Learning biomarkers from heterogeneous medical imaging data

Centralised and federated learning approaches

Jonas Richiardi

@TranslationalML W3 https://unil.ch/tml

CHUV

UNIL | Université de Lausanne

TRANSLATIONAL MACHINE LEARNING LAB

ML for image-based biomarkers

Image
acquisition

ML for image-based biomarkers

Image
acquisition

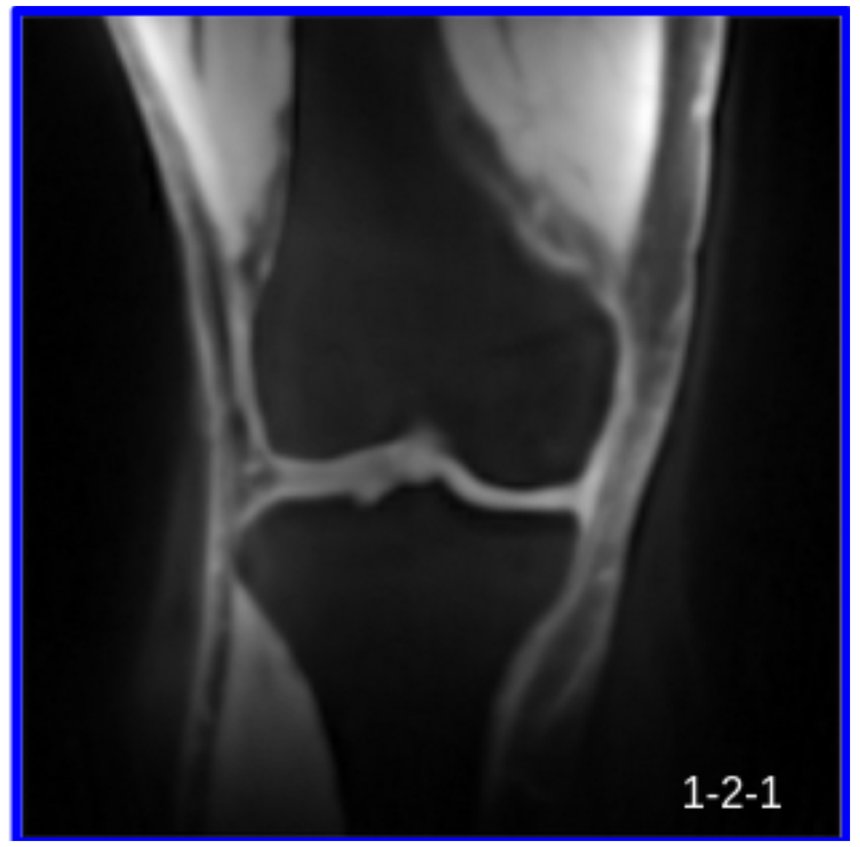

[Sieber et al. ISMRM 2022]

ML for image-based biomarkers

Image acquisition

Differential diagnosis

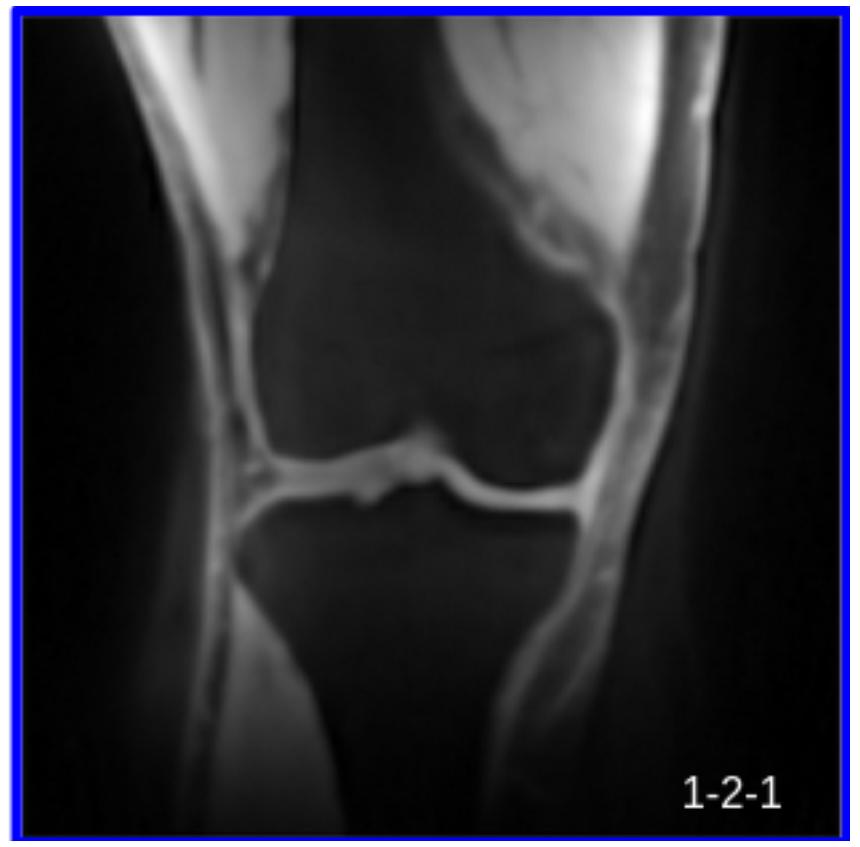

[Sieber et al. ISMRM 2022]

ML for image-based biomarkers

Image acquisition

Differential diagnosis

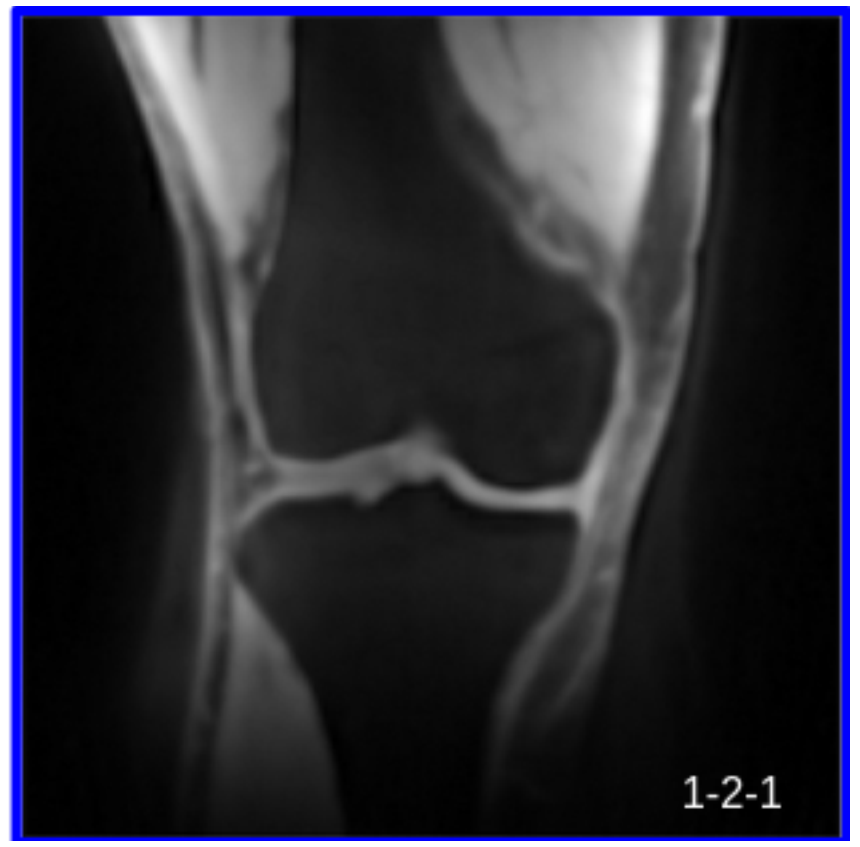

[Sieber et al. ISMRM 2022]

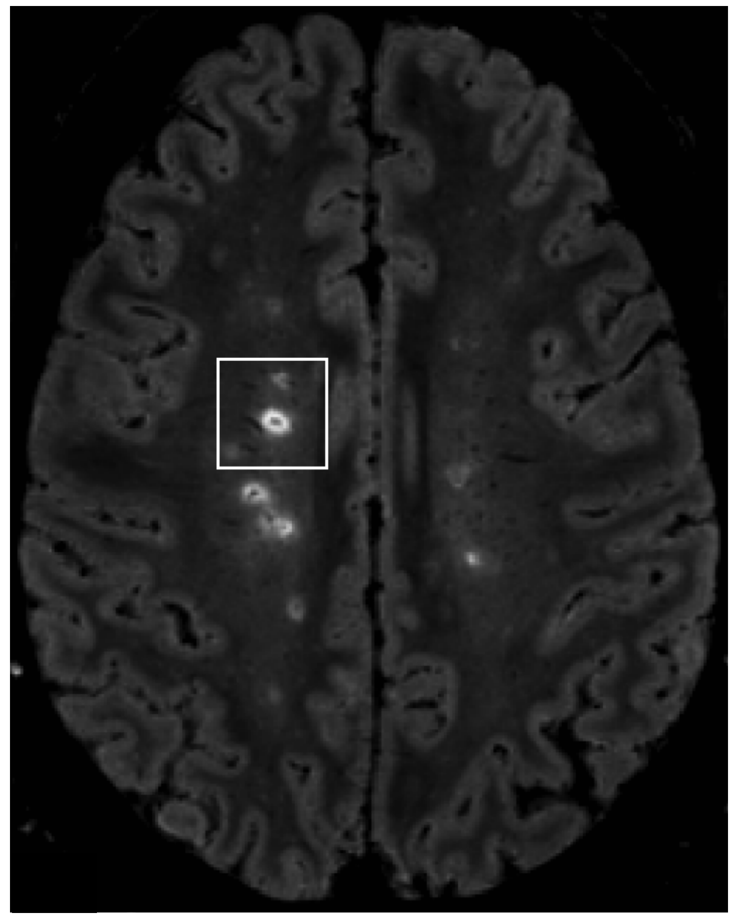

[Maggi, Fartaria, et al. 2020]

ML for image-based biomarkers

Image acquisition

Differential diagnosis

Intervention eligibility

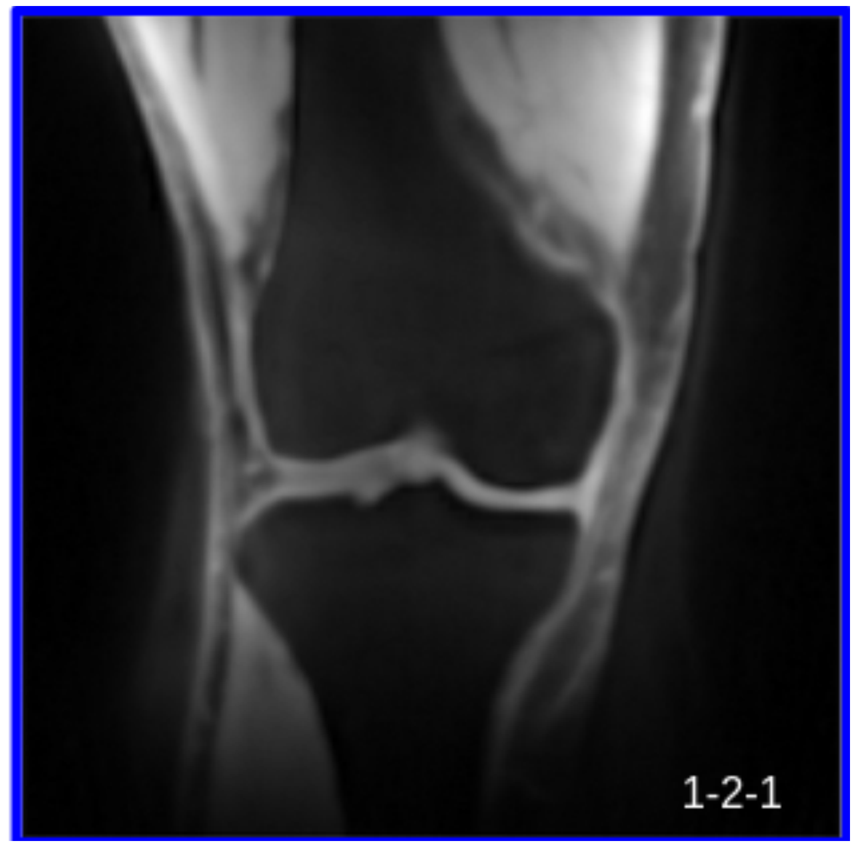

[Sieber et al. ISMRM 2022]

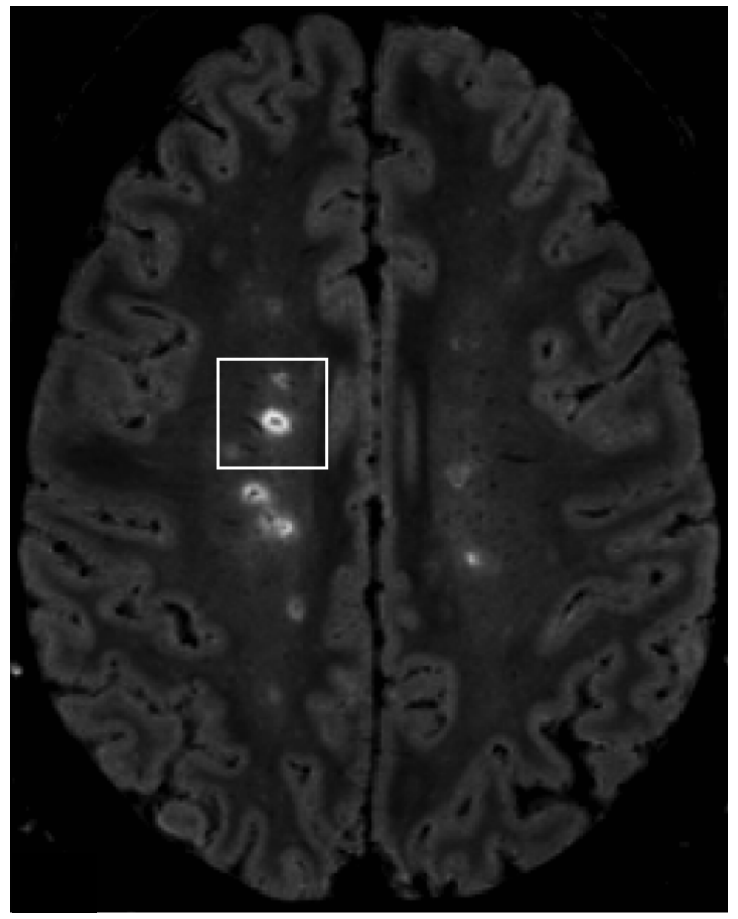

[Maggi, Fartaria, et al. 2020]

ML for image-based biomarkers

Image acquisition

Differential diagnosis

Intervention eligibility

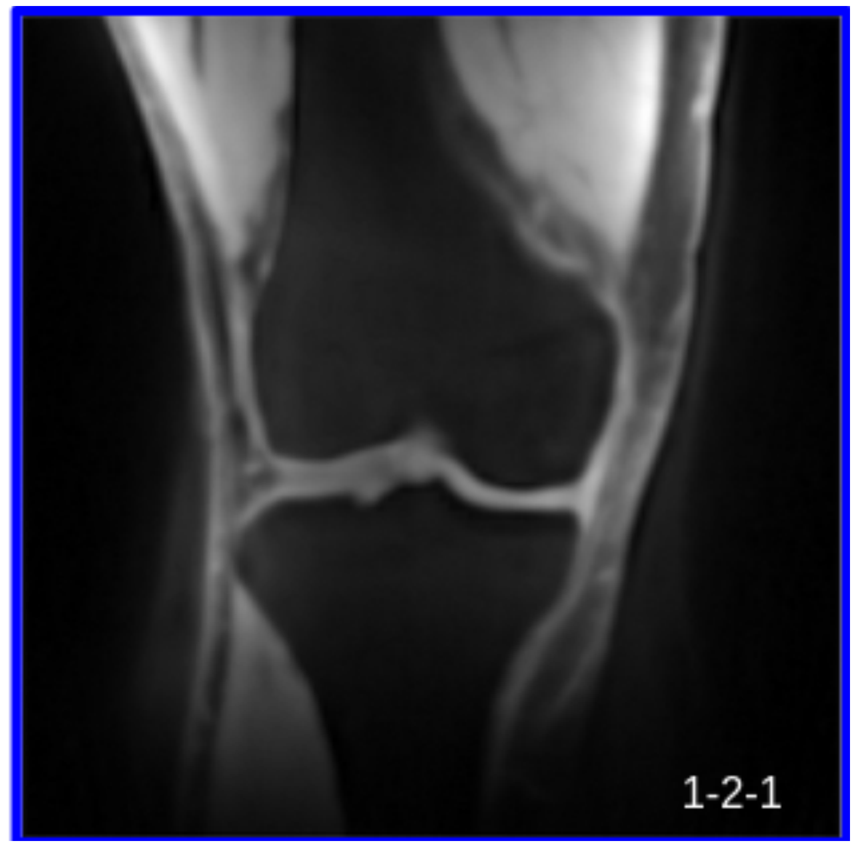

[Sieber et al. ISMRM 2022]

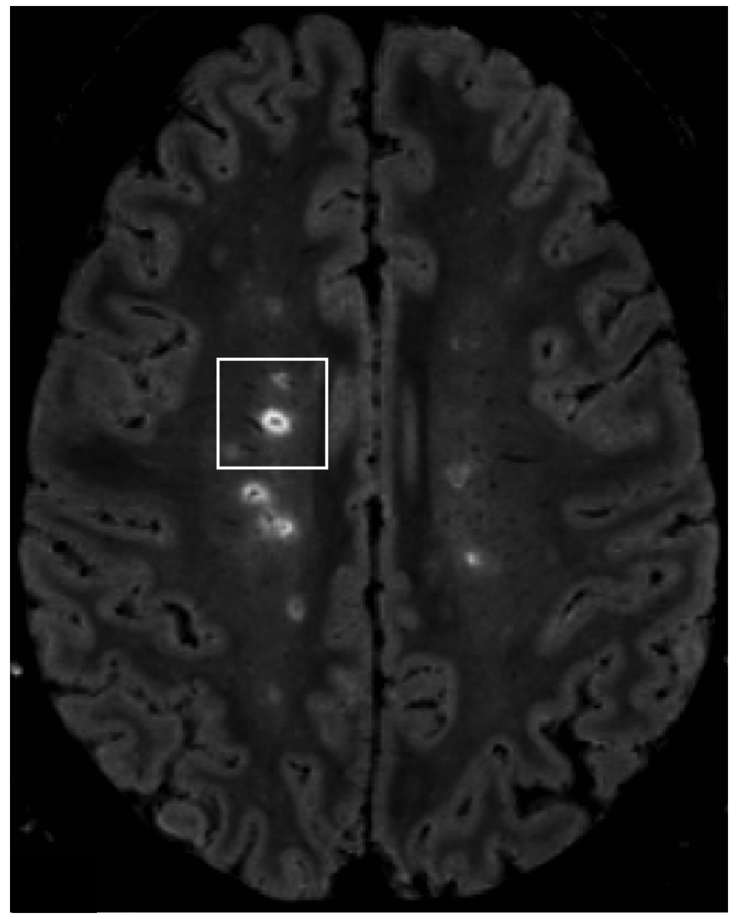

[Maggi, Fartaria, et al. 2020]

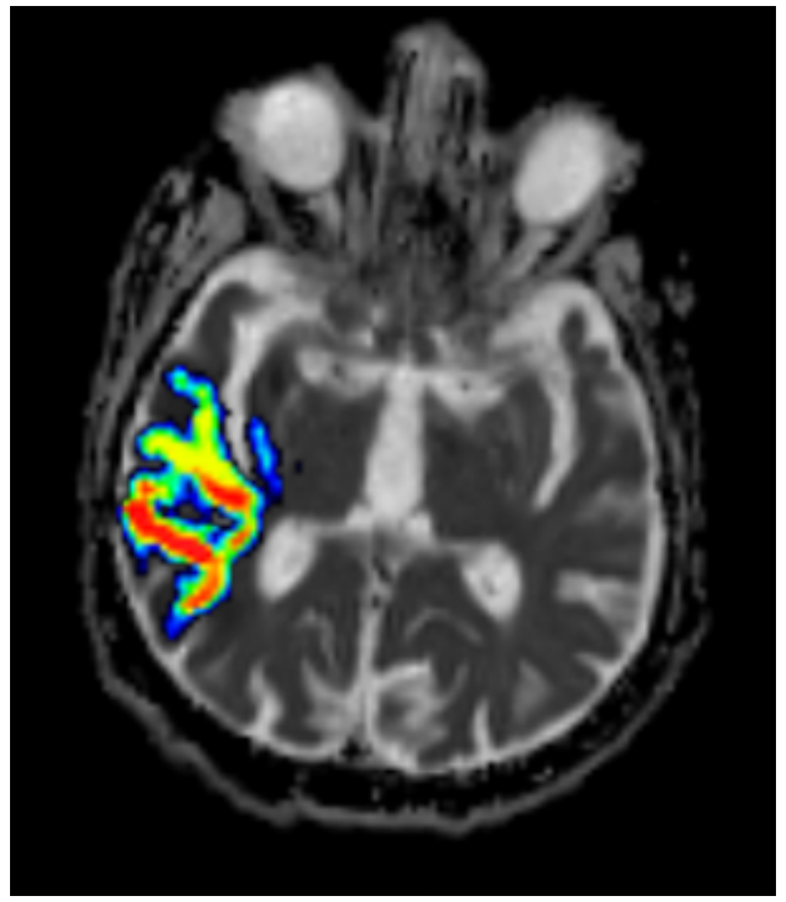

[Rafael-Patiño et al ISMRM 2022]

ML for image-based biomarkers

Image acquisition

Differential diagnosis

Intervention eligibility

treatment monitoring

...

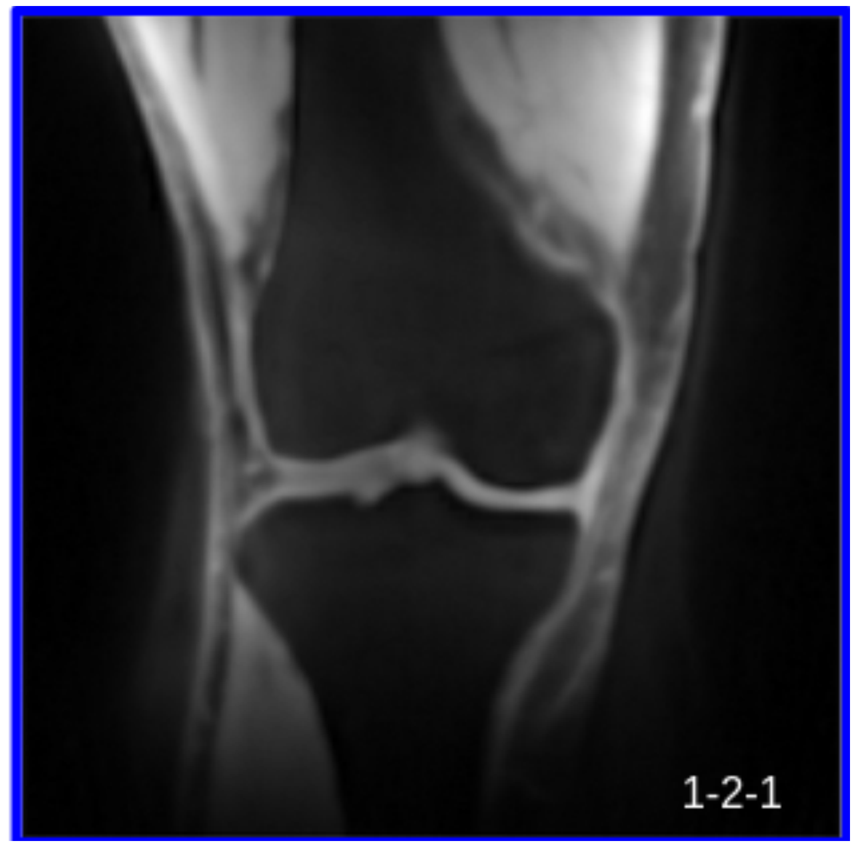

[Sieber et al. ISMRM 2022]

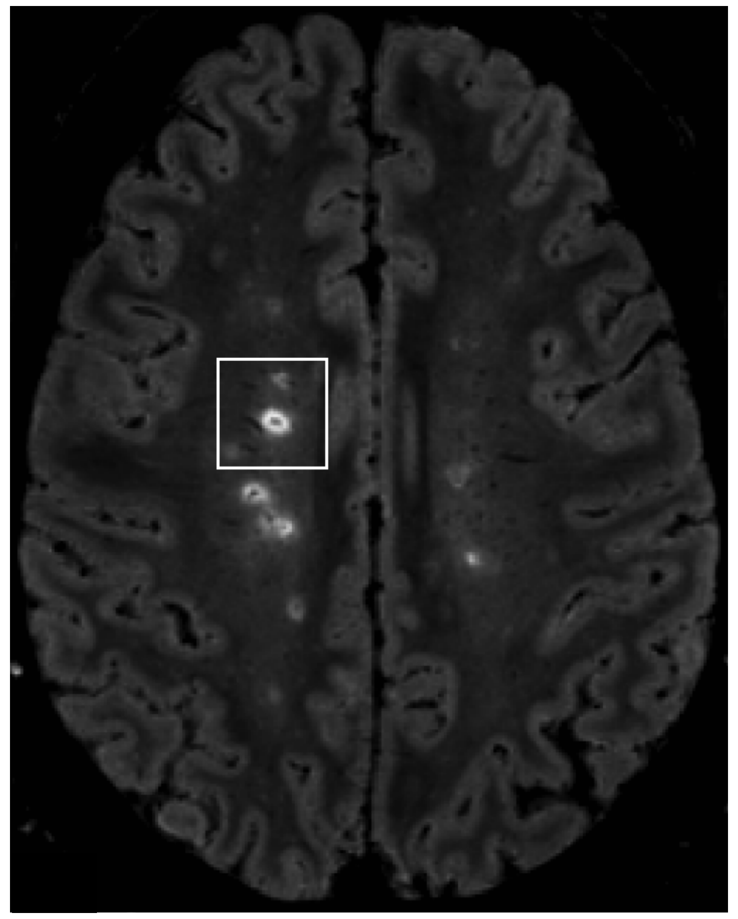

[Maggi, Fartaria, et al. 2020]

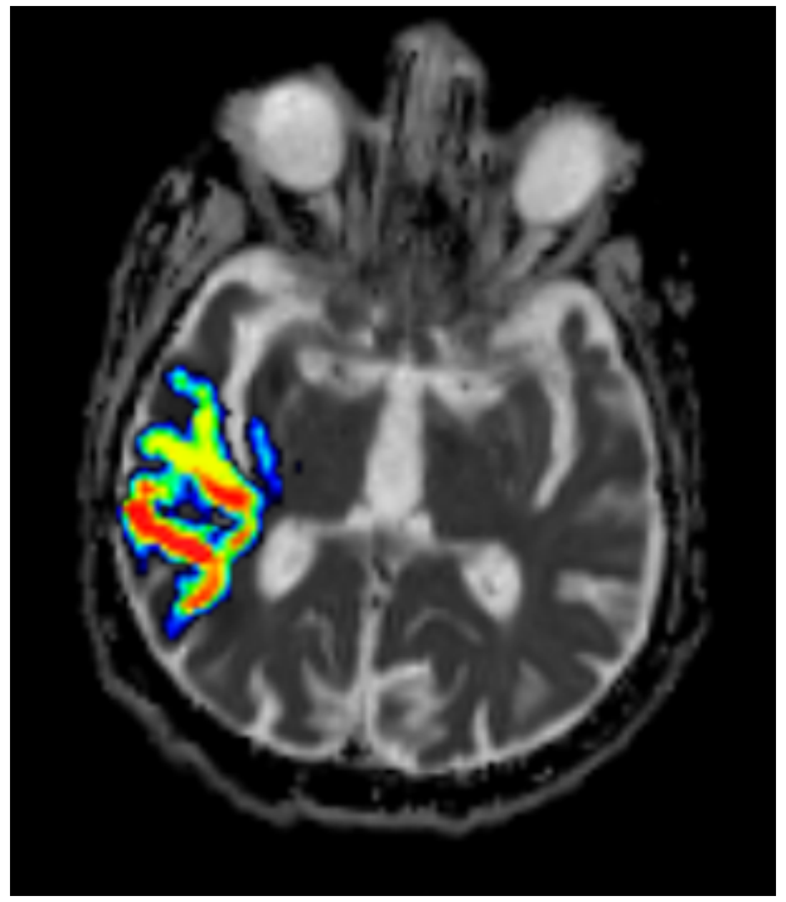

[Rafael-Patiño et al ISMRM 2022]

ML for image-based biomarkers

Image acquisition

Differential diagnosis

Intervention eligibility

treatment monitoring

...

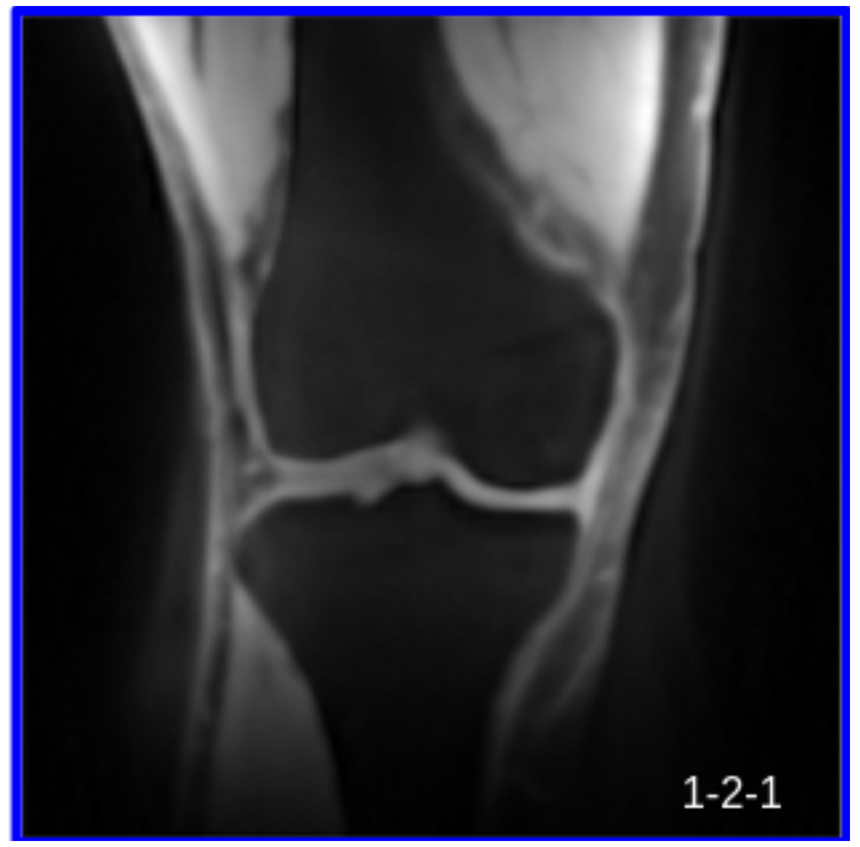

[Sieber et al. ISMRM 2022]

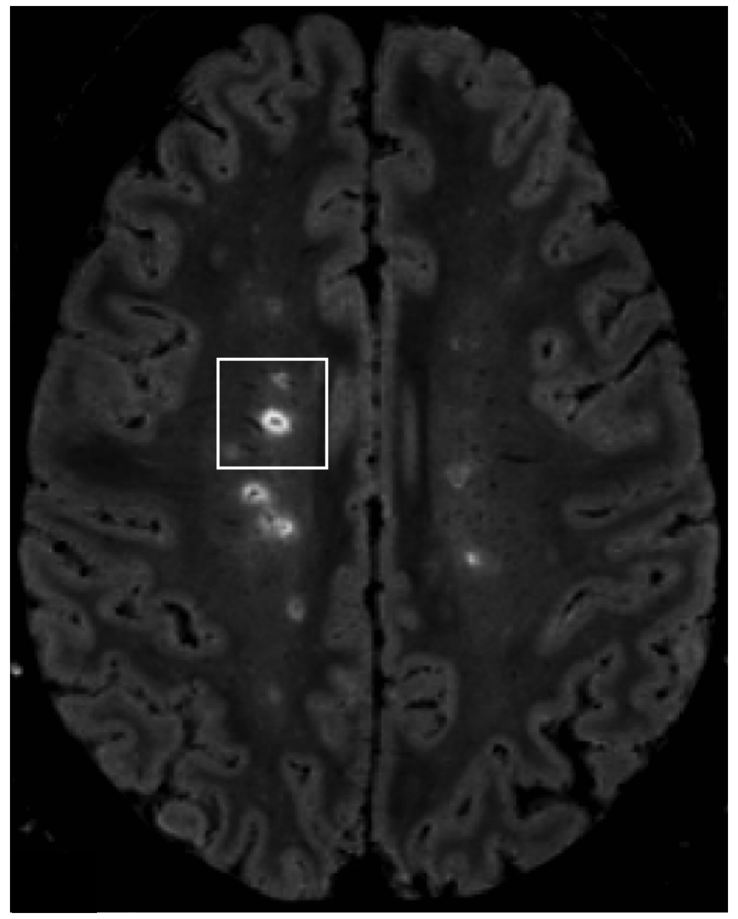

[Maggi, Fartaria, et al. 2020]

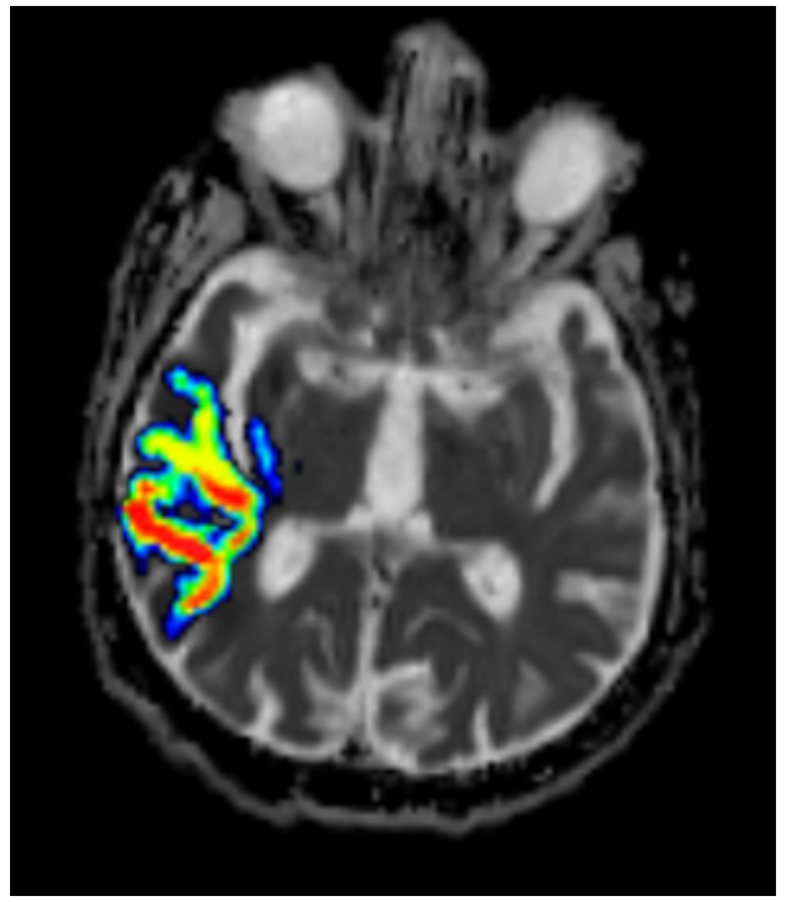

[Rafael-Patiño et al ISMRM 2022]

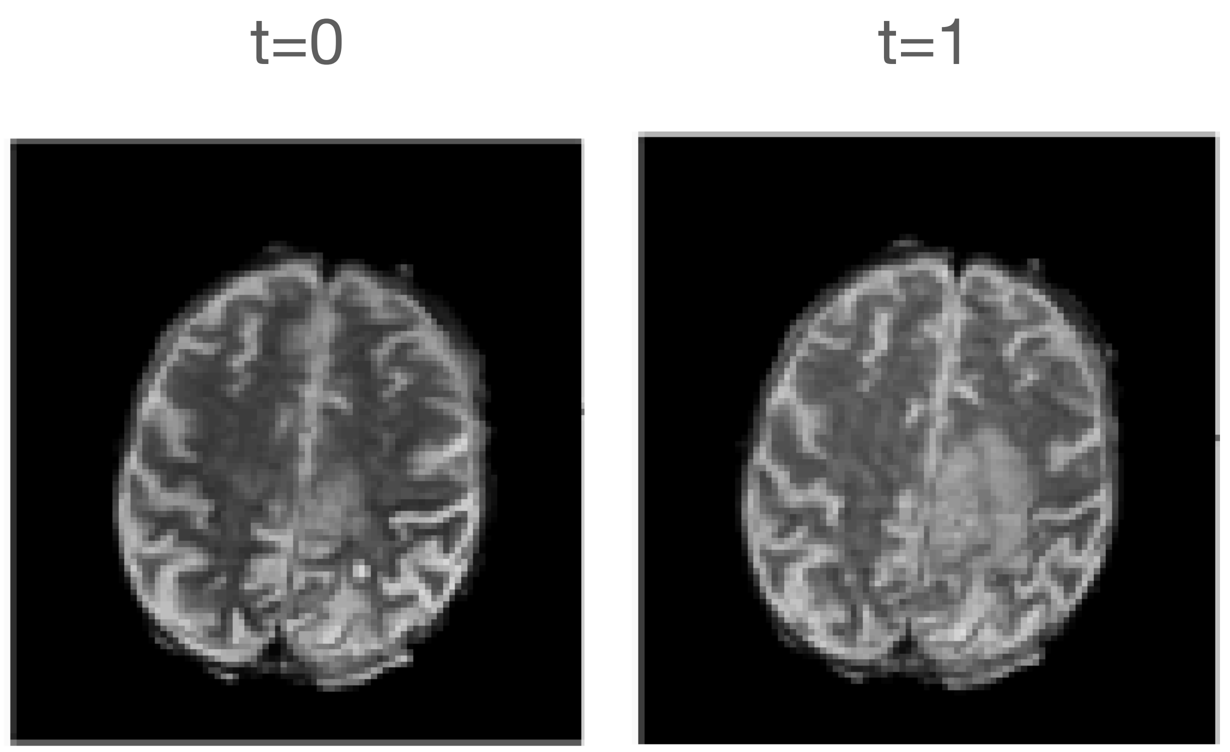

[Di Noto et al. ISMRM 2022]

Two learning scenarios for heterogeneous data

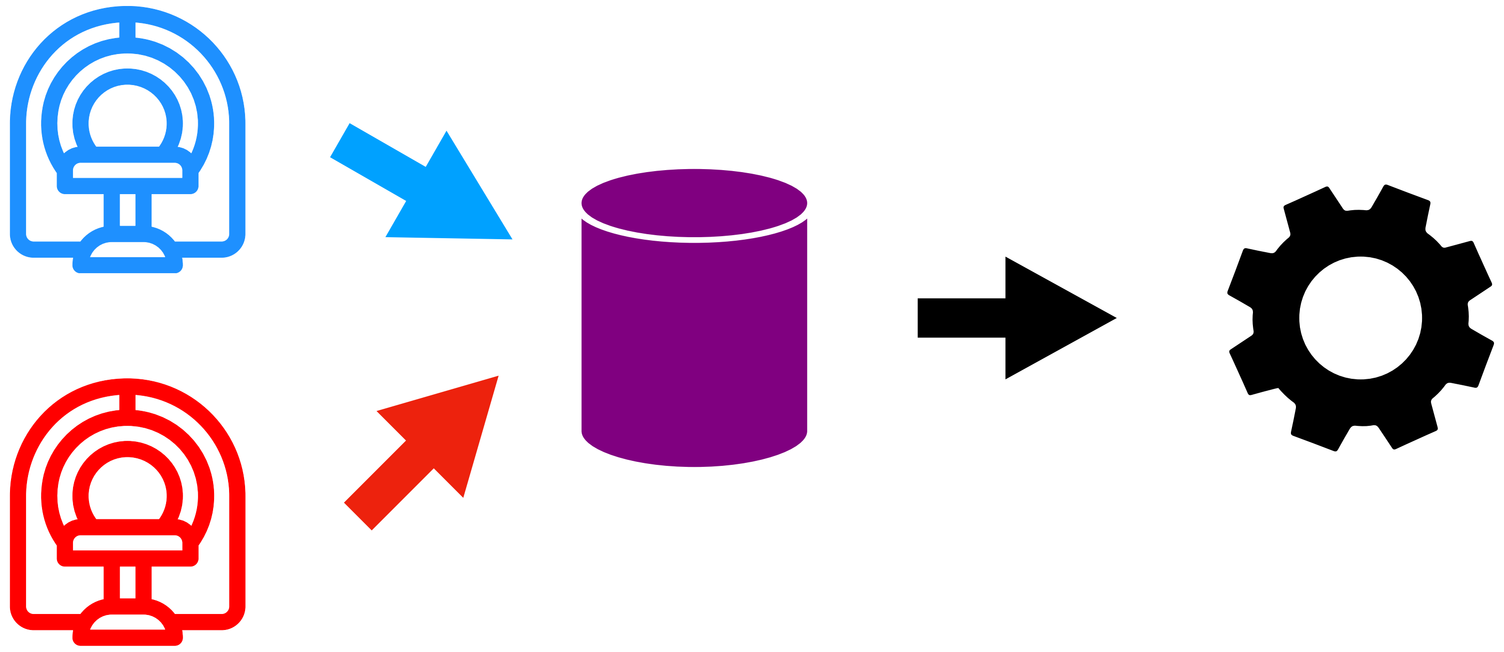

centralised

Two learning scenarios for heterogeneous data

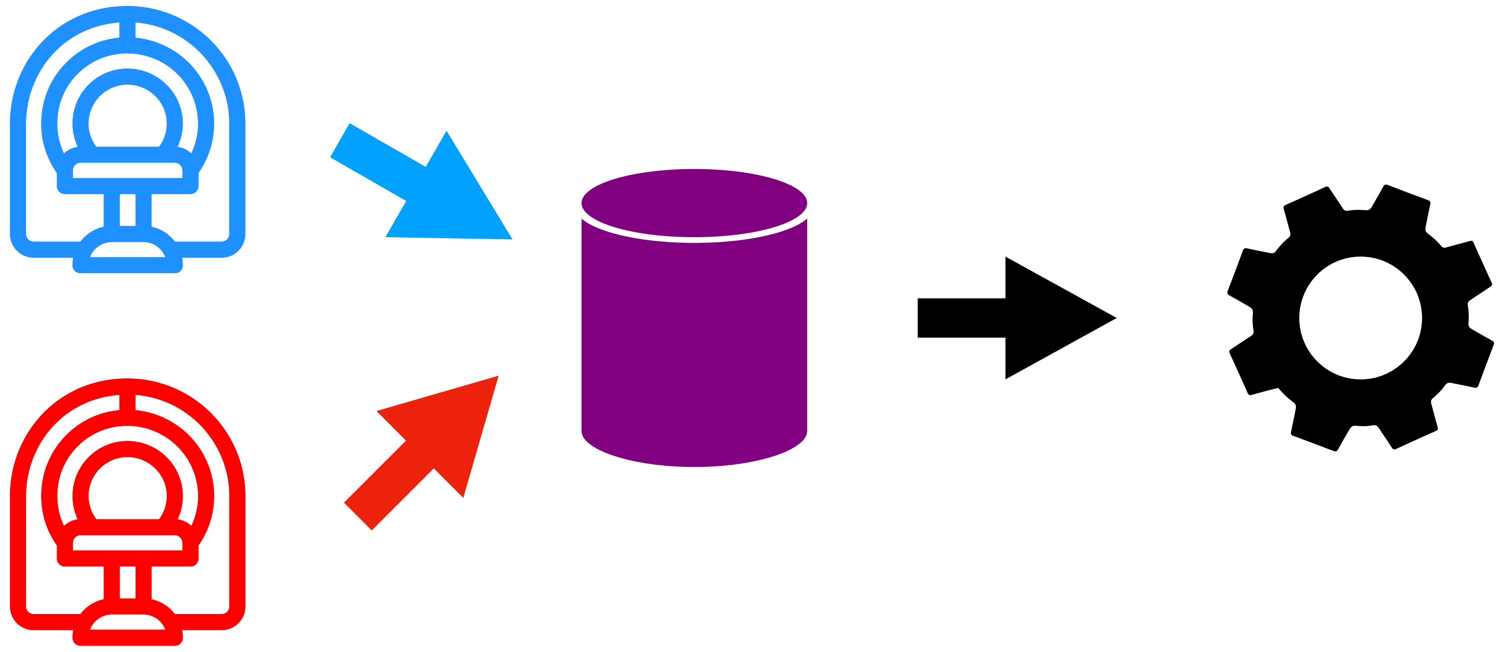

centralised

adapt the data, the training, or the model

Two learning scenarios for heterogeneous data

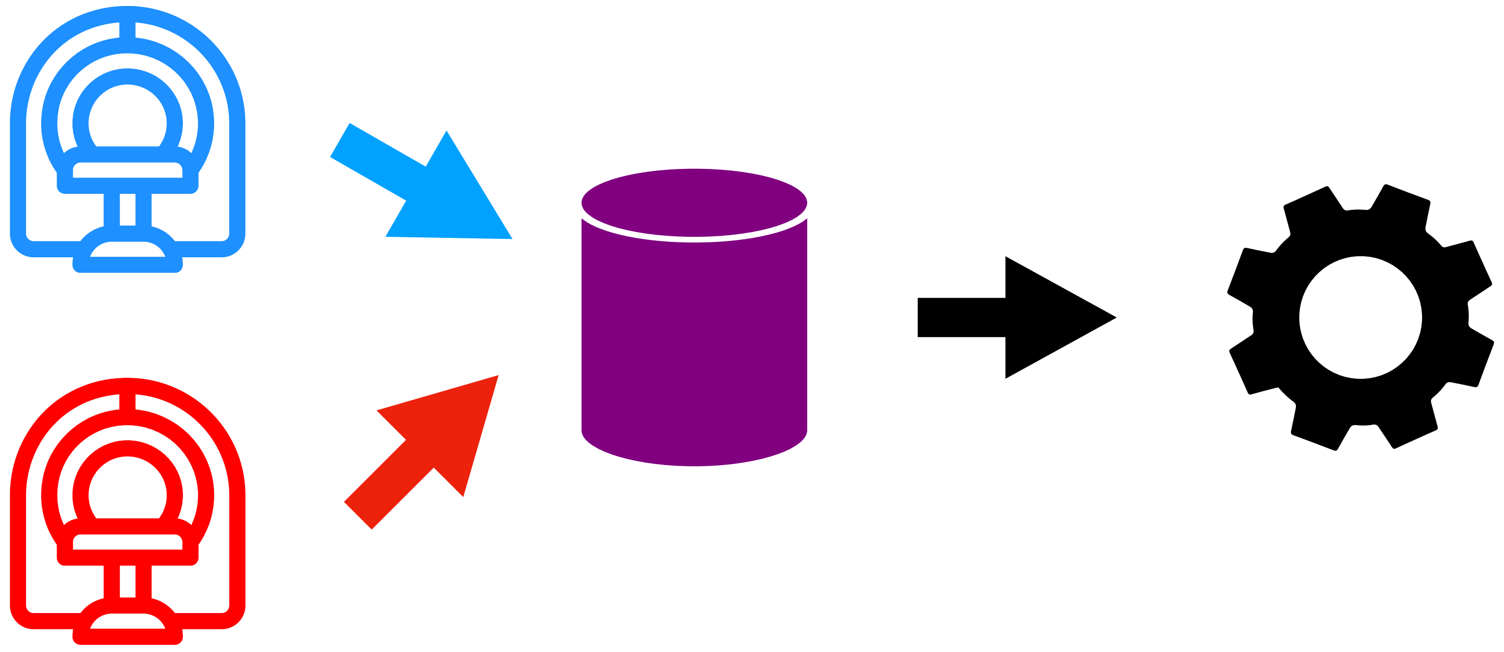

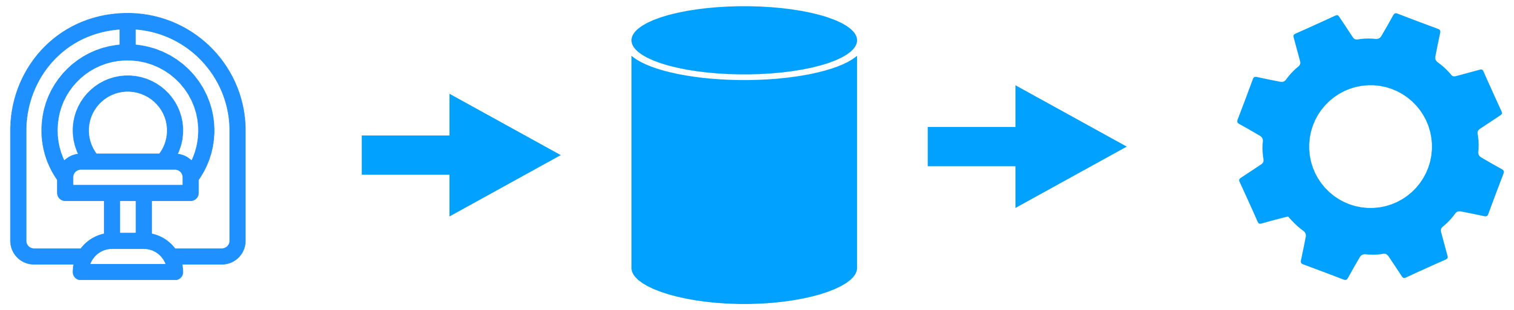

centralised

federated

adapt the data, the training, or the model

Two learning scenarios for heterogeneous data

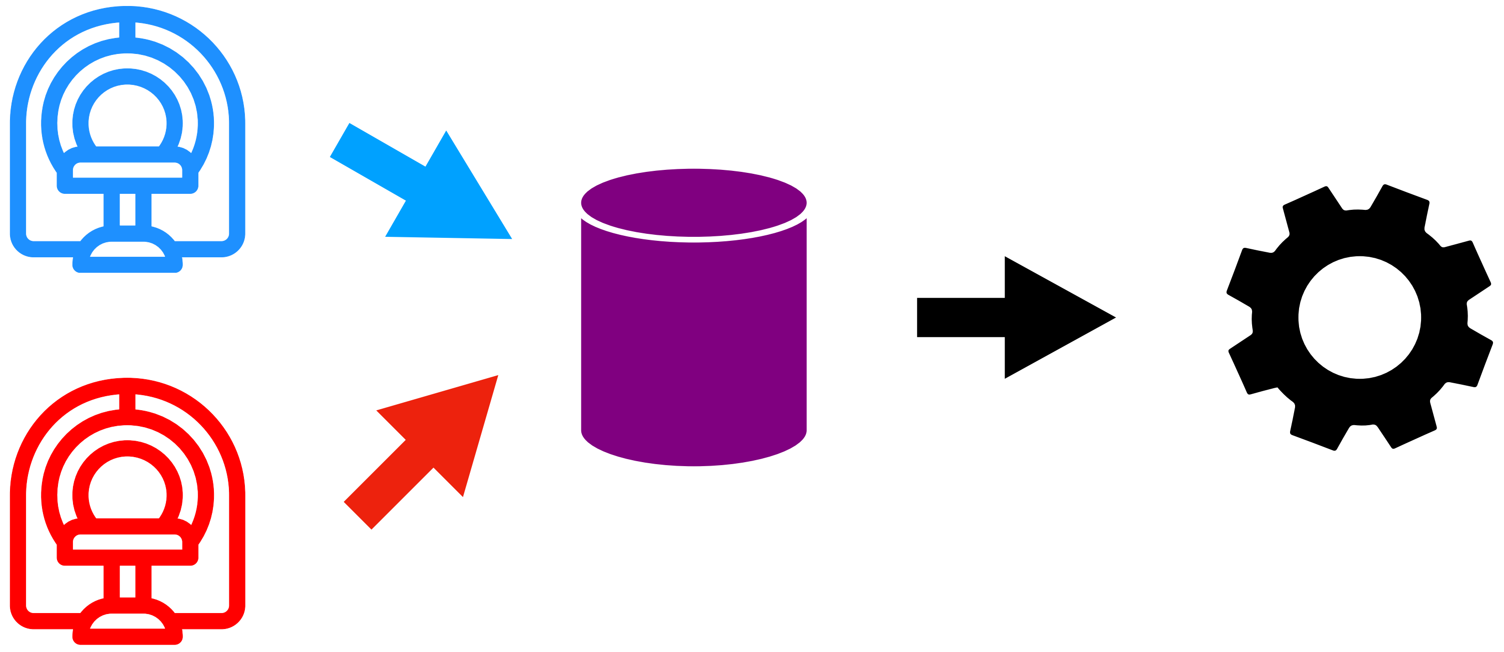

centralised

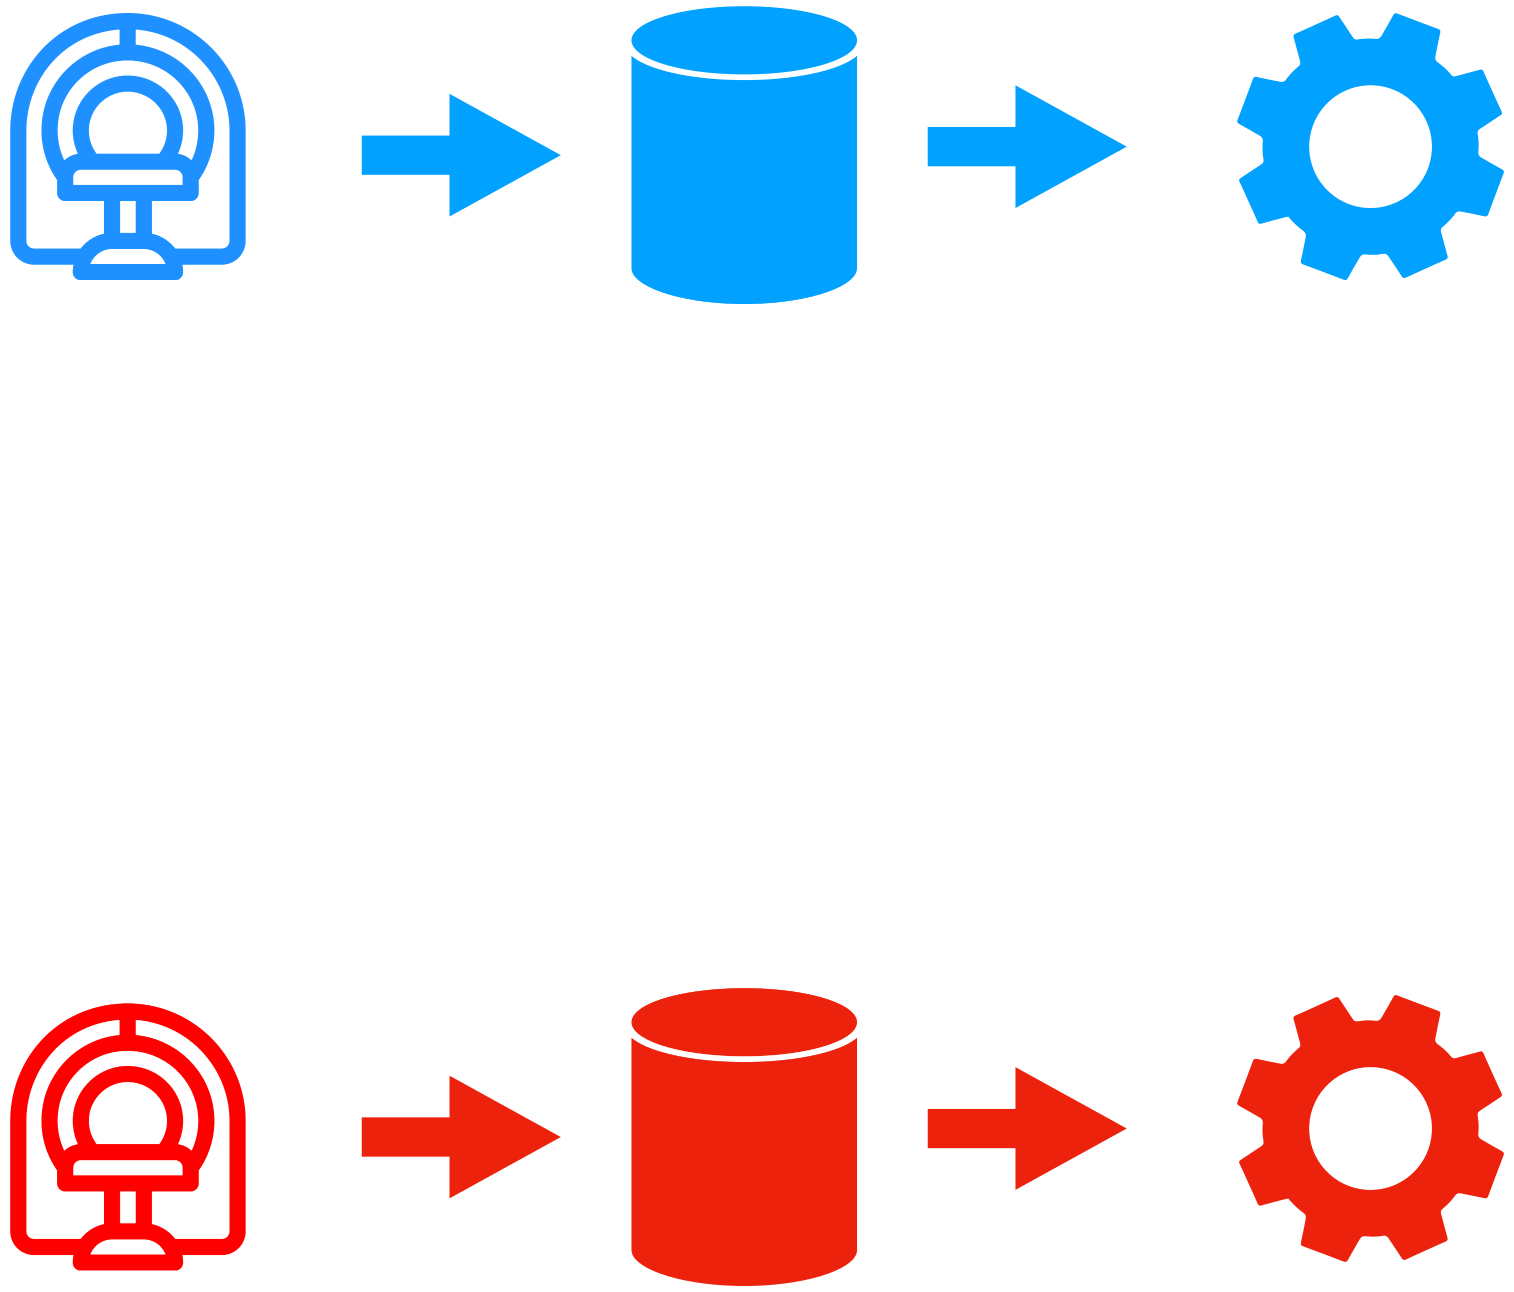

federated

adapt the data, the training, or the model

Two learning scenarios for heterogeneous data

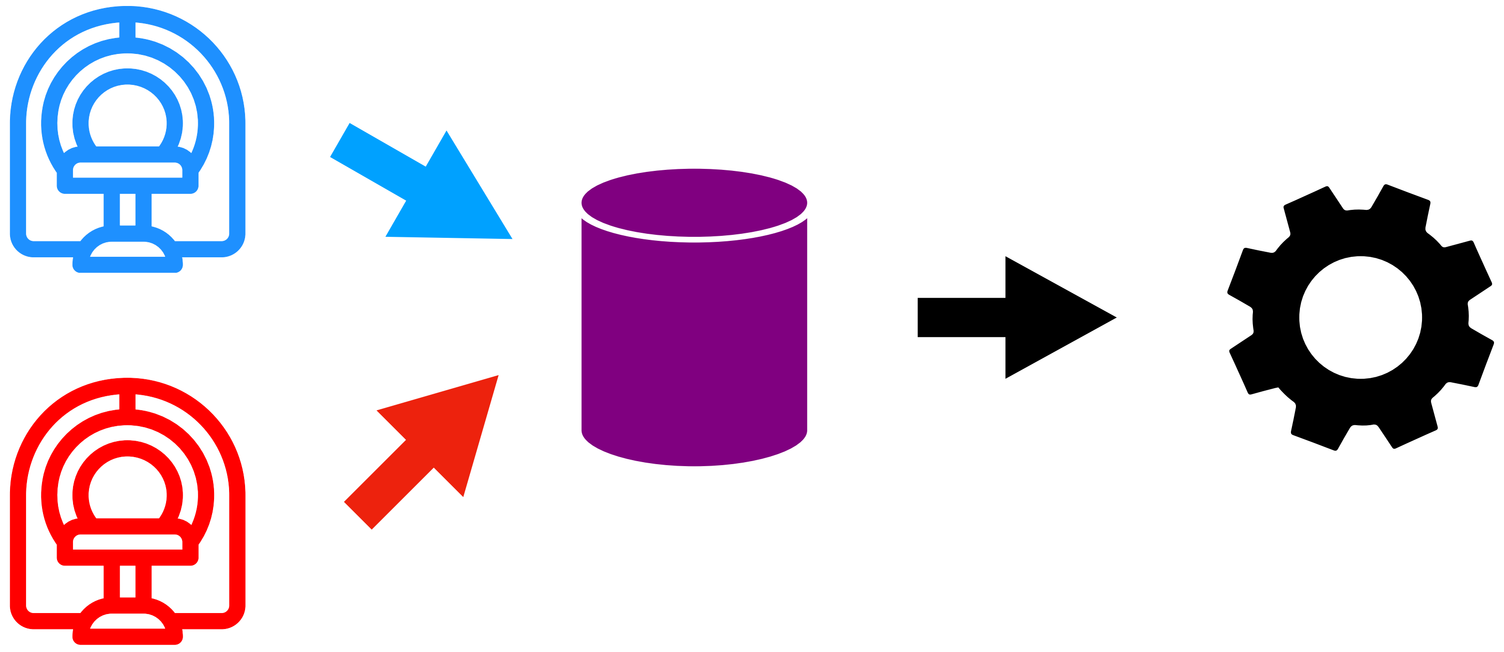

centralised

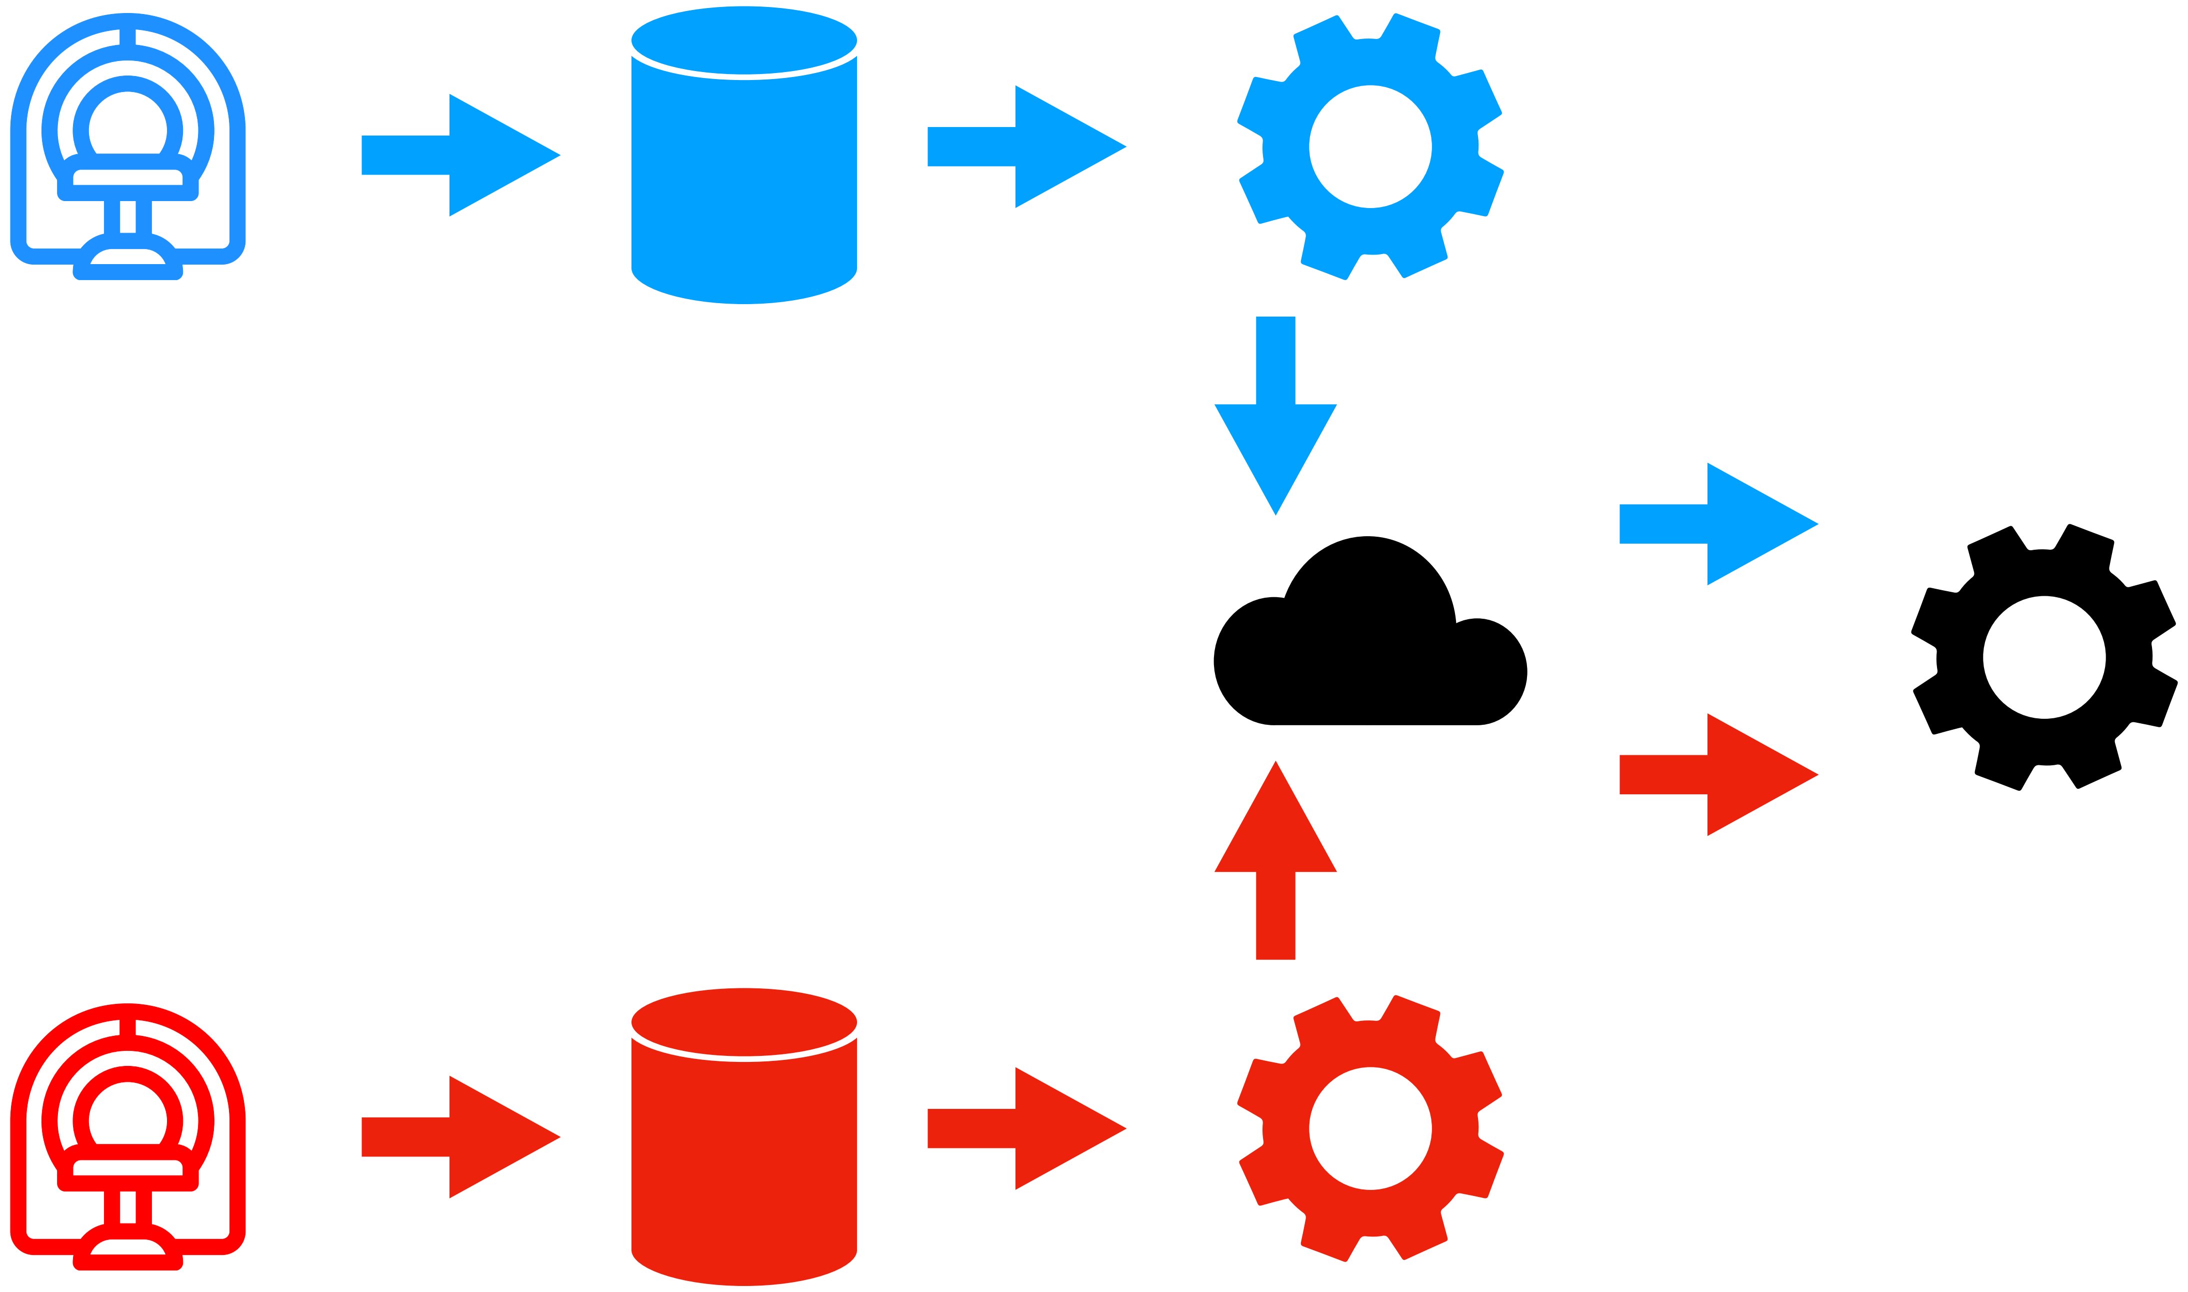

federated

adapt the data, the training, or the model

Two learning scenarios for heterogeneous data

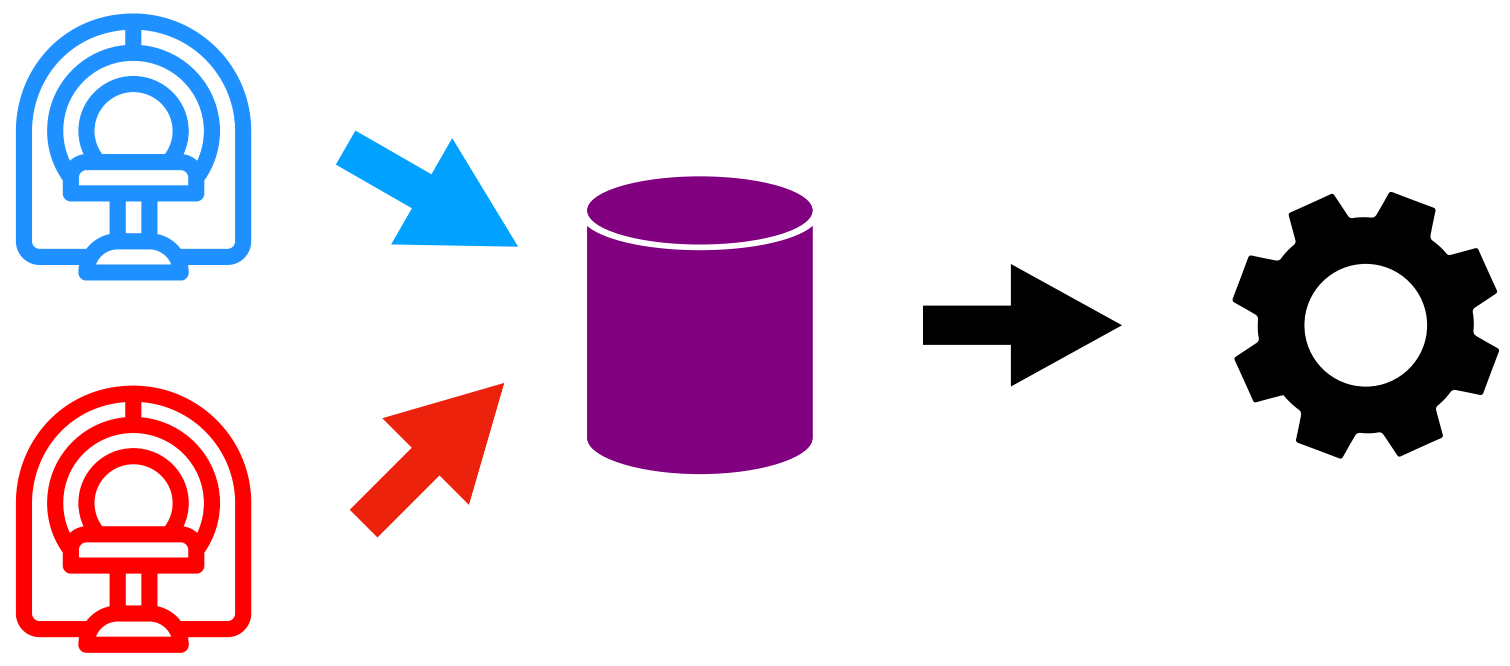

centralised

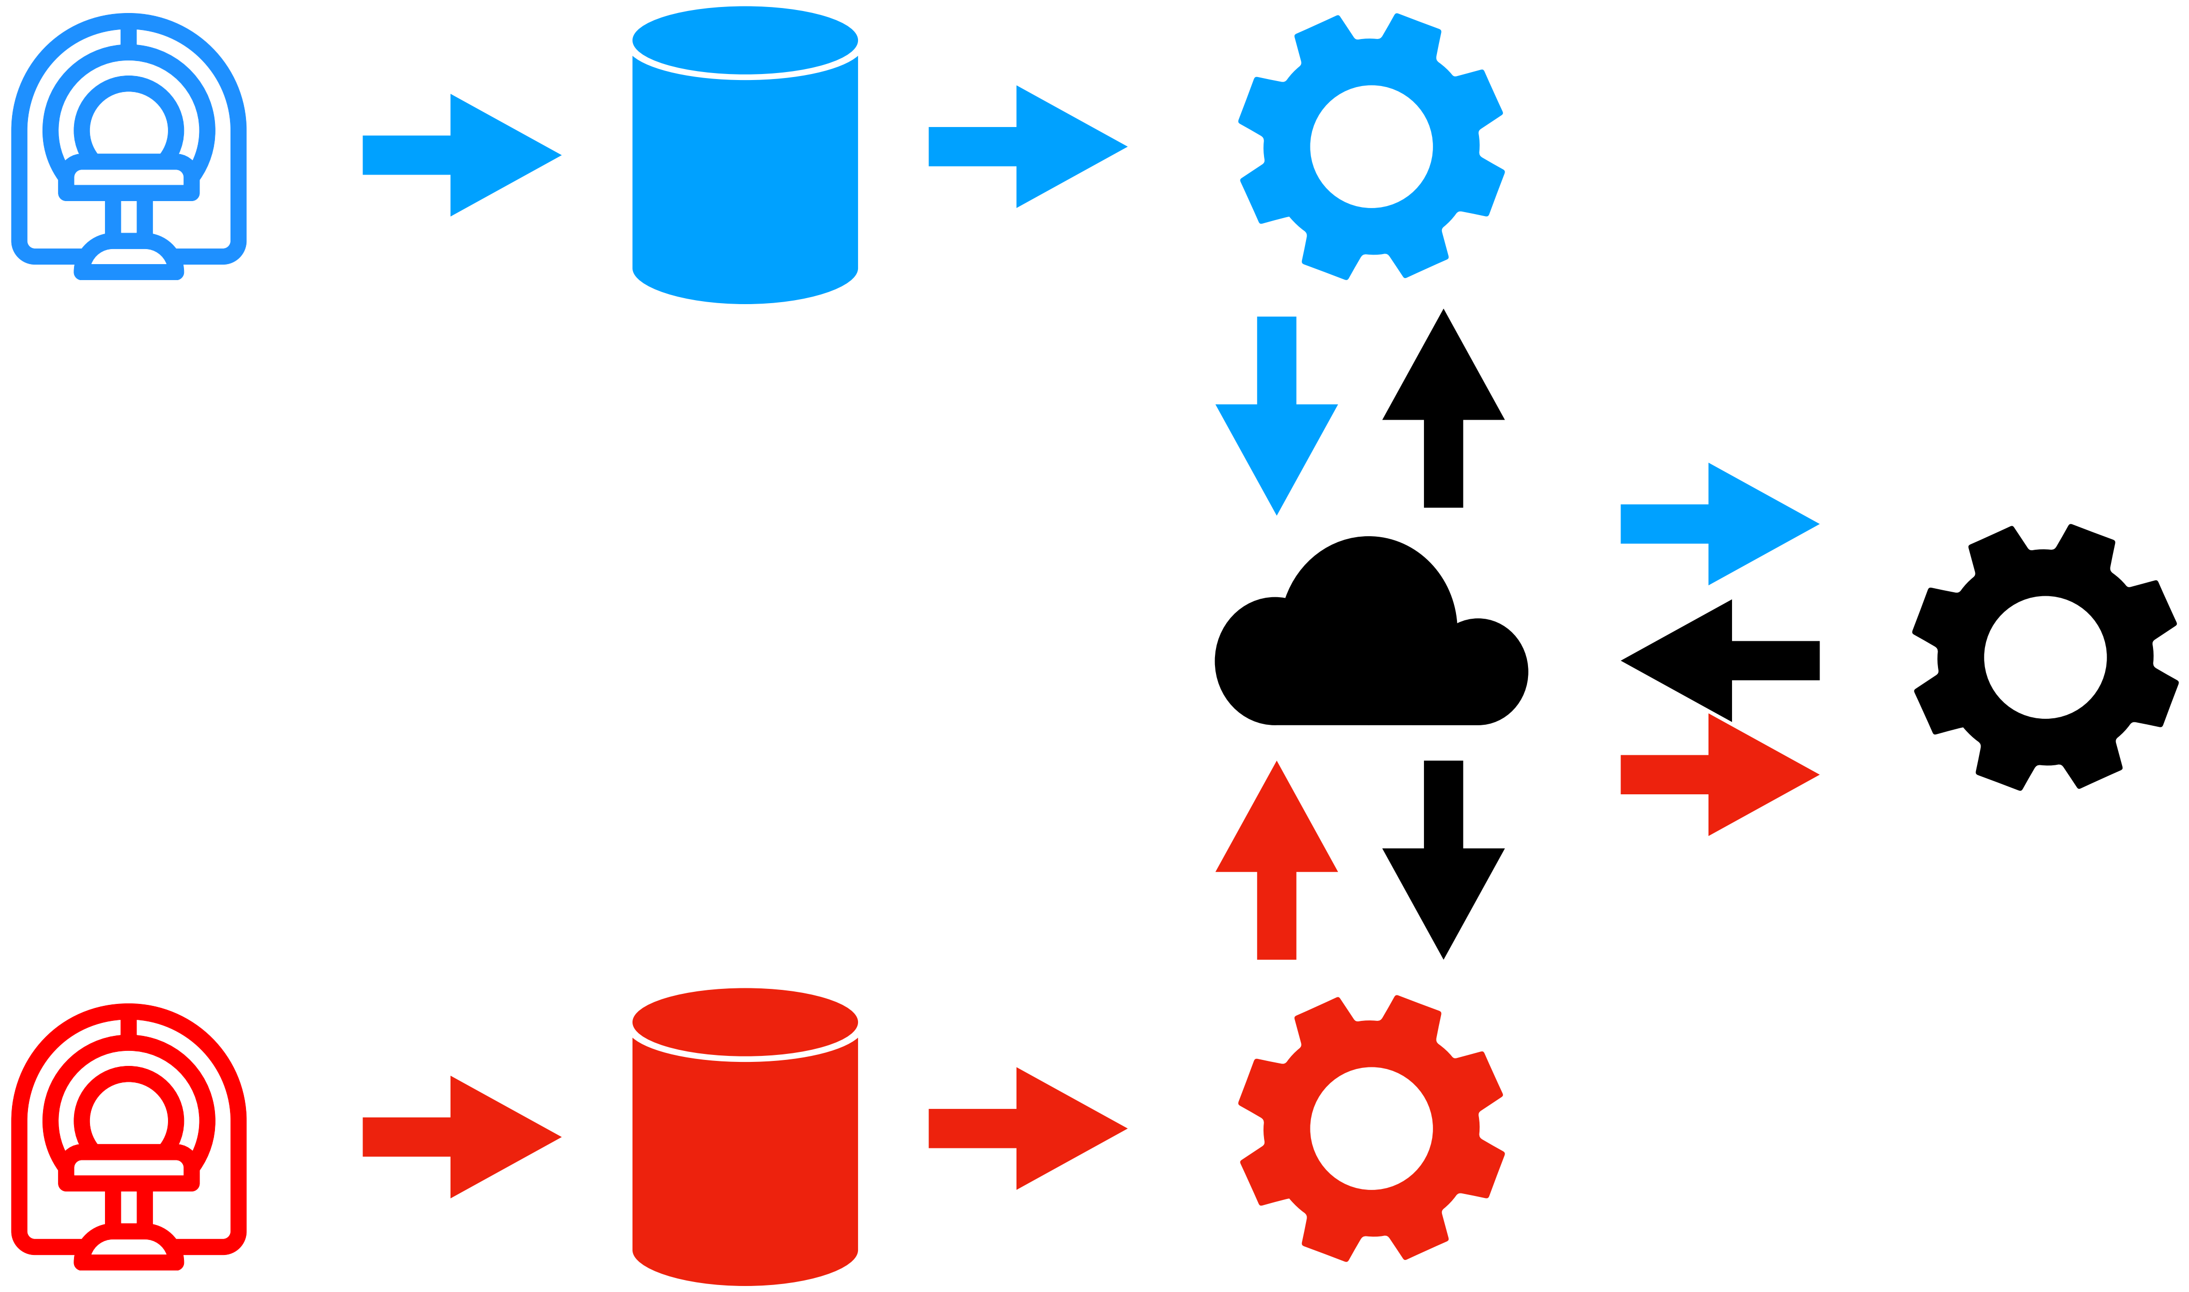

federated

adapt the data, the training, or the model

Two learning scenarios for heterogeneous data

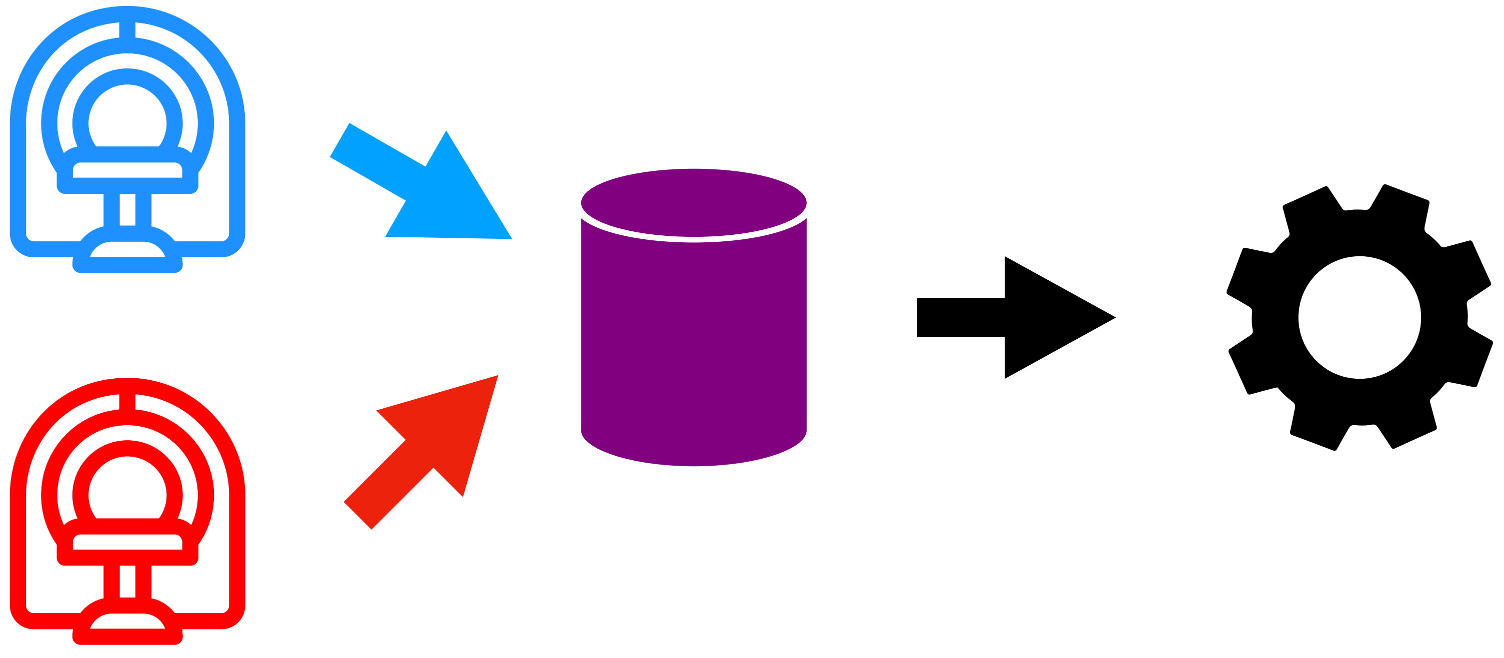

centralised

adapt the data, the training, or the model

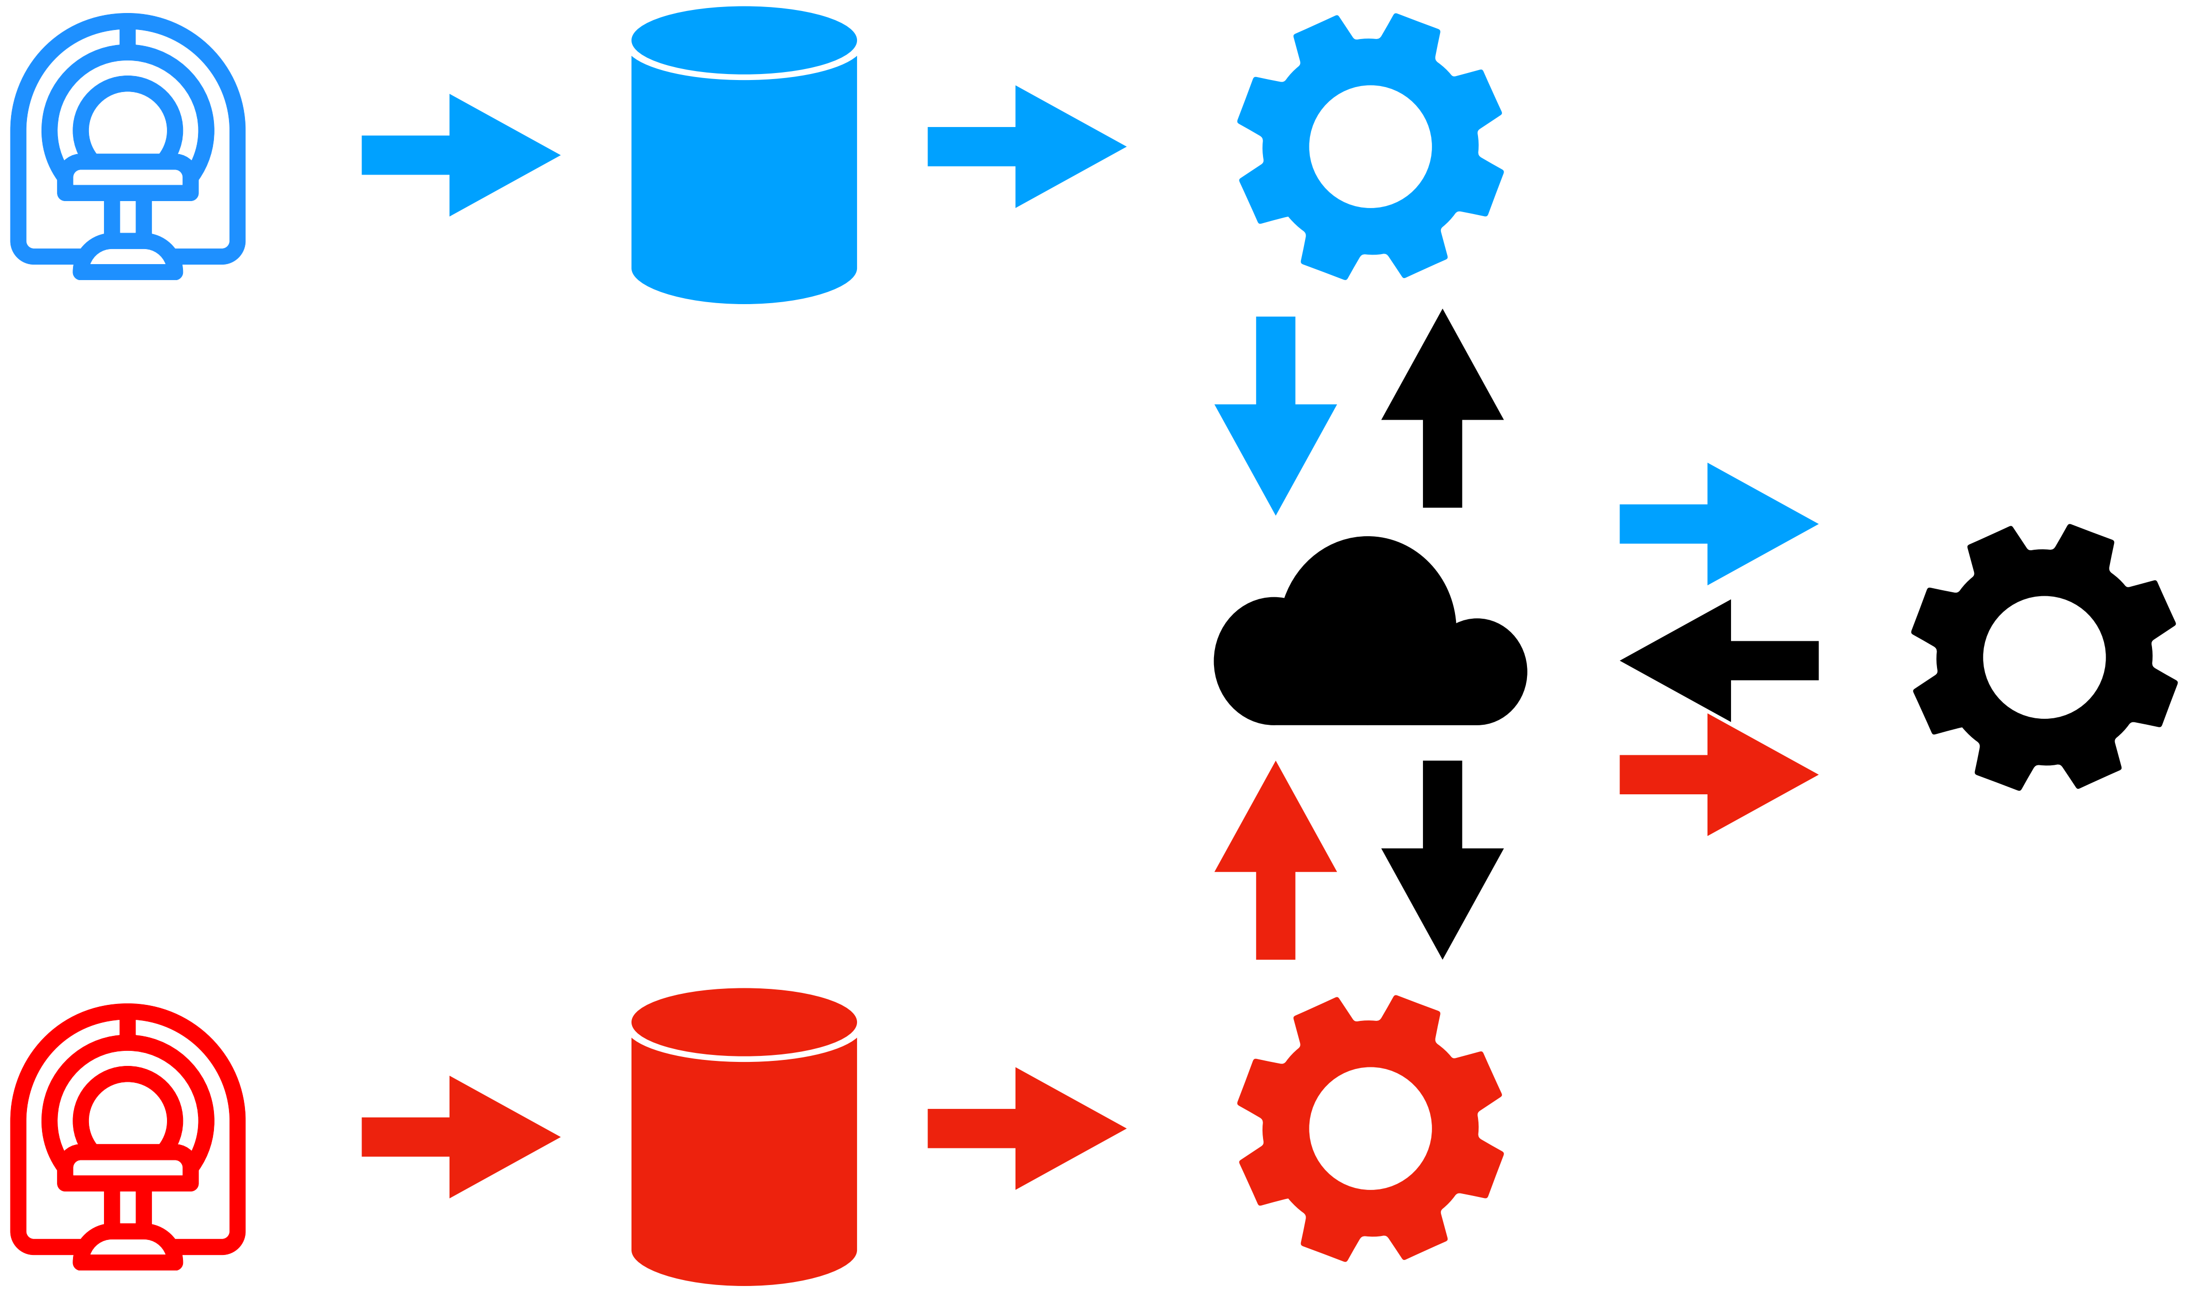

federated

adapt the training

Centralised case

Automated brain volumetry with heterogeneous data

[Stacy Jannis, alz.org]

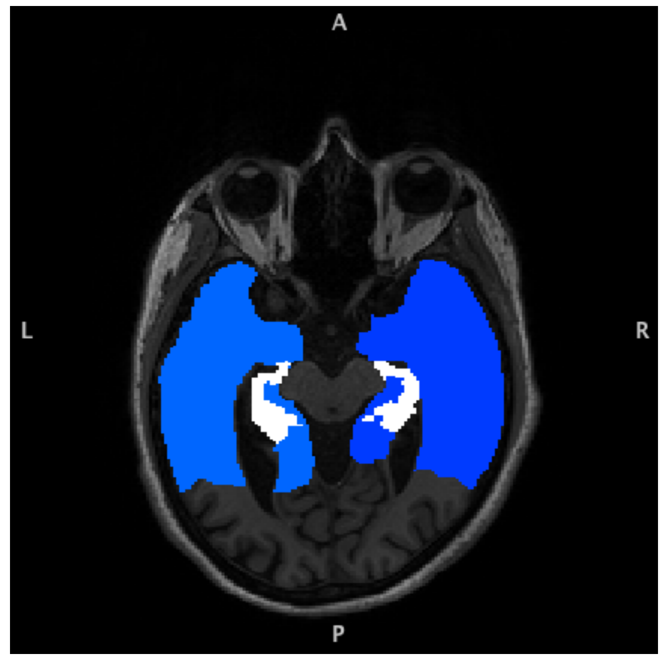

[Richiardi et al, sub.]

Veronica Ravano

Automated brain volumetry with heterogeneous data

[Stacy Jannis, alz.org]

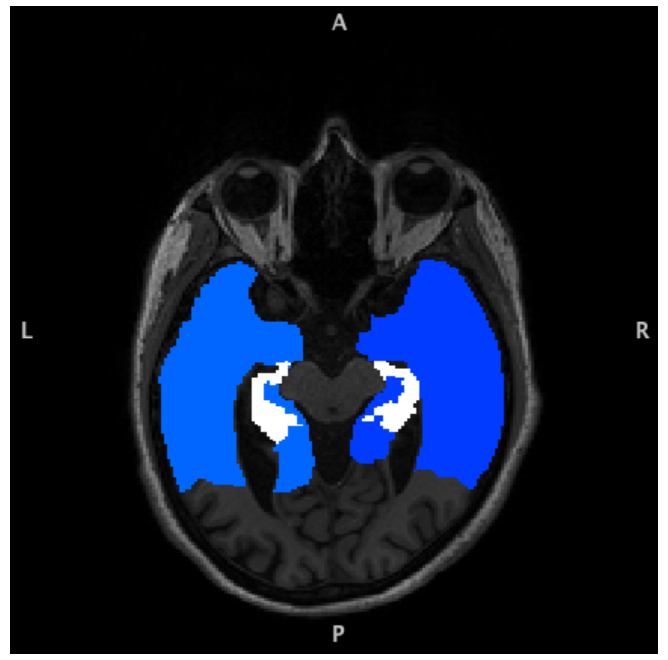

[Richiardi et al, sub.]

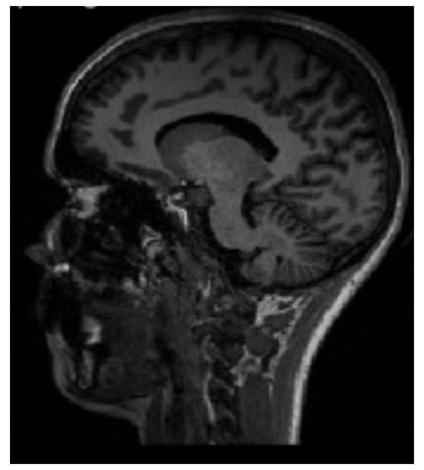

Veronica Ravano

Automated brain volumetry with heterogeneous data

[Stacy Jannis, alz.org]

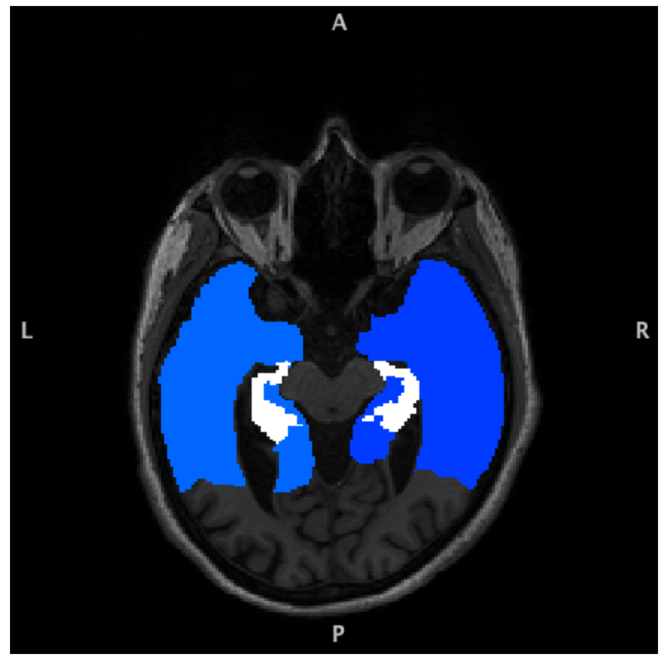

[Richiardi et al, sub.]

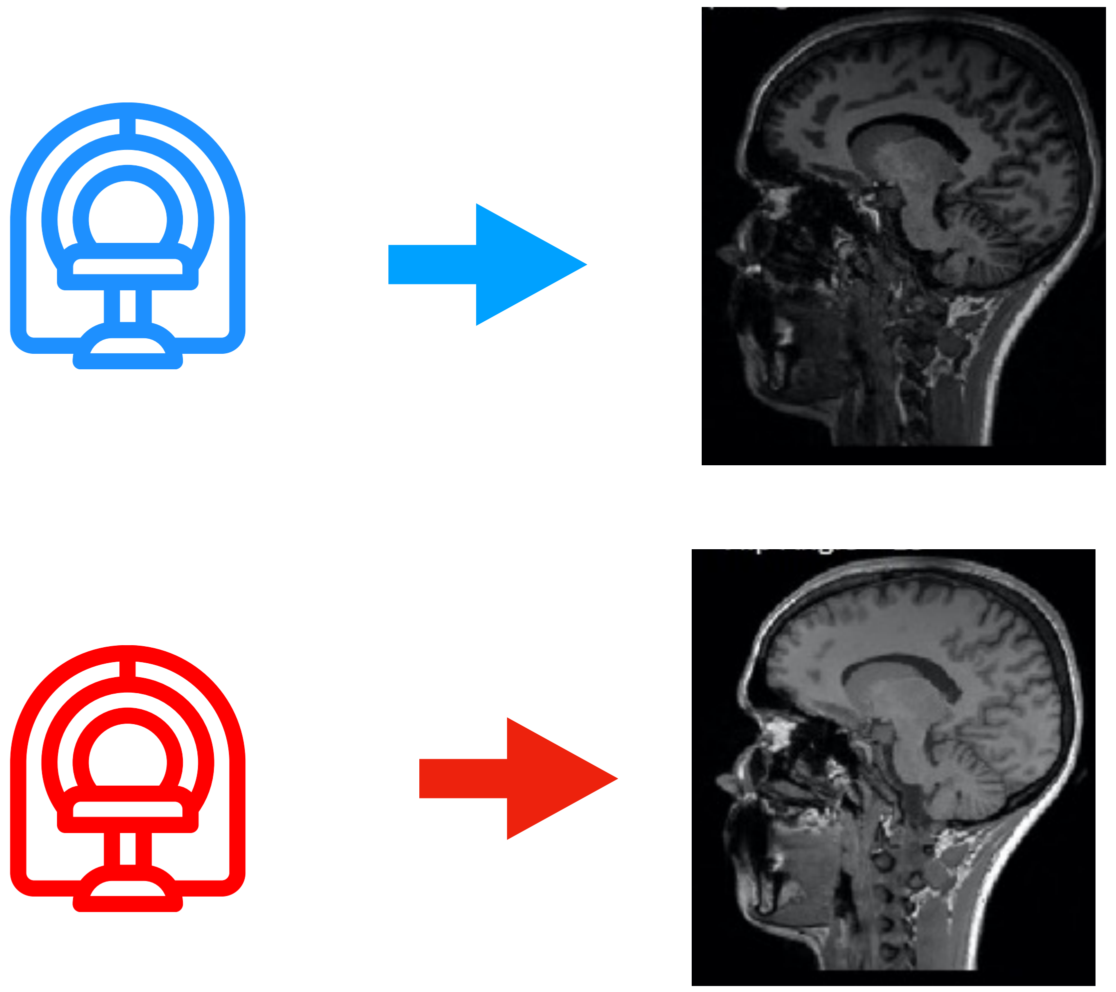

[Ravano et al, ISMRM 2021]

Veronica Ravano

Automated brain volumetry with heterogeneous data

[Stacy Jannis, alz.org]

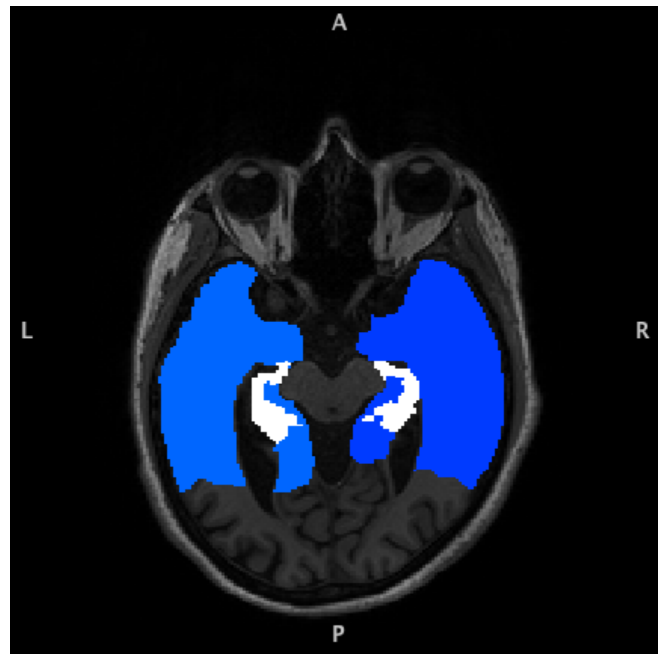

[Richiardi et al, sub.]

relative volume error due to protocol change

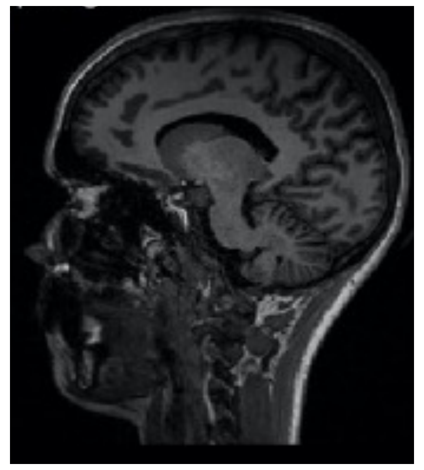

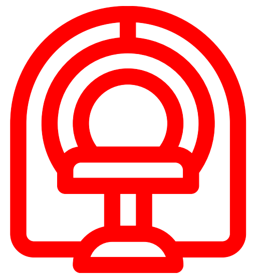

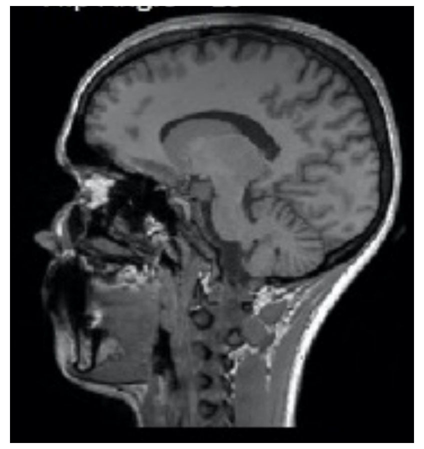

[Ravano et al, ISMRM 2021]

Veronica Ravano

Automated brain volumetry with heterogeneous data

[Stacy Jannis, alz.org]

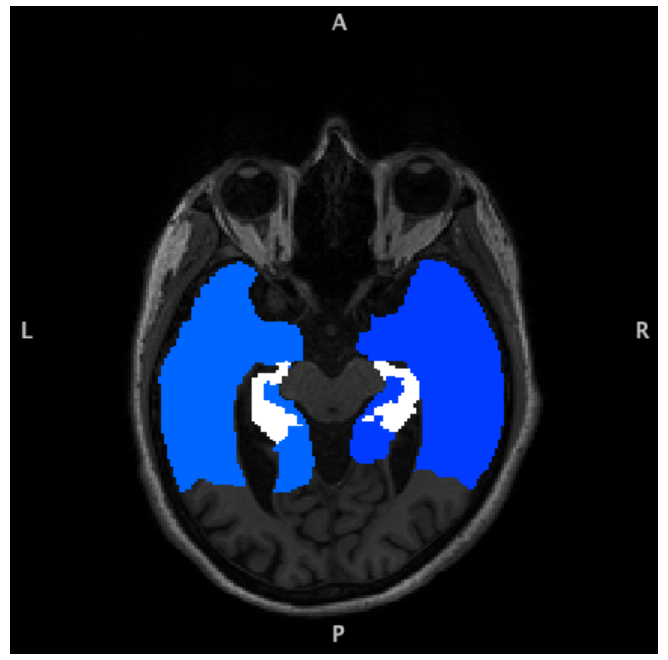

[Richiardi et al, sub.]

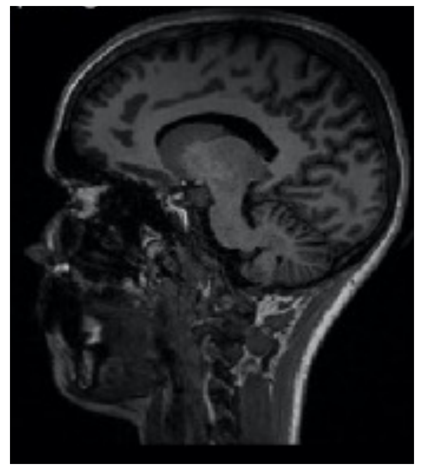

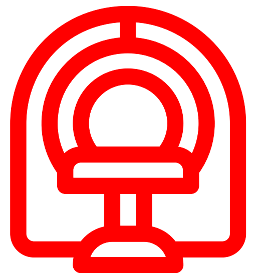

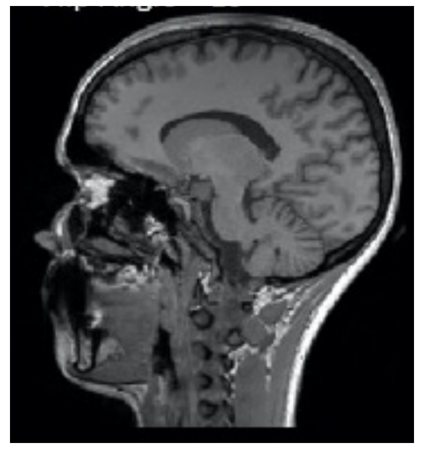

[Ravano et al, ISMRM 2021]

relative volume error due
to protocol change

Hippocampus: 5%
Cingulate: 8%
Deep white matter: 10%

Veronica Ravano

Automated brain volumetry with heterogeneous data

[Stacy Jannis, alz.org]

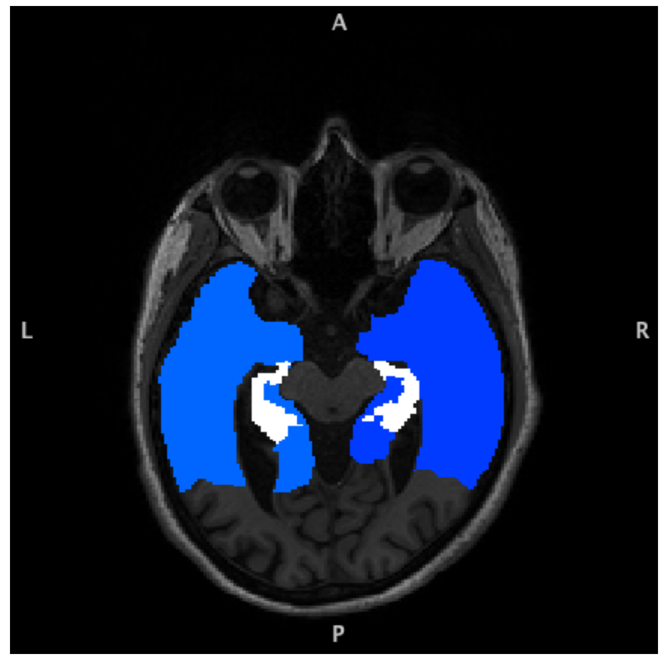

[Richiardi et al, sub.]

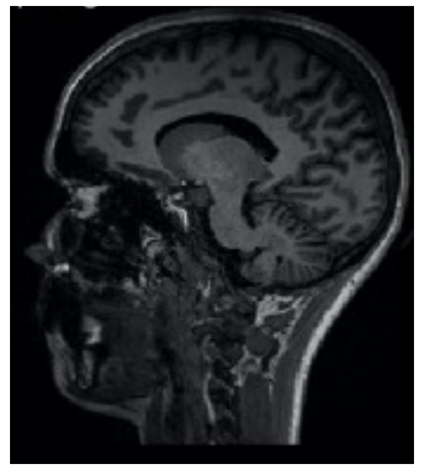

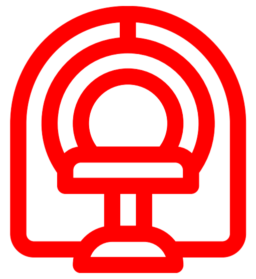

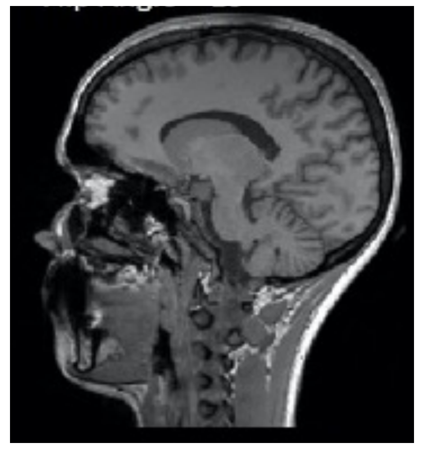

[Ravano et al, ISMRM 2021]

relative volume error due
to protocol change

Hippocampus: 5%
Cingulate: 8%
Deep white matter: 10%

annual hippocampus
atrophy in AD patients ~1%

Veronica Ravano

Centralised harmonization

Domain adaptation with image-to-image translation

CGAN [Isola et al., 2017]

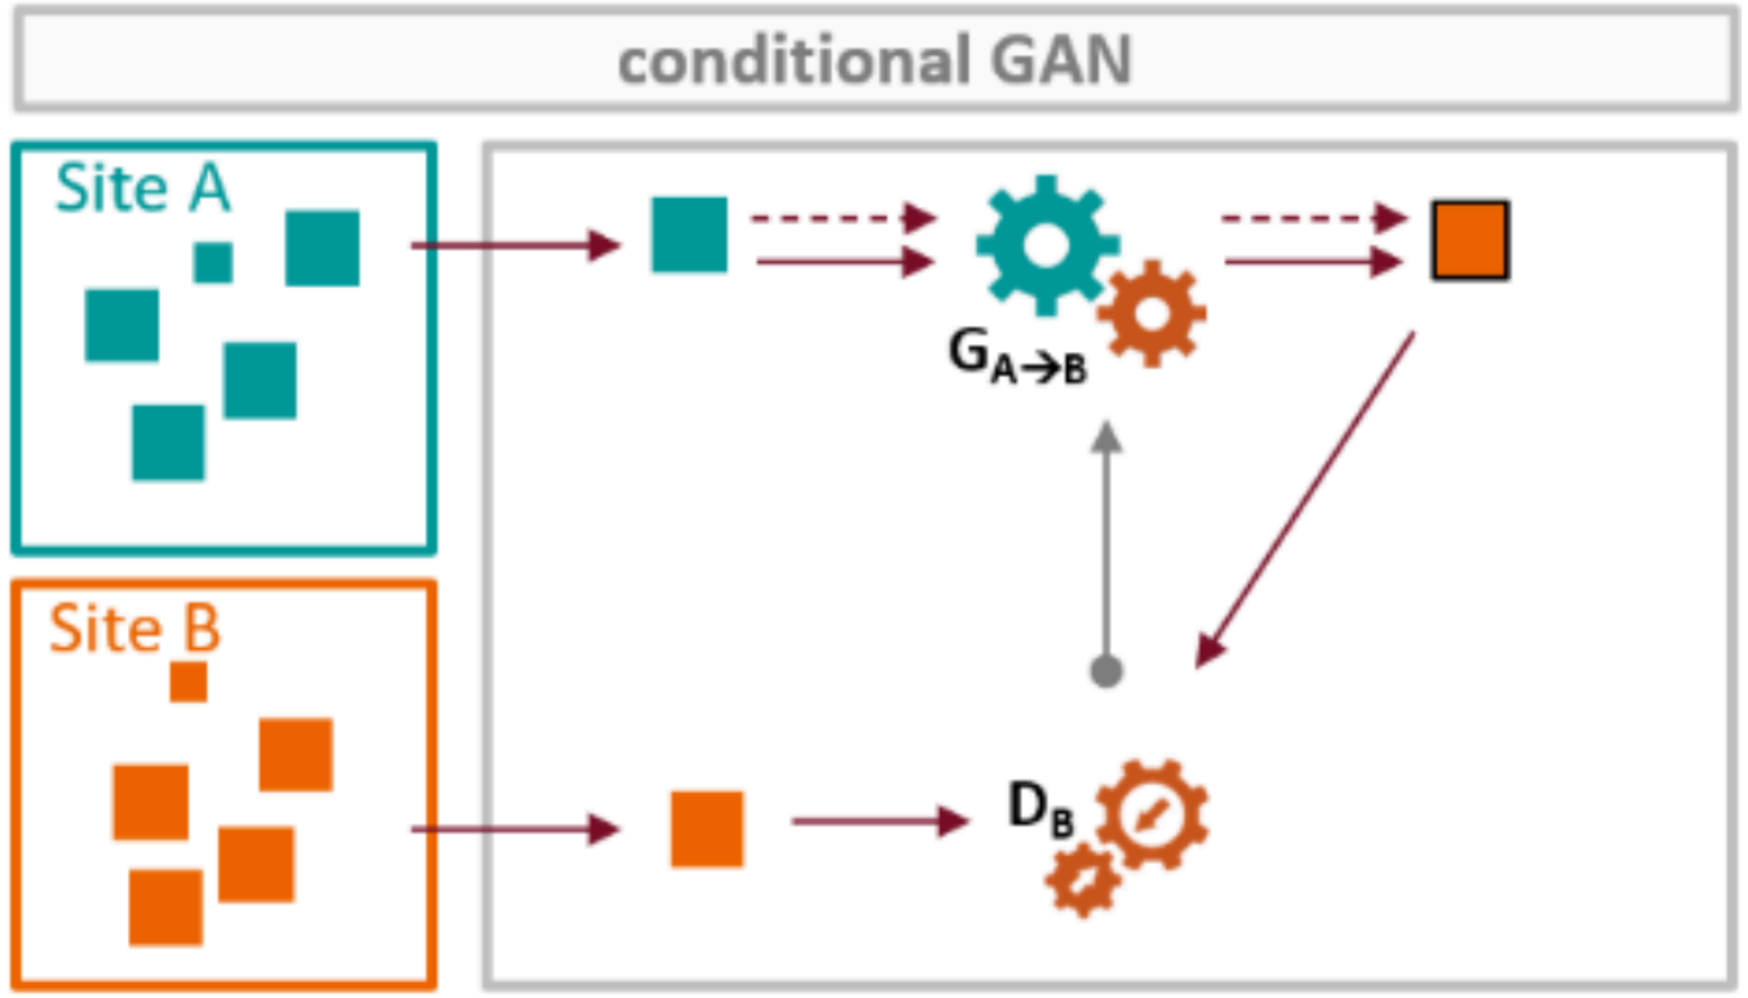

[Ravano et al., ISMRM 2022] [Ravano et al., sub.]

Centralised harmonization

Domain adaptation with image-to-image translation

CGAN [Isola et al., 2017]

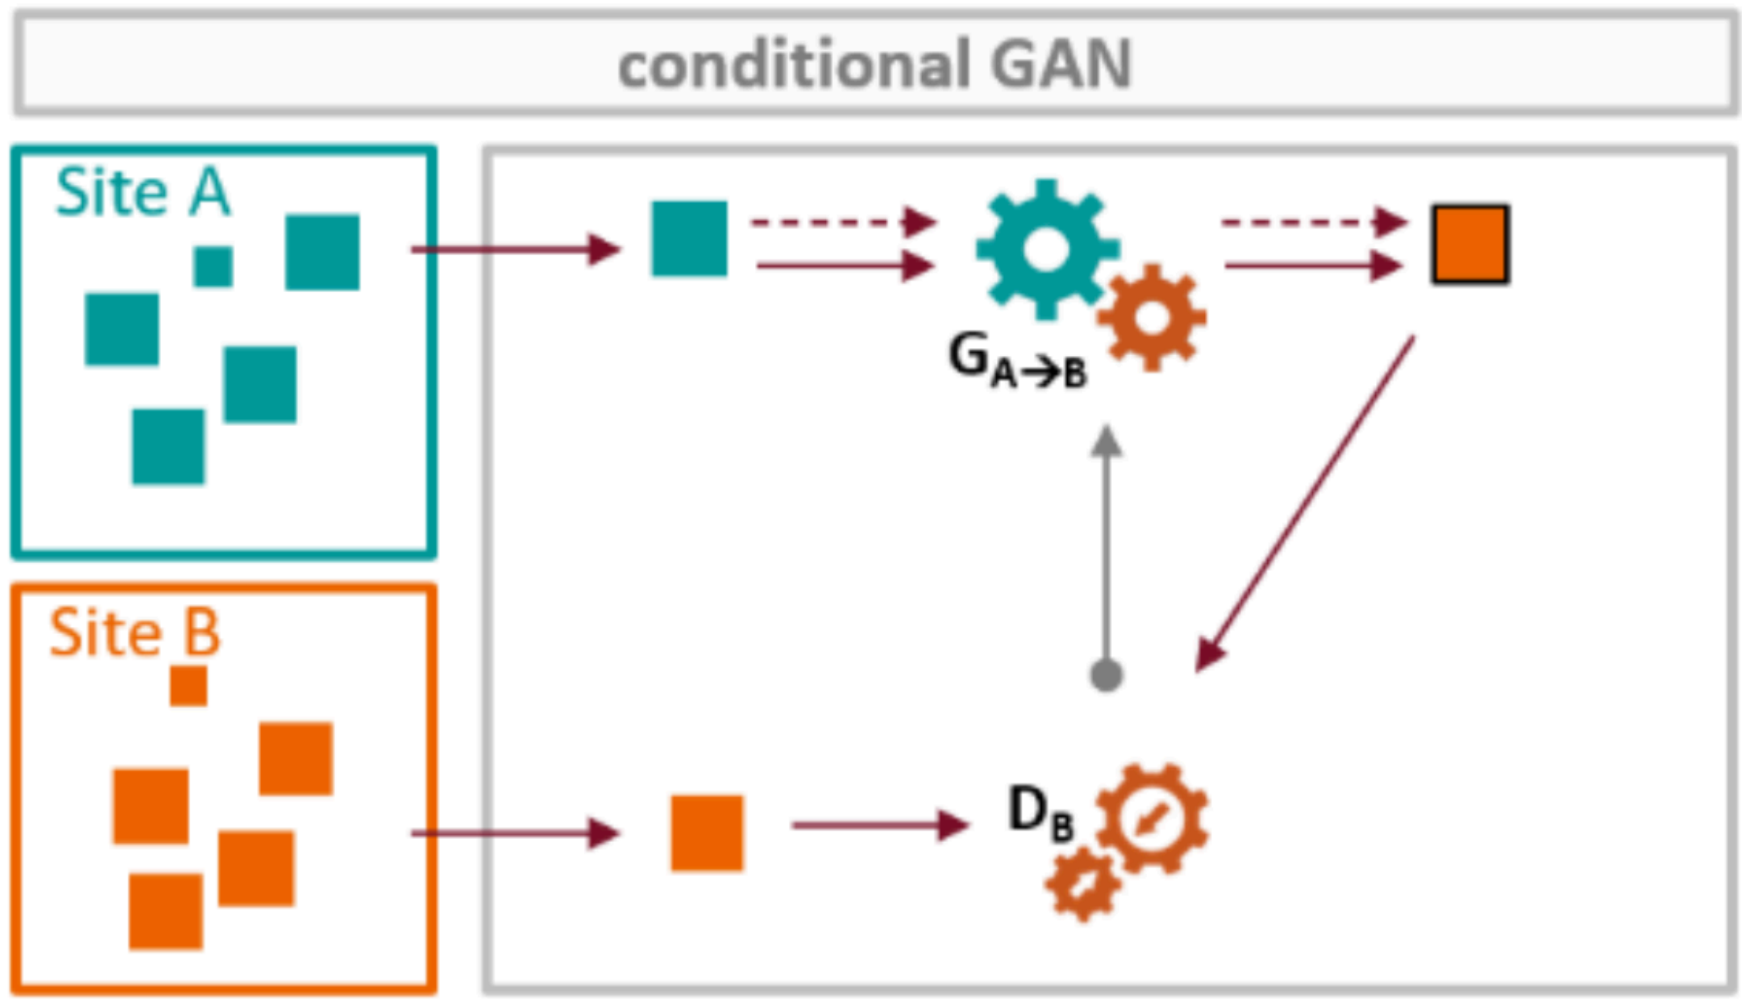

$$\mathcal{L}_G = \lambda_{L1} \sum_i^n \frac{|\hat{y}_i - y_i|}{n} + \lambda_{L2} \sum_i^n \frac{(\hat{y}_i - y_i)^2}{n} + \lambda_{LPIPS} LPIPS(\hat{\mathbf{y}}, \mathbf{y}) + \mathcal{L}_{adv}$$

[Ravano et al., ISMRM 2022] [Ravano et al., sub.]

Centralised harmonization

Domain adaptation with image-to-image translation

CGAN [Isola et al., 2017]

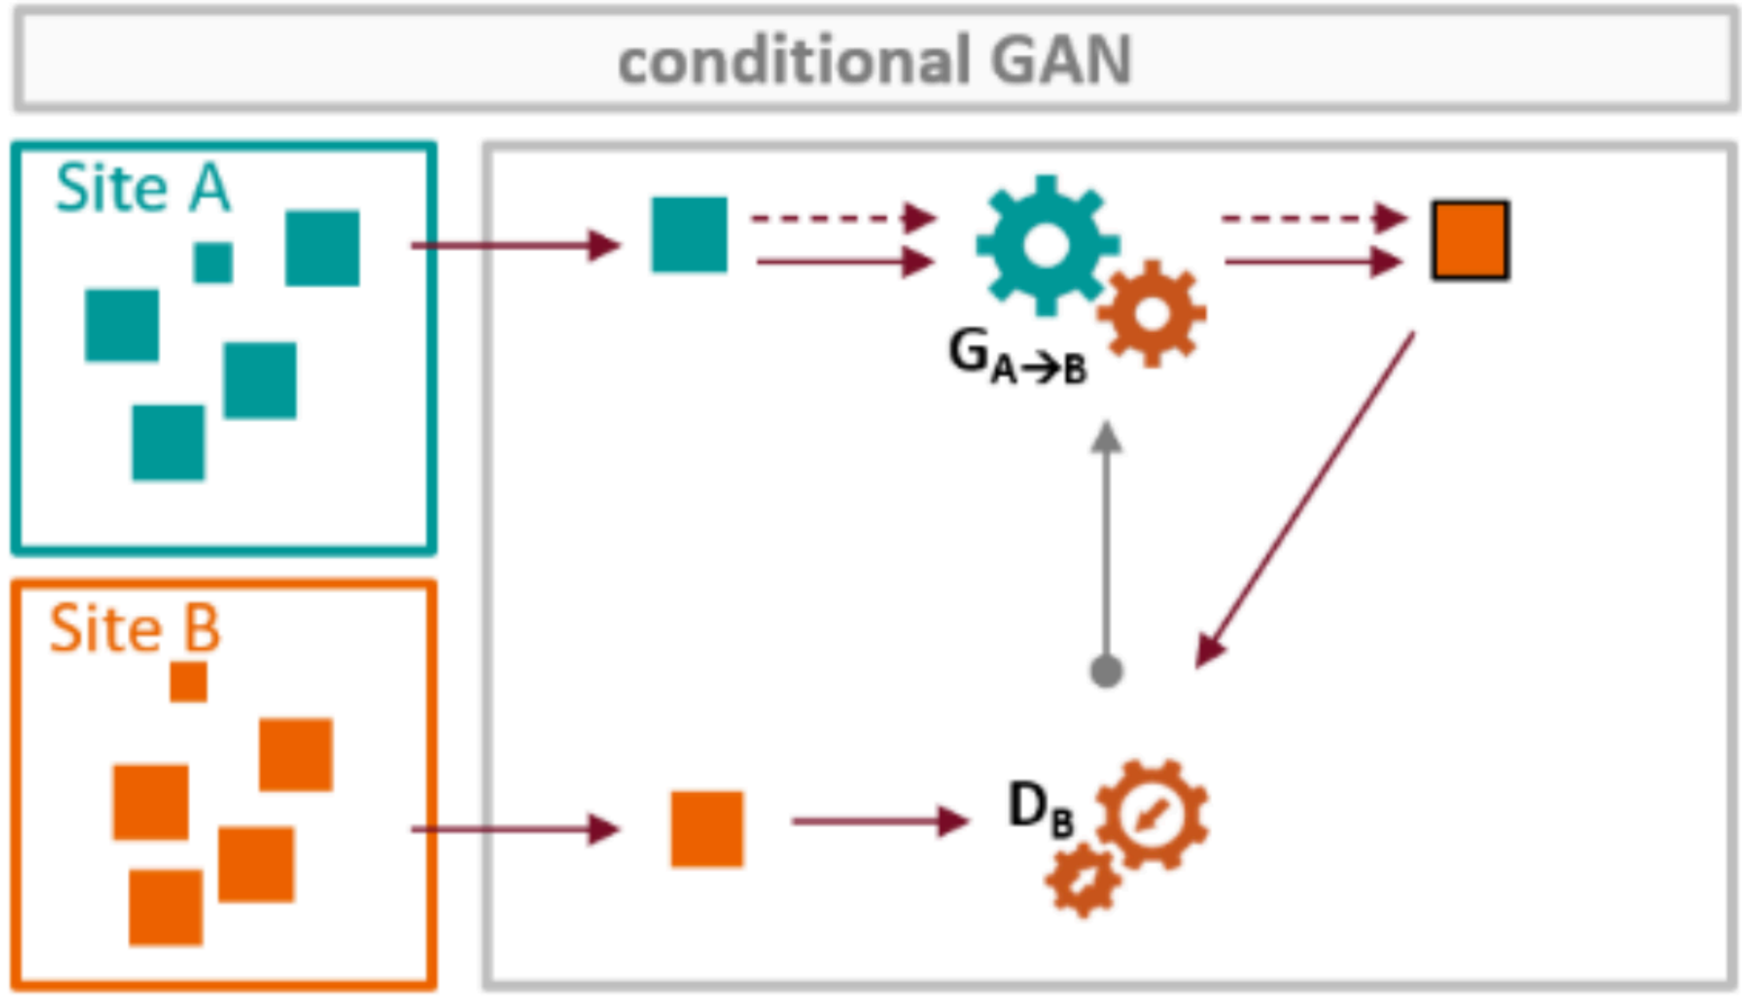

	Protocol 1	Protocol 2
TR/TE/TI [ms]	2300/2.98/900	1930/2.36/972
Resolution [mm^3]	1x1x1.1	0.87x0.87x0.9
Flip Angle [°]	9	8
Pixel readout bandwidth [Hz/ms]	240	200

N=74 paired (64 TR, 3x10 TE)

$$\mathcal{L}_G = \lambda_{L1} \sum_i^n \frac{|\hat{y}_i - y_i|}{n} + \lambda_{L2} \sum_i^n \frac{(\hat{y}_i - y_i)^2}{n} + \lambda_{LPIPS} LPIPS(\hat{\boldsymbol{y}}, \boldsymbol{y}) + \mathcal{L}_{adv}$$

[Ravano et al., ISMRM 2022] [Ravano et al., sub.]

Centralised harmonization

Domain adaptation with image-to-image translation

CGAN [Isola et al., 2017]

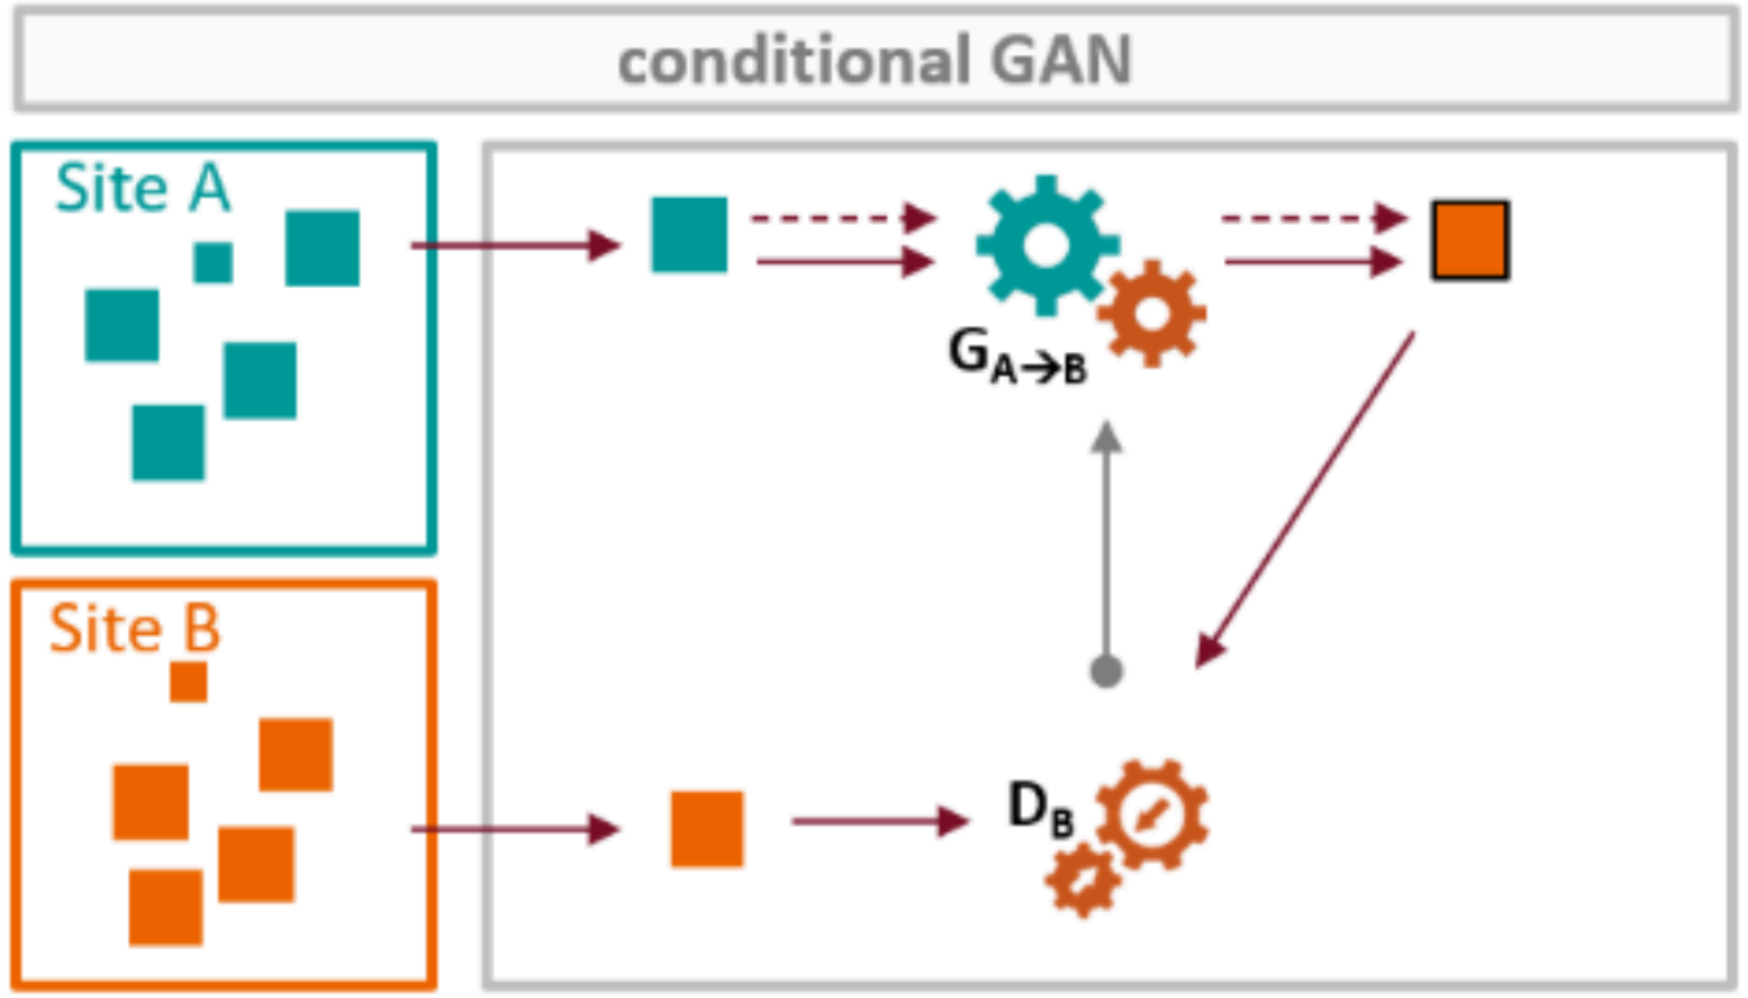

	Protocol 1	Protocol 2
TR/TE/TI [ms]	2300/2.98/900	1930/2.36/972
Resolution [mm^3]	1x1x1.1	0.87x0.87x0.9
Flip Angle [°]	9	8
Pixel readout bandwidth [Hz/ms]	240	200

N=74 paired (64 TR, 3x10 TE)

Model	MAE ↓	MSE ↓	WD ↓	SSIM ↑	PSNR ↑
baseline	128±18.0	18.56±4.80x10^3	127±18.1	0.75±0.04	14.5±1.5
cGAN (L1 + LPIPS)	**52.4±13.0**	**3.84±1.70x10^3**	**49.1±14.9**	0.85±0.04	21.6±2.7

$$\mathcal{L}_G = \lambda_{L1} \sum_i^n \frac{|\hat{y}_i - y_i|}{n} + \lambda_{L2} \sum_i^n \frac{(\hat{y}_i - y_i)^2}{n} + \lambda_{LPIPS} LPIPS(\hat{\boldsymbol{y}}, \boldsymbol{y}) + \mathcal{L}_{adv}$$

[Ravano et al., ISMRM 2022] [Ravano et al., sub.]

Centralised harmonization

Domain adaptation with image-to-image translation

CGAN [Isola et al., 2017]

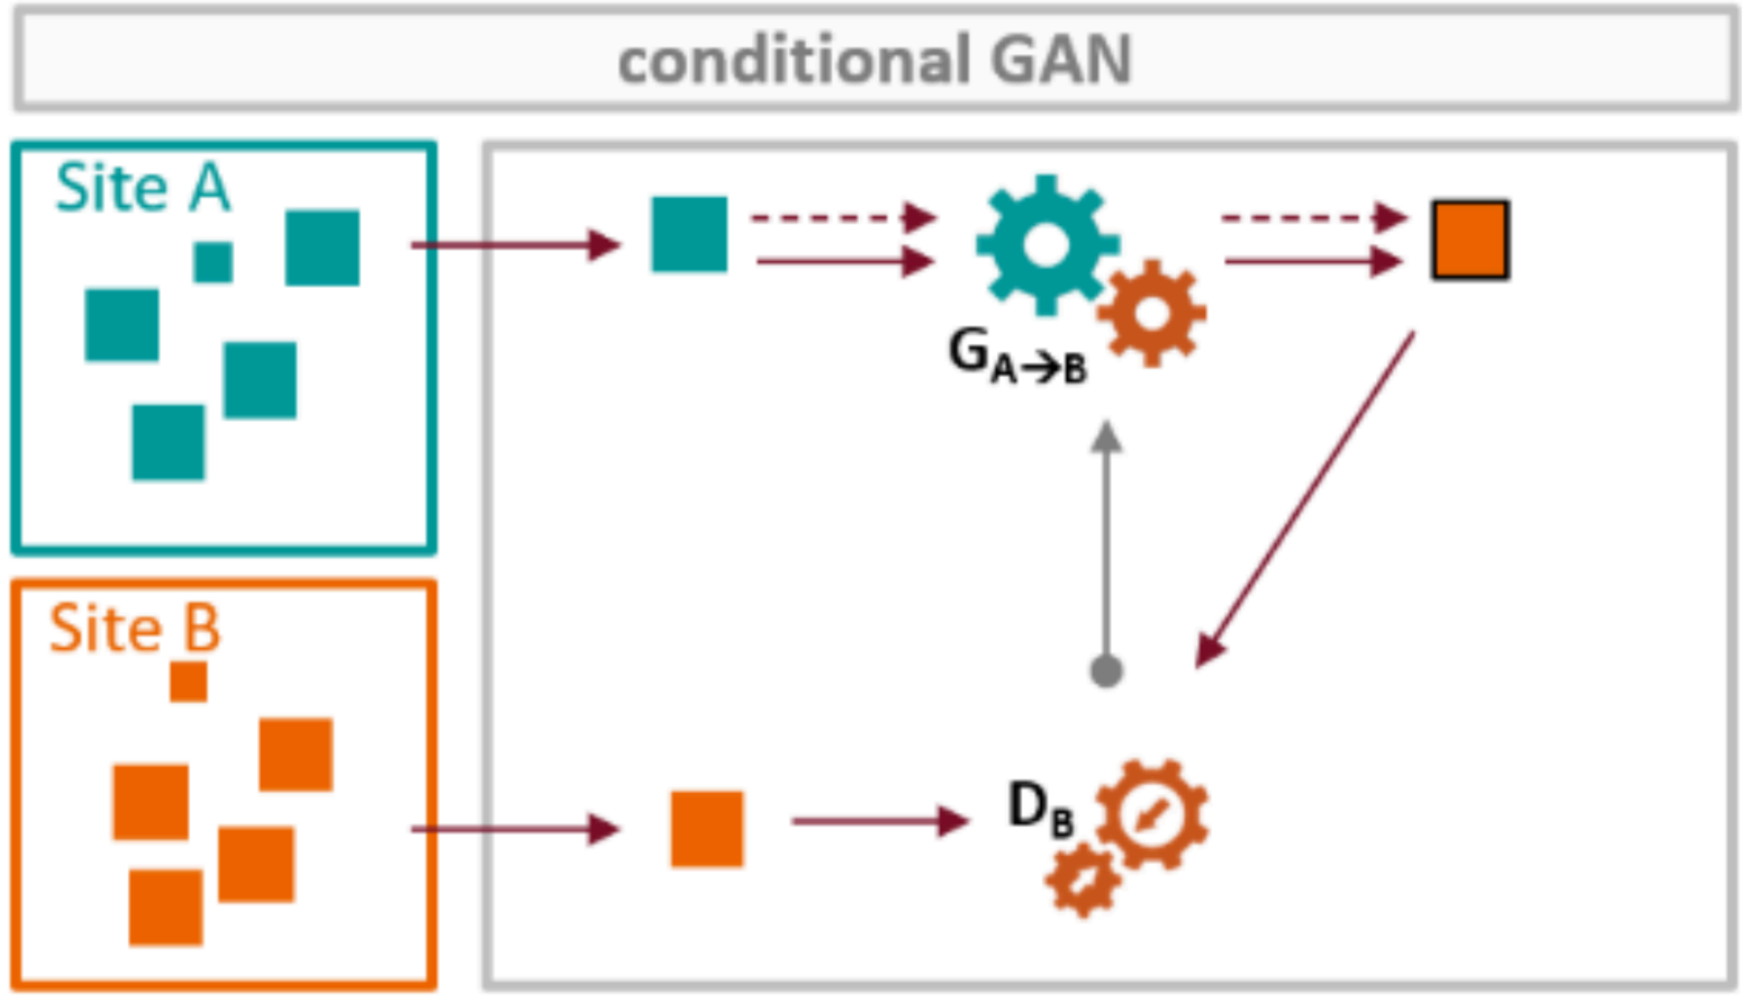

$$\mathcal{L}_G = \lambda_{L1} \sum_i^n \frac{|\hat{y}_i - y_i|}{n} + \lambda_{L2} \sum_i^n \frac{(\hat{y}_i - y_i)^2}{n} + \lambda_{LPIPS} LPIPS(\hat{\mathbf{y}}, \mathbf{y}) + \mathcal{L}_{adv}$$

	Protocol 1	Protocol 2
TR/TE/TI [ms]	2300/2.98/900	1930/2.36/972
Resolution [mm^3]	1x1x1.1	0.87x0.87x0.9
Flip Angle [°]	9	8
Pixel readout bandwidth [Hz/ms]	240	200

N=74 paired (64 TR, 3x10 TE)

Model	MAE ↓	MSE ↓	WD ↓	SSIM ↑	PSNR ↑
baseline	128±18.0	18.56±4.80x10^3	127±18.1	0.75±0.04	14.5±1.5
cGAN (L1 + LPIPS)	**52.4±13.0**	**3.84±1.70x10^3**	**49.1±14.9**	0.85±0.04	21.6±2.7

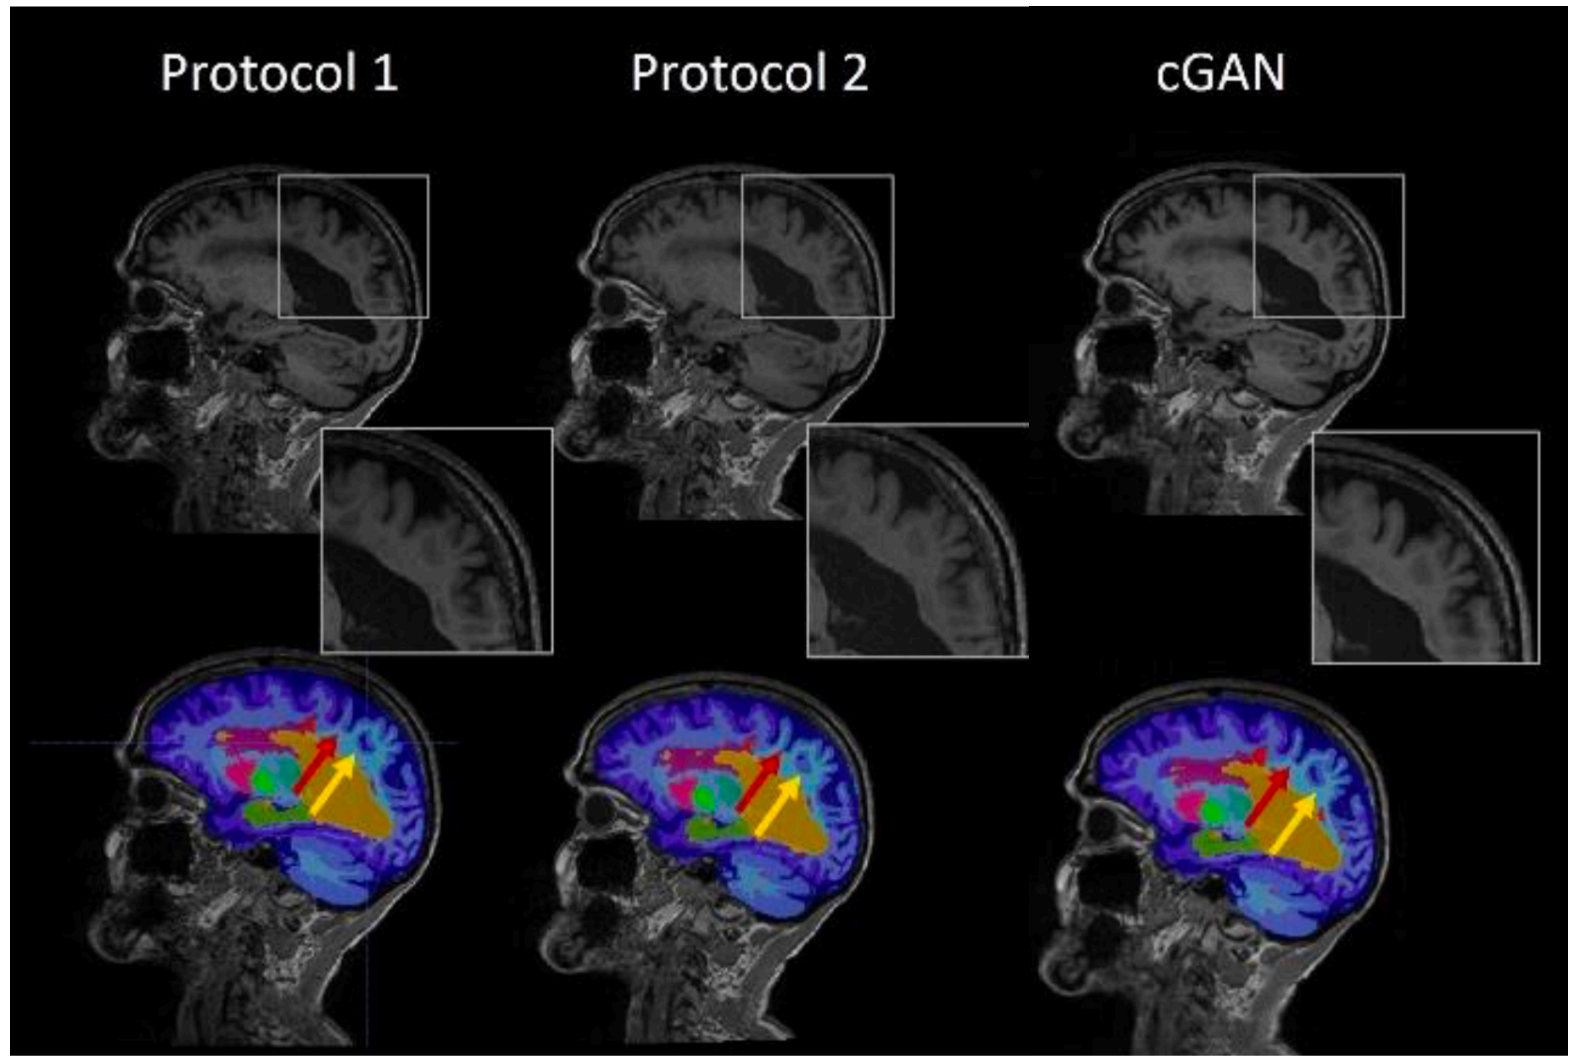

[Ravano et al., ISMRM 2022] [Ravano et al., sub.]

Centralised harmonization

Domain adaptation with image-to-image translation

CGAN [Isola et al., 2017]

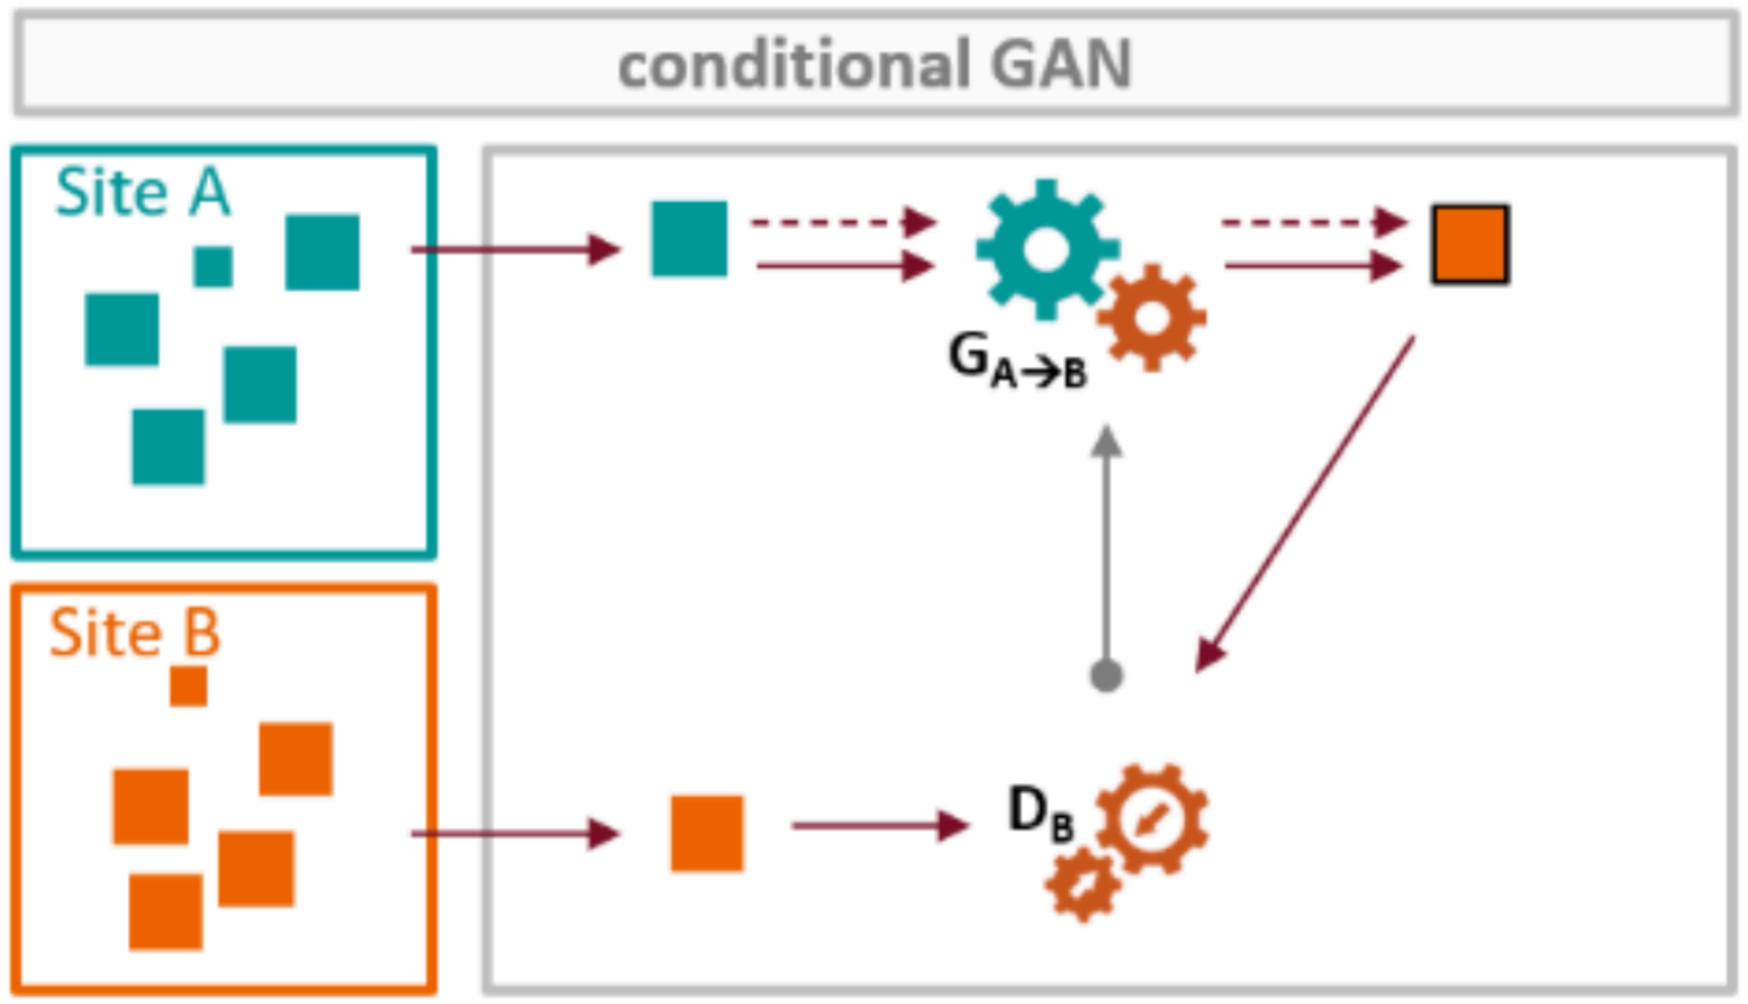

$$\mathcal{L}_G = \lambda_{L1} \sum_i^n \frac{|\hat{y}_i - y_i|}{n} + \lambda_{L2} \sum_i^n \frac{(\hat{y}_i - y_i)^2}{n} + \lambda_{LPIPS} LPIPS(\hat{\boldsymbol{y}}, \boldsymbol{y}) + \mathcal{L}_{adv}$$

	Protocol 1	Protocol 2
TR/TE/TI [ms]	2300/2.98/900	1930/2.36/972
Resolution [mm^3]	1x1x1.1	0.87x0.87x0.9
Flip Angle [°]	9	8
Pixel readout bandwidth [Hz/ms]	240	200

N=74 paired (64 TR, 3x10 TE)

Model	MAE ↓	MSE ↓	WD ↓	SSIM ↑	PSNR ↑
baseline	128±18.0	18.56±4.80x10^3	127±18.1	0.75±0.04	14.5±1.5
cGAN (L1 + LPIPS)	**52.4±13.0**	**3.84±1.70x10^3**	**49.1±14.9**	0.85±0.04	21.6±2.7

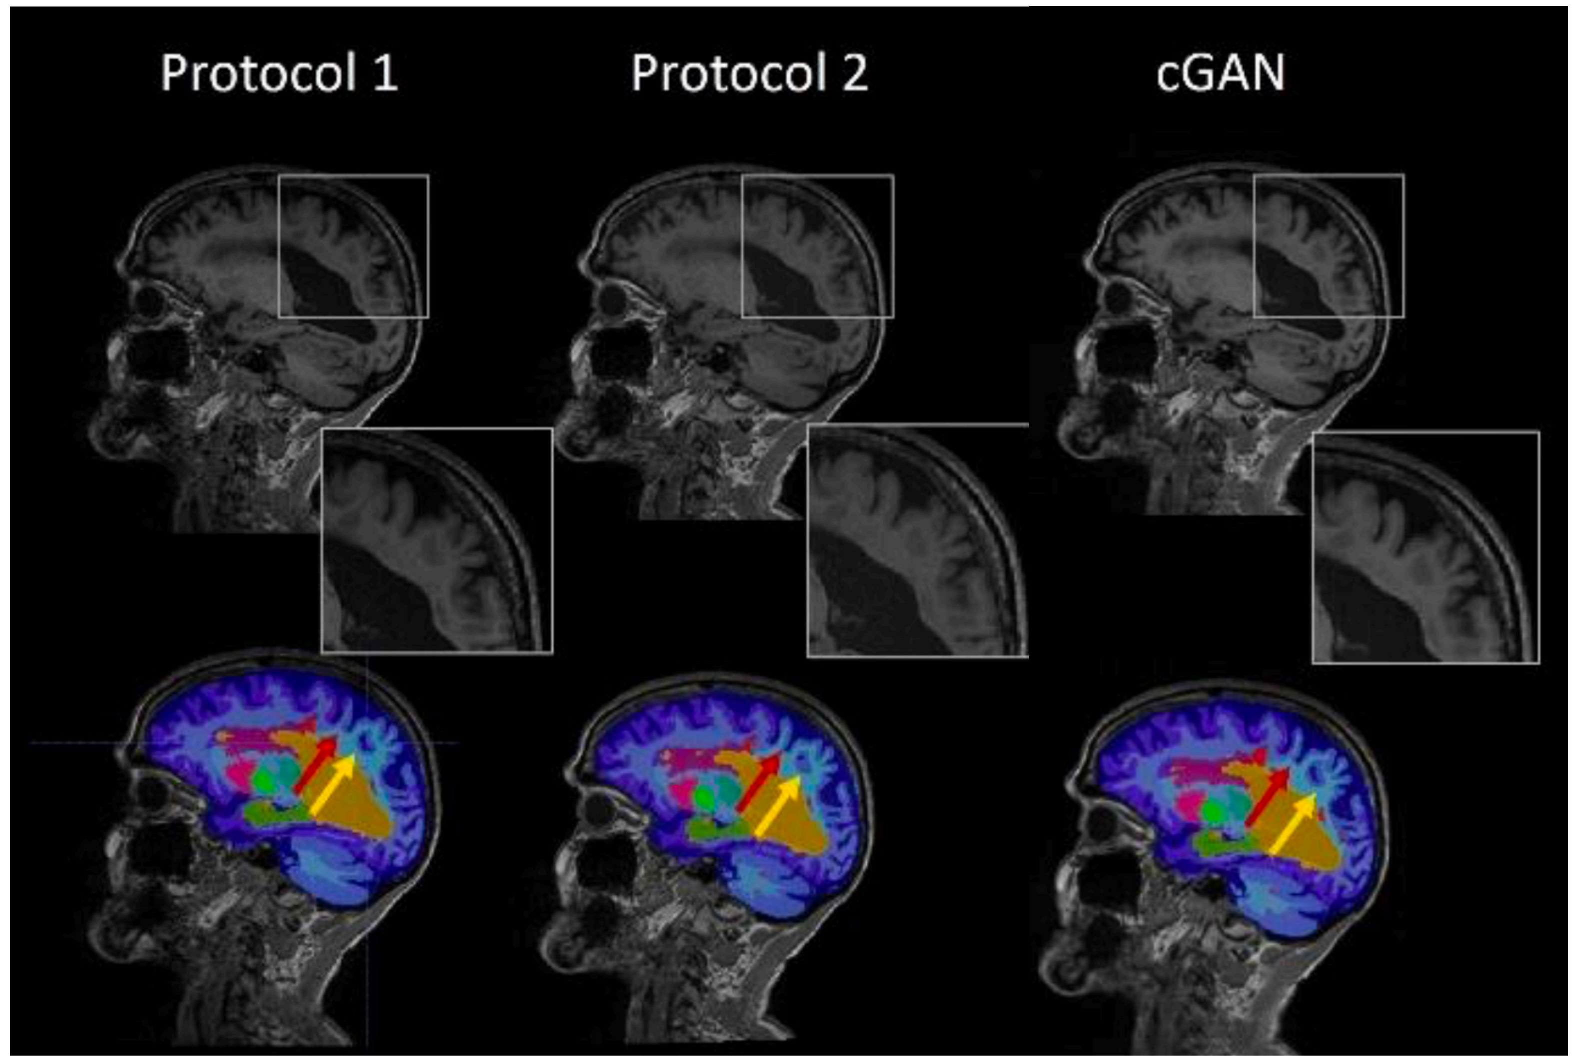

Error reduction:
Hippocampus: 5%→3% (n.s.)
Caudate: 6%→3%
Deep white matter: 10%→6%

[Ravano et al., ISMRM 2022] [Ravano et al., sub.]

Federated case

Stroke infarct core location with heterogeneous diffusion data

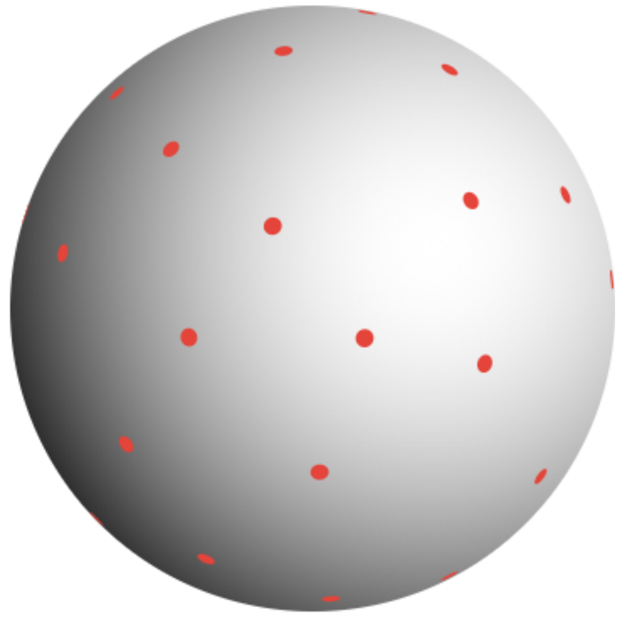

21-directions scheme
(precise but slow)

Diffusion imaging - sensitive to water diffusion direction in tissue

Dr J. Patiño-Lopez

Stroke infarct core location with heterogeneous diffusion data

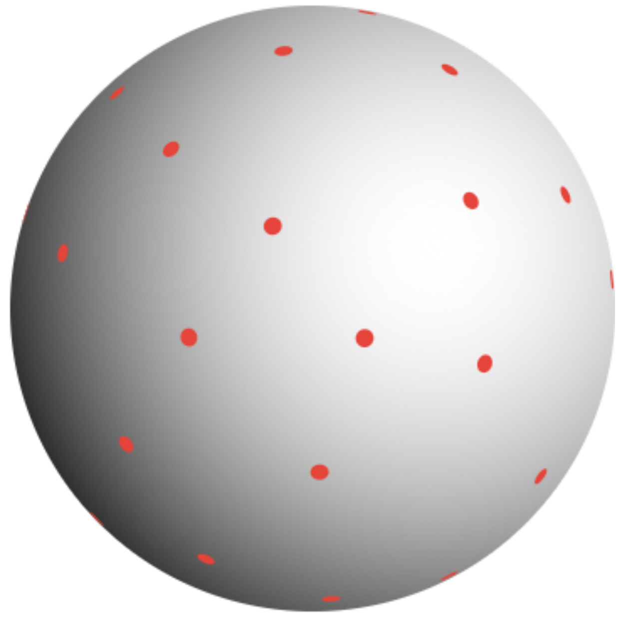

21-directions scheme
(precise but slow)

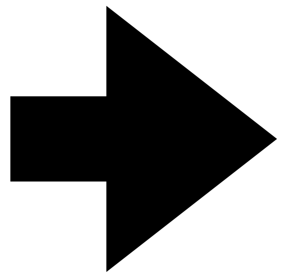

summarise

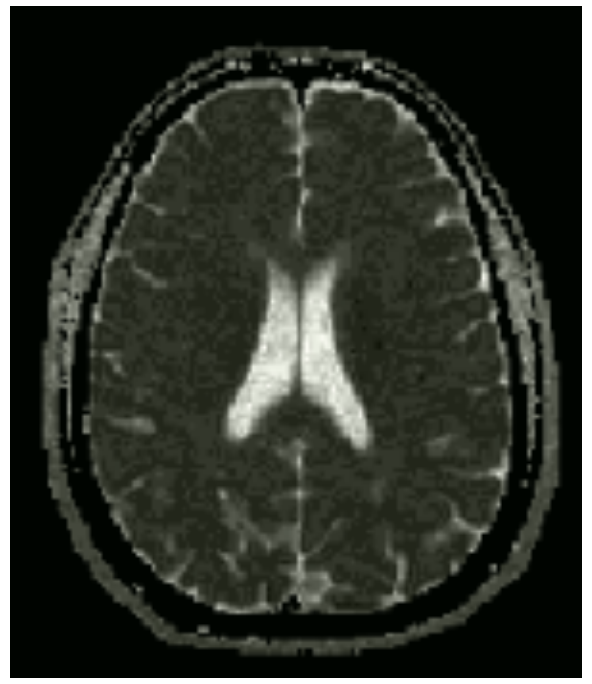

Apparent Diffusion Coefficient (ADC) map

Diffusion imaging - sensitive to
water diffusion direction in tissue

Dr J. Patiño-Lopez

Stroke infarct core location with heterogeneous diffusion data

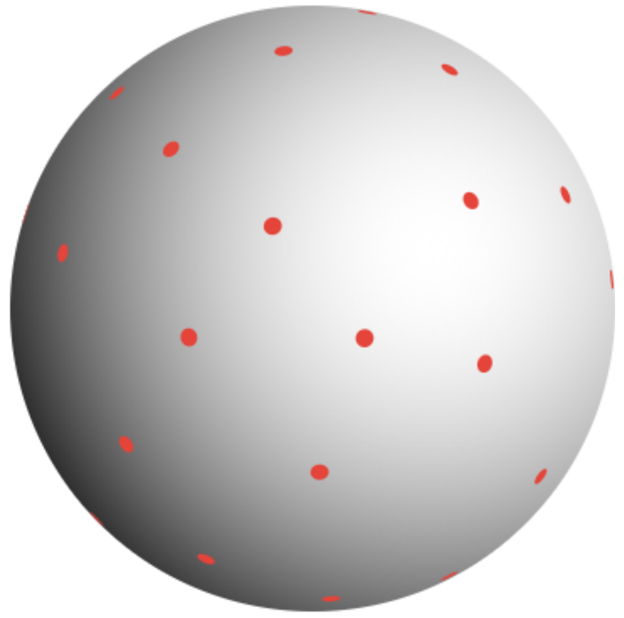

21-directions scheme
(precise but slow)

summarise

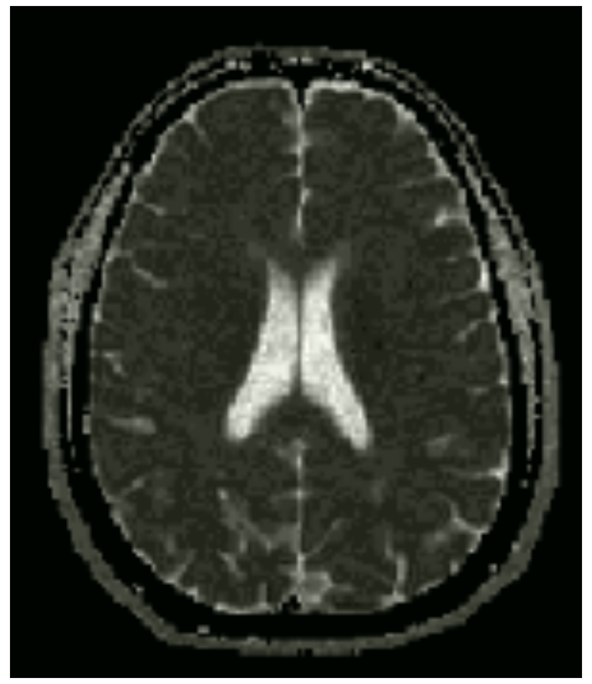

Apparent Diffusion Coefficient (ADC) map

segment

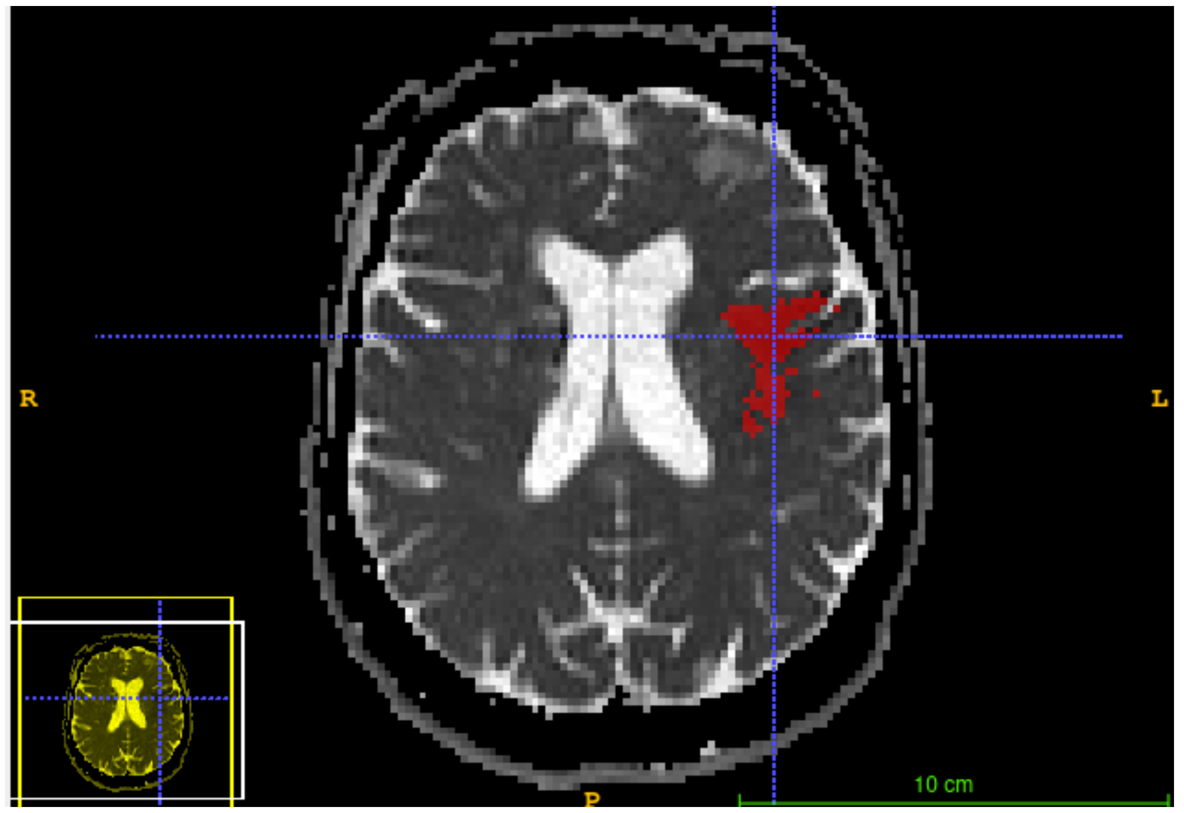

Diffusion imaging - sensitive to
water diffusion direction in tissue

Dr J. Patiño-Lopez

Stroke infarct core location with heterogeneous diffusion data

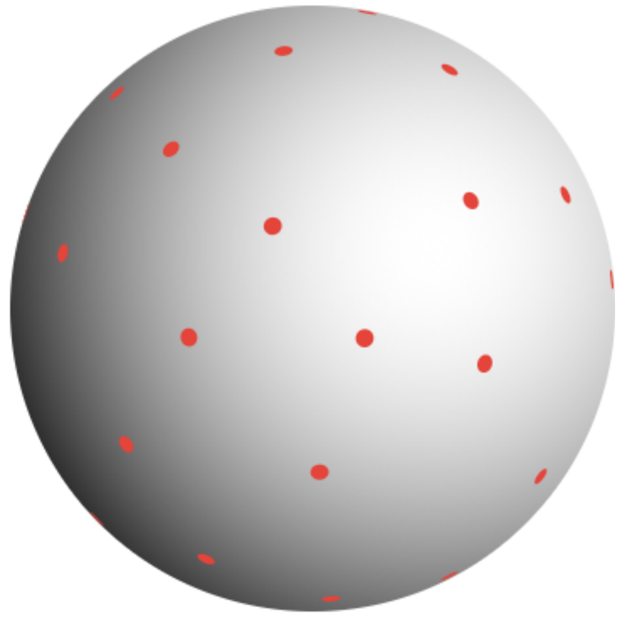

21-directions scheme
(precise but slow)

summarise

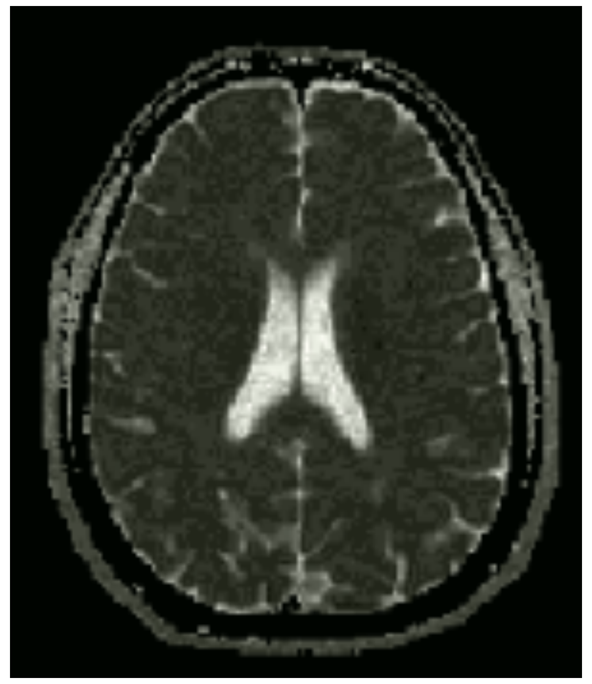

Apparent Diffusion Coefficient (ADC) map

segment

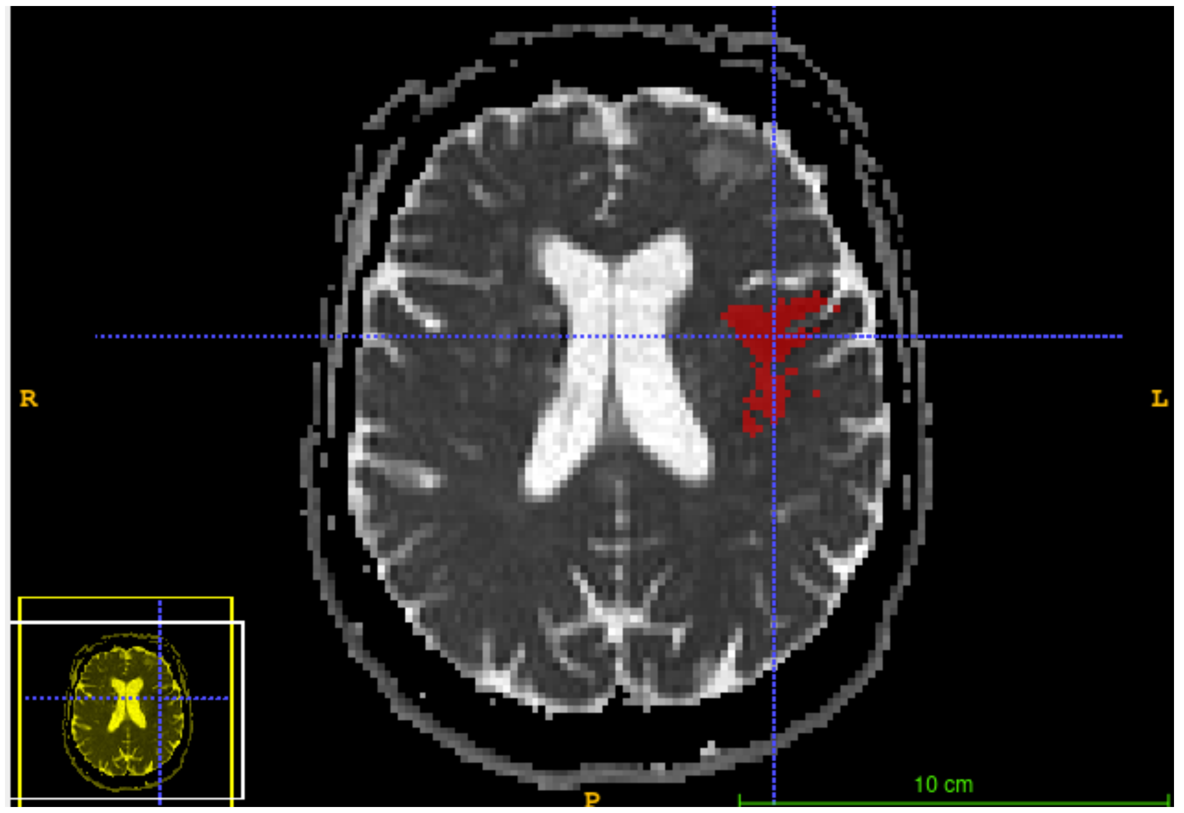

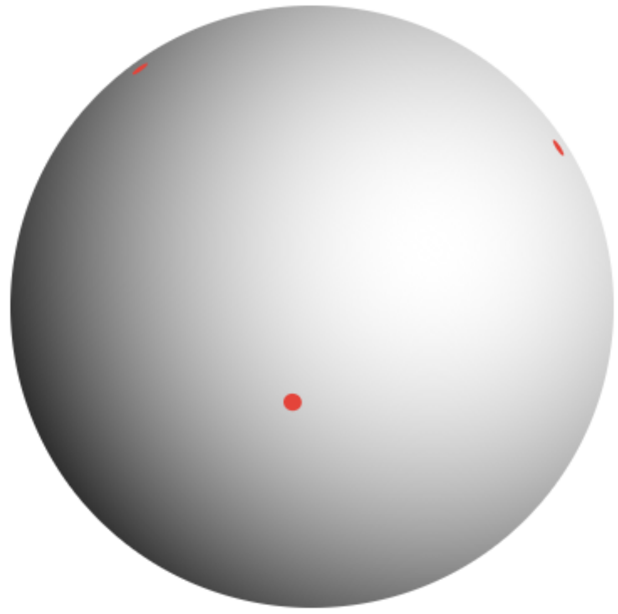

3-directions scheme
(coarse but fast)

Diffusion imaging - sensitive to
water diffusion direction in tissue

Dr J. Patiño-Lopez

Stroke infarct core location with heterogeneous diffusion data

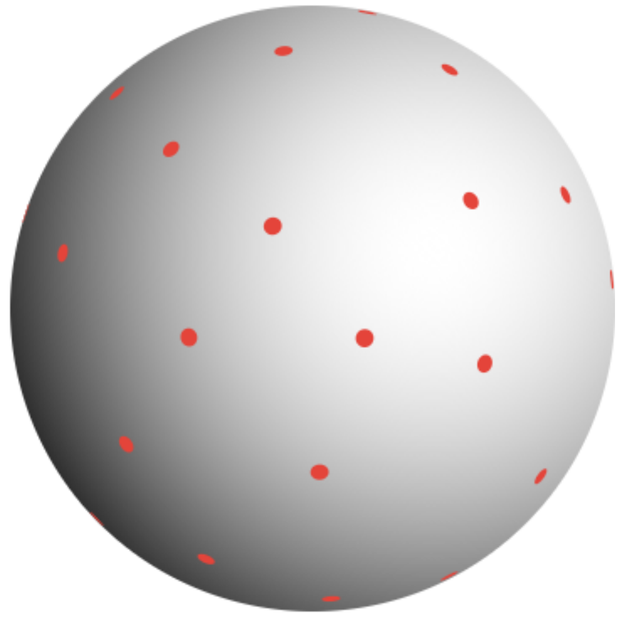

21-directions scheme
(precise but slow)

summarise

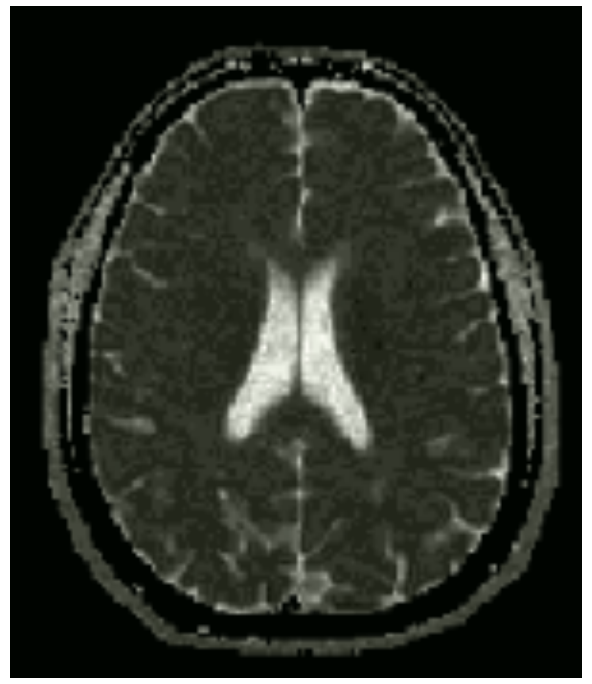

Apparent Diffusion Coefficient (ADC) map

segment

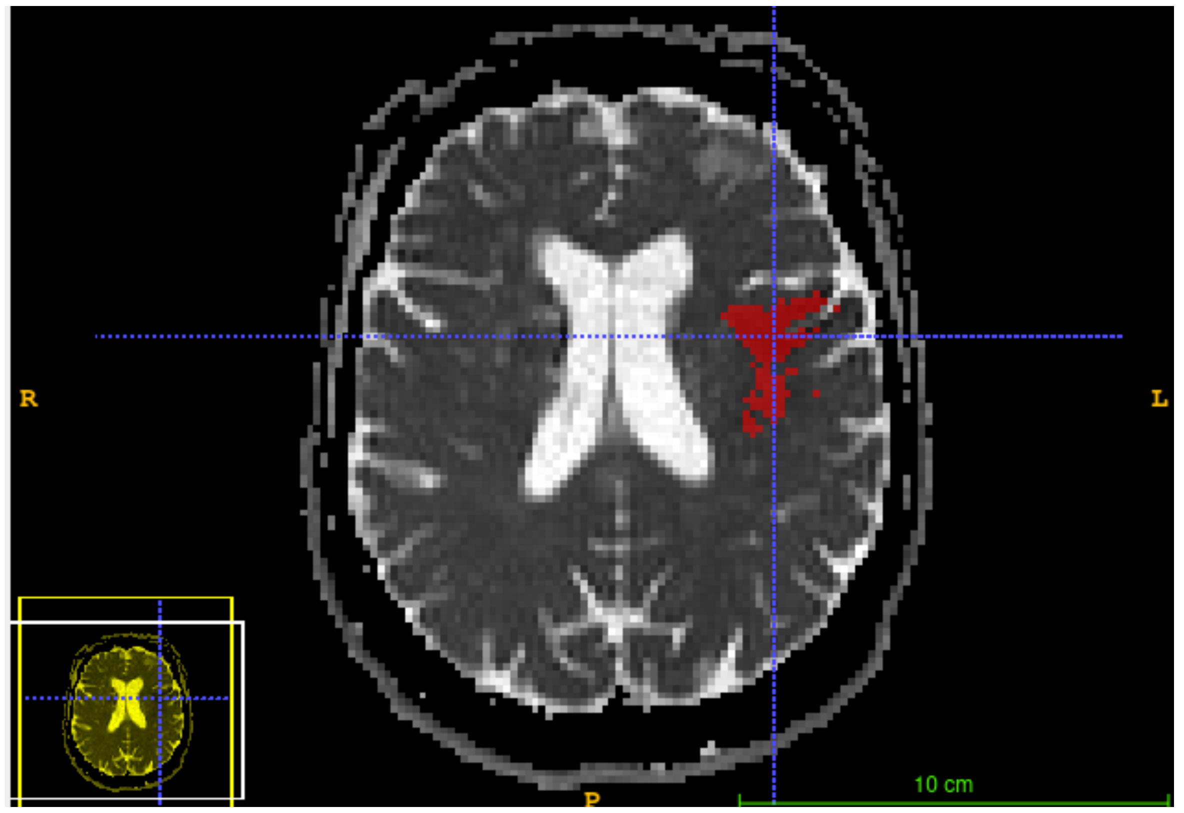

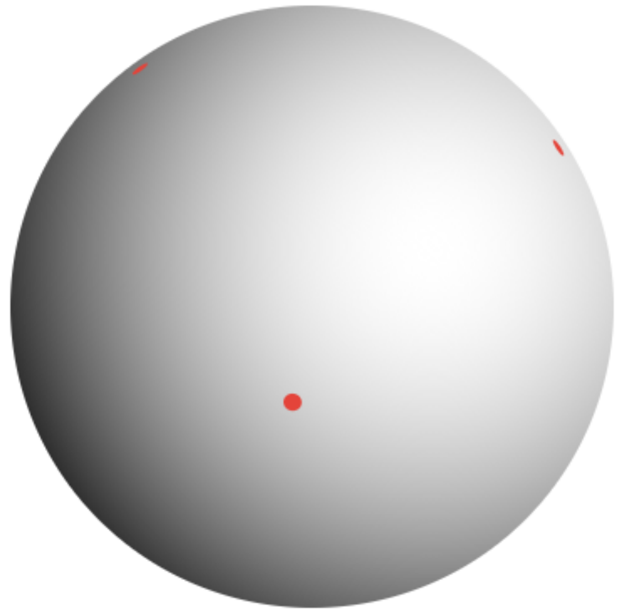

3-directions scheme
(coarse but fast)

summarise

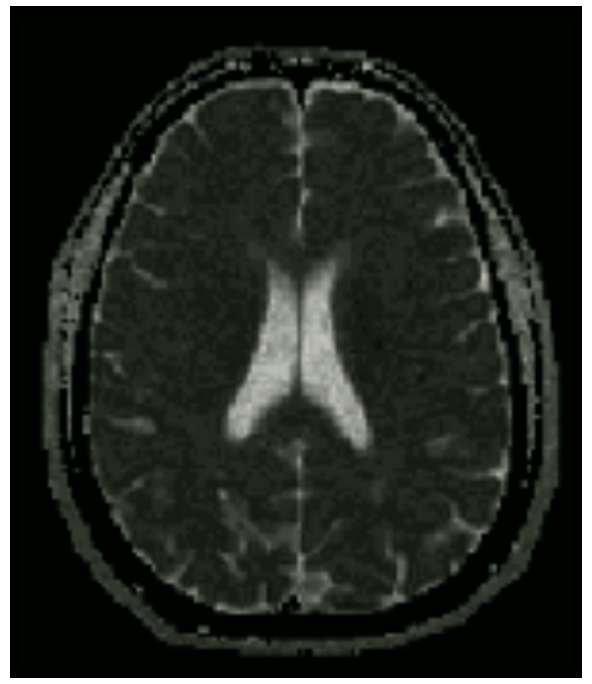

Apparent Diffusion Coefficient (ADC) map

Diffusion imaging - sensitive to
water diffusion direction in tissue

Dr J. Patiño-Lopez

Stroke infarct core location with heterogeneous diffusion data

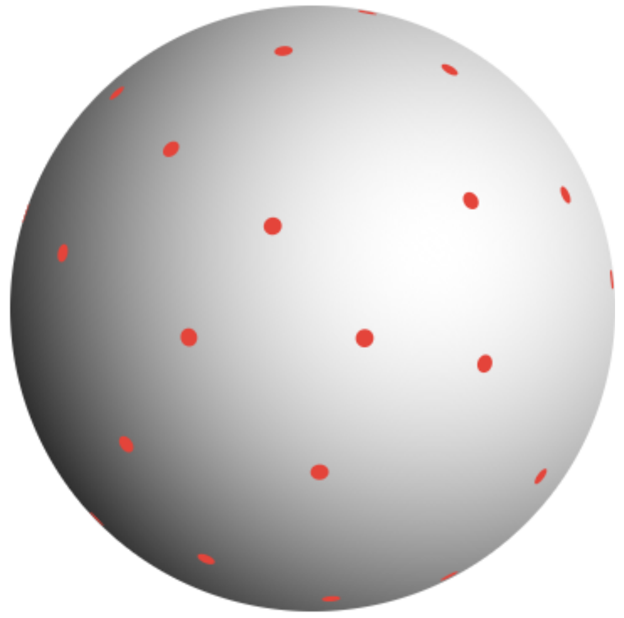

21-directions scheme
(precise but slow)

summarise

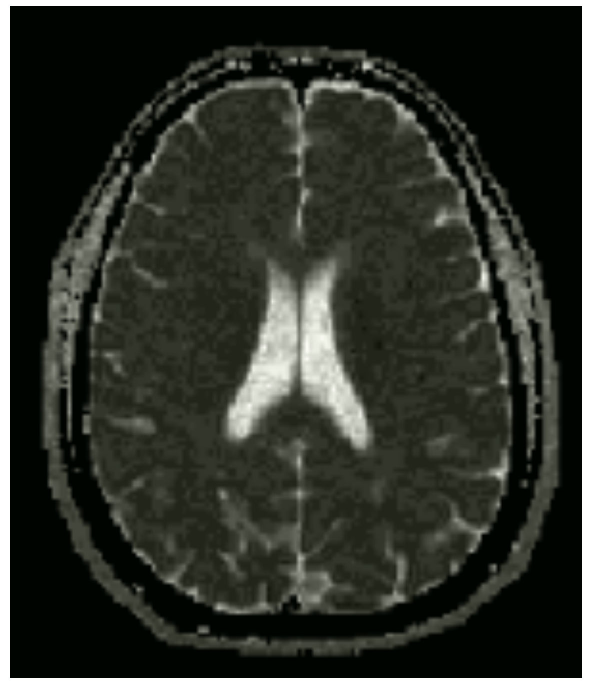

Apparent Diffusion Coefficient (ADC) map

segment

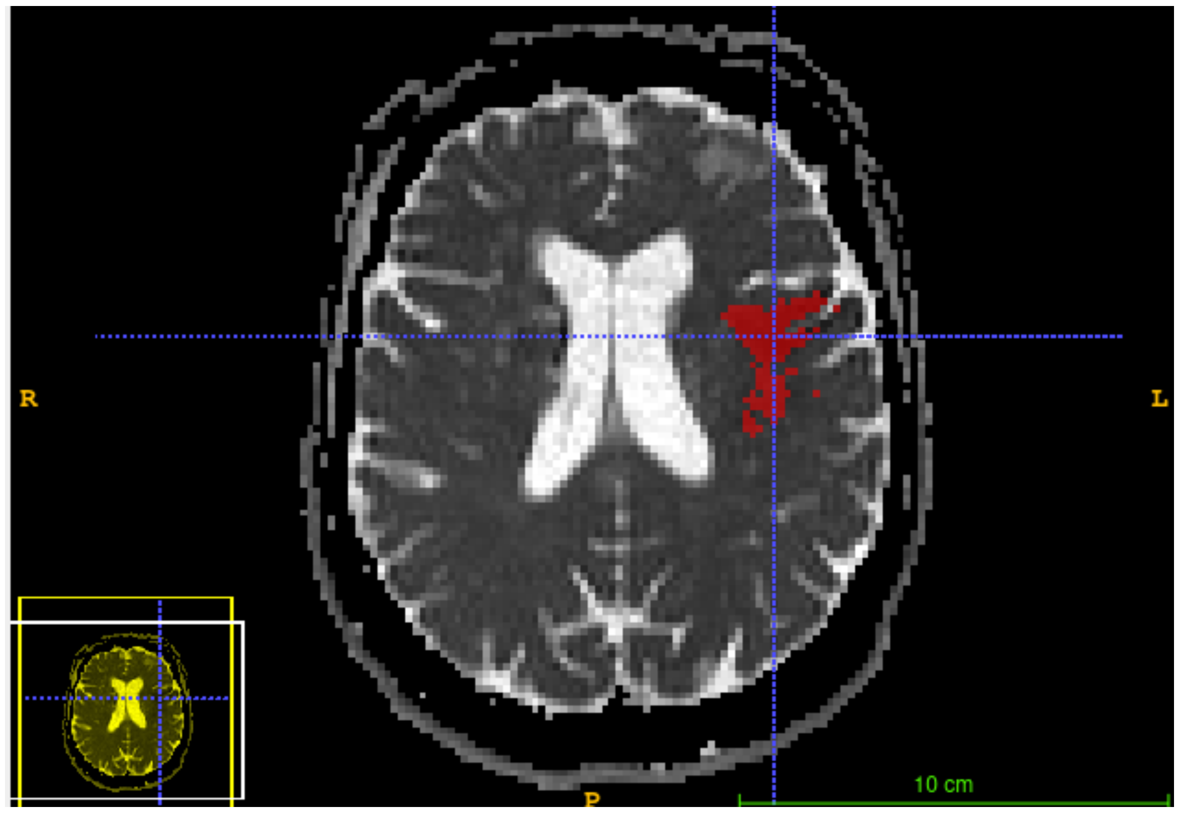

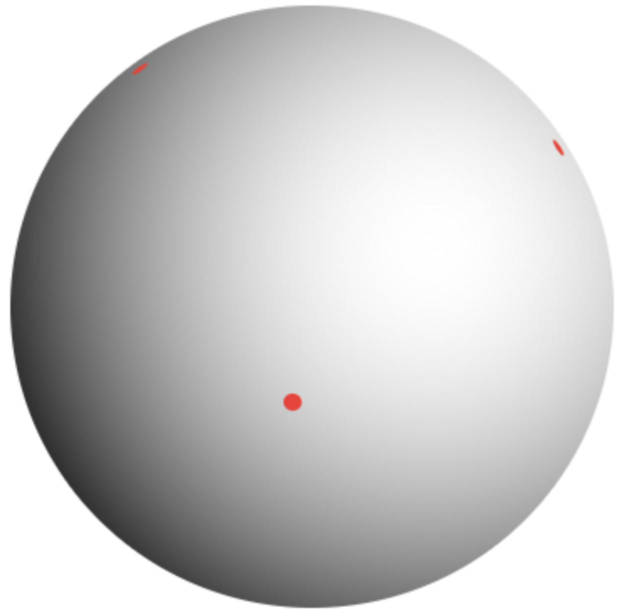

3-directions scheme
(coarse but fast)

summarise

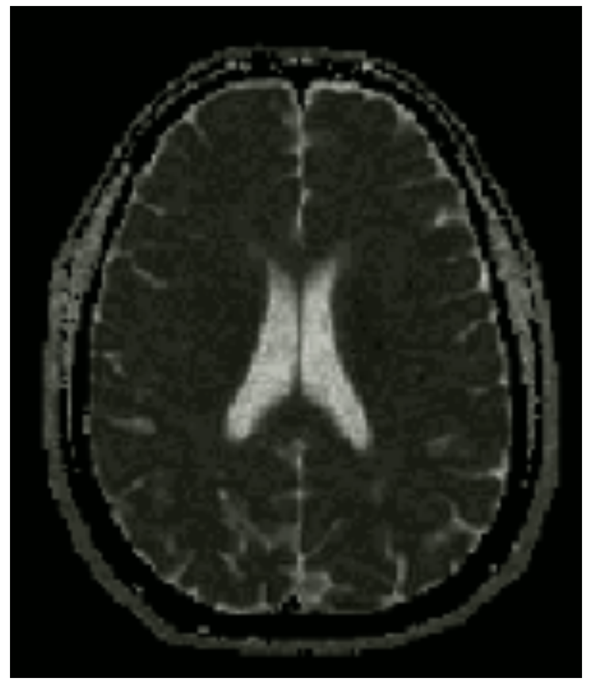

Apparent Diffusion Coefficient (ADC) map

segment

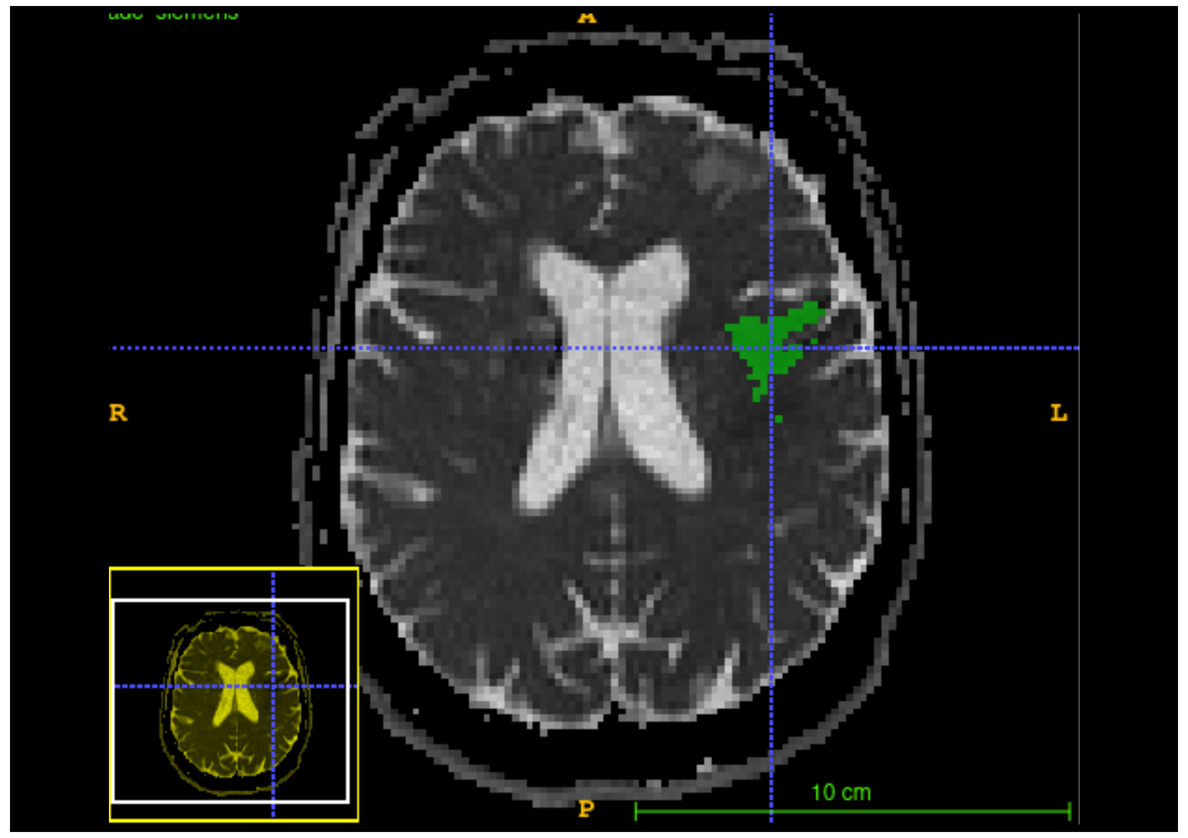

Diffusion imaging - sensitive to
water diffusion direction in tissue

Dr J. Patiño-Lopez

Diffusion protocol heterogeneity effects

1 patient,
14 different acquisition schemes

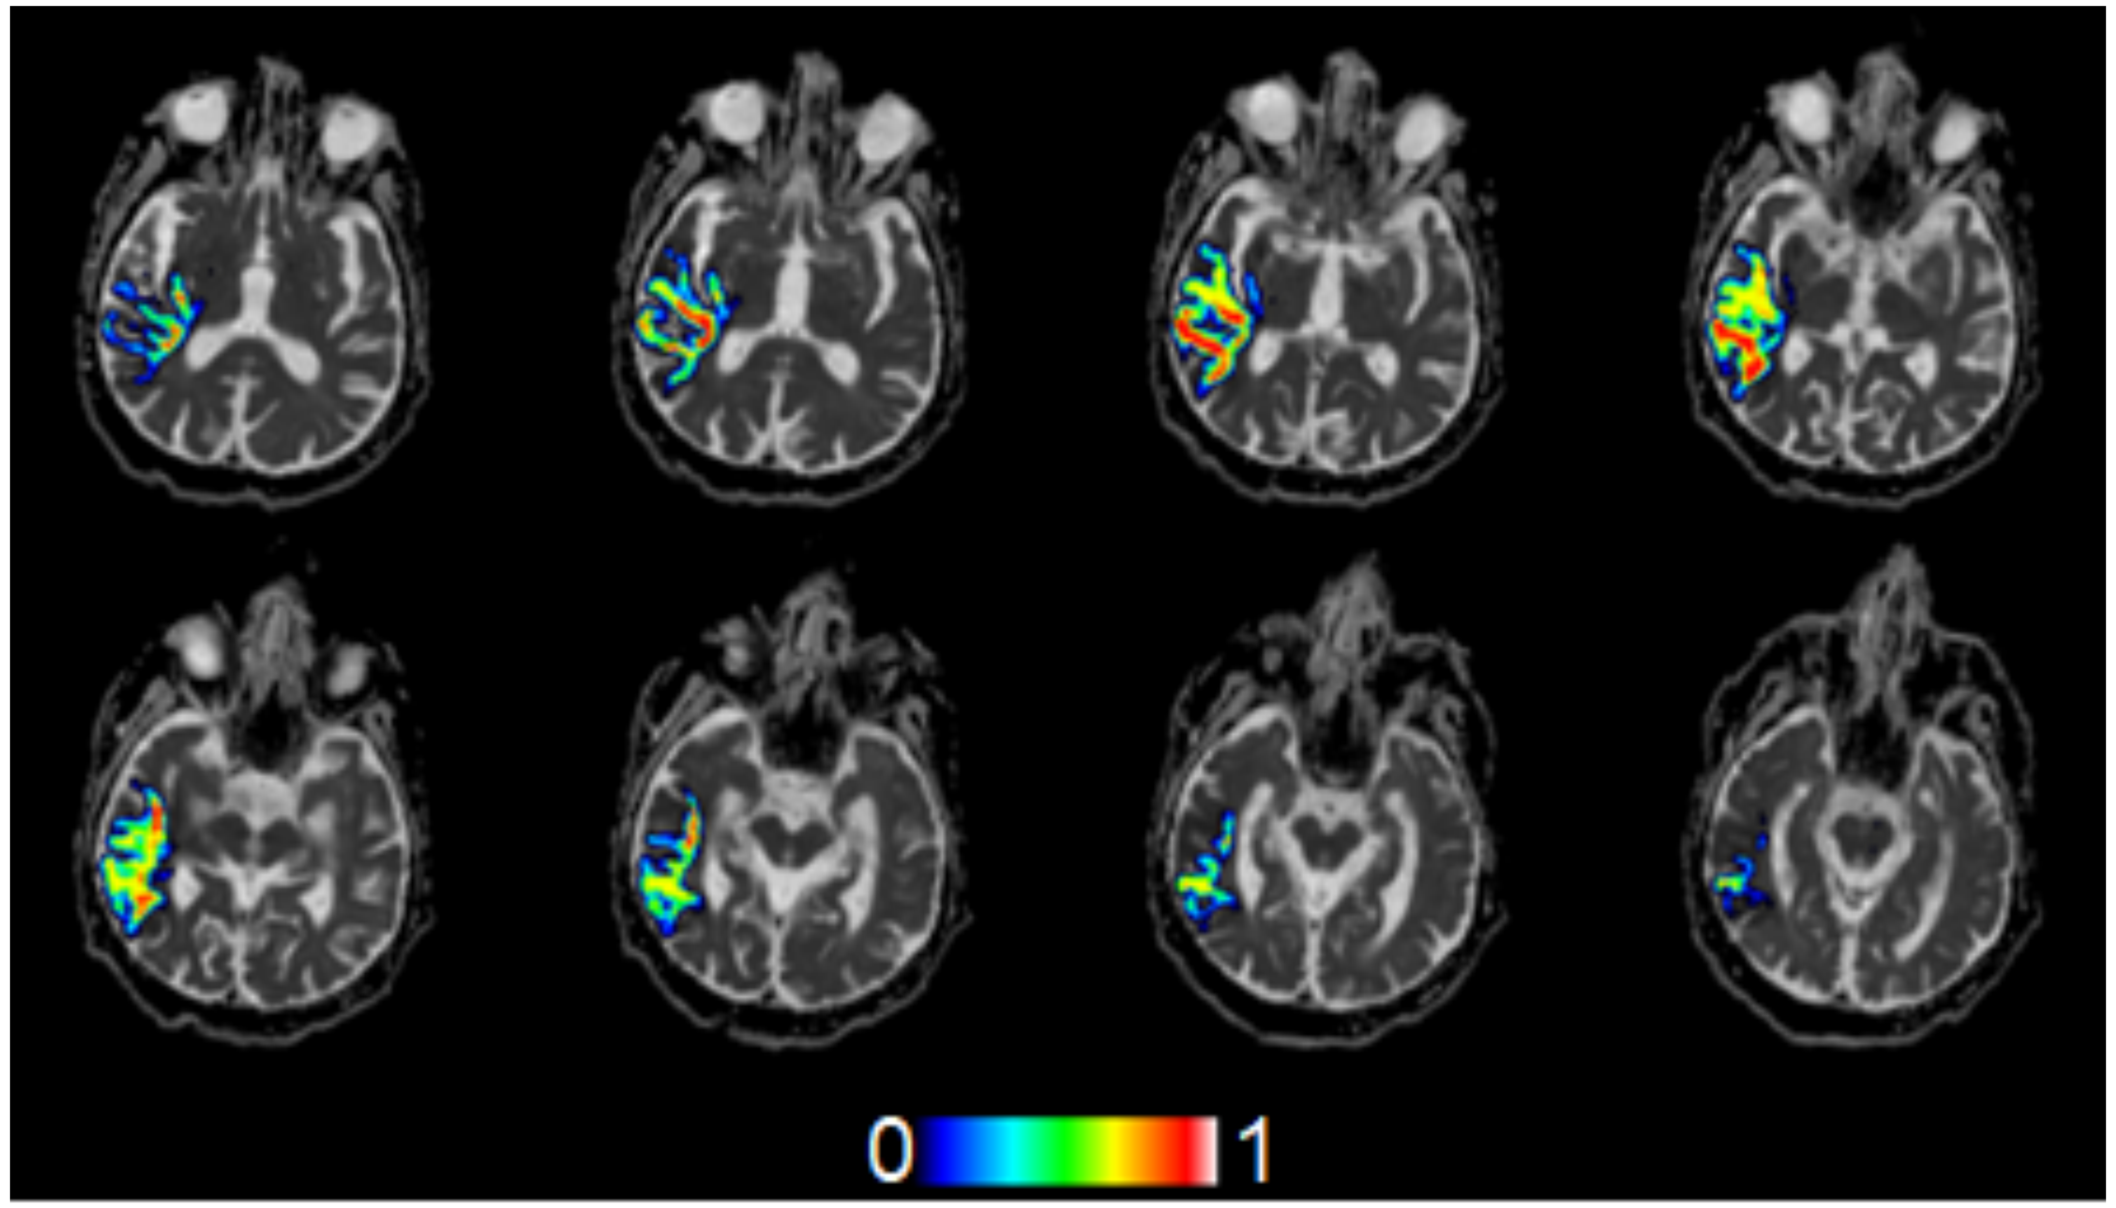

per-voxel percent agreement
between schemes

[Rafael-Patiño et al., ISMRM 2022]

Diffusion protocol heterogeneity effects

1 patient,
14 different acquisition schemes

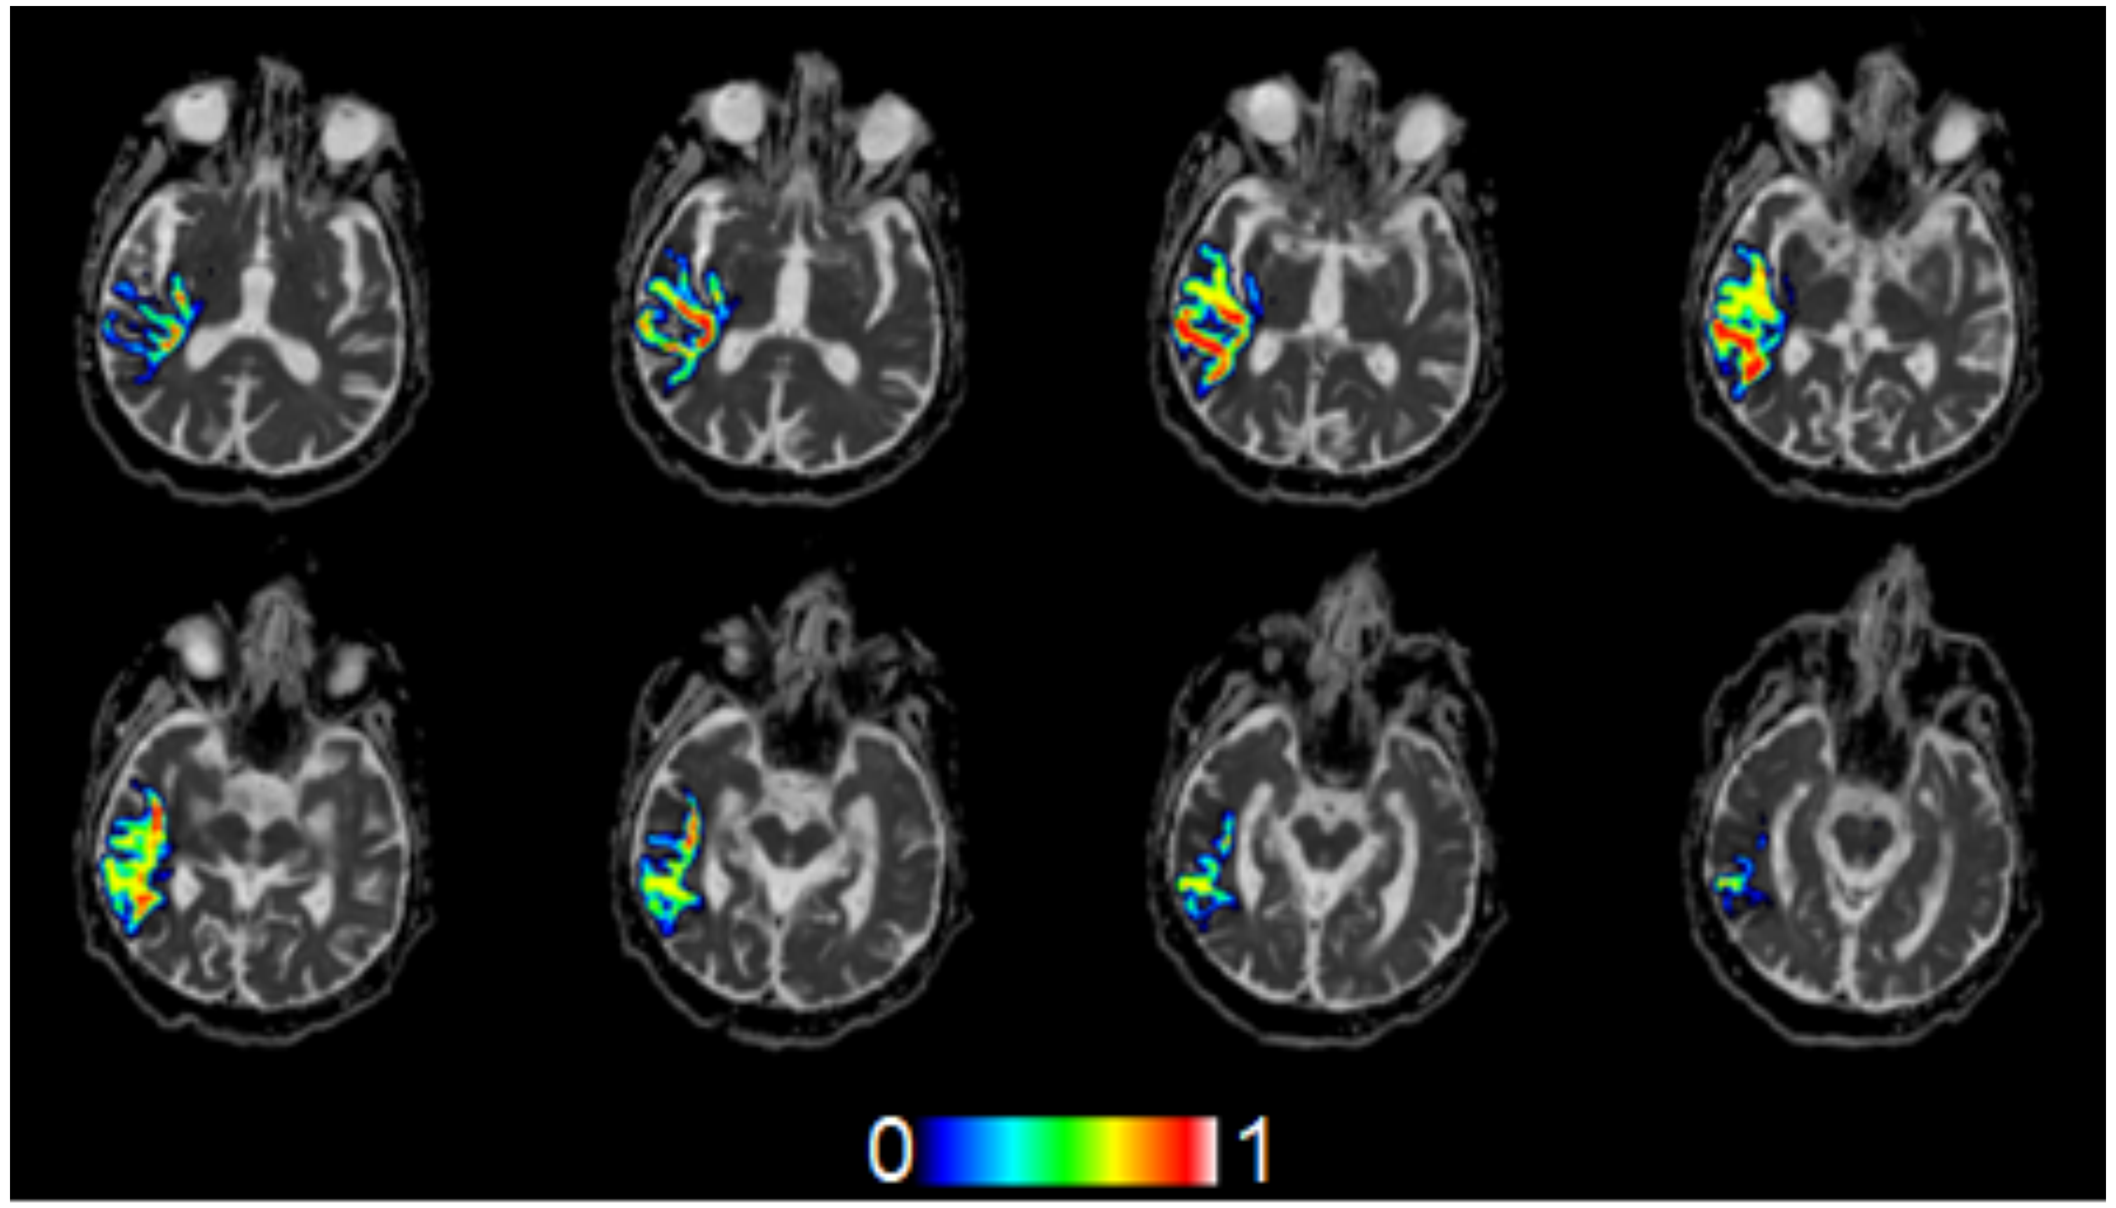

per-voxel percent agreement
between schemes

29 patients: Jaccard coefficient ~ 50% (range 5%-85%)

[Rafael-Patiño et al., ISMRM 2022]

Diffusion protocol heterogeneity effects

1 patient,
14 different acquisition schemes

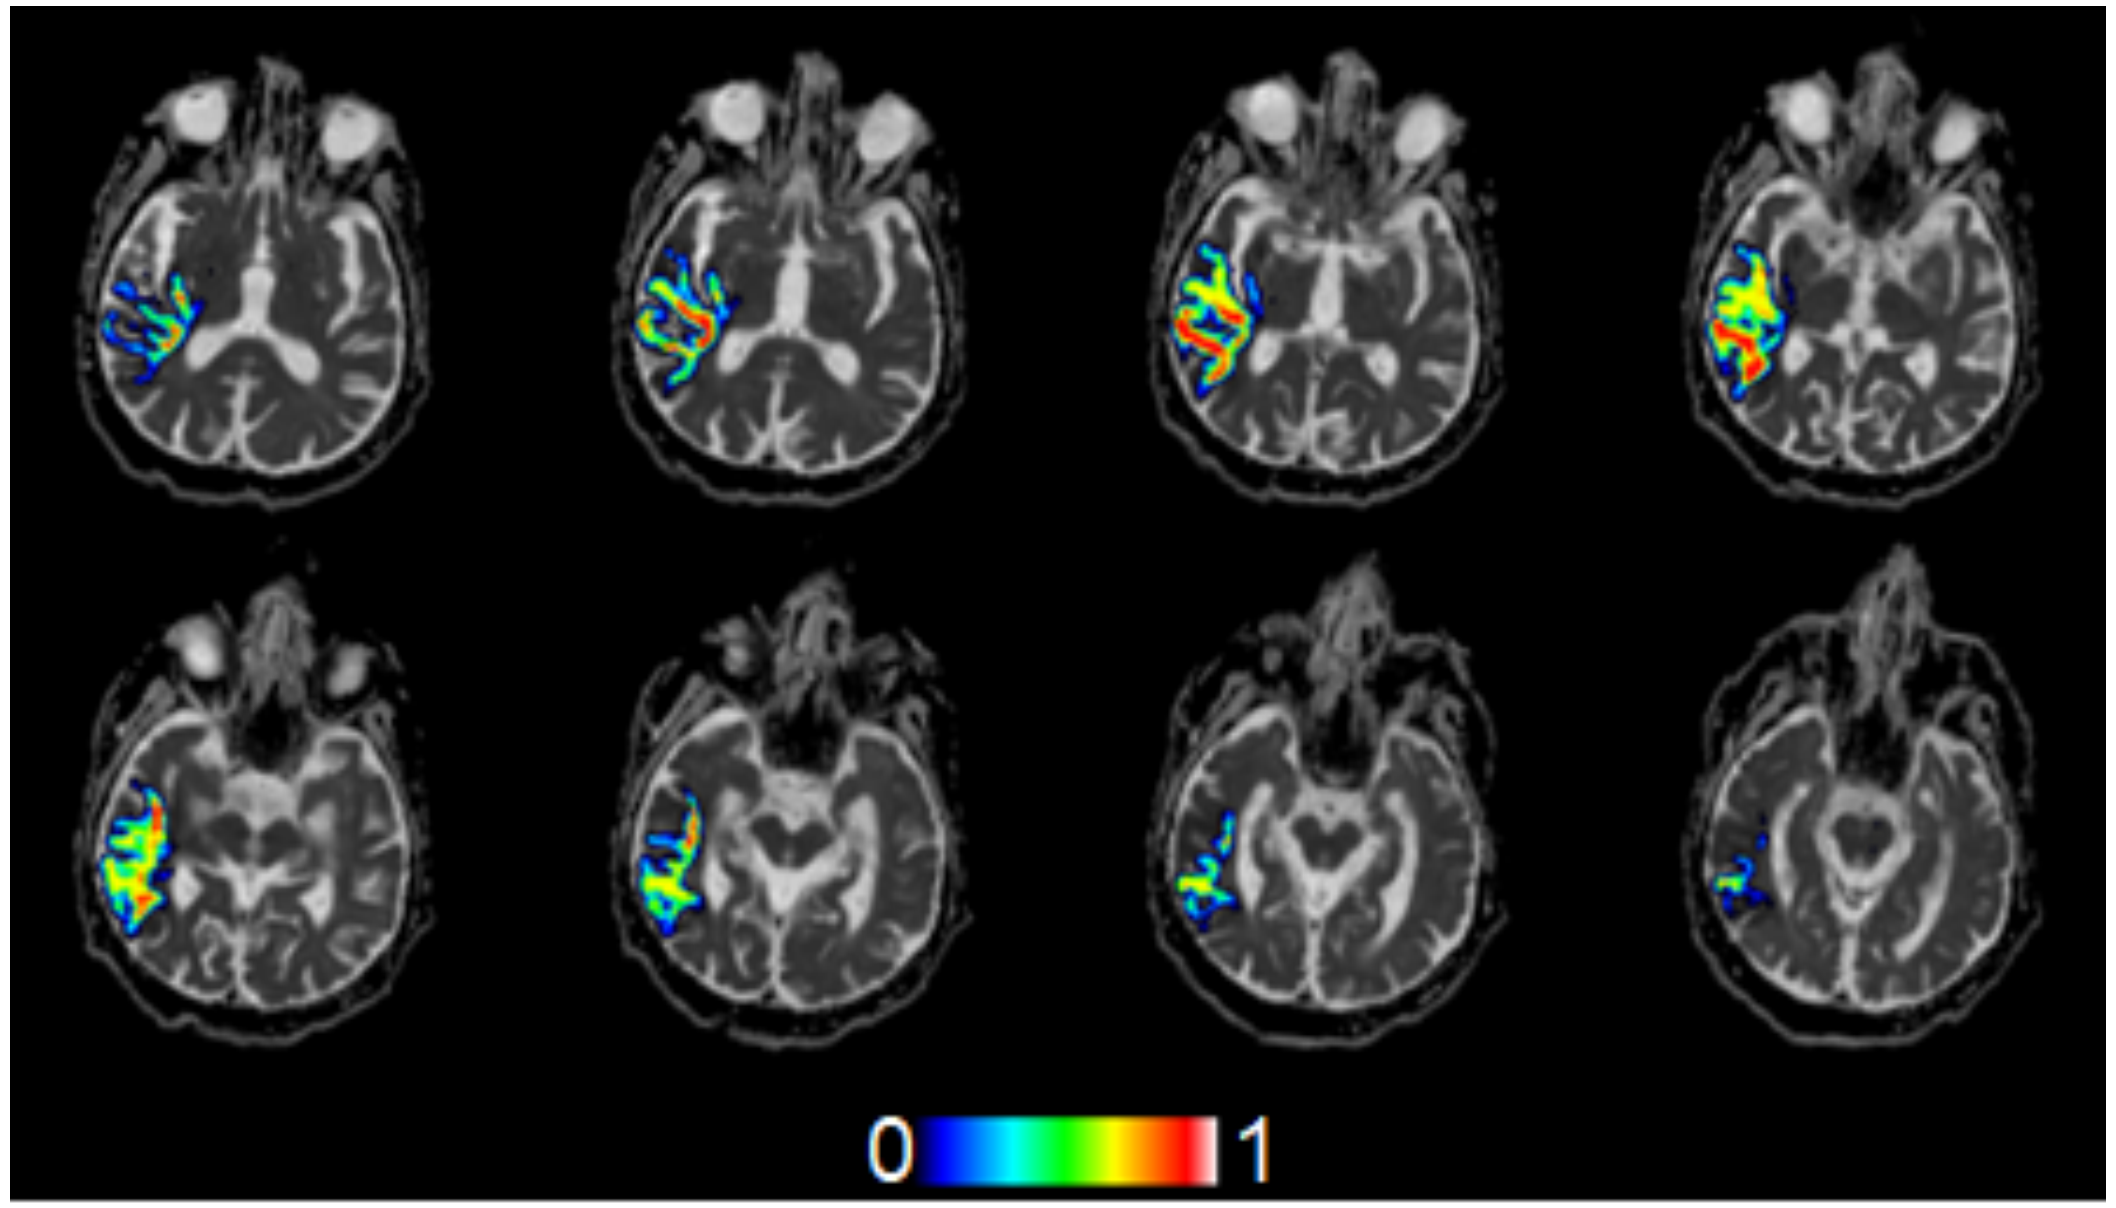

per-voxel percent agreement
between schemes

29 patients: Jaccard
coefficient ~ 50%
(range 5%-85%)

low agreement on
infarct core location

[Rafael-Patiño et al., ISMRM 2022]

Federated learning with heterogeneous data

Perfusion-based stroke segmentation across multiple hardware vendors

Federated average[1]

Initialize W_G
for each round $t = 1, 2, \ldots$ **do**
    Send W_G to each client
    for each client $k = 1, 2, \ldots, K$ **do**
        $w^k_{t+1} \leftarrow \mathrm{ClientUpdate}(k, W_G)$
    end for
    $W_G \leftarrow \frac{1}{Z} \sum_{k=1}^{K} f(n_k) w^k_{t+1}$
end for

Dr S. Otalora

Dr R. McKinley

[Otalora, Rafael-Patiño et al, sub.] [1][McMahan et al., AISTATS 2017] β-weighting [Cui et al CVPR 2019]

Many other approaches like SCAFFOLD [Karimireddy et al ICML 2020], FED-ROD [Chen et al ICLR 2022], FEDYOGI [Reddi et al ICLR 2021]…

Federated learning with heterogeneous data

Perfusion-based stroke segmentation across multiple hardware vendors

Federated average[1]

Initialize W_G
for each round $t = 1, 2, \ldots$ **do**
Send W_G to each client
for each client $k = 1, 2, \ldots, K$ **do**
$w^k_{t+1} \leftarrow \text{ClientUpdate}(k, W_G)$
end for
$W_G \leftarrow \frac{1}{Z} \sum_{k=1}^{K} \boxed{f(n_k)} w^k_{t+1}$
end for

Site-specific weight

Dr S. Otalora

Dr R. McKinley

[Otalora, Rafael-Patiño et al, sub.] [1][McMahan et al., AISTATS 2017] β-weighting [Cui et al CVPR 2019]
Many other approaches like SCAFFOLD [Karimireddy et al ICML 2020], FED-ROD [Chen et al ICLR 2022], FEDYOGI [Reddi et al ICLR 2021]…

Federated learning with heterogeneous data

Perfusion-based stroke segmentation across multiple hardware vendors

Federated average[1]

Initialize W_G
for each round $t = 1, 2, \ldots$ **do**
  Send W_G to each client
  for each client $k = 1, 2, \ldots, K$ **do**
    $w_{t+1}^{k} \leftarrow \mathrm{ClientUpdate}(k, W_G)$
  end for
  $W_G \leftarrow \frac{1}{Z} \sum_{k=1}^{K} \boxed{f(n_k)} w_{t+1}^{k}$
end for

Site-specific weight

Original formulation: n_k/N → favor majority site
Our approach: $(1-\beta)/(1-\beta^{n_k})$ → favor minority site


Dr S. Otalora


Dr R. McKinley

[Otalora, Rafael-Patiño et al, sub.]   [1][McMahan et al., AISTATS 2017]   β-weighting [Cui et al CVPR 2019]

Many other approaches like SCAFFOLD [Karimireddy et al ICML 2020], FED-ROD [Chen et al ICLR 2022], FEDYOGI [Reddi et al ICLR 2021]…

Federated learning with heterogeneous data

Perfusion-based stroke segmentation across multiple hardware vendors

Federated average[1]

Initialize W_G
for each round $t = 1, 2, \ldots$ **do**
  Send W_G to each client
  for each client $k = 1, 2, \ldots, K$ **do**
    $w_{t+1}^{k} \leftarrow \text{ClientUpdate}(k, W_G)$
  end for
  $W_G \leftarrow \frac{1}{Z} \sum_{k=1}^{K} \boxed{f(n_k)} w_{t+1}^{k}$
end for

Site-specific weight

Original formulation: n_k/N → favor majority site
Our approach: $(1-\beta)/(1-\beta^{n_k})$ → favor minority site

vendor A	vendor B	vendor C	vendor D
22	57	13	2

N=112 (94 TR, 18 TE)

Dr S. Otalora

Dr R. McKinley

[Otalora, Rafael-Patiño et al, sub.] [1][McMahan et al., AISTATS 2017] β-weighting [Cui et al CVPR 2019]
Many other approaches like SCAFFOLD [Karimireddy et al ICML 2020], FED-ROD [Chen et al ICLR 2022], FEDYOGI [Reddi et al ICLR 2021]…

Federated learning with heterogeneous data

Perfusion-based stroke segmentation across multiple hardware vendors

Federated average[1]

Initialize W_G
for each round $t = 1, 2, \ldots$ **do**
  Send W_G to each client
  for each client $k = 1, 2, \ldots, K$ **do**
    $w_{t+1}^k \leftarrow \mathrm{ClientUpdate}(k, W_G)$
  end for
  $W_G \leftarrow \frac{1}{Z} \sum_{k=1}^{K} \boxed{f(n_k)} w_{t+1}^k$
end for

Site-specific weight

Original formulation: n_k/N → favor majority site
Our approach: $(1-\beta)/(1-\beta^{n_k})$ → favor minority site

vendor A	vendor B	vendor C	vendor D
22	57	13	2

N=112 (94 TR, 18 TE)

	Centralized	FedAvg	β-Weighting
Dice	0.39 ± 0.01	0.29 ± 0.04	0.36 ± 0.01


Dr S. Otalora Dr R. McKinley

[Otalora, Rafael-Patiño et al, sub.] [1][McMahan et al., AISTATS 2017] β-weighting [Cui et al CVPR 2019]
Many other approaches like SCAFFOLD [Karimireddy et al ICML 2020], FED-ROD [Chen et al ICLR 2022], FEDYOGI [Reddi et al ICLR 2021]…

Wrap-up and take-home points

Heterogeneous medical imaging data…

severely increases biomarker variability for all applications

is a **fact of life for real clinical data**, also within single hospitals

can be tamed in a **centralised** fashion, although unpaired approaches are more difficult

can also be tamed in a **federated** setting using generic ML approaches, although much work remains

is a real test of model **generalisation** ability

Wrap-up and take-home points

Heterogeneous medical imaging data…

severely increases biomarker variability for all applications

is a **fact of life for real clinical data**, also within single hospitals

can be tamed in a **centralised** fashion, although unpaired approaches are more difficult

can also be tamed in a **federated** setting using generic ML approaches, although much work remains

is a real test of model **generalisation** ability

is the key to leveraging **very large retrospective datasets** (100K+) within and across hospitals

Thanks

Translational Machine Learning Lab

(back) Tommaso Di Noto**, Antoine Madrona, Xavier Sieber***, Dr. Jonathan Patiño-Lopez, Jonas Richiardi
(front) Veronica Ravano*, Dr Elda Fischi-Gomez, Costa Georgantas, Dr Jaume Banus Cobo
(and always looking for Master students…)

Collaborators

Inselspital
DrSc Richard McKinley
Prof. Roland Wiest
DrSc Sebastian Otalora
Dr Simon Jung

Siemens Healthcare
DrSc Tobias Kober
DrSc Bénédicte Maréchal

CHUV
Prof. Patrik Michel
Prof. Guillaume Saliou
Dr Steven Hajdu
Dr Silvia Pistocchi

Funding

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Swiss National Science Foundation

CHUV

co-supervisions: *Siemens Healthcare, **UNIL/MIAL, **CHUV/Connectomics Lab, ***CHUV/CVMR