Of networks and regularization: new developments in Lagrangian Particle Tracking

Daniel Schanz, Philipp Godbersen, Andreas Schröder and many more

Experiments in Fluids Seminar Series, Feb. 11 2025

DLR

Outline

- Improved particle image peak detection using convolutional neural networks
- 3D LPT evaluations: numerical cost and experiences with high-performance computing
- Advanced post-processing: Flow-field interpolation with different classes of approches (Binning, Data Assimilation, Neural Network)
- Closing advertisement: 2nd LPT and DA challenge

Peak detection in 3D reconstruction

- Multi-camera experiment
- Projections of 3D particle cloud on several cameras
- Aim: from 2D particle images back to 3D positions
- First step: Detection of particle image peaks on all cameras
- Second step Triangulation of 3D position by finding intersecting Lines-of-Sight form the different point clouds

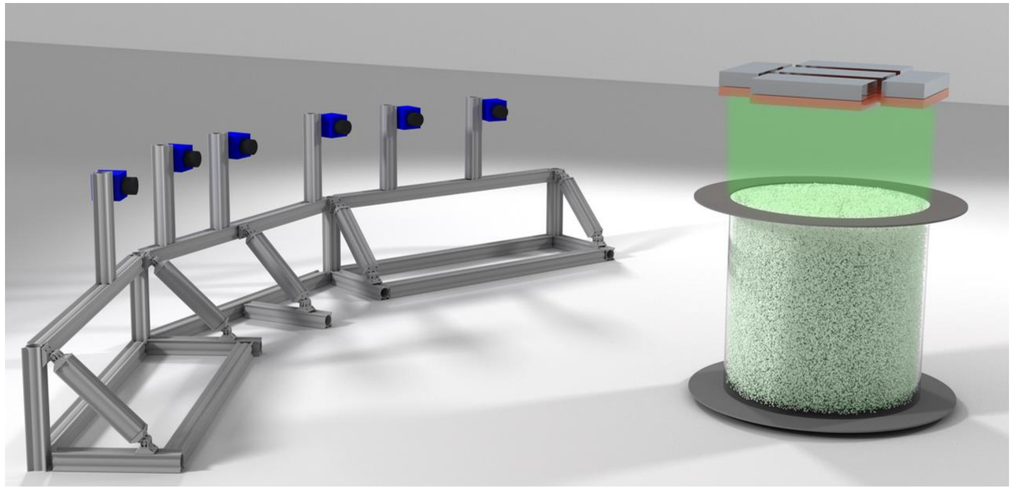

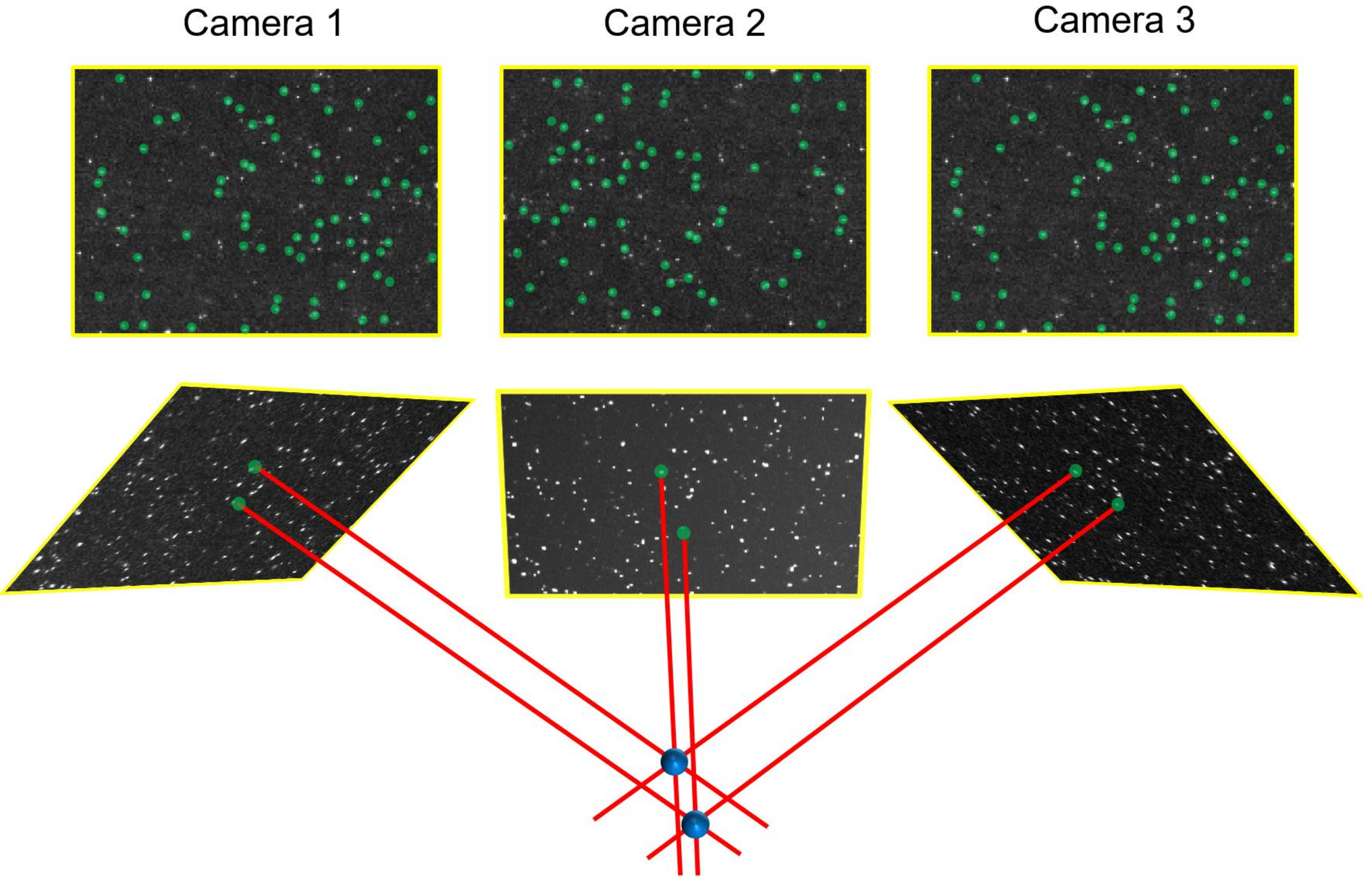

Iterative Particle Reconstruction (IPR) and Shake-The-Box (STB)

Focus of this part

IPR scheme

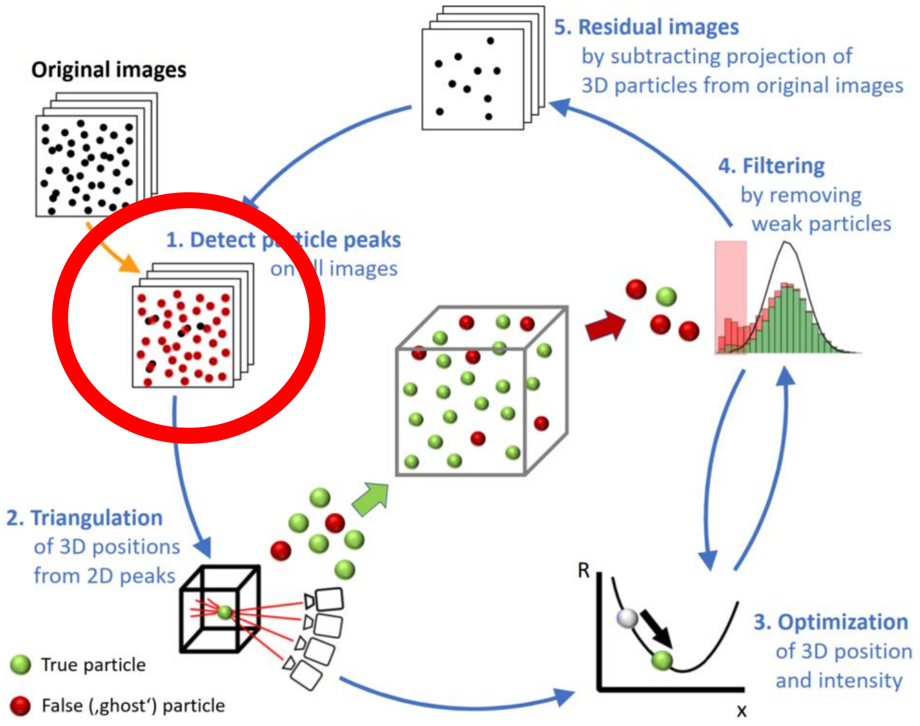

STB scheme

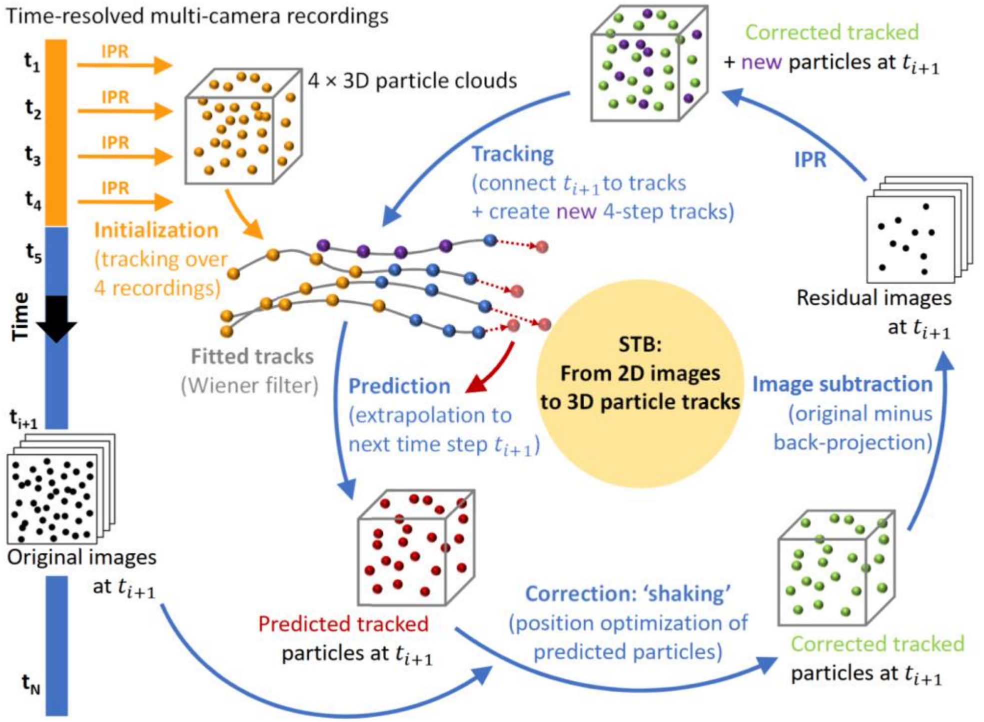

Wieneke (2013) Iterative Reconstruction Of Volumetric Particle Distribution, Meas. Sci. Technol. 24 024008

Jahn et al. (2021) Advanced Iterative Particle Reconstruction for LPT, Exp. Fluids 62, 179

Schanz et al (2016). Shake-The-Box: Lagrangian particle tracking at high particle image densities. ExpInFluids, 57(5), 70

Schröder,& Schanz, (2023) 3D Lagrangian particle tracking in fluid mechanics. Annual Review of Fluid Mechanics, 55(1)

What is Peakdetection?

- For a given image of a particle cloud:
  - determine the 2D position of the center of each particle image
  - Sub-pixel accuracy required

- Our conventional approach:
  - Cubic Interpolation (*CI*) of image
  - Identify local maxima
  - fit a gaussian for subpixel determination

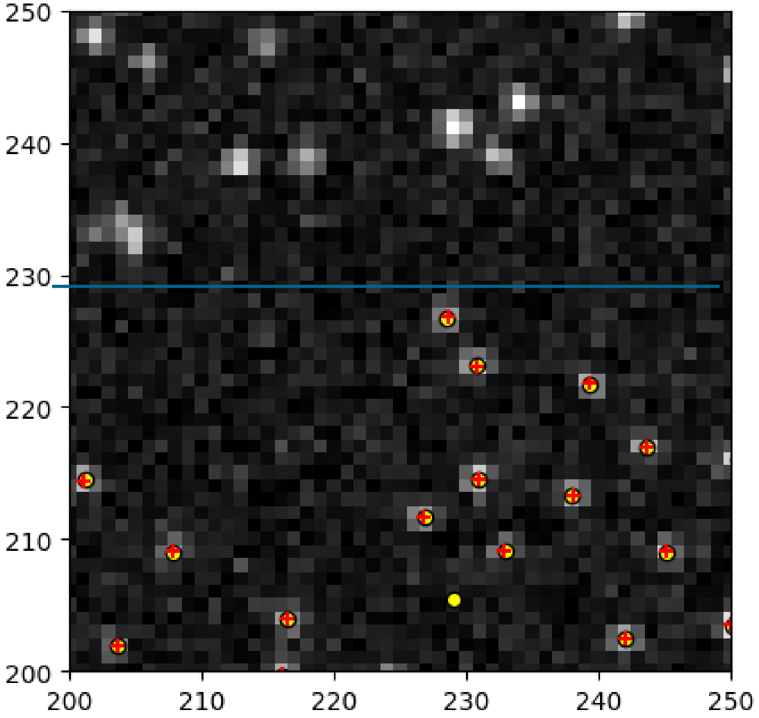

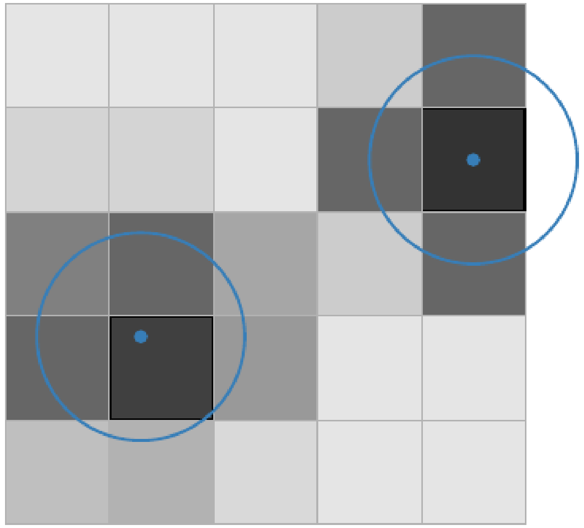

Camera Image containing two peaks

Example of peak detection on synthetic images

No markers drawn in top part

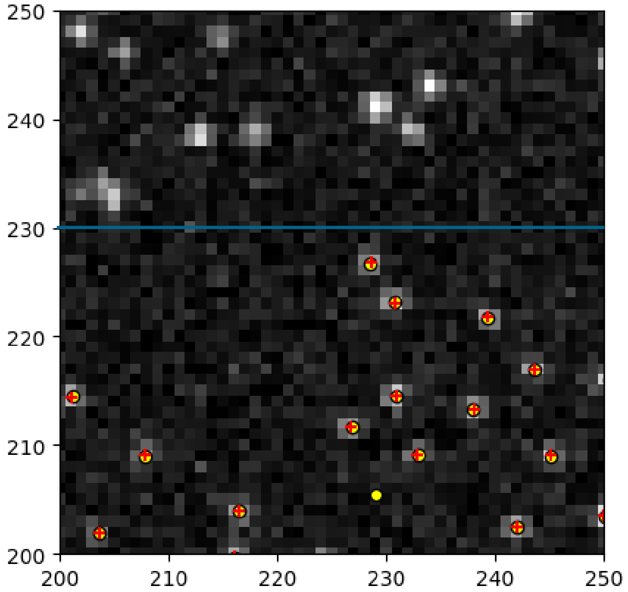

0.01 ppp

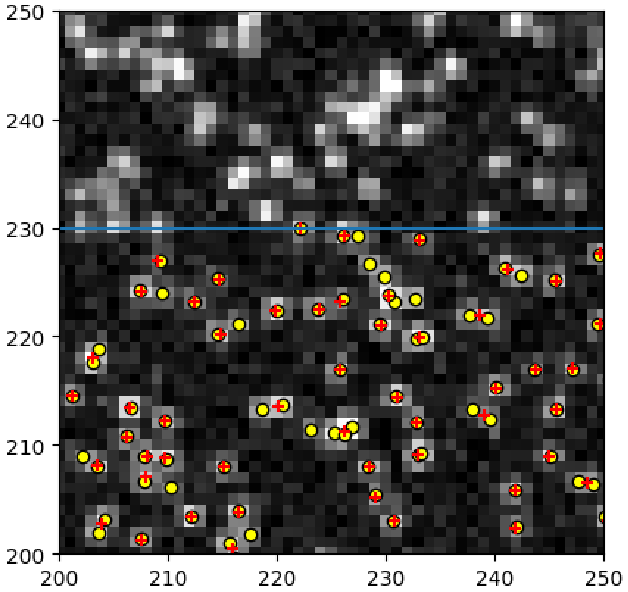

0.05 ppp

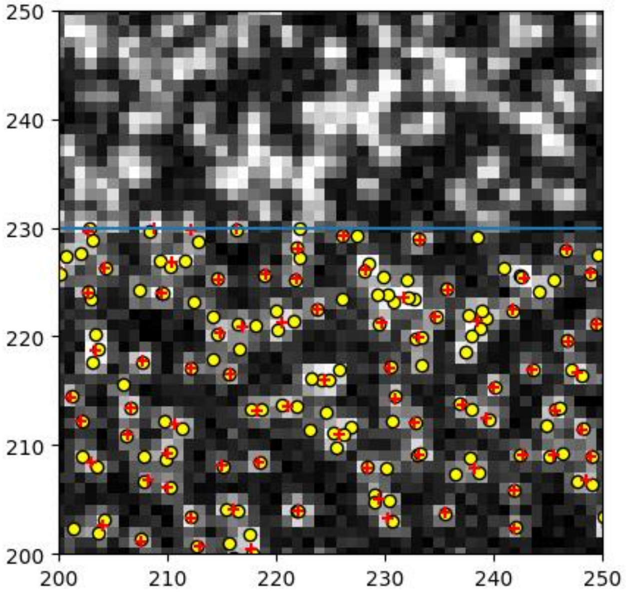

0.1 ppp

Ground truth

Detections

using conventional peakdetection

- Conventional approach works well on sparsely seeded images
- 0.01 ppp (particles per pixel): mostly solitary particles
- The higher the ppp, the more particle images overlap → problems in distinguishing peaks

Machine learned peak detection

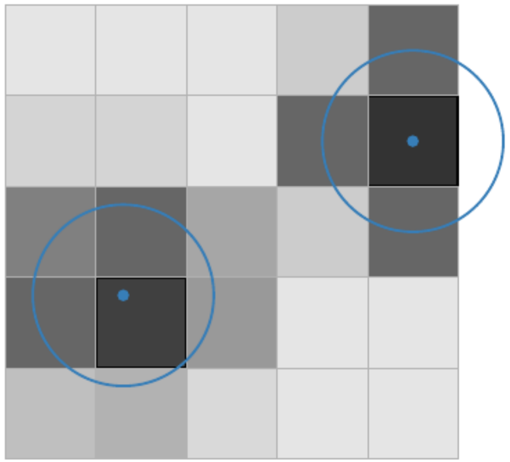

- Classification problem should be well suited to machine learing

- First consideration: Two-step approach
  - Step 1: Binary classification
  - Step 2: Subpixel-shift for each hit (individual model)

- Downsides:
  - Requires training of two models separately
  - Overhead due to individual treatment of each particle
  - Common ML approaches not designed for thousands of tiny 3x3 px objects

- Can we do this single stage?

Simultaneous classification and regression

Work by
Philipp Godbersen[1]

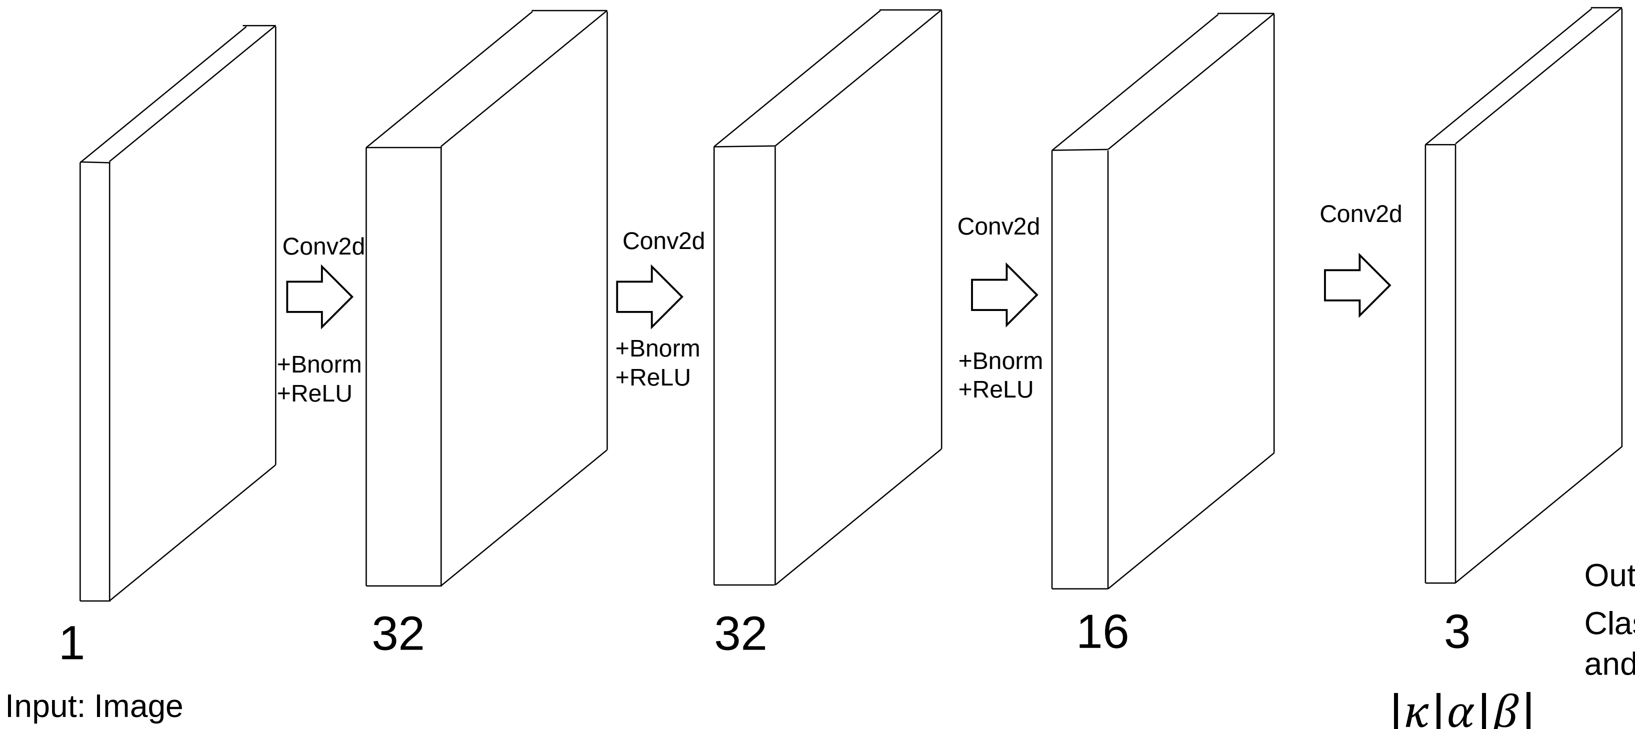

Inspired by
YOLO network
architecture[2]

Output for each pixel:
Classification κ between 0 and 1
and 2 subpixel offsets

$|\kappa|\alpha|\beta|$

Latin letters: Ground truth; Greek letters: Model output

$$\frac{1}{N_x N_y}\sum_{i}^{Nx}\sum_{j}^{Ny}\boxed{k_{ij}\log\left(\sigma\left(\kappa_{ij}\right)\right)+\left(1-k_{ij}\right)\log\left(1-\sigma(\kappa_{ij})\right)}$$

Classification loss

1 Godbersen P, Schanz D, Schröder A. (2024). Peak-CNN: improved particle image localization using single-stage CNNs. *Experiments in Fluids*, *65*(10), 153.
2 Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. IEEE conference on computer vision and pattern recognition (pp. 779-788).

Actual model architecture used

- U-net architecture*

- Starting from 256^2 px patches
  - More flexibility
  - Image size agnostic

- Downsampling to low-res patches
  - Better representation of low-frequency image content

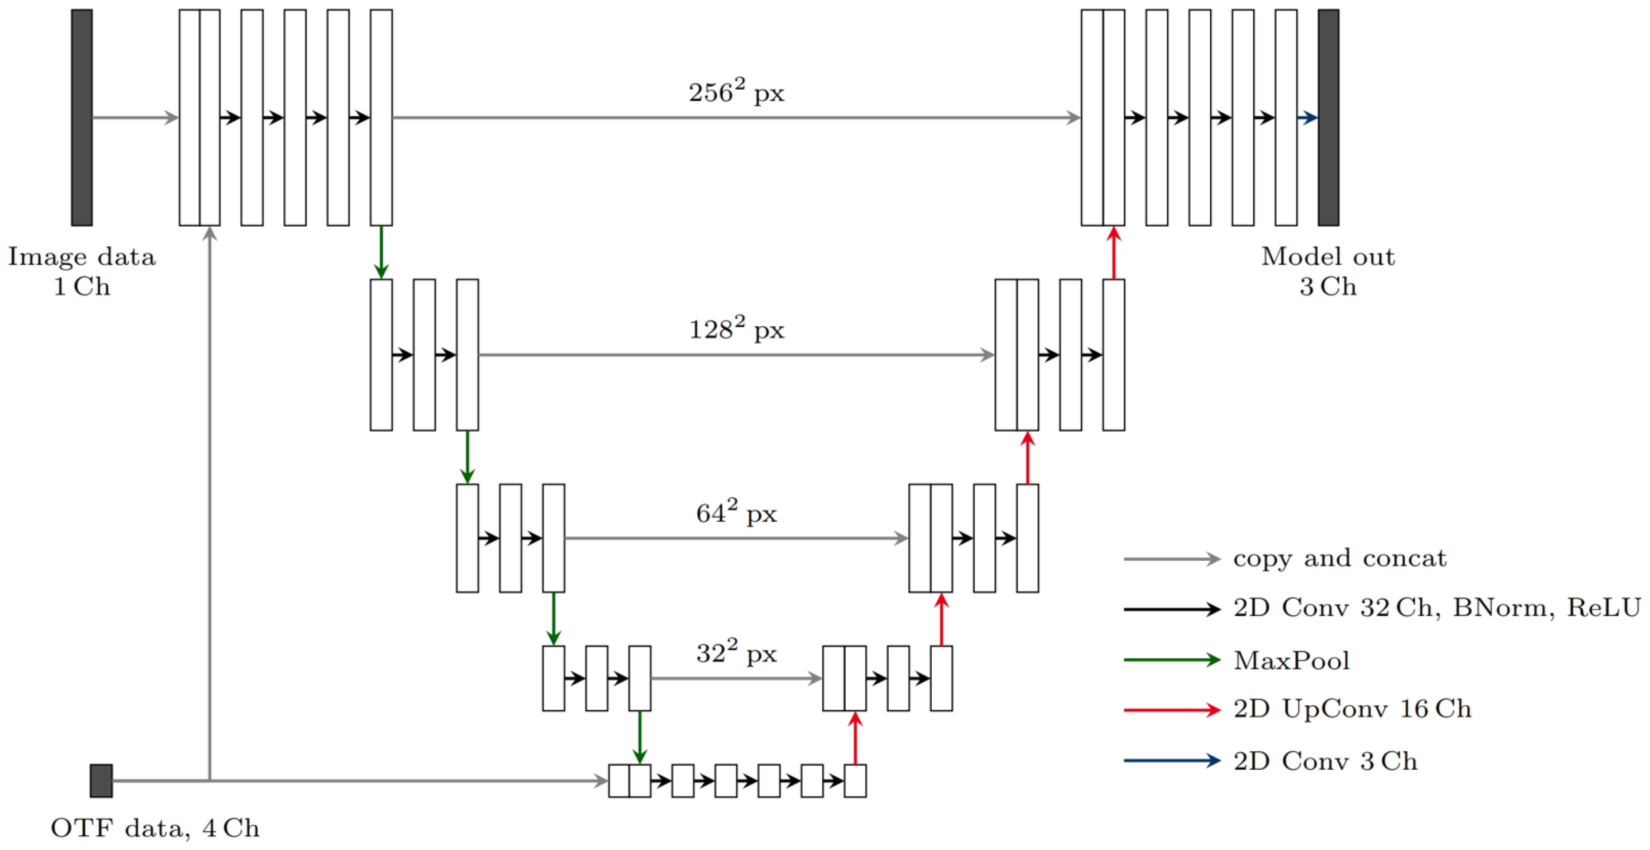

*Ronneberger O, Fischer P, Brox T (2015). U-net: Convolutional networks for biomedical image segmentation. *MICCAI 2015*

Peak-CNN applied to synthetic data

- Synthetic generation of images
- Training with images and known peaks positions

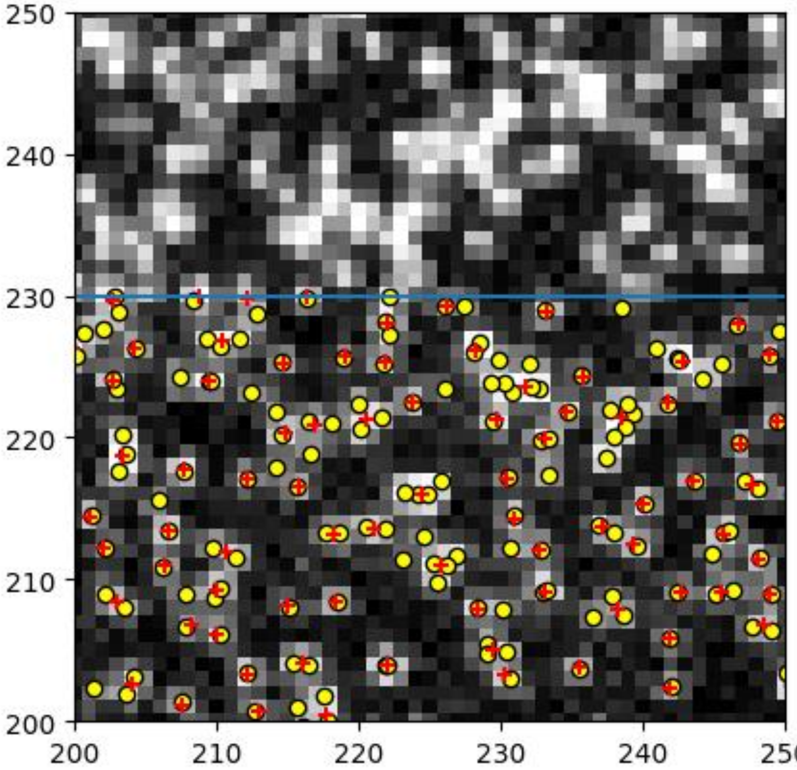

0.1 ppp conventional (CI)

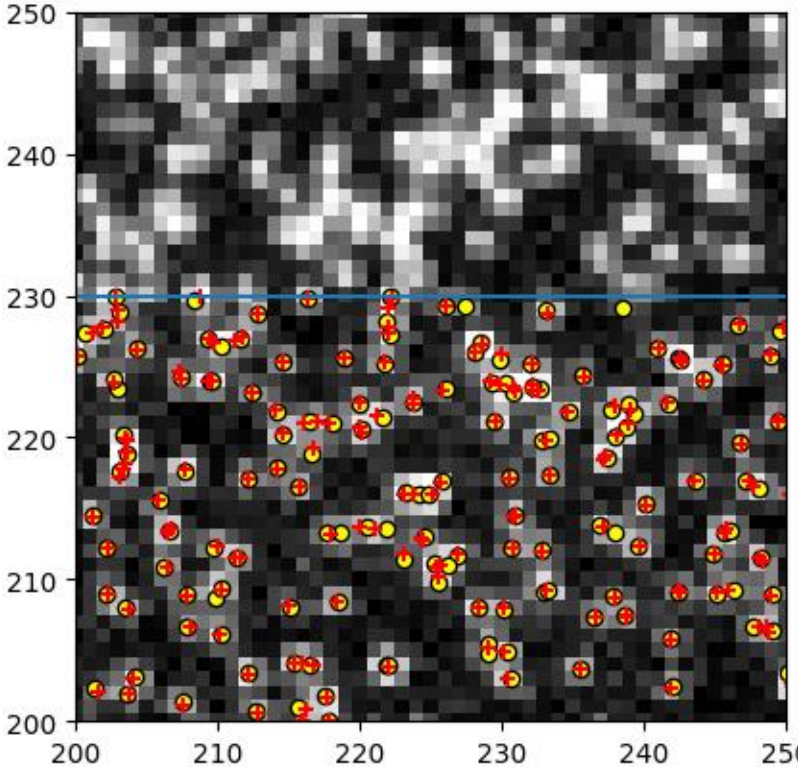

0.1 ppp Peak-CNN

Quantative evaluation of peak detection performance

Detection rate

Annotations of CI, TR, IE: pixel intensity threshold
PCNN: internal parameter

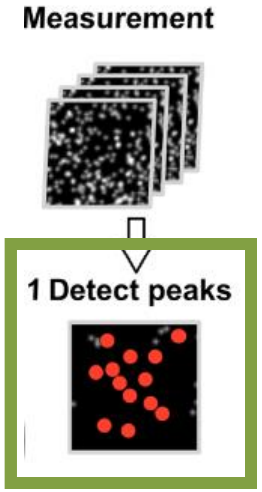

Peaks

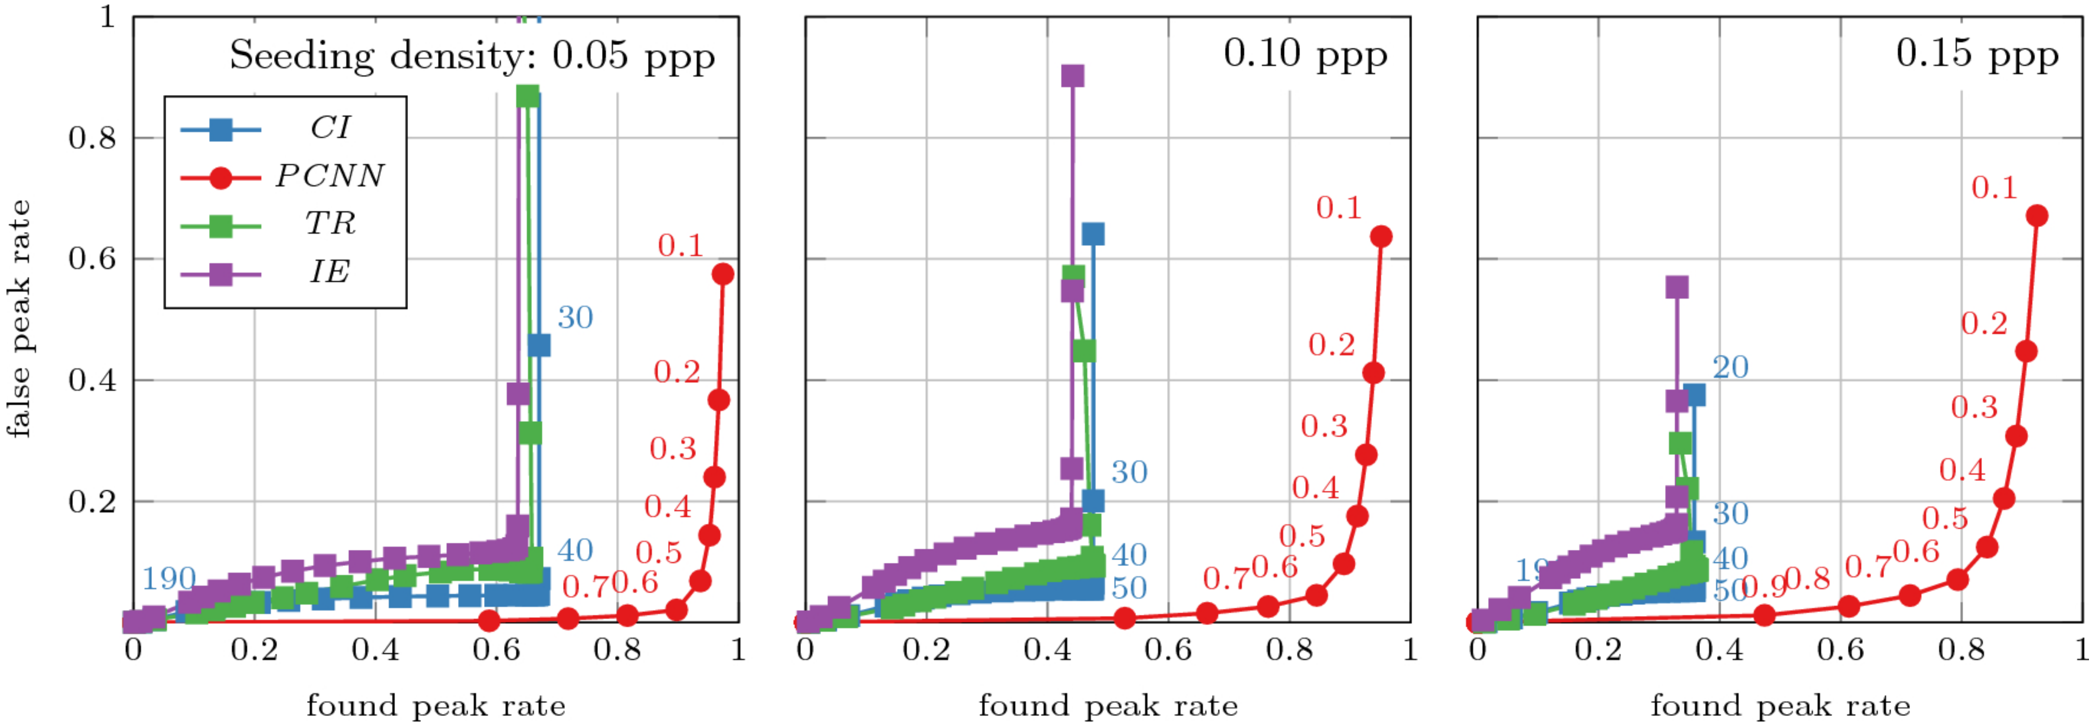

Godbersen P, Schanz D, Schröder A. (2024). Peak-CNN: improved particle image localization using single-stage CNNs. *Experiments in Fluids*, *65*(10), 153.

Quantative evaluation of peak detection performance

Triangulation results

Annotations: Used triangulation radius

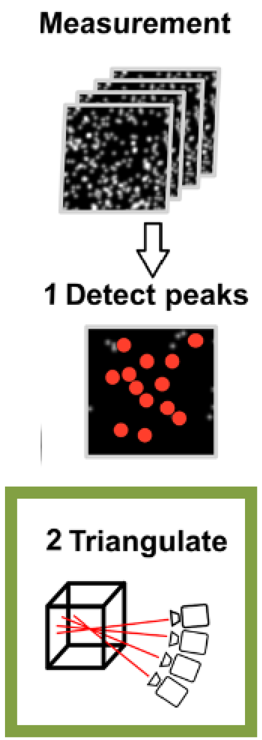

Particles

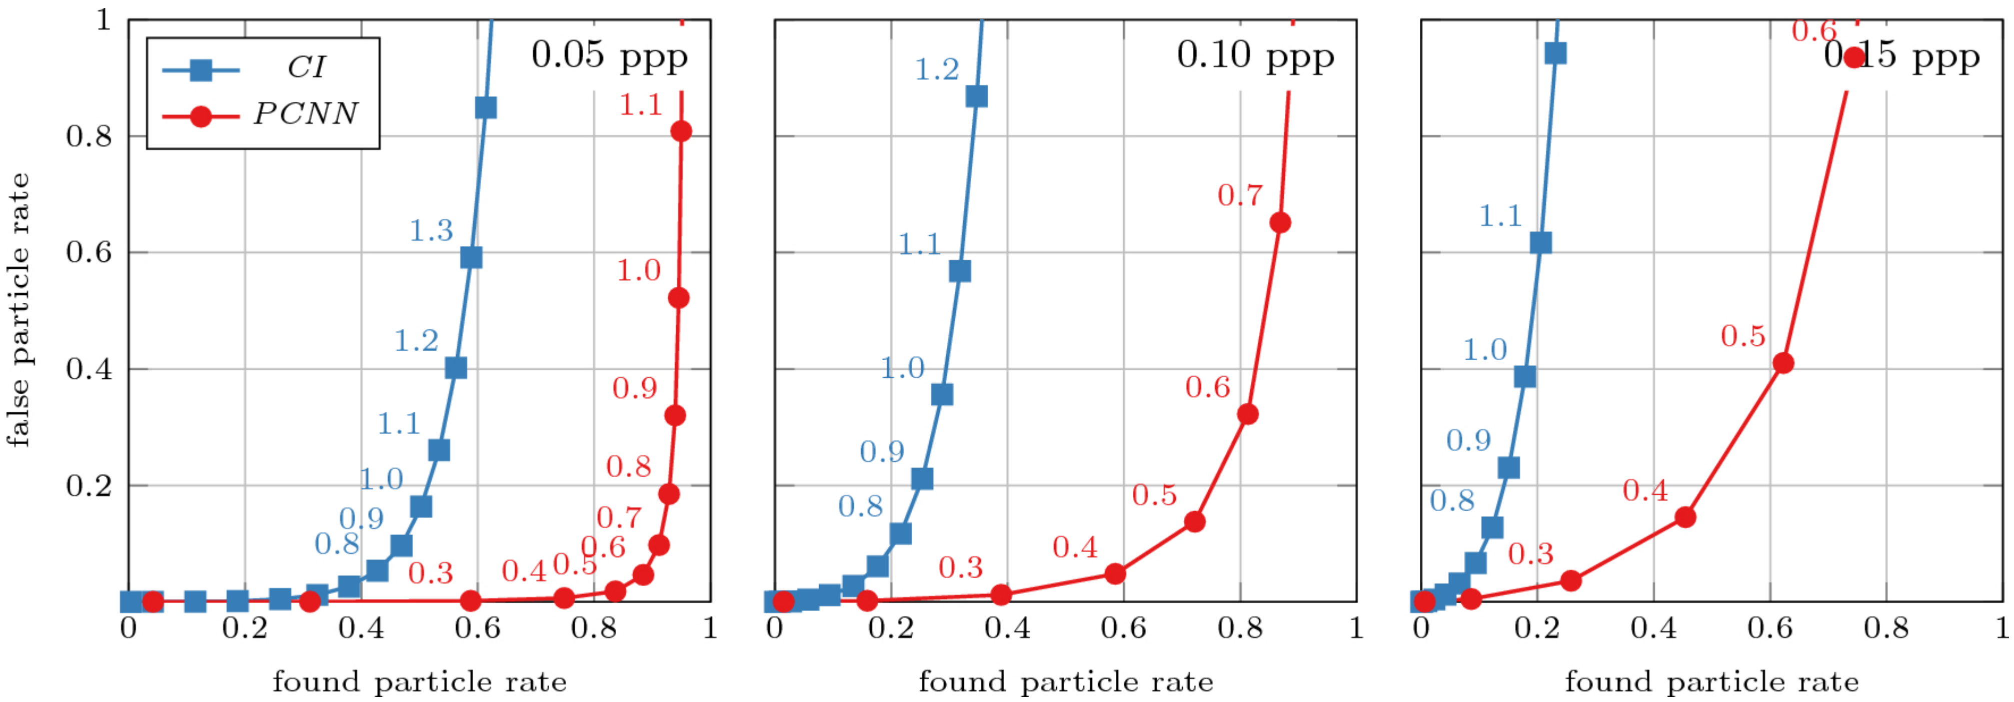

Godbersen P, Schanz D, Schröder A. (2024). Peak-CNN: improved particle image localization using single-stage CNNs. *Experiments in Fluids*, *65*(10), 153.

Evaluation on real world data

- 6-camera STB measurement data of Rayleigh–Bénard Cell using HFSB
- Reliable tracking of > 500,000 bubbles
- Backprojected 3d positions serve as pseudo ground truth data
- Closed volume: seeding density decreases as bubbles burst
  - 0.18 ppp to < 0.01 ppp (center)

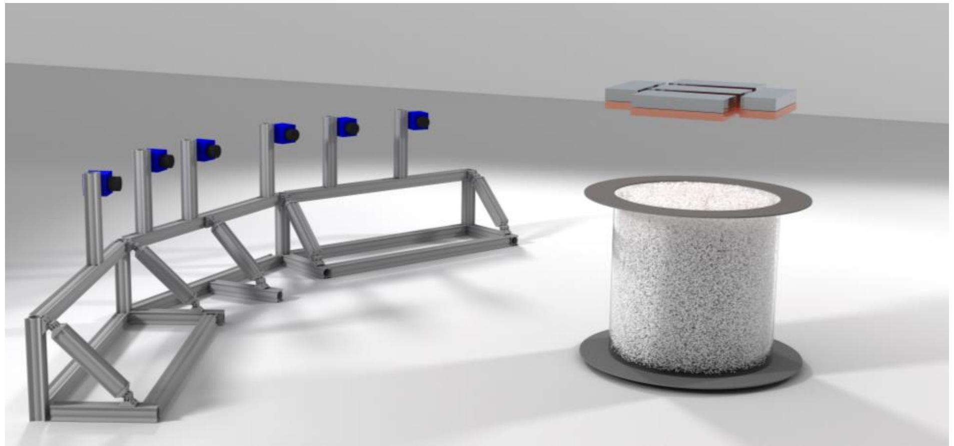

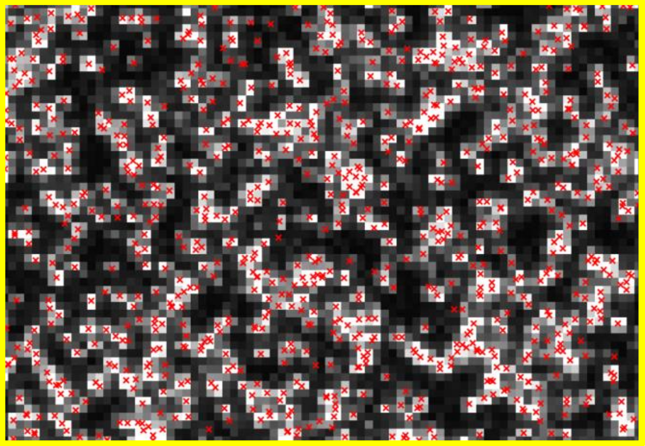

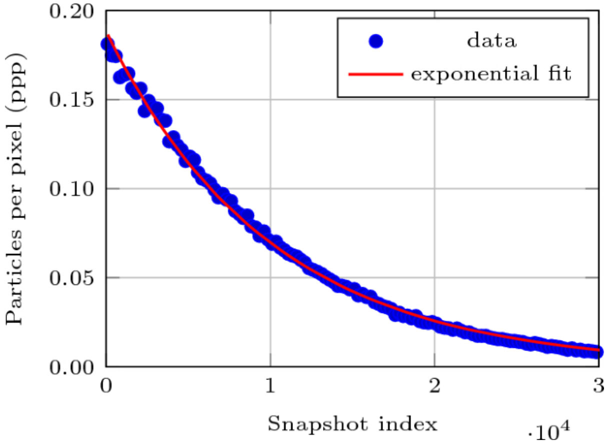

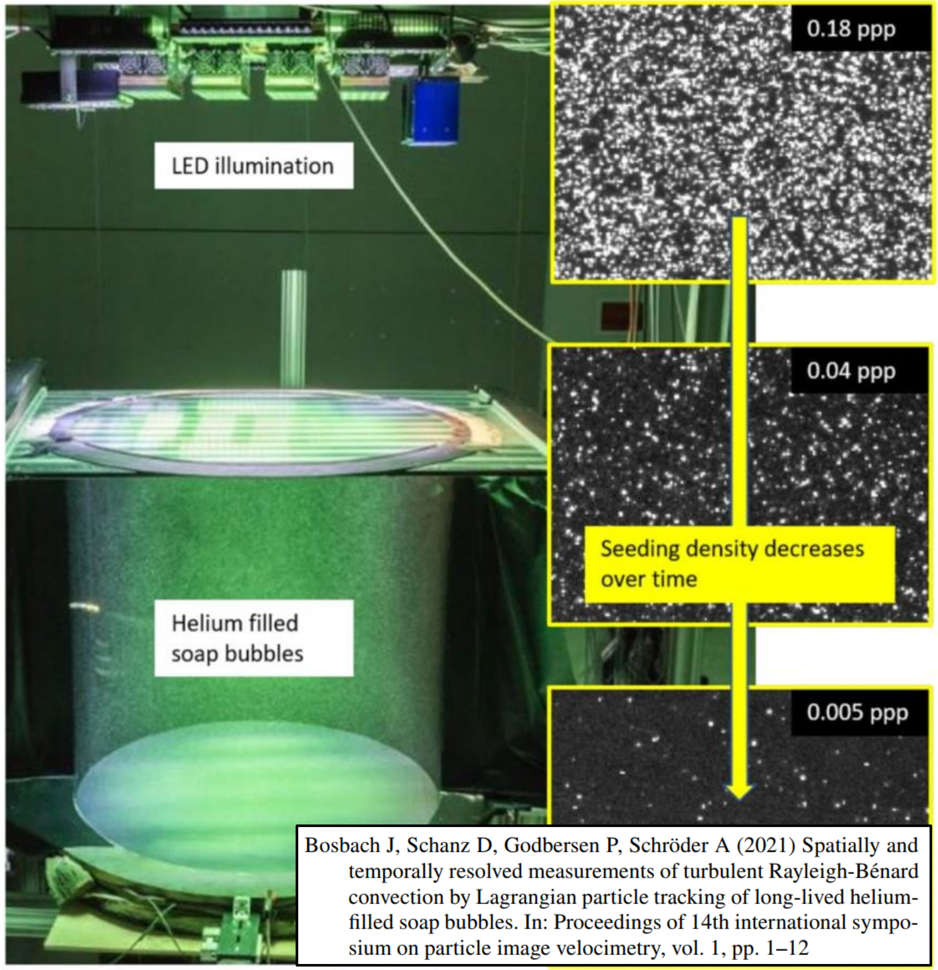

Bosbach J, Schanz D, Godbersen P, Schröder A (2021) Spatially and temporally resolved measurements of turbulent Rayleigh-Bénard convection by Lagrangian particle tracking of long-lived helium-filled soap bubbles. In: Proceedings of 14th international symposium on particle image velocimetry, vol. 1, pp. 1–12

Training data generation

- We currently train our model specifically for an individual experiment
- supervised learning approach → labeled training data needed. Two alternatives:

From existing particle tracks

- Backproject existing 3D tracks onto cameras
- Exploits high reliability of LPT measurements
- Chicken /Egg problem → apply at low ppp

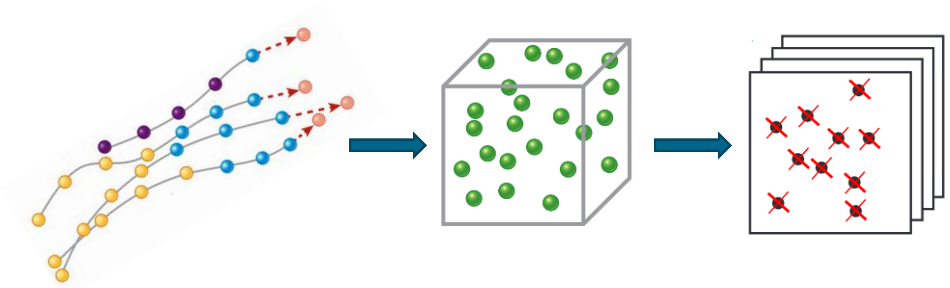

From conventional peakdetect

- Use conv. peak detector on lowly seeded images
- Training of ML-model on stacked images and detected peaks (e.g. 10-fold increase in image density)

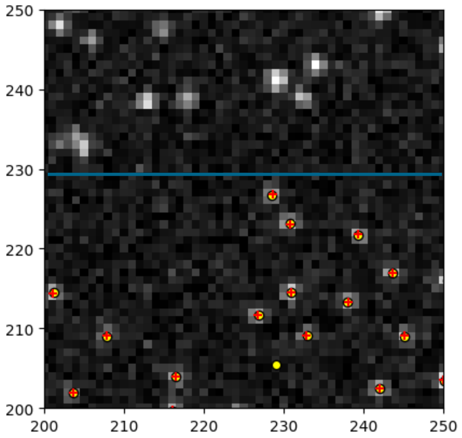

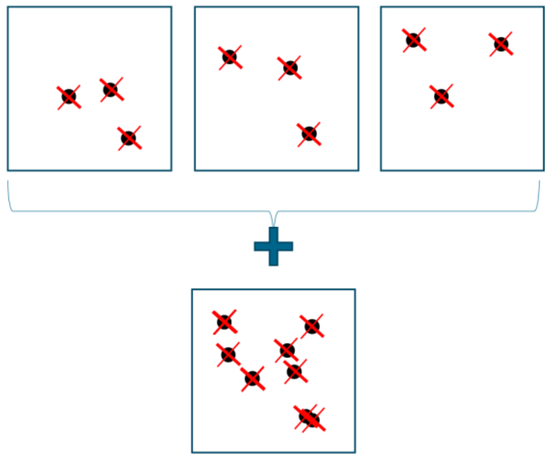

Performance of different labeling strategies

- Evaluation on RBC 0.12 ppp data
- Labels from tracking data give highest performance even if obtained at lower density
- Peak-only approach with stacking is slightly worse but still competitive
  - minimal user intervention required

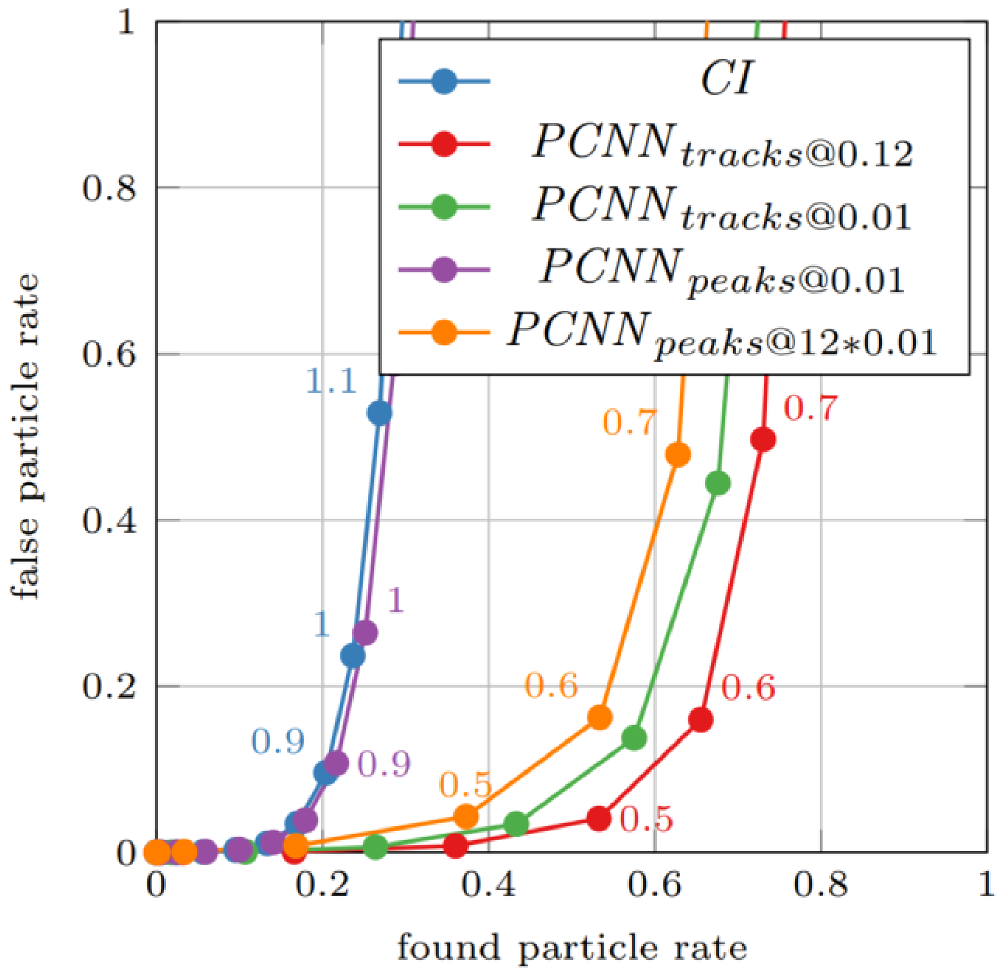

Godbersen P, Schanz D, Schröder A. (2024). Peak-CNN: improved particle image localization using single-stage CNNs. *Experiments in Fluids*, *65*(10), 153.

Peak-CNN: Conclusion and Outlook

- Peak-CNN clearly outperforms conventional approaches
- Significantly higher particle image densities can be handled
- Viable training strategies are available requiring only limited extra effort

- Strong gains expected for Two-Pulse, Four-Pulse and 2D Tracking methods

Godbersen P, Schanz D, Schröder A. (2024). Peak-CNN: improved particle image localization using single-stage CNNs. *Experiments in Fluids*, *65*(10), 153.

3D LPT evaluations: numerical cost and experiences with high-performance computing

- Sample Experiment:
  Geometry-induced separation of a TBL

- Performed as part of DLR project ADAMANT

- Setup in the large water tunnel at
  TU Braunschweig (GWB)

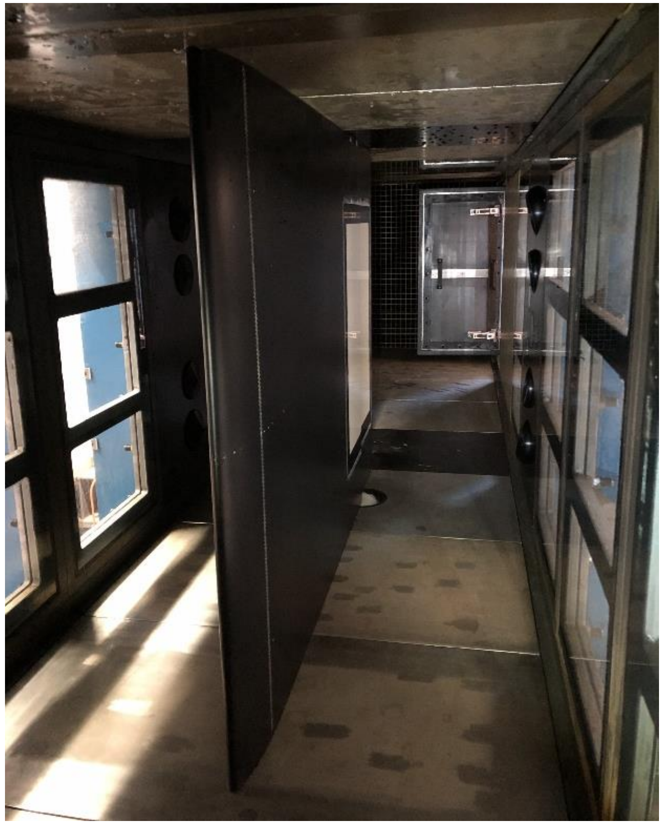

Details of backward-facing ramp experiment

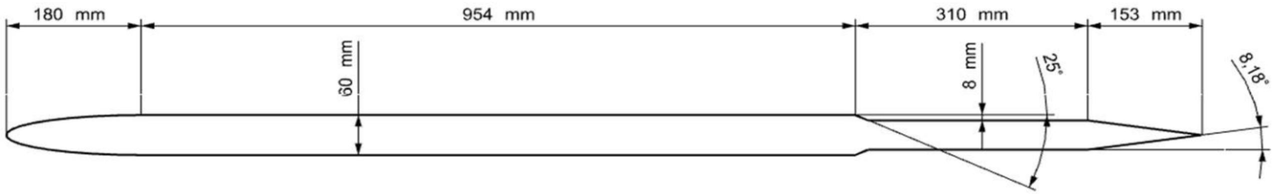

- Splitting plate (1.6 m length, 1.0 m height)

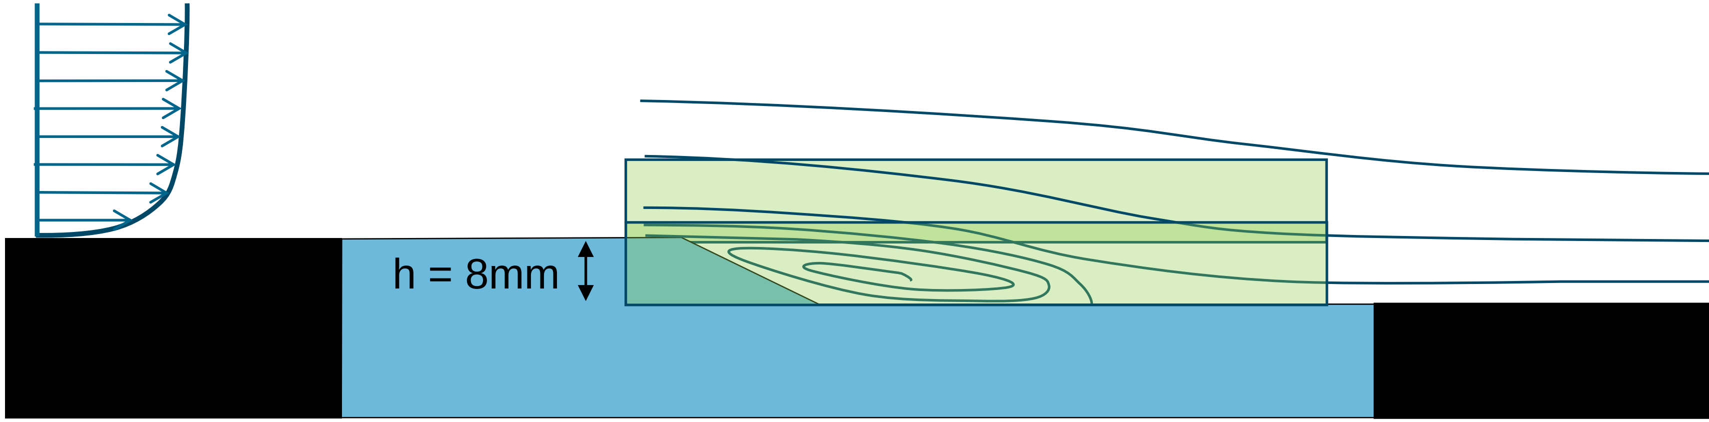

- Backward-facing ramp (25°) with incoming TBL
- Volumetric two-stage scanning laser illumination
  - Total volume: 90 × 90 × 16 mm^3 (streamwise × spanwise × wall-normal)
- Captured by five Phantom V2640 cameras

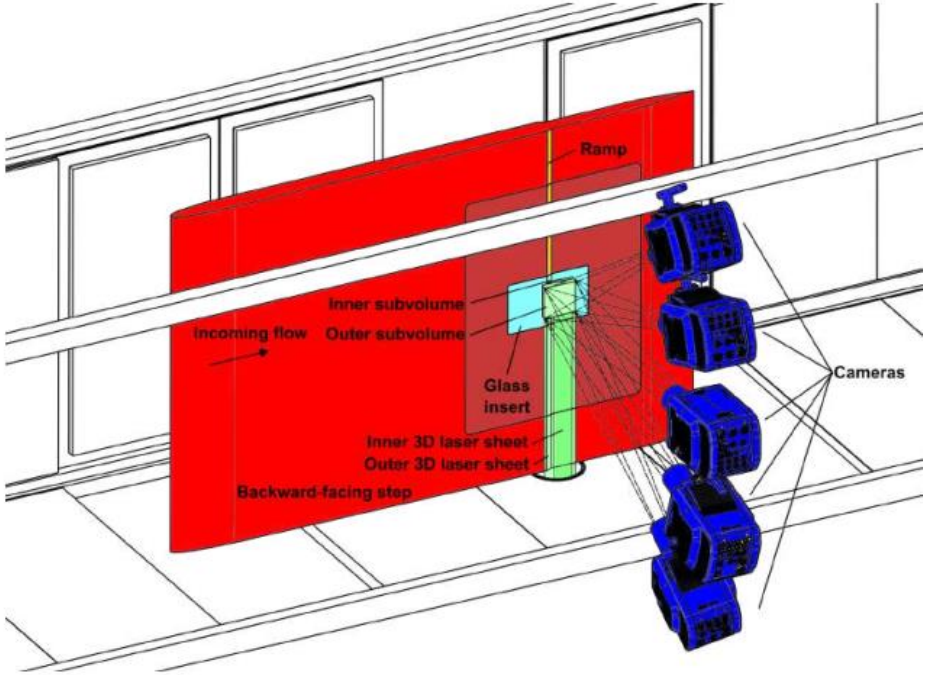

Details of BFR Experiment

- Four Reynolds-numbers (Re_x=1.2×10^6 – 2.7×10^6)
- recording rates 3.0 - 7.5 kHz per subvolume (6.0 - 15.0 kHz effectively)
- Two recording modes for each Re_x:
  - 3 fully time-resolved runs

    → 12.700 consecutive timesteps per subvolume and run: 76.200 3D reconstructions
  - 31 ‘chunked’ runs (for statistics)

    → 100 × 30 consecutive timesteps per subvolume and run: 186.000 3D reconstructions

Evaluation of BFR Experiment

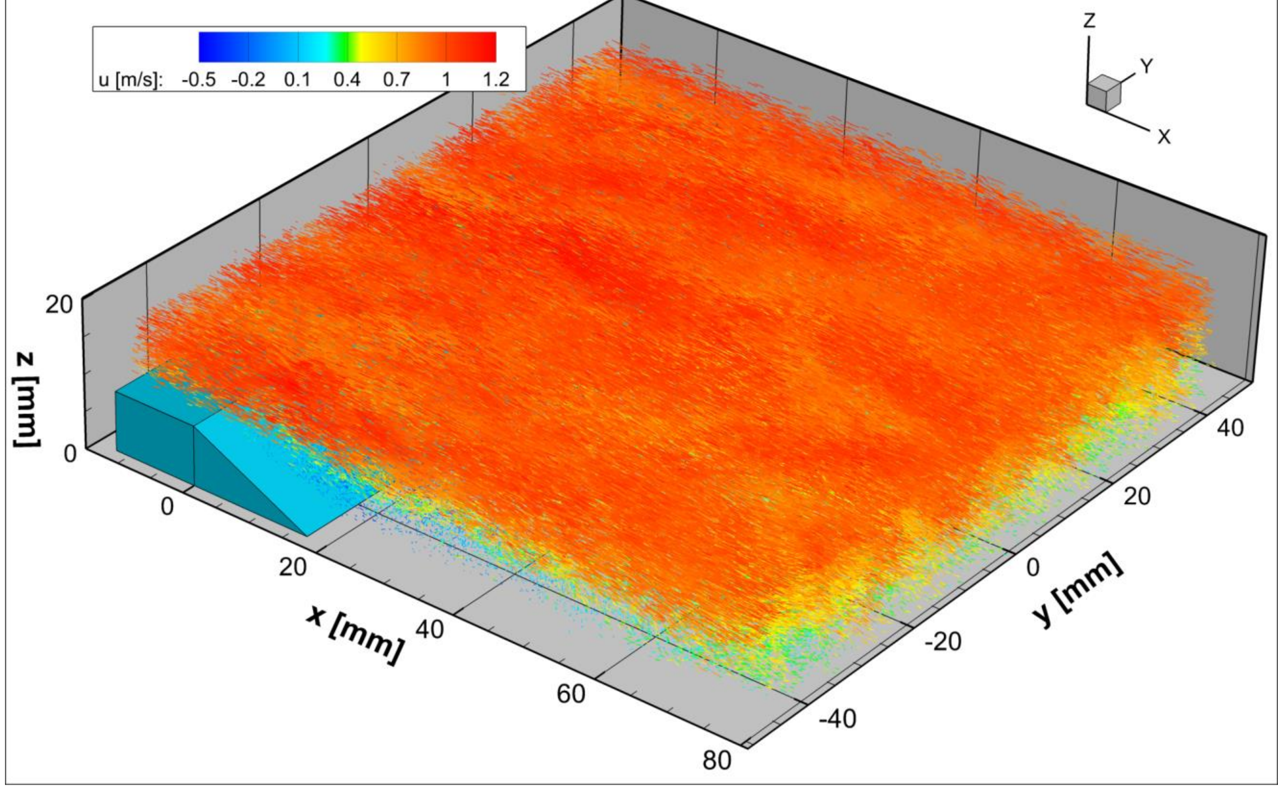

- STB Evaluation:
  - Separate STB processing for each of the two subvolumes
  - Multi-pass Variable-Time-Step processing (4 passes for lower volume, 2 for upper)
  - Tracking of ~90.000 and ~160.000 particles (lower and upper volume)
- Evaluation times ‚at home‘ on Ryzen 3950X (16 cores, 96GB RAM)
  - STB processing for each time-step and subvolume in converged state: **~3s**
  - Overhead for multi-pass processing and sum over subvolumes: **~20s** per time-step
  - Total evaluation time for each Re_x: around **39 days** (on single machine).

HPC Evaluation of BFR Experiment

- *CARO*: DLR HPC Cluster
- Located at GWDG in Göttingen
- 1364 Nodes
  - 2*64-cores (AMD EPYC 7702)
  - 256 GB Ram

18 Racks

1.364 Compute Nodes

175,744 CPU Cores

6 GPU nodes

100 Gbit/s Interconnect

3.46 PetaFlop/s

364 TB RAM

8,4 PiB Storage

- Time-resolved runs: All 12 parallely processed on single node
  - Evaluation time: ~ 3.5 days (compared to 35 days at home)
- Chunked runs: 20 jobs per node, 150 nodes started in parallel
  - All chunks for single Re_x evaluated instantaneously
  - Evaluation time per time-step approx. doubled (40 s)
  - Full processing finished in ~ **20 minutes**

Binning results

streamwise velocity

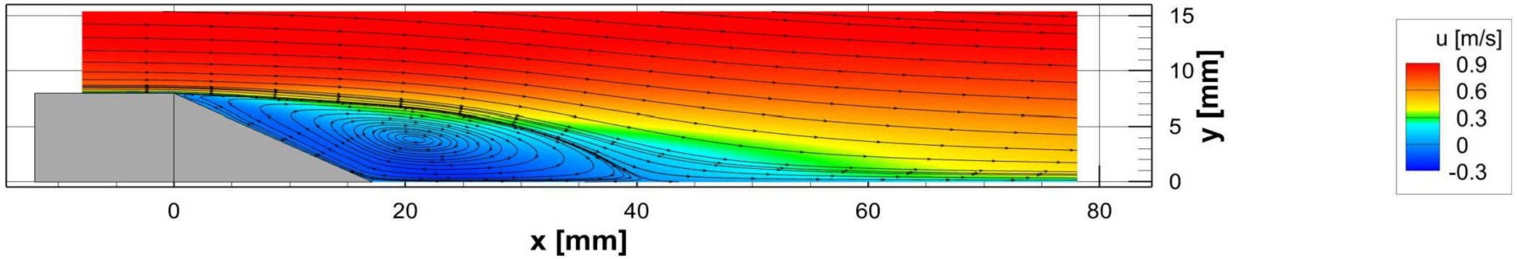

- Binning in 400 × 20 µm bins
- All data processed within ~ 2 minutes
- Quick access to averaged
  - Velocities
  - Acclerations
  - Reynolds stresses
  - Triple correlations
- HPC: Vastly improved turn-around time

  → easier identification of problems

Reynolds stresses

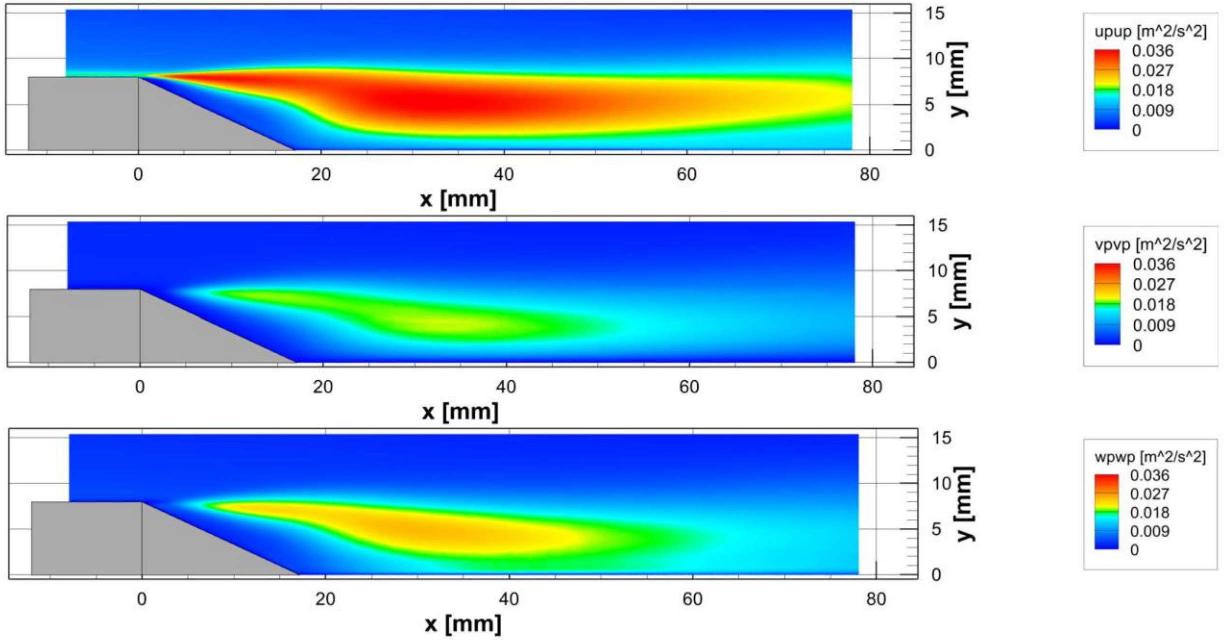

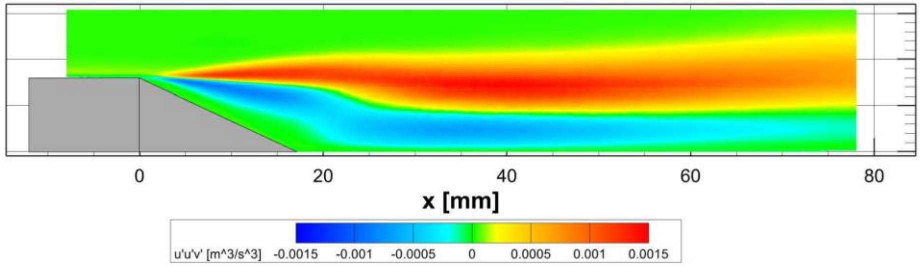

Triple correlations (here u‘u‘v‘)

Benchmarking Flowfield interpolation schemes

From discrete particle velocities and accelerations…

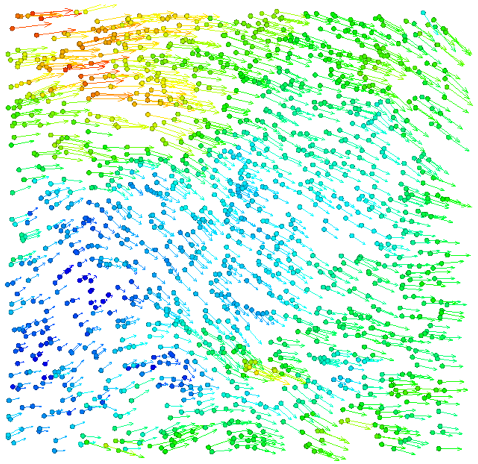

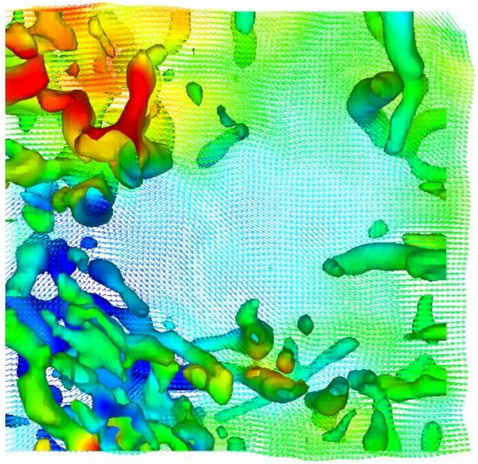

…to continous or gridded field variables

Ke Zhou[1], Sam Grauer[1], Daniel Schanz[2], Philipp Godbersen[2], Andreas Schröder[2,3], Thomas Rockstroh[4], Young Jin Jeon[4], Bernhard Wieneke[4]

[1] Pennsylvania State University, USA
[2] German Aerospace Center (DLR), Germany
[3] Brandenburgische Technische Universität, Germany
[4] LaVision GmbH, Germany

Zhou, K., Grauer, S.J., Schanz, D., Godbersen, P., Schröder, A., Rockstroh, T., Jeon, Y.J. and Wieneke, B., 2024. Benchmarking data assimilation algorithms for 3D Lagrangian particle tracking. Lisbon Laser Conference

Motivation

- Why interpolate?:
  - Getting access to Flowfield properties (i.e. the velocity gradient tensor, pressure, and more)
  - DA methods: Enhancing the spatial (and/or temporal) resolution of a measurement by feeding physical knowledge into the evaluation

- Aim here: compare the newest approaches on synthetic and experimental data

- Contenders:
  - Binning (baseline approach without regularization)
  - VIC# 3D and 4D
  - FlowFit 3
  - Physics-informed neural networks (PINNs; here: NIPA)

Methods *NIPA*

Physics-informed neural network *Raissi et al (2020)*

- Flow represented by a neural network
- Network parameters optimized to match LPT data, via a measurement operator
- Governing equations as soft physical constraints

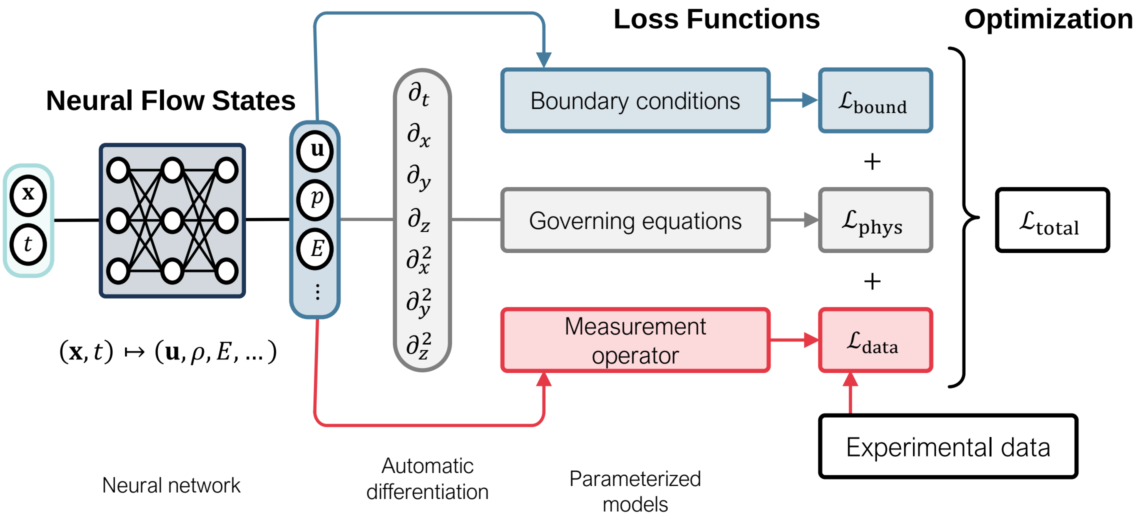

Neural-implicit particle advection *Zhou and Grauer (2023)*

- Dedicated particle model with LPT tracks as hard constraints
- Statistical treatment of noisy tracks for robust reconstruction
- Solves both Navier-Stokes and Maxey-Riley equations for fluids and particles
- Ability to handle nonideal tracers, (ballistic or buoyant particles)

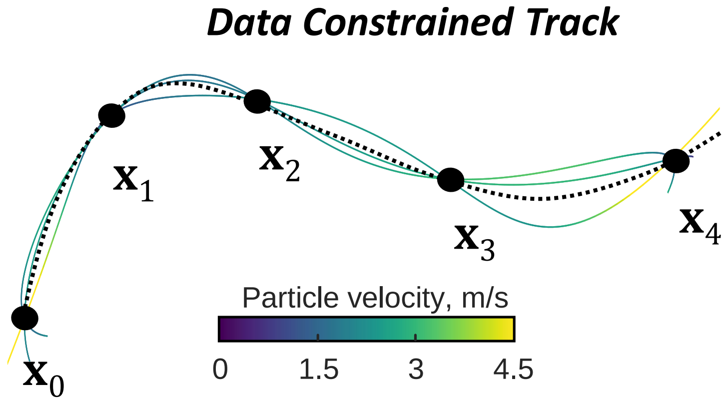

Raissi et al. **Hidden fluid mechanics:Learning velocity and pressure fields from flow visualizations** 2020 science 367 6481
Zhou & Grauer **Flow reconstruction and particle characterization from inertial Lagrangian tracks** 2023. *arXiv preprint*, 2311.09076

Methods *FlowFit 3*

Godbersen et al. 2024

Velocity field representation using **uniform B-splines**

Optimization of coefficients using L-BFGS

Staggered grid for hard constraint on divergence via orthogonal projection of gradients

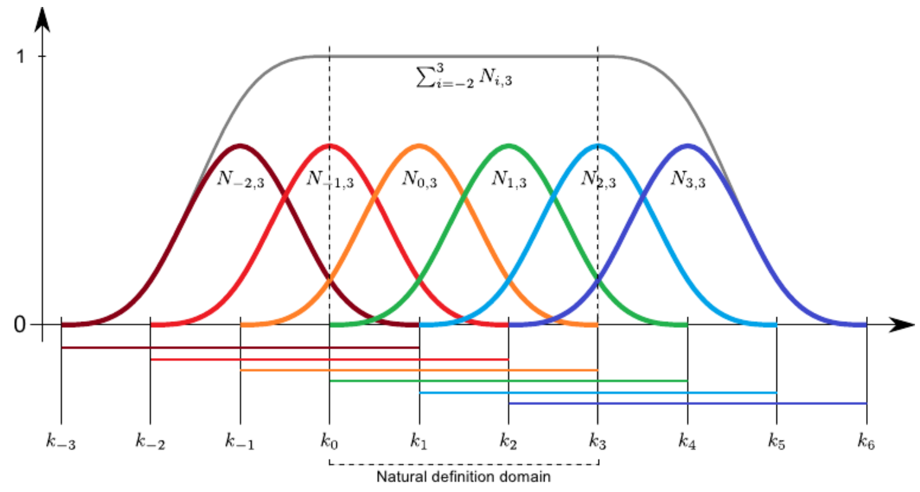

Linear Mode

- Uses **velocity** of particles from TrackFit
- L-BFGS on **velocity coefficients** to minimize **velocity reconstruction error**
- Physics: Continuity as **hard constraint** $\nabla \cdot \vec{u} = 0$

Nonlinear Mode

- Uses **velocity and acceleration** of particles from TrackFit
- L-BFGS on **velocity coefficients** to minimize **velocity and acceleration reconstruction error**
- Physics: Continuity as hard constraint, Pressure via **Poisson equation**

$$\nabla \cdot \vec{u} = 0 \qquad \Delta p + \nabla \cdot (\vec{u} \cdot \nabla \vec{u}) = 0 \qquad \vec{a} = \frac{D\vec{u}}{Dt} = -\nabla \bar{p} + \nu \Delta \vec{u}$$

Godbersen et al. 2024 **FlowFit 3: Efficient data assimilation of LPT measurements**. *Lisbon Symposium*

Methods VIC#

Vortex-in-Cell-plus (VIC+) *Schneiders and Scarano (2016)*

- Optimization of $\boldsymbol{\omega}$, $\mathbf{u}_{\partial\Omega}$ and $\partial\mathbf{u}/\partial t|_{\partial\Omega}$
- Calculation of $\mathbf{u}_{\Omega}$ and $\partial\mathbf{u}/\partial t|_{\Omega}$
- Cost function: disparity between VIC+ and STB

$$\nabla^2\mathbf{u} = -\nabla\times\boldsymbol{\omega} \text{ and } \nabla^2\frac{\partial\mathbf{u}}{\partial t} = -\nabla\times\left((\boldsymbol{\omega}\cdot\nabla)\mathbf{u} - (\mathbf{u}\cdot\nabla)\boldsymbol{\omega}\right)$$

$$J_{+} = \sum_n \left\|\mathbf{u}_n - \mathbf{u}_{\mathrm{STB},n}\right\|^2 + \alpha^2\sum_n\left\|\frac{\mathrm{D}\mathbf{u}}{\mathrm{D}t}\bigg|_n - \frac{\mathrm{D}\mathbf{u}}{\mathrm{D}t}\bigg|_{\mathrm{STB},n}\right\|^2$$

$$J_{\#} = J_{\mathbf{u}} + \alpha_{1\#}^2 J_{\mathrm{D}\mathbf{u}/\mathrm{D}t} + J_{\mathrm{G.E.}}$$

Vortex-in-Cell-sharp (VIC#) *Jeon et al. (2022)*

- **Additional constraints**
  - Proper correction of boundary
  - Stable convergence
  - Consideration of pressure
- **Coarse grid approximation**

$$J_{\#\mathrm{TM}}(D) = \sum_{t=t_0-D}^{t_0+D}\left[\, J_{\#}(t)\,\right]$$

$$J_{\mathrm{G.E.}} = f_{\mathrm{AC}}\left(J_{\mathrm{div}(\boldsymbol{\omega})} + J_{\mathrm{div}(\mathbf{u})} + \alpha_{2\#}^2\left(J_{\mathrm{div}(\partial\mathbf{u}/\partial t)} + f_{\mathrm{L}}\ldots(t-t_0)/\Delta t_{\mathrm{PTV}}\ldots\right) + f_{\mathrm{DN}}^2 J_{\mathrm{so}\ldots}\right)$$

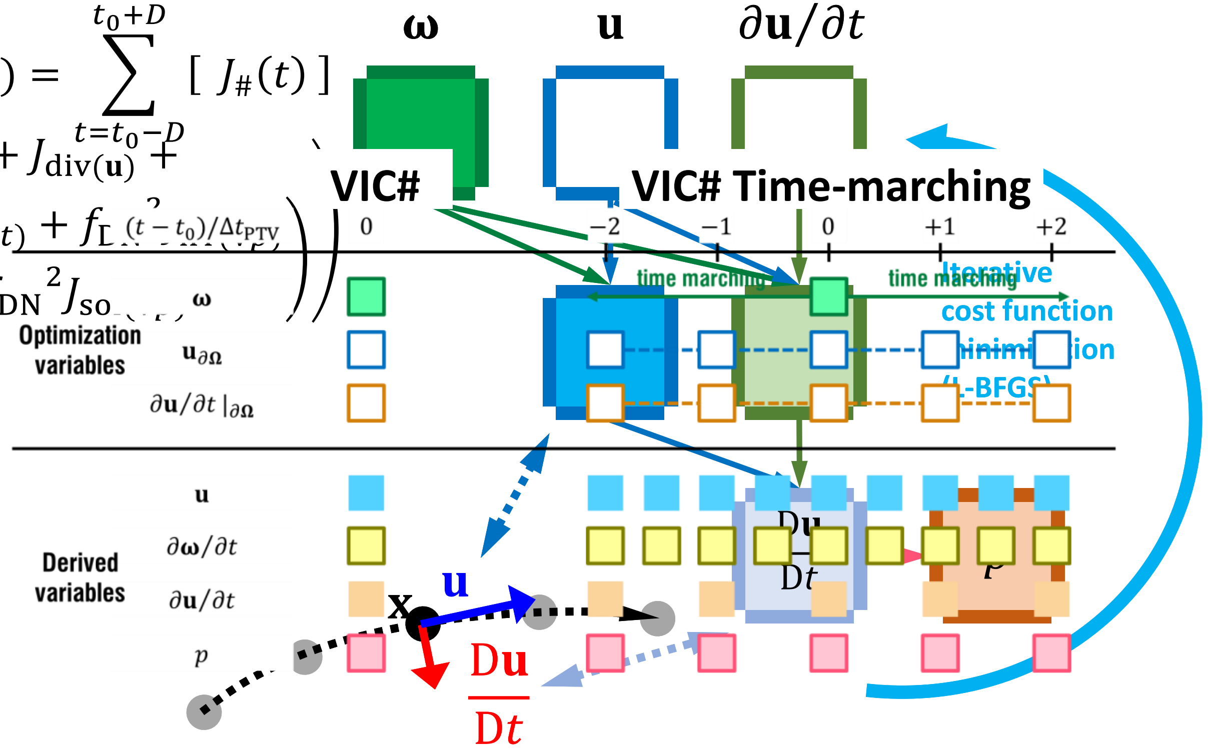

Extension to 4D *Jeon et al. (2023)*

- **Temporal evolution of flow field**
  - $\boldsymbol{\omega}(t+\Delta t) = f(\boldsymbol{\omega}(t),\ \dot{\boldsymbol{\omega}}(t),\ \ddot{\boldsymbol{\omega}}(t))$
  - $\dot{\boldsymbol{\omega}} = (\boldsymbol{\omega}\cdot\nabla)\mathbf{u} - (\mathbf{u}\cdot\nabla)\boldsymbol{\omega} + \nu\nabla^2\boldsymbol{\omega}$
  - $\ddot{\boldsymbol{\omega}} = (\dot{\boldsymbol{\omega}}\cdot\nabla)\mathbf{u} + (\boldsymbol{\omega}\cdot\nabla)\dot{\mathbf{u}} - (\dot{\mathbf{u}}\cdot\nabla)\boldsymbol{\omega} - (\mathbf{u}\cdot\nabla)\dot{\boldsymbol{\omega}} + \nu\nabla^2\dot{\boldsymbol{\omega}}$
- **Time integration via 2nd order Taylor expansion**

Schneiders and Scarano 2016 **Dense velocity reconstruction from tomographic PTV with material derivatives** 2016 *Exp Fluids* 57 9
Jeon et al. 2022 **Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+**. *Exp Fluids*, *63* 70
Jeon. 2023 **Vorticity time-marching method in Fine scale reconstruction (VIC#) for describing 4-D space-time**. ISPIV'23, San Diego, USA

Test cases *Flow configurations*

Synthetic HIT

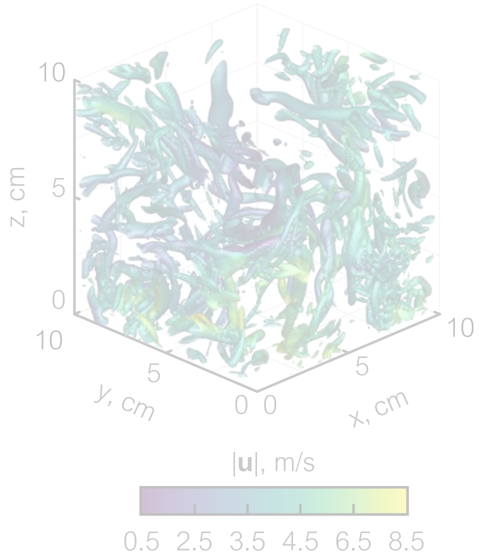

- JHTDB Perlman et al. (2007)
- Homog. Isotr. Turb. $Re_\lambda = 433$
- Artificial forcing term
- Particle spacing: 6l_η
- Particles: RK4 integration
- 50 time-steps
- ~105,000 particles

Synthetic TBL

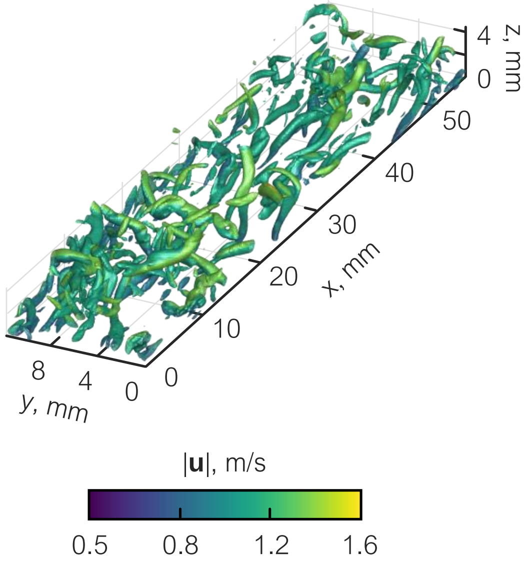

- JHTDB Perlman et al. (2007)
- turbulent boundary layer at $Re_\tau = 1000$
- Particle spacing: 6l_ν
- Particles: RK4 integration
- 50 time-steps
- ~276,000 particles

Experimental TBL

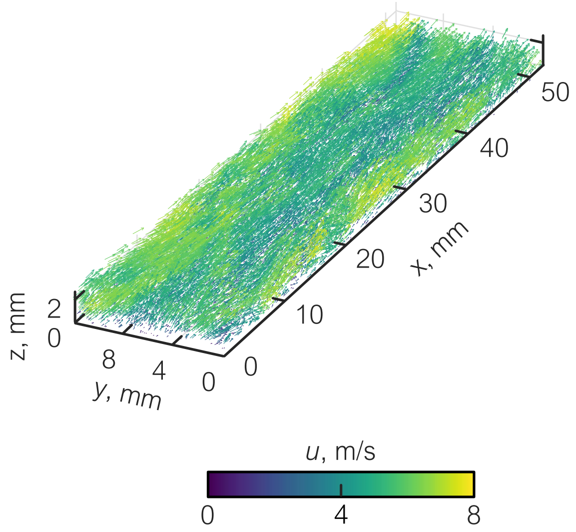

- DLR 1MG measurements Schröder et al. (2024)
- Zero-pressure-gradient turbulent boundary layer at $Re_\tau = 995$
- Particle spacing: ~7l_ν
- Particles: TrackFit Gesemann et al. (2016)
- 50 time-steps
- ~ 54,000 particles

Test cases *Parameter space*

Particle density

- From the initial number, the particle density is halved several times:
- Downsamplig factor: $k_{\mathrm{x}} = 2^{n}$ with $n = 1 \dots 6$
  - [1, ½, ¼, $1/8$, 1/16, 1/32, 1/64]
- Noise-free

Noise variations

- Noise is added to the particle positions.
- Standard case: $\sigma_{\mathrm{x}} = \sigma_{\mathrm{y}} = 0.01\ mm$ (~0.1 px) $\sigma_{\mathrm{z}} = 2 \cdot \sigma_{\mathrm{x}}$
- Noise variations: $\sigma_{\mathrm{z}} = [2, 4, 6, 8, 10] \cdot \sigma_{\mathrm{x}}$
- For these: k_{x} = 4, dt = 1

dt variations

- Downsampling of the temporal resolution
- Downsampling factor: $k_{\mathrm{t}} = 1, 2, 4, 6, 8$
- k_{x} = 4
- Noise-free

Results *Synthetic TBL*

Downselection factor

$k_x = 1$
276,000

Binning

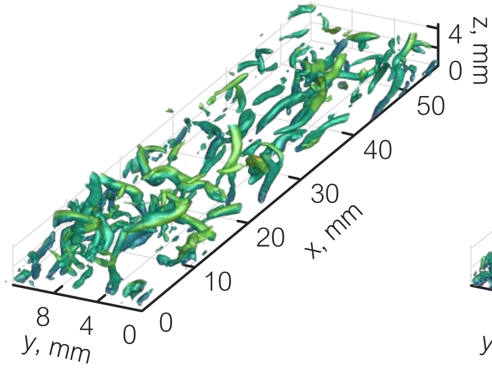

FF-nonlinear

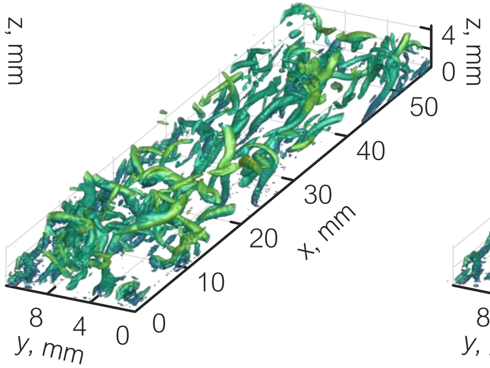

VIC#-4D

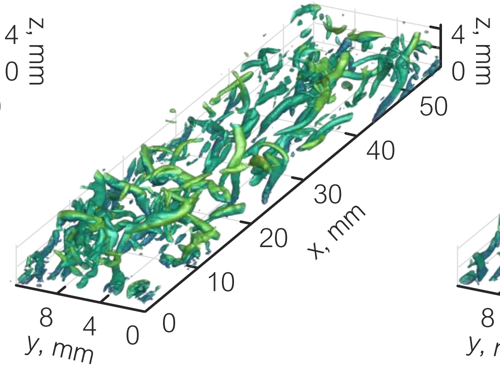

NIPA

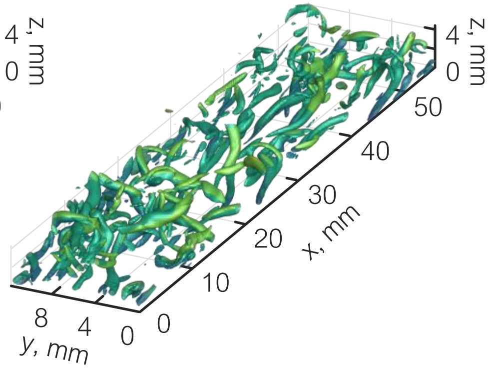

Exact

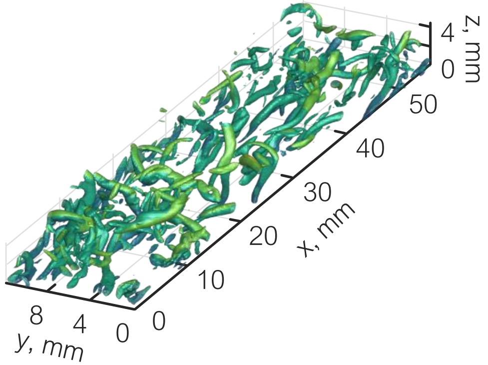

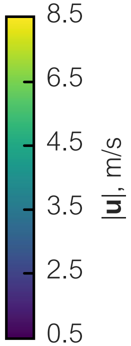

Results *Synthetic TBL*

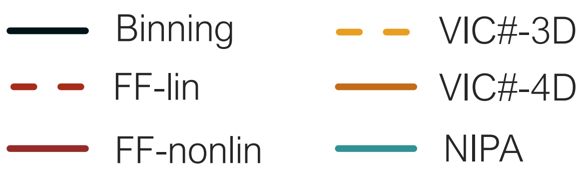

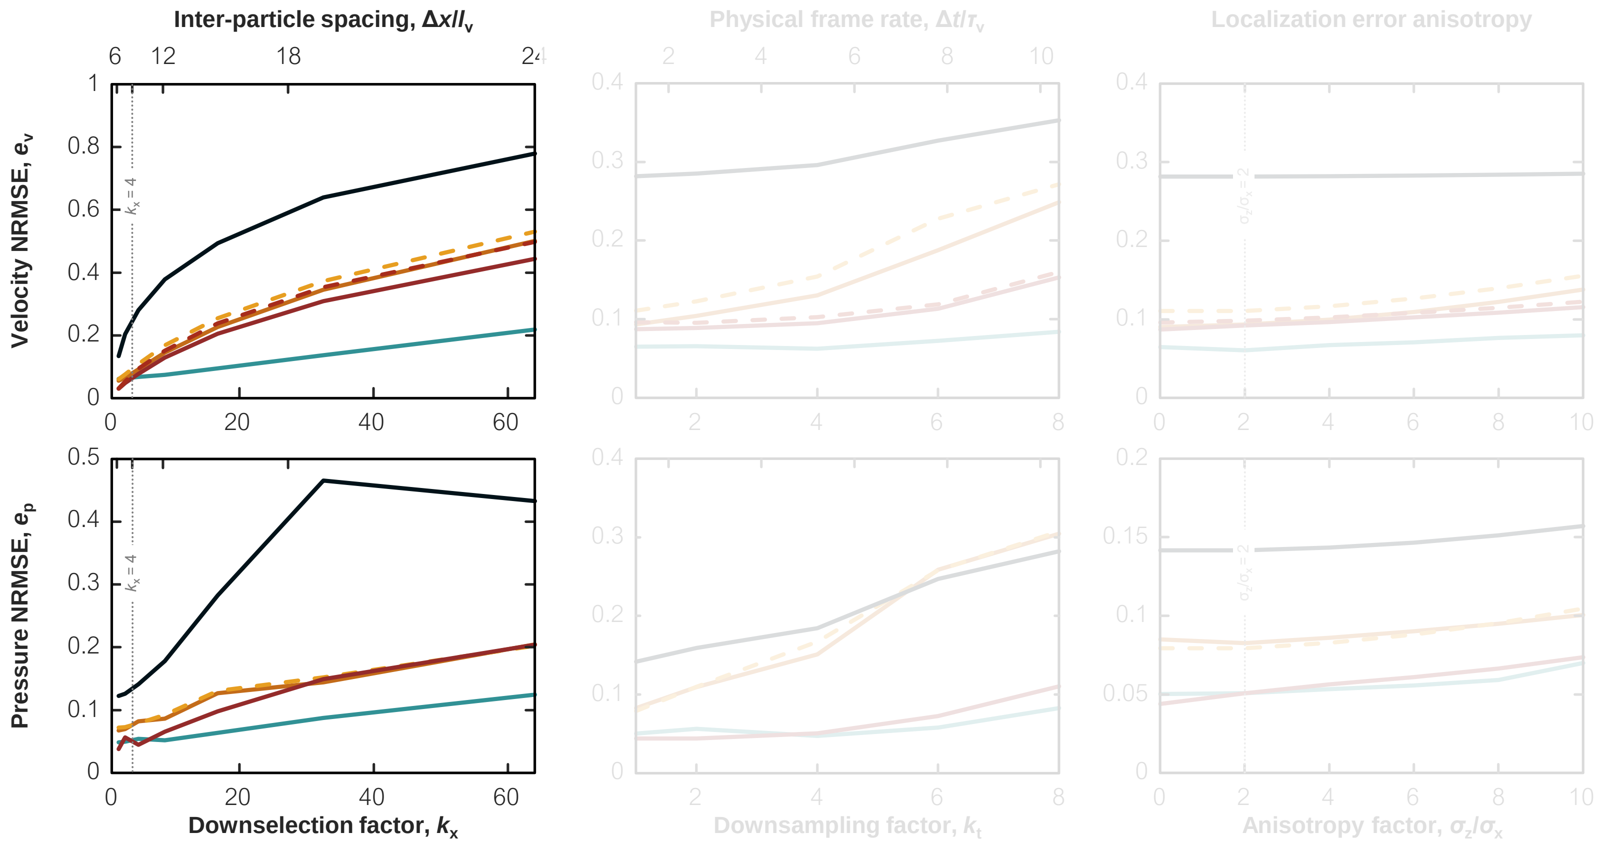

- At high particle densities, all DA methods perform comparably
- Binning shows in all cases the highest errors
- All methods suffer when decreasing particle
- FF3 linear performs comparably to VIC#-4D, FF3 non-linear is slightly better
- NIPA shows the least dependence on particle number, both for velocity and pressure

- NIPA and FF show little dependence on temporal sampling
- Worse results for VIC# (likely due to less accurate track filtering)
- Noise variations show little effect on all methods

Results *Synthetic TBL*

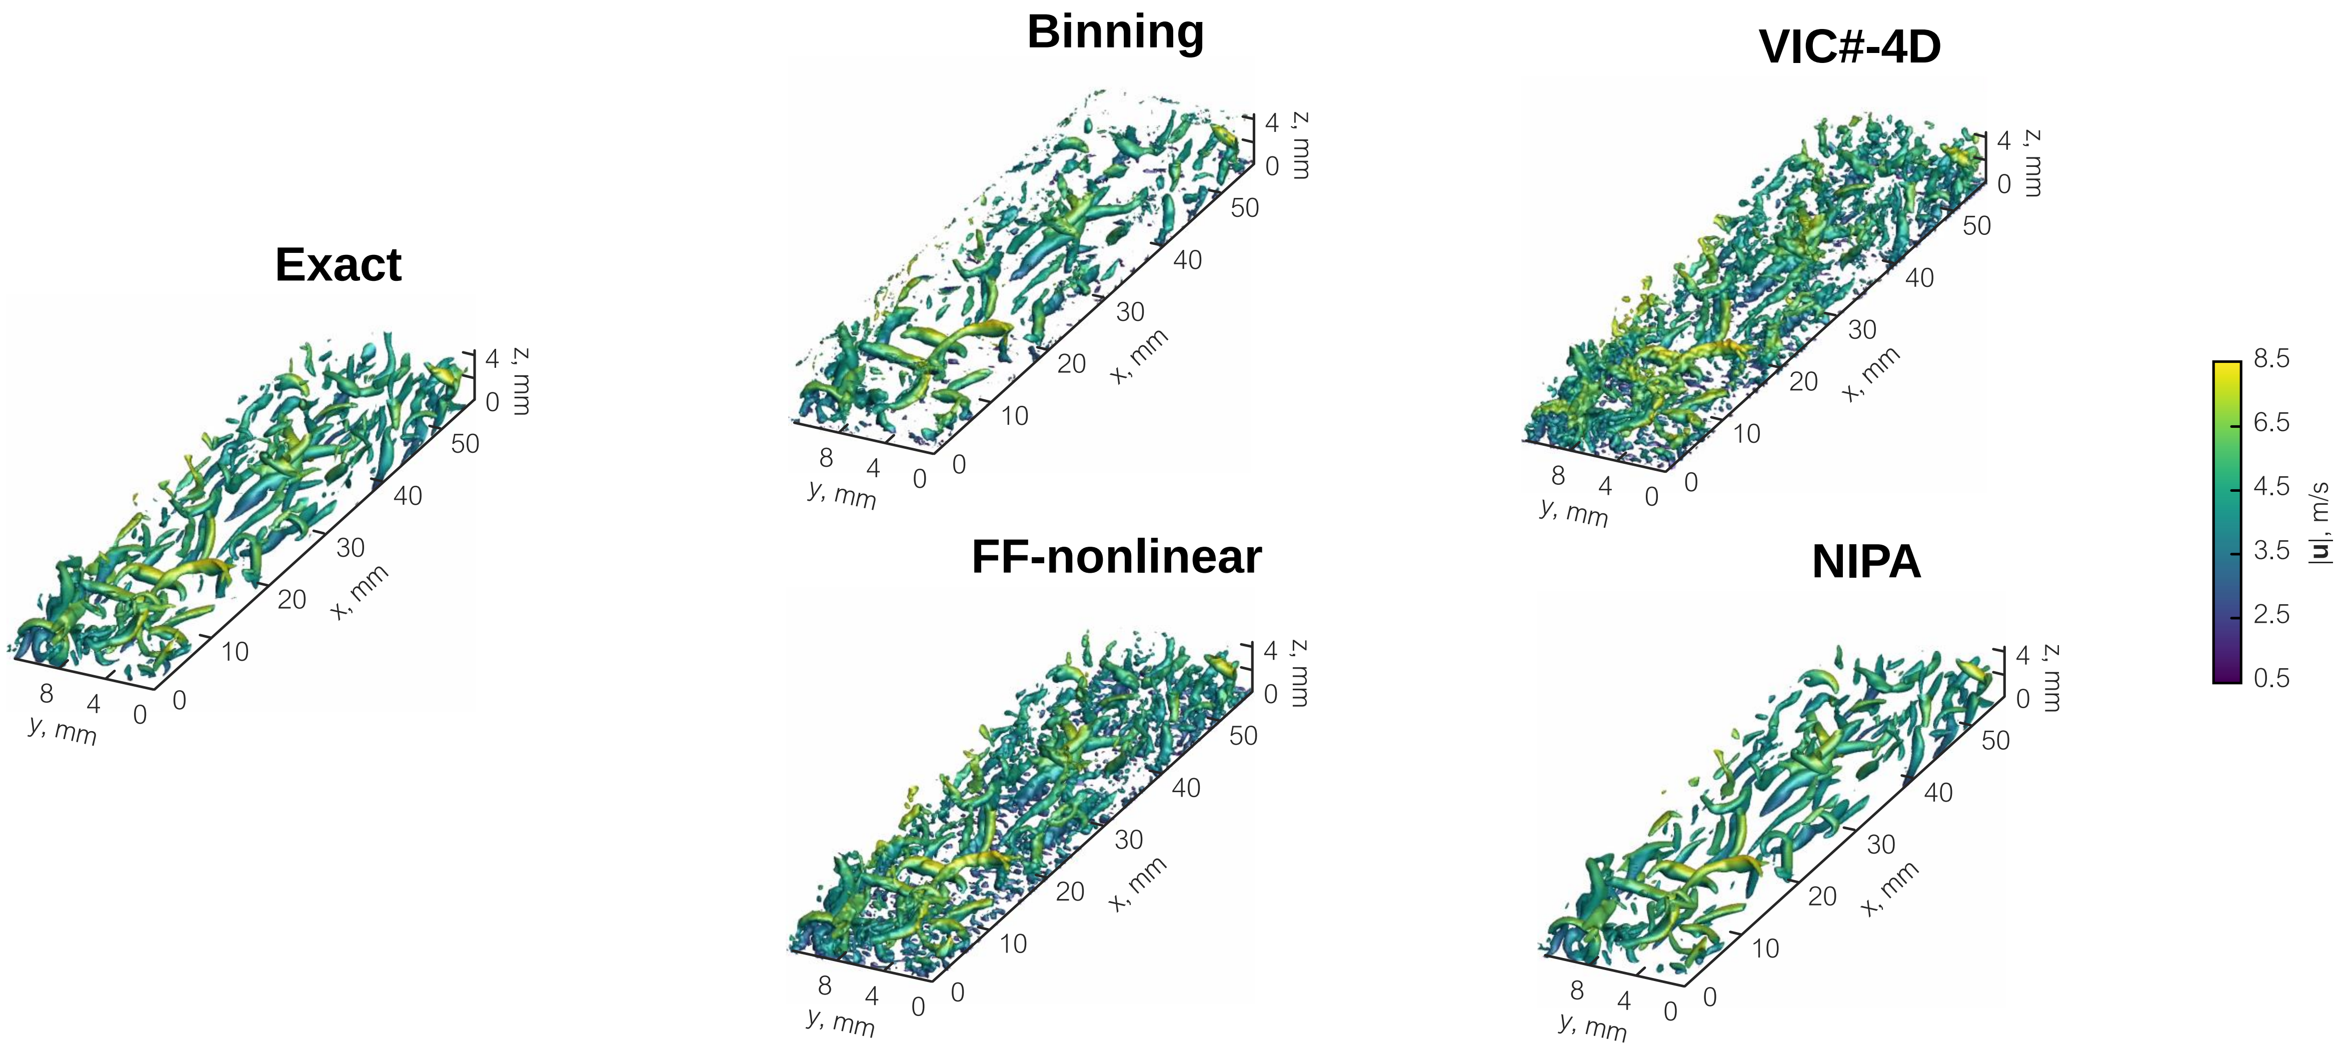

- Animations over all 50 time-steps for $k_{\mathrm{x}} = 8$

Results *Experimental TBL*

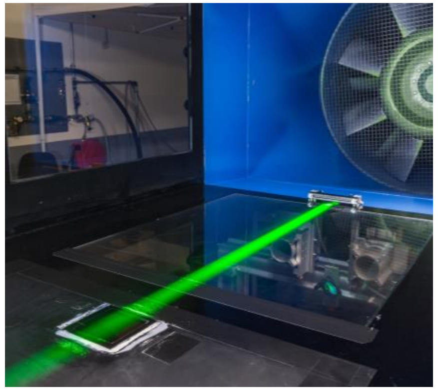

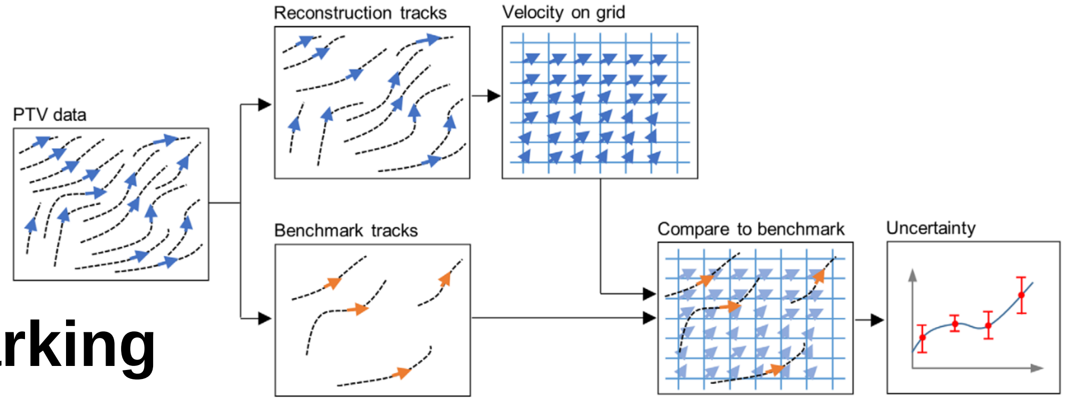

Schneiders & Sciacchitano (2017)

Validation:
Track benchmarking

separate track data in two groups

- 80% tracks for data assimilation (~43.200 particles at k_x=1)
- 20% tracks for validation (~10.800 positions)

Results *Experimental TBL*

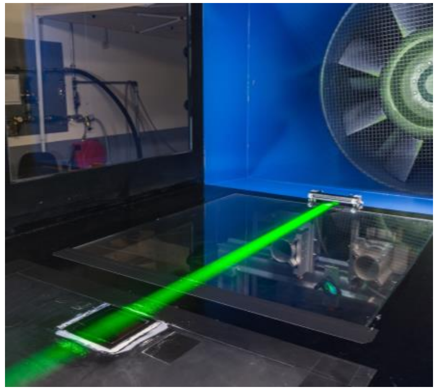

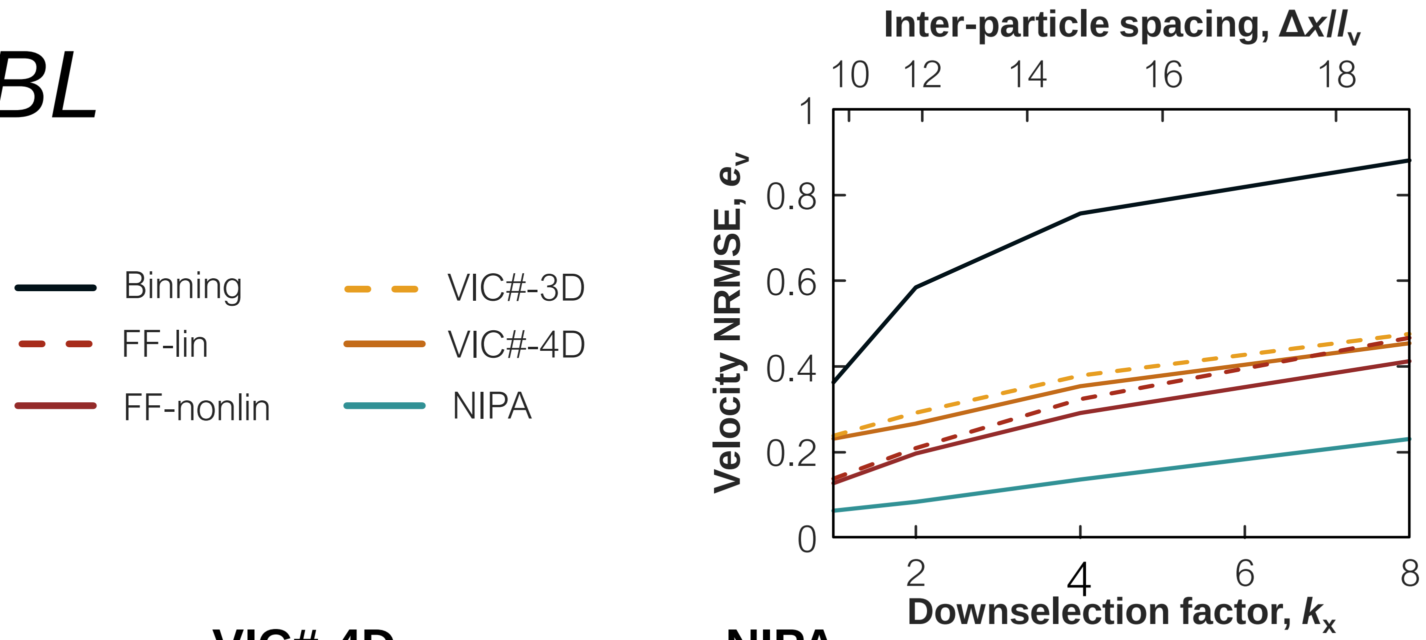

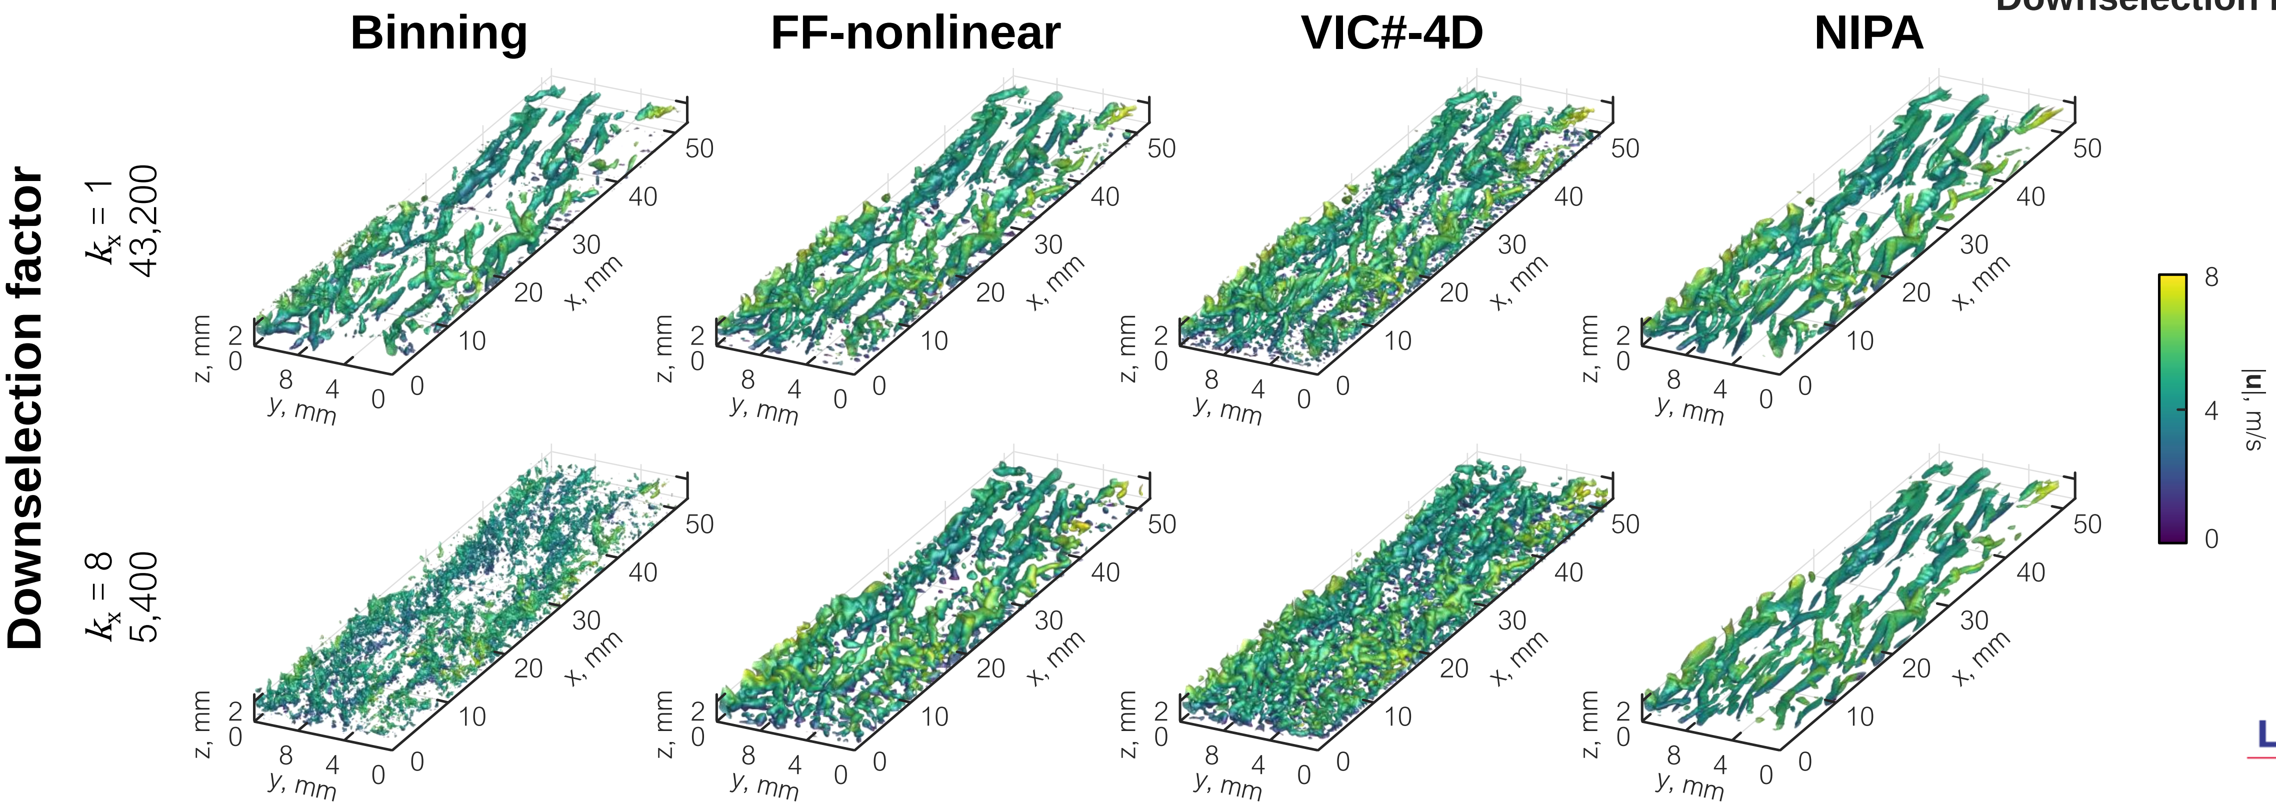

PennState

DLR

LAVISION
FOCUS ON IMAGING

Results *Computational cost*

Reconstruction time in seconds per time-step (k_x = 4)

	Binning	VIC#-3D	VIC#-4D	FF-linear	FF-nonlinear	NIPA
HIT	13	393	868	10	73	1420
TBL	6	174	585	6	22	850
	CPU Xeon w7-3465X (28 cores)	GPU RTX 4090	GPU Tesla V100	CPU Ryzen 5950X (16 cores)	CPU Ryzen 5950X (16 cores)	GPU RTX A6000

- 4D methods pay a signifcant computational toll
- Training of PINN within NIPA takes the longest time
- FlowFit 3.0 is very fast due to inherent constraint to divergence-free solutions and optimizing only the velocity coefficients

2nd LPT and DA Challenge: LIVE now!

- Organized by Onera,
  TU Delft, Penn State University and DLR

- Synthetic and experimental test cases for LPT and DA

- Data availabe since Jan. 31

- Deadline: Apr. 25 2025

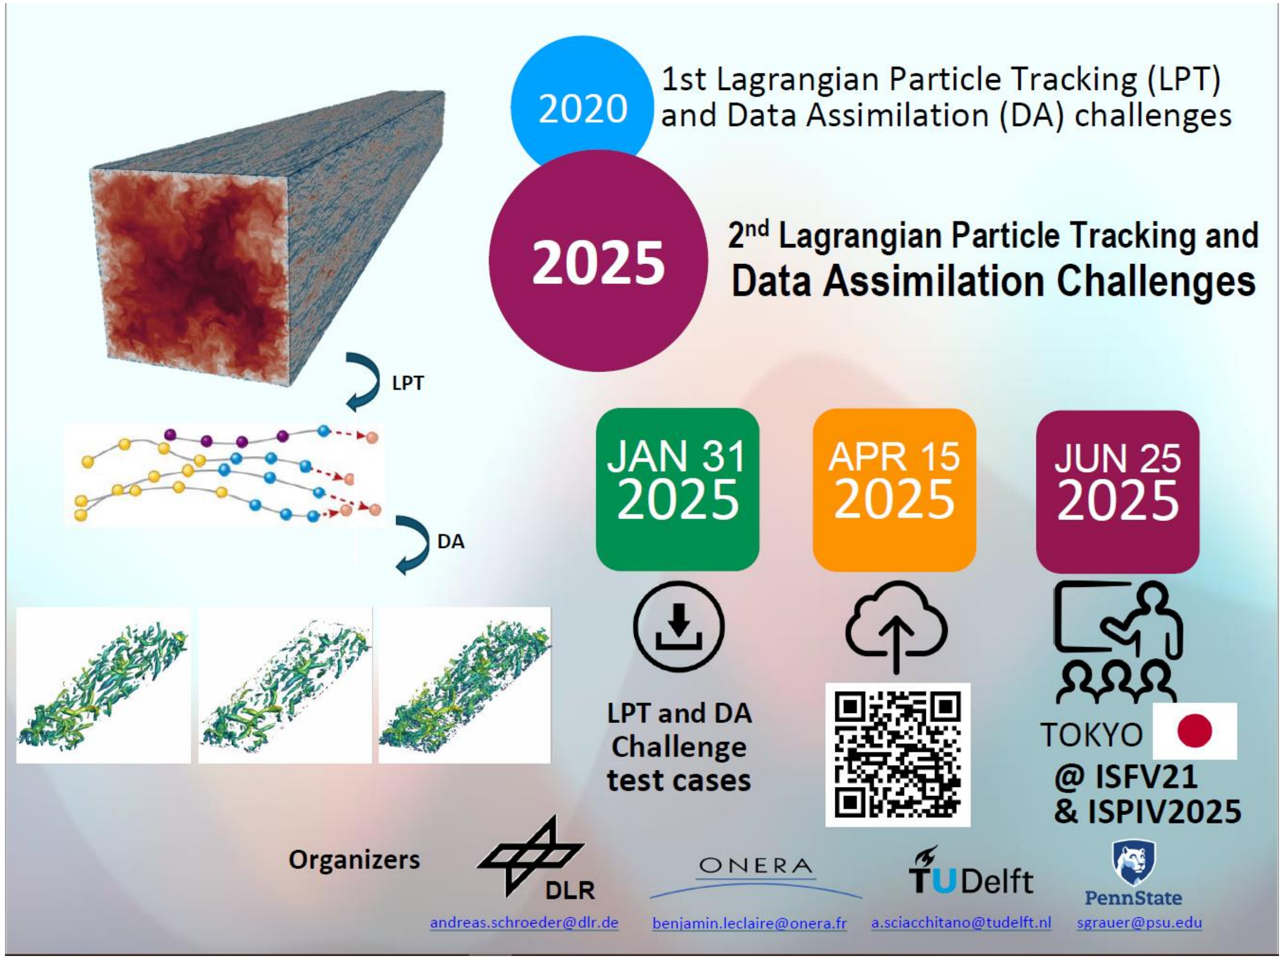

- https://w3-d8.onera.fr/flow-benchmarks/en/2ndLPTDAChallenge

LPT cases

2nd Lagrangian Particle Tracking and Data Assimilation Challenges

ONERA
TUDelft
PennState
DLR

- DNS of square duct flow

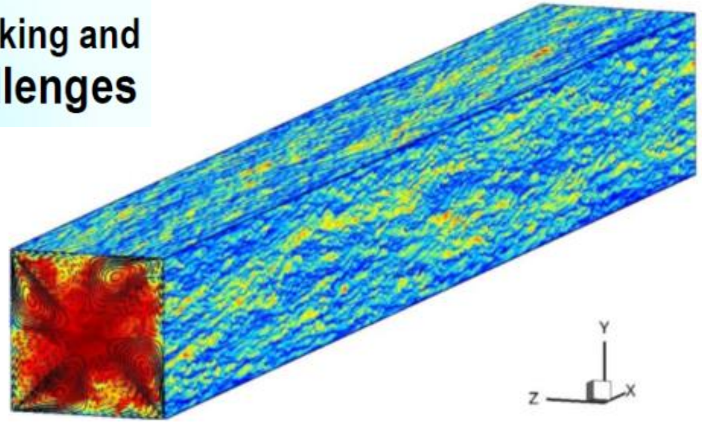

- LES of double cylinder flow

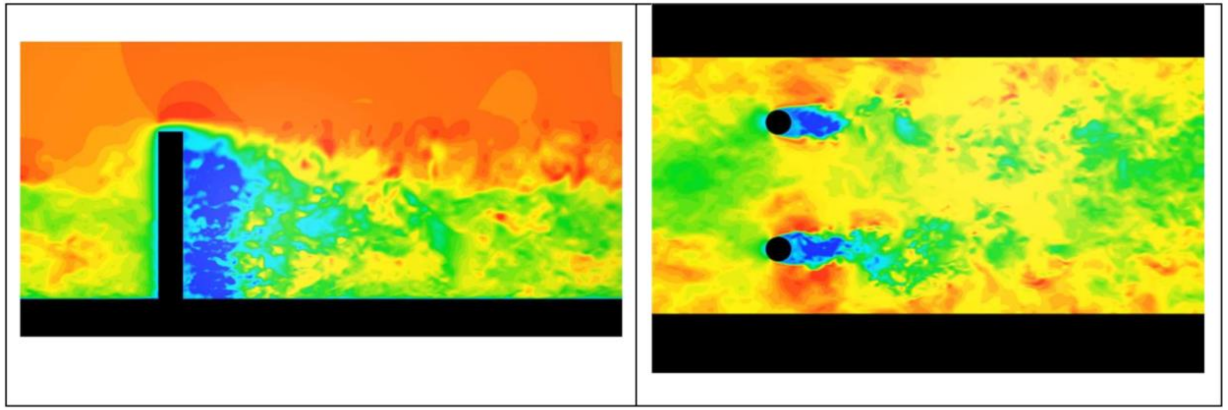

- Experimental RBC

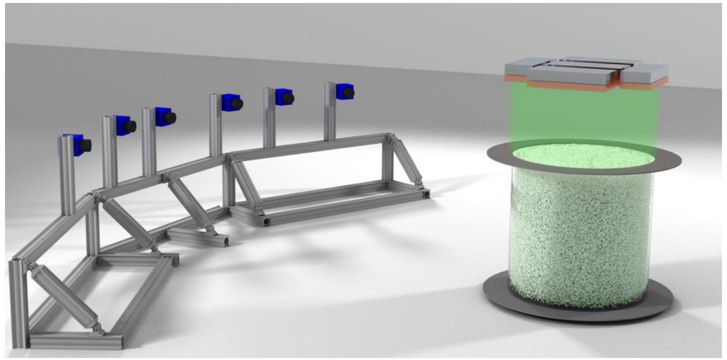

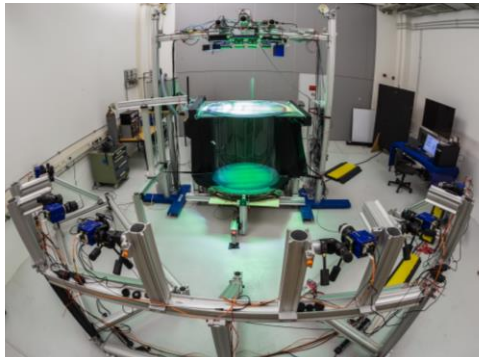

- Parameter variations:
  - Seeding density
  - Temporal sampling
  - Common volume of cameras
  - Mie scattering
- Time-Resolved and Two-Pulse cases

DA cases

2nd Lagrangian Particle Tracking and Data Assimilation Challenges

ONERA TU Delft PennState DLR

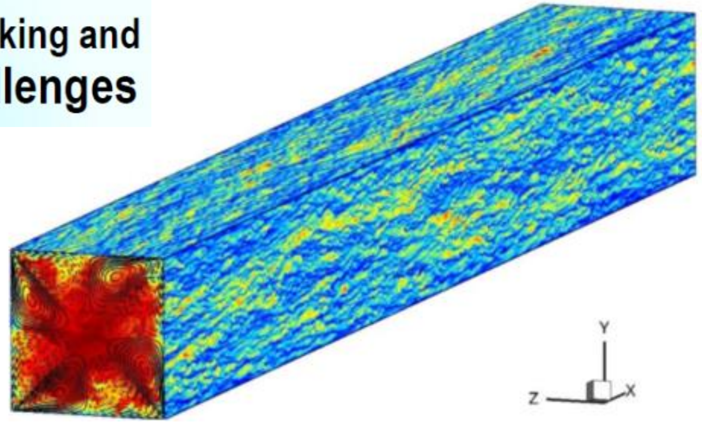

- DNS of square duct flow
- DNS of compressible turbulent boundary layer (TBL) @ Mach 4.9
  - Ideal and inertial tracers

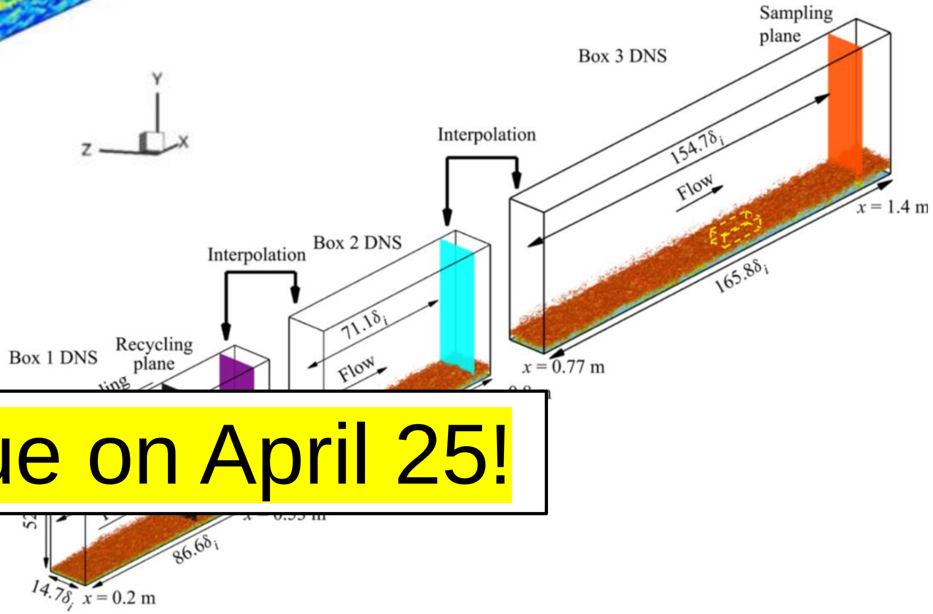

Download now! Results due on April 25!

- DNS of Rayleigh–Bénard convection
- Investigated properties:
  - Velocities and gradients
  - Pressure
  - Temperature
  - Density

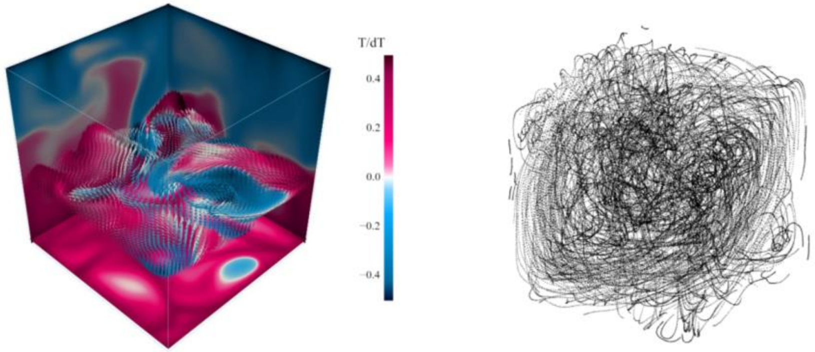

Thank you for your attention!