Deep Learning Estimation of Absorbed Dose for Nuclear Medicine Diagnostics

Luciano Melodia[0000−0002−7584−7287]

Chair of Information Science
University of Regensburg
93051 Regensburg, Germany
{luciano.melodia}@stud.uni-regensburg.de

Abstract. The distribution of absorbed dose from radiotherapy with Lu^{177} can be estimated by convolving an image of a time-integrated activity distribution with a dose-voxel kernel consisting of different tissue types. This fast and inaccurate approximation is unsuitable for personalised dosimetry, as it neglects tissue heterogeneity. The latter can be calculated using various imaging techniques such as computed tomography and single photon emission computed tomography in combination with a time-consuming Monte Carlo simulation.

The aim of this study is to estimate for the first time the dose of voxel kernels from density kernels derived from computed tomography by deep learning using neural convolutional networks. The proposed architecture achieved on a test data set of real patient data a ratio of the intersection over the union of $= 0.86$ after 308 epochs and a corresponding mean squared error $= 1.24 \cdot 10^{-4}$. This ability to generalise shows that the trained convolution network can actually learn the function from the density kernel to the dose-voxel kernel.

Future work will evaluate dose-voxel kernels estimated by neural networks using Monte Carlo simulations of whole-body computer tomography to predict patient-specific voxel dose maps.

Keywords: Deep learning · Dosimetry · Artificial intelligence · Nuclear medicine · Cancer therapy.

1 Motivation

Nuclear methods are becoming increasingly important in medicine, both in diagnostics and in cancer therapy. Cancer is a disease whose incidence in the population increases with age and which requires the most meaningful diagnosis and effective treatment possible.

The approach of hybrid imaging techniques such as SPECT/CT and PET/CT can provide information on tissue density and, at the same time, on the distribution of radioactive isotopes injected for predictive diagnosis in the body. The spatial resolution of such procedures is in the range of mm to

cm. Metabolically active tissues such as tumours can be easily recognised in the three-dimensional distributions.

Nuclear medicine is not only useful for detecting tumours, but can also be used for their treatment. For this purpose, the patient is injected with a radioactive compound, which then attaches itself to the tumour and irradiates it accordingly. During therapy with a lutetium-labelled radiopharmaceutical, β^- rays are emitted into the target regions. The high-energy particles with very short range must be precisely dosed. Patient specific dosimetry is required [5]. To estimate the patient-specific dose distribution, SPECT images are taken at different times (4, 24, 48 and 72 hours after injection).

To obtain a distribution of the number of decays per voxel from the activity distribution thus obtained, the images must be integrated using the time-activity curve. There are different approaches to integrate the statistically uncertain images, but will not be discussed in this paper.

In the following we assume that the decay distribution is known. First the established principles of dose estimation are explained and then a new approach is presented.

2   Principles of Dose Voxel Calculation

Based on the distribution of the decays, the dose distribution is estimated. A widely used method is the convolution with a dose voxel kernel. This creates a cube with - in this case - 9^3 voxels. A common method to calculate dose distributions is the use of s-values [49]. The s-values are the fraction of the average kinetic energy for each radioactive decay absorbed by a unit of volume. The energy that can be measured in a given volume unit is the linear superposition of each contributing voxel within the spatial distribution around a radiation source. The absorbed dose that can be measured as a point nearby a radiating isotropic source in a homogeneous medium is called the dose point kernel [57,3].

Monte Carlo simulations are used to calculate high-resolution voxel s-values. The source is defined as a voxel and the angle and location of the emitted radiation is determined randomly. The particles are calculated back to their emission origin (material with $p = 1.04\text{g/cm}^3$) and the absorbed energy is stored in each voxel, usually with a resolution of 5mm. The average absorbed dose for each simulated particle is stored in $\text{MeV} \cdot \text{cm}^{-3}$ and then in $\text{mGy} \cdot \text{MBq}^{-1} \cdot \text{s}^{-1}$ for use as a voxel s value.

The generated s-values can be used in a suitable size to calculate the absorbed radiation dose [4].

The dose for a given tissue sample $\mathbf{r}_T$ is the sum of the intrinsic activity and of surrounding tissue $\mathbf{r}_U$ according to

$$\langle D(\mathbf{r}_T)\rangle = \sum_{\mathbf{r}_T} \tilde{A}_{\mathbf{r}_U} \cdot S(\mathbf{r}_T \leftarrow \mathbf{r}_U). \tag{1}$$

The source and calculated target dose are voxels from the patient's tissue, according to the cumulative activity map from the SPECT detector field.

Accordingly, for a given voxel $\mathbf{v}_T$ the average absorbed dose is the sum of the contributions of all source voxels $\mathbf{v}_U^{(n)}$. The activity of $\mathbf{v}_U^{(n)}$ is calculated from the product of the temporally embedded activity and the s-values of voxel to voxel according to the formula

$$\langle D(\mathbf{v}_L) \rangle = \sum_{U=0}^{N-1} \tilde{A}_{\mathbf{v}_U} \cdot S(\mathbf{v}_T \leftarrow \mathbf{v}_U). \tag{2}$$

The resulting absorbed dose can be obtained by convolution. One possibility of implementation would be Fast Hartley Transform [9].

For each radionuclide used, a s-value matrix is calculated, which is used as a convolution kernel to determine the mean absorbed dose for each voxel. This method has a big disadvantage, because it cannot take into account inhomogeneities within the tissue.

The voxels have different classes of tissue, such as soft tissue, and thus an associated mass density of $1.004\mathrm{g/cm}^3$. In the voxel in the middle of the cube a radioactive Lu^{177} isotope is placed. With the help of a Monte Carlo simulation software (GAMOS) [12,53,64], decays of $1200 \cdot 10^7$ for the density kernels of the whole body of a patient or $20 \cdot 10^7$ for a dose voxel kernel are then simulated. The result is a dose distribution in the matrix in relation to a decay. This dose voxel kernel is convoluted with the decay distribution to obtain the dose distribution. The unit of absorbed dose is $\mathrm{Gray} = \mathrm{J/kg} = 6.25 \cdot 10^6$ TeV:

$$D = \frac{\partial \mathcal{L}}{\partial m} = \frac{1}{\rho} \frac{\partial \mathcal{L}(\mathbf{r}_T)}{\partial V}. \tag{3}$$

Here $\mathcal{L}(\mathbf{r}_T)$ denotes the deposited energy in a voxel, m the mass of the voxel, ρ the mass density within the voxel and V the volume of the voxel.

From this it can be concluded that the dose distribution in a dose voxel kernel is strongly dependent on the density. The convolution with the decay distribution assumes that the patient consists entirely of soft tissue, which leads to significant errors in the dose distribution.

3   Density Specific Dose Voxel Kernels

Convolving with a dose voxel kernel consisting of soft tissue is a time-effective method for estimating the radiation energy deposited in the body, but density differences of e.g. bone, soft tissue and air are not considered. Since a SPECT/CT has been taken from each patient, information on the density distribution is available.

To use this information, it would be necessary to convolve the different decay distributions with different kernels instead of using only one. First,

the CT image is decomposed into small volume units to obtain the cores of the density distribution. The CT image is scaled to the spatial resolution of the SPECT image. The absorbed radiation dose of the individual density distribution kernels is then calculated using Monte Carlo simulations.

Both the decay distribution and the density distribution from the CT are included in the Monte Carlo simulation, and the desired dose distribution is obtained. We summarise:

1. Whole-body computer tomography of a patient.
2. Separation of the computer tomography into image sections.
3. Calculation of the density distribution kernels.
4. Monte Carlo calculation of the dose voxel kernels.
5. Computation of decay distribution.
6. Convolution to maintain the dose distribution.

Since the calculation and simulation with the different density distribution kernels is immensely complex and time-consuming, but the convolution with only one kernel leads to insufficient results, a new approach is to be presented.

The aim is to learn the transfer, which is calculated in the Monte Carlo simulation, from the data of the different density distributions. Subsequently, the learned filters for the decay distribution are to be used for convolution to obtain the dose distribution. For this procedure, techniques from deep learning will be applied. The following section gives an overview of the most modern approaches and leads to the architecture designed for our experiment.

In the explanations of the neural networks the bold printed notation for vectors is used. The bold capital letters correspond to matrices.

4   Background: Deep Learning

In computer science, deep learning is understood to mean computing models that should be mentioned in the same breath as neural networks. These are multi-layered methods of information processing in which the representation of data with different levels of abstraction is to be learned. These are techniques of feature extraction [36].

These methods have had a strong impact on improving speech recognition, speech synthesis, graphics processing, genomics, object recognition, medical imaging and many other areas. Deep neural networks can recognise structures or even features in very large datasets by using a backpropagation algorithm to calculate how a machine must change its parameters to compute a desired representation in a layer from the previous one [36].

For speech processing, the recurrent neural network initially led to major breakthroughs, since within a processing layer not only afferent but also lateral couplings with backpropagation or anti-hebbian learning rules are trained [41], so that semantic information between morphosyntactic units can be encoded.

For graphics processing, it is the convolutional network that can encode not only sequential information but also spatial references [36].

Machine learning is divided into two broad areas: supervised and unsupervised learning. It is readily apparent that supervised learning is used for classification and regression problems and unsupervised learning is excellent for pre-processing data.

In this experiment, supervised learning is used for regression to learn the transfer function from one spatial representation to another. For deep learning, however, there are many technologies worth mentioning. The following chapter describes the algorithms used in this work. These include: principles from recurrent neural networks, principles of convolution, properties of convolution in combination with pooling and two optimisation techniques.

5   Supervised Learning with Neural Networks

Probably the best researched and most popular technique of machine learning is supervised learning. If a classification or regression task is to be performed, e.g. assigning images to different classes, a data set is first collected for each representative. During the training the machine is shown the images. It calculates a vector of target values, with each dimension corresponding to an assigned class. Either a vector is generated whose values can be zero or one or lie in the range of values between zero and one. The former specifies a fixed class affiliation for each image, the latter a probability.

For example, one would correspond to the membership of a class, zero would not. The aim of the task is to calculate the entered information in such a way that its value is maximised at the corresponding dimension of the vector [36]. Each value in this vector represents a class to which the entered information should be assigned.

Before starting the training, the weights are initialised randomly, so that it is unlikely that the corresponding class will result. Therefore an error function is used to calculate the distance between input and output. This function must be minimized. In neural networks, the machine modifies its parameters by backpropagation at several training steps to minimize this loss function [36]. The parameters are real numbers, often called weights.

The analogy to the dynamics of human brain cells is used as model for the nomenclature of these machines [48]. A single processing cell is therefore also called a neuron [22].

The activity of a layer is given by:

$$\mathbf{y} = h(\sum_n \mathbf{w}_{jn} \cdot \mathbf{x}_n) = h(\mathbf{W}\mathbf{x}). \tag{4}$$

$\mathbf{y}$ is the output signal as an explicit function of the input signals $\mathbf{x}_n$, $n \in \mathcal{D}$, wherein $\mathcal{X} = \{\mathbf{x}_1, ..., \mathbf{x}_n\}$ is the number of samples and $\mathcal{D}$ is the set of underlying data. $\mathbf{w}_{jn}$ is the weight in the j-th layer at the n-th neuron,

accordingly $\mathbf{W}$ is the weight matrix of the network. h is an activation function, e.g. Signum function, tangent hyperbolic or Fermi function.

If the Signum function is used for classification tasks, the condition $|\mathbf{y}| = 1$ follows and for each pattern μ there must be $\mathbf{t}^{(\mu)} = \text{sign}(\mathbf{w} \cdot \mathbf{x}^{(\mu)})$, where $\mathbf{t}$ means a vector pointing in the direction of the corresponding class. Consider $\mathbf{x}_*$ as an ideal example, then the following statements are equivalent $[\mathbf{x}_*^{(\mu)} = \mathbf{t}^{(\mu)} \cdot \mathbf{x}^{(\mu)}] \equiv [\mathbf{w} \cdot \mathbf{x}_*^{\mu} \geq 0, \forall \mu]$. Geometrically this is to be interpreted in such a way that with $\mathbf{w} \cdot \mathbf{x}_* = \mathbf{w} \cdot \mathbf{x} = 0$ a straight line, plane or hyperplane is defined in space, which divides the patterns in the input space into two classes. Thus, a single neuron can only solve linearly separable problems.

Neural networks work with an asymmetric forward coupling. Usually there is no feedback between the output and input layers and no lateral coupling to the neighbouring neurons. The algorithm receives patterns as input and strives for suitable weights.

For the adaptation of these weights a learning step size is defined, which is named η. This gives the δ- learning rule [40, p. 12]:

$$\mathbf{w}_{jn}^{(\text{new})} = \mathbf{w}_{jn}^{(\text{old})} + \eta \delta_j^{(\mu)} \cdot \mathbf{x}_n^{(\mu)}. \tag{5}$$

For the above equation there is always an exact solution if a direction $\mathbf{w}$ can be found for which all projections of the patterns $\mathbf{x} \in \mathcal{D}$ are positive, thus $\exists f : f(\mathbf{w}) \geq 0, \forall \mathbf{x}_*^{(\mu)}$.

For the difficulty of finding a separation level, the following function can be stated: $f(\mathbf{w}) = 1/|\mathbf{w}| \min_\mu (\mathbf{x} \cdot \mathbf{x}_*^{(\mu)})$.

A neuron of this type is suitable for simple binary classification tasks. Multilayer networks with two or more layers can also solve nonlinear separation problems.

6   Learning with Gradient Descent

By means of gradient descent, the weights of a neural network are iteratively adjusted until the algorithm converges, a stop criterion is met or the loss function reaches zero, which is rarely the case in practice.

The loss function is often formalised as distance or divergence between input and output vectors. Simple distance measures or metrics can be used, such as the Euclidean norm or the mean squared error. When using the logarithm (logarithmic compression), the loss function has fewer local minima in comparisson to the Euclidean norm. With the entropy like loss function, the error is directly proportional to the difference between target and actual output: $(\mathbf{y} - \mathbf{x})/(\mathbf{y}(1 - \mathbf{y}))$. The entropy like loss is given by:

$$-\sum_{(\mathbf{x},\mathbf{y}) \in \mathcal{D}} \mathbf{x} \log \mathbf{y} + (1 - \mathbf{x}) \log 1 - \mathbf{y}. \qquad \text{(Entropy like loss)}$$

Fig. 1: Illustration of forward propagation.

In the following, a list shall give an overview of loss functions mentioned above and used for classification problems [58,55,45].

$$\sum_{(\mathbf{x},\mathbf{y})\in\mathcal{D}} ||\mathbf{y}-\mathbf{x}||_2^2, \qquad \text{(Mean squared error)}$$

$$\sum_{(\mathbf{x},\mathbf{y})\in\mathcal{D}} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{p(\mathbf{y})}, \qquad \text{(Kullback-Leibler divergence)}$$

$$\sum_{(\mathbf{x},\mathbf{y})\in\mathcal{D}} \mathbf{x} \log \frac{\mathbf{x}}{\mathbf{y}} + (1-\mathbf{x}) \log \frac{1-\mathbf{x}}{1-\mathbf{y}}. \qquad \text{(Logarithmic loss)}$$

6.1 Forward Propagation

Forward propagation is the representation of a stimulus pattern in the input layer and the successive calculation of the weighted inputs of all subsequent layers up to the output. The algorithm described here was presented in [36].

Let $\mathbf{x}$ for the following examples be a vector of length 4 (see Fig. 1).

Each dimension represents a feature of the sample. The lines symbolise weights from one neuron to the next. Neurons add all the weighted patterns from the previous layer and transform them using the activation function. This process continues until output.

Here the activation function may change. For classification tasks it is necessary to decide whether an example belongs to a certain class or not. This is best done by using probabilities. The softmax function $\text{softmax}(\mathbf{x}) = e^{\mathbf{x}}/\sum_{\mathbf{x}\in\mathcal{D}} e^{\mathbf{x}}$, which outputs values in the interval $[0,1]$, to complete such a task. The output of a neural network with two layers for classification tasks will be denoted as $\mathcal{H}(\mathbf{x} \mid \mathbf{W})$.

6.2 Backpropagation

Backpropagation is an algorithm to correct the weights at each iteration step. The following explanations are taken from the 1987 seminal publication on computational methods using neural networks [40] and the 1986 essay by

Hinton et. al. [56] on learning how to represent data by backward propagated error calculation.

The algorithm of backpropagation corresponds to an iterative gradient descent for a loss function, here in the examples a metric. The losses are calculated from the actual output of the cumulative activity of a neural network and the desired target output. The non-linearity is guaranteed by a transformation using an activation function.

In the following $h(x)$ denotes an activation function.

The factor $1/2$ is prefixed because the activation function is then numerically easier to differentiate, but from an analytical point of view this term can be neglected. The δ-rule is used to update the weights. The total error between the actual output and the target pattern is given by:

$$\mathcal{L}(\mathbf{x}, \mathbf{y} \mid \mathbf{W}) = \frac{1}{2} \sum_{i,\mu} (\delta_i^{(\mu)})^2 = \frac{1}{2} \sum_{i,\mu} (\mathbf{y}_i^{(\mu)} - \mathbf{x}_i^{(\mu)})^2 \tag{6}$$

$$= \frac{1}{2} \sum_{i,\mu} (\mathbf{y}_i^{(\mu)} - h(\sum_j \mathbf{W}_{ij} \mathbf{v}_i^{(\mu)}))^2 \tag{7}$$

$$= \frac{1}{2} \sum_{i,\mu} (\mathbf{y}_i^{(\mu)} - h(\sum_j \mathbf{W}_{ij} h(\sum_n \mathbf{w}_{jn} x_n^{(\mu)})))^2. \tag{8}$$

$\mathbf{v}_i^{(\mu)}$ denotes the activity at the i-th neuron while the pattern μ is presented. It is assumed that the loss function is continuous and differentiable.

For the first layer $\delta\mathbf{w}_{jn}$ is calculated as follows:

$$\delta\mathbf{w}_{jn} = \eta \frac{\partial \mathcal{L}}{\partial \mathbf{w}_{jn}} = \eta \sum_\mu \frac{\partial \mathcal{L}}{\partial \mathbf{v}_j^{(\mu)}} \frac{\partial \mathbf{v}_j^{(\mu)}}{\partial \mathbf{w}_{jn}} \tag{9}$$

$$= \eta \sum_{i,\mu} (\mathbf{y}_i^{(\mu)} - \mathbf{x}_i^{(\mu)}) h'(\mathbf{v}_i^{(\mu)}) \mathbf{W}_{ij} h'(\mathbf{v}_j^{(\mu)}) \mathbf{x}_n^{(\mu)} \tag{10}$$

$$= \eta \sum_{i,\mu} \delta_i^{(\mu)} \mathbf{W}_{ij} h'(\mathbf{v}_j^{(\mu)}) \mathbf{x}_n^{(\mu)} = \eta \sum_n \delta_j \mathbf{x}_n^{(\mu)}. \tag{11}$$

In the output layer, the forward propagation error results as follows:

$$\delta_i^{(\mu)} = (\mathbf{y}_i^{(\mu)} - \mathbf{x}_i^{(\mu)}) h'(\sum_j \mathbf{W}_{ij} \mathbf{v}_i^{(\mu)}). \tag{12}$$

This error is the error propagated by the hidden layers:

$$\delta_j^{(\mu)} = h'(\sum_j \mathbf{W}_{ij} v_i^{(\mu)}) \sum_i \mathbf{W}_{ij} \delta_i^{(\mu)} \tag{13}$$

$$= h'(\sum_j \mathbf{W}_{ij} \mathbf{v}_i^{(\mu)}) \sum_i \mathbf{W}_{ij} (y_i^{(\mu)} - x_i^{(\mu)}) h'(\mathbf{v}_i^{(\mu)}). \tag{14}$$

A collection of training samples is called an ensemble, a complete arrangement of all samples is called an epoch. The selection of the samples during each optimisation run is random. The number of samples used for an iteration step is the batch size and a free parameter. The convergence of the algorithm is not proven, neither is a well-defined stop criterion. In general, the gradient $||\nabla_{\mathbf{w}}\mathcal{L}||_2$ should take a small value, or the weight vector should not change significantly over a certain number of iteration steps $||\sum_t \mathbf{w}(t+1)-\mathbf{w}(t)|| < \epsilon$.

In summary, the following algorithm results for a network of m layers, where $m \in \{1, ..., M\}$ and $\mathbf{v}_i^{(M)}$ denotes the activity of the i-th neuron in the m-th layer.

The activity in the input layer corresponds to $\mathbf{v}_i^{(M)} = \mathbf{x}_i$.

$\mathbf{w}_{ij}^{(M)}$ is the synaptic weight of the afferent coupling $(\mathbf{v}_j^{(m-1)} \to \mathbf{v}_i^{(m)})$:

First the weights are initialised with small values, e.g. $10^{-2} \leq \mathbf{w}_{ij} \leq 10^{-4}$ for all i, j. Then the first pattern $\mathbf{x}_i^{(\mu)}$ is presented so that each property of the pattern is encoded in a numeric value $(\mathbf{v}_i^{(0)} = \mathbf{x}_i^{(\mu)}, \forall i, \mu)$. Now the activities of the higher layers are calculated, with

$$\mathbf{v}_i^{(m)} = h(\sum_j \mathbf{w}_{ij}^{(m)}\mathbf{v}_j^{(m-1)}), \qquad (15)$$

until the last layer $\mathbf{v}_i^{(M)}$.

At this point the forward propagation is completed and the error calculation algorithm takes effect. Usually the error of the output layer is calculated first, so that the terms, which originate from the derivation by chain rule, can be reused for previous layers.

The error of the output is given by

$$\delta_i^{(M)} = h'(v_i^{(M)})(\mathbf{y}_i^{(\mu)} - \mathbf{v}_i^{(M)}), \qquad (16)$$

where $A_i^{(M)}$ represents the activity at the i-th neuron in the m-th layer. This also results in the defects of deeper layers:

$$\delta_i^{(m-1)} = h'(v_i^{(m-1)}) \sum_j \mathbf{w}_{ij}^{(M)}\delta_j^{(M)}. \qquad (17)$$

Finally, all weights are improved using the δ-rule and replaced:

$$\mathbf{w}_{ij}^{(\text{new})} = \mathbf{w}_{ij}^{(\text{old})} + \delta\mathbf{w}_{ij}. \qquad (18)$$

This process is repeated for all patterns.

7 Development and Variants of Recurrent Neural Networks

Recurrent neural networks are an architectural form that refers to a multitude of calculation models. Essentially, these models differ only slightly in architecture from classical neural networks.

Fig. 2: Illustration of the feed forward algorithm on the left side and the course of the backpropagation of the error on the right side.

However, this difference is manifested in their mathematical definition. In contrast to afferent neural networks with classical forward coupling, recurrent neural networks have recursive structures [44], which are calculated iteratively. Using lateral couplings, recurrent neural networks have also been the starting point for residual connections, which have a significant influence on the direction of the gradient during training. The residual connections first appeared in the residual neural networks and were later translated into U-networks [23,14]. These residual connections are used as bypasses of two layers during deep learning.

A recurrent neural network behaves in such a way that it develops its own dynamic activation behaviour within the recurrent structures. If we understand afferent networks as functions $f : U \to V, x \mapsto f(x)$ of any topological spaces U, V, a recurrent neural network is a dynamic system [44].

Furthermore, the recurrent network has a kind of nonlinear history to the already trained input patterns. It can be understood as short-term memory [44]. The couplings can be either a direct feedback (recursion within a neuron), an indirect feedback (recursion to the previous neuron), a lateral feedback (to the neighbouring neuron) or a completely connected graph.

Such models are based on a loss function, which is to be minimised by a stochastic process. The afferent and recurrent couplings are symmetrical [44].

The best known representatives of these networks are the Hopfield networks [29], the Boltzmann machines [2] or the recently discovered deep belief networks [24]. Typical applications in this area are associative storage, data compression, modelling of distributions and learning of time-dependent patterns [44,24].

Fig. 3: Simple recurrent neural network with one input, output and one hidden layer. The green dots correspond to associative memories in the form of sums.

7.1   Simple Recurrent Neural Networks

Simple recurrent neural networks work with backpropagation. In 1991 Elman described how this can be applied to recurrent couplings [20].

In the hidden layers in the neural network lateral couplings can be introduced, which lead to context units. These context units are neurons whose weight remains constant at one. The number of these context units corresponds to the number of units in the hidden layer with which the lateral coupling exists. The hidden units then reconnect to the same neurons from which they obtain their information. Each context neuron receives the weighted information from exactly one neural unit. The contextual units contain an exact copy of the information of the coupled neurons at iteration time t. At iteration time $t + 1$ the content is returned to the coupled layer. Each contextual neuron is afferently connected to each neuron of the hidden layer. Addition and weighting is done by forward propagation.

These contextual units are regarded as associative memory because they contain the information processed by the neural network during a single batch. The information at iteration time t causes the weighted output of the hidden layer, which in turn is stored in the contextual units for the next iteration step $t + 1$.

7.2   Learning of Sequences with Recurrent Neural Networks

Modified recurrent architectures are used to learn time series. A time series is any dependence on time, such as a share price, the learning of certain motion sequences or medical applications: Heartbeat, blood pressure or brain activity.

For instance, how can we learn if our blood pressure rises or falls depending on previous activity?

Let's assume that there is data from different patients. One of the patients was hospitalised and his values have been available for a long time,

unlike patients who take a 24-hour measurement. The first challenge is to standardize the dimension, as each subject has a different set of data [35].

One way to unify the dimensions would be to use the largest data matrix as the initial dimension. Then all other matrices are formatted to the same size by filling in the missing values with zeros.

It is advantageous to fill all samples with zeros in the same way.

Starting from a corner or the centre of the matrix. This procedure is called zero-padding. Then the different dimensional input of information can be reduced to a uniform size by convolution [35]. In order to project into lower dimensional subspaces and to consider the most important features, it is possible to chain a pooling layer after convolution. With pooling, however, the loss of information may be too high for image recognition. For the described scenario this method is rather unsuitable, because one wants to draw conclusions about values on the basis of precise global temporal information.

An alternative possibility to deal with input information of different size is the use of recurrent networks.

$\mathbf{x}_0, ..., \mathbf{x}_T$ are measuring points of the blood pressure at the time t of the total time T. $\mathbf{v}_0, ..., \mathbf{v}_T$ denote the activities in the neural network, often also called local receptive fields, similar to the notation of dense neural networks in the previous section. The weight from the input to the first deeper layer is called w, the weight between the deeper layers is called u and the weights from the deeper layer to the output are called v. These weights are scalars and equal for all neurons, so they are shared. After forward propagation the weights are optimised by gradient descent on the loss function.

The activities in the hidden layers can then be successively adjusted as follows, where h denotes the activity in the local receptive field:

$$\mathbf{y} = v\mathbf{v}_T, \tag{19}$$

$$\mathbf{v}_t = h(u\mathbf{v}_{t-1} + w\mathbf{x}_t), \text{ for } t = 1, ..., T, \tag{20}$$

$$\vdots \tag{21}$$

$$\mathbf{v}_0 = h(w\mathbf{x}_o). \tag{22}$$

8   Long Short-term Memory in Recurrent Neural Networks

Unfortunately, backpropagation over time is extremely problematic in convergence behaviour, as it often leads to exploding or disappearing gradients, due to the different spatial information density in the respective space dimensions of the input vectors [27].

Here is a brief example: for healthy people the blood pressure is usually stable during the night. In the morning, the person generally has a slightly

Fig. 4: Recurrent neural network. The input can have any number of dimensions. Missing values are filled up with zeros. The non-linear activity within the neurons with lateral connections – •-marked – is a sigmoid function. As output a code with information about the temporal correlations is generated. The output needs not to have the same dimension as the input.

elevated blood pressure, which fluctuates throughout the day. If the data correspond to this scenario, the system would always get a high value for each data point for training in the first dimension, i.e. in the morning.

This would correspond to a low gradient, which would hardly change during the optimisation process [35].

This behaviour is not desirable, as a similar weighting is expected for all times of the day.

Le formulates this as follows: If the gradient is extremely small in one dimension and extremely large in other dimensions, the landscape of the loss function looks like a valley, with very steep walls and a deep basin [35]. The reason for this behaviour of the gradient, Le continues, is the use of sigmoid activation functions.

The iterative improvement of the loss function is also multiplied by the derivative of the sigmoid function. This could quickly lead to saturation.

One way to prevent this are piecewise linear functions as activation [33]. This prevents the formation of extreme values within the computation of the gradient.

LSTM (Long-Short Term Memories) are probably the biggest improvement for recurrent learning [28,26,35]. The idea behind the LSTM is to change the recurrent structure so that the gradient remains stable during backpropagation. The recurrent couplings represent integration over time [35].

Assuming that the data have the values $\mathbf{x}_t$ at time t and the lower units had the activity $\mathbf{v}_{t-1}$ at the previous time step, the patterns that make up associative memory will then have the following values:

$$\mathbf{m}_t = \alpha \cdot \mathbf{m}_{t-1} + \beta \cdot h(\mathbf{x}_t, \mathbf{v}_{t-1}). \tag{23}$$

The symbol $(\,\cdot\,)$ denotes element-wise multiplication of two vectors. So the new weight of the associative memory is a weighted linear combination of $\mathbf{m}_{t-1}$ and $h(\mathbf{x}_t, \mathbf{v}_{t-1})$. This corresponds to two different paths which the

gradient can follow during backpropagation if one of the two components takes on extreme values. Often a sigmoid function like tanh is chosen for h to prevent exploding gradients. From this associative memory the state of the hidden layers can be calculated:

$$\mathbf{v}_t = \gamma \cdot h(\mathbf{m}_t), \tag{24}$$

where $\mathbf{v}_t$ and $\mathbf{m}_t$ are required for the subsequent time steps to calculate the activity.

This in turn gives the gradient more possibilities to take different paths.

The terms α, β and γ are often called gates. These gates are factors that weight the contribution of the respective term – the activity to the step t. They are implemented as follows:

$$\alpha_t = h(\mathbf{W_{x\alpha}x}_t + \mathbf{W_{v\alpha}v}_{t-1} + \mathbf{W_{m\alpha}m}_{t-1} + \mathbf{b}_\alpha), \tag{25}$$
$$\beta_t = h(\mathbf{W_{x\beta}x}_t + \mathbf{W_{v\beta}v}_{t-1} + \mathbf{W_{m\beta}m}_{t-1} + \mathbf{b}_\beta), \tag{26}$$
$$\gamma_t = h(\mathbf{W_{x\gamma}x}_t + \mathbf{W_{v\gamma}v}_{t-1} + \mathbf{W_{m\gamma}m}_{t-1} + \mathbf{b}_\gamma), \tag{27}$$
$$\mathcal{H}(\mathbf{x}_t \mid \mathbf{v}_{t-1}) = h(\mathbf{W_{xm}x}_t + \mathbf{W_{vm}v}_{t-1} + \mathbf{b_m}). \tag{28}$$

With these parameters and architecture the following variants are possible:

– The activation function can be used to scale the values to be added, e.g. between zero and one.
– The connections to the recurrent neurons can be used in the input layer so that the LSTM network behaves like a convolutional network. The gradients can converge to the values of each time step.
– The connections to the recurrent neurons can be initialised so that in extreme cases a simple recurrent network is trained.

Fig. 5: Graphical illustration of LSTM cells in a recurrent neural network. The curved path is intended to illustrate a non-linear path. The straight path should illustrate a linear path. •-market units have linear activation functions. •-market units behave as in Eq. 24. The •-marked units have sigmoid activation functions.

9 Convolutional Neural Networks

Convolutional neural networks are a technique from the field of deep learning, which have significantly influenced and developed image processing in recent years. Especially the image classification [59], as well as the object recognition [51,52] and the semantic mapping of image contents [43] have been improved. Convolutional networks show impressive performance in standardised benchmark tests with data sets such as ImageNet or MSCOCO.

The ability of semantic representation makes them equally valuable for learning a function for image-to-image transfer.

In the following the principles of convolutional neural networks are explained and common models are presented.

In the mathematical discipline of functional analysis, convolution is an operator that maps two functions $f(x)$ and $g(x)$ to a third $(f * g)(x)$.

Convolution is a very useful tool to find out how much information of one function is contained in another.

The convolution is defined as $(f * g)(x) = \int_0^x f(\tau) \cdot g(x - \tau)d$. The value of the given integral is a variable function of x, i.e. an integral function. The convolution is to be understood in such a way that the result function $(f * g)(x)$ indicates how much the value of the function $g(x)$ at the iteration time $x - \tau$ is contained in the function value of the weighting function $f(x)$ at the iteration step x.

Neural networks also use the principle of convolution. For the convolution of images one processes discrete quantities. For such a case the convolution of matrices with kernels is explained in the next section.

The information about a certain state is written into a convolution matrix. Convolution matrices are quadratic with an odd number of column and/or row vectors. For a set of pixels in an image or any discrete set of data points, the convolution operation can be formalised as follows

$$\mathbf{T}^{(x,y)} = \sum_n \sum_m \mathbf{K}_{nm} \mathbf{X}_{(x+n-\hat{z})(y+m-\hat{z})}, \tag{29}$$

$\mathbf{T}$ is the convolution matrix, where the coordinates $x, y \in \mathbb{N}$ define the respective entry. $\mathbf{K}$ denotes the kernel used for the convolution. $\mathbf{X}$ denotes the matrix of data points of the original data set.

If an image is convolved, $\mathbf{X}$ corresponds to the original image. $\hat{z}$ indicates the matrix centre of the kernel.

9.1 Translation Invariance of Convolution

Deep learning is used for large amounts of data. It is thus crucial that the amount of data remains manageable for the computing device.

Further it may happen, as it is often the case in medical imaging, that samples have very high dimension.

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} * \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 4 & 3 & 4 & 1 \\ 1 & 2 & 4 & 3 & 3 \\ 1 & 2 & 3 & 4 & 1 \\ 1 & 3 & 3 & 1 & 1 \\ 3 & 3 & 1 & 1 & 0 \end{pmatrix}$$

$$\mathbf{X} \qquad\qquad \mathbf{K} \qquad\qquad \mathbf{T}^{(x,y)}$$

Fig. 6: Graphical illustration of the convolution operation. The convolution is performed with a filter $- \bullet -$ with filter size 3×3. The convolution is first performed from left to right. The values within the $\bullet$-marked matrix are summed up and written into the $\bullet$-marked field.

Accordingly, for each sample, each neuron would take a sum of the weighted inputs for each dimension. Assuming there would be an image of 150px·150px, then the input space would have dimension 22500 and each neuron would add up as many components for the activity.

A possible solution to the problem of parameter size is to share the weights. An example will illustrate the principle. Starting from any layer and nine incoming weights $\mathbf{w}_1, ..., \mathbf{w}_9$, three of which have the same (or a similar) value:

$$\mathbf{w}_1 = \mathbf{w}_2 = \mathbf{w}_3, \tag{30}$$

$$\mathbf{w}_4 = \mathbf{w}_5 = \mathbf{w}_6, \tag{31}$$

$$\mathbf{w}_7 = \mathbf{w}_8 = \mathbf{w}_9. \tag{32}$$

This is intended to take the activities to the next layer. As the weights of the three connections are equal, they can share the value. This means that not the entire data set has to be processed. Instead of saving all weights $\mathbf{w}_1, ..., \mathbf{w}_9$, it is sufficient to process one weight from each input, e.g. $\mathbf{w}_1, \mathbf{w}_4$ and $\mathbf{w}_7$.

This idea of dividing weights corresponds to the operation of convolution, in which a matrix, i.e. a set of weights, is applied to several positions of the input signals of a layer.

These matrices are also called filters. Filters have another useful property for data processing. Within the pooling layer, the order of the information in the neurons does not matter. For natural data, translation variance is the most common cause of noisy data.

In max-pooling, only one of a defined number of neurons with the highest activity is transferred to the next layer. Therefore, such layers are particularly suitable for natural data, even though care has to be taken for outliers. In this process of convolution and pooling, the dimension in the data is drastically reduced, taking into account the most important information.

Many recent architectures also exhibit another type of transformation, namely local contrast normalisation [35]. This layer receives as input the activity of the max-pooling layer. The mean value is subtracted and the standard deviation of the incoming neurons is divided. This enables a brightness invariance that corresponds to a statistical normalisation, which is a useful tool for image segmentation and object recognition.

This construction is trained for forward propagation analogous to conventional neural networks. For backpropagation the weights shared by the grouped neurons are averaged:

$$\mathbf{w}_1^{(\text{new})} = \mathbf{w}_1^{(\text{old})} - \alpha(\frac{\partial \mathcal{L}}{\partial \mathbf{w}_1^{(\text{old})}} + \frac{\partial \mathcal{L}}{\partial \mathbf{w}_4^{(\text{old})}} + \frac{\partial \mathcal{L}}{\partial \mathbf{w}_7^{(\text{old})}}), \tag{33}$$

$$\mathbf{w}_4^{(\text{new})} = \mathbf{w}_4^{(\text{old})} - \alpha(\frac{\partial \mathcal{L}}{\partial \mathbf{w}_1^{(\text{old})}} + \frac{\partial \mathcal{L}}{\partial \mathbf{w}_4^{(\text{old})}} + \frac{\partial \mathcal{L}}{\partial \mathbf{w}_7^{(\text{old})}}), \tag{34}$$

$$\mathbf{w}_7^{(\text{new})} = \mathbf{w}_7^{(\text{old})} - \alpha(\frac{\partial \mathcal{L}}{\partial \mathbf{w}_1^{(\text{old})}} + \frac{\partial \mathcal{L}}{\partial \mathbf{w}_4^{(\text{old})}} + \frac{\partial \mathcal{L}}{\partial \mathbf{w}_7^{(\text{old})}}). \tag{35}$$

9.2   Pooling

In the literature so far, the most popular approaches are max-, min- and average pooling as well as stochastic pooling. More complex pooling layers, such as spatial pyramid-shaped pooling, have recently become the focus of investigations [7,8,66,38].

They are used after convolution to identify conspicuous features along certain dimensions of the processed matrix. A convolution network calculates filters that encode features of the presented samples, stores them in the kernel matrices and uses them in the prediction. In contrast to the traditional activation function, where all values of the matrix are processed separately, pooling allows learning the activity along a subspace of the matrix.

Pooling is a technique in which a fixed window size is predefined, which normally cannot be changed. This size can correspond to a vector or matrix for each dimension that is pooled. Normally a n^2 matrix is used for two-dimensional pooling.

With max-pooling this predefined size is created like a mask at the first row and column input of the input matrix. From this area the largest activity is selected, i.e. the maximum. This activity corresponds to the first entry of the output matrix. This process is repeated iteratively for the entire sample.

A step size must be defined, which specifies by how many pixels the mask is moved.

10   Stochastic Gradient Descent

The stochastic gradient descent is one of many optimisation methods used in neural networks to realize backpropagation. Up to now, different techniques for the composition of neural networks have been treated. However,

the optimisation of the parameters also plays a central role. In contrast to the classical gradient descent, the stochastic gradient descent is a strong simplification regarding the requirements for its computation but also regarding accuracy.

There is no free lunch [6]. Instead of calculating the gradient $\nabla_{\mathbf{w}}(f_t)$ over the whole dataset, the gradient is computed over a set of subsamples. In this case we choose a single $\mathbf{x}$ for explanatory purpose:

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \eta_t \nabla_{\mathbf{w}} \mathcal{L}(\mathbf{x}, \mathbf{w}_t). \tag{36}$$

In contrast the stochastic gradient descent can be written as:

$$\mathbf{w}_{t+1} = \mathbf{w}_t + \eta_t \frac{1}{n} \sum_{i=1}^{n} \nabla_{\mathbf{w}} \mathcal{L}(\mathbf{x}_i, \mathbf{w}_t). \tag{37}$$

For the gradient descent in Eq. 37 the argument of the gradient is replaced by a stochastic one, so that Eq. 36 behaves in the same way as Eq. 37, i.e. taking the limit of samples the same result is obtained [6].

Since the stochastic gradient descent has no relation to examples from previous iteration steps, it can be calculated online.

The stochastic gradient descent is a direct minimisation of the expected risk or the expected losses and is calculated from the actual distribution of the samples. The convergence behaviour of stochastic gradient descent has been extensively investigated in the literature. Understandably, it is limited by the noisy estimation of the actual gradient. If the gradient descends too slowly, the variance of the parameter $\mathbf{w}_t$ decreases as slowly.

Under sufficient regularisation, the best possible convergence behaviour corresponds to $\eta_t \sim t^{-1}$.

The expectation of the residual error converges at $\mathcal{L}\rho \sim t^{-1}$ [6,46].

10.1 Stochastic Gradient Descent
with Adaptive Lower Order Momentum

The growing interest in machine learning requires robust solutions to optimise loss functions of non-stationary data. The method was designed to combine two very successful approaches.

These include the `AdaGrad` algorithm [18] and `RMSProp` algorithm [25]. `AdaGrad` showed good performance with sparingly coded data up to this point, while `RMSProp` was very good with online and stationary data [32]. $\mathbf{g}_t = \nabla_{\boldsymbol{\Theta}} \mathcal{L}_t(\boldsymbol{\Theta})$ denotes the gradient of $\mathcal{L}(\boldsymbol{\Theta})$ calculated at the iteration time t. $\mathcal{L}_1, \ldots, \mathcal{L}_T$ denote the respective loss function at the time step $t \in \{1, \ldots, T\}$. β_1 and β_2 are regularisation terms [32].

When calculating the gradient for a new time step, it is moved towards the moments of first and second order.

Some advantages of the `ADAM` algorithm:

Algorithmus 1 `ADAM` algorithm from the paper by Kingma and Lei Ba [32], $\mathbf{g}_t^2$ is the element-wise squaring of $\mathbf{g}_t \cdot \mathbf{g}_t$. In their essay the two authors give as good initial parameters $\alpha = 0.001$, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$. In the following algorithm all vector operations are to be understood element-wise.

1: **Proc.:** ADAPTIVE LOWER ORDER MOMENTUM
2: **Param.:** α: step-size
3: **Param.:** $\beta_1, \beta_2 \in [0, 1]$: exponential decays for momentum terms
4: **Param.:** $\mathcal{L}(\mathbf{\Theta})$: stochastic loss function with parameters $\mathbf{\Theta}$
5: **Param.:** $\mathbf{\Theta}_0$: (*initialize parameters*)
6:      $\mathbf{m}_0 \leftarrow 0$ (*initialize vector for first order momentum*)
7:      $\mathbf{v}_0 \leftarrow 0$ (*initialize vector for second order momentum*)
8:      $t \leftarrow 0$(iteration)
9:      **while** $\mathbf{\Theta}_t$ not converged **do**
10:         $t \leftarrow t + 1$
11:         $\mathbf{g}_t \leftarrow \nabla_{\mathbf{\Theta}} \mathcal{L}_t(\mathbf{\Theta}_{t-1})$ (*compute gradient at time t*)
12:         $\mathbf{m}_t \leftarrow \beta_1 \cdot \mathbf{m}_{t-1} + (1-\beta_1) \cdot \mathbf{g}_t$ (*update estimate of first order momentum*)
13:         $\mathbf{v}_t \leftarrow \beta_2 \cdot \mathbf{v}_{t-1} + (1-\beta_2) \cdot \mathbf{g}_t^2$ (*update estimate of second order momentum*)
14:         $\widehat{\mathbf{m}}_t \leftarrow \mathbf{m}_t / (1 - \beta_1^t)$ (*error correction of first order momentum*)
15:         $\widehat{\mathbf{v}}_t \leftarrow \mathbf{v}_t / (1 - \beta_2^t)$ (*error correction of second order momentum*)
16:         $\mathbf{\Theta}_t \leftarrow \mathbf{\Theta}_{t-1} - \alpha \cdot \widehat{\mathbf{m}}_t / (\sqrt{\widehat{\mathbf{v}}_t} + \epsilon)$
        return θ_t

- The number of update steps is invariant with respect to scaling.
- The change of the gradient is almost completely bound to step size α.
- No stationary data are required.
- The algorithm also works with sparse coded matrices.
- The algorithm performs a stepwise approximation (simulated annealing).

In Alg. 1 the term $\mathcal{L}(\mathbf{\Theta})$ corresponds to a loss function that depends on a parameter vector $\mathbf{\Theta}$ [32]. Altogether the expected value of the function $\mathcal{L}(\mathbf{\Theta})$ shall be minimized.

The algorithm is described as stochastic, using randomly selected subsets of samples from the data set from which then the gradient is calculated. Another way of understanding the algorithm as stochastic would be the inherent noise of the function to be optimised [32].

Kingma and Lei Ba also tested the algorithm under experimental conditions with convolutional neural networks. Of special interest for our project is the integration of the Nesterov momentum into the optimisation algorithm. We discuss next this extension, which will eventually be used for the experiment.

10.2   Stochastic Gradient Descent with Nesterov Adaptive Lower Order Momentum

If one wishes to improve an existing deep learning system, the literature suggests a number of approaches [17]: The network is made deeper, common

---

Algorithm 2 Gradient descent with momentum [17].

1: $\mathbf{g}_t \leftarrow \nabla_{\boldsymbol{\Theta}_{t-1}} \mathcal{L}(\boldsymbol{\Theta}_{t-1})$ (*computation of the gradient* $\mathbf{g}$ *at time t*)
2: $\mathbf{m}_t \leftarrow \mu \mathbf{m}_{t-1} + \mathbf{g}_t$ (*computation of the momentum* $\mathbf{m}$ *with cost* μ)
3: $\boldsymbol{\Theta}_t \leftarrow \boldsymbol{\Theta}_{t-1} - \eta \mathbf{m}_t$ (*momentum with step-size* η)

---

recurrent units are replaced by LSTM cells or methods of preprocessing are used to present the data as noise-free as possible [17].

Furthermore, it could be shown that a careful selection of the distribution during the initialization of convolutional neural networks is crucial for the early development of the gradient and thus also for the performance of the system [63]. Thereby the parameters of the system are adjusted by optimisation algorithms in each iteration step.

For the acceleration of the gradient descent a momentum term was suggested [50]. Similar to the physical understanding of momentum, a kind of velocity vector, which indicates a corresponding spatial direction, is used for this purpose. This vector m is multiplied by a cost term (usually a constant, here called μ). This has the advantage that the gradient is drastically slowed down in spatial directions that show a rapid change and a strong fall and vice versa accelerated in directions in which the gradient oscillates [17].

Sutskever et al. already proposed the integration of the Nesterov momentum into the algorithm of gradient descent [17,63]. Thus, if we look at the first equation of the momentum calculation, we see that the update step is equivalent to a gradient descent in the direction of the momentum to time step $t-1$ and a further step is made in the direction of the gradient to time t (see Alg. 1):

$$\boldsymbol{\Theta}_t = \boldsymbol{\Theta}_{t-1} - (\mu \mathbf{m}_{t-1} + \alpha_t \mathbf{g}_t). \tag{38}$$

The momentum term $\mu \mathbf{m}_{t-1}$ does not depend on the current gradient ($\mathbf{m}_t \leftarrow \mu \mathbf{m}_{t-1} + \alpha_t \mathbf{g}_t$), but only on the gradient of the last iteration step (see Alg. 1). Sutskever et al. therefore suggested [63] that the parameters should be adjusted even before the gradient is calculated in order to achieve a quality improvement compared to the classical gradient descent:

$$\mathbf{g}_t \leftarrow \nabla_{\boldsymbol{\Theta}_{t-1}} \mathcal{L}_t(\boldsymbol{\Theta}_{t-1} - \mu \mathbf{m}_{t-1}), \tag{39}$$

$$\mathbf{m}_t \leftarrow \mu \mathbf{m}_{t-1} + \alpha_t \mathbf{g}_t, \tag{40}$$

$$\boldsymbol{\Theta}_t \leftarrow \boldsymbol{\Theta}_{t-1} - \mathbf{m}_t. \tag{41}$$

While for the classical gradient descent the momentum can be understood as a weighted sum of the last update steps, for ADAM it corresponds to a weighted average of the gradients of the last iterations [32,17]:

$$\mathbf{m}_t \leftarrow \mu \mathbf{m}_{t-1} + (1 - \mu) \mathbf{g}_t, \tag{42}$$

$$\boldsymbol{\Theta}_t \leftarrow \boldsymbol{\Theta}_{t-1} - \alpha_t \frac{\mathbf{m}_t}{(1 - \mu^t)}. \tag{43}$$

Algorithmus 3 Estimate of the nesterov adaptive momentum [17].

1: **Proc.:** NESTEROV ADAPTIVE LOWER ORDER MOMENTUM
2: **Param.:** $\alpha_0, \ldots, \alpha_T; \mu_0, \ldots, \mu_T; \nu; \eta$: initialize parameters
3:     $\mathbf{m}_0, \mathbf{n}_0 \leftarrow 0$ (*vectors of first / second order momentum*)
4:     **while** $\boldsymbol{\Theta}_t$ not converged **do**
5:         $\mathbf{g}_t \leftarrow \nabla_{\boldsymbol{\Theta}_{t-1}} \mathcal{L}_t(\boldsymbol{\Theta}_{t-1})$
6:         $\mathbf{m}_t \leftarrow \mu_t \mathbf{m}_{t-1} + (1 - \mu_t)\mathbf{g}_t$
7:         $\mathbf{n}_t \leftarrow \nu \mathbf{n}_{t-1} + \mathbf{g}_t^2(1 - \nu)$
8:         $\hat{\mathbf{m}} \leftarrow \mu_{t+1}\mathbf{m}_t/(1 - \prod_{i=1}^{t+1} \mu_i) + (1 - \mu_t)\mathbf{g}_t/(1 - \prod_{i=1}^{t} \mu_i)$
9:         $\hat{\mathbf{n}} \leftarrow \nu \mathbf{n}_t/(1 - \nu^t)$
10:         $\boldsymbol{\Theta}_t \leftarrow \boldsymbol{\Theta}_{t-1} - \alpha_t/(\hat{\mathbf{m}}_t \sqrt{\hat{\mathbf{n}}_t + \epsilon})$
         return $\boldsymbol{\Theta}_t$

The use of gradients instead of updates allows the algorithm to change direction steadily, even if the learning rate is already very low, leading to more refined convergence [32,17]. The algorithm also corrects the initialisation error caused by initialising the moments with zero values by the denominator $(1 - \mu^t)$ [17]. We would like to follow the paper by Dozat at this point [17] and explain how to integrate these modifications into the ADAM algorithm.

μ denotes the cost of a direction in which the gradient should be steered during an update step. Instead of first calculating the gradient, then returning to the initial values of the parameters, and then again taking a step towards the momentum, the momentum is calculated only once at the time step $t+1$, during the update of the time step t. This is done as follows [17,47]:

$$\mathbf{g}_t \leftarrow \nabla_{\boldsymbol{\Theta}_{t-1}} \mathcal{L}_t(\boldsymbol{\Theta}_{t-1}), \tag{44}$$

$$\mathbf{m}_t \leftarrow \mu_t \mathbf{m}_{t-1} + \alpha_t \mathbf{g}_t, \tag{45}$$

$$\boldsymbol{\Theta}_t \leftarrow \boldsymbol{\Theta}_{t-1} - (\mu_{t+1}\mathbf{m}_t + \alpha_t \mathbf{g}_t). \tag{46}$$

The same modification can now also be used for ADAM (see Alg. 1):

$$\boldsymbol{\Theta}_t \leftarrow \boldsymbol{\Theta}_{t-1} - \alpha_t \left(\frac{\mu_t \mathbf{m}_{t-1}}{(1 - \prod_{i=1}^{t} \mu_i)} + \frac{(1 - \mu_t)\mathbf{g}_t}{(1 - \prod_{i=1}^{t} \mu_i)} \right), \tag{47}$$

$$\boldsymbol{\Theta}_t \leftarrow \boldsymbol{\Theta}_{t-1} - \alpha_t \left(\frac{\mu_{t+1} \mathbf{m}_t}{(1 - \prod_{i=1}^{t+1} \mu_i)} + \frac{(1 - \mu_t)\mathbf{g}_t}{(1 - \prod_{i=1}^{t} \mu_i)} \right). \tag{48}$$

This results in the Nesterov adaptive lower order momentum. In Alg. 3 the complete algorithm is summarised once again.

The algorithm was empirically compared with two other optimisation algorithms, including ADAM and the classical SGD. One hundred training runs were tested for each algorithm. For a minimum improvement within 15 epochs of $\epsilon \geq 10^{-6}$ on the loss function Adam needed on average 300 ± 24 epochs with corresponding standard deviation, SGD 400 ± 38 epochs and the NADAM 290 ± 30 epochs. Since this algorithm has outperformed both SGD and ADAM

in convergence by determining the number of epochs required, we decided to use the same algorithm.

11   Numerical Experiments

For the experiment a transfer has to be learned from the tissue density kernels to dose voxel kernels. The starting point are previously calculated matrix kernels of the different tissue densities of a patient's tissue. Using Monte Carlo simulation, the absorbed radiation dose of the tissue was computed as described in Sect. 1. The decay distribution from imaging is concolved with the corresponding matrix kernels to obtain a new map of dose distribution. The status quo is a convolution according to tissue classes. This means that a filter exists for each tissue class.

A central problem of this approach is mixed tissue, which is present in the body but cannot be adequately represented by the fragmentation of the CT image into density distribution kernels. The dose distribution of a tissue with different density is thus estimated with a high error.

This experiment shows how machine learning methods can be used to use neural networks to obtain a better estimate of the absorbed radiation dose for a given tissue and isotope. Thereby, neural network architectures from image recognition and image segmentation are used. We briefly cover the state of the art and present related approaches.

11.1   Related Work on Image Segmentation

In image segmentation, the task is to recognise and delimit a semantically related part of an image. Various scientific disciplines use image segmentation and image recognition for analysis. Especially in medicine, these techniques are often used as diagnostic tools. Particular progress has been made with the introduction of convolutional neural networks, which have exceeded the previous standard for visual tasks in recent years [21,34]. At that time, convolutional neural networks were mainly used for classification tasks, but image segmentation covers more. By assigning each pixel to a certain class, which corresponds to semantic detection within an image, one can rephrase the segmentation problem in terms of classification [14]. This made convolutional neural networks interesting for image segmentation.

From this point on there was a great rush to exceed the benchmarks. Finally, fully convolutional networks, in short `FCNN`, which were purely convolutional networks without dense layers or other architectural elements, were declared the standard [11,15,39]. `FCNN`s are widely used for depth estimation [43], for image restoration [19], for reconstruction of images to improve resolution [16] and last but not least for estimating densities [42], which is why they are the focus of this study.

The combination of `LSTM` cells with convolutional neural networks led to even more powerful architectures for extraction of features [31]. A special architectural form of convolutional neural networks is the U-net, which was described first by Ronneberger et. al. [14,54]. The architecture consists of a symmetrical structure. The information is first processed by a decoding layer. The dimension in the data is successively reduced by convolution and pooling. At the same time, however, the number of filters in each convolution layer is increased. This is the contractive path for the localisation of the patterns [14]. It is followed by a layer to encode what has been learned. Each layer contains the same number of filters as in decoding, but in reverse order. By upsampling, the dimension of the data is brought back to the original one. The special feature of this network, which is also the decisive difference to conventional autoencoders, is the concatenation between decoding and encoding layers. Each decoding layer is connected to an encoding layer with the same dimension. The concatenation is done via a path with the simplest possible activation, e.g. the identity.

It could be shown that this type of neural networks in segmentation exceeds the previous standards by far [14]. Furthermore, these networks have a particularly favourable convergence behaviour [14]. For this reason, this architectural form is adapted for our reconstruction.

11.2   Measurements for Reconstruction

Several measurements are used in the experiment to interpret the result correctly. This study is the first of its kind, which is why, with regard to other problems, we try to give the state of art.

For the following distance functions it may be favourable to denote three indices, having a tensor of density values, as our data is stored in a data structure of shape $(9, 9, 9$. For this purpose the first dimension is called i, the second j and the third k, so that

$$\forall i, j, k \text{ with } 1 \leq i, j, k \leq I, J, K \in \mathbb{N} \text{ and } \mathbf{X}, \mathbf{Y} \in \mathbb{R}^{I \times J \times K}. \qquad (49)$$

Mean absolute and squared error The mean squared error is defined as the expected value of the squared error of the input and target,

$$\text{MSE}(\mathbf{X}, \mathbf{Y}) = \frac{1}{N} \sum_{i,j,k} (\mathbf{Y}_{ijk} - \mathbf{X}_{ijk})^2. \qquad (50)$$

This type of error measurement weights errors $|e| > 1$ squared higher and errors $|e| < 1$ squared lower. This measure is a common for regression problems, but not an ideal metric for measuring the quality of the reconstruction because it is not restricted to a closed interval. Due to a lack of research in this area, there are also no benchmarks that can be referred to. Nevertheless,

the mean squared error is measured to derive properties of the data from their relationship to the other two utilized measurements.

The MSE is compared with its absolute counterpart. For normalised data in the interval $[0, 1]$ the error is expected to be significantly smaller than the mean absolute error. For data with large variance it would be correspondingly larger. The MAE is defined as the amount of difference between input and output:

$$\text{MAE}(\mathbf{X}, \mathbf{Y}) = \frac{1}{N} \sum_{i,j,k} |\mathbf{Y}_{ijk} - \mathbf{X}_{ijk}|_1. \tag{51}$$

The MAE is obviously larger than the MSE for data normalized to $[0, 1]$. However, even this metric does not offer the possibility of a standardised comparison, because MAE is not restricted to any closed interval either. Instead we use this error to validate the relative behaviour of MSE to MAE and thus we can easily validate the implementation of our architecture.

Intersection over union The Jaccard coefficient or intersection over union is used to measure the learned overlap of two sets. It is for two sets X, Y:

$$\text{Jaccard}(X, Y) = \frac{|X \cap Y|}{|X \cup Y|}. \tag{52}$$

The following implementation of an intersection over union based measurement is used for the neural network as a loss function:

$$\text{Jaccardish}(\mathbf{X}, \mathbf{Y}) = \frac{\sum_{i,j,k} \min(\mathbf{X}_{ijk}, \mathbf{Y}_{ijk})}{\sum_{i,j,k} \max(\mathbf{X}_{ijk}, \mathbf{Y}_{ijk})}. \tag{53}$$

The smallest and largest value is taken in pairs for each pixel between the prediction and target matrix. Then the smaller value is divided by the larger one. With this implementation we get an error normed to the interval $[0, 1]$. For the calculation of the error the spatial positions in the matrix are also taken into account. Furthermore, this measure is widely used for segmentation methods and can thus provide information about the quality of the produced prediction in comparison to other applications [10,65].

In the evaluation, physical interpretability plays a central role as a second factor. About 60% of the total deposited energy is located in the centre where the simulated isotope was deposited. For this reason it is useful to calculate a loss function which weights the error with the proportion of deposited energy. By weighting the minimum of both matrices in pairs with the inverse of the maximum, this effect can be achieved $(\min(x, y) \cdot \max(x, y)^{-1})$. The intersection over union thus provides a distance function which indicates the error proportional to the deposited energy.

Normalisation Before processing, the data is normalised to a certain interval. Since a sigmoid function is applied in the last layer of the neural network, which also returns values in the interval $(0, 1)$, a min-max normalisation is particularly suitable. However, to prevent the gradient from saturating, the values are normalised to a slightly smaller interval $[0.1, 0.9]$. Thus, the activation it not reached close to the limits zero and one. These areas have almost no slope or curvature, which in return means saturation of the gradient. The normalisation for a value of a matrix $\mathbf{X}$ into any interval $[a, b]$ is given by:

$$\text{Norm}(\mathbf{X}_{ijk}) = (b - a) \cdot \frac{\mathbf{X}_{ijk} - \min(\mathbf{X})}{\max(\mathbf{X}) - \min(\mathbf{X})} + a. \tag{54}$$

11.3   Covariate Shift

There are many challenges in training convolutional neural networks as a large number of parameters must be optimised. Therefore the filters can take on very different distributions during gradient descent from layer to layer. This makes the correct initialisation of weights and parameters extremely difficult and has also consequences for the chosen learning step size.

Ioffe and Szegedy define the internal shift of the covariance as the change in the distribution of the activities of the network as the parameters of the network are adapted in the training process [30]. LeCun, Wiesler and Ney published early on their findings that neural networks with centered data, linearly transformed to zero mean and variance one, and decorrelated features can achieve better results [37,67].

In this section we would like to discuss the normalisation of batches and explain how this can prevent covariate shift. The batch normalisation can be used for each affine transformation followed by a nonlinearity of the form $\mathbf{y} = h(\mathbf{W}\mathbf{x} + \mathbf{b})$.

Unlike Ioffe and Szegedy, the batch normalisation in our experiment is performed after the activation function. As batch normalisation centres the data from a convolution layer according to the learned patterns, we have due to the usage of `LeakyReLU` activation functions that images of the function with negative sign can no longer be considered. Batch normalisation after activation normalises the positive data without statistically distorting the negative data. The idea here is that these features will be scaled accordingly in the next convolution anyway. This results in the following normalisation equation for some outputs:

$$\mathbf{y} = \text{BN}(h(\mathbf{W}\mathbf{x})) \tag{55}$$

For convolution layers, it is also important to consider that features from the same filter, on different spatial dimensions of the input vector, receive the same normalisation. To ensure this, a mini-batch normalisation is performed for all spatial dimensions [30].

Algorithmus 4 Batchnorm for convolutional neural networks [30].

1: **Proc.:** BATCH NORMALISATION
2: **Param.:** $\mathcal{B} = \{\mathbf{x}_1, ..., \mathbf{x}_m\}$: values for $\mathbf{x}$
3: **Param.:** $\boldsymbol{\Theta} = \{\gamma, \beta\}$: parameters
4:    $\bar{\mathbf{x}}_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^{m} \mathbf{x}_i$ (*mean value for a batch*)
5:    $\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^{m} (\mathbf{x}_i - \bar{\mathbf{x}}_{\mathcal{B}})^2$ (*variance for a batch*)
6:    $\hat{\mathbf{x}}_i \leftarrow \frac{\mathbf{x}_i - \bar{\mathbf{x}}_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}}$ (*normalisation*)
7:    $\mathbf{y}_i \leftarrow \gamma \hat{\mathbf{x}}_i + \beta \equiv \text{BN}(\mathbf{x}_i, \boldsymbol{\Theta})$ (*scale and shift*)
8: **return** $\mathbf{y}_i$

The parameters are learned iteratively with the initial model. During training the backpropagation of the loss function $\mathcal{L}$ must be performed according to the normalised batches. Thus, the chain rule is [30]:

$$\frac{\partial \mathcal{L}}{\hat{\mathbf{x}}_i} = \gamma \frac{\partial \mathcal{L}}{\partial \mathbf{y}_i}, \quad \frac{\partial \mathcal{L}}{\partial \sigma_{\mathcal{B}}^2} = \sum_{i=1}^{m} \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{x}}_i} (\mathbf{x}_i - \bar{\mathbf{x}}_{\mathcal{B}}) - \frac{1}{2} \left(\sigma_{\mathcal{B}}^2 + \epsilon \right)^{-\frac{3}{2}}, \tag{56}$$

$$\frac{\partial \mathcal{L}}{\partial \bar{\mathbf{x}}_{\mathcal{B}}} = \sum_{i=1}^{m} \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{x}}_i} \frac{-1}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}}, \tag{57}$$

$$\frac{\partial \mathcal{L}}{\partial \mathbf{x}_i} = \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{x}}_i} \frac{1}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} + \frac{\partial \mathcal{L}}{\partial \sigma_{\mathcal{B}}^2} \frac{2(\mathbf{x}_i - \bar{\mathbf{x}}_{\mathcal{B}})}{m} + \frac{1}{m} \frac{\partial \mathcal{L}}{\partial \bar{\mathbf{x}}_{\mathcal{B}}}, \tag{58}$$

$$\frac{\partial \mathcal{L}}{\partial \gamma} = \sum_{i=1}^{m} \frac{\partial \mathcal{L}}{\partial \mathbf{y}_i} \hat{\mathbf{x}}_i, \quad \frac{\partial \mathcal{L}}{\partial \beta} = \sum_{i=1}^{m} \frac{\partial \mathcal{L}}{\partial \mathbf{y}_i}. \tag{59}$$

11.4   Residual Neural Networks

Residual networks use neurons to learn an additional residual function $\mathcal{H}$ depending on the information input $h(\mathbf{x}_n)$ in a layer n [23]. From this the output from a residual unit can be formalised as follows:

$$\mathbf{y}_n = h(\mathbf{x}_n) + \mathcal{H}(\mathbf{x}_{n-1} \mid \mathbf{W}_{n-1}), \tag{60}$$

$$\mathbf{x}_{n+1} = h(\mathbf{y}_n), \tag{61}$$

where $\mathbf{x}_n$ refers to the input in the n-th layer and $\mathcal{H}$ is a residual function and $h(\mathbf{x}_n) = \mathbb{I}(\mathbf{x}_n)$ is the identity. In this way the entire content of a bridged part of the neural network can be represented by the sum of the propagated unweighted component and a residual function.

This technique is realised by a connection that bridges a part of the network and propagates the information from one part of the network to a deeper one [23]. When U-networks are used, the content is bridged by layers that provide the same dimension when processing the data as output. There are many names for this modification in the literature, most recently deep learning systems were called residual networks, which offer a linear bridging or even the identity between each layer [23].

Thus, theoretically the content can be propagated linearly backwards through the whole network.

The connections themselves are called residual connections and should be called bridging. Recently, various activation functions for propagation by bridging have been tested [28,62,61], with the best results achieved with identity mapping or linear activation. Other activation functions showed higher error rates in the empirical experiment.

Thus, if h is an identity map this gives:

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \mathcal{H}(\mathbf{x}_{n-1} \mid \mathbf{W}_{n-1}). \tag{62}$$

This applies to the layer with the bridging n and the layers in between. The characteristics of deeper activities $\mathbf{x}_{n+1}$ of some layer $n \in \{1, 2, \ldots, N-1\}$ can be represented by the input of layer $\mathbf{x}_n$ and a residual function $\sum_{i=1}^{n-1} \mathcal{H}(\mathbf{x}_i \mid \mathbf{W}_i)$ between any two neurons n and $n+1$. The features $\mathbf{x}_n + \sum_{i=1}^{n-1} \mathcal{H}(\mathbf{x}_i \mid \mathbf{W}_i)$ of each lower unit $n \in \{n, \ldots, N\}$ is the sum of the initial information of all previous residual functions and $\mathbf{x}_0$. This is in contrast to conventional networks in which a $\mathbf{x}_N$ is the product of the weighted inputs $\mathbf{Wx}$ (neglecting normalisation and regularisation) [23].

These connections also have a favourable influence on the backward propagated error, by the modification of the chain rule [23]:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{x}_n} = \frac{\partial \mathcal{L}}{\partial \mathbf{x}_N} \frac{\partial \mathbf{x}_N}{\partial \mathbf{x}_n} = \frac{\partial \mathcal{L}}{\partial \mathbf{x}_N} \left(1 + \frac{\partial \sum_{i=1}^{N-1} \mathcal{H}(\mathbf{x}_i, \mathbf{W}_i)}{\partial \mathbf{x}_i} \right). \tag{63}$$

The gradient $\nabla_{\mathbf{x}_n} \mathcal{L}(\mathbf{x}_n, \mathbf{W}_n)$ can be represented by two terms: $\partial \mathcal{L}/\partial \mathbf{x}_N$ which propagates the information directly without any weighting, and a second term $1 + \partial \sum_i^{N-1} \mathcal{H}(\mathbf{x}_i, \mathbf{W}_i)/\partial \mathbf{x}_n$ which is propagated by the weighting. It follows that the signal can be propagated directly into the bridged layer.

In this experiment the technique is embedded after the last convolution layer and before dropout and pooling. Thus the gradient can access the learned information of the first part even after further convolution operations. In U-nets this method already proved to be very effective [14,54].

11.5   Dropout

Neural networks with many layers are increasingly difficult to train. Using dropout is a simple way to prevent the phenomenon of overfitting, where the training data is memorised and the ability to generalise is lost.

Dropout thins out the network by using a random variable from a Bernoulli distribution to determine which neurons within a layer remain active and are thus propagated. A neural network with n neurons can be understood as a collection of 2^n thinned out subnets. These nets are further designed in such a way that all weights from the previous layer are divided so that the number of parameters is $\leq O(n^2)$.

$$\mathbf{w}_{n+1} \longrightarrow \text{ReLU} \longrightarrow \text{BN} \longrightarrow \mathbf{w}_{n+2} \longrightarrow \text{ReLU} \longrightarrow \text{BN}$$

$$\mathbf{x}_n \longrightarrow \text{Concat} \longrightarrow \mathbf{y}$$

Fig. 7: Illustration of a residual connection of two layers in a neural network with two hidden layers, a `ReLU` activation function and batch normalisation. The operation to unify both paths is the concatenation of the matrices. `BN` refers to batch normalisation, `ReLU` is the rectifying linear unit activation function and `Concat` refers to the concatenation of matrices. Of course, concatenation creates a matrix of which the size is the sum of the sizes of its components.

Given is a neural network with N hidden layers and $n \in \{1, \ldots, N\}$ is the index of the hidden layers. Furthermore I is the number of neurons in one of the deeper layers. In analogy, $i \in \{1, \ldots, I\}$ is the index of the neuron within one layer. Furthermore $\mathbf{y}_n$ denotes the information output $\mathbf{y}$ from layer n, where for $\mathbf{y}_0 = \mathbf{x}$ the unweighted pattern $\mathbf{x}$ is the input. $\mathbf{b}_n$ stands for a bias term and denotes the statistical distortion in the n-th layer. $\mathbf{w}_n$ is the weight vector $\mathbf{w}$ for the n-th layer. Furthermore $h(\mathbf{x})$ is a non-linear activation function and $(\,\cdot\,)$ is the element-wise multiplication of vectors. So the feedforward algorithm becomes:

$$\mathbf{z}_n = \mathbf{w}_n \cdot \mathbf{y}_{n-1} + \mathbf{b}_n, \tag{64}$$

$$\mathbf{y}_n = h(\mathbf{z}_n). \tag{65}$$

With dropout, the operation changes as follows:

$$\mathbf{r}_n \sim \text{Bernoulli}(p), \tag{66}$$

$$\tilde{\mathbf{y}}_n = \mathbf{r}_n \cdot \mathbf{y}_n, \tag{67}$$

$$\mathbf{z}_n = \mathbf{w}_n \cdot \tilde{\mathbf{y}}_n + \mathbf{b}_n, \tag{68}$$

$$\mathbf{y}_n = h(\mathbf{z}_n). \tag{69}$$

The motivation driving this process comes from the theory of evolution. In sexual reproduction, half of the genes of two parents are recombined with a low probability of mutation [60]. This forms the genome of the offspring.

An alternative in biology is asexual replication from one parent, where a copy of the genetic material is produced, which also mutates with a small amount of change. It seems plausible that asexual reproduction is the better alternative because it ensures that a functioning gene is passed on without detours. Furthermore, sexual reproduction can break up functioning genes through recombination [60]. Nevertheless, sexual reproduction is the most common form of reproduction in more complex organisms.

One possible explanation for the superiority of sexual reproduction is the advantage of the ability of genes to mix. Looking at an evolving population

Fig. 8: Architecture of the U-network for the estimation of dose voxel kernels. All convolution layers except the last one are provided with the `LeakyReLU` activation function and a layer for batch normalisation. The lateral coupling from the last convolution layer of the first block is a concatenation of the matrices without activation.

over a longer period of time, the fitness of an individual could be secondary. It would be advantageous for the population's survivability if a gene were to be matched with various other genes. The ability of a gene to function with a random other gene makes it very robust. According to this theory, it would be important not only to spread an advantageous gene as widely as possible within the population, but to ensure its robustness.

A similar idea is pursued by dropout. Each neuron should learn to generate the required traits with a randomly chosen amount of other units from the previous layer. This makes the unit more robust and ensures that errors are more likely to be corrected by other neurons.

11.6 Learning of Density-specific Dose Voxel Kernels by means of U-Residual Networks

The model of the U-network with lateral couplings between the layers implements a kind of associative memory, which is directly related to the extracted information of a convolution layer in the network.

This is realized as a hierarchical plateau in which the dimension of the processed data is not drastically changed or remains the same, i.e. no downsampling is performed (see Fig. 8).

The implementation is done with `Keras v2.1.2` [13] and `TensorFlow v1.5` [1]. `NADAM` (see Sect. 10.2) was used as optimisation algorithm.

Neural Network Architecture In total, the U-net has 45 layers. The filters in the convolution layers were initialised with LeCun uniform distribution [37]. Each convolution layer is associated with a `LeakyReLU` activation function, as well as a layer for normalisation with respect to each batch. The constant for multiplication by negative values has been set to 5.5 for the `LeakyReLU` function. Each convolution layer has been regularised. The first layer received a regularisation term of the L_1 norm with a factor $\gamma = 0.005$ and each subsequent layer received a regularisation term of the L_2 norm with a factor $\gamma = 0.001$.

Data and Tissue Classes The data set consists of equal parts of mass densities or absorbed radiation dose from lungs, kidneys, liver, bones and spleen. For each tissue class, there are $2,000$ samples of mass densities and absorbed radiation dose, making a total of $10,000$ samples. Since the different tissue types differ greatly in their structure, the data set is shuffled first. Then the batches are drawn from them randomly.

There are two data sets in total. A training data set with 7.000 samples and a validation data set with 3.000 samples, which is completely independent from the neural network training process. The batch size for the training process is fixed at 128.

Experimental Results Overall, the neural network shows a good generalization ability with intersection over union of 0.86 after 308 epochs. Learning was done with a learning step size of 10^{-4}. The learning step size was adjusted once and halved after 15 epochs had passed without a change of more than $\epsilon \geq 10^{-6}$.
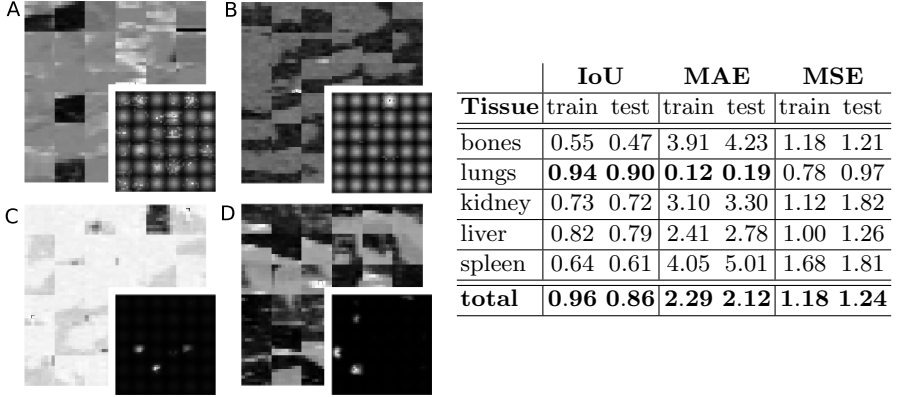
The normalised mass density was used as information input for the neural network and the normalised dose-voxel kernels were used as target data.

Tab. 9 shows the results for the different tissue classes. The best results were obtained in the lung tissue.

Fig. 10 shows different mass density kernels and dose-voxel kernels for the different tissue classes.

It was possible to construct a model which estimates the dose-voxel kernels with an accuracy sufficient to the current standard [10,65]. By using independent data it can be summarised that the presented architecture is capable of learning complex transfer functions as used in Monte Carlo simulation.

From the preparation of the data for the training run, knowledge was gained about how the variance in the data is distributed. For dimensional reduction the principal components have been tested. However, it was found that the variance in the data is equally distributed in all but one dimension and more than seventy percent of the absolute value of the eigenvalues were found within this dimension. This is not unusual for artificially generated data. For this reason, no dimension reduction was carried out for training.

	IoU		MAE		MSE	
Tissue	train	test	train	test	train	test
bones	0.55	0.47	3.91	4.23	1.18	1.21
lungs	**0.94**	**0.90**	**0.12**	**0.19**	0.78	0.97
kidney	0.73	0.72	3.10	3.30	1.12	1.82
liver	0.82	0.79	2.41	2.78	1.00	1.26
spleen	0.64	0.61	4.05	5.01	1.68	1.81
total	**0.96**	**0.86**	**2.29**	**2.12**	**1.18**	**1.24**

Fig. 9: *Left:* Cross section of mass densities and absorbed radiation dose along the i-th axis. In the large-backed picture, the mass densities of a specific tissue class are shown as a section of a mass kernel. The smaller picture shows the corresponding absorbed radiation dose. The picture displays squares that belong together. Each square is an independent section of the tissue that was randomly drawn from the data set. Example A is the density and absorbed radiation dose of the bone, B that of the kidney, C the liver, D the lung and E the spleen. For each picture 36 cross section pairs are shown. The darker the area, the denser the tissue. For absorbed radiation, the light area is the highest dose of radiation and the dark area is the lowest. *Right:* Validation of the neural network model after 308 epochs of training.

12   Discussion

It has been empirically shown that deep learning architectures are suitable to learn functions for the application in dosimetry. The developed method is able to learn weights for a linear combination which approximates a function of the Monte Carlo simulation and can therefore also process new, from the training data independent samples. By means of the selected error measurements it could be checked whether the pattern was actually learned.

Nevertheless, the measurement that reflects the error should be adapted to the physical properties of tissue. We provided an ad hoc example of such a metric, discussed and designed with the help of domain experts, but a medical evaluation is pending. The intersection over union gives good insight into the error distribution for individual samples. The positions within a DPK, however, are equally weighted, so that an error in the outer area of the kernel is just as decisive for the total error as an error in the central area. In order to verify whether the algorithm is suitable for clinical use, complete simulations must be generated by several people in further investigations and used for the training process. Then the neural network can be compared with the Monte Carlo simulation and the current standards for calculating dose distributions to draw a conclusion.

Fig. 10: Estimation results of the neural network. Cross section A is from the liver, B from the kidney tissue and C shows bone tissue. The sections of the density nuclei were taken along the first spatial direction of the data array on the fifth of nine cross sections. The sections of the density kernels were taken along the first spatial direction of the data array on the fifth of nine cross sections. The reason for this is that most of the activity takes place in the centre of the density kernel, where the radioactive nuclide was simulated. It can be seen particularly well that the scaling of the prediction and the outputs correspond. In A and C a good reconstruction of the distribution of the radiation energy on a very small scale was also successful..

For the patient, the exact calculation of the dose distribution in the centre of a DVK plays the most important role. At the place where the isotropic source was deposited, the radiation energy is also highest. For this reason, the error is particularly serious in this area.

Ideally, the weighting should be such that the centre has a particularly strong influence on the defect and is weighted less and less towards the periphery. A proposal for this implementation would be the weighting of the predicted kernel of the absorbed radiation dose with the calculated absorbed dose from the Monte Carlo simulation. Formalised this results in the following loss function, where · again denotes the element-wise multiplication:

$$\mathcal{L}(\mathbf{X}, \mathbf{Y}) = \sum_{i,j,k} \left(\mathbf{Y}_{ijk} - \mathbf{X}_{ijk} \right)^2 \cdot \frac{\mathbf{Y}_{ijk}}{\sum_{i,j,k} \mathbf{Y}_{ijk}} \tag{70}$$

Where $\mathbf{Y}$ denotes the matrix with the actual dose distribution (calculated with a Monte Carlo simulation) and $\mathbf{X}$ the prediction of the neural network. Eq. (70) is the mean squared error weighted by the relative proportion of radiation dose to the total DPK. In Fig. 11 this measurement is calculated for slightly more than 150 epochs. In fact, the error is much lower within the isotropic radiation source than at the boundary of the dose voxel kernels.

In further investigations, comparisons will be made with methods that are currently in use.

Fig. 11: *Left:* Evaluation of the training runs with mean squared error, mean absolute error and intersection over union. The U-net was trained with a training to validation ratio of $7 : 3$. A learning step size of 10^{-4} was chosen for initialisation. After 262 epochs a black vertical line in the plot marks the time when the learning step size was halved $(0.5 \cdot 10^{-4})$. As loss function the intersection over union was used. In total we trained for 308 epochs. A stop criterion was set for the minimum improvement within 15 epochs to be $\epsilon \geq 10^{-6}$. *Right:* Evaluation of the clinically motivated loss. The thickly highlighted lines of the measurements weighted with the inverse distribution correlate strongly with the original measurements. We continued training until epoch 453.

Code The code for the experiments is published on GitHub: https://github .com/karhunenloeve/MADVK.

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), https://www.tensorflow.org/, software available from tensorflow.org
2. Ackley, D., Hinton, G., Sejnowski, T.: A learning algorithm for boltzmann machines. Cognitive Science **9**(1), 147–169 (1985)

3. Bodei, L., Cremonesi, M., Ferrari, M., Pacifici, M., Grana, C., Bartolomei, M., Baio, S., Sansovini, M., Paganelli, G.: Long-term evaluation of renal toxicity after peptide receptor radionuclide therapy with 90 y-dotatoc and 177 lu-dotatate: the role of associated risk factors. European Journal of Nuclear Medicine and Molecular Imaging **35**(10), 1847–1856 (2008)

4. Bolch, W., Bouchet, L., Robertson, J., et al.: The dosimetry of nonuniform activity distributions: radionuclide s values at the voxel level. Journal of Nuclear Medicine **40**, 11S–36S (1999)

5. Botta, F., Mairani, A., Battistoni, G., Cremonesi, M., Di Dia, A., Fasso, A., Ferrari, A., Ferrari, M., Paganelli, G., Pedroli, G., et al.: Calculation of electron and isotopes dose point kernels with fluka monte carlo code for dosimetry in nuclear medicine therapy. Medical Physics **38**(7), 3944–3954 (2011)

6. Bottou, L.: Large-scale machine learning with stochastic gradient descent. In: Proceedings of the International Conference on Computational Statistics, pp. 177–186. Springer (2010)

7. Boureau, Y.L., Le Roux, N., Bach, F., Ponce, J., LeCun, Y.: Ask the locals: multi-way local pooling for image recognition. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2651–2658 (2011)

8. Boureau, Y.L., Ponce, J., LeCun, Y.: A theoretical analysis of feature pooling in visual recognition. In: Proceedings of the International Conference on Machine Learning. pp. 111–118 (2010)

9. Bracewell, R.: The fast hartley transform. Proceedings of the IEEE **72**(8), 1010–1018 (1984)

10. Bylinskii, Z., Judd, T., Oliva, A., Torralba, A., Durand, F.: What do different evaluation metrics tell us about saliency models? IEEE Transactions on Pattern Analysis and Machine Intelligence **41**(3), 740–757 (2018)

11. Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.: Semantic image segmentation with deep convolutional nets and fully connected crfs. In: International Conference on Learning Representations (2015)

12. Chiavassa, S., Aubineau-Lanièce, I., Bitar, A., Lisbona, A., Barbet, J., Franck, D., Jourdain, J., Bardiès, M.: Validation of a personalized dosimetric evaluation tool (oedipe) for targeted radiotherapy based on the monte carlo mcnpx code. Physics in Medicine & Biology **51**(3), 601 (2006)

13. Chollet, F., et al.: Keras. https://keras.io (2015)

14. Çiçek, Ö., Abdulkadir, A., Lienkamp, S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International Conference on Medical Image Computing and Computer-assisted Intervention. pp. 424–432. Springer (2016)

15. Dai, J., He, K., Sun, J.: Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1635–1643 (2015)

16. Dong, C., Loy, C., He, K., Tang, X.: Learning a deep convolutional network for image super-resolution. In: European Conference on Computer Vision. pp. 184–199. Springer (2014)

17. Dozat, T.: Incorporating nesterov momentum into adam (2015)

18. Duchi, J., Hazan, E., Singer, Y.: Adaptive subgradient methods for online learning and stochastic optimization. Journal of Machine Learning Research **12**(7) (2011)

19. Eigen, D., Krishnan, D., Fergus, R.: Restoring an image taken through a window covered with dirt or rain. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 633–640 (2013)
20. Elman, J.: Distributed representations, simple recurrent networks, and grammatical structure. Machine learning **7**(2-3), 195–225 (1991)
21. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 580–587 (2014)
22. Grossberg, S.: Nonlinear neural networks: Principles, mechanisms, and architectures. Neural Networks **1**(1), 17–61 (1988)
23. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: European Conference on Computer Vision. pp. 630–645. Springer (2016)
24. Hinton, G., Deng, L., Yu, D., Dahl, G., Mohamed, A.r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T., et al.: Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. IEEE Signal Processing Magazine **29**(6), 82–97 (2012)
25. Hinton, G., Srivastava, N., Swersky, K.: Neural networks for machine learning lecture 6a overview of mini-batch gradient descent. Toronto University (2012)
26. Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J., Kolen, J., Kremer, S.: A field guide to dynamical recurrent neural networks. In: Gradient Flow in Recurrent Nets: The Difficulty of Learning Long-Term Dependencies, pp. 237–243. Wiley-IEEE Press (2001)
27. Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J., et al.: Gradient flow in recurrent nets: the difficulty of learning long-term dependencies (2001)
28. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Computation **9**(8), 1735–1780 (1997)
29. Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. Proceedings of the National Academy of Sciences **79**(8), 2554–2558 (1982)
30. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015)
31. Keren, G., Schuller, B.: Convolutional rnn: an enhanced model for extracting features from sequential data. In: 2016 International Joint Conference on Neural Networks. pp. 3412–3419. IEEE (2016)
32. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. arXiv:1412.6980 (2014)
33. Klambauer, G., Unterthiner, T., Mayr, A., Hochreiter, S.: Self-normalizing neural networks. In: Advances in Neural Information Processing Systems. pp. 971–980 (2017)
34. Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. Communications of the ACM **60**(6), 84–90 (2017)
35. Le, Q., et al.: A tutorial on deep learning part 2: Autoencoders, convolutional neural networks and recurrent neural networks. Google Brain pp. 1–20 (2015)
36. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. Nature **521**(7553), 436–444 (2015)
37. LeCun, Y., Bottou, L., Orr, G.B., Müller, K.R.: Efficient backprop. In: Neural networks: Tricks of the Trade, pp. 9–48. Springer (2012)

38. Lee, C., Gallagher, P., Tu, Z.: Generalizing pooling functions in convolutional neural networks: Mixed, gated, and tree. In: Artificial Intelligence and Statistics. pp. 464–472 (2016)
39. Lin, G., Shen, C., Van Den Hengel, A., Reid, I.: Efficient piecewise training of deep structured models for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3194–3203 (2016)
40. Lippmann, R.: An introduction to computing with neural nets. IEEE Assp Magazine **4**(2), 4–22 (1987)
41. Lisman, J.: A mechanism for the hebb and the anti-hebb processes underlying learning and memory. Proceedings of the National Academy of Sciences **86**(23), 9574–9578 (1989)
42. Liu, F., Shen, C., Lin, G.: Deep convolutional neural fields for depth estimation from a single image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5162–5170 (2015)
43. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3431–3440 (2015)
44. Lukoševičius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. Computer Science Review **3**(3), 127–149 (2009)
45. Masnadi-Shirazi, H., Vasconcelos, N.: On the design of loss functions for classification: theory, robustness to outliers, and savageboost. In: Advances in Neural Information Processing Systems. pp. 1049–1056 (2009)
46. Murata, N.: A statistical study of on-line learning. Online Learning and Neural Networks pp. 63–92 (1998)
47. Nesterov, Y.: A method of solving a convex programming problem with convergence rate $o(\frac{1}{k^2})$. In: Soviet Mathematics Doklady. vol. 27 (1969)
48. Orbach, J.: Principles of neurodynamics. perceptrons and the theory of brain mechanisms. Archives of General Psychiatry **7**(3), 218–219 (1962)
49. Petoussi-Henss, N., Bolch, W., Zankl, M., Sgouros, G., Wessels, B.: Patient-specific scaling of reference s-values for cross-organ radionuclide s-values: what is appropriate? Radiation Protection Dosimetry **127**(1-4), 192–196 (2007)
50. Polyak, B.: Some methods of speeding up the convergence of iteration methods. USSR Computational Mathematics and Mathematical Physics **4**(5), 1–17 (1964)
51. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 779–788 (2016)
52. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in Neural Information Processing Systems. pp. 91–99 (2015)
53. Reynaert, N., Palmans, H., Thierens, H., Jeraj, R.: Parameter dependence of the mcnp electron transport in determining dose distributions. Medical Physics **29**(10), 2446–2454 (2002)
54. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-assisted Intervention. pp. 234–241. Springer (2015)
55. Rosasco, L., Vito, E.D., Caponnetto, A., Piana, M., Verri, A.: Are loss functions all the same? Neural Computation **16**(5), 1063–1076 (2004)
56. Rumelhart, D., Hinton, G., Williams, R.: Learning representations by back-propagating errors. Nature **323**(6088), 533–536 (1986)

57. Sgouros, G., Frey, E., Wahl, R., He, B., Prideaux, A., Hobbs, R.: Three-dimensional imaging-based radiobiological dosimetry. In: Seminars in Nuclear Medicine. vol. 38, pp. 321–334. Elsevier (2008)
58. Shen, Y.: Loss functions for binary classification and class probability estimation. Ph.D. thesis, University of Pennsylvania (2005)
59. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv:1409.1556 (2014)
60. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. The Journal of Machine Learning Research **15**(1), 1929–1958 (2014)
61. Srivastava, R.K., Greff, K., Schmidhuber, J.: Training very deep networks. In: Advances in Neural Information Processing Systems. pp. 2377–2385 (2015)
62. Srivastava, R.K., Greff, K., Schmidhuber, J.: Highway networks. arXiv:1505.00387 (2015)
63. Sutskever, I., Martens, J., Dahl, G., Hinton, G.: On the importance of initialization and momentum in deep learning. In: International Conference on Machine Learning. pp. 1139–1147 (2013)
64. Waters, L., Hendricks, J., McKinney, G.: Monte carlo n-particle transport code system for multiparticle and high energy applications. Los Alamos National Laboratory (2002)
65. Xu, Y., Li, J., Chen, J., Shen, G., Gao, Y.: A novel approach for visual saliency detection and segmentation based on objectness and top-down attention. In: 2017 2nd International Conference on Image, Vision and Computing. pp. 361–365. IEEE (2017)
66. Zeiler, M., Fergus, R.: Stochastic pooling for regularization of deep convolutional neural networks. arXiv:1301.3557 (2013)
67. Zeyer, A., Kulikov, I., Schlüter, R., Ney, H.: Faster sequence training. In: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 5285–5289. IEEE (2017)