



Conference paper

Oscar Björnham*, Håkan Grahn and Niklas Brännström

Reconstructing chemical plumes from stand-off detection data of airborne chemicals using atmospheric dispersion models and data fusion

<https://doi.org/10.1515/pac-2018-0101>

Abstract: Stand-off detection of airborne chemical compounds has proven to be a useful method that is gaining popularity following technical progress. There are obvious advantages compared to in situ measurements when it comes to the security aspect and the ability to measure at locations otherwise hard to reach. However, an inherent limitation in many of the stand-off detection techniques lies in the fact that the measured signal from a chemical depends nonlinearly on the distance to the detector. Furthermore, the measured signal describes the summation of the responses from all chemicals spatially distributed in the line of sight of the instrument. In other words, the three dimensional extension of the chemical plume is converted into a two-dimensional image. Not only is important geometric information per se lost in this process, but the measured signal strength itself depends on the unknown plume distribution which implies that the interpretation of the observation data suffers from significant uncertainty. In this paper we investigate different and novel approaches to reconstruct the original three-dimensional distribution and concentration of the plume by implementation of atmospheric dispersion models and numerical retrieval methods. In particular our method does not require a priori assumptions on the three-dimensional distribution of the plume. We also strongly advocate the use of proper constraints to avoid unphysical solutions being derived (or post-process ‘adjustments’ to correct unphysical solutions). By applying such a reconstruction method, both improved and additional information is obtained from the original observation data, providing important intelligence to the analysts and decision makers.

Keywords: atmospheric dispersion modeling; Chemical Weapons Convention 2017; hazardous substances; source term estimation; stand-off detection.

1 Introduction

For a week in April 2010 the ash emitted from the volcanic eruption at Eyjafjallajökull, Iceland, caused major disruptions in the European air traffic. There were in total 108.000 canceled flights affecting 10.5 million passengers and causing a \$1.7 billion loss of revenue for the airline industry alone [1]. Although not as severe, the following year in May, an eruption at Grimsvötn, Iceland, also led to flight cancelations. On the topic of Icelandic volcanoes, in September 2014 a gas fissure eruption of sulfuric compounds at Bardarbunga/Holuhraun triggered reports of foul smell in Norway, Sweden and Finland [2, 3]. While this release of sulfuric gases did not result in any casualties, the situation was different in Al-Mishraq, Iraq, in 2016. Here Daesh (ISIS) set fire to sulfur stockpiles, copying Saddam Hussein’s modus operandii from 2003 [4], as they used

Article note: A special issue containing invited papers on Innovative Technologies for Chemical Security, based on work done within the framework of the Chemical Weapons Convention.

*Corresponding author: Oscar Björnham, Swedish Defence Research Agency, Stockholm, Sweden,
e-mail: oscar.bjornham@foi.se

Håkan Grahn and Niklas Brännström: Swedish Defence Research Agency, Stockholm, Sweden

scorched earth tactics while retreating in the battle of Mosul. The sulfur dioxide released in the combustion caused ~1000 persons to seek medical attention in the cities south of Al-Mishraq and reportedly killed two Iraqis [5–7]. To make matters worse the nearby Al Qayyarah Oil Field was also set ablaze spreading plumes of thick black toxic smoke over the area [8, 9]. A more sinister use of chemical warfare was the sarin attack in Ghouta, Syria, 21 August 2013 resulting in a large number of deaths [10]. Sarin was again used in the attack on Khan Shaykhun, Syria, on April 4 2017 killing or seriously injuring at least 300 inhabitants [11, 12]. OPCW concluded that chlorine was likely used against the civilian population in Saraqib, Syria, on 4 of February 2018 [13]. Furthermore, the allegedly attack on Douma, Syria, on April 7 2018 involved dispersion of highly toxic chemicals, killing 40 people according to reports conveyed by the World Health Organization (WHO) [14].

As the examples above show, deliberate or undeliberate, injection of gases or particles into the atmosphere may cause downwind hazards in terms of pollution, damage, discomfort, irrevocable injuries and in worst case death. Fortunately, many of these risks may be mitigated or reduced by taking precautionary safety measures such as donning gas masks, evacuating people or canceling flights as appropriate to the situation. Now, wearing gas masks or canceling flights when no danger is imminent carries a high cost too, hence there is a need for the ability to predict where these hazardous substances are transported following a release, and to be able to assess whether they still constitute a hazard. Predicting downwind concentrations and arrival times are done using atmospheric dispersion models, and translating these downwind concentrations to potential hazard is done using toxicological models. The dispersion model will require three categories of inputs: weather, topography, and a source. This implies that even though it is actually the plume that is of interest, we need to find these three inputs. When this is accomplished we are able to predict the plume. Now, the weather and the topography are most likely known with an acceptable certainty. The source, on the other hand, is often considerably more challenging to estimate. Properties such as the substance, location, time-dependent release rate, and initial plume dynamics need to be quantified which requires information provided by measurements or a priori information.

Measuring the flow and injection height of a volcano, or accessing conflict zones like Al-Mishraq and Khan Shaykhun are either difficult, dangerous or both. In lieu of measurements at the source and downwind point sensors, data in the form of stand-off detections may be used to inform us on the characteristics of the plume and, indirectly, the source. Stand-off detections may come from satellites (e.g. MetOp-A/B, Suomi, Aura, the French-German Space Lidar Mission MERLIN [15] and soon Sentinel-5) carrying equipment to measure certain gases (such as O_3 , SO_2 , NO_2 , BrO, HCHO), from airplanes or helicopters carrying similar equipment, e.g. the DLR Falcon 20 or the Aerospace Corporation [16], or ground based stand-off detectors, see e.g. [17]. An illustration of this concept is shown in Fig. 1. The drawback with many of these stand-off techniques is that the measurements are presented as column densities, i.e. the sensor will integrate the response from all mass of the substance along its line of sight and return this value. In this paper we use the expression ‘area concentration’ to refer to the mass per area following such an integration of a concentration field along the line of sight, while the term ‘concentration’ refers to mass per volume in the traditional sense. To make matters worse, these measurements are sensitive to the spatial distribution of the substance as nearby substance will in general generate a stronger response than substance in the distant. Are stand-off measurements of area concentration sufficient to directly assess the downwind hazard? The answer is no for several reasons. First, without further information on the distribution of the cloud, the stand-off detections will come as area concentrations representing snap shots of highly uncertain line integrated mass. Secondly, we are interested in predicting the exposure for individuals in the region, and these individuals are not experiencing the average concentration along the line of sight since they are located at discrete points in the non-uniform concentration field. Thirdly, many stand-off detection systems suffer from inadequate revisiting rate, i.e. the frequency of measurements of the same area. Take observations from the sky; typically the individual is located on the ground, hence to assess the risk that he or she is exposed we need to know the ground level concentration, and to accurately assess the dose (or toxic load) we need to know how the plume behaves until the plane or satellite comes back to make the next measurement. That is, we need concentration data with high temporal resolution at the location of interest. Once again, dispersion models constitute the required tool as we will show. It is noted that the first two issues may be addressed by using LIDAR and DIAL detectors




Fig. 1: A typical example when stand-off measurements provide column densities, i.e. area concentrations, of the toxic plume. The area concentrations may be considered an image of the plume produced by the sensor of the stand-off detector. We are interested in reconstructing the plume that gave rise to the sensor response. Reconstruction of the plume is an ill-defined operation which is inherent to the type of problem that we are studying. The objective when reconstructing the plume is to find the plume that best explains the combined sensor response.

which may provide the concentration along line of sight [18, 19] or triangulation of several detection devices which would provide the missing spatial information. However, this advanced setup is not often available and will not be discussed furthermore in this paper, which is focused on single stand-off detectors measuring area concentrations.

Many stand-off detectors provide us with an area concentration field, often in low spatial and temporal resolution. We want to increase the resolution and uncover the missing spatial dimension, in other words we want to use the low-resolution area concentration to reconstruct the high-resolution concentration fields of the toxic cloud. To do so, dispersion simulation is introduced and the source term is therefore required. Fusing sensor data to recover the best estimate of the source term is a maturing field of atmospheric dispersion modeling and is an inverse modeling problem. Most analyses found in the literature address the problem of estimating the source term based on in situ point measurements using either a probabilistic approach (Bayesian) or optimization techniques. There are results based on stand-off detections, in particular satellite observations of volcanic emissions, e.g. [20–22]. A detail that is often overlooked in these studies is however the sensitivity of the stand-off detections on the plume height/width and its location relative to the detector. In this paper we describe a method that utilizes a dispersion model to reduce this uncertainty. The method we present relies on the source location (or source locations) being known a priori. However, the method does not rely on any a priori assumptions on the properties of the plume, rather the dispersion simulation informs us on how to interpret the observation data to determine the duration of release, how the release rate varies with time, the vertical profile of the release, and ultimately the concentration field of the plume and how it varies over time.

To summarize, we present a data fusion method that improves on inadequate information regarding concentration fields. The method is fed with empirical area concentration data in low resolution with no information regarding the time-dependent release rate or the initial plume dynamics (e.g. buoyancy effects). As an output the method provides the lacking source information as well as a high resolution concentration field. The key problem is to retrieve the source term which is the main focus in this work. Given the source

term provided by the data fusion method, detailed analysis of the plume in the past and prediction of the plume in the future in an ongoing event may be obtained by a straightforward dispersion simulation.

2 Data fusion methods

2.1 Overview

The main idea described in this work is to fuse simulated concentration data with stand-off observations to acquire more information about the dispersing cloud than is provided by the observation data alone. The simulated data is provided as an output of an atmospheric dispersion model. The atmospheric dispersion model is based on the predominant physical phenomena that govern the dispersion of a chemical plume with the aim of emulating the reality. As mentioned above, the source term includes the following properties: substance, location, time-dependent release rate, and initial plume dynamics. The first two are known in many cases and for the remainder of this study it will be assumed that this is the case. Some key information is typically missing when studying a real case, in particular the time-dependent release rate, i.e. released mass per time unit, of the source and the coupled initial plume dynamics, i.e. the vertical distribution of the plume close to the source. However, since these properties scale in a well described manner with respect to the final concentration fields, they lend themselves to be tuned post-simulation to match the simulated plume with the observations of the real plume. This tuning constitutes the data fusion and is the core of the method presented here. While saving the details just for now, we establish that the main ingredients required to perform data fusion are:

1. Design a source term for simulations
2. Perform dispersion simulations
3. Formulate commensurable data sets
4. Solve for source term parameters

The reason why the third issue is raised is technical but important: to be able to conduct the data fusion the simulated data set and the measured data set have to be commensurable. To achieve this we have to bear in mind that the observation data can often be burdened with strong and unavoidable limitations. These limitations may typically lie in the temporal and spatial resolution of the measured data set. In contrast, the simulated data can be extracted in virtually any desired resolution in both the temporal and spatial domains which implies that it should be chosen to match the observation data. It is therefore the observation data that constitutes the constraints of the system and determines the construction of the commensurable data sets.

Let us now turn to the details, i.e. the purpose, options, implications, and possible problems related to each of the steps outlined above.

2.2 Source design

To be able to simulate the dispersion process at hand we need to design a suitable source term. In a typical synthetic dispersion simulation all properties of the source are defined as input to the simulation model. This is not possible in the type of cases studied here since the release rate and the vertical distribution of the released substance are unknown a priori. The problem with lacking source information is addressed by designing a source term that includes a set of parameters whose ‘true’ values will be obtained as the result of the data fusion process. This means that dispersion simulation will initially be conducted using guessed, or normalized, values for some of the parameters that may be tuned (altered) post-simulation. The tuning of these parameters is handled in the source parameter solving algorithm which is discussed later. The properties that need to be defined in the source design are: the known substance, the known source location, the

known time interval of the release, a ‘temporal distribution method’, and a ‘vertical distribution method’. In the context of this work, we assume that the first property, the substance, is known. The location of the source is indeed often well known. One exception may be forest wildfires where the source location and extension actually changes with time which is a special case we will not discuss in detail here. The time interval of an entire release event refers to the time points when the release started and ended. In an analysis of an ongoing event, the time interval ends at the present time or at a time that relates to the last available observation. As with the location, this property is also often known with fair certainty. Even so, this is actually not pivotal information when the data fusion method is applied since the release rate is obtained as a result for all time points, hence the time interval can ideally be inferred with good accuracy from the results. However, all knowledge that guides the source term parameter solver as constraints will speed up the algorithm, improve the results and decrease the uncertainties. It is therefore recommended to provide the method with this information if available. We are then left with two important and unknown properties: the temporal distribution method and the vertical distribution method.

2.2.1 Temporal distribution method

The term temporal distribution method refers to how the source is designed to manage the dynamics of the release rate, i.e. how the release rate changes over time. In reality the release rate is a continuous function with an autocorrelation over time. The autocorrelation may vary from case to case and over time but the source term becomes unrealistic as the autocorrelation approaches zero since that implies that the release rate may change freely from one second to the next. An estimate of a lower limit of the autocorrelation may be acquired from previous experience of similar events. There is no need to set an upper limit of the autocorrelation. There are two principally different methods that may be applied to the source design. The first method is to let the release rate vary continuously using a sum of chosen subfunctions, $f_i(t)$, weighted by a set of parameters, p_i , that together allow for the shape of the release rate function to mimic all realistic developments of the release. The release rate function, $R(t)$, therefore resembles a truncated Fourier series with a limited number, N , of terms, see Eq. (1) and Fig. 2.

$$R(t) = \sum_{i=1}^N p_i f_i(t) \quad (1)$$

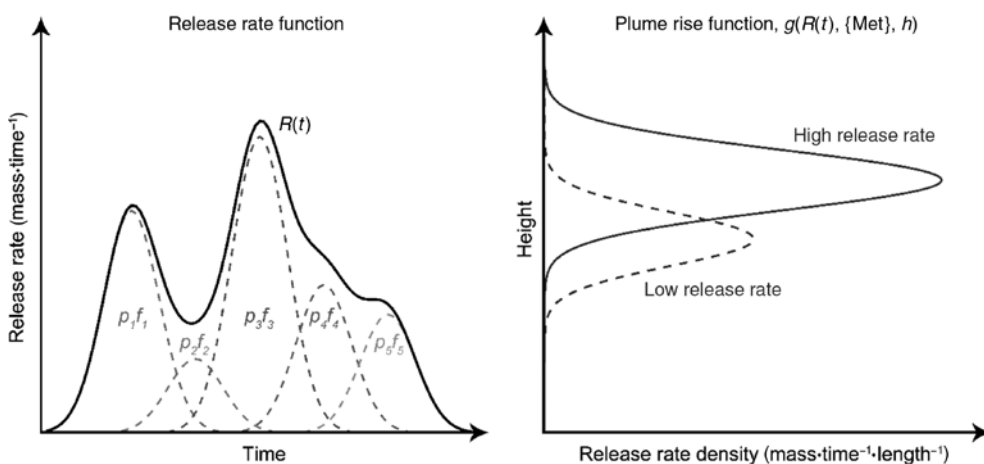


Fig. 2: Left, a small set of five subfunctions is shown as well as the release rate they add up to as equation 1 shows. Normalized Gaussian subfunctions have been chosen in this case. Right, illustration of the result from an arbitrary analytic plume rise function that provides the release rate density as a function of the current release rate, meteorology and height. The altitude of the initial plume is in general linked to the release rate, i.e. a high release rate implies a high altitude. The vertical integrals of the curves equal the release rates, i.e. the amount of released substance per time unit.

The second method applies discretization of the time into N intervals and assumes a constant release rate in each interval, i , but allows for variation between the intervals, see Eq. (2).

$$R(t) = \sum_{i=1}^N p_i \theta(t - t_{i-1}) \theta(t_i - t) \quad (2)$$

$\theta(t)$ is the Heaviside step function and t_i indicates the time point ending the time interval i (defining t_0 to equal the start time of the release). Equation (2) is merely a trivial stepwise constant release function. Even though this method obviously misrepresents the continuousness of nature, it may be argued that the present uncertainties renders this approach ‘good enough’. In both methods it is strongly recommended that the autocorrelation over time is taken into consideration. This may be implicitly incorporated in the case of the continuous approach by the choice of subfunctions since they may eliminate all abrupt changes in release rate. In the method with a discretized timeline, the autocorrelation may be applied by the use of a regularization term which is discussed later.

2.2.2 Vertical distribution method

As in the case with the temporal distribution of the release rate, the vertical distribution of the released substance is a result from the data fusion process but the method of representing the distribution must be defined in advance. It is well known that hot sources, such as large fires and volcanic eruptions, cause significant buoyancy which initiates a plume rise of the released chemical [4, 23–26]. This means that the main bulk of the substance will relatively quickly elevate to heights where the meteorological conditions may be completely different to those close to the ground. This phenomenon has been described as the ‘initial plume dynamics’ above. The issue of plume rise must be addressed in the simulation to obtain meaningful results. Depending on the capabilities of the dispersion model, plume rise may be incorporated as part of the dispersion process. However, should the dispersion model not include a plume rise functionality, the effect of the buoyancy could be incorporated directly in the source term by asserting a suitable vertical distribution of the released material [27, 28]. This approach has the benefit that the dispersion simulation will only disperse passive tracers and hence become linear, which is discussed further in the chapter Dispersion simulation. It is assumed, and strongly recommended, in this work that the dispersion model does not utilize buoyancy. That means that the initial vertical distribution of the plume is considered to be a property of the source which is a reasonable simplification since the plume rise mainly occurs close to the source itself. Similarly to the case with the temporal distribution method, there are two approaches to include the vertical distribution of the released material into the source design. We quantify the vertical distribution by the parameter w that is referred to as the ‘release rate density’, which means the released mass per time unit and per length unit in the vertical direction. The first method incorporates an analytic function g that provides the release rate density at the height h , given the current release rate R at the physical source.

$$w(t, h) = g(R(t), \{\text{Met}\}, h) \quad (3)$$

The analytic function usually incorporates a set of meteorological conditions, $\{\text{Met}\}$, that may change during the event but are always considered known, or deductible from the given meteorological input data. The standard analytic function for plume rise is the Briggs function [29, 30] and has been used frequently for many years, e.g. [31, 32]. The exact formulation of the continuous plume rise function is, though, of little interest here. The function can generically be formulated as in Eq. (3) and be visualized as the example in Fig. 2. Note that the actual plume rise distribution is in general not known in detail and that the plume rise function provides an analytic estimate based on theoretical assumptions and meteorological observations.

The second method includes a definition of a limited number, K , of discrete vertical layers that are indexed with k . Within each such vertical layer, the distribution of the released chemicals is assumed to be uniform and the release rate density is therefore constant with height. The amount of released chemicals at the different layers may vary individually over time representing a dynamically changing plume rise which

may occur as a result of variation in the total amount of released chemicals at the actual source and due to changing meteorological conditions [2, 22, 26, 33–35]. The parameter set h_k states how the current release is distributed between the layers. To maintain the same unit of the release rate density as above, the factor L_k^{-1} is introduced where L_k is the vertical height of layer k . The release rate density at layer k is then

$$w_k(t) = \frac{1}{L_k} h_k(t) R(t) \quad (4)$$

The set h_k is normalized and thereby does not alter the total release rate.

$$\sum_{k=1}^K h_k(t) = 1 \quad (5)$$

In case of a continuous temporal method, $h_k(t)$ needs either to be described by a plume rise function or to be formulated similar to the continuous release rate function, i.e. using a series of time dependent terms multiplied by parameters. In contrast, in case of a discrete temporal method, $h_k(t)$ may be formulated as a matrix H .

$$w_{i,k} = \frac{1}{L_k} H_{i,k} R_i \quad (6)$$

Here R_i is the release rate at the time interval i . The matrix $H_{i,k}$ now holds the parameters for the source where the index i increases with time and the index k increase with height in the fully discretized case. The sum of $H_{i,k}$ over the index k , for any i , is unity, similarly to Eq. 4. This case is illustrated in the bottom right in Fig. 3.

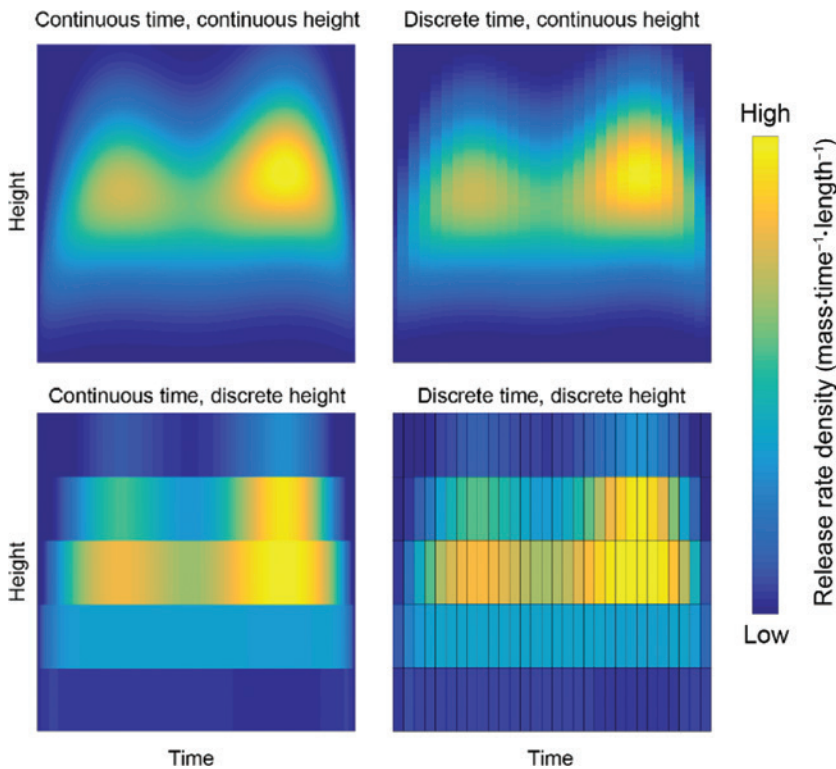


Fig. 3: The four different combinations of continuous and discretized time and heights. Yellow and blue colors indicate high and low release rate densities, respectively. Discretization in the vertical distribution is here made in five layers while the temporal discretization holds 30 different time intervals. These panels illustrate how the source term may be designed differently. The total masses in the fields are equal in all cases, i.e. the total release remains unchanged and independent on the choice of source term design.

The main principal difference between these two methods is the lack of direct dependence between total release rate and plume height in the discretized method, rather it is assumed that the dependence will emerge automatically from the source parameter solving algorithm. In other words, it is assumed that the actual plume rise will, due to the vertical variations in the meteorological conditions, influence the transport of the plume which will be observed. The source parameter solving algorithm will therefore force the vertical distribution in the source term to capture the correct initial plume rise for the simulated plume to mimic the observed data. To avoid physically unsound solutions, it is recommended that constraints and a regularization term are applied in the source parameter solving algorithm to smooth out too large variations in the relative release rate densities between the vertical layers of the source.

2.2.3 Combination of the methods

It might be worth to notice that the choices of temporal and vertical distribution methods show great resemblance. Either a continuous function or a discretization of the domains is applied. Even though all choices are valid, care must be taken in the choice of subfunctions, discretization, regularizations and constraints.

The four different combinations of continuous and discrete temporal and vertical distribution methods are illustrated in Fig. 3. The release rate densities in the discretized dimensions equal the mean of each line or area segment and the total amount of released material is equal in all cases. The four combinations of time and height representations imply different advantages and disadvantages that should be considered when designing the source. The top left panel presents a continuous release rate with a coupled continuous release rate density. The release rate and the vertical distribution, i.e. release rate density, thereof at each time point in this panel correspond to similar curves as shown in Fig. 2. This is by principal the closest representation of the actual real-life event since nature is inherently continuous. The drawback of this method is that the release rate subfunctions must be defined and also a plume rise function coupled to the release rate function must be specified. The temporal dimension has been discretized in the top right. Even though sudden jumps in release rate is a somewhat crude emulation of reality it provides an easier setup of the source. There is no need to define release rate subfunctions. However, this method will most likely require a regularization term in the source parameter solving algorithm. Instead of discretizing the temporal dimension, the vertical distribution is discretized in the bottom left panel. In this example, five vertical layers are used. Note that the temporal dimension is left in a continuous state. This means that release rate subfunctions must be defined including a plume rise function. This particular combination of representation methods in the source design can only be justified in case no trustworthy plume rise function is available. Finally, in the bottom right panel the fully discretized method is presented. This method deviates arguably most from reality which is a drawback and it also requires the incorporation of a regularization term. However, this choice of source design allows for a different approach in the source parameter solving algorithm. This new opportunity becomes available due to the fact that this method results in spatiotemporal area segments with constant release rate densities, i.e. each area segment in the panel in Fig. 3 has a constant color, which is a property that is useful and will be discussed later.

With the source designed, regardless of method, a set of parameters (p_i and possibly h_k or $H_{i,k}$) is defined. This set of parameters will be targeted in the source parameter solving algorithm where the best possible values of the parameters will be identified. Having determined the substance, the location of the source, the time interval of the release, a temporal distribution method, and a vertical distribution method, the source design is complete and dispersion simulation may be conducted.

2.3 Dispersion simulation

Atmospheric dispersion simulation is the process of simulating the fate of released substances as they propagate in the atmosphere by the conveying wind field and disperse by turbulence. There are many different

dispersion models that each handles meteorological information but also phenomena such as chemical degradation and dry and wet deposition differently. What they all have in common though is that they use the designed source term to deliver concentration fields at desired time points and at desired spatial resolutions. To simplify the nomenclature ahead we will assume that the dispersion model uses model particles (Lagrangian model). These model particles can represent both physical particles, aerosols and gases.

Here it is assumed that the dispersion process is linear which means that the evolution of the released chemicals is independent of the amount of released mass. Each model particle moves and evolves independently from the other model particles. Note that the linear assumption is not to be confused with the fact that the initial vertical distribution of the released mass do depend on the release rate. This relation lies within the source design and not the atmospheric dispersion model itself. The dependence of the vertical distribution upon the released mass will influence how much chemicals that are released at different heights, but not their movements in the air. This is a crucial property in the data fusion process since it allows for a direct recalculation of the concentration fields as a result of tuning the release rate and vertical distribution. That is, the position of each model particle at any given point in time will not change due to change in the parameters of the source.

Any set of source parameters give rise to a certain spatiotemporal release rate density as depicted in Fig. 3. The corresponding concentration fields can be found by releasing model particles that correspond to this release rate density. However, it is not practical to adjust the number of released model particles to the current release rate density since that would imply that the entire dispersion simulation must be rerun every time the source parameters change in the solving algorithm. Instead, each model particle is attributed a ‘weight’ that defines how much physical mass it represents. The temporal and vertical distributions of the released mass of the source can therefore be altered after the simulation has been executed in a post-process by changing only the weights of the model particles. Again, since the dispersion is linear, the weights do not influence the movement of the model particles. There is no need to rerun the dispersion simulation, the concentration fields can readily be updated by the new weights of the model particles. One master simulation is conducted using a large set of model particles that covers the entire spatiotemporal domain of the source. The weights are then calculated to acquire the desired release rate density, i.e. to represent the release rate density in Fig. 3. As a remark we note that this approach does not necessarily require a uniform distribution of model particles. The resolution and performance may be improved by initiating a higher density of model particles in regions where we expect the highest density in the two-dimensional domain. Of course the model particle distribution then needs to be accounted for in the calculation of the weights so that the release rate density remains correct.

2.4 Formulate the data sets

Measurement data of an area concentration field is normally provided as data points on a two-dimensional grid or as scattered data points. In the latter case, the data points may be regularly or irregularly distributed over the geometric region of interest, i.e. the data set provides pointwise area concentration values distributed arbitrarily over the domain. Scatter data can be interpolated to a two-dimensional grid using methods such as bi-linear or bi-cubic interpolation [36] or Kriging [37]. The use of a grid enables ubiquitous compatibility with input data. Furthermore the grid representation provides a pattern that may readily be utilized to quantify the agreement between observed and simulated data. Depending on the data set it is possible that the grid becomes disjoint or that there are sets of grid cells with unknown area concentration. However, with the method proposed in this work such situations constitute no serious complications, rather they introduce an increased uncertainty due to a lack of information. When using two-dimensional data sets obtained by spatial integration along the line of sight of the instrument (i.e. area concentration), the unit of the dataset becomes mass/area. As a standard, the Dobson unit scale is used for this physical quantity where one Dobson unit is defined as 2.687×10^{20} molecules per square meter. Commensurable area concentration fields can readily be rendered from simulated data by integrating the field column-wise as has been done by

the detector. This implies that whenever a simulation provides a concentration field, the corresponding area concentration field is always available as well.

For the data sets to be truly commensurable they must be spatially represented in an equal manner as described above. Additionally, the sets must also contain data corresponding to area concentrations valid at the same time points. That is, the simulated area concentration fields must be collected at the same point in time as the measurements were conducted. This is often a trivial matter, but in cases where the time scale of the data measurements is on the same, or higher, order as the time scale of the dynamics of the plume, this issue must be addressed. When this occurs, it is possible that the plume has changed position and distribution from the beginning to the end of the measurement of one single area concentration field. Since area concentration measurements typically are conducted by sweeping measurements of parts of the area, each such part is measured at different time points. The solution is therefore to collect the simulated area concentration field at diverse time points that correspond to the time points of the measurements, i.e. the different parts of the simulated area concentration field are not to be collected at the same time point, rather their time points should mimic the sweeping mechanism of the detector. However, when the plume dynamics are significantly slower than the measurement procedure, this detailed data collection is unnecessary.

2.4.1 Grid resolution

If the observed area concentration data is provided as scatter data points, a grid is constructed onto which the data is inserted. A two-dimensional (Cartesian) grid leaves few degrees of freedom. Any spatial translation of the grid will in general have little influence of the result and is ignored in this discussion. A more important setting is the grid resolution, i.e. the size of the grid cells. There are a couple of issues to consider while choosing the grid resolution. The plume must be resolved on the grid which sets a lower limit on the grid cell resolution. Though this limitation might be quite obvious, the other limit of the resolution scale might be less trivial to spot. In the source parameter solving algorithm proposed in this paper, the agreement between observed and simulated data is evaluated on the basis of comparison between individual grid cell area concentrations with no consideration to adjacent grid cells. This means that uncertainties in the prevailing winds may easily translate the simulated field one or more grid cells if the grid resolution is high. If this is the case, the source parameter solving algorithm will struggle to guide the solution in the right direction since there is a spatial mismatch in the field comparison. A high grid resolution may also require a large number of model particles depending on which method that is used to compile the simulated area concentration field. Common methods include straightforward box-counting [38, 39] or kernel density estimators [40, 41]. Box-counting is sensitive to the grid resolution since the model particles only contribute to their particular grid cell while the opposite applies to kernel density estimators. Any direct measure of an optimal grid resolution is not available but a recommendation is that the grid should resolve the main features of the plume and at the same time avoid to punish mismatch on a very detailed level.

2.5 Source parameter solving algorithm

2.5.1 Optimization algorithm

In this final step of the data fusion process the observed and simulated area concentration data sets are compared with the aim of identifying the optimal values of the source parameters. That is, the values of the source parameters that provide simulation fields that best reassemble the observed fields are found. This set of parameters will then define the source term, which, as input to the dispersion model, will provide the full spatio-temporal development of the complete release event. Note that the final concentration field in high spatiotemporal resolution is directly available by asserting particle weight according to the optimal source parameters as discussed in Section 2.3. The reconstruction of the observed plume will then be complete.

The mechanism to find the optimal solution is principally straightforward. Any set of parameters will provide area concentration fields for all time points where observed data is available. Observed and simulated data sets are conformed on a suitable grid, rendering them commensurable, and the degree of disagreement is quantified using a scalar statistical measure φ . Now, there are different options of statistical measures. A strong contender is to apply the sum of the mean square errors (MSE) of the fields over all M time points

$$\varphi \equiv \sum_{m=1}^M \text{MSE}(V_{\text{obs},m}, V_{\text{sim},m}) \quad (7)$$

where $V_{\text{obs},m}$ and $V_{\text{sim},m}$ are the observed and simulated area concentration fields, respectively, on the grid for the time point m . The statistical measure φ describes quantitatively the degree of mismatch between the simulated and observed fields and depends on the choice of source parameters. Since the aim is to match the simulated field as closely as possible to the observed field, the optimal solution is defined by the global minimum of φ . The statistical measure is actually a continuous scalar field of dimension equal to the number of unknown parameters. A high order dimensionality may render the localization of the global minimum to be numerically challenging. This may pose a problem especially for long lasting events and needs to be considered in the source design, i.e. try to avoid unnecessary parameters.

As a final note, we will discuss the timesaving trick referred to in the section Source design. When the source term parameters are changed it implies new weights of the model particles which translates into new area concentration fields. It may be time consuming to recalculate the area concentration fields over and over again as the search for the global minimum is under way. However, there is a shortcut in case of a source design utilizing discretization in both the temporal and vertical domains. In this specific setup there are many model particles with identical weights. One might capitalize on this fact by calculating one separate area concentration field for each area segment in the lower right panel of Fig. 3. When the weights are altered, each such area concentration field may be linearly adjusted by the change in the weights of the corresponding model particles. This means that there is no need to recalculate the area concentration field from the particle data. The total area concentration field is instead compiled by summing up all area concentration fields representing the area segments. This is a feature that may speed up the iterative calculations of the area concentration field significantly and thereby in many cases be the decisive factor in the choice of the source design.

2.5.2 Probabilistic interpretation

As an alternative to using an optimization algorithm we may view the source estimation problem as a probabilistic one and cast the problem as such, that is, instead of searching for an optimum set of source parameters we look to assign a probability to each combination of parameters. Guided by these probabilities the source parameters may be chosen. Making this choice could be done in several ways, e.g. by picking the maximum likelihood estimate or the expected value. By far the most common method of assigning probabilities to parameters in source estimation problems is using Bayes method [42]. Using Bayes theorem the ‘posterior’ distribution $P(\text{source} | \text{obs})$ giving the conditional probability of having source parameters ‘source’ conditional on having made observations *obs* is given by

$$P(\text{source} | \text{obs}) = \frac{P(\text{obs} | \text{source})P(\text{source})}{P(\text{obs})}$$

Bayes theorem thus relates the posterior distribution to the conditional probability, the ‘likelihood distribution’, $P(\text{obs} | \text{source})$ of making observation ‘obs’ given source parameters ‘source’ and the marginal probabilities $P(\text{source})$ of having these source parameters, the ‘prior distribution’, and the probability $P(\text{obs})$ of making the observation ‘obs’, the ‘evidence’. The likelihood distribution is a source-detector relationship and can be calculated using a dispersion model, and the prior distribution gives us the opportunity of weighing the posterior distribution with any external/prior information that we may have regarding the source, e.g. by

intelligence. The evidence distribution $P(\text{obs})$ only serves to normalize the distribution, and since we are mainly interested in relative probabilities this entity is normally not needed. As a remark, see [43] for details, we note that the dispersion problem that is studied in this paper, indeed in nearly all papers, only involves a linear relationship between the source and the sensor which allows us to write the likelihood distribution $P(\text{obs} | \text{source})$ on a form that involves the MSE cost function mentioned in the optimization algorithm above. Hence, for these problems, the main difference between the optimization algorithms and the probabilistic approach lies in how the problem is interpreted.

2.5.3 Regularizations and constraints

Non-detailed references to regularization terms and constraints have been given at several occasions above. A regularization term [44] can be viewed as a tunable weight on a cost function (the statistical measure in the present context). The idea is to be able to punish unwanted behavior and thereby to guide the source parameter solving algorithm heuristically towards the best realistic solution. For example, it could in the type of situations dealt with here be used to increase the statistical measure for parameter sets that are judged to represent too small autocorrelations in the release rate. Such a regularization will tilt the parameter space of the statistical measure, resulting in a shift of the global minimum towards a smoother release rate with regard to time. Regularization terms influence the scoring of parameter sets but does not exclude any subdomain of the parameter space. This is instead the purpose of constraints [45]. By including constraints we are limiting the domain of possible parameter values. This is mainly recommended to exclude unphysical solutions but can also be used to speed up the source parameter solving algorithm by leaving a smaller parameter space to examine in the search for the global minimum. An obvious example of an unphysical solution that can be excluded with the use of constraints are negative values of the source parameters. Such values would represent the annihilation of mass at the source, a phenomenon that is unphysical and thus unwanted.

3 Case study – the great sulfur fire at Al-Mishraq 2016

Please note that this case study is described in full detail in the article ‘The 2016 Al-Mishraq sulfur plant fire: Source and health risk area estimation’ [26]. In this section, we give a brief overview of the event, the source design and the method applied to estimate the source term. This is a special case of the more general description provided in this paper.

October 20 2016, retreating Daesh troops ignite the extensive sulfur stock piles of the Al-Mishraq plant in northern Iraq. The fire spreads rapidly and a release of sulfur dioxide at the size of a minor volcanic eruption is soon a reality. The Swedish Defence Research Agency launches an investigation into the matter by employing satellite data from optical spectrometers measuring backscattered solar irradiance in the ultraviolet part of the spectrum to monitor, among other substances, airborne SO_2 . As the satellites orbit the planet, they provide rapid measurements of SO_2 in Dobson units, see Fig. 4. That is, the measurement data is cast as area concentrations as described above, and illustrated in Fig. 1, with all its limitations. The measurements are strongly dependent on the vertical position of the plume which means that without further information it is difficult to draw any quantitative conclusions from the satellite data. The orbiting path of each satellite gives rise to measurement bands over the globe that leave gaps as is shown in Fig. 4. To overcome this limitation, satellites cooperate in pairs in such a way that the second satellite measures in the gaps of the first satellite. This complementary data is also closely positioned in time. The pair of satellites provides a satisfying coverage of the region once per day. To improve upon the low temporal frequency of observation data, two such pairs of satellites are employed in this study. This implies that there are two observation fields available per day and in total 12 observation fields over 6 days of dispersion.

It is clear from photographs by Al-Sudani/Reuters [46] at the location that the extensive heat, generated by the exothermal reactions, causes a significant plume rise. To be able to capture this property it is necessary

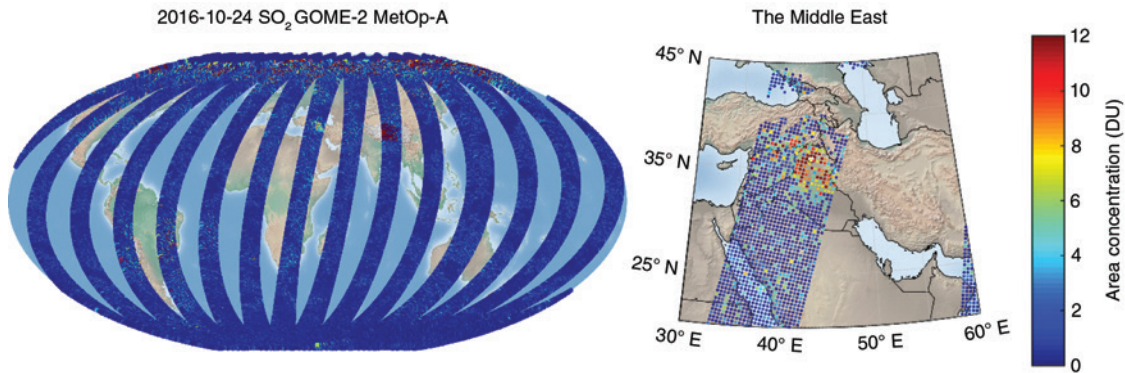


Fig. 4: Measurement data from one instrument (MetOp-A) for October 24. Left, each band on the world map represents one polar orbit and the entire earth is covered in 24 h by two satellites working in pair. The plume over the Middle East is visible in the global map but even more clearly in the zoomed view over the Middle East to the right. Some data points are removed due to cloud coverage. Note that these values need to be scaled according to the actual height of the plume.

to include heterogeneity to the vertical distribution of the source term. A fully discretized source term, as depicted in the lower right panel in Fig. 3, is defined using 6 h intervals covering in total 6 days and 6 h. The source term utilizes three vertical layers whose limits, 0–200 m, 200–1000 m and 1000–4000 m, are based on the meteorological conditions during the event. In total, this means that the source is divided into 75 different area segments, here referred to as ‘virtual sources’. Each virtual source has a constant release rate density of sulfur dioxide which implies that the corresponding released mass for the virtual source is uniformly distributed over the time interval and the ‘virtual source volume’. The term ‘virtual source volume’ is defined as the area of the sulfur plant multiplied with the vertical height of the virtual source. The source term is thereby defined by 75 parameters to be set by the optimization algorithm that is created as the source parameter solving algorithm. The discretized source term is illustrated in the right panel of Fig. 5 where the width and height of each area segment correspond to the discretization in time and vertical direction, respectively.

The optimization algorithm is built using the scalar statistical measure that is defined as the sum of the mean square error between the observed and simulated fields for all 12 time points, compare Eq. (7). This measure now depends on the 75 source parameters, i.e. the release rate densities for the virtual sources. Since

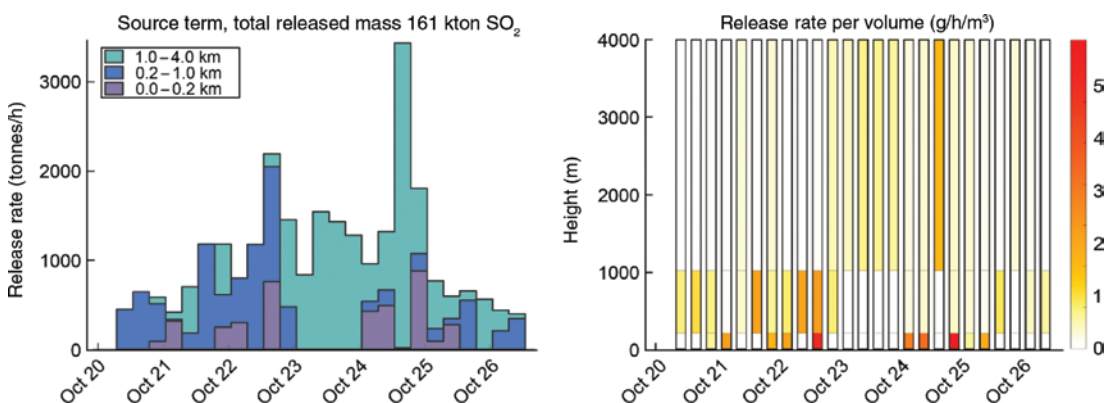


Fig. 5: The resulting source term presented in two different ways. Left, the release rate for the three different vertical layers. Right, this presentation uses the time intervals and the vertical heights of the virtual sources and indicates the release rate per volume. The two panels show the same information where the right panel includes information on the virtual source volumes. The right panel also functions as a depiction of the source design utilized in this case study. 25 time slots and three vertical layers result in 75 different virtual sources and an equal number of source parameters to be found. This figure is a reproduction from [26] with permission provided by the license type Creative Commons Attribution-NonCommercial-No Derivatives License (CC BY NC ND).

the source term is defined using full discretization, individual plumes for the entire scenario can be produced for each virtual source as discussed above. Alteration of the source parameters are directly implemented as linear scaling of the plumes for each virtual source which allows for a relatively effective source parameters solving algorithm. A complication arises as the observed fields depends on information of the plume height, which is unknown a priori. This plume height is therefore estimated from the simulated plume given the current source parameters. This implies that both the simulated and the observed area concentrations become dependent on the source term parameters. Constraints on the daily difference in total airborne mass between the observed and simulated plume are applied together with a requirement of non-negative release rates. Moreover, a regularization term is implemented to moderate any unrealistic fluctuations of the release rate that may originate as artifacts from the low frequency of observation data. The regularization term and other case specific details are available in Appendix A coupled to the paper on this case study [26]. To solve the system, a global minimum search algorithm [47] is implemented. The algorithm uses the derivative-free Nelder-Mead simplex algorithm [48]. The source parameters that give rise to the global minimum for the scalar statistical measure are identified and are illustrated in two different manners in Fig. 5.

4 Conclusions

A variety of stand-off measurement instruments and systems provide area concentration data that represents column-wise integration of the concentration field along the line of sight. This is a problematic limitation for two reasons. First, the data is lacking information regarding the spatial distribution of the substance measured in the direction of the observation. Secondly, the data signal is often dependent on this unknown spatial distribution and this property gives rise to a significant and troublesome uncertainty for the measured area concentration. A general description of a method to address this issue is presented in this paper. The method allows for different choices in the approach both with regard to the source description and the source parameter solving algorithm. There is, however, a distinct path to follow regardless of the specific choices made by the data analyst. Implementation of such a method will provide extensive additional information of the physical reality hidden in the measurement data.

Four combinations for the source term representation have been discussed. It can be argued that the completely continuous or completely discrete setups are to be recommended. The case with a completely continuous source term, i.e. continuous in both the temporal and the spatial dimensions, is intrinsically the closest representation to a physical event. As a drawback this approach requires a well-working analytic plume rise model and a carefully thought-out set of release rate subfunctions. As a contrast, the completely discretized source term is a somewhat crude representation of the reality. However, compensation is to be found in the fact that this type of source term is easy to set up, fast to execute and can arguably be a sufficiently sophisticated design given the influences of uncertainties and errors from other factors in the analysis. A general note is that it might be computational heavy to locate a global minimum in a parameter space of high dimension, which is worth keeping in mind when considering the design of the source term.

It is important to consider, and probably include, constraints and regularization terms to avoid unphysical solutions. These solutions may arise due to sparse measurement data and biases in the simulated distribution due to errors in the meteorological predictions and/or the dispersion model. A generic constraint is to disqualify all negative release rates. It is worth noticing that this limitation should be included as a constraint in the source parameter solving algorithm. It might be tempting to take a shortcut and instead set all negative values to zero as post-processing. However, there is no guarantee that the actual global minimum will be identified and preserved using this shortcut.

In this paper we have put the source design method on a mathematical footing discussing limitations and possibilities of discretizing vs. keeping aspects of the source continuous, which stand in contrast to much of the literature where the discretized setup tends to be the starting point. In contrast to the literature covering satellite based stand-off detections to estimate source terms, the method presented here requires no prior assumptions on thickness/width of the plume – a property that highly affects how to interpret the

measurements, instead this method iteratively determines these properties as part of the process of matching observed and simulated data. Finally we impose our constraints on the problem up-front instead of making post-processing ‘adjustments’. A key limitation of the method is that it subsumes the location of the source or the sources. Another limitation is that the method does not include a description of how to choose the best regularization term. Also, we have not performed a thorough study of under which conditions the method converges to the solution, but the Al-Mishraq case study suggests the method is fairly robust.

We note that atmospheric dispersion simulations target several shortcomings that are present using stand-off detection data. The highly resolved fields provided by the simulations help the analyst to interpret measured column data, increase the temporal resolution, adds an additional dimension to the concentration field and may be used in an ongoing event to forecast the future development of the plume and thereby assist authorities to establish health risk areas for the coming days.

References

- [1] L. Budd, S. Griggs, D. R. Howarth, S. G. Ison. *Mobilities* **6**, 31 (2011).
- [2] H. Grahn, P. von Schoenberg, N. Brännström. *J. Volcanol. Geoth. Res.* **303**, 187 (2015).
- [3] K. M. Hovland, J. Rossi, L. Hilmarsdóttir. Iceland Volcanic Eruption Sending Toxic Gases Throughout Region, in *The Wall Street Journal* 2014.
- [4] S. Carn, A. Krueger, N. Krotkov, M. Gray. *Geophys. Res. Lett.* **31**, (2004).
- [5] A. Jazeera. Sulphur cloud from torched plant kills two Iraqis, in *Al Jazeera*, 22 October 2016.
- [6] B. Dehghanpisheh. Burning sulfur near Mosul sends hundreds to hospital, U.S. troops don masks, in *Reuters*, 22 October 2016.
- [7] A. Sultany. WHO and Ministry of Health accelerate medical intervention to manage over 1000 cases of suffocation in Qayarra, Ijhala, Haj Ali, and Makhmour. 2016, WHO.
- [8] Reliefweb, Fires at the Al Qayyarah Oil Field, Nineveh Governorate, between 18 July 2016 and 7 January 2017, in *Reliefweb*, 12 January 2017.
- [9] Joint UNEP/OCHA Environment Unit. A rapid overview of Environmental and Health Risks Related to Chemical Hazards in the Mosul Humanitarian Response. 2016, J.U.O.E. Unit.
- [10] Å. Sellström, S. Cairns, M. Barbeschi. United Nations Mission to Investigate Allegations of the Use of Chemical Weapons in the Syrian Arab Republic 2013, United Nations.
- [11] UN-OPCW JIM, Seventh report of the Organisation for the Prohibition of Chemical Weapons – United Nations Joint Investigative Mechanism, United Nations Security Council.
- [12] OPCW Technical Secretariat, Report of the OPCW fact-finding mission in Syria regarding an alleged incident in Khan Shaykhun, Syrian Arab Republic April 2017. 2017, OPCW.
- [13] OPCW Technical Secretariat, Report of the OPCW fact-finding mission in Syria regarding an alleged incident in Saraqib, Syrian Arab Republic on 4 February 2018. 2018, OPCW.
- [14] WHO Media Centre. WHO concerned about suspected chemical attacks in Syria. 2018 [11 April 2018]; Available from: <http://www.who.int/mediacentre/news/statements/2018/chemical-attacks-syria/en/>.
- [15] G. Ehret, P. Bousquet, C. Pierangelo, M. Alpers, B. Millet, J. B. Abshire, H. Bovensmann, J. P. Burrows, F. Chevallier, P. Ciais, C. Crevoisier, A. Fix, P. Flamant, C. Frankenberg, F. Gibert, B. Heim, M. Heimann, S. Houweling, H. W. Hubberten, P. Jöckel, K. Law, A. Löw, J. Marshall, A. Agusti-Panareda, S. Payan, C. Prigent, P. Rairoux, T. Sachs, M. Scholze, M. Wirth. *Remote Sens.* **9**, 1052 (2017).
- [16] K. N. Buckland, E. R. Keim, J.L. Hall, T.L. Volquarts, K. R. Westberg, D.M. Tratt. Validating Urban-Scale Transport Models using Airborne LWIR HSI Sensing. 2017, The Aerospace Corporation, Los Angeles.
- [17] J. L. Hall, F. M. D’Amico, S. J. Kolodzey, J. Qian, M. L. Polak, K. Westerberg, C. S. Chang. Characterization of aerosol-containing chemical simulant clouds using a sensitive, thermal infrared imaging spectrometer. in *Chemical, Biological, Radiological, Nuclear, and Explosives (CBRNE) Sensing XII*. 2011.
- [18] G. G. Gimmetstad. *Differential-absorption Lidar for Ozone and Industrial Emissions*, in *Lidar*. Springer (2005). pp. 187–212.
- [19] P. Gaudio, M. Gelfusa, A. Malizia, S. Parracino, M. Richetta, L. De Leo, C. Perrimezzi, C. Bellecci. in *Journal of Physics: Conference Series*. 2015. IOP Publishing.
- [20] S. Eckhardt, A. J. Prata, P. Seibert, K. Stebel, A. Stoh. *Atmos. Chem. Phys.* **8**, 3881 (2008).
- [21] P. Seibert. *Inverse modelling of sulfur emissions in Europe based on trajectories*, *Inverse Methods*, in *Inverse Methods in Global Biogeochemical Cycles*. 2013, American Geophysical Union, Washington, DC. pp. 147–154.
- [22] A. Stohl, A. J. Prata, S. Eckhardt, L. Clarisse, A. Durant, S. Henne, N. I. Kristiansen, A. Minikin, U. Schumann, P. Seibert, K. Stebel, H. E. Thomas, T. Thorsteinsson, K. Tørseth, B. Weinzierl. *Atmos. Chem. Phys.* **11**, 4333 (2011).

- [23] G. A. Briggs, *Plume Rise USAEC Critical Review Series TID-25075*. National Technical Information Service, Springfield, (1969) 22161.
- [24] M. Bursik. *Geophys. Res. Lett.* **28**, 3621 (2001).
- [25] M. Woodhouse, A. J. Hogg, J. C. Phillips, R. S. J. Sparks. *J. Geophys. Res. Solid Earth* **118**, 92 (2013).
- [26] O. Björnham, H. Grahn, P. von Schoenberg, A. Waleij, B. Liljedahl, N. Brännström. *Atmos. Environ.* **169**, 287 (2017).
- [27] G. Ooms, A. Mahieu. *Appl. Sci. Res.* **36**, 339 (1981).
- [28] M. Marro, P. Salizzoni, F. X. Ciercol, K. E. D. L. Soulhac. *Environ. Fluid Mech.* **14**, 201 (2014).
- [29] G. Briggs. “Some recent analysis of plume rise observation”, in *Proc. Second International Clean Air Congress*, Academic Press, New York (1971).
- [30] G. Briggs. *Atmos. Environ.* **6**, 507 (1972).
- [31] M. Sofiev, T. Ermakova, R. Vankevich. *Atmos. Chem. Phys.* **12**, 1995 (2012).
- [32] S. Goodrick, S. R. Freitas, C. Kottmeier, I. Kraut, D. Rieger, H. Vogel, B. Vogel. *Int. J. Wildland Fire* **12**, 83 (2013).
- [33] C. Walter, S. R. Freitas, C. Kottmeier, I. Kraut, D. Rieger, H. Vogel, B. Vogel. *Atmos. Chem. Phys.* **16**, 9201 (2016).
- [34] S. Beirle, C. Hörmann, M. Penning de Vries, S. Dörner, C. Kern, T. Wagne. *Atmos. Chem. Phys.* **14**, 8309 (2014).
- [35] T. A. Mather, R. G. Harrison, V. I. Tsanev, D. M. Pyle, M. L. Karumudi, A. J. Bennett, G. M. Sawyer, E. J. Highwood. *Observations of the plume generated by the December 2005 oil depot explosions and prolonged fire at Buncefield (Hertfordshire, UK) and associated atmospheric changes*. in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*. 2007.
- [36] I. Ashraf, S. Hur, Y. Park. *IEEE Access* **5**, 8250 (2017).
- [37] M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Science & Business Media, New York (2012).
- [38] S. Raza, R. Avila, J. Cervantes. *Nucl. Eng. Des.* **208**, 15 (2001).
- [39] G. Tinarelli, L. Mortarini, S. T. Castelli, G. Carlino, J. Moussafir, C. Olry, P. Armand, D. Anfossi. “Review and validation of MicroSpray, a Lagrangian particle model of turbulent dispersion”, in *Lagrangian Modeling of the Atmosphere*. pp. 311–328, American Geophysical Union, Washington D. C. (2012).
- [40] N. Kljun, M. Rotach, H. Schmid. *Boundary-Layer Meteorol.* **103**, 205 (2002).
- [41] P. de Haan. *Atmos. Environ.* **33**, 2007 (1999).
- [42] A. M. Stuart. *Acta Num.* **19**, 451 (2010).
- [43] I. G. Enting. *Inverse Problems in Atmospheric Constituent Transport*, Cambridge University Press, Cambridge (2005).
- [44] A. Tikhonov. On the solution of incorrectly put problems and the regularisation method. *Outlines Joint Sympos. Partial Differential Equations*, 1963: p. 261–265.
- [45] S. Nash, A. Sofer. *Linear and Nonlinear Programming*, McGraw-Hill College, New York (1996).
- [46] S. Jones. ISIS Sets Sulfur Plant Ablaze In Northern Iraq, Choking The Air With Deadly Chemicals, in *Huffington Post*. 2016.
- [47] Z. Ugray, L. S. Lasdon, J. C. Plummer, R. Marti. *INFORMS J. Comput.* **19**, 328 (2007).
- [48] J. Lagarias, J. Reeds, M. Wright. *SIAM J. Optim.* **9**, 112 (2007).