https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Urban Atmospheric Inverse Modeling via Hybrid Reconstruction Methods based on Compressed Sensing

Tobias Grasberger [1], Jia Chen [1], Benjamin Zanger [1], Alessandro Lupoli [1], Moritz Oliveira Makowski [1], Josef Stauber [1], and Felix Krahmer [2]

[1]Professorship of Environmental Sensing and Modeling, Technical University of Munich (TUM), Munich, Germany
[2]Professorship of Signal Theory and Statistical Learning, Technical University of Darmstadt (TU Darmstadt), Darmstadt, Germany

Correspondence: Tobias Grasberger (tobias.grasberger@tum.de) and Jia Chen (jia.chen@tum.de)

Abstract. Being able to trace and pinpoint greenhouse gas sources in urban environments in the face of climate change is more urgent than ever. Current atmospheric inversion approaches typically rely on Bayesian Inversion (BI) with Gaussian priors, which perform well for wide-range area sources but systematically smooth sharp emission peaks. This limits their ability to detect unknown point sources. An emerging method addressing this limitation is Sparse Reconstruction (SR), which enables the detection of localized emitters. However, they often fail to adequately represent spatially diffuse emissions.

In this paper, we combine these two complementary methods by introducing a novel hybrid atmospheric inversion framework based on the mathematical theory of compressed sensing. Our approach separates the recovery of concentrated point sources and dispersed area sources into two sequential steps. Using an observing system simulation experiment, we compare the performance of conventional BI and the proposed hybrid approach. The results demonstrate that the hybrid method significantly improves the recovery of mixed emission fields, mitigates over-smoothing of point sources, and enables inversions without a spatial prior in situations where BI using a Gaussian prior alone becomes infeasible. These findings highlight the potential of sparse and smooth hybrid inversion strategies for more reliable urban-scale greenhouse gas emission monitoring.

1 Introduction

Tackling climate change is one of the most pressing and complex challenges of the 21st century. As we continue to emit greenhouse gases (GHG) into the atmosphere, the Earth's temperature will gradually rise, causing irreversible damage to ecosystems and triggering feedback loops that further accelerate global warming (IPCC, 2023). To combat climate change, it is crucial to accurately track where greenhouse gases are being emitted. While we already have a solid understanding of carbon emissions and sinks in oceans and rural areas, measuring emissions in urban environments remains a challenge, particularly when trying to attribute those emissions to specific sources. As targeted urban sensor networks are not yet widely available, policy makers rely mainly on bottom-up estimations, called inventories, rather than direct, real-world measurements. While modern inventories are highly sophisticated, capable of incorporating data such as vehicle counts, power and heat plant intake flows, and even human respiration (Kühbacher et al., 2024, 2025; van der Gon et al., 2017, 2019; Poupkou et al., 2007; Kuenen et al., 2014), they are only based on reported or estimated emission sources. Thus, it is essential to validate and improve them using

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

GHG measurements over time. This creates opportunities to identify previously unknown emission sources that may not be captured by inventory-based approaches. Furthermore, it also yields more accurate estimates of the total emissions, which can help strengthen trust in reported emissions.

Using atmospheric Bayesian Inversion (BI) with a Gaussian prior (Rodgers, 2000; Jones et al., 2021), one can take an emission inventory as an initial starting point and improve or validate it by incorporating wind data coupled with atmospheric observations from permanent sensor networks (Jones et al., 2021; Dietrich et al., 2021), mobile in-situ measurement campaigns (Forstmaier et al., 2023; Chen et al., 2020; Vogel et al., 2024), and satellite measurements (Guanter et al., 2025; Taylor et al., 2020). Bayesian Inversion has been applied in various contexts, including urban-scale studies (Jones et al., 2021; Staufer et al., 2016; Forstmaier et al., 2023) and regional/global applications (Thompson and Stohl, 2014; Sijikumar et al., 2023; Jacob et al., 2016). BI has many favorable properties, such as having a closed-form solution, enabling computationally fast solutions, while also providing credibility intervals, i.e., a posterior probability distribution of the inversion. BI was initially applied primarily at continental and global scales, largely due to the availability of coarse-resolution data and computational limitations. Adapting this method to urban environments introduces new challenges, such as higher spatial resolution requirements, complex terrain, and less uniform emission patterns, which demand careful methodological adjustments.

A main drawback of BI using a Gaussian prior is its inability to accurately detect unknown point sources. The inversion process tends to distribute sharp emission peaks across the surrounding area — effectively 'smoothing' over them — rather than isolating them. This behavior is by design: In regions with sparse measurements or limited sensor coverage, it is often impossible to resolve emissions at high resolution.

An alternative approach to circumvent this limitation uses Sparse Reconstruction (SR) (Ray et al., 2015; Zanger et al., 2022; Hase et al., 2017; Cusworth et al., 2018), and is based on and supported by the mathematical theory of compressed sensing (Rauhut et al., 2008).

This method exploits the fact that GHG emissions in an urban context can often be traced back to very specific local emitters such as power plants, roads, landfills, or gas leaks. The localized nature, in turn, corresponds to a sparse representation of these sources in the emission field. To ensure that the reconstructed solution appropriately reflects this sparsity structure, the SR-based methods utilize a sparsity-promoting reconstruction scheme such as L1 regularization. When applied correctly, SR allows to capture the main drivers of an urban emission field. In particular, SR is well-suited for detecting unknown local point sources as it naturally favors solutions with concentrated, isolated emission signals rather than spatially distributed ones.

At the same time, dispersed area sources, e.g. CO_2 emissions in the heating sector, are not well represented in the sparsity-based framework, so one encounters a lack of resolution for these types of contributions and emissions get "over-sparsified". For example, a dispersed source gets modeled via three point sources in the sparse inversion. A first step towards addressing this issue has been taken in the discrete wavelet-based simulation studies by Ray et al. (2015); Zanger et al. (2022), which show how both point And area sources can be represented with the discrete wavelet transformation for low-resolution and low-noise synthetic data.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

However, this method can introduce unwanted artifacts in high-noise and urban scenarios, as discrete wavelets are specifically designed to encode images via their visible edges, whereas urban GHG emission fields typically exhibit smoother features and fewer sharp edges.

In this paper, we circumvent this issue using a hybrid approach, which incorporates both point source recovery and dispersed area source recovery. To do this, we first extract large point sources via a Sparse Reconstruction approach. Then, in a second step, we reconstruct the underlying dispersed sources using either Bayesian Inversion with a Gaussian prior or a Sparse Reconstruction method adapted to non-localized components, based on the Discrete Cosine Transform. Our study compares our hybrid approach to BI with a Gaussian prior across a range of selected case studies. These illustrate the respective strengths and limitations of each method in reconstructing different emission patterns from simulated sensor observations. We find that our hybrid approach performs particularly well in the presence of additional unknown point sources. The hybrid approach also allows for an emission field recovery without prior knowledge of the spatial distributions of emissions, even in the presence of point sources that would make the application of standard BI with a Gaussian prior infeasible.

The paper is structured as following: Section 2 introduces the atmospheric inversion problem and explains the algorithmic principles of Bayesian Inversion with a Gaussian Prior, Sparse Reconstruction, and our hybrid approach. In Section 3, we show experimental results across a series of synthetic experiments, illustrating the behavior of different models when confronted with a mix of focused and dispersed emission sources. And in Sections 4 and 5, we discuss the limitations of the case-studies, draw conclusions, and investigate directions for further research.

2 Methodology

We consider an urban region of interest and its rural surroundings and divide it up into a uniform $k_1 \times k_2$ grid of approximately one-kilometer-square tiles. The resulting 2D emission field for the $n = k_1 \cdot k_2 \in \mathbb{N}$ grid cells is represented as a 1D vector $\boldsymbol{x} \in \mathbb{R}^n$, where each entry corresponds to the annual emission flux of a specified greenhouse gas (GHG) from a single grid cell. In our case, we rely on artificial measurements from a limited number of sensor locations, such as those from the MUCCnet network (Dietrich et al., 2021) based on the differential column measurement principle (Chen et al., 2016), visualized in Fig. 1.

Using wind measurements and simulations, it is possible to link the flux increase or decrease for each grid cell with the GHG concentration measurements via the forward model equation

$$F(\boldsymbol{x}) = \boldsymbol{y}\,, \tag{1}$$

where $\boldsymbol{y} \in \mathbb{R}^m$ are the $m \in \mathbb{N}$ observed measurements. Depending on the transport model, the forward operator $F : \mathbb{R}^n \to \mathbb{R}^m$ can be linear, non-linear or even a stochastic function. In our case, we use the Lagrangian particle dispersion model STILT (Lin et al., 2003) driven by ERA5 (Store et al., 2023) wind data, which results in a linear forward model i.e. $F(\boldsymbol{x}) := \boldsymbol{A}\boldsymbol{x}$ with a *footprint matrix* $\boldsymbol{A} \in \mathbb{R}^{m\times n}$. The forward model is constructed by simulating the path of particles in the wind backwards in time from each sensor location and for each measurement time, producing a *footprint*, that represents the sensitivity of each

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

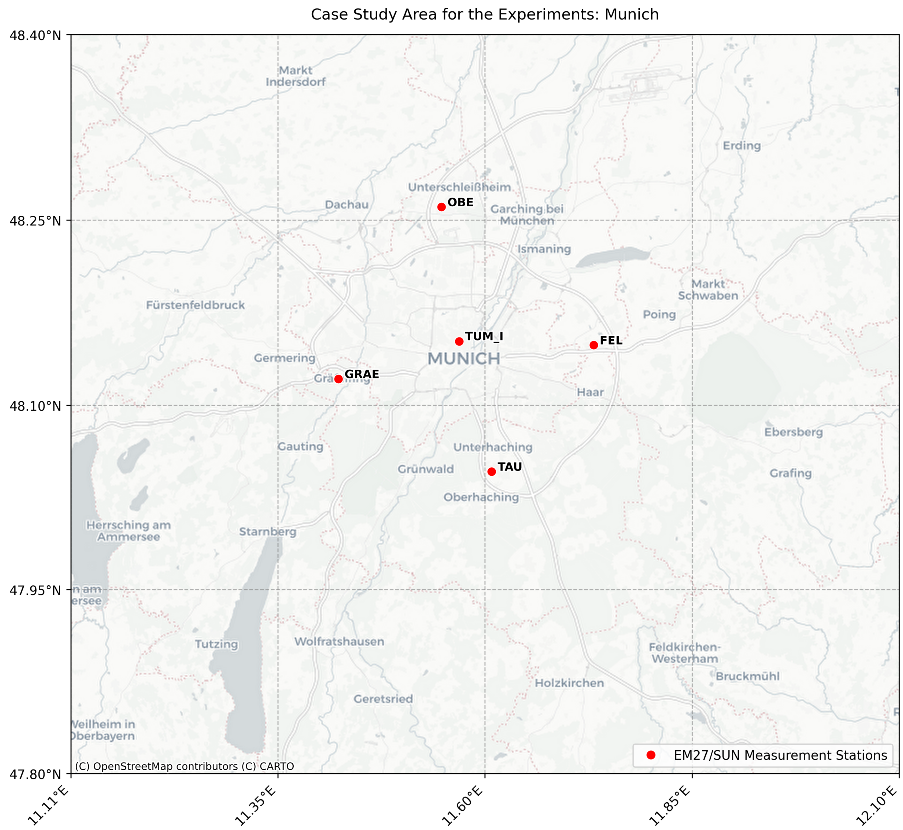

Figure 1. The MUCCnet (Dietrich et al., 2021) sensor locations are marked in red, overlayed onto a map of the broader Munich area. The five sensors are placed around and in the city to measure primarily urban emissions. Map: © CARTO (CARTO, 2024)

measurement to emission fluxes from every grid cell. The elements of the footprint matrix contain an estimated value of how much the GHG measurement in ppm should rise for a change in GHG fluxes in each specific grid cell.

Each row of the footprint matrix $\boldsymbol{A}$ corresponds to a vectorized footprint for a single measurement at a specific location and time, while each column of $\boldsymbol{A}$ represents the influence of one grid cell on all measurements. We illustrate typical footprint patterns we can expect from STILT and ERA5 wind data in Fig. 2.

Using the footprint matrix, we can set up a linear model that links the observed GHG enhancements $\boldsymbol{y} \in \mathbb{R}^m$ with the emission field $\boldsymbol{x}$ via

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \varepsilon \, . \tag{2}$$

ε captures all the unknown noise and model transport errors in the measurements, typically modeled to be i.i.d. Gaussian random variables $\varepsilon_i \sim N(0, \sigma^2)$ for $i = 1, \ldots, m$.

This setup defines a typical atmospheric inversion problem: We aim to recover the unknown emission field $\boldsymbol{x}$ based only on indirect measurements $\boldsymbol{y}$ and the forward model A. Note that this problem is typically *ill-posed*, especially when the number

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

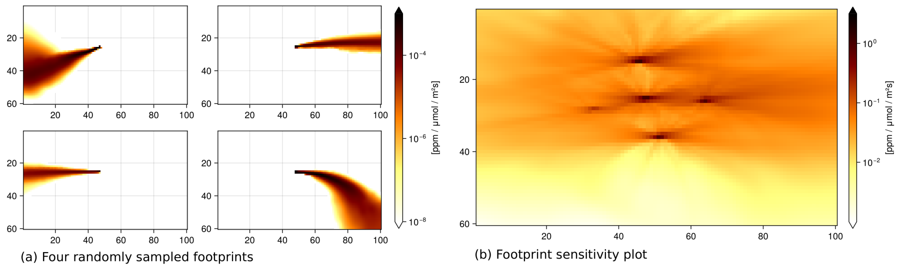

Figure 2. Footprint plots: (a) shows randomly sampled footprints for the middle station (TUM_I) of the MUCCnet sensor network. The typical plume shapes illustrate the sensitivity of the measurement with respect to specific grid-cell emissions. Essentially, the plots show the 'origin' of the wind mass that is being measured at one specific moment in time at the station. Each of the four plots shows the information contained in a row of the footprint matrix. (b) shows the cumulative sensitivity, i.e. the column-wise sum of 5422 sampled footprints from summer 2022 and 2023. The sensitivity is highest around the measurement stations. We can see that the wind is mostly coming from the west and east of Munich.

of grid cells n far exceeds the number of measurements m. In such cases, there is just not enough data to really recover the true emission field, and infinitely many solutions exist that can explain the data. Naive approaches, such as least-squares estimation, by solving

$$\underset{\boldsymbol{x}}{\operatorname{argmin}} \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}\|_2^2 \tag{3}$$

often lead to overfitting and extreme over-sensitivity to noise. We can add more measurements by measuring over longer periods or add more sensor sites to the measuring network to improve the recovery conditions, however this is usually very costly or infeasible. Instead, we can also circumvent this problem by adding more prior information via constraints and regularization into the model. In the following, we will review the current standard approach, Bayesian Inversion with a Gaussian prior (BI), and then introduce our alternative approaches based on Sparse Reconstruction (SR).

2.1 Bayesian Inversion with a Gaussian Prior

The key idea of Bayesian Inversion (Rodgers, 2000; Jones et al., 2021) is to utilize an emission inventory as an initial starting point. One penalizes the deviation from the prior via an additional regularization term, while also minimizing the model observation error term in Eq. (3) at the same time. Under the hood, we assume a probabilistic uncertainty distribution around the data points of the inventory, which is referred to as the *prior distribution* in Bayesian statistics. We represent the prior with the vectorized version $\boldsymbol{x}_a \in \mathbb{R}^n$, where each element in the prior vector stands for one grid cell. When the prior distribution is Gaussian, it is fully represented via its covariance matrix $\boldsymbol{S}_a \in \mathbb{R}^{n \times n}$, where the diagonal elements represent the variances of the entries of $\boldsymbol{x}_a$. The off-diagonal elements encode time-independent spatial covariance across all measurements: If they take non-zero values, this reflects, for example, that neighboring cells are correlated, are influenced by each other, or one observes similar behavior across specific emission sectors. It is common to assume that each entry of $\boldsymbol{x}_a$ has the same variance σ_a and

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

the spatial correlation at $\ell > 0$ km distance is given by

$$\boldsymbol{S}_{a;i,j} = \sigma_a^2 \exp\left(\frac{|d_{i,j}|}{\ell}\right) \tag{4}$$

where $d_{i,j}$ is the distance in kilometers between the grid cell elements i and j for $i,j \in [1,\dots,n]$.

By balancing the deviations from the observations and from the prior, we end up solving the following optimization problem to get the *Maximum A Posteriori (MAP)* estimator

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \quad (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})^T \boldsymbol{S}_{\boldsymbol{\varepsilon}}^{-1} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}) + (\boldsymbol{x} - \boldsymbol{x}_a)^T \boldsymbol{S}_a^{-1} (\boldsymbol{x} - \boldsymbol{x}_a). \tag{5}$$

$\boldsymbol{S}_\varepsilon \in \mathbb{R}^{m \times m}$ is the covariance matrix of the noise $\boldsymbol{\varepsilon}$ and reflects the observation uncertainty or the model data mismatch, which is also assumed to be Gaussian. The diagonal elements of $\boldsymbol{S}_\varepsilon$, the variance of the measurements reflect, for example, the measurement error or transport error introduced by the forward model $\boldsymbol{A}$. The off-diagonal elements allow for modeling either the correlation between different measurement stations or the temporal correlation of consecutive measurements at the same station. In this paper, however, we follow the common simplifying approach to model the noise as independent identically distributed, which corresponds to a rescaled identity as covariance matrix, $\boldsymbol{S}_\varepsilon := \boldsymbol{I}_m \cdot \hat{\sigma}^2$, where $\hat{\sigma}^2$ is an estimation of the observational noise. A more detailed derivation on how to solve Eq. (5) can be found in Appendix A.

2.2 Sparse Reconstruction

Sparse Reconstruction (SR) as applied in (Hirst et al., 2013; Ray et al., 2014, 2015; Cusworth et al., 2018; Zanger et al., 2022; Hase et al., 2017), emerged from the idea that certain emission fields can be assumed to be inherently sparse or compressible, that is, dominated by few large entries. For example, when recovering methane emissions in an urban environment, one typically expects to recover mostly point sources. SR addresses this phenomenon by replacing the quadratic regularization term in Eq. (5) with the L1 norm $\|\cdot\|_1$, yielding the *LASSO* (Least Absolute Shrinkage and Selection Operator) (Tibshirani, 1996) optimization problem derived as

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \quad (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})^T \boldsymbol{S}_\varepsilon^{-1} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}) + \lambda \|\boldsymbol{x} - \boldsymbol{x}_a\|_1 \,, \tag{6}$$

here stated in a more generalized form with the covariance matrix S_ε. The LASSO has been applied in various research areas, as a core tool in machine learning and compressed sensing.

The reason is that L1 norm minimization is known to effectively force the model to find a sparse solution. More precisely, compared to a quadratic penalty, L1 regularization penalizes larger deviations from the prior $\boldsymbol{x}_a$ less, while imposing a much higher cost on the model for small deviations from the prior. As a result, the minimizer of Eq. (6) will keep most of the state vector elements in $\hat{\boldsymbol{x}}$ to be zero and try to explain the observations with as few non-zero state vector elements as possible.

As a consequence, the LASSO is able to perform variable selection. If given two highly correlated variables as candidates for a point source, the LASSO is incentivized to select the one that better explains the target variable, while shrinking the other's coefficient to zero. A quadratic penalty or L2 regularization, on the other hand, as in effect in Bayesian inversion with a

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

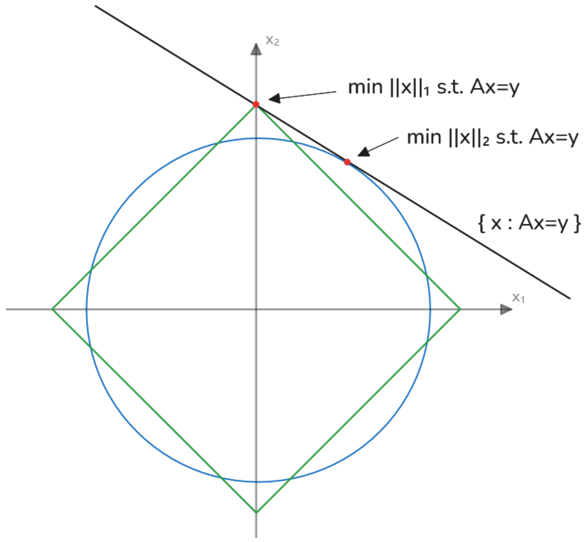

Figure 3. The visualization shows the effect of L1 (green) or L2 (blue) regularization on the solution of a two-dimensional equation. The solution that minimizes the L1 norm touches the solution space at the $\boldsymbol{x}_2$ axis, which represents a sparse solution. Illustration adapted from (Kutyniok, 2014).

Gaussian prior, would try to distribute the influence of each coefficient between both variables as balanced as possible to avoid big peaks. Fig. 3 provides further intuition on how the L1 norm affects optimal solutions in the two-dimensional toy case.

Another formulation of the optimization problem in Eq. (6) is called Basis Pursuit Denoising (BPDN) (Chen and Donoho, 1994; Chen et al., 2001). We formulate it here in a more generalized version with

155 $$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \quad \|\boldsymbol{x} - \boldsymbol{x}_a\|_1 \quad \text{s.t.} \quad (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})^T \boldsymbol{S}_\varepsilon^{-1} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}) \leq \eta \,. \tag{7}$$

Note that both BPDN and LASSO are equivalent, i.e. for every η there exists a λ_η such that both the LASSO and BPDN compute the same solution and the other way around as long as the solution is unique (Foucart and Rauhut, 2013, cpt. 15). However, the algorithms for solving both optimization problems and the strategies for selecting the hyperparameters η, λ differ. We go further into this in section 2.4.1. Note that the terms LASSO and BPDN are sometimes used interchangeably in the literature.

2.2.1 The DCT Transform

The selective behavior of L1 regularization is extremely helpful for recovering point sources in an emission field. The model is biased towards pinpointing the exact locations of local sources with distinct peaks, rather than distributing the signal into its surroundings. However, such a model will also "over-sparsify" more dispersed emission sources into single point sources, and as such be detrimental to a holistic inversion method that can deal with focused and dispersed sources at the same time. To increase the chances of a successful inversion using Sparse Reconstruction methods, it is often beneficial—if not essential—to apply a linear transformation $\boldsymbol{W} \in \mathbb{R}^{n \times n}$ onto the parameter space. In this way, the emission field $\boldsymbol{x}$ can now be described by its coefficients in the new transformed space via $\hat{\boldsymbol{x}}_{DCT} = \boldsymbol{W}\boldsymbol{x}$ and vice versa by $\boldsymbol{x} = \boldsymbol{W}^{-1}\hat{\boldsymbol{x}}_{DCT}$ with the inverse transformation matrix $\boldsymbol{W}^{-1}$.

In this paper, we model area sources with Sparse Reconstruction by applying a Discrete Cosine Transform (DCT) (Strang, 1999) onto the parameter space. By solving the inversion in this transformed space, we do not promote sparsity in the original grid-based representation, but in the introduced DCT encoding, which favors more 'cosine-like' locally smooth structures. Such a transformation can be applied via the inverse DCT transformation matrix $\boldsymbol{W}_{\mathrm{DCT}}^{-1}$ by multiplying it from the right to the footprint matrix to get a new forward model $\hat{\boldsymbol{A}}$ with

$$\hat{\boldsymbol{A}} = \boldsymbol{A} \cdot \boldsymbol{W}_{\mathrm{DCT}}^{-1} \,. \tag{8}$$

After solving Eq. (6) with the new forward model and yielding the solution $\hat{\boldsymbol{x}}_{DCT}$ in the transformed space, we can recover the inversion $\boldsymbol{x}$ of the emission field by multiplying $\hat{\boldsymbol{x}}_{DCT}$ with the inverse DCT transform $\boldsymbol{W}_{\mathrm{DCT}}^{-1}$ again.

2.3 Hybrid Inversion Method

The core idea of the hybrid method is to split point-source and area-source recovery into two separate steps.

To do this, we first compute a Sparse Reconstruction of the emission field without a transformation to yield an overly sparse recovery. This type of inversion captures point sources very well, but condenses area sources into multiple smaller single-point sources scattered across the range of the area source. Then we apply a threshold to the sparse reconstruction to keep only the large point sources that are true point sources with high likelihood, while discarding the smaller ones.

There are several options for choosing this cutoff value. The easiest is to just look at a histogram of the recovered point sources to find clear outlier values. Such a histogram is visualized for the last case study of the numerical experiments in Appendix D1. Another more sophisticated approach is to choose the cutoff value in advance by formulating a confidence interval (CI) around the prior. This CI specifies the maximum deviation from the prior that we can reasonably expect without the influence of an unknown point source. For example, for a component of a prior CO_2 inventory like a highway, we can expect, on average, at most double emissions over a monthly inversion. We keep any point sources recovered with SR that exceed this confidence interval. Then we model any other recovered sources within this confidence interval using a smooth inversion method, such as BI or SR, with a DCT transform.

We summarize the Hybrid Inversion approach in Algorithm 1. A similar approach, that also separates smooth and sparse structures in an atmospheric inversion, has already been shown by Chung et al. (2024). Their method performs the inversion in just one step. However, their method focuses on large-scale inversions with lower-resolution grids. Hirst et al. (2013) also shows a similar approach, but on a much more local scale with aircraft measurements.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Algorithm 1 Hybrid Inversion Method with separate Sparse and Smooth Recovery

Input: Measurements $\boldsymbol{y} \in \mathbb{R}^m$, forward operator $\boldsymbol{A} \in \mathbb{R}^{m\times n}$, prior $\boldsymbol{x}_a \in \mathbb{R}^n$

Noise estimate $\hat{\eta}$, threshold parameter $\tau \in \mathbb{R}$, inverse DCT transformation matrix $\boldsymbol{W}_{DCT}^{-1}$. Covariance Matrices $\boldsymbol{S}_\varepsilon^{-1}, \boldsymbol{S}_a^{-1}$

Output: Emission field reconstruction $\hat{\boldsymbol{x}}^{(2)}$

1: **Recover point sources via SR**

$$\hat{\boldsymbol{x}}^{(1)} = \underset{\boldsymbol{x}\geq 0}{\operatorname{argmin}} \quad \|\boldsymbol{x} - \boldsymbol{x}_a\|_1 \quad \text{s.t.} \quad (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})^T \boldsymbol{S}_\varepsilon^{-1} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}) \leq \hat{\eta}$$

2: **Update prior to include detected point sources above the threshold**

Cut elements in $\hat{\boldsymbol{x}}^{(1)}$ below threshold τ:

$$\hat{\boldsymbol{x}}_i^{(1)} = 0 \quad \text{for all } i = 1, \ldots, n \text{ where } \hat{\boldsymbol{x}}_i^{(1)} < \tau$$

Add recovered point sources to prior x_a:

$$\boldsymbol{x}_a^+ = \boldsymbol{x}_a + \hat{\boldsymbol{x}}^{(1)}$$

3: **Smooth recovery via either SR DCT or BI**

Inversion via BI using a Gaussian prior

$$\hat{\boldsymbol{x}}^{(2)} = \underset{\boldsymbol{x}\geq 0}{\operatorname{argmin}} \quad (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})^T S_\varepsilon^{-1} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}) + (\boldsymbol{x} - \boldsymbol{x}_a^+)^T \boldsymbol{S}_a^{-1} (\boldsymbol{x} - \boldsymbol{x}_a^+)$$

Or: Inversion via SR DCT

$$\hat{\boldsymbol{x}}_{DCT}^{(2)} = \underset{\boldsymbol{x}\geq 0}{\operatorname{argmin}} \quad \|\boldsymbol{x} - \boldsymbol{x}_a^+\|_1 \quad \text{s.t.} \quad (\boldsymbol{A}\boldsymbol{W}_{DCT}^{-1}\boldsymbol{x} - \boldsymbol{y})^T \boldsymbol{S}_\varepsilon^{-1} (\boldsymbol{A}\boldsymbol{W}_{DCT}^{-1}\boldsymbol{x} - \boldsymbol{y}) \leq \hat{\eta}$$

$$\hat{\boldsymbol{x}}^{(2)} = \boldsymbol{W}_{DCT}^{-1} \hat{\boldsymbol{x}}_{DCT}^{(2)}$$

4: **return** $\hat{\boldsymbol{x}}^{(2)}$

2.4 Experimental Setup

2.4.1 Solving the LASSO and BPDN

There exists a multitude of different algorithms to solve the LASSO, which all have their advantages and disadvantages. In this paper, we opted for a "Proximal Gradient Descent" or FISTA implementation in the programming language Julia (Antonello et al., 2018), which has shown a good balance of speed and high accuracy, even in our extremely high-dimensional setting. There are other algorithms, which do not solve the LASSO but directly find a sparse reconstruction, such as OMP and CoSAMP (Pati et al., 1993; Needell and Tropp, 2009). These may yield faster computation and granular control over the sparsity of the solution, but have generally weaker recovery guarantees and often perform empirically worse.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

A major disadvantage of FISTA is its inability to efficiently incorporate positivity constraints into the model. In such a case, where we want to restrict the solution space to positive solutions, it is more efficient to directly solve the BPDN problem with an additional positivity constraint via a direct convex solver framework such as Gurobi (Gurobi Optimization, LLC, 2024). While the implemented underlying interior-point method does not scale well with the number of measurements, it is still feasible to solve reasonably quickly when restricting the solution space to be positive.

2.4.2 Morozov Discrepancy Principle

In the results section, we will compare the performance of different reconstruction algorithms, including our Hybrid method and BI. However, depending on the strength of the regularization, the performance of each method may differ drastically. To mimic the conditions where we have no prior knowledge about which hyperparameters are the best for each inversion method, we apply the Morozov Discrepancy Principle (Nair et al., 2003; Scherzer, 1993) to make our models comparable. The core idea is to choose the hyperparameters so that the truth-observation mismatch is similar to or equal to the model-observation mismatch. At the and we want the equation

$$\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}\|_2 = \|\boldsymbol{\varepsilon}\|_2 := \eta \approx \|\boldsymbol{A}\hat{\boldsymbol{x}} - \boldsymbol{y}\|_2 \tag{9}$$

to hold true.

This method generally requires knowledge about the noise ε, which is fine for our synthetic tests but requires more careful consideration and fine-tuning when applied to real data. Another alternative to the Morozov Principle would have been to optimize the hyperparameters for each method to obtain the best possible performance for each model. However, this is not reproducible as soon as the true emission field is unknown.

BPDN has the advantage that it is easier to pick a good candidate for the noise parameter η in contrast to the LASSO. In our case, η can be estimated directly from the noise we add to the measurements with $\bar{\eta} := \|\boldsymbol{\varepsilon}\|_2$. For real data, we need to estimate η from, for example, sensor noise or approximated transport error.

For LASSO and BI, we are required to perform a grid search or cross-validation to find the right hyperparameters, such that the model-observation mismatch approximately matches the measurement noise.

2.4.3 Metrics

To ensure a fair comparison of all inversion methods, we use the same metrics across all scenarios.

Choosing appropriate metrics is challenging because, for example, large emitters such as power plants dominate the CO_2 emission field and contribute most of the total emissions in the emission field within a single grid cell. Consequently, if we apply a sparse recovery algorithm that is capable of recovering localized point sources and evaluate the result against the ground truth using, for example, the relative L2 error, the metrics may indicate a nearly flawless inversion—even though only the power plants were recovered and the remaining emission field was neglected.

Therefore, when a small number of big point sources would distort and skew the metrics towards them, we exclude the point source from the metrics. This allows us to assess whether the underlying emission structure is captured beyond the dominant

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

point sources. Note that if the focused point sources are very dominant, then sparse reconstruction methods will recover them very reliably in our synthetic setting. As such, we do not report any metrics on how well we recover the singular point sources. We compare all models under five different criteria listed in Table 1.

Table 1. Summary of error metrics and similarity measure. $\boldsymbol{x}$ is the true value, while $\hat{\boldsymbol{x}}$, stands for the prediction and $\bar{\boldsymbol{x}}$ is the mean of $\boldsymbol{x}$.

Metric	Formula	Explanation
R-Squared	$1 - \frac{\lVert \boldsymbol{x}-\hat{\boldsymbol{x}}\rVert_2^2}{\lVert \boldsymbol{x}-\bar{\boldsymbol{x}}\rVert_2^2}$	Proportion of variance in $\boldsymbol{x}$ explained by the prediction $\hat{\boldsymbol{x}}$. A value closer to 1 means better fit.
Local Relative L2 Error	$\frac{\lVert \hat{\boldsymbol{x}}-\boldsymbol{x}\rVert_2}{\lVert \boldsymbol{x}\rVert_2}$	Measures the relative magnitude of the prediction error over the entire domain using the L2 norm.
Regional Relative L2 Error	$\frac{\lVert \hat{\boldsymbol{x}}*\square-\boldsymbol{x}*\square\rVert_2}{\lVert x\rVert_2}$	Convolve the emission field with a filter $\square$ to diffuse spikes in the true or reconstructed emission field. In this paper, we use Gaussian kernels over 5x5 and 13x13 squares. This essentially compares the true emission field and the reconstruction after applying a Gaussian blur.
Total Relative Error	$\frac{\sum_{i=1}^{n}\hat{\boldsymbol{x}}_i-\boldsymbol{x}_i}{\sum_{i=1}^{n}\boldsymbol{x}_i}$	Total relative error over the whole domain.
SSIM	—	Structural Similarity Index (Wang et al., 2004). Measures perceptual similarity between images based on similar luminance (regional means), contrast (similar standard deviations), and structure (regional covariance being similar to the product of stds.). Values range from -1 to 1, with 1 being identical.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

3 Results

3.1 Point source recovery

We begin with a scenario that demonstrates how sparse recovery (SR) excels at recovering isolated point sources, whereas Bayesian inversion (BI) using a Gaussian prior struggles to recover such features. To this end, we set up a synthetic case study in which SR is used to recover point sources embedded in an otherwise known emission field. Modern breweries produce methane via anaerobic digestion (Arantes et al., 2017) of their wastewater, and as such are new potential candidates for possible undocumented methane sources. Using this concept, we have constructed a simple artificial methane emission field that simulates singular gas leak point sources at brewery locations around Munich over a static background emission field, which consists of the TNO methane inventory of Munich (Dellaert et al., 2019). We plotted the inventory in Appendix B1. The strength of the artificial leaks are randomly distributed via a Gaussian distribution with a mean of $0.013\,\mu\mathrm{mol\,m^{-2}s^{-1}}$ over a 1×1 kilometer grid cell and a standard deviation of $0.01\,\mu\mathrm{mol\,m^{-2}s^{-1}}$. This case study assumes perfect knowledge about the rest of the emission field and showcases the isolated behavior of BI and SR when dealing with unknown point sources. We generate the artificial measurements by applying the forward model onto the emission field and applying noise. The experimental setup is explained in a flowchart in Appendix C1. We use 5422 footprints from summer days in 2022 and 2023, where the real MUCCnet sensor network was actually measuring. We have added 15dB and 30dB of relative noise to the artificial 5422 measurements. All model hyperparameters were tuned according to the Morozov Discrepancy Principle. As shown in Fig. 4, the unknown point sources have good recovery conditions as they lie well between the measurement stations and have a signal strength that is clearly above the underlying methane emission field. As such, we expect a good inversion algorithm to be able to recover them reliably.

Figure 4 shows the precise locations of the unknown point sources in the plot of the emission field minus the prior, and illustrates how accurately the different models recover them. As we can see BI clearly "oversmoothes" the point sources even under low noise (30dB) and even recovers some unwanted negative artifacts in the emission field, while LASSO is capable of almost perfect recovery of the point sources. BI avoids directly assigning emissions to single pixels and instead spreads out the fluxes into the surrounding area. Table 2 shows how the different methods compare in terms of the metrics. We can see that the metrics agree with the inversion plots, as the Sparse Reconstruction method outperforms BI when the task is simply to isolate point sources. Further comparisons are shown in Appendix F1 and G1, that go into detail on how precisely the focused sources are recovered.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

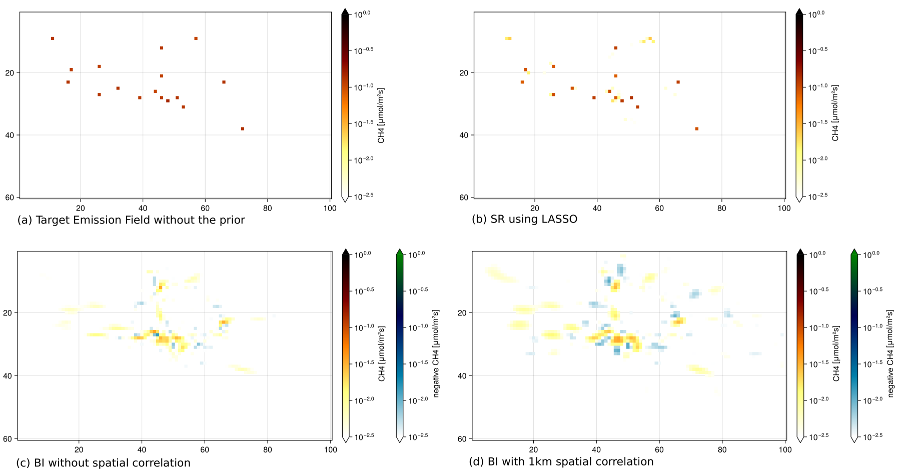

Figure 4. The target emission field with the prior subtracted vs. the respective recoveries minus the prior for the point sources recovery case study for 30dB noise. (a) shows the target ground truth. In (b), SR is able to recover the point sources almost perfectly, while BI in (c-d) underestimates the point sources and elevates the fluxes in the area around them. Adding spatial correlation only amplifies this effect.

Table 2. Metrics for different SNR levels for the point source recovery case study. The metrics are calculated with respect to the deviation from the prior, i.e. how well does the model recover the point sources.

	SNR = 15						SNR = 30					
Method	SSIM	Total Rel. Error	R^2	Regional Rel. 13×13	Regional Rel. 5×5	Local Rel. Error	SSIM	Total Rel. Error	R^2	Regional Rel. 13×13	Regional Rel. 5×5	Local Rel. Error
BI no corr.	0.9542	-0.3744	0.050	0.9871	2.0003	0.9736	0.9625	-0.0735	0.172	0.6828	1.7488	0.9084
BI 1.0 km	0.9543	**-0.3678**	0.028	1.0847	1.6843	0.9846	0.9572	**-0.0498**	0.123	1.0847	2.7583	0.9352
LASSO	**0.9643**	-0.4690	**0.314**	**0.5830**	**0.9952**	**0.8270**	**0.9942**	-0.1523	**0.877**	**0.1740**	**0.2550**	**0.3500**

3.2 Area source recovery

For the second case study, we are interested in recovering dispersed low-intensity area sources. A simple SR without any modifications, for example via a transformation, is very ill-equipped to reconstruct these types of sources. As such, this experiment is more aimed at showcasing how SR using a smoothness-based transformation is a viable alternative for Bayesian inversion. This is especially relevant for the second step in the hybrid inversion method presented in Algorithm 1. We have extracted and regridded a landcover map of the peat-/wetlands around Munich (Tegetmeyer et al., 2025) into our 60km×100km grid format. Peat- and wetlands are the largest natural methane emitters in the world (Saarnio et al., 2009). For the purpose of this case study, we assume that the entire peatland area is not drained and emits the same amount of methane. We scaled their approximate uniform emission rate to $0.05\,\mu\mathrm{mol\,m^{-2}s^{-1}}$. As before, we set the number of artificial measurements to 5422 and

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

we add 15dB and 30dB of noise.

We recover the emission field using both our presented approaches for recovering smooth, dispersed sources: Bayesian Inversion with a Gaussian Prior and Sparse Reconstruction using the DCT transform as a dictionary that represents smooth patterns in the emission field relatively well.

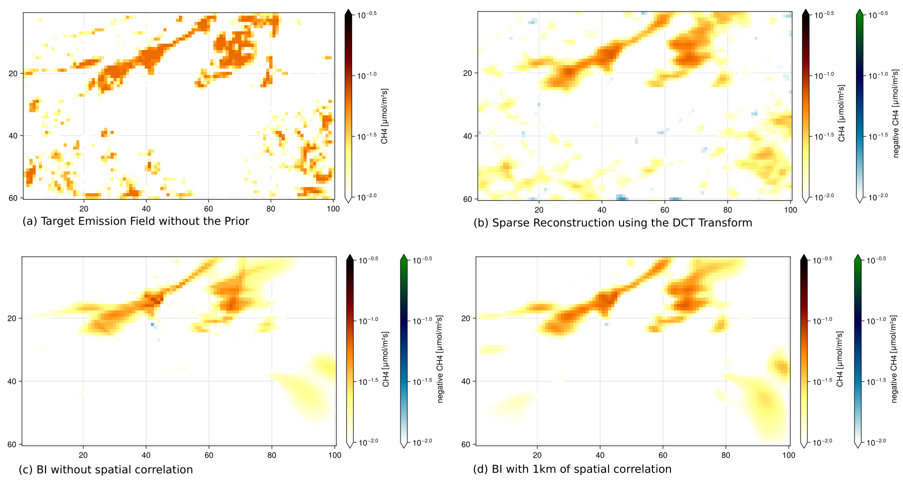

Figure 5. (a) shows the target emission field with the prior subtracted versus the respective recoveries for the area recovery case study with 30dB noise in (b-d). The sparse recovery using a smooth DCT transform in (b) performs captures the smooth structures comparably well, but introduces more artifact in the bottom emission field area.

Table 3. Comparison of inversion methods for the area recovery case study.

	SNR = 15						SNR = 30					
Method	SSIM	Total Rel. Error	R^2	Regional Rel. 13×13	Regional Rel. 5×5	Local Rel. Error	SSIM	Total Rel. Error	R^2	Regional Rel. 13×13	Regional Rel. 5×5	Local Rel. Error
Bayesian (no corr.)	0.8338	-0.2775	0.313	0.4083	0.6794	0.7383	0.8927	-0.1398	0.498	0.2568	0.5043	0.6310
Bayesian (1.0 km)	**0.8448**	-0.2684	**0.351**	**0.4069**	**0.6705**	**0.7179**	**0.9065**	-0.0661	**0.572**	**0.1779**	**0.4428**	**0.5825**
BPDN DCT	0.8264	**-0.0909**	0.347	0.3525	0.7188	0.7198	0.8567	**-0.0105**	0.454	0.2061	0.5839	0.6581
BPDN (no dict)	0.3467	0.6793	-13.909	1.3441	1.6864	3.4396	0.2806	1.2624	-18.553	1.6051	1.8914	3.9391

As shown in Fig. 5, both the sparse reconstruction method and Bayesian Inversion using a Gaussian prior perform well in recovering broad, spatially dispersed patterns. According to the metrics illustrated in Table 3, BI performs a little bit better across the board, especially with spatial correlation included. Nonetheless, the Sparse Reconstruction method using the discrete cosine transform (DCT) achieves a comparable performance, as the cosine basis is locally smooth and therefore well suited to representing smoothly varying emission patterns. However, this choice of basis also introduces its own artifacts, which become

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

more apparent in regions with weaker signals. In Appendix E1 we show how applying a Sparse Reconstruction method leads to an over-sparse emission field recovery in this case study.

Consequently, we expect substantial improvements in sparse reconstruction performance when using a more appropriately tailored dictionary. Possible approaches include learning dictionaries via methods such as K-SVD (Aharon et al., 2006), or employing generative neural networks to learn a compressed latent subspace that captures typical emission fields using only a small number of parameters.

3.3 Full CO_2 Inventory Recovery

In the final case study, we are investigating the full recovery of the CO_2 yearly average German national inventory from 2024 (here referred to as UBA inventory) (Gniffke et al., 2024) and we test the recovery of this emission field with and without spatial prior information. For the models that use prior information, we provide a European CO_2 inventory from 2018 (hereafter referred to as the TNO inventory) (Dellaert et al., 2019). Both inventories have been regridded from a 0.008 latitude × 0.016 longitude to a 0.01 lat/lon grid to match the shape of our generated footprint grid and are visualized in the top two plots of Fig. 6. Both look quite similar in structure and emission strength; however, most notably, the southern power plant present in the UBA inventory is missing from our version of the prior inventory to test how the inversion methods behave under the influence of a strong, unknown point source. Furthermore, the emissions in the UBA Inventory are generally lower than in the TNO inventory, especially in the city center. Only the main ring road around Munich is marked to have higher emissions in comparison.

3.3.1 Inversion with a spatial prior

The reconstructions of the different inversion methods when using a non-zero prior are illustrated in Fig. 6. The hybrid methods substantially outperform the pure Bayesian Inversion variants and yield the most accurate reconstructions for this emission field. They recover the inner urban area well, while placing greater reliance on the prior in the outer regions, where observational information is sparse.

It is noteworthy that the airport located to the northeast, visualized in the prior as a black dot in Fig. **??**, is retained as a point source in the reconstructions, even though it appears as an area source in the true emission field. This behavior can be attributed to the limited sensitivity of the forward model to emissions originating from the outer regions, as already shown briefly in the point-source recovery case study.

Table 4 quantifies how effectively the different models improve upon the prior toward the true emission field. The values overall agree with the emission field plots, as the hybrid methods are outperforming the standard Bayesian Inversion with a Gaussian prior. In addition, Appendix H1 presents difference plots between the final inversion and the prior emission field, highlighting how well each method captures the transition from the prior to the true emission field.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

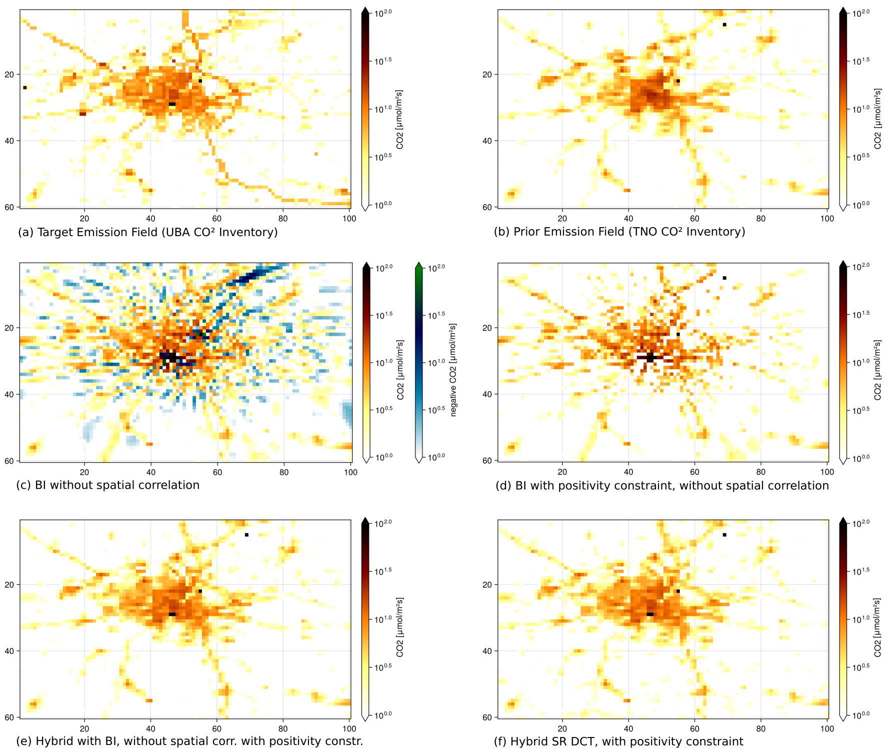

Figure 6. (a) shows the target emission field and (b) the used prior for the inversion the TNO CO_2 inventory. (c-f) show the results of the respective recovery methods for 30 dB noise. The methods that use just BI clearly oversmooth the point source in the southern powerplant, while the hybrid methods capture the point source well.

Table 4. Metrics for selected methods at 15 dB and 30 dB SNR, computed for the emission field with the power plants being excluded from the metric calculations as the sparse methods recover those very well. The metrics are evaluated by comparing the 'true emission field minus the prior' with the 'inversion minus the prior'. In this way, we can quantify how effectively each method improves upon the prior towards the true emission field.

	SNR = 15 dB						SNR = 30 dB					
Method	SSIM	Total Rel. Error	R^2	Regional Rel. Error 13×13	Regional Rel. Error 5×5	Local Rel. Error	SSIM	Total Rel. Error	R^2	Regional Rel. Error 13×13	Regional Rel. Error 5×5	Local Rel. Error
Hyb-BI (no corr.) Pos	0.1219	-0.3631	0.357	0.3704	0.4884	0.7085	0.2675	-0.2506	0.631	0.2518	0.3466	0.5369
Hyb-BI (1.0 km) Pos	0.1205	-0.2781	0.390	0.3129	**0.4606**	0.6899	**0.2772**	-0.1708	**0.644**	0.1835	**0.3182**	**0.5270**
Hyb-DCT Pos	**0.1264**	**0.0457**	**0.489**	**0.2698**	0.4934	**0.6315**	0.2748	**0.0208**	0.634	**0.1287**	0.3192	0.5344
BI (no corr.) Pos	0.0934	0.1493	-5.826	0.4205	0.7836	2.3079	0.2117	0.0720	-1.483	0.1944	0.4529	1.3920
BI (1.0 km) Pos	0.0809	0.1957	-6.895	0.4929	0.9323	2.4821	0.1594	0.0891	-2.330	0.2238	0.5436	1.6121
BI (no corr.)	0.0499	0.1117	-7.155	0.5571	1.1465	2.5227	0.1290	0.1107	-7.757	0.3090	0.7536	2.6142
BI (1.0 km)	0.0252	0.1381	-13.713	1.0023	2.5713	3.3884	0.0390	0.1134	-18.344	0.7731	2.6553	3.8853

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

3.3.2 Recovery without a spatial prior

For this case study, we set the prior emission field $\boldsymbol{x}_a$ to be zero. Such a scenario is especially relevant for countries or cities where high-quality, spatially gridded inventories are unavailable or not very trustworthy; i.e., the risk of adding incorrect information via a prior emission field is deemed greater than simply recovering the emission field from zero.

As shown in Fig. 7, the hybrid method with a positivity constraint reconstructs the emission field quite convincingly, whereas BI with a Gaussian prior and a positivity constraint struggles to resolve the power plant and the underlying structure. The metrics in Table 5 underline these observations. Adding spatial correlation to the Bayesian model only worsens the performance, as the power plants can not be correctly recovered. Here, for low noise, the hybrid method using BI in the second step outperforms the Hybrid DCT method slightly, as the latter produces some unwanted artifacts in the outer Munich area. The metrics table also shows that for high noise, none of the methods are able to recover the emission field, while the hybrid methods still achieve the best results overall.

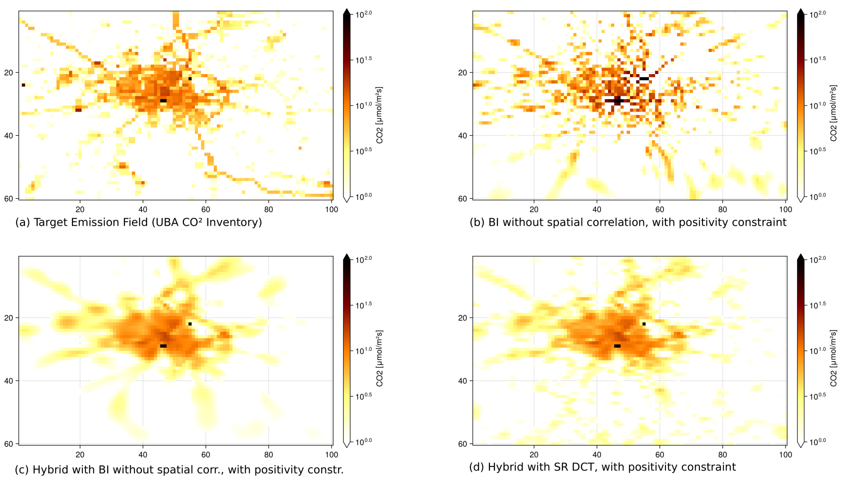

Figure 7. Recovery with zero mean prior for 30 dB noise. (a) shows the target emission field, while (b-d) show the respective recoveries for different methods.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Table 5. Priorless inversion: Metrics for selected methods at 15 dB and 30 dB SNR with power plants excluded.

	SNR = 15 dB						SNR = 30 dB					
Method	SSIM	Total Rel. Error	R^2	Regional Rel. Error 13×13	Regional Rel. Error 5×5	Local Rel. Error	SSIM	Tolal Rel. Error	R^2	Regional Rel. Error 13×13	Regional Rel. Error 5×5	Local Rel. Error
Hyb-BI (no corr.) Pos	0.0464	-0.5733	**0.024**	**0.7051**	**0.8860**	**0.9878**	0.1476	-0.3467	**0.853**	**0.3680**	**0.6001**	**0.3839**
Hyb-BI (1.0 km) Pos	**0.0827**	1.5435	-0.202	1.0224	1.2062	1.0962	**0.1572**	**0.0024**	0.852	0.4106	0.6320	0.3847
Hyb-DCT Pos	0.0515	1.0636	-0.054	0.7545	0.9610	1.0267	0.1482	0.6103	0.846	0.4455	0.6897	0.3918
BI (no corr.) Pos	0.0731	2.1349	-0.374	1.3379	1.6556	1.1721	0.1312	1.2042	0.436	0.6962	1.0266	0.7512
BI (1.0 km) Pos	0.0601	2.8307	-0.757	1.7688	2.1959	1.3253	0.1216	2.0210	0.265	0.9348	1.4038	0.8574
BI (no corr.)	0.0289	**0.0119**	-0.598	1.6811	2.1163	1.2639	0.0661	0.5909	-0.574	1.0539	1.5550	1.2547
BI (1.0 km)	0.0224	0.1381	-0.900	2.4101	3.3189	1.3784	0.0381	0.5079	-1.417	2.2342	4.1562	1.5545

4 Discussion

We introduce a new approach for urban atmospheric emission field inversions, which treats point and area sources in two separate steps. In this way, we show how we can overcome the limitation of Bayesian inversion using a Gaussian prior, which behaves poorly under the influence of big deviations from the prior. However, in this study, we deliberately ignored several factors relevant to real-world inversion problems to focus solely on the influence of the recovery algorithm. These would be, for example, errors arising from the transport model or background concentration influences that go beyond the scope of this paper, but need to be addressed in subsequent research. As a result, our analysis represents a controlled comparison of inversion algorithms under idealized conditions.

Our results demonstrate that the choice of regularization plays a critical role in the reconstruction quality. In particular, the blind use of a quadratic penalty, as is the case with Bayesian Inversion using a Gaussian prior, already leads to substantial reconstruction errors when dealing with bigger point sources that are not represented in inventory-based priors. This highlights the importance of carefully testing and validating inversion methods under simplified settings before applying them to more realistic cases. Only by doing so can one reliably identify which components of the inversion pipeline are responsible for specific errors when performance degrades under less favorable recovery conditions.

For smooth area source recoveries, BI using a Gaussian prior and SR using the DCT transform perform equally well overall. However, there is still significant room for improvement in the choice of transformation used for sparse reconstruction. While the discrete cosine transform provides a simple and effective baseline, more advanced approaches such as learned dictionaries using K-SVD (Aharon et al., 2006) or generative models could provide representations better suited for smoothly varying emission fields.

The proposed hybrid method is largely non-intrusive and can be integrated into existing Bayesian inversion pipelines with minimal modifications.

In contrast to BI, a major limitation of sparse reconstruction methods is the treatment of uncertainty (Basu et al., 2021). While Bayesian formulations of sparsity-promoting methods, such as the Bayesian LASSO (Park and Casella, 2008), do exist,

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

inference in these models can be computationally challenging and warrants further investigation. As a practical alternative, bootstrap-based uncertainty estimation or a Gaussian Bayesian inversion on the support of the LASSO inversion, as shown in (Basu et al., 2021) and utilized in (Zanger et al., 2022; Hase et al., 2017), can be applied directly and provides useful first approximations.

5 Conclusions

In this work, we introduced the Hybrid Inversion approach as a novel atmospheric greenhouse gas inversion framework. The method is designed to handle both smooth and sparse emission field structures commonly encountered in urban environments by recovering focused and dispersed sources in two separate steps.

First, we compare Bayesian inversion using a Gaussian prior (BI) and Sparse Reconstruction (SR) for atmospheric emission inversions across a range of controlled observing system simulation experiments. Our results demonstrate that Bayesian inversion with a Gaussian prior performs well at recovering smooth, spatially extended emission patterns but struggles to resolve localized point sources. In contrast, Sparse Reconstruction methods are well-suited for identifying such point sources, even when they account for only a fraction of the total emission field.

Consequently, our Hybrid Inversion approach combines the strengths of both methods and consistently delivers the most robust inversion performance when reconstructing typical urban emission fields. Crucially, it enables the possibility of an inversion without a spatial prior, i.e. an inversion which does not make use of prior knowledge given by greenhouse gas inventories in the presence of large unknown point sources. Further, the hybrid framework seamlessly integrates into existing Bayesian inversion workflows by appending SR as an initial step to address point sources, thereby mitigating the performance limitations of BI alone.

Future work should focus on improving sparse representations via learned dictionaries or generative models, as well as on developing more reliable uncertainty quantification methods for sparsity-based inversion.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix A: More on Bayesian Inversion with Gaussian prior

For Bayesian Inversion using a Gaussian prior, we solve for the *Maximum A Posteriori (MAP)* estimator by calculating the minimizer

$$\hat{\boldsymbol{x}} = \underset{x}{\operatorname{argmin}} \quad (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y})^T \boldsymbol{S}_{\varepsilon}^{-1} (\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}) + (\boldsymbol{x} - \boldsymbol{x}_a)^T \boldsymbol{S}_a^{-1} (\boldsymbol{x} - \boldsymbol{x}_a). \tag{A1}$$

As this minimization formula is purely a quadratic function, it is possible to calculate the unique solution $\hat{x}$ analytically by computing the derivative and solving for zero. This results in the solution formula

$$\hat{\boldsymbol{x}} := \boldsymbol{x}_a + \boldsymbol{G}(\boldsymbol{y} - \boldsymbol{A}\boldsymbol{x}_a) \tag{A2}$$

with

$$\boldsymbol{G} := \left(\boldsymbol{A}^T \boldsymbol{S}_{\varepsilon}^{-1} \boldsymbol{A} + \boldsymbol{S}_a^{-1}\right)^{-1} \boldsymbol{A}^T \boldsymbol{S}_{\varepsilon}^{-1} = \boldsymbol{S}_a A^T \left(\boldsymbol{A}\boldsymbol{S}_a \boldsymbol{A}^T + \boldsymbol{S}_{\varepsilon}\right)^{-1}. \tag{A3}$$

The matrix $\boldsymbol{G}$ is called *gain matrix* and both derivations we provide above are equivalent.
Furthermore, since we use a Gaussian distribution for the prior and observation error, the MAP also has a Gaussian distribution with the posterior covariance matrix $\hat{S}$. The posterior covariance can be seen as an update to the prior covariance, possibly reducing uncertainties at locations with higher footprint sensitivity. The posterior covariance matrix can be computed with

$$\hat{S} = \left(A^T S_{\varepsilon}^{-1} A + S_a^{-1}\right)^{-1} \tag{A4}$$

In summary, the posterior can be modeled as a multivariate Gaussian distribution with

$$\hat{x} \sim \mathcal{N}(x_a + G(y - Ax_a), \hat{S}) \tag{A5}$$

and we can make use of this distribution to for example provide a credibility interval around the flux estimates for each cell.

For more intuition and information about this method, we refer to (Rodgers, 2000; Jones et al., 2021).

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix B: TNO Methane Inventory

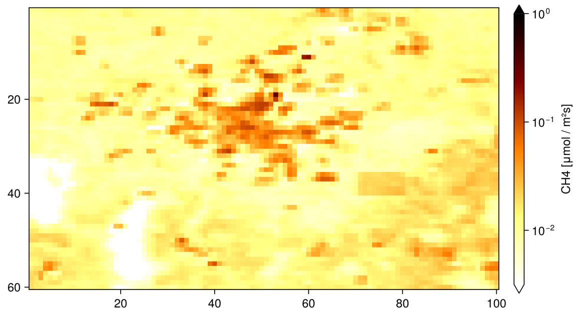

Figure B1. We use the 2018 TNO methane inventory (Dellaert et al., 2019) in two of our case studies as the prior.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix C: Experimental Setup

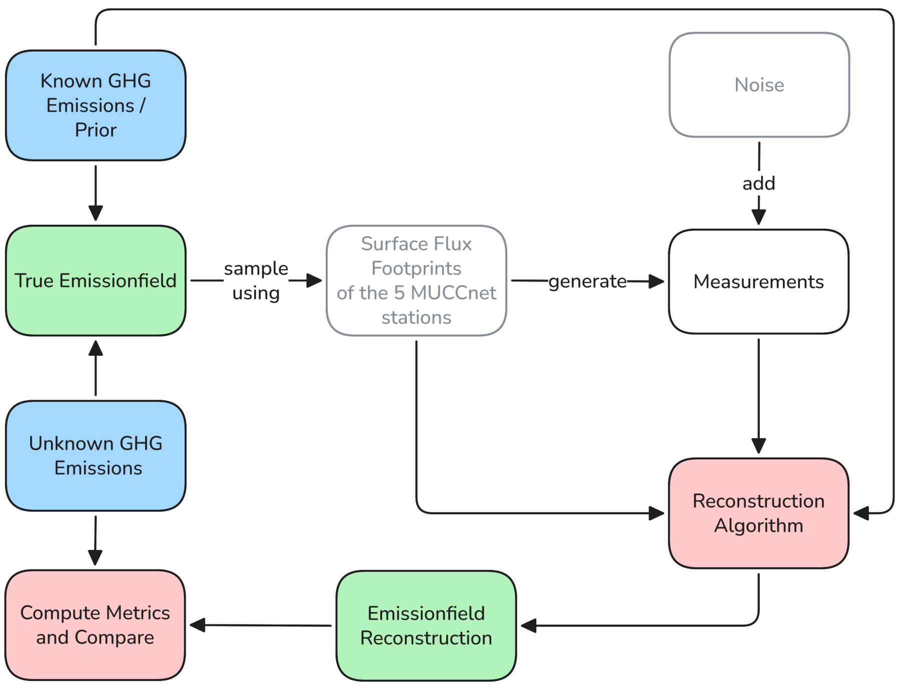

Figure C1. This figure shows our setup for all the numerical experiments in the paper.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix D: Choosing the cutoff value for the Hybrid Method

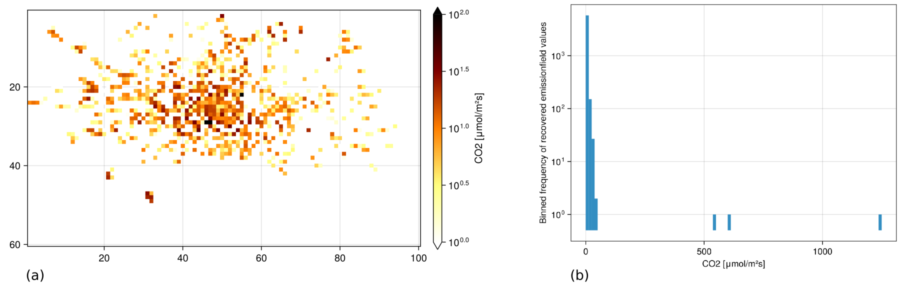

Figure D1. With these figures, we visualize the results after the initial first step of the hybrid method explained in Algorithm 1, when applied to the last case study, where we try to recover the UBA inventory without a spatial prior. (a) shows the over-sparse recovery of the emission field after applying a plain L1 regularized inversion. We can see that the power plants (marked with black color in the inversion) are recovered well, in contrast to the rough background. To separate both, we can plot a histogram of the sparse recovery, shown in (b). We can see that the power plants are nicely separated in terms of emission strength. Note that one power plant is recovered by two emission field grid cells, each with a bit more than $500 \mu\mathrm{mol}/\mathrm{m}^2\mathrm{s}$. This power plant is lying directly on the edge of both grid cells. Using the histogram, we can choose an arbitrary cutoff value of e.g. $250 \mu\mathrm{mol}/\mathrm{m}^2\mathrm{s}$ to separate the point sources from the background.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix E: Too sparse area source recovery

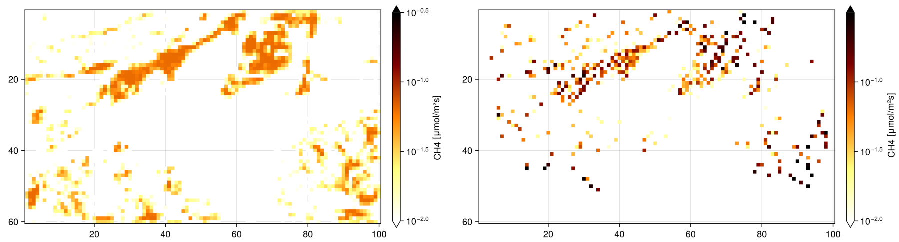

(a) Target Emission Field minus the Prior (b) Sparse Reconstruction without a Transformation

Figure E1. (a) shows the target emission field from case study 2 for reference. In (b), we show how a sparse recovery suffers from an over-sparse recovery, as we do not apply any transformation that makes a smooth area source recovery possible.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix F: Barplots on the point recovery case Study

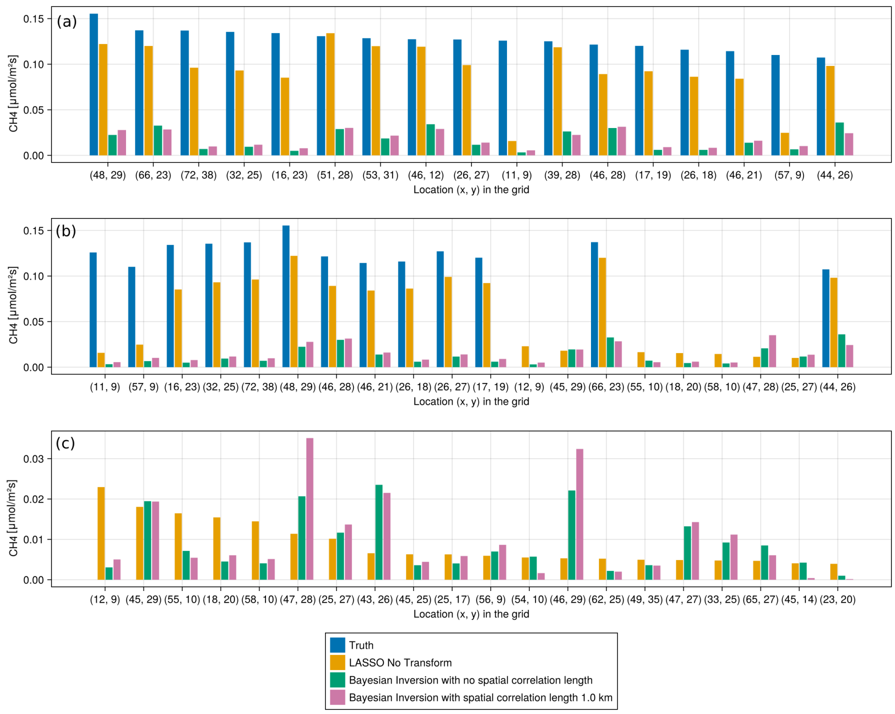

Figure F1. Barplots comparing absolute values of SR and BI point source recoveries. The prior has been subtracted from all values, and we add the coordinates of the observed 100x60 grid on the x-axis. The barplot in (a) is sorted by the biggest point sources in the true emission field. We can see that the best performance seems to be at the point sources that are close to the measurement locations or in the inner city, while the distant point source at (11,9) is only very barely recovered. (b) is ordered by the biggest error deviations of LASSO from the truth, and (c) by the biggest false positives of the SR via the LASSO method. The biggest false positive at (12,9) is an emission right next to a true actual source at (11,9), suggesting that perfect recovery in low-sensitivity areas is not guaranteed.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix G: Barplots for specific squares in the Point Recovery Case Study

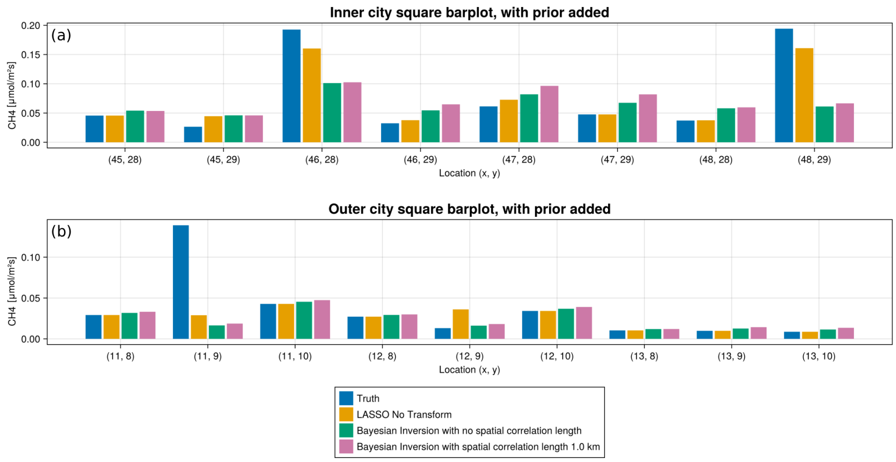

Figure G1. These barplots visualize the behavior of SR and BI in the close vicinity of point sources in the inner and outer city areas. The prior is added to the bar plots. BI smoothes over point sources and increases fluxes in the general surrounding area. SR, on the other hand, captures the location of the point source well under good conditions but can misattribute sources when they are located in the outer area, where recovery conditions are less favorable.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Appendix H: Difference Plots

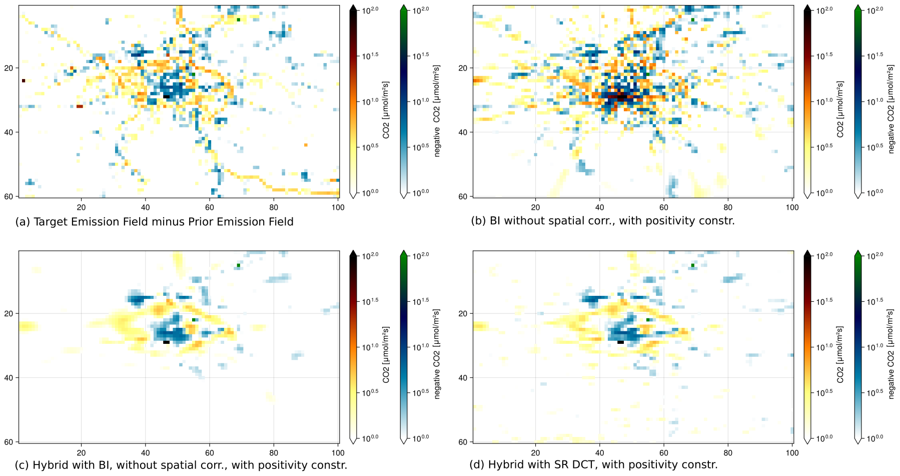

Figure H1. In this figure, we show how well the inversions transition from prior to the true emission field for 30dB noise. To do this we subtract the prior from the true emission field and the inversions. The hybrid methods are both able to capture the big point source and the surrounding area well, while BI using a Gaussian prior is unstable. (a) shows the target emission field minus the prior. (b-d) shows the respective recovered emission fields minus the prior for the shown different methods. We can see that BI in (b) behaves unstably around the power plant point sources, while both hybrid methods in (c-d) can recover the big point sources and the smooth underlying area sources well.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Code and data availability. The code and the footprints to perform the inversions is available on zenodo: DOI 10.5281/zenodo.20285333 (Grasberger et al., 2026) and only the code in the Github repository https://github.com/tum-esm/hybrid-sparse-inversion. The TNO_GHGco_v1.1 inventory (Dellaert et al., 2019) and the UBA 2024 inventory (Gniffke et al., 2024) are not publicly available.

Author contributions. JC conceived the concept. MM provided the footprints, and JS provided data preprocessing. TG devised the method and performed the analysis supervised by JC, FK, BZ and AL. TG wrote the paper with help and feedback from JC, FK, BZ, AL, MM and JS. TG revised the paper according to the reviewers' comments. JC and FK acquired the funding and guided the project.

Competing interests. There are no competing interests.

Acknowledgements. This work received funding by the ERC Consolidator Grant CoSense4Climate (grant no. 101089203, PI: Jia Chen), and from the Munich Data Science Institute (MDSI) at the Technical University of Munich through its Kickstarter Research Program under the project name MES. This work has received partly funding by the ICOS PAUL project: PAUL, Pilot Applications in Urban Landscapes – Towards integrated city observatories for greenhouse gases (ICOS Cities), funded by the European Union's Horizon 2020 Research and Innovation Programme under grant agreement no. 101037319.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

References

Aharon, M., Elad, M., and Bruckstein, A.: K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation, IEEE Transactions on signal processing, 54, 4311–4322, 2006.

Antonello, N., Stella, L., Patrinos, P., and Van Waterschoot, T.: Proximal gradient algorithms: Applications in signal processing, arXiv preprint arXiv:1803.01621, 2018.

Arantes, M. K., Alves, H. J., Sequinel, R., and da Silva, E. A.: Treatment of brewery wastewater and its use for biological production of methane and hydrogen, International Journal of Hydrogen Energy, 42, 26 243–26 256, 2017.

Basu, T., Einbeck, J., and Troffaes, M. C.: Uncertainty quantification in lasso-type regularization problems, Optimization Under Uncertainty with Applications to Aerospace Engineering, pp. 81–109, 2021.

CARTO: CARTO Basemaps, https://carto.com/basemaps, basemap tiles. Accessed January 15, 2026, 2024.

Chen, J., Viatte, C., Hedelius, J. K., Jones, T., Franklin, J. E., Parker, H., Gottlieb, E. W., Wennberg, P. O., Dubey, M. K., and Wofsy, S. C.: Differential column measurements using compact solar-tracking spectrometers, Atmospheric Chemistry and Physics, 16, 8479–8498, 2016.

Chen, J., Dietrich, F., Maazallahi, H., Forstmaier, A., Winkler, D., Hofmann, M. E., Denier van der Gon, H., and Röckmann, T.: Methane emissions from the Munich Oktoberfest, Atmospheric Chemistry and Physics, 20, 3683–3696, 2020.

Chen, S. and Donoho, D.: Basis pursuit, in: Proceedings of 1994 28th Asilomar Conference on Signals, Systems and Computers, vol. 1, pp. 41–44, IEEE, 1994.

Chen, S. S., Donoho, D. L., and Saunders, M. A.: Atomic decomposition by basis pursuit, SIAM review, 43, 129–159, 2001.

Chung, J., Jiang, J., Miller, S. M., and Saibaba, A. K.: Hybrid projection methods for solution decomposition in large-scale Bayesian inverse problems, SIAM Journal on Scientific Computing, 46, S97–S119, 2024.

Cusworth, D. H., Jacob, D. J., Sheng, J.-X., Benmergui, J., Turner, A. J., Brandman, J., White, L., and Randles, C. A.: Detecting high-emitting methane sources in oil/gas fields using satellite observations, Atmospheric Chemistry and Physics, 18, 16 885–16 896, 2018.

Dellaert, S., Super, I., Visschedijk, A., and Denier van der Gon, H.: High resolution scenarios of CO2 and CO emissions, 2019.

Dietrich, F., Chen, J., Voggenreiter, B., Aigner, P., Nachtigall, N., and Reger, B.: MUCCnet: Munich urban carbon column network, Atmospheric Measurement Techniques, 14, 1111–1126, 2021.

Forstmaier, A., Chen, J., Dietrich, F., Bettinelli, J., Maazallahi, H., Schneider, C., Winkler, D., Zhao, X., Jones, T., Van Der Veen, C., et al.: Quantification of methane emissions in Hamburg using a network of FTIR spectrometers and an inverse modeling approach, Atmospheric chemistry and physics, 23, 6897–6922, 2023.

Foucart, S. and Rauhut, H.: A Mathematical Introduction to Compressive Sensing, Applied and Numerical Harmonic Analysis, Springer, New York, NY, https://doi.org/10.1007/978-0-8176-4948-7, 2013.

Gniffke, P., Kotzulla, M., and Hausmann, K.: German Informative Inventory Report 2024, Web report, Umweltbundesamt (UBA), Berlin, 2024.

Grasberger, T., Chen, J., Zanger, B., Lupoli, A., Oliveira Makowski, M., Stauber, J., and Krahmer, F.: Supplements for Urban Atmospheric Inverse Modeling via Hybrid Reconstruction Methods based on Compressed Sensing, https://doi.org/10.5281/zenodo.20285334, 2026.

Guanter, L., Warren, J., Omara, M., Chulakadabba, A., Roger, J., Sargent, M., Franklin, J. E., Wofsy, S. C., and Gautam, R.: Remote sensing of methane point sources with the MethaneAIR airborne spectrometer, EGUsphere, 2025, 1–22, 2025.

Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual, https://www.gurobi.com, 2024.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Hase, N., Miller, S. M., Maaß, P., Notholt, J., Palm, M., and Warneke, T.: Atmospheric inverse modeling via sparse reconstruction, Geoscientific Model Development, 10, 3695–3713, 2017.

Hirst, B., Jonathan, P., del Cueto, F. G., Randell, D., and Kosut, O.: Locating and quantifying gas emission sources using remotely obtained concentration data, Atmospheric environment, 74, 141–158, 2013.

IPCC: Climate Change 2023: Synthesis Report, https://doi.org/10.59327/IPCC/AR6-9789291691647, 2023.

Jacob, D. J., Turner, A. J., Maasakkers, J. D., Sheng, J., Sun, K., Liu, X., Chance, K., Aben, I., McKeever, J., and Frankenberg, C.: Satellite observations of atmospheric methane and their value for quantifying methane emissions, Atmospheric Chemistry and Physics, 16, 14 371–14 396, 2016.

Jones, T. S., Franklin, J. E., Chen, J., Dietrich, F., Hajny, K. D., Paetzold, J. C., Wenzel, A., Gately, C., Gottlieb, E., Parker, H., et al.: Assessing urban methane emissions using column-observing portable Fourier transform infrared (FTIR) spectrometers and a novel Bayesian inversion framework, Atmospheric Chemistry and Physics, 21, 13 131–13 147, 2021.

Kuenen, J., Visschedijk, A., Jozwicka, M., and Denier Van Der Gon, H.: TNO-MACC_II emission inventory; a multi-year (2003–2009) consistent high-resolution European emission inventory for air quality modelling, Atmospheric Chemistry and Physics, 14, 10 963–10 976, 2014.

Kühbacher, D., Chen, J., Baertschi, J., Khan, A. A., Wenzel, A., Aigner, P., Grange, S., Rubli, P., and Emmenegger, L.: SCOUT: Street-Level Carbon Observatory for Urban Terrain, in: ICOS Science Conference 2024, https://www.icos-cp.eu/news-and-events/science-conference/icos2024sc/all-abstracts, 2024.

Kühbacher, D., Chen, J., Aigner, P., Ilic, M., Super, I., and Denier van der Gon, H.: DRIVE v1. 0: a data-driven framework to estimate road transport emissions and temporal profiles, Geoscientific Model Development, 18, 9967–9990, 2025.

Kutyniok, G.: Compressed Sensing, Mitteilungen der Deutschen Mathematiker-Vereinigung, 22, 24–29, https://doi.org/doi:10.1515/dmvm-2014-0014, 2014.

Lin, J., Gerbig, C., Wofsy, S., Andrews, A., Daube, B., Davis, K., and Grainger, C.: A near-field tool for simulating the upstream influence of atmospheric observations: The Stochastic Time-Inverted Lagrangian Transport (STILT) model, Journal of Geophysical Research: Atmospheres, 108, 2003.

Nair, M. T., Schock, E., and Tautenhahn, U.: Morozov's discrepancy principle under general source conditions, Zeitschrift für Analysis und ihre Anwendungen, 22, 199–214, 2003.

Needell, D. and Tropp, J. A.: CoSaMP: Iterative signal recovery from incomplete and inaccurate samples, Applied and computational harmonic analysis, 26, 301–321, 2009.

Park, T. and Casella, G.: The bayesian lasso, Journal of the american statistical association, 103, 681–686, 2008.

Pati, Y. C., Rezaiifar, R., and Krishnaprasad, P. S.: Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition, in: Proceedings of 27th Asilomar conference on signals, systems and computers, pp. 40–44, IEEE, 1993.

Poupkou, A., Symeonidis, P., Ziomas, I., Melas, D., and Markakis, K.: A spatially and temporally disaggregated anthropogenic emission inventory in the southern Balkan region, Water, Air, and Soil Pollution, 185, 335–348, 2007.

Rauhut, H., Schnass, K., and Vandergheynst, P.: Compressed sensing and redundant dictionaries, IEEE Transactions on Information Theory, 54, 2210–2219, 2008.

Ray, J., Yadav, V., Michalak, A., van Bloemen Waanders, B., and McKenna, S. A.: A multiresolution spatial parameterization for the estimation of fossil-fuel carbon dioxide emissions via atmospheric inversions, Geoscientific Model Development, 7, 1901–1918, 2014.

https://doi.org/10.5194/egusphere-2026-3161
Preprint. Discussion started: 9 September 2026

Ray, J., Lee, J., Yadav, V., Lefantzi, S., Michalak, A., and van Bloemen Waanders, B.: A sparse reconstruction method for the estimation of multi-resolution emission fields via atmospheric inversion, Geoscientific Model Development, 8, 1259–1273, 2015.

Rodgers, C. D.: Inverse methods for atmospheric sounding: theory and practice, vol. 2, World scientific, 2000.

Saarnio, S., Winiwarter, W., and Leitao, J.: Methane release from wetlands and watercourses in Europe, Atmospheric Environment, 43, 1421–1429, 2009.

Scherzer, O.: The use of Morozov's discrepancy principle for Tikhonov regularization for solving nonlinear ill-posed problems, Computing, 51, 45–60, 1993.

Sijikumar, S., Raju, A., Valsala, V., Tiwari, Y., Girach, I., Jain, C. D., and Ratnam, M. V.: High-Resolution Bayesian Inversion of Carbon Dioxide Flux Over Peninsular India, Atmospheric Environment, 308, 119 868, 2023.

Staufer, J., Broquet, G., Bréon, F.-M., Puygrenier, V., Chevallier, F., Xueref-Rémy, I., Dieudonné, E., Lopez, M., Schmidt, M., Ramonet, M., et al.: The first 1-year-long estimate of the Paris region fossil fuel CO 2 emissions based on atmospheric inversion, Atmospheric Chemistry and Physics, 16, 14 703–14 726, 2016.

Store, C. C. D. et al.: ERA5 hourly data on single levels from 1940 to present, Copernicus Climate Change Service (C3S), 2023.

Strang, G.: The discrete cosine transform, SIAM review, 41, 135–147, 1999.

Taylor, T. E., Eldering, A., Merrelli, A., Kiel, M., Somkuti, P., Cheng, C., Rosenberg, R., Fisher, B., Crisp, D., Basilio, R., et al.: OCO-3 early mission operations and initial (vEarly) XCO2 and SIF retrievals, Remote Sensing of Environment, 251, 112 032, 2020.

Tegetmeyer, C., Kaiser, M., Chobanova, M., Tanneberger, F., and Barthelmes, A.: The European Peatland Map 2025 (EPM2025): Geodata, geospatial dataset, 2025.

Thompson, R. L. and Stohl, A.: FLEXINVERT: an atmospheric Bayesian inversion framework for determining surface fluxes of trace species using an optimized grid, Geoscientific Model Development, 7, 2223–2242, 2014.

Tibshirani, R.: Regression shrinkage and selection via the lasso, Journal of the Royal Statistical Society Series B: Statistical Methodology, 58, 267–288, 1996.

van der Gon, H. D., Kuenen, J. J., Janssens-Maenhout, G., Döring, U., Jonkers, S., and Visschedijk, A.: TNO_CAMS high resolution European emission inventory 2000–2014 for anthropogenic CO 2 and future years following two different pathways, Earth System Science Data Discussions, 2017, 1–30, 2017.

van der Gon, H. D., Kuenen, J., Boleti, E., Muntean, M., Maenhout, G., Marshall, J., and Haussaire, J.: Emissions and natural fluxes Dataset, 2019.

Vogel, F., Ars, S., Wunch, D., Lavoie, J., Gillespie, L., Maazallahi, H., Ro¨ckmann, T., Necki, J., Bartyzel, J., Jagoda, P., et al.: Ground-based mobile measurements to track urban methane emissions from natural gas in 12 cities across eight countries, Environmental Science & Technology, 58, 2271–2281, 2024.

Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P.: Image quality assessment: from error visibility to structural similarity, IEEE transactions on image processing, 13, 600–612, 2004.

Zanger, B., Chen, J., Sun, M., and Dietrich, F.: Recovery of sparse urban greenhouse gas emissions, Geoscientific Model Development, 15, 7533–7556, 2022.