

# Water Resources Research®

## RESEARCH ARTICLE

10.1029/2023WR036159

### Key Points:

- Unmixing sources and magnitudes of tracers from nested concentration measurements in a river basin is a convex optimisation problem
- We provide new software that solves this problem very quickly given a D8 raster and some point observations of concentration
- This method can be used to apportion pollution in a river basin and in any geochemistry study that uses a nested sampling design

### Correspondence to:

A. G. Lipp and R. Barnes,  
alexander.lipp@merton.ox.ac.uk;  
rijard.barnes@gmail.com

### Citation:

Barnes, R., & Lipp, A. G. (2024). Using convex optimization to efficiently apportion tracer and pollutant sources from point concentration observations. *Water Resources Research*, 60, e2023WR036159. <https://doi.org/10.1029/2023WR036159>

Received 30 AUG 2023  
Accepted 15 APR 2024

# Using Convex Optimization to Efficiently Apportion Tracer and Pollutant Sources From Point Concentration Observations

Richard Barnes<sup>1</sup>  and Alex G. Lipp<sup>2</sup> 
<sup>1</sup>Lawrence Berkeley National Lab, Berkeley, CA, USA, <sup>2</sup>Merton College, University of Oxford, Oxford, UK

**Abstract** Rivers transport elements, minerals, chemicals, and pollutants produced in their upstream basins. A sample from a river is a mixture of all of its upstream sources, making it challenging to pinpoint the contribution from each individual source. Here, we show how a nested sample design and convex optimization can be used to efficiently unmix downstream samples of a well-mixed, conservative tracer in a steady state system into the contributions of their upstream sources. Our approach is significantly faster than previous methods. We represent the river's sub-catchments, defined by sampling sites, using a directed acyclic graph. This graph is used to build a convex optimization problem which, thanks to its convexity, can be quickly solved to global optimality—in under a second on desktop hardware for data sets of ~100 samples or fewer. Uncertainties in the upstream predictions can be generated using Monte Carlo resampling. We provide an open-source implementation of this approach in Python. The inputs required are straightforward: a table containing sample locations and observed tracer concentrations, along with a D8 flow-direction raster map. As a case study, we use this method to map the elemental geochemistry of sediment sources for rivers draining the Cairngorms mountains, UK. This method could be extended to non-conservative and non-steady state tracers. We also show, theoretically, how multiple tracers could be simultaneously inverted to recover upstream run-off or erosion rates as well as source concentrations. Overall, this approach can provide valuable insights to researchers in various fields, including water quality, geochemical exploration, geochemistry, hydrology, and wastewater epidemiology.

**Plain Language Summary** Rivers transport solid and dissolved elements, chemicals, and pollutants from upstream to downstream. Knowing where these elements are coming from, and in what quantities, is useful in targeting environmental protection or learning about the geology and hydrology of an area. However, a single sample of a river can't provide much useful information because all of the different sources have been mixed together by the river. If multiple samples are taken upstream and downstream, we show that a computational technique called convex optimization can be used to “unmix” the river's waters or sediments revealing where the elements have come from. This approach is accurate and fast enough that we can run it thousands of times with small changes to the sampled values to get a sense of the uncertainty. Though we speak about our method in terms of rivers here, it could be useful in other contexts, such as determining where pollutants or tracers of disease outbreaks are entering sewage networks.


## 1. Introduction

Rivers transport the elements, minerals, chemicals, and pollutants that are produced in their upstream basins; we refer to these, collectively, as tracers. When a water or sediment sample is collected from a specific point in a river, we can analyze the concentrations of these tracers to estimate their upstream production rates. This information is valuable in different ways depending on the type of tracer being measured. For pollutants, it can help us assess and target mitigation strategies, while for elemental and isotopic tracers, it can provide insights into geochemical processes such as chemical weathering (Bufe et al., 2021; Garzanti et al., 2014; Lizaga et al., 2019; Munro et al., 2019).

If a single point on a river is sampled, the water's composition is an integration of all the tracer sources within the river basin. Since upstream areas can be vast, this information is often not specific enough to be useful. If two samples are taken, one downstream of the other, the change in composition between them is a function of the tracer production rate specifically in that stretch of the river. However, if a tributary enters the river between the sample sites we can no longer say if the change in composition is due to sources feeding into the main channel or

© 2024. The Authors.

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.



**Figure 1. Intuition of nested sampling on a river** (a) One downstream sample,  $D$ , is taken at the mouth of the river (blue line). The concentration of a tracer in the sample tells us about the average production rate in the basin. (b) A new sample is taken upstream,  $U$ . Now, sample tracer concentrations tells us about production rates in the two sub-basins indicated by a gray-dashed line (c) A further sample  $T$  is taken on a tributary, further sub-dividing the catchment.

from the tributary's water mixing in. Adding a third sample site in the tributary can help us “unmix” the water, but the mathematics and computation required become increasingly complicated as the number of tributaries and sample sites increases. Apportioning downstream observations into the contribution from each sub-basin is thus recognised as challenging (e.g., in the field of geochemical mapping; Spadoni et al., 2004).

Lipp et al. (2021) previously worked to unmix observations of sediment geochemistry to determine the geochemical composition of source regions. First, the concentration of an element in downstream sediments was described as an integral of the concentration of the element in the upstream area. This forward model was then inverted. To overcome non-uniqueness in the solutions, the model was regularized by seeking the smoothest upstream concentration map that best fit observations downstream. Whilst this scheme was successfully used for geochemical mapping in northern Australia and for quantifying the sources of heavy metal pollution in the River Clyde, UK (Eschenfelder et al., 2023; Lipp et al., 2023), it has a number of limitations.

The primary limitation is that this scheme was solved as a general non-linear optimization problem. Solvers for such problems are unable to guarantee that they have found a globally optimal solution. Additionally, these solvers (e.g., Nelder–Mead, Powell's Method) can be computationally expensive relative to those used to solve other types of problems such as convex optimization, linear-programs, and least squares (Boyd & Vandenberghe, 2004; Press et al., 1992). This computational slowness (hours for tens of samples on standard desktop hardware) makes it impractical to explore hyper-parameters or quantify uncertainty via Monte Carlo methods.

Here, we present a new method for solving this unmixing problem. We consider only conservative tracers (i.e., tracers only affected by mixing) that are well-mixed in a steady state system. We model nested sampling sites (i.e., sites that lie upstream and downstream of each other on both the main channel and in its tributaries) as a directed acyclic graph (p. 574; Sedgewick & Wayne, 2011) constrained by the river network. We use this graph to pose unmixing as a convex optimization problem. The speed of our method makes it easy to quantify uncertainty through Monte Carlo resampling and to explore the impact of changing parameters such as regularization strength. The implementation, available as an open-source Python package, requires only a geospatial raster of “D8” flow-directions (which are widely available or can be generated from digital elevation data using GIS tools) and a data table of samples giving observed tracer concentrations and the location of sample sites. Our software can be used to analyze any data set of nested, conservative tracers, in both sediment and water, across a drainage network.

## 2. Methods

### 2.1. Intuition

If the concentration of elements, minerals, chemicals, isotopes, or pollutants (herein referred to generically as tracers) is sampled at a single point on a river, what can we say about the upstream processes that generated this sample? If we have observed a tracer concentration  $d$  in the water or sediment then, assuming we have no other information, a parsimonious interpretation is that the source concentration,  $c$ , is the same as the observed concentration,  $d$ , at the sample site (i.e.,  $c = d$ ) and that the tracer is generated uniformly across the whole upstream area (Figure 1a).

We consider a situation where we've taken two samples from the river, one downstream of the other. We call the upstream Sample U and the downstream Sample D (Figure 1b). As before, we can't infer much about the production rates upstream of Sample U: whatever the value of the concentration  $d_U$  is, it is assumed to be generated uniformly by all the area upstream of Sample U, and so  $c_U = d_U$ . However, we can now say more about the source concentration in the area between the two samples,  $c_D$ . The observed concentration at Sample D is:

$$d_D = \frac{c_U q_U + c_D q_D}{q_U + q_D}, \quad (1)$$

a weighted mixture of the sources upstream of U and in the area between U and D. In this expression,  $q$  refers to the flux of water or sediment added in the sub-basin of a given sample. For example,  $q_D$  is the flux generated between Sample U and Sample D. It's also important to note that *this relation is only true if the tracer is conservative and well-mixed in a steady state system*; that is, this is only true if the tracer does not get removed from the water as it flows downstream (via, e.g., photodegradation, consumption, flocculation and settling, and/or absorption) and the production and transport of the tracer is in steady-state.

Increasing the number of samples we collect means we can identify where tracers are coming from in the catchment in greater spatial detail. For example, imagine we take a sample on the tributary, Sample T, that joins the river between Sample U and Sample D (Figure 1c). We can see again that  $c_T = d_T$ , and again  $c_U = d_U$ . However, now, Sample D combines information from the areas defined by Samples U and D as well as the new Sample T. Therefore,

$$d_D = \frac{c_U q_U + c_D q_D + c_T q_T}{q_U + q_D + q_T}. \quad (2)$$


The question we are interested in is, given downstream observations, how can we solve these expressions for the upstream tracer sources? If  $d_U$ ,  $d_D$ , and  $d_T$  were exact measurements, we could find  $c_D$  by solving the above equation using the substitutions  $c_U = d_U$  and  $c_T = d_T$  and simple algebra. However, in all real scenarios none of the observations are exact due to analytical or sampling uncertainties. Additionally, our underlying assumptions (e.g., of the tracer being well-mixed, fully conservative, and in steady state) may not fully hold. These reasons mean that the system of equations implied by Equation 2 is unlikely to permit an exact solution.

We show in Appendix A that the method of least squares can be used to solve systems where exact solutions cannot be obtained by minimizing the squared distance between predicted and observed concentrations. However, least squares is an inappropriate way to measure differences for this situation as real concentration measurements typically follow log-normal distributions spanning orders of magnitude (e.g., Gaillardet et al., 1999). This effect is a particular concern in situations where pollutants from a point source are diluted as they flow downstream. For example, the addition of the tributary's flow  $q_T$  to the main river in Figure 1 may increase  $q_D$  significantly. More generally, as we move a distance  $L$  downstream, Hack's Law says that we should expect the basin area (and its flow) to increase as  $\sim L^{1.6}$  (Hack, 1957). Therefore any tracer discharged from a point source upstream, will naturally dilute as a function of  $L^{-1.6}$  as it is transported downstream. If we use least squares to solve such a case we will end up weighting the undiluted, upstream samples much more heavily than diluted, downstream samples. Therefore, it is more sensible to solve these equations by instead considering the relative difference (e.g., the log-ratio) between the observed and actual concentrations. This is similar to the well-known need to normalize data in machine learning (LeCun et al., 2012). Since the relative differences are all of a similar magnitude, they will be weighted equally by our solution procedure. If we needed to account for variable uncertainties across measurements, different weights could be applied to each relative difference term analogously to a weighted least squares.

We now formally describe how to set up and solve this problem for any set of observations on a drainage network. We first consider how to build a "forward" model that translates a given set of upstream concentrations into a suite of concentrations at sample sites. We then demonstrate how to "invert" this forward model to translate a set of downstream observations into a set of predicted upstream concentrations.

## 2.2. Forward Model

We assume that we have made nested (upstream, downstream, and tributary) observations of concentrations of a well-mixed, conservative, steady state tracer at sample-sites within a basin (e.g., Figure 2a). Each sample site



**Figure 2.** Sample sites on a drainage network and their representation as a directed acyclic graph. (a) Theoretical drainage network with six sample sites, 1 through 6. Each sample defines a sub-basin indicated by dashed lines. Color indicates the concentration ( $d$ ) of some tracer, which is sourced in the sub-basin of sample 2 then transported and diluted downstream. (b) The sample network in panel a, represented as a directed acyclic graph.

becomes a node in a directed acyclic graph (DAG) which we term a “sample network”. The topology of this graph (Figure 2b) reflects the topology of the drainage basin (Figure 2a). As detailed below, for each node in this graph we observe or assume the flux  $q$  (of water or sediment) generated by the sub-catchment node. We make observations of the tracer concentration  $d$  at the outlet of each sub-catchment node; from these we seek to recover the generated concentration  $c$ .

Each node of the DAG has parameters we need to specify and a variable we wish to solve for.

- $a$  [ $L^2$ ] for example,  $km^2$ : the sub-basin area, a constant. The area of each sub-basin uniquely defined by the sample. In Figure 2 this is indicated with dashed lines. This is implicit to  $q$ , below.
- $\phi$  [ $M L^{-2} T^{-1}$ ] for example,  $kg/km^2/yr$ : the material export-rate, a parameter. This quantity is the total amount of water or sediment produced in the sub-basin *that reaches the most downstream node in the network*. If there is no storage of material in the network this can be considered equivalent to the average run-off or sediment generation rate of the sub-basin. This is implicit to  $q$ , below.
- $q$  [ $M T^{-1}$ ] for example,  $kg/yr$ : the material flux of the sub-basin, given by  $q = a\phi$ , a parameter. Since we often do not know  $\phi$ , it is easier to treat  $a\phi$  as a combined quantity. For most purposes it suffices to treat this as being equal to the volume of water per time passing by the sample site associated with a node. Because we assume steady state,  $q$  is constant as tracers pass through the system.
- $d$ , unitless for example,  $mg/kg$ : the observed concentration, a parameter. This is the concentration of the tracer observed in the water or sediment at the sample site. This is a function of *all* the nodes in the network upstream of the sample point.
- $c$ , unitless for example,  $mg/kg$ : the source concentration, a variable, which we would like to estimate by solving the unmixing problem. This quantity is the average concentration of the tracer in material being exported uniquely by the sub-basin, excluding upstream basins. For example, in Figure 2,  $c$  is uniformly low for all nodes except for node 2. The tracer exported from node two is transported, and diluted downstream meaning that the observed concentration  $d$  is highest at node 2 and decreases downstream at sites 4 and 6. Identifying  $c$  for each node in a network, given point observations of  $d$  is the goal of this study.

As the tracer is assumed to be conservative and in steady state we can build expressions for the components of  $\mathbf{d}$  as flux-weighted mixtures of all the nodes upstream in the network:

$$d_i = \frac{\sum_{j \in U_i} q_j c_j}{\sum_{j \in U_i} q_j}, \quad (3)$$

where  $U_i$  is the set of nodes upstream of node  $i$ , including itself. Note that for “leaf” nodes (a node with no upstream nodes) the only contributor to the flux is the unique upstream area of that node and so the concentration at the sample site is simply the same as the source concentration. As an example, the concentration at node four in Figure 2 is:

$$d_4 = \frac{c_1 q_1 + c_2 q_2 + c_4 q_4}{q_1 + q_2 + q_4}.$$

Thus, assuming that  $\mathbf{a}$ , and  $\phi$  are known, it is possible to translate a vector of source concentration values  $\mathbf{c}$  into a vector of predicted concentrations at sample sites  $\mathbf{d}_{\text{pred}}$ .

### 2.3. Inverse Model

We now consider the inverse problem: what is the vector of source concentrations,  $\mathbf{c}$  that best fits a set of downstream tracer observations  $\mathbf{d}_{\text{obs}}$ , given a sample network,  $N$ . We pose this question as an optimization problem, seeking to minimize the total misfit between  $\mathbf{d}_{\text{obs}}$  and  $\mathbf{d}_{\text{pred}}$  (calculated using Equation 3) with respect to  $\mathbf{c}$  across all nodes in the network. We therefore solve

$$\begin{aligned} & \underset{\mathbf{w.r.t} \ \mathbf{c}}{\text{minimize}} \quad \sum_{i \in N} \text{misfit}(d_{\text{obs},i}, d_{\text{pred},i}) \\ & = \underset{\mathbf{w.r.t} \ \mathbf{c}}{\text{minimize}} \quad \sum_{i \in N} \text{misfit}\left(d_{\text{obs},i}, \frac{\sum_{j \in U_i} q_j c_j}{\sum_{j \in U_i} q_j}\right) \end{aligned} \quad (4)$$

given a misfit function between two concentration values, and where  $i$  is the index of nodes in the network  $N$ . A naive choice for the misfit function would be the Euclidean distance  $\sqrt{(d_{\text{pred}} - d_{\text{obs}})^2}$ . In fact, if such a Euclidean misfit function is chosen, the minimization problem can be solved using a linear least-squares approach as discussed above and in Appendix A. However, concentration data typically span many orders of magnitude, and as such the Euclidean distance over-weights observations of high concentration relative to observations of low concentration. In addition to this empirical observation, concentration data are technically compositional, and thus the information they contain is argued by some authors to be *relative* not absolute (Aitchison, 1986; Pawlowsky-Glahn & Egozcue, 2006).


As a result, to better extract these relative differences over ranges of many orders of magnitude, a more appropriate misfit function would be the absolute value of the log-ratio of the two concentrations:  $|\log(d_{\text{pred}}/d_{\text{obs}})|$  (Aitchison, 1986). This function is minimized at  $d_{\text{pred}} = d_{\text{obs}}$ , and is symmetrical with respect to relative changes in value. However, this also has a drawback: it is a non-convex function. Functions which are non-convex are challenging to minimize because there is no general algorithm for finding their global minimum, though local minima can be identified (Boyd & Vandenberghe, 2004). In contrast, for convex functions the global minimum is the same as the local minimum and an extensive set of robust tools are available to find it. We therefore define the following convex misfit function that penalizes relative differences between  $d_{\text{pred}}$  and  $d_{\text{obs}}$ :

$$\text{misfit}(d_{\text{pred}}, d_{\text{obs}}) = \max\left(\frac{d_{\text{obs}}}{d_{\text{pred}}}, \frac{d_{\text{pred}}}{d_{\text{obs}}}\right). \quad (5)$$

This misfit function is visualized in Figure 3. As  $\frac{d_{\text{obs}}}{d_{\text{pred}}}$  and  $\frac{d_{\text{pred}}}{d_{\text{obs}}}$  are convex and affine respectively for  $d_{\text{pred}} > 0$ , their maximum is also convex over this range. We note that Equation 5 is non-differentiable at  $d_{\text{pred}} = d_{\text{obs}}$  but this is not relevant for our optimisation approach. Note that Equation 5 takes its minimum value at the same point as  $|\log(d_{\text{pred}}/d_{\text{obs}})|$ ; thus, any solution which minimizes the convex function Equation 5 would also minimize  $|\log(d_{\text{pred}}/d_{\text{obs}})|$ . Although this function is undefined if either  $d_{\text{pred}}$  or  $d_{\text{obs}}$  are equal to zero, concentration observations in nature are rarely truly equal to zero; instead, they mostly represent values below detection limits. However, this limitation does mean that any zero values in a data set ought to be removed or replaced with small values prior to analysis (e.g., Martín-Fernández et al., 2003).

We can therefore pose a convex minimization problem to return the source concentrations,  $\mathbf{c}$  that best fit a set of observations downstream.

$$\underset{\mathbf{w.r.t} \ \mathbf{c}}{\text{minimize}} \quad \sum_{i \in N} \max\left(\frac{d_{\text{obs},i}}{d_{\text{pred},i}}, \frac{d_{\text{pred},i}}{d_{\text{obs},i}}\right) \quad (6)$$



**Figure 3. Visualizing two relative misfit terms.** Red curve shows log-ratio misfit between predicted concentrations  $d_{pred}$  and an arbitrarily chosen observed concentration  $d_{obs}$  value of 30. Blue curve is the misfit between the same values but using the convex misfit term chosen in this study. The log-ratio curve is non-convex. Note that the convex misfit function has a minimum value of 1, not 0 like the log-ratio misfit.

$$= \underset{w.r.t \mathbf{c}}{\text{minimize}} \sum_{i \in N} \max \left( \frac{d_{obs,i} \sum_{j \in U_i} q_j}{\sum_{j \in U_i} q_j c_j}, \frac{\sum_{j \in U_i} q_j c_j}{d_{obs,i} \sum_{j \in U_i} q_j} \right). \quad (7)$$

## 2.4. Regularization

Measurements of concentrations are always accompanied by a degree of uncertainty that reflects both analytical noise, and also cases where the modeling assumptions are not fully satisfied, for example, imperfect mixing of the water column, hydrodynamic sorting of sediment, or non-conservative behavior (e.g., Bouchez et al., 2010; Stutenbecker et al., 2023). To prevent the optimization from over-fitting this noise we regularize our solution by penalizing relative deviations in  $\mathbf{c}$  from the mean observed concentration. We opt to penalize differences from the geometric as opposed to the arithmetic mean of the observed concentrations,

$$\bar{d} = \left( \prod_{i=1}^n d_{obs,i} \right)^{\frac{1}{n}}.$$

We use the geometric mean because, as discussed above, concentration observations typically span many orders of magnitude and the arithmetic mean, if chosen, would overweight high concentrations. This regularization prevents solutions from containing implausibly extreme concentrations. We utilize the same misfit function as above to quantify relative deviations from the mean. Thus, we can add regularization terms to Equation 7:

$$\underset{w.r.t \mathbf{c}}{\text{minimize}} \sum_{i \in N} \max \left( \frac{d_{obs,i}}{d_{pred,i}}, \frac{d_{pred,i}}{d_{obs,i}} \right) + \lambda \sum_{i \in N} \max \left( \frac{\bar{d}}{c_i}, \frac{c_i}{\bar{d}} \right), \quad (8)$$

where  $\lambda$  is a hyper-parameter controlling the strength of regularization. The value of  $\lambda$  can be calibrated using various methods. One approach is to analyze the trade-off between data misfit and the total deviations of the model from the mean observed value. Then, the model with the smallest deviations that still provides an adequate fit to the data is selected (the Occam's or “elbow” approach; Constable et al., 1987). Another method is to utilize an unseen “validation” subset of the data to calibrate the value of  $\lambda$ , or similar cross-validation approaches.


Previous unmixing approaches regularized by penalizing the “roughness” of the source concentration (Lipp et al., 2021). Roughness was calculated as the first derivative of  $\mathbf{c}$  with respect to spatial coordinates. We opt not to use this approach here for two reasons. First, we haven't yet found a convex relative difference term between two variables that is compatible with the other constraints of our problem; Equation 5 is only convex if one argument is a constant. Thus, calculating the relative difference between two spatially adjacent  $\mathbf{c}$  variables is not, at present, possible without violating the convexity of our optimization problem. Second, prior authors solved for  $\mathbf{c}$  on a regular grid over the modeled domain for which roughness has a more intuitive interpretation (Lipp et al., 2021). By contrast, we now solve for  $\mathbf{c}$  defined on discrete sub-basins which are irregular in shape (Figure 2). Whilst it is possible to conceive of definitions of roughness on such a domain (e.g., normalizing the difference between adjacent nodes by the length of shared boundaries) it is not clear that these always represent sensible terms to penalize.

Note that neither of the optimization problems described by Equations 7 and 8 are sensitive to whether or not the directed acyclic network is fully connected. If two disjoint river networks with fully separate basins are combined in the same optimization problem the solution to that problem is the same as if the networks were solved separately, assuming no regularization. If regularization is applied, differences in solving the disjoint basins separately or combined may arise if the observations from the two basins have different means.

## 2.5. Transient Conditions

Our method will give the most accurate results when the system being sampled is in a “steady state”, meaning that tracer production and transport rates are constant and that there is good mixing and no dispersion or extinction of





**Figure 4. Flow-chart of the inverse approach presented in this study.** From a D8 raster and the coordinates of sample sites a ‘SampleNetwork’ data structure is built. This structure is a directed acyclic graph where each node corresponds to a sub-catchment induced by a sample site. By combining the SampleNetwork with sample concentration data (and optionally the export-rate of each sub-catchment i.e., run-off or erosion rate) the observations are unmixed to recover (1) the optimal source concentrations and (2) the fitted concentrations for each sample.

the tracers. An example of a steady state system is one in which we are sampling geochemical tracer production rates using rivers in the absence of recent rainfall. In such a system, samples could be taken in a fully desynchronized manner because the downstream state at *any* given time is representative of the upstream state at that same time.

In a transient system, flow conditions can lead to under- or over-estimates of production. For instance, consider a high-precipitation storm event. After the event has ended, simultaneous upstream and downstream samples are taken. If the upstream flow rate has dropped to normal while the downstream flow rate is still elevated then the observed downstream concentration will be lower due to the high flow condition. As a result, the upstream concentration will be inferred to be higher than it actually is. If a time series of concentration observations and flow rates is available for each sample site, this could be used to find appropriate samples at disparate times to infer the actual concentrations, though we do not explore this here.

### 3. Implementation

We provide an open-source python implementation of the above method at [github.com/r-barnes/faster-unmixer](https://github.com/r-barnes/faster-unmixer) and archived at [zenodo.org/records/8299812](https://zenodo.org/records/8299812) (Barnes & Lipp, 2023). The code is well-documented with unit tests, examples, and installation instructions. A package “funmixer” is generated which can be used to easily create and solve instances of the above problem given a D8 raster of flow-directions and a data table of sample locations and concentrations. A flow-chart illustrating our approach and the required inputs is shown in Figure 4. Here we describe our implementation in more detail.

#### 3.1. Building Sample Networks

The first stage for solving the inverse problem is building the sample network data structure. To build this, we require a D8 raster (O’Callaghan & Mark, 1984) of the modeled flow-directions for the study-region and a data table containing the spatial coordinates of the sample sites. The D8 hydrological model assumes only convergent flow across a network, that is, each cell in the raster is able to donate flow to only one neighboring node. As a result, our method is unable to consider networks which have strongly divergent flow, though it could be modified to do so since even divergent flow still forms a directed acyclic network.

To build a sample network, we iterate upstream from sink cells, which are assumed to occur at the boundary of the D8 and at natural sinks within the domain (e.g., lakes). At each cell, we check to see if our location corresponds to one of the provided sample-sites. If it does, we generate a new node in the sample-network. We then iterate through the neighbors of this cell, and if they donate flow to it, we assign them a label to indicate that these nodes donate their flow uniquely to that sample-site. We then recursively continue to search upstream of these donor cells labeling cells as we go. If we reach a cell which corresponds to a sample site we change the label we apply to the cells upstream of that site and the process continues.

To allow this approach to work even for very large D8 rasters, we implement it in C++ using tools from the RichDEM terrain analysis software (Barnes, 2016b); however, the resulting sample network is passed back to Python where it is represented using the networkx package (Hagberg et al., 2008). Conceptually, our method is general to any acyclic network; the use of rasters and D8 is only a convenience meant to make the method accessible for the significant use-case of analyzing data gathered on river networks.

The geographic locations of river sample sites and hydrologic stations are often misaligned to channels extracted from digital elevation models (Lindsay et al., 2008). As a result, observations are often automatically or manually snapped to the nearest channel cell in a digital elevation model (DEM), for example, using GIS software. We strongly recommend performing such snapping prior to creating the sample network and using our method. If samples are not aligned to channel cells beforehand, samples may not produce a connected network or instead have unrealistically small drainage areas. We also recommend that the D8 raster be hydrologically conditioned prior to building the network such that modeled channel locations align closely with real observed channels. This hydrological conditioning often removes artificial depressions that prevent continuous flow-lines across a landscape. Algorithms to perform this sink-filling are widely available (e.g., Barnes et al., 2014; Barnes, 2016a). Hydrologically conditioned D8 and DEMs are also increasingly publicly available at a variety of resolutions (e.g., Hydrosheds; [www.hydrosheds.org](http://www.hydrosheds.org)).

### 3.2. Solving the Unmixing Problem

To solve the optimization problem posed by Equation 8 we take the networkx representation of the sample network from §3.1 and associate each node with a *variable*  $c$  for its concentration, which we wish to solve for. We also associate each node with a *parameter*  $\phi$  for its export rate and a *constant*  $a$  for its area. We therefore distinguish between variables, values we wish to solve for, and parameters, values which we set but would like to modify, and constants, whose values are fixed. We also add a dummy variable  $d$ , the predicted concentration at the sample site, which is initially set to zero, but which we describe in terms of the concentrations of the upstream nodes, as in Equation 3. We then implement the forward model from §2.2 using a topological sort (p. 575; Sedgewick & Wayne, 2011) on the networkx sample network. As the sort moves from upstream to downstream we propagate concentrations and fluxes downstream building Equation 3 for each node as it is reached. Finally, we extract the dummy variables  $d$  into a vector  $\mathbf{d}_{\text{pred}}$  and combine these with the observed values held in the  $\mathbf{d}_{\text{obs}}$  parameters to build the objective function from Equation 7. Optionally, regularization terms are built by creating a list of the difference (using the convex relative misfit function in Equation 5) between the  $c$  parameters and the mean observation parameters. If regularization is chosen, this requires the regularization strength  $\lambda$  to be defined giving us the optimization problem defined by Equation 8.

To improve numerical accuracy, we internally normalize the values of the areas, export-rates, and observed concentrations such that they are all approximately equal in absolute value. We do this by dividing each of these values by the geometric mean of all the values before calculations. The final calculated concentrations are then multiplied by the geometric mean again before being returned. As our problem is only sensitive to relative differences in these values this has no impact on the final results (e.g., mixing the chemistry of two rivers in proportions 1:1 is the same as mixing them in the proportions 100:100). This normalization does however increase numerical accuracy when the problem is passed to solvers, below. Another consequence of this is that concentrations, areas and export rates can be provided in any units and our method will still produce the same results.

We implement the variables, parameters, objective function, and optimization problem from above using Python's CVXPY package, a domain-specific language for disciplined convex programming (Diamond & Boyd, 2016; Grant et al., 2006). The introduction of parameters means that we need to use CVXPY's Disciplined Parametrized Programming extensions (Agrawal et al., 2019).

Once the problem has been expressed to CVXPY, CVXPY must perform an operation called canonicalization. Canonicalization ensures that the problem is convex, and therefore solvable, while converting the many variables and parameters in the problem into an intermediate representation suitable for passing to a *solver*. At present canonicalization in CVXPY is a very time-consuming operation relative to solving. Fortunately, if only the values of the parameters (in this case, the observation values, export-rates and regularization strength) are changed the intermediate representation can be updated without having to do canonicalization again. In practice this means that, once a problem has been composed, the export rates, observed concentration values, and regularization strength can be adjusted and subsequent solves occur rapidly. This allows for very efficient uncertainty propagation using Monte Carlo inversion as well as efficient calibration of the regularization strength parameter. An example of this is detailed below.

Solvers are the programs that find the solutions to a given optimization problem. A variety of solvers can be used for a given problem and CVXPY facilitates this by providing a unified way of writing problems. For our problem,



we experiment with the SCS, ECOS, CLARABEL and Gurobi solvers (Domahidi et al., 2013; Goulart & Chen, 2023; Gurobi, 2023; O'Donoghue et al., 2016). A comparison of the performance of these solvers is detailed in §3.4. The solvers typically complete their work quickly, even for large problems, and CVXPY translates their solutions back into a convenient form. Solvers often have a number of settings that trade accuracy for speed. In the tests described in §3.3 the accuracy requirements for ECOS had to be reduced from its default settings for it to converge. SCS was run with its default settings and Gurobi was run in a high-accuracy mode. Note that solvers often perform preprocessing which allows them to separate out the components of disjoint problems. In practice this means that several disconnected river networks can be solved at once without having a significant impact on solution times (as shown in the case study below).

### 3.3. Method Verification


To verify the numerical accuracy of our method we ran it on randomly generated sample networks with known concentration values. Using the forward model (§2.2), we generated synthetic observations. Using our method, we inverted these observations to try to recover the original concentrations. We found that we can recover the original concentrations to within a specified tolerance thereby verifying that our method is accurate. We performed this test for each of three types of networks: random trees, full R-ary trees, and balanced trees. These three tree types cover a variety of structures likely to be represented in real data sets. The latter two networks required a branching factor to be specified, which determines how many upstream nodes flow into each downstream node. This branching factor was also chosen randomly. We assigned each node in each network a random concentration and random sub-basin area, both of which were drawn from log-uniform distributions. Subsequently, we generated a suite of synthetic concentration observations by mixing along the network using the forward model from §2.2. Finally, we solved the optimization problem from Equation 7 (without regularization) and compared what we recovered to the initial synthetic input. The software accompanying this paper implements the above test using Python's Hypothesis testing library (MacIver et al., 2019). Unlike a standard unit test, which checks that a given input produces an expected output, Hypothesis implements property-based testing. In this form of testing a large number of random inputs are passed to a function and we test whether some known property holds for each of the inputs. This form of testing can often identify errors that a single, manually created unit test would miss. In this case, the property we test is that, for any randomly generated network and set of concentrations and areas, we can recover these values after simulating how they would propagate downstream through the network.

We draw concentrations and areas uniformly from log-uniform distributions over the range [1, 100] (i.e., over two orders of magnitude), the network size uniformly from the range [2, 100], and the branching factor uniformly from the range [1, 5]. For each trial scenario we used both the ECOS and GUROBI solvers to solve the optimization problem posed by Equation 7. We find that under these conditions we are able to successfully recover the upstream concentration for each upstream node to an accuracy of better than 1%. This test is included in the code repository.

Because we normalize absolute values by their geometric mean (§3.2) this test is valid for any data sets where sub-basin area and source concentration range across two orders of magnitude, regardless of the absolute values. Additionally, because the flux from each node is the product of sub-basin area and source concentration (for a fixed export rate), the above test corresponds to networks where the tracer *flux* from the node varies over four orders of magnitude. For sample campaigns where sub-basin areas are approximately equal in size, we find that we can draw source concentrations from an even greater range and still achieve the same level of accuracy. The concentration values we expect to encounter in a given problem can vary widely, and may vary beyond the specific range tested here. We therefore recommend conducting more specific simulated tests that match a specific problem's characteristics in order to better quantify the accuracy of the recovered results—this is a best practice for any statistical analysis.

### 3.4. Benchmarking Runtime

Our method performs quickly across a range of network sizes, including sizes larger than any sampling campaign we are aware of. To demonstrate this, we generate and solve random problems, as described in §3.3. For simplicity we consider only full R-ary trees with a branching factor of three (i.e., ternary trees). We run our benchmark for the SCS, ECOS, GUROBI and CLARABEL solvers on networks up to a maximum size of 500 nodes. For each network size, 10 trials were performed and the average and standard deviation of the runtimes



**Figure 5. Scaling of algorithm solve time against network size.** Runtime tests performed on synthetic R-ary networks of branching factor three up to 500 nodes in size. Areas and concentrations of each node randomly varied between one and 100. For each network size, 10 repeats were performed and averages and standard deviations were calculated. (a) Total runtime for different network sizes. Different colored lines indicate different solvers. Solid colors with solid lines indicate total runtimes for the first solve of a problem including problem compilation. Faded colors with solid lines are total runtimes for all subsequent solves of a precompiled problem (relevant for Monte Carlo inversions and sensitivity analysis). (b) Time spent in the solver itself, same colors as panel (a). The time spent in the solver is a small fraction of the total time to solution (note different axes vs. panel a).

were calculated. The results are shown in Figure 5. Calculations were performed using standard desktop hardware. A script to reproduce a version of this figure is provided in our code repository.

Our results indicate that on a standard desktop it is possible to compile and solve the unmixing problem on networks of 500 nodes—significantly larger than most geochemical datasets—in under a minute. For networks of less than 100 nodes, solutions can be found in under a second. On networks of 100 nodes, ECOS, CLARABEL

and GUROBI show approximately linear scaling on their first solve, with log-log slopes of 1.04 and 1.02 respectively. For networks larger than this, the compilation time begins to increase at a greater rate. SCS shows a more complex scaling.

Although the time taken for the first solve is thousands of times faster than existing methods, subsequent solves of the same pre-compiled problem are one to two orders of magnitude faster still (faded lines in Figure 5b). These solves are those for which only parameter values are changed. Our benchmark shows that even for the largest networks we consider, subsequent solves take on the order of a second or less. For the networks we consider, the subsequent solves of ECOS, CLARABEL and GUROBI scale approximately linearly with respect to network size, with log-log slopes of 0.98 and 0.97 respectively. ECOS and SCS are faster than GUROBI. SCS shows a more complex scaling and is less consistent than the other solvers.

The total runtimes in Figure 5 include both the time spent in canonicalization as well as the time spent in the solvers. As discussed in §3.2 and demonstrated in the figure by the solver times (dashed lines), canonicalization dominates the runtime and canonicalizing the problem for the initial solve takes significantly longer than subsequent solves of the same problem using modified parameters. The amount of time spent in the solvers themselves is orders of magnitude less than the time spent building the problem. This shows that the compilation of the problem itself is limiting.

We can decrease the overhead of problem construction by using faster optimization languages such as Julia's Convex.jl library (Udell et al., 2014). Similarly to CVXPY, Convex.jl simply translates a given convex optimization problem into a form suitable for a solver; however, Julia essentially eliminates the time spent building the problem. With this limitation overcome, the solver is now the only bottleneck and very large problems can be solved. We test this on randomly generated ternary trees and observe that Convex.jl, using Gurobi as a solver, takes 32 s to solve a 59,049 node tree and 334 s to solve a 177,147 node tree. Note that for these much larger networks the solver is taking exponentially more time to find a solution. Although Julia's performance is attractive, the Python code we've used for the rest of this study is much more flexible and easier to use than its Julia counterpart and it is the code we recommend using. We expect most networks will be small enough that paying the performance cost to use Python is worth the benefits of accessibility.

Runtime is not the only factor that should be considered in finding solutions to an optimization problem. We notice that Gurobi almost always finds optimal solutions to the problems we give it. In contrast, ECOS often finds solutions which it cannot guarantee are accurate, but which, in practice, are close to the values Gurobi finds in the cases we've checked. SCS is often unable to find optimal accurate solutions and quits early if it determines this. In Figure 5 this effect accounts for the wide variance of SCS run-times. For its accuracy, consistent convergence, and speed, we prefer to use the Gurobi solver for our method. However, ECOS is the default solver we use with our implementation because it is freely available whereas Gurobi requires a license (academic licenses are free at the time of writing).

## 4. Case Study: Cairngorms, Scotland, UK

To demonstrate the method we apply it to a river sediment chemistry data set collected from rivers draining the Cairngorms mountains, UK (Figure 6a). This data set was first presented in Lipp et al. (2020, 2021), where it was used to test the unmixing method discussed in §1. The data set consists of fine-grained (<150  $\mu\text{m}$  fraction) samples gathered from 63 sites on five rivers (Spey, Dee, Deveron, Don, and Tay) draining NE Scotland (Figure 6b). The sample networks corresponding to this data set are shown in Figure 6c. The region contains a diverse range of sedimentary, metamorphic and igneous rocks resulting in a diversity of geochemistry in source regions (Figure 6d). An advantage of applying the method in this region is that the source regions of the sampled rivers are already geochemically very well constrained by the G-BASE geochemical survey (Johnson et al., 2005) which gathered first order stream sediment samples at a density of  $\sim 1$  sample per  $2 \text{ km}^2$  (Figure 6e). The G-BASE data can therefore be used as an independent data set to test our predictions of upstream concentrations. The primary control on source region chemistry is the composition of the underlying lithology (see Figures 6d and 6e).

### 4.1. Regularization

As this real data set incorporates some sampling noise, we opt to regularize the problem to prevent overfitting. This regularization will result in solutions that are stable and meaningful despite noise in the input data. However,

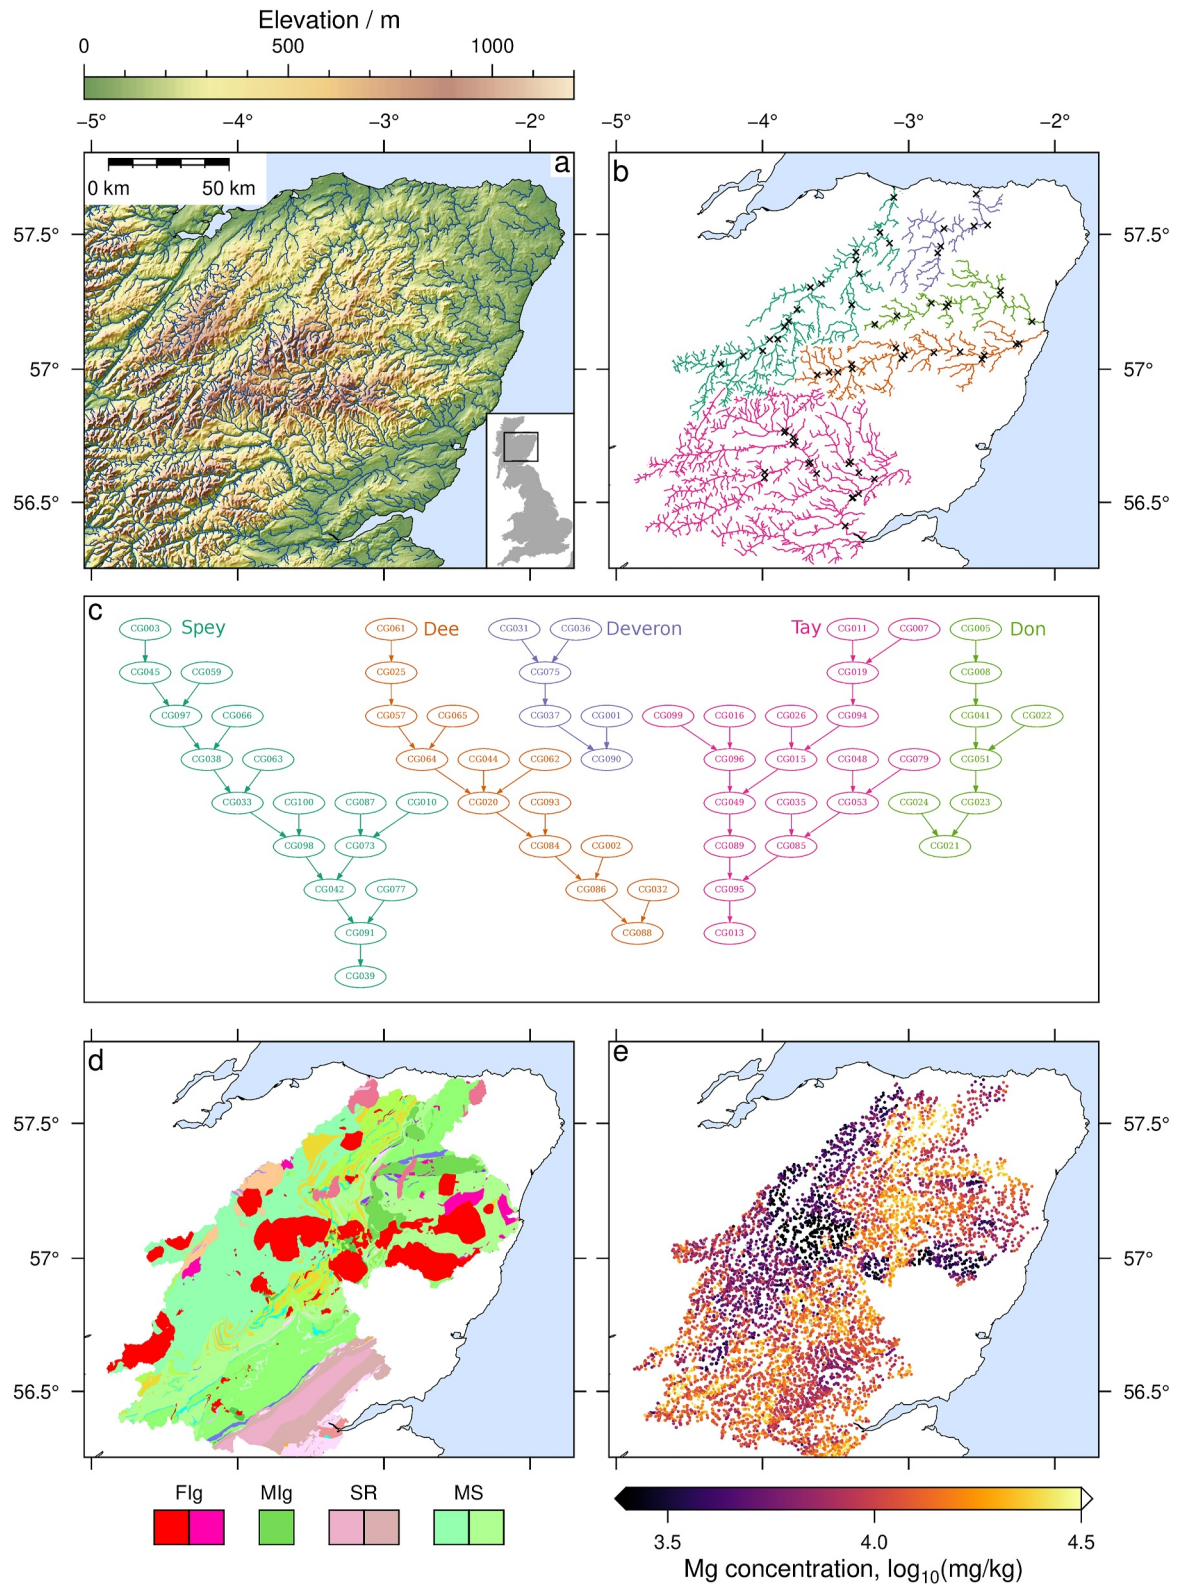



Figure 6.



**Figure 7. Calibrating the regularization parameter.** Plot of root mean square (RMS) error of the predicted downstream concentrations (see Figure 8f) against regularization parameter  $\lambda$ . Chosen value is  $\lambda = 10^{-1}$  which lies at the center of a trade-off curve between regularization strength and data-misfit. Note logarithmic  $x$  axis.

regularization does require us to calibrate the regularization parameter  $\lambda$ . Identifying a regularization parameter is in general a challenging task across all fields, and a full discussion is beyond the scope of this manuscript. We take an “Occam’s razor” style approach seeking the set of upstream concentrations with the smallest number of deviations from the mean observed value that is still able to adequately fit all of the observations downstream. We determine this by performing 100 inversions each with a different regularization parameter sampled equally between  $10^{-3}$  and  $10^1$ . For each  $\lambda$  we calculate the global root mean squared (RMS) error between the observed and predicted concentrations (i.e., between  $\mathbf{d}_{obs}$  and  $\mathbf{d}_{pred}$ ). Plotting this RMS against  $\lambda$  results in the trade-off plot shown in Figure 7. We chose a value ( $\lambda = 10^{-1}$ ) that lies at the “elbow” or “knee” of the trade-off plot, which is a standard, albeit not universally accepted, approach to choosing the  $\lambda$  parameter (James et al., 2013; Salvador & Chan, 2004; Schubert, 2023; Sugar & James, 2003). We could perform a sensitivity analysis with respect to this parameter, but opt not to do so here in order to focus on presenting the new method.

## 4.2. Monte Carlo Inversion

The uncertainty in the downstream observations was constrained by Lipp et al. (2020) by gathering duplicate samples at four randomly chosen sample sites. Concentration differences between the duplicate samples are controlled by the heterogeneity within a sample-site and also uncertainty introduced by the sampling process. Analytical uncertainty introduced in the laboratory process was noted to be negligible in comparison. For Magnesium, the du-

plicates on average showed a relative uncertainty (i.e.,  $\sigma/\mu$ , where  $\sigma$  is the standard deviation and  $\mu$  is the mean) of 3.53%. To propagate this downstream uncertainty upstream we performed uncertainty quantification with Monte Carlo sampling by resampling the downstream data 1,000 times from a distribution defined by the observed value and the relative uncertainty. We invert each of these resampled data sets and calculate the geometric mean and standard deviation of the upstream predictions.

## 4.3. Results

The central estimate and uncertainty for Mg are shown in Figures 8a and 8b. We can assess the accuracy of these predictions by comparison to the independent G-BASE geochemical survey. The predicted upstream concentration in each subcatchment represents the mixture of all sources of the element within the subcatchment. To provide an independent estimate of this mixture, we calculate the arithmetic mean concentration of all G-BASE samples (Figure 6e) within each sub-basin, resulting in the map shown in Figure 8c. We define the uncertainty on the G-BASE estimate for sub-basin concentration as the standard error of the calculated mean (i.e.,  $\sigma/\sqrt{N}$ ). Only one small sub-basin in the lower reaches of the Tay basin contained no G-BASE samples and therefore the predictions for this sub-basin cannot be independently validated. Visual comparison of the upstream predictions (Figure 8a) and the independent G-BASE data set (Figure 8c) indicate a strong similarity. Visualizing the misfit spatially (Figure 3d) shows close accordance for most of the study area barring a few sub-catchments to the south. The cross-plot is shown in Figure 8e. A linear-regression through these data, weighted by upstream-area, returns a slope sub-parallel to the 1:1 line with an  $R^2$  value of 0.63 and a root mean squared misfit of 0.13. This model is able to closely fit the concentration downstream shown in Figure 8f. Potential reasons for the small areas of misfit are discussed in greater detail in other work (i.e., Lipp et al., 2021) but it has not been possible to identify any

**Figure 6. Case-study area in the Cairngorms mountains, UK** (a) Topographic map of region overlain with rivers from Jager and Vogt (2010). Inset shows location of the area on the island of Great Britain. (b) Colored lines show the five studied river catchments. Black crosses show river sediment sampling sites. (c) The directed acyclic graph (DAG) representation of the five sample networks. Colors are same as in panel (a) (d) Geological map of drainage basins. Reproduced with the permission of the British Geological Survey UKRI, all rights reserved. Major lithologies indicated: Flg = Ordovician/Devonian felsic intrusions; Mlg = Ordovician/Silurian mafic intrusions; SR = Sedimentary rocks, mostly Devonian sandstones; MS = Metasedimentary rocks, mostly Neoproterozoic psammites. (e) Points are first-order stream sediment samples from the G-BASE geochemical survey colored by magnesium concentration (Johnson et al., 2005).



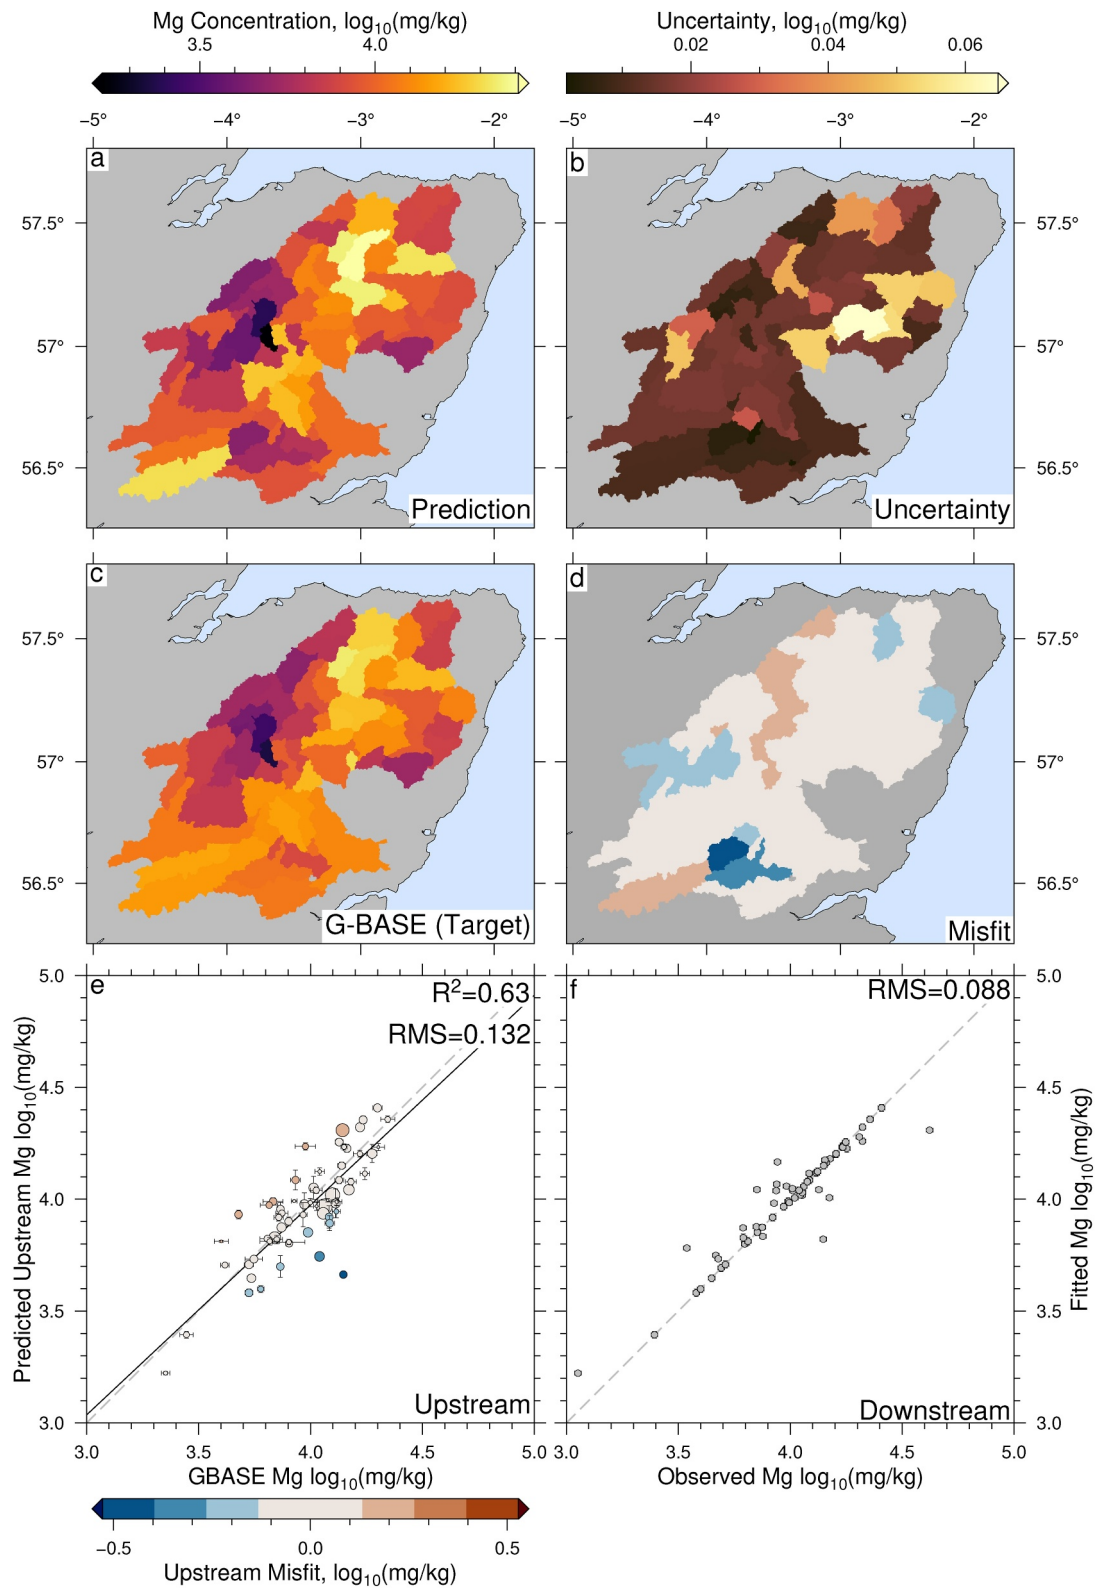
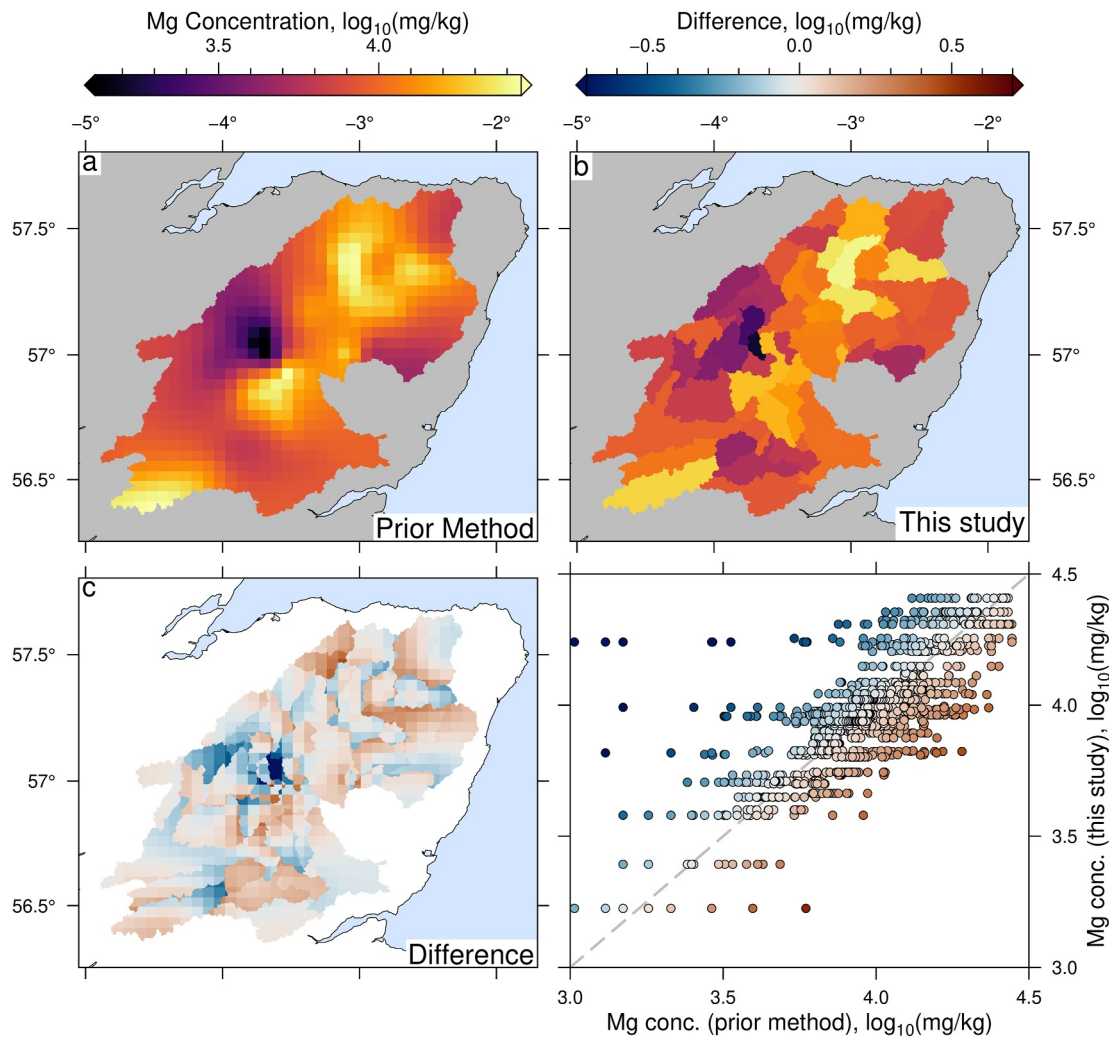


Figure 8.





**Figure 9.** Comparing the case-study results from this study to the results from using the unmixing method presented in Lipp et al. (2021). (a) Source concentration map of Mg from unmixing the data presented in Figure 8 in the main manuscript, using the previous method introduced by Lipp et al. (2021). (b) Same as panel (a) but using the convex unmixing method introduced in this study. The map presented here is the same as shown in Figure 8a (c) Difference between the maps shown in panels a and (b) (d) Cross-plot between the two predictions colored by difference. Gray dashed line indicates 1:1 line.

single reason. Any process that violates the assumptions of our model, for example, unsteady mixing or in-transit geochemical modification may introduce misfit.

In Figure 9 we benchmark the results presented here for this case-study against those calculated using the previous unmixing method introduced by Lipp et al. (2021). This comparison exercise indicates that both methods, when applied to the same data, produce acceptably similar results. As both methods are fundamentally solving different optimisation problems, a perfect correlation is not expected. Most of the differences can be accounted for by the smoothing method present in Lipp et al. (2021). For example, at the center of the studied region, a small sub-

**Figure 8.** Results from inverting Mg concentrations in Cairngorms river sediments. (a) Predicted upstream concentration of Mg in each sub-basin calculated using inversion procedure on downstream river sediments. (b) Uncertainty in prediction calculated via Monte Carlo resampling of downstream data. (c) Target upstream concentration calculated by taking arithmetic mean of all G-BASE samples within each sub-basin (see Figure 6e). (d) Misfit between predicted and target upstream concentrations. (e) Cross-plot of predicted and target concentration, colored by misfit. Symbol size is proportional to sub-basin area. Solid black line is linear regression through points. The regression,  $R^2$ , and RMS are weighted by upstream area. Vertical error bars are those in panel (b) Horizontal error bars are standard error of mean. If not displayed, errors are smaller than symbol. Gray dashed-line = 1:1. (f) Cross plot of predicted and observed downstream concentration.

catchment is identified by our method to have a significantly higher Mg concentration than the surrounding sub-catchments (Figure 9b). This short-wavelength geochemical variability cannot be resolved by the smoothed inversion in Figure 9a. The differing methods result in a relatively higher difference for this sub-catchment as shown in Figure 9c.

Overall we conclude that in this particular example the inversion scheme is able to generate accurate upstream concentrations with associated uncertainties, and is acceptably benchmarked against prior methods. We make no assertion that such accurate results will always be generated using this approach. The model's predictions will only be valid so long as its underlying assumptions are applicable. It is, of course, the responsibility of any scientist to confirm these prior to use. We encourage the inverse results to be independently validated (to the extent that this is possible; Oreskes et al., 1994) using approaches such as cross-validation. Potential variability in predictions should be explored using sensitivity analysis. Indeed, a major advance of this study is the ability to efficiently invert the problem repeatedly for different parameter and data values. For example, the entire procedure described above including the calibration of the regularization parameter and the Monte Carlo calculation of uncertainties was performed in  $\sim 30$  s on a standard desktop computer, despite performing over 1,000 separate inversions. In studies incorporating larger data sets, each of these wholly independent inversions could be performed in parallel for further speed-ups.

## 5. Discussion

### 5.1. Applications

The method we present is designed to be deliberately general and can be applied to any constituent in sediment or water that is well-mixed, conservative, and in steady state. The possibility of extending to non-conservative and time-varying tracers is discussed below. As a consequence of this generality, our method could be applied to a number of problems in different disciplines. For example, in the field of geochemical exploration, nested sampling campaigns are frequently conducted for river and stream sediments. These samples are then used to infer the geochemistry of the upstream lithology for mineral exploration or for setting environmental baselines (Garrett et al., 2008). This unmixing procedure can be applied to such data to accurately reconstruct the geochemistry of the sediment source regions (e.g., Lipp et al., 2023).

The method could also be used in pollutant apportionment studies. If a pollutant is measured at multiple nested sites across a basin, the presented scheme can be used to apportion the pollutant into the contribution from each sub-basin. This approach could be particularly useful for characterizing diffuse pollutant sources, which are typically challenging to identify and quantify (Bowes et al., 2008). This inverse approach to apportionment contrasts with the more common forward modeling approaches to estimating pollutant contributions for example, by assuming a relationship between land-use and pollutant production (e.g., “export coefficients”; Johnes, 1996).

The elemental and isotopic chemistry of river water and sediments is also frequently measured to understand controls on chemical erosion rates in upstream regions (e.g., Bufe et al., 2021); however, it is known that mixing of chemically heterogeneous source regions obscures the primary source region signature (Baronas et al., 2017; Torres et al., 2017). The scheme presented here can be used to account for these mixing effects allowing for more accurate estimates of upstream chemical weathering rates. This in turn can be used to better constrain their underlying driving mechanisms. Additionally, the concentration of cosmogenic nuclides (e.g.,  $^{10}\text{Be}$ ,  $^{26}\text{Al}$ ) in river sands are used to quantify the average erosion rate (von Blanckenburg, 2005). By unmixing nested cosmogenic nuclide samples, erosion rates could therefore be quantified across catchments at a greater spatial resolution.

Finally, whilst we consider only river networks in this study, our approach can be applied to unmixing tracers along any other structure that can be represented as a directed acyclic graph, such as municipal sewerage networks. For example, waste-water epidemiology uses observations of disease tracers such as the genetic remnants of SARS-Cov-2 in sewerage systems to identify the severity and timing of outbreaks (Sims & Kasprzyk-Hordern, 2020). If a suite of samples were taken on a sewerage network (e.g., Nourinejad et al., 2021) our approach could be used to isolate where in the network SARS-Cov-2 is entering the system and, therefore, to estimate where, geographically, the outbreak is occurring. A similar strategy could be applied to other diseases such as polio, malaria, HIV, tuberculosis, measles, and Zika (Kilaru et al., 2023). We will explore this further in future work.

## 5.2. Tributaries, Drainage Areas and Designing Sampling Campaigns

Our method works by solving for the source contribution from the sub-basin defined by individual samples. As a result, our method is agnostic to whether and where any specific tributaries may be joining other channels within that sub-basin. For example, in Figure 1, Samples 3 and 5 lie on a tributary upstream of Sample 6. If these samples are excluded, the sub-basin defined by Sample six includes the entire catchment of this tributary. As a result, the calculated source concentration for sub-basin six is the average production rate of a large area. Such a large sub-basin may not provide useful information.

This result has implications for how to sensibly design sampling campaigns to recover source concentrations. If the goal is equal coverage across a basin, samples should be gathered such that the resulting sub-basins are approximately equal in area. Such a campaign would likely involve nested sampling into tributary basins. Simply sampling along the main channel could result in irregular, large, sub-basins which are not diagnostic. Consequently it would be a challenge to disambiguate between what is generated on the main channel versus what is added to it by tributaries. Algorithms to identify suitable sample sites for finite sampling resources given these constraints would be an interesting further avenue of research.

## 5.3. Solving for Optimal Export-Rates

To solve the optimization problem discussed in this manuscript (i.e., Equation 7) the relative export rate of water or sediment from each sub-basin needs to be known or assumed. By default, our implementation assumes that the export rates of all the sub-basins are the same. If estimates of the run-off rate or sediment export-rate are known for each sub-basin (e.g., from gauging data) these values can instead be used for more accurate results. However, we note that previous studies have found that synthetic variations in export rate have a limited impact on the results of the inversion (Lipp et al., 2021, 2023). This insensitivity was attributed to the fact that within a sampled river basin, geochemical concentrations typically show much greater variability than sediment/water export rates. Nonetheless, an interesting question is whether optimal estimates of the export rate parameters,  $\phi$  can also be recovered from downstream observations. In other words, can a set of concentration observations be jointly inverted for *both* sub-basin source concentration *and* the export-rate. We now explore how this could be posed as an inverse, convex optimization problem.

The convex minimization expression in Equation 7 can be reformulated as a misfit function of the export rates, and a vector of tracer observations:

$$F(\phi, \mathbf{d}_{\text{obs}}) = \min_{\phi} \sum_{i \in N} \text{misfit} \left( d_{\text{obs},i}, \frac{\sum_{j \in U_i} \phi_j a_j c_j}{\sum_{j \in U_i} \phi_j a_j} \right), \quad (9)$$

where “min.” indicates the minimum value of the expression. This function,  $F$ , therefore maps a vector of export-rates and a vector of tracer concentrations onto a single value representing the misfit to the data. Note that because Equation 7 is convex, we can conclude that  $F$  must map onto a single value for the vectors  $\phi, \mathbf{d}_{\text{obs}}$  (i.e., for  $n$  observations,  $F : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ ).

For a particular data set each sample-site may have observations of many different tracers. For example, methods for determining elemental concentrations in water and sediment typically return the concentrations of many elements at once. We can therefore construct a matrix of concentration observations,  $\mathbf{D}_{n \times m}$ , of  $m$  different tracers at  $n$  different sample-sites. Each column of  $\mathbf{D}$  is a vector equivalent to  $\mathbf{d}_{\text{obs}}$ . However, in each of these cases, whilst  $\mathbf{d}_{\text{obs}}$  is different for each tracer, the export rate of the sub-basin is always the same. This is because whilst the particular value of  $d$  will be controlled by the tracer sources (e.g., lithology, pollutants, etc.) the total material export is controlled by independent processes such as erosion or rainfall rate. Collectively, all of these independent tracers provide some constraint on the upstream export rates. We are looking for the vector of export rates,  $\phi$ , that minimizes this expression, which gives, finally, the minimization problem:

$$\underset{\text{w.r.t } \phi}{\text{minimize}} \quad \sum_{j=1}^m F(\phi, \mathbf{D}_j). \quad (10)$$

Solving this expression could, potentially, be a novel approach to yield important geomorphic and hydrologic information. For example, if tracers in sediment are analyzed, the solution could provide, jointly, the relative upstream erosion rate in each sub-catchment as well as the source region geochemistry just from a set of point observations in a network. In the case of water, it would return upstream run-off as well as hydrochemical information. Equation 10 could be considered as a special instance of the general “explicit” mixing problem of Weltje (1997), which solves for mixture proportions *and* source region concentrations, applied to mixtures upon a network.

Unfortunately we have not yet found a convex approach to solving Equation 10, suggesting it may be challenging to optimize. Additionally, the mapping  $F$ , and the resulting misfit space, depends on the structure of the sample network as well as the data set of geochemical observations. Consequently, a full exploration of this optimization problem, including instances when it is easily solved, is beyond the scope of this manuscript. Nonetheless, we state it here as a potential interesting avenue for future work.

#### 5.4. Comparison to Prior Mixing Models

Methods to do some form of unmixing of hydrochemical and sedimentary geochemical observations do already exist in a variety of forms (e.g., Christophersen et al., 1990; Kemeny & Torres, 2021; Stock et al., 2018; Weltje, 1997). The most common form is to assume that the observations are mixtures of known source compositions (generally referred to as “endmembers”). For example, Kemeny and Torres (2021) describe the chemistry of river water as a mixture of water derived from weathering of a discrete set of lithologies, each with a distinct chemical composition (e.g., carbonates, silicates, and evaporites). This formulation of the unmixing operation can be solved with least squares. This type of approach is widely used in hydrology (e.g., Christophersen et al., 1990) and for determining the sources of sediment pollution, where it is termed “sediment fingerprinting” (Collins et al., 2020; Walling, 2005). However, the requirement for a priori knowledge of likely endmembers can be a significant burden. This either requires making significant assumptions about sources or the gathering of further data to constrain the endmembers. Approaches have therefore been developed to determine both the likely endmembers and mixing proportions concurrently (Carrera et al., 2004; Shaughnessy et al., 2021; Valder et al., 2012; Weltje, 1997).

However, none of these approaches, unlike the one presented here, consider the spatial distribution of the samples. Our approach is therefore unique in considering the topology of a drainage network when applied to observations nested on a network. The method presented here is most similar to the methods developed by Carraro et al. (2018, 2020) for mapping the sources of downstream eDNA observations in river basins. However, Carraro's approach requires continuous data sets of auxiliary variables across the studied region such as lithology or temperature, on which it trains an optimal transfer function. In cases where auxiliary data is absent, or not relevant to the target tracer (such as a point-source pollutant), such an approach cannot be used. By contrast, our approach requires only very limited upstream information.

#### 5.5. Further Work

The approach as defined here is only able to consider tracers which can be considered conservative on the timescales of mixing. However, it could be extended to take into account non-conservative tracers under certain assumptions. For example, if the concentration of a non-conservative tracer can be modeled to decay exponentially with time spent in transport, this can be added as a prefactor to the mixing calculation, making assumptions regarding transit times. Carraro et al. (2018) take such an approach for determining the sources of non-conservative eDNA in drainage basins. This approach would require the calibration of this decay hyper-parameter potentially using field measurements (e.g., Guillet et al., 2019). Elsewise, it could be calibrated using a held-out subset similar to the regularization parameter. Such an approach could be used for determining the sources of a number of significant non-conservative pollutants such as nitrate and phosphate.

At present we consider only well-mixed, conservative tracers at steady state (i.e., their concentration does not vary significantly in time). This assumption is valid only when the sources of a tracer vary slowly with respect to the tracer travel-time. However, hydrochemical data is observed to sometimes vary transiently over storm-events (Knapp et al., 2020). Additionally, a pollutant released from a point-source for a small duration may also result in transient increases in concentration downstream (e.g., waste water released from sewage overflows; Munro et al., 2019). Consequently, a scheme able to partition the sources of tracers in space and time given nested

time-series observations would be desirable (see Benettin et al., 2022). This development would require both time-series hydrochemical observations (e.g., using in-situ sensors; Rode et al., 2016) and a model of the drainage network hydrodynamics. Alternatively, composite samples could be gathered over an extended period of time, averaging out temporal variability (e.g., Cassidy et al., 2018). Unlike the static model presented in this manuscript, such a spatio-temporal inversion could also be used to objectively probe the pathways of solute transport rather than just their source. By modeling catchment transport using travel time distributions, the relative importance of different hydrological pathways could be explored (Husic et al., 2020; Kirchner et al., 2023; Stenger et al., 2024).

## 6. Conclusions

We describe a very efficient method for unmixing nested concentration data on drainage networks to determine the location and magnitude of the tracer sources. Our approach is applicable to well-mixed, conservative, steady state tracers in both river sediment or water. We show that the unmixing problem can be posed as a convex optimization problem which can thus be solved very quickly while being guaranteed to find a single global solution. This method also scales well to large problems. The rapid solve time allows for uncertainties in the upstream concentration to be calculated using Monte Carlo sampling. Our method can be used for apportioning pollutant sources within river basins, for more accurately constraining geochemical processes, or for generating geochemical maps. Beyond rivers, this approach could be used to analyze nested concentration data on other kinds of networks such as sewerage networks in the field of wastewater epidemiology. We provide an open-source Python implementation of our method that is benchmarked, unit-tested, and accompanied by examples. This implementation requires only a raster of hydrological flow-directions and a data-table giving sample-site locations and associated concentration data. We believe our network-based approach can be extended to consider tracers that are non-conservative, and not in steady-state. Our method therefore provides a strong framework for a wider inverse approach for analyzing geochemical data gathered on drainage networks, as well as in other settings.

## Appendix A: Linear Least-Squares Solution to Unmixing on a Network

In the main manuscript we use our unmixing method to recover the magnitude of a set of source concentrations from their observed counterparts using convex optimization, penalizing *relative* differences between predicted and observed concentrations. In some scenarios, it may be valid to instead penalize *absolute* differences between tracer observations. For example, isotopic geochemical data typically has very low log-variances and so absolute differences between observations may be an appropriate approximation (Blondes et al., 2016). Here we show that if absolute differences are penalized, the unmixing problem can be posed as a linear least squares matrix inversion.

We consider the directed acyclic graph shown in Figure 6 with sample sites 1–6. This network of sample sites,  $G$ , has observed concentrations  $\mathbf{d} = [d_1, d_2, \dots, d_6]$ , source concentrations  $\mathbf{c} = [c_1, c_2, \dots, c_6]$ , and total fluxes  $\mathbf{q} = [q_1, q_2, \dots, q_6]$ . As described in Equation 3, we can construct expressions for  $\mathbf{d}$  as mixtures of the source concentrations  $\mathbf{c}$ , weighted by the fluxes  $\mathbf{q}$ :

$$\begin{aligned} d_1 &= c_1 \\ d_2 &= c_2 \\ d_3 &= c_3 \\ d_4 &= \frac{q_1 c_1 + q_2 c_2 + q_4 c_4}{q_1 + q_2 + q_4} \\ d_5 &= \frac{q_3 c_3 + q_5 c_5}{q_3 + q_5} \\ d_6 &= \frac{q_1 c_1 + q_2 c_2 + q_3 c_3 + q_4 c_4 + q_5 c_5 + q_6 c_6}{q_1 + q_2 + q_3 + q_4 + q_5 + q_6} \end{aligned}$$

This set of linear equations can be represented in matrix form:

$$\begin{pmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ \frac{q_1}{\sum_{i=1,2,4} q_i} & \frac{q_2}{\sum_{i=1,2,4} q_i} & 0 & \frac{q_4}{\sum_{i=1,2,4} q_i} & 0 & 0 \\ 0 & 0 & \frac{q_3}{\sum_{i=3,5} q_i} & 0 & \frac{q_5}{\sum_{i=3,5} q_i} & 0 \\ \frac{q_1}{\sum_{i=1\dots 6} q_i} & \frac{q_2}{\sum_{i=1\dots 6} q_i} & \frac{q_3}{\sum_{i=1\dots 6} q_i} & \frac{q_4}{\sum_{i=1\dots 6} q_i} & \frac{q_5}{\sum_{i=1\dots 6} q_i} & \frac{q_6}{\sum_{i=1\dots 6} q_i} \end{pmatrix} \begin{pmatrix} c_1 \\ c_2 \\ c_3 \\ c_4 \\ c_5 \\ c_6 \end{pmatrix},$$

which is the linear matrix equation  $\mathbf{d} = \mathbf{M}\mathbf{c}$ , that can be inverted for  $\mathbf{c}$ , if  $\mathbf{M}$  is known. We therefore seek an expression for each entry in the matrix  $\mathbf{M}$ .

Note that  $M_{ij}$  is non-zero only when either  $i = j$  or node  $j$  is upstream of node  $i$ . We therefore define a 0–1 matrix  $\Theta$  which indicates whether an entry in  $\mathbf{M}$  is zero or not.  $\Theta$  is simply the path matrix of the graph  $G$  where every node also has a self-edge; that is  $A_{ij} = 1$  if a path exists from node  $j$  to node  $i$  and every node has a path to itself. With  $\Theta$  defined we can thus build the following expression for  $\mathbf{M}$ :

$$M_{ij} = \frac{\Theta_{ij}q_j}{\sum_{k=1}^n \Theta_{ik}q_k}. \quad (\text{A1})$$

We thus have an expression in the simple matrix form:

$$\mathbf{M}\mathbf{c} = \mathbf{d} \quad (\text{A2})$$

Therefore, assuming that we have a vector of downstream observations  $\mathbf{d}$ , and that  $\mathbf{M}$  can be calculated using Equation A1 we want to solve Equation A2 for  $\mathbf{c}$ . Since we have a linear system and  $\mathbf{M}$  has at least as many rows as columns, we have satisfied the conditions for using least squares to solve this system. As each entry in  $\mathbf{c}$  is strictly bounded between 0 and 1 this is a constrained least squares minimization problem:

$$\begin{aligned} &\underset{\mathbf{c} \in \mathbb{R}^n}{\text{minimize}} \quad \|\mathbf{M}\mathbf{c} - \mathbf{d}\|_2 \\ &\text{subject to } 0 < c_i < 1. \end{aligned}$$

## Data Availability Statement

Our Python package and relevant data is available from [github.com/r-barnes/faster-unmixer](https://github.com/r-barnes/faster-unmixer) and archived at [zenodo.org/record/8299812](https://zenodo.org/record/8299812) (Barnes & Lipp, 2023).

## References

- Agrawal, A., Amos, B., Barratt, S., Boyd, S., Diamond, S., & Kolter, J. Z. (2019). *Differentiable convex optimization layers advances in neural information processing systems* (Vol. 32). Curran Associates, Inc.
- Aitchison, J. (1986). *The statistical analysis of compositional data*. Chapman and Hall.
- Barnes, R. (2016a). Parallel Priority-Flood depression filling for trillion cell digital elevation models on desktops or clusters. *Computers and Geosciences*, 96, 56–68. <https://doi.org/10.1016/j.cageo.2016.07.001>
- Barnes, R. (2016b). RichDEM: Terrain analysis software. Retrieved from <http://github.com/r-barnes/richdem>
- Barnes, R., Lehman, C., & Mulla, D. (2014). Priority-flood: An optimal depression-filling and watershed-labeling algorithm for digital elevation models. *Computers and Geosciences*, 62, 117–127. <https://doi.org/10.1016/j.cageo.2013.04.024>
- Barnes, R., & Lipp, A. (2023). Identifying tracer and pollutant sources in drainage networks from point observations using an efficient convex unmixing scheme [dataset]. *Zenodo*. <https://doi.org/10.5281/zenodo.8299812>
- Baronas, J. J., Torres, M. A., Clark, K. E., & West, A. J. (2017). Mixing as a driver of temporal variations in river hydrochemistry: 2. Major and trace element concentration dynamics in the andes-amazon transition. *Water Resources Research*, 53(4), 3120–3145. <https://doi.org/10.1002/2016WR019729>

## Acknowledgments

RB was funded by the Grace Hopper Postdoctoral Fellowship in Computing Sciences from Lawrence Berkeley National Lab, USA. AGL was funded by a Junior Research Fellowship from Merton College, University of Oxford, UK. This study benefited from feedback from Gareth Roberts, Kajetan Chrapkiewicz, Megan Goodell and the comments of anonymous peer reviewers. The authors met at the 2022 CSDMS Annual Meeting in Boulder, CO.



- Benettin, P., Rodriguez, N. B., Sprenger, M., Kim, M., Klaus, J., Harman, C. J., et al. (2022). Transit time estimation in catchments: Recent developments and future directions. *Water Resources Research*, 58(11), e2022WR033096. <https://doi.org/10.1029/2022WR033096>
- Blondes, M. S., Engle, M. A., & Geboy, N. J. (2016). A practical guide to the use of major elements, trace elements, and isotopes in compositional data analysis: Applications for deep formation brine geochemistry. In J. A. Martín-Fernández & S. Thió-Henestrosa (Eds.), *Compositional data analysis* (pp. 13–29). Springer International Publishing.
- Bouchez, J., Lajeunesse, E., Gaillardet, J., France-Lanord, C., Dutra-Maia, P., & Maurice, L. (2010). Turbulent mixing in the Amazon River: The isotopic memory of confluences. *Earth and Planetary Science Letters*, 290(1), 37–43. <https://doi.org/10.1016/j.epsl.2009.11.054>
- Bowes, M. J., Smith, J. T., Jarvie, H. P., & Neal, C. (2008). Modelling of phosphorus inputs to rivers from diffuse and point sources. *Science of the Total Environment*, 395(2), 125–138. <https://doi.org/10.1016/j.scitotenv.2008.01.054>
- Boyd, S. P., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.
- Bufe, A., Hovius, N., Emberson, R., Rugenstein, J. K. C., Galy, A., Hassenruck-Gudipati, H. J., & Chang, J.-M. (2021). Co-variation of silicate, carbonate and sulfide weathering drives CO<sub>2</sub> release with erosion. *Nature Geoscience*, 14(4), 211–216. <https://doi.org/10.1038/s41561-021-00714-3>
- Carraro, L., Hartikainen, H., Jokela, J., Bertuzzo, E., & Rinaldo, A. (2018). Estimating species distribution and abundance in river networks using environmental DNA. *Proceedings of the National Academy of Sciences*, 115(46), 11724–11729. <https://doi.org/10.1073/pnas.1813843115>
- Carraro, L., Mächler, E., Wüthrich, R., & Altermatt, F. (2020). Environmental DNA allows upscaling spatial patterns of biodiversity in freshwater ecosystems. *Nature Communications*, 11(1), 3585. <https://doi.org/10.1038/s41467-020-17337-8>
- Carrera, J., Vázquez-Suñé, E., Castillo, O., & Sánchez-Vila, X. (2004). A methodology to compute mixing ratios with uncertain end-members. *Water Resources Research*, 40(12). <https://doi.org/10.1029/2003WR002263>
- Cassidy, R., Jordan, P., Bechmann, M., Kronvang, B., Kyllmar, K., & Shore, M. (2018). Assessments of composite and discrete sampling approaches for water quality monitoring. *Water Resources Management*, 32(9), 3103–3118. <https://doi.org/10.1007/s11269-018-1978-5>
- Christophersen, N., Neal, C., Hooper, R. P., Vogt, R. D., & Andersen, S. (1990). Modelling streamwater chemistry as a mixture of soilwater end-members — A step towards second-generation acidification models. *Journal of Hydrology*, 116(1), 307–320. [https://doi.org/10.1016/0022-1694\(90\)90130-P](https://doi.org/10.1016/0022-1694(90)90130-P)
- Collins, A. L., Blackwell, M., Boeckx, P., Chivers, C.-A., Emelko, M., Evrard, O., et al. (2020). Sediment source fingerprinting: Benchmarking recent outputs, remaining challenges and emerging themes. *Journal of Soils and Sediments*, 20(12), 4160–4193. <https://doi.org/10.1007/s11368-020-02755-4>
- Constable, S. C., Parker, R. L., & Constable, C. G. (1987). Occam's inversion; a practical algorithm for generating smooth models from electromagnetic sounding data. *Geophysics*, 52(3), 289–300. <https://doi.org/10.1190/1.1442303>
- Diamond, S., & Boyd, S. (2016). Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research : JMLR*, 17, 83.
- Domahidi, A., Chu, E., & Boyd, S. (2013). Ecos: An SOCP solver for embedded systems. In *2013 European control conference (ECC)* (pp. 3071–3076). <https://doi.org/10.23919/ECC.2013.6669541>
- Eschenfelder, J., Lipp, A. G., & Roberts, G. G. (2023). Quantifying excess heavy metal concentrations in drainage basins using conservative mixing models. *Journal of Geochemical Exploration*, 248, 107178. <https://doi.org/10.1016/j.gexplo.2023.107178>
- Gaillardet, J., Dupré, B., Louvat, P., & Allègre, C. J. (1999). Global silicate weathering and CO<sub>2</sub> consumption rates deduced from the chemistry of large rivers. *Chemical Geology*, 159(1), 3–30. [https://doi.org/10.1016/S0009-2541\(99\)00031-5](https://doi.org/10.1016/S0009-2541(99)00031-5)
- Garrett, R., Reimann, C., Smith, D., & Xie, X. (2008). From geochemical prospecting to international geochemical mapping: A historical overview. *Geochemistry: Exploration, Environment, Analysis*, 8(3–4), 205–217. <https://doi.org/10.1144/1467-7873/08-174>
- Garzanti, E., Padoan, M., Setti, M., López-Galindo, A., & Villa, I. M. (2014). Provenance versus weathering control on the composition of tropical river mud (southern Africa). *Chemical Geology*, 366, 61–74. <https://doi.org/10.1016/j.chemgeo.2013.12.016>
- Goulart, P., & Chen, Y. (2023). *Clarabel*. University of Oxford: Oxford Control Group. Retrieved from <https://oxfordcontrol.github.io/ClarabelDocs/>
- Grant, M., Boyd, S., & Ye, Y. (2006). (Vol. 84, pp. 155–210). *Disciplined convex ProgrammingGlobal optimization*. Kluwer Academic Publishers. (Series Title: Nonconvex Optimization and Its Applications). [https://doi.org/10.1007/0-387-30528-9\\_7](https://doi.org/10.1007/0-387-30528-9_7)
- Guillet, G., Knapp, J. L. A., Merel, S., Cirpka, O. A., Grathwohl, P., Zwiener, C., & Schwientek, M. (2019). Fate of wastewater contaminants in rivers: Using conservative-tracer based transfer functions to assess reactive transport. *Science of the Total Environment*, 656, 1250–1260. <https://doi.org/10.1016/j.scitotenv.2018.11.379>
- Gurobi. (2023). *Gurobi optimizer reference manual*. Gurobi Optimization LLC. Retrieved from <https://www.gurobi.com>
- Hack, J. T. (1957). *Studies of longitudinal stream profiles in Virginia and Maryland*. USGS. USGS Numbered Series No. 294B).
- Hagberg, A. A., Schult, D. A., & Swart, P. J. (2008). Exploring network structure, dynamics, and function using NetworkX. In G. Varoquaux, T. Vaught, & J. Millman (Eds.), *Proceedings of the 7th Python in science conference* (pp. 11–15).
- Husic, A., Fox, J., Mahoney, T., Gerlitz, M., Pollock, E., & Backus, J. (2020). Optimal transport for assessing nitrate source-pathway connectivity. *Water Resources Research*, 56(10), e2020WR027446. <https://doi.org/10.1029/2020WR027446>
- Jager, A. L. d., & Vogt, J. V. (2010). Development and demonstration of a structured hydrological feature coding system for Europe. *Hydrological Sciences Journal*, 55(5), 661–675. <https://doi.org/10.1080/02626667.2010.490786>
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An introduction to statistical learning* (Vol. 112). Springer.
- Johnes, P. J. (1996). Evaluation and management of the impact of land use change on the nitrogen and phosphorus load delivered to surface waters: The export coefficient modelling approach. *Journal of Hydrology*, 183(3), 323–349. [https://doi.org/10.1016/0022-1694\(95\)02951-6](https://doi.org/10.1016/0022-1694(95)02951-6)
- Johnson, C. C., Breward, N., Ander, E. L., & Ault, L. (2005). G-BASE: Baseline geochemical mapping of Great Britain and northern Ireland. *Geochemistry: Exploration, Environment, Analysis*, 5(4), 347–357. <https://doi.org/10.1144/1467-7873/05-070>
- Kemeny, P. C., & Torres, M. A. (2021). Presentation and applications of mixing elements and dissolved isotopes in rivers (MEANDIR), a customizable MATLAB model for Monte Carlo inversion of dissolved river chemistry. *American Journal of Science*, 321(5), 579–642. <https://doi.org/10.2475/05.2021.03>
- Kilaru, P., Hill, D., Anderson, K., Collins, M. B., Green, H., Kmush, B. L., & Larsen, D. A. (2023). Wastewater surveillance for infectious disease: A systematic review. *American Journal of Epidemiology*, 192(2), 305–322. <https://doi.org/10.1093/aje/kwac175>
- Kirchner, J. W., Benettin, P., & van Meerveld, I. (2023). Instructive surprises in the hydrological functioning of landscapes. *Annual Review of Earth and Planetary Sciences*, 51(1), 277–299. <https://doi.org/10.1146/annurev-earth-071822-100356>
- Knapp, J. L. A., von Freyberg, J., Studer, B., Kiewiet, L., & Kirchner, J. W. (2020). Concentration–discharge relationships vary among hydrological events, reflecting differences in event characteristics. *Hydrology and Earth System Sciences*, 24(5), 2561–2576. <https://doi.org/10.5194/hess-24-2561-2020>

- LeCun, Y. A., Bottou, L., Orr, G. B., & Müller, K.-R. (2012). Efficient BackProp. In G. Montavon, G. B. Orr, & K.-R. Müller (Eds.), *Neural networks: Tricks of the trade* (2nd ed., pp. 9–48). Springer. [https://doi.org/10.1007/978-3-642-35289-8\\_3](https://doi.org/10.1007/978-3-642-35289-8_3)
- Lindsay, J. B., Rothwell, J. J., & Davies, H. (2008). Mapping outlet points used for watershed delineation onto DEM-derived stream networks. *Water Resources Research*, 44(8). <https://doi.org/10.1029/2007WR006507>
- Lipp, A. G., de Caritat, P., & Roberts, G. G. (2023). Geochemical mapping by unmixing alluvial sediments: An example from northern Australia. *Journal of Geochemical Exploration*, 248, 107174. <https://doi.org/10.1016/j.gexplo.2023.107174>
- Lipp, A. G., Roberts, G. G., Whittaker, A. C., Gowing, C. J. B., & Fernandes, V. M. (2020). River sediment geochemistry as a conservative mixture of source regions: Observations and predictions from the Cairngorms, UK. *Journal of Geophysical Research: Earth Surface*, 125(12), e2020JF005700. <https://doi.org/10.1029/2020JF005700>
- Lipp, A. G., Roberts, G. G., Whittaker, A. C., Gowing, C. J. B., & Fernandes, V. M. (2021). Source region geochemistry from unmixing downstream sedimentary elemental compositions. *Geochemistry, Geophysics, Geosystems*, 22(10), e2021GC009838. <https://doi.org/10.1029/2021GC009838>
- Lizaga, I., Gaspar, L., Blake, W. H., Latorre, B., & Navas, A. (2019). Fingerprinting changes of source apportionments from mixed land uses in stream sediments before and after an exceptional rainstorm event. *Geomorphology*, 341, 216–229. <https://doi.org/10.1016/j.geomorph.2019.05.015>
- MacIver, D. R., Hatfield-Dodds, Z., & Contributors, M. O. (2019). Hypothesis: A new approach to property-based testing. *Journal of Open Source Software*, 4(43), 1891. <https://doi.org/10.21105/joss.01891>
- Martín-Fernández, J. A., Barceló-Vidal, C., & Pawłowsky-Glahn, V. (2003). Dealing with zeros and missing values in compositional data sets using nonparametric imputation. *Mathematical Geology*, 35(3), 253–278. <https://doi.org/10.1023/A:1023866030544>
- Munro, K., Martins, C. P. B., Loewenthal, M., Comber, S., Cowan, D. A., Pereira, L., & Barron, L. P. (2019). Evaluation of combined sewer overflow impacts on short-term pharmaceutical and illicit drug occurrence in a heavily urbanised tidal river catchment (London, UK). *Science of the Total Environment*, 657, 1099–1111. <https://doi.org/10.1016/j.scitotenv.2018.12.108>
- Nourinejad, M., Berman, O., & Larson, R. C. (2021). Placing sensors in sewer networks: A system to pinpoint new cases of coronavirus. *PLoS One*, 16(4), e0248893. <https://doi.org/10.1371/journal.pone.0248893>
- O’Callaghan, J. F., & Mark, D. M. (1984). The extraction of drainage networks from digital elevation data. *Computer Vision, Graphics, and Image Processing*, 28(3), 323–344. [https://doi.org/10.1016/S0734-189X\(84\)80011-0](https://doi.org/10.1016/S0734-189X(84)80011-0)
- O’Donoghue, B., Chu, E., Parikh, N., & Boyd, S. (2016). Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169(3), 1042–1068. <https://doi.org/10.1007/s10957-016-0892-3>
- Oreskes, N., Shrader-Frechette, K., & Belitz, K. (1994). Verification, validation, and confirmation of numerical models in the earth Sciences. *Science*, 263(5147), 641–646. <https://doi.org/10.1126/science.263.5147.641>
- Pawłowsky-Glahn, V., & Egozcue, J. J. (2006). Compositional data and their analysis: An introduction. *Geological Society, London, Special Publications*, 264(1), 1–10. <https://doi.org/10.1144/gsl.sp.2006.264.01.01>
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). *Numerical recipes in FORTRAN: The art of scientific computing* (2nd ed.). Cambridge University Press.
- Rode, M., Wade, A. J., Cohen, M. J., Hensley, R. T., Bowes, M. J., Kirchner, J. W., et al. (2016). Sensors in the stream: The high-frequency wave of the present. *Environmental Science and Technology*, 50(19), 10297–10307. <https://doi.org/10.1021/acs.est.6b02155>
- Salvador, S., & Chan, P. (2004). Determining the number of clusters/segments in hierarchical clustering/segmentation algorithms. In *16th IEEE international conference on tools with artificial intelligence* (pp. 576–584). <https://doi.org/10.1109/ICTAI.2004.50>
- Schubert, E. (2023). Stop using the elbow criterion for k-means and how to choose the number of clusters instead. *SIGKDD Explor*, 25(1), 36–42. <https://doi.org/10.1145/3606274.3606278>
- Sedgewick, R., & Wayne, K. D. (2011). *Algorithms* (4th ed. ed.). Addison-Wesley.
- Shaughnessy, A. R., Gu, X., Wen, T., & Brantley, S. L. (2021). Machine learning deciphers CO<sub>2</sub> sequestration and subsurface flowpaths from stream chemistry. *Hydrology and Earth System Sciences*, 25(6), 3397–3409. <https://doi.org/10.5194/hess-25-3397-2021>
- Sims, N., & Kasprzyk-Hordern, B. (2020). Future perspectives of wastewater-based epidemiology: Monitoring infectious disease spread and resistance to the community level. *Environment International*, 139, 105689. <https://doi.org/10.1016/j.envint.2020.105689>
- Spadoni, M., Cavarretta, G., & Patera, A. (2004). Cartographic techniques for mapping the geochemical data of stream sediments: The “sample catchment basin” approach. *Environmental Geology*, 45(5), 593–599. <https://doi.org/10.1007/s00254-003-0926-7>
- Stenger, R., Park, J., & Clague, J. (2024). Routine stream monitoring data enables the unravelling of hydrological pathways and transfers of agricultural contaminants through catchments. *Science of the Total Environment*, 912, 169370. <https://doi.org/10.1016/j.scitotenv.2023.169370>
- Stock, B. C., Jackson, A. L., Ward, E. J., Parnell, A. C., Phillips, D. L., & Semmens, B. X. (2018). Analyzing mixing systems using a new generation of Bayesian tracer mixing models. *PeerJ*, 6, e5096. <https://doi.org/10.7717/peerj.5096>
- Stutenbecker, L., Scheuven, D., Hinderer, M., Hornung, J., Petschick, R., Raila, N., & Schwind, E. (2023). Temporal variability of fluvial sand composition: An annual time series from four rivers in SW Germany. *Journal of Geophysical Research: Earth Surface*, 128(6), e2023JF007138. <https://doi.org/10.1029/2023JF007138>
- Sugar, C. A., & James, G. M. (2003). Finding the number of clusters in a dataset. *Journal of the American Statistical Association*, 98(463), 750–763. <https://doi.org/10.1198/016214503000000666>
- Torres, M. A., Baronas, J. J., Clark, K. E., Feakins, S. J., & West, A. J. (2017). Mixing as a driver of temporal variations in river hydrochemistry: 1. Insights from conservative tracers in the andes-Amazon transition. *Water Resources Research*, 53(4), 3102–3119. <https://doi.org/10.1002/2016WR019733>
- Udell, M., Mohan, K., Zeng, D., Hong, J., Diamond, S., & Boyd, S. (2014). Convex optimization in Julia. In *2014 first workshop for high performance technical computing in dynamic languages* (pp. 18–28). <https://doi.org/10.1109/HPTCDL.2014.5>
- Valder, J. F., Long, A. J., Davis, A. D., & Kenner, S. J. (2012). Multivariate statistical approach to estimate mixing proportions for unknown end members. *Journal of Hydrology*, 460–461, 65–76. <https://doi.org/10.1016/j.jhydrol.2012.06.037>
- von Blanckenburg, F. (2005). The control mechanisms of erosion and weathering at basin scale from cosmogenic nuclides in river sediment. *Earth and Planetary Science Letters*, 237(3), 462–479. <https://doi.org/10.1016/j.epsl.2005.06.030>
- Walling, D. E. (2005). Tracing suspended sediment sources in catchments and river systems. *Science of the Total Environment*, 344(1), 159–184. <https://doi.org/10.1016/j.scitotenv.2005.02.011>
- Weltje, G. J. (1997). End-member modeling of compositional data: Numerical-statistical algorithms for solving the explicit mixing problem. *Mathematical Geology*, 29(4), 503–549. <https://doi.org/10.1007/BF02775085>