Geoscientific
Model Development
Discussions

Open Access

A robust method for inverse transport modelling of atmospheric emissions using blind outlier detection

M. Martinez-Camara[1], B. Béjar Haro[1], A. Stohl[2], and M. Vetterli[1]

[1]School of Computer and Communication Sciences, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland
[2]Norwegian Institute for Air Research (NILU), 2027 Kjeller, Norway

Correspondence to: M. Martinez-Camara (marta.martinez-camara@epfl.ch)

Abstract

Emissions of harmful substances into the atmosphere are a serious environmental concern. In order to understand and predict their effects, it is necessary to estimate the exact quantity and timing of the emissions, from sensor measurements taken at different locations. There exists a number of methods for solving this problem. However, these existing methods assume Gaussian additive errors, making them extremely sensitive to outlier measurements. We first show that the errors in real-world measurement datasets come from a heavy-tailed distribution, i.e., include outliers. Hence, we propose to robustify the existing inverse methods by adding a blind outlier detection algorithm. The improved performance of our method is demonstrated on a real dataset and compared to previously proposed methods. For the blind outlier detection, we first use an existing algorithm, RANSAC, and then propose a modification called TRANSAC, which provides a further performance improvement.

1 Introduction

1.1 Motivation

Emissions of harmful substances into the atmosphere occur all the time. Examples include nuclear power plant accidents, volcano eruptions, and releases of greenhouse gases. However, these emissions are difficult to quantify. Depending on the scenario, measurement networks on scales from local to global may be needed. A robust technical framework to estimate the emissions properly from such measurements is also necessary.

This technical framework consists of three elements: measurements, atmospheric dispersion models, and inverse methods tailored to this specific linear inverse problem.

There has been a clear effort in deploying more reliable, precise, and extended sensor networks (CTBTO, 2014). Also, there has been an evident development of precise

atmospherical dispersion models (Holmes and Morawska, 2006). However, inverse methods are still in a relatively early stage of development.

These inverse methods are technically complex, and require a multidisciplinary approach; collaboration among researchers from different fields is necessary for further advances.

1.2 Related work

Atmospheric dispersion models such as Eulerian or Lagrangian particle dispersion models (LPDMs) (Zannetti, 1990) allow us to relate the source to the measurements in a linear way:

$$y = \overline{\mathbf{A}}x + n, \tag{1}$$

where y is the measurement vector, x is the source term, $\overline{\mathbf{A}}$ is the transport matrix, and n is the measurement error.

LPDMs have some advantages with respect to the Eulerian ones: they can have infinite temporal and spatial resolution; they avoid the artificial initial diffusion of a point source in the corresponding cell and the advection numerical errors; and they are computationally more efficient (Zannetti, 1990).

There are only a few freely available, open source implementations of LPDMs. FLEXPART (Stohl et al., 2005) is one of them. It has been used and validated in a large number of studies on long-range atmospheric transport (Stohl et al., 1998). Here we use it to derive $\mathbf{A}$, which is an estimate of the true transport matrix $\overline{\mathbf{A}}$.

It is clear from Eq. (1) that estimating the source means solving a linear inverse problem. Most environmental scientists use a least-squares approach with the Tikhonov (ℓ_2-norm) regularization, or variants of this method, to recover an estimate $\hat{x}$ of the source:

$$\hat{x} = \arg\min_{x} \|\mathbf{A}x - y\|_2 + \lambda\|x\|_2, \tag{2}$$

A robust method for inverse transport modelling

M. Martinez-Camara et al.

Title Page

Abstract	Introduction
Conclusions	References
Tables	Figures

◄	►
◄	►
Back	Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

where $\lambda \geq 0$ is the regularization parameter.

For example, in Seibert (2001), the Tikhonov regularization is combined with a smooth first derivative constraint:

$$\hat{x} = \arg\min_{x} \|\mathbf{A}x - y\|_2 + \lambda \|x\|_2 + \beta \|Dx\|_2. \tag{3}$$

Also, an a priori solution can be introduced to the Tikhonov regularization such as in (Stohl et al., 2012):

$$\hat{x} = \arg\min_{x} \|\mathbf{A}x - y\|_2 + \lambda \|x - x_a\|_2, \tag{4}$$

In Winiarek et al. (2012), the Tikhonov regularization is used with a non-negative constraint. A slightly different approach is the use of a sparsity constraint together with a non-negative constraint as in Martinez-Camara et al. (2013). Yet another point of view is given in Bocquet (2007), where both the source and the error distributions are estimated at the same time.

All these approaches minimize the energy of the disagreement between the model and the observations, while at the same time keeping the energy of the solution in check. While this is a reasonable approach, no metrics of real performance are (or can be) given in most of these studies, simply because no knowledge of the ground truth is available. This fact made it impossible to evaluate the true performance of any of these approaches.

However, a few controlled tracer experiments have been performed – the most important ones in Europe and in the US. They present exceptional opportunities to study model and measurement errors, as well as to develop and test the various source recovery algorithms.

The European Tracer EXperiment (ETEX) (Nodop et al., 1998) was established to evaluate the validity of long-range transport models. Perfluorocarbon (PFC) tracers were released into the atmosphere in Monterfil, Brittany, in 1994. Air samples were taken at 168 stations in 17 European countries for 72 h after the release. The data

collected in the ETEX experiment and the correspondent matrix estimated by FLEX-PART (hourly source discretisation, 120 unknowns) are used for several purposes in this paper. We will refer to this data as the *ETEX dataset*.

1.3 Contributions

In this paper we show that the errors present in a source estimation problem come from a heavy-tailed distribution, which implies the presence of outliers in the measurement dataset. Typical source estimation algorithms like Eq. (2) assume Gaussian additive errors. This incorrect assumption makes them highly sensitive to outliers. In fact, if the outliers are removed, the source estimation using Eq. (2) improves substantially.

Hence, we propose to combine Eq. (2) with algorithms to detect and remove outliers blindly, i, e. without any knowledge of the ground truth. First we use a well-known algorithm for this task, RANdom SAmple Consensus (RANSAC) (Fischler and Bolles, 1981), and study its performance. Next, we propose a new algorithm which overcomes some of the weaknesses of RANSAC, and test its performance. The efficiency of both algorithms is demonstrated on a real-world dataset, and their performance is evaluated and compared to other existing methods.

2 Non-Gaussian noise

Given $\mathbf{A}$, the estimate of the transport matrix produced by FLEXPART, the forward model (Eq. 1) now becomes

$$y = \mathbf{A}x + e \tag{5}$$

where e is an additive error term that encompasses both the model and measurement errors.

In the ETEX experiment we have access to the measurements y, the true source x, and the estimated transport matrix $\mathbf{A}$. This permits us to study the errors e. Let us

model the components e_i of the vector $\boldsymbol{e}$ as random and independent and identically distributed. We can approximate their empirical probability distribution by plotting the histogram of the elements e_i.

Figure 1 shows graphically that the error has a heavy-tailed distribution. The distribution clearly deviates from a Gaussian one. This is confirmed by calculating the excess kurtosis of the sample distribution. The value of $g = 123.64$ indicates that the underlying distribution is strongly super-Gaussian.

Using the ℓ_2 norm in the loss function in Eq. (2) is optimal when the additive errors are Gaussian – which is not our case. Even worse, this loss function is very sensitive to outliers, just like the ones present in the heavy-tailed distribution shown in Fig. 1. Hence, the performance of Eq. (2) and its variants could be improved by additional processing, aimed at removing and/or marginalizing the outliers. In the present paper we propose and demonstrate a novel scheme for this additional processing.

3 Outlier detection

Imagine that we have an oracle which reveals to us the measurements corresponding to the largest errors (i.e. the outliers). If we remove these measurements from the dataset, the performance of Eq. (2) in terms of the reconstruction error or mean square error (MSE) improves significantly[1]. In order to illustrate this, we remove the measurements associated with the largest errors (sorted by magnitude) and observe the effect on the MSE. Figure 2 shows how the MSE decreases as more and more outliers are removed. Some oscillations may occur due to outlier compensation effects.

However, in a real-world problem, we do not have such an oracle. The question becomes: how could one locate the outliers blindly?

---

[1]The MSE is defined as $\frac{1}{n}\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2^2$, where $\hat{\boldsymbol{x}}$ is the estimated source, $\boldsymbol{x}$ is the real source (ground truth), and n is the number of elements in $\boldsymbol{x}$.

3.1 RANSAC

One of the simplest and most popular algorithms to localize outliers blindly is RANSAC. RANSAC has been widely and successfully used, mainly by the computer vision community. Figure 3 illustrates the operation of RANSAC, and Algorithm 1 describes it in pseudocode.

Given a dataset y with m measurements, select randomly a subset y' containing p measurements. Typically, $n < p < m$, where n is the number of unknowns in the problem. In Fig. 3, $m = 8$ and $p = 2$, and the subset is shown in red diamonds. Using Eq. (2) and y', estimate $\hat{x}$, and then compute the residual $r = A\hat{x} - y$. Now we can count how many of the original samples are *inliers*. For a given tolerance η, the set of inliers is defined as $\mathcal{L} = \{q \in \{1, 2, \ldots, m\} \mid \eta \geq (r[q])^2\}$. Repeat this process N times and declare the final solution x^* to be that estimate $\hat{x}$ which produced the most inliers. In Fig. 3, $N = 2$.

Note that other regularizations can be used instead of Eq. (2). Here we use the Tikhonov regularization because it is simple, general, and most other existing approaches are based on it. Nevertheless, if some properties of the source are known a priori (e.g., sparsity or smoothness), this step of the algorithm can be adapted accordingly.

At the stage where the N possible solutions $\hat{x}$ have been generated, what RANSAC actually tries to do is to select the solution x^* with the smallest MSE. However, in a real world problem the ground-truth is unknown, so we do not have access to the MSE itself. So, as mentioned above, RANSAC overcomes this difficulty by using an indirect metric of the MSE: it assumes that the number of inliers is inversely proportional to the MSE. Figure 3 depicts the intuition behind this in a simple 1-D problem: the superior solution (subset 1) produces more inliers than the inferior solution (subset 2). Thus, RANSAC maximizes the number of inliers, in hopes that this also minimizes the estimation error.

As we will see in the following sections, if the optimal value for the threshold parameter η is known and used, using RANSAC as a pre-processing stage for outlier removal

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

before applying Eq. (2) significantly improves the overall performance (compared to using only Eq. 2 with no outlier removal pre-processing). Unfortunately, the performance of RANSAC depends strongly on the parameter η, and finding the optimal value of η is an open problem.

Furthermore, the assumed inverse proportionality between the number of inliers and the MSE does not always hold in the presence of *critical measurements*. This is the case in the ETEX dataset, as we can see in Fig. 4a.

3.2 Critical measurements

We identify critical measurements as those which have the largest influence in the source estimation process. A quantitative measure of this influence is the Cook's distance (Cook, 1977). Figure 5 shows the Cook's distance of the ETEX measurements. It is easy to observe the peak that identifies the critical measurements.

Let us consider again the ETEX dataset, the set of N solutions $\hat{x}$ that RANSAC generates, and their corresponding residuals r. It is interesting to note that the solutions $\hat{x}$ with most inliers (the superior solutions according to RANSAC) have high residuals at exactly the critical measurements. This is shown in Fig. 6. In other words, by considering the critical measurements as outliers, these solutions achieve more inliers.

RANSAC assumes that all the measurements have the same influence: it just wants to maximize the number of inliers, and does not care about which exact measurements are the inliers. This is why it fails in this case and the inverse proportionality between the number of inliers and the MSE does not hold.

In summary, RANSAC operates reliably when all the measurements are of similar importance, because the inverse proportionality between MSE and the number of inliers holds. However, when critical measurements are present, this proportionality does not hold, and RANSAC fails.

3.3 TRANSAC

In order to avoid the weakness of the standard RANSAC algorithm, we propose an alternative indirect metric to discriminate solutions with small MSE: the total residual $\epsilon = \|\mathbf{A}\hat{x} - y\|_2$. By replacing the number of inliers by the total residual metric, we create the first step of the Total residual RANdom SAmple Consensus (TRANSAC) algorithm. The second step consists in a *voting* stage. Both are described in Algorithm 2 in pseudocode.

The total residual is directly proportional to the MSE of reconstruction. Unlike the number of inliers, this proportionality is conserved also when critical measurements are present in the dataset (Fig. 4c and d). In a real-life problem, where we do not have access to the ground truth, we do not know if critical measurements are present. Hence, we need a robust algorithm like TRANSAC. In addition, TRANSAC does not depend on the threshold η.

The proportionality between the total residual and the reconstruction error is not perfect, as we can see in the scatter plot of Fig. 4d. Even if a candidate solution has the smallest total residual, it is not guaranteed to be the solution with the smallest MSE. The intention of the *voting* stage is, using the candidate solutions with a total residual under a certain threshold, to come up with the best possible final solution.

Intuitively, the solutions with the smallest total residual (i.e., smallest MSE) are generated using almost outlier-free random subsets of measurements y'. We refer to these as the *good* subsets. Outliers can appear sporadicly in some of these good subsets, but the same outlier is extremely unlikely to appear in all of them. Hence, in the voting stage we select the measurements that all the good subsets have in common, or in other words, exclude any measurements that appear very infrequently.

Thus, we first select the subsets y' associated with candidate solutions with a total residual smaller than a certain threshold, $\epsilon < \beta$. Then, for each measurement we count how many times it appears in these good subsets. Finally, we select the M measurements with the largest frequency of occurrence.

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

|◄ | ►|

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

4 Results

4.1 TRANSAC

We now perform two experiments to demonstrate various aspects of TRANSAC.

4.1.1 Sanity check

In Sect. 3.3 we confirmed the expected behaviour of the first stage of TRANSAC: we showed that the total residual is directly proportional to the MSE. Let us check now the second stage, the voting. To do so, let us suppose that during the voting we have access to the MSE of every candidate solution $\hat{x}$. Then, we would of course select the solutions which in fact have the smallest MSE, and use them to build the histogram. We run this modified TRANSAC with the dataset without critical measurements. Figure 7a shows the MSE obtained for different values of the parameter M.

The dashed line on the right indicates the maximum possible value of M, such that $M = m$, which corresponds to using the whole measurement dataset. The dashed line on the left indicates the minimum possible value, $M = n$, and corresponds to using as many measurements as unknowns. The red horizontal line indicates the MSE of the solution obtained by using just the Tikhonov regularization without TRANSAC, i.e., when $M = m$.

We can observe that the MSE of the solution increases as M increases. This is to be expected: as M grows, more outliers are included in the dataset that is used to obtain x^*, and its MSE increases. We note that the results curve is non-decreasing, because in this particular experiment we have access to the MSE and the histogram h is built from the actual best candidate solutions.

4.1.2 Actual ETEX

In this subsection, the performance of the complete TRANSAC algorithm is examined. Let us consider first the dataset without critical measurements. As in the sanity check

above, TRANSAC is run for different values of M. The results are shown in Fig. 7b. We observe that the MSE increases as M increases, as before, and the maximum MSE still occurs at $M = m$. This is reassuring: even if we do not find the optimal value for the parameter M, we will improve the solution (with respect to using only the Tikhonov regularization) by taking any $n < M < m$. Notice that the minimum MSE again occurs when $M = n$.

Let us examine now the whole dataset, including the critical measurements. Figure 7c shows the results. We can observe that, again, the maximum MSE occurs at $M = m$. On the other hand, the minimum MSE does not occur at n, but rather at $M = 330$.

These results show that TRANSAC clearly improves the performance of the Tikhonov regularization in both cases: with and without critical measurements.

4.2 Outlier removal

As explained in Sect. 3, RANSAC and TRANSAC are blind outlier detection algorithms that can be combined with different regularizations in order to improve their results. In this section we combine RANSAC and TRANSAC with two different regularizations previously used in the literature, Eqs. (2) and (3), and study their performance. As before, we use the ETEX dataset with and without the critical measurements. The results are shown in Fig. 8. It is important to note that all these results were generated using the optimal values of all the parameters (λ, η, β, M) that were found experimentally. Figure 9 gives a more qualitative assessment of the results.

First, with and without critical measurements, the outlier removal stage improves the performance of both regularizations. Hence, our idea of removing outliers, outlined in Sect. 2, indeed does lead to improved performance, regardless of critical measurements or type of regularization. Next, in all cases TRANSAC shows higher performance than RANSAC. Therefore, our proposed modification of the metric, and the addition of the voting stage result in improved performance, as expected.

5 Conclusions

In this work we showed that the additive errors present in the ETEX dataset come from a heavy-tailed distribution. This implies the presence of outliers. Existing source estimation algorithms typically assume Gaussian additive errors. This assumption makes such existing algorithms highly sensitive to outliers. We showed that, if the outliers are removed from the dataset, the estimation given by these algorithms improves substantially.

However, in a real life problem, we do not know which of the measurements are outliers. Hence, we do have to remove them in a blind fashion. For this purpose we proposed RANSAC, a well-known blind outlier detection algorithm. We then showed that RANSAC unfortunately strongly depends on the chosen tolerance parameter, and it is sensitive to critical measurements. To overcome these difficulties, we created TRANSAC, a modification of RANSAC which also includes a voting stage.

To demonstrate the efficiency of these methods in a real-world problem, we used the ETEX tracer experiment dataset. The source was recovered first with two previously proposed source estimation algorithms that assume Gaussian additive errors (Eqs. 2 and 3). Then it was recovered again with our algorithms that use RANSAC and TRANSAC. The results clearly display how the source estimation improves if an outlier detection algorithm is used. They also show that the performance of our proposed algorithm TRANSAC clearly exceeds the performance of RANSAC in every case.

References

Bocquet, M.: High-resolution reconstruction of a tracer dispersion event: application to ETEX, Q. J. Roy. Meteorol. Soc., 133, 1013–1026, 2007. 3196

Cook, R. D.: Detection of influential observation in linear regression, Technometrics, 19, 15–18, 1977. 3200

CTBTO: Preparatory Commision for the Comprehensive Nuclear-Test-Ban Treaty Organization, available at: www.ctbto.org (last access: 2 May 2014), 2014. 3194

Fischler, M. A. and Bolles, R. C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography, Commun. ACM, 24, 381–395, 1981. 3197

Holmes, N. and Morawska, L.: A review of dispersion modelling and its application to the dispersion of particles: an overview of different dispersion models available, Atmos. Environ., 40, 5902–5928, 2006. 3195

Martinez-Camara, M., Dokmanic, I., Ranieri, J., Scheibler, R., Vetterli, M., and Stohl, A.: The Fukushima Inverse Problem, 38th International Conference on Acoustics, Speech, and Signal Processing, 2013. 3196

Nodop, K., Connolly, R., and Giraldi, F.: The field campaigns of the European Tracer Experiment (ETEX): overview and results, Atmos. Environ., 32, 4095–4108, 1998. 3196

Seibert, P.: Inverse Modelling with a Lagrangian Particle Disperion Model: Application to Point Releases Over Limited Time Intervals, Springer US, 381–389, doi:10.1007/0-306-47460-3_38, 2001. 3196

Stohl, A., Hittenberger, M., and Wotawa, G.: Validation of the Lagrangian particle dispersion model FLEXPART against large-scale tracer experiment data, Atmos. Environ., 32, 4245–4264, 1998. 3195

Stohl, A., Forster, C., Frank, A., Seibert, P., and Wotawa, G.: Technical note: The Lagrangian particle dispersion model FLEXPART version 6.2, Atmos. Chem. Phys., 5, 2461–2474, doi:10.5194/acp-5-2461-2005, 2005. 3195

Stohl, A., Seibert, P., Wotawa, G., Arnold, D., Burkhart, J. F., Eckhardt, S., Tapia, C., Vargas, A., and Yasunari, T. J.: Xenon-133 and caesium-137 releases into the atmosphere from the Fukushima Dai-ichi nuclear power plant: determination of the source term, atmospheric dispersion, and deposition, Atmos. Chem. Phys., 12, 2313–2343, doi:10.5194/acp-12-2313-2012, 2012. 3196

Title Page

Abstract | Introduction

Conclusions | References

Tables | Figures

◄ | ►►

◄ | ►

Back | Close

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Winiarek, V., Bocquet, M., and Mathieu, O. S. A.: Estimations of errors in the inverse modeling of accidental release of atmospheric pollutant: application to the reconstruction of the cesium-137 and iodine-131 source terms from the Fukshima Dai-ichi power plant, J. Geophys. Res., 117, D05122, doi:10.1029/2011JD016932, 2012. 3196

5 Zannetti, P.: Lagrangian Dispersion Models, in: Air Pollution Modeling, Springer US, 185–222, doi:10.1007/978-1-4757-4465-1_8, 1990. 3195

A robust method for inverse transport modelling

M. Martinez-Camara et al.

INPUT: $\mathbf{y} \in \mathbb{R}^m$, $\mathbf{A} \in \mathbb{R}^{m \times n}$, λ, η, N, p
Require: $\lambda \geq 0, N > 0, \eta \geq 0, p \leq m$
  $\mathcal{L}^* \leftarrow \emptyset$
  $\mathbf{x}^* \leftarrow \mathbf{0} \in \mathbb{R}^n$
  $\mathbf{r} \leftarrow \mathbf{0} \in \mathbb{R}^m$
  $\mathbf{k} \leftarrow \mathbf{0} \in \mathbb{N}^p$
  $\mathbf{y}' \leftarrow \mathbf{0} \in \mathbb{R}^p$
  $\mathbf{A}' \leftarrow \mathbf{0} \in \mathbb{R}^{p \times n}$
  for $s = 1$ **to** N **do**
    $\mathbf{k} \leftarrow p$ unique random integers from $[1, m]$
    $\mathbf{y}' \leftarrow \mathbf{y}[\mathbf{k}]$
    $\mathbf{A}' \leftarrow \mathbf{A}[\mathbf{k}, :]$
    $\hat{\mathbf{x}} \leftarrow \underset{\mathbf{x}}{\arg\min} \|\mathbf{A}'\mathbf{x} - \mathbf{y}'\|_2^2 + \lambda\|\mathbf{x}\|_2^2$
    $\mathbf{r} \leftarrow \mathbf{A}\hat{\mathbf{x}} - \mathbf{y}$
    $\mathcal{L} \leftarrow \{q \in \{1, 2, \cdots, m\} \mid \eta \geq (\mathbf{r}[q])^2\}$
    if $\#\mathcal{L} > \#\mathcal{L}^*$ **then**
      $\mathcal{L}^* \leftarrow \mathcal{L}$
      $\mathbf{x}^* \leftarrow \hat{\mathbf{x}}$
    end if
  end for
  return $\mathbf{x}^*$

Algorithm 1. RANSAC

INPUT: $\mathbf{y} \in \mathbb{R}^m$, $\mathbf{A} \in \mathbb{R}^{m \times n}$, λ, N, p, M, β

Require: $\lambda \geq 0, N > 0, p \leq m, n \leq M \leq m, \beta \geq 0$

 $\boldsymbol{\epsilon} \leftarrow \mathbf{0} \in \mathbb{R}^N$

 $\mathbf{k} \leftarrow \mathbf{0} \in \mathbb{N}^p$

 $\mathbf{K} \leftarrow \mathbf{0} \in \mathbb{N}^{p \times N}$

 $\mathbf{y}' \leftarrow \mathbf{0} \in \mathbb{R}^p$

 $\mathbf{A}' \leftarrow \mathbf{0} \in \mathbb{R}^{p \times n}$

 $\mathcal{G} \leftarrow \emptyset$

 $\mathbf{h} \leftarrow \mathbf{0} \in \mathbb{R}^m$

 $\mathbf{b} \leftarrow \mathbf{0} \in \mathbb{R}^M$

 for $s = 1$ **to** N **do**

  $\mathbf{k} \leftarrow p$ unique random integers from $[1, m]$

  $\mathbf{y}' \leftarrow \mathbf{y}[\mathbf{k}]$

  $\mathbf{A}' \leftarrow \mathbf{A}[\mathbf{k}, :]$

  $\hat{\mathbf{x}} \leftarrow \arg\min_{\mathbf{x}} \|\mathbf{A}'\mathbf{x} - \mathbf{y}'\|_2^2 + \lambda\|\mathbf{x}\|_2^2$

  $\boldsymbol{\epsilon}[s] \leftarrow \|\mathbf{A}\hat{\mathbf{x}} - \mathbf{y}\|_2$

  $\mathbf{K}[\,:, s] \leftarrow \mathbf{k}$

 end for

 $\mathcal{G} \leftarrow \{q \in \{1, 2, \cdots, N\} \mid \boldsymbol{\epsilon}[q] \leq \beta\}$

 $\mathbf{K}_{\mathcal{G}} \leftarrow \mathbf{K}[:, \mathcal{G}]$

 $\mathbf{h}[k] \leftarrow$ how many times k appears in $\mathbf{K}_{\mathcal{G}}$, $\forall k \in \{1, 2, \cdots, m\}$

 $\mathbf{b} \leftarrow$ indices of the M largest elements of $\mathbf{h}$

 $\mathbf{y}^* \leftarrow \mathbf{y}[\mathbf{b}]$

 $\mathbf{A}^* \leftarrow \mathbf{A}[\mathbf{b}, :]$

 $\mathbf{x}^* \leftarrow \arg\min_{\mathbf{x}} \|\mathbf{A}^*\mathbf{x} - \mathbf{y}^*\|_2^2 + \lambda\|\mathbf{x}\|_2^2$

 return $\mathbf{x}^*$

Algorithm 2. TRANSAC
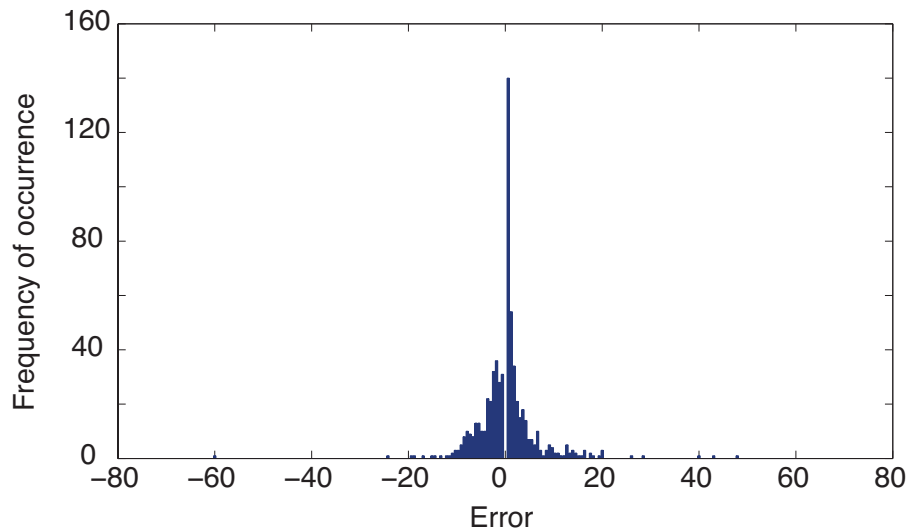
Figure 1. Histogram of the additive error e. For clarity, the zero-error bin has been omitted here.

A robust method for inverse transport modelling

M. Martinez-Camara et al.
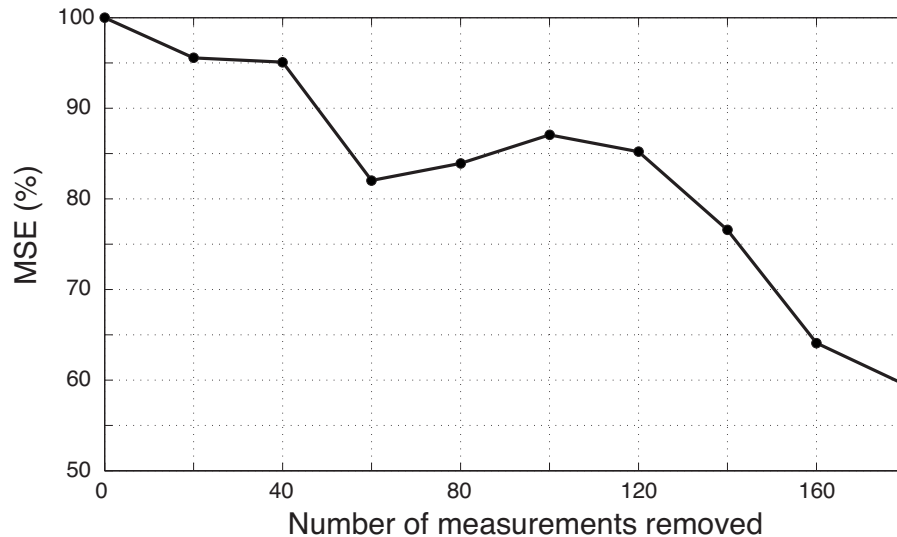
Figure 2. MSE of reconstruction obtained using Eq. (2). The strongest outlier measurements (the ones associated with the largest errors) have been removed manually. Notice that the MSE decreases as more outliers are removed.
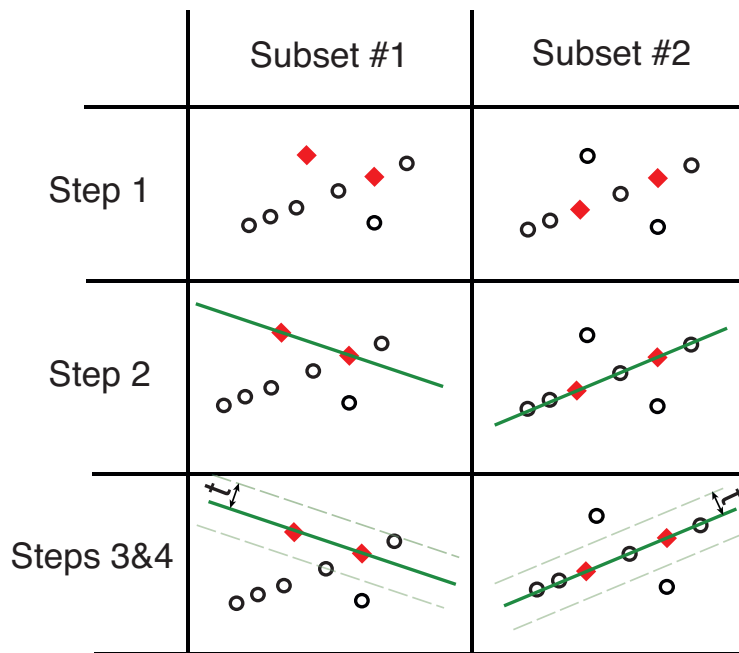
Figure 3. Visual representation of the functioning of RANSAC. Subset 1 and 2 represent two RANSAC iterations. The subset of measurements selected by RANSAC in each iteration is represented with red diamonds. Subset 1 contains one outlier. Hence, the solution corresponding with this subset generates fewer inliers than subset 2, which is free of outliers.
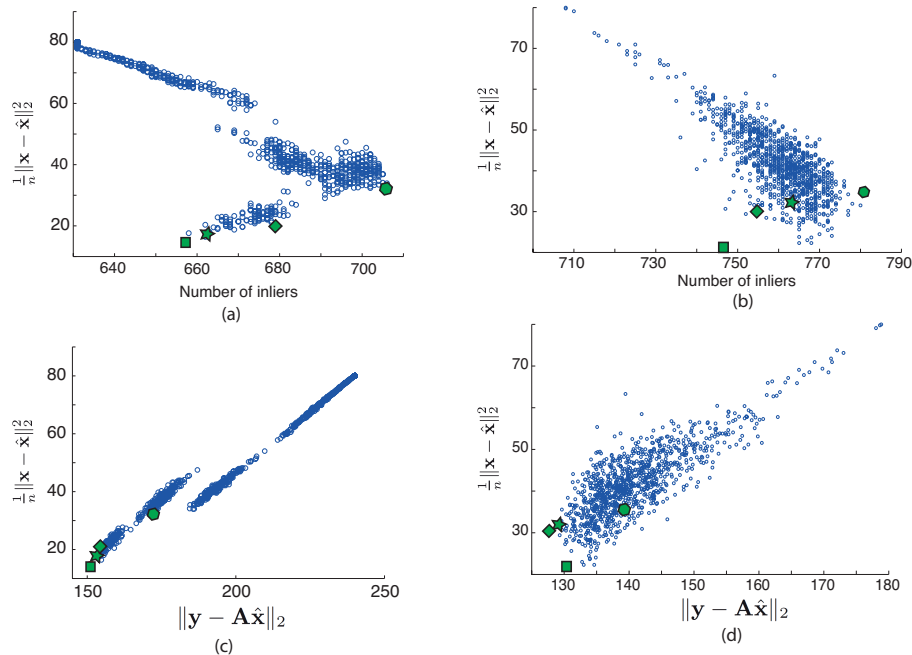
Figure 4. Performance of RANSAC and TRANSAC. **(a)** and **(b)** show graphically the correlation between MSE of reconstruction and the number of inliers. **(c)** and **(d)** show graphically the correlation between MSE of reconstruction and the total residual. To build **(a)** and **(c)** the complete dataset was used, to build **(b)** and **(d)** the dataset without critical measurements was used. The diamond indicates the solution obtained by the traditional Tikhonov regularization in Eq. (2), the star indicates the solution chosen by TRANSAC *before* the voting stage, the square indicates the final solution of TRANSAC, and the hexagon the solution chosen by RANSAC.
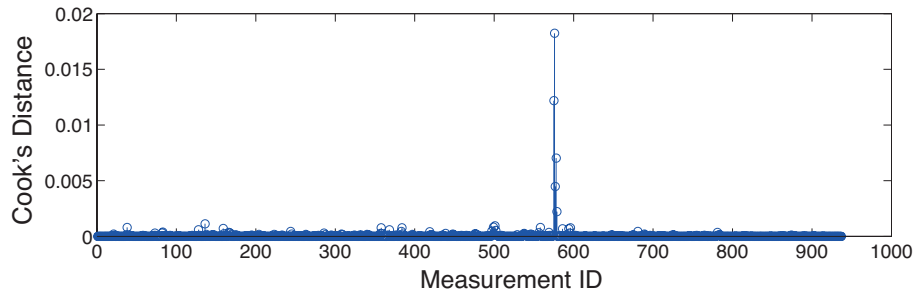
Figure 5. Cook's distance of the measurements in the ETEX dataset.

Figure 6. Residuals of two different source estimations: the blue peaks correspond to the residual produced by the solution $\hat{x}$ with the largest number of inliers in Fig. 4a. The black arrows on the top indicate where the two most critical measurements are localised. Clearly, the residual corresponding to these two measurements is much larger than the rest. The red peaks corresponds to the residual produced by the solution $\hat{x}$ with the smallest MSE in Fig. 4a.
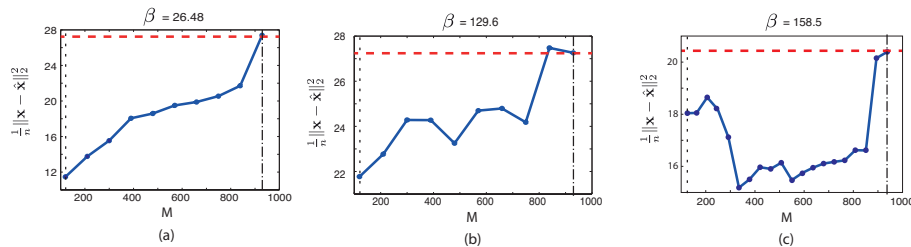
Figure 7. Performance of TRANSAC combined with Tikhonov regularization. In the three plots, the red dashed line indicates the estimation error given by typical Tikhonov (Eq. 2). The dashed line on the right indicates $M = m$, the one on the left indicates $M = n$. **(a)** shows the results of the sanity check. As the selected number of measurements M increases, the MSE of the estimation increases. Notice that the maximum MSE corresponds with $M = m$. **(b)** shows the results of applying TRANSAC to the ETEX dataset without critical measurements. Again, the MSE increases in general with M, and the maximum MSE appears in $M = m$. **(c)** shows the results of applying TRANSAC to the whole ETEX dataset, critical measurements included. In this case the MSE not always increases with M, but the maximum MSE still corresponds with $M = m$.
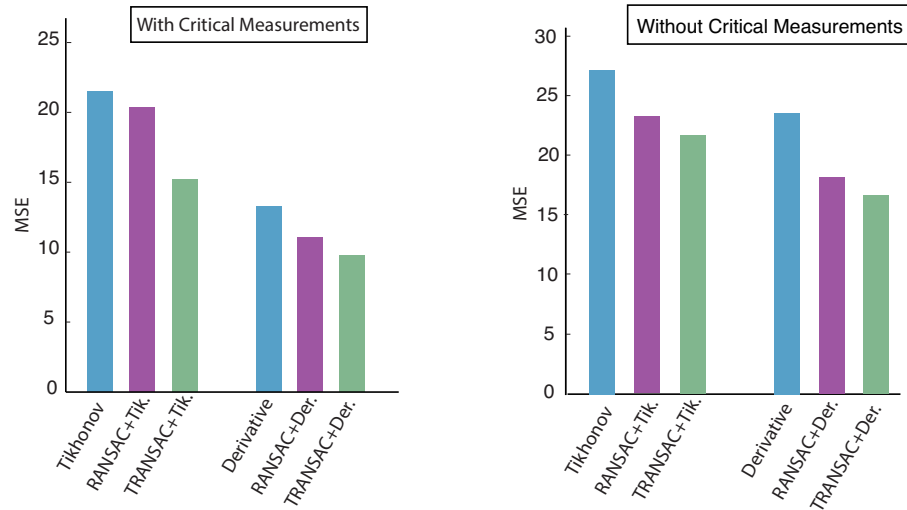
Figure 8. MSE of source estimated by different algorithms. The blue bars correspond with the original algorithms (Eqs. 2 and 3). The violet bars indicates that RANSAC is used for outlier removal, and the green ones that TRANSAC is used for outlier removal. The plot on the left was generated using the whole ETEX dataset. The plot on the right was generated using the ETEX dataset without critical measurements.
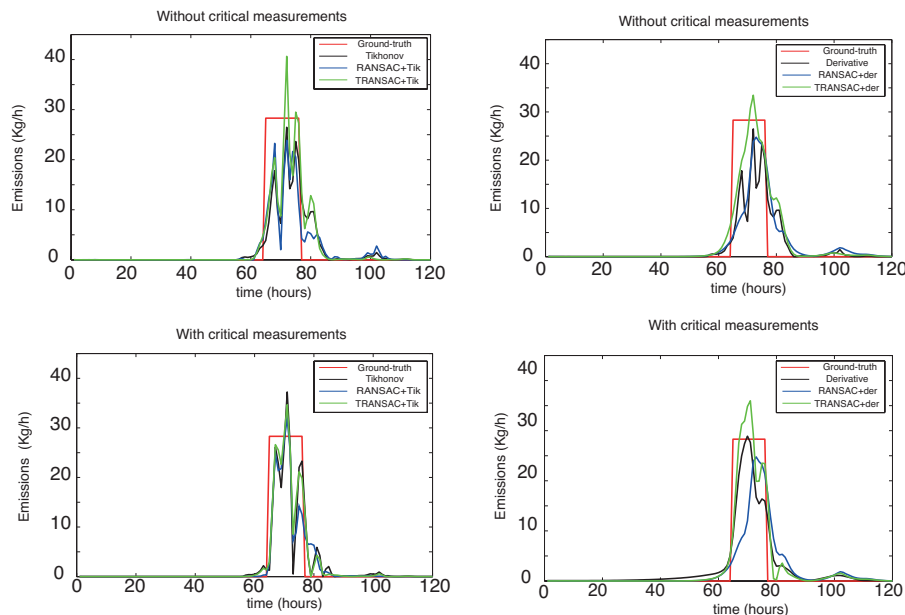
Figure 9. Source reconstructions given by the different algorithms. The plots on the left were generated combining Eq. (2) with RANSAC and TRANSAC. The plots on the right were generated combining Eq. (3) with TRANSAC and RANSAC. The plots on the top were generated using the ETEX dataset without critical measurements. The plots on the bottom were generated using the whole ETEX dataset.