enjoy Vienna

EUROGRAPHICS 2021

Tutorial:

Inverse Computational Spectral Geometry

3/4

Simone Melzi **Luca Cosmo** Emanuele Rodolà Maks Ovsjanikov Michael Bronstein

Outline

- **The problem of shape from sound**
- **Isospectralization:** numerical optimization technique
- **Applications:** matching, style transfer and universal adversarial attacks
- **Data driven approach**

Shape from sound

- Can we infer the boundary of a flat membrane just from the frequencies it emits?

1966

CAN ONE HEAR THE SHAPE OF A DRUM?

MARK KAC, The Rockefeller University, New York

To George Eugene Uhlenbeck on the occasion of his sixty-fifth birthday

"La Physique ne nous donne pas seulement l'occasion de résoudre des problèmes . . . , elle nous fait presentir la solution." H. POINCARÉ.

Before I explain the title and introduce the theme of the lecture I should like to state that my presentation will be more in the nature of a leisurely excursion than of an organized tour. It will not be my purpose to reach a specified destination at a scheduled time. Rather I should like to allow myself on many

Shape from sound

- Can we infer the boundary of a flat membrane just from the frequencies it emits?

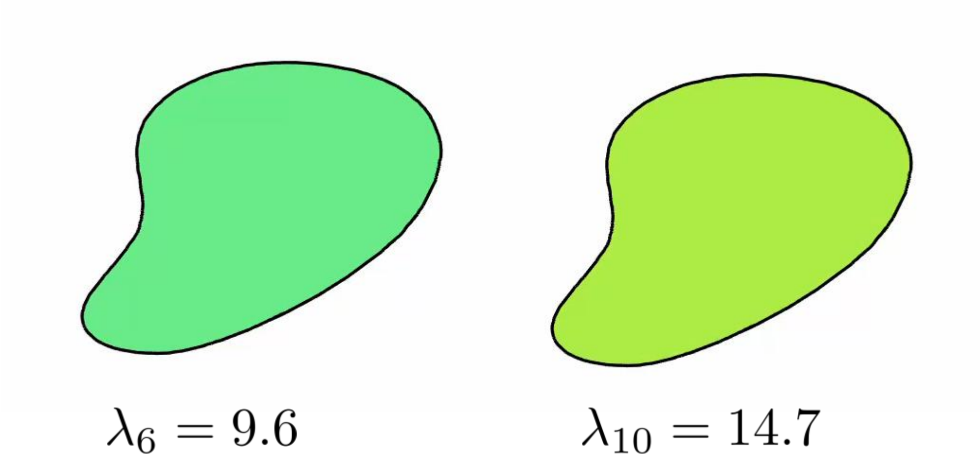

- Described by the wave equation $z = f\,(\mathrm{x,y,t})$

$$\frac{\partial^2 f}{\partial t^2} = c^2 \nabla^2 f$$

- Spatial frequencies are the eigenvalues of the Laplacian

Shape from sound

- Can we reconstruct a 3D mesh from the eigenvalues sequence of its Laplace Beltrami Operator?

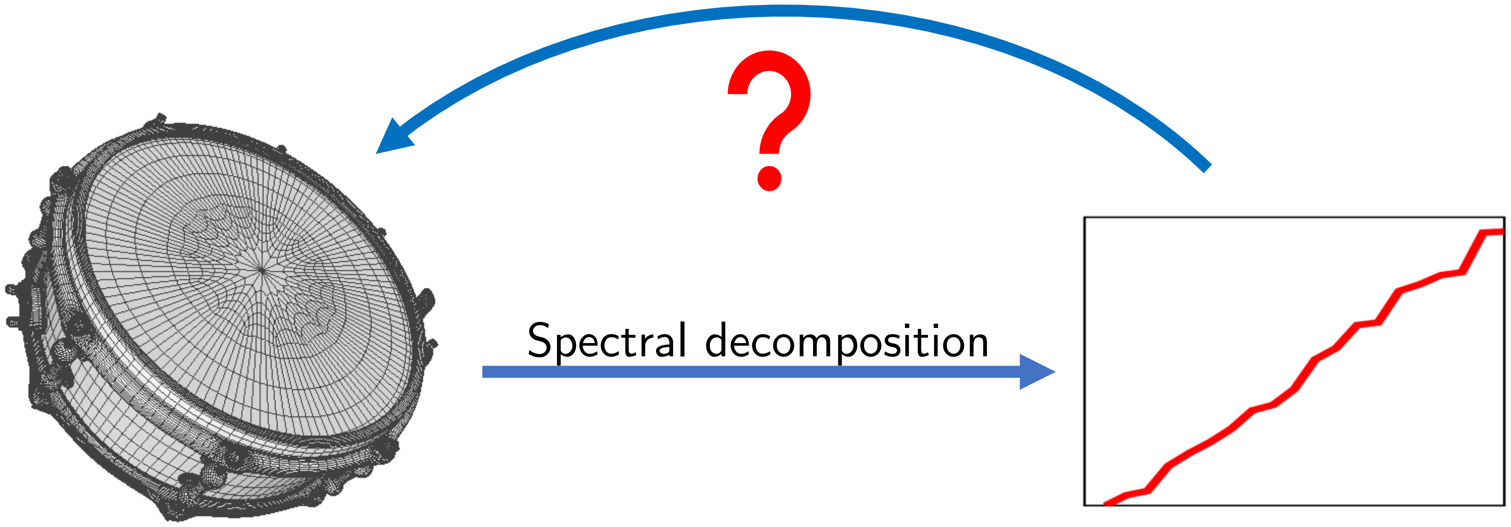

Shape from sound

- Can we reconstruct a 3D mesh from the eigenvalues sequence of its Laplace Beltrami Operator?

$$\arg\min_{\mathcal{X}} \left\|\boldsymbol{\lambda}\left(\boldsymbol{\Delta}_{\mathcal{X}}\right) - \boldsymbol{\mu}\right\|_{\omega}$$

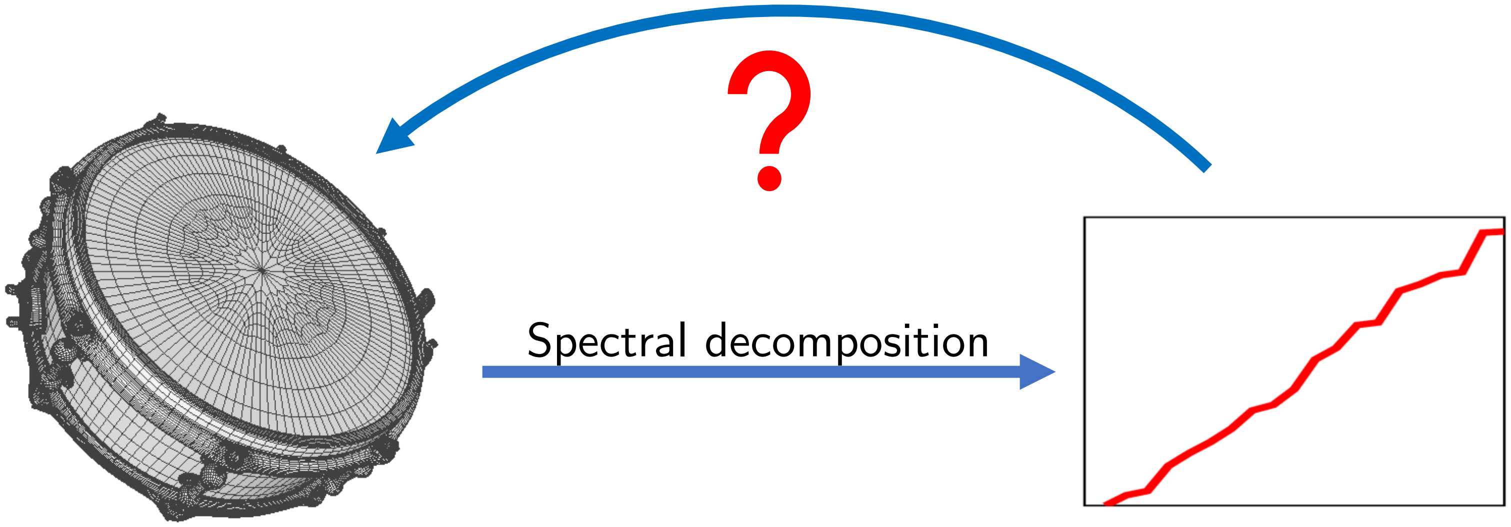

Shape from sound

- Shape from sound: toward new tools for quantum gravity *(Aesen et al. 2013)*
  - Manifold discretized as a star-shaped polyhedra

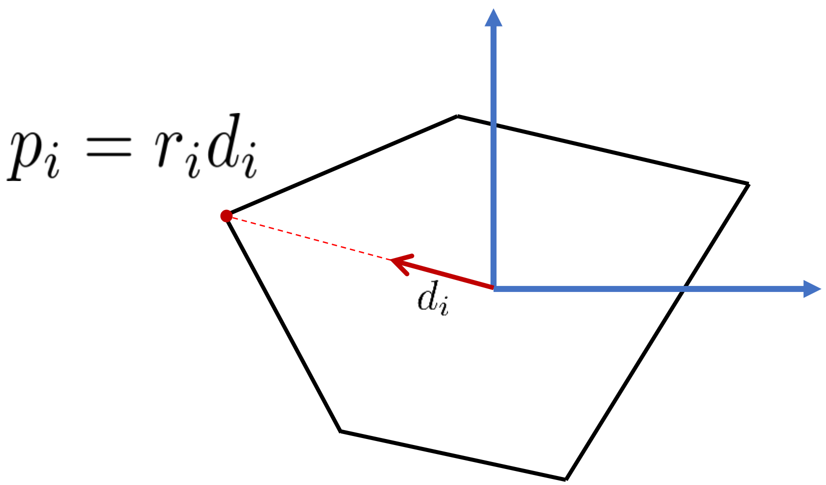

$$\underset{\{r_i\}}{\arg\min} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}(\{r_i\})) - \boldsymbol{\mu}\|_2^2$$

Shape from sound

- Shape from sound: toward new tools for quantum gravity *(Aesen et al. 2013)*
  - Manifold discretized as a star-shaped polyhedra
  - Metric of a 2-dimesnional manifold can be differentiated w.r.t. its eigenvalues

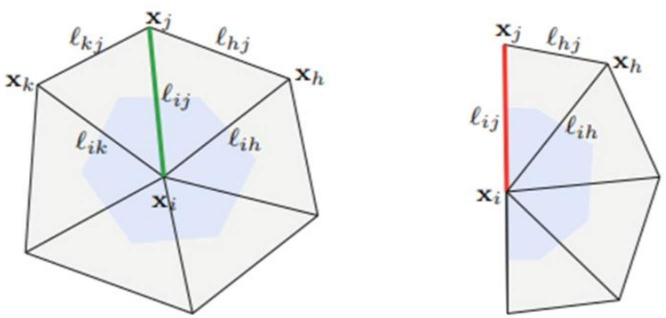

Cotangent Laplacian $\mathbf{\Delta} = \mathbf{A}^{-1}\mathbf{W}$ expressed in terms of discrete metric $\ell_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$ where

$$w_{ij} = \begin{cases} \frac{-\ell_{ij}^2+\ell_{jk}^2+\ell_{ki}^2}{8A_{ijk}} + \frac{-\ell_{ij}^2+\ell_{jh}^2+\ell_{hi}^2}{8A_{ijh}} & \text{if } e_{ij} \in \mathcal{E}_\text{i} \\ \frac{-\ell_{ij}^2+\ell_{jh}^2+\ell_{hi}^2}{8A_{ijh}} & \text{if } e_{ij} \in \mathcal{E}_\text{b} \\ -\sum_{k \neq i} w_{ik} & \text{if } i = j \end{cases} \qquad \mathbf{A} = \begin{pmatrix} a_1 & & \\ & \ddots & \\ & & a_n \end{pmatrix}$$

where A_{ijk} is area of triangle ijk and $a_i = \frac{1}{3} \sum_{ijk:ij,ik \in \mathcal{E}} A_{ijk}$

$$\underset{\{r_i\}}{\arg\min} \|\boldsymbol{\lambda}(\mathbf{\Delta}(\{r_i\})) - \boldsymbol{\mu}\|_2^2$$

- Gradient descent step:

$$r_j^{(t+1)} = r_j^{(t)} - \gamma \frac{\nabla}{\nabla r_j} \left\| \boldsymbol{\lambda}\left(\mathbf{\Delta}(\{r_i^{(t)}\})\right) - \boldsymbol{\mu} \right\|_2^2$$

Shape from sound

- Shape from sound: toward new tools for quantum gravity *(Aesen et al. 2013)*
  - Manifold discretized as a star-shaped polyhedra
  - Metric of a 2-dimesnional manifold can be differentiated w.r.t. its eigenvalues

$$x_i = r_i d_i$$

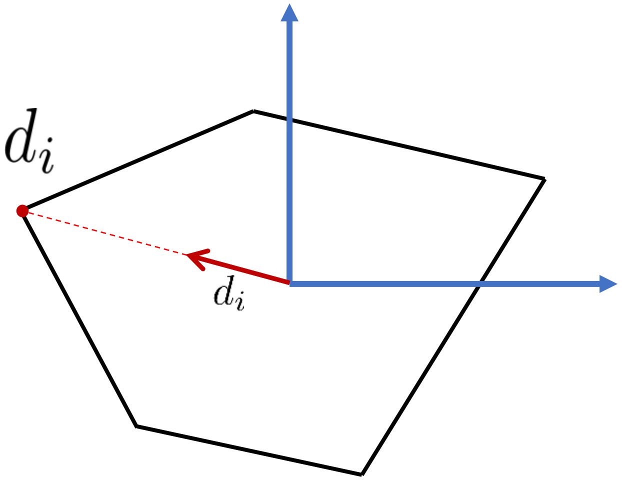

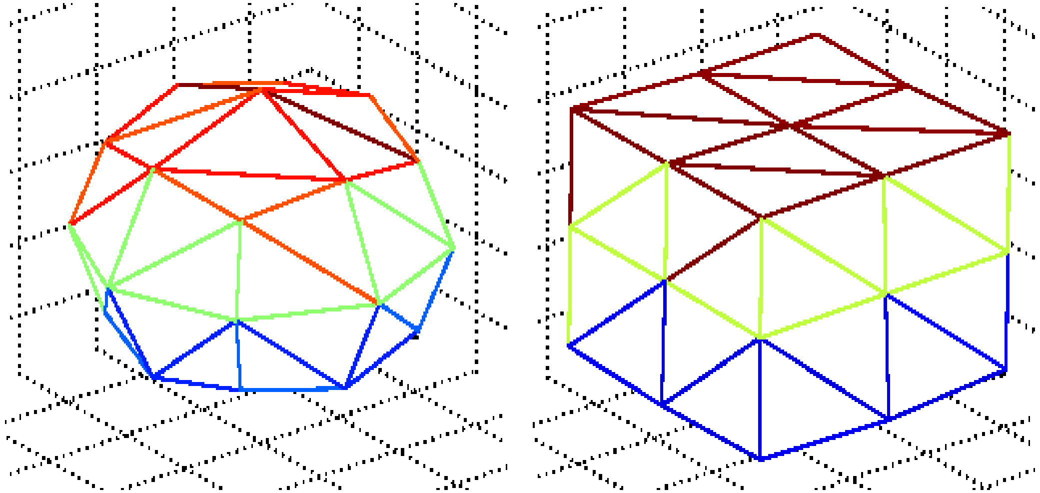

Shape from sound

- Isospectral != Isometric

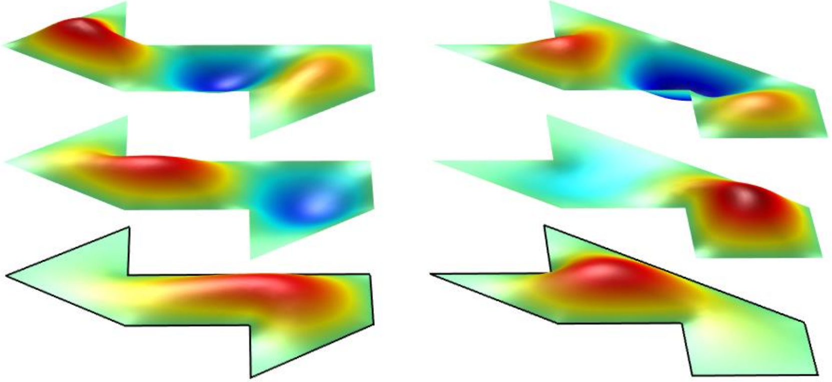

- Metric priors are not enough

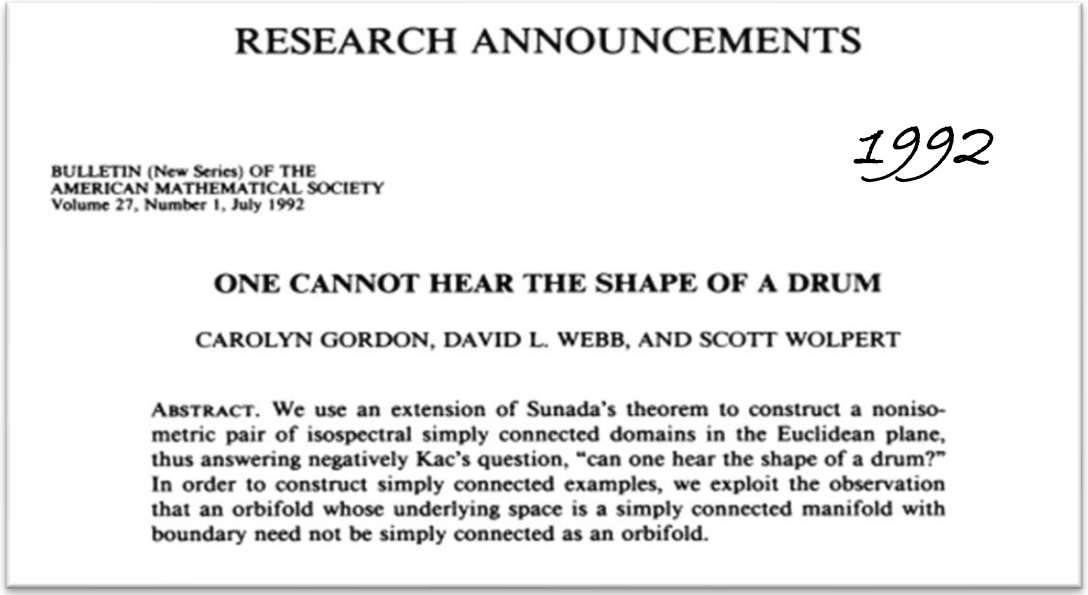

RESEARCH ANNOUNCEMENTS

1992

BULLETIN (New Series) OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 27, Number 1, July 1992

ONE CANNOT HEAR THE SHAPE OF A DRUM

CAROLYN GORDON, DAVID L. WEBB, AND SCOTT WOLPERT

ABSTRACT. We use an extension of Sunada's theorem to construct a nonisometric pair of isospectral simply connected domains in the Euclidean plane, thus answering negatively Kac's question, "can one hear the shape of a drum?" In order to construct simply connected examples, we exploit the observation that an orbifold whose underlying space is a simply connected manifold with boundary need not be simply connected as an orbifold.

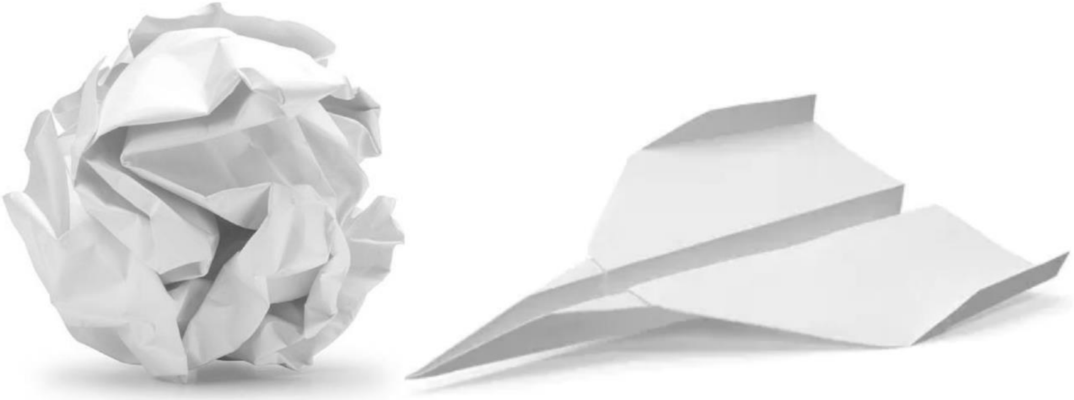

Isospectralization*

$$\min_{\mathbf{X}\in\mathbb{R}^{n\times d}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}(\mathbf{X})) - \boldsymbol{\mu}\|_\omega + \rho_X(\mathbf{X})$$

- Optimization directly on the 3D coordinates
- Data term: Weighted norm (frequency balancing)

$$\|\boldsymbol{\lambda} - \boldsymbol{\mu}\|_\omega^2 = \sum_{i=1}^{k} \frac{1}{\mu_i^2}(\lambda_i - \mu_i)^2$$

* Cosmo et al. *Isospectralization, or how to hear shape, style, and correspondence. CVPR 2019*

Isospectralization

$$\min_{\mathbf{X}\in\mathbb{R}^{n\times d}} \|\boldsymbol{\lambda}(\boldsymbol{\Delta}(\mathbf{X})) - \boldsymbol{\mu}\|_{\omega} + \rho_X(\mathbf{X})$$

- Regularizers to promote smoothness / maximize volume

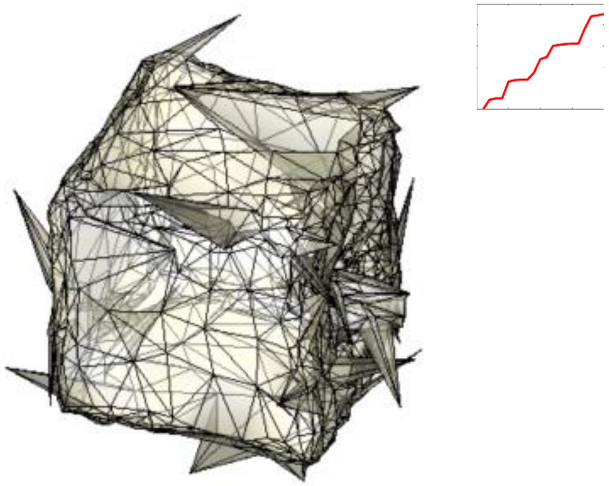

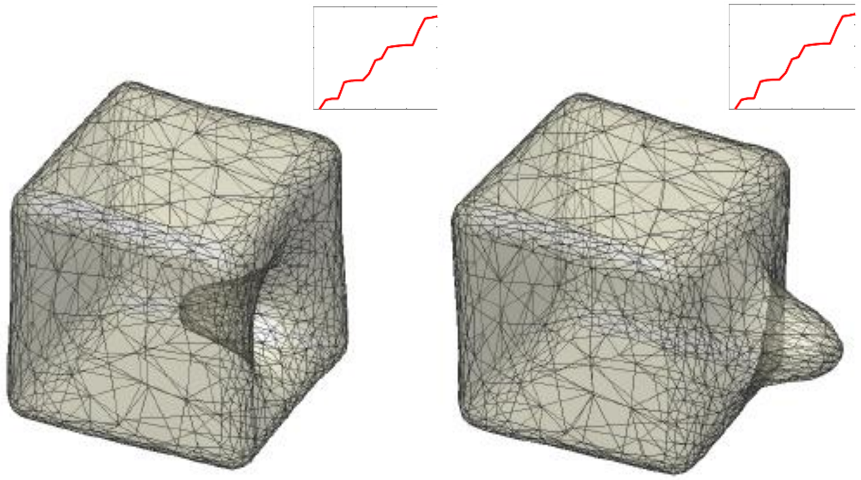

Isospectralization

- 2D shapes:

$$\rho_{X,1}(\mathbf{X}) = \sum_{e_{ij} \in E_b} \ell_{ij}^2(\mathbf{X}) \qquad \rho_{X,2}(\mathbf{X}) = \left(\sum_{ijk \in F} \left(\mathbf{R}_{\frac{\pi}{2}} \left(\mathbf{x}_j - \mathbf{x}_i \right) \right)^\top \left(\mathbf{x}_k - \mathbf{x}_i \right) \right)$$

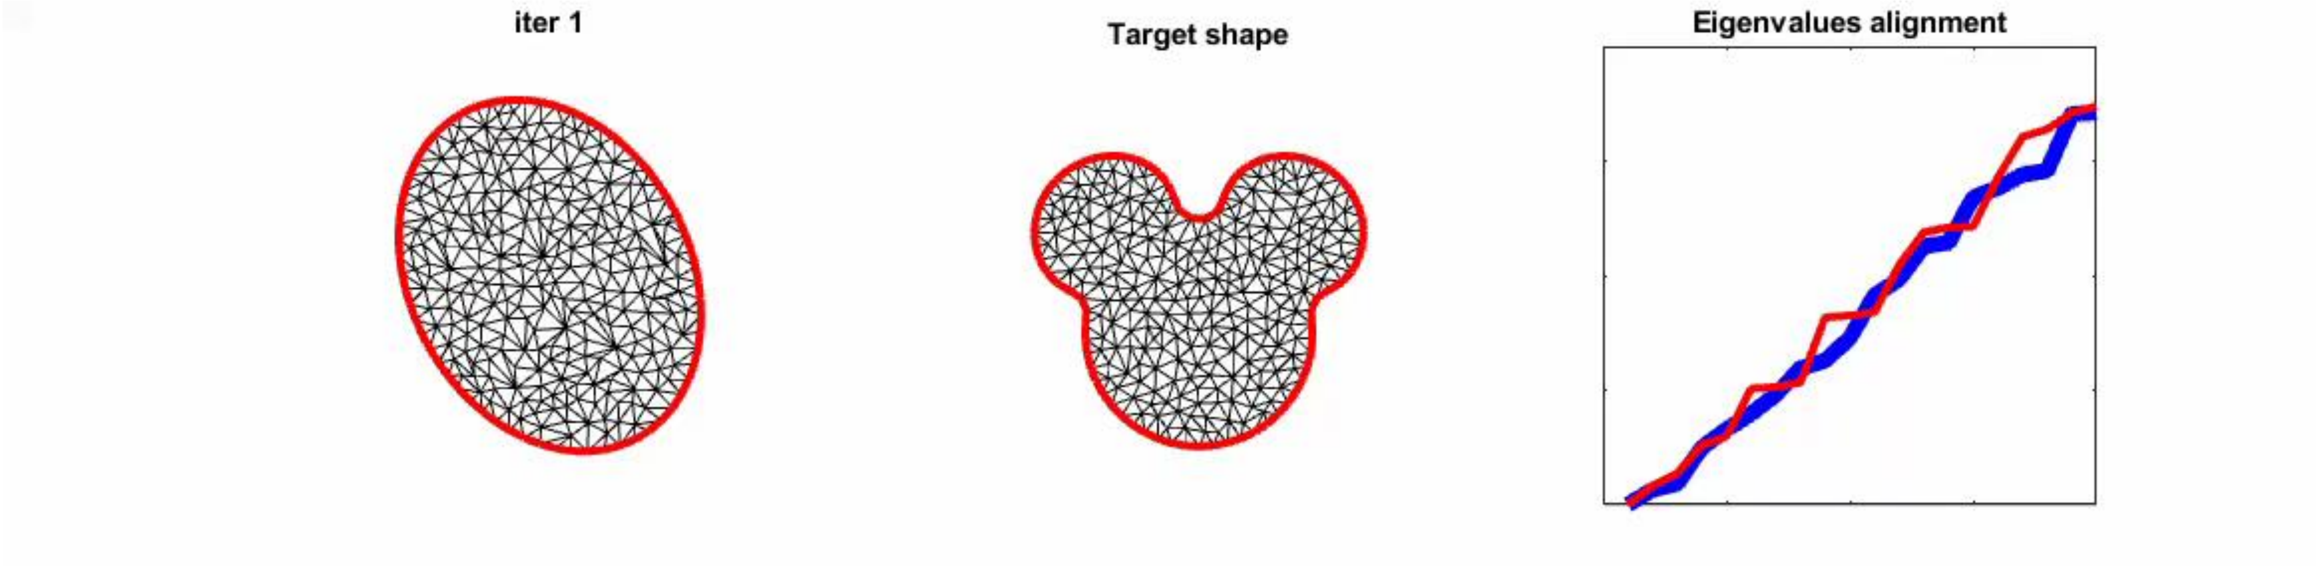

Isospectralization

- 3D shapes:

$$\rho_{X,1}(\mathbf{X}) = \|\boldsymbol{\Delta}(\mathbf{X})\mathbf{X}\|_F^2$$

$$\rho_{X,2}(\mathbf{V}) = -\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}^\top \sum_{ijk \in F} \left((\mathbf{x}_j - \mathbf{x}_i) \times (\mathbf{x}_j - \mathbf{x}_k) \right) (\mathbf{x}_i + \mathbf{x}_j + \mathbf{x}_k)$$

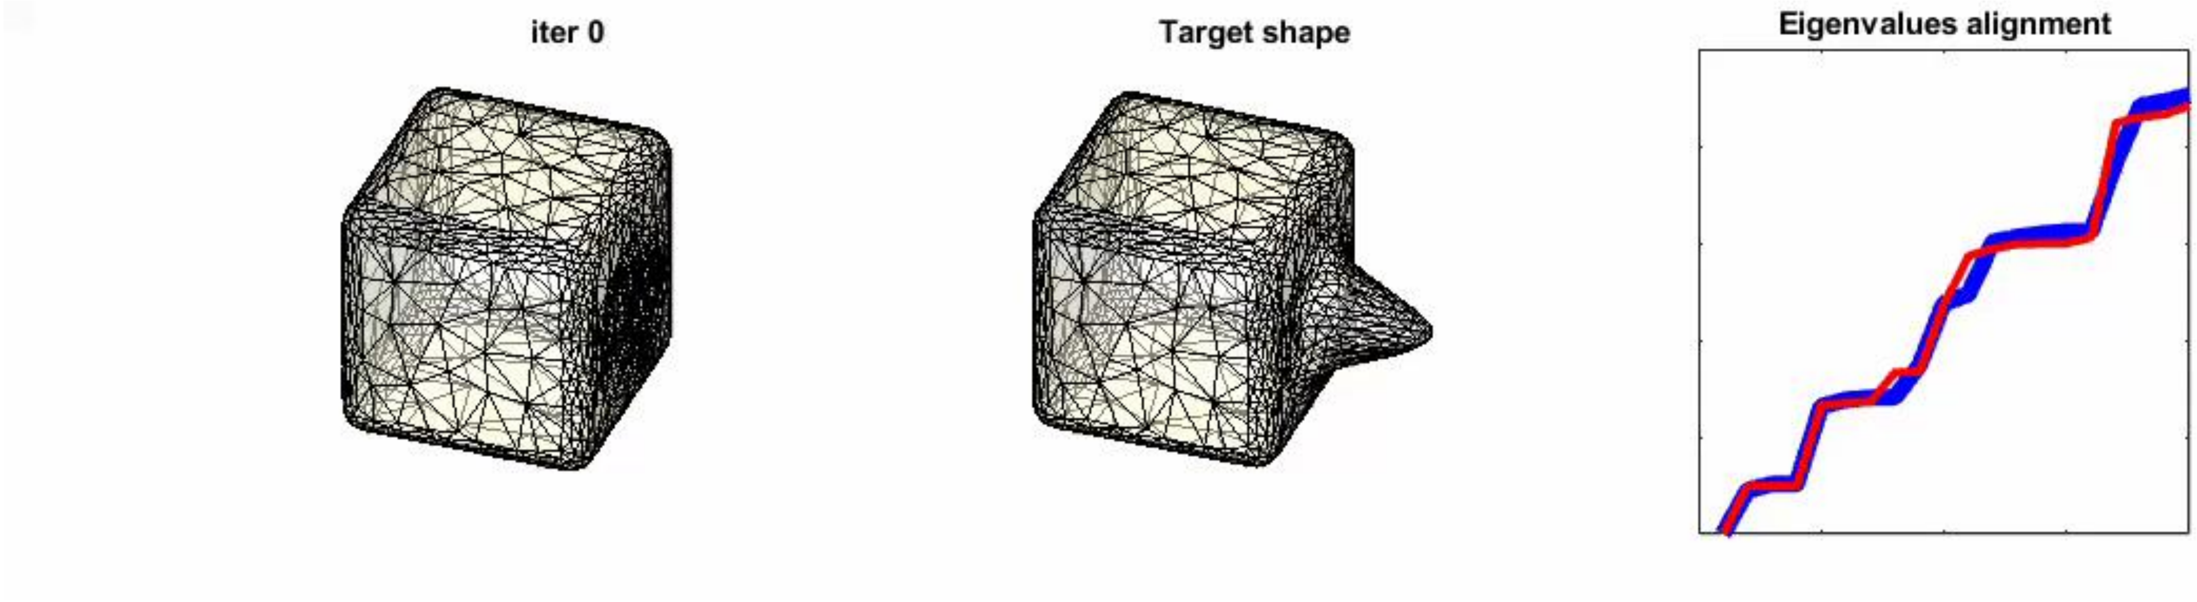

Isospectralization: Applications

- Preprocessing step in Functional Map based matching algorithms for non-isometric shapes

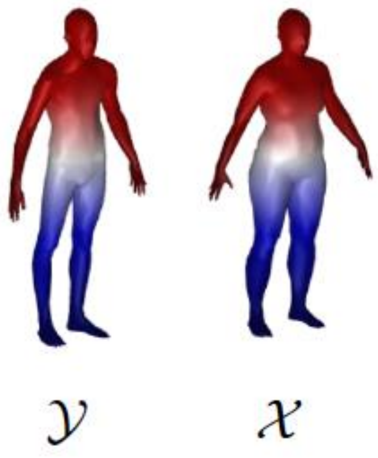

Isospectralization: Applications

- Preprocessing step in Functional Map based matching algorithms for non-isometric shapes

Isospectralization: Applications

- Preprocessing step in Functional Map based matching algorithms for (highly) non-isometric shapes

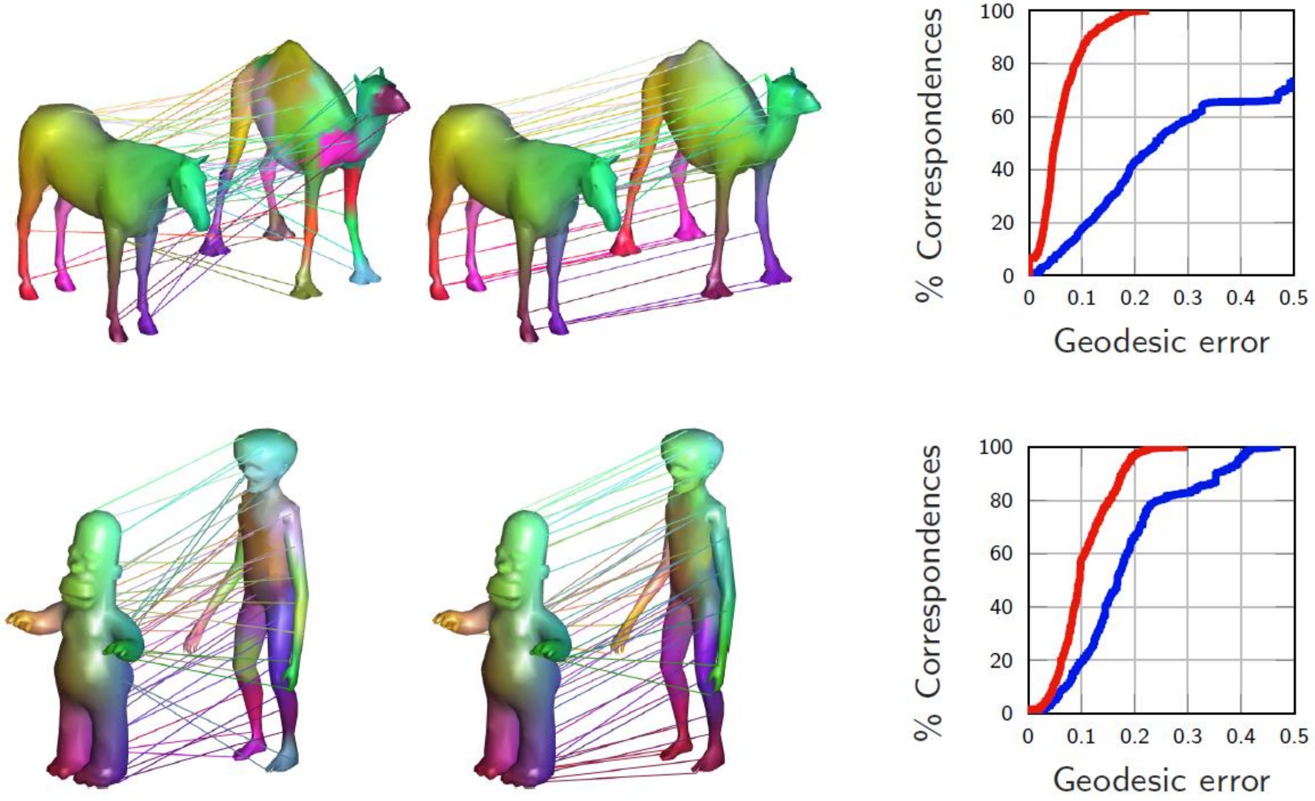

Isospectralization: Applications

- Isospectralization induces isometry

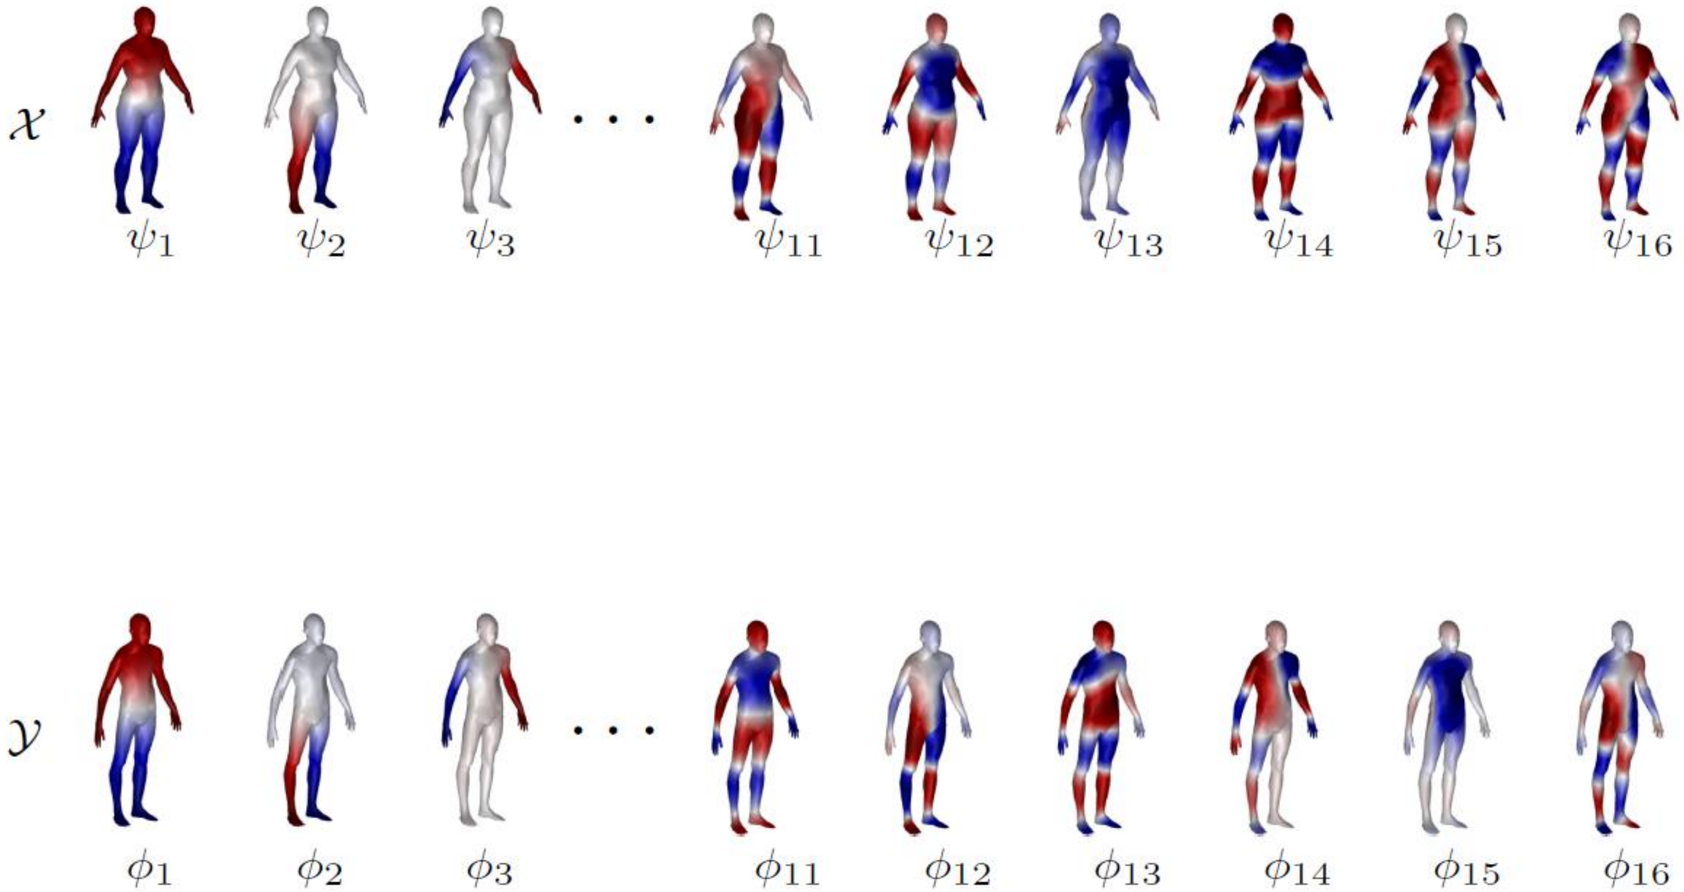

Isospectralization: Applications

- Style transfer
  Eigenvalues do not encode pose information

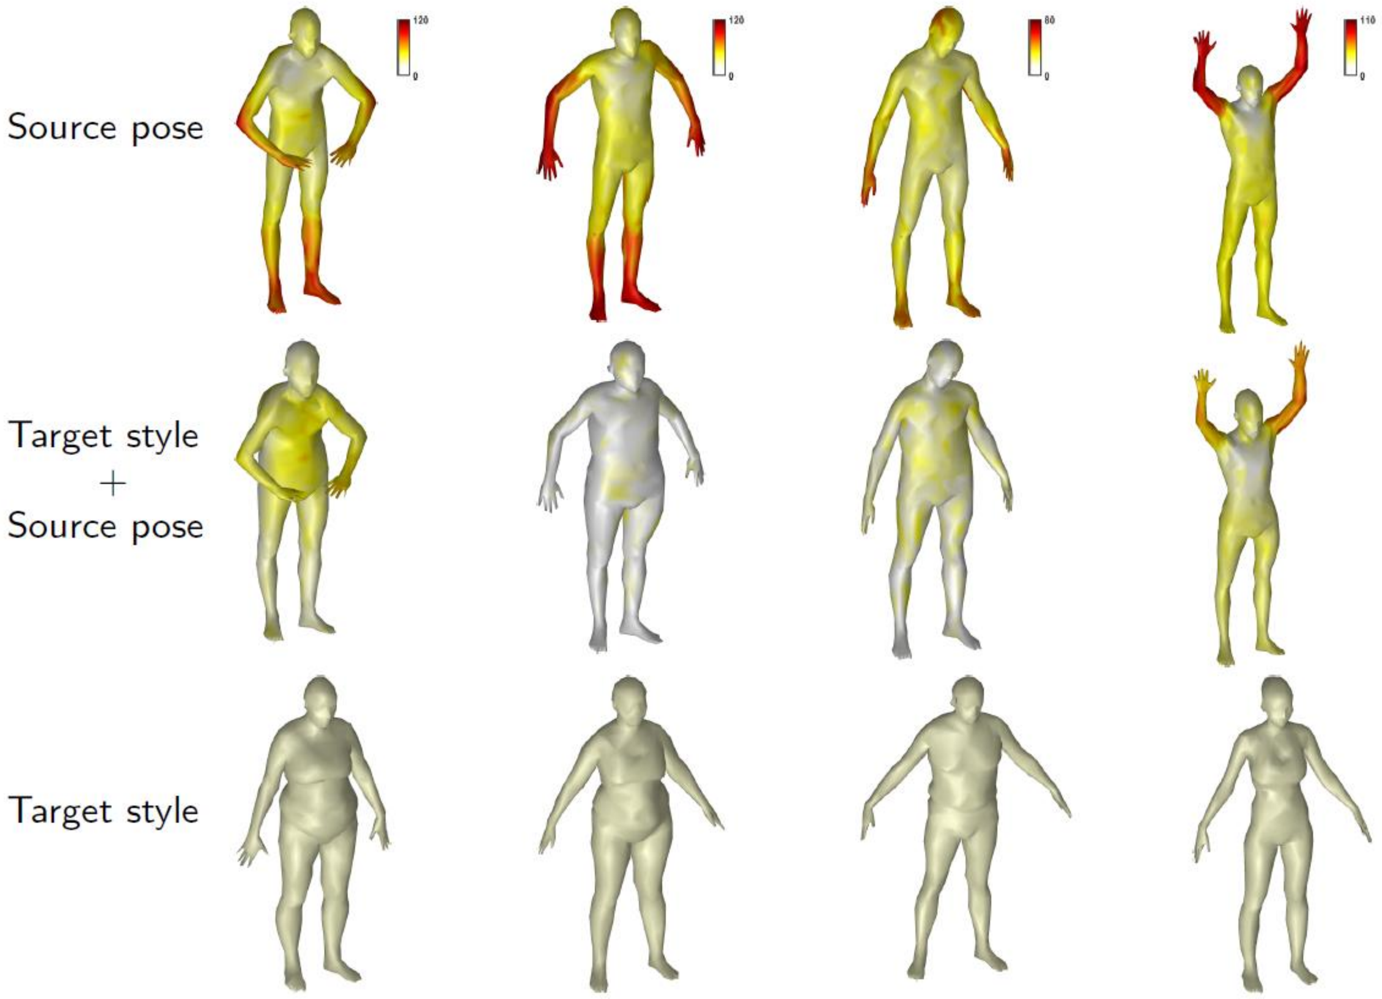

Isospectralization: Applications

Universal Spectral Adversarial Attacks for Deformable Shapes

Rampini et al. CVPR 2021

- Spectrum as a proxy for Universal Deformations

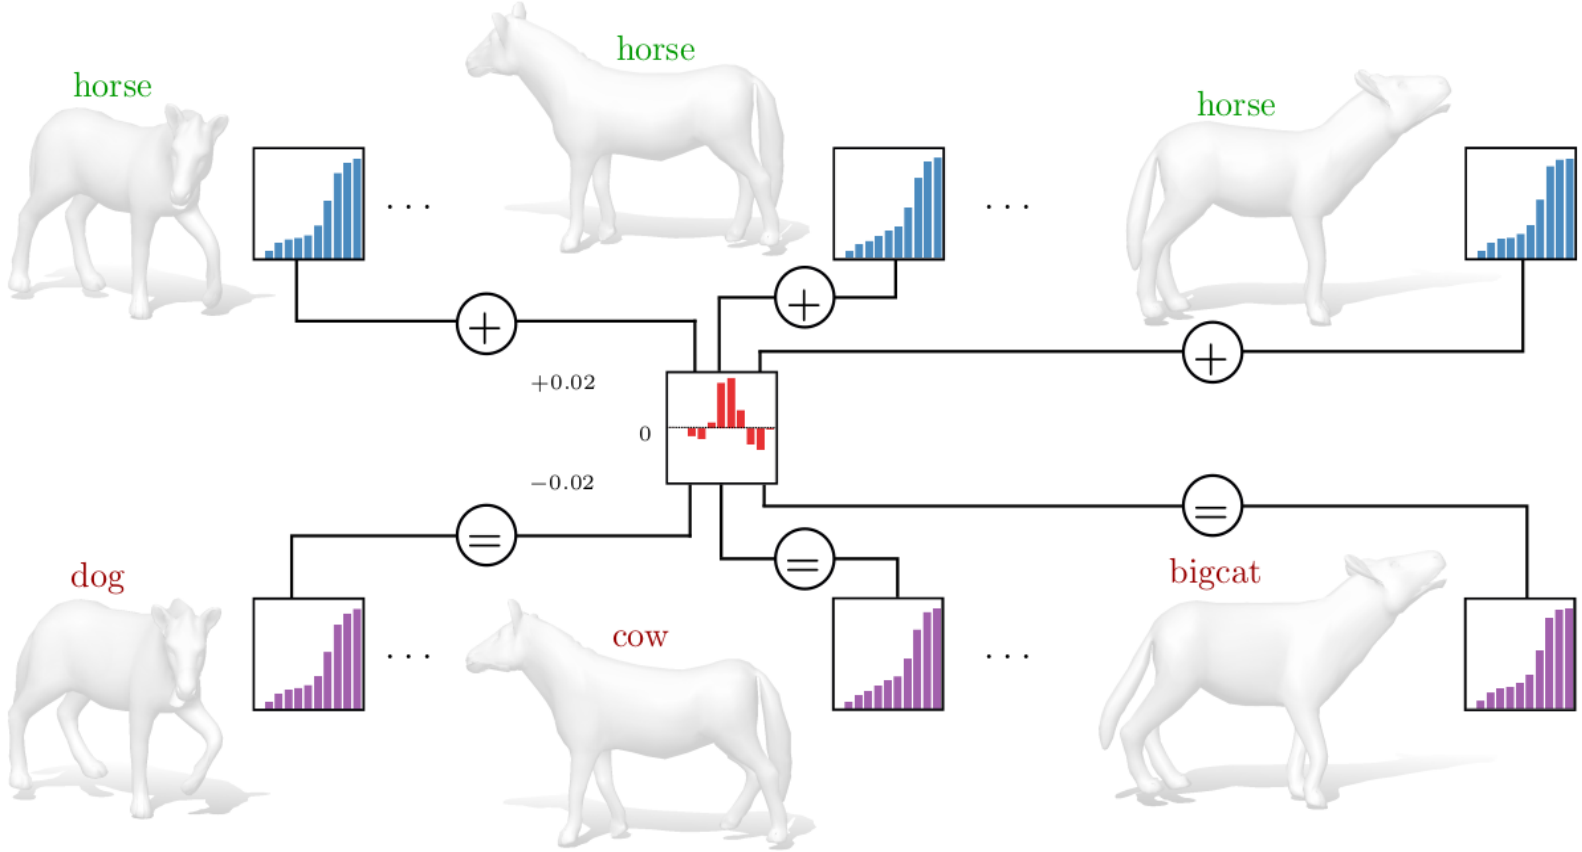

Isospectralization: Applications

Universal Spectral **Adversarial Attacks** for Deformable Shapes

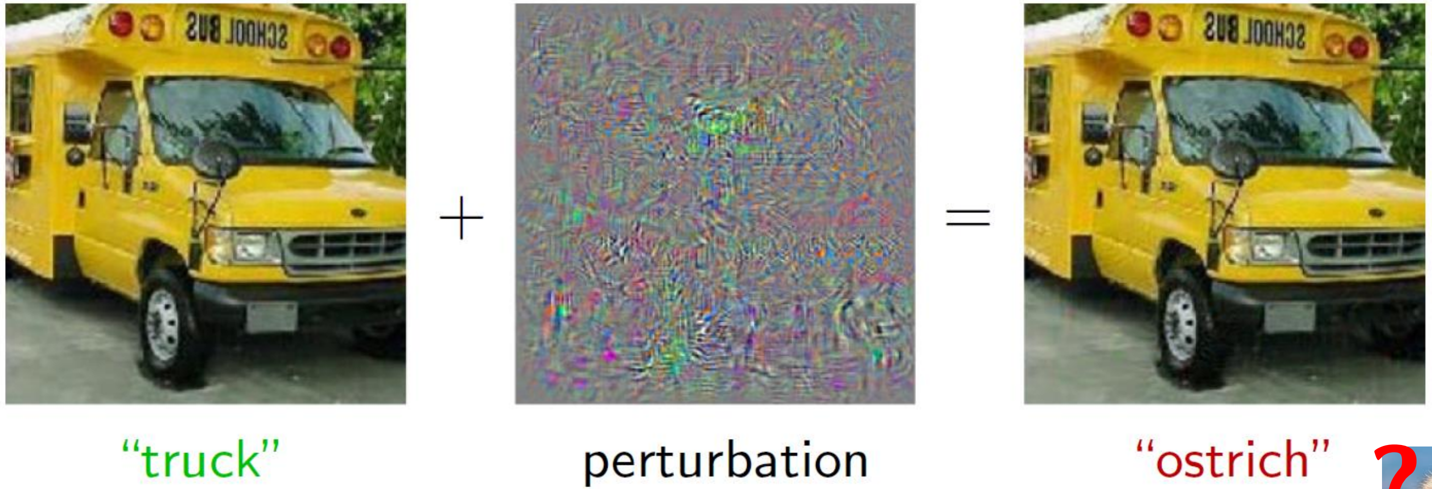

The perturbation should be undetectable and can be explicitly optimized for.

Szegedy et al. *Intriguing properties of neural networks.* ICLR 2014
Goodfellow et al. *Explaining and Harnessing Adversarial Examples.* ICLR 2015

Isospectralization: Applications

Universal Spectral **Adversarial Attacks** for Deformable Shapes

- Given an input shape x, a classifier C, and possibly a target class t, *consider:*

$$\min_{\mathbf{x}' \in [0,1]^n} \|\mathbf{x} - \mathbf{x}'\|_2^2$$

$$\text{s.t. } C(\mathbf{x}') = t \;\; \mathit{or} \;\; C(\mathbf{x}') \neq C(\mathbf{x})$$

We call x' an adversarial attack.

- Miss-classification constraint relaxed to a penalty term

$$\min_{\mathbf{x}' \in [0,1]^n} \|\mathbf{x} - \mathbf{x}'\|_2^2 + c\, L(\mathbf{x}', t)$$

Isospectralization: Applications

Universal Spectral **Adversarial Attacks** for Deformable Shapes

- A more general approach is given by:

$$\min_{\delta \in [0,1]^n} d(\mathbf{x}, \mathbf{x} + \boldsymbol{\delta}) + c f(\mathbf{x} + \boldsymbol{\delta})$$

where the perturbation δ appears explicitly, and d is some distane

f is such that $C(\mathrm{x}+\delta)=\mathrm{t}$ if and oly if $f(\mathrm{x}+\delta) \leq 0$.

$$f_4(x') = (0.5 - F(x')_t)^+$$
$$f_5(x') = -\log(2F(x')_t - 2)$$
$$f_6(x') = (\max_{i \neq t}(Z(x')_i) - Z(x')_t)^+$$
$$f_7(x') = \text{softplus}(\max_{i \neq t}(Z(x')_i) - Z(x')_t) - \log(2)$$

See:

Carlini and Wagner, 2016
"Towards evaluating the robustness of neural networks"

Isospectralization: Applications

Universal Spectral Adversarial Attacks for Deformable Shapes

Image-agnostic perturbations are known to exist *.

$$\min_{\delta\in[0,1]^n} \sum_i d(\mathbf{x}_i, \mathbf{x}_i + \boldsymbol{\delta}) + cf(\mathbf{x}_i + \boldsymbol{\delta})$$

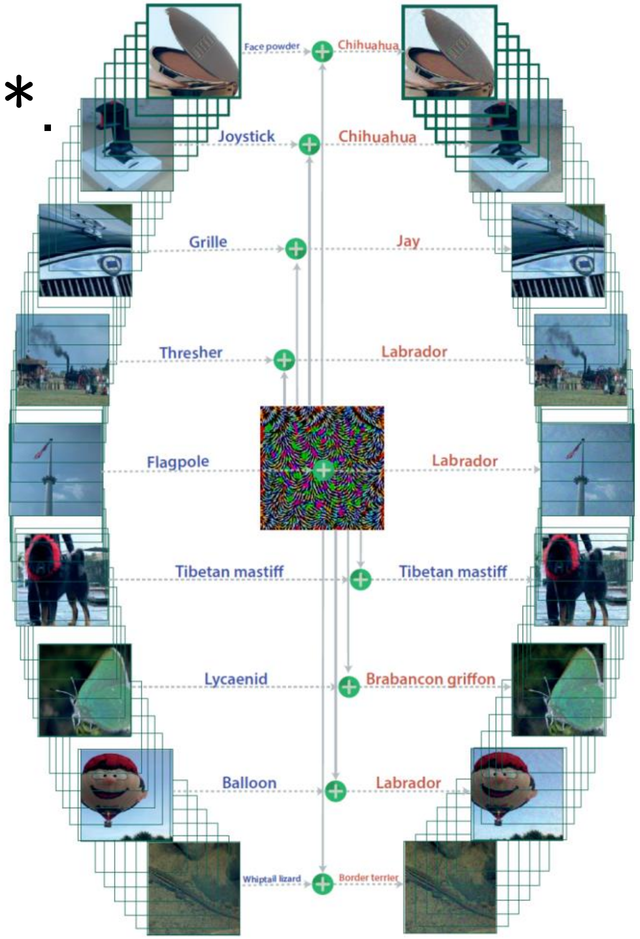

What about surfaces and point clouds?

Can we even define a single spatial perturbation for an entire collection of shapes?

* Moosavi-Dezfooli et al. *Universal adversarial perturbations*. CVPR 2017

Isospectralization: Applications

Universal Spectral Adversarial Attacks for **Deformable Shapes**

- We do not always have shapes in **correspondence**
- Spatial transformations are **not invariant** to isometries.

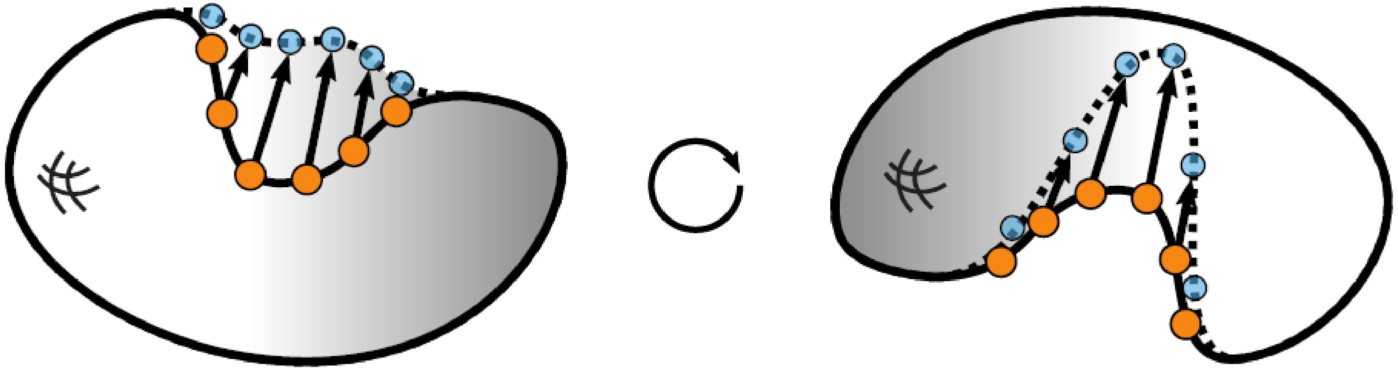

Isospectralization: Applications

Universal Spectral Adversarial Attacks for **Deformable Shapes**

- Let $\sigma(X) = (\lambda_1, \lambda_1, \ldots, \lambda_k)$ be the shape spectrum

$$\min_{\substack{\rho \in \mathbb{R}^k \\ \mathcal{P}_i}} \sum_{X_i \in \mathcal{S}} \|\sigma(X_i)(1+\rho) - \sigma(\mathcal{P}_i(X_i))\|_2^2$$

$$\text{s.t.} \quad \mathcal{C}(\mathcal{P}_i(X_i)) \neq \mathcal{C}(X_i) \quad \forall X_i \in \mathcal{S}$$

ρ shape-agnostic, universal perturbation.

$\mathcal{P}_i$ shape-specific, extrinsic (acting on $\mathbb{R}^3$) perturbations for each shape

$$\begin{array}{ccc} \boldsymbol{X}_i & \xrightarrow{\sigma} & (\boldsymbol{\lambda}^i) \\ \downarrow \mathcal{P}_i & & \downarrow \rho \\ \tilde{\boldsymbol{X}}_i & \xrightarrow{\sigma} & (\tilde{\boldsymbol{\lambda}}^i) \end{array}$$

Isospectralization: Applications

Universal Spectral Adversarial Attacks for **Deformable Shapes**

- Perturbation expressed as a linear combination of smooth vector fields (eigenvectors of LBO)*:

$$\mathcal{P}_i\left(X_i\right) = X_i + \Phi_i \alpha_i$$

- Resulting in the optimization problem:

$$\min_{\substack{\boldsymbol{\rho} \in \mathbb{R}^k \\ \{\alpha_i\}_i}} \sum_{X_i \in \mathcal{S}} \|\sigma(X_i)(1+\boldsymbol{\rho}) - \sigma(X_i + \Phi_i \alpha_i)\|_2^2$$

$$\text{s.t.} \quad \mathcal{C}(X_i + \Phi_i \alpha_i) \neq \mathcal{C}(X_i) \quad \forall X_i \in \mathcal{S}$$

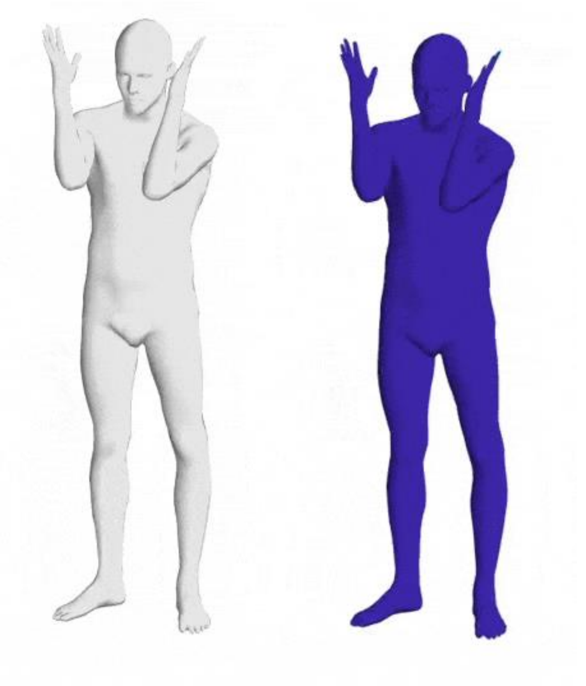

[*] Mariani et al. 2020. "Generating adversarial surfaces via band-limited perturbations". CGF 2020

Isospectralization: Applications

Universal Spectral Adversarial Attacks for Deformable Shapes

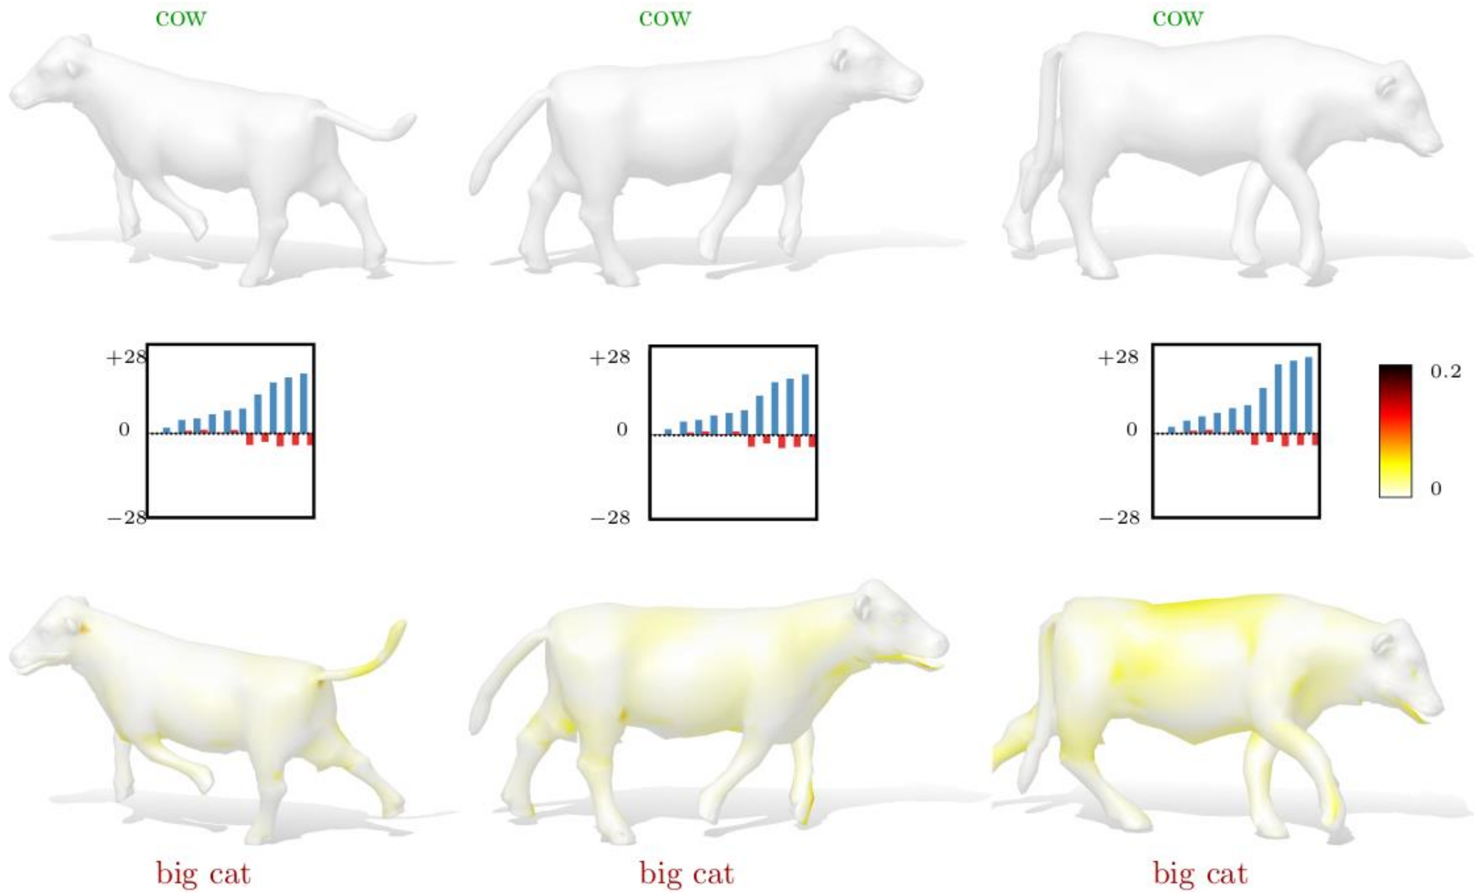

Isospectralization: Applications

Universal Spectral Adversarial Attacks for Deformable Shapes

Generalization: the deformation can be transferred to unseen shapes and cause misclassification.

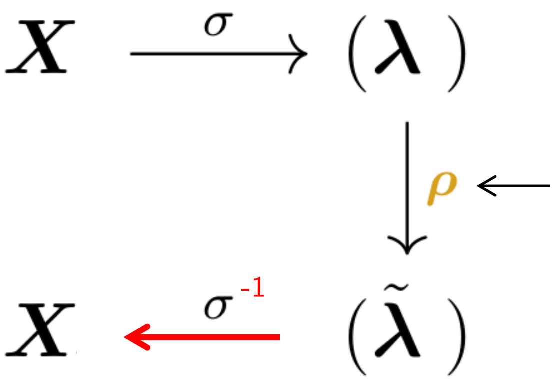

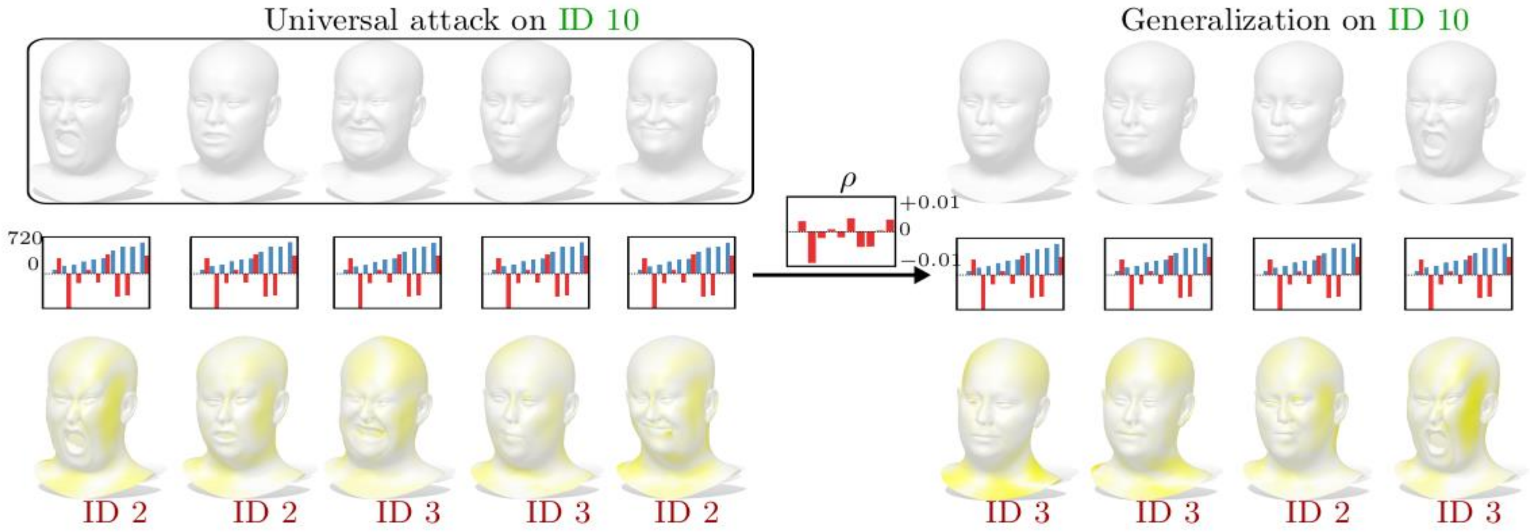

$$\min_{\substack{\rho \in \mathbb{R}^k \\ \alpha \in \mathbb{R}^k}} \sum_{X \in \mathcal{S}} \left\| \sigma(X)(1+\boldsymbol{\rho}) - \sigma(X + \Phi\alpha) \right\|_2^2$$

Isospectralization!

Isospectralization

Drawbacks of optimization strategy:

- Slow and tedious
  - At most 30 eigenvalues
  - Alternate optimization of boundary and interior points every 10 iterations
  - Re-sampling step is performed once every 200 iterations
  - Advanced optimization algorithms to escape local minima (Adam)
- Not straightforward to define priors/regularizers for specific domains

Data driven approach

AE-based learning model. (Marin et al. *Instant recovery of shape from spectrum via latent space connections.* 3DV 2020)

- Latent space connections

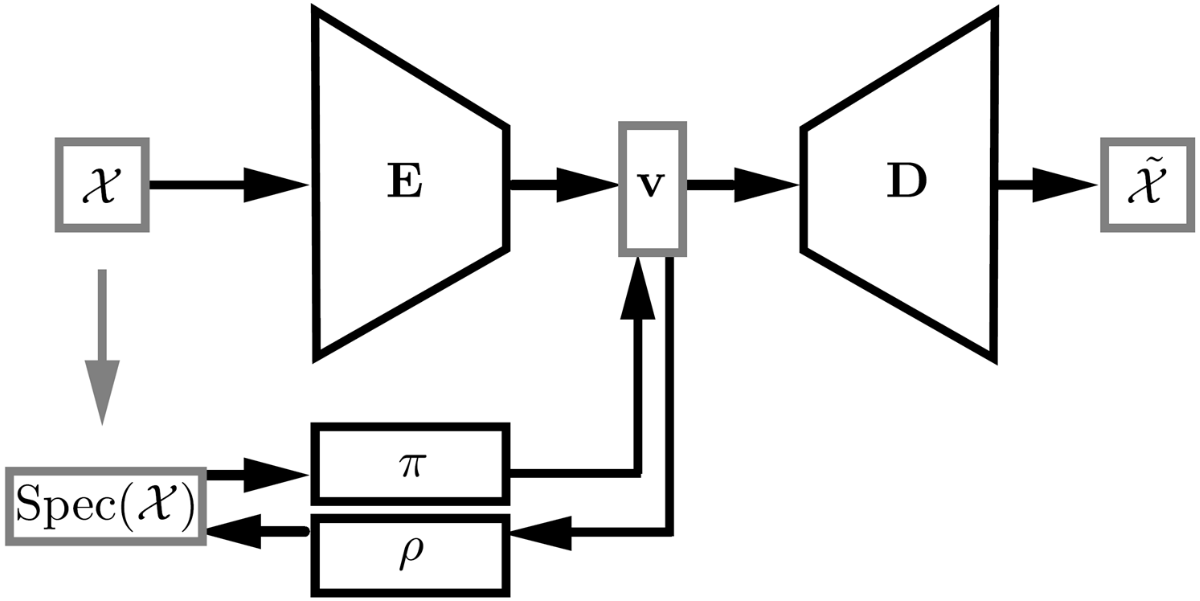

Data driven approach

AE-based learning model. (Marin et al. *Instant recovery of shape from spectrum via latent space connections.* 3DV 2020)

- Latent space connections

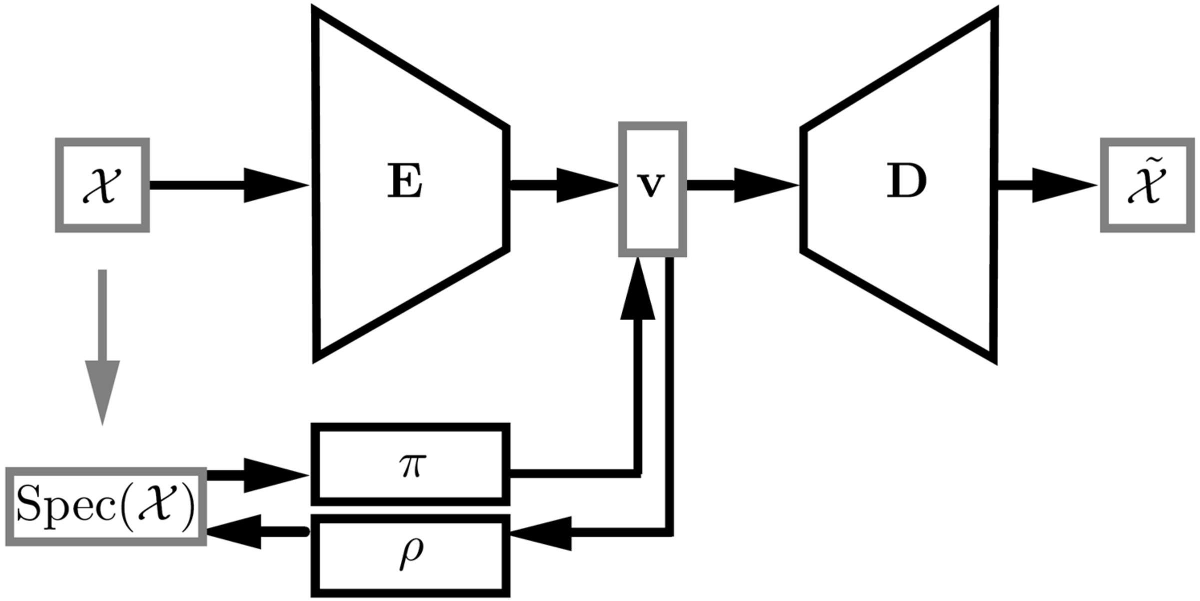

$$\ell = \ell_{\mathcal{X}} + \alpha \ell_{\lambda}\,, \quad \text{with}$$

$$\ell_{\mathcal{X}} = \frac{1}{n}\|D(E(\mathbf{X})) - \mathbf{X}\|_F^2$$

$$\ell_{\lambda} = \frac{1}{k}(\|\pi(\boldsymbol{\lambda}) - E(\mathbf{X})\|_2^2 + \|\rho(E(\mathbf{X})) - \boldsymbol{\lambda}\|_2^2)$$

Data driven approach

AE-based learning model. (Marin et al. *Instant recovery of shape from spectrum via latent space connections.* 3DV 2020)

- Latent space connections

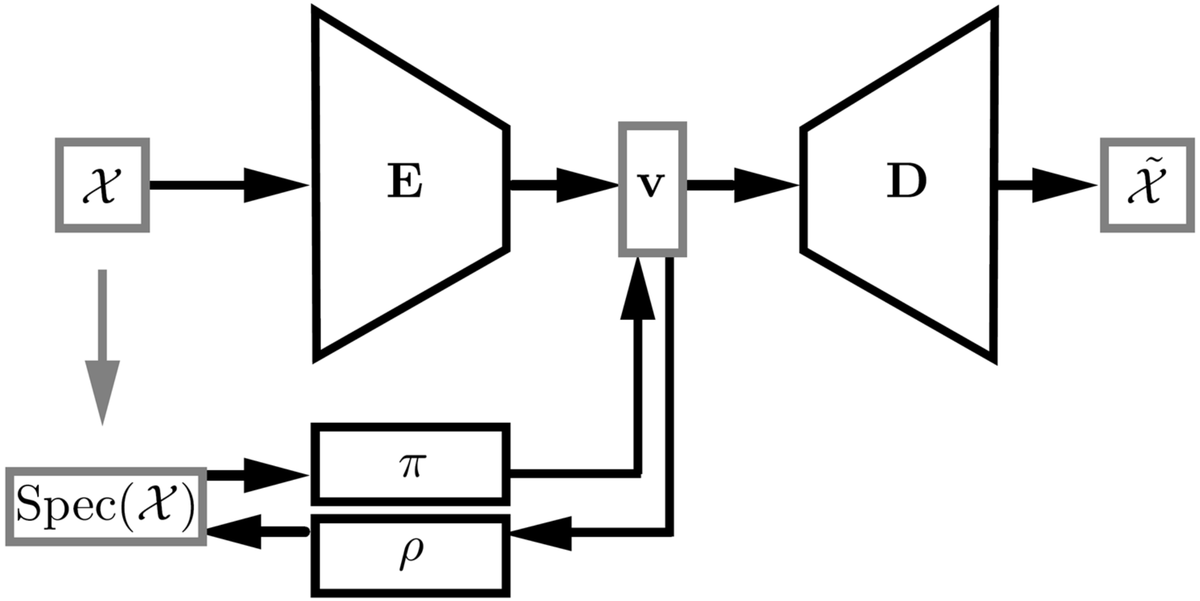

$$\ell = \ell_{\mathcal{X}} + \alpha \ell_{\lambda}\,, \quad \text{with}$$

$$\ell_{\mathcal{X}} = \frac{1}{n}\|D(E(\mathbf{X})) - \mathbf{X}\|_F^2$$

$$\ell_{\lambda} = \frac{1}{k}(\|\pi(\boldsymbol{\lambda}) - E(\mathbf{X})\|_2^2 + \|\rho(E(\mathbf{X})) - \boldsymbol{\lambda}\|_2^2)$$

The spectral loss enforces:

$$\rho \approx \pi^{-1}$$

Data driven approach

AE-based learning model. (Marin et al. *Instant recovery of shape from spectrum via latent space connections.* 3DV 2020)

- Latent space connections

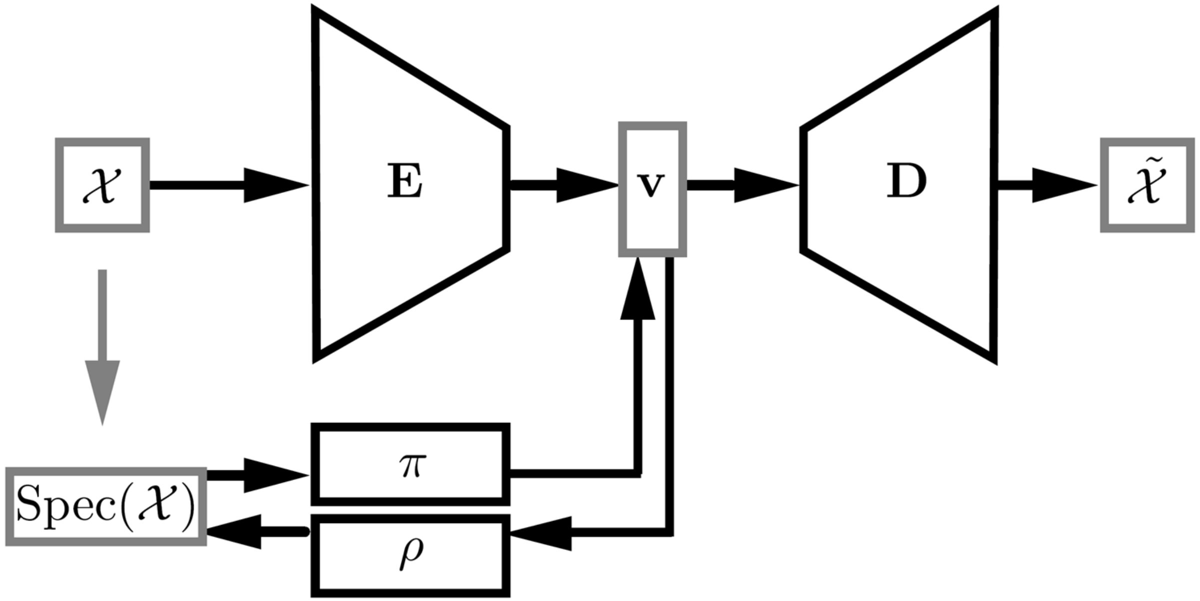

Remarks:

- **No back-propagation** through the eigen-decomposition

Data driven approach

AE-based learning model. (Marin et al. *Instant recovery of shape from spectrum via latent space connections.* 3DV 2020)

- Latent space connections

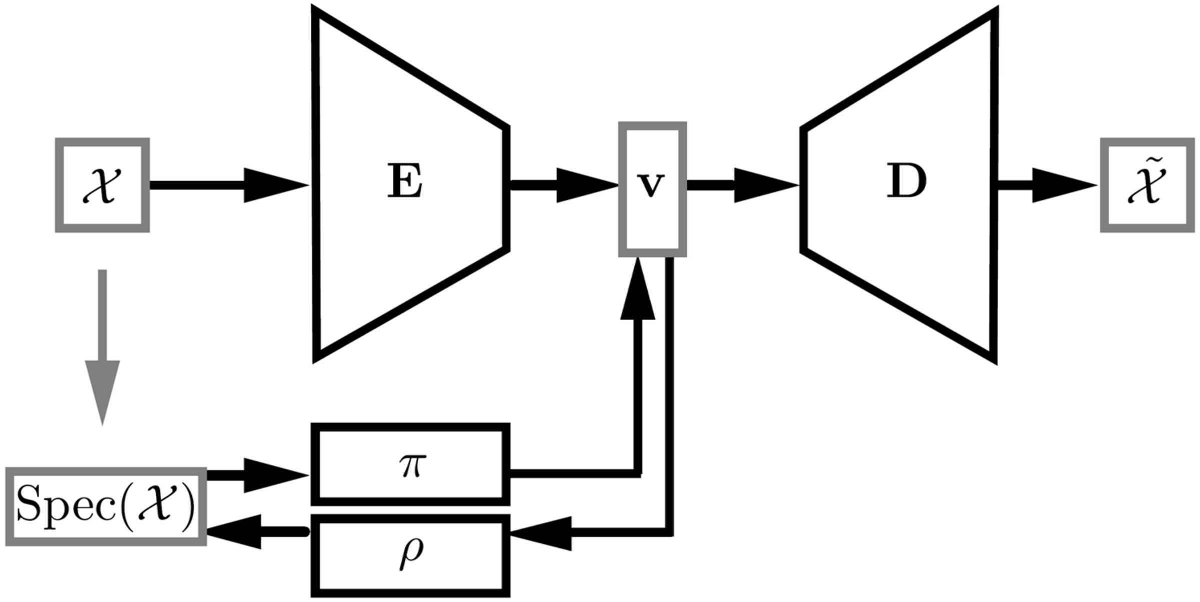

Remarks:

- **No back-propagation** through the eigen-decomposition
- The input spectrum can be **arbitrarily accurate**

Data driven approach

AE-based learning model. (Marin et al. *Instant recovery of shape from spectrum via latent space connections.* 3DV 2020)

- Latent space connections

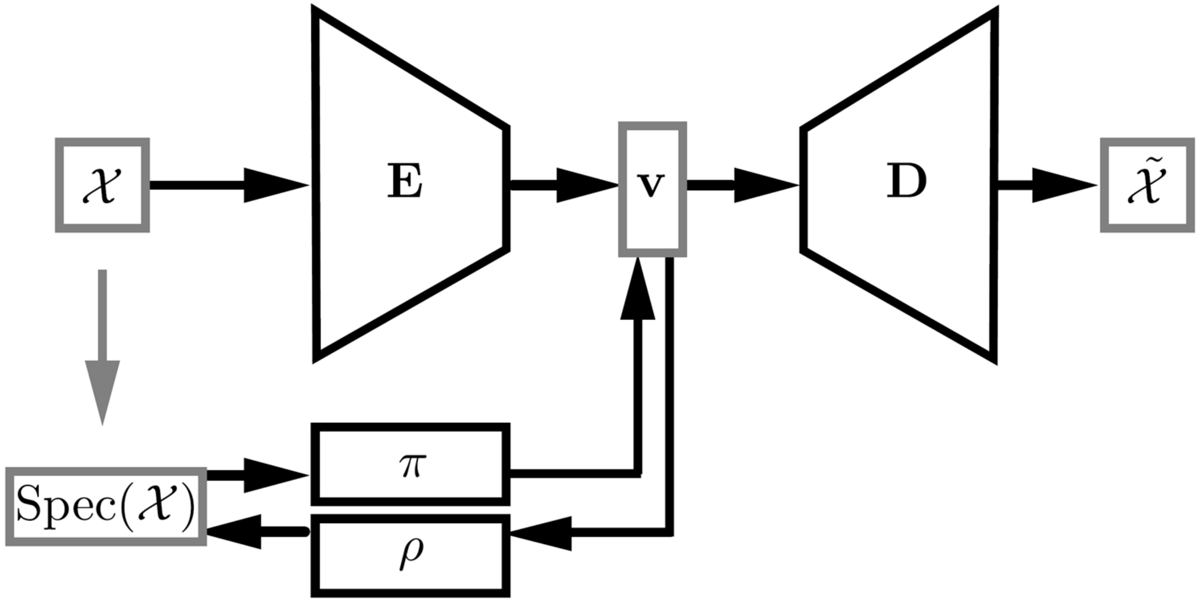

Remarks:

- **No back-propagation** through the eigen-decomposition
- The input spectrum can be **arbitrarily accurate**
- Admits **any AE** model (e.g. for point clouds, meshes, etc.)

Data driven approach

- Shape-from-spectrum reconstruction

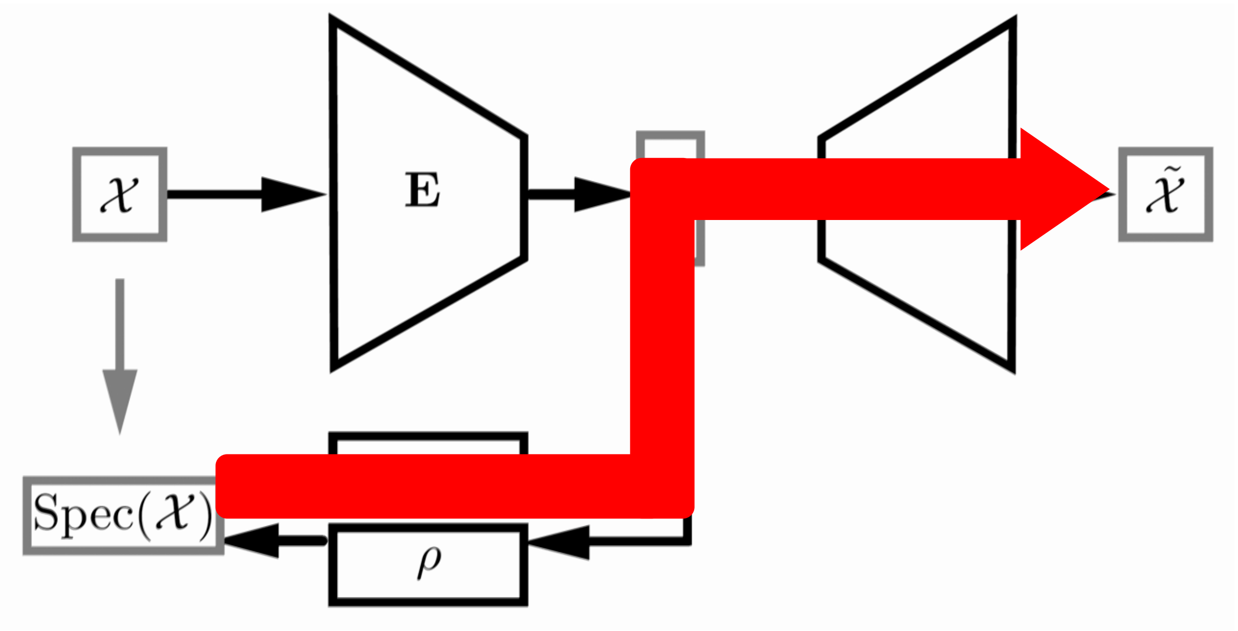

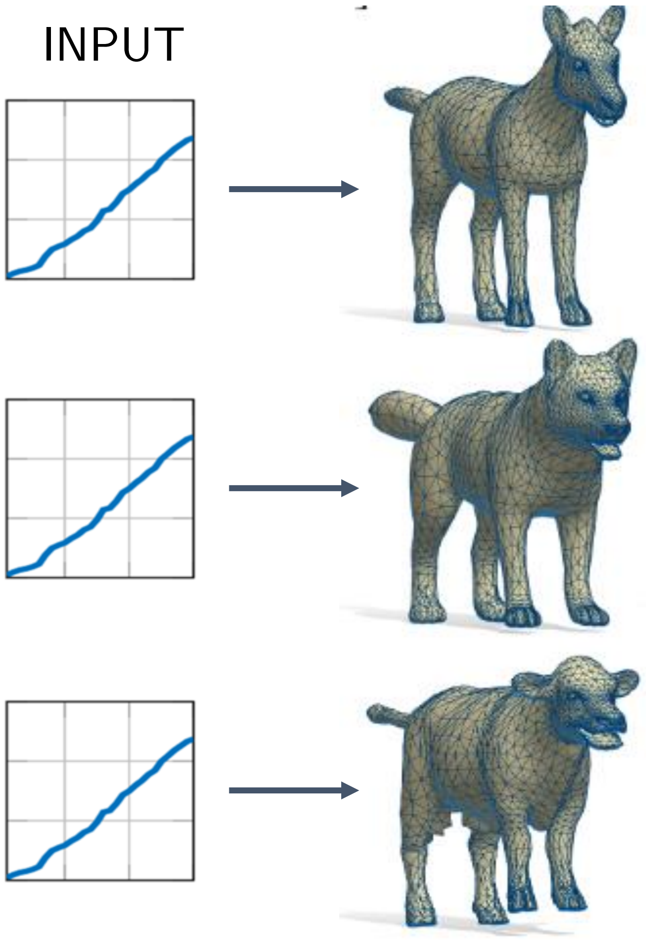

Data driven approach

- Style transfer

$$\min_{\mathbf{v}} \|\mathrm{Spec}(\mathcal{X}_{\mathrm{style}}) - \rho(\mathbf{v})\|_2^2 + w\|\mathbf{v} - E(\mathcal{X}_{\mathrm{pose}})\|_2^2$$

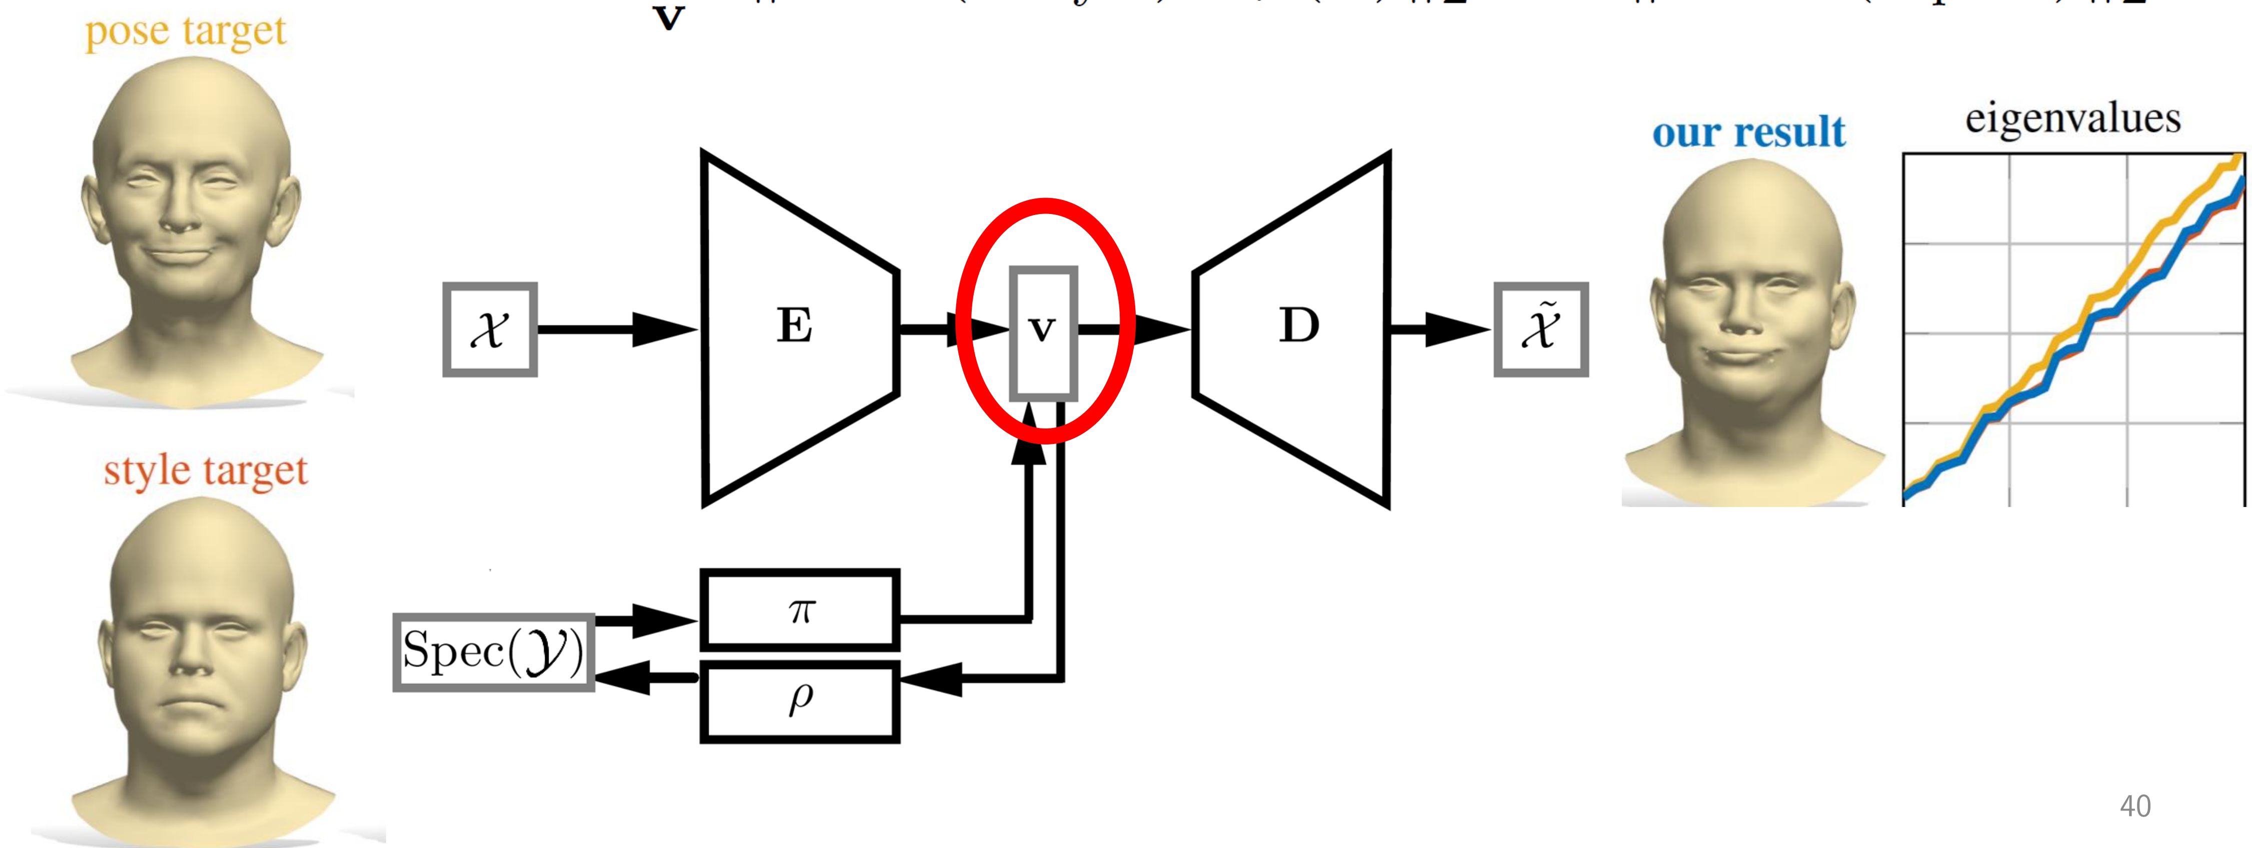

Data driven approach

- Shape exploration

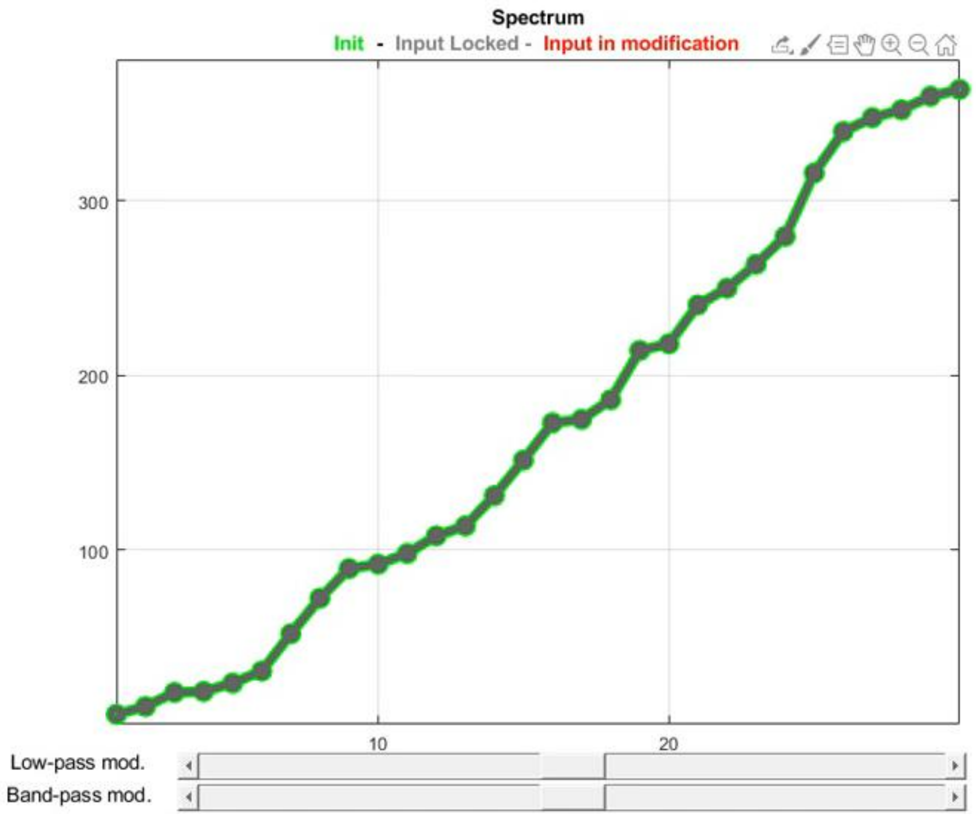

Data driven approach

- Spectra estimation

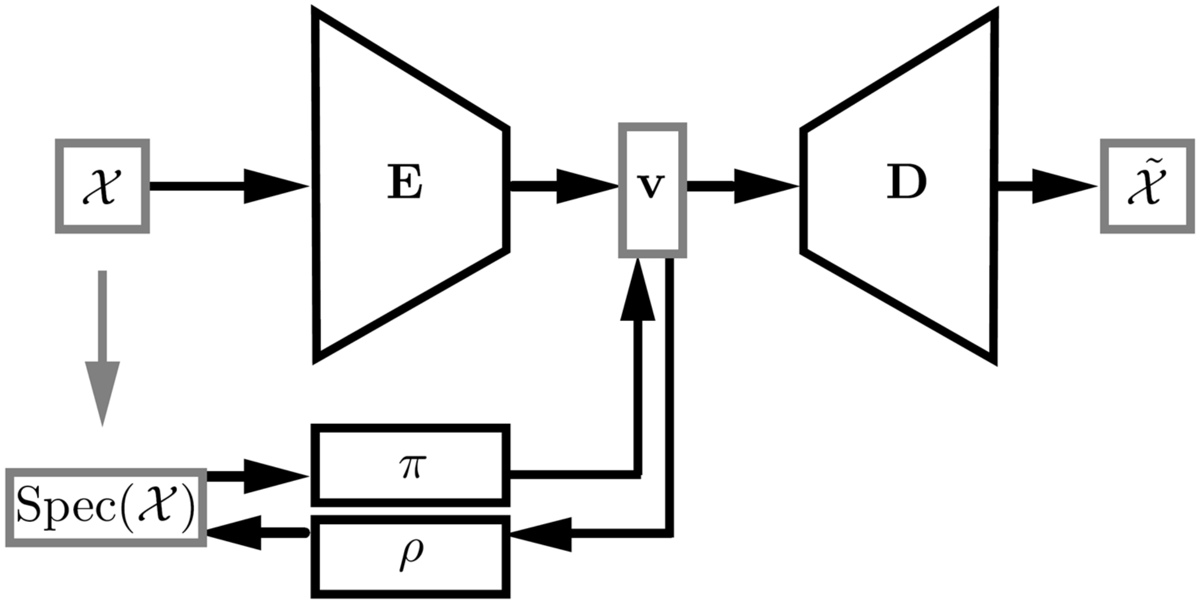

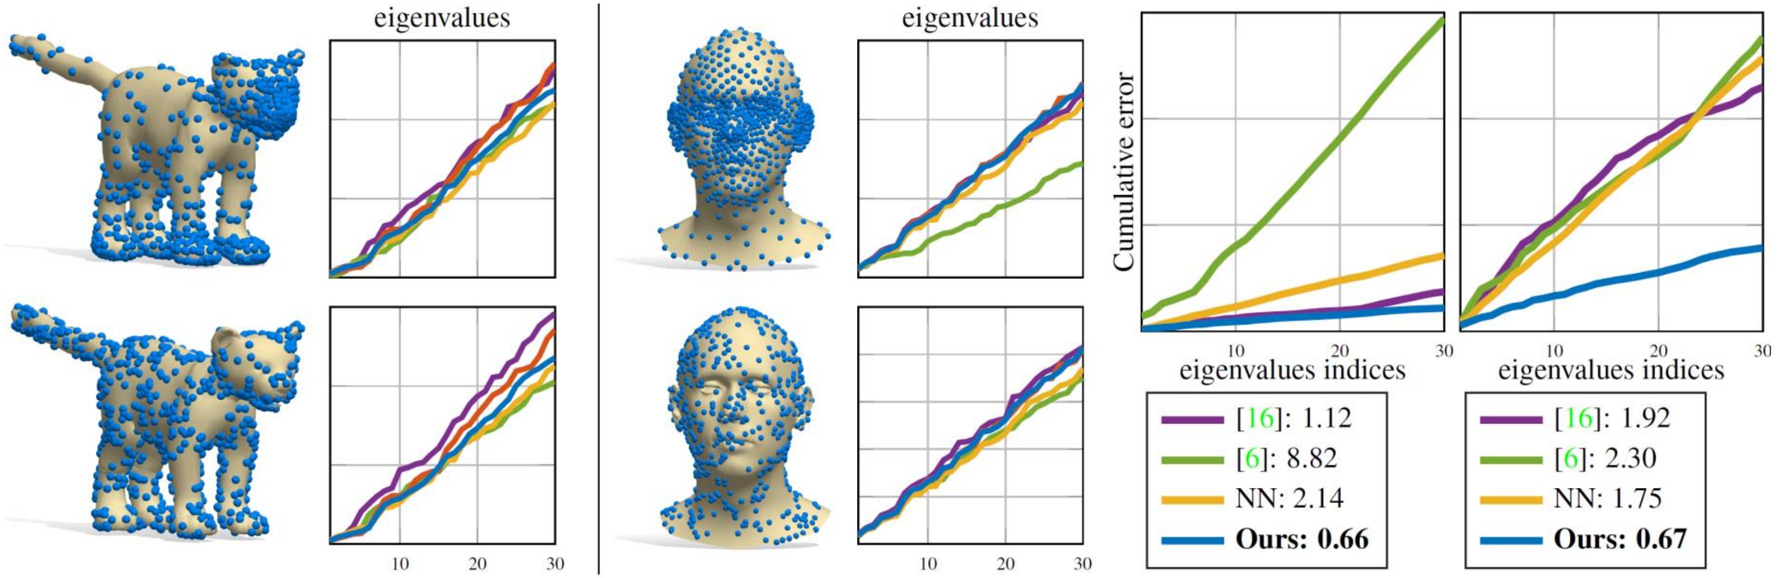

Summary

- Eigenvalues are
  - Isometry invariant
  - Discretization invariant
  - Correspondence free
- Enables a lot of applications in the shape analysis field:
  - Shape compression and reconstruction
  - Style transfer
  - Shape correspondence
- Physically meaningful (latent) space for shape exploration