3D Imaging of Atmospheric Dispersion Processes with Dial

Robert Lung[0], Nick Polydorides

University of Edinburgh

May 2023

[0] Supported by the James Clerk Maxwell Foundation, UK

Differential Absorption based Imaging

Basics Working Principle

Problem: Determine the 3D spatial concentration profile of a known trace gas using differential absorption Lidar.

- Pulsed laser used as light source
- Measure (back-)scattered light binned based on time-of-flight
- 3D imaging requires scan of a cone. ($\rightarrow$ Lidar cube)

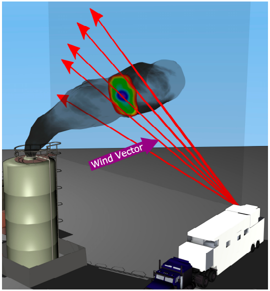

Mobile 2D Lidar scanning a plume cross section[1]

[1] Illustration taken from Innocenti, F and Robinson, R and Gardiner, T and Finlayson, A and Connor, A. Differential Absorption Lidar (DIAL) measurements of landfill methane emissions, *Remote Sensing*, 2017.

Differential Absorption based Imaging

Assumptions and Problems

- Gas of interest must be known a priori!
- Source is tuned to emit pulses at two wavelengths λ_{on} and λ_{off} which are chosen such that:
  - λ_{on} is absorbed by the target gas more than λ_{off}
  - Their scattering behaviour can be assumed identical
- Additional atmospheric data is sometimes necessary or useful.
- Simple approach: Trivial inverse problem but requires good signal quality which typically makes it **impractical for 3D reconstruction**
- Goal: Make better use of measurable data and prior knowledge.

- We consider the advection-diffusion equation given by

$$\frac{\partial}{\partial t}u + \nabla \cdot (\eta u) - Q + \frac{1}{2}\nabla \cdot (\kappa \nabla u) = 0 \tag{1}$$

with $Q = \rho_Q \cdot \delta(\vec{x} - \vec{q})\delta(t)$ is an instantaneous source term at $\vec{q}$ while η, κ model drift and diffusion respectively and shall be functions of time only.

- The plume can be modelled as a superposition of puffs ϕ

$$\sum_{j=1}^{N} w_j \phi \left(\frac{\|x - m_j\|_2}{h_j} \right) \tag{2}$$

for w_j, h_j and m_j which depend on the dispersion quantities and **regularise the inverse problem by imposing PDE based constraints**.

The first ingredient

Animated turbulent plume

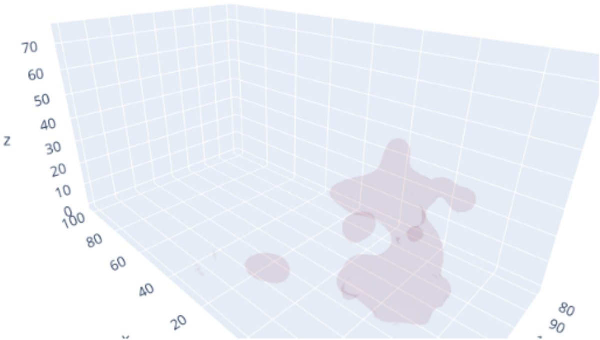

The first ingredient

Approximation error in the dispersion process

When it comes to the actual image the reconstruction is essentially just a **low-resolution** (blurry) version of the true image that preserves certain features

- There is some empirical/experimental justification for the approximation.
- The true generating process is not tractable.

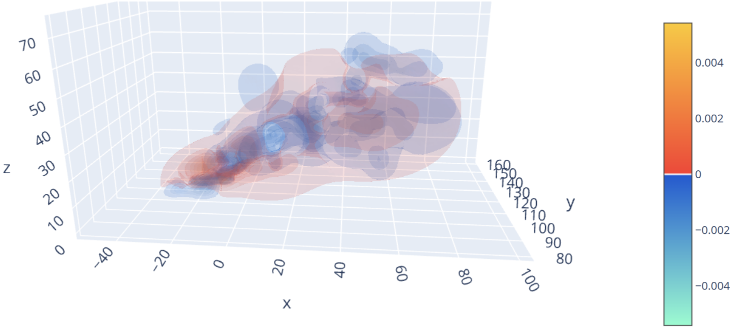

Low- vs. High-res difference ≈ 0.5 relative L_1 error in image

The first ingredient
Some notes on the approximation

- The parameters ultimately used in order to approximate the plume can be thought of as width and position for each down-wind cross section.
- The reconstructed distributions will approximately match in these features (and capture those well) but have significantly higher entropy and mismatched higher moments.

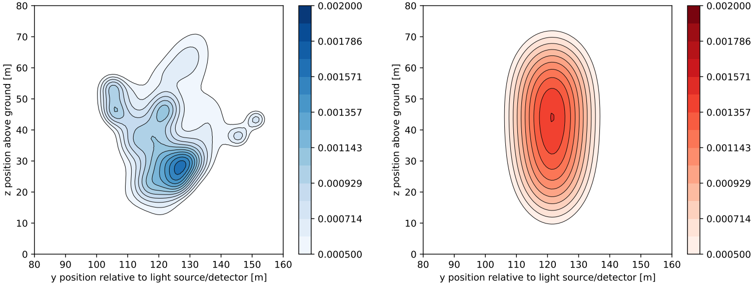

Cross sections of the gas distribution at x=35m. Note that the oval shape is due to the upwards drift of the gas.

The second ingredient

Radiative Transfer: Off-beam Signals

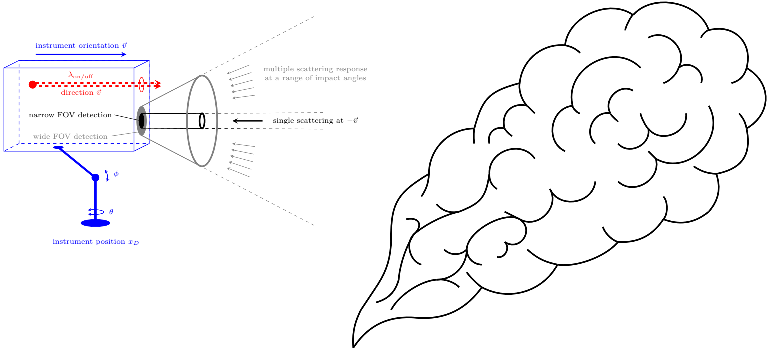

Laser pulses of two different wavelengths, $\lambda_{\mathrm{on/off}}$ (red), are released at the source, located at x_D, in direction $\vec{v}$. A detector, also located at x_D, with a narrow narrow FOV (black) captures the single scattering response incident at $-\vec{v}$ whereas a wide FOV captures light from a range of directions which have scattered multiple times. After a measurement in direction $\vec{v}$ has been taken the instrument is re-oriented by adjusting θ or ϕ (blue). This procedure is carried out for azimuthal angle and polar angles at a fixed instrument location x_D.

The second ingredient

Radiative Transfer: The scattering dilemma

- Optical remote sensing mostly relies on scattering
  - For “hard“ scattered light (from surfaces) the trajectory is uniquely determined by the outgoing angle and time of flight.
  - The same is true for single-scattered light (in narrow FOVs) when the scattering is caused by airborne particles.
- Gases can be measured by absorption at certain wavelengths
  - We can use Beer-Lambert law when the light’s trajectory is known.
  - Using a narrow FOV means we make very little error in the light paths at the cost of excluding possibly useful data.

Challenge

Make the other photons useful despite not knowing exact paths!

The second ingredient
Radiative Transfer: The forward model

The dynamics of light in heterogeneous scattering media can be modelled via the Radiative Transfer Equation (RTE)

$$\left(\frac{\partial}{\partial t} + v \cdot \nabla_x + \sigma_a^{\mathrm{on/off}} + \sigma_s\right) H^{\mathrm{on/off}} = \sigma_s \int_{\mathbb{S}^2} H^{\mathrm{on/off}} f_p \mathrm{d}v'$$

where $\sigma_s, \sigma_a^{\mathrm{on/off}}$ are heterogeneous scattering/absorption parameters and f_p is a phase function.

- The source term is $\delta(v - v_j)\delta(t)$ and differs for each direction v_j within the scanned cone.
- The measurement is taken **at a single point** on the boundary separately for each v_j.

The second ingredient

Neumann Series: Extension of the Lidar-Equation

- The solution of the RTE can (often) be written as the sum of contributions from all orders of scattering

$$H = \underbrace{H_1}_{\substack{\text{single scattering (narrow FOV)} \\ \text{=standard Lidar}}} + \underbrace{\sum_{j=2}^{\infty} H_j}_{\text{multiple scattering (wide FOV)}}$$

- When multiple scattering is considered there is no more closed form solution for the inverse problem.
- When the FOVs can be separated we gain access to H_1 and $\sum_{j=2}^{\infty} H_j$ individually.

The second ingredient

Radiative Transfer: The forward model

- RTE is well understood but the inverse problem is typically studied for “more complete“ measurements.
    - σ_a and σ_s control the rate of absorption and scattering.
    - f_p is a probability density that determines the direction after scattering events.
- There are many ways to “solve“ the RTE
    - Simplifying approximations for some regimes such as diffusion, single-scattering, etc.
    - Accurate solutions are typically expensive
- The RTE and dispersion should interact as seamlessly as possible (different “natural grids“).

The second ingredient

Approximation error for Radiative Transfer

Observation - where it really might go wrong!

The parameters of the RTE cannot be fully reconstructed but are necessary to evaluate the (forward-) model and compute the image.

- We must develop a theory regarding the information contained within the measured data.
- The central idea is to exchange temporal with spatial averaging.
- Use evaluations of the optical transport model that can be thought of as a high entropy approximation relative to a reference distribution.

Parameter Uniqueness under RTE Assumption

Single vs. Multiple Scattering

For functions such as (2) we can exploit the existence of a "first impact point" and use that single scattering is more singular and can be measured earlier than higher order scattering to show:

Theorem (first attempt uniqueness - informal)

Assuming the optical forward model is governed by the RTE, then the differential absorption field $\sigma_a^{\mathrm{on}} - \sigma_a^{\mathrm{off}}$ *and* σ_s *both akin to the form* (2) *are uniquely determined by the on and off intensities regardless of the FOV, provided the other parameters are given.*

In other words, given a subset of RTE parameters, there is a difference between wide and narrow FOVs iff we consider noisy data:

- Discrepancies between the average model used in the inverse problem and the true concentration profile
- Optical noise due to limited photon counts in each bin

Computational Trade-off

A semi-parametric approach

- Using a large number of kernels as in (2) we can account for (turbulence induced) variation in the plume at the cost of a high-dimensional inverse problem.
- With a small number of kernels RTE solutions become cheaper but the resulting model error/discrepancy can become large quickly.
- **Trade-off:** Use a semi-parametric model
  - High-dimensional parameterisation for intensity H^{off}
  - Low-dimensional form through dispersion parameters for absorption $\frac{H^{\mathrm{on}}}{H^{\mathrm{off}}}$

This is technically a relaxation!

Previous uniqueness result no longer valid in that generality.

Parameter Uniqueness under Relaxed Assumptions

A semi-parametric approach

Having made a relaxation to the RTE we must settle with a uniqueness result that makes further assumptions on the nature of the scattering and absorption functions.

Theorem (relaxed uniqueness - informal)

Assume that f_p is known while $\sigma_a^{\mathrm{off}}, \sigma_s$ and $\Delta_a := \sigma_a^{\mathrm{on}} - \sigma_a^{\mathrm{off}}$ are as in (2). If further $\sigma_a^{\mathrm{off}}, \sigma_s$ and Δ_a have common mid-points as well as widths, their base kernels satisfy a certain regularity condition and kernel weights are equal up to proportionality (with a constant shared between all summands), then the absorption $\frac{H^{\mathrm{on}}}{H^{\mathrm{off}}}$ uniquely determines σ_a, σ_s and Δ_a.

- Differential absorption sufficient to identify correct parameters if the scattering is “well aligned“.
- The role of the optical scattering parameters is different.

Detection of low concentrations

A toy problem with wider implications

- In general we assume a Poisson noise model for the optical data $\boldsymbol{m}_{t_i,v_j}, \boldsymbol{n}_{t_i,v_j}$ binned at mid-points t_i and directions v_j
- For low differential absorption and known plume shape the distribution of $\log\left(\frac{\boldsymbol{n}_{t_i,v_j}}{\boldsymbol{m}_{t_i,v_j}}\right)$ can be approximated by a Gaussian and (regardless of the FOV) used to test the Hypothesis

  $$H_0 : \text{No gas present} \qquad vs. \qquad H_1 : \text{Absorbing gas present}$$

  when the “shape“ of the gas is known
- For narrow FOVs tests constructed this way are essentially optimal whereas for wider FOVs their quality depends on the alignment of the true and estimated scattering behaviour

Detection of low concentrations

Worst case analysis under various conditions

- Detection in the case of a single kernel for varying distributions of scattering particles and phase functions[2]
- Plots show **worst case behaviour** relative to the (known) equivalent narrow FOV quantity subject to different constraints

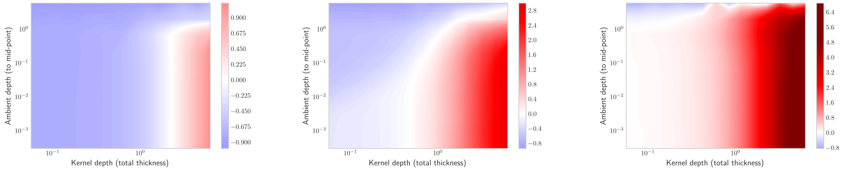

(a) unconstrained (b) known phase function (c) full knowledge

Figure (a): If nothing about the distribution of ambient particles is known then wide FOVs will improve the reconstruction only for optically thick plumes. Figure (b) and (c) show the degree of improvement from approximate knowledge and full of the scattering parameters.

[2] Henyey-Greenstein with range g=0 to g=0.7

Detection of low concentrations

Worst case analysis under various conditions

- Knowing that ambient scattering is limited, i.e. the feature of interest is well aligned with the scattering particles, is virtually equivalent to full knowledge

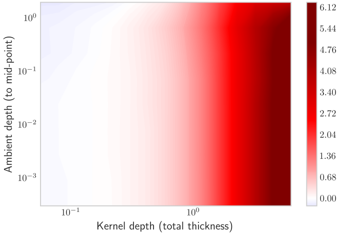

(a) limited ambient scattering

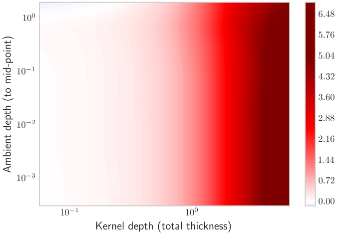

(b) full knowledge

Figure (a) and (b) show very similar results indicating that full knowledge of the scattering parameters does not yield considerably better results than a mere limit on ambient scattering.

Detection of low concentrations
Ambient scattering effects

- Light along paths that do not reach the region of interest bear no information about the parameter of interest
- Large amounts of ambient scattering render photons detected by the wide FOV a nuisance and reduce the sensitivity of the measurement

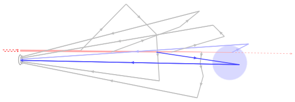

(a) ambient scattering
top/left in previous graphic

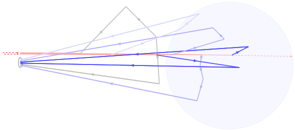

(b) plume scattering
bottom/right in previous graphic

Figure (a) and (b) show why the absorption in the wide FOV is heavily dependent on ambient scattering. Only blue trajectories are sensitive to the patch of interest. Dark blue patches are strongly affected whereas photons along grey paths behave like noise.

Detection of low concentrations

Generalisations from the toy problem

- The attempt at sensing a small feature enclosed within a larger plume will suffer from essentially the same effects as caused by ambient particles

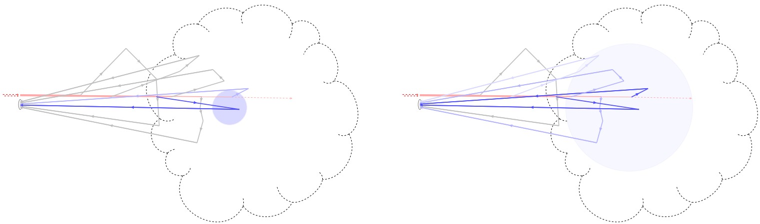

(a) small feature (b) large feature

Figure (a) and (b) show why the absorption in the wide FOV is heavily dependent on feature size. Only blue trajectories are sensitive to the patch of interest. Dark blue patches are strongly affected whereas photons along grey paths behave like noise.

Detection of low concentrations

Narrow FOVs

- Light collected by a narrow FOV will follow a fixed trajectory and pass through the section of interest as long as it is aimed in the right direction

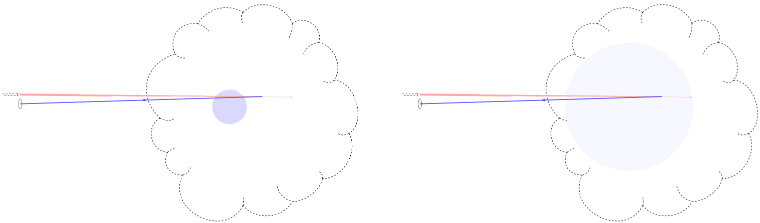

(a) small feature (b) large feature

Figure (a) and (b) show why the absorption in the narrow FOV is less dependent on feature size.

Qualitative conclusions and focus of simulation

Where can we expect a benefit and how much better can it get?

- **Main limitation: We cannot improve imaging resolution!**
  - Lack of information about the scattering parameters.
  - Even if σ_a and σ_s were known the "smoothing" of multiple scattering makes the data much less useful.
  - RTE solutions become insufficiently accurate and computationally problematic.
- Due to non-linearity there will be some bias in the estimator but the main contribution from the unknown RTE parameters is to be expected in low order scattering.
- Our main contribution is making sense of low-order scattering (e.g. up to 4-5 events), i.e. the "difficult" case in-between a diffusion approximation and single-scattering.

Solving the Coupled Inverse Problem
Likelihood and optical noise model

The likelihood can be expressed as

$$\begin{aligned}\mathsf{L}(\theta \mid \boldsymbol{m}, \boldsymbol{n}) = \sum_{i,j} & H^{\text{on}}_{t_i,v_j} + H^{\text{off}}_{t_i,v_j} \\ & - \boldsymbol{m}_{t_i,v_j} \log(H^{\text{on}}_{t_i,v_j}) - \boldsymbol{n}_{t_i,v_j} \log(H^{\text{off}}_{t_i,v_j})\end{aligned}$$

where $\theta = (\psi, H^{\text{off}})$ and $H^{\text{on}} = H^{\text{off}} \mathbb{E}_{p \sim Q_\psi}[\alpha_\psi(p)]$

- The parameters α, Q are suitably defined functions parameterised by low-dimensional dispersion related parameters ψ.
- Closed form solutions for H^{off} alongside low-dimensionality of profile likelihood lead to tractable reconstruction method.

Solving the Coupled Inverse Problem

Parameter Fitting

Maximum of $\mathsf{L}(\cdot \mid \boldsymbol{m}, \boldsymbol{n})$ w.r.t. H^{off} is at $H^{\text{off}}_{\psi} = \frac{\boldsymbol{m}_{v_j,t_i} + \boldsymbol{n}_{v_j,t_i}}{1+\mathbb{E}_{p\sim Q_{\psi}}[\alpha_{\psi}(p)]}$ so we can find ψ by iterating

$$\psi_{r+1} = \psi_r + \mathcal{I}(\psi_r)^{-1} \partial_{\psi} \mathsf{L}(\alpha_{\psi_r}, Q_{\psi_r}, H^{\text{off}}_{\psi_r} \mid \boldsymbol{m}, \boldsymbol{n})$$

and $\mathcal{I}(\psi)$ approximates the Hessian and is of the form

$$\mathcal{I}(\psi) = \sum_{i,j} (\boldsymbol{m}_{v_j,t_i} + \boldsymbol{n}_{v_j,t_i}) \frac{\partial_{\psi} P_{i,j}(\psi) \partial_{\psi} P_{i,j}(\psi)^{\mathsf{T}}}{P_{i,j}(\psi)(1 - P_{i,j}(\psi))}$$

- Only first derivatives! Limits number of RTE evaluations.
- In practice we will typically also have a prior for ψ which doesn't change the structure or complexity.

Solving the Coupled Inverse Problem

Interpreting scattering parameters

Intuitive control over the likely photon trajectories

- H^{off} acts as a back-scattering coefficient
- $\sigma_a^{\text{off}}(\psi), \sigma_s(\psi)$ control the forward propagation via $Q(\psi)$

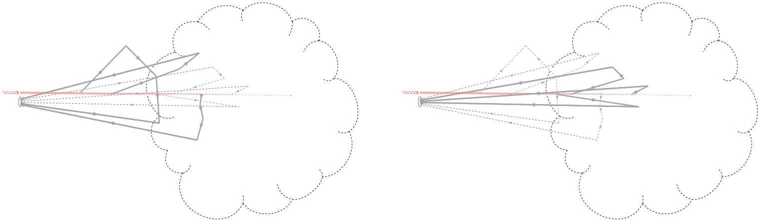

(a) diffuse paths (high σ_s)

(b) spiky paths (low σ_s)

Increasing the scattering rate σ_s results in more diffuse paths, thick lines in figure (a), while a reduction puts weight on more straight paths as shown in figure (b). In the semi-parametric model the magnitude of the signal is not affected.

- Simulated reconstruction from $30 \times 10 \times 50$ Lidar scan of 14 parameter dispersion which can be recovered when conventional reconstruction fails due to the low SNR.
- Scattering caused largely by particles around the gas plume and effective resolution $\approx$ granularity of absorbing gas within scatterer!
- Different phase function f_p used for simulation and reconstruction to emulate the complexity of real conditions!
- Fixed system parameters used are (approximately):
  - Detector: 3cm lens with 4% detection rate
  - Methane amount: 50mol or 0.8kg
  - Distance: 100m
  - Wavelength (absorbing): 1645.55nm
  - Ambient intensity: 0.025W uniformly over hemisphere, much less than peak signal but not entirely negligible.

Simulations

Reconstruction of Smooth Image and Parameters of Interest

- Low concentration means that the noise is large relative to the logarithmic differential absorption
- If the measurement isn't increasing regularisation is needed to avoid negative concentration values

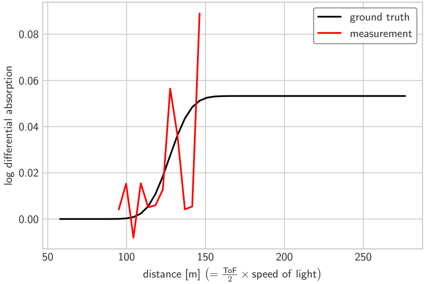

Figure (measured narrow FOV data in the region of interest): The level of noise in the measurement results in data that isn't monotone for the observed concentration levels. The proposed parameterisation through dispersion related quantities is one way of dealing with this issue.

Simulations

Reconstruction from $30 \times 10 \times 50$ Lidar scan

L_1 errors: $\frac{\|u_{\text{est}} - u_{\text{true}}\|_1}{\|u_{\text{true}}\|_1}$ and release amount errors: $\left|\frac{\|u_{\text{est}}\|_1 - \|u_{\text{true}}\|_1}{\|u_{\text{true}}\|_1}\right|$ where u_{true} denotes the ground truth concentration field and u_{est} the estimation from the data.

Turbulent	Counts (nFoV:wFOV)	L_1 errors (nFOV|wFOV|mFOV)	Release amount errors (nFOV|wFOV|mFOV)
Yes	7.2	(56% | 49% | 44%)	(13% | 15% | 14%)
Yes	3.9	(49% | 37% | 32%)	(16% | 12% | 10%)
Yes	2.4	(50% | 44% | 42%)	(20% | 13% | 9%)
Yes	1.7	(50% | 46% | 49%)	(19% | 20% | 17%)
No	7.2	(39% | 28% | 30%)	(14% | 9% | 12%)
No	3.9	(39% | 23% | 25%)	(18% | 9% | 10%)
No	2.4	(34% | 23% | 23%)	(13% | 9% | 9%)
No	1.7	(46% | 27% | 25%)	(20% | 16% | 14%)

Table: 10 plumes with 1 data set each, presented in increasing SNR. "nFOV" denotes narrow, "wFOV" wide and "mFOV" multiple (i.e. separately measured) fields of view respectively.

Simulations
Dispersion parameter reconstruction

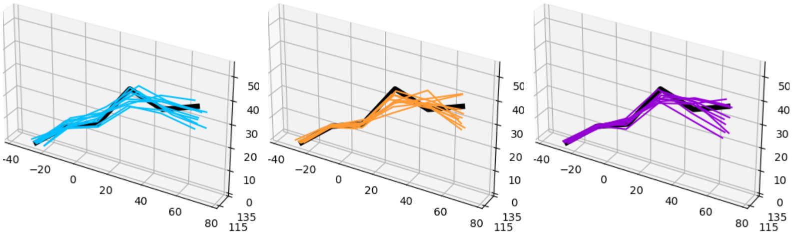

(a) narrow FOV (b) wide FOV (c) multiple FOVs

Reconstructed plume centre-lines for a smooth plume and a scattering particle concentration corresponding to a 2.4 nFOV:wFOV ratio in the measured data (line 7 in the previous table). Wide FOVs are more beneficial near the source due to higher optical thickness of the plume.

Simulations

Dispersion parameter reconstruction

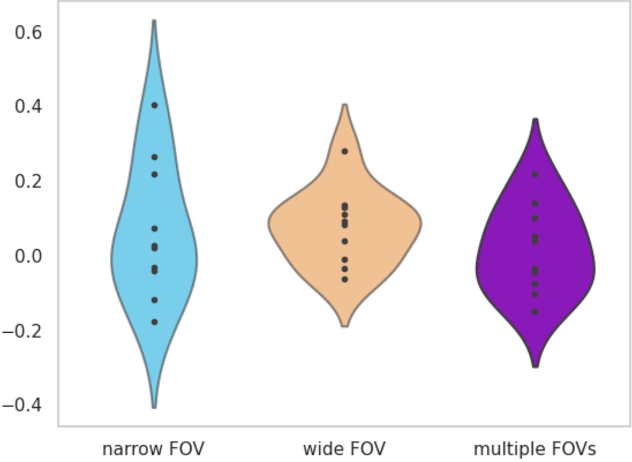

Relative deviation of reconstructed release rates corresponding to the same instance as shown in line 7 in the previous table

- Quadratic expansion involving $\mathcal{I}$

(pro) captures complex correlations relatively well

(con) under-estimates errors due to turbulence

- MCMC based approaches can work but require RTE evaluations for high dimensional parameters.

- Replace $\frac{H^{on}}{H^{off}} = E_{p \sim Q_\psi}[\alpha_\psi(p)]$ with $\frac{H^{on}}{H^{off}} \approx E_{p \sim Q_\psi}[\alpha_\psi(p)]$ in order to "correct" $\mathcal{I}$

(pro) Laplace approximations of marginal posterior may be obtained more quickly than MCMC samples.

(con) Hyper-parameters for the distribution of Q_ψ, α_ψ to "match" a prior for dispersion are hard to determine.

Thank You!

- For details and more rigorous results can be found in (arXiv) ***Imaging of atmospheric dispersion processes with Differential Absorption Lidar***

- Any further questions?

Simulations

Dispersion parameter reconstruction

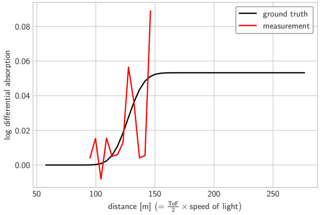

(a) low concentration

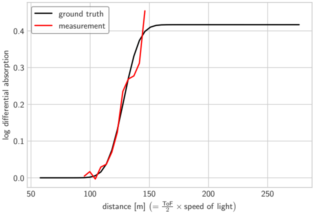

(b) high concentration

Higher concentrations are easier to recover due to larger gradients which are more robust to noisy measurements. Identical system parameters result in a much more well behaved differential absorption curve.

Simulations

Reconstruction of Smooth Image and Parameters of Interest

- On the left all scanned directions (red dots correspond to grid mid-points) and right single direction with wide FOV (in green). The FOV is still rather narrow thus not too much ambient light!
- The cone is about 10m wide at 100m distance and photons are measured in the wider FOV is the last scattering even occurs in the green cone. The narrow FOV is assumed extremely narrow (infinitesimal) and thus captures exclusively single scattering.

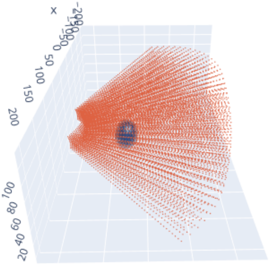

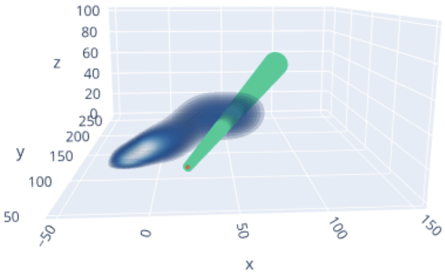

Solving the Coupled Inverse Problem

(Matrix) Concentration Inequalities

RTE evaluations can be done in parallel for each direction and need not be accurate!

- Monte Carlo integration (path tracing) provides a straightforward way to obtain RTE evaluations as well as gradients.
- Arguably the most difficult quantity to compute is the Hessian approximation of the form

$$\mathcal{I}(\psi) = \boldsymbol{A}(\psi)^T \boldsymbol{A}(\psi) + B$$

- structurally we have $\boldsymbol{A} \in \mathbb{R}^{m \times \dim(\psi)}$ is a random matrix with $m \gg \dim(\psi)$ consisting of independent blocks.
- Matrix concentration inequalities (Bernstein, Chernoff) ensure that $\mathcal{I}(\psi)$ is close to its expected value even when a small number of paths is traced in each direction

Simulations

Dispersion parameter reconstruction (incl. turbulence)

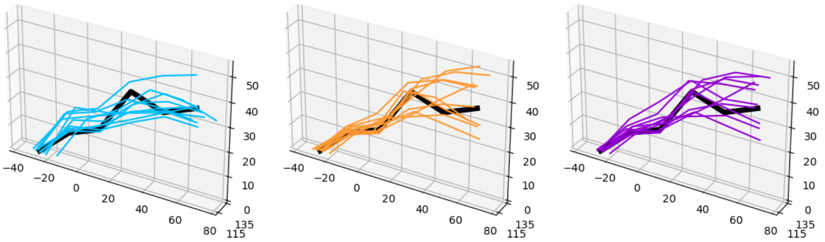

(a) narrow FOV (b) wide FOV (c) multiple FOVs

Reconstructed plume centre-lines for a turbulent plume and a scattering particle concentration corresponding to a 1.7 nFOV:wFOV ratio in the measured data (line 4 in the previous table)

Simulations

Dispersion parameter reconstruction (incl. turbulence)

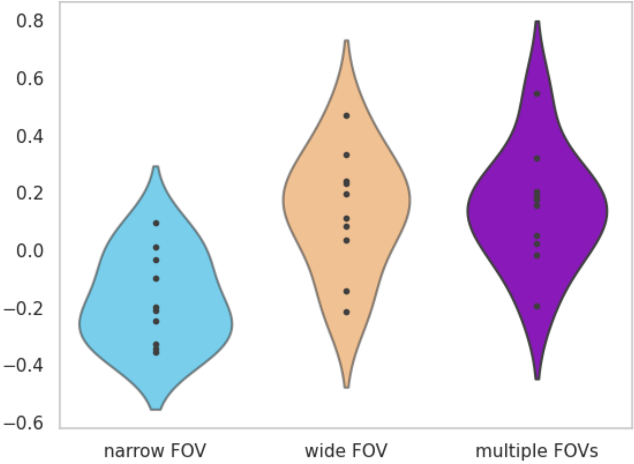

Relative deviation of reconstructed release rates corresponding to the same instance as shown in line 4 in the previous table