

# Understanding when spatial transformer networks do not support invariance, and what to do about it

Lukas Finnveden\*, Ylva Jansson\* and Tony Lindeberg

Computational Brain Science Lab, Division of Computational Science and Technology  
KTH Royal Institute of Technology, Stockholm, Sweden

**Abstract**—Spatial transformer networks (STNs) were designed to enable convolutional neural networks (CNNs) to learn invariance to image transformations. STNs were originally proposed to transform *CNN feature maps* as well as input images. This enables the use of more complex features when predicting transformation parameters. However, since STNs perform a *purely spatial transformation*, they do not, in the general case, have the ability to align the feature maps of a transformed image with those of its original. STNs are therefore unable to support invariance when transforming CNN feature maps. We present a simple proof for this and study the practical implications, showing that this inability is coupled with decreased classification accuracy. We therefore investigate alternative STN architectures that make use of complex features. We find that while deeper localization networks are difficult to train, localization networks that share parameters with the classification network remain stable as they grow deeper, which allows for higher classification accuracy on difficult datasets. Finally, we explore the interaction between localization network complexity and iterative image alignment.

## I. INTRODUCTION

Spatial transformer networks (STNs) [1] constitute a widely used end-to-end trainable solution for CNNs to learn invariance to image transformations. This makes it part of a growing body of work concerned with developing CNNs that are invariant or robust to image transformations. The key idea behind STNs is to introduce a trainable module – the spatial transformer (ST) – that applies a data dependent spatial transformation of input images or CNN feature maps before further processing. If such a module successfully learns to align images to a canonical pose, it can enable invariant recognition. However, when transforming *CNN feature maps*, such alignment is, in general, not possible, and STNs can therefore not enable invariant recognition. The reasons for this are that: (i) STs perform a purely spatial transformation, whereas transforming an image typically also results in a shift in the channel dimension of the feature activations (Figure 1), (ii) The shapes of the receptive fields of the individual neurons are not invariant (Figure 2). These problems have, to our knowledge, not been discussed in the STN literature. In the original paper and a number of subsequent works, STNs are presented as an option for achieving invariance also when applied to intermediate feature maps [1]–[5].

\*The first and second authors contributed equally to this work.

The support from the Swedish Research Council (contract 2018-03586) is gratefully acknowledged.




Fig. 1. A spatial transformation of a CNN feature map cannot, in general, align the feature maps of a transformed image with those of its original. Here, the network  $\Gamma$  has two feature channels "W" and "M", and  $T_g$  corresponds to a 180° rotation. Since different feature channels respond to the rotated image as compared to the original image, it is not possible to align the respective feature maps by applying the inverse spatial rotation to the feature maps. This implies that *spatially transforming feature maps* cannot enable invariant recognition by the means of aligning a set of feature maps to a common pose.

Our first contribution is to present a simple proof that STNs do not enable invariant recognition when transforming CNN feature maps. We do not claim mathematical novelty of this fact, which is in some sense intuitive and can be inferred from more general results (see e.g. [6]), but we present a simple alternative proof directly applicable to STNs. Our second contribution is to explore the practical implications of this result. Is there a point in transforming intermediate feature maps if this cannot enable invariant recognition? To investigate this, we compare different STN architectures on the MNIST [7], SVHN [8] and Plankton [9] datasets. We show that STNs that transform the feature maps are, indeed, worse at compensating for rotations and scaling transformations, while they – in accordance with theory – handle pure translations well. We also show that the inability of STNs to fully align CNN feature maps is coupled with decreased classification performance. Our third contribution is to explore alternative STN architectures that make use of more complex features when predicting image transformations. We show that: (i) Using more complex features can significantly improve performance, but most of this advantage is lost if transforming CNN feature maps, (ii) Sharing parameters between the localization network and the classification network makes training of deeper localization networks more stable and (iii) Iterative image alignment can be complementary to, but is no replacement for, using more complex features.

In summary, this work clarifies the role and functioning

of STNs that transform input images vs. CNN feature maps and illustrates important tradeoffs between different STN architectures that make use of deeper features.

#### A. Related work

Successful *image alignment* can considerably simplify a range of computer vision tasks by reducing variability related to differences in object pose. In classical computer vision, Lukas and Kanade [10] developed a methodology for image alignment, which estimates translations iteratively. This approach was later generalized to more general parameterized deformation models [11]. Affine shape adaptation of affine Gaussian kernels to local image structures – or equivalently, normalizing image structures to canonical affine invariant reference frames [12] – has been an integrated part in frameworks for invariant image-based matching and recognition [13]–[15]. Such classical approaches always align *the input images*.

Lately, the idea to combine structure and learning has given rise to the subfield of *invariant neural networks*, which add structural constraints to deep neural networks to enable e.g. scale or rotation invariant recognition [16]–[19]. *Spatial transformer networks* [1] are based on a similar idea of combining the knowledge about the structure of image transformations with learning. An STN, however, does not hard code invariance to any specific transformation group but learns input dependent image alignment from data. The original work [1] simultaneously claims *the ability to learn invariance from data* and that ST modules can be inserted at “*any depth*”. This seems to have left some confusion about whether STNs can enable invariance when transforming CNN feature maps. A number of subsequent works advocate to perform image alignment by transforming CNN feature maps [2]–[5] including e.g. pose alignment of pedestrians [5] and to use a spatial transformer to mimic the kind of patch normalization done in SIFT [2]. Additional works transform CNN feature maps without giving any specific motivation [20], [21]. As we will show, transforming CNN feature maps is *not equivalent* to extracting features from a transformed input image. Other approaches dealing with pose variability by transforming neural network *feature maps or filters* are e.g. spatial pyramid pooling [22] and dilated [23] or deformable convolutions [24]. Our results imply that these approaches have limited ability to enable e.g. scale invariance.

Weight sharing between the classification and the localization network has previously been considered in [25], primarily as a way of regularizing CNNs. Here, we are instead interested in it as a way to make use of deeper features when predicting image transformations. [26] combines STNs with iterative image alignment. In this paper, we investigate whether such iterative alignment is complimentary to using deeper features.

Previous theoretical work [6] characterizes all equivariant (covariant) maps between homogeneous spaces using the theory of fibers and fields. We, here, aim for a different perspective on the same theory. We present a simple proof for the special case of *purely spatial transformations* of CNN feature maps together with an *experimental evaluation* of

different STN architectures. A practical study [27] indicated that approximate alignment of CNN feature maps can be possible if allowing for a full transformation, as opposed to the purely spatial transformations that we analyze in this paper.

## II. THEORETICAL ANALYSIS OF INVARIANCE PROPERTIES

Spatial transformer networks [1] were introduced as an option for CNNs to learn invariance to image transformations by transforming *input images or convolutional feature maps* before further processing. A spatial transformer (ST) module is composed of a localization network that predicts transformation parameters and a transformer that transforms an image or a feature map using these parameters. An STN is a CNN with one or several ST modules inserted at arbitrary depths.

#### A. How STNs can enable invariance

We will here work with a continuous model of the image space. We model an image as a measurable function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$  and denote this space of images as  $V$ . Let  $\{\mathcal{T}_h\}_{h \in H}$  be a family of image transformations corresponding to a group  $H$ .  $\mathcal{T}_h$  transforms an image by acting on the underlying space

$$(\mathcal{T}_h f)(x) = f(\mathcal{T}_h^{-1} x) \quad (1)$$

where  $\mathcal{T}_h : \mathbb{R}^n \rightarrow \mathbb{R}^n$  is a linear map. We here consider affine image transformations, but the general argument is also valid for non-linear invertible transformations such as e.g. diffeomorphisms. Let  $\Gamma : V \rightarrow V^k$  be a (possibly non-linear) translation covariant feature extractor with  $k$  feature channels.  $\Gamma$  could e.g. correspond to a sequence of convolutions and pointwise non-linearities.  $\Gamma$  is *invariant* to  $\mathcal{T}_h$  if the feature response for a transformed image is equal to that of its original

$$(\Gamma \mathcal{T}_h f)_c(x) = (\Gamma f)_c(x), \quad (2)$$

where  $c \in [1, 2, \dots, k]$  corresponds to the feature channel. An ST module can *support invariance* by learning to transform input images to a canonical pose, before feature extraction, by applying the inverse transformation

$$(\Gamma \text{ST}(\mathcal{T}_h f))_c(x) = (\Gamma \mathcal{T}_h^{-1} \mathcal{T}_h f)_c(x) = (\Gamma f)_c(x). \quad (3)$$

We will in the following assume such a *perfect ST* that always manages to predict the correct pose of an object<sup>1</sup>. We now show that even a perfect ST cannot support invariance if instead applied to *CNN feature maps*.

#### B. The problems of transforming CNN feature maps

An advantage of inserting an ST deeper into the network is that the ST can make use of more complex features shared with the classification network. When using STNs that transform feature maps, as opposed to input images, the key question is whether it is possible to undo a transformation of an image

<sup>1</sup>There is no hard-coding of invariance in an STN and no guarantee that the predicted object pose is correct or itself invariant. We here assume the ideal case where the localization network does learn to predict the transformations that would align a relevant subset of all images (e.g. all images of the same class) to a common pose.




Fig. 2. For any transformation that includes a scaling component, the field of view of a feature extractor with respect to an object will differ between an original and a rescaled image. Consider a simple linear model that performs template matching with a single filter. When applied to the original image, the filter matches the size of the object that it has been trained to recognize and thus responds strongly. When applied to a rescaled image, the filter never covers the full object of interest. Thus, the response cannot be guaranteed to take even the same set of values for a rescaled image and its original.

after feature extraction. Is there a  $\mathcal{T}_g$  dependent on  $\mathcal{T}_h$  such that (applying the same transformation in each feature channel)

$$(\mathcal{T}_g \Gamma \mathcal{T}_h f)_c(x) = (\Gamma \mathcal{T}_h f)_c(\mathcal{T}_g^{-1}x) \stackrel{?}{=} (\Gamma f)_c(x) \quad (4)$$

holds for all  $f, \Gamma$  and  $h$ ? If this is possible, we refer to it as *feature map alignment*. An ST that transforms CNN feature maps could then support invariance by the same mechanism as for input images. We here present the key intuitions and the outline of a proof that this is, in the general case, *not possible*. We refer to [28] for a mathematically rigorous proof. Note that for any translation covariant feature extractor, such as a continuous or discrete CNN, feature map alignment for *translated images* is, however, possible by means of a simple translation of the feature maps.

1) Using  $\mathcal{T}_g = \mathcal{T}_h^{-1}$  is a necessary condition to align CNN feature maps with an ST: The natural way to align the feature maps of a transformed image to those of its original would be to apply the *inverse spatial transformation* to the feature maps of the transformed image i.e.

$$\mathcal{T}_h^{-1}(\Gamma \mathcal{T}_h f)_c(x) = (\Gamma \mathcal{T}_h f)_c(\mathcal{T}_h x). \quad (5)$$

For example, to align the feature maps of an original and a rescaled image, we would, after feature extraction, apply the inverse scaling to the feature maps. Using  $\mathcal{T}_h^{-1}$  is, in fact, a *necessary condition* for (4) to hold to hold [28]. To see this, note that the value for each spatial position  $x$  must be computed from *the same region* in the original image for the right hand and left hand side. Clearly, features extracted from different image regions cannot be guaranteed to be equal. Assume that  $(\Gamma f)_c(x)$  is computed from a region  $\Omega$  centered at  $x$  in the original image. Using the definition of  $\mathcal{T}_h$  and the fact that  $\Gamma$  is translation covariant, we get that  $(\mathcal{T}_g \Gamma \mathcal{T}_h f)_c(x)$  is computed from a region  $\Omega'$  centered at  $\mathcal{T}_g^{-1} \mathcal{T}_h^{-1} x$ . Now,

$$\Omega = \Omega' \implies \mathcal{T}_g^{-1} \mathcal{T}_h^{-1} x = x \implies \mathcal{T}_g = \mathcal{T}_h^{-1} \quad (6)$$

and thus the only candidate transformation to align CNN feature maps with an ST is  $\mathcal{T}_g = \mathcal{T}_h^{-1}$ . We, next, show that using  $\mathcal{T}_g = \mathcal{T}_h^{-1}$  is, however, not a *sufficient condition* to enable feature map alignment for two key reasons.

2) *Transforming an image typically implies a shift in the channel dimension of the feature map*: When transforming an input image, this typically causes not only a spatial shift in its feature representation but also a shift in the *channel dimension*. This problem is illustrated in Figure 1. Since an ST performs a *purely spatial transformation*, it cannot correct for this. A similar reasoning is applicable to a wide range of image transformations. An exception would be if the features extracted at a specific layer are themselves invariant to  $H$ . An example of this would be a network built from rotation invariant filters  $\lambda$ , where  $\lambda(x) = \lambda(\mathcal{T}_h x)$  for all  $\lambda$ . For such a network, or a network with more complex (learned or hardcoded) rotation invariant features at a certain layer, feature map alignment of rotated images would be possible.

3) *Receptive field shapes of neural networks are not invariant*: A second problem is that the receptive fields of neural networks are typically not invariant to the relevant transformation group. This will be the case for any filters of finite support together with a transformation that contains a uniform or non-uniform scaling component. In that case,  $\mathcal{T}_h^{-1}(\Gamma \mathcal{T}_h f)_c(x)$  will not only differ from  $(\Gamma f)_c(x)$  because it might be computed from differently oriented image patches, but also because it will be computed from *not fully overlapping* image patches. This problem is illustrated in Figure 2. For a scaling transformation, a convolutional filter, adapted to detecting a certain feature at one scale, will for a larger scale never fully cover the relevant object. An inverse scaling transformation cannot align the feature maps, since even *the set of values* in the feature maps is not guaranteed to be equal.

4) *Conclusion*: Our arguments show that a purely spatial transformation cannot align the feature maps of a transformed image with those of its original for general affine transformations. This implies that while an STN transforming feature maps will support invariance to translations, it will *not enable invariant recognition for more general affine transformations*. The exception is if the features in the specific network layer are themselves invariant to the relevant transformation group.

### III. STN ARCHITECTURES

We test four different ways to structure STNs, all depicted in Figure 3. By comparing these four architectures, we can separate out the effects of (i) whether it is good or bad to transform feature maps, (ii) whether it is useful for localization networks to use deep features when predicting transformation parameters, and (iii) whether it is better for a localization network to make use of representations from the classification network, or to train a large localization network from scratch.

1) *ST transforming the input*: The localization network and the ST are placed in the beginning of the network and transform the input image. This approach is denoted by STN-C0.

2) *ST transforming a feature map*: The localization network takes a CNN feature map from the classification network as input, and the ST transforms the feature map. This architecture does not support invariance. A network with the ST after  $X$  convolutional layers is denoted by STN-CX.

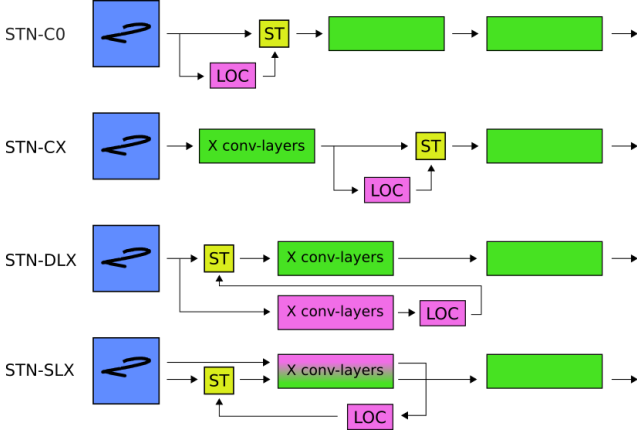


Fig. 3. Depiction of four different ways to build STNs. LOC denotes the localization network, which predicts the parameters of a transformation. ST denotes the spatial transformer, which takes these parameters and transforms an image or feature map according to them. In **STN-C0**, the ST transforms the input image. In **STN-CX**, the ST transforms a feature map, which prevents proper invariance. **STN-DLX** transforms the input image, but makes use of deeper features by including copies of the first X convolutional layers in the localization network. This is not fundamentally different from (1), but acts as a useful comparison point. **STN-SLX** is similar to STN-DLX, but shares parameters between the classification and localization networks.

3) *Deeper localization*: The localization network is placed in the beginning, but it is made deeper by including copies of some layers from the classification network. In particular, an STN where the localization network includes copies of the first X convolutional layers is denoted by STN-DLX. STN-DLX is not fundamentally different from STN-C0, since both architectures place the ST before any other transformations, but it acts as a useful comparison point to STN-CX: Both networks can make use of equally deep representations, but STN-DLX does not suffer any problems with achieving invariance. In addition, the deep representations of STN-DLX are *independently trained* from the classification network. This is beneficial if different filters are useful to find transformation parameters than to classify the image, but requires more parameters. If the training signal becomes less clear when propagated through the ST, it could also make the network harder to train. These differences motivate our fourth architecture.

4) *Shared localization*: As with STN-DLX, we place a deeper localization network in the beginning. However, for each of the copied layers, we *share parameters* between the classification network and the localization network. An STN where the first X layers are shared will be denoted by STN-SLX. STN-SLX solves the theoretical problem, uses no more parameters than STN-CX, and like STN-CX, the localization network makes use of layers trained directly on the classification loss.

#### A. Iterative methods

An iterative version of STNs known as IC-STN was proposed in [26]. It starts from an architecture that has multiple successive STs in the beginning of the network, and develops it in two crucial ways: (i) Instead of letting subsequent STs

transform an already-transformed image, all predicted transformation parameters are remembered and composed, after which the composition is used to transform the original image. Note that each localization network still uses the transformed image when predicting transformation parameters: the composition is only done to preserve image quality and to remove artefacts at the edges of the image. (ii) All STs use localization networks that share parameters with each other.

Both of these improvements can be generalized to work with STN-SLX. The simplest extension is to use several STs at the same depth, where all STs share localization network parameters with each other in addition to sharing parameters with the classification network. Moreover, the first improvement can be used with STN-SLX even when there are multiple STs *at different depths*, since all of them will transform the input image, regardless. In this case, the final layers of their localization networks remain separate, but whenever they predict transformation parameters, the parameters are composed with the previously predicted transformations, and used to transform the input image. This is illustrated in Figure 4.

## IV. EXPERIMENTS

### A. MNIST

1) *Datasets*: MNIST is a simple dataset containing grayscale, handwritten digits of size 28x28 [7]. To see how different STN architectures compensate for different transformations, we compare them on 3 different variations of MNIST (the first two constructed as in [1]): In Rotation (R), the digits are rotated a random amount between  $\pm 90^\circ$ . In Translation (T), each digit is placed at a random location on a 60x60-pixel background; to make the task more difficult, the background contains clutter generated from random fragments of other MNIST-images. In Scale (S), the digits are scaled a random amount between 0.5x and 4x, and placed in the middle of a 112x112-pixel background cluttered in a similar way to (T). Additional details about the experiments, networks and datasets are given in the supplementary materials.

2) *Networks*: We use network architectures similar to those in [1]. On all three datasets, the baseline classification network is a CNN that comprises two convolutional layers with max pooling. On (R), the localization network is a FCN with 3 layers. Since the images in (T) and (S) are much larger, their localization networks instead comprise two convolutional layers with max pooling and a single fully connected layer. Like in [1], the ST used with (T) produces an image with half the pixel width of its input. This is done because a perfect localization network should be able to locate the digit and scale it 2x, so all 60x60 pixels would not be needed.

MNIST is an easy dataset, so there is a risk that strong networks could learn the variations without using the ST. To make the classification accuracy dependent on the ST's performance, we intentionally use quite small networks. The networks of (R) and (T) have 50 000-70 000 learnable parameters in total, while those of (S) have around 135 000. All networks are trained using SGD, with cross-entropy loss. Networks trained



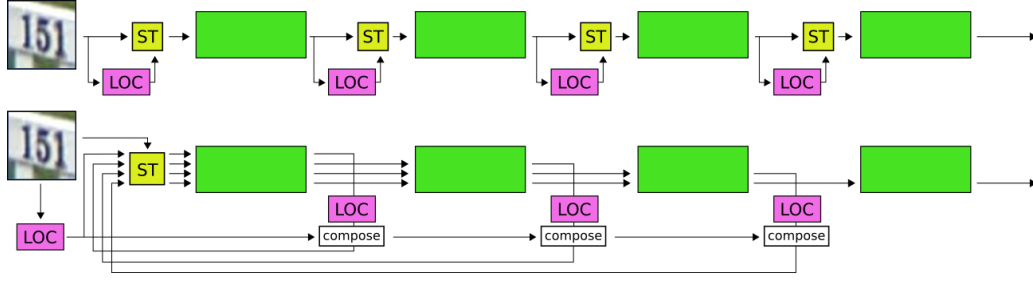


Fig. 4. Depiction of how an STN with STs at different depths can be transformed into an iterative STN with shared layers. The first network transforms CNN feature maps, while the second network always transforms the input image.

on the same dataset have approximately the same number of parameters and use the same hyperparameters.

The tested architectures are STN-C0, STN-C1 (which places the ST after the first convolutional layer and max pooling), STN-DL1, and STN-SL1. A baseline CNN is also tested. Note that because of the transformations present in the training dataset, this equals a standard CNN trained with *data augmentation*. All networks are trained with 10 different random seeds, and each architecture is evaluated on 100 000 images generated by random transformations of the MNIST test set.

3) *Results*: As a first investigation of the different architectures, we study how well the STs learn to transform the digits to a canonical pose, when the networks are trained to predict the MNIST labels.

Figure 5 shows examples of how STN-C1 and STN-SL1 perform on (R), (T), and (S). As predicted by theory, STN-C1 can successfully localize a translated digit, but STN-SL1 is better at compensating for differences in rotation and scale. The difference between the networks' abilities to compensate for rotations is especially striking. Figure 6 shows that STN-SL1 compensates for rotations well, while STN-C1 barely rotates the images at all. The reason for this is that a rotation is not enough to align deeper layer feature maps.

To quantify the STs' abilities to find a canonical pose, we measure the standard deviation of the digits' final poses after they have been perturbed and the ST has transformed them. For the rotated, translated, and scaled dataset, the final pose is measured in units of degrees, pixels translated, and logarithm of the scaling factor, respectively.<sup>2</sup> Table I displays this value for each network.

As can be seen, STN-SL1 is the best network on (R) and STN-DL1 is the best network on (T) and (S). Both these architectures use deeper localization networks, which gives them the potential to predict transformations better than STN-C0. As opposed to STN-C1, which also uses deeper representations for predicting transformation parameters, STN-SL1, STN-DL1, and STN-C0 all transform the input. This

<sup>2</sup>As measure of the rotation of an affine transformation matrix, we compute  $\arctan((a_{21} - a_{12})/(a_{11} + a_{22}))$  determined from a least-squares fit to a similarity transformation.  $a_{13}$  and  $a_{23}$  are used to measure the translation. As a measure of the scaling factor, we use the  $\log_2$  of the determinant  $a_{11}a_{22} - a_{12}a_{21}$ . The standard deviation (or in the case of (T), the standard distance deviation) of the final pose is measured separately for each label, with Table I reporting the average across all labels and across 10 different random seeds.

TABLE I  
AVERAGE STANDARD DEVIATION OF THE FINAL ANGLE (R), FINAL DISTANCE (T), AND FINAL SCALING (S).

Network	R (degrees)	T (pixels)	S ( $\log_2(\det)$ )
STN-C0	23.2	1.16	0.319
STN-C1	47.2	1.15	0.508
STN-DL1	26.8	<b>1.08</b>	<b>0.291</b>
STN-SL1	<b>18.7</b>	1.32	0.330

allows them to perform better on (R) and (S), as predicted by theory. However, STN-C1 performs adequately on (T).

STN-SL1 performing well on (R) can be explained by it sharing parameters between the localization network and classification network. The localization network processes images before the ST transforms them, and the classification network processes them after the transformation. Thus, we should expect parameter sharing to be helpful if the filters needed before and after the transformation are similar. This is true for (R), since edge- and corner-detectors of multiple different orientations are likely to be helpful for classifying MNIST-digits independently of digit orientation. In contrast, on both (T) and (S), the scale and resolution of the digits vary significantly before and after the transformation. Since the images contain scales and resolutions before the transformation that are not present after the transformation, the localization network requires filters that the classification network does not need. STN-DL1 allows for the networks to use different filters, and consequently, it does better than STN-SL1 on (T) and (S).

Do these differences in ST-performance affect the classification performance? Table II shows that they do. STN-SL1 remains the best network on (R), while STN-DL1 remains best on (T) and (S). The architectures that transform the input remain better than STN-C1 on (R) and (S). One difference is that STN-SL1 is better than STN-C1 on (T), which is the opposite from their relation in Table I. Since the differences in both Table I and Table II are small, this could be caused by STN-SL1 being better at compensating for rotation and scale differences inherent in the original MNIST dataset.

4) *Concluding remarks*: As predicted by theory, it is significantly better to transform the input image when compensating for differences in rotation and scale, while transforming intermediate feature maps works reasonably well when compensating for differences in translation. STN-SL1 benefits from




Fig. 5. Illustration of how STN-C1 and STN-SL1 compensate for different perturbations. The top row shows three digits rotated (first image), translated (second image), or scaled (third image) in three different ways. The middle row and bottom row show how STN-C1 and STN-SL1 transform the digits in the top row. STN-C1 does not compensate for rotations at all, but it successfully localizes and zooms in on translated digits. It only compensates somewhat for scaling. STN-SL1 finds a canonical pose for all perturbations. Note that STN-C1 does not transform the input image, so the middle row is just an illustration of the transformation parameters that are normally used to transform its CNN feature map.

TABLE II

AVERAGE CLASSIFICATION ERROR AND STANDARD DEVIATION ON THE MNIST TEST-SET, ACROSS 10 DIFFERENT RUNS.

Network	R ( <i>std</i> )	T ( <i>std</i> )	S ( <i>std</i> )
CNN	1.71%(0.07)	1.61%(0.06)	1.38%(0.04)
STN-C0	1.08%(0.05)	1.10%(0.11)	0.85%(0.06)
STN-C1	1.32%(0.04)	1.16%(0.03)	0.96%(0.04)
STN-DL1	1.05%(0.02)	<b>1.08%</b> (0.07)	<b>0.77%</b> (0.04)
STN-SL1	<b>0.98%</b> (0.06)	1.13%(0.04)	0.82%(0.06)

sharing parameters on the rotation task, but in the presence of large scale variations, it is better for the localization and the classification network to be independent of each other.

### B. Street View House Numbers

1) *Dataset*: In order to learn how the different STN architectures perform on a more challenging dataset, we evaluate them on the Street View House Numbers (SVHN) [8]. The photographed house numbers contain 1-5 digits each, and we preprocess them by taking 64x64 pixel crops around each sequence, as done in [1]. This allows us to compare our results with [1], who achieved good results on SVHN with an iterative ST that transforms feature maps. Since the dataset benefits from the use of a deeper network, it also allows us to test how the architectures' performance varies with their depth.

2) *Comparison with [1]*: We use the same baseline CNN as [1]. The classification network comprises 8 convolutional layers followed by 3 fully connected layers. Since the images can contain up to 5 digits, the output consists of 5 parallel softmax layers. Each predicts the label of a single digit.

As variations of the baseline CNN, [1] considers STNs with two different localization networks: One with a large localization network in the beginning, here denoted by STN-C0-large, and one with a small localization network before each of the first 4 convolutional layers, here denoted by STN-C0123. The large localization network uses two convolutional layers and two fully connected layers, while each of STN-C0123's localization networks uses two fully connected layers. As a further variation, we consider the network STN-SL0123,




Fig. 6. The rotation angle predicted by the ST module (y-axis) as a function of the rotation angle applied to the input image (x-axis). The black points constitute a heatmap of 100 000 datapoints generated by random rotations of the MNIST test set, and reports the rotations done by a single ST. The red line corresponds to the best fit with orthogonal regression. STN-C1 cannot compensate for the rotations in a useful way, since it transforms CNN feature maps, while STN-SL1 directly counteracts the rotations by rotating the input.

which is similar to STN-C0123 but always transforms the input, as described in Section III-A. STN-C0123 and STN-SL0123 are illustrated in Figure 4.

The average classification errors of these networks are shown in Table III. STN-SL0123 achieves the lowest error, which shows that it is better to transform the input than to transform CNN feature maps, also on the SVHN dataset.

TABLE III

CLASSIFICATION ERRORS ON THE SVHN DATASET, AVERAGED OVER 3 RUNS, COMPARING OUR IMPLEMENTATION WITH [1]

Network	Error [1]	Error (ours)
CNN	4.0%	3.88%
STN-C0-large	3.7%	3.69%
STN-C0123	3.6%	3.61%
STN-SL0123	-	<b>3.49%</b>

3) *Comparing STs at different depths*: If an ST is placed deeper into the network, it can use deeper features to predict transformation parameters, but the problem with spatial transformations of CNN feature maps may get worse. STN-DLX and STN-SLX would not suffer from the latter, but might benefit from the former. We test this by placing STs at depths 0, 3, 6, or 8 in our base classification network, using the small localization network from the previous section.

The results are shown in Table IV. STN-C3, STN-DL3, and STN-SL3 perform better than STN-C0, indicating that STN-C0's inability to find the correct transformation parameters causes more problems than STN-C3 causes by transforming the feature map at depth 3. However, at depths 6 and 8, STN-CX becomes worse than not using any ST at all (see Table III) and STN-DLX performs worse than at depth 3. This stands in sharp contrast to STN-SLX, where the classification error keeps decreasing, reaching 3.26% for STN-SL8.

This shows that localization networks benefit from using deep representations, if they use filters from the classification network, while still transforming the input. It is not achievable by spatially transforming CNN feature maps, since

TABLE IV  
MEAN SVHN CLASSIFICATION ERROR OF STN-CX, STN-DLX AND STN-SLX AT 4 DIFFERENT DEPTHS, ACROSS 3 RUNS.

Depth	STN-CX	STN-DLX	STN-SLX
$X = 0$	3.81%	-	-
$X = 3$	3.70%	3.48%	3.54%
$X = 6$	3.91%	3.75%	3.29%
$X = 8$	4.00%*	3.76%	<b>3.26%</b>

\* One run diverged to  $> 99\%$  classification error. The error is the average of three runs where that did not happen.

transforming deep feature maps causes too much distortion. Just using deeper localization networks does not always work, either, as these results show that they fail to find appropriate transformation parameters at greater depths.

Note that the larger localization networks of STN-DLX and STN-SLX take more time to train. STN-SL0123 takes 1.8 times as long to train as STN-C0123, while STN-SL8 takes 1.6 times as long as STN-C0123.

4) *Comparing iterative STNs*: Given that the use of deeper localization networks helps performance, and that [26] showed that iterative methods improve performance, a natural question is whether using iterative methods is a replacement for using deeper features or if iterations and deeper features are complimentary. To answer this, we train STN-C0, STN-DL3, STN-SL3, STN-SL6, and STN-SL8 with two iterations.

The second iteration does not change STN-C0’s performance (mean error 3.80%). However, STN-SL3 improves significantly (mean error 3.38%), and STN-SL6 improves slightly (mean error 3.26%). STN-SL8 and STN-DL3 both perform worse with a second iteration. If a third iteration is added to STN-SL3 and STN-SL6 *during testing*, the errors further decrease to 3.33% and 3.18%, respectively.

This shows that multiple iterations do not improve performance when the localization network is too weak (as with STN-C0). Thus, multiple iterations is not a replacement for using deeper features, but can confer additional advantages. Here, parameter-shared networks are the only ones that benefit.

### C. Plankton

1) *Dataset*: Finally, we test the STN architectures on PlanktonSet [9]. This dataset is challenging, because it contains 121 classes and only about 30 000 training samples.

2) *Networks*: We use network architectures inspired by [29]. As base classification network, we use a CNN consisting of 10 convolutional layers with maxpooling after layers 2, 4, 7, and 10; followed by two fully connected layers before the final softmax output. As base localization network, two fully connected layers are used. All hyperparameters were chosen through experiments on the validation set.

3) *Results*: Table V displays the results, evaluated on the test set. The best architectures are STN-SL2 and STN-SL4. As on SVHN and MNIST, networks that transform the input are better than STN-CX. Also similar to the results on SVHN is that STN-DLX is significantly better for low values of  $X$ .

Increasing the depth of STN-SLX does not decrease performance as much as it does for STN-DLX, but neither does




Fig. 7. A rotated plankton transformed by various networks. The top row contains the input images, the middle row contains the top images transformed by STN-C2, STN-DL2, and STN-SL2, and the bottom row contains the top images transformed by STN-C7, STN-DL7, and STN-SL7. In the middle row, all networks find a canonical angle. However, in the bottom row, only STN-SL7 finds a canonical angle, while STN-C7 suffer from not being able to transform the input directly, and STN-DL7’s localization network is too deep.

TABLE V  
MEAN PLANKTONSET CLASSIFICATION ERROR OF STN-CX, STN-DLX, AND STN-SLX AT 5 DIFFERENT DEPTHS, ACROSS 4 RUNS.

Depth	STN-CX	STN-DLX	STN-SLX
$X = 0$	22.1%	-	-
$X = 2$	22.3%	21.6%	<b>21.5%</b>
$X = 4$	22.5%	21.7%	<b>21.5%</b>
$X = 7$	22.9%	22.7%	21.6%
$X = 10$	22.7%	22.7%	21.9%

it increase performance, as it does on SVHN. This is likely caused by differences between the datasets. Identifying the interesting objects in SVHN is not a trivial task, because the digits must be distinguished from the background. Since the classification network must be able to identify the digits, the localization network can be expected to benefit from sharing more layers. However, on PlanktonSet, the interesting objects are easily identifiable against the black background, which makes fewer layers sufficient. In addition, the easy separability of planktons and background allows for large-scale data augmentation on PlanktonSet, which makes the STs task both easier and less important for classification accuracy.

4) *Rotational invariance*: Figure 7 shows how the architectures transform a plankton rotated in different ways. In contrast to STN-C1 on MNIST (see Figure 5 and Figure 6), STN-C2 has learned to transform the example image to a canonical angle. Despite this, STN-C0, STN-DL2 and STN-SL2 perform substantially better in Table V. This shows that STN-CX’s problem is not that it is difficult to find correct transformation parameters; it is the problem that even a “correct” transformation does not yield proper invariance. In this case, it is nonetheless beneficial for STN-C2 to rotate the plankton to a canonical angle, while transformations made by the deeper STN-C7 introduce enough distortions that it is better to do nothing.

5) *Comparing iterative STNs*: Training STN-C0 with a second iteration improves it somewhat, yielding mean error 21.8%, but training it with a third iteration decreases performance. STN-SL2, STN-SL7, and STN-DL2 are all slightly improved by a second iteration (mean error 21.4%, 21.5%, and 21.5%), while the performance of STN-SL4 and STN-

DL4 decreases.

We, again, observe that multiple iterations is not a substitute for deeper features when predicting transformations, as the deeper networks remain better than STN-C0.

We also observed that the *training error* decreases much more than the test error, when a second iteration is added. For example, STN-SL2, STN-DL4, and STN-SL7 all reduce their training error with 1.7% or more. This suggests that the second iterations make the STs substantially more effective, but that this mostly leads to overfitting.

## V. SUMMARY AND CONCLUSIONS

We have presented a theoretical argument why STNs cannot transform CNN feature maps in a way that enables invariant recognition for other image transformations than translations. Investigating the practical implications of this result, we have shown that this inability is clearly visible in experiments and negatively impacts classification performance. In particular, while STs transforming feature maps perform adequately when compensating for differences in translation, STs transforming the input perform better when compensating for differences in rotation and scale. Our results also have implications for other approaches that spatially transform CNN feature maps, pooling regions, or filters. We do not argue that such methods cannot be beneficial, but we do show that these approaches have limited ability to achieve invariance.

Furthermore, we have shown that STs benefit from using deep representations when predicting transformation parameters, but that localization networks become harder to train as they grow deeper. In order to use representations from the classification network, while still transforming the input, we have investigated the benefits of sharing parameters between the localization and the classification network. Our experiments show that parameter-shared localization networks remain stable better as they grow deeper, which has significant benefits on complex datasets. Finally, we find that iterative image alignment is not a substitute for using deeper features, but that it can serve a complimentary role.

## REFERENCES

- [1] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial transformer networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2015, pp. 2017–2025.
- [2] C. B. Choy, J. Gwak, S. Savarese, and M. Chandraker, "Universal correspondence network," in *Advances in Neural Information Processing Systems*, 2016, pp. 2414–2422.
- [3] J. Li, Y. Chen, L. Cai, I. Davidson, and S. Ji, "Dense transformer networks," *arXiv preprint arXiv:1705.08881*, 2017.
- [4] S. Kim, S. Lin, S. R. JEON, D. Min, and K. Sohn, "Recurrent transformer networks for semantic correspondence," in *Advances in Neural Information Processing Systems*, 2018, pp. 6126–6136.
- [5] Z. Zheng, L. Zheng, and Y. Yang, "Pedestrian alignment network for large-scale person re-identification," *IEEE Transactions on Circuits and Systems for Video Technology*, 2018.
- [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 (NIPS)*, 2019, pp. 9142–9153.
- [7] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [8] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," in *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011.
- [9] R. K. Cowen, S. Sponaugle, K. Robinson, J. Luo, O. S. University, and H. M. S. Center, "PlanktonSet 1.0: Plankton imagery data collected from F.G. Walton Smith in Straits of Florida from 2014-06-03 to 2014-06-06 and used in the 2015 National Data Science Bowl (NCEI accession 0127422)." [Online]. Available: <https://accession.nodc.noaa.gov/0127422>
- [10] B. D. Lukas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Image Understanding Workshop*, 1981.
- [11] J. R. Bergen, P. Anandan, K. J. Hanna, and R. Hingorani, "Hierarchical model-based motion estimation," in *European Conference on Computer Vision (ECCV)*. Springer, 1992, pp. 237–252.
- [12] T. Lindeberg and J. Gårding, "Shape-adapted smoothing in estimation of 3-D depth cues from affine distortions of local 2-D structure," *Image and Vision Computing*, vol. 15, pp. 415–434, 1997.
- [13] A. Baumberg, "Reliable feature matching across widely separated views," in *Proc. Computer Vision and Pattern Recognition (CVPR'00)*, Hilton Head, SC, 2000, pp. 1:1774–1781.
- [14] K. Mikolajczyk and C. Schmid, "Scale & affine invariant interest point detectors," *International Journal of Computer Vision*, vol. 60, no. 1, pp. 63–86, 2004.
- [15] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. van Gool, "A comparison of affine region detectors," *International Journal of Computer Vision*, vol. 65, no. 1–2, pp. 43–72, 2005.
- [16] L. Sifre and S. Mallat, "Rotation, scaling and deformation invariant scattering for texture discrimination," in *Proc. Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 1233–1240.
- [17] T. Cohen and M. Welling, "Group equivariant convolutional networks," in *International Conference on Machine Learning*, 2016, pp. 2990–2999.
- [18] R. Kondor and S. Trivedi, "On the generalization of equivariance and convolution in neural networks to the action of compact groups," in *International Conference on Machine Learning (ICML)*, 2018, pp. 2752–2760.
- [19] T. Lindeberg, "Provably scale-covariant continuous hierarchical networks based on scale-normalized differential expressions coupled in cascade," *Journal of Mathematical Imaging and Vision*, vol. 62, no. 1, pp. 120–148, 2020.
- [20] A. Arcos-García, J. A. Alvarez-García, and L. M. Soria-Morillo, "Deep neural network for traffic sign recognition systems: An analysis of spatial transformers and stochastic optimisation methods," *Neural Networks*, vol. 99, pp. 158–165, 2018.
- [21] T. V. Souza and C. Zanchettin, "Improving deep image clustering with spatial transformer layers," in *International Conference on Artificial Neural Networks (ICANN)*. Springer, 2019, pp. 641–654.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," in *European Conference on Computer Vision*. Springer, 2014, pp. 346–361.
- [23] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," in *Int. Conf. on Learning Representations (ICLR)*, 2016.
- [24] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, "Deformable convolutional networks," in *Proc. International Conference on Computer Vision (ICCV)*, 2017, pp. 764–773.
- [25] B.-I. Cirstea and L. Likforman-Sulem, "Tied spatial transformer networks for digit recognition," in *2016 15th International Conference on Frontiers in Handwriting Recognition (ICFHR)*. IEEE, 2016, pp. 524–529.
- [26] C.-H. Lin and S. Lucey, "Inverse compositional spatial transformer networks," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2568–2576.
- [27] K. Lenc and A. Vedaldi, "Understanding image representations by measuring their equivariance and equivalence," in *Proc. Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 991–999.
- [28] Y. Jansson, M. Maydanskiy, L. Finnveden, and T. Lindeberg, "Inability of spatial transformations of CNN feature maps to support invariant recognition," *arXiv preprint:2004.14716*, 2020.
- [29] J. Burms, P. Buteneers, J. Degraeve, S. Dieleman, I. Korshunova, A. van den Oord, and L. Pigou, "Classifying plankton with deep neural networks." [Online]. Available: <https://benanne.github.io/2015/03/17/plankton.html>