Atmospheric Environment 44 (2010) 3787–3796

Contents lists available at ScienceDirect

Atmospheric Environment

journal homepage: www.elsevier.com/locate/atmosenv

Characterization of atmospheric contaminant sources using adaptive evolutionary algorithms

Guido Cervone [a,b,*], Pasquale Franzese [b], Adrian Grajdeanu [c]

[a] Dept. of Geography and Geoinformation Science, George Mason University, United States
[b] Center for Earth Observing and Space Research, George Mason University, United States
[c] Dept. of Computer Science, George Mason University, United States

ARTICLE INFO

Article history:
Received 18 December 2009
Received in revised form
25 June 2010
Accepted 28 June 2010

Keywords:
Source characterization
Evolutionary algorithms
Monte Carlo
Dispersion modeling

ABSTRACT

The characteristics of an unknown source of emissions in the atmosphere are identified using an Adaptive Evolutionary Strategy (AES) methodology based on ground concentration measurements and a Gaussian plume model. The AES methodology selects an initial set of source characteristics including position, size, mass emission rate, and wind direction, from which a forward dispersion simulation is performed. The error between the simulated concentrations from the tentative source and the observed ground measurements is calculated. Then the AES algorithm prescribes the next tentative set of source characteristics. The iteration proceeds towards minimum error, corresponding to convergence towards the real source.

The proposed methodology was used to identify the source characteristics of 12 releases from the Prairie Grass field experiment of dispersion, two for each atmospheric stability class, ranging from very unstable to stable atmosphere. The AES algorithm was found to have advantages over a simple canonical ES and a Monte Carlo (MC) method which were used as benchmarks.

1. Introduction

Atmospheric contaminant hazards typically include traffic emissions, forest fires, volcano ash plumes, and intentional or accidental releases of toxic chemical, biological, or radiological agents. An unknown source of emissions can be effectively identified by source detection algorithms, which in general can be divided in two categories: backward and forward simulation techniques. Backward techniques are based on a transport and dispersion simulation in the reverse direction (i.e., from the receptor to the source) and may include Kalman filtering, adjoint and tangent linear models, and variational data assimilation (Rao, 2007). Instead, forward techniques perform multiple transport and dispersion simulations iteratively from different candidate sources, and compare the resulting concentrations to the available measurements. The algorithms are designed to find the characteristics of the source that minimize the error between simulated and measured concentrations. These methods can be used with any type of dispersion model, can be implemented independently of the amount and type of available data, and can be applied to non-linear processes as well. In this paper we will be using a forward simulation technique.

The simplest implementation of the source detection forward simulation technique employs a canonical Monte Carlo (MC) algorithm. The basic MC method may often converge to sub-optimal solutions, especially when the initial guess is far from the real source. This drawback is exacerbated for large domain sizes, creating potential difficulties for large scale dispersion problems (Cervone and Franzese, 2010). Additional difficulties associated with MC algorithms include a significant number of possible design decisions and parameters to be chosen such as, e.g., the annealing schedule, the burn-in period, and various classes of variations or proposals (Johannesson et al., 2004)

More refined source detection methods are obtained by coupling MC or Markov chain Monte Carlo stochastic sampling with Bayesian updating and inference methods, where probability distributions for the parameters are iteratively updated. The Bayesian Monte Carlo (BMC) methods are significantly more powerful, overcome most of the limitations of the MC method, and have been tested in several different settings. Sohn et al. (2002) describe the application of a BMC method along with an indoor airflow and pollutant transport model to the characterization of air

* Corresponding author. Dept. of Geography and Geoinformation Science, George Mason University, 4400 University Dr, Fairfax, VA 22030, United States.
E-mail addresses: gcervone@gmu.edu (G. Cervone), pfranzes@gmu.edu (P. Franzese), agrajdea@gmu.edu (A. Grajdeanu).

1352-2310/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.atmosenv.2010.06.046

pollution sources in a building. Johannesson et al. (2004) successfully apply BMC methods to a set of synthetic data to reconstruct source location, release rate and start time. Lundquist et al. (2005) apply BMC methods to designing more accurate and cost-effective sensor networks. The study focus on the sensitivity of a source location algorithm to different sensor network configurations in terms of number of sensors, detection sensitivities, false-alarm rates, frequency of data collection and sensor detection range. Neumann et al. (2006) and Chow et al. (2006, 2008) apply the BMC approach of Johannesson et al. (2004) to the source characterization of urban dispersion problems at building scales. Senocak et al. (2008) simulate dispersion using an analytical Gaussian plume model and extend the BMC methodology described in Johannesson et al. (2004) and Chow et al. (2006, 2008) introducing a likelihood function that takes into account both zero and non-zero concentration measurements, and which includes random parameters estimated using data and prior probabilities to avoid tuning.

Delle Monache et al. (2008) build on the work of Johannesson et al. (2004, 2005); Lundquist et al. (2005), Neumann et al. (2006), Chow et al. (2008) and includes Markov chains, prior likelihoods, cooling-off procedures, adaptive steps, burn-in, sophisticated convergence criteria, and tuned error parameters. Delle Monache et al. (2008) were able to locate with remarkable accuracy the source of a real accidental release of radioactive material on a continental scale, which was detected by only a few sensors, using a Lagrangian particle dispersion model to perform the iterative forward transport and dispersion simulations.

An alternative approach is the forward source detection method based on genetic algorithms (GA) (e.g., Holland, 1975; De Jong, 2008). Unlike Monte Carlo algorithms which maintain one candidate solution, GA maintain populations of candidate solutions. GA are thus likely to be less affected by ill initial conditions and more adept at escaping premature convergence to local optima. The promise of the *canonical* GA as initially introduced by Holland (1975) was that they would be problem- and domain-independent. For this reason, real-valued variables were encoded using a binary representation. Standard genetic operators afforded variation by means of bit-flip mutation and crossover at either bit or gene boundary. However, such implementations suffered from poor performance due to the limited precision with which the binary representation encodes real values and, in no lesser measure, due to disruptive effects induced by standard genetic operators. In contrast, Evolutionary strategies (ES) were developed specifically for continuous domain optimization (Rechenberg, 1971; Schwefel, 1974). This characteristic makes them particularly suited to the source detection problem, where real-valued attributes are most common. In recent times, GA and ES converged (De Jong, 2008): it is quite common now to encounter GA operating on real-valued parameters such as the continuous parameter GA described in Haupt (2005), Haupt et al. (2007), Allen et al. (2007) and Long et al. (2010), as well as ES with parallel search streams and even crossover operator (e.g., De Jong, 2008).

For example, a continuous parameter GA was used in the context of a source apportionment exercise (Haupt, 2005); coupled with a Gaussian plume model in a source characterization problem using synthetic data as receptor data (Haupt et al., 2007; Allen et al., 2007); and to assess the sensitivity of a source detection method to the number of sensors (Long et al., 2010). In all these applications, the evolutionary approach performed well, demonstrating its suitability as optimization technique. A variant of the evolutionary algorithm approach was proposed by Cervone et al. (2010), where new candidate solutions were created through a reasoning process guided by machine learning, rather than through the Darwinian operators of mutation and crossover.

In this paper we present an implementation of an Adaptive ES (AES) methodology to identify the characteristics of a source of atmospheric emissions. We will employ the AES method to find the source characteristics of several releases from the Prairie Grass controlled field experiment (Barad, 1958). The releases were selected so as to sample all atmospheric stability classes. The algorithm was tested against the results obtained by a simple MC program and by a traditional basic ES, which represent baseline methods used as benchmarks.

The methodology and details on the algorithms used are shown in Section 2. Section 3 discusses the prairie Grass experiment and the Gaussian plume model used in the experiments. Section 4 describes the numerical experiments performed and presents the comparative study.

2. Methodology

The methodology presented in this paper is based on a class of evolutionary algorithms called Evolutionary strategies (Rechenberg, 1971; Schwefel, 1974). While evolutionary algorithms in general are heuristics based on biologically inspired iterative processes, the Evolutionary strategies (ES) address continuous parameter optimization problems in particular. Formally, if n is the number of optimized parameters p_i, and D_i are their domains of interest, then the evolutionary algorithm attempts to minimize a goal function $f(p_{1:n})$.

Central to the terminology of an evolutionary algorithm is the concept of *potential solution*, or *candidate solution*. Each solution contains a value for every one of the optimized parameters p_i, as well as values for meta-parameters if any. Like all evolutionary algorithms, ES maintain a *population* of potential solutions and attempt to improve on it iteratively. The particular flavor used in this paper is known as ES($\mu + \lambda$) in which μ current solutions act as parents and are used to produce λ offspring that compete with their parents for survival. Table 1 gives the outline of this algorithm. At each iterative step, only the best μ solutions (parents and offspring) are maintained, the rest being discarded.

Producing offspring from parents involves cloning, followed by the application of a perturbing (mutation) operator that induces minor stochastic variation. When optimizing on a continuous domain, most of the time a Gaussian mutation is used. The standard deviation of this operator, the meta parameter σ, tunes the degree of variation induced by the operator. One of the most delicate aspects is quantifying and controlling the magnitude of the stochastic variations. On a few standard search landscapes, Rechenberg (1971) formally proved the *one-fifth rule*, of which a more modern and relaxed form is as follows. If the ratio of offspring improving on their parental performance is less than 1/5, then σ should decrease; otherwise σ should increase. This rule has been empirically shown to perform well on many other landscapes, so it is still used in some of its modern incarnations.

In the case of multi-dimensional optimization, Schwefel (1974) introduced a mechanism that affords each dimension, corresponding to one parameter, a separate meta parameter σ that is

Table 1
The ES($\mu + \lambda$) algorithm.

Step 1:	Randomly generate initial population of $\mu + \lambda$ candidate solutions
Step 2:	Evaluate each solution in the population
Step 3:	Sort solutions in decreasing order of performance
Step 4:	Truncate the solution population to only μ best (parents)
Step 5:	If termination criteria met, exit, reporting best (first) parent
Step 6:	For each of the μ parents, generate λ/μ offspring and add to the population
Step 7:	Resume with *Step 2*

allowed to vary independently of other parameters. The *one-fifth rule* changes in such a way that it no longer indicates the direction of modifying the σ meta parameters, but merely a weak bias of the change. Specifically, if σ_i are all the variance parameters, r is the ratio of offspring improving on their parental fitness, and $N(0,1)$ is a random variable with normal distribution, of mean zero and standard deviation 1, then the updating rule is:

$$\beta = N(0,1) + \gamma(r - 0.2) \tag{1}$$

$$\sigma_i \leftarrow \sigma_i e^{\alpha[\beta + N(0,1)]} \tag{2}$$

where $\alpha = 0.2$ and $\gamma = 5$ are modulating parameters empirically chosen. Updating the optimized parameters p_i follows as

$$p_i = p_i + N(0, \sigma_i) \tag{3}$$

with the understanding that p_i will be constrained to the domain of interest D_i.

The process of randomly generating a solution in the initial population consists of assigning to each of the optimized parameters p_i random values uniformly sampled from D_i. The global parameters σ_i are initialized to 1/6 the size of D_i.

2.1. Monte Carlo simulation

Monte Carlo methods represent a family of optimization techniques where many potential solutions are *stochastically* generated and tried. This principle is applied to generating many starting points for a ES$(1+1)$. Note that ES$(1+1)$ is mostly a new name for an algorithm also known as stochastic hill climbing (Russell and Norvig, 1995, p. 115). The only significant addition to hill climbing is the use of the one-fifth rule that acts as a self-adaptive mechanism to speed up convergence. After every 100 births the ratio of offspring improving on parental performance is computed and the meta parameters σ_i are updated.

The ES$(1+1)$ process used in this paper iterates for 1000 steps and then is restarted from a new stochastically generated point. The restart aspect is salient because a greedy heuristic, while quickly convergent to an optimum, is also very much prone to identifying a *local* optimum. The restarting process mitigates this weakness and affords greater chances to pinpointing the *global* optimum. This heuristic is hereby called a Monte Carlo simulation because the exploratory power of this variant stems from the Monte Carlo principle of stochastically generating multiple starting points. As detailed in the introduction, more advanced and effective methodologies based on Bayesian inference have been developed and successfully implemented (e.g., Johannesson et al., 2004; Lundquist et al., 2005; Delle Monache et al., 2008). In this paper, we are interested only in the most basic implementation as a reference benchmark for comparison.

2.2. Evolutionary strategy algorithm

A different strategy of reducing the chance of convergence to a *local* optimum is the use of *parallel* search. This paper analyzes a canonical ES$(10+90)$ with deterministic parental selection: each parent produces the exact same number of $\lambda/\mu = 9$ offspring. The values of $\mu = 10$ and $\lambda = 90$ are empirical. The same mechanism of adapting the meta parameters σ_i according to the one-fifth rule every 100 births is used as in the case of Monte Carlo simulations. A very basic implementation of a strictly canonical ES is used in this study as a comparison benchmark. A comparison with more sophisticated and powerful evolutionary algorithms (e.g., Haupt, 2005; Haupt et al., 2007; Allen et al., 2007) is outside the scope of this paper.

2.3. Adaptive evolutionary strategy

Other than the one-fifth rule technique, there is the possibility of *evolving* the meta parameters σ_i independently and along with the actual parameters under optimization. This heuristic differs from the canonical ES in that the process of perturbing a current solution first alters the σ_i encoded in the solution itself, and then uses them to alter the actual optimized parameters using equations (1)–(3). Unlike the case of canonical ES with the one-fifth rule, there is no global set of σ_i, as each solution has and uses its own copy.

The intuition behind this adaptive mechanism is that for a solution to perform well (i.e., be better than other potential solutions generated) it needs to have been generated both from a good combination of optimized parameters *and* using an adequate set of meta parameters. Furthermore, this solution will pass on to its offspring (with variation) both its own good combination of optimized parameters *and* a good working set of meta parameters.

The same $\mu = 10$ and $\lambda = 90$ parameters were used as in the case of canonical ES. The $\gamma(r - 0.2)$ term in equation (1) is always zero in this case.

3. Dispersion experiment and simulations

3.1. Prairie grass experiment

The search algorithms are applied to identify the characteristics of the source in the Prairie Grass field experiment (Barad, 1958). The experiment consisted of 68 consecutive releases of trace gas SO_2 of 10 min each from a single source. The mean concentration was measured at sensors positioned along arcs radially located at distances of 50 m, 100 m, 200 m, 400 m and 800 m from the source. The number of available concentration samplers is not constant in all experiments, because only sensors that recorded values above a minimum threshold were considered reliable. The Prairie Grass dataset includes detailed information on the atmospheric conditions at the time of each release. It is then possible to classify each experiment according to Pasquill's atmospheric stability classes (Pasquill and Smith, 1983; Hanna et al., 1990). The experiments were conducted over all stability classes, ranging from very unstable (class A) to neutral (class D) to stably stratified atmosphere (class F). Note that the atmospheric stability strongly affects the characteristics of the concentration field. In unstable atmosphere spread is enhanced, with lower ground level concentrations, whereas stable atmosphere determines narrower plumes, and higher ground concentrations.

3.2. Gaussian plume model

The dispersion simulations are performed using a Gaussian reflected dispersion model, which predicts the mean concentration C_p at a location x, y and z generated by a source located at x_s, y_s, and z_s as:

$$C_p(x,y,z,x_s,y_s,z_s) = \frac{Q g_y g_z}{2\pi U\left[(\sigma_s^2+\sigma_y^2)(\sigma_s^2+\sigma_z^2)\right]^{1/2}} \tag{4}$$

with

$$g_y = \exp\left[-\frac{(y-y_s)^2}{2(\sigma_s^2+\sigma_y^2)}\right]; \tag{5}$$

$$g_z = \exp\left[-\frac{(z-z_s)^2}{2(\sigma_s^2+\sigma_z^2)}\right] + \exp\left[-\frac{(z+z_s)^2}{2(\sigma_s^2+\sigma_z^2)}\right] \tag{6}$$

Table 2
Relevant characteristics of the 12 Prairie Grass experiments considered in this study (left panel), along with absolute errors Δ in the source characteristics identified by AES. For any variable ξ, $\Delta\xi = \xi_{obs} - \xi_{pred}$. Averages over the 12 experiments of the results obtained by AES, ES and MC are also reported at the bottom.

Exp. ID	ψ	θ (deg)	U (ms^{-1})	Q (gs^{-1})	ΔD (m)	$\Delta\sigma_s^2$ (m^2)	ΔQ (gs^{-1})	$\Delta\theta$ (deg)	NRMSE
15	A	209	2.90	95.5	39.67	5.88	91.84	−4.56	0.47
47	A	243	3.02	103.1	9.18	0.00	93.81	3.33	0.46
2	B	100	1.74	83.9	60.45	0.00	116.10	56.25	1.35
7	B	188	4.02	89.9	73.02	0.00	110.10	−81.78	2.08
5	C	176	5.15	77.8	16.63	0.00	42.65	−1.02	0.68
8	C	184	4.06	91.1	8.58	1.69	20.28	1.47	0.53
11	D	184	6.77	95.9	14.47	0.00	24.97	−0.16	0.28
22	D	176	6.39	48.4	16.80	2.00	16.88	−5.02	0.27
18	E	187	2.68	57.6	14.71	0.04	-2.66	−5.70	0.45
68	E	174	2.19	42.8	7.96	2.06	13.83	−0.38	1.04
32	F	171	1.60	41.4	35.62	2.60	9.22	0.94	0.45
58	F	178	1.65	40.5	38.53	1.99	3.04	2.47	0.81
Avg. AES					27.97	1.36	45.01	−2.85	0.74
Avg. ES					71.06	3.62	84.07	2.41	2.12
Avg. MC					32.93	3.12	86.99	4.40	1.63

where Q is the mass emission rate, U is the wind speed, $\sigma_y(x, x_s;\psi)$ and $\sigma_z(x, x_s;\psi)$ are the crosswind and vertical dispersion coefficients (i.e. the plume spreads) where ψ describes the atmospheric stability class (i.e., $\psi = A$ to $\psi = F$), and $\sigma_s^2 = \sigma_y^2(x_s, x_s, \psi) = \sigma_z^2(x_s, x_s, \psi)$ is a measure of the area of the source. The dispersion coefficients were computed from the tabulated curves of Briggs (Arya, 1999). The result of the simulation is the concentration field generated by the release along an arbitrary wind direction θ.

In this study, U and ψ are assumed to be known, and their values set according to the observations reported in the Prairie Grass dataset. Each candidate solution is thus comprised of the 6 variables x_s, y_s, z_s, σ_s, Q, and θ.

3.3. Error function

Central to every evolutionary algorithm is the definition of the error function, often called fitness or objective function. The error function evaluates each candidate solution quantifying the error between the observed concentrations and the corresponding simulated values. This information is used by the search algorithm to drive the stochastic iterative process.

The definition of error, or uncertainty, in the model predictions is not univocal, and depending on the case, different measures of accuracy of a dispersion calculation can be adopted (Hanna et al., 1993). For example, the error may refer to the comparison of

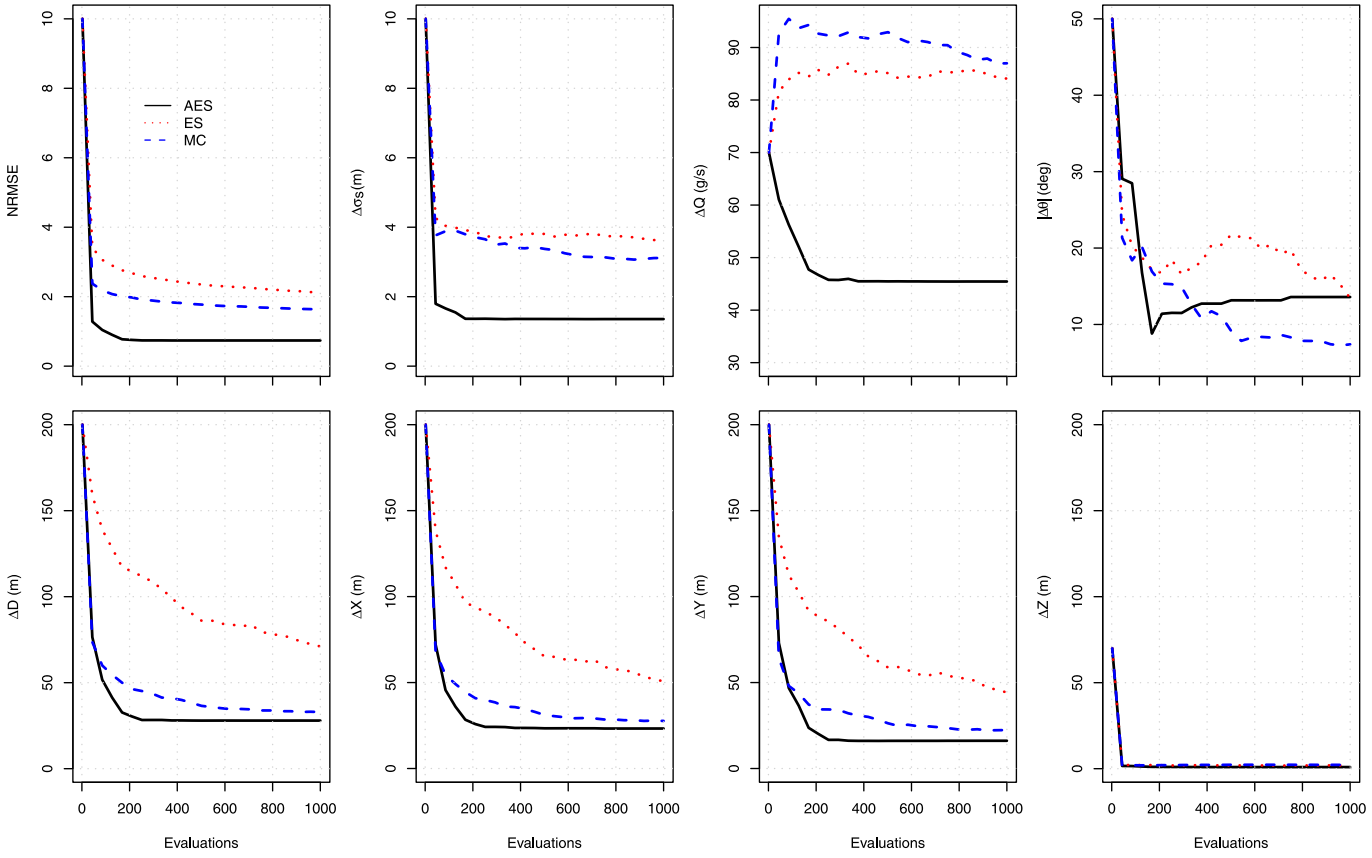

Fig. 1. Evolution of NRMSE, $\Delta\sigma_s$, ΔQ, $|\Delta\theta|$, ΔD, ΔX, ΔY, and ΔZ averaged over all the releases as functions of the number of iterations obtained with AES (solid line), ES (dotted line), and MC (dashed line).

Table 3
Summary of performance gains of AES over ES and MC for each release. The performance gain over ES and over MC for each variable is defined as the best value achieved by ES and MC respectively, divided by the best value achieved by AES.

ID	D		σ_s^2		Q		θ		NRMSE	
	$\frac{ES}{AES}$	$\frac{MC}{AES}$	$\frac{ES}{AES}$	$\frac{MC}{AES}$	$\frac{ES}{AES}$	$\frac{MC}{AES}$	$\frac{ES}{AES}$	$\frac{MC}{AES}$	$\frac{ES}{AES}$	$\frac{MC}{AES}$
15	0.94	0.73	0.67	0.64	0.92	0.96	0.97	1.00	2.59	2.04
47	3.55	2.47	5.62	4.38	0.87	0.95	0.96	0.98	2.30	1.72
2	0.88	0.92	5.24	3.71	0.87	0.95	0.99	1.01	1.06	1.03
7	1.05	1.03	6.39	5.95	0.90	0.96	1.85	1.79	1.05	1.03
5	1.13	1.14	3.78	3.74	1.34	1.38	1.00	1.00	1.68	1.59
8	3.62	1.55	1.30	1.43	1.36	1.55	0.99	1.00	2.45	2.47
11	2.15	1.28	3.27	3.14	1.29	1.45	1.01	1.00	6.68	4.43
22	1.94	0.98	1.54	1.18	1.98	2.10	1.00	0.99	9.19	6.89
18	3.97	1.83	3.65	3.71	2.55	2.36	0.92	1.00	6.33	4.20
68	7.61	3.03	1.38	1.39	2.27	2.39	1.03	1.00	2.66	1.97
32	5.10	1.42	1.43	1.30	2.94	2.58	1.14	1.04	7.18	4.82
58	6.19	1.15	1.74	1.30	3.71	2.65	0.87	1.00	4.84	3.36

simulated and observed peak concentration over the entire sensor network, regardless of time and location of occurrence of the peak. In other cases, the only realistic expectation is to calculate the predictions that fall within a certain factor (e.g., 2 or 10) of the observations. For this reason, inter-model comparison studies always include several performance metrics (e.g., Chang et al., 2003, 2005; Chang and Hanna, 2004).

In this study, we will use the Normalized Root Mean Square Error (NRMSE), which was shown to be a suitable metric for source detection algorithms in a comparative study of characteristics of several error functions (Cervone and Franzese, 2010). The fitness of each candidate solution is computed using the NRMSE between observed and simulated concentrations:

$$\mathrm{NRMSE} = \sqrt{\frac{\overline{(C_o - C_p)^2}}{\overline{C_o} \cdot \overline{C_p}}} \tag{7}$$

where C_o is each sensor's observed mean concentration, C_p is the corresponding simulated value, and the bar indicates an average over all the observations.

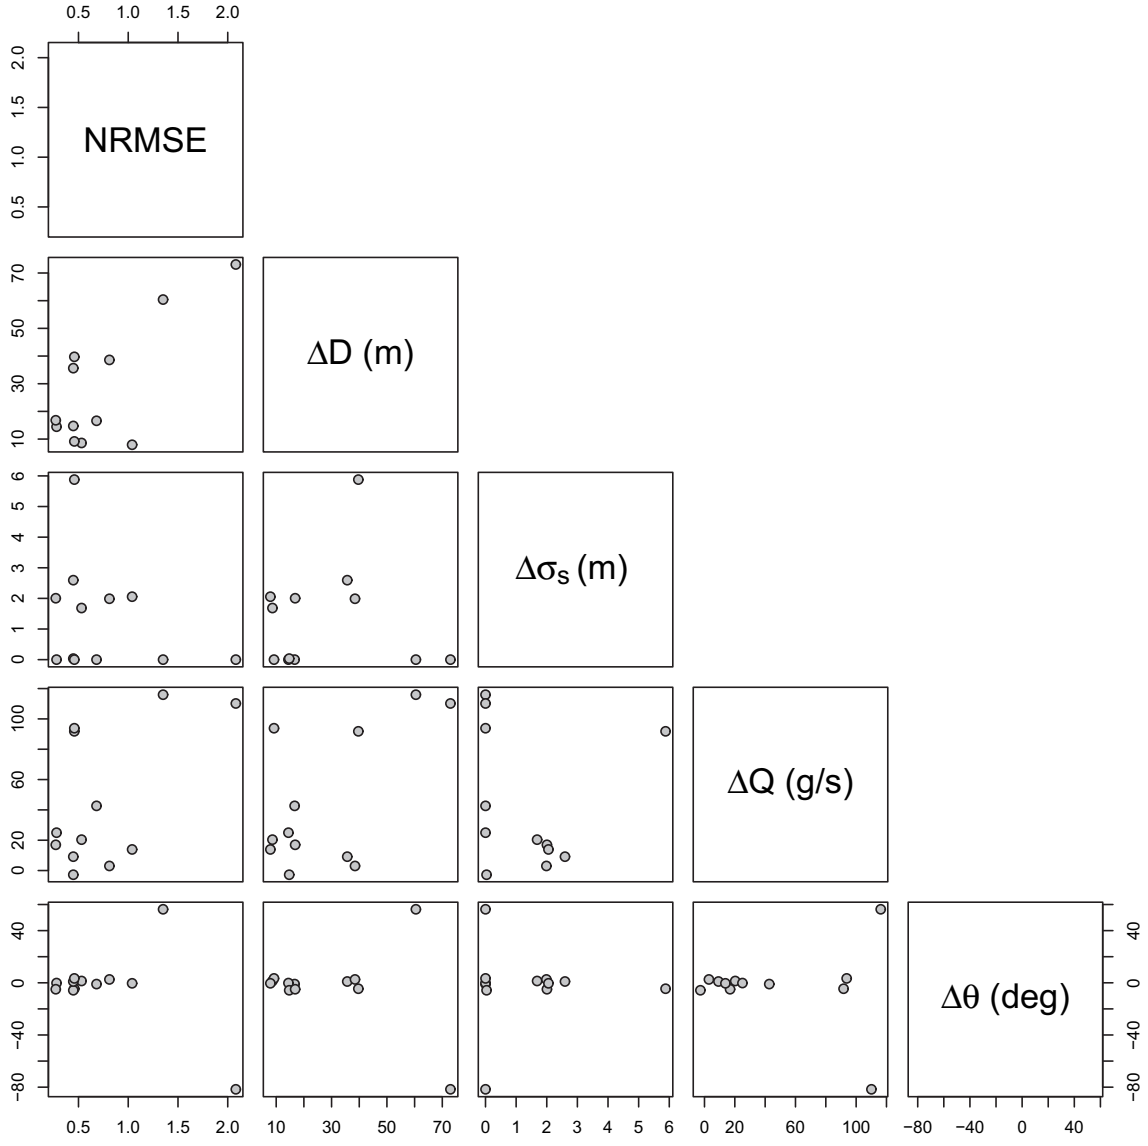

Fig. 2. Pairwise graph generated by AES showing the results of all 12 experiments, where each variable is plotted against all others.

4. Results

The search methodologies are applied to find the characteristics of the sources of 12 Prairie Grass tracer release experiments, two for each atmospheric stability class. The search space was defined as follows: x and y both range from −4 km to 4 km, defining an area of 64 km^2 centered at the source; z ranges from 0 to 200 m; Q from 0 to 200 gs^{-1}, and θ from 0 to 360 degrees. The distance from the source is defined as $D = \sqrt{(x-x_s)^2 + (y-y_s)^2 + (z-z_s)^2}$. Since the size of the source was negligible compared to the plume spread over the sensors, the emissions can be assumed to be released from a point source, and the source size was assumed to be $\sigma_s^2 = 1$ m^2.

For each of the 12 experiments, a series of 30 runs was performed by each of the three algorithms used, i.e., Adaptive Evolutionary Strategy (AES), Evolutionary Strategy (ES), and Monte Carlo (MC), totaling 1080 runs. For each set of 30 runs, all parameters were kept constant except the random seed, which is responsible for the generation of the initial set of candidate solutions and for guiding the stochastic search.

Table 2 summarizes relevant characteristics of the field dispersion experiments used in this study, along with the results in terms of absolute errors obtained by AES. For any variable ξ, the absolute error Δ is defined as $\Delta\xi = \xi_{obs} - \xi_{pred}$. As reported in Table 2, AES is able to detect the source location within an average of about 30 m from the real location, with a maximum error of about 73 m in the case of experiment 7, which was conducted in unstable conditions (i.e., stability class B). The two experiments conducted under atmospheric stability class B appear to display the largest errors for all variables. This may be due to a combination of dispersion model inadequacy and poor data quality. For example, most of the experiments in the stability class B did not pass quality control and were not included in the dataset reported by Hanna et al. (1990). The emission rate Q is captured with an average error of about 45 gs^{-1}, or within about a factor 2 of the real Q. The wind direction is generally predicted with good accuracy, except for the cases in stability class B, which will be discussed below. NRMSE is also in general quite low, showing that part of the errors reported in the table are due to the multi-variable nature of the solution, which admits low NRMSE for specific combinations of the variables within a finite range. Overall, the results indicate the suitability of the Gaussian reflected model for this application. Table 2 also reports the results averaged over the 12 experiments obtained by AES, along with ES and MC, which are in general less accurate.

A more complete comparison between the algorithms is shown in Fig. 1, which shows the evolution of NRMSE, $\Delta\sigma_s$, ΔQ, $|\Delta\theta|$, ΔD, ΔX, ΔY, and ΔZ as functions of the number of iterations for AES, ES, and MC. The results are averaged over all the releases, namely the averages are taken over 30 runs performed for each release, over all releases. Each run of the algorithm differs only in the random initial

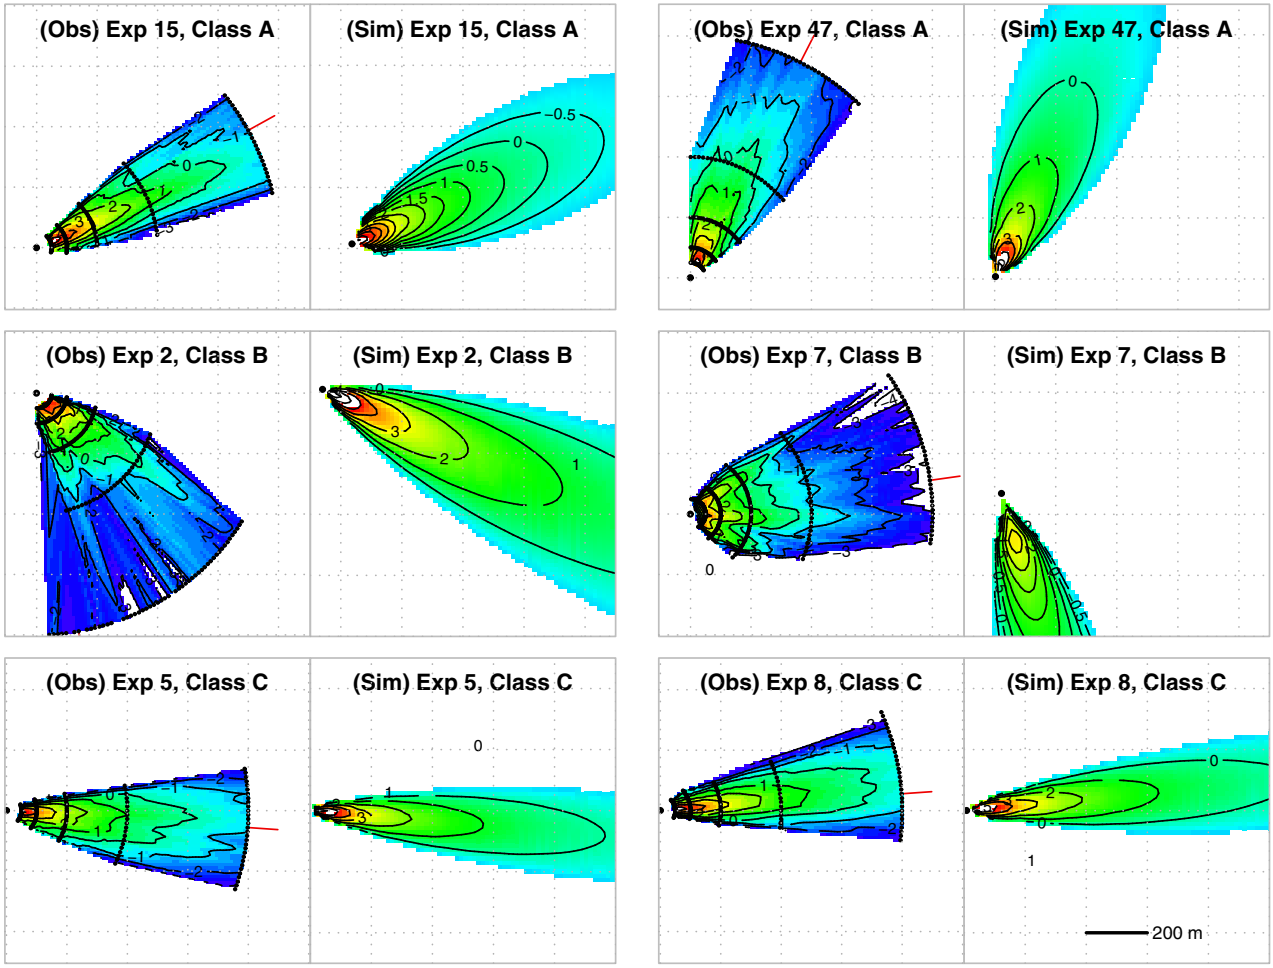

Fig. 3. Contour plot of concentrations for the Prairie Grass experiments with stability classes A, B, and C. The left column of each panel shows the field reconstructed from observations; the right column shows the corresponding field generated by the Gaussian model [Equation (4)] using the source characteristics identified by AES. The scale indicated at the bottom right is the same in all the graphs.

population. Fig. 1 (top-left panel) shows that convergence for NRMSE is faster with AES, namely on average AES determines a solution with a smaller error than both ES and MC by a wide margin in a smaller number of evaluations. MC is also seen to outperform ES (i.e., smaller error in a shorter time). Note that all curves start with NRMSE well above the range of the vertical axis. The first few points have been omitted in order to reduce overplotting. Fig. 1 also shows that ΔQ calculated by either ES and MC is about a factor of 2 larger than AES. The other variables have in general better accuracy with AES, with the exception of θ, which displays a smaller error when calculated by MC (the AES $|\Delta\theta|$ is strongly influenced by release 7, which is an outlier and will be separately analyzed below). AES slightly outperforms MC in the estimation of D. (The best result of AES is a solution about 8 meters of the source). However, we emphasize that the location of the source is only one of the characteristics identified by the algorithm. A better localization of the source does not necessarily correspond to a better solution, and to a smaller NRMSE. In other words, a solution which predicts a shorter distance from the source may display a larger error than a solution which predicts a larger distance. This phenomenon is due to the multi-dimensionality of the problem, where all characteristics of the source are simultaneously optimized, leading to a multi-modal solution space, and also to possible approximations in the dispersion model and errors in the field observations.

Table 3 provides a summary of the performance gains of the AES algorithm over ES and MC obtained for each release. The performance gain over ES and over MC for each variable is defined as the best value achieved by ES and MC respectively, divided by the best value achieved by AES. For example, the NRMSE performance gain over ES and over MC is defined as the best NRMSE (i.e., the minimum error) achieved by ES and MC respectively, divided by the best NRMSE achieved by AES.

AES achieves consistently a lower NRMSE, followed by MC: the average NRMSE for AES is 0.74, while the average NRMSE for ES and MC are 1.63 and 2.12 respectively (as reported in the two bottom rows of Table 2). The average NRMSE performance gain for AES is thus 2.86 over ES, and 2.2 over MC.

In terms of distance D, AES achieves better results in most of the 12 releases, and in general MC outperforms ES. The performance gains for D averaged over all releases of AES over ES and MC corresponds to 2.54 and 1.17 respectively.

The comparative advantage for source location and emission rate identification of AES is more pronounced for increasing atmospheric stability (i.e., from class A to F), which corresponds to smaller plume spread and consequently to fewer sensors sampling non-zero concentration. This may be an indication that the AES search strategy is less sensitive to the amount of relevant information input into the model.

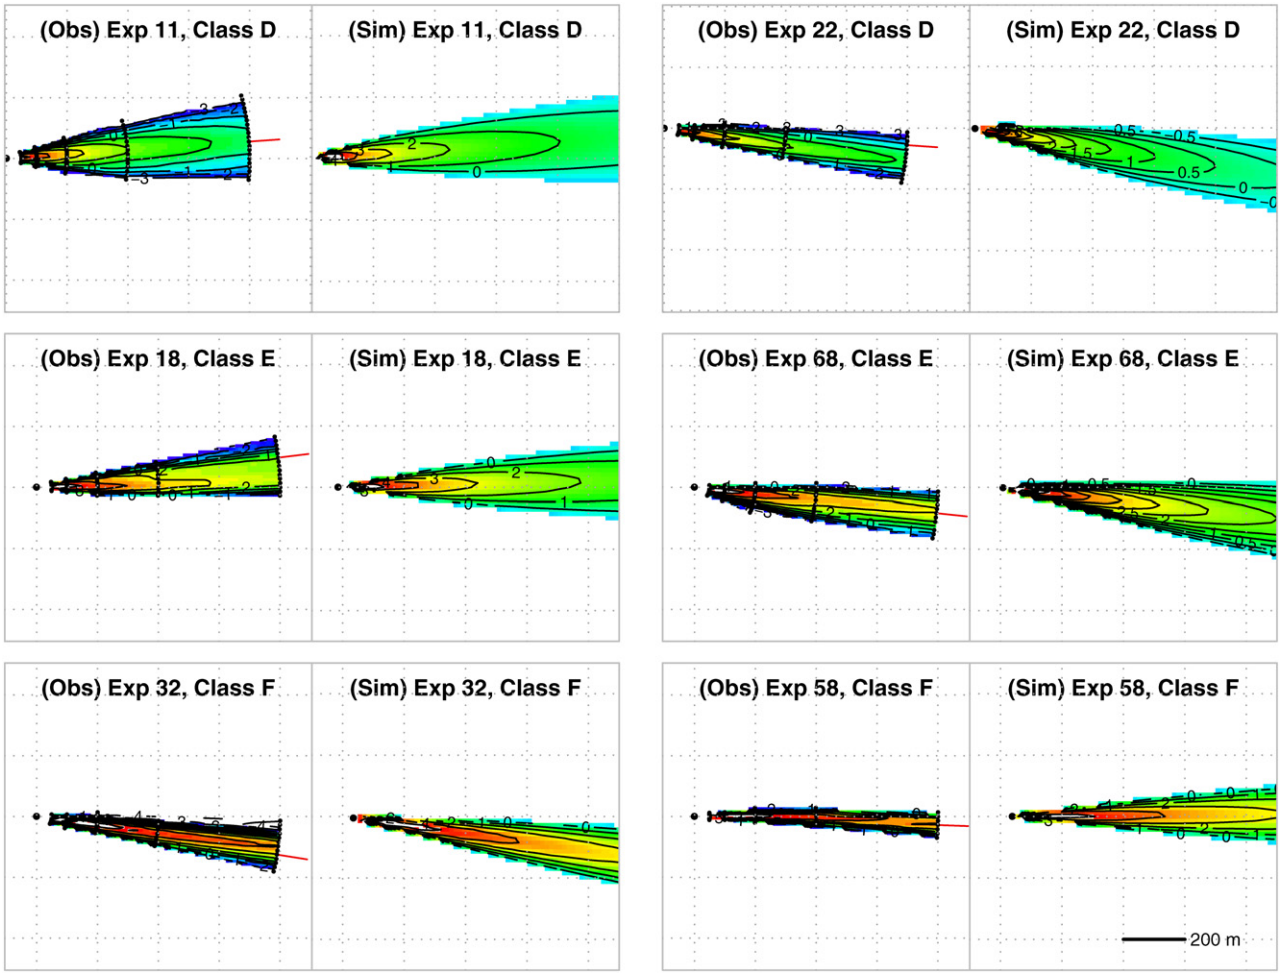

Fig. 4. Contour plot of concentrations for the Prairie Grass experiments with stability classes D, E, and F. The left column of each panel shows the field reconstructed from observations; the right column shows the corresponding field generated by the Gaussian model [Equation (4)] using the source characteristics identified by AES. The scale indicated at the bottom right is the same in all the graphs.

Fig. 2 shows a pairwise plot generated by AES comparing the results of all 12 experiments performed, where each variable is plotted against all others. The pairwise plots are useful to detect inter-variable correlations. The plot shows that NRMSE is a good predictor for *D*, because a correlation can be detected between the two variables. The predicted wind direction is relatively very accurate for most cases over the range of values of NRMSE. The two cases with larger errors in the predicted wind direction correspond also to the highest NRMSE values. A consistent behavior is observed when comparing the predicted wind direction with all other variables. This indicates that over- or under-estimating the wind direction can lead to very large NRMSE.

Figs. 3 and 4 show the observed and predicted mean concentration field for the 12 experiments considered. The plumes were reconstructed using the Akima spline interpolation algorithm (Akima, 1978). The observed concentration field was computed by interpolating the values measured at the sensors, which are represented by full dots. Note that the concentration footprint decreases as stability increases, consistent with the notion that stability inhibits dispersion. In our calculations a smaller footprint implies a smaller number of sensors measuring concentrations, and therefore a potentially higher error in the solution. In general, source location, plume direction and concentration field are reproduced accurately, except for experiment 7, where the predicted wind direction differs by almost 90 degrees from the observations. An analysis of the anomaly of this case is useful to elucidate the range of possible behaviors of the multi-parameter solutions identified by the search algorithm.

Fig. 5 shows the concentration observed at each sensor for Prairie Grass release 7, as a function of the sensor number. The sensors are positioned along five concentric arcs, located at 50 m, 100 m, 200 m, 400 m and 800 m from the source. In the figure, each arc is represented in one of the panels which are indicated by alternate white and grey background. For each arc, the sensor numbers are sorted counterclockwise. The concentration simulated by Equation (4) using the source characteristics *identified by AES* is also plotted in the top figure using a dashed line; the concentration simulated by Equation (4) using the source characteristics reported in the Prairie Grass dataset is plotted in the bottom figure using a dashed line.

The crosswind profile of concentration is tri-modal (this is especially visible at the two closest arcs, and is consistent with Fig. 3), suggesting that the distribution of wind direction had a similar shape and caused the plume to disperse unevenly. The ground measurements are in fact concentrations over a fixed period of time, and as such they reflect the spatio-temporal variation of the wind plume. The multi-modal distribution is likely to be attributed to wind shifts occurred during the experiments. This case illustrates the limits of a simple analytical dispersion model such as the Gaussian plume model. More complex simulations may provide a more accurate, and possibly time-dependent estimate of θ.

The theoretical concentration calculated by the Gaussian model consists of a single peak much higher than any of the observed values (Fig. 5, bottom). The poor agreement between the observed and simulated values is responsible for a high NRMSE of 2.3. On the other hand, the concentration simulated using the AES parameters (including a wind direction off by about 90 degrees) shows a better agreement, leading to a lower NRMSE equal to 0.9 (this value does not correspond to NRMSE = 2.08 reported in Table 2, which is the average over 30 runs). In other words, a simulated plume with a wind direction off by about 90 degrees is a better fit (at least in terms of NRMSE) than the theoretical plume.

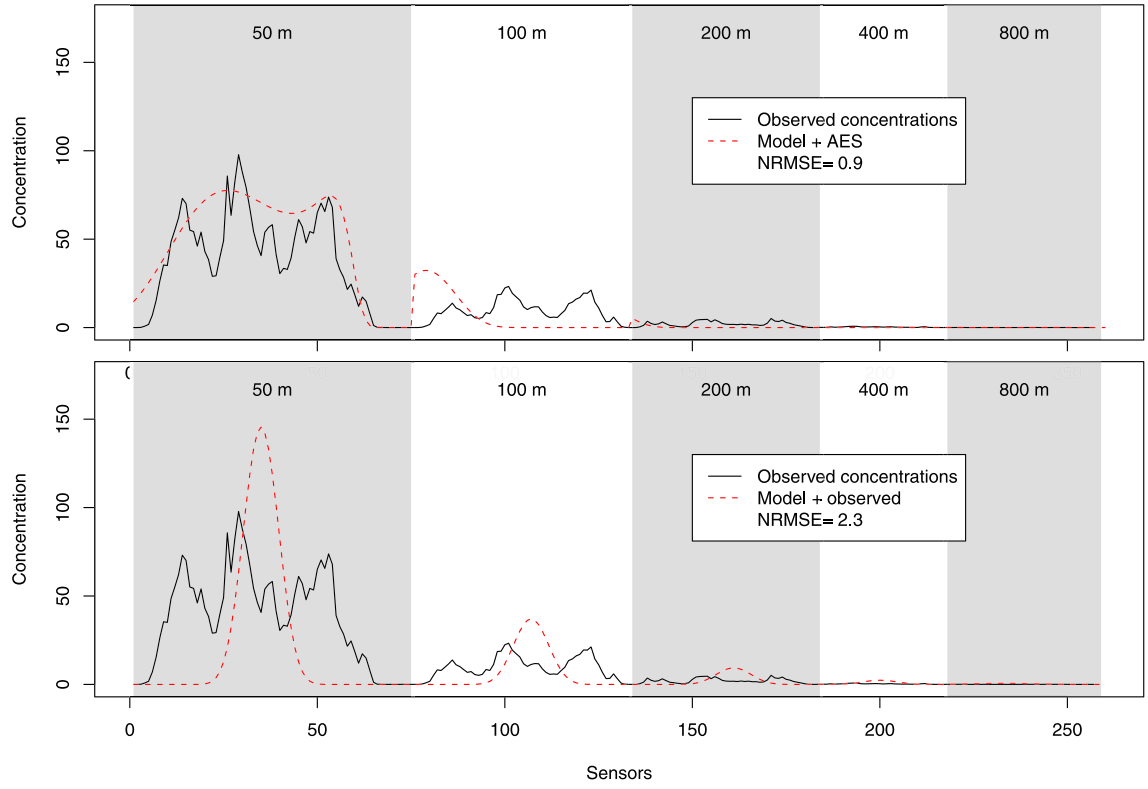

Fig. 5. Cross-wind profiles of concentration for Prairie Grass release 7, plotted as functions of the sensor number. The sensors are positioned along five concentric arcs, located at 50 m, 100 m, 200 m, 400 m and 800 m from the source. Each arc is represented in one of the panels which are indicated by alternate white and grey background. For each arc, the sensor numbers are sorted counterclockwise. Solid line: measured concentration; dashed line top figure: concentration simulated by the Gaussian model [Equation (4)] using the source characteristics identified by AES; dashed line bottom figure: concentration simulated by Equation (4) using the source characteristics reported in the Prairie Grass dataset.

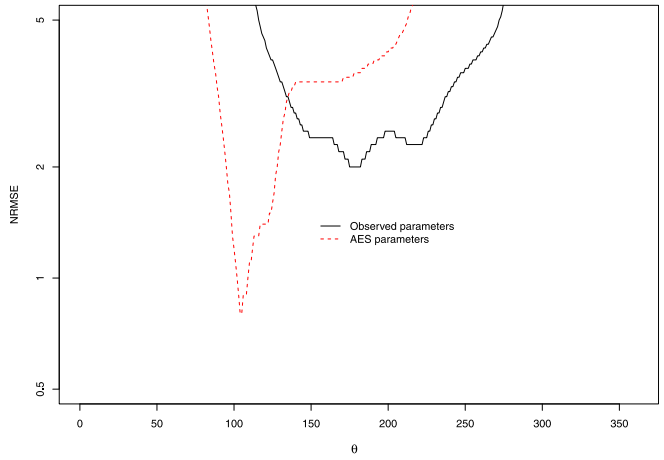

Fig. 6. Sensitivity analysis of NRMSE with respect to wind direction θ for Prairie Grass experiment 7 using the observed (solid line) and the AES-derived (dashed line) source characteristics.

Fig. 6 shows a sensitivity analysis of NRMSE with respect to wind direction θ for Prairie Grass release 7 using the observed (solid line) and the AES-estimated (dashed line) source characteristics. The NRMSE minima correspond to 188 degrees (roughly corresponding to the observed plume direction in Fig. 3) using the theoretical parameters, and 106 degrees using the AES-estimated parameters. Note the sharp gradient of the AES-estimated NRMSE as a function of wind direction, which explains the very small error generally associated with the identified wind direction, as shown in Table 2.

5. Conclusions

We used the Adaptive Evolutionary Strategy (AES) algorithm, which is a modified version of a standard Evolutionary Strategy (ES) algorithm, to identify six variables characterizing the sources of 12 releases from the Prairie Grass field experiment. The results indicate the suitability of AES to be employed in problems of source characterization of atmospheric releases. All the experiments were repeated using a simple canonical ES algorithm, and a canonical Monte Carlo (MC) simulation, which were used as benchmarks. The experiments show that AES has achieved both higher accuracy compared to the benchmarks and a much faster rate of convergence.

In a multi-dimensional problem, the distance from the source is only one of the variables optimized by the algorithm. Therefore a lower error may not always correspond to a solution closer to the source. Although the results are generally good, there are cases where the error may be significant. The ground concentrations generated by the reflected Gaussian model may not adequately represent the plume, therefore creating an inconsistency between the model and the observations. This can be overcome by employing more sophisticated models. The function used to compute the error between sensor measurements and simulated concentrations may lead to convergence to local optima. The NRMSE function is only one of the possible feedback mechanisms which can be employed as error function (Cervone and Franzese, 2010). It is possible that for different problems other error functions could provide a better feedback to the search algorithm. Finally, a certain amount of noise is to be expected in all experimental data. The Prairie Grass experiment provides good and reliable data, which have been found to approach a Gaussian distribution, and thus can be accurately modeled by a Gaussian plume model. However, there are of course inconsistencies and departures from the predicted values, as detailed in our discussion of release 7 and in Figs. 5 and 6.

In this study we performed 30 runs (each with a different initial condition) for each of the 12 Prairie Grass releases considered. Each run includes at least 10^6 evaluations and stops as soon as possible after reaching this threshold. Typical run-time durations (on a common Linux PC) for one source estimation, which includes 30 runs, are about 17 min for AES, 21 min for ES, and 17 min for Monte Carlo. However, the codes were not specifically optimized for performance (beyond standard compiler optimization) so the straight execution times could be reduced. For example, the above run-time values include significant amounts of input/output operations in order to display performance traces (e.g., best-so-far solutions at intermediate times). If the sole final solution is of interest, the computational times would be shorter.

Future work will concentrate on extending the methodology to real episodes, as opposed to controlled field experiments, which are typically conducted at much smaller scales. The challenges typically presented by large-scale dispersion problems are the lack of concentration data and meteorological observations, and the need of sophisticated, and computationally expensive, numerical transport and dispersion models. In this respect, the good computational efficiency of AES may be a valuable factor.

Acknowledgments

This material is partly based upon work supported by the National Science Foundation under Grant AGS 0849191 and by George Mason University Summer Research Funding.

References

Akima, H., 1978. A method of bivariate interpolation and smooth surface fitting for irregularly distributed data points. In: ACM Transactions on Mathematical Software, vol. 4. Association for Computing Machinery, Inc. 148–159.
Allen, C.T., Young, G.S., Haupt, S.E., 2007. Improving pollutant source characterization by better estimating wind direction with a genetic algorithm. Atmospheric Environment 41 (11), 2283–2289.
Arya, P.S., 1999. Air Pollution Meteorology and Dispersion. Oxford University Press.
Barad, M., 1958. Project Prairie Grass, a field program in diffusion. Tech. Rep. Geophysical Research Paper, No. 59, Report AFCRC-TR-58-235, Air Force Cambridge Research Center, 218pp.
Cervone, G., Franzese, P., 2010. Monte Carlo source detection of atmospheric emissions and error functions analysis. Computers & Geosciences 36 (7), 902–909.
Cervone, G., Franzese, P., Keesee, A.P., 2010. Algorithm quasi-optimal (AQ) learning. WIREs: Computational Statistics 2 (2), 218–236.
Chang, J.C., Franzese, P., Chayantrakom, K., Hanna, S.R., 2003. Evaluations of CALPUFF, HPAC, and VLSTRACK with two mesoscale field datasets. Journal of Applied Meteorology 42 (4), 453–466.
Chang, J.C., Hanna, S.R., 2004. Air quality model performance evaluation. Meteorology and Atmospheric Physics 87, 167–196.
Chang, J.C., Hanna, S.R., Boybeyi, Z., Franzese, P., 2005. Use of Salt Lake City urban 2000 field data to evaluate the urban hazard prediction assessment capability (HPAC) dispersion model. Journal of Applied Meteorology 44 (4).
Chow, F., Kosović, B., Chan, T., 2006. Source inversion for contaminant plume dispersion in urban environments using building-resolving simulations. In: Proceedings of the 86th American Meteorological Society Annual Meeting, Atlanta, GA, pp. 12–22.
Chow, F., Kosović, B., Chan, T., 2008. Source inversion for contaminant plume dispersion in urban environments using building-resolving simulations. Journal of Applied Meteorology and Climatology 47, 1553–1572.
De Jong, K., 2008. Evolutionary computation: a unified approach. In: Proceedings of the 2008 GECCO Conference on Genetic and Evolutionary Computation. ACM, New York, NY, USA, pp. 2245–2258.
Delle Monache, L., Lundquist, J., Kosović, B., Johannesson, G., Dyer, K., Aines, R., Chow, F., Belles, R., Hanley, W., Larsen, S., Loosmore, G., Nitao, J., Sugiyama, G., Vogt, P., 2008. Bayesian inference and Markov Chain Monte Carlo sampling to reconstruct a contaminant source on a continental scale. Journal of Applied Meteorology and Climatology 47, 2600–2613.
Hanna, S.R., Chang, J.C., Strimaitis, G.D., 1990. Uncertainties in source emission rate estimates using dispersion models. Atmospheric Environment 24A (12), 2971–2980.
Hanna, S.R., Chang, J.C., Strimaitis, G.D., 1993. Hazardous gas model evaluation with field observations. Atmospheric Environment 27A, 2265–2285.
Haupt, S.E., Dec. 2005. A demonstration of coupled receptor/dispersion modeling with a genetic algorithm. Atmospheric Environment 39 (37), 7181–7189.

Haupt, S.E., Young, G.S., Allen, C.T., 2007. A genetic algorithm method to assimilate sensor data for a toxic contaminant release. Journal of Computers 2 (6), 85–93.

Holland, J., 1975. Adaptation in Natural and Artificial Systems. The MIT Press, Cambridge, MA.

Johannesson, G., Chow, F., Glascoe, L., Glaser, R., Hankley, W., Kosović, B., Krnajajic, M., Larsen, S., Lundquist, J.K., Mirin, A., Nitao, J., Sugiyama, G., 2005. Sequential Monte-Carlo based framework for dynamic data-driven event reconstruction for atmospheric release. In: Proceedings of the American Statistical Association, ASA Section on Bayesian Statistical Science, Joint Statistical Meeting, Minneapolis, MN, vol. 1, Paper 441. pp. 73–80.

Johannesson, G., Hanley, B., Nitao, J., 2004. Dynamic bayesian models via monte carlo - an introduction with examples -. Tech. Rep. UCRL-TR-207173, Lawrence Livermore National Laboratory.

Long, K.J., Haupt, S.E., Young, G.S., 2010. Assessing sensitivity of source term estimation. Atmospheric Environment 44 (12).

Lundquist, J.K., Kosović, B., Belles, R., 2005. Synthetic event reconstruction experiments for defining sensor network characteristics. Tech. Rep. UCRL-TR-217762, Lawrence Livermore National Laboratory.

Neumann, S., Glascoe, L., Kosović, B., Dyer, K., Hanley, W., Nitao, J., Gordon, R., 2006. Event reconstruction for atmospheric releases employing urban puff model UDM with stochastic inversion methodology. In: Proceedings for the American Meterological Society, Annual meeting, Atlanta, GA. No. J4.6.

Pasquill, F., Smith, F., 1983. Atmospheric Diffusion. Ellis Horwood.

Rao, S.K., 2007. Source estimation methods for atmospheric dispersion. Atmospheric Environment 41 (33), 6964–6973.

Rechenberg, I., 1971. Evolutionstrategie - optimierung technischer systeme nach prinzipien der biologischen evolution. Ph.D. thesis, Technical University of Berlin.

Russell, S., Norvig, P., 1995. Artificial Intelligence, A Modern Approach. Prentice Hall, Upper Saddle River, New Jersey.

Schwefel, H.-P., 1974. Numerische optimierung von computer-modellen. Ph.D. thesis, Technical University of Berlin.

Senocak, I., Hengartner, N., Short, M., Daniel, W., 2008. Stochastic event reconstruction of atmospheric contaminant dispersion using Bayesian inference. Atmospheric Environment 42 (33), 7718–7727.

Sohn, M., Reynolds, P., Singh, N., Gadgil, A.J., 2002. Rapidly locating and characterizing pollutant releases in buildings. Journal of Air and Waste Management 52, 1422–1432.