Comparison of Inverse Methods for Reconstructing the Release History of a Groundwater Contamination Source

Roseanna M. Neupauer[1], Brian Borchers[2], and John L. Wilson[1]

[1]Dept. of Earth and Environ. Sci., New Mexico Institute of Mining and Technology, Socorro, NM, 87801 USA
[2]Dept. of Mathematics, New Mexico Institute of Mining and Technology, Socorro, NM, 87801 USA

Accepted for publication in *Water Resources Research*

ABSTRACT

Inverse methods can be used to reconstruct the release history of a known source of groundwater contamination from concentration data describing the present-day spatial distribution of the contaminant plume. Using hypothetical release history functions and contaminant plumes, we evaluate the relative effectiveness of two proposed inverse methods, Tikhonov regularization (TR) and minimum relative entropy (MRE) inversion, in reconstructing the release history of a conservative contaminant in a one-dimensional domain [*Skaggs and Kabala*, 1994; *Woodbury and Ulrych*, 1996]. We also address issues of reproducibility of the solution and the appropriateness of models for simulating random measurement error. The results show that if error-free plume concentration data is available, both methods perform well in reconstructing a smooth source history function. With error-free data, the MRE method is more robust than TR in reconstructing a non-smooth source history function; however, the TR method is more robust if the data contain measurement error. Two error models were evaluated in this study, and we found that the particular error model does not affect the reliability of the solutions. The results for the TR method have somewhat greater reproducibility because, in some cases, its input parameters are less subjective than those of the MRE method; however, the MRE solution can identify regions where the data give little or no information about the source history function, while the TR solution cannot.

1 Introduction

Groundwater contamination is a widespread problem that can affect the utility of an aquifer. To minimize the deleterious effects of this contamination, the aquifer must be remediated to acceptable levels. Groundwater remediation is expensive, and the costs should be distributed among the parties responsible for the contamination. If a known contamination source is located at a facility that has changed ownership in the past, the release history of the source must be determined to identify the responsible parties. Often the present spatial distribution of the contaminant concentration is the only information available. These concentration data can be used in an inverse model, with reasonable estimates of the transport parameters, to reconstruct the release history from the contamination source.

Liu and Ball [1999] classify source history reconstruction problems into two broad categories: full-estimation problems and function-fitting problems. Full-estimation problems reproduce the time-varying source history at a specific location; while function-fitting problems assume a functional form of the source history and estimate the parameter values describing the function. The source history reconstruction problem has been addressed as a function-fitting problem by *Gorelick et al.* [1983], who used linear programming; and by *Wagner* [1992], who used a maximum likelihood method. The problem has been addressed using full-estimation methods including Tikhonov regularization [*Skaggs and Kabala* 1994, 1998; *Liu and Ball* 1999], the method of quasi-reversibility [*Skaggs and Kabala* 1995]; minimum relative entropy inversion [*Woodbury and Ulrych*, 1996, 1998a; *Woodbury et al.*, 1998]; and a geostatistical approach [*Snodgrass and Kitanidis*, 1997].

Two full-estimation methods, Tikhonov regularization (TR) and minimum relative entropy inversion (MRE), have received much interest in recent literature. *Woodbury and Ulrych* [1996, 1998b] and *Kabala and Skaggs* [1998] presented several arguments concerning the relative merits of MRE and TR, respectively. The relative merits of the two methods were not properly assessed because the methods were not evaluated with the same test cases. *Woodbury and Ulrych* [1996] and *Skaggs and Kabala* [1994] both reconstructed the same source history function (shown in Figure 1a of this paper and in Eq. (25) in *Skaggs and Kabala* [1994]); however, they used different quantities of data points and different data sets in their evaluations of the respective methods. The purpose of this paper is to provide an unbiased comparison of Tikhonov regularization and minimum relative entropy inversion for solving the hypothetical source history reconstruction problem as presented by *Skaggs and Kabala* [1994] and *Woodbury and Ulrych* [1996]. For this comparison, we prescribe a source history at a point source of contamination at a known location in a one-dimensional flow field. The present spatial distribution of the contaminant concentration that results from this source release is sampled, and these data are used in the inverse methods to reconstruct the temporal distribution of the contaminant concentration at the source. We assume that the contaminant is conservative, that transport can be modeled by the advection dispersion equation, and that the transport parameters are known.

2 Inverse Methods

Tikhonov regularization and minimum relative entropy inversion are two full-estimation methods of solving inverse problems. Tikhonov regularization [*Tikhonov and Arsenin*, 1977] is a regularized least-square method that replaces the ill-posed inverse problem with a well-posed minimization problem, given by

$$\min \left[||\mathbf{y} - \mathbf{G}\mathbf{s}||^2 + \alpha^2 ||\mathbf{L}\mathbf{s}||^2 \right] \ , \tag{1}$$

where $\mathbf{y}$ is the vector of plume data, $\mathbf{y} = [y(x_1), y(x_2), \ldots, y(x_{N_y})]$; x_n is the n^{th} sample location; $\mathbf{s}$ is the solution vector describing the source concentration over time, $\mathbf{s} = [s(t_1), s(t_2), \ldots, s(t_{N_g})]$; t_j is the j^{th} source input time; $\mathbf{G}$ is a matrix of values of the weighted kernel function describing the contaminant transport, α is the regularization weight, $\mathbf{L}$ is the regularization operator matrix, and $||\cdot||$ denotes the Euclidean norm. The kernel function is the solution of the advection dispersion equation for an instantaneous point source. For the problem considered here, the elements of the $\mathbf{G}$ matrix are given by

$$G_{n,j} = \frac{x_n a}{\sqrt{4\pi D t_j^3}} \exp\left[-\frac{(x_n - v t_j)^2}{4 D t_j} \right] \ , \tag{2}$$

where a is a weight based on the temporal discretization, D is the dispersion coefficient, and v is the groundwater velocity. The regularization operator, $\mathbf{L}$, is a k^{th}-derivative operator, where k is the order of regularization. For zero-order regularization, $\mathbf{L}$ is the identity matrix. For $\alpha = 0$, the TR solution would be the least-squares solution; however, since the problem is ill-posed, the least-squares solution is unstable. The regularization term in (1) stabilizes the problem. The TR method is a trade-off between matching the data (minimizing the first term on the right-hand side of (1)) and stabilizing the problem (minimizing the k^{th}-derivative of the solution). The degree of the trade-off depends on the value of the regularization weight, α; several methods are available for selecting the optimal value of α. The reader is referred to *Skaggs and Kabala* [1994] for more information on using the TR method for the source history reconstruction problem. To run the TR inversions, we used CONTIN [*Provencher*, 1982a, b], a general purpose computer program that performs Tikhonov regularization. CONTIN, which was also used by *Skaggs and Kabala* [1994], employs the F-test criteria [*Provencher*, 1982a; *Obenchain*, 1977] to select the optimal value of the regularization weight, α. Unless otherwise noted, we selected second-order regularization, which minimizes the second derivative of the source history function, or maximizes its smoothness. *Skaggs and Kabala* [1994] also selected second-order regularization.

Minimum relative entropy inversion treats each element of the solution (source history) vector, $\mathbf{s}$, as a random variable. Using prior information and the measured data, the MRE method obtains a multivariate probability density function (PDF) for the random variables; the mean of this PDF is the solution to the inverse problem. Information on the upper bounds, lower bounds, and expected values of the source concentration is used to create a prior PDF, which is a truncated exponential distribution. The final posterior distribution is chosen so that the relative entropy between the prior and posterior distributions is minimized, subject to the data constraint that the posterior mean solution matches the measured data to within a tolerance related to the standard deviation of the random measurement error. The functional form of this data constraint depends on the model used to represent the random measurement error, and the standard deviation of the measurement error is a parameter in the data constraint. The reader is referred to *Woodbury and Ulrych* [1993, 1996] for more information on the MRE method. Following their approach, we wrote a MATLAB routine to implement the MRE method. Details of the implementation can be found in *Neupauer* [1999]. Unless otherwise noted, we used a lower bound of 0.0, an upper bound of 1.1, and a prior expected value of 0.8 for all elements of the source history solution vector $\mathbf{s}$. All units are dimensionless.

3 Comparison of the Two Inverse Methods

Woodbury and Ulrych [1996] raised several concerns with the TR approach of *Skaggs and Kabala* [1994]; these issues were further discussed in *Kabala and Skaggs* [1998] and *Woodbury and Ulrych* [1998b]. We have grouped the issues into four categories: a proper comparison of the two methods, the appropriate error model, uniqueness of the solutions, and confidence intervals on the solution. In this section, we address each of these issues. For all simulations shown here, the parameters values were $v = 1.0$ and $D = 1.0$, and the source of contamination was at $x = 0$. We used two different source history functions–the three-peaked smooth function used by *Skaggs and Kabala* [1994], *Woodbury and Ulrych* [1996] and others (hereafter called Source A, shown in Figure 1a) and a step function with $C_{in}(t) = 1.0$ for $125 < t < 225$, where $C_{in}(t)$ is the source concentration (hereafter called Source B, shown in Figure 1b). The reconstructed source history (solution) vector has 100 elements uniformly spaced between $t = 0.01$ and $t = 250$. Because of algorithmic constraints with CONTIN, the solution vector cannot contain $t = 0$ or $t = t_s$, where t_s is the sampling time ($t_s = 300$ for this problem). All units are dimensionless.

The plumes generated by these source history functions were sampled at $t = t_s = 300$ at 25 locations ($\mathbf{x} = [0.01, 25.05, 50, 60, 70, 80, \ldots, 230, 240, 250, 275, 300]$). For some simulations presented here, we added random measurement error to the sampled concentrations using either a multiplicative error model or an additive error model. The multiplicative error model is

$$C^*(x_n, t) = C(x_n, t) + \epsilon_m \delta_n C(x_n, t) \ , \tag{3}$$

where $C^*(x_n, t)$ is the measured concentration at location x_n at time t, $C(x_n, t)$ is the true concentration, x_n is the location of the n^{th} sample, ϵ_m is the standard deviation of the random error (error level), and δ_n is the n^{th} random deviate (standard normal). The additive error model is

$$C^*(x_n, t) = C(x_n, t) + \epsilon_a \delta_n \ , \tag{4}$$

where ϵ_a is the standard deviation of the random error. We used $\epsilon_m = 0.05$ for the multiplicative error model and $\epsilon_a = 0.009$ for the additive error model. The true plume and sampled data are shown in Figure 1c,d.

3.1 Proper Comparison of the Methods

Skaggs and Kabala [1994] showed the results of several examples using the TR method to reproduce Source A. *Woodbury and Ulrych* [1996] presented the results of several similar, although slightly different, examples using the MRE method. The two sets of examples differed in the number of data points (sample locations) used in the inversion, and in the error models used to generate hypothetical random measurement error. Because of the slight differences in the examples, a proper comparison of the two methods was not made. *Kabala and Skaggs* [1998] suggested that to properly compare the methods, the number of data points and the error model must be the same for each method, and also suggested that the reconstruction of other source history functions be evaluated. To address these issues, we compared the MRE and TR methods for reproducing Sources A and B using an equal number of data points (25) and the same error model (multiplicative).

Figures 2 and 3 show the results for Source A and Source B, respectively. For these and all remaining figures, the horizontal axis has units of dimensionless time, and the vertical axis has units of relative concentration. Both Figures 2 and 3 show results for error-free data and for data with random measurement error (error level of 0.05; the data are shown as triangles in Figure 1). For Source A (Figure 2), the MRE and TR methods perform very well with both data sets. For Source B (Figure 3), MRE performs better than TR when the error-free data set is used. Recall that with second-order regularization, the TR method maximizes the smoothness of the solution; since the true source history is not smooth, the TR solution oscillates around the plateau. When the data containing measurement error are used, both MRE and TR appear to perform equally well. We also tested the two methods using 60 data points and obtained similar results which are not shown here.

3.2 Error Models

Skaggs and Kabala [1994] and *Woodbury and Ulrych* [1996] used different error models to generate random measurement error on the sampled data. *Skaggs and Kabala* [1994] used the multiplicative error model in (3), and *Woodbury and Ulrych* [1996] used the additive error model in (4). Both models are physically plausible [*Kabala and Skaggs*, 1998; *Woodbury and Ulrych*, 1996, 1998b]. We cannot justify that either error model is more appropriate than the other, but we have found that the results of the inverse methods are similar for both error models. The results are shown in Figure 4 where random measurement error was simulated using the multiplicative model ($\epsilon_m = 0.05$), and Figure 5 where the additive error model ($\epsilon_a = 0.009$) was used. The data used in these simulations are shown in Figure 1c. With these values of ϵ_m and ϵ_a, the norms of the noise vectors are equal in both models, and the inversion results from the two data sets can be compared. The noise vector is the vector of the difference between the measured and exact concentrations. By equating the norms of the noise vectors, we obtain $\epsilon_m ||C||/\sqrt{N_y} = \epsilon_a$. Using the exact data shown in Figure 1c, $||C||/\sqrt{N_y} = 0.1716$. With this result and $\epsilon_m = 0.05$, the appropriate value for the additive noise level is $\epsilon_a = 0.009$. Note that for all of the results presented here, we used only one set of random numbers δ_n, $n = 1, 2, \ldots, N_y$. Similar results were obtained for other (> 50) sets of random numbers.

Figures 4 and 5 show that both inverse methods produce reasonable results, regardless of the error model used. Also, with the MRE method, similar results are obtained when using the correct data constraint (i.e., Figure 4b, where the random measurement error and the data constraint both follow the multiplicative error model, and Figure 5c, where the random measurement error and the data constraint both follow the additive model) and the incorrect data constraint (i.e., Figure 4c, where the random measurement error follows the multiplicative error model and the data constraint assumes that the additive error model is followed, and Figure 5b, where the random measurement error follows the additive model and the data constraint assumes that the multiplicative model is followed). The solution is insensitive to the particular error model, provided consistent error levels, ϵ_m and ϵ_a, are used.

In a field situation, the true error level is not known, but the MRE method requires an error level to be specified in the data constraint. We found that the MRE method is sensitive to the specified error level. Figure 6 shows the effects of using the wrong error level in the MRE data constraint. The correct error level (Figure 6b) produces reasonable results, while the under- and over-estimated error levels (Figures 6a and c) produce less accurate results. In fact, the solution shown in Figure 6a does not reproduce the data. *Woodbury and Ulrych* [1996] suggested pre-filtering the data to eliminate the random error; however, in following their described method, we obtained very different results. The TR method can also be made sensitive to the measurement error if the regularization parameter is chosen using a method that depends on the noise level [*Groetsch*, 1994]. With either method, an experienced user might find the flexibility of a user-adjusted error level useful, while a naive user would not.

3.3 Reproducibility of the Solution

The source history reconstruction problem is ill-posed, and therefore, a unique solution cannot be obtained. In the context of inverse problems, a solution would be "unique" if it were the only solution that could reproduce the data. Any inverse problem is non-unique. *Woodbury and Ulrych* [1996, 1998b] state that, given the set of information available, the MRE method uses all available information to produce a "unique" posterior PDF from which the solution is obtained. In this context, "unique" refers to the reproducibility of the inverse model result. In other words, given a set of measured data and prior information about the source history, the MRE method will always produce the same posterior distribution. To avoid confusion with the term "uniqueness" in the context of the ill-posedness of the problem, we will describe the uniqueness of the PDF using the term "reproducibility". In other words, the solution has a high degree of reproducibility if two modelers, given the same set of available information, obtain the same solution. To our knowledge, the distinction between uniqueness and reproducibility of the solution has not been made before.

Woodbury and Ulrych [1998b] argue that Tikhonov regularization involves two subjective factors, the regularization order (k) and the regularization weight (α), which can decrease the degree of reproducibility. Several methods are available for selecting the optimal regularization weight. The choice of methods can be subjective, but once a method is chosen, a unique value of the regularization weight is obtained. We evaluated the sensitivity of the TR solution to two methods of selecting the optimal regularization weight, the F-test method [*Provencher*, 1982a; *Obenchain*, 1977] and generalized cross-validation [*Wahba*, 1977], and found essentially no difference in the TR solutions [*Neupauer*, 1999]. We also evaluated the sensitivity of the TR results to the order of regularization. The results are shown in Figure 7 for zero-, first-, and second-order regularization using error-free data to reproduce Source B. The results of first- and second-order regularization are essentially indistinguishable; while the zero-order-regularization result is oscillatory. Along the plateau, the zero-order solution vector alternates between values above the true solution and values below the true solution; these large fluctuations are unrealistic. If the temporal spacing of the solution vector were reduced, zero-order regularization would produce a similar alternating pattern with a higher frequency. Since the frequency of the oscillations depends on the temporal discretization of the source history, the oscillating behavior is clearly a numerical artifact and not supported by the data; therefore, zero-order regularization is obviously unreliable in this case and can be so identified by a user in application. In comparing the results of first- and second-order regularization, the solution is insensitive to the regularization order. We found this same pattern for Source A. Thus, the TR solution is relatively insensitive to the two subjective factors (k and α), providing strong reproducibility of the TR results. We also evaluated results using higher-order regularization (up to order five). The results for Source A are relatively indistinguishable for regularization orders of one or higher; and for Source B, the results become less accurate as the regularization order increases beyond two.

With the MRE method, the prior PDF is developed using prior information including lower and upper bounds on the solution, and its prior expected value. For the source history reconstruction problem, the lower bound is typically zero, since concentration is non-negative and the release concentration is likely to be zero during part of the recovery interval. A reasonable estimate for the upper bound is the solubility limit of the contaminant; however, the actual source concentration might be far below this level.

We evaluated the sensitivity of the MRE solution to the upper bound, and found that, in some cases, the results are quite sensitive. Figure 8 shows the results of the MRE method for Source B with three different values of the upper bound, indicating that for this source, the solution becomes less accurate as the upper bound increases. Suppose we use the solubility as the upper bound and that the maximum release concentration in Source B were approximately equal to the solubility limit. The resulting MRE solution is shown in Figure 8a. If the maximum release concentration were one half of the solubility limit, the resulting MRE solution would be the solution shown in Figure 8c. Obviously, the solution is more accurate when the upper bound is near the maximum release concentration. The oscillations continue to magnify as the maximum release concentration decreases relative to the solubility limit. Thus, if the upper bound is not chosen wisely, it is possible to obtain a poor solution with the MRE method for Source B.

In contrast, the MRE solution for Source A was relatively insensitive to the value of the upper bound [*Neupauer*, 1999]. We also evaluated the sensitivity of the MRE method to the prior expected value, and found the solutions to be nearly indistinguishable [*Neupauer*, 1999].

3.4 Confidence Intervals

The MRE method produces a multivariate posterior PDF for the source concentration; and *Woodbury and Ulrych* [1996] show how probability levels of this posterior PDF can be used as confidence intervals. *Kabala and Skaggs* [1998] and *Woodbury and Ulrych* [1998b] debated the interpretation of confidence intervals defined by these probability levels. Since we are comparing the effectiveness of the two methods, we will not address the interpretation of the confidence intervals; however, some features of the confidence intervals affect the comparison. As noted by *Woodbury and Ulrych* [1998b], the confidence intervals for the TR solution are affected by the choice of the regularization weight, α. Unbiased confidence intervals can only be obtained when $\alpha = 0$, i.e., with no regularization. In general, the solution is regularized to improve stability, and the resulting confidence intervals are biased. With the MRE method, the confidence intervals are based on the posterior PDF, which depends on the measured data and the dimensionality of the solution vector. Assuming that the confidence intervals are correctly interpreted (i.e., Bayesian interpretation), the MRE confidence intervals are meaningful.

Woodbury and Ulrych [1998b] also explained that if the measured data contain no information to update the prior PDF, the posterior distribution will default to the prior distribution. In this case, the MRE solution will default to the prior expected value, with wide confidence intervals. With this feature, we can identify regions for which the data contain no information about the solution. In these regions, TR will produce a solution, and there is no way to determine if the solution is supported by the data. The MRE solution will also be equivalent to the prior expected value if the data confirm the prior. We can distinguish this case from the case where the data provide no information by evaluating the confidence intervals. If the data confirm the prior, the confidence intervals will be narrower than the spread of the prior distribution; and if the data provide no information, the confidence intervals will equal the spread of the prior distribution.

4 Discussion and Conclusions

Woodbury and Ulrych [1996, 1998b] and *Kabala and Skaggs* [1998] debated several issues about the relative merits of the TR and MRE methods for the source history reconstruction problem, including the appropriate error model, the reproducibility of the solution, the interpretation of confidence intervals, and a comparison of consistent test cases. In this paper, we addressed these four issues. We found that the results of both inverse methods are reasonable when either the multiplicative error model of *Skaggs and Kabala* [1995] or the additive error model of *Woodbury and Ulrych* [1996] is used; therefore, the solution does not depend on the particular model used to generate random measurement error.

Both TR and MRE require some subjective inputs. For Tikhonov regularization, the user must choose the order of regularization and the method for selecting the optimal regularization weight. We found that, for the cases we evaluated, similar results were obtained using first- and second-order regularization; however zero-order regularization produced unrealistic results because the solution oscillated at a frequency that depended on the temporal discretization. We also found that two methods for selecting the regularization weight (F-test and generalized cross-validation) produced nearly identical results, and once a method is chosen, a unique value of the regularization weight is obtained. Therefore, for the cases we evaluated, the TR method contains little subjectivity. The subjective inputs for the MRE method include the upper bounds and prior expected values of the source concentration. We assume that the lower bound is always zero, and therefore is not subjective. We found that the MRE solution is relatively insensitive to the prior expected value, sensitive to the upper bound for Source B, and insensitive to the upper bound for Source A. Based on these results, for the source history function we evaluated, the inputs to the MRE method are slightly more subjective that the inputs to the TR method. Subjectivity is neither good nor bad. An expert may properly demand more subjectivity, while the naive user deserves almost none.

Both methods produce confidence intervals for the solution. Because of regularization, the confidence intervals obtained with the TR method are biased, and therefore are not accurate. The probability levels obtained with the MRE method are meaningful if they are interpreted as "the probability of a parameter being in an interval that is conditional on the observed data in the current experiment" [*Woodbury and Ulrych*, 1998b]. These probability levels are not classical confidence intervals.

The results of the comparisons show that both TR and MRE are effective in reconstructing the smooth source history function (Source A), but since second-order Tikhonov regularization attempts to fit a smooth function, the TR method is less effective than MRE in reproducing the step source history function (Source B). The MRE method requires that an error level be specified in the data constraint. We found that if the noise level is known exactly, the MRE method performs as well as the TR method; however, if the noise level is underestimated, the MRE method does not perform as well. These findings show that the TR method is more robust when the data contain measurement error of unknown magnitude.

In summary, we have compared the relative effectiveness of the TR and MRE methods in reproducing two source history functions–a three-peaked Gaussian and an step function. For these input functions, the MRE method is more robust than the TR method if the step (non-smooth) source history function; while the TR method is more robust if the data contain measurement error of unknown magnitude. In general, for any input function, the MRE method has the advantage of identifying regions where a solution cannot be determined, because the posterior distribution defaults to the prior distribution in regions where the data give no information about the source history function. The TR method cannot make this distinction.

Acknowledgments

This research was supported in part by the Environmental Protection Agency's STAR Fellowship program under Fellowship No. U-915324-01-0. This work has not been subjected to the EPA's peer and administrative review and therefore may not necessarily reflect the views of the Agency and no official endorsement should be inferred. Todd Skaggs and Zbigniew Kabala provided their code and input files for the Tikhonov regularization method. Allan Woodbury provided valuable information about his implementation of the MRE method. Zbigniew Kabala and Allan Woodbury also reviewed the test results in *Neupauer* [1999] prior to our preparation of this paper, and Jose Capilla and two anonymous reviewers provided valuable comments on the manuscript.

References

Gorelick, S.M., B. Evans, and I. Remson, Identifying sources of groundwater pollution: an optimization approach, *Water Resour. Res.*, 19(3), 779–790, 1983.

Groetsch, C.W.,*The Theory of Tikhonov Regularization for Fredholm Integral Equations of the First Kind*, Pitman Publishing Inc., Marshfield, Mass., 1984.

Kabala, Z.J. and T.H. Skaggs, Comment on "Minimum relative entropy inversion: Theory and application to recovering the release history of a groundwater contaminant" by Allan D. Woodbury and Tadeusz J. Ulrych, *Water Resour. Res.*, 34(8), 2077–2079, 1998.

Liu, C. and W.P. Ball, Application of inverse methods to contaminant source identification from aquitard diffusion profiles at Dover AFB, Delaware, *Water Resour. Res.*, 35(7), 1975–1985, 1999.

Neupauer, R.M., *A comparison of two methods for recovering the release history of a groundwater contamination source*, M.S. Thesis, New Mexico Institute of Mining and Technology, Socorro, NM, 1999.

Obenchain, R.L., Classical F-tests and confidence regions for ridge regression, *Technometrics*, 19(4), pp. 429–439, 1977.

Provencher, S.W., A constrained regularization method for inverting data represented by linear algebraic or integral equation, *Computer Physics Communications*, 27, 213–227, 1982a.

Provencher, S.W., CONTIN: A general purpose constrained regularization program for inverting noisy linear algebraic and integral equations, *Computer Physics Communications*, 27, 229–242, 1982b.

Skaggs, T.H. and Z.J. Kabala, Recovering the release history of a groundwater contaminant, *Water Resour. Res.*, 30(1), 71–79, 1994.

Skaggs, T.H. and Z.J. Kabala, Recovering the history of a groundwater contaminant plume: Method of quasi-reversibility, *Water Resour. Res.*, 31(11), 2669–2673, 1995.

Skaggs, T.H. and Z.J. Kabala, Limitations in recovering the history of a groundwater contaminant plume, *J. Contam. Hydrol.*, 33, 347–359, 1998.

Snodgrass, M.R. and P.K. Kitanidis, A geostatistical approach to contaminant source identification, *Water Resour. Res.*, 33(4), 537–546, 1997.

Tikhonov, A.N., and V.Y. Arsenin, *Solution of Ill-Posed Problems*, Winston and Sons, New York, 1977.

Wagner, B.J., Simultaneous parameter estimation and contaminant source characterization for coupled groundwater flow and contaminant transport modeling, *Journal of Hydrology*, 135, 275–303, 1992.

Wahba, G., Practical approximate solutions to linear operator equations when the data are noisy, *SIAM Journal on Numerical Analysis*, 14, 651–667, 1977.

Woodbury, A.D., E. Sudicky, T.J. Ulrych, and R. Ludwig, Three-dimensional plume source reconstruction using minimum relative entropy inversion, *Journal of Contaminant Hydrology*, 32, 131–158, 1998.

Woodbury, A.D. and T.J. Ulrych, Minimum relative entropy: Forward probabilistic modeling, *Water Resour. Res.*, 29(8), 2847–2860, 1993.

Woodbury, A.D. and T.J. Ulrych, Minimum relative entropy inversion: Theory and application to recovering the release history of a groundwater contaminant, *Water Resour. Res.*, 32(9), 2671–2681, 1996.

Woodbury, A.D. and T.J. Ulrych, Minimum relative entropy and probabilistic inversion in groundwater hydrology, *Stochastic Hydrology and Hydraulics*, 12, 317–358, 1998a.

Woodbury, A.D. and T.J. Ulrych, Reply, *Water Resour. Res.*, 34(8), 2081–2084, 1998b.

Figure Captions

Figure 1: True release history for (a) Source A and (b) Source B. True plume and sampled data for (c) Source A (d) Source B.

Figure 2: Comparison of TR and MRE results with Source A. (a) TR method with error-free data. (b) MRE method with error-free data. (c) TR method with noisy data. (d) MRE method with noisy data. The multiplicative error model with $\epsilon_m = 0.05$ is used in (c) and (d).

Figure 3: Comparison of TR and MRE results with Source B. (a) TR method with error-free data. (b) MRE method with error-free data. (c) TR method with noisy data. (d) MRE method with noisy data. The multiplicative error model with $\epsilon_m = 0.05$ is used in (c) and (d).

Figure 4: Solutions of the two inverse methods using data generated using the multiplicative error model with $\epsilon_m = 0.05$. (a) TR results. (b) MRE results with multiplicative error model used in the data constraint ($\epsilon_m = 0.05$). (c) MRE results with additive error model used in the data constraint ($\epsilon_a = 0.009$).

Figure 5: Solutions of the two inverse methods using data generated using the additive error model with $\epsilon_a = 0.009$. (a) TR results. (b) MRE results with multiplicative error model used in the data constraint ($\epsilon_m = 0.05$). (c) MRE results with additive error model used in the data constraint ($\epsilon_a = 0.009$).

Figure 6: Solutions of the MRE method with different error levels used in the data constraint. The data were generated using the multiplicative error model with $\epsilon_m = 0.05$. MRE results with (a) $\epsilon_m = 0.01$, (b) $\epsilon_m = 0.05$, and (c) $\epsilon_m = 0.25$ in the data constraint.

Figure 7: Solutions of the TR method with different regularization orders using exact data. (a) zero-order regularization, (b) first-order regularization, and (c) second-order regularization.

Figure 8: Solutions of the MRE method with different upper bounds and no measurement error. Upper bounds of (a) 1.02, (b) 1.5, and (c) 2.0.

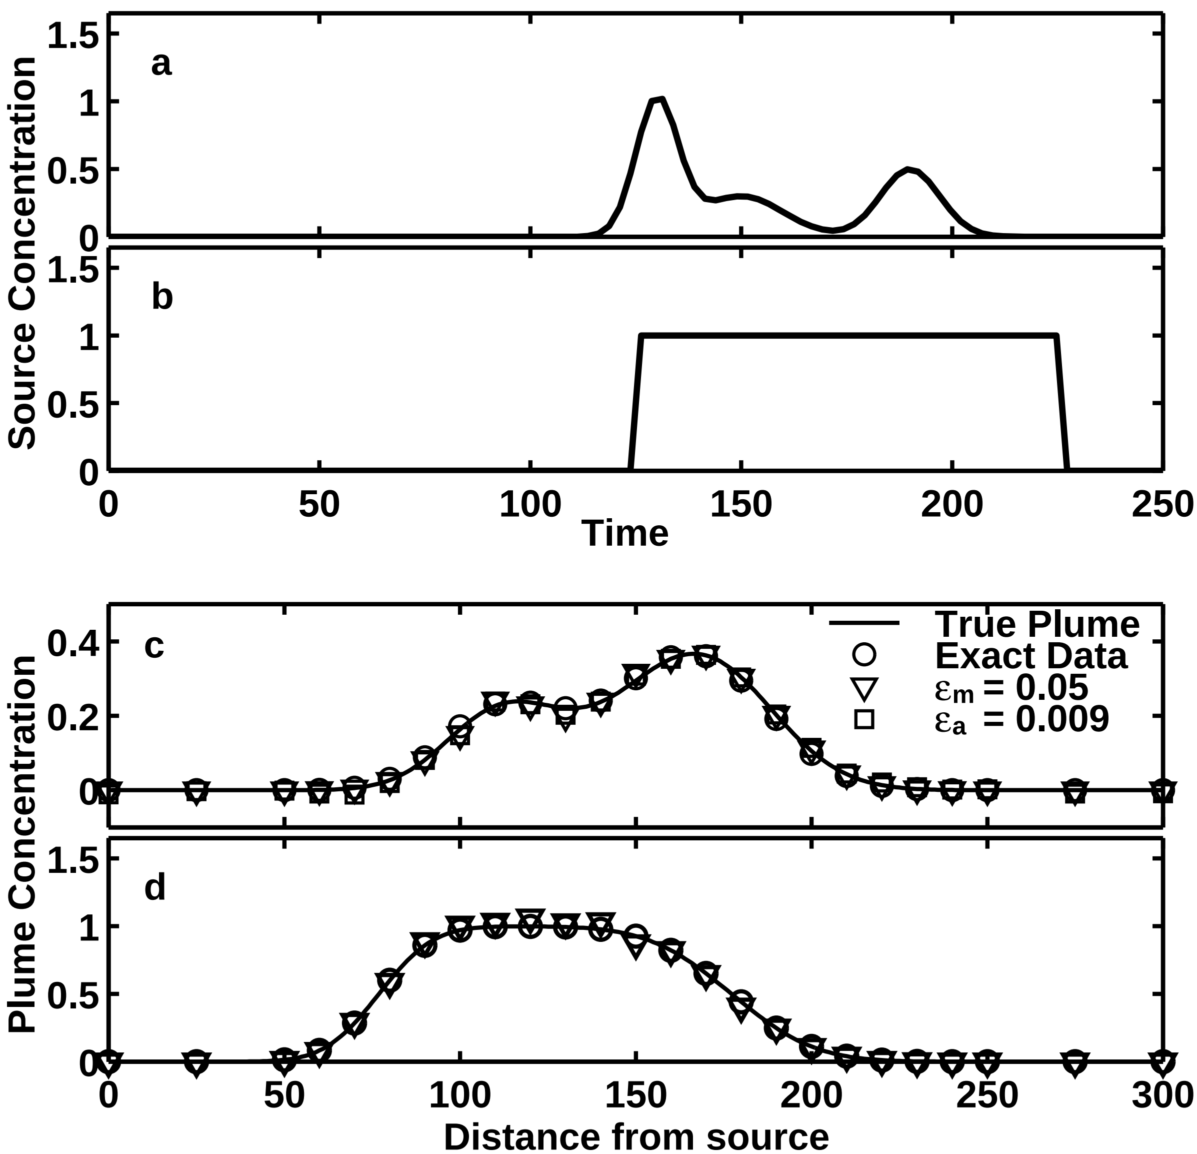

a
b
Source Concentration
Time
c
d
True Plume
Exact Data
ε_m = 0.05
ε_a = 0.009
Plume Concentration
Distance from source

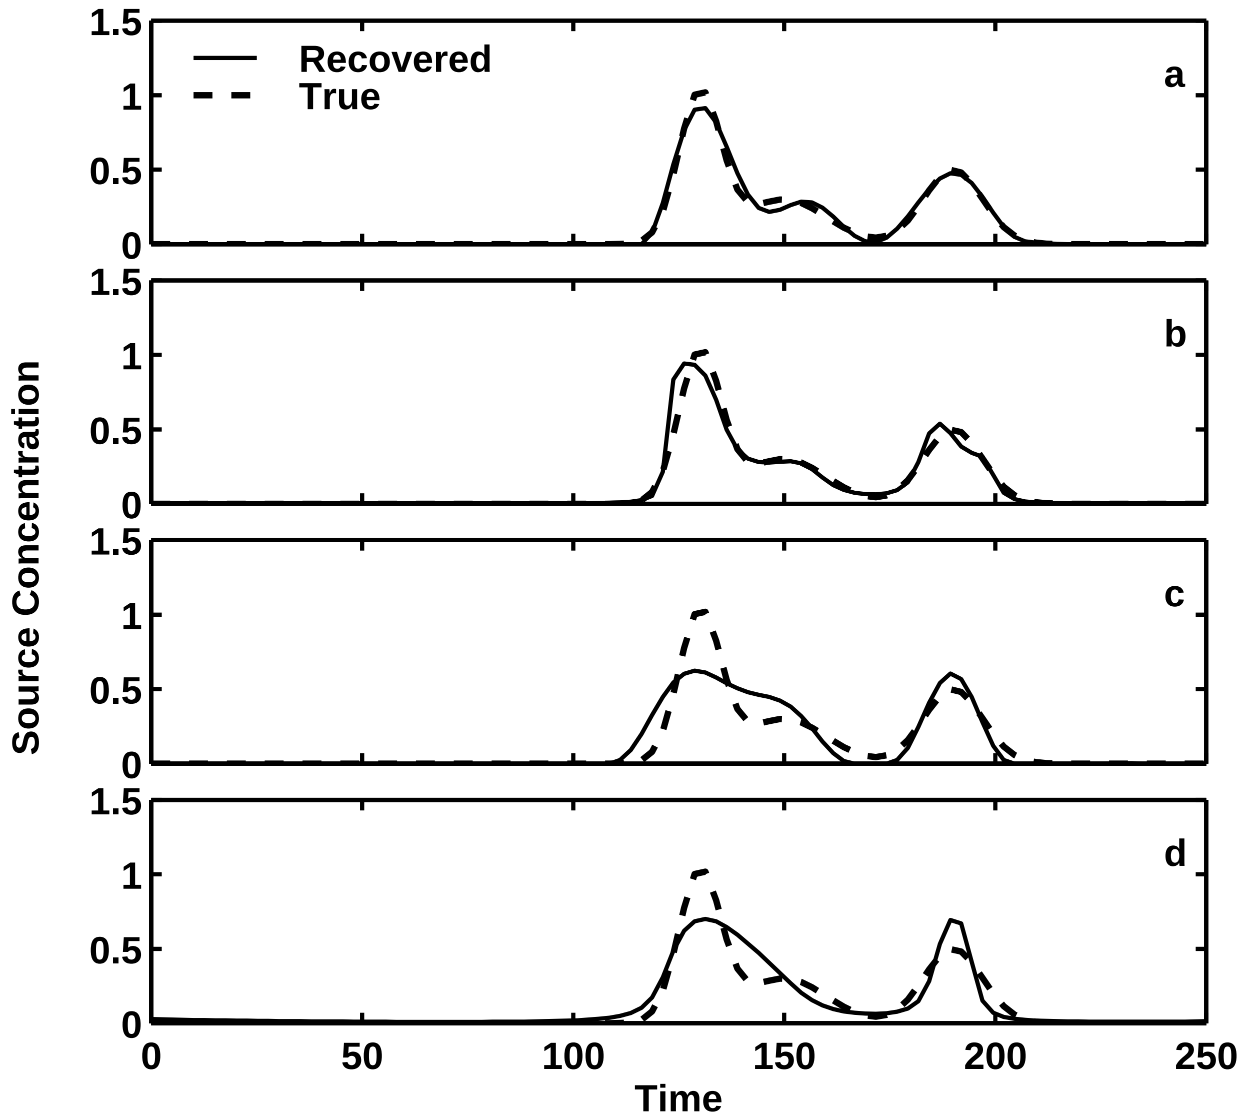

Recovered
True
a
b
c
d
1.5
1
0.5
0
0
50
100
150
200
250
Time
Source Concentration

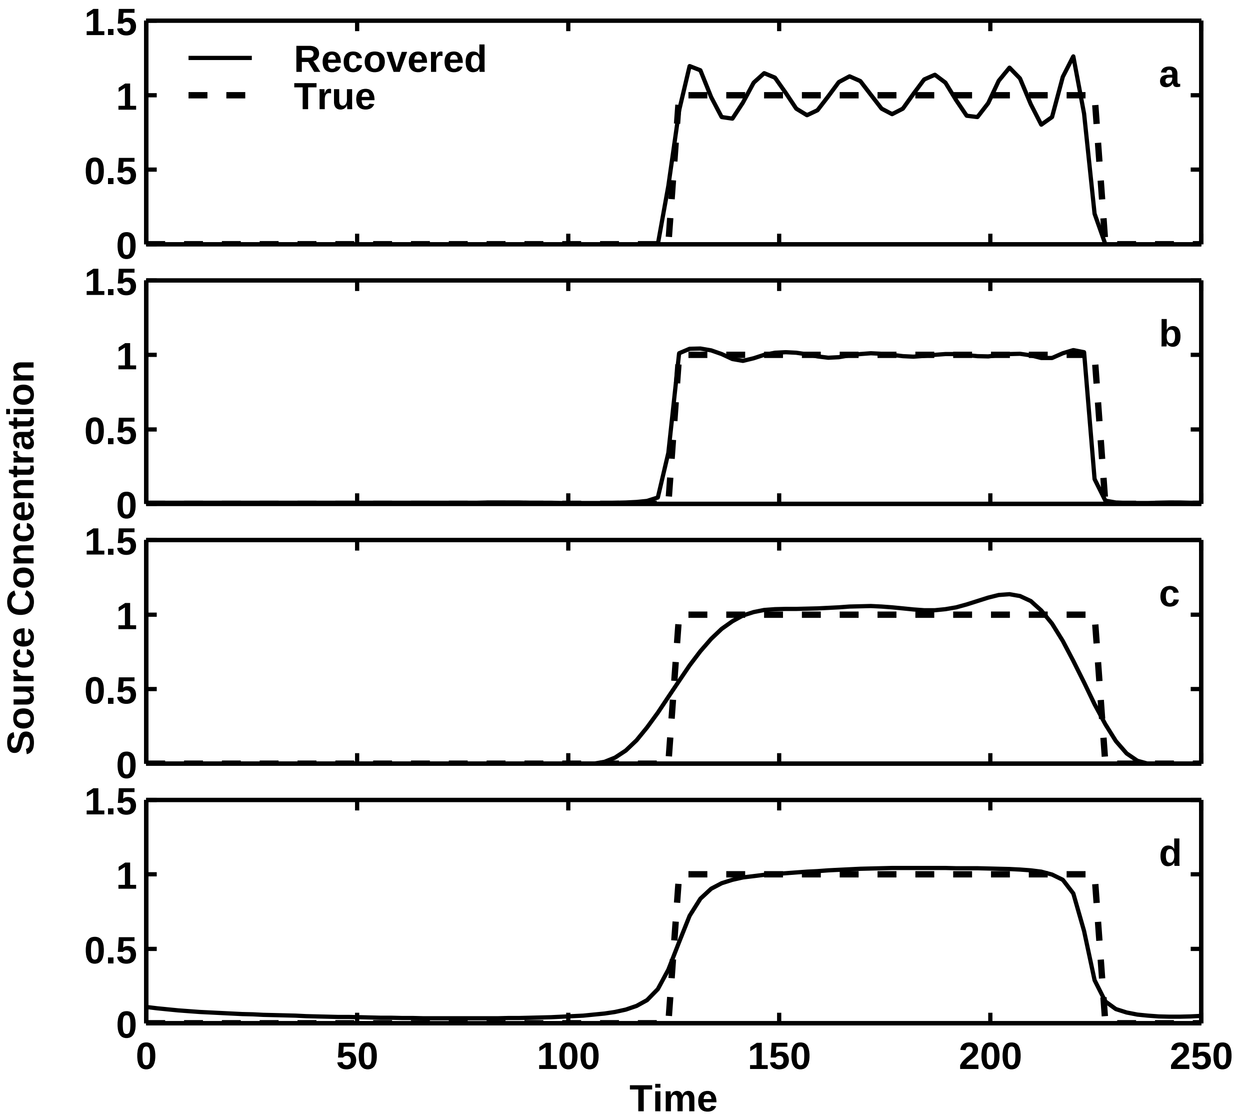

Recovered
True
a
b
c
d
1.5
1
0.5
0
0
50
100
150
200
250
Time
Source Concentration

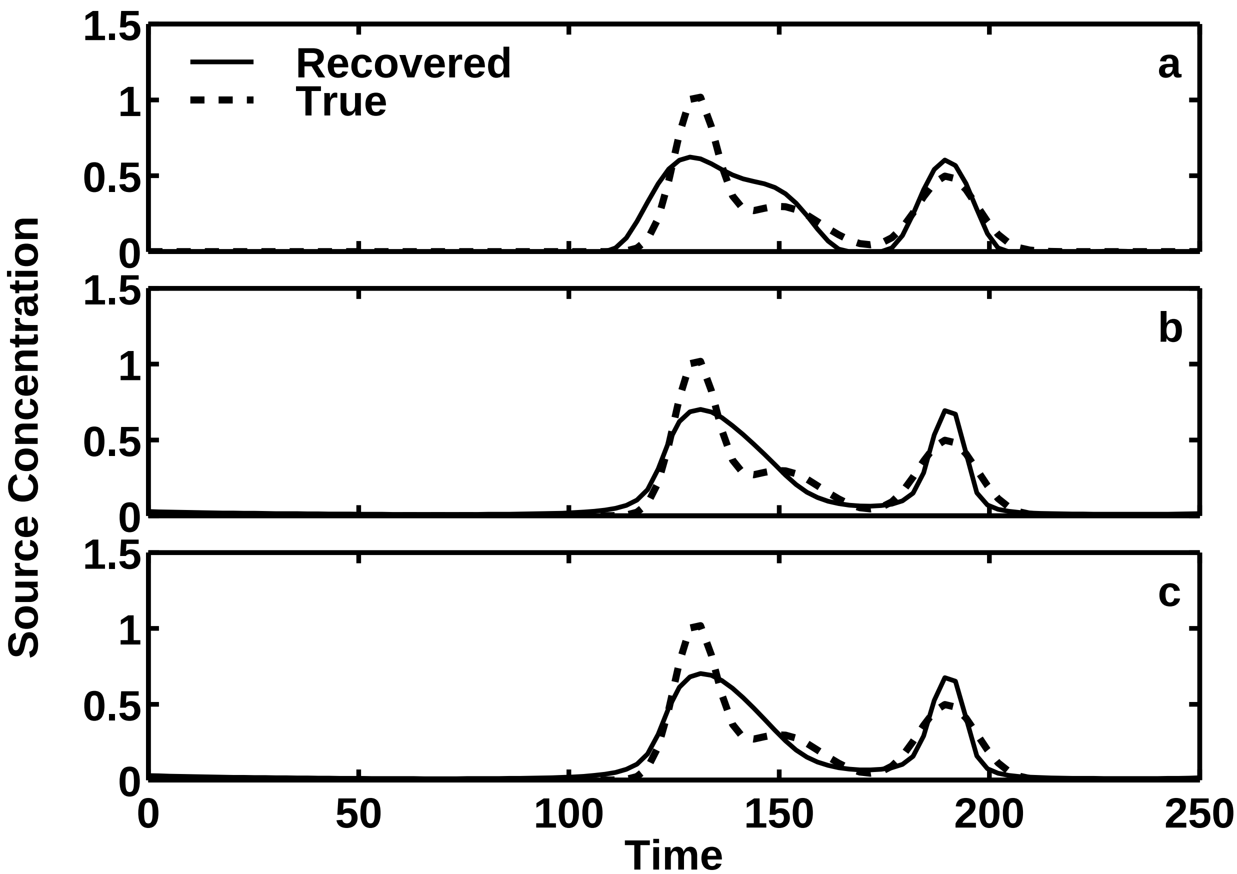

Recovered
True
a
b
c
1.5
1
0.5
0
0
50
100
150
200
250
Time
Source Concentration

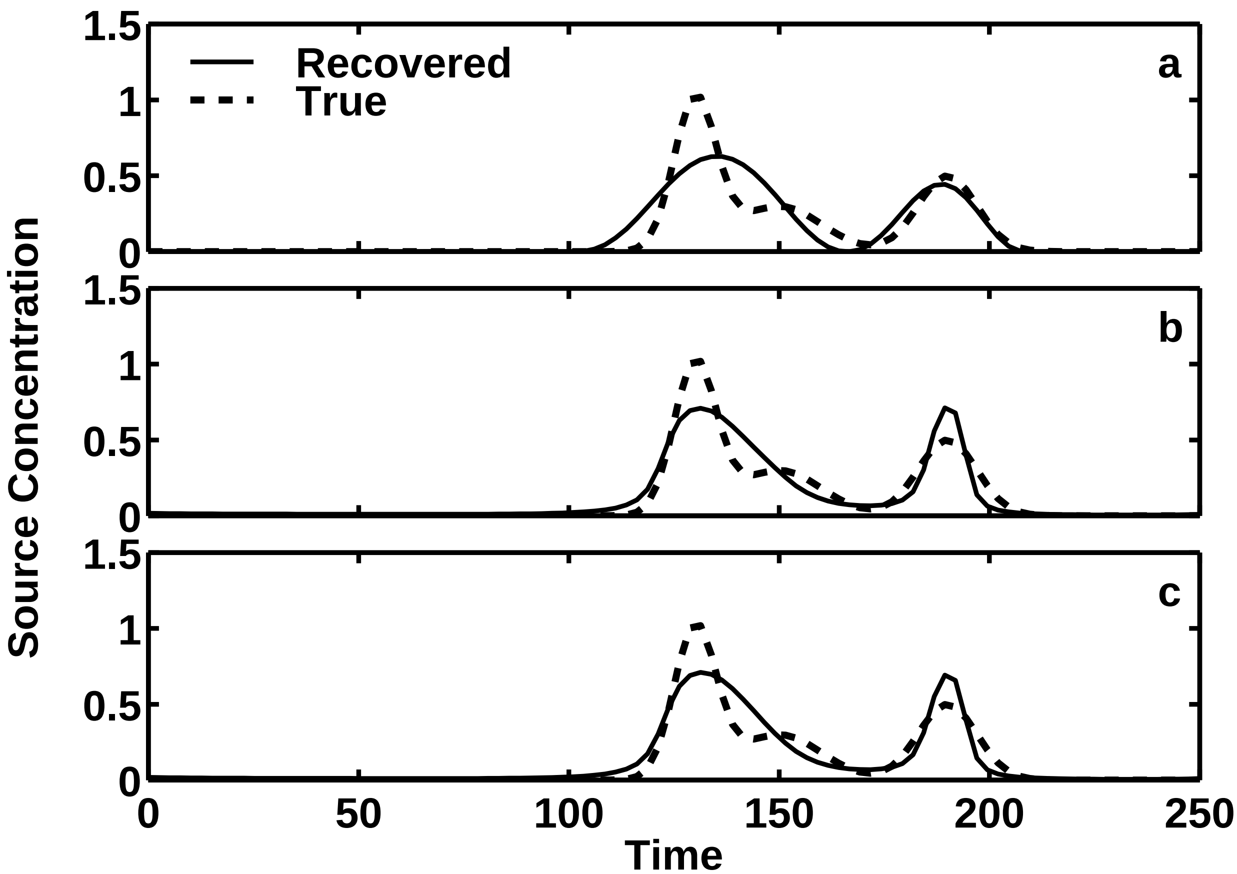
1.5
1
0.5
0
Recovered
True
a
b
c
Source Concentration
0
50
100
150
200
250
Time

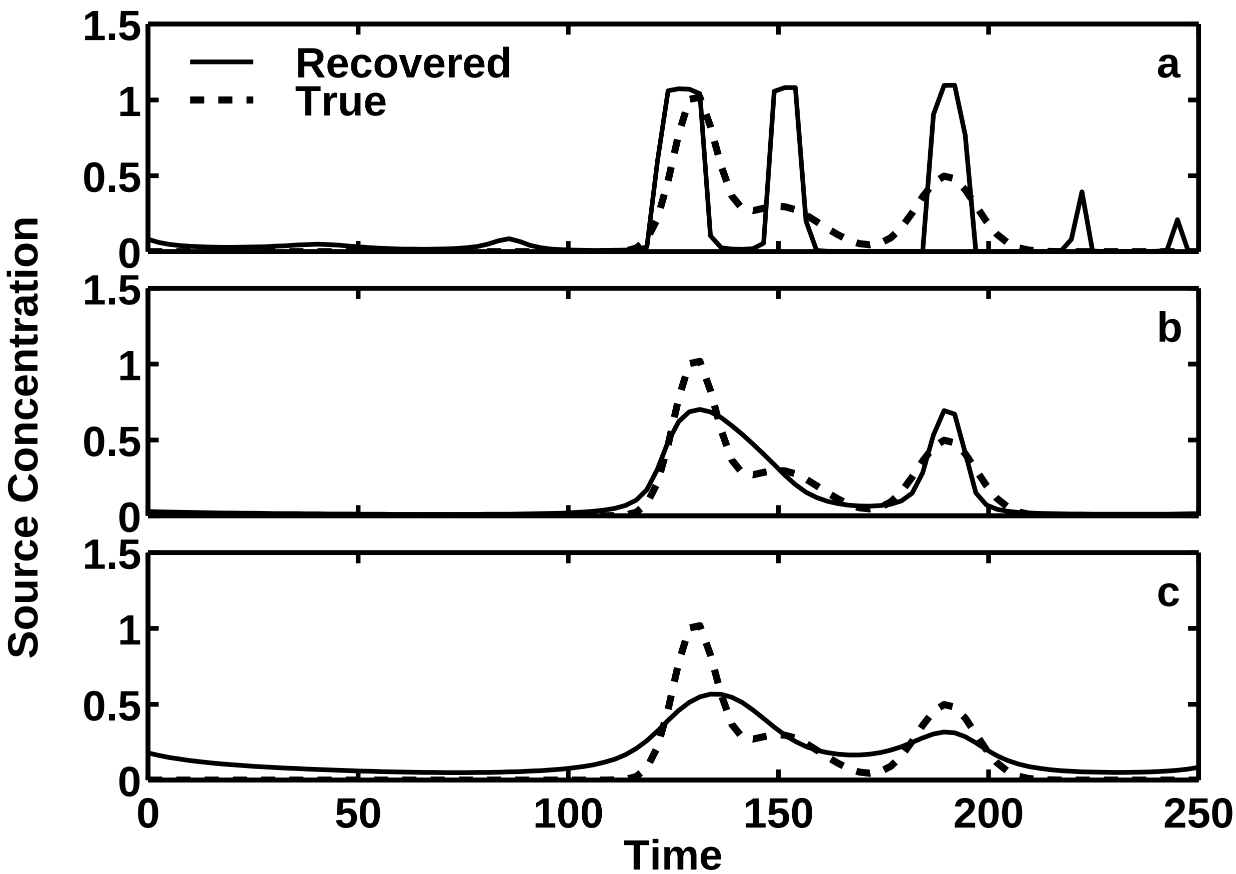

Source Concentration
Recovered
True
a
b
c
1.5
1
0.5
0
0
50
100
150
200
250
Time

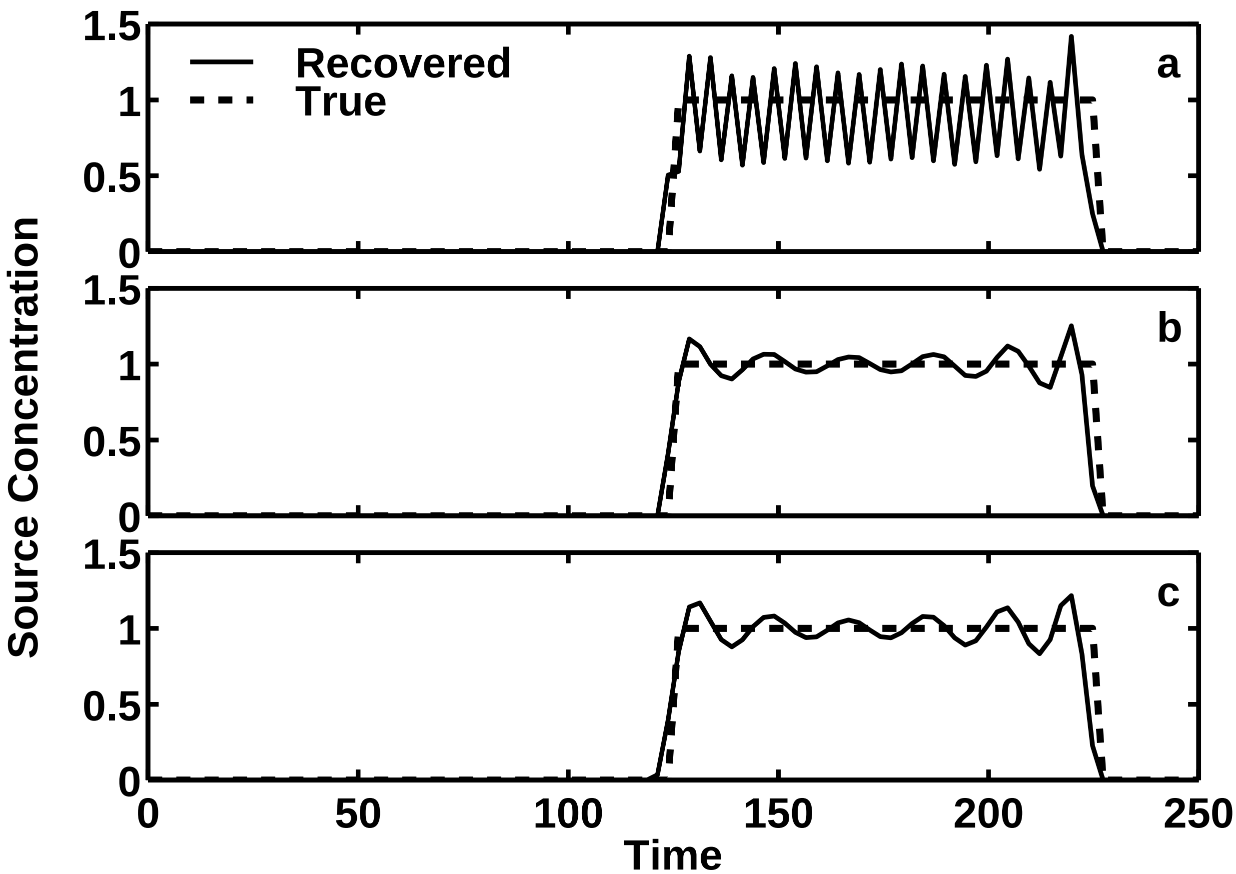

Source Concentration
Recovered
True
a
b
c
1.5
1
0.5
0
0
50
100
150
200
250
Time

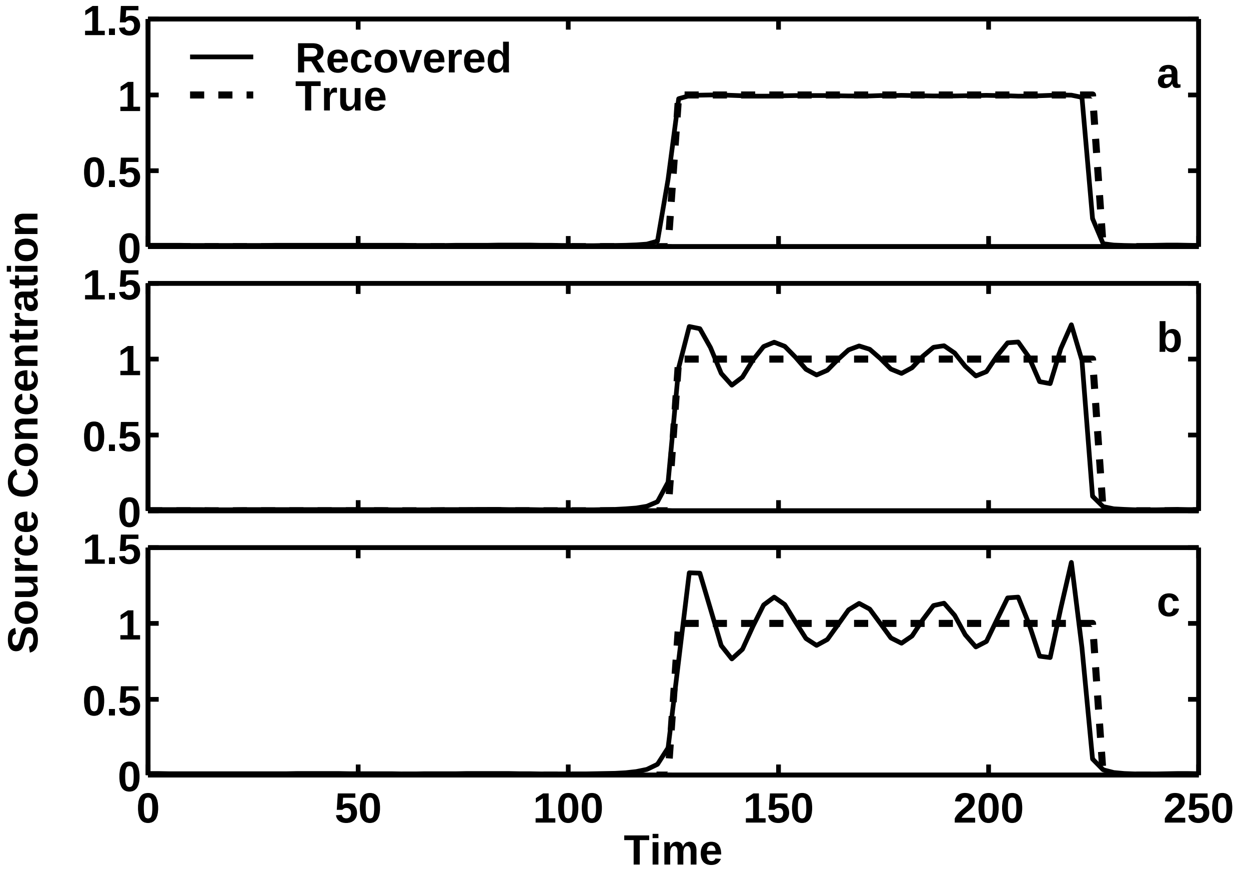

Recovered
True
a
b
c
1.5
1
0.5
0
0
50
100
150
200
250
Time
Source Concentration