



Diffeomorphism-Equivariant Neural Networks



Josephine Elisabeth Oettinger^{1,2}, Zakhar Shumaylov², Peter Zaika², Jan Lellmann¹, Carola-Bibiane Schönlieb²


¹Institute of Mathematics and Image Computing, University of Lübeck and ²Department of Applied Mathematics and Theoretical Physics, University of Cambridge

17. March 2026

GAMM: Mathematical signal and image processing

Invariance and Equivariance

G-Invariant




$$f(g(x)) = f(x);$$

data/images: [LeCun et al., 1998], [RSUA, 2023]

G-Equivariant

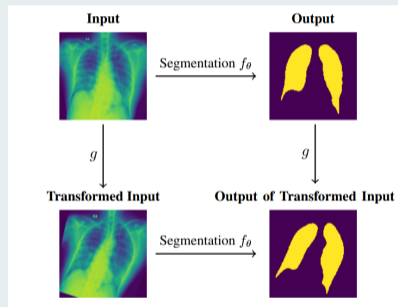
Invariance and Equivariance

G-Invariant



$$f(g(x)) = f(x);$$

G-Equivariant



$$f(g(x)) = g(f(x));$$

data/images: [LeCun et al., 1998], [RSUA, 2023]

Di eomorphisms

Definition (Diffeomorphism)


X and Y di erentiable manifolds.

A map $g : X \rightarrow Y$ is called a di eomorphism if

- g is bijective
- g is di erentiable
- inverse $g^{-1} : Y \rightarrow X$ is di erentiable


When $X = Y$, this set forms an infinite-dimensional group under composition, $D(X)$.


Training Data X_E

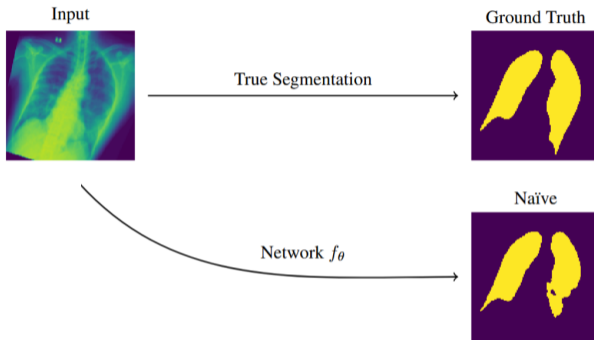


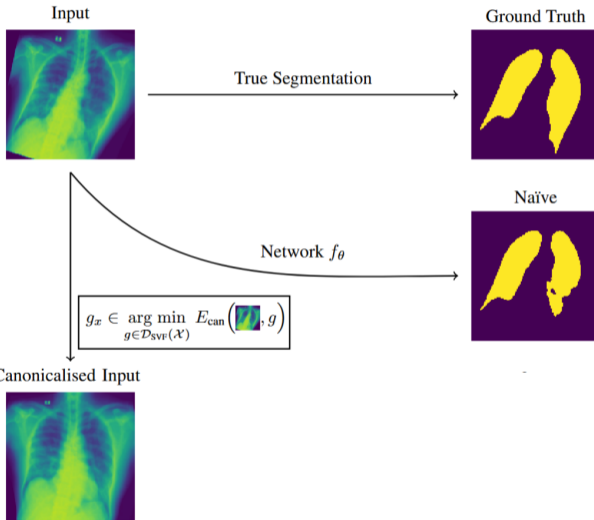
data: [Karmakar, 2024]

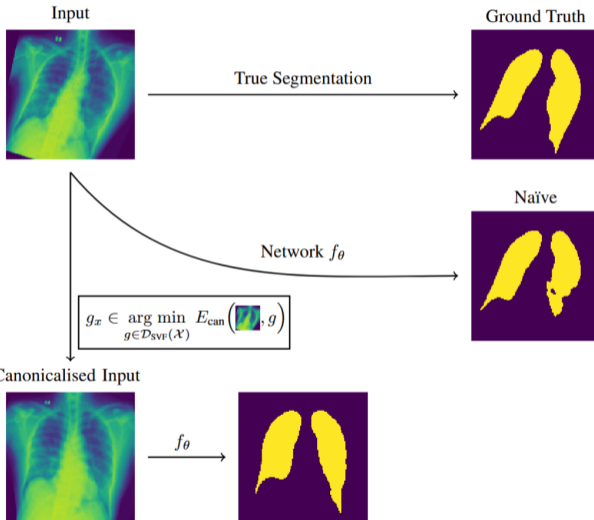
Diffomorphically Transformed Images

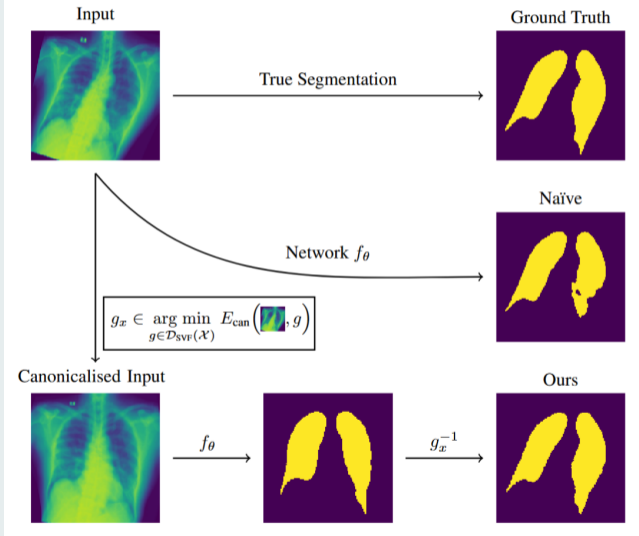














Canonicalisation

We have: input x .



Canonicalisation

We have: input x .

We want: canonical representation x_C , that is “close” to the training data X_E .

Canonicalisation

We have: input x .

We want: canonical representation x_c , that is “close” to the training data X_E .

Idea: [energy-based canonicalisation](#) [Shumaylov et al., 2024],
canonicalise input to $x_c := g_x \cdot x$ with

$$g_x \in \arg \min_{g \in D_{SVF}(X)} E_{\text{can}}(x; g):$$

Energy minimisation through gradient-based optimisation following [Bostelmann et al., 2024].



Canonicalisation Energy E_{can}

Intuitively, energy $E_{\text{can}}(x; g)$ small when $g \circ x$ is “close” to training data X_E and transformation g physically plausible, large otherwise.

Canonicalisation Energy E_{can}

Intuitively, energy $E_{\text{can}}(x; g)$ small when $g \circ x$ is “close” to training data X_E and transformation g physically plausible, large otherwise.

$$E_{\text{can}} : X \times D_{\text{SVF}}(X) \rightarrow \mathbb{R}; \quad E_{\text{can}}(x; g) := \underbrace{E_{X_E}(g \circ x)}_{\text{image similarity}} + \underbrace{E_{\text{reg}}(g)}_{\text{regularisation}};$$

Canonicalisation Energy E_{can}

Intuitively, energy $E_{\text{can}}(x; g)$ small when $g \circ x$ is “close” to training data X_E and transformation g physically plausible, large otherwise.

$$E_{\text{can}} : X \times D_{\text{SVF}}(X) \rightarrow \mathbb{R}; \quad E_{\text{can}}(x; g) := \underbrace{E_{X_E}(g \circ x)}_{\text{image similarity}} + \underbrace{E_{\text{reg}}(g)}_{\text{regularisation}};$$

with

$$E_{X_E}(g \circ x) := E_{\text{VAE}}(g \circ x) + E_{\text{adv}}(g \circ x);$$

Canonicalisation Energy E_{can}

Intuitively, energy $E_{\text{can}}(x; g)$ small when $g \circ x$ is close to training data X and transformation g physically plausible, large otherwise.

$$E_{\text{can}} : X \times \text{D}_{\text{SVF}}(X) \rightarrow \mathbb{R}; \quad E_{\text{can}}(x; g) := \underbrace{E_{X_E}(g \circ x)}_{\text{image similarity}} + \underbrace{E_{\text{reg}}(g)}_{\text{regularisation}};$$

with

$$E_{X_E}(g \circ x) := \text{VAE } E_{\text{VAE}}(g \circ x) + \text{adv } E_{\text{adv}}(g \circ x);$$

$$E_{\text{reg}}(g) := \underbrace{\int_{\text{grad}} \sum_{i=1}^d (r_{v_i}(p))^2}_{\text{gradient loss}} + \underbrace{\int_{\text{jac}} \max(0; \det(J_g(p)))}_{\text{Jacobian determinant loss}};$$

SVF-Based Diffeomorphisms

Stationary Velocity Field (SVF): vector field $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$,

$\hat{v}(p)$: velocity of a fluid at position p .

Flow: $\gamma_t(p) : \mathbb{R} \rightarrow \mathbb{R}^d$,

defined as solution to ODE:

$$\frac{d}{dt} \gamma_t(p) = v(\gamma_t(p)); \quad \gamma_0(p) = p; \quad t \in [0; 1]:$$

SVF-Based Diffeomorphisms

Stationary Velocity Field (SVF): vector field $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$,

$\hat{v}(p)$: velocity of a fluid at position p .

Flow: $\gamma_t(p) : \mathbb{R} \rightarrow \mathbb{R}^d$,

$\hat{\gamma}_t$ defined as solution to ODE:

$$\frac{d}{dt} \gamma_t(p) = v(\gamma_t(p)); \quad \gamma_0(p) = p; \quad t \in [0, 1]:$$

SVF-Based Diffeomorphisms

Stationary Velocity Field (SVF): vector field $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$,

$\hat{v}(p)$: velocity of a fluid at position p .

Flow: $\gamma_t(p) : \mathbb{R} \rightarrow \mathbb{R}^d$,

$\hat{\gamma}_t$ defined as solution to ODE:

$$\frac{d}{dt} \gamma_t(p) = v(\gamma_t(p)); \quad \gamma_0(p) = p; \quad t \in [0, 1]:$$

SVF-Based Diffeomorphisms

Stationary Velocity Field (SVF): vector field $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$,

$\hat{v}(p)$: velocity of a fluid at position p .

Flow: $\gamma_t(p) : \mathbb{R} \rightarrow \mathbb{R}^d$,

$\hat{\gamma}_t$ defined as solution to ODE:

$$\frac{d}{dt} \gamma_t(p) = v(\gamma_t(p)); \quad \gamma_0(p) = p; \quad t \in [0, 1]:$$

SVF-Based Diffeomorphisms

Stationary Velocity Field (SVF): vector field $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$,

$\hat{v}(p)$: velocity of a fluid at position p .

Flow: $\phi_t(p) : \mathbb{R}^d \rightarrow \mathbb{R}^d$,

defined as solution to ODE:

$$\frac{d}{dt} \phi_t(p) = v(\phi_t(p)); \quad \phi_0(p) = p; \quad t \in [0; 1]:$$

Diffeomorphism:

$\hat{\phi}_1$ solution of ODE at $t = 1$: $\hat{\phi}_1 = \exp(v)$,

$\hat{\exp} : \text{SVF} \rightarrow \text{D}_{\text{SVF}}(X)$,

structured and computationally efficient subset of full group of diffeomorphic transformations

Di eomorphism-Equivariant Lung Segmentation

Input x

Canon.
Input x_c

Segm.
of x_c

Output
Di eoNN

Output
Naïve

Output
Aug.

Ground
Truth

data: [Karmakar, 2024]

Diffomorphism-Equivariant Lung Segmentation

| Model | IoU " | Dice " | Acc. " |
|----------------|--------|--------|--------|
| Naïve | 0.9147 | 0.9547 | 0.9755 |
| DiffoNN (ours) | 0.9214 | 0.9586 | 0.9777 |
| Aug. | 0.9251 | 0.9606 | 0.9790 |

Table: Mean Intersection-over-Union (IoU), Dice coefficient (Dice), and pixel-wise accuracy (Acc.) on diffeomorphically transformed chest X-rays.

Di eomorphism-Invariant Homology Classi cation

Naïve

Aug.

Di eoNN

Ground
Truth

Naïve

Aug.

Di eoNN

Ground
Truth

data: [LeCun et al., 1998]

Di eomorphism-Invariant Homology Classification

| Model | Acc. " |
|-----------------|--------|
| Naïve | 0.68 |
| DiffeoNN (ours) | 0.82 |
| Aug. | 0.82 |

Table: Mean accuracy (Acc.) on di eomorphically transformed MNIST.

Bounding the Generalisation Error

For energy-based canonicalisation in LieLAC setup [Shumaylov et al., 2024]:
(compact and finite-dimensional setting):

$$\mathbb{E}_{\tau} [L(f(g_x, x)); g_x, y_x^{\text{gt}}] \leq \hat{L} + (L_{\max} \epsilon);$$

with

- \hat{L} bounds training error,
- L_{\max} maximal possible loss,
- ϵ small, when X_E well sampled (probability mass on poorly-sampled group orbits),
- ϵ under mild assumption.

Proof Idea

$$\begin{aligned}
 & \mathbb{E}_{\tau} [L(f_{\mathcal{Z}}(g_x(x)); g_x, y_x^{\text{gt}})] \\
 &= \int_{\mathcal{Z}^X} L(f(g_x(x)); y_{g_x(x)}^{\text{gt}}) d_{\tau}(x) \\
 &= \int_{\mathcal{Z}} L(f(g_x(x)); y_{g_x(x)}^{\text{gt}}) d_{\tau}(x) \\
 & \quad \left| \frac{A}{\mathcal{Z}} \right. \left. \frac{\{ \mathcal{Z} \}}{(1 - \tau(A^c))} \right. \\
 & \quad + \int_{\mathcal{Z}} L(f(g_x(x)); y_{g_x(x)}^{\text{gt}}) d_{\tau}(x) \\
 & \quad \left| \frac{A^c}{\tau(A^c)L_{\max}} \right. \left. \frac{\{ \mathcal{Z} \}}{\tau(A^c)L_{\max}} \right. \\
 & + (L_{\max})
 \end{aligned}$$

Bounding the Generalisation Error

$$E_{\tau} [L(f(g_x, x)); g_x, y_x^{\text{gt}}] \\ + (L_{\text{max}})$$

Bounding the Generalisation Error

$$E_{\tau} [L(f(g_x, x)); g_x, y_x^{\text{gt}}] \\ + (L_{\text{max}})$$

Bounding the Generalisation Error

$$E_{\tau}[L(f(g_x x)); g_x y_x^{gt}] + (L_{\max})$$

$$E_{\tau}[L(g_x^{-1}(f(g_x x)); y_x^{gt})] \\ \leq C E_{\tau}[L(f(g_x x); g_x y_x^{gt})]:$$

Di eomorphism-Equivariant Neural Networks (Di eoNN)

- ^ turn any network
di eomorphism-equivariant
- ^ no data augmentation or retraining
necessary (only training on \mathcal{E})
- ^ bounded generalisation error








Bounding Generalisation Error - Measures

Lung Segmentation: Failed Canonicalisation

| Input x | Can. Input x_c | Seg. of x_c | Output Diform-NN | Output Naïve | Output Aug. | Ground -Truth |
|-----------|---------------------|------------------|---------------------|-----------------|----------------|------------------|
| | | | | | | |

References

-  [Shumaylov et al., 2024] Zakhar Shumaylov, Peter Zaika, James Rowbottom, Ferdia Sherry, Melanie Weber, and Carola-Bibiane Schönlieb. “Lie algebra canonicalization: Equivariant neural operators under arbitrary lie groups.”, 2024. arXiv preprint arXiv:2410.02698.
-  [Bostelmann et al., 2024] Johannes Bostelmann, Ole Gildemeister, and Jan Lellmann. “Stationary velocity fields on matrix groups for deformable image registration.”, 2024. arXiv preprint arXiv:2410.10997.
-  [LeCun et al., 1998] LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. “Gradient-based learning applied to document recognition.” Proc. IEEE, 86:2278–2324, 1998. URL <https://api.semanticscholar.org/CorpusID:14542261>.
-  [RSUA, 2023] RSUA, R. RSUA Chest X-Ray Dataset. Version V1. 2023. doi: 10.17632/2jg8vfdmpm. url: <https://doi.org/10.17632/2jg8vfdmpm.1>.
-  [Karmakar, 2024] Karmakar, Tapendu. Chest X-ray Dataset for Lung Segmentation. Kaggle, 2024, <https://www.kaggle.com/datasets/iامتapendu/chest-x-ray-lungs-segmentation>. Accessed 13 Feb. 2026.