AI for CMR across scanners, sites and populations

Prof. Sotirios A. Tsaftaris
(Sotos)

Canon Medical / Royal Academy of Engineering Research Chair in Healthcare AI
Professor in Machine Learning and Computer Vision
Turing Fellow & ELLIS Fellow

http://vios.science
S.Tsaftaris@ed.ac.uk
@STsaftaris

Declaration of Conflicts

- Receive grant support from Canon Medical

A visual bio

Dipl. ECE @ Aristotle University

PhD & MSc @ EECS/ECE Northwestern

Faculty @ Northwestern
EECS/Radiology

Director of PRIAn @ IMT

@ University of Edinburgh

- Canon Medical/RAEng
- Prof/Chair in ML, CV
- Turing Fellow
- ELLIS Fellow

At VIOS we do interdisciplinary AI

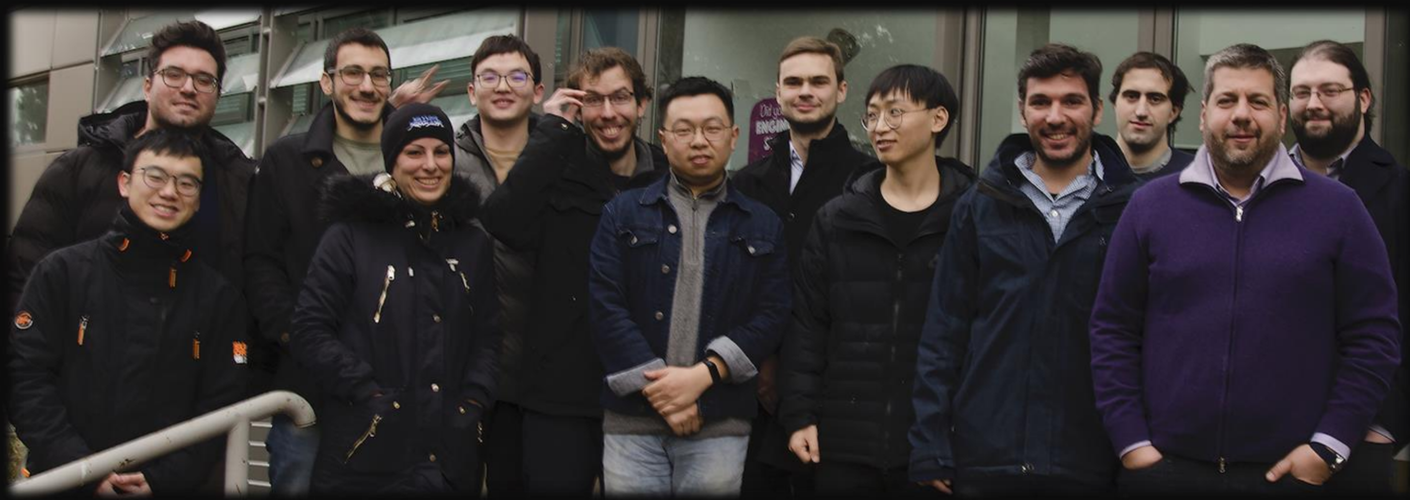

CMR a key tool in clinical worklows

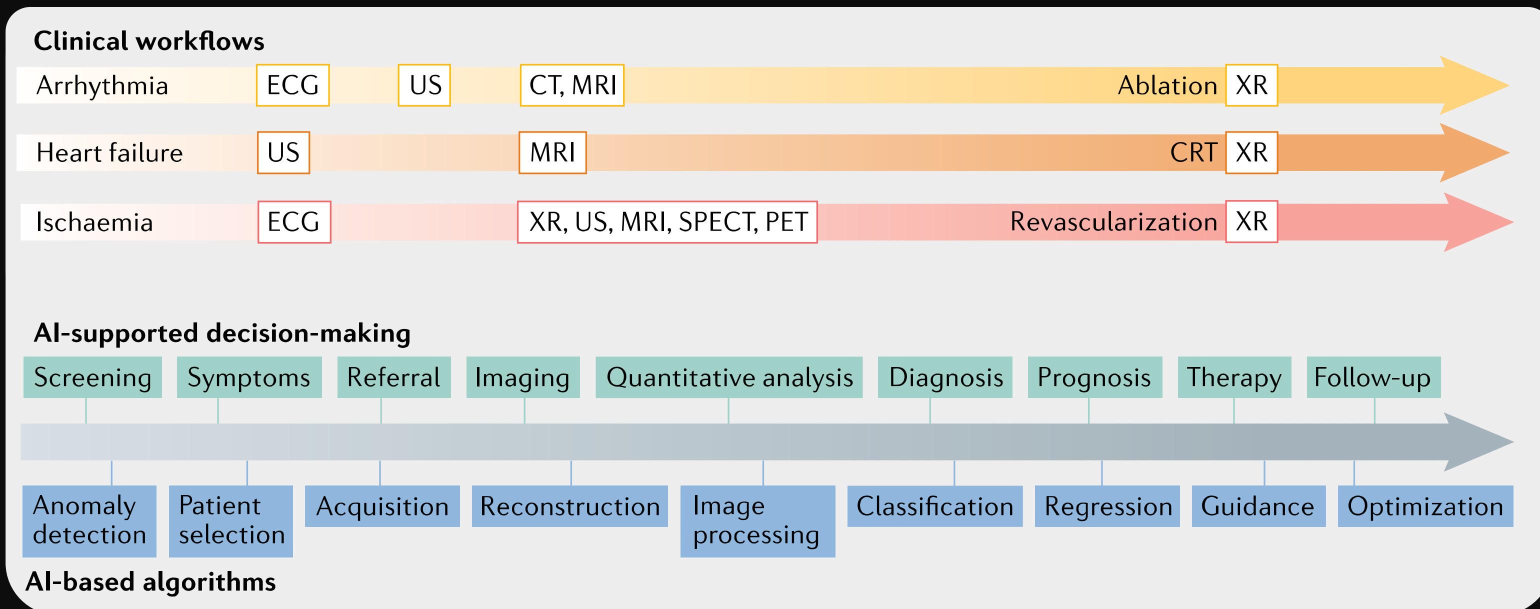

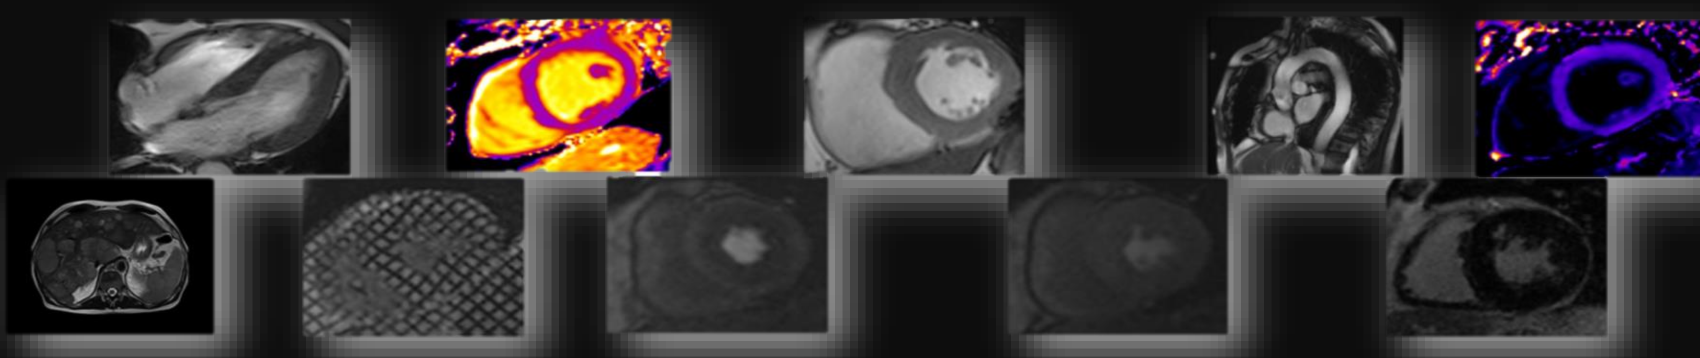

CMR

Deep learning

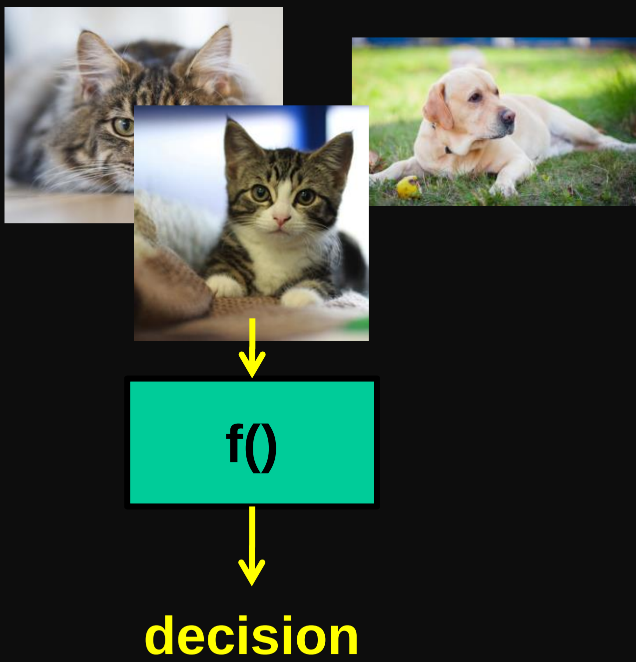

- Needs examples

(, cat)

(, dog)

Scanners, Sites and Populations

Origins of **variation**:

- Scan parameters
  (e.g. resolution)
- Scanner
  (e.g. 1.5T vs 3T, vendor)
- Sequence used
  (e.g. T2w, cine, T1)
- Population
  (e.g. ethnicity, gender)
- Site
  (e.g. protocols)

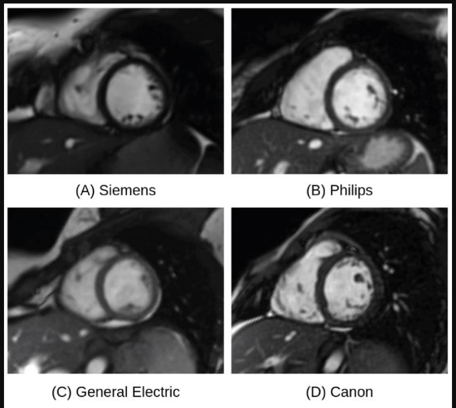

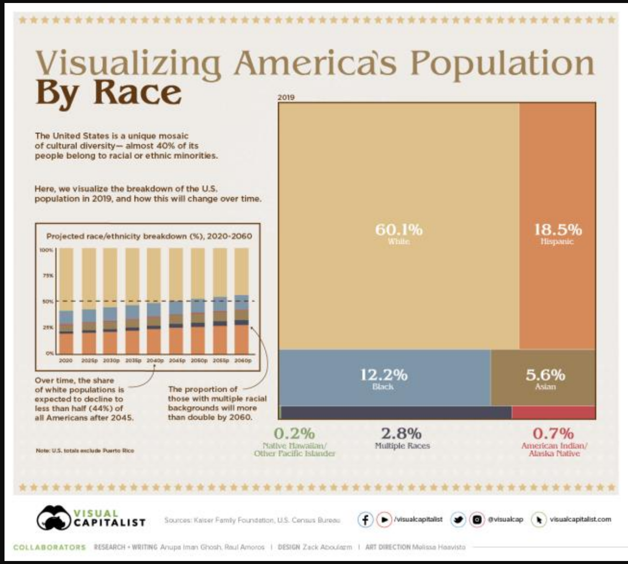

Figure from: Campello et al, Multi-Centre, Multi-Vendor and Multi-Disease Cardiac Segmentation : The MMs Challenge, IEEE TMI 2021

Scanners, Sites and Populations

Origins of **variation**:

- Scan parameters
  (e.g. resolution)
- Scanner
  (e.g. 1.5T vs 3T, vendor)
- Sequence used
  (e.g. T2w, cine, T1)
- Population
  (e.g. ethnicity, gender)
- Site
  (e.g. protocols)

Type of shift(s):

- Scan parameters, Scanner, Sequence used → Acquisition
- Population → Population, prevalence, manifestation
- Site → Annotation, population

Castro, Walker & Glocker. Causality matters in medical imaging. *Nat Commun* **11,** 3673 (2020).

The naïve solution

- ~~**Classical ML**~~
  ~~**more** data variation ➔ better performance~~

Vendor A

Vendor B

Vendor C

✓ Better AI generalise to
  - **Sensor** variation
  - **Population** variation

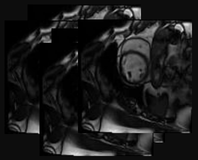

Vendor A

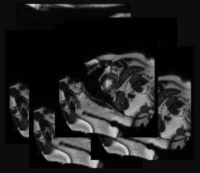

Vendor B

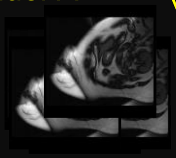

Vendor C

The impossible gold standard

Gold standard to isolating effects of e.g. scanner: the same patient "*travelling*" to be imaged by all vendors

Vendor A

We will have to repeat this for **any** effect we want to isolate.

Vendor B

Vendor C

But what if we can separate influence of imaging from anatomy/pathology?

How to achieve generalisation?

- During training:
  - Harmonisation
  - Domain generalization
    - Invariance
    - Adaptation
  - Data diversity by synthesis

- During application (deployment)
  - Test-time adaptation / training

LEARN REPRESENTATIONS

Disentangled representations

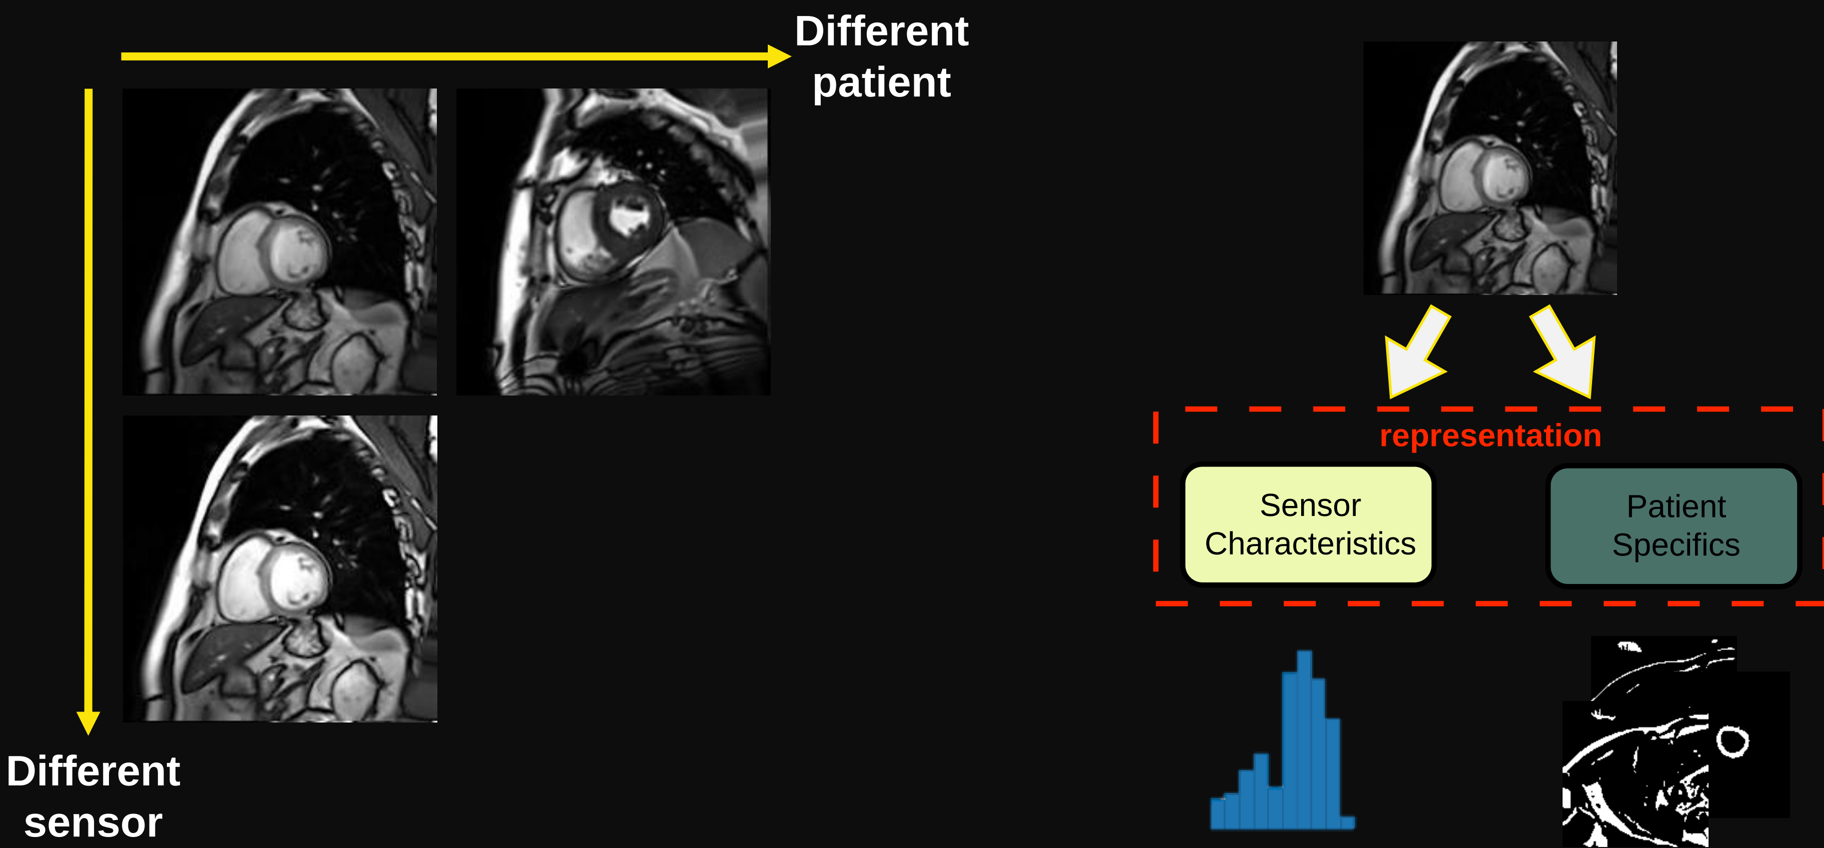

- ✓ Separate & organize factors of variation!
- ✓ Such separation works in images, timeseries, text, …
- ✓ A paradigm shift…

To learn more

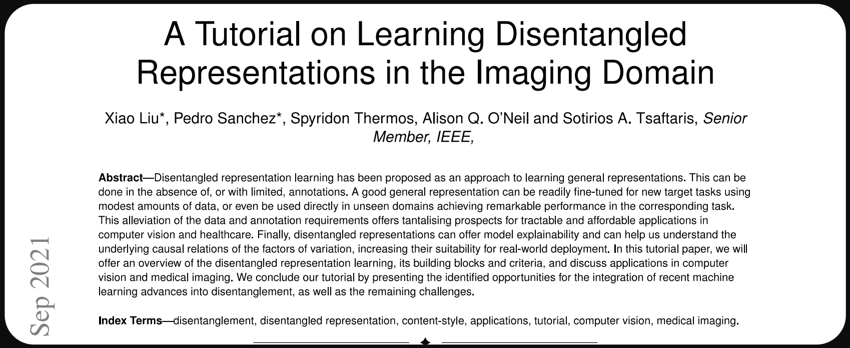
A Tutorial on Learning Disentangled Representations in the Imaging Domain

Xiao Liu*, Pedro Sanchez*, Spyridon Thermos, Alison Q. O'Neil and Sotirios A. Tsaftaris, *Senior Member, IEEE,*

Abstract—Disentangled representation learning has been proposed as an approach to learning general representations. This can be done in the absence of, or with limited, annotations. A good general representation can be readily fine-tuned for new target tasks using modest amounts of data, or even be used directly in unseen domains achieving remarkable performance in the corresponding task. This alleviation of the data and annotation requirements offers tantalising prospects for tractable and affordable applications in computer vision and healthcare. Finally, disentangled representations can offer model explainability and can help us understand the underlying causal relations of the factors of variation, increasing their suitability for real-world deployment. In this tutorial paper, we will offer an overview of the disentangled representation learning, its building blocks and criteria, and discuss applications in computer vision and medical imaging. We conclude our tutorial by presenting the identified opportunities for the integration of recent machine learning advances into disentanglement, as well as the remaining challenges.

Index Terms—disentanglement, disentangled representation, content-style, applications, tutorial, computer vision, medical imaging.

Sep 2021

2nd MICCAI 2021 tutorial

vios.science/tutorials/dream2021

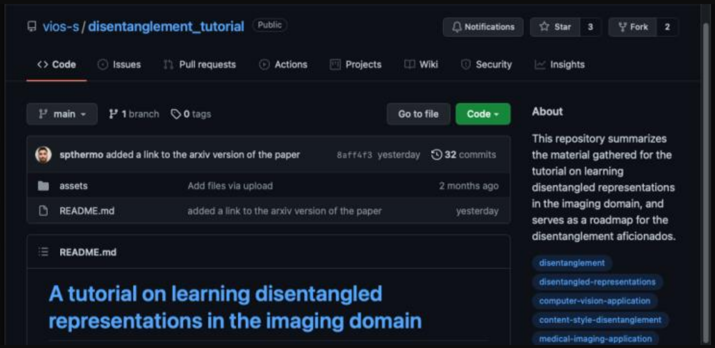

MAD MICCAI Workshop 2022

1st MICCAI 2022 workshop on disentanglement

https://mad.ikim.nrw

Disentangling the learning of anatomy

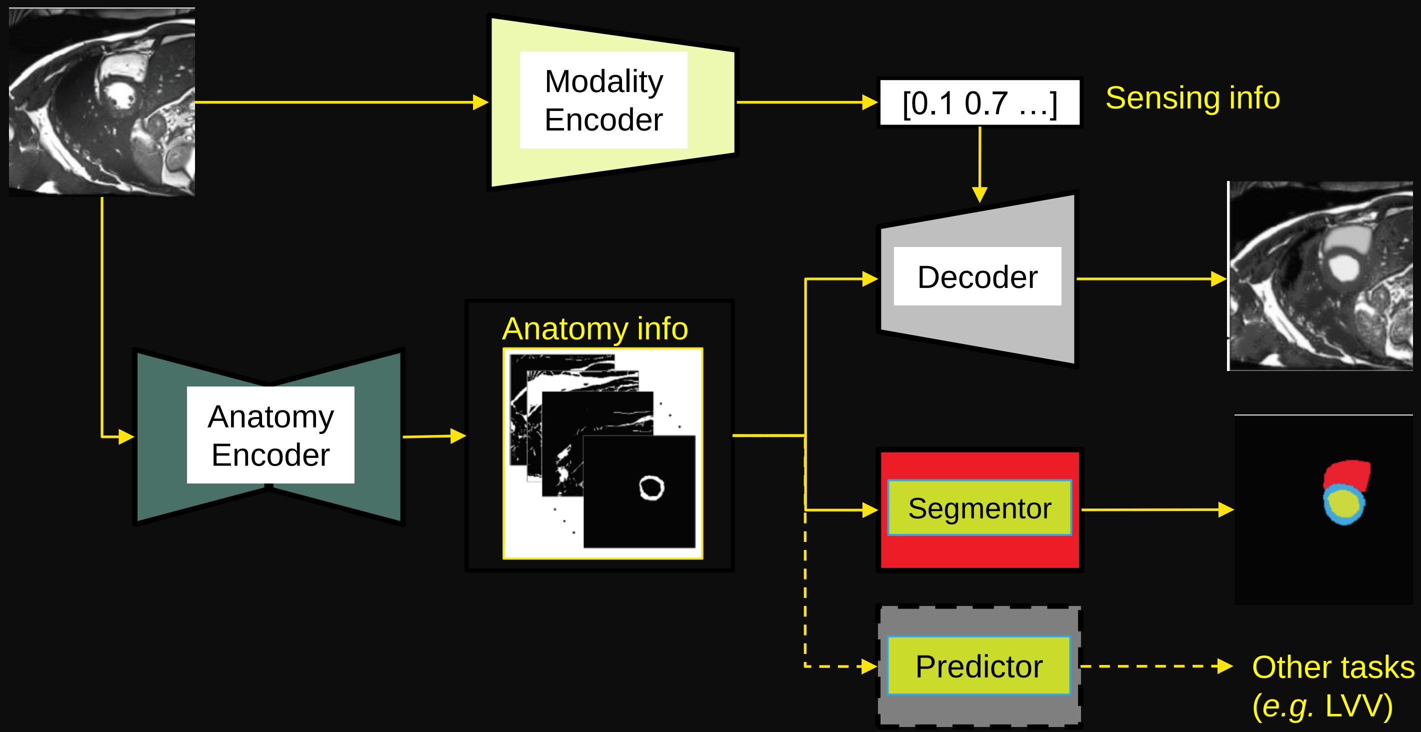

Chartsias, Joyce, Papanastasiou, Semple, Williams, Newby, Dharmakumar & Tsaftaris (2018 and 2019) Disentangled Representation Learning in Cardiac Image Analysis. Medical Image Analysis 2019 and MICCAI 2018
Code: https://github.com/agis85/anatomy_modality_decomposition

Use unlabeled images

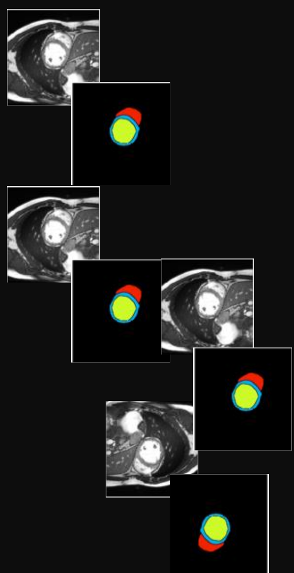

Annotated images
and masks

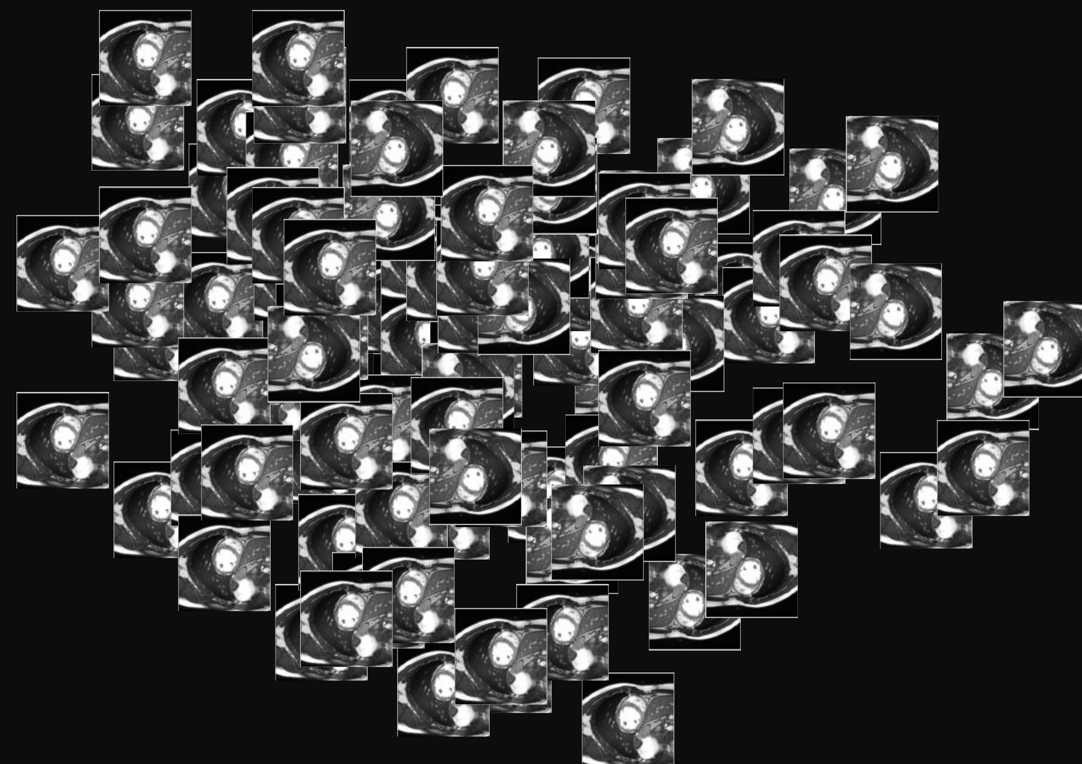

Other images

Only with 6% of annotated data: 75.6 (Dice)

[all annotated data 84 Dice]

What else we can do

- ✓ **MRI** and **CT**[1]
- ✓ **Leverage EHR**[1]
- ✓ Exploit **temporal** relationships[2]
- ✓ Expand to **pathology** [4,5]
- ✓ Use **multi-input**[6]

MR data	CT data	MR Dice	CT Dice
		78	80
		74	-

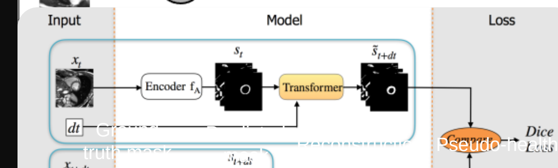

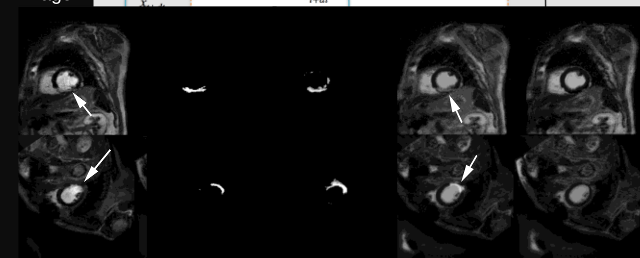

1. Chartsias, Joyce, Papanastasiou, Semple, Williams, Newby, Dharmakumar, Tsaftaris (2018 & 2019) Disentangled Representation Learning in Cardiac Image Analysis. MedIA and MICCAI
2. Valvano, Chartsias, Leo, Tsaftaris Temporal Consistency Objectives Regularize the Learning of Disentangled Representations DART @MICCAI 2019
3. Chartsias, Papanastasiou, Wang, Semple, Newby, Dharmakumar, Tsaftaris, Disentangle, align and fuse for multimodal and zero-shot image segmentation. STACOM @ MICCAI 2019, arXiv 1911.04417 , & IEEE TMI 2020
4. Jiang, Wang, Chartsias, Tsaftaris, Max-Fusion U-Net for Multi-Modal Pathology Segmentation with Attention and Dynamic Resampling," 2020 MyoPS Challenge, STACOM, a MICCAI 2020 workshop
5. Jiang, Chartsias, Zhang, Papanastasiou, Semple, Dweck, Semple, Dharmakumar, Tsaftaris, "Semi-supervised Pathology Segmentation with Disentangled Representations," DART, a MICCAI 2020 workshop
6. Chartsias, Papanastasiou, Wang, Semple, Newby, Dharmakumar, Tsaftaris, Disentangle, align and fuse for multimodal and zero-shot image segmentation. STACOM @ MICCAI 2019, arXiv 1911.04417, & IEEE TMI 2020

How we review CMR images in a protocol

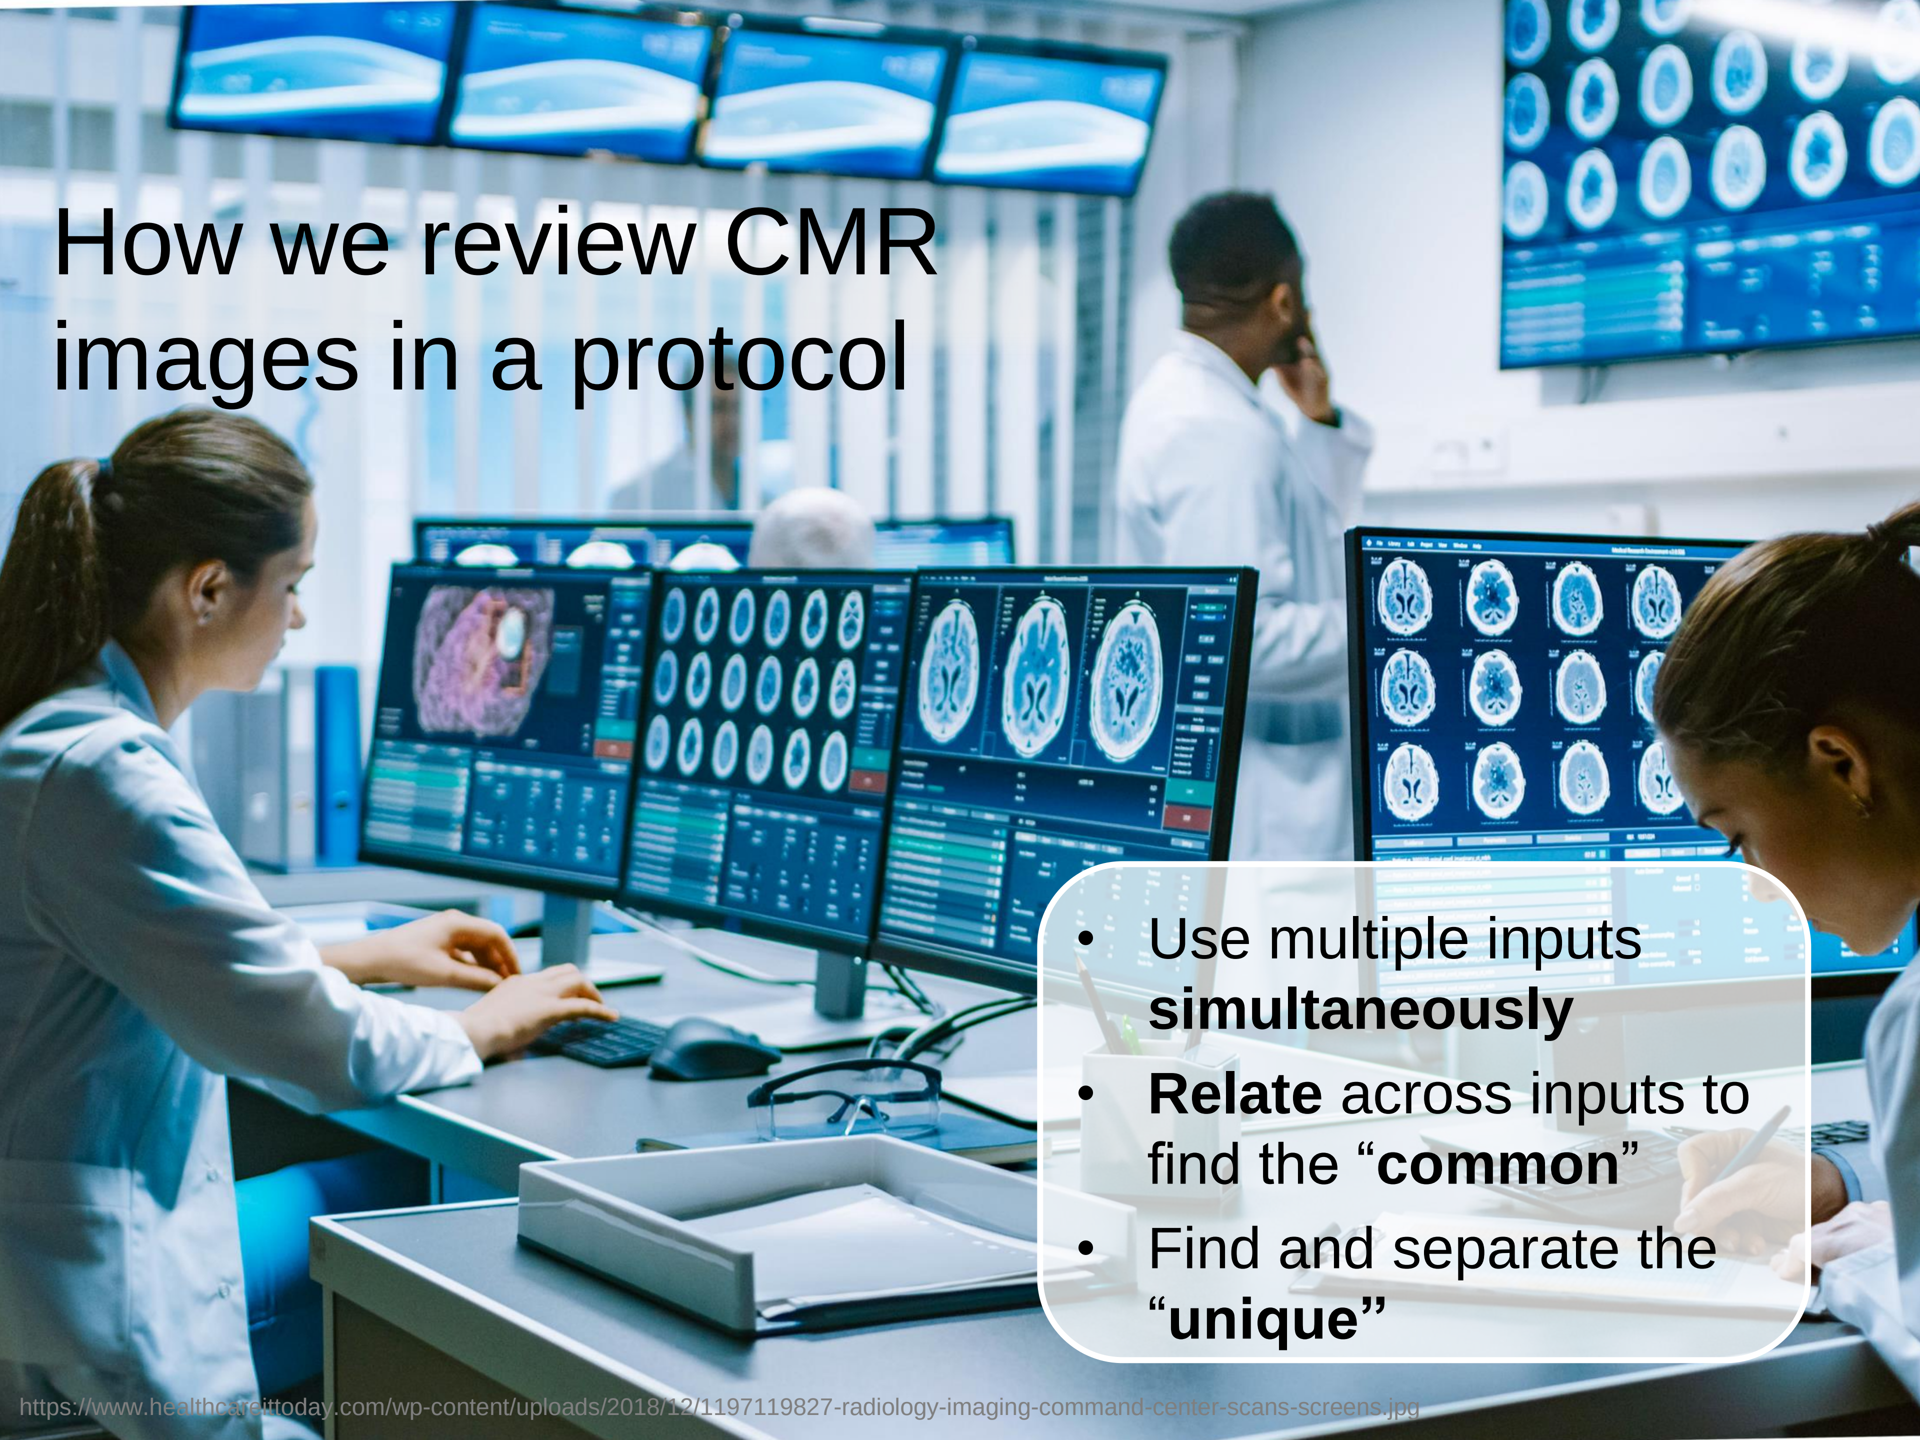

- Use multiple inputs **simultaneously**
- **Relate** across inputs to find the "**common**"
- Find and separate the "**unique**"

https://www.healthcareittoday.com/wp-content/uploads/2018/12/1197119827-radiology-imaging-command-center-scans-screens.jpg

DAFnet: Disentangle, Align and Fuse

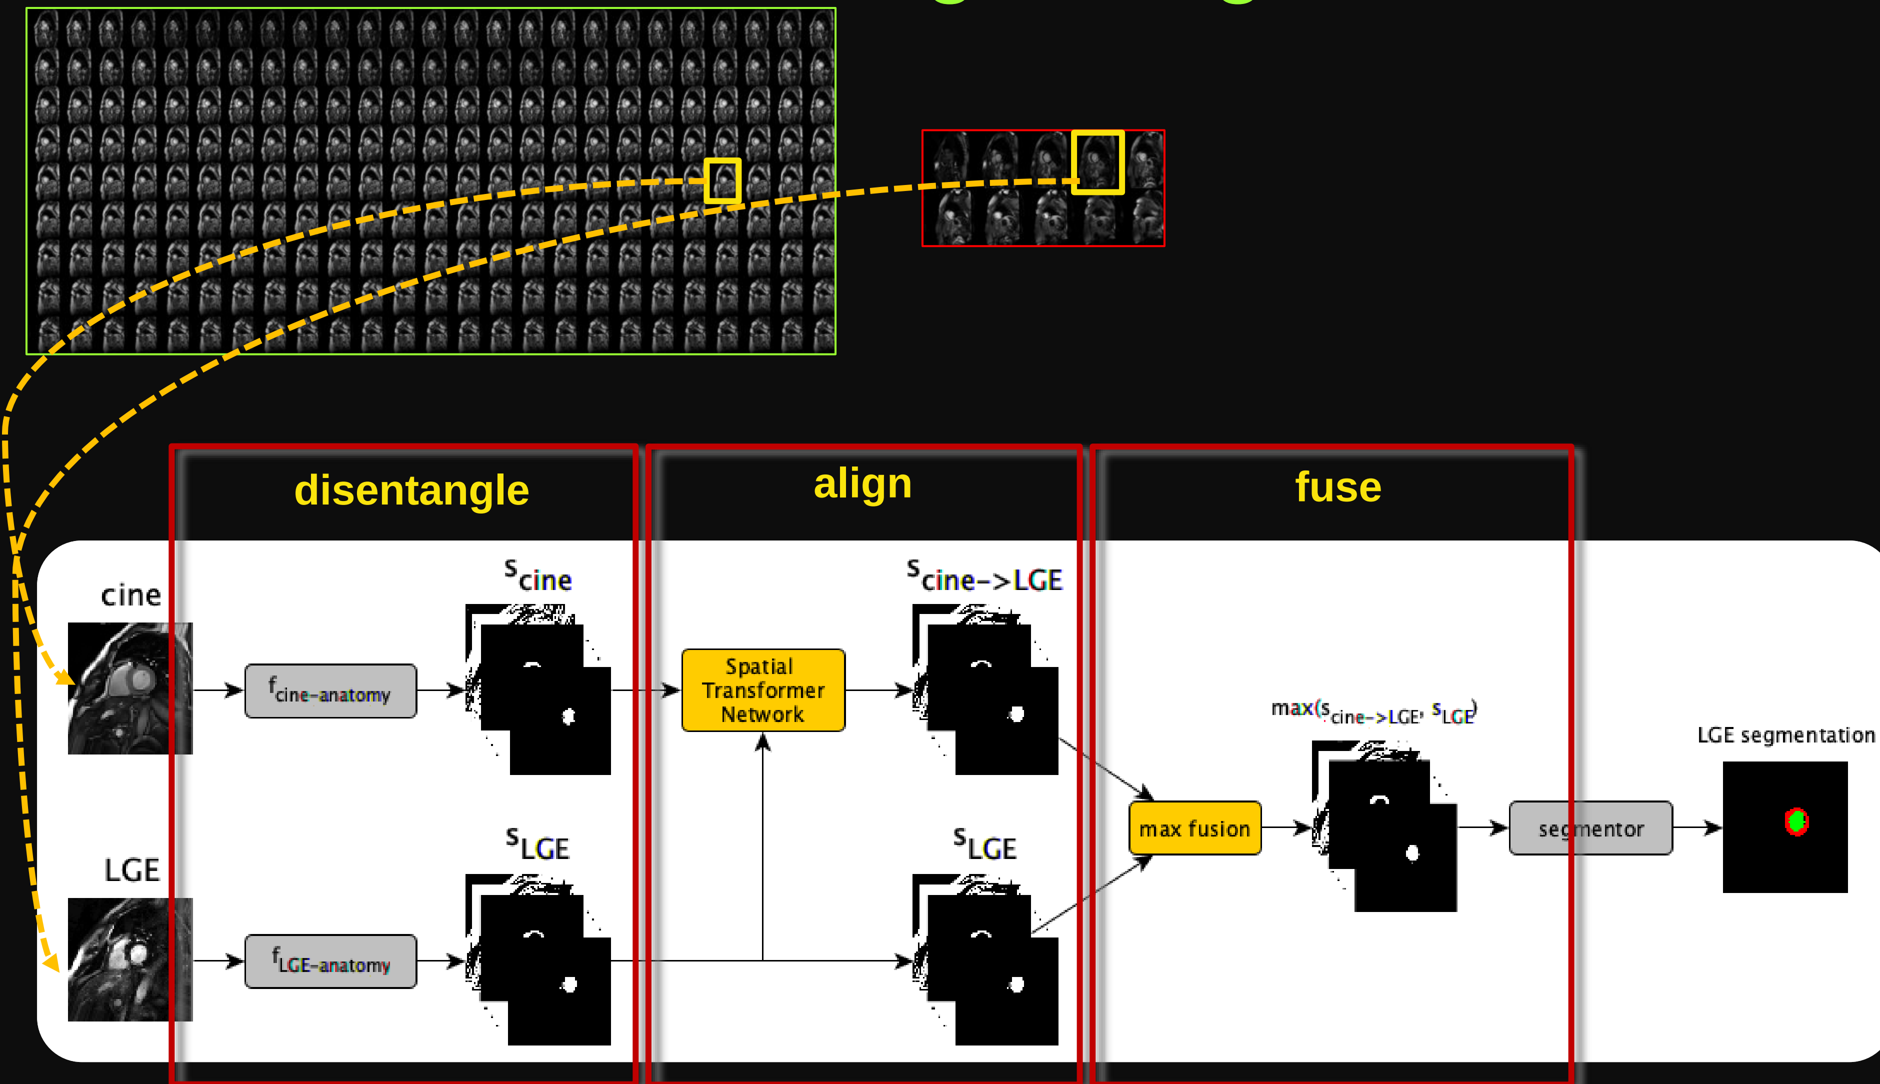

Chartsias, Papanastasiou, Wang, Semple, Newby, Dharmakumar, & Tsaftaris, Disentangle, align and fuse for multimodal and zero-shot image segmentation. STACOM @ MICCAI 2019, arXiv 1911.04417, & IEEE TMI 2020 code: https://github.com/vios-s/multimodal_segmentation

ACHIEVING GENERALISATION

Because of better representations

Remember this!

Gold standard to isolating effects of e.g. scanner: the same patient “*travelling*” to be imaged by all vendors

Vendor A

Now we can simulate this!

Vendor B

Vendor C

New combinations of factors

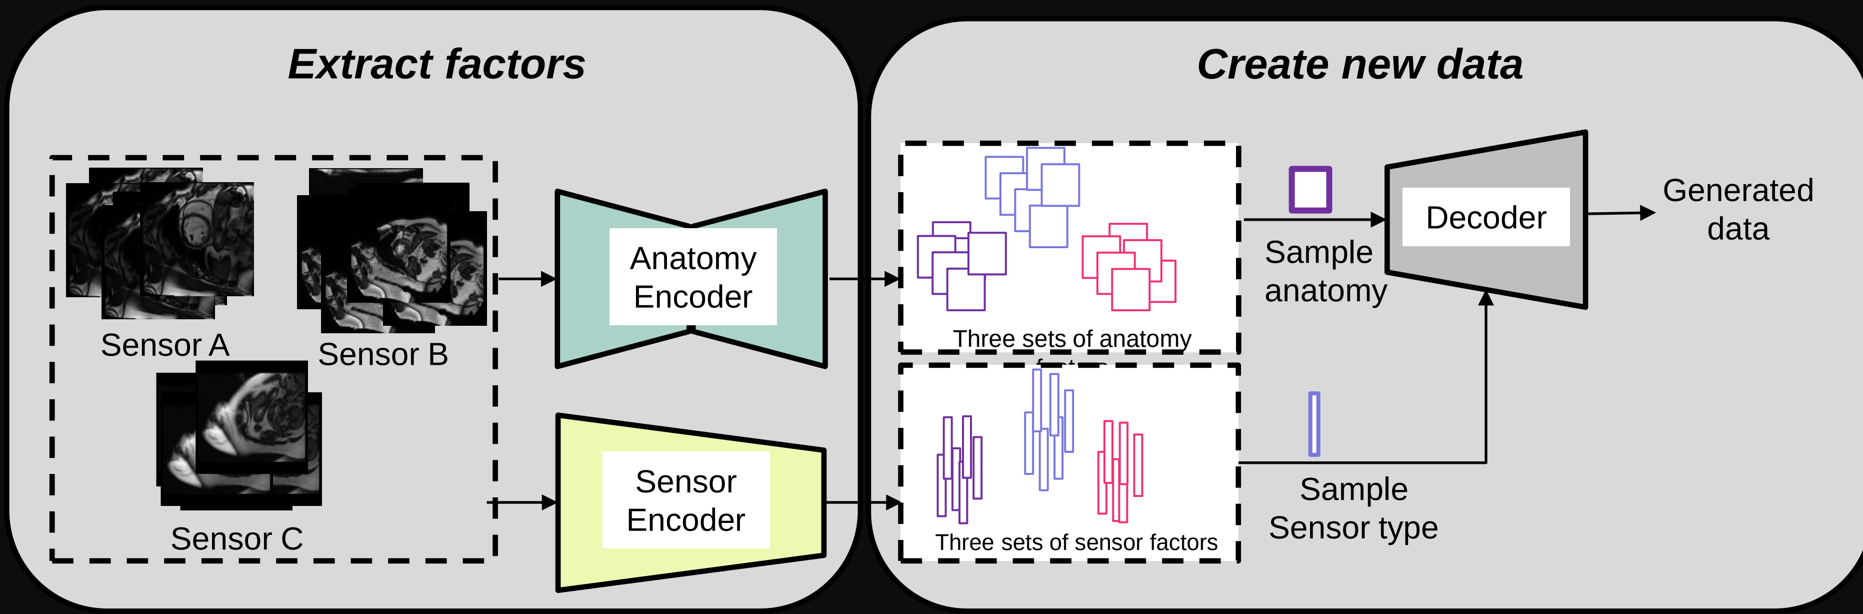

1. Take images from different sites/scanners
2. Separate anatomy from sensor variation
3. Create new virtual combinations

Liu, Thermos, Chartsias, O'Neil, & Tsaftaris (2020) Disentangled Representations for Domain-generalized Cardiac Segmentation. arXiv 2008.11514 M&Ms STACOM@MICCAI 2020 Challenge

We can improve the representations

- Penalise if separation of variation is not good
  - Meta-learning domain generalization

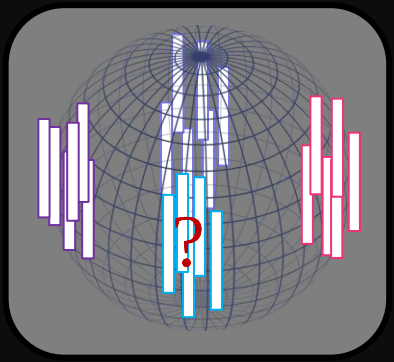

	Fully supervised (nnUnet)	Random mixing (semi-supervised + disentanglement)	Fully supervised meta-learning (w/o disentanglement)	**Semi-supervised meta-learning**
20%	64.85 (5.2)	76.73 (11)	73.5 (12)	**79.6 (11)**
100%	71.51 (5.4)	81.37 (11)	80.95 (13)	**82.25 (11)**

Liu, Thermos, O'Neil, & Tsaftaris (2021) Semi-supervised Meta-learning with Disentanglement for Domain-generalised Medical Image Segmentation. MICCAI 2021

We can add bespoke anatomical diversity

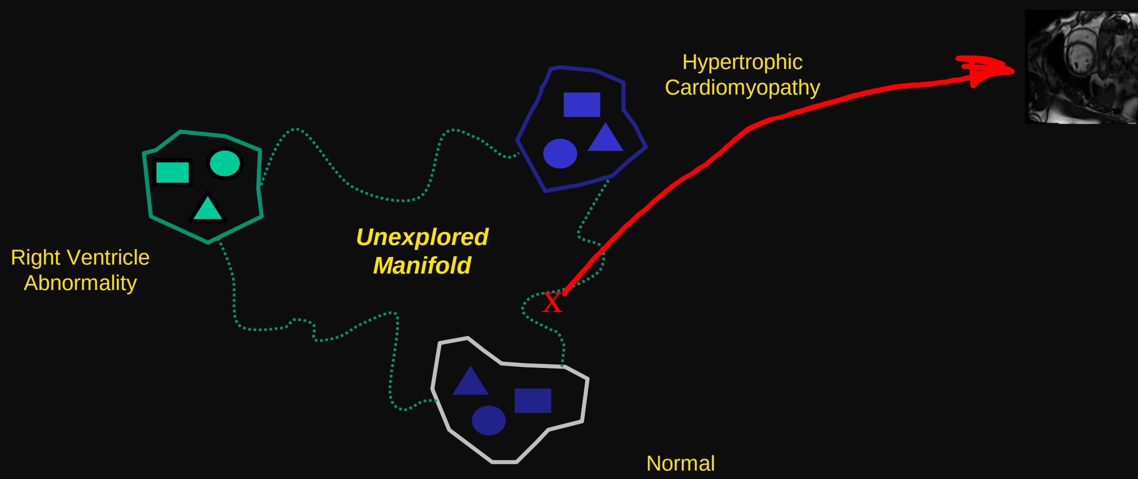

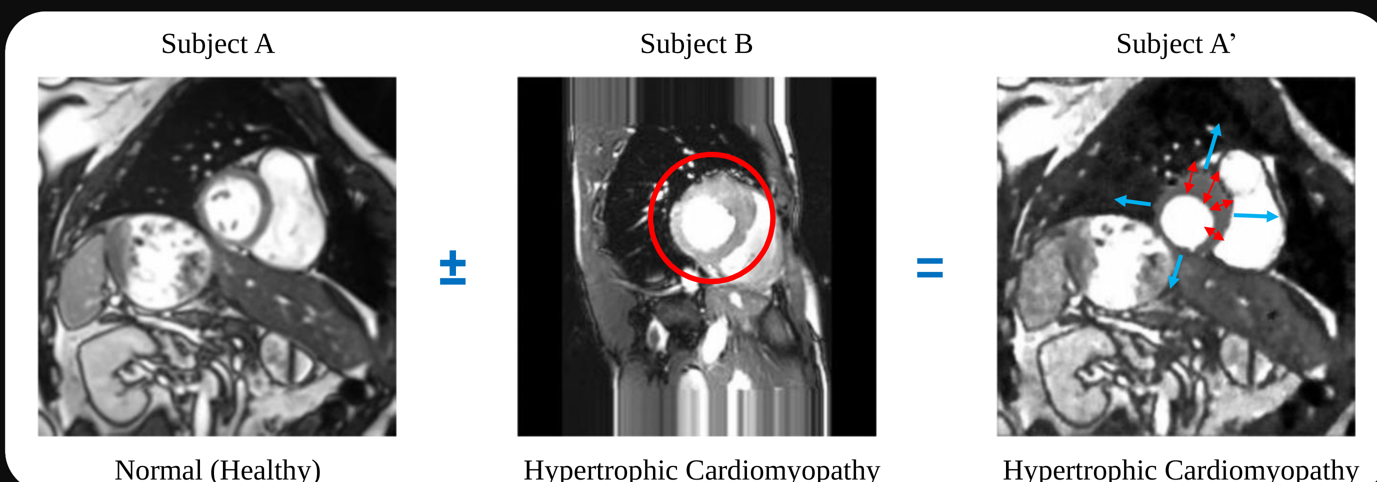

Thermos, Liu, O'Neil, Tsaftaris (2021) Controllable cardiac synthesis via disentangled anatomy arithmetic. MICCAI 2021
Patent pending.

To correct for prevalence shifts

- ✓ Augment training sets:
  - Rare pathologies
  - Under-represented sites/hospitals
  - Gender/race imbalance

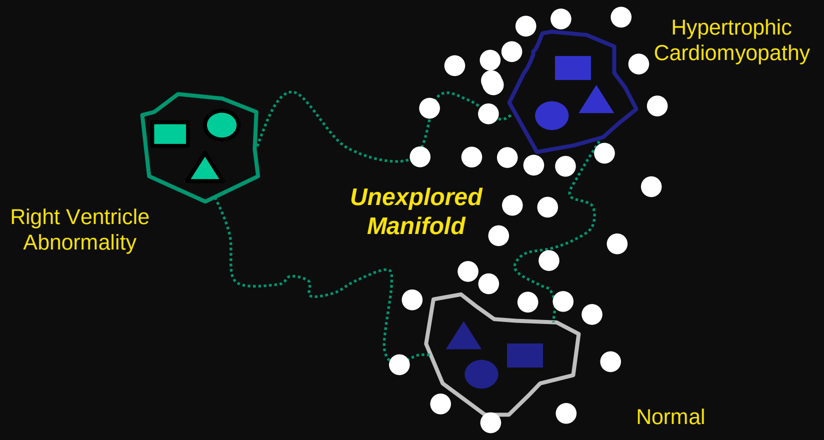

- ✓ Data can be generated as **desired**

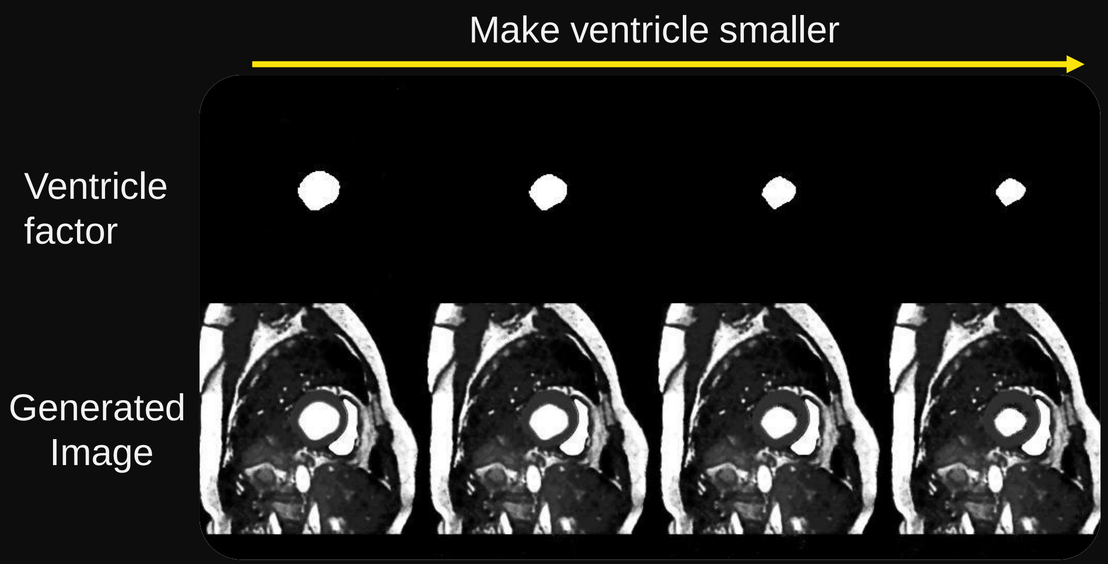

An example of editing

Correcting age distribution

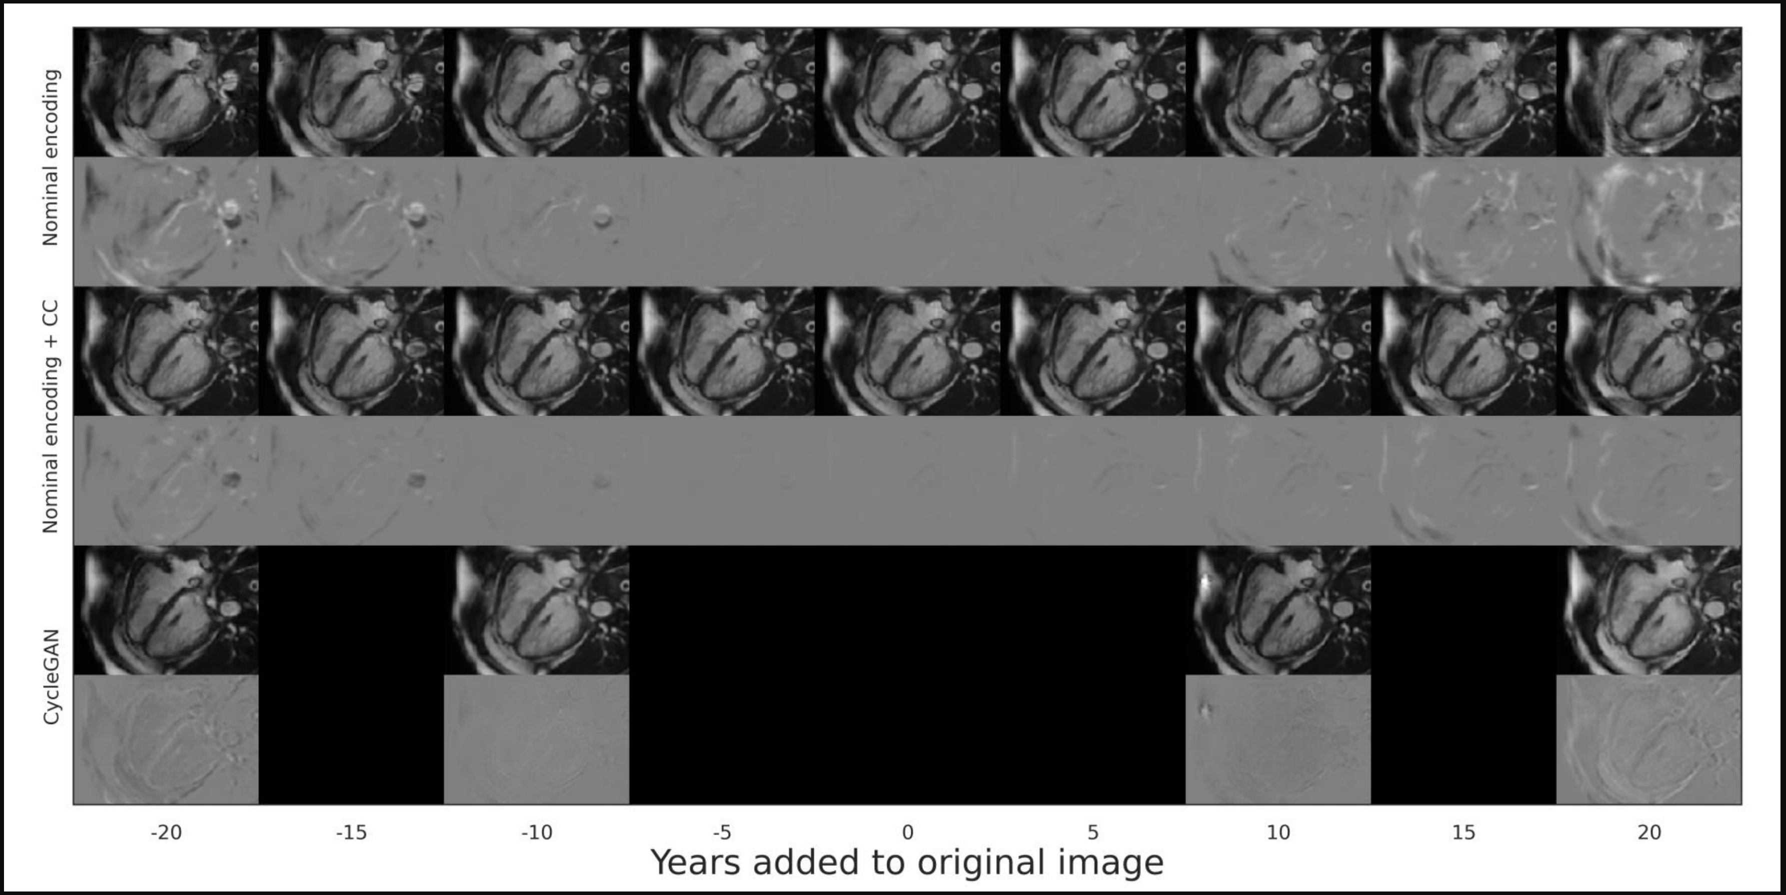

Campello et al, Cardiac aging synthesis from cross-sectional data with conditional generative adversarial networks, work in progress

Generalisation by test-time training

- TTT: adapt model on the fly without annotations
- Encoders, decoders, discriminators all capture priors
- At test time we **throw** components away
- Stop! **Reuse** them [1,2]

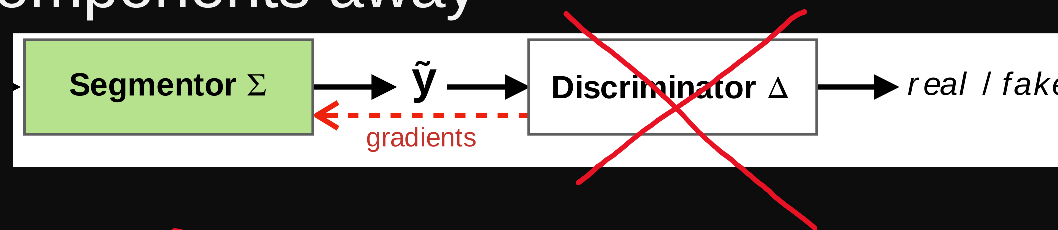

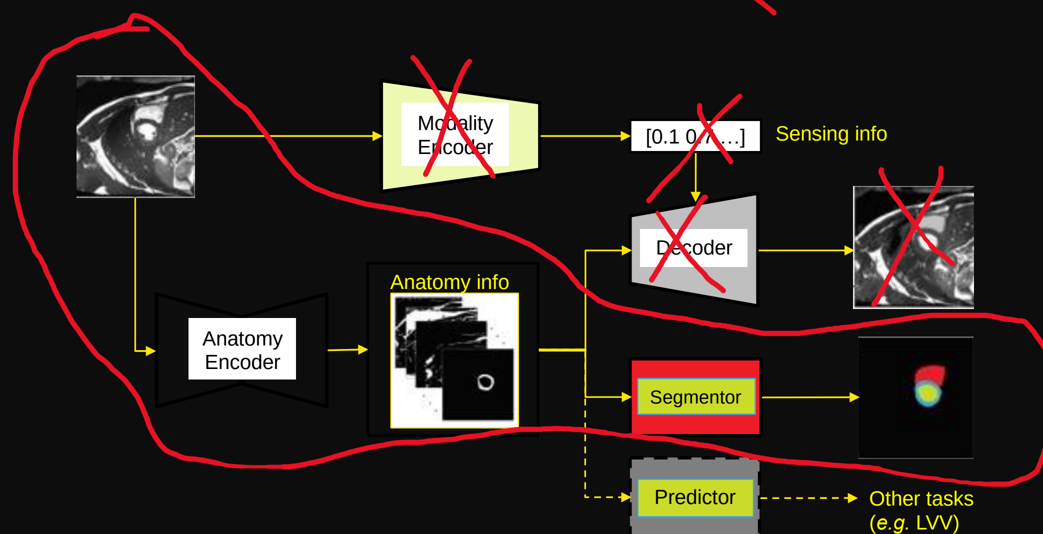

1. Valvano, Leo & Tsaftaris, Stop Throwing Away Discriminators! Re-using Adversaries for Test-Time Training, DART@ MICCAI 2021
2. Valvano, Leo, & Tsaftaris Re-using Adversarial Mask Discriminators for Test-time Training under Distribution Shift, arXiv, MELBA 2022

Gains

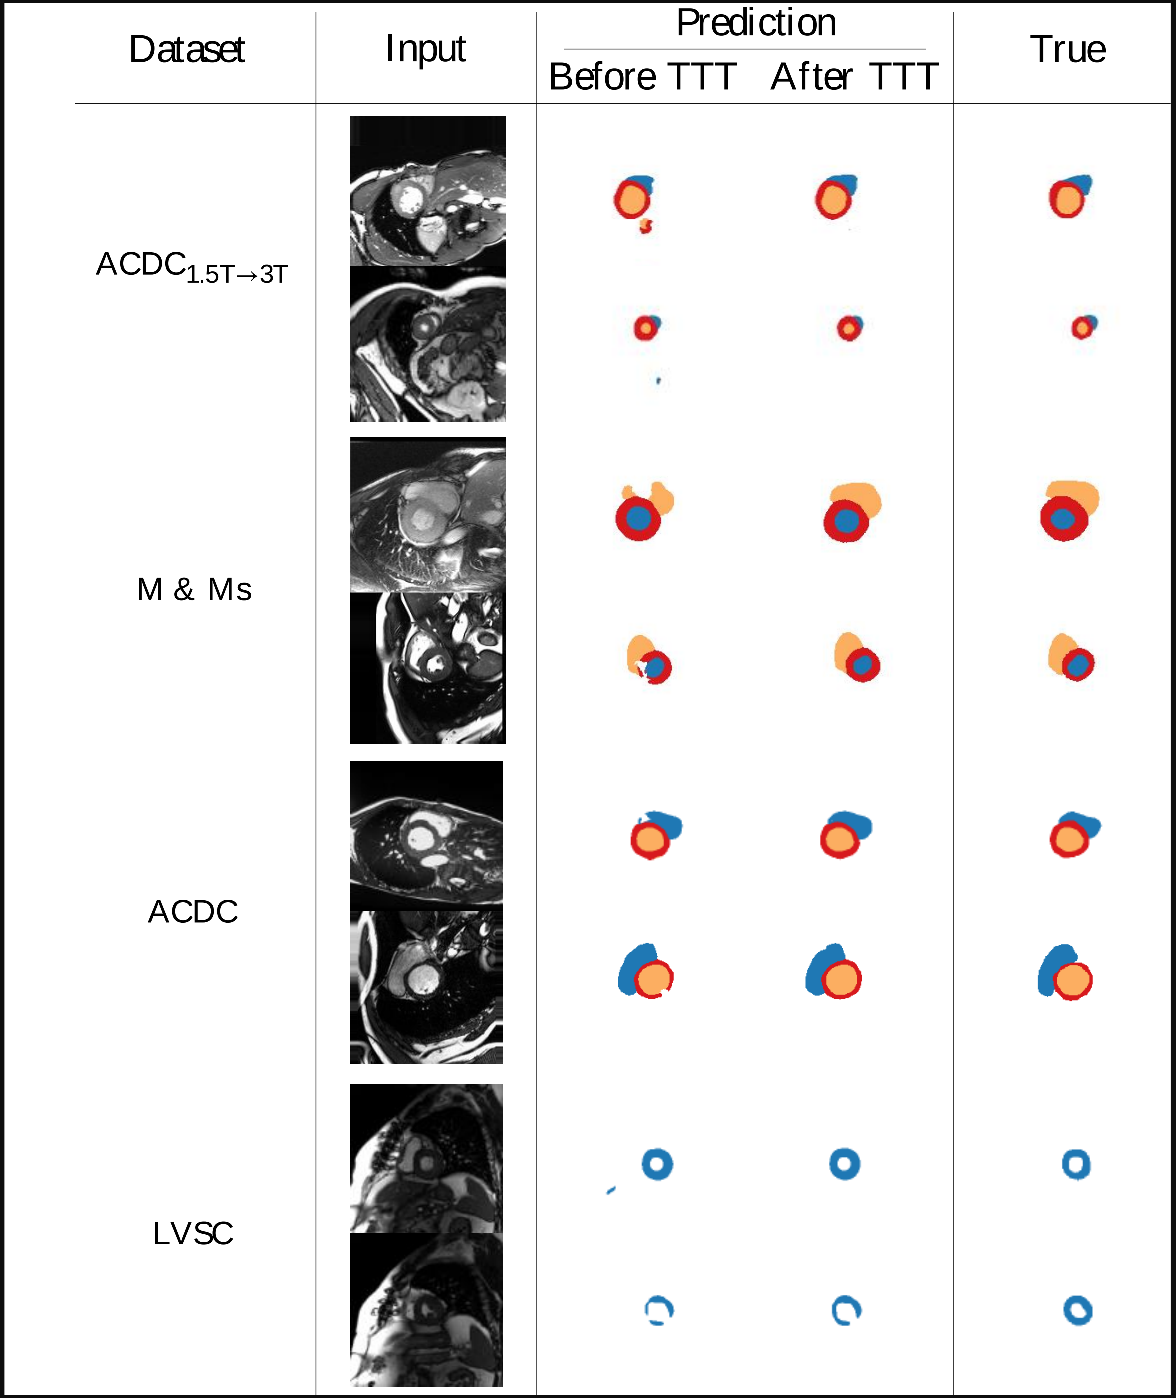
Dataset
Input
Prediction
Before TTT
After TTT
True
ACDC1.5T→3T
M & Ms
ACDC
LVSC

AI and Health

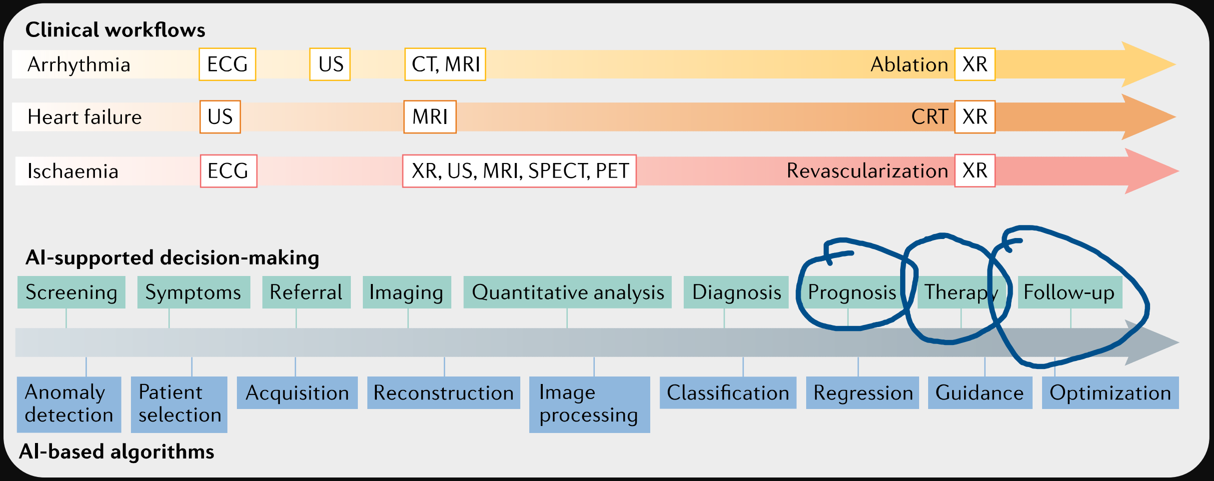

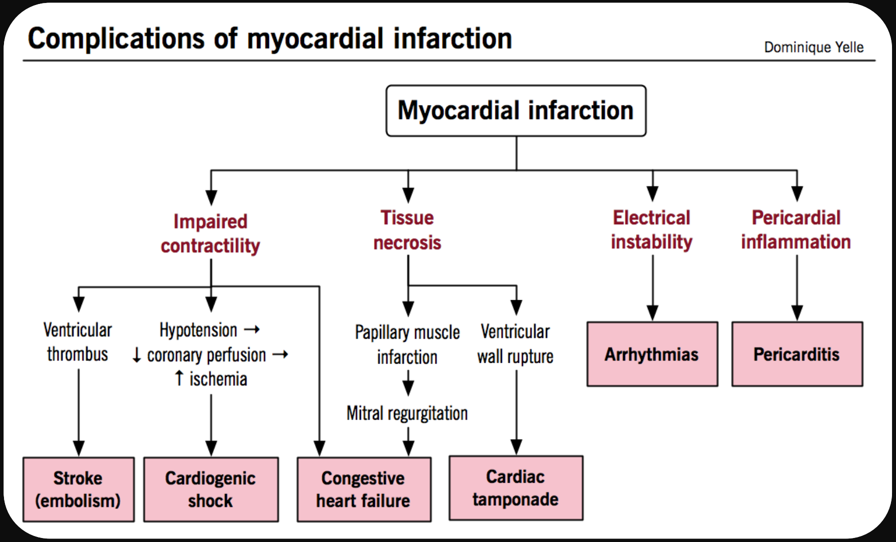

Sanchez & Tsaftaris, (2022) Diffusion Causal Models for Counterfactual Estimation. Conference on Causal Learning and Reasoning

Sermesant, Delingette, Cochet, Jais, Ayache, Applications of artificial intelligence in cardiovascular imaging. Nature Rev Cardiology 2021

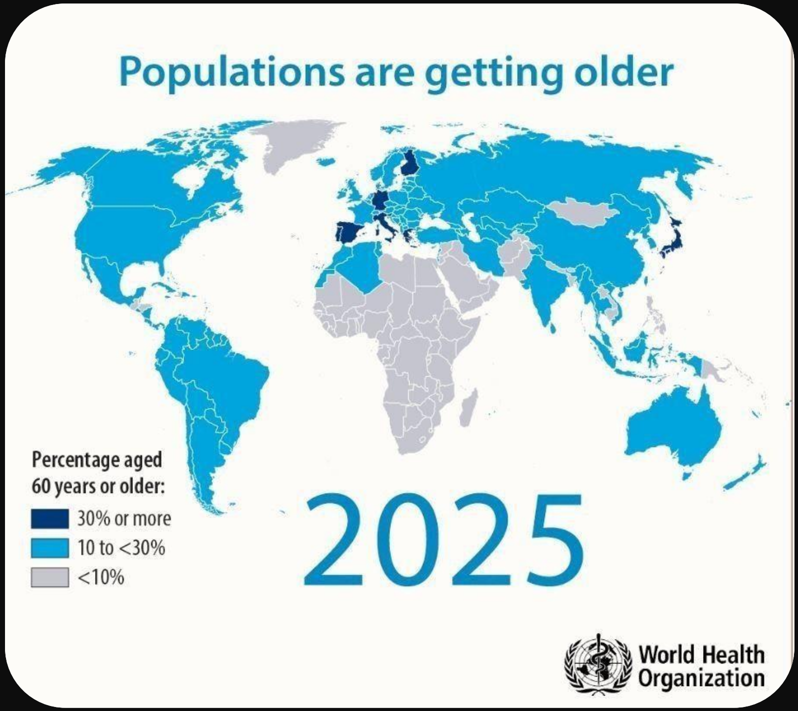
Populations are getting older
Percentage aged 60 years or older:
30% or more
10 to <30%
<10%
2025
World Health Organization

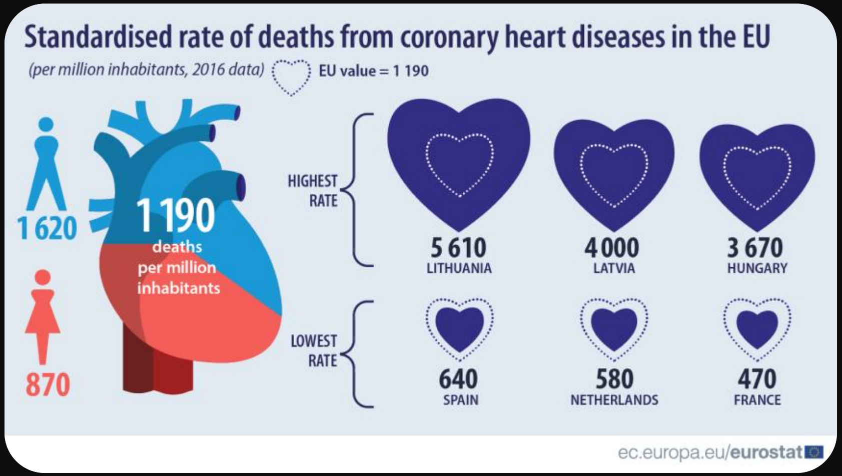
Standardised rate of deaths from coronary heart diseases in the EU
(per million inhabitants, 2016 data)
EU value = 1 190
1 620
870
1 190 deaths per million inhabitants
HIGHEST RATE
5 610 LITHUANIA
4 000 LATVIA
3 670 HUNGARY
LOWEST RATE
640 SPAIN
580 NETHERLANDS
470 FRANCE
ec.europa.eu/eurostat

Thanks to my team…

We have several PhD/RA openings if you want to join us!

vios.science

…my collaborators…

UK

- S. Weir
- A. Smout
- A. O'Neil
- A. Frangi
- D. Newby
- S. Semple
- G. Papanastasiou
- M. Williams

World

- R. Dharmakumar
- X. Papademetris
- N. Merchant
- H. Scharr
- P. Perata
- S. Bakas
- T. Arbel
- L. Maier-Hein

…funders and you!

Funding:

NIH 4R01HL091989
NIH 1R01HL136578

EP/N510129/1

EP/P022928/1

BB/N02334X/1
BB/P023487/1

Sotirios A. Tsaftaris, PhD

Email: S.Tsaftaris@ed.ac.uk

URL: http://vios.science