

# Learning from and with persistent homology 

Roland Kwitt

## Talk outline

$\triangleright$ Quick recap of the learning framework (supervised learning)
$\triangleright$ Neural networks
$\triangleright$ Learning from persistent homology
$\triangleright$ Learning with persistent homology

## Problem setting (of supervised learning)

> Domain set
> Label set
> Hypothesis class

Distribution over domain \& labels
$\mathcal{X}\left(\right.$ e.g., $\left.\mathbb{R}^{\text {d }}\right)$
$y(e . g .,\{0,1\})$
$\mathcal{H}$

Training data $S=\left(\left(x_{1}, y_{1}\right), \ldots,\left(x_{\mathfrak{m}}, y_{\mathfrak{m}}\right)\right) \sim \mathcal{P}^{m}$

## Problem setting (of supervised learning)

$$
\begin{array}{rc}
\text { Domain set } & \mathcal{X}\left(e . g ., \mathbb{R}^{\mathrm{d}}\right) \\
\text { Label set } & y(\mathrm{e} . \mathrm{g},\{0,1\}) \\
\text { Hypothesis class } & \mathcal{H} \\
\text { domain \& labels } & \left(x_{i}, y_{\mathfrak{i}}\right) \sim \mathcal{P} \\
\text { Training data } & S=\left(\left(x_{1}, y_{1}\right), \ldots,\left(x_{\mathfrak{m}}, y_{\mathfrak{m}}\right)\right) \sim \mathcal{P}^{\mathfrak{m}} \\
\hline
\end{array}
$$

Distribution over domain \& labels

A learner (upon receiving training data) needs to output a hypothesis

$$
\mathcal{H} \ni h: X \rightarrow y
$$

## Problem setting (of supervised learning)

$$
\begin{array}{rc}
\text { Domain set } & x\left(e . g ., \mathbb{R}^{\mathrm{d}}\right) \\
\text { Label set } & y(\mathrm{e} . \mathrm{g},\{0,1\}) \\
\text { Hypothesis class } & \mathcal{H} \\
\text { r domain \& labels } & \left(x_{i}, y_{i}\right) \sim \mathcal{P}
\end{array}
$$

Distribution over domain \& labels

A learner (upon receiving training data) needs to output a hypothesis

$$
\mathcal{H} \ni h: X \rightarrow y
$$

Such a hypothesis should have small risk, defined as

$$
\mathrm{L}_{\mathcal{P}}(\mathrm{h})=\operatorname{Pr}_{(x, y) \sim \mathcal{P}}[\mathrm{h}(\mathrm{x}) \neq \mathrm{y}]
$$

## Problem setting (of supervised learning)

However, we can only measure the empirical risk

$$
L_{S}(h)=\frac{\left|i \in\{1, \ldots, m\}: h\left(x_{i}\right) \neq y_{i}\right|}{m}
$$

## Problem setting (of supervised learning)

However, we can only measure the empirical risk

$$
L_{S}(h)=\frac{\left|i \in\{1, \ldots, m\}: h\left(x_{i}\right) \neq y_{i}\right|}{m}
$$

Classic learning paradigm: minimize empirical risk

$$
h \in \arg \min _{h \in \mathscr{H}} L_{S}(h)
$$

## Problem setting (of supervised learning)

However, we can only measure the empirical risk

$$
L_{S}(h)=\frac{\left|i \in\{1, \ldots, m\}: h\left(x_{i}\right) \neq y_{i}\right|}{m}
$$

Classic learning paradigm: minimize empirical risk

$$
h \in \arg \min _{h \in \mathcal{H}} L_{S}(h)
$$

Example:

$$
\begin{aligned}
& \mathcal{X}=\mathbb{R}^{\mathrm{d}}, \boldsymbol{y}=\{+1,-1\} \\
& \mathcal{H}=\left\{\boldsymbol{x} \mapsto \operatorname{sgn}\langle\boldsymbol{x}, \boldsymbol{w}\rangle: \boldsymbol{w} \in \mathbb{R}^{\mathrm{d}}\right\}
\end{aligned}
$$

(aka halfspace classifiers)


## Problem setting (of supervised learning)

An important aspect is that, typically, inputs are of fixed size!

## Problem setting (of supervised learning)

An important aspect is that, typically, inputs are of fixed size!
Other types of data, such as
$\triangleright$ sets,
$\triangleright$ multi-sets,
$\triangleright$ graphs, or
$\triangleright$ point clouds
are (or were) - lets put it this way - more challenging to handle!

## Problem setting (of supervised learning)

An important aspect is that, typically, inputs are of fixed size!
Other types of data, such as
$\triangleright$ sets,
$\triangleright$ multi-sets,
$\triangleright$ graphs, or
$\triangleright$ point clouds
are (or were) - lets put it this way - more challenging to handle!
General recipe: Find a reasonable way to vectorize!

## Neural networks

Typical (feed-forward) neural networks compose maps of the form

$$
\begin{aligned}
f: & \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R}^{e} \\
& \boldsymbol{x} \mapsto \sigma(\boldsymbol{A x})
\end{aligned}
$$

i.e., a linear map $\boldsymbol{A}$, followed by a (component-wise) activation, e.g.,


## Neural networks

Typical (feed-forward) neural networks compose maps of the form

$$
\begin{aligned}
\mathrm{f}: & \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R}^{e} \\
& \boldsymbol{x} \mapsto \sigma(\boldsymbol{A} \boldsymbol{x})
\end{aligned}
$$

i.e., a linear map $\boldsymbol{A}$, followed by a (component-wise) activation, e.g.,


Composition of such "building blocks" gives

$$
\begin{aligned}
\mathrm{F}: & \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R} \\
& \boldsymbol{x} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \boldsymbol{x}\right) \cdots\right)\right)
\end{aligned}
$$

i.e., the hypothesis class is parametrized by $\left(\boldsymbol{A}_{1}, \ldots, \boldsymbol{A}_{\mathrm{L}}, \boldsymbol{w}\right)$.

## Barcodes as input?

So, what if the input, $\boldsymbol{x}$, to

$$
F: \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R}, \quad \boldsymbol{x} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \boldsymbol{x}\right) \cdots\right)\right)
$$

## Barcodes as input?

So, what if the input, $x$, to

$$
F: \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R}, \quad \boldsymbol{x} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \boldsymbol{x}\right) \cdots\right)\right)
$$

is now a persistence barcode, G, i.e., a multi-set of (birth, death) tuple?

## Barcodes as input?

So, what if the input, $x$, to

$$
F: \mathbb{R}^{\mathscr{\alpha}} \rightarrow \mathbb{R}, \quad \chi \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma(\boldsymbol{A} / 1 \boldsymbol{x}) \cdots\right)\right)
$$

is now a persistence barcode, G, i.e., a multi-set of (birth, death) tuple?

## Barcodes as input?

So, what if the input, $x$, to

$$
F: \mathbb{R}^{\mathscr{A}} \rightarrow \mathbb{R}, \quad \chi \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma(\boldsymbol{A} / 1 \mathrm{x}) \cdots\right)\right)
$$

is now a persistence barcode, G, i.e., a multi-set of (birth, death) tuple?
Question: How can we deal with this?

## Barcodes as input?

So, what if the input, $x$, to

$$
F: \mathbb{P}^{\mathscr{A}} \rightarrow \mathbb{R}, \quad \chi \not \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma(\boldsymbol{A} / 1 \mathrm{x}) \cdots\right)\right)
$$

is now a persistence barcode, G, i.e., a multi-set of (birth, death) tuple?
Question: How can we deal with this?
A pragmatic approach [Bendich et al., 2014]図:

take the lengths of the N -longest bars $\rightarrow$ gives a N -dim. vectorization

## Barcodes as input?

So, what if the input, $x$, to

$$
F: \mathbb{R}^{\boldsymbol{\alpha}} \rightarrow \mathbb{R}, \quad \boldsymbol{\chi} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma(\boldsymbol{A} / 1 \boldsymbol{x}) \cdots\right)\right)
$$

is now a persistence barcode, G, i.e., a multi-set of (birth, death) tuple?
Question: How can we deal with this?
A pragmatic approach [Bendich et al., 2014]

take the lengths of the N -longest bars $\rightarrow$ gives a N -dim. vectorization
Question: Why should we care about "how" we vectorize?
Well, it would be desirable to preserve stability wrt. $d_{B}, d_{W_{p, q}}$.

## Prior art

## Vectorization techniques

Persistence landscapes
Persistence silhouettes
Persistence images
Template functions
ATOL ${ }^{\dagger}$
［Bubenik，2015］囚
［Chazal et al．，2014］図 ［Adams et al．，2017］図
［Perea et al．，2019］囚
［Royer et al．，2019］区

## Prior art

## Vectorization techniques

Persistence landscapes
Persistence silhouettes
Persistence images
Template functions
ATOL ${ }^{\dagger}$
Kernel－based techniques
Persistence scale－space kernel
Sliced Wasserstein kernel
Persistence－weighted Gaussian kernel
Kernel for multi－parameter persistent homology
［Reininghaus et al．，2014］国 ［Carrière et al．，2017］区
［Kusano et al．，2016］因
［Corbet et al．，2019］
Theoretical results related to metric distortion［Carrière \＆Bauer，2019］조

## Prior art

## Vectorization techniques

Persistence landscapes
Persistence silhouettes
Persistence images
Template functions
ATOL ${ }^{\dagger}$

## Kernel－based techniques

Persistence scale－space kernel
Sliced Wasserstein kernel
Persistence－weighted Gaussian kernel
Kernel for multi－parameter persistent homology
［Reininghaus et al．，2014］且 ［Carrière et al．，2017］区 ［Kusano et al．，2016］因 ［Corbet et al．，2019］션 Theoretical results related to metric distortion［Carrière \＆Bauer，2019］［

This is，by far，not an exhaustive listing！

## Can we obtain task-optimal vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

## Can we obtain task-optimall vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

This motivates learnable vectorization schemes:
[Hofer et al., 2017,2019] 図, [Carrière et al., 2019] 조

## Can we obtain task-optimal vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

This motivates learnable vectorization schemes:
[Hofer et al., 2017,2019] 図, [Carrière et al., 2019] 조
Example (for a vectorization into $\mathbb{R}^{k}, k=2$ ):


## Can we obtain task-optimal vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

This motivates learnable vectorization schemes:
[Hofer et al., 2017,2019] 図, [Carrière et al., 2019] 조
Example (for a vectorization into $\mathbb{R}^{k}, k=2$ ):


## Can we obtain task-optimal vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

This motivates learnable vectorization schemes:

## [Hofer et al., 2017,2019]因, [Carrière et al., 2019]図

Example (for a vectorization into $\mathbb{R}^{k}, k=2$ ):

$$
\xrightarrow[\text { Birth }]{\text { cole }}
$$

## Can we obtain task-optimal vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

This motivates learnable vectorization schemes:

## [Hofer et al., 2017,2019]図, [Carrière et al., 2019] 因

Example (for a vectorization into $\mathbb{R}^{k}, k=2$ ):


$$
\begin{aligned}
& c_{1}=\sum p \in G S_{\theta_{1}}(p) \\
& c_{2}=\sum p \in G S_{\theta_{2}}(p)
\end{aligned}
$$

Birth

## Can we obtain task-optimal vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

This motivates learnable vectorization schemes:

## [Hofer et al., 2017,2019]因, [Carrière et al., 2019]図

Example (for a vectorization into $\mathbb{R}^{k}, k=2$ ):


$$
\left.\begin{array}{l}
c_{1}=\sum_{p \in G} s_{\theta_{1}}(p) \\
c_{2}=\sum_{p \in G} s_{\theta_{2}}(p)
\end{array}\right\}\binom{c_{1}}{c_{2}}
$$

In general ${ }^{\dagger}: \mathrm{G} \mapsto \mathcal{V}_{\Theta}(\mathrm{G})$
$\Theta=\left(\theta_{1}, \theta_{2}\right)$

Birth

## Can we obtain task-optimal vectorizations?

In fact, most vectorization strategies are task-agnostic!
Question: Shouldn't the vectorization be informed by the learning task?

This motivates learnable vectorization schemes:

## [Hofer et al., 2017,2019]因, [Carrière et al., 2019] 区

Example (for a vectorization into $\mathbb{R}^{k}, k=2$ ):


$$
\left.\begin{array}{l}
c_{1}=\sum_{p \in G} s_{\theta_{1}}(p) \\
c_{2}=\sum_{p \in G} s_{\theta_{2}}(p)
\end{array}\right\}\binom{c_{1}}{c_{2}}
$$

In general ${ }^{\dagger}: \mathrm{G} \mapsto \mathcal{V}_{\Theta}(\mathrm{G})$
$\Theta=\left(\theta_{1}, \theta_{2}\right)$
$\dagger$ plus some technicalities to ensure stability
Learnable means that we can optimize the $\boldsymbol{\theta}_{i}$ 's for a given task/criterion!

## Can we obtain task-optimal vectorizations?

Overall, this changes

$$
F: \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R}, \quad \boldsymbol{x} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \boldsymbol{x}\right) \cdots\right)\right)
$$

to

$$
F: \mathcal{B} \rightarrow \mathbb{R}, \quad \mathrm{G} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \mathcal{V}_{\Theta}(\mathrm{G})\right) \cdots\right)\right)
$$

## Can we obtain task-optimal vectorizations?

Overall, this changes

$$
F: \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R}, \quad \boldsymbol{x} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \boldsymbol{x}\right) \cdots\right)\right)
$$

to

$$
F: \mathcal{B} \rightarrow \mathbb{R}, \quad \mathrm{G} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \mathcal{V}_{\Theta}(\mathrm{G})\right) \cdots\right)\right)
$$

Upon the definition of a suitable loss function

$$
\ell: \mathcal{H} \times \mathcal{X} \times y \rightarrow \mathbb{R}
$$

we can compute, for a training sample, $\left(G_{i}, y_{i}\right)$, the parameter update ${ }^{\dagger}$

$$
\Theta^{t+1}=\Theta^{t}-\eta \frac{\partial l\left(F,\left(G_{i}, y_{i}\right)\right)}{\partial \Theta}
$$

## Can we obtain task-optimal vectorizations?

Overall, this changes

$$
F: \mathbb{R}^{\mathrm{d}} \rightarrow \mathbb{R}, \quad \boldsymbol{x} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \boldsymbol{x}\right) \cdots\right)\right)
$$

to

$$
F: \mathcal{B} \rightarrow \mathbb{R}, \quad \mathrm{G} \mapsto \boldsymbol{w}^{\top} \sigma\left(\boldsymbol{A}_{\mathrm{L}} \sigma\left(\boldsymbol{A}_{\mathrm{L}-1} \cdots \sigma\left(\boldsymbol{A}_{1} \mathcal{V}_{\Theta}(\mathrm{G})\right) \cdots\right)\right)
$$

Upon the definition of a suitable loss function

$$
\ell: \mathcal{H} \times \mathcal{X} \times \mathcal{y} \rightarrow \mathbb{R}
$$

we can compute, for a training sample, $\left(G_{i}, y_{i}\right)$, the parameter update ${ }^{\dagger}$

$$
\Theta^{t+1}=\Theta^{t}-\eta \frac{\partial \ell\left(F,\left(G_{i}, y_{i}\right)\right)}{\partial \Theta}
$$

"Easy" because of automatic differentiaton (e.g., using PyTorch).


## Transitioning to learning with PH

In ML, we have, for long, degraded PH to a "fancy" feature extractor.

## Transitioning to learning with PH

In ML, we have, for long, degraded PH to a "fancy" feature extractor.

Question: What if we want to control topological properties?

## Transitioning to learning with PH

In ML, we have, for long, degraded PH to a "fancy" feature extractor.

Question: What if we want to control topological properties?

Example:


## Transitioning to learning with PH

In ML, we have, for long, degraded PH to a "fancy" feature extractor.

Question: What if we want to control topological properties?

Example:

e.g., control the lifetime of 0-dim. features (from Vietoris-Rips PH)

## Transitioning to learning with PH

Example (contd.):


## Transitioning to learning with PH

## Example (contd.):


persistence barcode of 0-dim. features

## Transitioning to learning with PH

## Example (contd.):

Connectivity loss (ConnLoss)

persistence barcode of 0-dim. features

## Transitioning to learning with PH

## Example (contd.):

Connectivity loss (ConnLoss)


Importantly,
$\triangleright$ the $l_{i}$ 's depend on the $\boldsymbol{A}_{i}$ 's (as they influence the $\boldsymbol{z}_{\mathfrak{i}}$ 's)

## Transitioning to learning with PH

## Example (contd.):

Connectivity loss (ConnLoss)


Importantly,
$\triangleright$ the $l_{i}$ 's depend on the $\boldsymbol{A}_{i}$ 's (as they influence the $\boldsymbol{z}_{i}$ 's)
$\triangleright$ minimizing the (joint) loss, requires gradients wrt. all $\boldsymbol{A}_{i}$ 's

## Transitioning to learning with PH

## Example (contd.):



Importantly,
$\triangleright$ the $l_{i}$ 's depend on the $\boldsymbol{A}_{i}$ 's (as they influence the $\boldsymbol{z}_{i}$ 's)
$\triangleright$ minimizing the (joint) loss, requires gradients wrt. all $\boldsymbol{A}_{i}$ 's

## Transitioning to learning with PH

## Example (contd.):



Importantly,
$\triangleright$ the $l_{i}$ 's depend on the $\boldsymbol{A}_{i}$ 's (as they influence the $\boldsymbol{z}_{i}$ 's)
$\triangleright$ minimizing the (joint) loss, requires gradients wrt. all $\boldsymbol{A}_{i}$ 's

## Transitioning to learning with PH

## Example (contd.):



Importantly,
$\triangleright$ the $l_{i}$ 's depend on the $\boldsymbol{A}_{i}$ 's (as they influence the $\boldsymbol{z}_{i}$ 's)
$\triangleright$ minimizing the (joint) loss, requires gradients wrt. all $\boldsymbol{A}_{i}$ 's
$\triangleright$ The good news is that this can be done
[Hofer et al., 2019] [Carrière et al., 2020] 区
[Brüel-Gabrielsson et al., 2019] 囚

## Transitioning to learning with PH

Lets look at some toy data first.


## Transitioning to learning with PH

Lets look at some toy data first.


Here's what we aim to do:
$\triangleright$ Compute 0-dim. Vietoris-Rips PH
$\triangleright$ Minimize ConnLoss wrt. the $x_{i}($ for a desired $\eta>0)$

## Transitioning to learning with PH

Lets look at some toy data first.


Here's what we aim to do:
$\triangleright$ Compute 0-dim. Vietoris-Rips PH
$\triangleright$ Minimize ConnLoss wrt. the $x_{i}($ for a desired $\eta>0)$
Notably, this controls the length of the minimal spanning tree (MST).
[Robins, 2000] 因

## Transitioning to learning with PH



## Transitioning to learning with PH



MST (after optimization)


## Some self-advertisement :)

Embedding into the PyTorch framework:

```
import torch
import numpy as np
from torchph.pershom import vr_persistence_ll
device = "cuda"
toy_data = np.random.rand(300, 2)
X = torch.tensor(toy_data, device=device, requires_grad=True)
opt = torch.optim.Adam([X], lr=0.01)
for i in range(1,100+1):
    pers = vr_persistence_l1(X, 1, 0)
    h_0 = pers[0][0]
    lt = h_0[:, 1] # HO lifetimes
    loss = (lt - 0.1).abs().sum()
    opt.zero_grad()
    loss.backward()
    opt.step()
```

Note that this uses our own PH implementation (works on GPU), see

## Why would this be useful?

In [Hofer et al., 2019] 因, we study ConnLoss with autoencoders.

## Why would this be useful?

In [Hofer et al., 2019] 因, we study ConnLoss with autoencoders.


## Why would this be useful?

In [Hofer et al., 2019] 因, we study ConnLoss with autoencoders.

persistence barcode of 0-dim. features

## Why would this be useful?

In [Hofer et al., 2019] 因, we study ConnLoss with autoencoders.

persistence barcode of 0-dim. features
Why? You might want to do kernel density estimation in $\mathcal{Z}\left(=\mathbb{R}^{n}\right)$

## Why would this be useful?

In [Hofer et al., 2019]図, we study ConnLoss with autoencoders.

persistence barcode of 0-dim. features
Why? You might want to do kernel density estimation in $\mathcal{Z}\left(=\mathbb{R}^{n}\right)$


Can be problematic, due to scale differences $\rightarrow$ we can impose scale via $\eta$

## Application: One-class learning

## Training (step I)

Trained only once using unlabeled data


Notably, [Moor et al., 2019] 図 follow similar ideas to learn a representation space (Z) that preserves the input space topology.

## Application: One-class learning

Training (step II)


## Application: One-class learning

## Training (step II)



## Evaluation protocol

Computation of a one-class score


Count \#samples falling into balls of radius $\eta / 2$, anchored at the one-class instances $\quad \square$

## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]因

## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]
Question: Can we control topological properties for generalization

## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]
Question: Can we control topological properties for generalization

Key idea: encourage "densification" of learned representations

## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]
Question: Can we control topological properties for generalization

Key idea: encourage "densification" of learned representations


## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]
Question: Can we control topological properties for generalization

Key idea: encourage "densification" of learned representations

${ }^{\dagger} \mathrm{c}: x \rightarrow y$ is an unknown labeling function

## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]
Question: Can we control topological properties for generalization

Key idea: encourage "densification" of learned representations


## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]
Question: Can we control topological properties for generalization

Key idea: encourage "densification" of learned representations

${ }^{\dagger} \mathrm{c}: x \rightarrow y$ is an unknown labeling function

## Application: Topological regularizers

How about neural classifiers? [Hofer et al., 2020]
Question: Can we control topological properties for generalization

Key idea: encourage "densification" of learned representations

${ }^{\dagger} \mathrm{c}: X \rightarrow y$ is an unknown labeling function

One aspect of the generalization puzzle in deep learning:

> Generalization in spite of memorization

## Application: Topological regularizers

One aspect of the generalization puzzle in deep learning:

## Generalization in spite of memorization

In fact, we can typically fit the training data without error, i.e., $L_{S}(F)=0$. (even under random labels [Zhang et al., 2017] (区)

## Application: Topological regularizers

One aspect of the generalization puzzle in deep learning:

## Generalization in spite of memorization

In fact, we can typically fit the training data without error, i.e., $L_{S}(F)=0$. (even under random labels [Zhang et al., 2017] (⿶)

Consider


In [Hofer et al., 2020]図, we study how the distribution around representations of training samples, $\varphi_{*}(\mathcal{P})$, affects generalization.

## Application: Topological regularizers

Lets decompose F as $\mathrm{F}=\gamma \circ \varphi: \mathcal{X} \rightarrow z \rightarrow \mathcal{y}$ with $\gamma(\boldsymbol{x})=\operatorname{sgn}\left(\boldsymbol{w}^{\top} \boldsymbol{x}\right)$.

Application: Topological regularizers

Lets decompose F as $\mathrm{F}=\gamma \circ \varphi: \mathcal{X} \rightarrow z \rightarrow \mathcal{y}$ with $\gamma(\boldsymbol{x})=\operatorname{sgn}\left(\boldsymbol{w}^{\top} \boldsymbol{x}\right)$.

$\triangleright \mathcal{Z}$ is the codomain of $\varphi, \gamma^{-1}(i)$ the decision region of class $i$

Application: Topological regularizers

Lets decompose F as $\mathrm{F}=\gamma \circ \varphi: \mathcal{X} \rightarrow z \rightarrow \mathcal{y}$ with $\gamma(\boldsymbol{x})=\operatorname{sgn}\left(\boldsymbol{w}^{\top} \boldsymbol{x}\right)$.

$\triangleright \mathcal{Z}$ is the codomain of $\varphi, \gamma^{-1}(i)$ the decision region of class $i$
$\triangleright$ Label-wise distribution, $\mathrm{Q}_{\mathfrak{i}}$ (restriction of $\varphi_{*}(\mathcal{P})$ to class $\mathfrak{i}$ ), in $\mathcal{Z}$

Application: Topological regularizers

Lets decompose F as $\mathrm{F}=\gamma \circ \varphi: \mathcal{X} \rightarrow z \rightarrow \mathcal{y}$ with $\gamma(\boldsymbol{x})=\operatorname{sgn}\left(\boldsymbol{w}^{\top} \boldsymbol{x}\right)$.

$\triangleright \mathcal{Z}$ is the codomain of $\varphi, \gamma^{-1}(i)$ the decision region of class $i$
$\triangleright$ Label-wise distribution, $\mathrm{Q}_{\mathfrak{i}}$ (restriction of $\varphi_{*}(\mathcal{P})$ to class $\mathfrak{i}$ ), in $\mathcal{Z}$
We aim for a densification of $Q_{i}$ via regularization of $\varphi$.

## Application: Topological regularizers

Lets take a closer look at densification.

## Application: Topological regularizers

Lets take a closer look at densification.
Consider, for a reference set $M \subset \mathcal{Z}$, its metric extension ${ }^{\dagger}$

$$
M_{\epsilon}=\bigcup_{x \in M} B(x, \epsilon), \quad \epsilon>0
$$

$$
{ }^{\dagger} B(x, e)=\{u \in z: d(x, u) \leqslant e\}
$$

## Application: Topological regularizers

Lets take a closer look at densification.
Consider, for a reference set $M \subset \mathcal{Z}$, its metric extension ${ }^{\dagger}$

$$
M_{\epsilon}=\bigcup_{x \in M} B(x, \epsilon), \quad \epsilon>0
$$

Question: How much mass is in the $\epsilon$-belt?

$$
{ }^{\dagger} B(x, \epsilon)=\{u \in z: d(x, u) \leqslant \epsilon\}
$$

## Application: Topological regularizers

Informally, densification means:
For a given mass in the reference set $M$, increase the mass concentrated in its $\epsilon$-extension!


## Application: Topological regularizers

The idea is to exert control over connectivity properties!

## Application: Topological regularizers

The idea is to exert control over connectivity properties!
Consider the (Euclidean) minimal spanning tree (MST) ${ }^{\dagger}$ :

## Application: Topological regularizers

The idea is to exert control over connectivity properties!
Consider the (Euclidean) minimal spanning tree (MST) ${ }^{\dagger}$ :


## Application: Topological regularizers

The idea is to exert control over connectivity properties!
Consider the (Euclidean) minimal spanning tree (MST) ${ }^{\dagger}$ :

## Application: Topological regularizers

The idea is to exert control over connectivity properties!
Consider the (Euclidean) minimal spanning tree (MST) ${ }^{\dagger}$ :


$$
\operatorname{len}\left(e_{i}\right)=d\left(z_{i_{1}}, z_{i_{2}}\right)
$$

## Application: Topological regularizers

The idea is to exert control over connectivity properties!
Consider the (Euclidean) minimal spanning tree (MST) ${ }^{\dagger}$ :


## Application: Topological regularizers

The idea is to exert control over connectivity properties!
Consider the (Euclidean) minimal spanning tree (MST) ${ }^{\dagger}$ :


The idea is to exert control over connectivity properties!
Consider the (Euclidean) minimal spanning tree (MST) ${ }^{\dagger}$ :
as $\varphi$ is parametrized by a neural network with parameters $\theta$

$$
\operatorname{len}\left(e_{i}\right)=\mathrm{d}\left(\varphi_{\theta}\left(x_{i_{1}}\right), \varphi_{\theta}\left(x_{i_{2}}\right)\right)
$$

Differentiable in $\theta$
$\Rightarrow$ we can control the edge lengths of the MST (as mentioned earlier)

$$
{ }^{\dagger} d(x, y)=\|x-y\|
$$

## Application: Topological regularizers

We call $z_{1}, \ldots, z_{\mathrm{b}} \in \mathcal{Z} \beta$-connected if all edges in the corresponding MST are not longer than $\beta$.


## Application: Topological regularizers

We call $z_{1}, \ldots, z_{\mathrm{b}} \in \mathcal{Z} \beta$-connected if all edges in the corresponding MST are not longer than $\beta$.


Application: Topological regularizers

We call $z_{1}, \ldots, z_{\mathrm{b}} \in Z \beta$-connected if all edges in the corresponding MST are not longer than $\beta$.


This allows us to talk about properties of $z_{1}, \ldots, z_{\mathrm{b}} \sim \mathrm{Q}$, i.e., b iid draws from Q.

## Application: Topological regularizers

Let $\mathrm{b} \in \mathbb{N}$. We call $\mathrm{Q} \operatorname{a} c_{\mathrm{b}}^{\beta}$-connected distribution if

$$
c_{\mathrm{b}}^{\beta} \leqslant \operatorname{Pr}\left[Z_{1}, \ldots, \mathrm{Z}_{\mathrm{b}} \text { are } \beta \text {-connected }\right]
$$

holds for $\mathrm{Z}_{1}, \ldots, \mathrm{Z}_{\mathrm{b}} \stackrel{\text { iid }}{\sim} \mathrm{Q}$ with $\beta>0, \mathrm{c}_{\mathrm{b}}^{\beta}>0$.

## Application: Topological regularizers

Let $\mathrm{b} \in \mathbb{N}$. We call $\mathrm{Q} \operatorname{ac} c_{b}^{\beta}$-connected distribution if

$$
c_{\mathrm{b}}^{\beta} \leqslant \operatorname{Pr}\left[Z_{1}, \ldots, \mathrm{Z}_{\mathrm{b}} \text { are } \beta \text {-connected }\right]
$$

holds for $\mathrm{Z}_{1}, \ldots, \mathrm{Z}_{\mathrm{b}} \stackrel{\text { iid }}{\sim} \mathrm{Q}$ with $\beta>0, \mathrm{c}_{\mathrm{b}}^{\beta}>0$.
This is a property of the product measure $Q^{b}$.

Application: Topological regularizers

Let $\mathrm{b} \in \mathbb{N}$. We call $\mathrm{Q} \operatorname{ac} c_{b}^{\beta}$-connected distribution if

$$
c_{\mathrm{b}}^{\beta} \leqslant \operatorname{Pr}\left[Z_{1}, \ldots, Z_{b} \text { are } \beta \text {-connected }\right]
$$

holds for $\mathrm{Z}_{1}, \ldots, \mathrm{Z}_{\mathrm{b}} \stackrel{\text { iid }}{\sim} \mathrm{Q}$ with $\beta>0, \mathrm{c}_{\mathrm{b}}^{\beta}>0$.
This is a property of the product measure $Q^{b}$.
Example: five draws from $\mathrm{Q}^{\mathrm{b}}$ with $\mathrm{b}=4$


Application: Topological regularizers

Let $\mathrm{b} \in \mathbb{N}$. We call $\mathrm{Q} \operatorname{a} c_{\mathrm{b}}^{\beta}$-connected distribution if

$$
c_{\mathrm{b}}^{\beta} \leqslant \operatorname{Pr}\left[Z_{1}, \ldots, Z_{b} \text { are } \beta \text {-connected }\right]
$$

holds for $\mathrm{Z}_{1}, \ldots, \mathrm{Z}_{\mathrm{b}} \stackrel{\text { iid }}{\sim} \mathrm{Q}$ with $\beta>0, \mathrm{c}_{\mathrm{b}}^{\beta}>0$.
This is a property of the product measure $Q^{b}$.
Example: five draws from $\mathrm{Q}^{\mathrm{b}}$ with $\mathrm{b}=4$


Application: Topological regularizers

Let $\mathrm{b} \in \mathbb{N}$. We call $\mathrm{Q} \operatorname{a} c_{\mathrm{b}}^{\beta}$-connected distribution if

$$
c_{b}^{\beta} \leqslant \operatorname{Pr}\left[Z_{1}, \ldots, Z_{b} \text { are } \beta \text {-connected }\right]
$$

holds for $\mathrm{Z}_{1}, \ldots, \mathrm{Z}_{\mathrm{b}} \stackrel{\text { iid }}{\sim} \mathrm{Q}$ with $\beta>0, \mathrm{c}_{\mathrm{b}}^{\beta}>0$.
This is a property of the product measure $Q^{b}$.
Example: five draws from $\mathrm{Q}^{\mathrm{b}}$ with $\mathrm{b}=4$


区 $\beta$-connected
$x$ not $\beta$-connected

## Application: Topological regularizers

1. We can show that controlling connectivity properties ( $\beta$-connectedness) of $Q^{b}$ leads to densification of $Q$.
2. We can show that densification directly relates to generalization.

## Application: Topological regularizers

Some results for a neural classifier ${ }^{\ddagger}$ on MNIST (10 classes) in a small sample-size regime (250 samples):
Vanilla $7.1+/-1.0$

## Application: Topological regularizers

Some results for a neural classifier ${ }^{\ddagger}$ on MNIST ( 10 classes) in a small sample-size regime (250 samples):

| Vanilla | $7.1+/-1.0$ |
| :--- | :--- |
| + Jacobian reg. | $6.2+/-0.8$ |
| + DeCov | $6.5+/-1.1$ |
| + VR | $6.1+/-0.5$ |
| + cw-CR | $7.0+/-0.6$ |
| + cw-VR | $6.2+/-0.8$ |

## Application: Topological regularizers

Some results for a neural classifier ${ }^{\ddagger}$ on MNIST (10 classes) in a small sample-size regime (250 samples):

| Vanilla | $7.1+/-1.0$ |
| :--- | :--- |
| + Jacobian reg. | $6.2+/-0.8$ |
| + DeCov | $6.5+/-1.1$ |
| + VR | $6.1+/-0.5$ |
| + cw-CR | $7.0+/-0.6$ |
| + cw-VR | $6.2+/-0.8$ |
| + ConnLoss (best) | $5.6+/-0.7$ |
| + ConnLoss |  |

${ }^{\dagger} \beta$ chosen via cross-validation on a small validation set

## Application: Topological regularizers

Some results for a neural classifier ${ }^{\ddagger}$ on CIFAR10 (10 classes) in a small sample-size regime ( 500 samples):
Vanilla 39.4 +/- 1.5

## Application: Topological regularizers

Some results for a neural classifier ${ }^{\ddagger}$ on CIFAR10 (10 classes) in a small sample-size regime ( 500 samples):

| Vanilla | $39.4+/-1.5$ |
| :--- | :--- |
| + Jacobian reg. | $39.7+/-2.0$ |
| + DeCov | $38.2+/-1.5$ |
| + VR | $38.6+/-1.4$ |
| + cW-CR | $39.0+/-1.9$ |
| + cw-VR | $38.5+/-1.6$ |

## Application: Topological regularizers

Some results for a neural classifier ${ }^{\ddagger}$ on CIFAR10 (10 classes) in a small sample-size regime ( 500 samples):

| Vanilla | $39.4+/-1.5$ |
| :--- | :--- |
| + Jacobian reg. | $39.7+/-2.0$ |
| + DeCov | $38.2+/-1.5$ |
| + VR | $38.6+/-1.4$ |
| + cw-CR | $39.0+/-1.9$ |
| + cw-VR $^{+}$ConnLoss (best) | $36.5+/-1.6$ |
| + ConnLoss $^{\dagger}$ | $36.8+/-1.2$ |

${ }^{\dagger} \beta$ chosen via cross-validation on a small validation set

## What＇s ahead of us？

There is so much exciting stuff that is going on right now！
Here are some examples ．．．
$\triangleright$ Theory for for optimizing PH－based functions
$\triangleright$ Studying learning behavior of neural networks
［Carrière et al．，2020］囚 ［Rieck et al．，2018］囚
$\triangleright$ PH for learning with graphs［Hofer et al．，2019；Rieck et al．2021］$\sqrt{\text { 人 }}$
$\triangleright$ Using simplicial complexes for message passing［Bodnar et al．，2021］⿴囗
$\triangleright$ Differentiable topology layers ［Brüel－Gabrielsson et al．，2019］囚
$\triangleright$ Topological attention for time－series forecasting
$\triangleright$ Topology－preserving image segmentation
$\triangleright$ Topological regularization of decision boundaries
［Zeng et al．，2021］因
［Hu et al．，2019］囚
［Chen et al．，2019］区

Again，this is，by far，not an exhaustive listing！

## What I（personally）find interesting

Continuing work along the lines of［Bianchini \＆Scarselli，2014］国，i．e．， using concepts from topology to study hypothesis set complexity．

## What I（personally）find interesting

Continuing work along the lines of［Bianchini \＆Scarselli，2014］国，i．e．， using concepts from topology to study hypothesis set complexity．

```
see also [Ramamurthy et al., 2019]园
    [Guss & Salakhutdinov, 2018]図
```

Can we possibly come up with other／better measures of quantifying hypothesis set complexity（similar to VC－dim．，or Rademacher complexity）？

## What I（personally）find interesting

Continuing work along the lines of［Bianchini \＆Scarselli，2014］囚，i．e．， using concepts from topology to study hypothesis set complexity．

```
see also [Ramamurthy et al., 2019]园
    [Guss & Salakhutdinov, 2018]図
```

Can we possibly come up with other／better measures of quantifying hypothesis set complexity（similar to VC－dim．，or Rademacher complexity）？

With differentiable layers for NN＇s that compute PH，we have a great tool －but，we do not really know what to do with it（yet）．

## Collaborators




Florian Graf
Univ. Salzburg


Chris Hofer
Univ. Salzburg


Stefan Huber IST Austria (back then) @shuber 3


