



HAL
open science

Can Generalised Divergences Help for Invariant Neural Networks?

Santiago Velasco-Forero

► **To cite this version:**

Santiago Velasco-Forero. Can Generalised Divergences Help for Invariant Neural Networks?. Geometric Science of Information, 14071, Springer Nature Switzerland, pp.82-90, 2023, Lecture Notes in Computer Science, 10.1007/978-3-031-38271-0_9 . hal-04303522v2

HAL Id: hal-04303522

<https://hal.science/hal-04303522v2>

Submitted on 8 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Can generalised divergences help for invariant neural networks?

Santiago Velasco-Forero

MINES Paris - PSL-Research University - Centre de Morphologie Mathématique

Abstract. We consider a framework including multiple augmentation regularisation by generalised divergences to induce invariance for non-group transformations during training of convolutional neural networks. Experiments on supervised classification of images at different scales not considered during training illustrate that our proposed method performs better than classical data augmentation.

1 Introduction

Deep neural networks are the primary model for learning functions from data, in different tasks ranging from classification to generation. Convolutional neural networks (CNNs) have become a widely used method across multiple domains. The *translation* equivariance of convolutions is one of the key aspects to their success [23]. This equivariance is induced by applying the same convolutional filter to each area of an image producing learned weights that are independent of the location. This mechanism is called *weight sharing*. Ideally, CNNs should perform equally well regardless of input scale, rotation or reflections. Numerous attempts have been made to address using the formalism of *group-convolutions*[12], *steerable filters*[8], *moving frames*[31], *wavelet*[2], *partial differential equations* [36], *Gaussian filters*[26], *Elementary Symmetric Polynomials*[27] among others. Despite all these recent advances, it is still unclear what is the most adequate way to adapt these methods for the case of more general transformations that cannot be considered as a group [38,30]. The most commonly used solution is to take advantage of *data augmentation*, where the inputs are randomly transformed during training to induce an output (which is) insensible to some given transformations [34]. But still data augmentation implies neither equivariance nor invariance. An more elaborated path would be to apply the *weight sharing* mechanism for each discretisation of the transformation, followed by an integration in the sense of [22]. In this paper, we study the use of contrastive based regularisation on a set of transformations during training. Surprisingly, our proposition presents the best performance considering the power of generalisation outside the interval of values where the transformation has been sampled during training. This phenomenon is illustrated in the case of supervised classification on aerial images and traffic signs at different scales.

2 Proposition

2.1 Motivation

Data augmentation is nowadays one of the main components of the design of efficient training for deep learning models. Initially proposed to improve over-sampling on class-imbalanced datasets [9] or to prevent overfitting when the model contains more parameters than training points [24]. Recent research has shown its interest in increasing generalization ability especially when augmentations yield samples that are diverse [15]. We restrict our study to augmentations which act on a single sample and do not modify labels, this means that we do not consider *mixup augmentations* [40]. Namely, we study augmentations which can be written as $(t(\mathbf{x}), y)$, where (\mathbf{x}, y) denotes an input-label pair, and $t \in \mathcal{T}$ is a random transformation sampled from a set of possible transformations \mathcal{T} . Let f denotes a projection from the input space to a *latent space*. The latent space is said to be *invariant* to \mathcal{T} if for any input \mathbf{x} and any $t \in \mathcal{T}$, $f(t(\mathbf{x})) = f(\mathbf{x})$. Practitioners recommend to use *data augmentation* to induce invariance *by training*[3]. Usually, data augmentation consists of randomly applying an element of the set of \mathcal{T} during training. An alternative to data augmentation is possible when \mathcal{T} is a group. One can construct an invariant function $f_\theta(\mathbf{x}; \eta)$ from a non-invariant function $g_\theta(\mathbf{x})$ by integrating over all the group actions. This concept is referred to as *insensitivity* [37], *soft-invariance* [4], or *deformation stability* [6]. The special case in which there exist a subgroup H where the computation can be reduced to summing over H is called *Reynolds design* [29]. For topological groups, there is a non-zero, translation invariant measure called *Haar measure* that can be used to define invariant convolution on a group [20,28] or invariance by integration of kernels [17,25]. An alternative to define an invariant function, is to use composition of *equivariant* functions followed by an invariant pooling in the *Geometric Deep Learning Blueprint* [6]. In both cases, the invariance is defined *by structure*[3] which can be seen as a constraint in the model that one is learning.

However, in many applications, the transformations under study is not a group, so that the above arguments are not easily generalizable. In this paper we are interested in using data augmentation to induce invariance during training in deep learning models for the case where the set \mathcal{T} is not a group. The idea is to include a regularisation term that takes into account K realisations of the transformation family, i.e, in the loss function, we will include $\text{Loss}_{\mathcal{T}}(\mathbf{x}, t_1(\mathbf{x}), t_2(\mathbf{x}), \dots, t_K(\mathbf{x}))$ where the t_i denotes a random value of transformations \mathcal{T} . To apply this method we do not need any requirements on the transformations \mathcal{T} .

2.2 Related work

The idea of using multiple random augmentations (K is our case) during training is also found in the following methods:

Semi-supervised learning In a semisupervised case, [41] proposed to learn a classifier penalised for quick changes in its predictions.

Self-training Self-training also known as decision-directed or self-taught learning machine, is one of the earliest approach in semi-supervised learning [14,32]. The idea of these approaches is to start by learning a supervised classifier on the labelled training set, and then, at each iteration, the classifier selects a part of the unlabelled data and assigns pseudo-labels to them using the classifier’s predictions. These pseudo-labeled examples are considered as additional labeled examples in the following iterations. The function loss includes a trade-off term to balance the influence of pseudo labels.

Self-supervised learning Most of these works are placed in a *joint-embedding framework* [10,18], where augmented views (usually two) are generated from a source image. These two views are then projected to an encoder, giving representations, and then through a projection back to an embedding space. Finally, a loss minimises the distance between the embeddings, i.e. makes them invariant to the augmentations, and is combined with a regularisation loss to spread embeddings in space.

Data augmentation regularisations A negative aspect of data augmentation has been illustrated in [13] which is the slow down of training speed and a minimal effect on the variance of the model. The idea of using multiple augmentation per image in the same minibatch has been used to solve that problem, and it has been used to improve at the same time the classifier’s generalisation performance [13]. This simple modification computes an average of the minibatch on different augmentations that asymptotically approaches a Reynolds operator [29] when the number of considered augmentations gets as large as possible. Recently, [39,5] proposed the use of a regularisation term for multiple augmentations, which is the mechanism that we will evaluate in this paper.

2.3 Supervised regularisation by generalised divergences

We propose to use multiple data augmentations in the target transformation, and use a generalised divergence as a regularisation term. The idea follows those presented in [11], and is contrary to the usual mechanism of *data augmentation*, where the network is trained to classify in the same class each of the augmentations, but never considers a term related to the divergence produced by the transformation. Since we use K augmentations, we must consider a divergence from multiple probability distributions, which is called *generalised divergences*. We consider the classical framework of training deep learning models from N

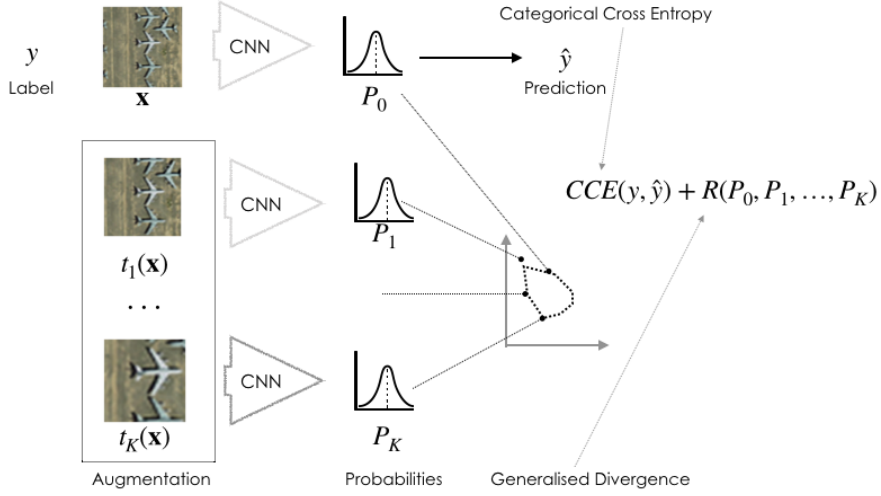


Fig. 1. Scheme of our proposition. We propose to use a regularisation that considers multiple realisations of the transformation family, this regularisation uses generalised divergences. Since you want to evaluate the invariance of a classification problem, the model uses only the classification of the original image (not of the transformations). The probability distributions are obtained in the output of a softmax layer.

samples $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ and as objective minimising the following loss function:

$$\text{Loss}(\mathbf{x}, y) = \sum_i^N \text{Loss}_{class}(y_i, \hat{y}_i) + \alpha \sum_i^N \text{Loss}_{\mathcal{T}}(\mathbf{x}_i, t_0(\mathbf{x}_i), t_1(\mathbf{x}_i), \dots, t_K(\mathbf{x}_i)), \quad (1)$$

where \hat{y}_i denotes the prediction of the model, and y_i the ground-truth class of the i -th sample \mathbf{x}_i .

The first term is a supervised classification term, and the second term $\text{Loss}_{\mathcal{T}}$ is the main interest of our proposition. We propose to use statistical divergences to compare the outputs produced by model f applied to the original data \mathbf{x} and $K + 1$ random augmentations of \mathbf{x} , i.e $\{\mathbf{x}_i, t_0(\mathbf{x}_i), t_1(\mathbf{x}_i), \dots, t_K(\mathbf{x}_i)\}$. In our supervised case, we use the last layer of the model f , which is usually a sum-one layer (softmax) indicating the probability of belonging to a given class. For two probability distributions P, Q , the most renowned statistical divergence rooted in information theory [21] is the *Kullback–Leibler divergence*,

$$D_{KL}(P||Q) = \sum P(x) \log \left(\frac{P(x)}{Q(x)} \right).$$

Defining divergence between more than two distributions has been studied for many authors called often *generalised divergences* or *dissimilarity coefficient* in [33]. Let $K \geq 1$ be a fixed natural number. Each generalised divergence R that we consider here, satisfies the following properties:

1. $R(P_0, P_1, P_2, \dots, P_K) \geq 0$
2. $R(P_0, P_1, P_2, \dots, P_K) = 0$ whenever $P_0 = P_1 = \dots = P_K$
3. R is invariant to permutation of input components.

These three properties are important for the minimisation of this divergence to induce the invariance during training in the case we are studying. Accordingly, we consider the following two generalised divergences, the *Average Divergence* [33]

$$R_1(P_0, P_1, P_2, \dots, P_K) = \frac{1}{K(K+1)} \sum_{i,j=0, i \neq j}^K D_{KL}(P_i || P_j) \quad (2)$$

the *Information radius* [35] which is the generalised mean of the Rényi’s divergences between each of the P_i ’s and the generalised mean of all the P_i ’s,

$$R_2(P_0, P_1, P_2, \dots, P_K) = \frac{1}{K+1} \sum_i^K D_{KL}((K+1)^{-1} \sum_j^K P_j || P_i) \quad (3)$$

In the following section, we compare the use of (1), considering as Loss_T the average divergence in (2) or the information radius (3).

3 Experiments

In this experimental section, we have followed the training protocol presented in [1] on two datasets, *Aerial* and *Traffic Signs*, which contains images 64×64 RGB-color images on 48 different scales. The objective is to obtain a scale and translation invariant model for supervised classification on nine (resp. 16) classes on *Aerial* (resp. *Traffic Signs*) dataset. Keen readers are referred to [16,19,38,30,1] for a deeper understanding of different propositions for scale invariant convolutional networks. Following [1] the model is a CNN with two layers using categorical cross-entropy as a supervised term in (1). An example per dataset at different scales is shown in Figure 3. The models are trained on the middle interval of the transformation parameterisation and the performance of the models are evaluated outside this interval. This is called *Mid2Rest* scenario in [1]. The value of K in (3) and (2) has been set equal to three in our experiments. Quantitative comparison of results are found in Table 3 for both considered datasets. The reported result is the average and standard deviation of performance on the scales and images that were not considered during training. On the considered datasets, the information radius (3) presents better results in terms of performance over the unseen scales, with respect to both the average divergence (2), and the classical data augmentation method. Finally, for a better illustration of the difficulty of the task, the best value of the lambda and regularisation function is compared with the data augmentation in five random training runs, and compared across the different scales for the two databases in Figure 3.



Fig. 2. Examples of images at different parameter transformation in the two considered datasets. From left to right: Scale 1, 3, 23, 25, 45 and 47. Training is done considering only images of intermediate scales (17 to 32) in both training and validation. Evaluation is performed on both small (0 to 16) and large (33-48) scales. In first row: An example of *Traffic Sign* dataset. In second row: An example of *Aerial* dataset.

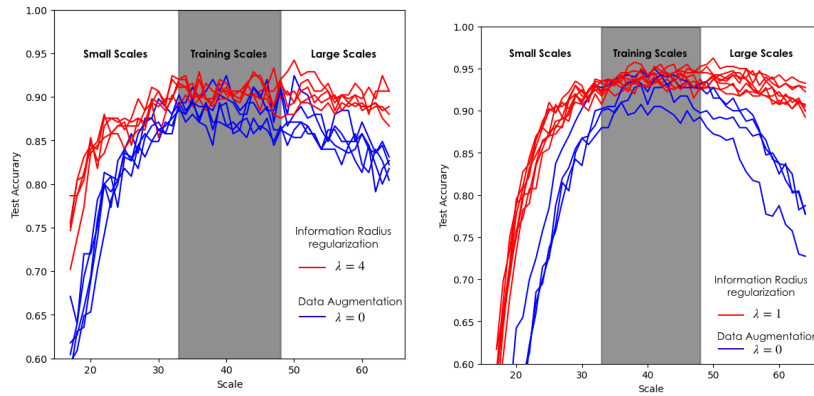


Fig. 3. Detailed plots of scale generalisation on *Mid2Rest* scenario in *Aerial* datasets (Left) and *Traffic Sign* dataset (Right). Five repetitions of the training is illustrated per method. Our proposition performs clearly better than classical data augmentation.

Table 1. Results of the Generalised Divergence in Aerial and Traffic Sign dataset on scales non-considered during training. A visual comparison of results are shown in Fig.3

Aerial		Small Scales		Large Scales	
Method	λ	test acc.	\pm std	test acc.	\pm std
Data Aug.	0.0	0.776	± 0.014	0.845	± 0.008
Av. Div.(2)	0.5	0.854	± 0.013	0.889	± 0.011
	1.0	0.852	± 0.016	0.888	± 0.009
	1.5	0.858	± 0.013	0.889	± 0.011
	2.0	0.841	± 0.028	0.880	± 0.020
	2.5	0.846	± 0.015	0.881	± 0.009
	3.0	0.834	± 0.023	0.877	± 0.018
	3.5	0.833	± 0.008	0.873	± 0.017
	4.0	0.822	± 0.021	0.864	± 0.016
	5.0	0.828	± 0.016	0.872	± 0.022
	10.0	0.824	± 0.019	0.865	± 0.012
Inf. Rad.(3)	0.5	0.845	± 0.016	0.878	± 0.017
	1.0	0.847	± 0.012	0.885	± 0.011
	1.5	0.853	± 0.010	0.881	± 0.009
	2.0	0.853	± 0.009	0.884	± 0.011
	2.5	0.850	± 0.014	0.887	± 0.005
	3.0	0.859	± 0.013	0.885	± 0.010
	3.5	0.841	± 0.008	0.886	± 0.009
	4.0	0.839	± 0.018	0.877	± 0.016
	5.0	0.845	± 0.017	0.883	± 0.010
	10.0	0.829	± 0.014	0.870	± 0.011

Traffic Sign		Small Scales		Large Scales	
Method	λ	test acc.	\pm std	test acc.	\pm std
Data Aug.	0.0	0.721	± 0.021	0.824	± 0.024
Av. Div. (2)	0.5	0.821	± 0.014	0.898	± 0.020
	1.0	0.829	± 0.012	0.921	± 0.018
	1.5	0.820	± 0.015	0.902	± 0.012
	2.0	0.806	± 0.027	0.886	± 0.025
	2.5	0.797	± 0.044	0.883	± 0.034
	3.0	0.789	± 0.025	0.882	± 0.018
	3.5	0.754	± 0.026	0.842	± 0.028
	4.0	0.743	± 0.026	0.832	± 0.042
Inf. Rad.(3)	0.5	0.806	± 0.012	0.908	± 0.017
	1.0	0.825	± 0.022	0.921	± 0.016
	1.5	0.829	± 0.014	0.918	± 0.015
	2.0	0.830	± 0.016	0.918	± 0.007
	2.5	0.815	± 0.020	0.906	± 0.016
	3.0	0.828	± 0.013	0.909	± 0.018
	3.5	0.818	± 0.019	0.898	± 0.014
	4.0	0.823	± 0.020	0.917	± 0.016

4 Conclusions

In this paper we present a proposal for the use of regularisation from multiple data augmentation with generalised divergences. Quantitative results show the interest of our method in the case of generalisation to scales that have not been considered during training. Future studies may include the study of multi-parametric transformations, as these are used to avoid overfitting in large neural networks. Additionally, generalised divergence considering barycenters for probability distributions in [7] seems a promising direction to generalise the results of this article.

Acknowledgements This work was granted access to the Jean Zay supercomputer under the allocation 2022-AD011012212R2.

References

1. Altstidl, T., Nguyen, A., Schwinn, L., Köferl, F., Mutschler, C., Eskofier, B., Zanca, D.: Just a matter of scale? reevaluating scale equivariance in convolutional neural networks. arXiv preprint arXiv:2211.10288 (2022)

2. Andén, J., Mallat, S.: Deep scattering spectrum. *IEEE Transactions on Signal Processing* 62(16), 4114–4128 (2014)
3. Barnard, E., Casasent, D.: Invariance and neural nets. *IEEE Transactions on neural networks* 2(5), 498–508 (1991)
4. Benton, G., Finzi, M., Izmailov, P., Wilson, A.G.: Learning invariances in neural networks from training data. *Advances in neural information processing systems* 33, 17605–17616 (2020)
5. Botev, A., Bauer, M., De, S.: Regularising for invariance to data augmentation improves supervised learning. *arXiv preprint arXiv:2203.03304* (2022)
6. Bronstein, M.M., Bruna, J., Cohen, T., Velicković, P.: Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478* (2021)
7. Cazelles, E., Tobar, F., Fontbona, J.: A novel notion of barycenter for probability distributions based on optimal weak mass transport. *Advances in Neural Information Processing Systems* 34, 13575–13586 (2021)
8. Cesa, G., Lang, L., Weiler, M.: A program to build E(N)-equivariant steerable CNNs. In: *International Conference on Learning Representations* (2022)
9. Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16, 321–357 (2002)
10. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *International conference on machine learning*. pp. 1597–1607. PMLR (2020)
11. Chen, W., Tian, L., Fan, L., Wang, Y.: Augmentation invariant training. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*. pp. 0–0 (2019)
12. Cohen, T., Welling, M.: Group equivariant convolutional networks. In: *International conference on machine learning*. pp. 2990–2999. PMLR (2016)
13. Fort, S., Brock, A., Pascanu, R., De, S., Smith, S.L.: Drawing multiple augmentation samples per image during training efficiently decreases test error. *arXiv preprint arXiv:2105.13343* (2021)
14. Fralick, S.: Learning to recognize patterns without a teacher. *IEEE Transactions on Information Theory* 13(1), 57–64 (1967)
15. Geiping, J., Goldblum, M., Somepalli, G., Shwartz-Ziv, R., Goldstein, T., Wilson, A.G.: How much data are augmentations worth? an investigation into scaling laws, invariance, and implicit regularization. *arXiv preprint arXiv:2210.06441* (2022)
16. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 580–587 (2014)
17. Haasdonk, B., Vossen, A., Burkhardt, H.: Invariance in kernel methods by haar-integration kernels. In: *Image Analysis: 14th Scandinavian Conference, SCIA 2005, Joensuu, Finland, June 19-22, 2005*. *Proceedings* 14. pp. 841–851. Springer (2005)
18. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9729–9738 (2020)
19. Kokkinos, I.: Pushing the boundaries of boundary detection using deep learning. *arXiv preprint arXiv:1511.07386* (2015)
20. Kondor, R., Trivedi, S.: On the generalization of equivariance and convolution in neural networks to the action of compact groups. In: *International Conference on Machine Learning*. pp. 2747–2755. PMLR (2018)

21. Kullback, S.: Information theory and statistics. Courier Corporation (1997)
22. Laptev, D., Savinov, N., Buhmann, J.M., Pollefeys, M.: Ti-pooling: transformation-invariant pooling for feature learning in convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 289–297 (2016)
23. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* 521(7553), 436–444 (2015)
24. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11), 2278–2324 (1998)
25. Mroueh, Y., Voinea, S., Poggio, T.A.: Learning with group invariant features: A kernel perspective. *Advances in neural information processing systems* 28 (2015)
26. Penaud-Polge, V., Velasco-Forero, S., Angulo, J.: Fully trainable Gaussian derivative convolutional layer. In: 2022 IEEE International Conference on Image Processing (ICIP). pp. 2421–2425. IEEE (2022)
27. Penaud-Polge, V., Velasco-Forero, S., Angulo, J.: Genharris-resnet: A rotation invariant neural network based on elementary symmetric polynomials. In: *Scale Space and Variational Methods in Computer Vision*. pp. 149–161. Springer International Publishing, Cham (2023)
28. Procesi, C.: Lie groups: an approach through invariants and representations, vol. 115. Springer (2007)
29. Rota, G.C.: Reynolds operators. In: *Proceedings of Symposia in Applied Mathematics*. vol. 16, pp. 70–83. American Mathematical Society Providence, RI (1964)
30. Sangalli, M., Blusseau, S., Velasco-Forero, S., Angulo, J.: Scale equivariant neural networks with morphological scale-spaces. In: *DGMM*. pp. 483–495. Springer (2021)
31. Sangalli, M., Blusseau, S., Velasco-Forero, S., Angulo, J.: Moving frame net: SE (3)-equivariant network for volumes. In: *NeurIPS Workshop on Symmetry and Geometry in Neural Representations*. pp. 81–97. PMLR (2023)
32. Scudder, H.: Adaptive communication receivers. *IEEE Transactions on Information Theory* 11(2), 167–174 (1965)
33. Sgarro, A.: Informational divergence and the dissimilarity of probability distributions. *Calcolo* 18(3), 293–302 (1981)
34. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. *Journal of big data* 6(1), 1–48 (2019)
35. Sibson, R.: Information radius. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 14(2), 149–160 (1969)
36. Smets, B.M., Portegies, J., Bekkers, E.J., Duits, R.: Pde-based group equivariant convolutional neural networks. *Journal of Mathematical Imaging and Vision* 65(1), 209–239 (2023)
37. Wilk, M.v.d., Bauer, M., John, S., Hensman, J.: Learning invariances using the marginal likelihood. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. pp. 9960–9970 (2018)
38. Worrall, D., Welling, M.: Deep scale-spaces: Equivariance over scale. *Advances in Neural Information Processing Systems* 32 (2019)
39. Yang, S., Dong, Y., Ward, R., Dhillon, I.S., Sanghavi, S., Lei, Q.: Sample efficiency of data augmentation consistency regularization. *AISTATS* (2023)
40. Zhang, H., Cissé, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: *ICLR* (2018)
41. Zhu, X., Ghahramani, Z., Lafferty, J.D.: Semi-supervised learning using Gaussian fields and harmonic functions. In: *Proceedings of the 20th International conference on Machine learning (ICML-03)*. pp. 912–919 (2003)