TUHH

Bayesian Inference of Spatial Vascularity Parameters from Ultrasound Data

Fynn Bensel, Sophie Externbrink, Dimitre Hristov, Ahmed El Kaffas, Sebastian Götschel

TUHH
Technische
Universität
Hamburg

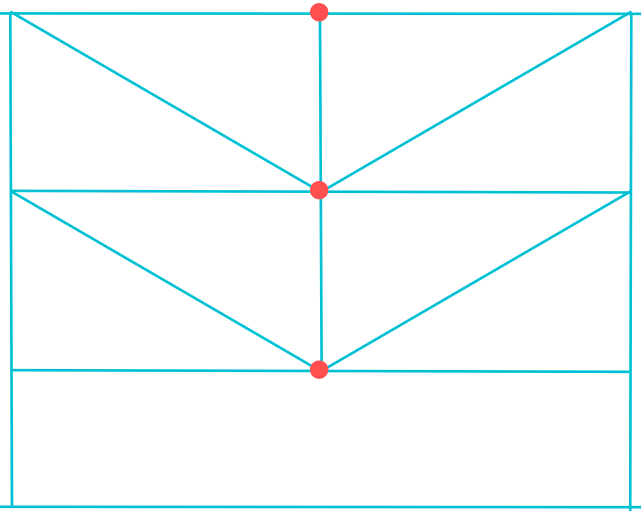

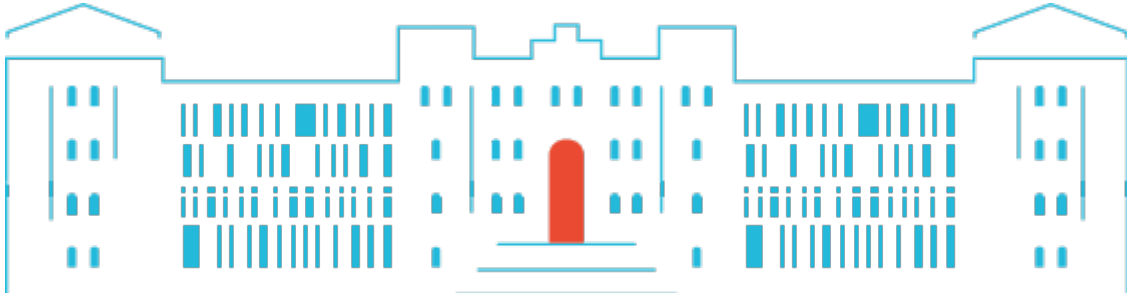

Go 20, Marsalforn, Gozo

21/05/26

Overview

1. Characterizing liver cancer treatment
2. Blood flow and perfusion model
3. Bayesian inference
4. First results
5. Conclusion

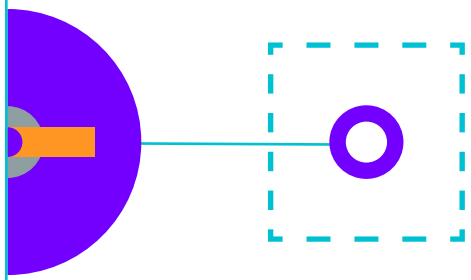

Characterizing liver cancer treatment

Project description

- Perfusion of a liver tumor as indicator for the success of the cancer treatment (chemotherapy)
  - Blood flow/perfusion visualized via a microbubble contrast agent
- Patient-specific characterization of tumor perfusion from dynamic contrast-enhanced ultrasound (DCE-US) data
  - Broadly available, cost-effective and easily accessible, bedside-imaging
  - Very low resolution, high noise

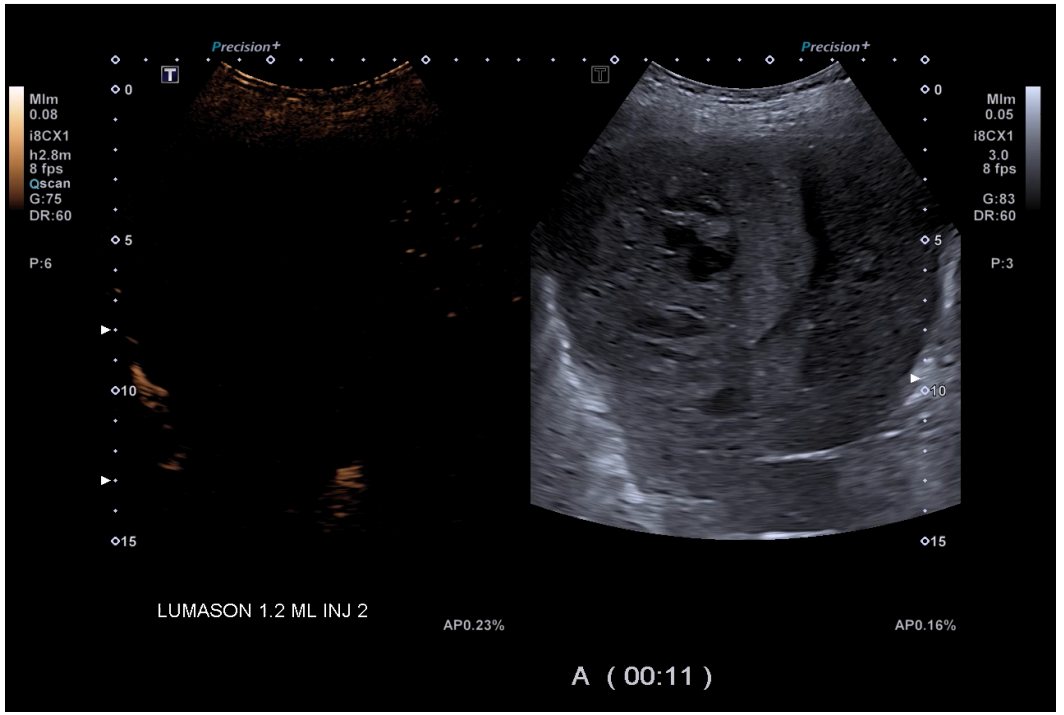

B-Mode (right) and micro bubble contrast agent visualization, data courtesy of UCSD

Characterizing liver cancer treatment

Project aim

- Reconstruction of contrast flow field (CFF) and characterization of *perfusion* based on US data
- Characterization framework for US data could enable frequent checkups and individualization of treatments
- Collaboration with Translational Ultrasound Lab of the University of California San Diego (UCSD)

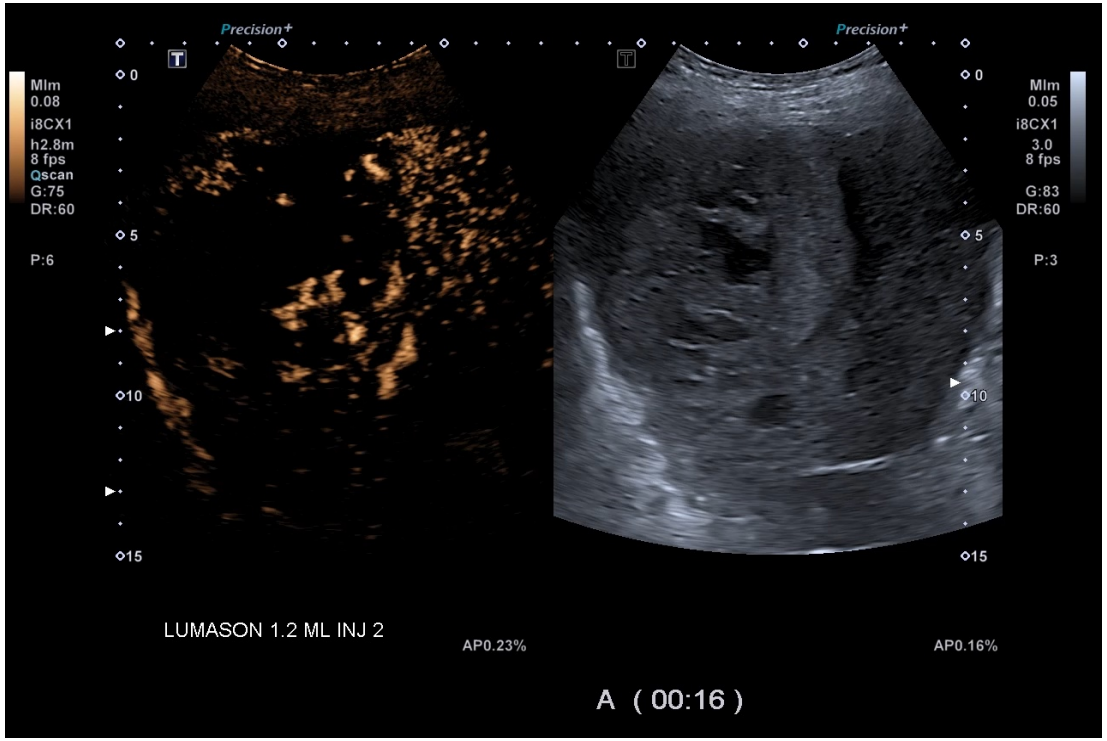

B-Mode (right) and micro bubble contrast agent visualization, data courtesy of UCSD

[1-3]

Blood flow and perfusion model

Advection diffusion

- First approach in *Hristov et al. (2022)*
- Framework for CFF based on advection-diffusion model
  - Parameter identification via minimizing regularized least-square problem

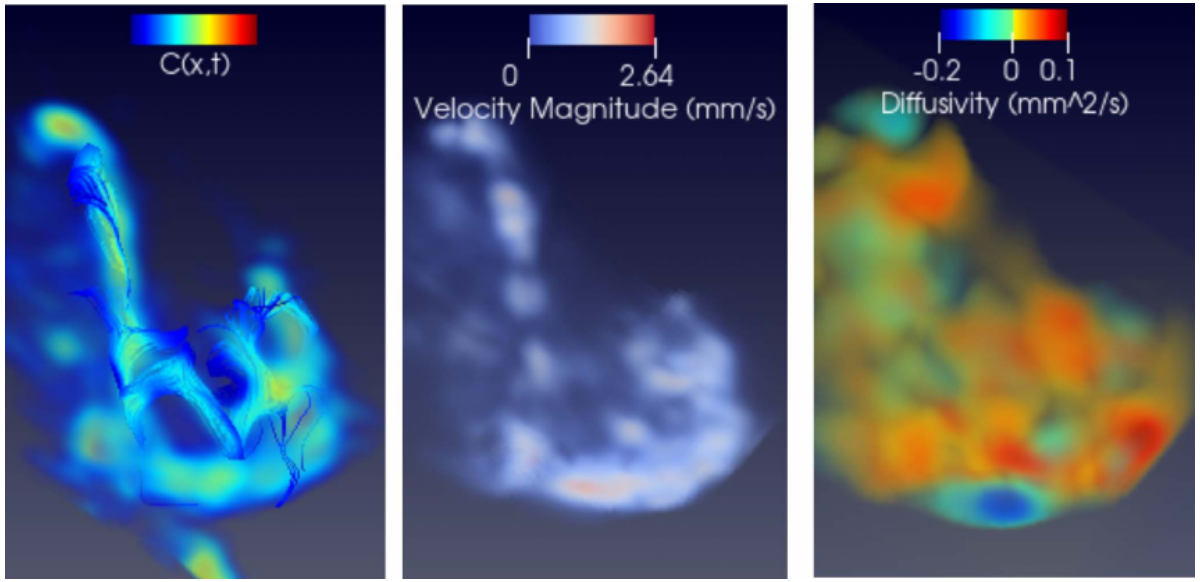

Volume renderings of reconstructed fields for concentration (left), blood-flow velocity (center) and diffusion (right) from mouse ultrasound data [3]

[1-3]

Blood flow and perfusion model

Advection diffusion

- First approach in *Hristov et al. (2022)*
- Framework for CFF based on advection-diffusion model
  - Parameter identification via minimizing regularized least-square problem
- Drawbacks:
  - Underlying model lacks interpretability of perfusion based on computed diffusion
    - (*Pseudo*-) Diffusion ≠Perfusion
    - Areas of negative diffusion

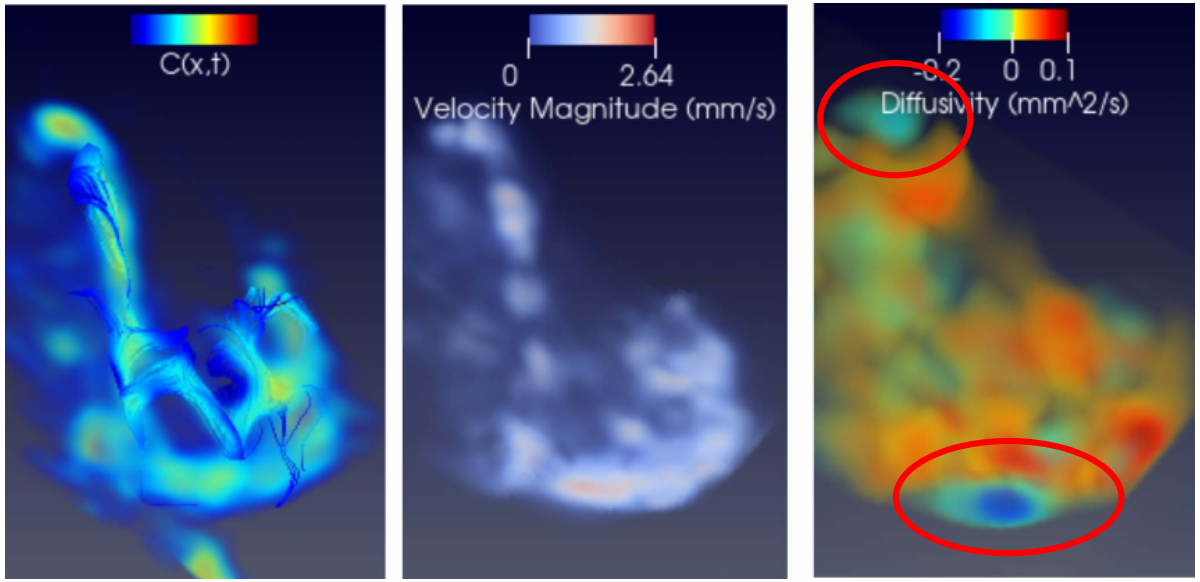

Volume renderings of reconstructed fields for concentration (left), blood-flow velocity (center) and diffusion (right) from mouse ultrasound data [3]

[1-3]

Blood flow and perfusion model

Two compartment model

- For a direct parametrization of the perfusion, we directly separate arterial and venous blood
  - Perfusion = direct transition of tracer concentration from arteries into veins

$$\partial_t C^a(\boldsymbol{x}, t) = -\nabla \cdot (\boldsymbol{u}^a(\boldsymbol{x})\, C^a(\boldsymbol{x}, t)) - K(\boldsymbol{x}) C^a(\boldsymbol{x}, t)$$

$$\partial_t C^v(\boldsymbol{x}, t) = -\nabla \cdot (\boldsymbol{u}^v(\boldsymbol{x})\, C^v(\boldsymbol{x}, t)) + K(\boldsymbol{x}) C^a(\boldsymbol{x}, t)$$

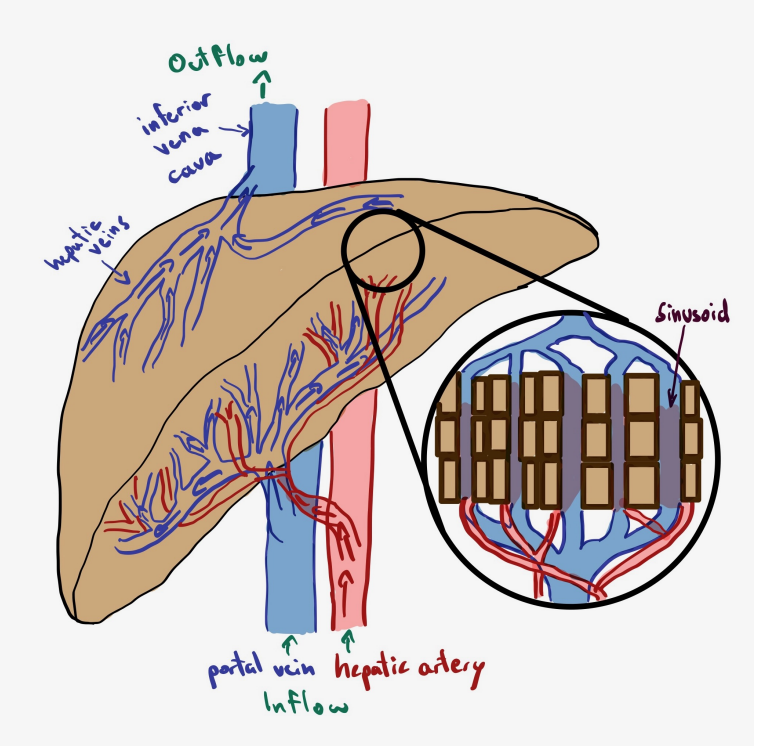

Sketch of blood flow inside the liver [3]

[1-3]

Blood flow and perfusion model

Two compartment model

- For a direct parametrization of the perfusion, we directly separate arterial and venous blood
  - Perfusion = direct transition of tracer concentration from arteries into veins

$$\partial_t C^a(\boldsymbol{x},t) = -\nabla \cdot (\boldsymbol{u}^a(\boldsymbol{x})\, C^a(\boldsymbol{x},t)) - K(\boldsymbol{x}) C^a(\boldsymbol{x},t)$$

$$\partial_t C^v(\boldsymbol{x},t) = -\nabla \cdot (\boldsymbol{u}^v(\boldsymbol{x})\, C^v(\boldsymbol{x},t)) + K(\boldsymbol{x}) C^a(\boldsymbol{x},t)$$

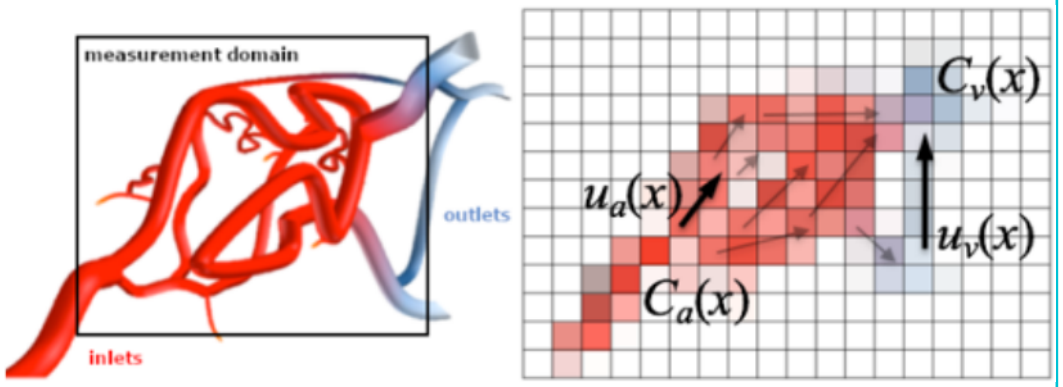

Sketch of processes in two compartment model [3]

[1-5]

Blood flow and perfusion model

Two compartment model

- Low resolution of ultrasound prohibits resolving individual vessels
  - Description of quantities on tissue level
  - Blood and tracer concentration is separated into arterial and venous shares
  - Each pixel contains information for both concentrations and local values for each stationary parameter field

$$\text{2D: } N_\text{P} = 5 * N_{pixel} = \mathcal{O}(10^5)$$

$$\text{3D: } N_\text{P} = 7 * N_{pixel} = \mathcal{O}(10^6)$$

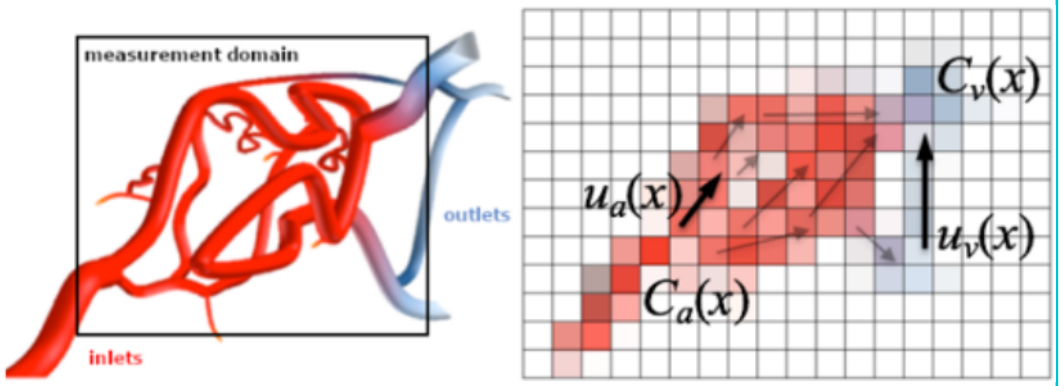

Sketch of processes in two compartment model [3]

[1-5]

Blood flow and perfusion model

TUHH

Benefits of two compartment model:

- Physiologically more meaningful (direct information of perfusion in each pixel)
- Explicit time integration (RK3)
- Linear problem, still requires nonlinear WENO-solver (WENO3)

Drawbacks of two compartment model:

- Need to infer velocities for each compartment from measurements
- Only possible to track tissue concentration $C(\boldsymbol{x},t) = C^a(\boldsymbol{x},t) + C^v(\boldsymbol{x},t)$
  - How to set initial and boundary conditions ?

[4-5]

Bayesian inference

Patient-specific model parameters

- Reconstruction of CFF, perfusion and velocity parameters such that concentration predicted by our model “match” the measurement

$$\arg\min_{\boldsymbol{\theta}} \mathcal{L}(D, C) \quad \text{with } \mathcal{M}(\boldsymbol{s}; \boldsymbol{\theta}) = C \qquad \boldsymbol{\theta} = \{\boldsymbol{K}, \boldsymbol{u}^{\mathrm{a}}, \boldsymbol{u}^{\mathrm{v}}\}$$

- Several methods available for inverse optimization (e.g. adjoint gradient computation, DL-approaches, Bayesian Inference, …)

[1-3]

Bayesian inference

Bayesian parameter estimation

- Solve inverse problem of finding the model parameters in a probabilistic setting

$$\pi\left(\theta|D\right) \propto \pi\left(D|\theta\right)\pi\left(\theta\right)$$

Bayesian inference

Bayesian parameter estimation

- Solve inverse problem of finding the model parameters in a probabilistic setting

$$\pi\left(\theta|D\right) \propto \pi\left(D|\theta\right)\pi\left(\theta\right)$$

Posterior distribution
- distribution of parameters w.r.t the measurements

Likelihood function
- error/mismatch function
- depends on noise model/distribution

$$\pi\left(D|\theta\right) = \mathcal{L}_{\pi}\left(D, \mathcal{M}(\theta)\right)$$

Prior distribution
- distribution before measurements
- chance to include **prior knowledge** here
- good prior has significant impact on predictions and inference

[6-7]

Bayesian inference

Speckle noise model and likelihood

- Interpretation of parameters as *random* – inferred parameters are means of posterior

$$\log\left[\pi\left(\theta|D\right)\right] \propto \log\left[\mathcal{L}_\pi\left(D, \mathcal{M}(\theta)\right)\right] + \log\left[\pi\left(\theta\right)\right]$$

- Likelihood depends on noise model (here: speckle noise)

$$D_i = C_i + \sqrt{C_i}\, n_i\,, \text{ with } n_i \sim \mathcal{N}(0, \sigma_n^2)\,, i = 1, \ldots, N_{\text{pixel}}$$

- Speckle noise translates to log-likelihood being sum of errors

$$\log\left[\mathcal{L}_\pi(D, \mathcal{M}(\theta))\right] \propto \sum_{i=1}^{N_{\text{pixel}}} \sum_{k=1}^{N_{\text{time}}} \| \frac{D_{i,k} - C_{i,k}}{\sqrt{C_{i,k}}} \|^2$$

[6-7]

Bayesian inference

TUHH

Solution strategies

Variational inference

- Assume a distribution to approximate the posterior

$$\tilde{\pi}(\theta|D;\xi) \approx \pi(\theta|D)$$

- Minimize distance to *true* posterior

$$\arg\min_{\xi} \mathcal{D}\left(\tilde{\pi}(\theta|D;\xi), \pi(\theta|D)\right)$$

- Posterior is easy handleable
- Choice of distribution heavily impacts the final solution
- Generally used for field inference

Sampling

- Generating initial sample from prior and shifting them to the posterior (e.g. Metropolis-Hastings)
- MCMC most prominent

- Prefered since samples approximate the true posterior
- Computational hard/infeasible in high-dimensional parameter spaces
- MCMC algorithms require serial evaluation of samples

[6-8]

Bayesian inference

Transitional Markov Chain Monte Carlo (TMCMC)

- Use fractional powers to smoothly transition from prior to posterior (simulated annealing)

$$\pi\left(\theta|D\right) \propto \mathcal{L}_{\pi}\left(D,C\right)^{\beta_j} \pi\left(\theta\right) \ , \quad 0 = \beta_0 < \beta_1 < \ldots < \beta_m = 1$$

[6]

Bayesian inference

Transitional Markov Chain Monte Carlo (TMCMC)

- Use fractional powers to smoothly transition from prior to posterior (simulated annealing)

$$\pi\left(\theta|D\right) \propto \mathcal{L}_{\pi}\left(D,C\right)^{\beta_j}\,\pi\left(\theta\right)\ , \quad 0=\beta_0<\beta_1<\ldots<\beta_m=1$$

- Incremental particle-based method
  - Iteratively cycles 4 stages
  - Naively parallelizable due to independent samples
  - Number of samples grows (exponentially) with problem-dimension

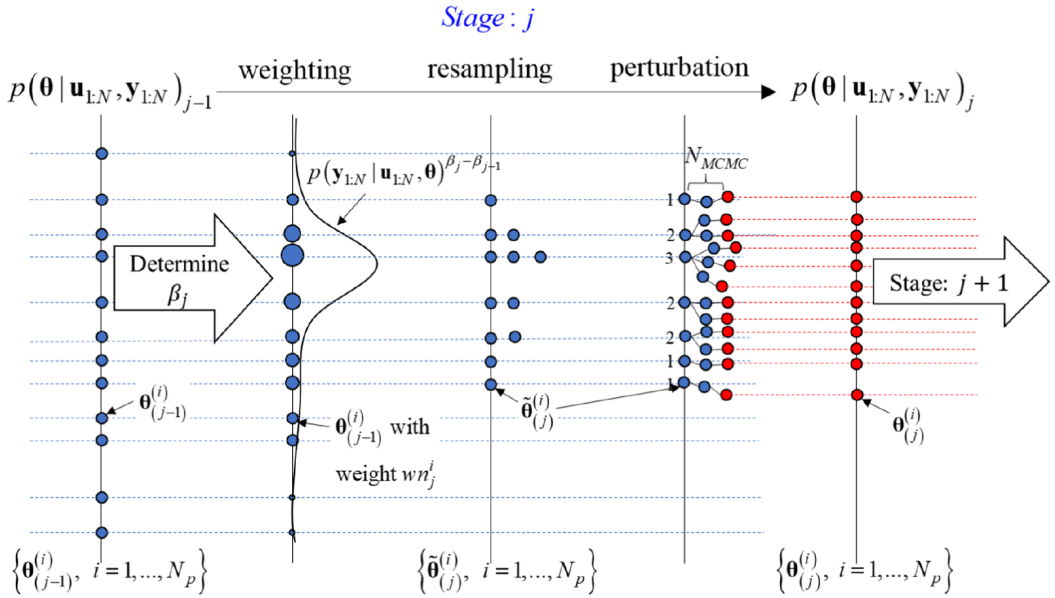

Single stage of TMCMC algorithm for a single parameter problem, adapted from [6]

[6]

Bayesian inference

Transitional Markov Chain Monte Carlo (TMCMC)

- Example for 2 parameter eigenvalue problem

$$\begin{bmatrix} \theta_1 + \theta_2 & -\theta_2 \\ -\theta_2 & \theta_2 \end{bmatrix} \Rightarrow \{(\psi_1, \lambda_1)\,,\ (\psi_2, \lambda_2)\}$$

$$\tilde{\lambda}_1 = \lambda_1 + n_1\ , \quad n_1 \sim \mathcal{N}(0.0, 1.0)$$

$$\tilde{\lambda}_2 = \lambda_2 + n_2\ , \quad n_2 \sim \mathcal{N}(0.0, 0.1)$$

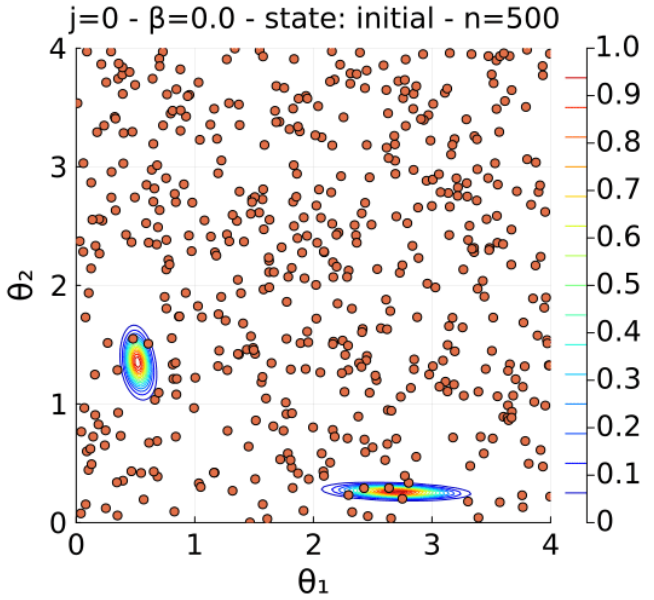

TMCMC posterior approximation for 2 parameter eigenvalue problem

[9]

Bayesian inference

Lower dimensional representation of parameter fields

- Probabilistic setting, describe parameter fields as Gaussian Random Fields

$$K(\boldsymbol{x};\mu_K,\boldsymbol{\xi}) = \mu_K + \sum_{\mathrm{i}=1}^{N_{\mathrm{DOF}}} \sqrt{\lambda_{\mathrm{i}}}\psi_{\mathrm{i}}(\boldsymbol{x})\xi_{\mathrm{i}} \qquad \xi_i \sim \mathcal{N}(0,1)$$

- Decomposition of correlation matrix into Eigenmode, Eigenvalue-pairs

$$K(\boldsymbol{x}) \sim \mathcal{N}(\boldsymbol{\mu}_K, \boldsymbol{\Sigma}_K)\ , \quad \boldsymbol{\Sigma}_K \rightarrow \{\lambda_i, \psi_i\}$$

- Random field depends on type and parameters of correlation function $\phi(\boldsymbol{x})$

$$\boldsymbol{\Sigma}(\boldsymbol{x_1},\boldsymbol{x_2}) = \sum_{k=1}^{\infty} \lambda_k \phi_k(\boldsymbol{x_1})\phi_k(\boldsymbol{x_2})$$

$$\int_\Omega \boldsymbol{\Sigma}(\boldsymbol{x_1},\boldsymbol{x_2})\phi_k(\boldsymbol{x_2})\mathrm{d}\boldsymbol{x_2} = \lambda_k \psi_k(\boldsymbol{x_1})$$

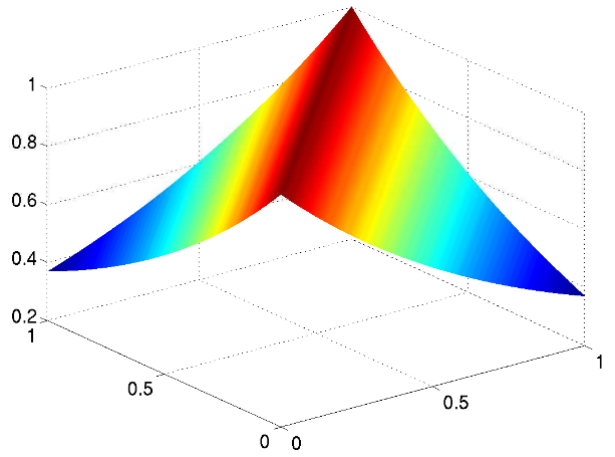

Example correlation function

[7]

Bayesian inference

TUHH

Lower dimensional representation of parameter fields

- Truncated KL expansion gives lower-dimensional functional representation
- $$K(\boldsymbol{x};\boldsymbol{\theta}^{(K)}) = \sum_{\mathrm{i}=0}^{N_{\mathrm{TO}}} \hat{\psi}_{\mathrm{i}}(\boldsymbol{x})\theta_{\mathrm{i}}^{(K)} \qquad \hat{\psi}_{\mathrm{i}} = \begin{cases} 1, & \text{if i} = 0 \\ \sqrt{\lambda_{\mathrm{i}}}\psi_{\mathrm{i}}, & \text{else} \end{cases}$$

  - log-normal:

$$K(\boldsymbol{x};\boldsymbol{\theta}^{(K)} = \prod_{\mathrm{i}=0}^{N_{\mathrm{TO}}} \hat{\psi}_{\mathrm{i}}(\boldsymbol{x})^{\theta_{\mathrm{i}}^{(K)}} \qquad \hat{\psi}_{\mathrm{i}} = \begin{cases} \exp(1), & \text{if i} = 0 \\ \exp(\sqrt{\lambda_{\mathrm{i}}}\psi_{\mathrm{i}}), & \text{else} \end{cases}$$

- Solve Bayesian inference to find certain random field realizations
- Two big obstacles:
  - Number of parameters: $\mathcal{O}(N_{\mathrm{TO}}) = \mathcal{O}(10^2)$
  - How to choose N_{TO}?

Bayesian inference

Incremental model updating

- Incrementally increase fidelity of parameter fields by adding additional mode per field
- Solve N_{TO} consecutive Bayesian inverse problems

$$\pi(\mathcal{M}_{\mathrm{i}}|D, \mathcal{M}_{\mathrm{i}-1}) \propto \mathcal{L}(D, \mathcal{M}_{\mathrm{i}}|\mathcal{M}_{\mathrm{i}-1})\, \pi(\mathcal{M}_{\mathrm{i}}|\mathcal{M}_{\mathrm{i}-1})$$

- 3 big benefits
  - 1) Number of parameters to solve is $\mathcal{O}(1)$
  - 2) Don’t need to predefine N_{TO}
  - 3) Easily exclude single parameter fields if previous update was below a certain threshold

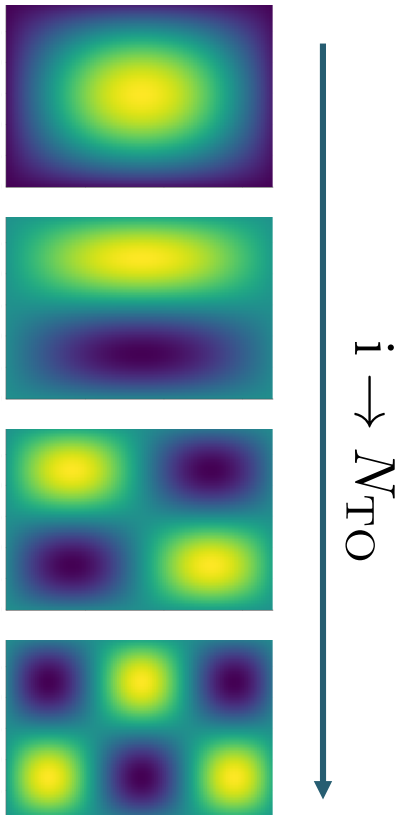

Bayesian inference

TUHH

Why stop there?

- Iterating over single parameters

  - Would hold true if parameter fields are all uncorrelated

- Fit coefficients for each mode and parameter separately
  → several 1D problems, only linear scaling of number of samples

Algorithm 1 Incremental Bayesian Model Parameter Inference

Input: $\Psi \in \mathbb{R}^{N_{\mathrm{DOF}} \times N_{\mathrm{TO}}}$ ▷ Set of functions
Input: $\boldsymbol{\Theta}$ ▷ Parameter fields
Input: ϵ, N_{it} ▷ tolerance, max. iterations

for $m = 0 : N_{\mathrm{TO}}$ **do**
  $\tilde{\mathcal{M}} \leftarrow \mathcal{M}_{\mathrm{m}-1}$
  if `isempty`($\boldsymbol{\Theta}$) **then** ▷ check if parameters left to update
    `break`
  end if
  while $i < N_{\mathrm{it}}$ **do** ▷ Iterate over single field updates
    for $\theta_j \in \boldsymbol{\Theta}$ **do**
      $\theta_j^{(i)} \leftarrow$ `TMCMC`($\tilde{\mathcal{M}}$)
      $\tilde{\mathcal{M}} \leftarrow$ `update_model`($\tilde{\mathcal{M}}, \theta_j^{(i)}$)
    end for
  end while
  $\mathcal{M}_{\mathrm{m}} \leftarrow \tilde{\mathcal{M}}$ ▷ Update new model and check tolerance
  `check_update`($\mathcal{M}_{\mathrm{m}}, \mathcal{M}_{\mathrm{m-1}}, \epsilon$)
end for

First results

2D Benchmark

- 121 x 81 x 21 grid

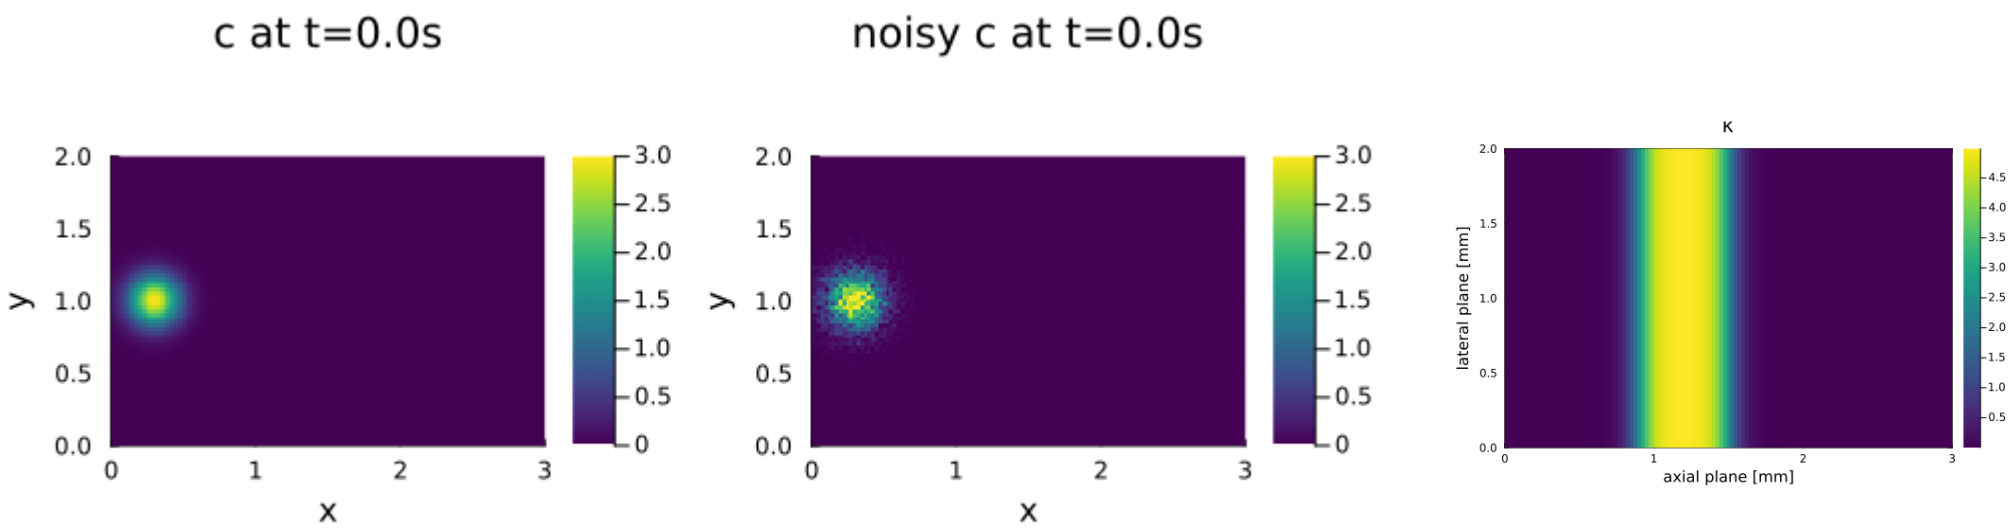

Reference concentration (left) and noisy (center) synthetic concentration data and the reference perfusion field (right)

[6]

First results

2D Benchmark

- 121 x 81 x 21 grid
- 35 Functions, 2 Iterations only for mean, 1% tolerance
- 53 fitted coefficients
- MRMS = 5.1%

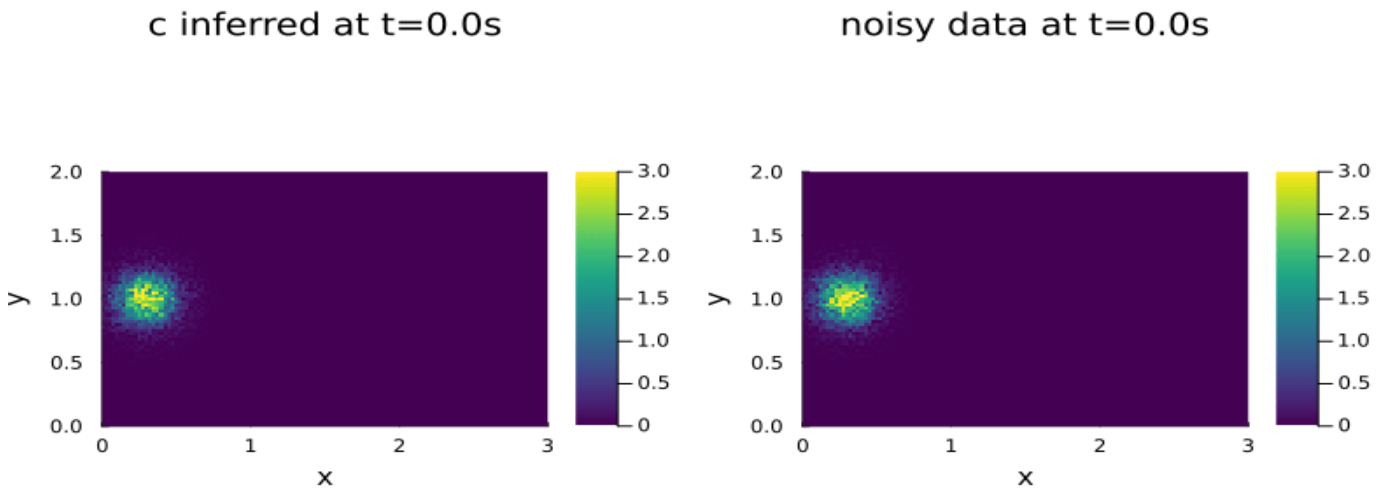

Reference concentration (left) and noisy (center) synthetic concentration data and the reference perfusion field (right)

[6]

First results

2D Benchmark

- 121 x 81 x 21 grid
- 35 Functions, 2 Iterations only for mean, 1% tolerance
- 53 fitted coefficients
- MRMS = 5.1%

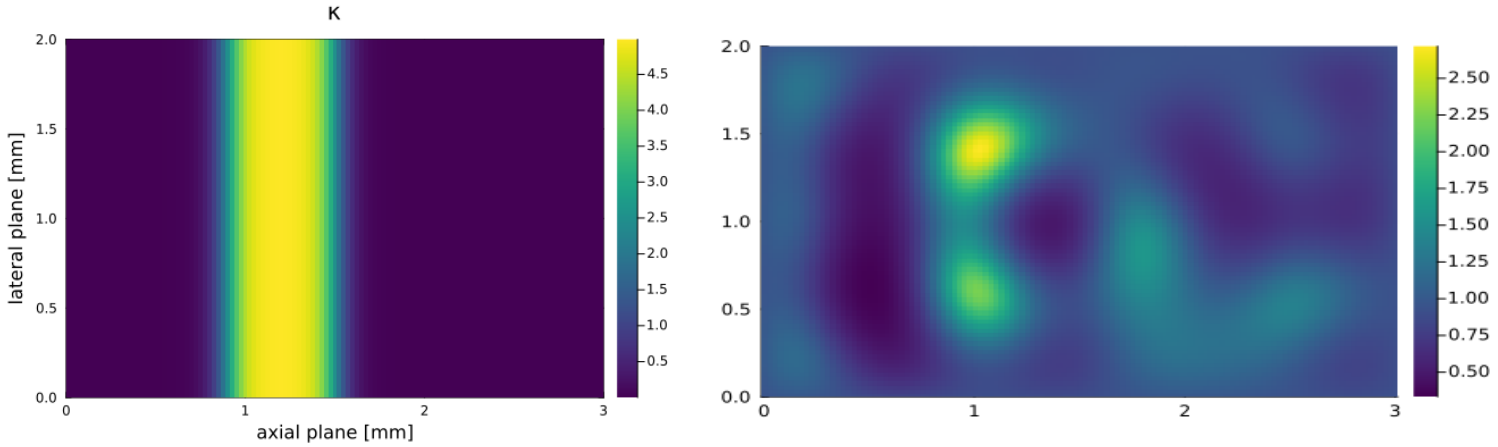

Reference perfusion field (left) and inferred perfusion field (right)

[6]

First results

2D Benchmark

- 121 x 81 x 21 grid
- Specific representation based on analytical function of known fields
- 15 fitted coefficients
- MRMS = 0.8%

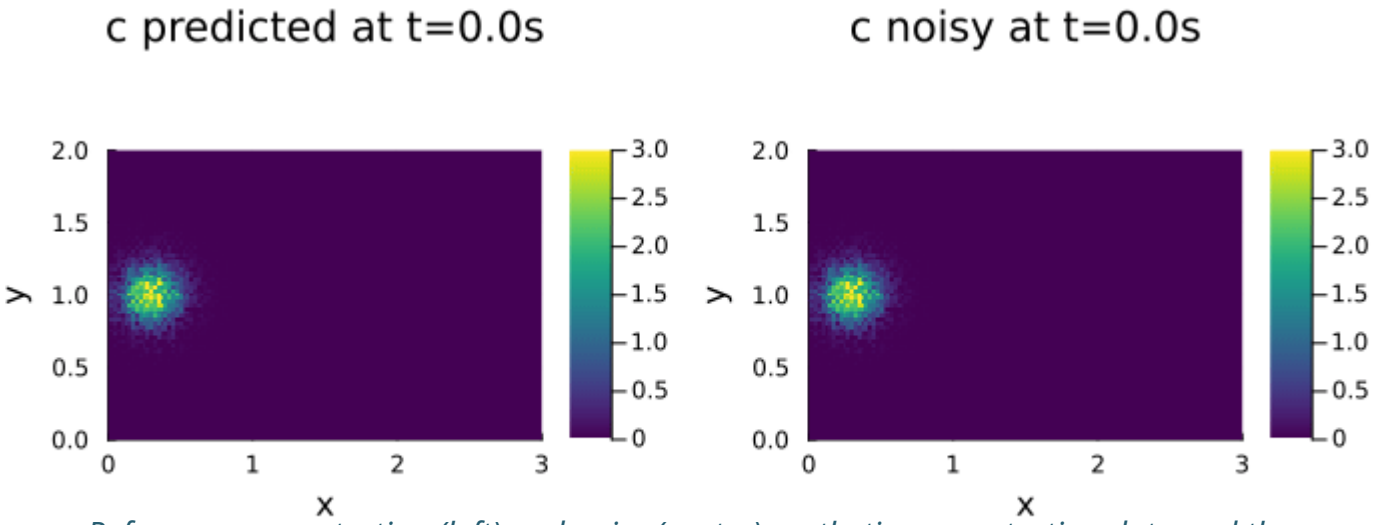

Reference concentration (left) and noisy (center) synthetic concentration data and the reference perfusion field (right)

[6]

First results

2D Benchmark

- 121 x 81 x 21 grid
- Specific representation based on analytical function of known fields
- 15 fitted coefficients
- MRMS = 0.8%

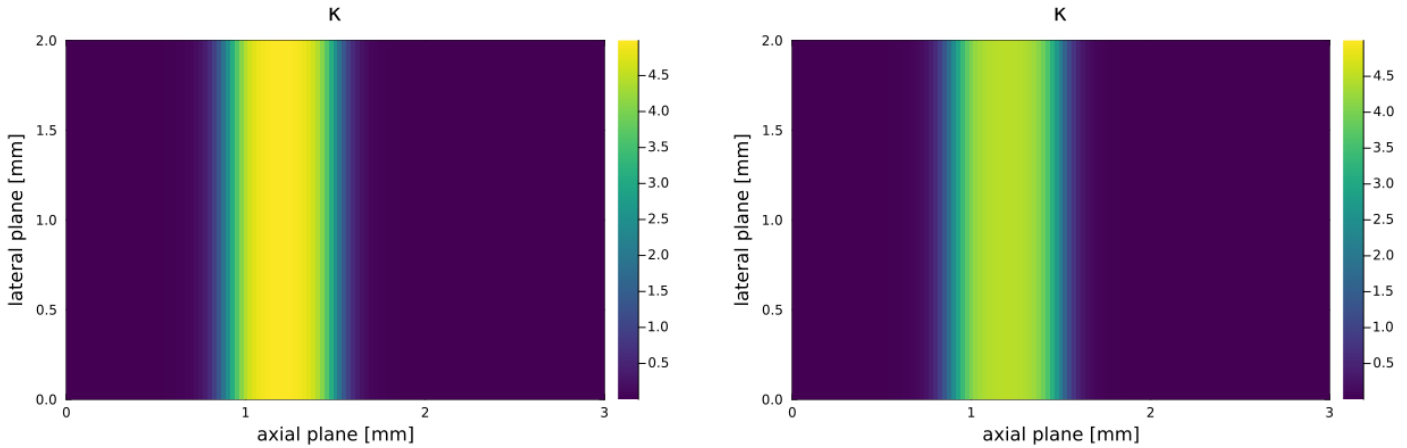

Reference perfusion field (left) and inferred perfusion field (right)

[6]

Real data application

2D Patient data

- 4 different measurements
  (2 patients, pre and 12 weeks post)
- Grid ~150 x ~130
- MRMS between 15 – 20% (compared to measurements)
- ~250 fitted coefficients (1% tolerance) in ~14h on a single node
- Still heavy effects due to breathing motion

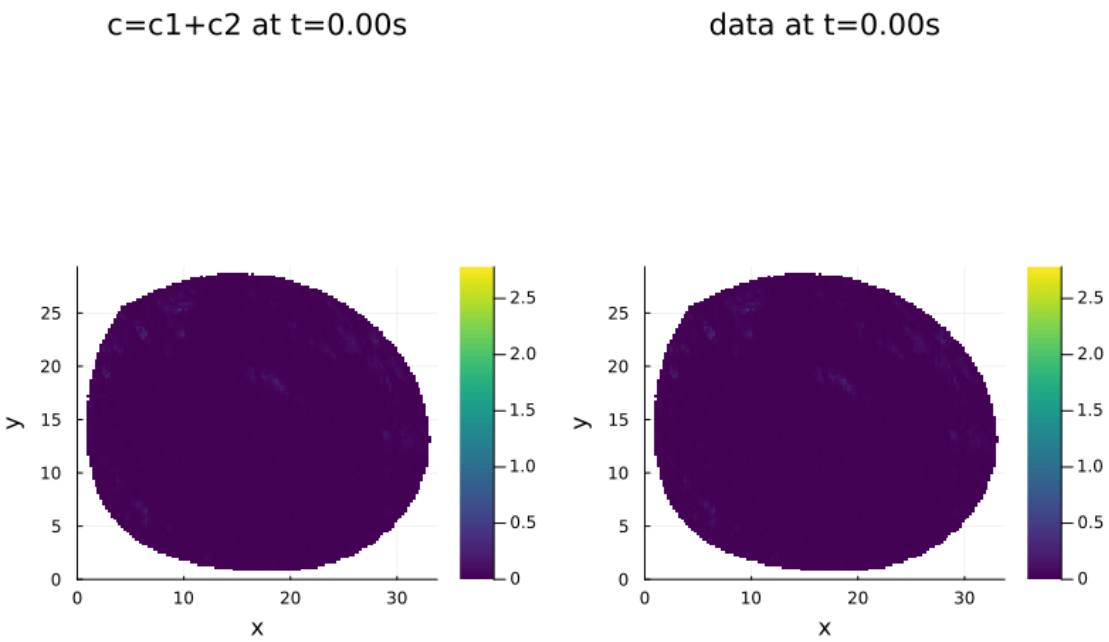

Inferred concentration (left) and measured concentration (right)

Real data application

2D Patient data

- Setting IC and BC only allows to run reconstructions for inflow phase
  - Data as IC and BC for arterial concentration
  - Zero IC and BC for venous concentration
- Need to add additional source-term to account for out-of-plan inflow

$$\partial_t C^a(\boldsymbol{x},t) = -\nabla \cdot (\boldsymbol{u}^a(\boldsymbol{x})\, C^a(\boldsymbol{x},t)) - K(\boldsymbol{x}) C^a(\boldsymbol{x},t) + S^a(\boldsymbol{x})$$

$$\partial_t C^v(\boldsymbol{x},t) = -\nabla \cdot (\boldsymbol{u}^v(\boldsymbol{x})\, C^v(\boldsymbol{x},t)) + K(\boldsymbol{x}) C^a(\boldsymbol{x},t)$$

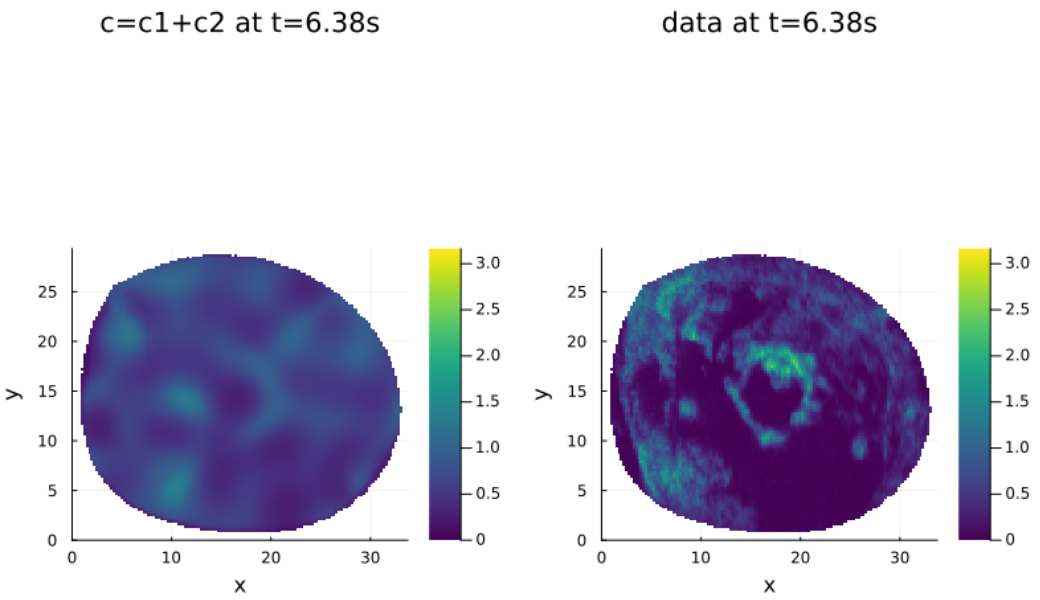

Inferred concentration (left) and measured concentration (right)

Real data application

3D Mouse data

- >60 different measurements (treated and control group)
- Grid ~90 x ~60 x ~45
- MRMS between 10 – 13% (compared to measurements)
- ~168 fitted coefficients (1% tolerance) in ~24h on a single node
- Missing frames (up to 3 consecutive) - directly impacts BC

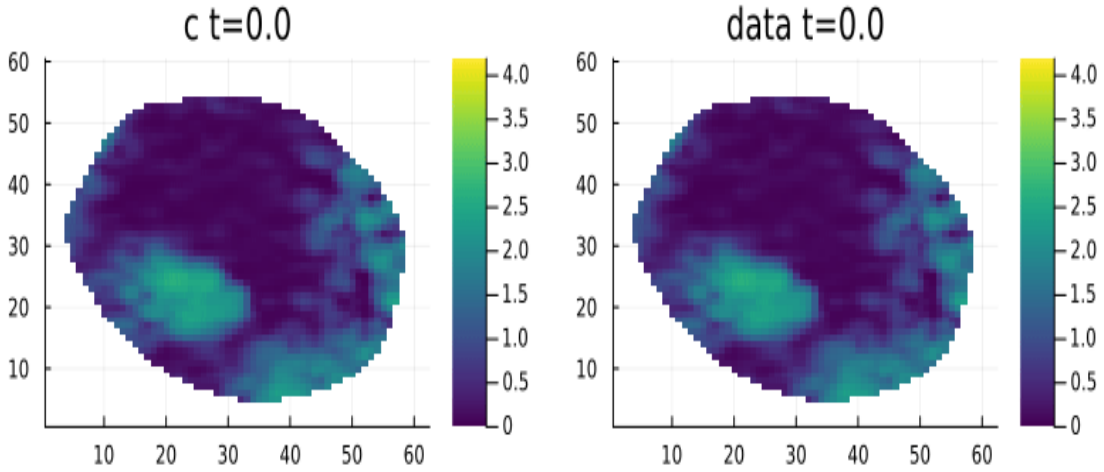

Inferred concentration (left) and measured concentration (right)

Real data application

TUHH

3D Mouse data

- >60 different measurements (treated and control group)
- Grid ~90 x ~60 x ~45
- MRMS between 10 – 13% (compared to measurements)
- ~168 fitted coefficients (1% tolerance) in ~24h on a single node
- Missing frames (up to 3 consecutive) - directly impacts BC

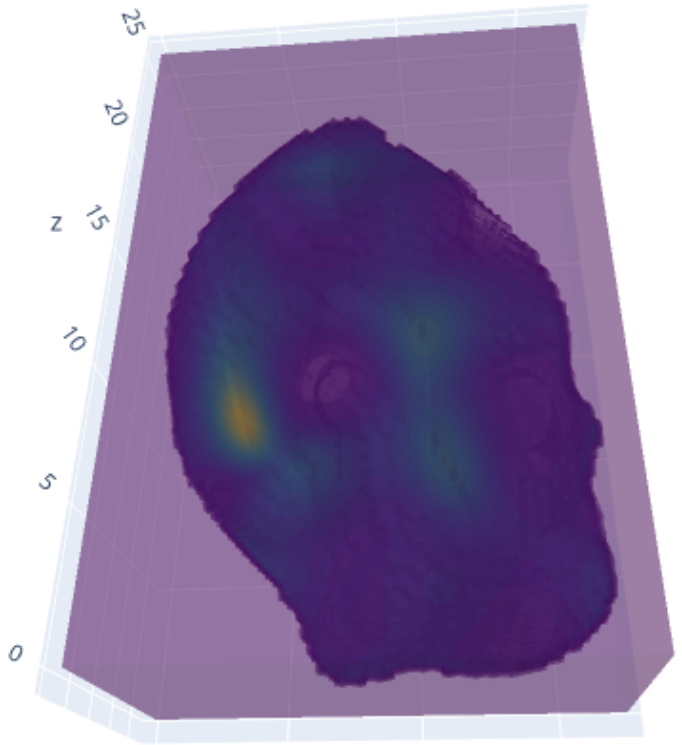

Inferred perfusion

Conclusion

TUHH

Summary

- 2C-model gives physiological interpretable parameter field for the perfusion
- Increased complexity compared to 1C-model (more parameters, only tissue concentration measurable, boundary conditions)
- 2D patient data problems need additional source for out-of-plan influx and motion-correction
- Incremental Bayesian inference for random fields gives accurate contrast-flow-field reconstruction in reasonable time
- Quality of inferred fields and total number of needed coefficients highly depends on quality of modes

Conclusion

TUHH

Current work

- More test on additional benchmark
  - Single parameter updates vs. combined mode update
  - Tolerance, number of iterations
- Further test and physiological validation of first 2D and 3D results
- Better description of prior and correlation (Ultrasound Localization Microscopy)
- Other methods for lower dimensional representation
- Simulation of full in- and outflow of tracer agent
  - Out-of-plan outflux, split measured data on arterial and venous concentration for BC

TUHH
Technische
Universität
Hamburg

Contact info

Hamburg University of Technology (TUHH)
Institute of Mathematics - Chair
Computational Mathematics
Am Schwarzenberg-Campus 3
21073 Hamburg

Bensel Fynn
fynn.bensel@tuhh.de

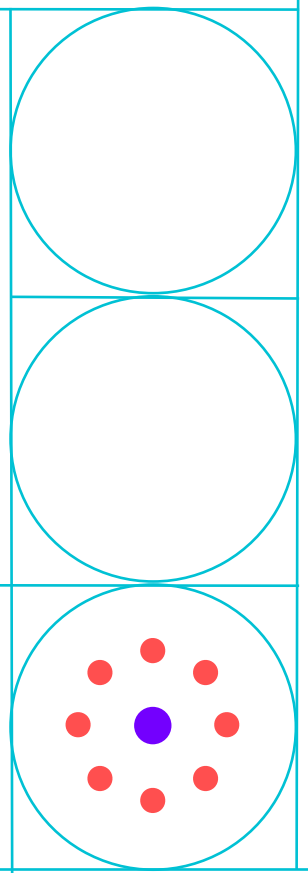

We acknowledge support by the US National Institutes of Health, National Cancer Institute, award 1R01CA286505-01A1.

INSTITUTE OF
MATHEMATICS

TRANSLATIONAL
ULTRASOUND LAB
UNIVERSITY OF CALIFORNIA, SAN DIEGO

Stanford
MEDICINE
School of Medicine

Sources

1 El Kaffas, Ahmed et al. *"Quantitative Three-Dimensional Dynamic Contrast-Enhanced Ultrasound Imaging: First-In-Human Pilot Study in Patients with Liver Metastases."* Theranostics vol. 7,15 3745-3758. 2017, doi:10.7150/thno.20329

2 El Kaffas, Ahmed et al. *"Spatial Characterization of Tumor Perfusion Properties from 3D DCE-US Perfusion Maps are Early Predictors of Cancer Treatment Response."* Scientific reports vol. 10,1 6996, 2020, doi:10.1038/s41598-020-63810-1

3 Hristov, Dimitre et al. *"Dynamic Contrast-Enhanced Ultrasound Modeling of an Analog to Pseudo-Diffusivity in Intravoxel Incoherent Motion Magnetic Resonance Imaging."* IEEE transactions on medical imaging vol. 41,12, 3824-3834, 2022, doi:10.1109/TMI.2022.3197363

4 Sourbron, Steven. *"A tracer-kinetic field theory for medical imaging."* IEEE transactions on medical imaging vol. 33,4 935-46, 2014, doi:10.1109/TMI.2014.2300450

5 Externbrink, Sophie. *"Two-Component Model for Tracer Simulation."* Hamburg University of Technology, 2022, doi:10.15480/882.4765

Sources

6 Ramancha, Mukesh K., et al. *"Bayesian Updating and Identifiability Assessment of Nonlinear Finite Element Models."* Mechanical Systems and Signal Processing, vol. 167, 108517, 2022, doi:10.1016/j.ymssp.2021.108517

7 Uribe, Felipe et al. *"Bayesian inference of random fields represented with the Karhunen–Loève expansion."* Computer Methods in Applied Mechanics and Engineering, 358:112632, 2020, doi: 10.1016/j.cma.2019.112632

8 Westerkamp, Margret et al. *"Towards a Field Based Bayesian Evidence Inference from Nested Sampling Data"*, arxiv, 2024, doi: 10.3390/e26110930

9 Behrensdorf, Jasper et al. *"Friesischscott/uncertaintyquantification.jl: V0.12.0."* v0.12.0, Zenodo, 2025, doi:10.5281/zenodo.14901342