Equivariant Neural Networks for Indirect Measurements*

Matthias Beckmann[†] and Nick Heilenkötter[‡]

Abstract. In recent years, deep learning techniques have shown great success in various tasks related to inverse problems, where a target quantity of interest can only be observed through indirect measurements by a forward operator. Common approaches apply deep neural networks in a post-processing step to the reconstructions obtained by classical reconstruction methods. However, the latter methods can be computationally expensive and introduce artifacts that are not present in the measured data and, in turn, can deteriorate the performance on the given task. To overcome these limitations, we propose a class of equivariant neural networks that can be directly applied to the measurements to solve the desired task. To this end, we build appropriate network structures by developing layers that are equivariant with respect to data transformations induced by well-known symmetries in the domain of the forward operator. We rigorously analyze the relation between the measurement operator and the resulting group representations and prove a representer theorem that characterizes the class of linear operators that translate between a given pair of group actions. Based on this theory, we extend the existing concepts of Lie group equivariant deep learning to inverse problems and introduce new representations that result from the involved measurement operations. This allows us to efficiently solve classification, regression or even reconstruction tasks based on indirect measurements also for very sparse data problems, where a classical reconstruction-based approach may be hard or even impossible. We illustrate the effectiveness of our approach in numerical experiments and compare with existing methods.

Key words. Equivariant Deep Learning, Operator Equivariance, Inverse Problems

AMS subject classifications. 47B38, 65J22, 94A05

1. Introduction. Indirect measurements naturally occur in various applications, where a quantity of interest cannot be observed directly but only through effects of certain processes leading to an *inverse problem*. One prominent example is computerized tomography (CT), where the goal consists in imaging the interior of a scanned object by measuring and processing the attenuation of X-rays travelling through the object under investigation. Mathematically, the common setting in inverse problems is the observation of a signal x through indirect measurements $y = \mathscr{A}x$ with forward operator \mathscr{A} . One particular challenge is that inverse problems are typically ill-posed in the sense that no continuous inverse \mathscr{A}^{-1} exists so that small noise in the measured data y can lead to large errors in the reconstructed signal x.

Recently, deep learning approaches based on neural networks have shown remarkable results in various tasks in the context of inverse problems, like classification or regression tasks. Such methods typically include the solution of the inverse problem. In contrast to this, the idea presented in this work is to construct data-efficient neural network structures that operate directly on the measured data y and avoid the need for classical reconstruction or backprojection

^{*}Submitted to the editors February 28, 2024.

Funding: N. Heilenkötter acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG) – Project number 281474342.

[†]Center for Industrial Mathematics, University of Bremen, Germany & Department of Electrical and Electronic Engineering, Imperial College London, UK (research@mbeckmann.de).

[‡]Center for Industrial Mathematics, University of Bremen, Germany (heilenkoetter@uni-bremen.de).



Figure 1. The setting of our approach, illustrated for 2-angle fan-beam Radon measurements of tubes: symmetries in \mathcal{Y} (right) are induced by well-known group representations in the physical space \mathcal{X} (left). However, the data is measured only at a finite set of sensor points V, depicted by gray circles.

into the original signal domain. For this purpose, we follow the concept of equivariant neural networks that leverage the symmetries that may be present in the measurements y.

This might become useful especially in settings where only very sparse measurement data is available and, as a result, good reconstructions without additional data-knowledge are infeasible. A motivating example for our work is the thickness measurement of tubes using a sparse two-angle fan-beam CT geometry, as displayed in Figure 1. In this context, one cannot expect to obtain good reconstructions (which could then be handled by learned regression models) by classical methods. However, a full reconstruction is not needed, only scalar values for the minimal and maximal thickness are desired. The idea we follow in this paper is then to circumvent any classical reconstruction and construct an appropriate learned regression model that takes the sparse sinogram data as an input.

On this path, several theoretical and numerical challenges arise. First of all, to construct equivariant networks, a characterization of sensible symmetry transforms on the measured signals y is required. Due to the discrete nature of measurements in practice, symmetries often occur only in an approximate way, requiring a modelling scheme that allows to characterize the underlying continuous symmetries. Finally, one has to construct neural networks that respect these symmetries while operating on the discrete measured data. In this work, we will tackle all these challenges and demonstrate a successful construction of data-efficient network structures that operate on indirect measurements.

1.1. Literature Review. The exploitation of symmetries to improve the generalization performance of deep learning models has been a large field of research during the last years. Approaches include data augmentation [24], specialized training schemes or loss functions [4, 11] and specially tailored equivariant network structures, on which we focus in the following.

Motivated by the success of convolutional neural networks (CNNs), the research on generalized group-equivariant neural network architectures has led to network structures that show promising results on various types of data [14, 22]. Group-equivariance is commonly perceived as a strong inductive bias for the construction of data-efficient network architectures. Its success is believed to rely on two basic properties [22]: On the one hand, equivariance incorporates prior knowledge on the data to provide guarantees for the extrapolation of results to unseen transformed data points and enables parameter-sharing. On the other hand, the non-trivial kernel size of convolutions together with downsampling layers allow to learn local features and combine those features on coarser scales.

In 2016, Cohen and Welling [7] constructed generalized CNNs for image processing that are equivariant with respect to discrete groups. Since then, several extensions and variants have been proposed. For example, in [8] the authors develop memory-efficient steerable convolution kernels based on group representation theory. Lately, Finzi et al. [14] constructed neural networks that are equivariant with respect to arbitrary Lie groups. To handle the continuous groups on point clouds, they parametrized convolution kernels by small fully-connected networks. Since we aim to treat arbitrary discretizations, we also follow this approach for the construction of equivariant networks.

In inverse problems, a vast range of classical theory and numerical methods concerned with the reconstruction of x from possibly noisy data $y^{\delta} = y + \eta^{\delta}$ exists [13, 20]. Symmetries have been a central tool in the development of methods, ranging from steerable wavelets [5] to regularization by equivariance [29].

During the last years, diverse deep learning methods have shown to improve the quality of reconstructions [1, 26]. The incorporation of symmetries in such reconstructions turned out to be a valuable prior and has been studied in various approaches [2, 4, 9, 29, 3].

In [2], the authors use an equivariant neural network architecture as proximal operator to improve the reconstruction quality obtained by learned iterative schemes. In this context, possible issues caused by non-equivariant measurement operators that partially regard information are shortly discussed. On the contrary, the reconstruction approach in [4] makes explicit use of this property to design specialized loss functions for a self-supervised learning strategy for end-to-end reconstructions. In our approach, this effect is handled by the separation of continuous forward measurement operators and their partial evaluation, e.g., due to discretization.

Other deep learning based reconstruction approaches include symmetry knowledge by the direct use of equivariant networks for post-processing of classical reconstructions [33] or in the construction of learned regularizations in variational methods [29]. To the best of our knowledge, none of the existing works has tackled the problem of directly handling induced symmetries in indirect measurements.

1.2. Overview of the Paper. The outline of this paper is as follows. In Section 2 we explain the general setting and introduce the necessary mathematical preliminaries for the analysis of group symmetries in indirect measurements in a compact way. Section 3 is the theoretical main part of this paper and contains two mathematical results that investigate the connection between measurement operators and the symmetries that may appear in the measured data. In particular, Theorem 3.2 sharply characterizes all linear operators that

translate between given symmetry transforms. For illustration, we apply this result to the classical Radon operator that models the measurement process in computerized tomography. Extending on this knowledge, in Section 4 we build equivariant neural networks with respect to new types of transforms that have not been tackled before and treat the challenges that appear in the context of applied inverse problems, where only partial measurements are available. Finally, in Section 5 we apply our constructed networks that can operate directly on the measured data and perform numerical experiments on classification and regression tasks, where we are able to show improved generalization properties when compared to existing methods that are based on learned post-processing of classical reconstructions.

2. Mathematical Setting and Preliminaries.

2.1. General Setting. In this paper, we deal with linear inverse problems $y = \mathscr{A}x$ that can be modelled by a continuous linear forward operator

$$\mathscr{A}:\mathcal{X}\to\mathcal{Y}$$

mapping between spaces of test functions $\mathcal{X} = \mathcal{D}(\Omega_{\mathcal{X}}, \mathbb{K}^{n_{\mathcal{X}}})$ on some domain $\Omega_{\mathcal{X}} \subseteq \mathbb{R}^{d_{\mathcal{X}}}$ and $\mathcal{Y} = \mathcal{D}(\Omega_{\mathcal{Y}}, \mathbb{K}^{n_{\mathcal{Y}}})$ on some domain $\Omega_{\mathcal{Y}} \subseteq \mathbb{R}^{d_{\mathcal{Y}}}$, where $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$. Since we focus on scalar data in the application, we will abbreviate notation and mostly omit the signal range dimension, i.e. we write $\mathcal{X} = \mathcal{D}(\Omega_{\mathcal{X}}), \mathcal{Y} = \mathcal{D}(\Omega_{\mathcal{Y}})$. Keep in mind that we implicitly also treat vector field transforms and, hence, kernels are matrix-valued. In Section 4 we will add a discretization scheme to this continuous setting to deal with discrete and potentially noisy measurements.

The ultimate aim of our work is to construct artificial neural networks

$$\mathcal{N}_{\theta}: \mathcal{Y} \to \mathcal{Z}$$

that can be efficiently applied to the measured data and optimized to solve several postprocessing tasks while avoiding explicitly solving the inverse problem in a first step. To this end, we will make use of symmetries in the measurement space \mathcal{Y} that are known in the underlying source space \mathcal{X} and transformed by the forward operator \mathscr{A} .

2.2. Symmetries and Group Representations. In many cases, continuous local or global symmetries of a source signal $x \in \mathcal{X}$ can be described by means of a known representation \mathscr{P} , i.e., a homomorphism $\mathscr{P} : G \to \operatorname{GL}(\mathcal{X})$, of a Lie group G on \mathcal{X} . To emphasize the underlying space \mathcal{X} we also write $\mathscr{P} \equiv \mathscr{P}_{\mathcal{X}}$. For vector spaces \mathcal{V}, \mathcal{W} , we call an operator $\mathscr{B} : \mathcal{V} \to \mathcal{W}$ equivariant with respect to a representation $\mathscr{P}_{\mathcal{V}}$ on \mathcal{V} if a representation $\mathscr{P}_{\mathcal{W}}$ on \mathcal{W} exists such that for all $x \in \mathcal{V}$ and $g \in G$ we have

$$\mathscr{P}_{\mathcal{W}}[g]\mathscr{B}(x) = \mathscr{B}(\mathscr{P}_{\mathcal{V}}[g]x).$$

A common type of representations \mathscr{P} on function spaces $\mathcal{V} \subseteq \{\mathbb{R}^d \supseteq \Omega \to \mathbb{R}^n\}$ are socalled *domain transforms*, which are induced by a proper Lie group action π , i.e., a group homomorphism $\pi : G \to \{\Omega \to \Omega\}$ with $\pi[e] = I_d$, acting on the functions $x \in \mathcal{V}$ by (possibly non-linearly) transforming their domain Ω so that

$$(\mathscr{P}[g]x)(u) = x(\pi[g]^{-1}(u)).$$

Here, a group action is called proper if the mapping $G \times \Omega \to \Omega \times \Omega$, $(g, u) \mapsto (\pi[g](u), u)$ is proper, i.e., pre-images of compact sets are compact.

To deal with more general transforms that can appear in measurements, we introduce the notion of *generalized domain transforms*, which are representations that can additionally act locally and linearly on the range of the functions via

$$(\mathscr{P}[g]x)(u) = p[g](u) x(\pi[g]^{-1}(u)),$$

where $p: G \to \mathcal{C}^{\infty}(\Omega, \operatorname{GL}(\mathbb{K}^n))$ satisfies $p[gh](u) = p[g](u) p[h](\pi[g]^{-1}(u))$ and $p[e] = I_n$, which is also the sufficient and necessary condition for the induced mapping \mathscr{P} to define a group representation on \mathcal{V} . In this case we also write $\mathscr{P} \equiv \mathscr{P}(\pi, p)$ to emphasize the mappings $\pi: G \to \{\Omega \to \Omega\}$ and $p: G \to \mathcal{C}^{\infty}(\Omega, \operatorname{GL}(\mathbb{K}^n))$. While representations in the range of functions are also considered for vector fields e.g. in [6], the symmetries that appear in indirect measurements are new in the way that they can contain local variations also for scalar signals. This behaviour is properly modelled by the introduced class of representations.

All (generalized) domain transforms are primarily characterized by the action of π on the domain Ω . If for all $u, u_0 \in \Omega$ there exists $g \in G$ such that $\pi[g](u_0) = u$, the set Ω is called a *homogeneous space* of π . In this case, the domain Ω can be parametrized by the group elements once an origin $u_0 \in \Omega$ has been fixed. To this end, we consider the *stabilizer subgroup*

$$N_{\pi}(u_0) = \{ n \in G \mid \pi[n](u_0) = u_0 \}.$$

of all elements that keep the origin fixed. Then, the set $G/N_{\pi}(u_0) = \{gN_{\pi}(u_0) \mid g \in G\}$ of *left* cosets can be identified with the domain, i.e., we have $G/N_{\pi}(u_0) \cong \Omega$. This observation will be an important tool to establish a relation between equivariant operators and convolutions. For the sake of brevity, in the following we write g_u for a representative of the group elements that move the origin u_0 to u, i.e., $\pi[g_u](u_0) = u$, and specify its choice once necessary.

3. Mathematical Results. As the overall goal of this work is to define equivariant neural networks on the measurements $\mathscr{A}x, x \in \mathcal{X}$, it is desirable to describe the symmetries in the measurement space \mathcal{Y} . Given a symmetry transform $\mathscr{P}_{\mathcal{X}}$ on the source space \mathcal{X} , we consider the transforms that are induced by the measurement, i.e., for all $g \in G$, we search for a mapping $\mathscr{P}_{\mathcal{Y}}[g] : \operatorname{Rg}(\mathscr{A}) \to \operatorname{Rg}(\mathscr{A})$ that satisfies

$$\mathscr{P}_{\mathcal{Y}}[g]\mathscr{A}x = \mathscr{A}\mathscr{P}_{\mathcal{X}}[g]x$$

for all $x \in \mathcal{X}$. The following visibility condition gives a sharp criterion for the existence and properties of such a mapping.

Theorem 3.1 (visibility condition). For a group representation $\mathscr{P}_{\mathcal{X}}$ on \mathcal{X} and a measurement operator $\mathscr{A} : \mathcal{X} \to \mathcal{Y}$ the following statements are equivalent.

- (i) The mapping $\mathscr{P}_{\mathcal{Y}}[g] : \operatorname{Rg}(\mathscr{A}) \to \operatorname{Rg}(\mathscr{A})$ given by (3.1) is well-defined for all $g \in G$.
- (ii) Equation (3.1) defines a group representation $\mathscr{P}_{\mathcal{Y}}: G \to (\operatorname{Rg}(\mathscr{A}) \to \operatorname{Rg}(\mathscr{A}))$ on \mathcal{Y} .
- (iii) We have $\operatorname{Ker}(\mathscr{AP}_{\mathcal{X}}[g]) = \operatorname{Ker}(\mathscr{A})$ for all $g \in G$.

Proof. (i) \Longrightarrow (ii): Let $\mathscr{P}_{\mathcal{Y}}: G \to (\operatorname{Rg}(\mathscr{A}) \to \operatorname{Rg}(\mathscr{A}))$ be defined via (3.1). Then, $\mathscr{P}_{\mathcal{Y}}$ is a group homomorphism as, for $y = \mathscr{A}_{\mathcal{X}} \in \operatorname{Rg}(\mathscr{A})$, we have $\mathscr{P}_{\mathcal{Y}}[e]y = \mathscr{P}_{\mathcal{Y}}[e]\mathscr{A}_{\mathcal{X}} = \mathscr{A}_{\mathcal{X}} = y$

and, for arbitrary $g, h \in G$, direct computations show that

$$\begin{split} \mathscr{P}_{\mathcal{Y}}[gh]y &= \mathscr{P}_{\mathcal{Y}}[gh]\mathscr{A}x = \mathscr{A}\mathscr{P}_{\mathcal{X}}[gh]x = \mathscr{A}\mathscr{P}_{\mathcal{X}}[g]\mathscr{P}_{\mathcal{X}}[h]x \\ &= \mathscr{P}_{\mathcal{Y}}[g]\mathscr{P}_{\mathcal{Y}}[h]\mathscr{A}x = \mathscr{P}_{\mathcal{Y}}[g]\mathscr{P}_{\mathcal{Y}}[h]y. \end{split}$$

Moreover, for $g \in G$ the mapping $\mathscr{P}_{\mathcal{Y}}[g] : \operatorname{Rg}(\mathscr{A}) \to \operatorname{Rg}(\mathscr{A})$ is linear as for $\lambda_1, \lambda_2 \in \mathbb{R}$ and $y_1 = \mathscr{A}x_1, y_2 = \mathscr{A}x_2 \in \operatorname{Rg}(\mathscr{A})$ holds that

$$\begin{aligned} \mathscr{P}_{\mathcal{Y}}[g](\lambda_{1}y_{1}+\lambda_{2}y_{2}) &= \mathscr{P}_{\mathcal{Y}}[g]\mathscr{A}(\lambda_{1}x_{1}+\lambda_{2}x_{2}) = \mathscr{A}\mathscr{P}_{\mathcal{X}}[g](\lambda_{1}x_{1}+\lambda_{2}x_{2}) \\ &= \lambda_{1}\mathscr{A}\mathscr{P}_{\mathcal{X}}[g]x_{1}+\lambda_{2}\mathscr{A}\mathscr{P}_{\mathcal{X}}[g]x_{2} = \lambda_{1}\mathscr{P}_{\mathcal{Y}}[g]y_{1}+\lambda_{2}\mathscr{P}_{\mathcal{Y}}[g]y_{2}. \end{aligned}$$

(ii) \implies (iii): Let $\mathscr{P}_{\mathcal{Y}} : G \to (\operatorname{Rg}(\mathscr{A}) \to \operatorname{Rg}(\mathscr{A}))$ be the above group representation and let $g \in G$ such that

$$\operatorname{Ker}(\mathscr{A}\mathscr{P}_{\mathcal{X}}[g]) \neq \operatorname{Ker}(\mathscr{A}).$$

Then, there is $x \in \text{Ker}(\mathscr{A})$ with $x \notin \text{Ker}(\mathscr{AP}_{\mathcal{X}}[g])$ or $\hat{x} \in \text{Ker}(\mathscr{AP}_{\mathcal{X}}[g])$ with $\hat{x} \notin \text{Ker}(\mathscr{A})$. In the first case, however, we have

$$\mathscr{A}\mathscr{P}_{\mathcal{X}}[g]x = \mathscr{P}_{\mathcal{Y}}[g]\mathscr{A}x = 0$$

in contradiction to $x \notin \operatorname{Ker}(\mathscr{AP}_{\mathcal{X}}[g])$. In the second case, we similarly get the contradiction

$$\mathscr{A}\hat{x} = \mathscr{A}\mathscr{P}_{\mathcal{X}}[g^{-1}]\mathscr{P}_{\mathcal{X}}[g]\hat{x} = \mathscr{P}_{\mathcal{Y}}[g^{-1}]\mathscr{A}\mathscr{P}_{\mathcal{X}}[g]\hat{x}.$$

(iii) \implies (i): Let $\operatorname{Ker}(\mathscr{AP}_{\mathcal{X}}[g]) = \operatorname{Ker}(\mathscr{A})$ for all $g \in G$ and let $y = \mathscr{A}x_1 = \mathscr{A}x_2$ for $x_1, x_2 \in \mathcal{X}$. Then, $x_1 - x_2 \in \operatorname{Ker}(\mathscr{A}) = \operatorname{Ker}(\mathscr{AP}_{\mathcal{X}}[g])$ and, consequently, for arbitrary $g \in G$ follows that

$$\mathscr{AP}_{\mathcal{X}}[g]x_1 = \mathscr{AP}_{\mathcal{X}}[g]x_2 \quad \Longrightarrow \quad \mathscr{P}_{\mathcal{Y}}[g]\mathscr{A}x_1 = \mathscr{P}_{\mathcal{Y}}[g]\mathscr{A}x_2,$$

which shows that $\mathscr{P}_{\mathcal{Y}}[g] : \operatorname{Rg}(\mathscr{A}) \to \operatorname{Rg}(\mathscr{A})$ is well-defined.

The sharp condition $\operatorname{Ker}(\mathscr{AP}_{\mathcal{X}}[g]) = \operatorname{Ker}(\mathscr{A})$ for all $g \in G$ can be interpreted in terms of the visibility of information in the measurements: If a feature in the source data can be made visible or invisible in the measurement by the application of a group transform, it is impossible to characterize $\mathscr{AP}_{\mathcal{X}}[g]x$ only in terms of x and g. Otherwise, the output can be determined solely from x and g and, in addition, behaves well in the sense of a group representation structure. An example for the violation of the visibility condition due to partial measurements is depicted in Figure 2. Due to the previous theorem, classical reconstructions that do not incorporate any additional assumptions on the data cannot be equivariant with respect to group representations in the source space. As a consequence, we will focus on applications in which the continuous full measurement operator can be modelled such that it satisfies the visibility condition. Violations of the visibility condition that are caused by the discretization of the operator are discussed in more detail in Section 4.

In a finite-dimensional setting, the authors of [4] state a related condition, where the non-fulfilment of the visibility condition is an explicit prerequisite for their approach to be of



Figure 2. Illustration how classical reconstructions (right) of an MNIST digit based on sparse measurements (center, visualized by circles) behave under transformations from $Aff^+(2) = T(2) \rtimes GL^+(2)$ (left). In this case, the resulting representation in the full sinograms (center) is a generalized domain transform.

advantage for reconstruction. This is in contrast to our approach, where equivariance is built explicitly into the network structure and requires the fulfilment of the visibility condition.

Our next goal is to exactly characterize the class of measurement operators that transform between certain group representations. For this purpose, we restrict ourselves to the case of generalized domain transforms as group representation in the source space. In practice, the representations in the source space are well-known and the notion of generalized domain transforms covers all symmetries that are commonly considered in the known physical domain, cf. [7, 14, 6]. In addition, we assume that the domain $\Omega_{\mathcal{X}}$ of the input signal is a homogeneous space of the corresponding group action on \mathcal{X} . This assumption, however, can be weakened to obtain a more general but weaker result by treating the domain as a union of orbits, which are homogeneous spaces on their own. To treat general linear measurement operators with a priori unknown induced symmetries, we do not restrict the representations $\mathcal{P}_{\mathcal{Y}}$ in any way.

In the following, for a distribution $a \in \mathcal{D}'(\Omega_{\mathcal{X}})$ and a function $x \in \mathcal{D}(\Omega_{\mathcal{X}})$, we use the notation

$$a(x) = \int_{\Omega_{\mathcal{X}}} a(u) x(u) \, \mathrm{d}u.$$

For a group representation $\mathscr{P}_{\mathcal{X}}: G \to (\mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{X}}))$ we define its *adjoint representation* $\mathscr{P}_{\mathcal{X}}^*: G \to (\mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{X}}))$ via

$$\int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{X}}^*[g]x)(u) \ \varphi(u) \ \mathrm{d}u = \int_{\Omega_{\mathcal{X}}} x(u) \ (\mathscr{P}_{\mathcal{X}}[g]\varphi)(u) \ \mathrm{d}u \quad \forall \ \varphi \in \mathcal{D}(\Omega_{\mathcal{X}})$$

and, with this, for $a \in \mathcal{D}'(\Omega_{\mathcal{X}})$ and $g \in G$ we define $\mathscr{P}_{\mathcal{X}}[g]a \in \mathcal{D}'(\Omega_{\mathcal{X}})$ via

$$\mathscr{P}_{\mathcal{X}}[g]a(x) = a(\mathscr{P}_{\mathcal{X}}^*[g]x) \quad \forall \ x \in \mathcal{D}(\Omega_{\mathcal{X}}).$$

Theorem 3.2 (Equivariant measurement operator is a convolution.). Let $\mathscr{P}_{\mathcal{X}} \equiv \mathscr{P}_{\mathcal{X}}(\pi_{\mathcal{X}}, p_{\mathcal{X}})$ be a generalized domain transform such that $\Omega_{\mathcal{X}}$ is a homogeneous space of $\pi_{\mathcal{X}}$ and let $\mathscr{A} : \mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{Y}})$ be a continuous linear measurement operator. Then, the equivariance condition (3.1) holds if and only if \mathscr{A} is given by a generalized convolution operator

(3.2)
$$(\mathscr{A}x)(v) = (a *_{\mathcal{X} \to \mathcal{Y}} x)(v)$$
$$\coloneqq \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{Y}}[g_u] a)(v) \ p_{\mathcal{X}}[g_u^{-1}](u_0) x(u) \ |\det(\mathrm{D}\pi_{\mathcal{X}}[g_u](u_0))|^{-1} \,\mathrm{d}u,$$

where the distributional kernel $a \in \mathcal{D}'(\Omega_{\mathcal{Y}})$ fulfils

(3.3)
$$(\mathscr{P}_{\mathcal{Y}}[g_u n] a)(v) = (\mathscr{P}_{\mathcal{Y}}[g_u] a)(v) \ p_{\mathcal{X}}[n](u_0) \ |\det(\mathrm{D}\pi_{\mathcal{X}}[n](u_0))|$$

in the sense of distributions for all $n \in N_{\pi_{\mathcal{X}}}(u_0)$.

Before we proceed to the proof of this theorem, we wish to pass the following remarks.

The condition (3.3) guarantees that the representation in equation (3.2) is independent of the choice of the representative g_u of the group elements that move u_0 to u.

The theorem has several implications for the handling of symmetries in inverse problems. For instance, in the setting of indirect measurements, it enables us to deduce the induced representations from integral measurement operators and to exactly state the set of operators that translate between these symmetries.

The expression *generalized convolution* is motivated by the following remark.

Remark 3.3. If $\mathscr{P}_{\mathcal{X}} \equiv \mathscr{P}_{\mathcal{X}}(\pi_{\mathcal{X}})$ and $\mathscr{P}_{\mathcal{Y}} \equiv \mathscr{P}_{\mathcal{Y}}(\pi_{\mathcal{Y}})$ are conventional domain transforms, representation (3.2) reads

$$(\mathscr{A}x)(v) = \int_{\Omega_{\mathcal{X}}} a(\pi_{\mathcal{Y}}[g_u^{-1}](v)) \ x(u) \ |\det(\mathrm{D}\pi_{\mathcal{X}}[g_u](u_0))|^{-1} \ \mathrm{d}u.$$

Moreover, condition (3.3) reduces to

$$a(\pi_{\mathcal{Y}}[g_u n]^{-1}v) = a(\pi_{\mathcal{Y}}[g_u]^{-1}(v)) \left| \det(\mathrm{D}\pi_{\mathcal{X}}[n](u_0)) \right| \quad \forall \ n \in N_{\pi_{\mathcal{X}}}(u_0)$$

and, hence, constrains the kernel along the orbits of the stabilizer of $\pi_{\mathcal{X}}$ in \mathcal{Y} . In particular, the special case $\Omega_{\mathcal{X}} = \mathbb{R}^d = \Omega_{\mathcal{Y}}$ with origin $u_0 = 0$ and translation group $G = \mathbb{R}^d$ with $g_u = u$ and group action $\pi[g](u) = u + g$ recovers the classical convolution operator

$$(\mathscr{A}x)(v) = \int_{\mathbb{R}^d} a(v-u) \ x(u) \ \mathrm{d}u$$

with convolution kernel $a \in \mathcal{D}'(\mathbb{R}^d)$.

For the construction of equivariant neural networks, Theorem 3.2 characterizes the structure of all possible linear layers. In this context, a special case of this theorem for discrete domain transforms is well-known in the literature and has been proven, e.g., in [22]. In [6], the authors embed a variant of this theorem into a more general and group theoretical framework of Mackey functions for equivariant neural networks, where their notion of transforms on vector fields can be viewed as a realization of our concept of generalized domain transforms. However, in contrast to this, we provide a rigorous proof in the context of function spaces, where we generalize to distributional kernels, which is more suitable for forward operators in inverse problems. In addition, compared to existing settings, we do not restrict the group transforms in the output space \mathcal{Y} in any way. This enables us to treat symmetries of all measurement operators (that satisfy the visibility condition) for known transforms in the source space \mathcal{X} . Example 3.4 below will illustrate the consequences of this theorem for the special case of the classical Radon transform of functions on \mathbb{R}^2 .

In Section 4, we will construct equivariant linear layers that are suited for the application to indirect measurements and characterized by our above theorem, which is formulated in the continuous setting.

Proof of Theorem 3.2. In the following, for the sake of brevity, we set

$$\omega[g](u) = |\det(\mathrm{D}\pi_{\mathcal{X}}[g](u))|$$

for $g \in G$ and $u \in \Omega_{\mathcal{X}}$.

For the proof of the 'if'-direction let $a \in \mathcal{D}'(\Omega_{\mathcal{X}})$ satisfy (3.3). Then, we can choose any representative $g_u \in g_u N_{\pi_{\mathcal{X}}}(u_0)$ and define $\mathscr{A} : \mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{Y}})$ via (3.2). To verify (3.1), we compute the image of a transformed $x \in \mathcal{D}(\Omega_{\mathcal{X}})$ under \mathscr{A} by a change of variables in the distribution and use the fact that $\omega[gh](u) = \omega[g](\pi_{\mathcal{X}}[h](u)) \omega[h](u)$. Then, direct calculations indeed show that

$$(a *_{x \to y} \mathscr{P}_{\mathcal{X}}[g]x)(v) = \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{Y}}[g_u] a)(v) \ (\mathscr{P}_{\mathcal{X}}[g_u^{-1}g]x)(u_0) \ \omega[g_u](u_0)^{-1} \, \mathrm{d}u$$
$$= \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{Y}}[gg_u] a)(v) \ (\mathscr{P}_{\mathcal{X}}[g_u^{-1}]x)(u_0) \ \omega[gg_u](u_0)^{-1} \ \omega[g](u) \, \mathrm{d}u$$
$$= \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{Y}}[g] \mathscr{P}_{\mathcal{Y}}[g_u] a)(v) \ (\mathscr{P}_{\mathcal{X}}[g_u^{-1}]x)(u_0) \ \omega[g_u](u_0)^{-1} \, \mathrm{d}u$$
$$= (\mathscr{P}_{\mathcal{Y}}[g](a *_{x \to y} x))(v).$$

For the proof of the 'only if'-direction let $\mathscr{A} : \mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{Y}})$ satisfy the equivariance condition (3.1). Due to the continuity of \mathscr{A} and the Schwartz kernel theorem [18], we have the representation

$$(\mathscr{A}x)(v) = \int_{\Omega_{\mathcal{X}}} A(u,v) x(u) \,\mathrm{d}u,$$

where $A(\cdot, v) \in \mathcal{D}'(\Omega_{\mathcal{X}})$ is a distribution for every $v \in \Omega_{\mathcal{Y}}$. Let $(\delta_k)_{k \in \mathbb{N}}$ be a Dirac sequence of

smooth functions $\delta_k \in \mathcal{D}(\Omega_{\mathcal{X}})$ and, for $k \in \mathbb{N}$, consider the approximate kernel

$$A_k(w,v) = \int_{\Omega_{\mathcal{X}}} A(u,v) \; (\mathscr{P}_{\mathcal{X}}[g_w] \,\delta_k)(u) \; p_{\mathcal{X}}[g_w^{-1}](u_0) \; \omega[g_w](u_0)^{-1} \, \mathrm{d}u,$$

which satisfies $A_k(\cdot, v) \in \mathcal{D}(\Omega_{\mathcal{X}})$ for all $v \in \Omega_{\mathcal{Y}}$ as well as $A_k(u, \cdot) \in \mathcal{D}(\Omega_{\mathcal{Y}})$ for all $u \in \Omega_{\mathcal{X}}$. The induced operator $\mathscr{A}_k : \mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{Y}})$ can be rewritten as

$$(\mathscr{A}_k x)(v) = \int_{\Omega_{\mathcal{X}}} A_k(w, v) \, x(w) \, \mathrm{d}w$$

= $\int_{\Omega_{\mathcal{X}}} \int_{\Omega_{\mathcal{X}}} A(u, v) \, (\mathscr{P}_{\mathcal{X}}[g_w] \, \delta_k)(u) \, p_{\mathcal{X}}[g_w^{-1}](u_0) \, \omega[g_w](u_0)^{-1} \, \mathrm{d}u \, x(w) \, \mathrm{d}w$
= $\int_{\Omega_{\mathcal{X}}} A(u, v) \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{X}}[g_w] \, \delta_k)(u) \, p_{\mathcal{X}}[g_w^{-1}](u_0) \, x(w) \, \omega[g_w](u_0)^{-1} \, \mathrm{d}w \, \mathrm{d}u,$

where we use Fubini's theorem for distributions [30, Theorem 40.4]. By considering

$$(\delta_k *_{\mathcal{X} \to \mathcal{X}} x)(u) = \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{X}}[g_w] \,\delta_k)(u) \ p_{\mathcal{X}}[g_w^{-1}](u_0) \ x(w) \,\omega[g_w](u_0)^{-1} \,\mathrm{d}w,$$

we obtain the representation

$$\mathscr{A}_k x = \mathscr{A}(\delta_k *_{\chi \to \chi} x).$$

Standard computations show the convergence $\delta_k *_{\chi \to \chi} x \to x$ in $\mathcal{D}(\Omega_{\chi})$ and by continuity of \mathscr{A} follows that

$$\mathscr{A}_k x \xrightarrow{k \to \infty} \mathscr{A} x.$$

Using the equivariance property (3.1) of \mathscr{A} and the equivariance of $*_{\mathcal{X}\to\mathcal{X}}$ that we have already shown in the first part of this proof, we also have

$$\mathscr{A}_{k}\mathscr{P}_{\mathcal{X}}[g]x = \mathscr{A}\left(\delta_{k} \ast_{x \to x} \mathscr{P}_{\mathcal{X}}[g]x\right) = \mathscr{A}\mathscr{P}_{\mathcal{X}}[g]\left(\delta_{k} \ast_{x \to x} x\right) = \mathscr{P}_{\mathcal{Y}}[g]\mathscr{A}\left(\delta_{k} \ast_{x \to x} x\right) = \mathscr{P}_{\mathcal{Y}}[g]\mathscr{A}_{k}x$$

and a change of variables implies that

$$\int_{\Omega_{\mathcal{X}}} A_k(\pi_{\mathcal{X}}[g](u), v) \ p_{\mathcal{X}}[g](\pi_{\mathcal{X}}[g](u)) \ x(u) \ \omega[g](u) \ \mathrm{d}u = \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{Y}}[g]A_k)(u, v) \ x(u) \ \mathrm{d}u.$$

Since this holds true for all $x \in \mathcal{D}(\Omega_{\mathcal{X}})$, we obtain for all $g \in G$ the condition

$$A_k(u,v) = (\mathscr{P}_{\mathcal{Y}}[g]^{-1}A_k)(\pi_{\mathcal{X}}[g](u),v) \ p_{\mathcal{X}}[g](\pi_{\mathcal{X}}[g](u)) \ \omega[g](u).$$

By choosing $g = g_u^{-1}$ and defining $a_k := A_k(u_0, \cdot) \in \mathcal{D}(\Omega_{\mathcal{Y}})$, we arrive at

$$A_k(u,v) = (\mathscr{P}_{\mathcal{Y}}[g_u] a_k)(v) \ p_{\mathcal{X}}[g_u^{-1}](u_0) \ \omega[g_u](u_0)^{-1},$$

applying the inverse function rule in ω . Since this holds true for all choices $g_u \in g_u N_{\pi_{\mathcal{X}}}(u_0)$, the a_k have to fulfil

(3.4)
$$(\mathscr{P}_{\mathcal{Y}}[g_un]a_k)(v) = (\mathscr{P}_{\mathcal{Y}}[g_u]a_k)(v) \ p_{\mathcal{X}}[g_u^{-1}](u_0) \ p_{\mathcal{X}}[(g_un)^{-1}](u_0)^{-1} \ |\det(\mathrm{D}\pi_{\mathcal{X}}[n](u_0))| = (\mathscr{P}_{\mathcal{Y}}[g_u]a_k)(v) \ p_{\mathcal{X}}[n](u_0) \ |\det(\mathrm{D}\pi_{\mathcal{X}}[n](u_0))|.$$

for all $n \in N_{\pi_{\mathcal{X}}}(u_0)$. Finally, taking the limit $k \to \infty$ gives

$$(\mathscr{A}x)(v) = \lim_{k \to \infty} \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{Y}}[g_u] a_k)(v) \ p_{\mathcal{X}}[g_u](u)^{-1} \ x(u) \ \omega[g_u](u_0)^{-1} \ \mathrm{d}u$$
$$= \int_{\Omega_{\mathcal{X}}} (\mathscr{P}_{\mathcal{Y}}[g_u] a)(v) \ p_{\mathcal{X}}[g_u](u)^{-1} \ x(u) \ \omega[g_u](u_0)^{-1} \ \mathrm{d}u,$$

where $a \in \mathcal{D}'(\Omega_{\mathcal{Y}})$ is the limit of $(a_k)_{k \in \mathbb{N}}$ in $\mathcal{D}'(\Omega_{\mathcal{Y}})$ and $u \mapsto (\mathscr{P}_{\mathcal{Y}}[g_u] a)(v)$ has to be understood as a distribution in $\mathcal{D}'(\Omega_{\mathcal{X}})$ for each $v \in \Omega_{\mathcal{Y}}$. Similarly, we obtain the condition (3.3) for a by taking the limit $k \to \infty$ in (3.4) in distributional sense.

To close this section we now discuss the Radon transform as an illustrating example.

Example 3.4. For $\Omega_{\mathcal{X}} = \mathbb{R}^2$ and $\Omega_{\mathcal{Y}} = \mathbb{R} \times [0, 2\pi)$ the Radon operator $\mathscr{R} : \mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{Y}})$ models the measurement process in computerized tomography, cf. [25], and is defined as

$$(\mathscr{R}x)(r,\varphi) = \int_{\mathbb{R}^2} \delta(u^T \vec{\varphi} - r) \, x(u) \, \mathrm{d}u,$$

integrating over the line $\{u \in \mathbb{R}^2 \mid u^T \vec{\varphi} = r\}$ with normal direction $\vec{\varphi} = (\cos(\varphi), \sin(\varphi))^T \in \mathbb{S}^1$. We consider the roto-translation group $G = \operatorname{SE}(2) = \operatorname{T}(2) \rtimes \operatorname{SO}(2)$, which we parametrize by $(s, \gamma) \in \mathbb{R}^2 \times [0, 2\pi)$ with operation

$$(s_1, \gamma_1) \cdot (s_2, \gamma_2) = (s_1 + R(\gamma_1)s_2, \gamma_1 + \gamma_2),$$

where addition of angles is performed modulo 2π and $R(\gamma) \in SO(2)$ is the rotation matrix

$$R(\gamma) = \begin{pmatrix} \cos \gamma & -\sin \gamma \\ \sin \gamma & \cos \gamma \end{pmatrix}.$$

A group action $\pi_{\mathcal{X}}$ on $\Omega_{\mathcal{X}} = \mathbb{R}^2$ can be defined via

$$\pi_{\mathcal{X}}[(s,\gamma)](u) = R(\gamma)u + s$$

so that $\Omega_{\mathcal{X}}$ is a homogeneous space of $\pi_{\mathcal{X}}$ and fixing the origin $u_0 = 0 \in \Omega_{\mathcal{X}}$ we obtain

$$N_{\pi_{\mathcal{X}}}(u_0) = \{ (0, \gamma) \mid \gamma \in [0, 2\pi) \}.$$

Straightforward computations show that

$$(\mathscr{R}(x \circ \pi_{\mathcal{X}}[(s,\gamma)^{-1}]))(r,\varphi) = (\mathscr{R}x)(r-s^T\vec{\varphi},\varphi-\gamma)$$

and, hence, we define the group action $\pi_{\mathcal{Y}}$ on $\Omega_{\mathcal{Y}} = \mathbb{R} \times [0, 2\pi)$ via

$$\pi_{\mathcal{Y}}[(s,\gamma)](r,\varphi) = (r + s^T(\overrightarrow{\varphi + \gamma}), \varphi + \gamma).$$

With this, $\Omega_{\mathcal{Y}}$ is a homogeneous space of $\pi_{\mathcal{Y}}$. For $u \in \Omega_{\mathcal{X}}$ we choose $g_u = (u, 0) \in G$ and obtain $\pi_{\mathcal{Y}}[g_u^{-1}](r, \varphi) = (r - u^T \vec{\varphi}, \varphi)$ as well as $\det(\mathrm{D}\pi_{\mathcal{X}}[g_u](u_0)) = 1$. As a result, Theorem 3.2 implies that any bounded linear operator $\mathscr{A} : \mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{Y}})$ that transforms translations and rotations exactly like the Radon operator \mathscr{R} can be written as

$$(\mathscr{A}x)(r,\varphi) = \int_{\mathbb{R}^2} a(r - u^T \vec{\varphi}, \varphi) x(u) \, \mathrm{d}u$$

with $a \in \mathcal{D}'(\Omega_{\mathcal{Y}})$. Furthermore, using $(g_u n)^{-1} = (-R(-\gamma)u, -\gamma)$ for $n = (0, \gamma) \in N_{\pi_{\mathcal{X}}}(u_0)$, the condition on a reads

$$a(r - (R(-\gamma)u)^T(\overrightarrow{\varphi - \gamma}), \varphi - \gamma) = a(r - u^T \vec{\varphi}, \varphi) \iff a(r - u^T \vec{\varphi}, \varphi - \gamma) = a(r - u^T \vec{\varphi}, \varphi)$$

for arbitrary $\gamma \in [0, 2\pi)$ so that $a \equiv a(r)$ is independent of the angle $\varphi \in [0, 2\pi)$. Therefore, all possible operators are given by

$$(\mathscr{A}x)(r,\varphi) = \int_{\mathbb{R}^2} a(r - u^T \vec{\varphi}) x(u) \, \mathrm{d}u$$

for some distribution $a \in \mathcal{D}'(\mathbb{R})$ and we obtain the Radon operator itself for $a = \delta$.

To also include anisotropic scaling, we consider the group $G = \text{Aff}^+(2) = \text{T}(2) \rtimes \text{GL}^+(2)$ of all invertible affine transforms with positive determinant, for which we define the group operation via

$$(s_1, A_1) \cdot (s_2, A_2) = (s_1 + A_1 s_2, A_1 A_2)$$

and the group representation $\mathscr{P}_{\mathcal{X}}$ on \mathcal{X} via

$$(\mathscr{P}_{\mathcal{X}}[(s,A)]x)(u) = x(\pi_{\mathcal{X}}[(s,A)]^{-1}(u)).$$

with

$$\pi_{\mathcal{X}}[(s,A)](u) = Au + s.$$

Then, $\Omega_{\mathcal{X}}$ is a homogeneous space of $\pi_{\mathcal{X}}$ and fixing the origin $u_0 = 0 \in \Omega_{\mathcal{X}}$ we obtain

$$N_{\pi_{\mathcal{X}}}(u_0) = \{(0, A) \mid A \in \mathrm{GL}^+(2)\}.$$

Standard computations show that

$$(\mathscr{RP}_{\mathcal{X}}[(s,A)]x)(r,\varphi) = \int_{\mathbb{R}^2} \delta(u^T \vec{\varphi} - r) \ x(A^{-1}(u-s)) \ \mathrm{d}u = \frac{\det(A)}{\alpha(A,\varphi)} (\mathscr{R}x) \Big(\frac{r - s^T \vec{\varphi}}{\alpha(A,\varphi)}, \theta(A,\varphi)\Big)$$

with

$$\alpha(A, \varphi) = \|A^T \vec{\varphi}\|_2$$
 and $\theta(A, \varphi) = \arccos\left(\frac{(A^T \vec{\varphi})_1}{\alpha(A, \varphi)}\right)$

so that

$$\vec{\theta}(A,\varphi) = \frac{A^T \vec{\varphi}}{\|A^T \vec{\varphi}\|_2}$$

Hence, we define the group representation $\mathscr{P}_{\mathcal{Y}}$ on \mathcal{Y} via

$$(\mathscr{P}_{\mathcal{Y}}[(s,A)]y)(r,\varphi) = p_{\mathcal{Y}}[(s,A)](r,\varphi) \ y(\pi_{\mathcal{Y}}[(s,A)]^{-1}(r,\varphi))$$

where

$$p_{\mathcal{Y}}[(s,A)](r,\varphi) = \frac{\det(A)}{\alpha(A,\varphi)} \quad \text{and} \quad \pi_{\mathcal{Y}}[(s,A)](r,\varphi) = \left(\frac{r+s^T A^{-T} \vec{\varphi}}{\alpha(A^{-1},\varphi)}, \theta(A^{-1},\varphi)\right)$$

so that $\mathscr{P}_{\mathcal{Y}}$ is a generalized domain transform and

$$\pi_{\mathcal{Y}}[(s,A)]^{-1}(r,\varphi) = \Big(\frac{r-s^T\vec{\varphi}}{\alpha(A,\varphi)}, \theta(A,\varphi)\Big).$$

For illustration, Figure 2 shows the action of $\mathscr{P}_{\mathcal{Y}}$ on sinograms of handwritten digits. For $u \in \Omega_{\mathcal{X}}$ we choose $g_u = (u, I_2) \in G$ and obtain $p_{\mathcal{Y}}[g_u](r, \varphi) = 1$, $\pi_{\mathcal{Y}}[g_u]^{-1}(r, \varphi) = (r - u^T \vec{\varphi}, \varphi)$ as well as $\det(\mathrm{D}\pi_{\mathcal{X}}[g_u](u_0)) = 1$. Consequently, Theorem 3.2 implies again that any bounded linear operator $\mathscr{A} : \mathcal{D}(\Omega_{\mathcal{X}}) \to \mathcal{D}(\Omega_{\mathcal{Y}})$ that transforms invertible affine transforms with positive determinant exactly like the Radon operator \mathscr{R} can be written as

$$(\mathscr{A}x)(r,\varphi) = \int_{\mathbb{R}^2} a(r - u^T \vec{\varphi}, \varphi) x(u) \, \mathrm{d}u$$

with $a \in \mathcal{D}'(\Omega_{\mathcal{Y}})$. Furthermore, using $(g_u n)^{-1} = (-A^{-1}u, A^{-1})$ for $n = (0, A) \in N_{\pi_{\mathcal{X}}}(u_0)$, the condition on a reads

$$\frac{\det(A)}{\|A^T\vec{\varphi}\|_2} a\Big(\frac{r-u^T\vec{\varphi}}{\|A^T\vec{\varphi}\|_2}, \theta(A,\varphi)\Big) = a(r-u^T\vec{\varphi},\varphi) \det(A)$$
$$\iff a\Big(\frac{r-u^T\vec{\varphi}}{\|A^T\vec{\varphi}\|_2}, \theta(A,\varphi)\Big) = \|A^T\vec{\varphi}\|_2 a(r-u^T\vec{\varphi},\varphi)$$

for arbitrary $A \in \mathrm{GL}^+(2)$ so that $a \equiv a(r)$ is independent of the angle $\varphi \in [0, 2\pi)$. Moreover, $a \in \mathcal{D}'(\mathbb{R})$ is positively homogeneous of degree -1 and, according to [16, Chap. 1, § 3.11], there are constants $c_1, c_2 \in \mathbb{R}$ such that

$$a = c_1 \, \chi^{-1} + c_2 \, \delta_z$$

where

$$\chi^{-1}(\psi) = \int_{\mathbb{R}} \frac{\psi(r) - \psi(0)}{r} \, \mathrm{d}r \quad \text{for } \psi \in \mathcal{D}(\mathbb{R}).$$

In summary, the symmetries of the Radon operator characterize most of its behaviour and leave only very few degrees of freedom for operators that share the same symmetry properties. This might become useful, e.g., in the construction of end-to-end reconstruction schemes, which could benefit from an encoded reverse conversion of symmetries. **4. Equivariant Layers.** Now that we have characterized the group representation $\mathscr{P}_{\mathcal{Y}}$ on the measurements in \mathcal{Y} , we aim for the construction of equivariant neural networks $\mathscr{N}_{\theta} : \mathcal{Y} \to \mathcal{Z}$ such that

$$\mathscr{P}_{\mathcal{Z}}[g] \, \mathscr{N}_{\theta}(y) = \mathscr{N}_{\theta}(\mathscr{P}_{\mathcal{Y}}[g] \, y) \quad \forall \, y \in \mathcal{Y}, \ g \in G.$$

Here, the space \mathcal{Z} and representation $\mathscr{P}_{\mathcal{Z}}$ have to be chosen depending on the application. In reconstruction or semantic segmentation tasks, the natural choice is $\mathcal{Z} = \mathcal{X}$ and $\mathscr{P}_{\mathcal{Z}} = \mathscr{P}_{\mathcal{X}}$. Learned preprocessing could be obtained by $\mathcal{Z} = \mathcal{Y}$ and $\mathscr{P}_{\mathcal{Z}} = \mathscr{P}_{\mathcal{Y}}$. Invariant classification tasks require $\mathscr{P}_{\mathcal{Z}}[g] = \mathrm{Id}$ independently of $g \in G$. Based on the findings in [7], it is however useful to construct invariance only in the last layer of the network and use more complex group representations in the inner hidden part.

In the continuous setting studied previously, the construction of equivariant linear neural network layers for generalized domain transforms is completely characterized by Theorem 3.2. In practical applications, however, we do not only need a discretization for numerical purposes but also have to deal with a finite and possibly very limited number of measurement points. We therefore model the whole indirect measurement process as $y_V^{\delta} = S_V \mathscr{A} x + \eta^{\delta}$, where η^{δ} models noise in the measurements of noise level δ and S_V defines a sampling scheme

$$S_V: \mathcal{Y} \to \mathcal{Y}_V$$

evaluating a continuous signal $y \in \mathcal{Y}$ at discrete measurement points $V = \{v_i\}_{i=1,\dots,n} \subset \Omega_{\mathcal{Y}}$.

In this section, we will construct neural networks that are (approximately) equivariant to generalized domain transforms as representations in the input and output of all layers and that can be applied directly to arbitrary data generated by indirect measurements. Figure 3 gives an overview of the setting and the involved spaces.

4.1. The Challenge of Discrete Equivariance. When trying to treat equivariance in discrete measurements, the first problem that occurs is the definition of appropriate group representations on the discrete data. Since the data is only partially captured, the discretized operators violate the visibility condition in Theorem 3.1. As a result, it is not clear how general group transforms in the input can be handled in the measurements. In addition, also classical reconstruction methods applied to sparse measurements cannot be exactly group equivariant, which limits the performance of all methods that are based on reconstruction schemes.

A possible way out that can also be implemented numerically is to be satisfied with equivariance only with respect to a discrete subgroup $G_V \subset G$, which would result in an approach similar to [22] and yields a direct generalization of classical CNNs. However, this requires the measurements to lie on a grid that fits the structure of the subgroup. More precisely, to handle the measurement points V it is necessary that we have the set equality

$$\pi_{\mathcal{V}}[G_V](v_0) = V$$

for a suitably chosen $v_0 \in \Omega_{\mathcal{Y}}$. In general, such a subgroup G_V does not exist and, hence, this idea can only be applied for very specific sampling schemes S_V . A numerically more complex but also more general solution is to construct networks that are *approximately* equivariant



Figure 3. Overview of our approach, including the measurement scheme as well as the learned processing by an equivariant neural network. The plots illustrate the modelling and symmetries of a two-angle fan-beam Radon operator on a rotated and translated elliptical ring.

to the representation $\mathscr{P}_{\mathcal{Y}}$ on the whole Lie group G and can operate on arbitrary point clouds, extending the work in [14]. Following this approach, only the continuous measurement operator \mathscr{A} has to be equivariant, i.e., fulfil the visibility condition in Theorem 3.1. For many applications (e.g. sparse CT or image inpainting), this can be achieved by a proper modelling of \mathscr{A} and S_V such that the definition of \mathscr{A} reveals the underlying symmetries that would appear in the measurement data analytically if one was able to measure the complete signal. However, discrete measurements also induce limitations which cannot be overcome by any advanced approach (without additional data knowledge): the equivariance of the resulting methods in the limit and its improvement for an increasing number of measurement points is constrained by the uniformity of the sensor locations. Nevertheless, we expect the resulting networks to provide an efficient architecture also in case of non-uniform sparse measurements.

4.2. Construction of Equivariant Networks. Neural networks are typically built based on linear layers and the localized application of non-linear activation functions. To construct equivariant networks, our approach is to concatenate equivariant layers $\mathcal{N}^{\ell} : \mathcal{Z}_{\ell-1} \to \mathcal{Z}_{\ell}$ and set

$$\mathscr{N}_{\theta} = \mathscr{N}^{N_L} \circ \cdots \circ \mathscr{N}^1$$

where $\mathcal{Z}_0 = \mathcal{Y}$ and $\mathcal{Z}_L = \mathcal{Z}$. As a consequence of Theorem 3.2, for given input and output generalized domain transforms, the structure of an equivariant linear layer is known to be convolutional and the layer can be fully parametrized by the (distributional) kernel $a \equiv a_{\theta}$.

For the implementation of such linear layers we have to approximate the integral in (3.2), where we solely consider generalized domain transforms for the input and output of the layers. Such transforms, however, cover the representations in \mathcal{Y} or \mathcal{Z} that may appear in our desired indirect measurement applications. As usual for convolutional networks, we use localized convolution kernels with compact support around the origin. We approximate the integral for linear layer $\mathcal{N}^{\ell+1}$ by

$$(\mathscr{N}^{\ell+1}z)(w'_{j}) = \int_{\Omega_{\mathcal{Z}_{\ell}}} (\mathscr{P}_{\mathcal{Z}_{\ell+1}}[g_{w}] a_{\theta})(w'_{j}) \ p_{\mathcal{Z}_{\ell}}[g_{w}^{-1}](w_{0}) \ z(w) \ |\det(\mathrm{D}\pi_{\mathcal{Z}_{\ell}}[g_{w}](w_{0}))|^{-1} \ \mathrm{d}w$$
$$\approx \sum_{i=1}^{M_{\ell}} c_{i} \ (\mathscr{P}_{\mathcal{Z}_{\ell+1}}[g_{w_{i}}] a_{\theta})(w'_{j}) \ p_{\mathcal{Z}_{\ell}}[g_{w_{i}}^{-1}](w_{0}) \ z(w_{i}) \ |\det(\mathrm{D}\pi_{\mathcal{Z}_{\ell}}[g_{w_{i}}](w_{0}))|^{-1},$$

for the parameterized continuous kernel a_{θ} and collocation points $W_{\ell} = \{w_i\}_{i=1,\ldots,M_{\ell}} \subset \Omega_{\mathcal{Z}_{\ell}}$, where $z \in \mathcal{Z}_{\ell}, w'_i \in \Omega_{\mathcal{Z}_{\ell+1}}$ and $C_{\ell} = \{c_i\}_{i=1,\ldots,M_{\ell}} \subset \mathbb{K}$ denote quadrature weights.

In the inner layers of the network, the w_i are sampled uniformly, such that we perform Monte-Carlo summation with constant $c_i = |\Omega_{\mathcal{Z}_\ell}|/M_\ell$. In the first layer, the input is known only at the measurement locations and, therefore, we have to choose $W_0 = V$ and perform numerical integration with fixed quadrature points. For this purpose, we build on the recent work in [12], where the authors report successful usage of learned quadrature weights, i.e., the c_i become parameters during the training of the network. Even though the approximation of the integrals may induce numerical errors, we expect that the resulting convolutional layer structure forms a good inductive bias in all layers while this scheme enables us to handle arbitrary measurement geometries. The inputs then consist of pairs $(y_i, v_i)_{v_i \in V}$, where y_i corresponds to the measured value at node v_i (e.g., $v_i = (r_i, \varphi_i)$ in the case of a sinogram).

The constraint (3.3) for the kernel along the orbits of the stabilizer can however be quite restrictive and limit the expressive power of the layers. For the symmetries that appear in indirect measurements, the limitations are especially severe, since these orbits can become unbounded, preventing localized kernels. To overcome this, one has to properly choose the group representations used throughout the neural network. In most existing discrete [22] and continuous [14] approaches, this is tackled by choosing $\mathcal{Z}_{\ell} = \mathcal{D}(G, \mathbb{R}^{c_{\ell}}), \ \ell = 1, \ldots, L-1$, and the canonical domain transform $\mathscr{P}_G = \mathscr{P}(\pi_G), \ \pi_G : G \to G, \ g \mapsto g$, on the group itself as the input and output representation of inner layers. As a homogeneous space of itself, G has a trivial stabilizer and the constraint (3.3) vanishes.

In the first layer, $\mathscr{N}^1: \mathscr{Y} \to \mathcal{D}(G, \mathbb{R}^{c_1})$, Theorem 3.2 gives

$$(\mathscr{N}^1 y)(g) = \int_{\Omega_{\mathcal{Y}}} a_{\theta}(g_v^{-1}g) \ p_{\mathcal{Y}}[g_v^{-1}](v_0) \ y(v) \ |\det(\mathrm{D}\pi_{\mathcal{Y}}[g_v](v_0))|^{-1} \ \mathrm{d}v,$$

where the kernel $a_{\theta} \in \mathcal{D}'(G)$ has to fulfil

 $a_{\theta}(n^{-1}g) = a_{\theta}(g) \ p_{\mathcal{Y}}[n](v_0) \ |\det(\mathrm{D}\pi_{\mathcal{Y}}[n](v_0))|.$

This can be realized by setting

$$a_{\theta}(g) = \tilde{a}_{\theta}(\pi_{\mathcal{Y}}[g^{-1}](v_0)) \ p_{\mathcal{Y}}[g](v_0)^{-1} \ |\det(\mathrm{D}\pi_{\mathcal{Y}}[g^{-1}](v_0))|$$



Figure 4. Classification accuracy for fan-beam CT measurements of randomly rotated or $GL^+(2)$ -transformed MNIST test dataset for different neural networks that are trained exclusively on upright digits.

for a kernel $\tilde{a}_{\theta} \in \mathcal{D}'(\Omega_{\mathcal{Y}})$ so that

$$(\mathscr{N}^1 y)(g) = \int_{\Omega_{\mathcal{Y}}} \tilde{a}_{\theta}(\pi_{\mathcal{Y}}[g^{-1}](v)) \ p_{\mathcal{Y}}[(g^{-1})](v_0)) \ y(v) \ |\det(\mathrm{D}\pi_{\mathcal{Y}}[g](v_0))|^{-1} \ \mathrm{d}v$$

As usual for neural networks, we construct non-linear layers using local non-linearities that act only on the values of activations and are independent of the spatial position. Since, in most cases, activation functions are applied only in the inner parts of the network (where all representations $\mathscr{P}_{\mathcal{Z}_{\ell}}$ are given by the canonical domain transform on G), we do not have to consider additional constraints on these functions.

5. Numerical Experiments. To obtain first insights on the effectiveness of our idea in practice, we implemented the previously defined layers and performed a series of experiments in the setting of classification and regression tasks based on (sparse) CT measurements, where we require equivariance/invariance from \mathcal{X} to \mathcal{Z} , while measurements are taken in \mathcal{Y} .

5.1. Implementation. We implemented a flexible PyTorch [27] framework that allows us to compute the generalized convolution in (4.2) for different input and output representations. Our code is publicly available on GitHub¹.

In each layer, we use a small fully-connected neural network with two layers, batch normalization [19] and Swish activations [28] to parameterize a basis of the convolution kernels a_{θ} . The kernels connecting all input and output channels are then additionally parameterized by individual linear factors. In this setting, we use the automatic optimization of Einstein summation in PyTorch to perform the efficient PointConv trick [14]. Where the constraint (3.3) allows, we enforce locality of the kernels a_{θ} respectively \tilde{a}_{θ} by an additional rapidly decaying factor, i.e., $a_{\theta}(w'_j) = \exp(-d(w'_j, w'_0)^2/r^2) \hat{a}_{\theta}(w'_j)$, where $d : \Omega_{\mathcal{Z}_{\ell}} \times \Omega_{\mathcal{Z}_{\ell}} \to \mathbb{R}_{\geq 0}$ is a metric that is invariant with respect to $\pi_{\mathcal{Z}_{\ell}}$. In addition, we compute the resulting sum only over a set of k = 27 nearest neighbours to minimize the computational effort.

 $^{^{1}} https://github.com/nheilenkoetter/equivariant-nn-for-indirect-measurements$



Figure 5. Classification accuracy for fan-beam CT projections of the RotMNIST dataset for different numbers of measurement angles when trained on RotMNIST and tested on RotMNIST or LinMNIST.

We use a residual network architecture [17] that is built of three residual blocks connected by downsampling layers (increasing the field of view and decreasing the amount of sampled points). In each block, we use three convolutional layers combined with batch normalization and ReLU activations [15]. We choose 22 channels in the initial layer and double the amount of channels in each downsampling step, while we halve the amount of sampled points. We obtain invariance by adding a global average pooling in the last layer, leaving us only with the constructed channel dimension.

In every iteration of the optimization, we sample a new set of inner collocation points. However, to reduce the computational effort, we sample only once for each batch and each residual block and use the same points along all layers with identical scale.

In all of our experiments on image data, we simulate fan-beam Radon measurements using the ASTRA [31, 32] toolbox. For optimization, we use Adam [21] with weight decay and a learning rate scheduling scheme.

5.2. Digit Classification. As an exemplary task on indirect measurements, we perform experiments on the classification of sparse Radon measurements of digits from the popular MNIST [10] dataset and its variants, cf. Figure 2. We study invariance with respect to the roto-translation group SE(2) and the group Aff⁺(2) of affine transforms with positive determinant we considered previously for the Radon operator in Example 3.4. To this end, we uniformly sample a rotation angle in $[0, 2\pi)$, scaling factors for each direction in [0.75, 1.25] and a horizontal shear factor in [-0.5, 0.5]. For the discretization of the integrals in the convolutions we always sample 2000 points in the first residual block.

5.2.1. Invariance. To verify our implementation as well as discretization approach, we perform tests on the actual numerical invariance of the constructed networks. Let us again stress that we can only expect *approximate* invariance due to the built-in approximations and sampling scheme. For this purpose, we train roto-translation equivariant neural networks on varying geometries of sparse-angle fan-beam Radon measurements of the standard MNIST dataset (containing only upright digits) and evaluate the models on transformed versions



(a) Samples from the LinMNIST dataset, which includes rotations, anisotropic scaling and shearing of MNIST digits.



(b) Geometry of the simulated tube measurements, where d_{max} and d_{min} denote the maximal and minimal thickness of the tube.

Figure 6. Visualizations of our generated datasets for the digit and ellipse experiments.

of the test dataset, where we consider arbitrary rotations and transforms from $GL^+(2)$. The transformed signals are computed using nearest-neighbour interpolation on the source images, before simulating the sparse measurements. We train for 90 epochs with a batch size of 14.

The results for a varying amount of measurement angles are depicted in Figure 4, where the angles are chosen equidistantly in $[0, 2\pi)$, except for two angles, where we use 0 and $\frac{\pi}{2}$. While the classification accuracy on the non-rotated test data is similar for all geometries, the constructed invariance is better for increasing number of measurement angles. Nevertheless, in all cases, we notice some extrapolation ability to unseen and transformed data, which underpins the intuition that the constructed network architecture serves as a good inductive bias even if exact equivariance cannot be guaranteed. Surprisingly, the roto-translation equivariant networks perform similarly to the networks that were built to be equivariant with respect to the larger group Aff⁺(2) even in terms of generalization to unseen affinely transformed data. This could indicate that the latter method requires additional fine-tuning.

5.2.2. RotMNIST. For further investigation we evaluate the classification accuracy of our approach on noisy indirect measurements of the rotated MNIST benchmark dataset, cf. [23]. We consider the same measurement geometries and network architectures as before and add 6% Gaussian noise to the simulated Radon data. For comparison, we also train a standard discrete T(2)-equivariant CNN on the sparse sinograms and compute classical reconstructions for post-processing using the ASTRA toolbox implementation of the SIRT reconstruction scheme. An example for the obtained reconstructions in the 2-angle case is depicted in Figure 2. On these reconstructions, we train a discrete group-equivariant (with respect to discrete translation and rotations in $\pi/2$ -steps) ResNet18 [17] based on the layer structure that was presented in [7] as well as the continuous Lie-group-equivariant network implemented in LieConv [14] for their experiments on RotMNIST. All models are trained for





(a) Classification accuracy of SE(2)-equivariant networks for randomly rotated and $GL^+(2)$ transformed MNIST test measurements when exclusively trained on upright digits.

(b) Classification accuracy for RotMNIST test measurements for varying numbers of angles and different classification approaches.

Figure 7. Invariance and classification accuracy for the case of limited-angle fan-beam CT geometries. While our network decreases in equivariance if the sensor locations do not cover the whole measurement space, it still provides an efficient bias when compared to other methods.

400 epochs with a batch size of 14. We evaluate the models on the RotMNIST test data as well as on a $GL^+(2)$ -transformed version of it, which we refer to as LinMNIST. The results are visualized in Figure 5, while samples from the LinMNIST dataset are depicted in Figure 6(a).

In all experiments, our approach outperforms the reconstruction-based approaches as well as the standard CNN, which does not encode proper symmetries in measurement space, in terms of generalization ability. While the discrete reconstruction-based network shows better convergence on the training data, our method consistently obtains better results on the test dataset, especially for sparse geometries where the reconstructions are very inexact and amplify the noise. These results show that we were able to construct a valuable prior for indirect measurements and that it can be advantageous to avoid reconstructions in certain cases.

5.2.3. Limited-angle CT. To test the performance of our methods in applications where the sensor locations do not fully cover the measurement space, we additionally run experiments in the context of limited-angle CT. To this end, we also use a fan-beam geometry with a varying number of equidistant measurement angles, but choose them only in the range $[0, \frac{\pi}{2}]$. In this case, we do not expect good equivariance properties for any method due to the violation of the visibility condition even in the limit case of infinitely many data. If a digit has been measured only from a certain view angle during training, it becomes impossible to classify rotated version of it without assuming additional data knowledge.

The results using our proposed SE(2)-equivariant neural network architecture are depicted in Figure 7, where, in a first experiment, we train our roto-translation equivariant neural networks on limited-angle fan-beam Radon measurements of standard MNIST digits and evaluate the resulting models on rotated and $GL^+(2)$ -transformed versions of the test dataset, and, in a second experiment, we train and evaluate various models on RotMNIST measurements.

_			
	ы		
l a	U	e.	н.
			_

Regression results on the generated ellipse dataset for varying amount of training data, where MSE refers to the mean squared error for both the minimal and maximal thickness. The methods 'GrouPy' and 'LieConv' act on previously computed classical reconstructions, while our methods directly handles the indirect measurement.

	$MSE \times 10^{-3}$		error at $d_{\min} \times 10^{-2}$			error at $d_{\max} \times 10^{-2}$			
dataset size	ours	GrouPy	LieConv	ours	Groupy	LieConv	ours	GrouPy	LieConv
1000	0.93	5.59	1.51	1.95	2.80	2.72	2.49	3.60	3.23
2000	0.81	0.92	1.56	1.80	2.01	2.68	2.35	2.40	3.30
4000	0.70	0.85	1.51	1.67	1.97	2.88	2.18	2.35	3.13
8000	0.67	0.51	1.44	1.63	1.50	2.65	2.17	1.82	3.15

In summary, we observe that especially for sparse angles the data efficiency of our proposed method still holds, while its equivariance cannot be guaranteed and does not improve with an increasing number of measurement angles in the limited range $[0, \frac{\pi}{2}]$.

5.3. Thickness Regression of Tubes. In addition to the above classification task, we also test a regression setting that is motivated by our initial example: the challenge of estimating the minimal and maximal thickness of a tube based on fan-beam data measured from only two angles, namely 0° and 85°. To simulate this task, we compute a dataset of elliptical rings, constructed from two ellipses with identical center point but randomly different rotation angles and radii. We add 5% Gaussian noise to the analytically computed sinogram, which is closer to real-world scenarios avoiding the simulation of forward measurements from given discrete data. As ground truth, we approximate the minimal and maximal thickness of the tube by computing the length of the intersection with a line going through the center point of the ellipses. For comparison, we compute reconstructed images of 129×129 pixels using the SIRT scheme and compare to group-equivariant models that are trained on these reconstructions. To evaluate the generalization performance, we evaluate different training dataset sizes, where all models are trained for 3000 epochs on a batch size of 8. In the initial residual block of our model, which is constructed to be equivariant to the induced representations of SE(2), we sample 2700 points for the estimation of the integrals.

Table 1 lists the obtained test errors. We observe that, for dataset sizes below 8000 training samples, our method outperforms both reconstruction-based approaches on the unseen data. These results give further indication towards improved generalization properties and show the data-efficiency of our approach.

6. Conclusion. We have introduced neural networks that incorporate symmetries present in indirect measurements and theoretically investigated the connection between operators and symmetries. Our key Theorem 3.2 precisely characterizes symmetries induced by given operators and, vice versa, determine all operators that share the same symmetries. Based on these insights we introduced data-efficient neural networks that are tailored to indirect measurements and can handle sparse inverse problems. In particular, building upon our definition of generalized domain transforms, we were able to consider a class of symmetries that has not been treated before and shows new characteristics. The effectivity of our approach is demonstrated on a classification and a regression task. We want to stress that our work is meant to serve as a first attempt at the treatment of equivariance on indirectly measured data. Several open questions and possible extensions arise both on the theoretical and the numerical side. One direction could be the application and analysis of our approach in end-to-end reconstructions for sparse measurements. In this setting, Theorem 3.2 indicates that a large part of the reconstruction operator is already encoded in the equivariance of our network. Another important aspect is the incorporation of uncertainties and inexactness in the forward operator or its action on underlying symmetries in the source space. In this context, a Bayesian setting could help to also treat symmetries in noisy scenarios in a rigorous way, which is beyond the scope of this work. Finally, an alternative to our approach of building equivariant network architectures on the whole measurement space \mathcal{Y} could be the construction of networks \mathcal{N}_{θ} for discrete measurements in \mathcal{Y}_V such that the whole process $\mathcal{N}_{\theta} \circ S_V \circ \mathscr{A} : \mathcal{X} \to \mathcal{Z}$ is approximately equivariant. We expect that this can be realized by training with augmented loss functions along with additional assumptions on the data distribution, which points to an interesting direction for future research.

REFERENCES

- S. ARRIDGE, P. MAASS, O. OKTEM, AND C.-B. SCHÖNLIEB, Solving inverse problems using data-driven models, Acta Numerica, 28 (2019), pp. 1–174, https://doi.org/10.1017/S0962492919000059.
- [2] E. CELLEDONI, M. J. EHRHARDT, C. ETMANN, B. OWREN, C.-B. SCHÖNLIEB, AND F. SHERRY, Equivariant neural networks for inverse problems, Inverse Problems, 37 (2021), p. 085006, https: //doi.org/10.1088/1361-6420/ac104f.
- [3] D. CHEN, M. E. DAVIES, M. J. EHRHARDT, C.-B. SCHÖNLIEB, F. SHERRY, AND J. TACHELLA, Imaging with equivariant deep learning: From unrolled network design to fully unsupervised learning, IEEE Signal Processing Magazine, 40 (2023), pp. 134–147, https://doi.org/10.1109/MSP.2022.3205430.
- [4] D. CHEN, J. TACHELLA, AND M. E. DAVIES, Equivariant imaging: Learning beyond the range space, in IEEE/CVF International Conference on Computer Vision (ICCV), 2021, pp. 4359–4368, https: //doi.org/10.1109/ICCV48922.2021.00434.
- N. CHENOUARD AND M. UNSER, 3D steerable wavelets in practice, IEEE Transactions on Image Processing, 21 (2012), pp. 4522–4533, https://doi.org/10.1109/TIP.2012.2206044.
- [6] T. S. COHEN, M. GEIGER, AND M. WEILER, A general theory of equivariant CNNs on homogeneous spaces, in Advances in Neural Information Processing Systems 32 (NeurIPS 2019), 2019.
- [7] T. S. COHEN AND M. WELLING, Group equivariant convolutional networks, in Proceedings of the 33th International Conference on Machine Learning, 2016, pp. 2990–2999.
- [8] T. S. COHEN AND M. WELLING, Steerable CNNs, in International Conference on Learning Representations, 2017.
- [9] M. DAX, S. R. GREEN, J. GAIR, M. DEISTLER, B. SCHÖLKOPF, AND J. H. MACKE, Group equivariant neural posterior estimation, in 10th International Conference on Learning Representations (ICLR), 2022, pp. 1–24.
- [10] L. DENG, The MNIST database of handwritten digit images for machine learning research, IEEE Signal Processing Magazine, 29 (2012), pp. 141–142.
- S. DITTMER, D. ERZMANN, H. HARMS, AND P. MAASS, SELTO: Sample-efficient learned topology optimization, 2022, https://arxiv.org/abs/2209.05098.
- [12] K. DOHERTY, C. SIMPSON, S. BECKER, AND A. DOOSTAN, QuadConv: Quadrature-based convolutions with applications to non-uniform PDE data compression, 2022, https://arxiv.org/abs/2211.05151.
- [13] H. W. ENGL, M. HANKE, AND A. NEUBAUER, Regularization of Inverse Problems, no. 375 in Mathematics and Its Applications, Kluwer Academic Publisher, 2000.
- [14] M. FINZI, S. STANTON, P. IZMAILOV, AND A. G. WILSON, Generalizing convolutional neural networks for equivariance to Lie groups on arbitrary continuous data, in Proceedings of the 37th International Conference on Machine Learning, 2020, pp. 3165–3176.

EQUIVARIANT NETWORKS FOR INDIRECT MEASUREMENTS

- [15] K. FUKUSHIMA, Cognitron: A self-organizing multilayered neural network, Biological Cybernetics, 20 (1975), pp. 121–136, https://doi.org/10.1007/BF00342633.
- [16] I. M. GEL'FAND AND G. E. SHILOV, Generalized Functions. Volume I: Properties and Operations, Academic Press, 1964.
- [17] K. HE, X. ZHANG, S. REN, AND J. SUN, Deep residual learning for image recognition, in IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 770–778.
- [18] L. HÖRMANDER, The Analysis of Linear Partial Differential Operators I: Distribution Theory and Fourier Analysis, no. 256 in Grundlehren der mathematischen Wissenschaften, Springer, 2 ed., 1990.
- [19] S. IOFFE AND C. SZEGEDY, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in Proceedings of the 32nd International Conference on Machine Learning, 2015, pp. 448–456.
- [20] J. P. KAIPIO AND E. SOMERSALO, Statistical and Computational Inverse Problems, vol. 160 of Applied Mathematical Sciences, Springer, 2005.
- [21] D. P. KINGMA AND J. BA, Adam: A method for stochastic optimization, 2014, https://arxiv.org/abs/ 1412.6980.
- [22] R. KONDOR AND S. TRIVEDI, On the generalization of equivariance and convolution in neural networks to the action of compact groups, in Proceedings of the 35th International Conference on Machine Learning, 2018, pp. 2747–2755.
- [23] H. LAROCHELLE, D. ERHAN, A. COURVILLE, J. BERGSTRA, AND Y. BENGIO, An empirical evaluation of deep architectures on problems with many factors of variation, in Proceedings of the 24th International Conference on Machine Learning, 2007, pp. 473–480, https://doi.org/10.1145/1273496.1273556.
- [24] Y. LECUN, L. D. JACKEL, L. BOTTOU, C. CORTES, J. S. DENKER, H. DRUCKER, I. GUYON, U. A. MÜLLER, E. SÄCKINGER, P. SIMARD, AND V. VAPNIK, *Learning algorithms for classification: A comparison on handwritten digit recognition*, in Neural Networks: The Statistical Mechanics Perspective, World Scientific, 1995, pp. 261–276.
- [25] F. NATTERER AND F. WÜBBELING, Mathematical methods in image reconstruction, SIAM Monographs on Mathematical Modeling and Computation, SIAM, 2001.
- [26] D. OTERO BAGUER, J. LEUSCHNER, AND M. SCHMIDT, Computed tomography reconstruction using deep image prior and learned reconstruction methods, Inverse Problems, 36 (2020), p. 094004, https://doi. org/10.1088/1361-6420/aba415.
- [27] A. PASZKE, S. GROSS, F. MASSA, A. LERER, J. BRADBURY, G. CHANAN, T. KILLEEN, Z. LIN, N. GIMELSHEIN, L. ANTIGA, A. DESMAISON, A. KOPF, E. YANG, Z. DEVITO, M. RAISON, A. TE-JANI, S. CHILAMKURTHY, B. STEINER, L. FANG, J. BAI, AND S. CHINTALA, *PyTorch: An imperative* style, high-performance deep learning library, in Advances in Neural Information Processing Systems 32 (NeurIPS 2019), 2019.
- [28] P. RAMACHANDRAN, B. ZOPH, AND Q. V. LE, Searching for activation functions, 2017, https://arxiv. org/abs/1710.05941.
- [29] J. TANG, Equivariance regularization for image reconstruction, 2022, https://arxiv.org/abs/2202.05062.
- [30] F. TREVES, Differentiable functions with values in topological vector spaces. Tensor product of distributions, in Topological Vector Spaces, Distributions and Kernels, vol. 25 of Pure and Applied Mathematics, Elsevier, 1967, ch. 40, pp. 411–419, https://doi.org/10.1016/S0079-8169(08)60289-5.
- [31] W. VAN AARLE, W. J. PALENSTIJN, J. CANT, E. JANSSENS, F. BLEICHRODT, A. DABRAVOLSKI, J. DE BEENHOUWER, K. J. BATENBURG, AND J. SIJBERS, Fast and flexible X-ray tomography using the ASTRA toolbox, Optics Express, 24 (2016), pp. 25129–25147, https://doi.org/10.1364/OE.24. 025129.
- [32] W. VAN AARLE, W. J. PALENSTIJN, J. DE BEENHOUWER, T. ALTANTZIS, S. BALS, K. J. BATENBURG, AND J. SIJBERS, The ASTRA toolbox: A platform for advanced algorithm development in electron tomography, Ultramicroscopy, 157 (2015), pp. 35–47, https://doi.org/10.1016/j.ultramic.2015.05.002.
- [33] M. WINKELS AND T. S. COHEN, Pulmonary nodule detection in CT scans with equivariant CNNs, Medical Image Analysis, 55 (2019), pp. 15–26, https://doi.org/10.1016/j.media.2019.03.010.