Checking up on the neighbors: Quantifying uncertainty in relative event location

Oleg V. Poliannikov and Alison E. Malcolm, MIT
Michael Prange and Hugues Djikpesse, Schlumberger-Doll Research

With high-permeability hydrocarbon reservoirs exhausting their potential, developing low-permeability reservoirs is becoming of increasing importance. In order to be produced economically, these reservoirs need to be stimulated to increase their permeability. Hydraulic fracturing is a technique used to do this. A mixture of water, additives, and proppants is injected under high pressure into the subsurface; this fluid fractures the rock, creating additional pathways for the oil or gas. Understanding the nature of the resulting fracture system, including the geometry, size, and orientation of individual fractures, as well as the distance from one fracture to the next, is key to answering important practical questions such as: What is the affected reservoir volume? Where should we fracture next? What are the optimal locations for future production wells?

The failure of the rock produces microseismic events. Because of the lack of virtually any other method of locating the induced fractures, the locations of these events are frequently used as proxies for the fracture locations. To obtain accurate microseismic event locations, one or more arrays of receivers are placed in boreholes or on the surface.

Nearly all event-location techniques rely on somewhat calibrated assumptions about the physical model of the subsurface (e.g., velocity and density), the physics of wave propagation, the recording system, etc. Because these assumptions are always approximate, so are the final event locations. We, therefore, seek to obtain not just point estimates of the event locations but also to quantify the uncertainty of these estimates.

Ideally, the goal is to obtain a multidimensional joint probability distribution of the locations of all recorded events. This distribution can then be used to measure the likelihood that microseismic events occurred at their estimated locations. In addition to providing estimates of the microseismic event locations, along with their individual uncertainties, the joint probability distribution captures the correlations between these locations.

All properties of the set of microseismic events, along with their associated uncertainties, can then be computed from this joint distribution. The location of an individual event, for example, is given by a marginal distribution, which is the integral of the joint distribution over the locations of all of the other events. The size of a fracture can be inferred from the maximum distance between events within that same fracture. The fracture spacing can be estimated using the average distance between events from two nearby fractures.

A complete description of the fracture system includes the absolute locations of the fractures as well as their relative geometry. This information is traditionally obtained from the inversion of traveltimes from the receivers to the event

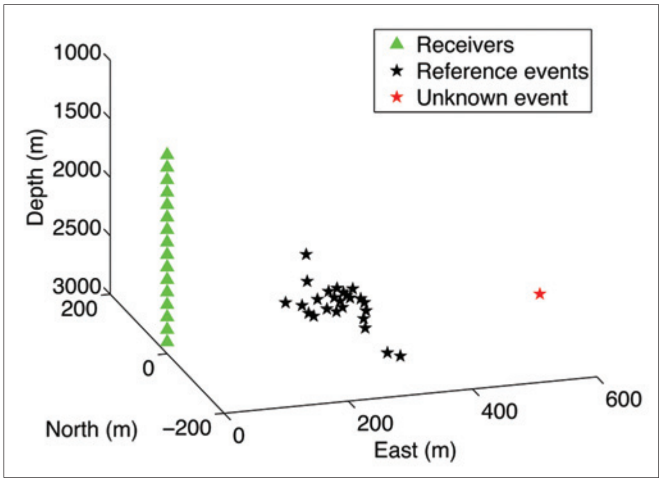

Figure 1. Events in a reference fracture (black stars) and an event (red star) in a subsequent fracture. The receivers are indicated by green triangles.

hypocenters and a hodogram analysis of waveform data around first arrivals. When the estimated velocity model has a large uncertainty, the absolute traveltimes may be biased, and/or may carry a large uncertainty. Important properties of the fracture system, such as the fracture size or the fracture spacing, do not depend on the absolute position of the fractures; if the entire fracture system were moved, these quantities would not change.

Obtaining the full description of the joint probability distribution is a complicated and unsolved problem. In this paper, we address an important subproblem of relative event location. Suppose that a number of microseismic events have already been located; we will call them reference events. The reference events may be perforation shots whose locations are known a priori; they can be events triggered and located at a previous fracturing stage, or they can simply be any events whose locations we may temporarily assume known. A natural question then arises: Can we locate other events relative to these reference events?

To give a concrete example, events that are closer to the receivers are typically located more accurately and thus can be located first. Then we can use the locations of these reference events to locate other events that are further away from the monitoring well (Figure 1). Errors in the absolute locations of the reference events may indeed propagate to the absolute locations of subsequently located events. Their relative position, however, would be substantially unaffected by these errors. Thus estimates of fracture size or the fracture spacing can be improved using reference events, even if the accuracy of the absolute locations of the corresponding fractures cannot. In this example, the question becomes: Can we use the

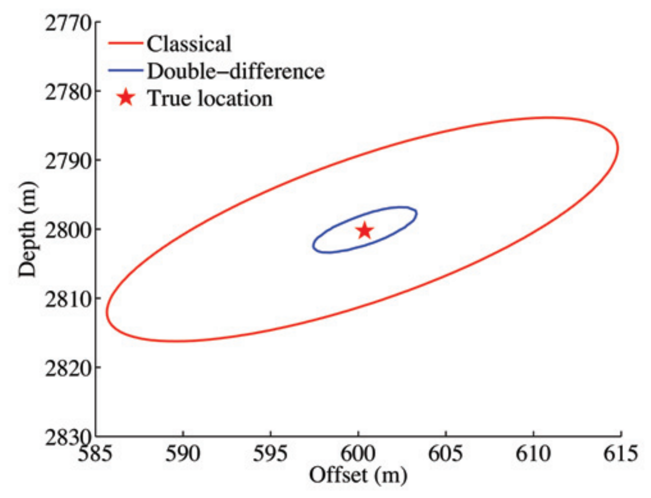

Figure 2. *95% confidence regions for the classical and double-difference location method in a known velocity model. The double-difference method results in much smaller location uncertainty than the classical method because traveltime uncertainty is mitigated not only by averaging over receivers, as in the classical method, but also by averaging over reference-source locations.*

locations of reference events to further constrain the locations of other events in order to derive fracture parameters that are less sensitive to absolute microseismic locations?

In what follows, we present a method of relative event location that uses available information about reference events to produce estimates of the locations of unknown events, along with their uncertainties. This method is general and applicable to any velocity model and any well geometry. By the probability chain rule, solving the problem of relative location in principle solves the problem of full joint location, but we do not address this.

Problem setup

For illustration purposes, we use a specific geometry, which is modeled after a real field experiment, and leads to simple numerical simulations. A vertical monitoring well is instrumented with 16 receivers. We assume that the velocity model is layered, and we provisionally divide it into multilayered overburden and a single production layer. We assume that the physical parameters, including the velocities, of the production layer are known. The characteristics of the overburden are less certain. The medium contains two fractures represented by microseismic events. We assume that we know the locations of events in the first (reference) fracture, and we will attempt to locate an event in the subsequent fracture (Figure 1). The basic methodology that we propose here is fully applicable to other geometries including those with fully 3D velocity models, deviated wells, microseismic events outside of the reservoir layer, etc. Those changes will affect the performance of each individual method of relative location that we discuss in this paper. However, they will not change the way in which that performance is evaluated and compared to other methods.

Our approach to solving the relative location problem is purely kinematic; it is based on fitting the picked traveltimes of seismic events, or on fitting the estimated differences between the traveltimes of seismic event pairs (called the correlation lag because waveform correlation is used to determine the travel-time delay between pairs of events). Both the picked traveltimes and the estimates of the correlation lags are always uncertain, which translates into an uncertainty in the computed event locations. Another primary contributor to uncertainty in the event location is uncertainty in the velocity model. For instance, assuming a layered subsurface model instead of a full 3D model might be a source of uncertainties in the velocity model. Similarly, assuming an isotropic velocity model while strong anisotropy is observed in the observed traveltimes is another source of uncertainties in the velocity model.

In order to quantify the uncertainty in seismic-event location, we need to first quantify both of these sources of uncertainty, namely the estimates of traveltime and the velocity model, by representing them in terms of manageable parameters. Errors in the event locations will then become a function of these simple parameters. Although it at first seems somewhat esoteric, such rigorous uncertainty analysis helps objectively evaluate the quality of the location results; it forces us to make otherwise vague assumptions explicit, and it yields a range of important benefits from identifying the main sources of error to providing clues to a better survey design.

In practice, velocity is estimated using surface seismic studies, perforation shots, and well-log data. For simplicity of presentation, we assume that the uncertainty in the overburden velocity is captured with a single parameter δ: $V = V_0(1 + \delta)$. Here V_0 is the true velocity model (P or S), and V is the velocity with some uncertainties. The velocity inside the production layer is assumed to be constant and known. These assumptions on the velocity model are not necessary but they lead to a simple example that is sufficient to highlight the strengths of the methods of relative location that are discussed below. Note that these methods are applicable to any model of the velocity and its uncertainty so long as they are explicitly provided, as is the case here. The performance of each of these methods for other models can be evaluated using the general framework presented here.

In order to infer wave propagation times from event time picks, we need to know the origin time of each microseismic event. Estimating the origin times and incorporating their uncertainty into the uncertainty of the event location is a problem that we do not attempt to fully tackle here. If the data quality is such that both P and S arrivals can be detected, then the P-S method can be employed to estimate the event origin time. For the velocity uncertainty assumed in this simple study, this method will produce a virtually unbiased estimate of the origin time, which is sufficient for our purposes. Once the origin time is determined, the traveltime from the event origin to the receiver equals the difference between the event picked time and the estimated origin time. If both the origin time and the event picked time are unbiased, then so is the traveltime. In summary, we assume that we have noisy measurements of traveltimes of the form $\hat{T}_{ik} = T(s_i, r_k) + n_{ik}$, where $\hat{T}_{ik}$ are measured time picks, $T(s_i, r_k)$

are true traveltimes in a given velocity, and n_{ik} are zero-mean Gaussian random perturbations.

We now compare three different microseismic event location methods: the classical method that uses the absolute travel times, the double-difference method that uses the time moveout lag between two events, and the recently proposed interferometric method that uses just the stationary phase information contained in that lag.

Classical event location

The classical approach to locating seismic events is to ray trace in the estimated velocity model from the receiver location for a given time. This procedure is repeated for every receiver combining all the traveltimes through triangulation. Assuming the errors in traveltimes are Gaussian with zero mean and standard deviation σ_j, the probability distribution of the event location can be written as

$$p_{CL}(s) \propto \exp\left[-\frac{1}{2}\sum_j\left(\frac{\hat{T}_j - T(s,r_j)}{\sigma_j}\right)^2\right]$$

where the constant of proportionality is chosen to ensure that the density integrates to unity. Averaging over a large number of receivers effectively mitigates the effect of uncertainty in the traveltime measurements, effectively reducing this uncertainty by the square root of the number of receivers (Figure 2). Note that the above formula assumes that the velocity model contains no uncertainty. Velocity uncertainty is treated later, through the marginalization of this formula over the domain of possible velocity models.

Additional information concerning event location, such as the polarization of the incoming wave, can be used to impose additional constraints that further reduce location uncertainty.

Relative location with correlograms

In order to tie the location of one event to the location of another, we use additional information that is inherently present in every recorded data set. For each receiver, we take two traces that correspond to two different events, and we crosscorrelate a window of the waveforms around the respective direct arrivals. Provided that the source mechanisms of both events are sufficiently similar, the correlation will contain a large event. The time of this event is approximately the difference of the traveltimes of the two seismic events. These time differences, or correlation lags, are new data that can be used together with or instead of the absolute traveltimes.

When the velocity is known, the measured correlation lags are unbiased estimates of the true lags. Using all of them in addition to the direct arrival times to constrain the location of the unknown event is referred to as the double-difference localization method. One can see in Figure 2 that the additional constraints significantly reduce the uncertainty of the location of the unknown event resulting in a superior estimate.

Localization in uncertain velocity. Measured traveltimes and lags have been successfully used jointly to localize unknown events relative to reference events in a known velocity model using a well-known approach called the double-difference method. However, if the velocity model is not known exactly then a bias may be introduced. A velocity too slow can underestimate distances, whereas, if it is too fast, then the distances

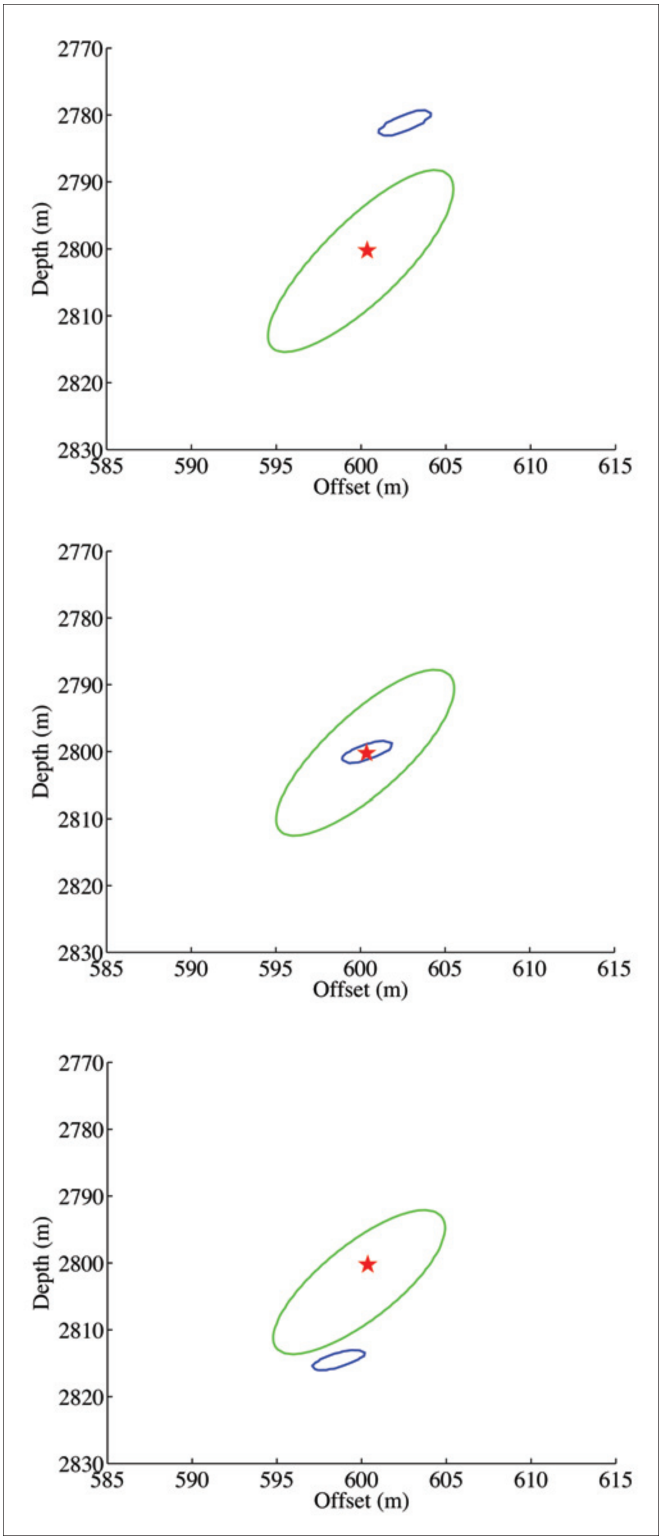

Figure 3. *Effect of overburden velocity model errors ($V = 0.8\ V_0$, $V = V_0$, $V = 1.2V_0$) on event location uncertainty estimates using double-difference (blue) and interferometry (green).*

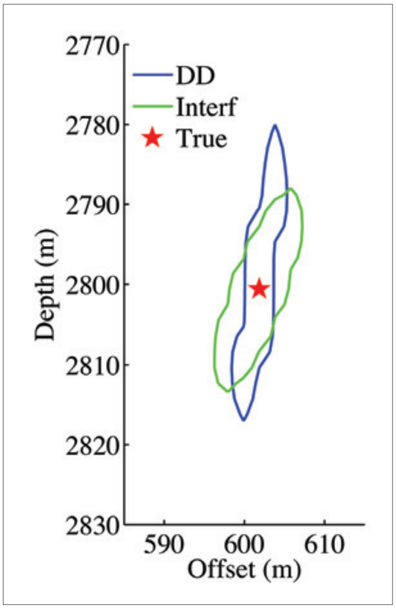

Figure 4. *95% confidence regions for double-difference (blue) and interferometric (green) localization averaged over all admissible velocity models.*

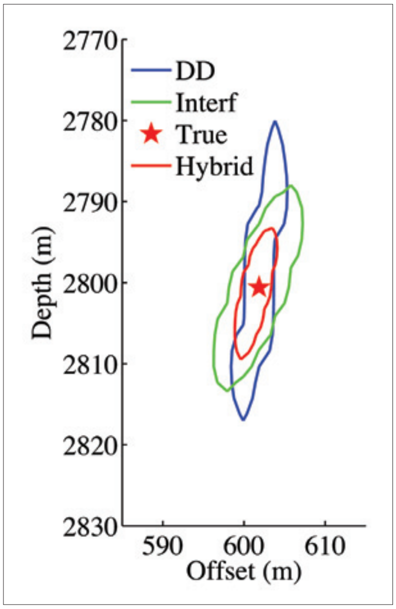

Figure 5. *95% confidence regions for double-difference (blue), interferometric (green), and hybrid localization (red).*

are likely to be overestimated. Although it is possible to imagine a pathological case in which combining different data with different biases would result in an unbiased estimate, we do not expect this to occur in most situations.

Direct traveltimes from the events to the receivers are typically more sensitive to perturbations in the velocity than the correlation lags. This is also true for our numerical model. This means that the rays that connect any two events with a receiver travel through similar media so much, but not all, of the velocity uncertainty is cancelled during crosscorrelation. Therefore, we will henceforth use only lags for relative location. Assuming that the errors in the lags are Gaussian, we can write the location estimate as follows:

$$p_{DD}(s) \propto \exp\left[-\frac{1}{2}\sum_i\sum_j\left(\frac{\hat{\tau}_{ij} - \tau(s_i,s,r_j)}{\xi_{ij}}\right)^2\right]$$

This location estimate is in the same vein as the double-difference method, but uses only the lags and not the traveltimes; we will call it the double-difference method for brevity. Results of this location procedure can be seen in Figures 2 and 3 (blue). In Figure 2, where we assume that the velocity model is known, the double-difference method yields much smaller location uncertainty than the classical method because it adds an average over the reference sources, which is not included in the classical method. In Figure 3 we observe the strong effect of velocity uncertainty, manifested as a bias in the estimate of location uncertainty.

The authors have previously proposed an alternative method of relative location. This method is based on seismic interferometry, and it aims to remove the bias caused by the velocity uncertainty for some experimental geometries. In this technique, instead of fitting all correlation lags to their predicted values, we fit the correlation lag only at the stationary receiver. For the two events, s and s_i, the receiver denoted $r_{i,*}$ is stationary if the lag between the two arrival times is maximal. Using seismic interferometry, we can write the location estimate as follows:

$$p_{INT}(s) \propto \exp\left[-\frac{1}{2}\sum_i\left(\frac{\hat{\tau}_{i,*} - \tau(s_i,s,r_{i,*})}{\xi_{i,*}}\right)^2\right]$$

It can be shown that, for a layered velocity model and a vertical well, the lag computed at the stationary receiver does not depend on the velocity in the overburden and hence is not affected by any uncertainty therein. On the other hand, because only one stationary lag can be used for each pair of events, the total number of lag measurements being averaged is significantly smaller, making interferometry less effective at reducing signal noise. The resulting estimate has a smaller bias but a larger uncertainty as can be seen in Figure 3 (green).

So far we have discussed two methods of relative location, double-difference and interferometry, and we have shown the results achieved by each method in the assumed geometry. These results are certainly specific to the assumptions of the experiment, i.e., the layered structure, the vertical well, and the particular form of the velocity uncertainty and signal noise. However, they serve to illustrate the general link between the uncertainty in the experiment, and the quality of location achieved by different methods.

No method can be a universal remedy. For our model, if the noise is strong but the velocity is known exactly, then the double-difference location provides a superior estimate. Alternatively, if the signal noise is small but the errors in the assumed velocity in the overburden are large, then the interferometric method is preferred.

Velocity marginalization. Comparing the results of the double-difference method to those of the interferometric method, we have identified cases where one method is clearly superior to another. Such a judgment is more difficult to make when the assumptions do not fall into either extreme. Suppose that there is significant uncertainty in the velocity inside the overburden and the recorded signal is also quite noisy. Which is better to have—a biased estimate with a small spread or an unbiased estimate with a large spread?

Adding to the challenge, in practice we may never know that we have underestimated (or overestimated) the velocity; if we did, we would simply correct our assumptions accordingly. We thus do not obtain a reliable estimate of the uncertainty by considering only a single velocity model. We must instead consider all velocity models that are within our estimated error. If we average the location uncertainties over an appropriate sample of velocity realizations, we obtain an estimate of the location uncertainty that is unbiased by the velocity uncertainty. Assume for example that δ defined above is Gaussian with mean 0 and standard deviation 10%. The velocity independent uncertainty regions are shown in Figure 4. Both velocity-independent estimators are unbiased so they can be properly compared.

Hybrid method. We have outlined a framework for evaluating the performance of two different methods of relative event location. This framework treats both velocity uncertainty and noise in the recorded signal, and allows us to evaluate location uncertainty for any given scenario. The double-difference and interferometric methods both use correlation lags. The double-difference method uses all available lags, and the interferometric method uses only the best (stationary) lag for each reference event.

Using these two approaches as extremes, we can construct a more general relative location estimator that uses only the best subset of measured lags and discards the rest. The performance of any location estimator of this kind can be rigorously evaluated following the framework outlined above. In particular, for a given set of reference events and a given unknown event, each method would produce a corresponding confidence region. The method that produces the smallest uncertainty region is the best (Figure 5).

Conclusions

We have put the problem of microseismic event location into a statistical framework. Ideally, we seek to find the joint distribution of the locations of all recorded events. The problem of finding the joint distribution of the event locations can be reduced to the problem of the relative location of one event given a set of already located reference events.

Relative location enables us to obtain more precise estimates of many important properties of the fracture system. For example, the fracture size and fracture spacing, as estimated from microseismic data, depend not on the absolute position of the events, but on the relative position of one microseismic event with respect to another.

The performance of any location method is measured by the uncertainty of the estimator that it produces, and must also consider the bias in that estimator. In order to rigorously analyze the uncertainty of reconstructed event locations, we must explicitly quantify all assumptions about the experiment, such as the geometry, the velocity model and its associated uncertainty, the noise in the signal, etc. We have presented a framework that allows us to analyze the performance of any location method that is based on travel time picks or correlation lags. This analysis is general and is applicable to any velocity model and any well geometry.

Our analysis does not point to a "best" method, but rather gives a framework for evaluating methods. Using this framework, for each set of assumptions, the best method of relative location can be derived by minimizing its uncertainty estimate. TLE

Corresponding author: oleg.poliannikov@gmail.com