer. Computation of the node's values is done layer per layer in which one starts at the input and subsequently walks through the

¹One should recall that other types of interconnections are possible: cycles, bi-directional, etc. Nevertheless, in my thesis work I will stick to these type of network, often called *feedforward* neural networks.

²Remember that the nodes in the input layer reflect copies from the injected data and do not need to be computed.

hidden layer in incremental order until reaching the output layer: the feedforward is then completed.

The interconnections, represented by f , normally consist of a *linear* and a *non-linear* part. The non-linear part is crucial if realistic properties must be learned [66]. As first step, a specific node is computed as an affine combination of the nodes that pass information to it. The second step forwards this intermediate output through a non-linear function which is equivalent for all nodes. By expanding eq. 5.1 as suggested the nodes are computed as follows

$$a_u^l = h \left(\sum_{v \in \mathcal{N}^{l-1}} w_{u,v}^l a_v^{l-1} + b_u^l \right) \quad (5.2)$$

in which the w 's and b are the weights and bias, and h an arbitrary predefined non-linear scalar function, called the *activation* function. The vector containing all node values simultaneously is easily retrieved as

$$\mathbf{a}^l = h \left(W^l \mathbf{a}^{l-1} + \mathbf{b}^l \right) \quad (5.3)$$

in which the matrix W^l and vector \mathbf{b}^l consist of all the weights and biases of the l^{th} -layer respectively. Having chosen the architecture and activation function, the state of the network depends only on the weights and biases. Updating these parameters allows the network to learn from samples that are accompanied with their ground truths, or *labels*.

A natural question arises on how to update the parameters such that the network gets increasingly better. Two steps remain: (1) evaluating the "goodness" of the current state of the network and (2) updating the parameters accordingly. Evaluation is done with a function that measures how the output relates to the ground truth. This *cost function*, f_C , obeys the logical convention that it generates a small value if truth and output are nearly similar and vice versa. Ideally, the ground truth and output are almost equivalent: in this case the prediction and/or estimation of the network is nearly perfect. The cost (value), C , is computed as

$$\begin{aligned} C &= f_C \left(\mathbf{a}^L, y \right) \\ &= f_C \left(\left\{ w^l, b^l \right\}_{l=1,2,\dots,L}, y \right). \end{aligned} \quad (5.4)$$

Here, y is the (set of) label(s) associated with the input data and L is the amount of layers³. If the cost function and activation function are continuously differentiable w.r.t. its parameters, a proper way to update them is in the direction of the negative gradient:

$$\begin{aligned} \boldsymbol{\theta}_{k+1} &= f_U \left(\nabla f_C \left(\boldsymbol{\theta}_k \right), \boldsymbol{\theta}_k \right) \\ &= \boldsymbol{\theta}_k - \epsilon \nabla f_C \left(\boldsymbol{\theta}_k \right). \end{aligned} \quad (5.5)$$

³Closely paying attention reveals that this is not entirely true: the first layer is by convention the *first hidden* layer. Therefore, the total amount of layers, including the input layer, becomes $L + 1$. Furthermore, the input layer is the 0^{th} -layer.

Here, k represents the current iteration, ϵ the *step size* or *learning rate*, and θ a *vectorized*-collection of all parameters. An ingenious approach to compute these gradients efficiently is based on the concept of *backpropagation* which is properly worked out in several good test books such as *Dive into Deep Learning* (Sec. 4.7 to be precise) [67]. The idea resides in the fact that the cost entirely depends on the last layer, the last layer only on the layer before that and so on. Therefore, without loss of information, the gradient of the last layer's parameters can be computed first. After having computed those derivatives, the gradient of the second to last layer's parameters can be computed. This procedure ends if the first layer is reached. It is efficient because the computed gradient in a certain layer is used in the computation of the next (thus posterior) layer. The full training procedure is depicted in Fig. 5.2.

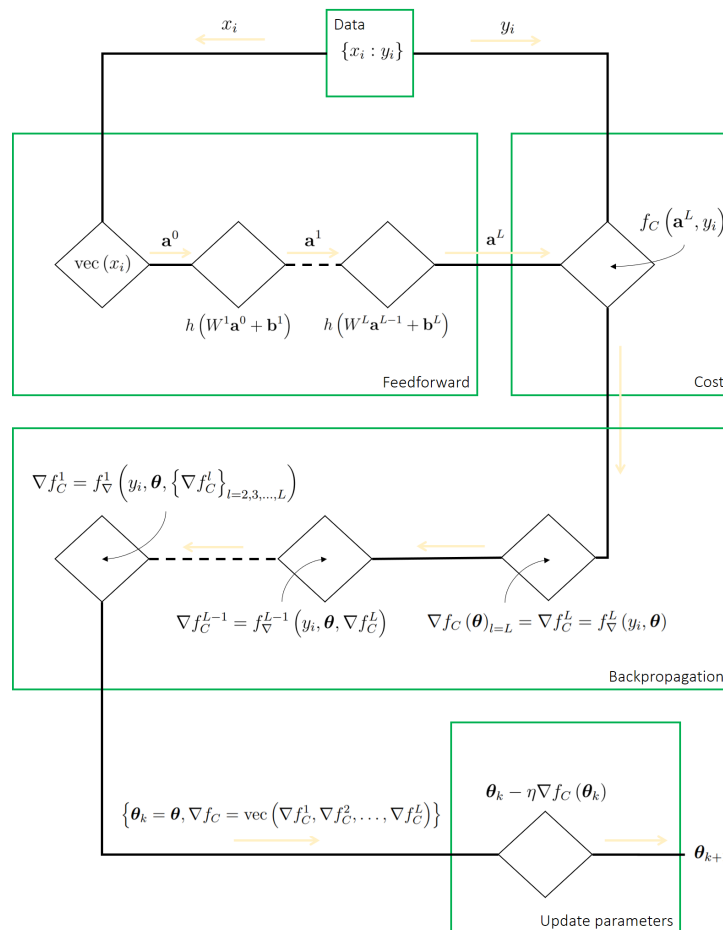


Figure 5.2: Update procedure for artificial feedforward neural networks.

In Ch. 6 more emphasis is laid on the choice of activation and cost function, *hyperparameters* such as the size of the network and the learning rate, and more advanced learning/updating procedures. In this section, I provided the most general organization of ANNs which is a basis for almost every network that is used elsewhere. Although these networks are able to learn complex tasks, they do not yet incorporate useful knowledge. The next sections will focus on the inclusion of this knowledge by interconnecting the weights in a smarter manner.

5.2 Equivariant operations

In a general ANN, nodes between consecutive layers are all connected. Therefore, *fully* connected layers exchange information among all nodes. One might think that such a property of the network is ideal, but in many cases it is nonetheless far from optimal. Arguably, in some cases it might be beneficial but in many domains it is not. At least not in the domain we are interested in: computer vision.

In this field, *data is tantamount to images*. Images are 2D-projected observations of the real, 3D, world. Mathematically, images *live* and are captured in a two-dimensional discrete world, mostly denoted as

$$\mathbb{Z}^2 = \left\{ (x, y) \mid x, y \in \{\dots, -2, -1, 0, 1, 2, \dots\} \right\} : \quad (5.6)$$

the two-dimensional lattice of the Euclidean space, \mathbb{R}^2 . Every such 2-tuple, or in understandable language *pixel* or *point*, in this grid contains a spectrum for which its heuristic depends on the way the data is retrieved: for monochrome one channel is associated with every pixel, for RGB three and for hyperspectral images this is N ($\gg 3$).

If such an image is fed into a general ANN, all inherent *structure* of the image is not used. It is important that you fully grasp why this is true. In this context, which is that of images, structure is the way in which the pixels are *organized*. For example, every pixel has 4 orthogonal neighbors (horizontal and vertical) and 4 diagonal neighbors. Moreover, some pixels are close to each other and some are far apart. The human ability to discriminate objects through vision is largely contributed by the fact that the brain mechanism *incorporates* the structure in the discrete scene it captures with its eyes. This provides, for example, the ability to learn how to associate structured signatures such as lines, circles, rectangles, and more complex shapes and combinations thereof with specific objects. Since consecutive layers in general ANNs are fully connected, they do *not* incorporate the hierarchical patterns that are present in the image. These networks are in many computer vision related tasks prone to overfitting and moreover need to be much wider and deeper to establish a similar performance as other *smarter* networks.

Next subsections describe and design layers in which some image patterns are successfully integrated. More specifically, layers will be designed such that realistic *transformations* are *preserved* by design. This domain is sometimes referred to as *geometric learning*. If a particular layer preserves a specific transformation, say a rotation, it is said to be *equivariant under* that transformation. The exact mathematical definition will be discussed later but it is this property that must be proved for the layers that will be explained. Preservation of translations is the feature of *convolutional layers* which are explained in 5.2.1. Extending by the inclusion of the preservation of 90-degree rotations and mirror reflections, *group convolutional layers* are explained in 5.2.2. 5.2.3 and 5.2.4 dive deeper in which the concept of *steerability* is used to fulfill the need for an increased number of rotations the layers are equivariant under.

5.2.1 Convolutions

A tool that is often used in the altering of images is the *convolution* operator. The two-dimensional discrete convolution between two function f and g , both specified on \mathbb{Z}^2 , $(f, g) : \mathbb{Z}^2 \rightarrow \mathbb{R}^C$, is defined as

$$[f \star g](x) = \sum_{y \in \mathbb{Z}^2} f(y) g(y - x) \quad (5.7)$$

in which x combines both spatial coordinates, and C is the amount of channels of the input image (f) and kernel (g) it convolves with (Sec. 6.12 in [67]). Furthermore, we use \star to indicate the convolution⁴. Discrete convolutions are best explained schematically which is done in Fig. 5.3.

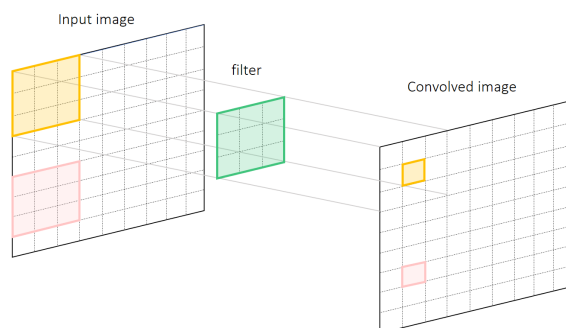


Figure 5.3: Two-dimensional discrete convolution. For every new pixel in the convolved image on the right a dot product between the input image (left) and filter (green) are taken: this scalar value is allocated at the central position of the filter. Moving around the filter and letting its center coincides with the whole domain of the image, the new altered image is retrieved.

Using convolutions in ANNs is relatively easy. The kernels (g) must be seen as carriers of the weights that are to be learned. Consequently, (1) the amount of weights decreases tremendously and (2) the weights are *shared* at nearly every position (up to edge effects). (1) Every *new* pixel (new node in the next layer) is only interconnected to its neighbouring pixels in the input image (depending on the size and shape of the kernel). (2) Since the kernel slides over the domain of the image, every new pixel is interconnected in equivalent fashion w.r.t. to its neighbouring pixels in the input image.

By design, a convolution produces a *new, altered version of the image* with only one channel: $[f \star g] : \mathbb{R}^K \rightarrow \mathbb{R}$. In networks that employ convolutions it is beneficial to have a feature space that has a multitude of channels: a *stack* of features (Ch. 6 in [67]). Obviously, a two-dimensional kernel will not work in this case and needs deepening (see the colored filters in 5.4). Every such filter generates one output feature. The amount of output features

⁴Precise readers may notice that the definition is not entirely consistent with literature: convolutions employ $x - y$ instead of $y - x$ which is formally known as the *correlation* operator. The use of the \star must become clear as well as it is mostly associated with the correlation operator.

is therefore equivalent to the amount of used filters. The set of filters is referred to as the *filter bank*. To summarize, the convolutional layer convolves the input feature map, $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^{C^{l-1}}$, with the filter bank, Ψ^l , a set of C^l filters $\psi^i : \mathbb{R}^{C^{l-1}} \rightarrow \mathbb{R}$ such that

$$[f \star \psi^i](x) = \sum_{y \in \mathbb{Z}^2} \sum_{c=1}^{C^{l-1}} f_c(y) \psi_c^i(y - x) \quad i \in 1, 2, \dots, C^l. \quad (5.8)$$

For visual clarity, Fig. 5.4 depicts what is happening.

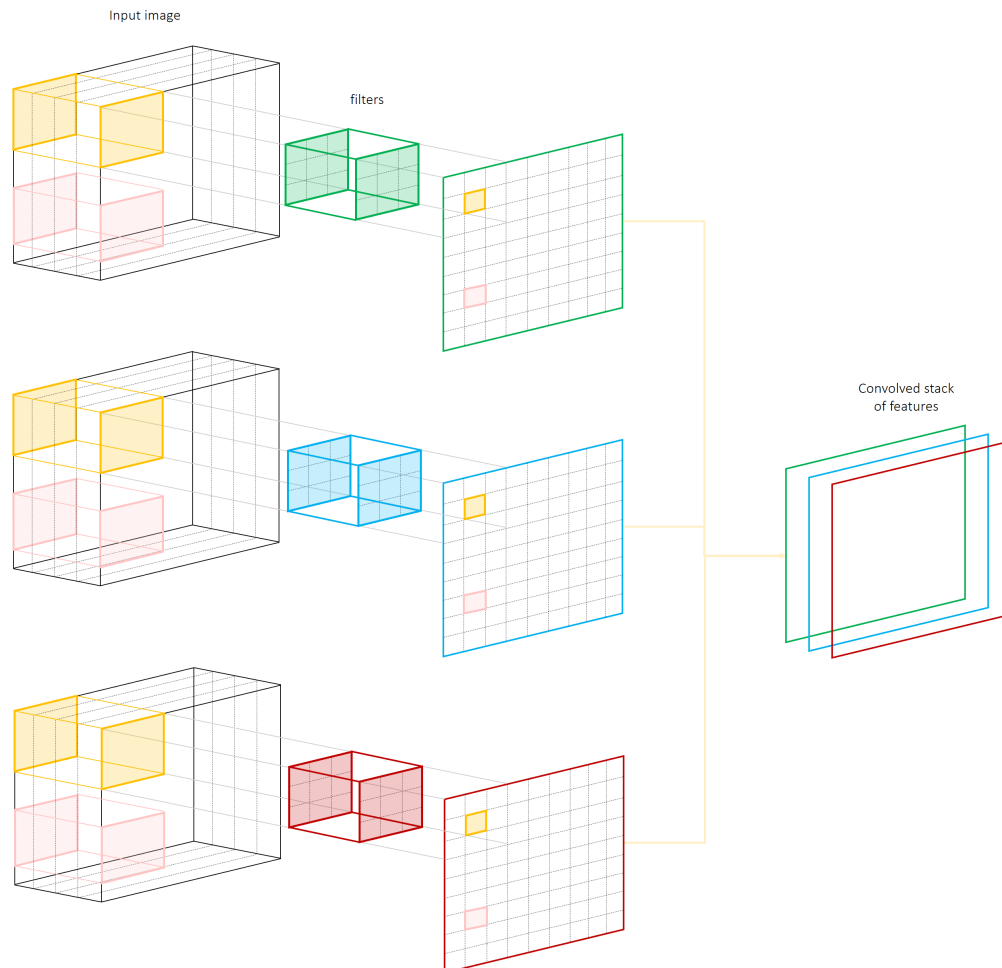


Figure 5.4: Convolutional layer in a CNN.

A nice property of convolutional ANNs, or *CNNs*, is their translational equivariance. First, we define the concept of equivariance. Equivariance is a notion for functions and is stated as follows. A function or map Φ is said to be equivariant under the transformation group G and the domain X if,

$$\Phi([\pi(g)](x)) = [\pi(g)](\Phi(x)) \quad \forall (x, g) \in (X, G) \quad (5.9)$$

in which $\pi(g)$ is the operator form of the group transformation g [68]. In words, it means that first transforming x and map it is equivalent to first map it and than transforming it.

$\pi(g)$ and $\pi(g)'$ do not necessarily be equivalent, they only have to linearly represent G , that is they have to obey

$$\pi(gh) = \pi(g)\pi(h) \quad \forall g, h \in G. \quad (5.10)$$

Since we are concerned with transformations of feature maps the group action is further defined as

$$[\pi(g)f](x) = [f \circ g^{-1}](x) = f(g^{-1}x). \quad (5.11)$$

Here it basically states that a transformation of the feature map is the same as the feature map that looks at the transformed coordinates, which is computable if the action is defined. For CNNs specifically, we define the translational or *shift* action acting on a feature map as

$$[\pi(t)f](x) = f(t^{-1}x) = f(x - t). \quad (5.12)$$

It is easy to prove that for every 2-tuple of translations (t, s) the combinatory action obeys Eq. 5.10 because

$$\begin{aligned} [\pi(t)\pi(s)f](x) &= [\pi(t)f](x - s) \\ &= f(x - s - t) \\ &= f(x - u) \quad u \leftarrow s + t \\ &= [\pi(u)f](x) \\ &= [\pi(st)f](x). \end{aligned} \quad (5.13)$$

It remains to prove that the convolutional layer in Eq. 5.8 is equivariant under shifts of the input. To reduce clutter, the amount of in- and output channels are held to a minimum. We show that

$$\begin{aligned} [[\pi(t)f] \star \psi](x) &= \sum_{y \in \mathbb{Z}^2} [\pi(t)f](y)\psi(y - x) \\ &= \sum_{y \in \mathbb{Z}^2} f(y - t)\psi(y - x) \\ &= \sum_{y \in \mathbb{Z}^2} f(y)\psi(y - (x - t)) \quad y + t \leftarrow y \\ &= [f \star \psi](t^{-1}x) \\ &= [[f \star \psi] \circ t^{-1}](x) \\ &= [\pi(t)[f \star \psi]](x) \end{aligned} \quad (5.14)$$

and hence the proof (see also Section 5 in [69]).

5.2.2 Group convolutions

Eventhough the conventional convolution is by far the most used operator in many models, networks still have to be huge (in depth and/or width) in order to achieve proper performance. Instead of changing the networks architecture (by somehow increasing the amount

of layers), the stages themselves can be adapted. This is exactly what T. S. Cohen and M. Welling did by introducing *group-equivariant convolutional networks* (G-CNNs) [69]. Besides translations, G-CNNs exploit other realistic symmetries of the sampling lattice as well. The working horses are the group convolutions, or G-convolutions, that again increase the degree of weight sharing. Depending on the specific implementation used, feature maps in G-CNNs are equivariant under those newly imposed symmetry transformations. Specifically, they designed G-convolutions that were able to learn features that equivary under 90-degree rotations and mirror reflections besides being equivariant under shifts only - that famous property of the CNN. This section lays emphasis on a visual interpretation first and finalizes with the mathematical proof. Especially the latter is a mere reformulation of that which Cohen and his colleagues describe in the earlier mentioned article. Furthermore, slightly extended versions of both visual interpretation and proof are found in the blog I wrote earlier this year [70].

Visual interpretation

Consider a cactus in Fig. 5.5. We will denote a (pure) rotation and a (pure) reflection as $r()$ and $m()$. For brevity we forget about the brackets and perform both separately on the object (see what happens in the figure). Extending this further, we allow every possible combination of r and m applied sequentially. For example we can apply a reflection first, followed by two rotations. Not surprisingly, some combinations lead to equivalent poses the object can adapt to (e.g. $rm = mr^3$ and $r = e$). In a Cartesian sampling grid, the possible orientations are limited to eight. The ingenious idea is to generate graphs in which the nodes denote the possible poses and the connections between them the transformations. These graphs are depicted in Fig. 5.5 on the right, in which a red arrow symbolizes a rotation and a blue line a reflection. We define the group $p4$ as the group that includes all rotations and the group $p4m$ that includes all rotations and mirror reflections.

Mathematicians call these graphs representations of the groups if they obey to some specific rules which are more subtle than will be explained [71]. The idea is that the graph is perfectly symmetrical in some sense. From your perspective the graph is identical no matter the location you are at: sequentially taking a red, a blue, a red and a blue path will lead you back to your initial position. Understanding how graphs can illustrate groups will help you understand the way G-convolutions work and thus are implemented. Moreover, it will help you understand the proof later on.

The procedure is as follows. Start with a simple square image as input for the first G-convolution: the group $p4$ is considered. First, the input is transformed to every possible pose (being four in this case) after which you perform conventional convolution to every one of them with the same filter, creating four transformation-dependent stacks of feature maps. The main idea is to allocate every such feature map in the aforementioned graph according to the transformation it had to undertake. To reduce computational overhead, the filter is transformed, generally being much smaller, instead of the input, providing a similar outcome. Fig. 5.6 sketches this map.

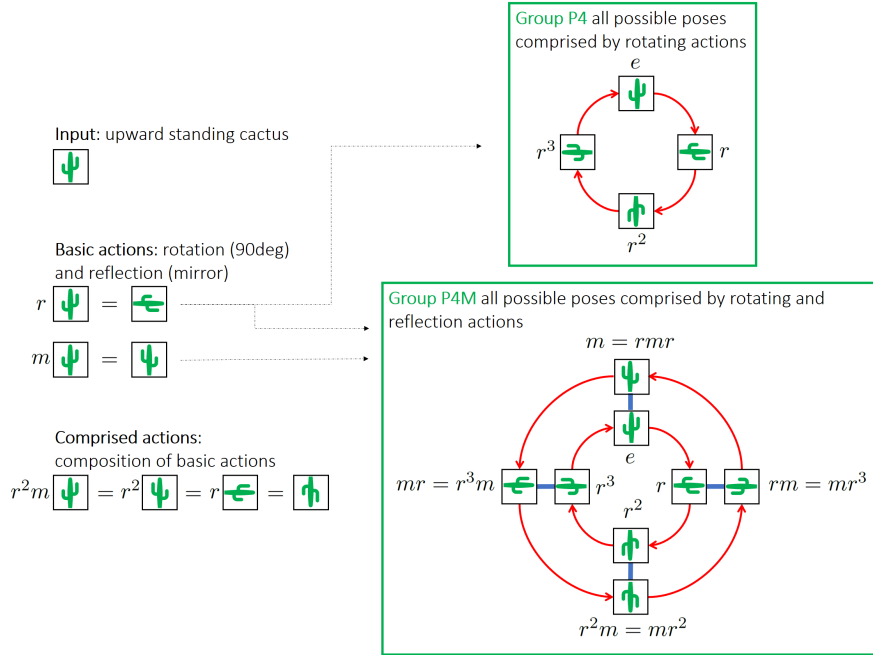


Figure 5.5: Visual interpretation of the group P4 and P4M. The group P4 comprises all possible combinations of 90-degree rotating actions (thus having dimension 4), while the group P4M includes the possibility of mirror reflections on top of that (thus having dimension 8). The graphs or webs on the right provide the roadmaps on how to move from one patch to the other: a 90-degree rotation in clockwise direction is achieved by following a red arrow and a mirror reflection by following the blue pathway.




Figure 5.6: First G-convolution.

Observe that the first G-convolution (P4ConvZ2) adds an extra dimension that resembles the group elements. Visually, it means that the individual feature maps will be placed at different locations. The overarching feature map, or output signal, is said to be structured. The specific locations at which stacks of features are placed are referred to as *patches*. Now, what happens if the inputs for the next group equivariant convolution is the structured feature map? Or, in other words, how do G-convolutions work on their group elements? The idea is similar in which an identically structured feature map is outputted. Two vital tools are needed to wholly understand what happens: the transformation of a structured object and the dot-product between two structured objects. The transformation of a structured object undergoes two actions in parallel: individual transformation of the data on every node

and following the arrows that indicate that transformation. In the case of a rotation, all data should be (1) individually rotated and (2) moved to the a different node by following one red arrow. People often call this final step permutation: the individual nodes permute locations. The dot-product of two (identically) structured objects is the pointwise summation of the results of the conventional convolutions that are taken if we overlay both there graphs.

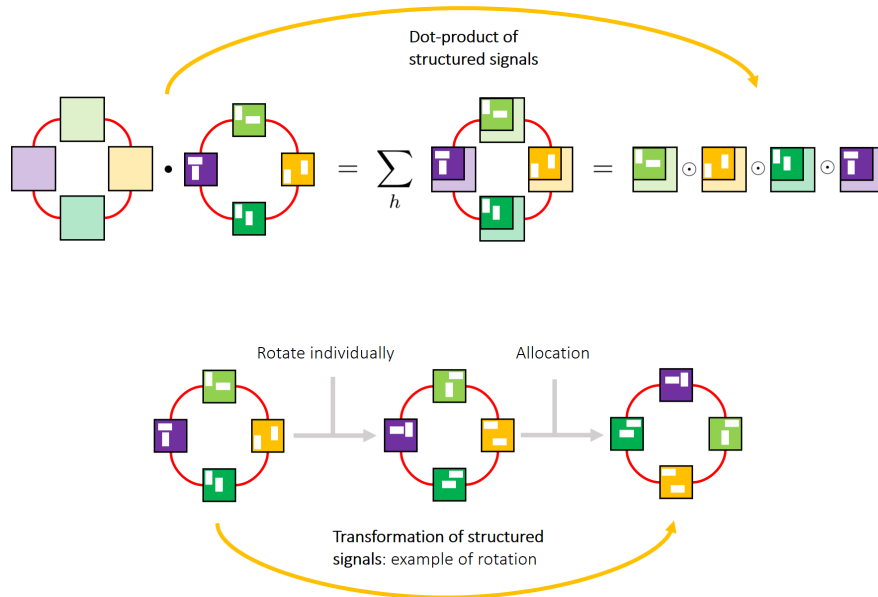


Figure 5.7: Two vital operations on group structured signals. Transformation of group structured signals (bottom) and the dot-product between two structured ones (top)

The full G -convolution (P4ConvP4) is defined next. First, a structured filter is created in which at every node the weights are independent (same used as in the previous two illustrations). We transform the structured filter for every group element, thus creating four differently transformed structured filters. We separately take the dot-product between these filters and the structured (input) feature map, thus creating four (unstructured) output objects. Finally we allocate them in a structured feature map according to the transformations the structured filter had to undertake. The P4ConvP4 is illustrated below.

Proof

So the idea is to find a group, whatever a group is, for which a function is defined, the G -convolution, that renders that function equivariant under all group actions. A symmetry transformation leaves the structure of the object invariant: e.g. if a square is rotated by 90 degrees, its structure is preserved and merely its orientation changed. For squares there is 8 such orientations. By applying one or more symmetry transformations, every orientation that preserves structure can be obtained. By definition, a symmetry group, G , is defined as a set of symmetry transformations, (a, b, \dots) , including a binary operation on $(g, h) \rightarrow gh$ that obeys the following rules:

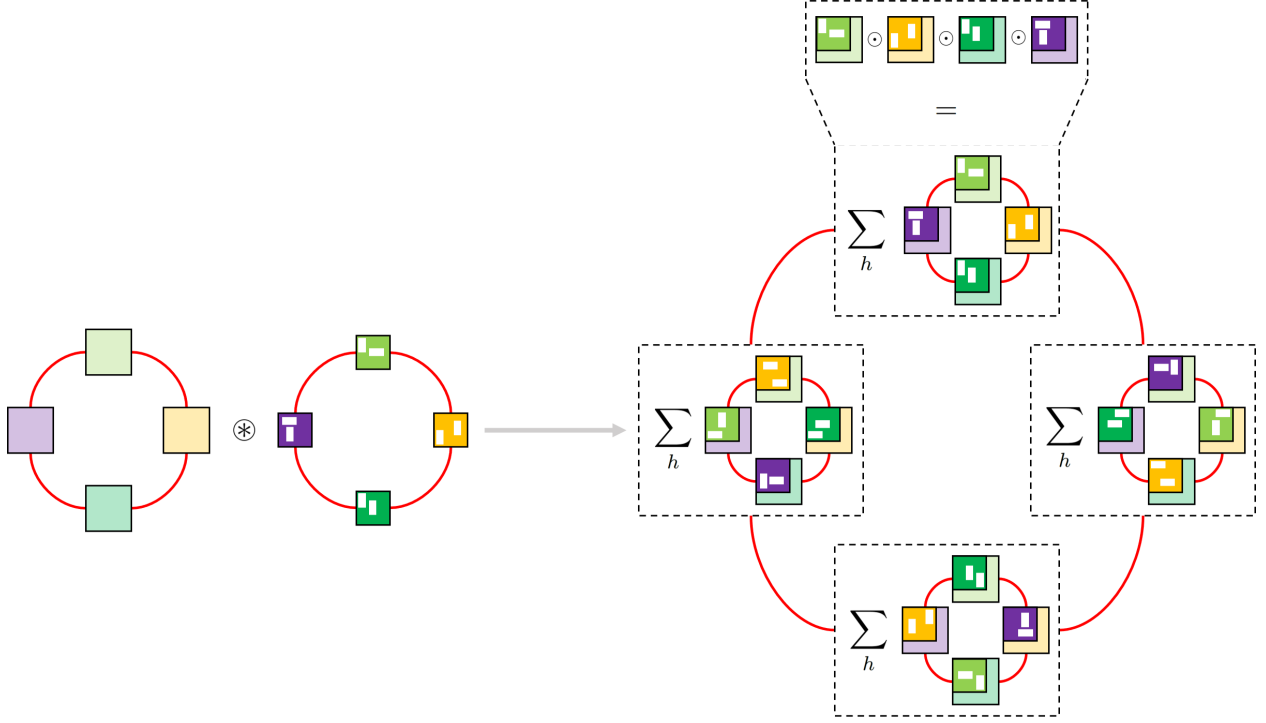


Figure 5.8: Full G-convolution.

1. **Closure:** if g and h belong to G , gh does as well
 2. **Associative:** $(gh)i = g(hi)$ for all g, h, i in G
 3. **Identity:** there exists an element e in G such that $ge = eg = g$ for all g in G
 4. **Inverse:** for every element in G there exists an inverse that is also an element of G
- (5.15)

In this context, convenient parameterizations for the group $p4$ and $p4m$ are,

$$g_{p4}(r, u, v) = \begin{bmatrix} \cos(r\pi/2) & -\sin(r\pi/2) & u \\ \sin(r\pi/2) & \cos(r\pi/2) & v \\ 0 & 0 & 1 \end{bmatrix} \quad (5.16)$$

and

$$g_{p4m}(m, r, u, v) = \begin{bmatrix} (-1)^m \cos(r\pi/2) & -(-1)^m \sin(r\pi/2) & u \\ \sin(r\pi/2) & \cos(r\pi/2) & v \\ 0 & 0 & 1 \end{bmatrix} \quad (5.17)$$

in which r is the parameter that expresses rotations, m reflection and (u, v) translations - $m, r, (u, v) \in \{0, 1\}, \{0, 1, 2, 3\}, \mathbb{Z}^2$. Observe that the group $p4m$ is simply an extension of the group $p4$. The group operation is defined as the matrix multiplication, that is

$$g_{p4}(r, u, v)x = \begin{bmatrix} \cos(r\pi/2) & -\sin(r\pi/2) & u \\ \sin(r\pi/2) & \cos(r\pi/2) & v \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u' \\ v' \\ 1 \end{bmatrix}. \quad (5.18)$$

in which (u, v) is the original location in the sampling grid. Please prove for yourself that they indeed obey the four rules. We will use these groups to define the G-convolutions. We have already shown in Eq. 5.14 that conventional convolutional layers are equivariant for the shift operations. As we will see in the definition of the full G-convolution, we must deal with group actions that work on group elements (remember that the sampling grid is part of that group, a subset). This is solved by replacing x , an element of the sampling grid, with h , an element of G . To understand what happens if a transformation works on a group element, visual interpretation becomes to some degree necessary, or at least helpful. Generally speaking (read for both the groups), h has 3 coordinates - 1 transformation coordinate (pose/orientation) and 2 translation coordinates (x). This new dimension was clarified with the introduction of a graph in which every node resembles a different pose or orientation which indicates the transformation coordinate. If a transformation is applied to a function on one of the groups, all three coordinates (can) change. As a matter of facts, I already provided an example with the transformation (specifically rotation) of a structured object in Fig. 5.7. In this case the four images allocate themselves by following 1 red arrow while simultaneously undergoing a 90-degree rotation. The G-convolution for the first layer is found by replacing the translation by a more general transformation, g , that belongs to G :

$$\begin{aligned}
[f \star \Phi^i](x) &= \sum_{y \in \mathbb{Z}^2} \sum_{k=1}^{K^l} f_k(y) \Phi_k^i(y - x) \\
&= \sum_{y \in \mathbb{Z}^2} \sum_{k=1}^{K^l} f_k(y) \Phi_k^i(t^{-1}x) \\
&= \sum_{y \in \mathbb{Z}^2} \sum_{k=1}^{K^l} f_k(y) [L_t \Phi_k^i](x).
\end{aligned} \tag{5.19}$$

The new feature map is a function on the discrete group G . For the full G-convolution we must replace y with h (an element of the group G). The full G-convolution is defined as

$$[f \star \Phi](x) = \sum_{h \in G} \sum_k f_k(h) \Phi_k(g^{-1}h). \tag{5.20}$$

For both layer types we can prove that they equivary for all group actions. The main step in the proof is the replacement of h by uh such that

$$\begin{aligned}
[L_u f] \star \Phi(g) &= \sum_{h \in G} \sum_k f_k(u^{-1}h) \Phi_k(g^{-1}h) \\
&= \sum_{h \in G} \sum_k f_k(h) \Phi_k(g^{-1}uh) \\
&= \sum_{h \in G} \sum_k f_k(h) \Phi_k((u^{-1}g)h) \\
&= [L_u [f \star \Phi]](g).
\end{aligned} \tag{5.21}$$

Similarly we can prove that individual feature maps for a conventional convolution are not equivariant under rotation because

$$\begin{aligned}
[L_r f] \star \Phi(x) &= \sum_{y \in \mathbb{Z}^2} \sum_k f_k(r^{-1}y) \Phi_k(y - x) \\
&= \sum_{y \in \mathbb{Z}^2} \sum_k f_k(y) \Phi_k(ry - x) \\
&= \sum_{y \in \mathbb{Z}^2} \sum_k f_k(y) \Phi_k(r(y - r^{-1}x)) \\
&= \sum_{y \in \mathbb{Z}^2} \sum_k f_k(y) L_{r^{-1}} \Phi_k(y - r^{-1}x) \\
&= [L_r[f \star L_{r^{-1}} \Phi]](x).
\end{aligned} \tag{5.22}$$

which states that the convolution of a rotated feature map is the same as rotating the convolution of the original feature map with an inverse rotated filter. It means that the stack of feature maps can be equivariant if the CNN learns rotated copies of all filters, but the individual feature maps will not.

5.2.3 Steerable convolutions

Steerable representations

This section is a reformulation of Cohen's work in his article *Steerable CNNs* [72]. For convolutional stages, scalar addition and multiplication of the input signal, $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^C$, are allowed. Therefore, the set of signals of this signature is a linear space \mathcal{F} ($f \in \mathcal{F}$). Note that every layer, depending on the amount of in- and output channels has its own feature space \mathcal{F}^l which will be often suppressed to reduce clutter. In order to grasp the concept of steerability and more specifically steerable representations, two distinct compositions of the feature space are mentioned. The most well-known comprises the space as a *stack of feature maps* while the other as a *bundle of feature fibers* (Fig. 5.9). Every such fiber, F_x , is the column of values at a specific position, x , having length C . Thus, without loss of information, the signal could be interpreted as composition of feature vectors $f(x)$ that live in the fibers F_x .




Figure 5.9: *Decomposition* of linear feature spaces in convolutional layers. It can be seen as a *depth-stack* of two-dimensional *feature maps* (green indicates one such feature) or as a *bundle of fibers* (yellow indicates one such fiber).

Again, we will be concerned with group transformations G that act on points in \mathbb{Z}^2 in which the operation will be denoted as $\pi(g)$ in similar fashion as previous sections. The group action is found in Eq. 5.11. Furthermore, the group actions should commute (thus obey Eq. 5.10). The feature space equipped with the set of linear commuting transformations is known as a *group representation* and its package denoted as the 2-tuple (F, π) . If we define a convolutional network $\Phi : \mathcal{F} \rightarrow \mathcal{F}'$, "the feature space F' is said to be (linearly) *steerable* w.r.t. G , if for all transformations $g \in G$, the features Φf and $\Phi \pi(g)f$ are related by a linear transformation $\pi'(g)$ that does not depend on f . So $\pi'(g)$ allows us to "steer" the features in \mathcal{F}' without referring to the input in \mathcal{F} from which they were computed." [72]. It is not hard to see that $\pi'(g)$ must also be a group representation as well⁵ since

$$\pi'(gh)\Phi f = \Phi \pi(gh)f = \Phi \pi(g)\pi(h)f = \pi'(g)\Phi \pi(h)f = \pi'(g)\pi'(h)\Phi f. \quad (5.23)$$

For clarification, Fig. 5.10 provides a visualization. The theory can be generalized, but for simplicity we will stick to discrete representation groups that thus consist of a finite amount of transformations. For purposes that will become clear, the group is divided in a part that comprises the newly added transformations (rotations and reflections), H , and a

⁵At least in the span of the *new* image that is created through the convolutional operator Φ .

part associated with the translations, Z , such that $G = \{H, Z\}$. It remains to show that filter banks constructing H -steerable output fibers that are fed as kernels in the convolution construct G -steerable feature spaces. Therefore, (H -)equivariant filter banks are developed first and subsequently I will prove by induction that the feature spaces constructed with such filter banks are steerable w.r.t. the whole representation group.




Figure 5.10: The principle of steerable filters

Equivariant filter banks

Is filter bank is intuitively seen as a linear map $\Psi : \mathcal{F} \rightarrow \mathbb{R}^{C'}$ in which C' is the feature-stack size of the next layer. If allocated at a specific⁶ position, x , it thus produces a C' -dimensional feature fiber. In this view, the correlation or convolution operator produces fiber per fiber which are retrieved by the dot-product of the filter bank with translated copies of the $f \in \mathcal{F}$ in which the filter bank is represented by a matrix of size $C' \times C \cdot k_{x_1} \cdot k_{x_2}$. Here k_{x_1} and k_{x_2} are the kernel sizes in both orthogonal directions of the sampling grid. Consequently, we have to find some representation ρ of H that acts on the output fibers such that the filter bank is H -equivariant such that

$$\rho(h)\Psi = \Psi\pi(h) \quad \forall h \in H. \quad (5.24)$$

If the representation π allows us to steer \mathcal{F} , the output of the convolution is steerable as well if some ρ is found that satisfies the equivariance constraint. It remains to find the space of maps that satisfy this equivariance constraint which is denoted as $\text{Hom}(\pi, \rho)$. Maps that live in such a space are also called *intertwiners*. The linear equivariance constraint renders this space of admissible filter banks to be a vector space. Therefore, any linear combination of intertwiners results in an intertwiner. A basis can thus be computed offline in which the coefficients are the parameters that are learned in the network:

$$\Psi = \sum_i \gamma_i \psi_i \quad (5.25)$$

6

in which ψ_i and γ_i are the basis elements and their corresponding weights respectively. In the next section (Sec. 5.2.4), specific H -equivariant filter banks are designed that work for any finite number of equally spaced rotations.

G-steerability of the whole feature space

It is shown how to parameterize filter banks such that π and ρ intertwine. At last, I will show that H -steerability of the individual fibers leads to G -steerability of the output feature space. Before the proof is shown, the group actions are specifically mentioned below, in which all actions are matrix multiplications with the augmented version of x (in which a homogeneous coordinate is added to make it applicable).

$$\begin{array}{ll}
\text{translation} & t \in Z \quad \begin{bmatrix} I & T \\ 0 & 1 \end{bmatrix} \xrightarrow{\cdot^{-1}} \begin{bmatrix} I & -T \\ 0 & 1 \end{bmatrix} \\
\text{reflection or rotation} & r \in H \quad \begin{bmatrix} R & 0 \\ 0 & 1 \end{bmatrix} \xrightarrow{\cdot^{-1}} \begin{bmatrix} R^{-1} & 0 \\ 0 & 1 \end{bmatrix} \\
\text{general element} & g \in G \quad \begin{bmatrix} R & T \\ 0 & 1 \end{bmatrix} = tr \xrightarrow{\cdot^{-1}} \begin{bmatrix} R^{-1} & -R^{-1}T \\ 0 & 1 \end{bmatrix} \\
\text{explicit translation by } x & \bar{x} \in Z \quad \begin{bmatrix} I & x \\ 0 & 1 \end{bmatrix} \xrightarrow{\cdot^{-1}} \begin{bmatrix} I & -x \\ 0 & 1 \end{bmatrix}.
\end{array} \tag{5.26}$$

It is immediately observed that the correlation operator is defined as:

$$[\Psi \star f](x) = \Psi \pi(\bar{x}^{-1}) f \tag{5.27}$$

which again states that the filter bank acts on translated copies of the feature maps. The following equality will be necessary and follows directly from the definition of the group actions:

$$\begin{aligned}
(tr)^{-1} \bar{x} r &= \begin{bmatrix} R^{-1} & -R^{-1}T \\ 0 & 1 \end{bmatrix} \begin{bmatrix} I & x \\ 0 & 1 \end{bmatrix} \begin{bmatrix} R & 0 \\ 0 & 1 \end{bmatrix} \\
&= \begin{bmatrix} R^{-1} & -R^{-1}T \\ 0 & 1 \end{bmatrix} \begin{bmatrix} R & 0 \\ 0 & 1 \end{bmatrix} \\
&= \begin{bmatrix} I & R^{-1}x - R^{-1}T \\ 0 & 1 \end{bmatrix} \\
&= \overline{R^{-1}x - R^{-1}T} \\
&= \overline{\begin{bmatrix} R^{-1} & -R^{-1}T \\ 0 & 1 \end{bmatrix}} \cdot x \\
&= (tr)^{-1} \cdot x.
\end{aligned} \tag{5.28}$$

Transforming the feature space by an element of the representation group and convolve it with an H -equivariant filter bank leads to

$$\begin{aligned}
[\Psi \star \pi(g) f](x) &= [\Psi \star \pi(tr) f](x) \\
&= \Psi \pi(\bar{x}^{-1}) \pi(tr) f \\
&= \Psi \pi(\bar{x}^{-1} tr) f \\
&= \Psi \pi(rr^{-1} \bar{x}^{-1} tr) f \\
&= \Psi \pi(r) \pi(r^{-1} \bar{x}^{-1} tr) f \\
&= \rho(r) \Psi \pi(r^{-1} \bar{x}^{-1} tr) f \\
&= \rho(r) \Psi \pi\left(\left((tr)^{-1} \bar{x} r\right)^{-1}\right) f \\
&= \rho(r) \Psi \pi\left(\overline{(tr)^{-1} \cdot x}^{-1}\right) \\
&= \rho(r) [\Psi \star f]\left((tr)^{-1} \cdot x\right) \\
&= \rho(r) [\Psi \star f]\left(g^{-1} \cdot x\right).
\end{aligned} \tag{5.29}$$

This result is remarkable and shows that if we define π' as

$$[\pi'(g) f](x) = \rho(r) [f(g^{-1}x)] \tag{5.30}$$

the output feature space is G -steerable: $\Psi \star \pi'(g) f = \pi'(g) \Psi \star f$. This has happened through an important construction known as the *induced representation* in which $\pi'(g)$ is naturally induced by the representation ρ of H . It is sometimes denoted as $\pi' = \text{Ind}_H^G \rho$. Even more remarkable is the fact that the induced representation is nearly similar as a simply transforming the output of convolution: the only difference resides in the presence of the factor $\rho(g)$ which tells you how to *mix* the channels in the end. Fig. 5.11 gives an illustrative example.

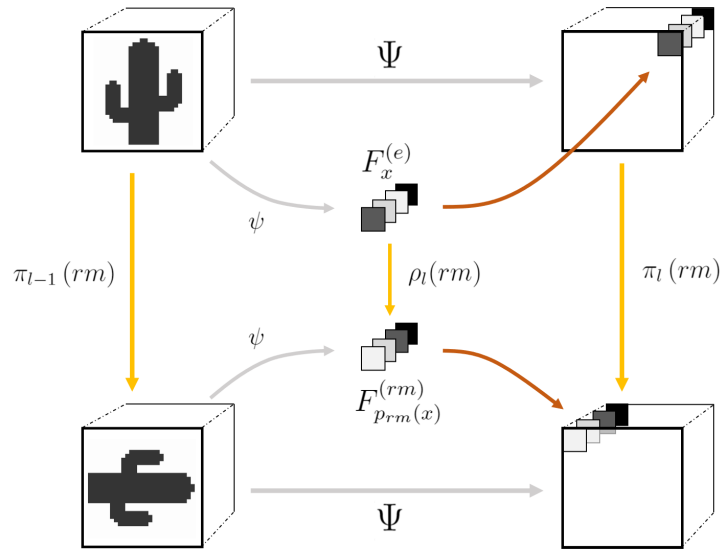


Figure 5.11: Visual clarification of steerability of the whole domain.

The reader may ask how to design such filters. The next section develops a specific implementation which was first designed and implemented in *Learning Steerable Filters for Rotation-Equivariant CNNs* [73].

5.2.4 Steerable filter design

By recapping the information of previous sections, it is concluded that if the observed images naturally adapt to specific poses for which the outcome of the network should be similar, incorporating this knowledge in the network is highly beneficial. The more of these poses are taken along in the symmetry group the network is equivariant under, or the output feature spaces have the ability to be steered in, the smaller the so called *hypothesis space* becomes. Say we aim to fit a model, $\mathcal{M} : \mathcal{I} \rightarrow \mathcal{L}$, in which \mathcal{I} and \mathcal{L} are the image and label space respectively, and if the model is structured such that its equivariant under the proposed symmetry group G , the hypothesis space shrinks to

$$\tilde{\mathcal{M}} : \mathcal{I}/G \rightarrow \mathcal{L}/G \quad (5.31)$$

in which $(\cdot)/G$ denotes the quotient space resulting from collapsing equivalent images in an orbit to a single point. This section conveys the explanation and prove of layers that develop rotationally steerable filters, being steerable in an arbitrary finite number of equally spaced rotations around an orbit which are subsequently fed into a group convolutional layer. These networks are referred to as *SFCNNs* and were first developed in [73]. The initialization of such a network is not in the scope of this research, but it should be noted that this is not trivial. A proper implementation is given in the paper above.

Learning steerable filters

In order to design rotationally steerable filters, a set of fixed fundamental functions $\{\Psi_k\}_{k=1}^K$ should span the filter bank space. By specifically expanding Eq. 5.25, the filter bank should satisfy

$$\rho(\theta) \Psi(x) = \sum_{k=1}^K \gamma_k(\theta) \Psi_k(x) \quad \forall \theta \in (-\pi, \pi] \quad (5.32)$$

in which $\rho(\theta)$ is the transformation action associated with a rotation of θ around the orbit and will for simplicity be denoted as ρ_θ . An easy basis is given by a set known as the circular harmonics:

$$\Psi_k(r, \phi) = \tau(r) e^{jk\phi}. \quad (5.33)$$

By design, it follows that

$$\rho_\theta \Psi_k = e^{-jk\theta} \Psi_k. \quad (5.34)$$

thus making the representation steerable for any arbitrary angle θ . It can only be utilized if the set is discretized in both k and τ . Furthermore, the Gaussian radial part $\tau(r)$ is further specified which results in the final set of admissible fundamental functions:

$$\begin{aligned}\Psi_{km} &= \tau_m(r) e^{jk\phi} \\ &= e^{-\frac{(r-m)^2}{2\sigma^2}} e^{jk\phi} \quad k \in \{1, 2, \dots, K\}, m \in \{1, 2, \dots, M\}\end{aligned}\tag{5.35}$$

limited to the point at which aliasing occurs: angular frequencies k (depending on m) are held therefor to a maximum. The filter bank is now defined as the stack of such a set associated with individual weights that are learnable:

$$\tilde{\Psi}(x) = \sum_{m=1}^M \sum_{k=1}^K w_{mk} \Psi_{mk}(x) \quad w_{mk} \in \mathbb{C}.\tag{5.36}$$

The complex phase allows rotating the atomic filters w.r.t. each other, but such a system can be steered as a whole by phase manipulation of the atoms via

$$\rho_\theta \tilde{\Psi}(x) = \sum_{m=1}^M \sum_{k=1}^K w_{mk} e^{-jk\theta} \Psi_{mk}(x)\tag{5.37}$$

in which we define the single orientation as its real part

$$\Psi(x) = \text{Re} \left[\tilde{\Psi}(x) \right].\tag{5.38}$$

Rotational equivariant layers

The concept is to feed these filter banks into a group convolutional layer which creates $|G|$ (:the heuristic of the group) patches that associate with every proposed rotation it should be equivariant under. It should be clear from previous section on GCNN that the first group convolutional convolves over the trivial field and adds a rotational coordinate, $\Phi_{0 \rightarrow 1} : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2 \times \Theta$, and that the full group convolution convolves over all patches, $\Phi : \mathbb{Z}^2 \times \Theta \rightarrow \mathbb{Z}^2 \times \Theta$. Here Θ is the set of equidistant orientations:

$$\Theta = \left\{ 0, 2\pi \frac{1}{\Lambda}, 2\pi \frac{2}{\Lambda}, \dots, 2\pi \frac{\Lambda-1}{\Lambda} \right\}\tag{5.39}$$

in which $\Lambda = |G|$ is the amount allowed orientations and a choice for the designer.

The un-activated output signal is defined as

$$\begin{aligned}z_{\hat{c}}^{(1)}(x, \theta) &= \sum_{c=1}^C \left(\rho_\theta \Psi_{\hat{c}c}^{(1)} \star f_c \right) (x) \\ &= \sum_{c=1}^C \left(\text{Re} \left[\sum_{m=1}^M \sum_{k=1}^K w_{\hat{c}cmk} e^{-jk\theta} \Psi_{mk} \right] \star f_c \right) (x) \\ &= \text{Re} \left[\sum_{c=1}^C \sum_{m=1}^M \sum_{k=1}^K w_{\hat{c}cmk} e^{-jk\theta} [\Psi_{mk} \star f_c] (x) \right].\end{aligned}\tag{5.40}$$

It states that the next pre-nonlinearity feature can be computed as a linear combination of all individual convolutions with the fundamental orientations. Rotational weight sharing is reflected by phase manipulation of the weights:

$$w_{\hat{c}cmk}e^{-jk\theta}. \quad (5.41)$$

To conclude, the activated output feature is computed by adding a bias and the introduction of a pointwise non-linearity as. Fig. 5.12 provides a visual clarification for both the fundamental filters as well as the computation of the un-activated output signal.

$$a_{\hat{c}}^{(1)}(x, \theta) = \sigma \left(z_{\hat{c}}^{(1)}(x, \theta) + b_{\hat{c}}^{(1)} \right) \quad (5.42)$$

in which σ and b denote the non-linear function and bias respectively.

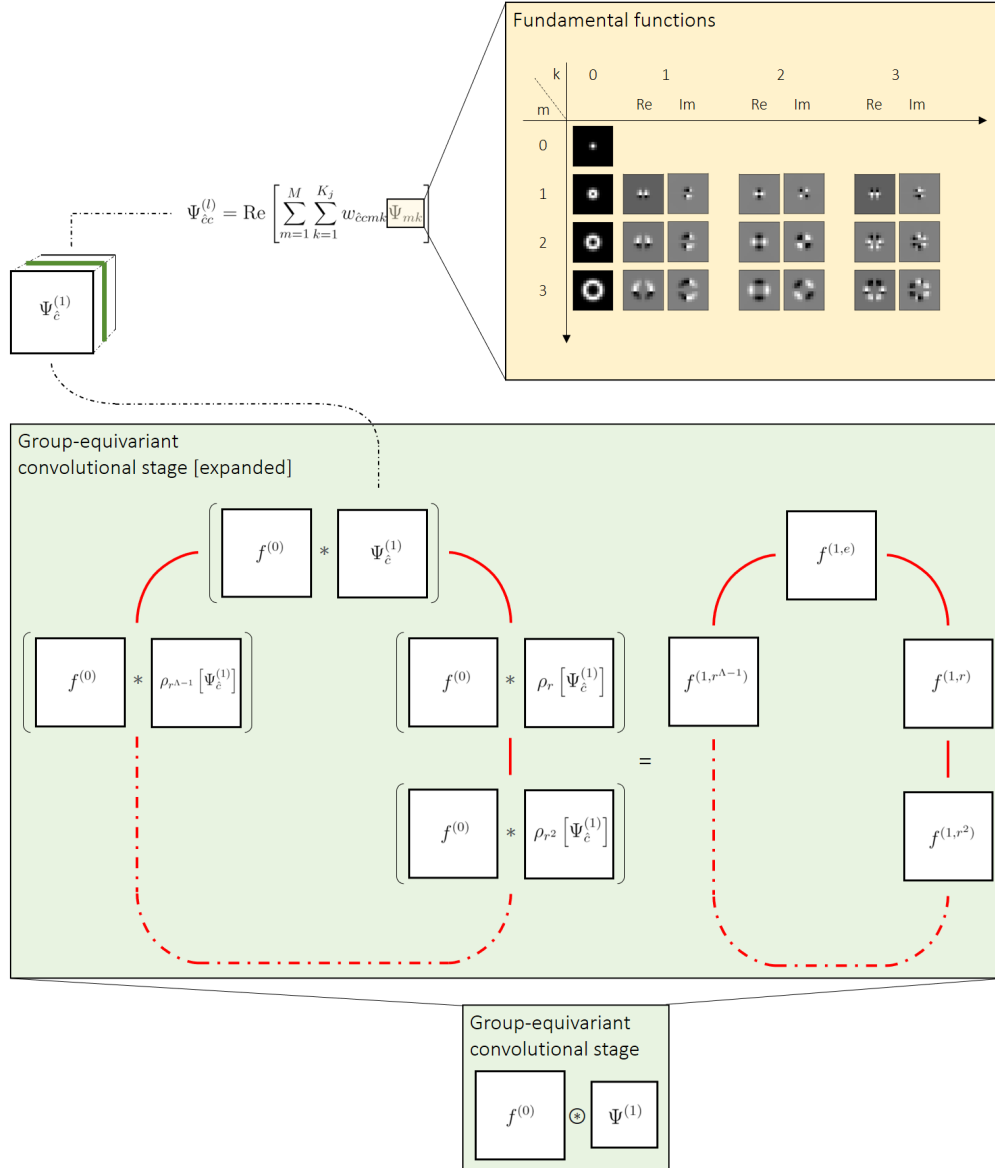


Figure 5.12: First group-equivariant convolutional stage of the SF-GCNN.

For the full group convolution one is concerned with maps that act on the semi-direct group $\mathbb{R}^2 \rtimes \theta$ that besides spatial coordinates is equipped with a coordinate that reflects the discrete set of rotations. A full group convolution in the continuous sense is given by

$$(\Phi \circledast \zeta)(g) = \int_G \zeta(h) \Psi(h^{-1}g) d\lambda(h) \quad (5.43)$$

in which $\lambda(h)$ is a Haar measure. The full group convolution that acts on the input feature space can be expanded as,

$$\begin{aligned} z_{\hat{c}}^{(l)}(x, \theta) &= \sum_{c=1}^C [\Psi_{\hat{c}c}^{(l)} \circledast a_c^{(l-1)}](x, \theta) \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \int_{\mathbb{R}^2} a_c^{(l-1)}(x, \theta) \Psi_{\hat{c}c}^{(l)}((u, \phi)^{-1}(x, \theta)) du \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \int_{\mathbb{R}^2} a_c^{(l-1)}(x, \theta) \Psi_{\hat{c}c}^{(l)}(\phi^{-1}(x-u), \theta - \phi) du \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \int_{\mathbb{R}^2} a_c^{(l-1)}(x, \theta) \rho_\phi \Psi_{\hat{c}c}^{(l)}(x-u, \theta - \phi) du \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} [\rho_\phi \Psi_{\hat{c}c}^{(l)}(\cdot, \theta - \phi) * a_c^{(l-1)}(\cdot, \phi)](x) \\ &=: \sum_{c=1}^C \sum_{\phi \in \Theta} [R_\phi \Psi_{\hat{c}c}^{(l)}(\cdot, \theta) * a_c^{(l-1)}(\cdot, \phi)](x). \end{aligned} \quad (5.44)$$

in which $R_\phi \Psi(x, \theta) := \rho_\phi \Psi(x, \theta - \phi)$. The reader must now be familiar with such operations. Again, on the group, the filters are defined as

$$\Psi_{\hat{c}c}^{(l)}(x, \theta) = \text{Re} \left[\sum_{m=1}^M \sum_{k=1}^K w_{\hat{c}cmk\theta} \Psi_{mk} \right](x) \quad (5.45)$$

in which θ in the weights reflects the orientation coordinate. Expanding Eq. 5.44 further concludes on the un-activated feature:

$$\begin{aligned} z_{\hat{c}}^{(l)}(x, \theta) &= \sum_{c=1}^C \sum_{\phi \in \Theta} \left[R_\phi \left[\text{Re} \left[\sum_{m=1}^M \sum_{k=1}^K w_{\hat{c}cmk\theta} \Psi_{mk}(\cdot) \right] \right] * a_c^{(l-1)}(\cdot, \phi) \right](x) \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \left[\text{Re} \left[\sum_{m=1}^M \sum_{k=1}^K w_{\hat{c}cmk\theta - \phi} e^{-jk\phi} \Psi_{mk}(\cdot) \right] * a_c^{(l-1)}(\cdot, \phi) \right](x) \\ &= \sum_{c=1}^C \sum_{\phi \in \Theta} \left[\text{Re} \left[\sum_{m=1}^M \sum_{k=1}^K w_{\hat{c}cmk\theta - \phi} e^{-jk\phi} \left[\Psi_{mk}(\cdot) * a_c^{(l-1)}(\cdot, \phi) \right] \right] \right](x). \end{aligned} \quad (5.46)$$

It becomes once more clear that implementation can be done in clear separate steps: spatial convolution (indicated in **red**), rotations (indicated in **orange**) and linear combinations (indicated in **olive green**). Adding a bias and point-wise linearity in similar fashion as before, the activated output feature is easily computed by

$$a_{\hat{c}}^{(l)}(x, \theta) = \sigma \left(z_{\hat{c}}^{(l)}(x, \theta) + b_{\hat{c}}^{(l)} \right). \quad (5.47)$$

6 | Design, training and implementation of deep networks

In order to *classify* thus *predict* if a segmented red blood cell is infected or not, it either has to be done manually or computer-based. The design of the latter is the core of this chapter in which a deep learning model is designed into which the operations described in the previous chapter are being implemented.

6.1 Deep learning: residual networks

The group convolutions will be implemented in a residual network architecture, popularly known as *ResNets* [74]. The main property of this type of architecture is the flowthrough of residuals that *prevents vanishing gradients* in networks that are (said to be) deep, providing the ability to compute these gradients at near precision. Furthermore, the architecture allows the *direct interchangeability* of conventional convolutions with layers that have an increased amount of transformations under which they equivary without the necessity to change other parts of the network while remaining invariant as a whole. Ergo, it makes it a suitable architecture to test performance for different operations that were developed in Ch. 5.

6.1.1 The residual block

ResNets consist of residual blocks in which the input data at the start splits and *parallelly flows* through two branches that merge again in the end. The branches are asymmetric in which one performs minimal action and the other provides all the essential operations that comprise all learnable weights¹. The residual block is the main building brick for every ResNet: the bricks are generally piled up, forming a (high) 1-brick-wide building in which the data falls from the uppermost sky lounge to the ground floor in feedforward mode and elevates back up to the top in the backward propagation step. Fig. 6.1 depicts the most general form. Since every residual block is equivariant if its individual operations have that feature, a stack of those remains equivariant as well.

Even though the bigger pictures is sketched, it needs specifications. Relevant questions could

¹As one will see, this is not entirely true as the branch that performs minimal action is sometimes equipped with a (1×1) convolution which is necessary if grid downsampling and/or channel expansion is wanted. In this case, the branch consists, although being nearly negligible, of some parameters.

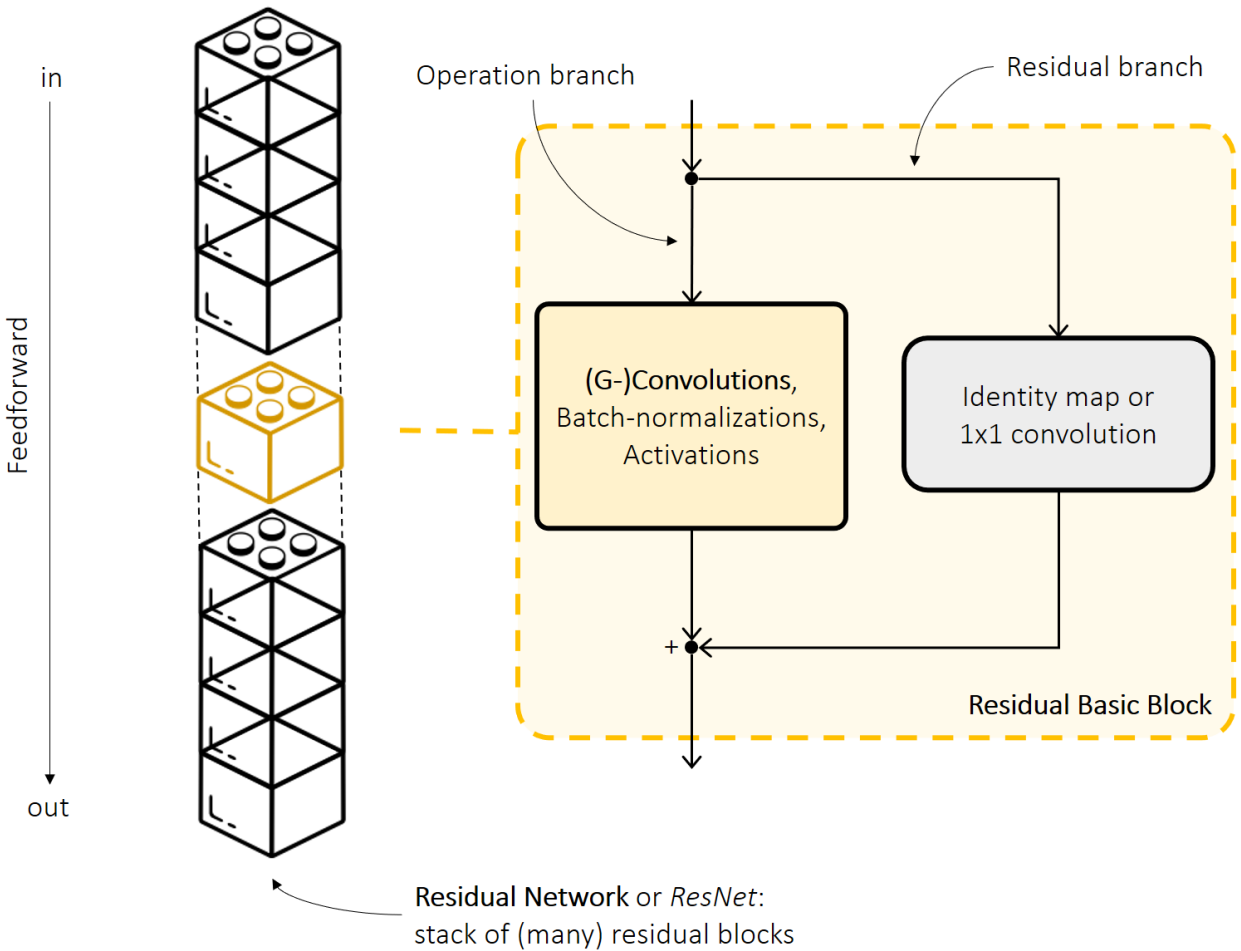


Figure 6.1: General form of the residual unit, or residual block (fenced by yellow rectangle). A ResNet is the consecutive follow-up of many residual blocks (left). The input parallelly flows through two independent lanes: the *operation branch* and *residual branch*. Nearly all learnable weights and thus operations are located in the operation branch while the data has negligible resistance in the other: the residual branch exclusively conveys trivial operations.

be for example about the amount of residual blocks, the amount of channels per convolution, the initiation and/or ending of the network, the type of batch-normalizations, activations, and pooling, expansion and/or downsampling throughout the network. All answers are reflected in the choice of both *architectural* and *function hyperparameters*. First of all, the answers are highly correlated with the application it is used for, the amount of computer memory and power you have access to, the maximum training time you wish the network to converge in and, unfortunately, the experience you have. In time, you will gain better intuition of what to use in what situation, and until this day it is hard to provide profound arguments how to do so. Secondly, and more important, this research is *not* dedicated to the choice and evaluation of such parameters. Therefore, model hyperparameters are chosen

that will *not* differ between implementations². In the remainder of this section, an explicit implementation is described that is used in this research.

Every operation branch in a residual block consists of N_{cl} convolutional layers, in which batch-normalization and (non-linear) activation always follow after in that respective order. The final, thus N_{cl} -th, activation is allocated slightly different in that the output of the residual branch is first added instead. In my case, the residual branch either performs an identity mapping or 1x1 convolutions if downsampling of the image domain and/or expansion of the amount of feature maps is needed.

6.1.2 Encoding the network

It has been shown that expanding the network, that is the increment of the amount of feature channels, at several points in the network benefits performance [74]. That part of the network that diverges in terms of width is called the *encoder*. In ResNets, those changes in width are reflected in *residual layers* which are stacks of equally wide residual blocks. Each residual layer conveys N_{rb} residual blocks and performs feature expansion and domain downsampling in its first block only. More specifically, except for the first residual layer, residual layers start with a residual block that reduces the size of every feature's domain (stride/pooling is 2) and increases the amount of channels (expansion is 2)³. The amount of consecutive residual layers is denoted as N_{rl} . To wrap up the *encoding* part, it remains to initiate the network. This is the follow up of an initial convolution (usually equipped with a relatively large kernel, e.g. 7x7 or 9x9), a batch-normalization and an activation. Here, the initial convolution acts on the image input field which, depending on the type of input (e.g. monochrome, RGB, hyperspectral), has N_{in} input features.

6.1.3 Decoding the network

As the network should have the ability to classify images, convergence at some points is mandatory. The part that does exactly that is commonly referred to as the *decoder*. Intuitively, the decoder can be seen as the part that learns what combinations of two-dimensional features (comprised in the output of the last residual layer) belong to which class: it *decodes* how the features should be *mixed* for every class. Therefore, the decoder typically consists of fully connected layer (allowing every linear (activated) combination of features). More specifically, it conveys two such layers which both consecutively activate their outputs. As the classification should be *invariant* under transformations of the input, the final output features must be shrunk to scalars before the first fully connected layer. To do so, it first undergoes *point-wise anti-aliased average pooling* of the whole sampling domain (global) and subsequently *point-wise maximum pooling* (global). The remaining stack of scalar-valued features can now be fed into the first fully connected layer.

²The amount of channels per convolutional layer will though differ per implementation. This accounts for the total amount of trainable parameters which should be held at the same value. A detailed discussion is provided at the end of this section (see 6.1.5).

³Observe that by definition of the entrance domain and width, those in the next residual layers follow directly: therefore, they do not need specification.

6.1.4 Overall architecture

An overview is given in 6.2. Furthermore, all model hyperparameters including their choices are summarized in Tab. 6.1.

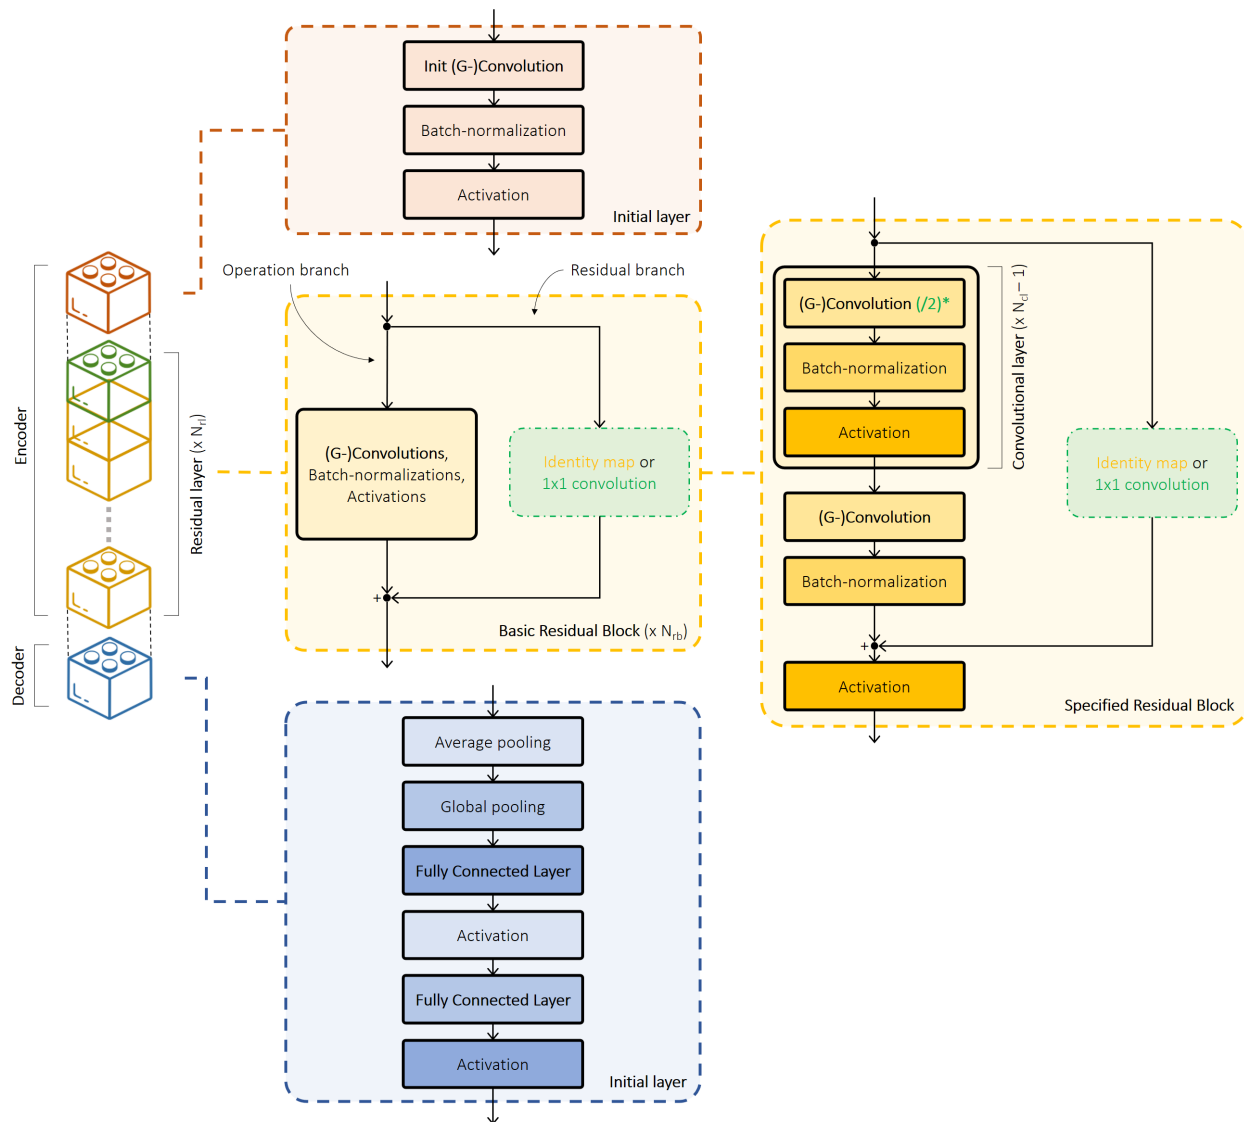


Figure 6.2: Overview of a residual network. The initial layer (red block) conveys an initial convolution that maps between the input field and the field over which the (G-)convolution convolves and subsequently normalizes the batches and activates the remaining signal through an pointwise activation function. The output of the initial layer is fed into the most crucial part of the network which consists of (many) consecutive residual blocks (green and yellow). After every N_{rb} residual blocks, the signal is expanded in terms of channels while simultaneously the domain is downsampled (green blocks). After exiting the last residual layer, the decoder lets the signal converge into a vector from which the predicted classes are extracted. This final element conveys two fully connected layers both directly followed by pointwise activations.

Table 6.1: Models hyperparameters

Table 6.2: Architecture

Explanation	Symbol	Value
kernel size of initial filter	k_x^I	7
kernel size of filters	k_x	5
Amount of hidden nodes in decoder	N_H	250
Amount of residual blocks per residual layer	N_{RB}	2
Amount of convolutions per residual block	N_{CL}	2
Amount of residual layers	N_{RL}	4
Amount of input fields	C^{in}	32
Spatial domain of the input	HW^{in}	varied

Table 6.3: Function

Explanation	Symbol	Value
Activation function	ν	ReLU/ELU
Batch-normalization	BN	
Average pooling	avgpool	Average anti-aliased pooling of whole spatial domain
Global pooling	globpool	Global pooling over all patches ($H = G/\mathbb{Z}^2$)

6.1.5 Remarks

Size of the networks

In order to evaluate the proposed operations, networks are tested that have the same architecture and the same *size*. I would like to emphasize that architecture and size are inherently different. Size is solely determined on the amount of learnable parameters. Size clearly depends on the architecture but it does as well on the type of operations. Recall that G-convolutions use filters that are structured and thus have an increased amount of weights compared to conventional convolutions (if the amount of channels and kernel sizes remain unchanged). In a conventional convolutional layer the amount of learnable parameters is retrieved as

$$S_l^{\text{CNN}} = k_x^2 \cdot C^{l-1} \cdot C^l \quad (6.1)$$

in which k_x is the size of the filter (usually 3 or 5). For a G-convolutional layer this amount linearly increases with the heuristic of G, $|G|$:

$$S_l^{\text{G-CNN}} = |G| \cdot k_x^2 \cdot C^{l-1} \cdot C^l. \quad (6.2)$$

To account for this increase, the amount of channels is instead reduced accordingly by $|G|^{1/2}$, as downsizing the filter domain is undesirable. In case of steerable filters, the amount of learnable parameters is computed as

$$S_l^{\text{SF-CNN}} = |G| \cdot |W| \cdot C^{l-1} \cdot C^l. \quad (6.3)$$

Here, $|W|$ connotes the measure for the amount of fundamental filters. Correction is done similarly in which the channels are reduced by $(|W|/(|G|k_x^2))^{1/2}$. Tab. 6.4 provides a summary of the above and provides the design choices and their consequences.

Table 6.4: Size and relations between operations.

	CNN	p4(m)-CNN	SF-CNN
Amount of parameters per layer, l	$k_x^2 \cdot C^{l-1} \cdot C^l$	$ G \cdot k_x^2 \cdot C^{l-1} \cdot C^l$	$ G \cdot W \cdot C^{l-1} \cdot C^l$
Relations between implementations	Choose: $k_x, \{C^l\}$	$C^l \rightarrow \lceil \left(\frac{C^l}{ G ^{1/2}} \right) \rceil$	$C^l \rightarrow \lceil \left(\frac{ W ^{1/2}}{ G k_x^2} \cdot C^l \right) \rceil$
Explicit choices	k_x, C^l : see Tab. 6.1	k_x, C^l : see Tab. 6.1 $ G $: see Sec. 5.2.2	k_x, C^l : see Tab. 6.1 $ G , W $: Sec. 5.2.4
Total parameters ResNet-18 or RN18	1.91M	1.90M (all)	1.90M (all)

Batch-normalizations

Depending on the input distribution of the mini-batches (e.g. differently retrieved data or non-normalized batches), machine learning algorithms behave and update differently. The *covariate shift* refers to this change. Accelerating the training process by reducing this covariate shift through normalizing the activations of each layer is the concept known as *batch-normalization*, or BN. Practically it computes both mean and variance of the mini-batch's signal while allowing scales and shifts of its normalized counterpart in its computation in the form of learnable parameters. The output, y_i , is retrieved as

$$y_i = \gamma \hat{x}_i + \beta := BN_{\gamma, \beta}(x_i) \quad i = 1, \dots, m \quad (6.4)$$

in which $\mathcal{B} = \{x_1, \dots, x_m\}$ represents the mini-batch comprised of m (batchsize) signals, \hat{x}_i the normalized version of x_i , and γ and β the learnable parameters for the scale and shift respectively. Although debate around "why BNs are working" is still in its midst, it is shown that it enhances the procedure in most cases [75].

6.2 Training the networks

Training of neural networks is a challenging task. It involves the choice in cost-function, optimizer and other training-hyperparameters, such as the learning rate and length of training. Following subsections are about those choices.

6.2.1 Cost function

Comparing and validation are essential for training. In 5.1 and especially Eq. 5.4, I explained how prediction and/or estimation can be measured w.r.t. the ground truth, or label, of the

sample using a cost function. In this research, cross-entropy is used as measure,

$$f_C(\mathbf{x} : \mathbf{a}^L, y) = -\log\left(\frac{e^{\mathbf{x}[y]}}{\sum_j e^{\mathbf{x}[j]}}\right) = -\mathbf{x}[y] + \log\sum_{j=1}^{|\mathbf{y}|} e^{\mathbf{x}[j]}, \quad (6.5)$$

which is a combination of (1) *SoftMax* that maps the logits (output feature vector of the ANN) to probabilities that associate with every class and (2) an entropy function that evaluates all contributing probabilities based on a logarithmic measure (see Fig. 6.3). This objective function is continuously differentiable which is necessary for the back-propagation, is not complex, and, most important, works appropriately for classification purposes [76].

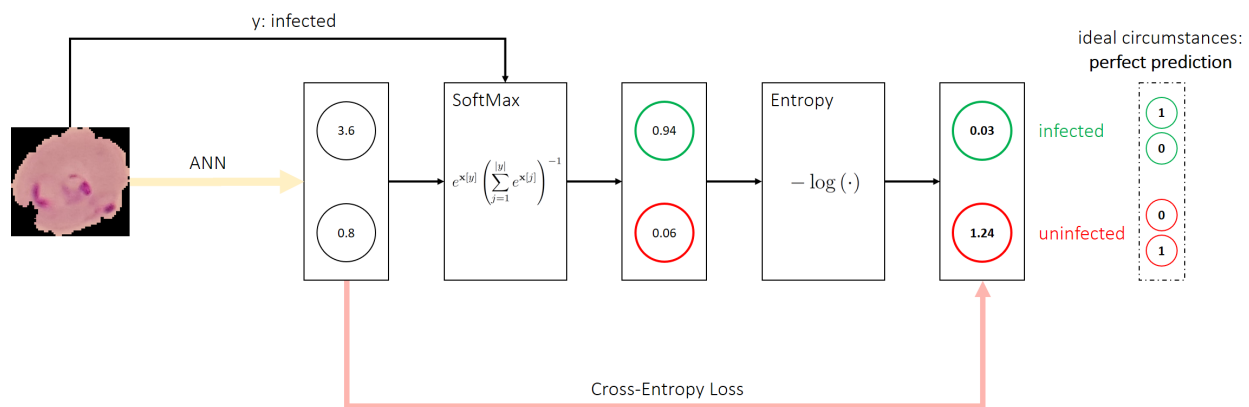


Figure 6.3: Visualization of the cross entropy loss: infected and uninfected red blood cells.

6.2.2 Activation function

Activation in CNNs make it possible to learn non-linear relationships between in- and output, therefore being crucial in the process. As specified in Tab. 6.1, the non-linear rectified linear unit, or *ReLU*, is used as activation function:

$$\text{ReLU}(x) = \max(x, 0). \quad (6.6)$$

The main benefits are immediately apparant: (1) easiness of computation and (2) gradient is very consistent without second-order effects. Therefore, it makes the network relatively fast and provides a useful gradient direction [77]. In theory, the major drawback is the uncontinuity of the gradient at 0 which is 0 if approached from the left and 1 if approached from the right. In practice though, one can safely disregard such phenomena as digital computers are prone to numerical errors: it is unlikely that the evaluation at the point 0 truly returns 0. There are many "mutations" of this simple function that correct for the non-differentiability or the fact that values below zero are treated as non-active immediately. Even though in some cases it could be slightly beneficial to use altered versions, its simple form mostly performs comparably well [77].

6.2.3 Sample and batch size

Learning needs observations, actions, facts, etc. In computer vision, networks learn on the basis of, mostly many, images. As we, humans, and all other lifeforms on earth that have the ability to *see*, learn on the basis of just a *few* images, the best artificial networks are not even close to achieve such remarkable perfection and need *many* training samples instead. Depending on the type and size, well-performing networks that have mere negligible error rates (<0.1 %, comparable to those of their biological ancestors,) need at least a few thousands training samples per class. The images upon which networks *train* are comprised in the trainset and are referred to as training samples. The same holds for the collection of images that are used for *testing*: the testset consists of testing samples. For deep models, both train- and testset must be huge to achieve state-of-the-art performances. More specifically, between four and ten thousand training samples is mostly enough to squeeze out optimal performance. On the other hand, the testset need not be that large: around a thousand samples render the variance of the outcome imperceptible.

Updating the network for every individual sample separately is time consuming and has the tendency to get networks stuck in local minima. *Bundling* many training samples into *batches* and update on their average contributions has shown to increase the available computational parallelism, reduce overfitting and increase steps and time of convergence [77]. On the contrary, some research, backed by Yann LeCunn, shows that smaller batch-sizes have improved generalizing properties [78].

6.2.4 Stochastic gradient descent

Optimizing the updating scheme plays an important role in many training procedures. A general form, as a function of the gradient, was given earlier in the first line of Eq. 5.5 and its most simple form in the second line: *gradient descent*. Gradient descent updates based on only one training sample at a time. In order to speed up the process, one could draw randomly picked (i.i.d.) minibatches and compute its average gradient. Statistically, one obtains lesser biased estimates which, loosely speaking, determine more realistic gradient. The updating scheme for *stochastic gradient descent*, or *SGD*, is given by

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \epsilon_k \frac{1}{m} \nabla_{\boldsymbol{\theta}_k} \left(\sum_{i=1}^m f_C(\boldsymbol{\theta}_k^{(i)}, \mathbf{y}^{(i)}) \right). \quad (6.7)$$

in which **orange** and **red** indicate the **averaging** and **summation** over all samples in the minibatch respectively. m is the batchsize (see previous subsection) and ϵ_k is the learning rate that depends on the current iteration. It is necessary to gradually decrease the learning rate over time since the SGD gradient estimator insurmountably introduces a source of noise. It means that the gradient will *never* vanish, even upon arrival at a minimum. Setting the learning rate "is more of an art than a science, and most guidance on this subject should be regarded with some skepticism" as *Ian Goodfellow* states in *Deep learning* [79]. Instead of following predefined protocols, he suggests to closely monitor the *learning curve*: the curve of the value of the *cost*, *training error* and/or *testing error* as a function of *time*. He elaborates that is good practice to closely investigate the rate at which is converges, the oscillatory

behavior and the general trend the curvature follows. A *healthy* learning curve in his, and many other's, eyes relates to near convergence in the range of 100 iterations, the allowance of gentle oscillations (huge oscillatory behavior is unwanted) and a declining decreases of the cost.

6.2.5 On the evaluation of the training process

The previous subsection already touched upon a way to monitor the learning process through investigation of the learning curve. In many cases, as batch-sizes may differ, time is measured in *epochs*. An epoch is defined as the moment at which the whole training set has been used for updating. Especially in deep learning, the epoch is a proper choice because their numbers do not get out of hand: between 25 and 300 epochs are common limits on its minimum and maximum. Setting the learning rate and schedule its value throughout the training procedure is hard. The learning rate can be held *constant*, dropped by a factor every few epochs (*step decay*), or being controlled by another function in epochs (*time-based decay*). Furthermore, more advanced algorithms that look at the "state" the model is in, can be implemented as well, among which are *Adagrad*, *RMSprop*, *Adadelta* and *Adam* [80, 81]. Usually though, step decay is chosen and proven to work really well in most cases [82]. In the previous subsection, I laid emphasis on the notion of a *healthy* learning curve which will be more specifically addressed in the following.

Convergence

The fact that the losses and error rates converge to minima, referred to as *convergence*, is a typical property learning curves (need to) have. In doing so they decrease the loss which (mostly) leads to increments in the models *accuracy*: the percentage of correct predictions. It is vital to use "mostly" here as will be clearer soon. Before I elaborate on this, three distinct curves need mentioning: (1) the loss, (2) the training accuracy/error, and (3) the testing accuracy/error. Training accuracy is the amount of correct "predictions" of the trainset and the testing accuracy is the amount of correct predictions of the testset. As the networks learn from training samples, one usually observes that training errors are lower than those of the testset. The true "goodness" of a network is therefore best measured on the basis of the testing accuracy as it is independent of anything the network has already *seen*. Typically, both have a similar *shape* in which that of the testset mimics that of the trainset (only slightly lower). In those cases, the shape alone should be investigated. (Too) fast convergence is mostly due to the convergence into "non-optimal" local minima emerging on the higher planes of the cost's landscape, while slow convergence leads to undesirable long training procedures. On the other hand, in cases in which both curves take different shapes, other tools and notions are needed to understand the network's behavior.

Overfitting, underfitting and the good fit

In many cases, networks learn *too well* or, more annoyingly, *not enough*. Observe the illustrative example in Fig. 6.4 in which a line should be found that "best" classifies which region belongs to which class. As *Robert M. Pirsig* beautifully describes in *Zen: and the art of*

Motorcycle Maintenance, we humans *know*, based on his abundantly mentioned term *Quality*, what should be the right fit, or at least have an idea how it could be sketched or what *not* must be chosen [83]. Networks, without the premise that they ever will, have no clue up till today. Features or properties that play significant roles in the discrimination between "good" and "bad" fits are therefore hard to explicitly write down.

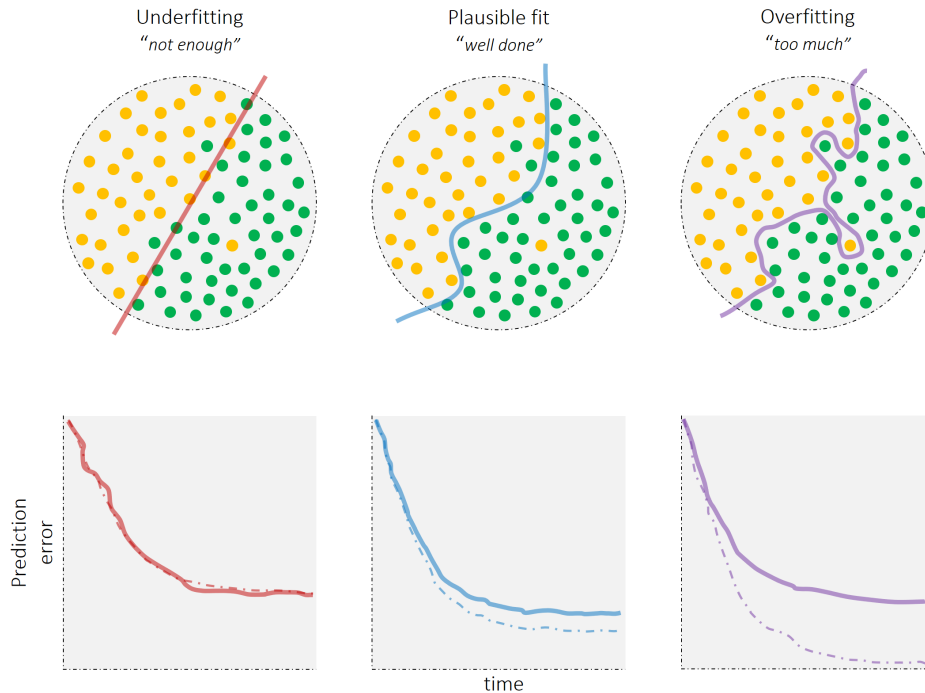


Figure 6.4: Overfitting and underfitting in machine learning.

A good start though is the division of "bad" into (1) "too well" and (2) "not enough". A network that learns too much or too specific is said to *overfit* the data and a network that does not learn enough is said to *underfit* it. By comparing the learning curves that associate with both "bad" types, it has been observed that networks that underfit the data have near identical curves for train and test while networks that overfit have the tendency to have a large, with time increasing, gap between the two [79]. Finding the *optimal* solution, in our case the *best* boundary, is one of the hardest tasks in machine learning and deserves much attention. In earlier parts of this report, I mentioned the term *hypothesis space* which is best seen as the space of allowed models the network can adapt to. On the one hand, one aims to have a large hypothesis space such that options are *open* and that nothing is missed. On the other hand, proper generalization could reduce the hypothesis space or make networks search in the space in a smarter manner. A common name for those strategies is *regularization* and is defined as "any modification we make to a learning algorithm that is intended to reduce its generalization error but not its training error" [79]. There are many ways to do so that include the incorporation of a priori knowledge, such as its embodiment in efficient sharing of weights, or include the extra addition of terms in the objective function that lay *soft* constraints on the weights [79]. In this context, it becomes even clearer what the capability of smart operations could be.

7 | Evaluation and results

This chapter is about the evaluation and results on the classifying models described in Ch. 6. First, the used databases are mentioned in Sec. 7.1. Sec. 7.2 and 7.3 give both the performance measures the networks are being validated on and the validation procedure itself. All results are included in the final section (Sec. 7.4).

7.1 Data sets and manipulations

Data is a mandatory ingredient for ANNs to improve on performance. Without it, learning is hopeless. Ideally, a data set is (1) large, having a few thousands of samples per class, (2) is properly representing the physical world or whatever it should describe¹, and (3) has sufficient resolution such that details are perceptible. In the case of infected-uninfected classification of red blood cells, many individual, or *segmented*, cells (samples) are necessary (for both classes). Before arrival upon a decent data set, several steps are to be taken upfront. The diagram in Fig. 7.1 successively conveys *data acquisition*, the retrieval of the blood films, *pre-processing*, such as calibration, noise reduction, and normalization, *segmentation*, the cut-out of individual cells, and finally *organization*, the development of structured data sets.

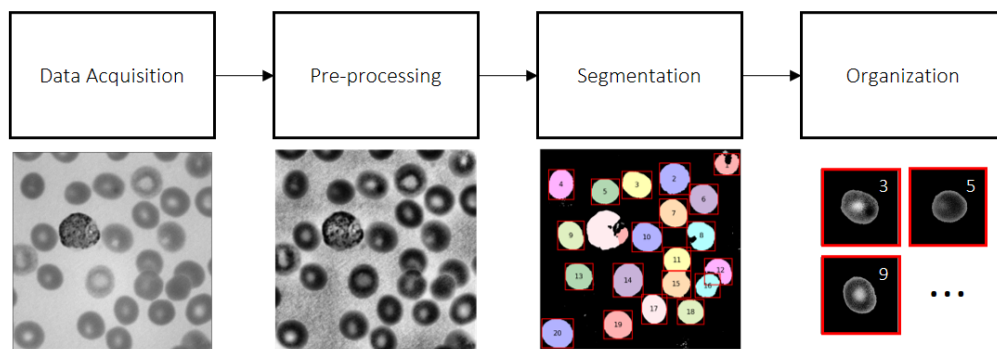


Figure 7.1: Processing steps in the development of data sets that are comprised of segmented red blood cells.

¹Here, I agree upon the obscurity of the statement. In general, researchers aim to use datasets that represent all features of certain objects or processes that are needed to wholly describe them such that the network has the *possibility* to see all that is potentially necessary. A network that should be able to discriminate apples from pears is ideally given data of different types of apples and pears, different stages (ripe, rotten), photographed from different angles, etc.

Data acquisition is described in more detail in Part I. Furthermore, other steps were briefly mentioned before in Ch. 2 and more in-depth discussions will not be done. The next subsections are about two data sets that are used to evaluate the proposed networks on (see Ch. 6 for the networks): the *Rajaraman* and *AiDx* data sets.

7.1.1 Rajaraman data set

An incredible open-source data set is that which is obtained by researchers at the Lister Hill National Center for Biomedical Communications (LHNCBC) which is part of the National Library of Medicine (NLM). In fact, they developed a mobile-integrated microscope device that runs on a standard Android smartphone. In the Chittagong Medical College Hospital in Bangladesh they undertook a massive task by capturing images from thin blood smears from 150 infected and 50 healthy patients which were subsequently examined and annotated by expert slide readers at the Mahidol-Oxford Tropical Medicine Research Unit in Bangkok, Thailand [3].

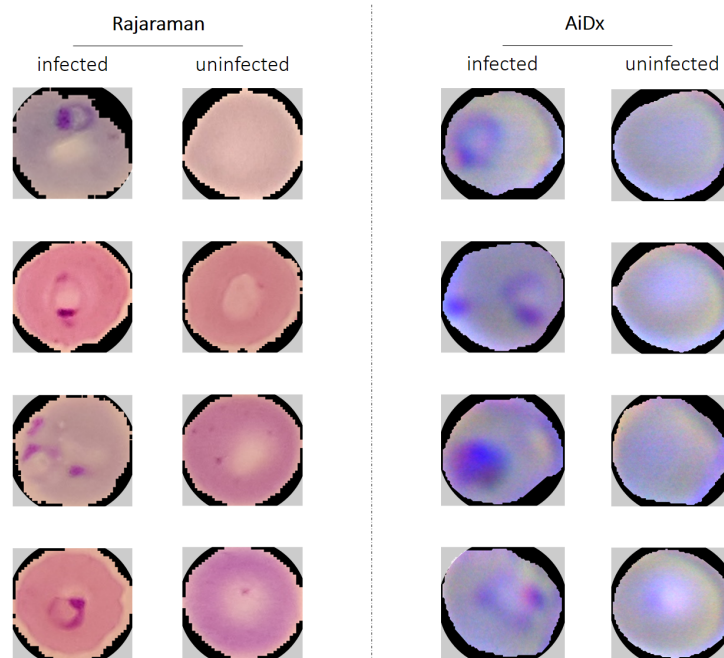


Figure 7.2: Several samples from the Rajaraman data set (left) and AiDx data set (right).

The thin blood smear data set contains an impressive total of 27,558 segmented and pre-processed cells. Approximately, there are as many infected as uninfected instances. The data set will be referred to as the *Rajaraman* data set: Sivaramakrishnan Rajaraman was one of the first researchers that evaluated this huge set on a multitude of the most famous deep learning algorithms that existed at the time, such as ResNet, VGG and AlexNet [84, 85]. Moreover, he stays predominant in the field of automated detection of malaria [65, 3]. Some samples of the data set are provided in Fig. 7.2. Although the exact equipment and configuration used by the researchers at the LHNCBC is not found, it looks as if high-end

microscopes were used in the setup providing homogeneous illumination while preserving high spatial resolution. Detection and segmentation of the cells were done by applying a level-set based algorithm [86].

7.1.2 AiDx data set

The AiDx data set was captured by the *AiDx Assist*, specifically the B.02 version, which is one of the products developed by *AiDx Medical* in Rijswijk, the Netherlands [4]. It is a low-cost, reliable and quick ML-assisted diagnostics microscope. It is field compatible and scalable to a variety of parasitic diseases [4]. The specifications are given in Tab. 7.1. Although the setup does not perfectly correct for chromatic (and spherical) aberrations, the relatively cheap equipped setup is designed smartly as can be seen in some segmented cells filtered from the blood films (see Fig. 7.2). Compared to high-end microscopes, the spatial resolution is reduced but seems sufficient for the detection of parasites. A U-net-based method is used to crop out the cells [87].

Table 7.1: Specifications of the AiDx Assist B.02.

Property	AiDx Assist B.02
Resolution (camera)	[4896 x 3680] or 18MP
Pixel size (camera) [μm]	5.5
FoV (sample) [μm]	[664 x 499]
Objective [mag.]	10(/20)
Source	LED: 400 - 650 nm
Immersion (type)	Oil-immersed sample-condenser
Costs [€]	1065
Remaining	weight: 3.2kg

7.2 Performance metrics

7.2.1 Accuracy, sensitivity and specificity

Performance can be measured in a multitude of ways and the measure’s relevance is application dependent. In deep learning, and especially classification in computer vision, the most abundantly used measure is *accuracy* - the amount of correct predictions. It is highly *global* as it quantifies the *total* amount of correct predictions as percentage of *all* tested samples, in which, by design, it discards information per individual class. In applications for which it is essential that some bound on performance *per class* must be guaranteed, solely accuracy as performance indication is not enough. In the diagnosis of malaria, *parasitemia* - the ratio of infected cells as of the total - in the range of 0.1 - 1%, which is 1 to 10 in thousand cells, tends to be the point at which proper treatment is already necessary [88]. As a consequence, if diagnosis is done at the lower limit of the regime, and say 10.000 cells are examined, it is crucial that the prediction of the amount of *infected* cells is near perfect. In this two-class world - infected-uninfected - two key performance measures are introduced: *sensitivity* and

specificity which were earlier mentioned in Ch. 2. Sensitivity is often referred to as *true-positive rate* (TPR), *recall*, *probability of detection*, or *power* and specificity as *true-negative rate* (TNR) or *selectivity*. In this report, naming stays consistent: only in the definition of the F1-score, recall will be used instead. This decomposition of accuracy is done often and serves as common measure in the development and testing of diagnostic tools. They are summarized in the *confusion matrix* in Fig. 7.3.

		Outcome of test		
		Predicted YES	Predicted NO	
Condition	Actual YES	True Positive tp	False Negative fn	Sensitivity "test-ability to correctly classify a positive individual" $\frac{tp}{tp + fn}$
	Actual NO	False Positive fp	True Negative tn	Specificity "test-ability to correctly classify a negative individual" $\frac{tn}{tn + fp}$
				Accuracy "global correctness" $\frac{tp + tn}{tp + fp + tn + fn}$

Figure 7.3: *Diagnosis* confusion matrix for the logical two-class system: positive-negative. P is the amount of positives and N the amount of negatives measured. The subscript indicates whether it is correctly (t) or falsely (f) predicted. In the yellow boxes, 3 distinct performance measures are denoted - (1) *sensitivity*, the ability of the test to correctly classify if an individual is positive, (2) *specificity*, a similar measure that does it for an individual that is negative, and (3) *accuracy*, the overall correctness.

7.2.2 The F-score

Another score that is often used in binary classification is the *F-score* - the harmonic mean between *precision*, the fraction of true positives to be correct, and *recall* (sensitivity). The balanced F-score, or *F1-score*, is computed as

$$F_1 = \frac{2}{r^{-1} + p^{-1}} = 2 \cdot \frac{pr}{p + r} = \frac{tp}{tp + \frac{1}{2}(tn + fn)}, \quad (7.1)$$

in which p and r are precision and recall. Its more generic form, which lets the designer choose upon the relative weight he or she lays on either precision or recall, expands as

$$F_\gamma = (1 + \gamma^2) \cdot \frac{pr}{\gamma^2 p + r} = \frac{(1 + \gamma^2) tp}{(1 + \gamma^2) tp + \gamma^2 tn + fp}. \quad (7.2)$$

Here, γ represents that weight - recall is considered γ times as important as precision. This trade-off between precision and recall is a relevant measure for low-detection-limited diagnostics, such as malaria, because it lays emphasis on the fact that class distribution can be highly imbalanced.

7.3 K-fold cross-validation

As must be clear by now, data sets are *never* abundant enough and are often more towards the other end of the spectrum - being too small and insufficient. Increasing the amount of training samples is paramount to better performance, but to be able to validate your model part of the data set is not to be used for training. *Cross-validation* is a method in which the validation set can be reduced in size without a desperate fall-back in uncertainty (Section 5.3.1 in [79]). It initializes by segmenting the data-set in k equally sized and randomly picked cross-sections, hence its reference *K-fold* cross-validation. The network is separately trained, using every subset for validation once and its remainder as training data. Although individually the results are more uncertain, its ensemble is clearly not. In fact, though depending on k , it actually increases the certainty of that what is measured and is less prone to large, unwanted, fluctuations during the training procedure. The method is illustrated in Fig. 7.4.

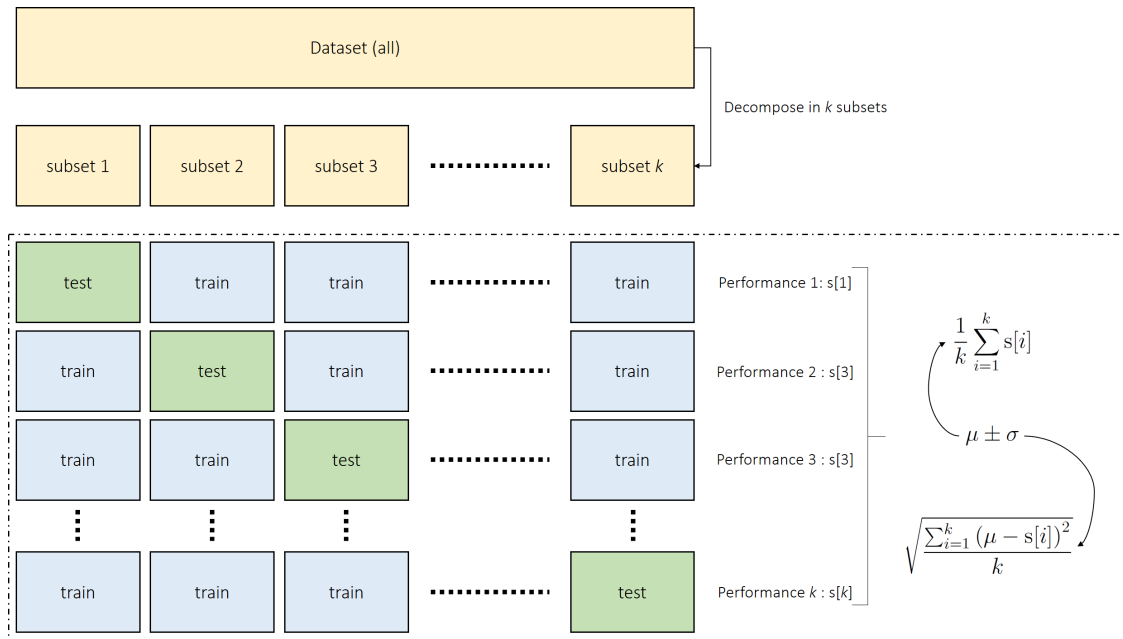


Figure 7.4: K-fold cross-validation.

7.4 Results

Networks will be tested that are specified in Tab. 6.4 and 6.1. Both data-sets have square-resized inputs, having dimension equal to 64 (in both width and height). The batch-size was chosen to be 64, which is common practice in deep learning. Note that changing the *size* of the input or batch does not influence the interior of the networks. On the other hand, they are, approximately quadratically in input dimension and linearly inverse in batch-size, correlated with both time of training and allocated memory needed. To increase diversity of the training set, small random transformations are applied to the training database: *data augmentation*. It is chosen to make minor alterations in the form of rotations (up to 30 degrees) and small grid allocations (shifts). Furthermore, random reflections are permitted as well (both vertically as well as horizontally). Finally all training data is normalized between equivalent bounds. As the training set is relatively large, online augmentation is preferred, without letting the set size explode. It is done for every mini-batch individually before being fed in to the network. The test set is only normalized. For clarity, an overview - hierarchy and naming - is given in Fig. 7.5.

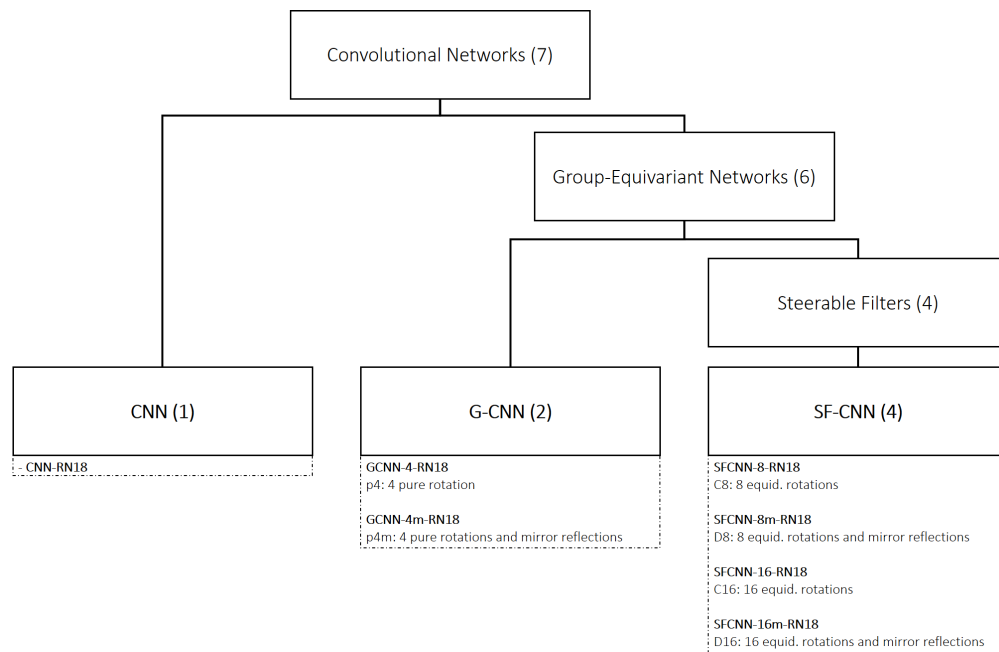


Figure 7.5: Hierarchy, naming and relevant specifications of the proposed convolutional networks. The common name, as earlier specified, of the architecture used is called *ResNet-18* (comprised of 18 operational stages), which is abbreviated as RN18. In between the type and used architecture, the amount of equidistant rotations is mentioned. Furthermore, an extra "m" reflects the equivariance under mirror reflections.

7.4.1 Rajaraman data set

Concerning the samples in the Rajaraman data set, the error rates (inverse of accuracy) for both training and test, are plotted in Fig. 7.7 during training for all 7 configurations at every

epoch, including their best scores. The initial learning rate was 0.05 and discretely decreased 10-fold every 50 epochs. From epoch 150 onward, the learning rate remained constant, being 0.00005, until its ending at epoch 300. SGD, with momentum equal to 0.9, and Cross-Entropy Loss were used as optimizer and cost functions respectively. Global max-pooling over the complete discrete set of transformation and global anti-aliased average-pooling along all features were done at the entrance of the decoder to make networks wholly *invariant* for all inherent transformations.

Table 7.2: Results on Rajarman data set for 5 performance metrics. Written in bold denote the highest scores - red for first, blue for second and black for third place.

	Accuracy	Sensitivity	Specificity	Precision	F1-score
CNN-RN18	95.62 ± 0.09	95.72 ± 0.11	95.54 ± 0.13	95.55 ± 0.11	95.64 ± 0.11
GCNN-4-RN18	96.87 ± 0.10	97.00 ± 0.10	96.45 ± 0.11	96.68 ± 0.09	96.28 ± 0.10
GCNN-4m-RN18	97.60 ± 0.11	98.21 ± 0.07	97.55 ± 0.08	97.42 ± 0.08	97.82 ± 0.10
SFCNN-8-RN18	96.97 ± 0.10	97.05 ± 0.10	96.84 ± 0.09	96.96 ± 0.06	97.01 ± 0.08
SFCNN-8m-RN18	97.10 ± 0.10	97.52 ± 0.07	97.02 ± 0.10	97.20 ± 0.09	97.36 ± 0.08
SFCNN-16-RN18	97.12 ± 0.07	97.56 ± 0.09	96.57 ± 0.10	97.05 ± 0.11	97.31 ± 0.10
SFCNN-16m-RN18	97.63 ± 0.08	98.34 ± 0.08	96.99 ± 0.09	97.45 ± 0.10	97.90 ± 0.10

Tab. 7.2 provides all outcomes for the (other) performance metrics. All measures were cross-validatedly computed for 5 equally-sized cross-sections of the data set - indicated by mean and standard deviation, $\mu \pm \sigma$. The (best) scores are given in bold. In addition, Tab. 7.3 shows how the scores are relatively improved compared to the CNN baseline². It becomes apparent that all networks perform significantly better among all metrics. More specifically, all networks improve on any metric by, relatively measured, at least 15%, having its average approximately between 30 and 40%. In general, the trend between metrics shows that sensitivity is highest, the F1-score is second highest, accuracy is somewhere in the middle, precision is slightly below average and specificity is, in most cases, lowest.

Table 7.3: Improvement indicator for group-equivariant operations. The baseline is the conventional CNN-RN18. The relative percentage increments between the network's and the baseline's error rates are given.

	Accuracy	Sensitivity	Specificity	Precision	F1-score
GCNN-4-RN18	+28.5%	+29.9%	+20.4%	+25.4%	+14.7%
GCNN-4m-RN18	+45.2%	+58.2%	+45.1%	+42.0%	+50.0%
SFCNN-8-RN18	+30.8%	+31.1%	+29.1%	+31.7%	+39.4%
SFCNN-8m-RN18	+33.8%	+42.0%	+33.2%	+37.1%	+38.3%
SFCNN-16-RN18	+34.2%	+43.0%	+23.1%	+33.0%	+38.3%
SFCNN-16m-RN18	+45.9%	+61.2%	+35.5%	+42.7%	+51.8%

The learning curves seem healthy, without a disastrous display of huge spikes. Furthermore, they show that for both G-CNN and SF-CNN, the best scores are commonly achieved before it reaches epoch 100. Especially in the case of SF-CNN, the scores tend to slightly worsen

²More precisely, it gives the percentage of how the error rates have improved: $100 \cdot \left(1 - \frac{\text{error}_{net}}{\text{error}_{base}}\right)$




Figure 7.6: Performance for increasing heuristic. Recall that $|p4m| = |c8| = 8$ and $|D8| = |C16| = 16$.

after the first update of the learning rate (epoch 50). The gaps between the training and test curves become obvious and grow larger from the primal change in learning rate onward and seem to grow in the heuristic of the group, thus the number of added transformations. For both G-CNN and SF-CNN, increasing the number of rotations, and more significantly by adding reflections, results in better performance. Although G-CNNs have the tendency to perform, so to say, more efficient per added transformation (see Fig. 7.6), SF-CNNs gets the better of it if enough weights are shared in its peripheries: the best network (except for specificity which is third best) is the SFCNN-16m-RN18, being 97.63% (+45.6% compared to baseline) accurate, having an F1-score of 97.90% (+51.8%), and achieves a sensitivity of 98.34% (+61%).




Figure 7.7: Learning curves (accuracy) for all different operations. Both training (dotted) and test (line) curves are shown. The best-scoring networks are highlighted in red (SF-CNN-16m-RN18, best) and blue (G-CNN-4m-RN18, 2nd best).

7.4.2 AiDx data set

The *pre-trained models* that were evaluated on the Rajaraman data set are used as potential classifiers for the AiDx data set. The results on all performance metrics are epitomized in Tab. 7.4. It should be mentioned that the whole data set is used for validation simultaneously because the size is small: no cross-validation is used. In line with the results on the Rajaraman data set, specificity is relatively low compared to sensitivity. Therefore, in combination with the unbalancedness of the AiDx data set - approximately 1/10th of the segmented cells are infected - low precisions and F1-scores are found. Again, the SFCNN-19m-RN18s outperforms all other models on all metrics.

Table 7.4: Results on AiDx data set for 5 performance metrics for the pre-trained models on the Rajaraman data set. Written in bold denote the highest scores - red for first, blue for second and black for third place.

	Accuracy	Sensitivity	Specificity	Precision	F1-score
CNN-RN18	61.1	84.7	58.5	18.3	30.1
GCNN-4-RN18	62.1	87.2	59.4	18.8	30.9
GCNN-4m-RN18	62.7	87.7	60.0	19.2	31.5
SFCNN-8-RN18	64.9	92.0	62.0	20.2	33.2
SFCNN-8m-RN18	75.0	95.0	72.9	26.8	41.8
SFCNN-16-RN18	70.4	92.9	68.1	22.9	36.7
SFCNN-16m-RN18	91.0	95.5	90.5	51.5	66.9

8 | Discussion and conclusions

Group-equivariant neural networks, comprised of convolutional stages that additively equivary under a discrete group of transformations (rotations and mirror reflections), are evaluated on the Rajaraman data set (Sec. 7.1.1) and subsequently used as classifiers for our own AiDx data set (Sec. 7.1.2). All are implemented in a residual network architecture (ResNet-18) and compared with the conventional convolutional stage. The sizes of the networks remained consistent among all operations.

In terms of overall performance on the Rajaraman data set, it is shown that all non-conventional networks significantly increase on all performance metrics - accuracy, sensitivity, specificity, precision and the F1-score (see Tab. 7.2 and 7.3). Although the sizes of the networks are relatively small, that is about 1.90 million parameters, best scores are all marginally above 97% and for the baseline just above 95.5%. Smaller-sized networks that have re-interpreted the way feature maps are connected can therefore be adapted in the process while maintaining a profound margin on top of the limit the WHO prescribes, which is 95% on both sensitivity and specificity (thus accuracy as well). Interesting research would be about the implementation of group-equivariant convolutional stages within much larger networks as well, such as the VGG-19, SqueezeNet or InceptionResNet-V2 which, for conventional convolutional stages, were evaluated by Rajaraman, Jaeger and Antani [65].

All training curves look healthy and converge mostly within 100 epoch. In most cases it saturates at a slightly higher platform in which overfitting starts to occur between epoch 50 and 100 (see Fig. 7.7). To boost performance without being prone to end up in relatively bad local minima, the training procedure could potentially be enhanced by proper choice of the learning rate scheduler. By many claimed as the most important hyperparameter to tune, it can significantly increase the time of convergence, the way it converges and the point it converges to. Grid-search methods to find an optimal starting point and more adaptive variants instead of predefined curvatures that alter the state of the learning rate throughout training could be interesting as next step.

In general, increasing the amount of rotations the network's layers equivary under improves performance. More significantly though is the addition of mirror reflections. It is for example observed that the GCNN-4m-RN18 ($|G|= 8$) performs much better than SFCNN-8-RN18 ($|G|= 8$) and SFCNN-16-RN18 ($|G|= 16$). Strangely, it outperforms the SFCNN-8m-RN18 (equivariant under mirror reflections) as well. It could be that GCNNs lays its connections more profoundly as its filters have lesser constraints while SFCNNs have the tendency that

their fundamental filters show aliasing effects. A more empirical explanation could be that SFCNNs, according to its developers, work best if large kernels are used. Instead of using filters of size 5, larger ones could be tested (e.g. 7 or 9).

Evaluating the pre-trained models on the AiDx data set is encouraging the use of more generalizing operations as well. Although all networks have increased performance on all metrics if compared to the baseline, inducing the convolutional stages to equivary under at least 16 transformations overly outperforms it (see Tab. 7.4). This is an indication that these operations are more robust toward varying states the blood films can be in or the way they are being acquisitioned and/or (pre-)processed. Even so, in this particular case, specificity is just above 90%, making the pre-trained models prone for having a large amount of false positives which, in the context of a low-needed limit of detection, could result in a large amount of misdiagnoses (many non-infected patients will be positively diagnosed). In contrary to what was found for the Rajarman data set, the effect of the way the stages have built-in equivariance (GCNN vs. SFCNN) seems negligible as performance always increases with the heuristic of the groups. This may imply that, although the method to induce equivariance differs, the amount of realistic transformations matters most in determining the robustness of the ResNet-18 classifier - the more transformations the more robust.

An interesting addition, if *too* specific data is to be predicted, is to (shortly) train the pre-trained models on that specific database. In doing so, one could either train on all parameters, or keep some parts fixed. For example, the encoder part, which is best seen as feature extractor, could be held fixed while only learning how the features are to be combined for this specific set, thus solely making the decoder part changeable. This method, also seen as a way to exploit *transfer learning*, could reduce on training time compared to the all-training variant while still improving on performance. As to implement, it is a prerequisite that training can be done at the location: devices being able to train at and personnel that is able to properly label.

A main disadvantage of using generalizing operations is their increased computational complexity. Especially if the models are to be implemented on micro-computers (such as a mobile-phone or Raspberry-Pie), complexity of its inherent operations is ideally held to a minimum. Finding these limits, or pushing the limits by efficient and redefined implementation of the proposed networks, would be interesting to delve into.

Introducing a-priori knowledge into neural networks through its manifestation within the networks operations (stages), has shown to improve on performance for many databases. In this report it is shown to be promising in the context of discriminating malaria-parasitized from uninfected red blood cells as well. For further exploration, the domain of implementation especially needs attention.

Bibliography

- [1] Delft Center for Systems and Control. URL <https://www.tudelft.nl/en/3me/about/departments/delft-center-for-systems-and-control>
- [2] Delft University of Technology. URL <https://www.tudelft.nl/en/>.
- [3] Rajaraman, S. *et al.* Pre-trained convolutional neural networks as feature extractors toward improved malaria parasite detection in thin blood smear images. *PeerJ* **6**, e4568 (2018). URL <https://peerj.com/articles/4568>. Publisher: PeerJ Inc.
- [4] BV, A. M. AiDx Medical. URL <https://www.aidx-medical.com/>.
- [5] Beck, T. Figure 1, [The malaria parasite life cycle...]. (2006). URL https://www.ncbi.nlm.nih.gov/sites/books/NBK5951/figure/malaria_LifeCycle/. Library Catalog: www.ncbi.nlm.nih.gov Publisher: National Center for Biotechnology Information (US).
- [6] Bosschaart, N., Edelman, G. J., Aalders, M. C. G., van Leeuwen, T. G. & Faber, D. J. A literature review and novel theoretical approach on the optical properties of whole blood. *Lasers in Medical Science* **29**, 453–479 (2014). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3953607/>.
- [7] Wang, Y. W., Reder, N. P., Kang, S., Glaser, A. K. & Liu, J. T. Multiplexed Optical Imaging of Tumor-Directed Nanoparticles: A Review of Imaging Systems and Approaches. *Nanotheranostics* **1**, 369–388 (2017). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5647764/>.
- [8] Marquardt, D. W. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *Journal of the Society for Industrial and Applied Mathematics* **11**, 431–441 (1963). URL <https://epubs.siam.org/doi/10.1137/0111030>. Publisher: Society for Industrial and Applied Mathematics.
- [9] WHO. World malaria report 2019. URL <https://www.who.int/publications-detail/world-malaria-report-2019>. Library Catalog: www.who.int.
- [10] The Malaria Atlas Project. URL <https://malariaatlas.org/>. Library Catalog: malariaatlas.org.

- [11] Ferrão, J. L., Mendes, J. M., Painho, M. & Zacarias, S. Malaria mortality characterization and the relationship between malaria mortality and climate in Chimoio, Mozambique. *Malaria Journal* **16**, 212 (2017). URL <https://doi.org/10.1186/s12936-017-1866-0>.
- [12] Steketee, R. W. *et al.* The effect of malaria and malaria prevention in pregnancy on offspring birthweight, prematurity, and intrauterine growth retardation in rural Malawi. *The American Journal of Tropical Medicine and Hygiene* **55**, 33–41 (1996). URL <https://www.ajtmh.org/content/journals/10.4269/ajtmh.1996.55.33>. Publisher: The American Society of Tropical Medicine and Hygiene.
- [13] Iqbal, S. A. *et al.* Hematological Differences among Malaria Patients in Rural and Urban Ghana. *Journal of Tropical Pediatrics* **62**, 477–486 (2016). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5141942/>.
- [14] Greenwood, B. M. *et al.* Mortality and morbidity from malaria among children in a rural area of The Gambia, West Africa. *Transactions of the Royal Society of Tropical Medicine and Hygiene* **81**, 478–486 (1987). URL <http://www.sciencedirect.com/science/article/pii/0035920387901702>.
- [15] Deressa, W., Ali, A. & Enqueslassie, F. Self-treatment of malaria in rural communities, Butajira, southern Ethiopia. *Bulletin of the World Health Organization* **81**, 261–268 (2003). URL <https://www.scielosp.org/article/bwho/2003.v81n4/261-268/en/>. Publisher: World Health Organization.
- [16] CDC - Parasites - Malaria (2020). URL <https://www.cdc.gov/parasites/malaria/index.html>. Library Catalog: www.cdc.gov.
- [17] LUMC. Malaria | LUMC. URL <https://www.lumc.nl/org/parasitologie/patient-en-zorg/Meer>. Library Catalog: www.lumc.nl.
- [18] Prevention, C.-C. f. D. C. a. CDC - Malaria - About Malaria - Biology (2019). URL <https://www.cdc.gov/malaria/about/biology/index.html>. Library Catalog: www.cdc.gov.
- [19] Miller, L. H., Good, M. F. & Milon, G. Malaria pathogenesis. *Science* **264**, 1878–1883 (1994). URL <https://science.sciencemag.org/content/264/5167/1878>. Publisher: American Association for the Advancement of Science Section: Articles.
- [20] Poostchi, M., Silamut, K., Maude, R. J., Jaeger, S. & Thoma, G. Image analysis and machine learning for detecting malaria. *Translational Research* **194**, 36–55 (2018). URL <http://www.sciencedirect.com/science/article/pii/S193152441730333X>.
- [21] Tangpukdee, N., Duangdee, C., Wilairatana, P. & Krudsood, S. Malaria Diagnosis: A Brief Review. *The Korean Journal of Parasitology* **47**, 93–102 (2009). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2688806/>.

- [22] Mathison, B. A. & Pritt, B. S. Update on Malaria Diagnostics and Test Utilization. *Journal of Clinical Microbiology* **55**, 2009–2017 (2017). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5483902/>.
- [23] Endeshaw, T. *et al.* Evaluation of light microscopy and rapid diagnostic test for the detection of malaria under operational field conditions: a household survey in Ethiopia. *Malaria Journal* **7**, 118 (2008). URL <https://doi.org/10.1186/1475-2875-7-118>.
- [24] Payne, D. Use and limitations of light microscopy for diagnosing malaria at the primary health care level. *Bulletin of the World Health Organization* **66**, 621–626 (1988).
- [25] Hassan, S. E.-D. H., Okoued, S. I., Mudathir, M. A. & Malik, E. M. Testing the sensitivity and specificity of the fluorescence microscope (Cyscopeó) for malaria diagnosis. *Malaria Journal* **9**, 88 (2010). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2860357/>.
- [26] Hassan, S. E.-D. H., Haggaz, A. E. D., Mohammed-Elhassan, E. B., Malik, E. M. & Adam, I. Fluorescence microscope (Cyscopeó) for malaria diagnosis in pregnant women in Medani Hospital, Sudan. *Diagnostic Pathology* **6**, 88 (2011). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3189106/>.
- [27] Hathiwala, R., Mehta, P. R., Nataraj, G. & Hathiwala, S. LED fluorescence microscopy: Novel method for malaria diagnosis compared with routine methods. *Journal of Infection and Public Health* **10**, 824–828 (2017). URL <http://www.sciencedirect.com/science/article/pii/S1876034117300011>.
- [28] Sousa-Figueiredo, J. C. *et al.* Investigating portable fluorescent microscopy (CyScopeó) as an alternative rapid diagnostic test for malaria in children and women of child-bearing age. *Malaria Journal* **9**, 245 (2010). URL <https://doi.org/10.1186/1475-2875-9-245>.
- [29] Lenz, D. *et al.* Assessment of LED fluorescence microscopy for the diagnosis of Plasmodium falciparum infections in Gabon. *Malaria Journal* **10**, 194 (2011). URL <https://doi.org/10.1186/1475-2875-10-194>.
- [30] Kawamoto, F. Rapid diagnosis of malaria by fluorescence microscopy with light microscope and interference filter. *The Lancet* **337**, 200–202 (1991). URL <http://www.sciencedirect.com/science/article/pii/014067369192159Y>.
- [31] Hänscheid, T. & Grobusch, M. P. How useful is PCR in the diagnosis of malaria? *Trends in Parasitology* **18**, 395–398 (2002). URL <http://www.sciencedirect.com/science/article/pii/S1471492202023486>.
- [32] Snounou, G. *et al.* High sensitivity of detection of human malaria parasites by the use of nested polymerase chain reaction. *Molecular and Biochemical Parasitology* **61**, 315–320 (1993).

- [33] Johnston, S. P. *et al.* PCR as a Confirmatory Technique for Laboratory Diagnosis of Malaria. *Journal of Clinical Microbiology* **44**, 1087–1089 (2006). URL <https://jcm.asm.org/content/44/3/1087>. Publisher: American Society for Microbiology Journals Section: PARASITOLOGY.
- [34] Kain, K. C. *et al.* Qualitative and semiquantitative polymerase chain reaction to predict Plasmodium falciparum treatment failure. *The Journal of Infectious Diseases* **170**, 1626–1630 (1994).
- [35] Lucchi, N. W. *et al.* Evaluation of the Illumigene Malaria LAMP: A Robust Molecular Diagnostic Tool for Malaria Parasites. *Scientific Reports* **6**, 1–8 (2016). URL <https://www.nature.com/articles/srep36808>. Number: 1 Publisher: Nature Publishing Group.
- [36] Paris, D. H. *et al.* Loop-Mediated Isothermal PCR (LAMP) for the Diagnosis of Falciparum Malaria. *The American Journal of Tropical Medicine and Hygiene* **77**, 972–976 (2007). URL <https://www.ajtmh.org/content/journals/10.4269/ajtmh.2007.77.972>. Publisher: The American Society of Tropical Medicine and Hygiene.
- [37] Notomi, T. *et al.* Loop-mediated isothermal amplification of DNA. *Nucleic Acids Research* **28**, E63 (2000).
- [38] Poon, L. L. *et al.* Sensitive and Inexpensive Molecular Test for Falciparum Malaria: Detecting Plasmodium falciparum DNA Directly from Heat-Treated Blood by Loop-Mediated Isothermal Amplification,. *Clinical Chemistry* **52**, 303–306 (2006). URL <https://academic.oup.com/clinchem/article/52/2/303/5626787>. Publisher: Oxford Academic.
- [39] Schoone, G. J., Oskam, L., Kroon, N. C. M., Schallig, H. D. F. H. & Omar, S. A. Detection and Quantification of Plasmodium falciparum in Blood Samples Using Quantitative Nucleic Acid Sequence-Based Amplification. *Journal of Clinical Microbiology* **38**, 4072–4075 (2000). URL <https://jcm.asm.org/content/38/11/4072>. Publisher: American Society for Microbiology Journals Section: PARASITOLOGY.
- [40] Smits, H. L. *et al.* Detection, identification and semi-quantification of malaria parasites by NASBA amplification of small subunit ribosomal RNA sequences. *Journal of Microbiological Methods* **28**, 65–75 (1997). URL <http://www.sciencedirect.com/science/article/pii/S0167701296009669>.
- [41] Wilson, M. L. Malaria Rapid Diagnostic Tests. *Clinical Infectious Diseases* **54**, 1637–1641 (2012). URL <https://academic.oup.com/cid/article/54/11/1637/321357>. Publisher: Oxford Academic.
- [42] Moody, A. Rapid Diagnostic Tests for Malaria Parasites. *Clinical Microbiology Reviews* **15**, 66–78 (2002). URL <https://cmr.asm.org/content/15/1/66>. Publisher: American Society for Microbiology Journals Section: REVIEWS.

- [43] Wongsrichanalai, C., Barcus, M. J., Muth, S., Sutamihardja, A. & Wernsdorfer, W. H. A Review of Malaria Diagnostic Tools: Microscopy and Rapid Diagnostic Test (RDT). *The American Journal of Tropical Medicine and Hygiene* **77**, 119–127 (2007). URL <https://www.ajtmh.org/content/journals/10.4269/ajtmh.2007.77.119>. Publisher: The American Society of Tropical Medicine and Hygiene.
- [44] Tek, F. B., Dempster, A. G. & Kale, I. Computer vision for microscopy diagnosis of malaria. *Malaria Journal* **8**, 153 (2009). URL <https://doi.org/10.1186/1475-2875-8-153>.
- [45] Mendelow, B. V. *et al.* Automated malaria detection by depolarization of laser light. *British Journal of Haematology* **104**, 499–503 (1999). URL <https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2141.1999.01199.x>. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1046/j.1365-2141.1999.01199.x>.
- [46] Pirstill, C. W. & Coté, G. L. Malaria Diagnosis Using a Mobile Phone Polarized Microscope. *Scientific Reports* **5**, 1–13 (2015). URL <https://www.nature.com/articles/srep13368>. Number: 1 Publisher: Nature Publishing Group.
- [47] Breslauer, D. N., Maamari, R. N., Switz, N. A., Lam, W. A. & Fletcher, D. A. Mobile phone based clinical microscopy for global health applications. *PloS One* **4**, e6320 (2009).
- [48] Skandarajah, A., Reber, C. D., Switz, N. A. & Fletcher, D. A. Quantitative Imaging with a Mobile Phone Microscope. *PLoS ONE* **9** (2014). URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4019540/>.
- [49] Breslauer, D. N., Maamari, R. N., Switz, N. A., Lam, W. A. & Fletcher, D. A. Mobile Phone Based Clinical Microscopy for Global Health Applications. *PLoS ONE* **4**, e6320 (2009). URL <https://dx.plos.org/10.1371/journal.pone.0006320>.
- [50] Scholl, P. F. *et al.* RAPID DETECTION OF MALARIA INFECTION IN VIVO BY LASER DESORPTION MASS SPECTROMETRY. *The American Journal of Tropical Medicine and Hygiene* **71**, 546–551 (2004). URL <https://www.ajtmh.org/content/journals/10.4269/ajtmh.2004.71.546>. Publisher: The American Society of Tropical Medicine and Hygiene.
- [51] Park, B. & Lu, R. (eds.) *Hyperspectral Imaging Technology in Food and Agriculture*. Food Engineering Series (Springer-Verlag, New York, 2015). URL <https://www.springer.com/gp/book/9781493928354>.
- [52] Chang, C.-I. *Hyperspectral Imaging: Techniques for Spectral Detection and Classification* (Springer US, 2003). URL <https://www.springer.com/gp/book/9780306474835>.
- [53] Lu, G. & Fei, B. Medical hyperspectral imaging: a review. *Journal of Biomedical Optics* **19**, 010901 (2014). URL <https://www.spiedigitallibrary.org/journals/journal-of-biomedical-optics/volume-19/i> Publisher: International Society for Optics and Photonics.

- [54] Hyperspectral Imaging System. URL <https://www.thorlabs.com/newgrouppage9.cfm?objectgroup>
- [55] The Feynman Lectures on Physics. URL <https://www.feynmanlectures.caltech.edu/>.
- [56] Paraxial Rays. URL <http://hyperphysics.phy-astr.gsu.edu/hbase/geoopt/paraxial.html>.
- [57] Cambridge Philosophical Society. *Transactions of the Cambridge Philosophical Society* (1833). URL <http://archive.org/details/transactionsofca05camb>.
- [58] Greivenkamp, J. E. *Field Guide to Geometrical Optics* (Society of Photo Optical, 2004). Google-Books-ID: 1YfZNWZAwCAC.
- [59] Light and Video Microscopy - 3rd Edition. URL <https://www.elsevier.com/books/light-and-video-microscopy/wayne/978-0-12-816501-0>.
- [60] Zeitschrift für wissenschaftliche Mikroskopie und mikroskopische Technik. URL <https://archive.org/stream/zeitschriftfrw10stut>.
- [61] Resolution. URL <https://www.microscopyu.com/microscopy-basics/resolution>.
- [62] aodunsi. Color Blindness. URL <http://chicago.medicine.uic.edu/departments/academic-departments>
- [63] Neitz, J., Carroll, J. & Neitz, M. Almost Reason Enough for Having Eye s. *COLOR VISION* 8.
- [64] *QED* (2014). URL <https://press.princeton.edu/books/paperback/9780691164090/qed>.
- [65] Rajaraman, S., Jaeger, S. & Antani, S. K. Performance evaluation of deep neural ensembles toward malaria parasite detection in thin-blood smear images. *PeerJ* **7**, e6977 (2019).
- [66] Hornik, K., Stinchcombe, M. & White, H. Multilayer feedforward networks are universal approximators. *Neural Networks* **2**, 359–366 (1989). URL <http://www.sciencedirect.com/science/article/pii/0893608089900208>.
- [67] Dive into Deep Learning Dive into Deep Learning 0.7 documentation. URL <https://classic.d2l.ai/index.html>.
- [68] Kondor, R. & Trivedi, S. On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups. *arXiv:1802.03690 [cs, stat]* (2018). URL <http://arxiv.org/abs/1802.03690>. ArXiv: 1802.03690.
- [69] Cohen, T. S. & Welling, M. Group Equivariant Convolutional Networks. *arXiv:1602.07576 [cs, stat]* (2016). URL <http://arxiv.org/abs/1602.07576>. ArXiv: 1602.07576.
- [70] Engelenburg, C. v. Geometric Deep Learning: Group Equivariant Convolutional Networks (2020). URL <https://medium.com/swlh/geometric-deep-learning-group-equivariant-convolutional-netw>

- [71] Smarandache, F. W. B. Vasantha Kandasamy e-mail: vasanthakandasamy@gmail.com web: <http://mat.iitm.ac.in/~wbv> www.vasantha.in 168.
- [72] Cohen, T. S. & Welling, M. Steerable CNNs. *arXiv:1612.08498 [cs, stat]* (2016). URL <http://arxiv.org/abs/1612.08498>. ArXiv: 1612.08498.
- [73] Weiler, M., Hamprecht, F. A. & Storath, M. Learning Steerable Filters for Rotation Equivariant CNNs. *arXiv:1711.07289 [cs]* (2018). URL <http://arxiv.org/abs/1711.07289>. ArXiv: 1711.07289.
- [74] He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. *arXiv:1512.03385 [cs]* (2015). URL <http://arxiv.org/abs/1512.03385>. ArXiv: 1512.03385.
- [75] Ioffe, S. & Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv:1502.03167 [cs]* (2015). URL <http://arxiv.org/abs/1502.03167>. ArXiv: 1502.03167.
- [76] Janocha, K. & Czarnecki, W. M. On Loss Functions for Deep Neural Networks in Classification. *arXiv:1702.05659 [cs]* (2017). URL <http://arxiv.org/abs/1702.05659>. ArXiv: 1702.05659.
- [77] Brownlee, J. A Gentle Introduction to the Rectified Linear Unit (ReLU) (2019). URL <https://machinelearningmastery.com/rectified-linear-activation-function-for-deep-learn>
- [78] Masters, D. & Luschi, C. Revisiting Small Batch Training for Deep Neural Networks. *arXiv:1804.07612 [cs, stat]* (2018). URL <http://arxiv.org/abs/1804.07612>. ArXiv: 1804.07612.
- [79] Deep Learning. URL <https://www.deeplearningbook.org/>.
- [80] Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]* (2017). URL <http://arxiv.org/abs/1412.6980>. ArXiv: 1412.6980.
- [81] Ruder, S. An overview of gradient descent optimization algorithms. *arXiv:1609.04747 [cs]* (2017). URL <http://arxiv.org/abs/1609.04747>. ArXiv: 1609.04747.
- [82] Ge, R., Kakade, S. M., Kidambi, R. & Netrapalli, P. The Step Decay Schedule: A Near Optimal, Geometrically Decaying Learning Rate Procedure For Least Squares. *arXiv:1904.12838 [cs, math, stat]* (2019). URL <http://arxiv.org/abs/1904.12838>. ArXiv: 1904.12838.
- [83] Zen and the Art of Motorcycle Maintenance (Phaedrus, #1). URL https://www.goodreads.com/work/best_book/175720-zen-and-the-art-of-motorcycle-maintenance
- [84] Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks 9.

- [85] Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv:1409.1556 [cs]* (2015). URL <http://arxiv.org/abs/1409.1556>. ArXiv: 1409.1556.
- [86] Ersoy, I., Bunyak, F., Higgins, J. M. & Palaniappan, K. Coupled edge profile active contours for red blood cell flow analysis. In *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*, 748–751 (2012). ISSN: 1945-8452.
- [87] Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv:1505.04597 [cs]* (2015). URL <http://arxiv.org/abs/1505.04597>. ArXiv: 1505.04597.
- [88] Murray, S. Determination of Parasitemia Protocol 4.