Invertible Neural Networks beyond Particle Physics

Speaker: Lynton Ardizzone

July 6, 2021

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Work in this talk by VLL & collaborators

VLL (Uni HD)
: Lynton Ardizzone, Jakob Kruse, Radek Mackowiak, Carsten Rother, Ullrich Köthe

DKFZ
: Tim Adler, Sebastian Wirkert, Lena Maier-Hein

ITA (Uni HD)
: Victor Ksoll, Ralf Klessen

Psychologisches Institut (Uni HD)
: Stefan Radev, Ulf Mertens, Andreas Voss

Institute for Wind Energy Systems (Uni Hannover)
: Pablo Noever-Castelos, Claudio Balzani

Methods

Inverse problems

Experimental design

Diverse generation

Anomaly detection

Generative classification*

*Skipped due to time

Methods

Inverse problems

Experimental design

Diverse generation

Anomaly detection

Generative classification*

*Skipped due to time

Fields

Medicine

Engineering

Chemistry

Computer vision

Astrophysics

Epidemiology

Medical imaging

Methods **Fields**

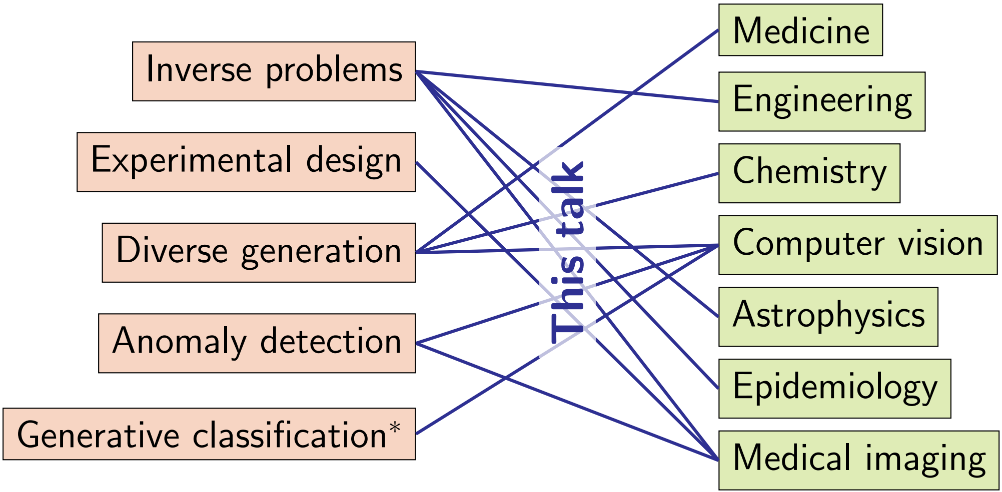

*Skipped due to time

Normalizing flows & INNs

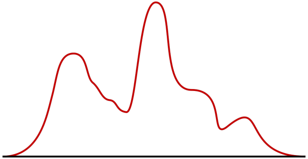

$q(x)$ (learned)

Normalizing flows & INNs

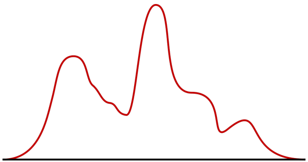

$q(x)$ (learned)

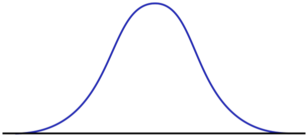

$p(z)$ (fixed)

Normalizing flows & INNs

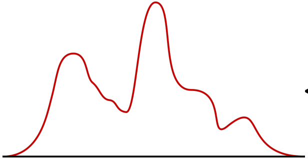

$$z = f(x)$$

$$x = f^{-1}(z)$$

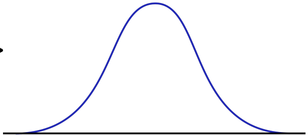

$q(x)$ (learned)

$p(z)$ (fixed)

Normalizing flows & INNs

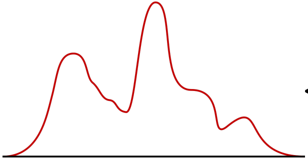

$$z = f(x)$$

$$x = f^{-1}(z)$$

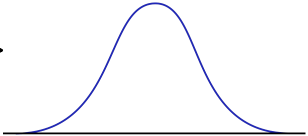

$q(x)$ (learned)

$p(z)$ (fixed)

$$q(x) =$$

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Normalizing flows & INNs

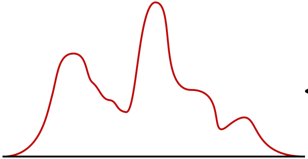

$q(x)$ (learned)

$$z = f(x)$$

$$x = f^{-1}(z)$$

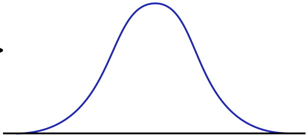

$p(z)$ (fixed)

$$q(x) = p\big(z = f(x)\big)\,|\det J|$$

Normalizing flows & INNs

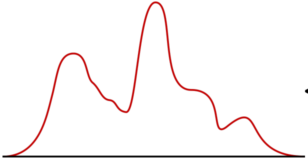

$q(x)$ (learned)

$$z = f(x)$$

$$x = f^{-1}(z)$$

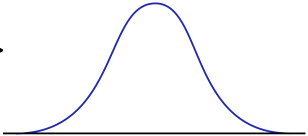

$p(z)$ (fixed)

$$q(x) = p\big(z{=}f(x)\big)\,|\det J|$$

INN

Normalizing flows & INNs

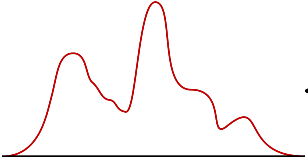

$q(x)$ (learned)

$$z = f(x)$$

$$x = f^{-1}(z)$$

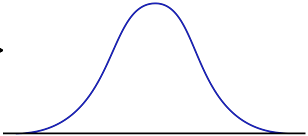

$p(z)$ (fixed)

$$q(x) = p\big(z{=}f(x)\big)\,|\det J|$$

INN

$|\det J|$, with $J = \dfrac{\partial f}{\partial x}$ (free with INNs)

Normalizing flows & INNs

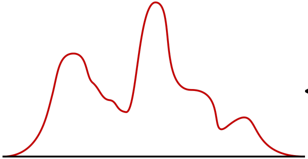

$$z = f(x) \longrightarrow$$
$$\longleftarrow x = f^{-1}(z)$$

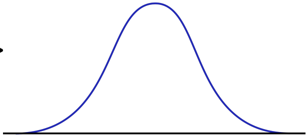

$q(x)$ (learned)

$p(z)$ (fixed)

$$q(x) = p(z{=}f(x))\,|\det J|$$

INN

$$J = \frac{\partial f}{\partial x}$$ (free with INNs)

- Max likelihood training: $\max_f \log q(x_{\text{train}})$

Normalizing flows & INNs

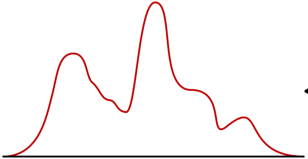

$$z = f(x)$$

$$x = f^{-1}(z)$$

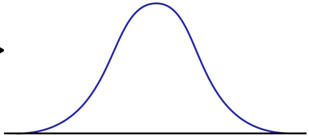

$q(x)$ (learned)

$p(z)$ (fixed)

$$q(x) = p\big(z{=}f(x)\big)\,|\det J|$$

INN

$$J = \frac{\partial f}{\partial x}$$ (free with INNs)

- Max likelihood training: $\max_f \log q(x_{\text{train}})$
- Cheap sampling: $x_{\text{samp}} = f^{-1}(z_{\text{samp}})$

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Normalizing flows & INNs

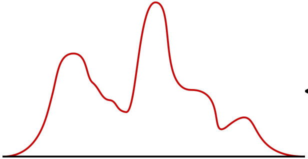

$q(x)$ (learned)

$$z = f(x)$$

$$x = f^{-1}(z)$$

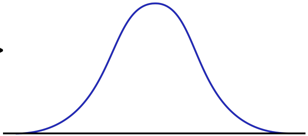

$p(z)$ (fixed)

$$q(x) = p\big(z = f(x)\big)\,|\det J|$$

INN

$J = \dfrac{\partial f}{\partial x}$ (free with INNs)

- Max likelihood training: $\max_f \log q(x_{\text{train}})$
- Cheap sampling: $x_{\text{samp}} = f^{-1}(z_{\text{samp}})$
- Various INN architectures (Dinh et al. (2017); Behrmann et al. (2019); Grathwohl et al. (2018),...)

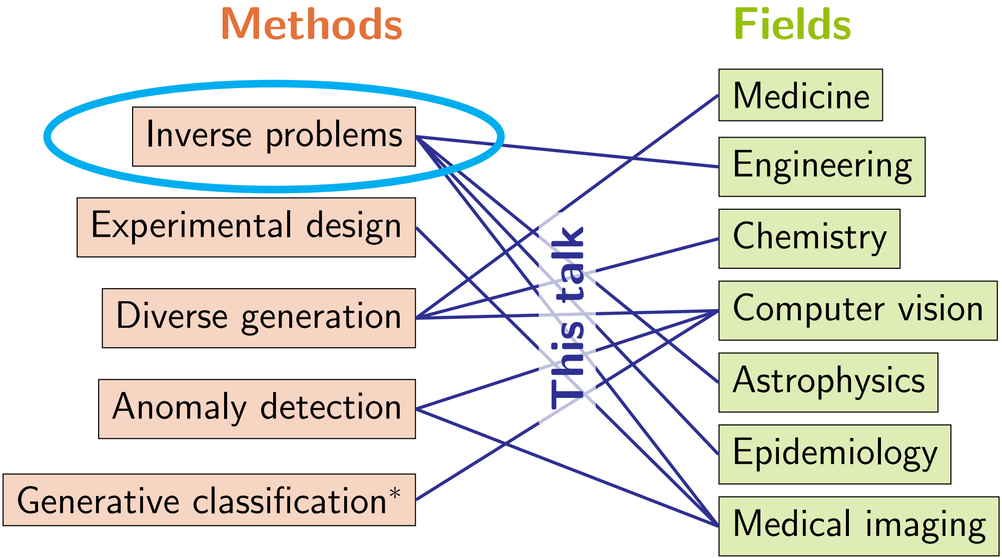

*Skipped due to time

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Inverse Problems

System parameter $\hat{x}$ $\longrightarrow$ known 'forward process' $\longrightarrow$ y observable

See also: Kruse et al. (2021b); Ardizzone et al. (2019a); Kruse et al. (2021a)

Inverse Problems

System parameter $\hat{x}$ $\longrightarrow$ known 'forward process' $\longrightarrow$ y observable

Uncertainty $q(x|y)$ $\longleftarrow$ **inverse problem** $\longleftarrow$ y observable

- Ambiguities
- Noisy process
- Finite cpacity

See also: Kruse et al. (2021b); Ardizzone et al. (2019a); Kruse et al. (2021a)

Conditional INNs (cINNs, Ardizzone et al., 2019b)

$$x \longleftrightarrow \boxed{\begin{array}{c}\text{INN}\\ z = f(x)\\ x = f^{-1}(z)\end{array}} \longleftrightarrow z$$

Standard: $q(x) = p\big(z = f(x)\big)\,|\det J|$

Conditional INNs (cINNs, Ardizzone et al., 2019b)

INN

$$x \longleftrightarrow \boxed{\begin{array}{c} z = f(x) \\ x = f^{-1}(z) \end{array}} \longleftrightarrow z$$

Standard: $q(x) = \; p\big(z{=}f(x)\big)\,|\det J|$

Conditional: $q(x \mid y) =$

Conditional INNs (cINNs, Ardizzone et al., 2019b)

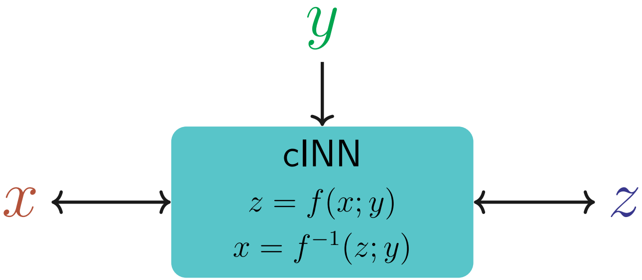

Standard: $q(x) = \; p\big(z{=}f(x)\big)\,|\det J|$

Conditional: $q(x \mid y) =$

Conditional INNs (cINNs, Ardizzone et al., 2019b)

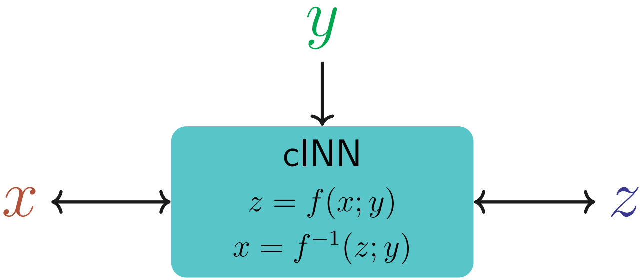

$$\text{Standard: } q(x) = p\big(z = f(x)\big)\,|\det J|$$

$$\text{Conditional: } q(x \mid y) = p\big(z = f(x; y)\big)\,|\det J|$$

Inverse problem – Stellar evolution (Ksoll et al., 2020)

Estimate physical properties of stars x from observed spectral lines y

Inverse problem – Stellar evolution (Ksoll et al., 2020)

Estimate physical properties of stars x from observed spectral lines y (not shown)

Luminosity L

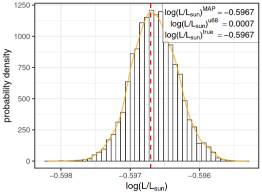

Mass M

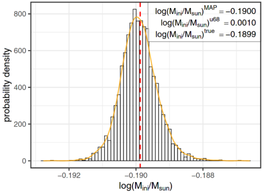

Surface gravity g

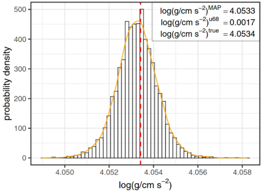

Inverse problem – Stellar evolution (Ksoll et al., 2020)

Estimate physical properties of stars x from observed spectral lines y (not shown)

Luminosity L Mass M Surface gravity g

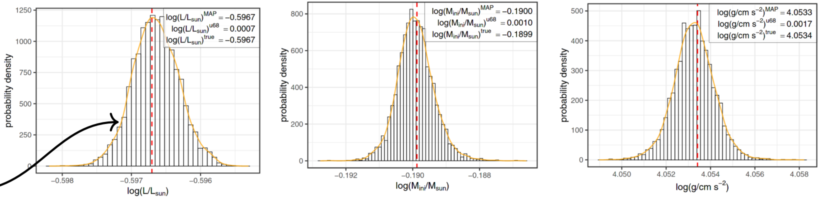

$q(L|y)$

Inverse problem – Stellar evolution (Ksoll et al., 2020)

Estimate physical properties of stars x from observed spectral lines y (not shown)

Luminosity L

True answer

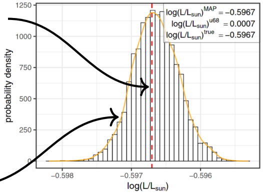

$q(L|y)$

Mass M

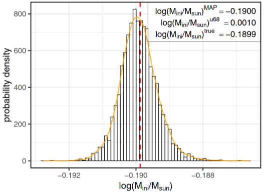

Surface gravity g

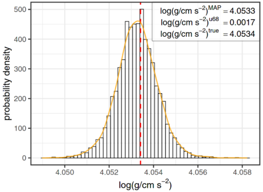

Inverse problem – Stellar evolution (Ksoll et al., 2020)

Estimate physical properties of stars x from observed spectral lines y (not shown)

Luminosity L | Mass M | Surface gravity g

True answer

$q(L|y)$

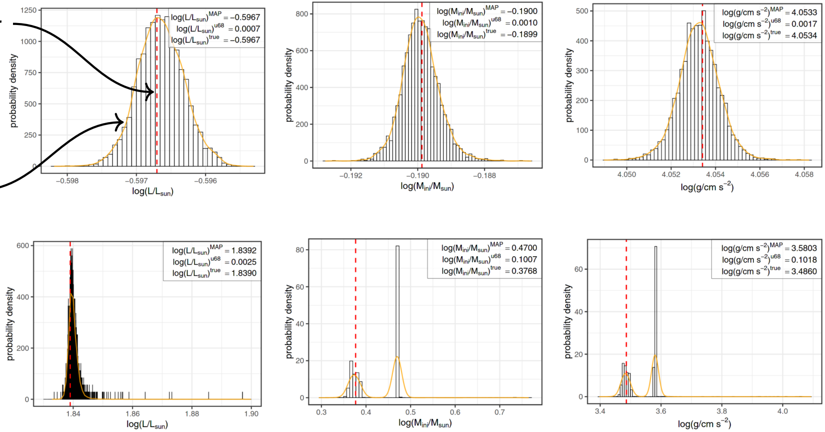

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Stellar evolution (Ksoll et al., 2020)

Full uncertainty distributions can **point to new physics**:

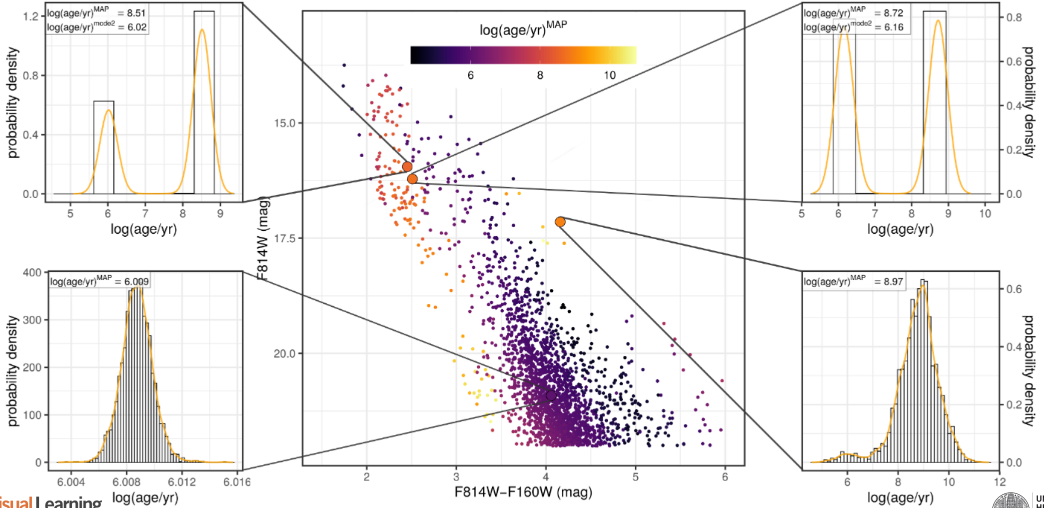

Inverse problem – Wind turbine design (Noever-Castelos et al., 2021)

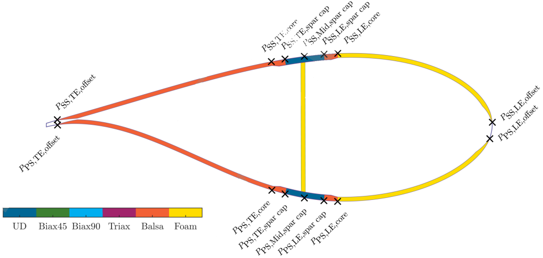

x: full physical model

y: diameters, weight

Inverse problem – Wind turbine design (Noever-Castelos et al., 2021)

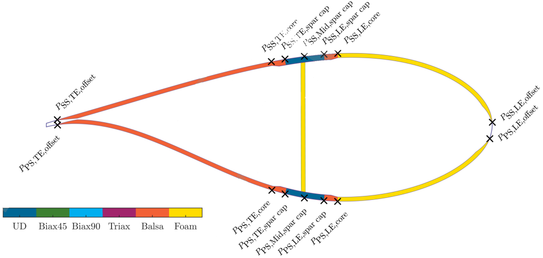

x: full physical model

y: diameters, weight

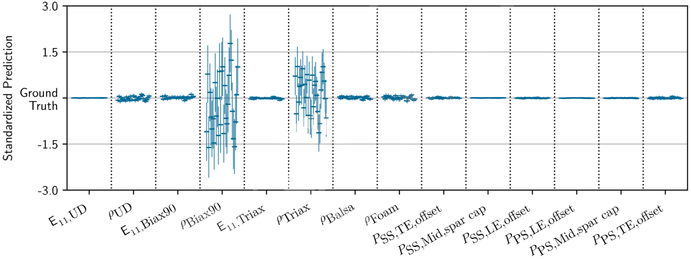

Inverse problem – Wind turbine design (Noever-Castelos et al., 2021)

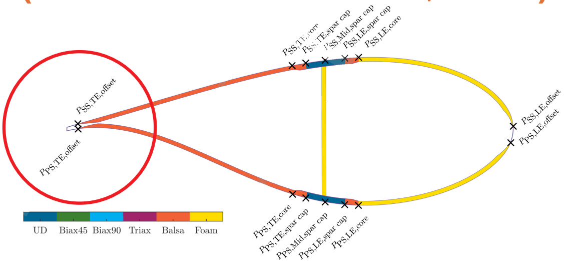

x: full physical model

y: diameters, weight

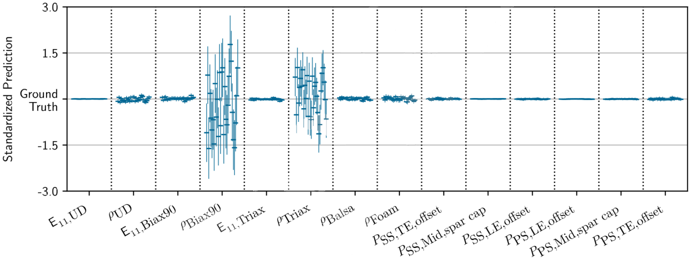

Inverse problem – Wind turbine design (Noever-Castelos et al., 2021)

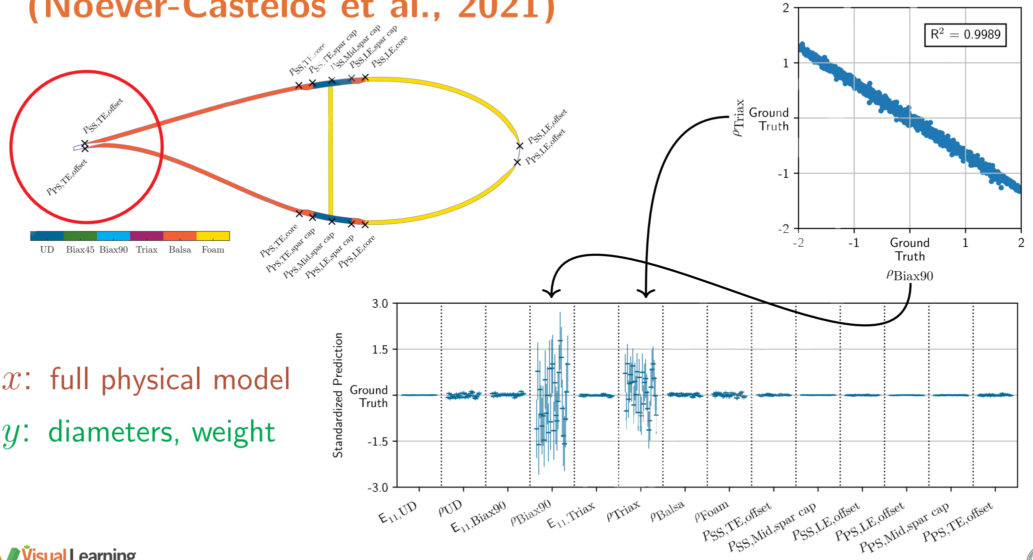

x: full physical model

y: diameters, weight

Inverse problem – Epidemiology (Radev et al., 2020)

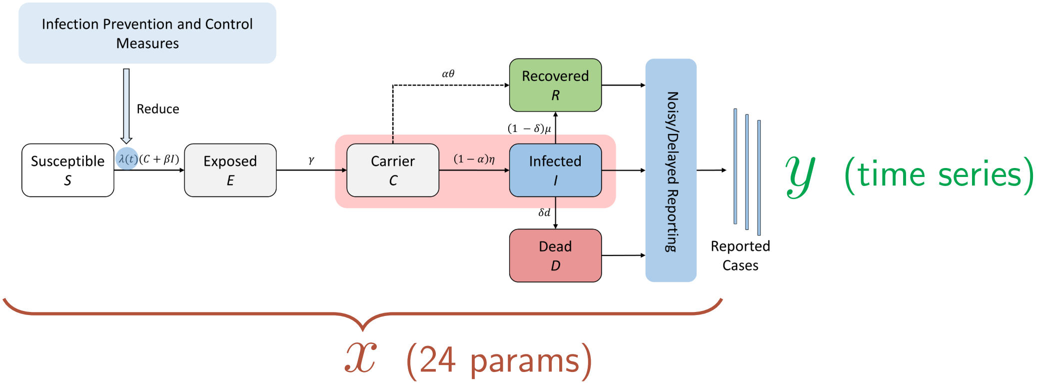

Inverse problem – Epidemiology (Radev et al., 2020)

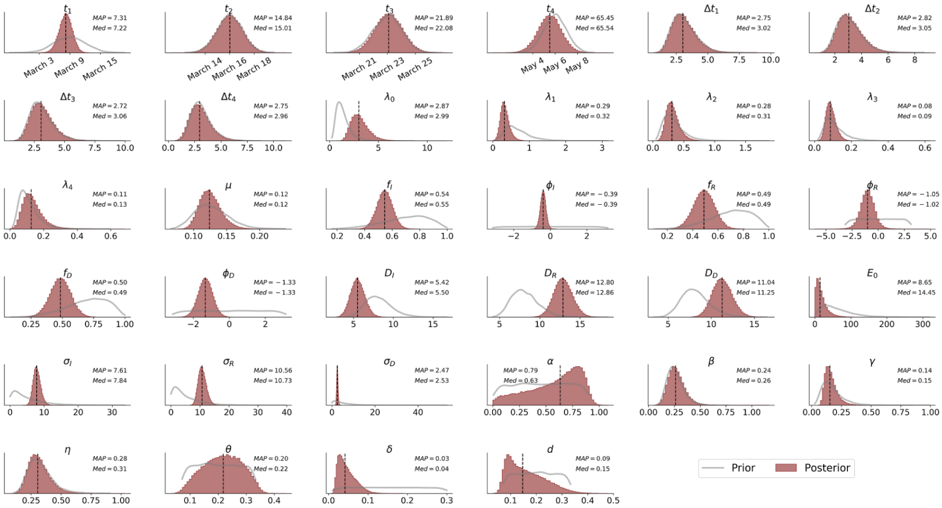

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Inverse problem – Epidemiology (Radev et al., 2020)

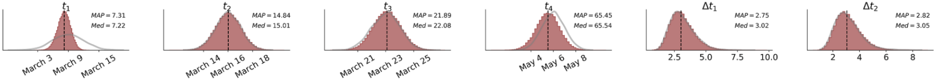

Simulation with predicted params matches obervations:

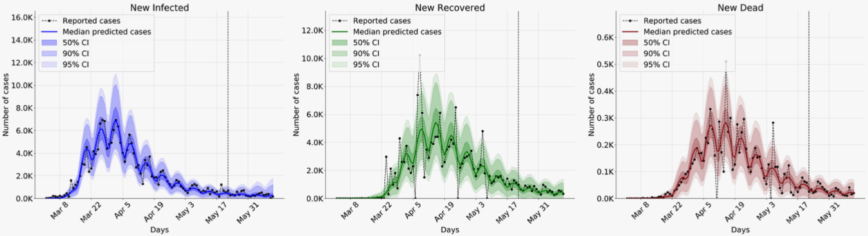

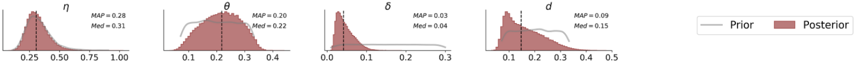

Methods Fields

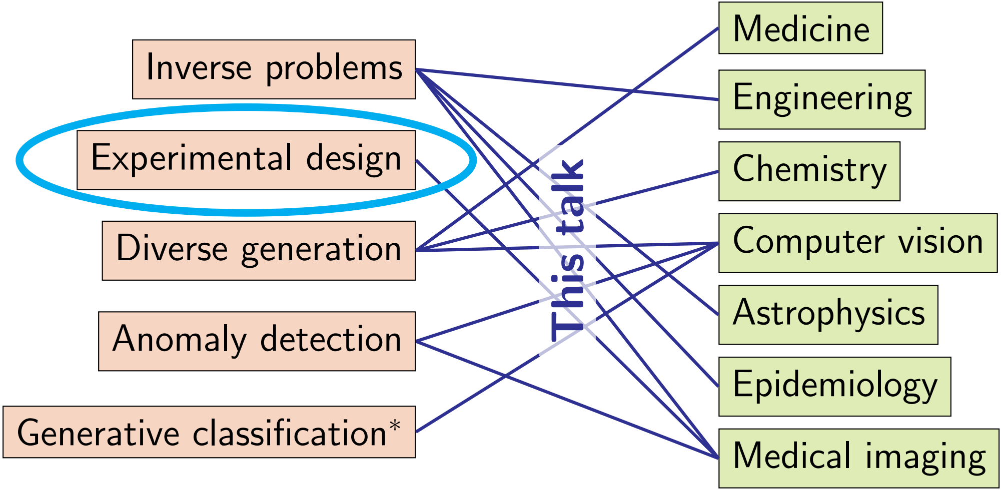

*Skipped due to time

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Experimental Design (Adler et al., 2019a)

(Multispectral) camera during surgery

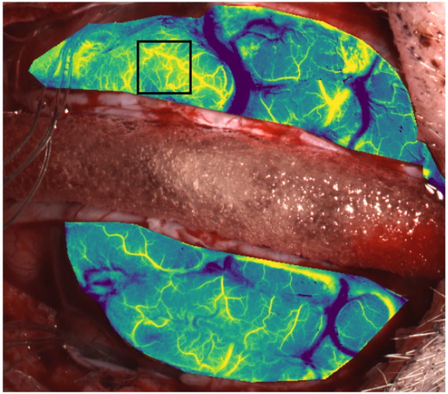

Experimental Design (Adler et al., 2019a)

(Multispectral) camera during surgery
Determine oxygen, blood flow, ...
from spectrum in each pixel

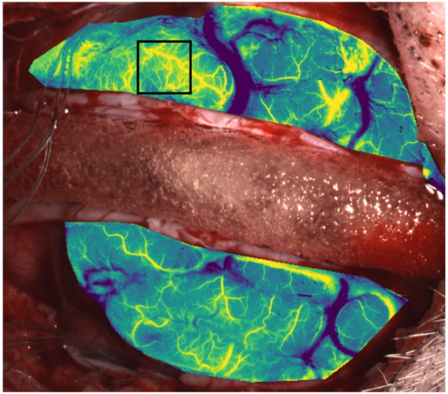

Experimental Design (Adler et al., 2019a)

(Multispectral) camera during surgery

Determine oxygen, blood flow, ...

from spectrum in each pixel

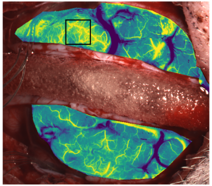

Experimental Design (Adler et al., 2019a)

(Multispectral) camera during surgery
Determine oxygen, blood flow, ...
from spectrum in each pixel

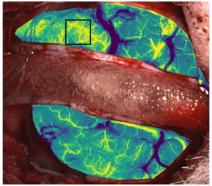

Simulate four camera models:

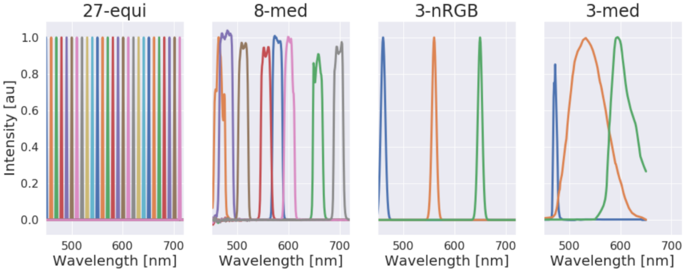

Experimental Design (Adler et al., 2019a)

(Multispectral) camera during surgery
Determine oxygen, blood flow, ...
from spectrum in each pixel

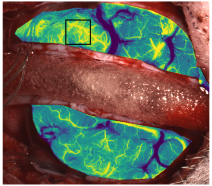

Simulate four camera models:

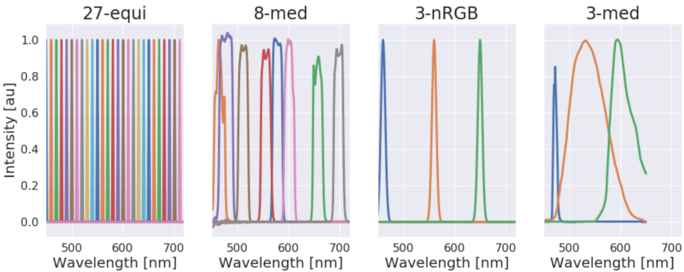

Percentage of ambiguous cases:

Experimental Design (Adler et al., 2019a)

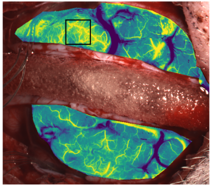

(Multispectral) camera during surgery
Determine oxygen, blood flow, ...
from spectrum in each pixel

Simulate four camera models:

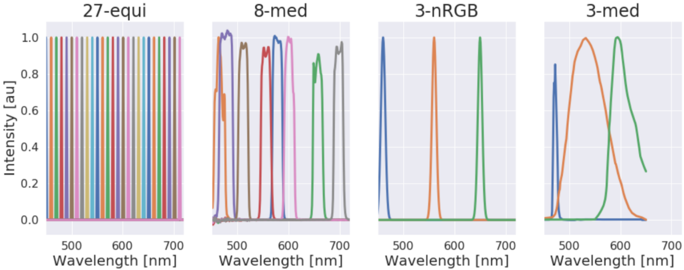

Percentage of ambiguous cases: 0.3% 0.5%

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Experimental Design (Adler et al., 2019a)

(Multispectral) camera during surgery
Determine oxygen, blood flow, ...
from spectrum in each pixel

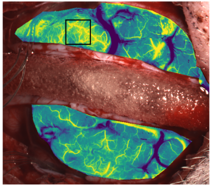

Simulate four camera models:

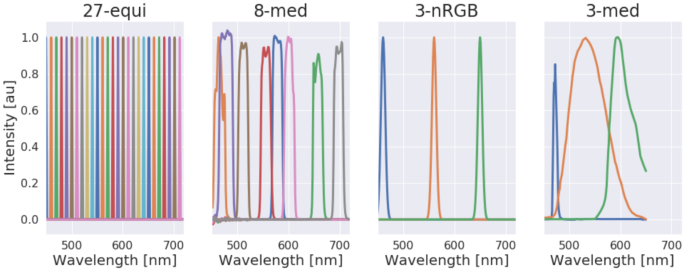

Percentage of ambiguous cases: 0.3% 0.5% 9.3%

Experimental Design (Adler et al., 2019a)

(Multispectral) camera during surgery
Determine oxygen, blood flow, ...
from spectrum in each pixel

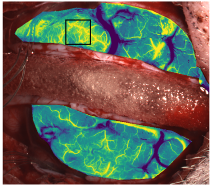

Simulate four camera models:

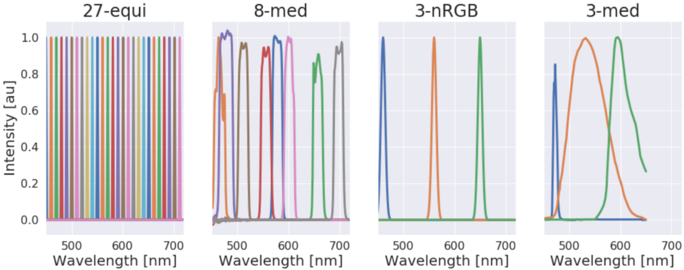

Percentage of ambiguous cases:

	27-equi	8-med	3-nRGB	3-med
Percentage of ambiguous cases	0.3%	0.5%	9.3%	3.6%

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

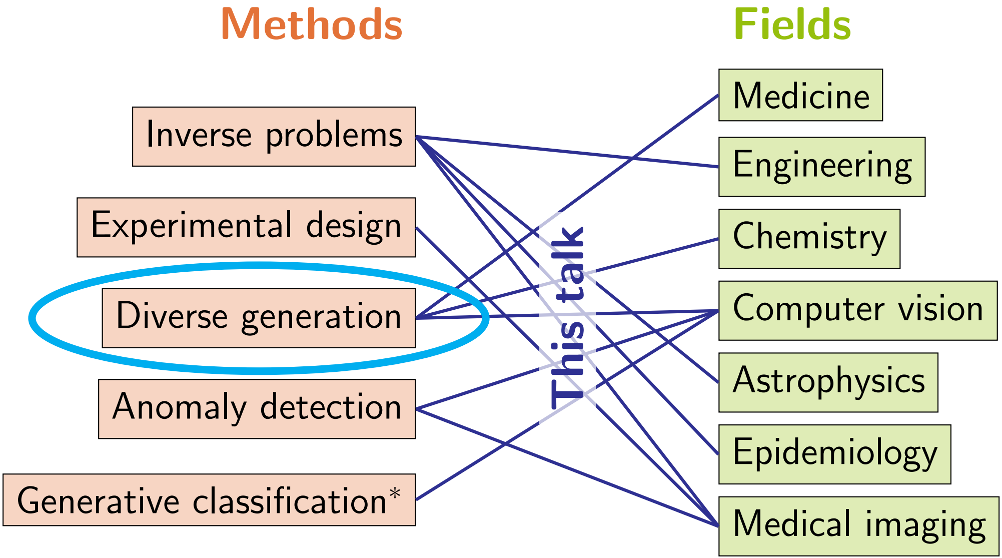

*Skipped due to time

Diverse Generation

Sampling from a distribution $p(x)$
is too hard if

Diverse Generation

Sampling from a distribution $p(x)$
is too hard if

- x very high dimensional

Diverse Generation

Sampling from a distribution $p(x)$
is too hard if

- x very high dimensional
- $p(x)$ not explicitly known

Diverse Generation

Sampling from a distribution $p(x)$
is too hard if

- x very high dimensional
- $p(x)$ not explicitly known

Use NF for $q(x) \approx p(x)$

Diverse Generation

Sampling from a distribution $p(x)$
is too hard if

- x very high dimensional
- $p(x)$ not explicitly known

Hypothesis generation

Use NF for $q(x) \approx p(x)$

Diverse Generation

Sampling from a distribution $p(x)$ is too hard if

- x very high dimensional
- $p(x)$ not explicitly known

Use NF for $q(x) \approx p(x)$

Hypothesis generation

Proposals

Diverse Generation

Sampling from a distribution $p(x)$
is too hard if

- x very high dimensional
- $p(x)$ not explicitly known

Use NF for $q(x) \approx p(x)$

- Hypothesis generation
- Proposals
- MCMC prior

Diverse Generation

Sampling from a distribution $p(x)$
is too hard if

- x very high dimensional
- $p(x)$ not explicitly known

Use NF for $q(x) \approx p(x)$

- Hypothesis generation
- Proposals
- MCMC prior
- ...

Diverse Generation

Sampling from a distribution $p(x)$ (or $p(x \mid y)$) is too hard if

- x very high dimensional
- $p(x)$ not explicitly known

Use NF for $q(x) \approx p(x)$
(or $q(x \mid y) \approx p(x \mid y)$)

- Hypothesis generation
- Proposals
- MCMC prior
- ...

Diverse Generation with cINNs (Ardizzone et al., 2021)

Condition y Sampled $x \sim q(x|y)$ Original $\hat{x}$

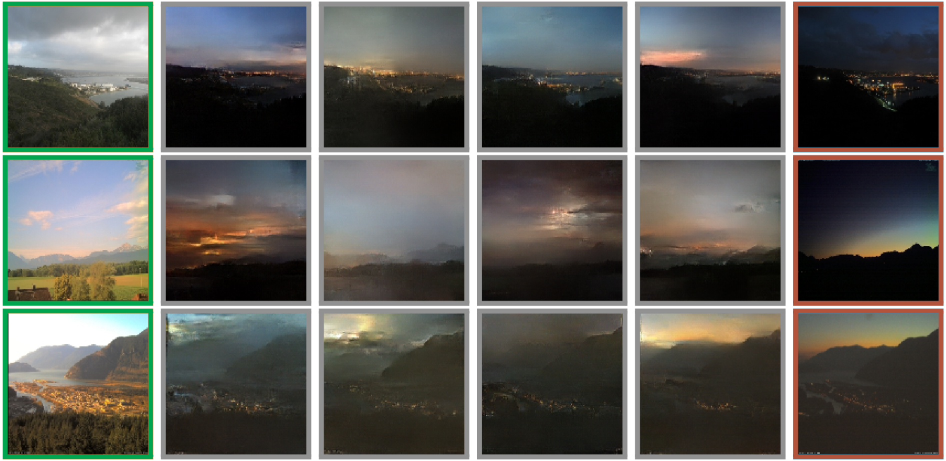

Generating plausible molecules (Satorras et al., 2021)

- Dataset of known moleclues $\rightarrow$ generate more
- $E(n)$-equivariant INN

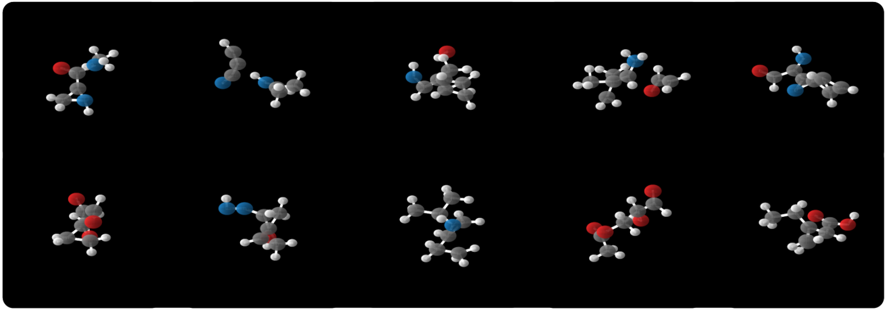

Adenoviruses as localized vectors (in progress)

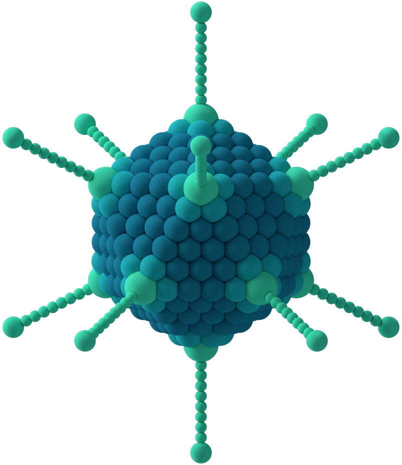

- Adenoviruses for vaccination, cancer treatment, gene therapy, ...

Adenoviruses as localized vectors (in progress)

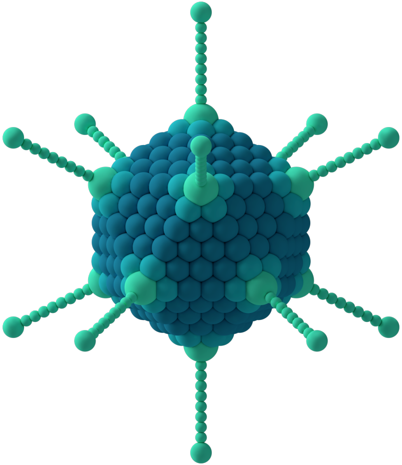

- Adenoviruses for vaccination, cancer treatment, gene therapy, ...
- Fibers (‘antennas’) to target certain organs
- 40^{10} possible fibers

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Adenoviruses as localized vectors (in progress)

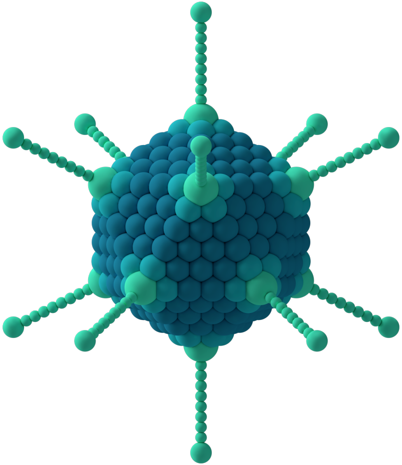

- Adenoviruses for vaccination, cancer treatment, gene therapy, ...
- Fibers ('antennas') to target certain organs
- 40^{10} possible fibers

$\rightarrow$ Generate candidates

Methods

Fields

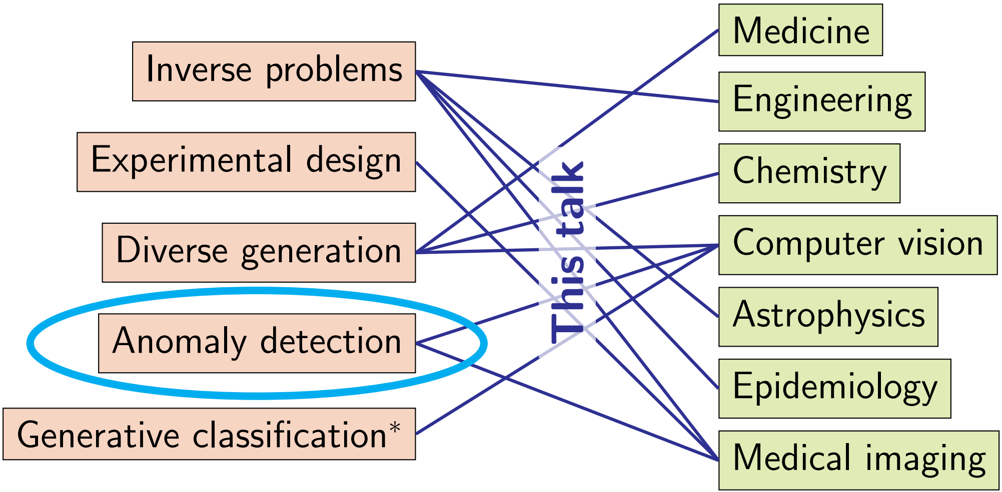

*Skipped due to time

Anomaly Detection with Normalizing Flows

Detect anomalies without knowing what they are?

Reading: Nalisnick et al. (2019); Choi et al. (2018); Serrà et al. (2019); Mackowiak et al. (2021)

Anomaly Detection with Normalizing Flows

Detect anomalies without knowing what they are?

- Train NF $q(x)$ on in-distribution data

Reading: Nalisnick et al. (2019); Choi et al. (2018); Serrà et al. (2019); Mackowiak et al. (2021)

Anomaly Detection with Normalizing Flows

Detect anomalies without knowing what they are?

- Train NF $q(x)$ on in-distribution data
- $q(x^*)$ very low $\implies$ x^* from out-of-distribution

Reading: Nalisnick et al. (2019); Choi et al. (2018); Serrà et al. (2019); Mackowiak et al. (2021)

Anomaly Detection with Normalizing Flows

Detect anomalies without knowing what they are?

- Train NF $q(x)$ on in-distribution data
- $q(x^*)$ very low $\implies$ x^* from out-of-distribution
- Hypothesis testing framework

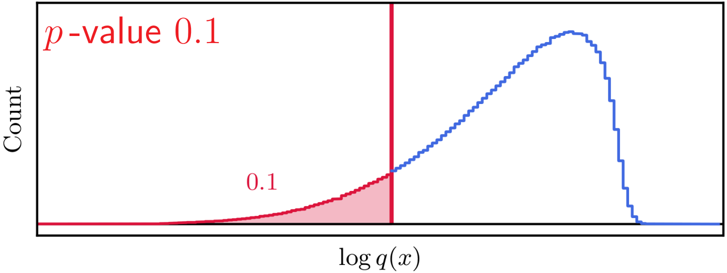

Reading: Nalisnick et al. (2019); Choi et al. (2018); Serrà et al. (2019); Mackowiak et al. (2021)

Anomaly Detection with Normalizing Flows

Detect anomalies without knowing what they are?

- Train NF $q(x)$ on in-distribution data
- $q(x^*)$ very low or high $\implies$ x^* from out-of-distribution
- Hypothesis testing framework

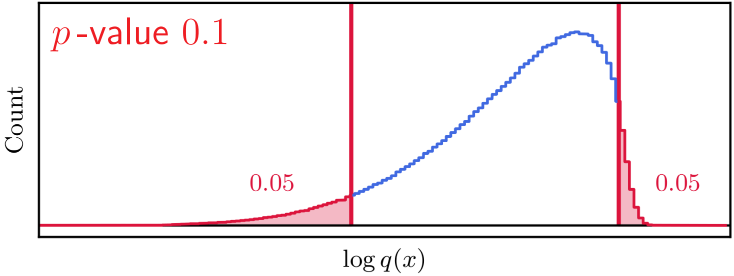

Reading: Nalisnick et al. (2019); Choi et al. (2018); Serrà et al. (2019); Mackowiak et al. (2021)

Anomaly detection in multispectral medical imaging (Adler et al., 2019b)

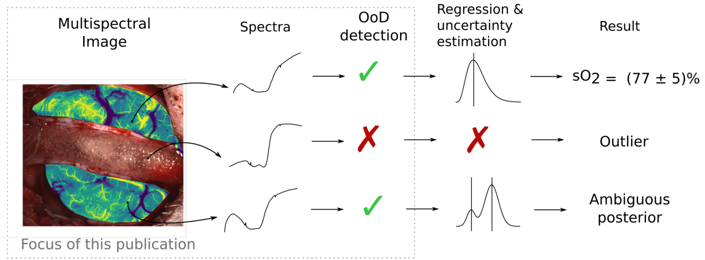

Anomaly detection in multispectral medical imaging (Adler et al., 2019b)

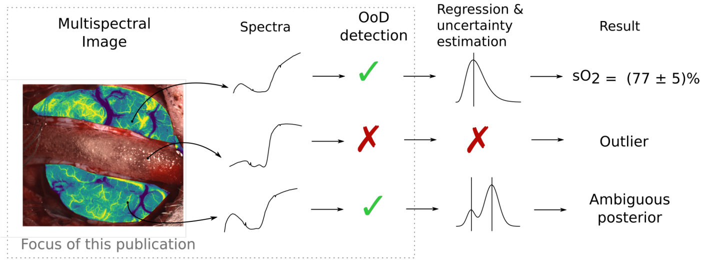

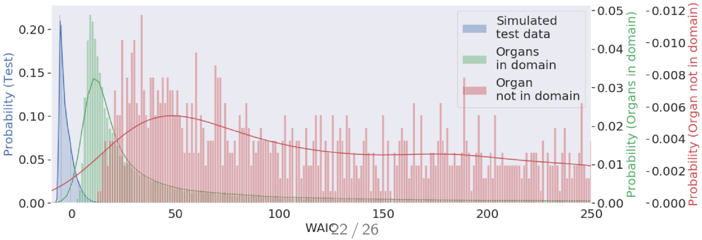

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

Methods | Fields

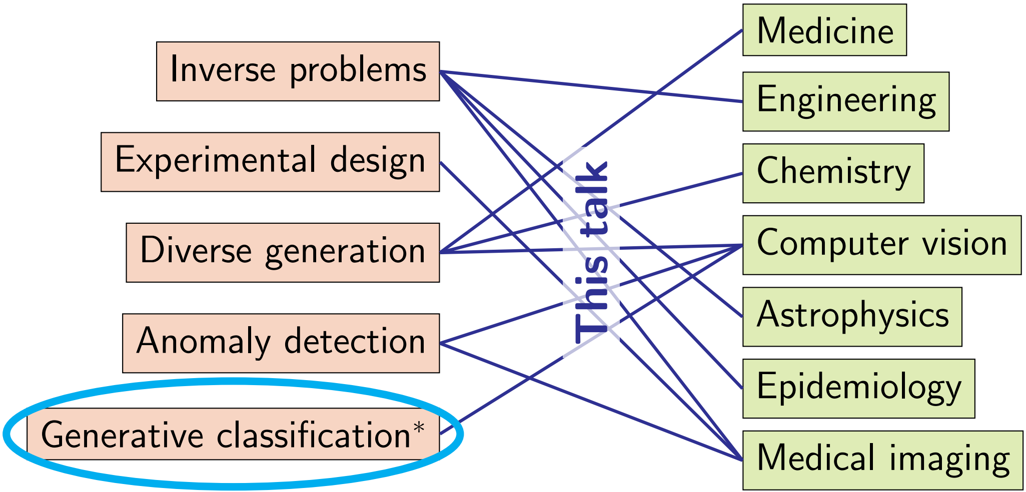

*Skipped due to time

Methods

- Inverse problems
- Experimental design
- Diverse generatio
- Anomaly detection
- Generative classification*

Fields

- Medicine
- Engineering
- Me ing

G.C. very exciting and useful!
See papers & talks:
Link to videos
Ardizzone et al. (2020)
Mackowiak et al. (2021)

*Skipped due to time

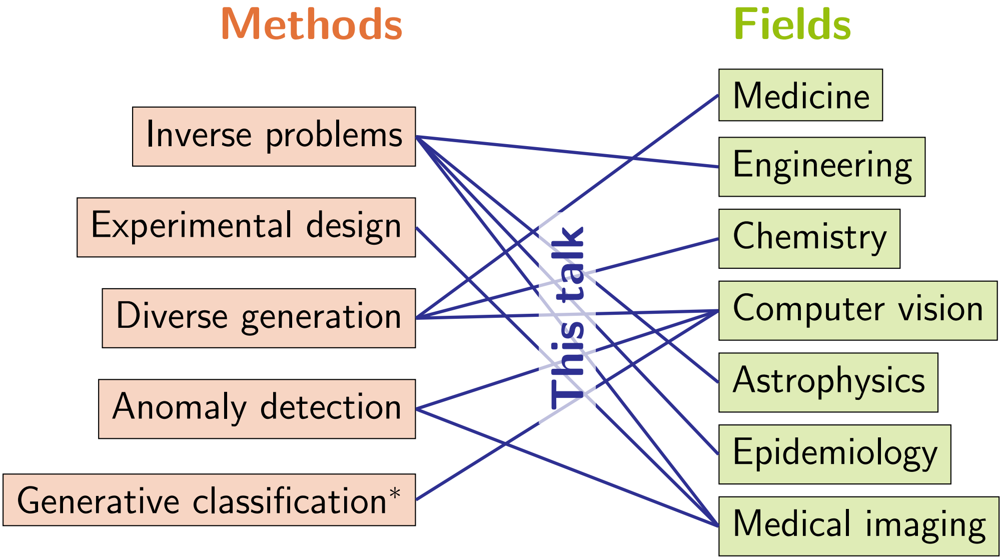

*Skipped due to time

Visual Learning Lab Heidelberg

UNIVERSITÄT HEIDELBERG ZUKUNFT SEIT 1386

References I

Adler, T. J., Ardizzone, L., Vemuri, A., Ayala, L., Gröhl, J., Kirchner, T., Wirkert, S., Kruse, J., Rother, C., Köthe, U., et al. Uncertainty-aware performance assessment of optical imaging modalities with invertible neural networks. *International journal of computer assisted radiology and surgery*, 14(6):997–1007, 2019a.

Adler, T. J., Ayala, L., Ardizzone, L., Kenngott, H. G., Vemuri, A., Müller-Stich, B. P., Rother, C., Köthe, U., and Maier-Hein, L. Out of distribution detection for intra-operative functional imaging. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging and Clinical Image-Based Procedures*, pp. 75–82. Springer, 2019b.

Ardizzone, L., Kruse, J., Rother, C., and Köthe, U. Analyzing inverse problems with invertible neural networks. In *Intl. Conf. on Learning Representations*, 2019a.

Ardizzone, L., Lüth, C., Kruse, J., Rother, C., and Köthe, U. Guided image generation with conditional invertible neural networks. *CoRR*, abs/1907.02392, 2019b. URL `http://arxiv.org/abs/1907.02392`.

Ardizzone, L., Mackowiak, R., Rother, C., and Köthe, U. Training normalizing flows with the information bottleneck for competitive generative classification. *Advances in Neural Information Processing Systems*, 33, 2020.

Ardizzone, L., Kruse, J., Lüth, C., Bracher, N., Rother, C., and Köthe, U. Conditional invertible neural networks for diverse image-to-image translation. In *Pattern Recognition: 42nd DAGM German Conference, DAGM GCPR 2020, Tübingen, Germany, September 28–October 1, 2020, Proceedings 42*, pp. 373–387. Springer, 2021.

Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D., and Jacobsen, J.-H. Invertible residual networks. In *International Conference on Machine Learning*, pp. 573–582. PMLR, 2019.

Choi, H., Jang, E., and Alemi, A. A. Waic, but why? generative ensembles for robust anomaly detection. *arXiv preprint arXiv:1810.01392*, 2018.

Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real NVP. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017. URL `https://openreview.net/forum?id=HkpbnH9lx`.

Grathwohl, W., Chen, R. T., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models. In *International Conference on Learning Representations*, 2018.

References II

Kruse, J., Ardizzone, L., Rother, C., and Köthe, U. Benchmarking invertible architectures on inverse problems. *arXiv preprint arXiv:2101.10763*, 2021a.

Kruse, J., Detommaso, G., Köthe, U., and Scheichl, R. Hint: Hierarchical invertible neural transport for density estimation and bayesian inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 8191–8199, 2021b.

Ksoll, V. F., Ardizzone, L., Klessen, R., Koethe, U., Sabbi, E., Robberto, M., Gouliermis, D., Rother, C., Zeidler, P., and Gennaro, M. Stellar parameter determination from photometry using invertible neural networks. *Monthly Notices of the Royal Astronomical Society*, 499(4): 5447–5485, 2020.

Mackowiak, R., Ardizzone, L., Kothe, U., and Rother, C. Generative classifiers as a basis for trustworthy image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2971–2981, 2021.

Nalisnick, E., Matsukawa, A., Teh, Y. W., and Lakshminarayanan, B. Detecting out-of-distribution inputs to deep generative models using a test for typicality. *arXiv preprint arXiv:1906.02994*, 2019.

Noever-Castelos, P., Ardizzone, L., and Balzani, C. Model updating of wind turbine blade cross sections with invertible neural networks. 2021.

Radev, S. T., Graw, F., Chen, S., Mutters, N. T., Eichel, V. M., Bärnighausen, T., and Köthe, U. Model-based bayesian inference of disease outbreak dynamics with invertible neural networks an application to the covid-19 pandemics in germany. 2020.

Satorras, V. G., Hoogeboom, E., Fuchs, F. B., Posner, I., and Welling, M. E (n) equivariant normalizing flows for molecule generation in 3d. *arXiv preprint arXiv:2105.09016*, 2021.

Serrà, J., Álvarez, D., Gómez, V., Slizovskaia, O., Núñez, J. F., and Luque, J. Input complexity and out-of-distribution detection with likelihood-based generative models. In *International Conference on Learning Representations*, 2019.