What moved where?: The impact of velocity uncertainty on microseismic location and moment-tensor inversion

Oleg Poliannikov[1], Alison Malcolm[2], and Michael Fehler[1]

Abstract

With the rise of unconventional resources, microseismic monitoring is becoming increasingly important because of its cost-effectiveness. This has led to significant research activity on how best to locate events and characterize their moment tensors. Locations tell us where fracturing is occurring, allow the tracking of fluid movement, and fracture propagation. Moment tensors help to determine the type of failure occurring, which is beneficial in planning and interpreting the results of hydraulic-fracturing jobs and in monitoring production. The rising number of methods to determine parameters raises important questions about how uncertainties in the input parameters are translated into uncertainties in the final locations and moment tensors. We present a framework for assessing these uncertainties and use it to demonstrate how velocity uncertainty — as well as uncertainties in arrival times and amplitudes — translates into uncertainties on the recovered quantities of location and moment-tensor parameters.

Introduction

As long as countries like the United States continue to move toward exploiting internal sources of energy, unconventional resources will be an important part of the world's energy portfolio. Microseismic monitoring is one of the most commonly used methods to estimate the reservoir volume, drainage area, and other characteristics of such unconventional reservoirs. Microseismic monitoring is attractive because it is less expensive than imaging, and it can potentially be done in real time. For multistage hydraulic fracturing, such real-time monitoring is advantageous because then the results of the previous stage of fracturing can be used to guide activities during subsequent stages. Most often, the typical output of a microseismic monitoring program is a collection of estimated event locations along with moment tensors (Eisner et al., 2011). Accurate location and origin-time estimations of microseismic events are important because they potentially allow us to monitor changes in the stress field due to the injection and where fractures have been generated or stimulated. Moment tensors give more information about the stress state than do locations alone. They also provide an indication of whether or not faults are active, fractures have opened, and they provide information about the physical fault parameters (strike, dip, and rake) in addition to the event size.

Locations are usually obtained by the processing of arrival times and moment tensors by analyzing the polarities or amplitudes of first-arriving P- and S-waves. There are methods (Li et al., 2011; Nakata and Beroza, 2016; Sharan et al., 2016; Kaderli et al., 2015 and references therein) that use the entire waveform to recover both this information and additional subsurface parameters, but we will not deal with those methods in this paper. All location and moment-tensor inversion (MTI) algorithms rely on velocity information, and thus uncertainty

in the velocity may have a strong impact on the recovered results. From a theoretical viewpoint, it would be best to use the information from arrival times, velocity analyses, and the entire recorded waveform to jointly reduce the uncertainty in event location, MTI, and in the velocity model itself. Doing all these things together is computationally intractable, but in this paper we set up a framework that takes a first step in this direction.

Our framework tells us how to go from estimates of uncertainties in the input data (e.g., arrival times, waveforms, velocities) to estimates of uncertainty in model parameters (e.g., locations, moment tensors, velocities) through application of Bayes' theorem. Bayes' theorem, which we describe in more detail to follow, tells us how to translate uncertainties on our input parameters into the resulting uncertainties on our output parameters. Here, we apply this methodology to determine what the uncertainties are on both locations and focal mechanisms for given uncertainties in velocity. The location part builds on our past work (Poliannikov et al., 2011a, 2011b, 2012, 2014, 2015, 2016; Poliannikov and Djikpesse, 2015; Poliannikov and Malcolm, 2016), which has focused primarily on the location problem and highlights how different location algorithms perform under different assumptions on the uncertainties in both arrival times and velocity. That work shares some similarities with that of Eisner et al. (2009). A recent overview of MTI methods can be found in Cesca et al. (2013) or Gu (2016).

Problem setup

Our goal is to provide a general framework through which uncertainties in velocity and arrival times can be incorporated into a final estimate of locations and moment tensors with uncertainties on each quantity. To illustrate our method, we use a simple problem of processing a single source. We can process many events jointly using straightforward generalizations of methods presented here. Our framework is applicable to both surface and borehole monitoring as long as the direct arrivals from the event are visible. From either recording geometry, we assume that (uncertain) arrival times and amplitudes are picked from recorded data either manually or automatically. When the signal-to-noise ratio (S/N) is low and arrivals are not easily picked, migration-based location algorithms may need to be used, which makes Bayesian analysis more challenging, though some progress can be made (Poliannikov and Malcolm, 2016). Figure 1 shows the source location and monitoring geometry we use for our examples.

Measuring data uncertainties

Velocity uncertainty. Although many velocity-analysis techniques do not come with an explicit error determination, any method that shows an imperfect fit to the data used to determine the velocity or has limitations in the resolution of the velocity due to data distribution can be used to estimate the velocity uncertainty

[1]Massachusetts Institute of Technology.
[2]Memorial University of Newfoundland.

Special Section: Velocity-model uncertainty

in some way. In each case, the Bayesian formulation requires us to make explicit which velocity models we consider possible before we estimate locations and moment tensors from the microseismic data. As the procedure continues, and locations are computed for various scenarios or a joint location-velocity inversion is performed, we may refine our velocity model and its associated uncertainty model. In all cases, however, some initial assumptions must be made about the velocity, and most conventional velocity-analysis methods can provide the necessary information about uncertainty. For example, if velocities are picked through some sort of semblance analysis, the associated semblance functions will have peaks with a width allowing different interpreters to arrive at different models. If velocities instead are taken from an analysis of well logs, then the log data may be blocked or averaged in different ways and must be transformed to compensate for potential differences in what the log measures and what the microseismic wave will see (e.g., different components or frequencies). These procedures can be used to give quantitative information about the range of plausible models. A third method of velocity analysis encompasses velocity-inversion algorithms that seek to minimize some sort of cost function. These algorithms must be stopped at what is often a somewhat arbitrarily chosen point, which means that, again, more than one feasible model can be estimated from this analysis.

Traditionally this uncertainty information is ignored, at least partially, because it is not clear what to do with this information; this is part of what we address in this paper. Quantifying velocity uncertainty means building a prior probability distribution $p(v)$ from which we can sample different velocity models, called velocity realizations, that each represent a plausible velocity model. Simply providing several possible velocity models is a particular case of velocity-uncertainty quantification. These models would make up $p(v)$, and sampling $p(v)$ would entail choosing one of these models at random.

Making the right assumptions about the best choice for $p(v)$ in a given situation is difficult. Underestimating the uncertainty may lead to a bias in the final results, as the true velocity model may not be included in the initial set. By contrast, overestimating the uncertainties is fine in theory but makes further processing more computationally expensive and may result in unrealistically large error bounds on the final estimated parameters.

Arrival time and amplitude-picking uncertainties. Our framework for location and moment-tensor inversion relies on direct arrival times and amplitudes, which may be picked automatically, manually, or through a combination of the two. The quality of both sets of picks depends on many factors, including the source time function, radiation pattern, signal noise, etc. It is often the case that picks are characterized as simply good or bad, which is insufficient for careful uncertainty analysis. Any estimated quantity such as the P-arrival time from a given event to a given station must come with some error bars, which imply a probability distribution for the pick, rather than a single number. One way to do this is to use characteristic functions (see Massin and Malcolm [2016] and references therein).

Estimating model uncertainties

As mentioned above, we will work with a single event to simplify the discussion. We assume that this event is located at x, with moment tensor m embedded in a material with velocities $v = (v_p, v_s)$. To each station, we associate an arrival time, $t = (t_p, t_s)$,
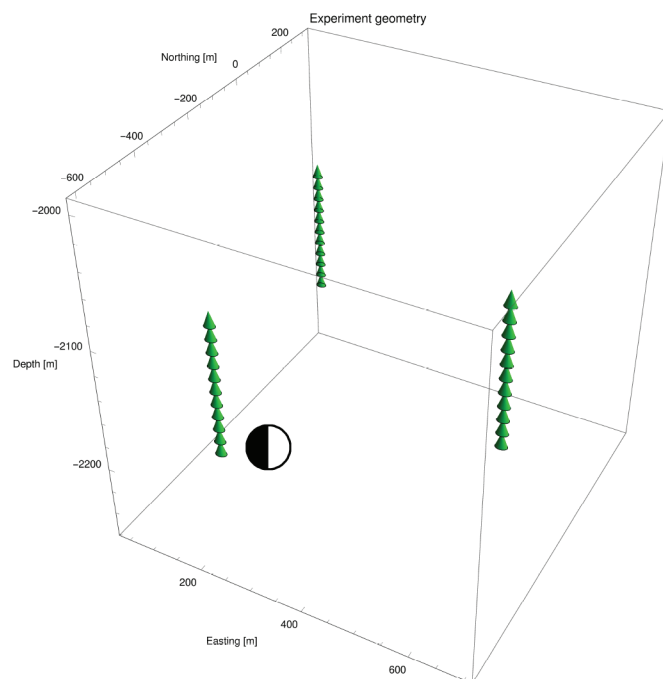

Figure 1. Experimental setup. The triangles are monitoring stations, and the beach ball represents the moment tensor we attempt to recover. We will also recover the location of this event.

and P/S amplitude ratio r. We treat (t, v, r) as random variables, which means they are conceptually characterized by a joint probability distribution. We then attempt to estimate the location and moment tensor and their associated uncertainties using this probabilistic description of the input information. To make this explicit, we briefly outline the theory first for locations given arrival times and then for moment tensors given arrival amplitude ratios.

As one example of a probability distribution, we may characterize each quantity by a mean and a standard deviation. The problem with this assumption is that it implies that the distributions are Gaussian, which we shall see is not the case here. As a result, we must define our probability distributions more generally.

Event location. As mentioned above, we must begin with a probabilistic statement of the problem. We formulate the problem by defining a likelihood function, $p(t|x,v)$, which is read as "the likelihood of observing arrival time t, given the input distributions on x and v." This is a modeled range of arrival times for a given range of locations and velocities. In other words, we calculate $p(t|x,v)$ by computing the arrival times that would be recorded for a microseism at x with velocity model v. For simplicity, we are fixing the origin time; it can be incorporated easily (Poliannikov et al., 2014) but adds to the notational clutter. From this, we would then like to compute the posterior, $p(x|t)$, or the probability of the event being located at x given the distributions of arrival times with the effects of the velocity uncertainty folded in. We will use two tools to obtain this final likelihood: Bayes' theorem and marginalization, described hereafter. We can use these tools to obtain different probabilities for different parameters by simply adjusting which parameters we consider as inputs and which we estimate. For example, we can first compute $p(x,v|t)$, which is the joint probability of observing the location x, and velocity v, given the arrival time data, t, and then, if desired, integrate (or marginalize) this joint posterior over v to get the posterior of just the location.

Bayes' theorem relates a likelihood function (forward model) to the posterior distribution (inversion). Explicitly it tells us that

$$p\left(model \mid data\right) = \frac{p\left(data \mid model\right) p\left(model\right)}{p\left(data\right)}, \quad (1)$$

so given a forward model $p(data|model)$ and a prior distribution of the input parameters $p(model)$, we can probabilistically invert for the input parameters model using the observed output data. To apply equation 1, we must simply define our model and data appropriately. For example, to estimate the posterior, $p(x,v|t)$, we set data = (t), model = (x,v) and obtain

$$p\left(x,v \mid t\right) = \frac{p\left(t \mid x,v\right) p\left(x,v\right)}{p\left(t\right)}. \quad (2)$$

In the expression above, the normalization constant $p(t)$ is the unconditional probability of the observed data that ensures that the left-hand side is a valid probability density that integrates to one. We can then reduce the joint posterior $p(x,v|t)$ to the smaller distribution that we are really interested in, namely $p(x|t)$, by marginalizing (integrating) over our auxiliary variable, v. Marginalizing simply means incorporating all possible outcomes for a particular variable. In our case, we may be interested only in the location and not the velocity, so we take the joint distribution estimated with equation 2, which measures jointly the location and velocity and their associated uncertainties, and estimate a distribution on only the location. To remove the dependence on the velocity, we then calculate

$$p\left(x \mid t\right) = \int p\left(x,v \mid t\right) dv. \quad (3a)$$

Similarly, we can calculate the uncertainty in the velocity model as constrained by the recorded microseismicity:
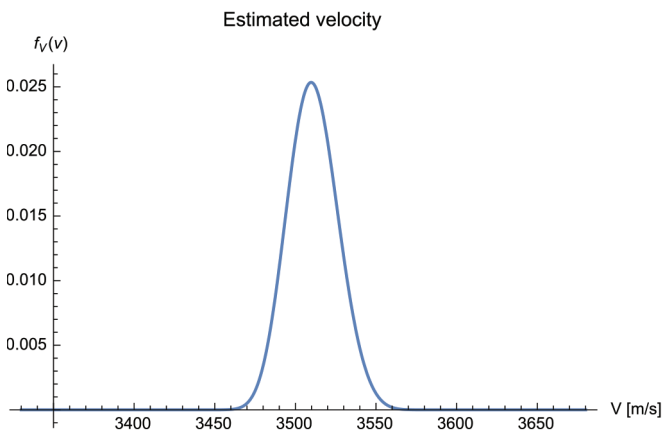
$$p\left(v \mid t\right) = \int p\left(x,v \mid t\right) dx. \quad (3b)$$

This distribution is shown in Figure 2.

The event location algorithm can be summarized as follows. First, we make initial assumptions about the velocity model, along with the associated uncertainty, using data unrelated to a microseismic survey. Second, we detect events and estimate arrival times and their errors from the recorded waveforms. We then model the likelihood in $p(t|x,v)$ using Monte Carlo sampling over the range of models provided in the probability distribution for v and the grid of possible locations. These quantities together allow us to apply equations 2 and then 3 to estimate the posterior distribution for the event location x given t. The results of this process are shown in Figure 3.

Observe that the peak of posterior distribution does not have to coincide with the true location of the event. If it did, it would imply that we could recover the event location precisely, which is generally not possible. As expected, the distributions are larger in the depth direction than in either of the lateral directions; we expect this because the monitoring array has larger extent in both lateral directions than it does in depth. Figure 3b shows 2D uncertainty ellipses as confidence regions. A 50% confidence region means that for 50% of the realizations of velocity and picked arrival times, the event is located somewhere within this ellipse.

Moment–tensor inversion. Having obtained a distribution for the

Figure 2. The posterior distribution of the P-velocity. The true value of the P-velocity is 3500 m/s.
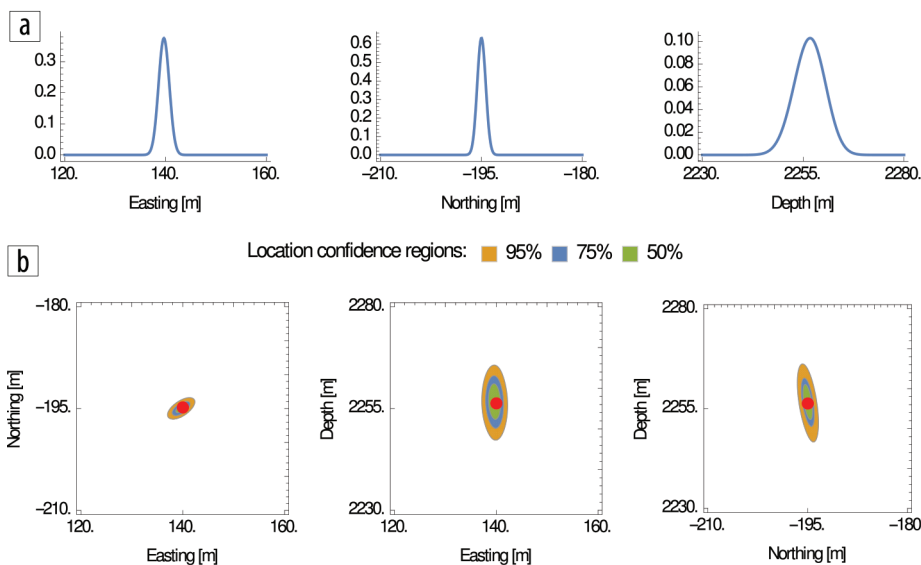

Figure 3. The posterior estimates of locations and their uncertainties, $p(x|t)$. (a) 1D slices of the resulting distribution; (b) 2D confidence ellipses for three different confidence levels. The red dot shows the true microseism location.

Special Section: Velocity-model uncertainty

locations, we now move on to estimate the moment tensor, m, characterized by the strike, dip, and rake of the fault, and calculate the uncertainty of this estimate, given the P/S amplitude ratio, r. We use the amplitude ratio because it is less dependent on velocity errors as well as the differences in the measurement response at different stations, which may come from different instrument types and qualities as well as differences in their immediate environment.

Our basic approach to moment-tensor estimation is the same as it was to the event-location problem: we first estimate our errors on the data, r, and construct a likelihood function (forward model), $p(r|m,x,v)$, that describes what measurements we are likely to observe given certain model parameters. This likelihood along with "priors" on the input parameters allow us to use Bayes'

theorem from equation 1 to calculate the joint posterior distribution, $p(m,x,v|r)$. By integrating over x and v, we obtain the final posterior distribution $p(m|r)$, which gives us the best estimate of the moment tensor as well as the associated uncertainty. Mathematically,

$$p(m|r) \propto \iint p(r|m,x,v)\,p(m)\,p(x,v)dx\,\,dv. \qquad (4)$$

Recall that we are estimating the moment tensor after we have already located the event. The probability density function $p(x,v)$ that acts as a prior here is actually the posterior distribution for the location problem given in equation 2. The moment tensor depends both on the event location and the velocity model, which in turn are correlated through arrival times that probabilistically tie them together. This means that to quantify the moment-tensor uncertainty, we first must solve, in principle, a joint location-velocity estimation problem with uncertainty quantification and then use the solution to that problem in moment-tensor inversion. In this work, we implicitly use an additional assumption that amplitude ratios, r, carry no information about event locations beyond what can already be inferred from arrival times. We believe that this assumption is approximately satisfied for most, if not all, practical setups.

From a workflow perspective, we first model the possible observations, r, given a range of moment tensors, $p(m)$, and locations $p(x,v|t)$, to obtain $p(r|m,x,v)$ using Monte Carlo sampling. Once we have this, we perform a weighted stack over all possible velocity models and moment tensors to obtain our final likelihood, which is that of a particular moment tensor, m, given the observed ratio, r, (at all stations). The weights in this stack are the prior, which consists of two parts. First there is $p(m)$, which is the a priori information (that we have before the experiment is performed); for example, one can look only over double-couple mechanisms, or choose to include a tensile component. In our example we use a flat prior allowing for any combination of strike, dip, and rake within shown ranges. The second part of the prior is the location information $p(x,v|t)$ discussed above. We would like to emphasize that this is not a single prior location and velocity, but a probability distribution that incorporates all of our available information about possible event locations. In this way, we incorporate all of the work done to obtain the location probability into our estimated moment tensor and its likelihood function.

Figure 4. True (left) and recovered (right) beach balls computed from the strike, dip, and rake recovered through the Bayesian process.
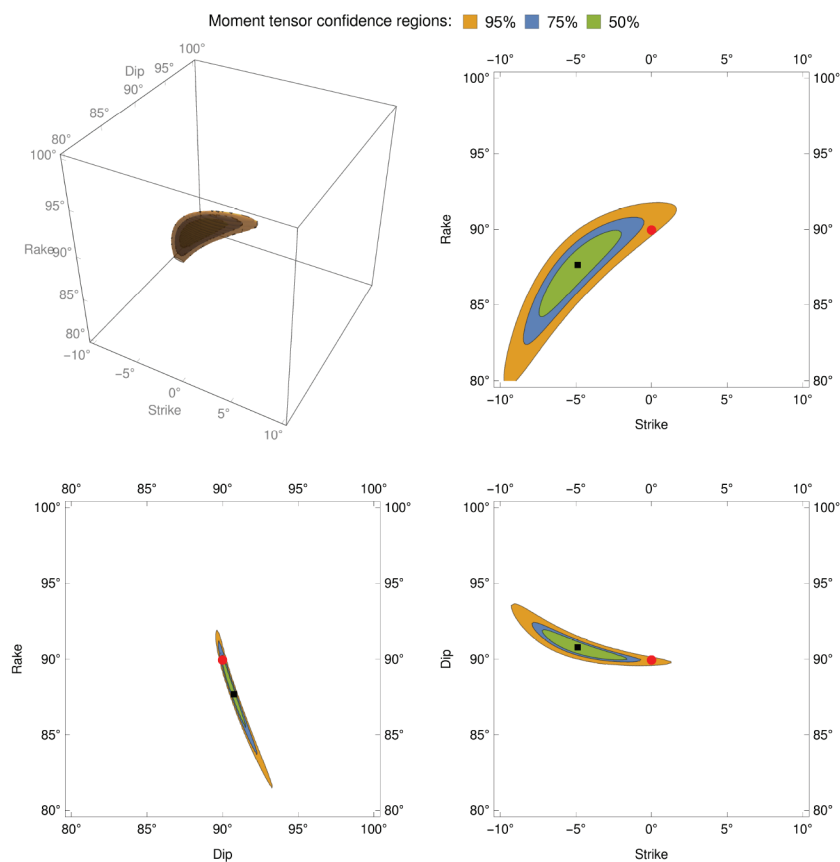

Figure 5. Confidence intervals for the strike, dip, and rake of the event. In this case, we note that the strike and dip are better recovered than is the rake.

We show the results of this procedure in Figures 4 and 5. In Figure 4, we show the true and the maximum likelihood estimate of the moment tensor. From this figure, we see that the recovered moment tensor is very similar to the true moment tensor, indicating that our method is indeed producing the correct answer as expected. However, it provides no information about estimation errors. In Figure 5, we see the additional information that becomes available after a careful uncertainty analysis. We see that we do get estimates of the strike, dip, and rake that are within the 90% confidence region of the true values. The inversion is done for a single realization of the noise in both the velocity and the arrival times, and thus some bias is expected. Now we can also compare the recovery of different parameters. We see that the strike and dip are better recovered than is the rake. This is likely a result of our particular acquisition geometry and should not be taken as a conclusion of the method itself. We also note in this figure that the resulting distribution is clearly not Gaussian. This is a key observation as it tells us that we cannot characterize the resulting uncertainties with a simple error bar or standard deviation, but instead must really visualize the entire probability distribution to understand the uncertainties in the resulting fault parameters.

Discussion and conclusions

We have presented a framework for quantitatively estimating the uncertainties in microseismic event locations and moment tensors beginning from simple assumptions and estimates of the errors on arrival times, velocities, and amplitude ratios. This information is readily available, when translated correctly from standard processing tools.

We show with simple examples that our method is able to correctly recover both location and moment-tensor estimates, but also gives much more information. This new information is the uncertainties in the recovered parameters, which can be used to better constrain decision making, improve our understanding of what our data truly tell us about these parameters, and design a better network for a proposed experiment. Although we have illustrated our framework with relatively simple examples, the basic ideas extend to more complicated situations including non-homogeneous velocities and multiple events. Dealing with the computational cost of moving beyond relatively simple models with a few events and a handful of layers may be challenging. Through smart sampling of a carefully chosen model space and fast algorithms, we believe that more complicated models can be handled within the framework we present.

Throughout the paper, we have treated event locations and moment tensors as ultimate quantities of interest. In reality, these often become inputs to more complicated research that seeks to understand variations in the stress field, fluid flow during and after the injection, and the flow characteristics of a producing reservoir. Understanding the uncertainty in basic properties of microseismic events is critical for understanding the quality and reliability of conclusions made down the line. **TLE**

Corresponding author: poliann@mit.edu

References

Cesca, S., A. Rohr, and T. Dahm, 2013, Discrimination of induced seismicity by full moment tensor inversion and decomposition: Journal of Seismology, **17**, no. 1, 147–163, http://dx.doi.org/10.1007/s10950-012-9305-8.

Eisner, L., P. M. Duncan, W. M. Heigl, and W. R. Keller, 2009, Uncertainties in passive seismic monitoring: The Leading Edge, **28**, no. 6, 648–655, http://dx.doi.org/10.1190/1.3148403.

Eisner, L., M. Thornton, and J. Griffin, 2011, Challenges for microseismic monitoring: 81st Annual International Meeting, SEG, Expanded Abstracts, 1519–1523, http://dx.doi.org/10.1190/1.3627491.

Gu, C., 2016, Ground motions and source mechanisms of earthquakes in multiscales — microseismicity to macroseismicity: PhD thesis, Massachusetts Institute of Technology.

Kaderli, J., M. D. McChesney, and S. E. Minkoff, 2015, Microseismic event estimation in noisy data via full waveform inversion: 85th Annual International Meeting, SEG, Expanded Abstracts, 1159–1164, http://dx.doi.org/10.1190/segam2015-5867154.1.

Li, J., H. S. Kuleli, H. Zhang, and M. N. Toksöz, 2011, Focal mechanism determination of induced microearthquakes in an oil field using full waveforms from shallow and deep seismic networks: Geophysics, **76**, no. 6, WC87–WC101, http://dx.doi.org/10.1190/geo2011-0030.1.

Massin, F., and A. Malcolm, 2016, A better automatic body-wave picker with broad applicability: 86th Annual International Meeting, SEG, Expanded Abstracts, 2617–2621, http://dx.doi.org/10.1190/segam2016-13818011.1.

Nakata, N., and G. C. Beroza, 2016, Reverse time migration for microseismic sources using the geometric mean as an imaging condition: Geophysics, **81**, no. 2, KS51–KS60, http://dx.doi.org/10.1190/geo2015-0278.1.

Poliannikov, O. V., and H. Djikpesse, 2015, What can microseismic data say about fluid flow?: The Leading Edge, **34**, no. 11, 1356–1358, http://dx.doi.org/10.1190/tle34111356.1.

Poliannikov, O. V., A. Malcolm, H. Djikpesse, and M. Prange, 2011a, How can we use one fracture to locate another?: The Leading Edge, **30**, no. 5, 551–555, http://dx.doi.org/10.1190/1.3589115.

Poliannikov, O. V., and A. E. Malcolm, 2016, The effect of velocity uncertainty on migrated reflectors: Improvements from relative-depth imaging: Geophysics, **81**, no. 1, S21–S29, http://dx.doi.org/10.1190/geo2014-0604.1.

Poliannikov, O. V., A. E. Malcolm, H. Djikpesse, and M. Prange, 2011b, Interferometric hydrofracture microseism localization using neighboring fracture: Geophysics, **76**, WC27–WC36, http://dx.doi.org/10.1190/geo2010-0325.1.

Poliannikov, O. V., A. E. Malcolm, M. Prange, and H. Djikpesse, 2012a, Checking up on the neighbors: Quantifying uncertainty in relative event location: The Leading Edge, **31**, no. 12, 1490–1494, http://dx.doi.org/10.1190/tle31121490.1.

Poliannikov, O. V., M. Prange, H. Djikpesse, A. E. Malcolm, and M. Fehler, 2015, Bayesian inversion of pressure diffusivity from microseismicity: Geophysics, **80**, no. 4, M43–M52, http://dx.doi.org/10.1190/geo2014-0374.1.

Poliannikov, O. V., M. Prange, A. E. Malcolm, and H. Djikpesse, 2014, Joint location of microseismic events in the presence of velocity uncertainty: Geophysics, **79**, no. 6, KS51–KS60, http://dx.doi.org/10.1190/geo2013-0390.1.

Sharan, S., F. Herrmann, R. Wang, and T. Van Leeuwen, 2016, Sparsity-promoting joint microseismic source collocation and source-time function estimation: 86th Annual International Meeting, SEG, Expanded Abstracts, 2574–2579, http://dx.doi.org/10.1190/segam2016-13871022.1.