

Precision Learning: Towards Use of Known Operators in Neural Networks

Andreas Maier*, Frank Schebesch*, Christopher Syben*, Tobias Würfl*, Stefan Steidl*,
Jang-Hwan Choi⁺, Rebecca Fahrig*

* Pattern Recognition Lab

Friedrich-Alexander-University Erlangen-Nuremberg, Germany

+ Medical Imaging Systems Lab, Division of Mechanical and Biomedical Engineering,

College of Engineering, Ewha Womans University, Seoul 03760, Korea

Email: andreas.maier@fau.de

Abstract—In this paper, we consider the use of prior knowledge within neural networks. In particular, we investigate the effect of a known transform within the mapping from input data space to the output domain. We demonstrate that use of known transforms is able to change maximal error bounds.

In order to explore the effect further, we consider the problem of X-ray material decomposition as an example to incorporate additional prior knowledge. We demonstrate that inclusion of a non-linear function known from the physical properties of the system is able to reduce prediction errors therewith improving prediction quality from SSIM values of 0.54 to 0.88.

This approach is applicable to a wide set of applications in physics and signal processing that provide prior knowledge on such transforms. Also maximal error estimation and network understanding could be facilitated within the context of *precision learning*.

I. INTRODUCTION

Neural networks and deep learning recently produced numerous exciting results in pattern recognition and machine intelligence [1], [2], [3], [4]. Most of these approaches use hand-crafted deep network architectures and all weights are trained starting from a random initialisation. One effect of this procedure is that for some applications hundreds of millions of training samples are required to get optimal performance [5], [6]. Interestingly, emerging networks share similarities with known operators [7], [8]. These approaches seem adequate for general perceptive tasks as prior knowledge is difficult to obtain. Doing so gives rise to the current trend of “network engineering”. However, in many slightly more restricted application problems, a broad body of prior knowledge is available, e.g. physics, signal processing, and other expert knowledge.

Only recently a few investigations using known operators in the neural network modelling process have been published. Würfl et al. demonstrated that back-projection and filtering of CT reconstruction can be embedded into a neural network context allowing the learning of redundancy weights [9]. Based on these observations Syben et al. showed that neural networks are also able to learn discrete filter versions from their continuous counterparts [10]. By design the network is restricted to a general reconstruction filter and by clever synthesis of training data, 10 observations are sufficient for the learning process. Fu et al. developed this concept even further

and demonstrated that even more complex operations such as the eigenvalue computation in Frangi’s vesselness can be modelled with mathematical equivalence in a neural Network design [11]. As all of these approaches fix parts of the network due to prior information about the underlying algorithm, the number of unknown parameters is dramatically reduced. Still neural Network training is employed using typical optimisation approaches as provided by libraries such as TensorFlow. Additionally they employ specialised layers that are able to compute the gradient of the network efficiently. We consider such approaches that mix known functions with unknown operators as examples of *precision learning* which typically requires less training examples and fewer training iterations. Figure 1 shows a graph of such an approach conceptually.

In this paper, we wish to explore the general bounds on such approaches and make comparisons to the universal approximation theorem [12]. Based on these observations, we pick an example in which machine learning has been successfully applied previously: X-ray material decomposition [13], [14]. We briefly investigate its theory and explore the effect of using known operators on the quality of the material decomposition to verify our theoretical observations.

II. BOUNDS FOR SEQUENCES OF OPERATORS

The universal approximation theorem proves that any continuous function $u(\mathbf{x})$ on a compact subset of \mathbb{R}^N can be approximated by a linear combination of sigmoid functions $s(x)$

$$u(\mathbf{x}) \approx U(\mathbf{x}) = \sum_i u_i s(\mathbf{w}^\top \mathbf{x} + w_0) \quad (1)$$

where \mathbf{w} and w_0 are hyper plane parameters that project \mathbf{x} and u_i are the weights of the linear superposition of the sigmoid curves. The resulting approximation error ϵ_u is an upper bound following

$$|U(\mathbf{x}) - u(\mathbf{x})| \leq \epsilon_u. \quad (2)$$

Following this concept, we can expand to a multidimensional $U(\mathbf{x})$ and $\mathbf{u}(\mathbf{x})$ in which every component follows

$$|U_j(\mathbf{x}) - u_j(\mathbf{x})| \leq \epsilon_{u_j}. \quad (3)$$




Fig. 1. Schematic of the idea of precision learning: One or more known operators are embedded into a network. Doing so, allows dramatic reduction of the number of parameters that have to be estimated during the learning process. The minimal requirement for the operator is that it must allow the computation of a sub-gradient for use in combination with the back-propagation algorithm. This requirement can be met by a large class of operations.

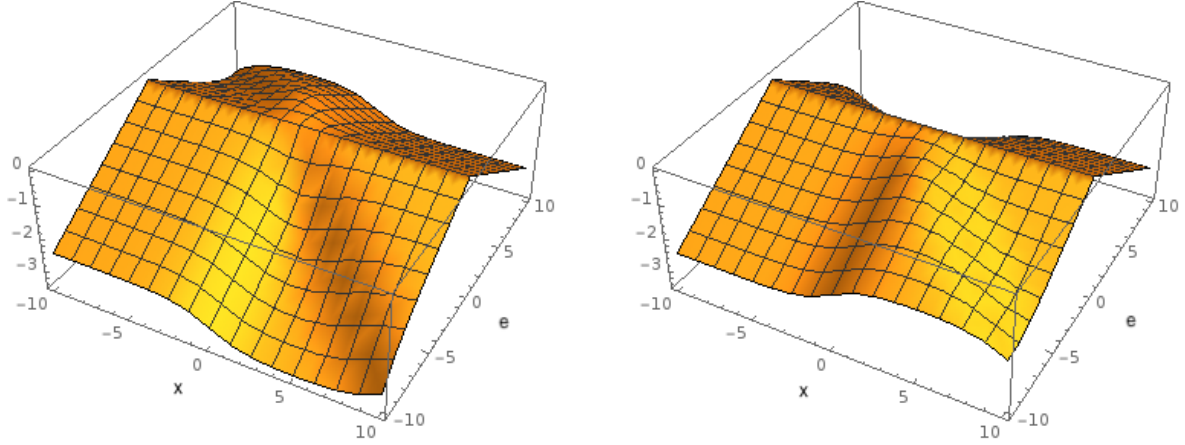


Fig. 2. Plots for the relation $g_j s(x+e) - g_j s(x) - |g_j| \cdot l_s \cdot |e| \leq 0$ (cf. Eq. 9) for $g_j > 0$ on the left hand side and for $g_j < 0$ on the right hand side, normalised by $|g_j|$. As can be seen in the plots, the inequality holds for both cases.

This enables us to explore the chain of operators $f(\mathbf{x}) = g(\mathbf{u}(\mathbf{x}))$. In the following, we will refer to the approximations of the original functions as capital letters, i.e., $G(\mathbf{x})$ is the approximation of $g(\mathbf{x})$. Furthermore, we link the plane projection $\mathbf{w}^\top \mathbf{x} + w_0$ only to the input layer which gives rise to a typical hidden layer network structure with alternating linear combinations and sigmoid activation functions. This leads to the following possible approximations

$$F_u(\mathbf{x}) = g(\mathbf{U}(\mathbf{x})) = f(\mathbf{x}) - e_u \quad (4)$$

$$F_g(\mathbf{x}) = G(\mathbf{u}(\mathbf{x})) = f(\mathbf{x}) - e_g \quad (5)$$

$$F(\mathbf{x}) = G(\mathbf{U}(\mathbf{x})) = f(\mathbf{x}) - e_f \quad (6)$$

where $F_u(\mathbf{x})$ uses the approximation of $\mathbf{u}(\mathbf{x})$, but knows $g(\mathbf{x})$, $F_g(\mathbf{x})$ knows $\mathbf{u}(\mathbf{x})$, but approximates $g(\mathbf{x})$, and $F(\mathbf{x})$ needs to approximate both. Respectively, e_u , e_g , and e_f are the errors that are introduced by these approximations. Each of these errors is bound by a positive corresponding ϵ , i.e., $|e_u| \leq \epsilon_u$. Of particular interest is the chain of operators

$$F(\mathbf{x}) = G(\mathbf{U}(\mathbf{x})) = \sum_j g_j s \left(\sum_i u_{i,j} s(\mathbf{w}^\top \mathbf{x} + w_0) \right) + g_0 \quad (7)$$

and bounds to this approximation. Based on our previous observations, we find

$$\begin{aligned} f(\mathbf{x}) &= g(\mathbf{u}(\mathbf{x})) = G(\mathbf{u}(\mathbf{x})) + e_g \\ &= \sum_j g_j s(u_j(\mathbf{x})) + g_0 + e_g \\ &= \sum_j g_j s(U_j(\mathbf{x}) + e_{u_j}) + g_0 + e_g. \end{aligned} \quad (8)$$

If we consider the case of overestimation, i.e. $f(\mathbf{x}) \leq F(\mathbf{x})$, we are interested in finding an upper bound for the approximation. As a broad range of functions, the sigmoid function is Lipschitz continuous, i.e. its slope is bounded by a positive value l . For the sigmoid function, this value is found at $s(0)$ with $l_s = 0.25$. Thus, a general upper bound for $s(x+e)$ is

$$s(x+e) \leq s(x) + l_s \cdot |e|.$$

When combined with a multiplicative constant g_j , this upper bound needs to be modified to

$$g_j s(x+e) \leq g_j s(x) + |g_j| \cdot l_s \cdot |e|. \quad (9)$$

This relation is visualised in Figure 2. Hence, we can find an upper bound of Eq. 8 as

$$\begin{aligned} f(\mathbf{x}) &\leq \underbrace{\sum_j g_j s(U_j(\mathbf{x})) + g_0 + \sum_j |g_j| \cdot l_s \cdot |e_{u_j}|}_{F(\mathbf{x})} + e_g \\ &\leq F(\mathbf{x}) + \sum_j |g_j| \cdot l_s \cdot |e_{u_j}| + e_g. \end{aligned} \quad (10)$$

Subtraction of $F(\mathbf{x})$ yields

$$\begin{aligned} \underbrace{f(\mathbf{x}) - F(\mathbf{x})}_{e_f} &\leq \sum_j |g_j| \cdot l_s \cdot |e_{u_j}| + e_g \\ e_f &\leq \sum_j |g_j| \cdot l_s \cdot |e_{u_j}| + e_g. \end{aligned} \quad (11)$$

Using $|e_g| \leq \epsilon_g$ leads to

$$e_f \leq \sum_j |g_j| \cdot l_s \cdot |e_{u_j}| + \epsilon_g. \quad (12)$$

For the case of underestimation with $f(\mathbf{x}) \geq F(\mathbf{x})$, a similar bound is found as

$$e_f \geq -\sum_j |g_j| \cdot l_s \cdot |e_{u_j}| - \epsilon_g \quad (13)$$

Thus, an absolute bound for e_f is found as

$$|e_f| \leq \sum_j |g_j| \cdot l_s \cdot |e_{u_j}| + \epsilon_g. \quad (14)$$

This observation is interesting as $\sum_j |g_j| \cdot l_s \cdot |e_{u_j}|$ is also an upper bounded for $|e_u|$, if $g(\mathbf{x})$ is a superposition of sigmoid functions. If we consider a more general $g(\mathbf{x})$, bound by the Lipschitz constant l_g , we obtain the following upper bound

$$g(\mathbf{x} + \mathbf{e}_u) \leq g(\mathbf{x}) + l_g \cdot \|\mathbf{e}_u\|_1 \quad (15)$$

where \mathbf{e}_u is a vector containing all e_{u_j} and $\|\cdot\|_1$ is the L₁ norm. Following the derivation above, it is easy to see, that

$$|e_u| \leq l_g \cdot \|\mathbf{e}_u\|_1. \quad (16)$$

Conceptually both bounds are similar as

$$l_g \cdot \|\mathbf{e}_u\|_1 = \sum_j l_g \cdot |e_{u_j}| \approx \sum_j |g_j| \cdot l_s \cdot |e_{u_j}|. \quad (17)$$

This means that $|e_f|$ is essentially bounded by ϵ_u and ϵ_g . Thus, the maximal estimation error achieved using known operators is lower or equal to the error achieved when training $F(\mathbf{x})$ from scratch without prior knowledge. Knowledge both on $g(\mathbf{x})$ and $\mathbf{u}(\mathbf{x})$ respectively reduces the maximal error in an additive fashion which leads us to the conclusion that any use of prior domain knowledge is useful to support the learning problem. This forms a theoretical foundation for the success of previous applications of *precision learning*.

III. X-RAY MATERIAL DECOMPOSITION

In this section, we want to explore an example for Eq. 14. In order to do so, we choose X-ray material decomposition as many properties of the transform are known. Yet analytical inversion requires accurate calibration of the X-ray spectrum $I_0(E)$ where E is the X-ray energy. For each energy, the line integral along the X-ray beam is observed

$$I(E) = I_0(E) \cdot e^{-\sum_i \mu(i,E)l_i} \quad (18)$$

where $\mu(i, E)$ is the X-ray absorption coefficient for energy E and material i and l_i is the path length through material i . Typically, X-ray detectors measure the integral over several energies

$$I_b(x, y) = \int_{b_l}^{b_u} I(x, y, E) dE \quad (19)$$

at each pixel (x, y) from the low energy limit b_l to the upper energy limit b_u of the respective energy bin b . Inversion can be performed iteratively, if the spectrum and the material coefficients are accurately known [15]. In practice, this process is difficult as generally all materials in the object under consideration have to be known and the spectral calibration is error prone. A common approach is to model the decomposition using a low-order polynomial [16] as the function is generally smooth. Lu et al. demonstrated that this approach is a general regression problem that predicts l_i given an approximation of the inverse function $F(\mathbf{I})$ where \mathbf{I} is the vector containing all measured energy bins [13]. Amongst several machine learning techniques, multi-layer perceptrons demonstrated the best results. Lu et al. already used a logarithmic transform on their data and expanded the coefficients to a polynomial space as they knew that both assumptions are supposed to be part of the regression problem.

IV. EXPERIMENTAL RESULTS

In order to explore our idea, we will introduce a $-\log$ transform given the calibrated I_0 in each energy bin for $\mathbf{u}(\mathbf{x})$ and a polynomial expansion for $g(\mathbf{x})$ following the ideas presented in [16]. The remaining estimation problem is left to a three-layer perceptron $F(\mathbf{x})$ which has a higher modelling ability than a single hidden layer network. We explore the effect of the transforms on the following four combinations: $F(\mathbf{I})$, $F(u(\mathbf{I}))$, $F(g(\mathbf{I}))$, and $F(g(u(\mathbf{I})))$. All implementations were performed in the CONRAD open source software package [17].

The experimental data for our experiment was collected at Stanford University using a Zeego C-arm device (Siemens Healthineers, Forchheim). Using prototype software, scans of a Sawbones pelvis phantom were acquired. For the three different energy settings, 41, 70, and 125kVP were selected. Image dimensions were 620×480 pixels. In order to get reference labels for the ground truth, we inserted a biopsy needle into the phantom, performed the three scans, and carefully removed the needle to cause minimal motion of the phantom. Subsequently, another series of three scans was performed. In order to generate the ground truth, the 125 kVp scans before




Fig. 3. Example images of the different data: In the left column, an X-ray image $I_b(x, y)$ of the phantom and the biopsy needle are shown. The center images show the data after transform $u(I_b(x, y))$, i.e. line integral domain. The right column shows the subtracted image after needle removal that was used as ground truth. Note that the needle grip is embedded in plastic which is of course a different material than metal.




Fig. 4. Even in narrow window and level $[0, 0.3]$, the material separation generally works well. Note that in particular, the artefacts in the background gradually disappear with increasing use of prior knowledge. In addition, note that the plastic grip in the bottom of the image generally disappears in the predicted data, as the grip does not consist of metal.

and after needle insertion were subtracted. Examples of the X-ray images are shown in Figure 3.

TABLE I
OVERVIEW ON THE RESULTS OF THE PREDICTION. PEARSON'S r IS GENERALLY HIGH, WHILE THE SSIM IS DRASTICALLY INCREASED WITH INCREASING PRIOR KNOWLEDGE.

	$F(\mathbf{I})$	$F(u(\mathbf{I}))$	$F(g(\mathbf{I}))$	$F(g(u(\mathbf{I})))$
Pearson's r [%]	95.0	95.2	95.1	95.5
SSIM [%]	54.1	63.1	73.8	88.4

In total, training and test contained 273.360 feature vectors. We evaluated the prediction quality using Pearson's r [18] and the Structural Similarity Index (SSIM) [19]. Table I gives an overview on the results. For both correlation as well as SSIM, the prediction quality increases with increasing use of prior knowledge.

Figure 4 shows the prediction result visually. The needle is visualised well in all four versions of the material decom-

position and the plastic grip is removed. No bone is shown in the image anymore. Still, there are some artefacts in the background of the image that decrease with increasing use of prior knowledge.

In Figure 5, we show the estimation vs. the ground truth of all four methods in comparison. With increasing use of prior knowledge, i.e., transforms $u(\mathbf{x})$ and $g(\mathbf{x})$, the prediction becomes more and more linear. Also note that the plastic in the reference introduces a bias in the estimation of the low path lengths through the needle.

V. DISCUSSION

In this paper, we show an upper bound for errors that are introduced by the approximation of functions using neural networks. In our derivation, we follow the universal approximation theorem using a single layer per function. Generally, the approximation introduces an error and a combination of multiple approximations increases the upper error bound. This approach could also be expanded to more layers. Using our approximation formula, this would only increase the maximal bound further. Hence, we consider one layer as demonstrated in the universal approximation theorem as sufficient. Furthermore, it is noteworthy to state that the universal approximation theorem always combines one hyper plane with one weighted superposition of sigmoid functions. Here, we only do that in the input layer, as we wanted our model to mimic traditional hidden layer approaches. Additional use of another hyper plane in the second function approximation would not change our results as this is already modelled in the linear combination of




Fig. 5. Display of the scatter plot ground truth vs. prediction of the four predictors. Note that the shape of the scattering becomes more and more linear with increasing use of prior knowledge.

$U(\mathbf{x})$. Yet, the notation would become less tangible. Therefore, we chose the current notation.

Our upper bounds share further similarities with the universal approximation theorem: Specification of the number of nodes per layer is not required and with increasing number of nodes, the maximal bounds for the error will shrink. Furthermore, we also do not show that our upper bound is tight. It is merely a worst-case analysis and does not describe the average or expected error. Further analysis could be performed based on the work of Barron [20].

Still, these observations are useful for understanding other applications of *precision learning* that already have been published [9], [10], [11]. With the bounds found in this paper, we are able to see that the maximal error margin decreases, independent of whether the transform is applied before or after the network, if that is the true sequence of operators. This confirms experimental findings that replace operators within the network to increase the prediction quality [11].

Using this framework enables us to mix operators with deep learning procedures and provides a new view and understanding of deep learning. The benefits are at least two-fold: First, the mixing allows us to reduce the number of parameters of the network and therewith a reduction in training data [10]. Second, it opens the view on neural networks as an ensemble of operators that we work with in terms of mathematical analysis. Technically, this also enables to “derive” new networks. Even operators that are typically more difficult to grasp in a mathematical sense can be used in this framework if they allow the computation of sub-gradients since the gradients can be processed in the back-propagation chain. This is also in line with our assumptions on $g(\mathbf{x})$ being

Lipschitz continuous. This even allows inclusion of sorting operators such as quantile operations [21]. As such this could potentially be a game changer for deep learning.

Furthermore, *precision learning* will also help to understand complicated network structures in more detail. As such, we can now try to interpret the U-Net [22] algorithm as a multi-scale dictionary learning approach with learned shrinkage operators in which the reconstruction dictionary is independent of the analyser dictionary. Similar observations were already done for CNNs in [23]. Hence, one could attempt to approximate its structure using traditional iterative shrinkage methods, wavelet and custom bases, and therewith reduce the number of parameters drastically.

The distinction of *precision learning* to general regularisation and deep learning is that we use precise domain knowledge expressed in terms of mathematical expressions to model our network instead of heuristically discovered network architectures. As such we interpret networks as a method to express a general objective function that describes an optimisation problem. Doing so, every block in the emerging objective function serves a particular purpose to which it is restricted to by design. Hence, the trained block can also be interpreted with respect to its use.

There are also links to classical pattern recognition theory [24], if we interpret $u(\mathbf{x})$ as the feature extractor and $g(\mathbf{x})$ as the classifier/regressor. In particular, the observation that the error caused by $F_u(\mathbf{x})$ has much higher bounds than the error produced by $F_g(\mathbf{x})$ might also be linked to the classical importance of feature extraction and feature engineering. Even the best classifier cannot recover information lost during feature extraction.

In our experiments, we demonstrate another application of *precision learning*: X-ray material decomposition. We demonstrate that the use of additional prior domain knowledge is able to reduce the prediction error. As such, we only explore feature domain transformation in this paper. Given the realistic setting of our experiment, we deemed this to be more convincing. In every case, the material decomposition that is learned is physically correct and the flawed ground truth is mapped to a physically plausible result. With a constrained simulation experiment, we would not have been able to observe this finding.

VI. CONCLUSION & OUTLOOK

We showed that *precision learning* is able to reduce maximal error bounds, if known operators are introduced into the learning process. This forms a theoretical basis for prior observations using the same approach. Our derivation shows a relationship between the number of approximation steps in their upper bounds. If replaced with the true transform this error is removed from the learning process.

In addition, we explored the effect also in real data for an X-ray material decomposition and confirmed prior work by Lu et al. The more prior knowledge that is introduced, the better the prediction gets. In our case improvements from 54.1% to 88.4% in SSIM could be obtained.

We believe that this approach will be applicable to a large set of problems in particular in physics and signal processing where a lot of prior constraints on the domain are known. Also medical or other high risk applications could also benefit from this approach, as known operators in conjunction with few network layers are easier to analyse in terms of maximal error. Lastly, this approach also allows to design networks that follow a particular algorithm that enforces desired criteria, e.g. safety bounds on a separate path through the network.

REFERENCES

- [1] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016.
- [2] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 779–788.
- [3] F. C. Ghesu, E. Krubasik, B. Georgescu, V. Singh, Y. Zheng, J. Hornegger, and D. Comaniciu, "Marginal space deep learning: efficient architecture for volumetric image parsing," *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1217–1228, 2016.
- [4] M. Aubreville, M. Krappmann, C. Bertram, R. Klopffleisch, and A. Maier, "A Guided Spatial Transformer Network for Histology Cell Differentiation," in *Eurographics Workshop on Visual Computing for Biology and Medicine*, T. E. Association, Ed., 2017, pp. 021–025. [Online]. Available: <https://www5.informatik.uni-erlangen.de/Forschung/Publikationen/2017/Aubreville17-AGS.pdf>
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 2009, pp. 248–255.
- [6] C. Sun, A. Shrivastava, S. Singh, and A. Gupta, "Revisiting unreasonable effectiveness of data in deep learning era," *arXiv preprint arXiv:1707.02968*, vol. 1, 2017.
- [7] V. Sindhwani, T. Sainath, and S. Kumar, "Structured transforms for small-footprint deep learning," in *Advances in Neural Information Processing Systems*, 2015, pp. 3088–3096.
- [8] E. J. Humphrey, J. P. Bello, and Y. LeCun, "Moving beyond feature design: Deep architectures and automatic feature learning in music informatics," in *ISMIR*, 2012, pp. 403–408.
- [9] T. Würfl, F. C. Ghesu, V. Christlein, and A. Maier, "Deep learning computed tomography," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2016, pp. 432–440.
- [10] C. Syben, B. Stimpel, K. Breininger, T. Würfl, R. Fahrig, A. Dörfler, and A. Maier, "A deep learning approach for reconstruction filter kernel discretization," *arXiv preprint arXiv:1710.06287*, 2017.
- [11] W. Fu, K. Breininger, T. Würfl, N. Ravikumar, R. Schaffert, and A. Maier, "Frangi-Net: A Neural Network Approach to Vessel Segmentation," *arXiv preprint arXiv:1711.03345*, nov 2017. [Online]. Available: <http://arxiv.org/abs/1711.03345>
- [12] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Mathematics of Control, Signals, and Systems (MCSS)*, vol. 2, no. 4, pp. 303–314, 1989.
- [13] Y. Lu, J. Geret, M. Unberath, M. Manhart, Q. Ren, R. Fahrig, J. Hornegger, and A. Maier, "Projection-based Material Decomposition by Machine Learning using Image-based Features for Computed Tomography," in *The 13th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine*, F. 2015, Ed., 2015, pp. 448–451. [Online]. Available: <https://www5.informatik.uni-erlangen.de/Forschung/Publikationen/2015/Lu15-PMD.pdf>
- [14] K. C. Zimmerman and T. G. Schmidt, "Experimental comparison of empirical material decomposition methods for spectral CT," *Physics in medicine and biology*, vol. 60, no. 8, p. 3175, 2015.
- [15] M. Manhart, A. Maier, J. Hornegger, and A. Dörfler, "Material Decomposition for Energy Resolving Detectors using Weighted Levenberg-Marquardt Optimization," in *Proceedings of the 1st Conference on Image Guided Interventions*, B. Preim, G. Rose, M. Skalej, and F. Wacker, Eds., 2014, pp. 23–24. [Online]. Available: <https://www5.informatik.uni-erlangen.de/Forschung/Publikationen/2014/Manhart14-MDF.pdf>
- [16] N. Maaß, S. Sawall, M. Knaup, and M. Kachelrieß, "Empirical multiple energy calibration (EMEC) for material-selective CT," in *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2011 IEEE*. IEEE, 2011, pp. 4222–4229.
- [17] A. Maier, H. G. Hofmann, M. Berger, P. Fischer, C. Schwemmer, H. Wu, K. Müller, J. Hornegger, J.-H. Choi, C. Riess et al., "Conrad — a software framework for cone-beam imaging in radiology," *Medical physics*, vol. 40, no. 11, 2013.
- [18] K. Pearson, "Correlation coefficient," in *Royal Society Proceedings*, vol. 58, 1895, p. 214.
- [19] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [20] A. R. Barron, "Approximation and estimation bounds for artificial neural networks," *Machine Learning*, vol. 14, no. 1, pp. 115–133, 1994.
- [21] F. Schirmacher, T. Köhler, L. Husvogt, J. G. Fujimoto, J. Hornegger, and A. K. Maier, "QuaSI: Quantile Sparse Image Prior for Spatio-Temporal Denoising of Retinal OCT Data," in *Medical Image Computing and Computer-Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings*, vol. 10434. Springer, 2017, p. 83.
- [22] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 234–241.
- [23] V. Pappayan, Y. Romano, and M. Elad, "Convolutional neural networks analyzed via convolutional sparse coding," *arXiv preprint arXiv:1607.08194*, 2016.
- [24] H. Niemann, *Pattern analysis and understanding*. Springer Science & Business Media, 2013, vol. 4.