Towards Diffeomorphism-Equivariant Neural Networks via Canonicalization

Josephine Elisabeth Oettinger[1,2], Zakhar Shumaylov[2], Johannes Bostelmann[1], Jan Lellmann[1], Carola-Bibiane Schönlieb[2]

[1] MIC, University of Luebeck and [2] DAMTP, University of Cambridge

Motivation: Diffeomorphism Equivariance

We have:
- sparse, "canonical" dataset X_E,
- model f_θ trained on a sparse "canonical" dataset.

Goal:
- model that generalizes to diffeomorphic variations of the "canonical" dataset $\rightarrow$ **diffeomorphism equivariance**

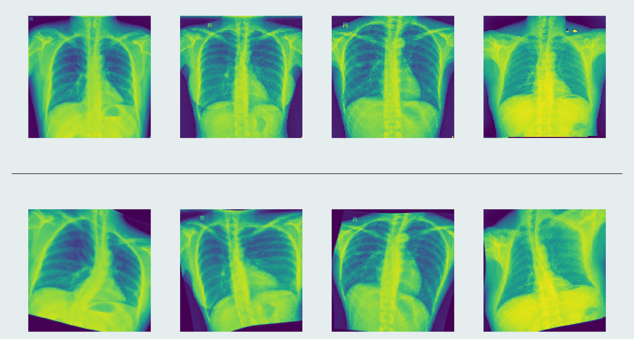

Idea:

CANONICALIZATION
- no extensive data augmentation
- no retraining
- can handle diffeomorphisms

Group Equivariance:

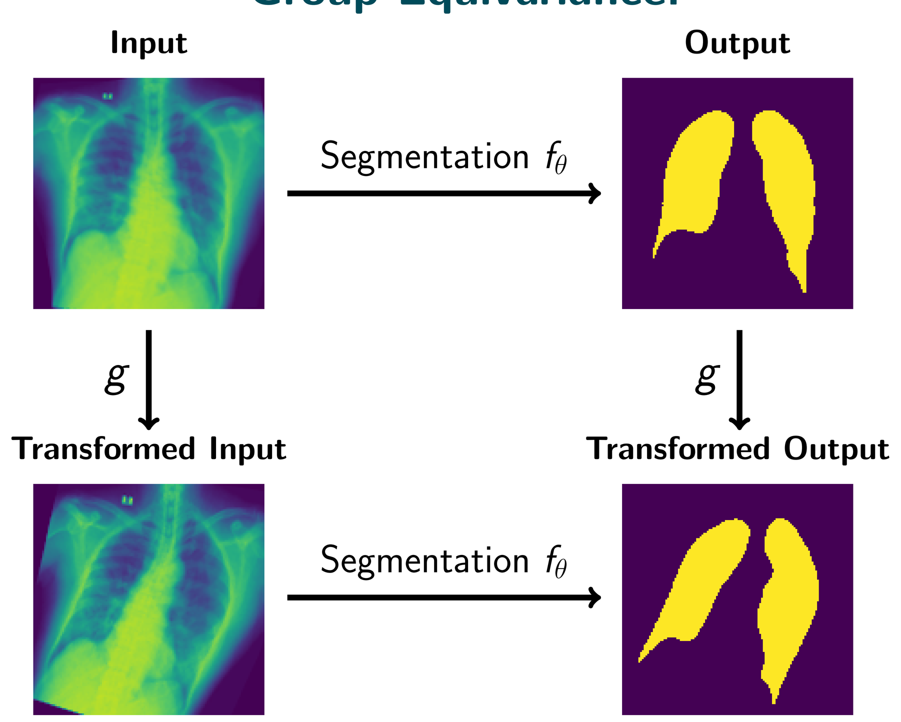

$f(g \cdot x) = g \cdot f(x)$

Group Invariance:

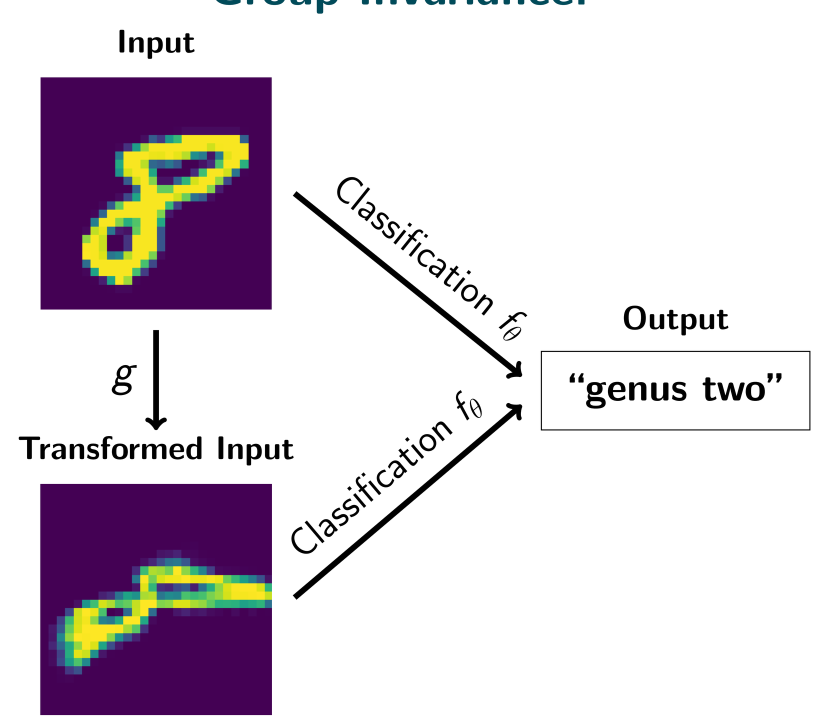

$f(g \cdot x) = f(x)$

Group of Diffeomorphisms:
The set of maps $g : \mathcal{X} \rightarrow \mathcal{X}$ which are **bijective**, **differentiable**, and have a **differentiable inverse**, forms an infinite-dimensional group $\mathcal{D}(\mathcal{X})$ under composition.

Challenges: Create a framework that
- handles the **infinite-dimensional** nature of the group of diffeomorphisms,
- is **computationally feasible** and
- guarantees (approximate) **diffeomorphism equivariance**.

References

[1] Shumaylov, Z., Zaika, P., Rowbottom, J., Sherry, F., Weber, M., and Schönlieb, C.-B. **Lie algebra canonicalization: Equivariant neural operators under arbitrary Lie groups**, 10 2024.

[2] Kaba, S.-O., Mondal, A. K., Zhang, Y., Bengio, Y., and Ravanbakhsh, S. **Equivariance with learned canonicalization functions**, 2023. URL https://arxiv.org/abs/2211.06489.

[3] Bostelmann, J., Gildemeister, O., and Lellmann, J. **Stationary velocity fields on matrix groups for deformable image registration.** Journal of Mathematical Imaging and Vision, 68:9, 2026. ISSN 1573-7683. doi:10.1007/s10851-026-01284-y. URL https://doi.org/10.1007/s10851-026-01284-y.

[4] Han, K., Sun, S., Yan, X., You, C., Tang, H., Naushad, J., Ma, H., Kong, D., Xie, X.: **Diffeomorphic image registration with neural velocity field.** IEEE Winter Conf Appl Comput Vis, 2023.

[5] Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J., and Dalca, A. V. **VoxelMorph: a learning framework for deformable medical image registration.** IEEE Transactions on Medical Imaging, 38(8):1788–1800, 2019. doi:10.1109/TMI.2019.2897538.

DiffeoNN (Diffeomorphism-Equivariant Neural Networks)

- Novel framework to achieve **diffeomorphism equivariance** via **energy-based canonicalization** [1, 2]

DiffeoNN Framework

Energy-Based Canonicalization $\rightarrow$ **"Inner" Model f_θ** $\rightarrow$ **Reverse Canonicalization**

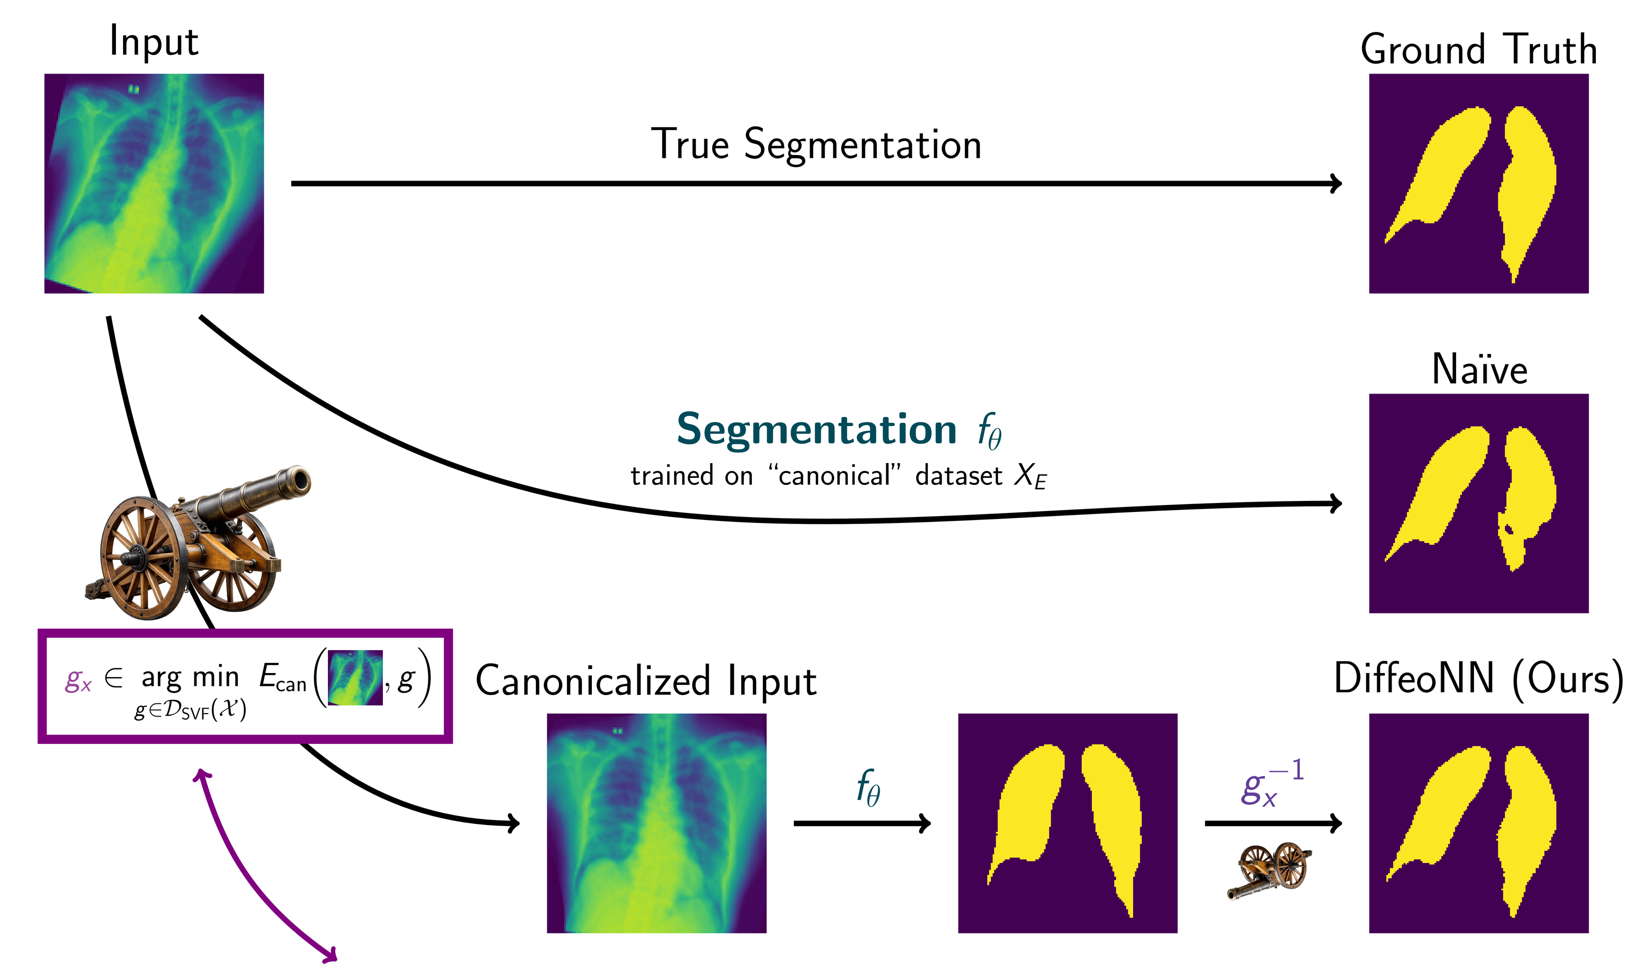

Energy-Based Canonicalization

We have: input x.

We seek: canonical representation x_c that is "close" to the training data X_E.

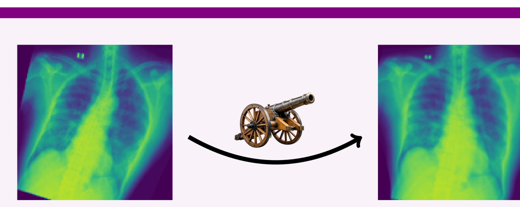

Idea: canonicalize input to $x_c := g_x \cdot x$ with $g_x \in \arg\min_{g \in \mathcal{D}_{SVF}(\mathcal{X})} E_{can}(x, g)$.

Parametrization via Stationary Velocity Fields (SVFs):
- diffeomorphisms as solutions of **flow equation** at time $t = 1$ for SVF $v : \Omega \rightarrow \mathbb{R}^d \in C_0^1$ [3, 4]:

$$\frac{\mathrm{d}g_t(p)}{\mathrm{d}t} = v(g_t(p)),$$
$$g_0(p) = p.$$

Gradient-Based Optimization:
- minimizes the canonicalization energy E_{can} to obtain canonicalizing diffeomorphism $g_x \in \mathcal{D}_{SVF}(\mathcal{X})$
- based on image registration methods from [3, 5]

Learned Canonicalization Energy $E_{can}(x, g)$:
- *small* when $g \cdot x$ is **"close" to training data** X_E and transformation g **physically plausible**,
- *large* otherwise.

$$E_{can}(x, g) := \underbrace{E_{X_E}(g \cdot x)}_{\text{image similarity}} + \underbrace{E_{reg}(g)}_{\text{regularization}}$$

$$E_{X_E}(g \cdot x) := \lambda_{VAE} E_{VAE}(g \cdot x) + \lambda_{adv} E_{adv}(g \cdot x),$$

$$E_{reg}(g) := \underbrace{\lambda_{grad} \sum_{p \in \Omega} \sum_{i=1}^{d} (\nabla v_i(p))^2}_{\text{gradient loss}} + \underbrace{\lambda_{jac} \sum_{p \in \Omega} \max(0, -\det(\mathcal{J}_g(p)))}_{\text{Jacobian determinant loss}}.$$

Results

Diffeomorphism-Equivariant Lung Segmentation:

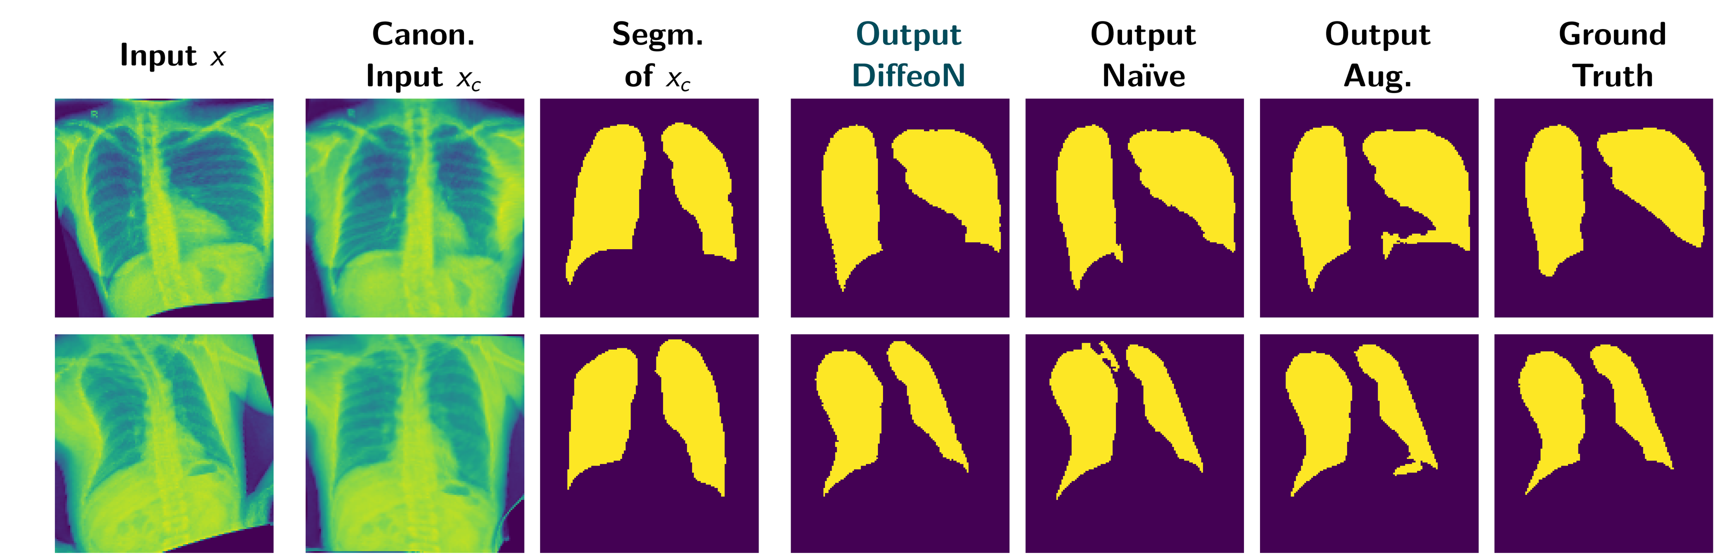

Performance of **DiffeoNN** compared to
- only the inner network of DiffeoNN (**Naïve**)
- data augmentation approach (**Aug.**) as baseline

Model	Acc. ↑
Naïve	0.68
DiffeoNN	**0.82**
Aug.	**0.82**

Table: Mean accuracy (Acc.) of homology classification on diffeomorphically transformed **MNIST**.

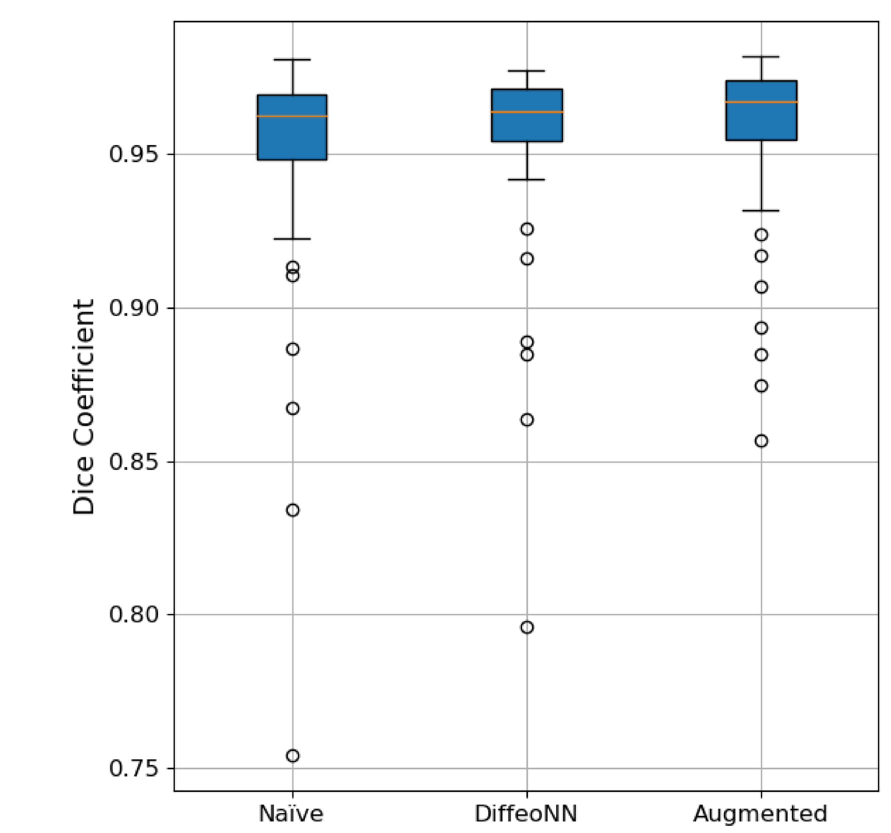

Figure: Dice coefficients of lung segmentation on chest X-rays.

Diffeomorphism-Invariant Homology Classification:

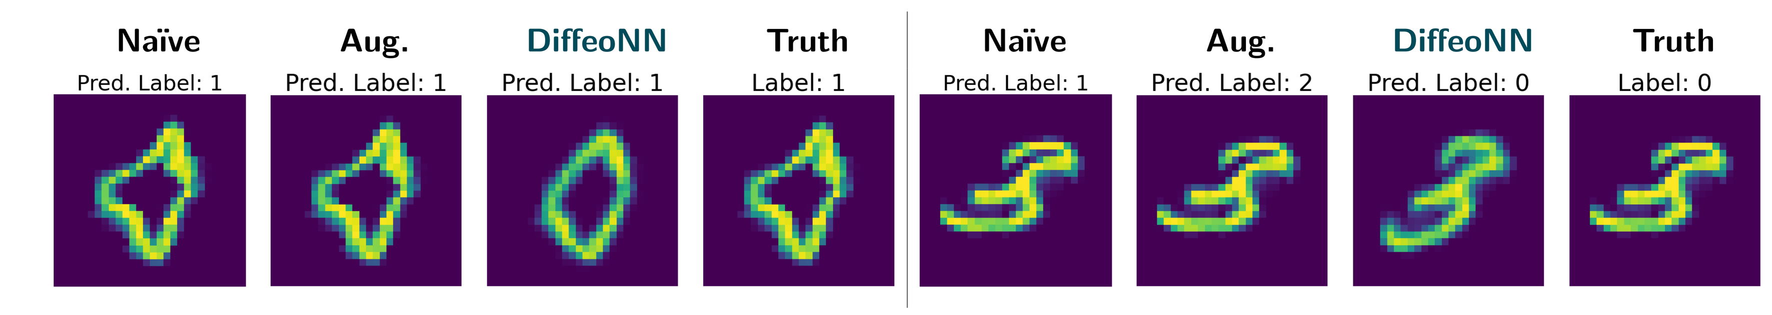

Conclusion

DiffeoNN: flexible and cheap approach using canonicalization to turn arbitrary network diffeomorphism-equivariant

- no extensive data augmentation and no retraining (only training of energy on sparse, "canonical" dataset)
- outperforms naïve and has comparable performance to augmented approach, while dealing slightly better with outliers
- Source code available: https://github.com/mic-uzl/DiffeoNN

https://openreview.net/forum?id=hyK30n4mra