Nonlinear motion separation via untrained generator networks with disentangled latent space variables and applications to cardiac MRI

Malena Sabaté Landman

Joint work with **Abdullah Abdullah**, **Martin Holler** and **Karl Kunisch**

Elevator Pitch - CMIH
July 2022

General setting: MRI as a discrete inverse problems

We are interested in finding solutions of linear inverse problems of the form:

$$Ax + e = b, \quad \text{where} \quad A \in \mathbb{R}^{m\times n}.$$

Here, we assume that:

- ▷ the linear operator A and the measurements b are known;
- ▷ the additive noise e is Gaussian;
- ▷ the input x is unknown.

Magnetic resonance imaging

$$b = M\mathcal{F}(\sigma_c u)_c + e;$$

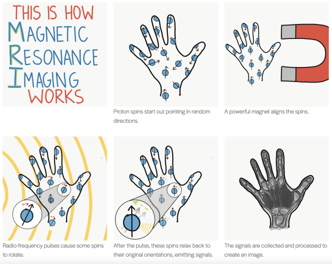

Proton spins start out pointing in random directions.

A powerful magnet aligns the spins.

Radio-frequency pulses cause some spins to rotate.

After the pulse, these spins relax back to their original orientations, emitting signals.

The signals are collected and processed to create an image.

General setting: discrete inverse problems

Since we deal with ill-posed problems we need regularization to obtain a meaningful solution. Since we are considering high-dimensional problems, we consider iterative methods, maybe in combination with variational regularization:

$$\min_x \|Ax - b\| + \lambda R(x).$$

We assume the images space to be the range of a generator:

$$\min_z \|AG_\theta(z) - b\| + \lambda R(G_\theta(z), z)$$

and, in particular, we would like to devise a method that can incorporate prior information on the solution in G_θ.

Moving from projections to generators

Can we use the structure of the latent space Z to add appropriate regularization?

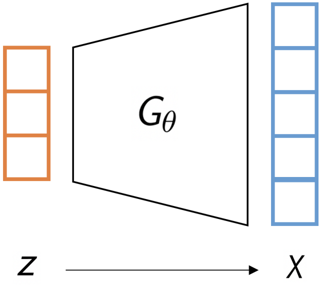

Applications to motion separation in video

We assume that each frame of the video is the output of a generator

$$x_t = G_\theta(z_t^1, z_t^2, z^3) \in \mathbb{R}^n, \tag{1}$$

where θ, z^3 are fixed over time, and z_t^1 and z_t^2 are time-dependent.

Applications to motion separation in video

We assume that each frame of the video is the output of a generator

$$x_t = G_\theta(z_t^1, z_t^2, z^3) \in \mathbb{R}^n, \tag{1}$$

where θ, z^3 are fixed over time, and z_t^1 and z_t^2 are time-dependent.

We "don't train!", instead, we minimize over all parameters

$$\theta^*, z^* = \arg\min_{\theta, z} \frac{1}{K} \sum_{k=0}^{K-1} \|y_k - G_\theta(z_k)\|^2, \quad \text{so } \ x^*(t_k) = G_{\theta^*}(z_k^*),$$

where $z_k = [z_k^1 \quad z_k^2 \quad z^3]$ and we sometimes assume prior 1d information on one of the movement types (e.g. z_k^1 is known).

Motivational example

The starting example for this work was dynamical MRI data of this form:

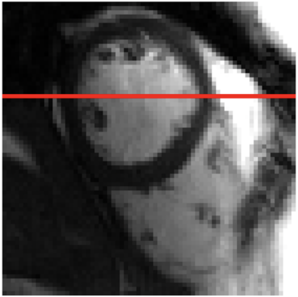

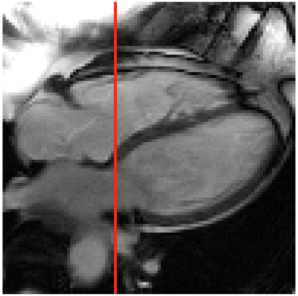

Data from the ISMRM reconstruction challenge 2014. 99 frames with 100x100 pixels. Generator: CNN with 7 layers. Latent space has 2 dynamic components. No inverse problem.

Experiments - four chamber view.

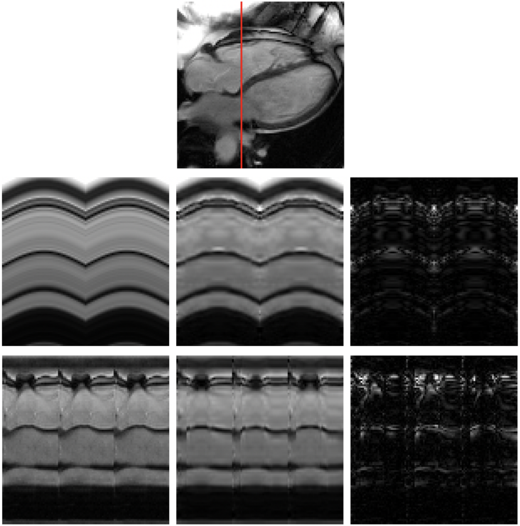

Experiments - short axis view

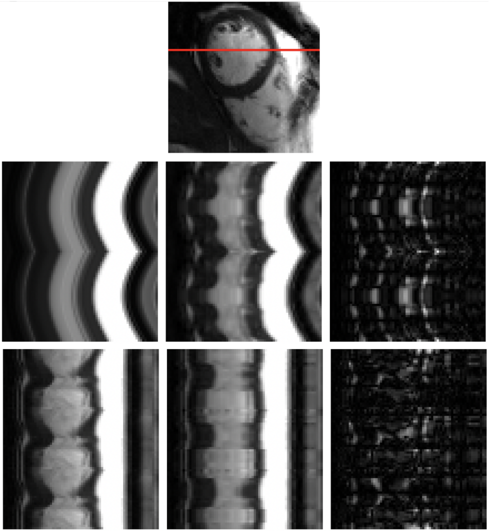

Experiments - real data

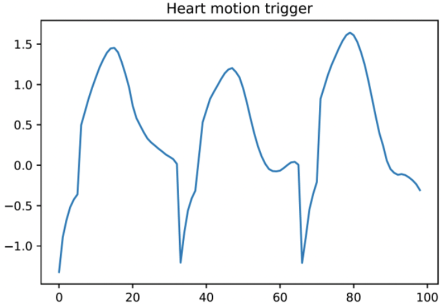

(a) Four-chamber view.

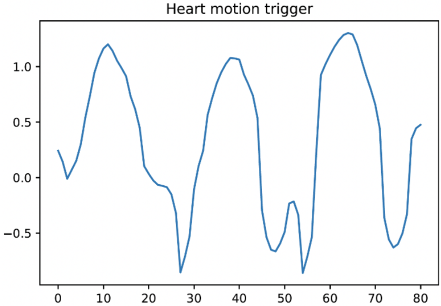

(b) Short-axis view.

Comments

Novelties:

- Untrained generators + latent space disentanglement.
- Full motion separation.

Considerations:

- Cons: initialization dependent, the dynamics splitting depends on the chosen freezing frame.
- Pros: No data bias (as there is not training!), good extrapolation to unseen combinations of states, no motion model.

Future:

- More applications: different types of movement and forward problems.
- Theoretical guarantees (by conditions on the generator).
- Regularization terms for specific latent space disentanglement.