Abstract

Nearfield acoustic holography is a powerful technique used in industry and research to retrieve vibro-acoustic source properties from non-contact measurements close to the source. A variety of reconstruction approaches based on different working principles has been proposed so far, all intended to model the measured sound field to then infer the underlying source distribution through inverse calculations. Due to the ill-posed and underdetermined nature of typical holographic systems, finding a meaningful solution poses many challenges and difficulties and almost always requires to apply some form of regularization to introduce additional constraints.

The traditional solution approach proposes a reformulation of the problem into an optimization task, with the goal of attaining an approximative solution that minimizes the mismatch between model and measurements while simultaneously minimizing some regularizer encoding prior knowledge or assumptions about the source. Choosing the right type and amount of regularization is most critical for reconstruction success. The solution accuracy rises and falls with quantity and quality of source knowledge, as well as the level of engineering skills and intuition on parameter tuning. In practical application of acoustic holography, there commonly is little to no prior knowledge about the source available and assumption accuracy cannot be globally guaranteed, which makes the reconstruction particularly prone to local misestimations and artifacts. Despite intensive research on new families of regularizers and optimal parameter choice methods in the last few decades, there are still many problems in this regard that current state-of-the-art methods cannot handle. What is more, optimization-based holography is usually not capable of real-time execution as demanded by industrial measurement and monitoring applications.

More recent advances in the field of image reconstruction suggest that deep(-learning-based) models might be the future way to effectively cope with these problems. Deep models utilize neural networks as computational units, whose parameters are trained on huge datasets, rather than being explicitly designed by hand to solve the inverse problem. Through observational learning procedures, deep models are able to relief users from the burden of crucial decision-making by automating critical parameter selections and regularization. In theory, trained deep models can approximate any mapping to arbitrary accuracy without the need of prior knowledge, just by looking at generic data samples, which makes them an extremely powerful tool for modern image reconstruction. Furthermore, by shifting most of the computational burden of optimization from the recontruction task to an offline training routine performed in advance, the actual retrieval may take place in real-time. Despite their remarkable performance in various holographic tasks, particularly in medical imaging, and their significant progress in most recent years, there is - to the best of my knowledge - no substantial research available thus far on the use of deep models to facilitate data-driven sound source retrieval. This PhD thesis is intended to close this research gap by proposing a generalized deep-learning-based framework for acoustic holography that is universally applicable, independent of the source's nature, and without the particular need of having any prior knowledge or beliefs about it.

The main objective of this study is to examine how to best leverage the great potential of data-driven modelling for acoustic holography to learn its most critical aspect of regularization from preexisting data. For this purpose, a representative training database of synthesized and experimental source-field pairs must be constructed in order to discover how to benefit best from different types and sizes of data subsets. This study shall also provide some guidance on how to find the right balance of embedding human and data-driven knowledge into network architecture and training in order to achieve optimal results. This requires a comprehensive analysis on the behaviour of deep neural networks in the presence of noise, artifacts and measurement imperfections, as well as an extensive study about new types of errors and challenges induced by the learning models itself. Eventually, all of this aims at making the reader aware of the possibilities and difficulties of data-driven source reconstruction and familiar with current deep learning best practices adapted to the specific needs of acoustic holography.

Exposé - Deep Learning In Acoustic Holography

Introduction

Nearfield acoustic holography is a powerful technique for the retrieval of vibro-acoustic source properties from contactless measurements in proximity to the source [Pag19]. It is widely used in many branches of industry and research for localization of noise sources or acoustic leaks and characterization of structure-borne sound sources. A variety of acoustic imaging techniques has been proposed so far, each of which based on different working principles and with its own advantages and drawbacks. They all can be used to analytically or numerically model a measured sound field radiated by arbitrary sources, to then infer the underlying source distribution and its parameters using inverse reconstruction approaches. One popular state-of-the-art model is the equivalent source method (ESM), which describes the sound field as a superposition of sound waves radiated by elementary sources of different strengths.

In this case, the sound field model takes the form of the linear equation system

$$\mathbf{p} = \mathbf{G}\mathbf{q} + \mathbf{n}, \tag{0.1}$$

where $\mathbf{p}$ contains the measured pressures, $\mathbf{G}$ is the propagation matrix containing Green's functions that describe the sound propagation between source and measurement points, $\mathbf{n}$ models the observation noise and $\mathbf{q}$ contains the unknown source strengths that we wish to evaluate. This is achieved by solving the equation system inversely for the source strengths $\mathbf{q}$.

In general, the solution of an inverse problem critically depends on the observation and model layout, i.e. the ratio of known and unknown parameters therein, as well as the spectral properties of the linear operator $\mathbf{G}$ and the inevitable corruption with noise. All these factors may cause the inverse problem to be ill-posed, meaning that small perturbations in the input data may lead to large errors in a possibly non-unique solution.

The equation systems encountered in practical application of acoustic holography are typically underdetermined, as there are fewer measurement points than unknown elementary sources, leading to a non-unique, infinite set of solutions. Moreover, a lack of evanescent waves in the measurements causes the singular values of the propagation matrix $\mathbf{G}$ to be very small, making the problem severely ill-conditioned and the solution numerically unstable. Solving such inverse problem poses many challenges and difficulties and almost always requires to apply regularization techniques to introduce additional constraints.

Generally speaking, there are currently three types of mathematical models that, in theory, can be used to accomplish the reconstruction task - these are: analytical or handcrafted models, data-driven deep (learning) models and combinatorial approaches of both techniques. The following section is dedicated to provide a compact overview of all three reconstruction models, to expound their distinct advantages and drawbacks and to highlight their individual potential and limitations.

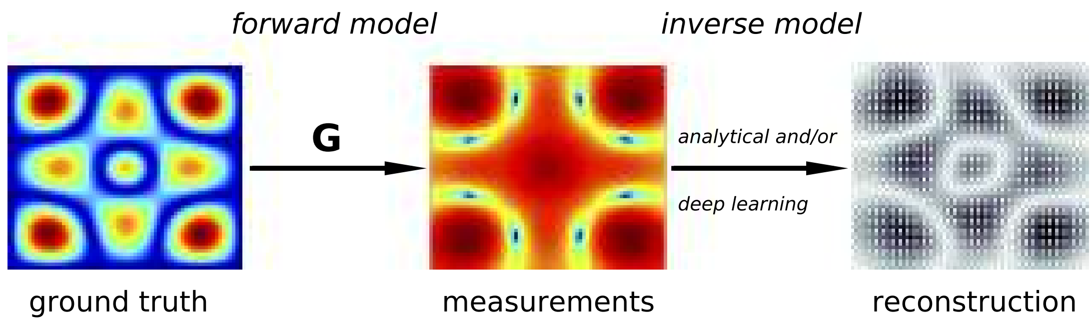

Schematic representation of the forward and inverse problem in acoustic holography.

Handcrafted, deep learning and combinatorial approaches

Handcrafted model: The traditional way of solving ill-posed inverse problems is through analytical methods that rely on explicit mathematical models merged with human knowledge or assumptions about the underlying distribution to be reconstructed [Pag19]. There is a wide variety of both deterministic and stochastic approaches available, all with a solid theoretical foundation and a broad history of research and development. The most common approach of solving the ill-posed inverse problem is to reformulate the problem into an optimization task. The goal, then, is to attain an approximative solution that minimizes the mismatch between the model and measured data while simultaneously minimizing some regularization term that encodes prior knowledge or belief about specific source properties in a certain domain. The optimization problem takes the form

$$\min_{\mathbf{q}} \Big\{ ||\mathbf{p} - \mathbf{G}\mathbf{q}||_2^2 + \lambda \mathcal{R}(\mathbf{W}\mathbf{q}) \Big\},$$

where $||\mathbf{p} - \mathbf{G}\mathbf{q}||_2^2$ is the model mismatch, $\mathcal{R}(\mathbf{W}\mathbf{q})$ is the regularizer to impose additional constraints on the solution, either directly ($\mathbf{W} = \mathbf{I}$) or in a transformed domain ($\mathbf{W} \neq \mathbf{I}$), and λ is the regularization parameter controlling the weight between the residual and constraint. Choosing the right type and amount of regularization is the most critical part of the modeling process as it mainly determines the accuracy of the reconstruction. That being said, the success of these analytical methods rises and falls with the amount and quality of prior knowledge about the source as well as the engineering skills and intuition on parameter tuning. Despite intensive research on new families of regularizers and optimal parameter choice methods during the last few decades, there are still many problems these handcrafted methods face in that regard.

- What if there is little to no prior knowledge about the source available and assumption accuracy can't be guaranteed (as it's often the case in practice)?
- What if there are multiple sources of different character distributed over the investigated surface patch? Most regularizers are designed to globally enforce a specific domain structure, hence they fail at handling miscellaneous source types locally in the same retrieval task.
- What if constraints applied in one domain contradict expected structures in another domain?
- What is the optimal amount of regularization and according to which metric is it best determined? There is no universally applicable automated parameter choice method, so it mostly has to be selected manually.

- Most analytical methods rely on optimization algorithms, whose hyper-parameters also need to be cautiously selected by hand to speed up the procedure while guarantee convergence. Adaptive determination of optimal parameters is often difficult or even impossible.

- The vast majority of existing optimization algorithms is not applicable to real-time implementation of acoustic holography which is substantial for an efficient integration in industrial measurement workflow.

- Some methods benefit from irrelgular array geometries to facilitate wideband holography while others are strictly limited by the spatial sampling limit due to sensor spacing.

Fully deep model: Recent advances in image reconstruction suggest that deep-learning-based models, or deep models, might be the future way to effectively cope with these problems [LIMK18, ZD19, MSLR18]. However, to the best of my knowledge, there is no current research available on the use of such deep models in acoustic holography to facilitate data-driven sound source reconstruction thus far. This thesis is intended to close this research gap.

Deep models utilize so-called neural networks as computational units, whose parameters are trained on huge datasets rather than being explicitly designed to solve the inverse problem. A neural network computes highly non-linear mappings between two spaces, where the mapping is acquired through an observational learning procedure in which the model is intended to reproduce given input-output data pairs by iteratively tweaking its parameters. In theory, a trained neural network can approximate virtually any function to any desired accuracy level. This makes deep models an extremely powerful tool for modern image reconstruction tasks.

The basic building unit of a neural network is a so-called neuron, which takes a weight vector $\mathbf{w}$ and a bias b and passes the biased weighted sum of its input vector coefficients p_i through a non-linear activation function $h(\cdot)$ to model the output. A conventional neural network consists of a successive composition of (hidden) layers of multiple individual neurons that are highly interconnected by combinatorial weights a. All parameterizable network weights can be summarized in the parameter set omega $\Theta = \{w_0, ..., w_m, a, b\}$. The number of neurons determines the width N of the network, whereas the number of layers determines its depth L. Both the geometrical properties of the network and the type of neuronal activation function and weighting factors will impact the models approximation abilities.

A shallow neural network consisting of a single hidden layer (L=1) thus can be expressed as a linear combination of N individually weighted neurons

$$\tilde{f}_{N,L=1}(\mathbf{p}, \Theta) = \sum_{i=1}^{N} a_i h(\mathbf{w}_i^T \mathbf{p} + b_i),$$

which realizes the mapping $\mathbb{R}^m \mapsto \mathbb{R}$. Due to special compositions of multiple such shallow layers, deep neural networks are also able to approximate arbitrary high-dimensional mappings $\mathbb{R}^m \mapsto \mathbb{R}^n$. The architecture of the neural network therefore determines how it transforms its input into output data.

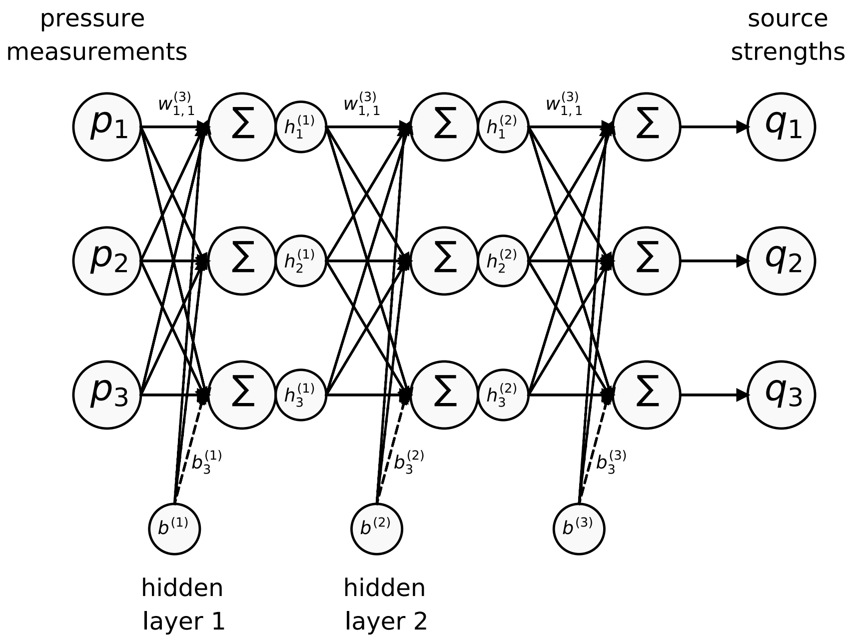

Example of a fully interconnected shallow neural network with 2 hidden layers.

In order to solve the inverse problem using deep learning, we can now introduce a function $\tilde{f}_{N,L}(\mathbf{p}, \Theta)$ that represents an arbitrary neural network with parameter set Θ, which shall be trained to act as a substitute for the inverse $\mathbf{G}^{-1}$ of the propagation matrix. The goal of training is to find the parameter set $\hat{\Theta}$ that minimizes the model mismatch $||\mathbf{q} - \hat{f}_{N,L}(\mathbf{p}, \hat{\Theta})||_2^2$ over a set of predefined training data $(\mathcal{P} \times \mathcal{Q})$ to obtain a direct mapping $\hat{f}_{N,L}(\mathbf{p}, \hat{\Theta})$ from the measurement $\mathbf{p} \in \mathcal{P}$ to the source distribution $\mathbf{q} \in \mathcal{Q}$. The scope and representativity of the training data set $(\mathcal{P} \times \mathcal{Q})$ crucially dictates the accuracy of the mapping and hence the final reconstruction error. Therefore, It should cover as many different source variations - in shape, position, intensity, frequency and character - as possible to make the deep model generally applicable, rather than being task-specific.

Given a set of model parameters Θ, the training procedure is performed by solving an optimization problem, that can be generalized as

$$\min_{\Theta} \Big\{ \frac{1}{\#(\mathcal{P} \times \mathcal{Q})} \sum_{(\mathbf{p},\mathbf{q}) \in (\mathcal{P} \times \mathcal{Q})} \mathcal{J}(\tilde{f}_{N,L}(\mathbf{p}, \Theta), \mathbf{q}) + \mathcal{R}(\Theta) \Big\},$$

where $\#(\mathcal{P} \times \mathcal{Q})$ is the cardinality of the training data set $(\mathcal{P} \times \mathcal{Q})$ which provides desired input-output data pairs $(\mathbf{p}, \mathbf{q})$, $\mathcal{J}(\cdot)$ is an arbitrary loss or cost function comparing the approximated model output $\tilde{f}_{N,L}(\mathbf{p}, \Theta)$ with the ground truth $\mathbf{q}$ and $\mathcal{R}(\cdot)$ is an optional regularization term that helps to address problems like overfitting or bias resulting from limited diversity in training data. The objective function is minimized by iteratively adjusting the weights Θ of the neural network using a suitable optimization algorithm. This network training must take place before the actual reconstructive task can be performed. Once the mapping $\hat{f}_{N,L}(\mathbf{p}, \hat{\Theta})$ is learned from the available data, it can be applied to unseen measurements to achieve the source retrieval.

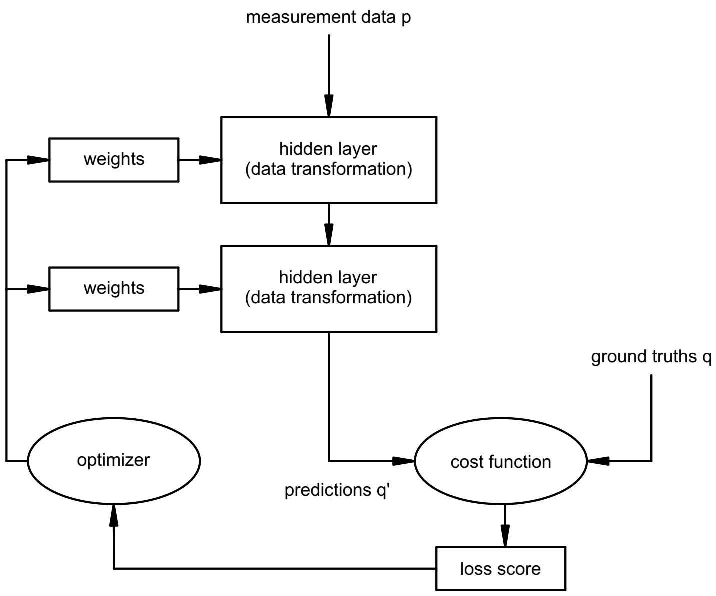

Interaction between neural network, training data, loss function and optimizer.

This data-driven approach for solving the inverse problem has some distinct advantages over traditional handcrafted methods.

- Deep learning automates the most critical aspects of convential reconstruction tasks that normally are adjusted analytically or empirically by humans.
- There essentially is no need of any prior knowledge about the source anymore and hence no necessity of choosing source-dependent regularization terms in order to achieve an accurate reconstruction. The deep model extracts all required source features from training data.
- Manual parameter tuning is not required anymore due to the discard of regularizers.
- Deep models shift the computational burden of optimization from the recontruction task to an offline training routine that is performed beforehand, which allows the actual retrieval tasks (simple algebraic operations) to take place in (nearly) real-time.
- In case important information in the measurements is either missing (loss of evanescent waves due to distance, higher-frequency information due to inadequate spatial sampling) or masked (measurement noise, array imperfections), deep models are able to compensate this lack of information by learning (parts of) it from available training data. Deep learning methods are proven to outperform conventional methods in this regard.
- Deep models were also found to be superior in suppressing artifacts induced by noise and measurement imprecision or incompletion.

On the downside, there are also several challenges and limitations with applying deep learning techniques for sound source reconstruction.

- Learning neural networks often requires huge amounts of training data to be successful. Currently, there is little to no open-source data available that could be utilized to develop deep models for acoustic holography.

- Although mostly in offline training, deep learning can be computationally demanding and therefore heavily relies on parallel computing architectures like GPUs to accelerate computations.
- Due to the black-box nature of deep neural networks, there is a current lack of solid theoretical foundation and clear interpretability of their learned representations. Due to their strong opacity, they have earned a bad reputation, despite their remarkable performance in various tasks and their significant progress in recent years.
- Yet, there are no universal standards or guidelines available on how to build, train and apply neural networks and how to benefit best from different types of datasets. The possible is what works. Thus, deep learning is often thought of as more of an art than an engineering discipline.
- Despite their good suppression of artifacts, deep models may introduce new kinds of artifacts attributed to small, non-representative datasets or architecture peculiarities.
- Certain architecutres of deep neural networks can have lots of hyperparameters which often have to be initiallized manually. Optimal configuration can be difficult.
- Learning the complete direct mapping from scratch may require exhaustive amounts of training data and network parameters to reach a satisfactory reconstruction accuracy.
- What if there is useful prior knowledge available about the source under investigation? Learning the direct mapping does not exploit any human knowledge or known model metadata like measurement distance or sensor spacing, as they are entirely learned from the data seen in training. Incorporation of prior knowledge into a neural network architecture is not straight forward but would definitely be of advantage in most cases.

Handcrafted + deep model: In order to counter some of the problems of deep modelling to make it more generally applicable, a great deal of recent research has gone into the development of methods that combine handcrafted models with deep learning techniques to capture the benefits of both approaches. Currently, there are two main approaches in this regard: post-processing and raw-to-image by unrolling dynamics [ZD19].

Post-processing is a quality enhancement procedure applied to an initial low-quality reconstruction. The actual retrieval task is delegated to the handcrafted, analytical model that provides a rather poor approximation of the source distribution which may contain heavy misestimations and artifacts due to incomplete data and noise. The deep model is used subsequently to upgrade the quality of the retrieved source distribution through artifact and noise removal. However, missing information cannot be reliably retrieved through post-processing.

With the unrolling dynamics approach, which has become one great focus of the more recent research, it is possible to unroll existing and well-studied iterative optimization algorithms to form a feed-forward-network, which is then trained on available data to determine optimal algorithm parameters. By unrolling algorithms for network construction, the trainable elements do not only include convergence-related parameters like step-sizes or proximal weights, but also morphological operators like the Gradient or Hessian, resolvent or proximity operators, adjoints, regularization parameters - almost every part of the algorithm that can be parameterized by specific functions can be approximated by neural networks. This cannot only lead to highly improved convergence rates, but it also enables handy incorporation of handcrafted rules or prior information about the source and/or model physics. This way, one has the freedom to decide which features are best to be extracted from data and which are explicitly known and can be appropriately modelled by handpicked regularization. Reconstruction may even benefit from the inclusion of multiple regularizers, some handcrafted, some parametrized for training to refine the results. Such unrolled algorithmic networks are able to achieve a significant performance increase over traditional fully-handcrafted methods. What is more, deep models generated through unrolling dynamics are generally much more transparent in their structure and approximation processes and hence better interpretable, manipulable and generalizable. Additionally, they typically require a lot less network parameters than conventional deep models which makes them more suitable for training on relatively small datasets. At first sight, given the facts, the unrolling dynamics approach is most likely the best way to utilize the full potential of modern deep-learning-based methods in acoustic holography.

Research questions and proceeding strategy

This PhD thesis is intended to answer the following research questions:

- How to best leverage the great potential of data-driven modelling in acoustic holography to relief users from the burden of crucial decision-making by automating critical parameter selections and regularization?
- How to best make use of existing synthesized or experimental data and how to extract representative data sets thereof? How to design efficient learning models that allow successful training on very limited quantity and diversity of data and that may even capitalize on task unspecific data?
- The accuracy of the learned mapping is always constrained by the data seen in training, and hence largely by the choice of the propagator $\mathbf{G}$ used in the forward model, which depends on distance, frequency and medium properties. Can the deep model also generalize and adapt well to unseen data with different physical and geometrical characteristics? Which type and amount of errors occur during the reconstruction process?
- Can deep learning possibly alleviate the sampling requirements, reducing the number and density of microphones needed to achieve accurate retrieval up to high frequencies? Is it possible to overcome the spatial sampling limit and allow wideband reconstruction independent of the source type?
- How to adress the issue of dealing with the complex-valued nature of sound wave data in neural network training? Do we have to handle real and imaginary parts or magnitude and phase seperately and if so, how?
- Provided that successful source retrieval is possible, can the estimated source strengths be subsequently used for accurate reconstruction of the sound field at any point in space? Which acoustic quantities can be recovered by degree of error?
- Is it possible to compensate the increasing loss of evanescent waves due to increasing measurement distances by extracting the missing information from training data through neural networks? Can this extend the measurement distance while still maintaining a high spatial resolution?
- Can deep models obtain a higher reconstruction accuracy in the presence of noise, artifacts and measurement imperfections? How to best incorporate intentionally-generated errors in the training data to increase the methods robustness and suppressive capabilities?
- What are the new types of errors that we face in deep-learning-based acoustic holography? Are we able to detect the origins and develop any feasible counter strategies?

On a larger picture, the main goals of this thesis are to ...

- ... propose a generalized deep-learning-based framework for acoustic holography that is universally applicable independent of the source's nature, completely or at least partly detached from any prior knowledge about the source.
- ... try to find the optimal balance between imposing engineering and data-driven knowledge on the holographic task and eventually come up with a set of rules or guidelines on how to achieve optimal results.
- ... construct a representative database of synthesized and experimental source-field pairs with open-source access to enable and promote future research on deep learning strategies for sound source reconstruction.
- ... pursue mathematical explanations and justifications for the possible success of deep models instead of just taking it for granted.

Strategy and roadmap to answer the research questions:

- Set up a (local or cloud-based) deep-learning workstation, choose a suitable deep learning framework (Keras, TensorFlow, Theano, Torch, Caffe, etc.) and become familiar with the syntax. The recommended programming language for research of deep learning is Python, in which I am already experienced, which eases the start-up tasks.
- If you think of deep learning as the possible engine of quality improvement in source reconstruction, then training data is the fuel. Thus, the next logical step is to set up a database of synthetic and experimental acoustic radiation data. Currently, there is no sound source radiation database available to the acoustic imaging community to facilitate research on deep learning for source retrieval. As a long-term goal, the construction of a representative database of synthesized and experimental source-field pairs, using established forward models like ESM or BEM, is therefore of utmost importance for progress in deep-learning-based sound source reconstruction. This database must include a wide variety of source types, including point, area, and line sources of different geometric shapes and at several distinct positions, radiating at different frequencies and velocities, singly or in arbitrary combinations. The resulting sound fields must be evaluated at different positions within a certain distance to the source and may be corrupted by different realizations of noise to simulate measurement errors. Organizing the collected data with focus on optimal workflow integration using data management tools may also be of advantage.
- The next step is to explore the fundamental theory and principles of deep learning needed to fully understand the aspects necessary for model building and evaluation in the context of source retrieval. Additionally, it may be beneficial to conduct a comprehensive review about the current limitations and drawbacks of traditional acoustic holography as a segue into how deep learning may replace or supplement the current state-of-the-art to overcome these challenges and shortcomings.
- Deep learning for image reconstruction is a new area of research that only recently became more popular in various types of imaging systems due to an increased availability of data and information on the task, but progress is being made very rapidly. It is expected to revolutionize the field of image reconstruction as it has done in many other areas of image and signal processing, which provide a rich computational toolbox from which we can choose freely. Since, to the best of my knowledge, there is no available research on deep-learning-based acoustic holography, this thesis shall also present an extensive survey of existing literature on deep learning techniques from recent years with main focus on the field of image reconstruction. The goal is to extract what is useful, discard what is useless, and seeking to bring in new ideas in order to eventually come up with a framework that is best suitable for acoustic imaging.

- Exploring different types of neural network architectures, different parametrizations, loss functions used for measuring training success, optimizers used to find the optimal net parameters by minimizing the loss, activation functions used to model the neurons' behaviour, initialization techniques, strategies to prevent over- and underfitting, the different types of testing and validation of the model, etc. to become familiar with current deep learning best practices.

- Since training is the most crucial and difficult part of deep models that determines their success or failure, this is presumably where the main focus will lie on. In this regard, priority will be given to studying the effects of different training data in terms of versatility, sample size, geometrical setup, degradation, etc. Other key things to address are training algorithms, regularization and acceleration strategies, as well as training metrics, data prepocessing and augmentation.

- Then comes the task-dependent selection of suitable deep models both from a direct mapping perspective and by exploring possible links between existing iterative optimization algorithms used for image reconstruction and special structures of neural networks to maybe benefit from a fused implementation.

- At last, a comprehensive simulation study and experimental validation must be conducted to prove the effectiveness of the proposed deep learning techniques by adressing the main research questions. In addition, a computational performance analysis may be provided to show real-time capabilities of proposed algorithms. This may also include a convergence study of learning algorithms used for network training.

Surveys of ML/DL for image reconstruction

[BR18] M. Bakator and D. Radosav, “Deep learning and medical diagnosis: A review of literature,” *Multimodal Technologies and Interaction*, vol. 2, p. 47, 08 2018.

[KMR18] F. Knoll, A. Maier, and D. Rueckert, *Machine Learning for Medical Image Reconstruction*. Springer, 2018.

[LIMK18] A. Lucas, M. Iliadis, R. Molina, and A. Katsaggelos, “Using deep neural networks for inverse problems in imaging: Beyond analytical methods,” *IEEE Signal Processing Magazine*, vol. 35, pp. 20–36, 01 2018.

[LJC+17] J.-G. Lee, S. Jun, Y.-W. Cho, H. Lee, G. Kim, J. Seo, and N. Kim, “Deep learning in medical imaging: General overview,” *Korean Journal of Radiology*, vol. 18, p. 570, 07 2017.

[LKEB+17] G. Litjens, T. Kooi, B. Ehteshami Bejnordi, A. Setio, F. Ciompi, M. Ghafoorian, J. van der Laak, B. Ginneken, and C. SÃ¡nchez, “A survey on deep learning in medical image analysis,” *Medical Image Analysis*, vol. 42, 02 2017.

[MSLR18] A. Maier, C. Syben, T. Lasser, and C. Riess, “A gentle introduction to deep learning in medical image processing,” 2018.

[WYMF18] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler, “Image reconstruction is a new frontier of machine learning,” *IEEE Transactions on Medical Imaging*, vol. 37, pp. 1289–1296, 2018.

[ZD19] H. Zhang and B. Dong, “A review on deep learning in medical image reconstruction,” 2019.

Deep learning in medical imaging

[AMJ17] H. Aggarwal, M. Mani, and M. Jacob, “Modl: Model based deep learning architecture for inverse problems,” *IEEE Transactions on Medical Imaging*, vol. PP, 12 2017.

[AÖ17] J. Adler and O. Öktem, “Solving ill-posed inverse problems using iterative deep neural networks,” *Inverse Problems*, vol. 33, no. 12, p. 124007, Nov 2017. [Online]. Available: http://dx.doi.org/10.1088/1361-6420/aa9581

[AÖ18] ——, “Learned primal-dual reconstruction,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, p. 1322?1332, Jun 2018. [Online]. Available: http://dx.doi.org/10.1109/TMI.2018.2799231

[CCA+18] J. Cheng, F. Chen, M. Alley, J. Pauly, and S. Vasanawala, “Highly scalable image reconstruction using deep neural networks with bandpass filtering,” 05 2018.

[DCE+18] M. Dedmari, S. Conjeti, S. Estrada, P. Ehses, T. Stoecker, and M. Reuter, “Complex fully convolutional neural networks for mr image reconstruction,” 07 2018.

[GSB+19] V. Ghodrati, J. Shao, M. Bydder, Z. Zhou, W. Yin, K.-L. Nguyen, Y. Yang, and P. Hu, “Mr image reconstruction using deep learning: evaluation of network structure and loss functions,” *Quantitative Imaging in Medicine and Surgery*, vol. 9, pp. 1516–1527, 09 2019.

[HKK+17] K. Hammernik, T. Klatzer, E. Kobler, M. Recht, D. Sodickson, T. Pock, and F. Knoll, “Learning a variational network for reconstruction of accelerated mri data,” *Magnetic Resonance in Medicine*, vol. 79, 04 2017.

[KHK+17] F. Knoll, K. Hammernik, E. Kobler, T. Pock, M. Recht, and D. Sodickson, “Assessment of the generalization of learned image reconstruction and the potential for transfer learning,” *Magnetic Resonance in Medicine*, vol. 81, 12 2017.

[KKHP17] E. Kobler, T. Klatzer, K. Hammernik, and T. Pock, “Variational networks: Connecting variational methods and deep learning,” pp. 281–293, 08 2017.

[KMA17] B. Kelly, T. Matthews, and M. Anastasio, “Deep learning-guided image reconstruction from incomplete data,” 09 2017.

[KMC+18] E. Kobler, M. Muckley, B. Chen, F. Knoll, K. Hammernik, T. Pock, D. Sodickson, and R. Otazo, “Variational deep learning for low-dose computed tomography,” pp. 6687–6691, 04 2018.

[LIMK18] A. Lucas, M. Iliadis, R. Molina, and A. Katsaggelos, “Using deep neural networks for inverse problems in imaging: Beyond analytical methods,” *IEEE Signal Processing Magazine*, vol. 35, pp. 20–36, 01 2018.

[OKC+18] C. Oh, D. Kim, J.-Y. Chung, Y. Han, and H. Park, *ETER-net: End to End MR Image Reconstruction Using Recurrent Neural Network: First International Workshop, MLMIR 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings*, 09 2018, pp. 12–20.

[WYMF18] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler, “Image reconstruction is a new frontier of machine learning,” *IEEE Transactions on Medical Imaging*, vol. 37, pp. 1289–1296, 2018.