Let M_T denote the maximum travel time for any phase. Initially, we start with an event window of size W from $t_0 = 0$ to $t_1 = W$, and an arrival window of size $W + M_T$ from $t_0 = 0$ to $t_2 = W + M_T$. Then we perform a series of local moves that add or update events in the event window, delete existing events, or classify (as true arrival, false arrival, or coda arrival) the arrivals in the arrival window. Next, the windows are moved forward by a step size, S. At this point, events older than $t_0 - M_T$ become stable: none of the moves modify either the events or arrivals associated with them. These events are then output. While in theory this algorithm never needs to terminate, our experiments continue until the test dataset is fully consumed.

The algorithm's initial hypothesis is that all new arrivals added to the arrival window are false arrivals. This is refined by classifying any arrival a (at station k), with the immediately preceding arrival a^-, as a coda arrival if

$$P_{\gamma,d}(a^-)P_\gamma(a|a^-) > [1 - P_{\gamma,d}(a^-)]P^k_\omega(a).$$

In simple terms, the condition expressed in the previous equation states that it is more likely that the arrival a^- generated a coda and this coda was a than that a^- did not generate a coda and a was a false arrival. This default classification for an arrival is retained whenever it is no longer associated with an event. For convenience we define

$$\Upsilon^k(a) = \max\{P_{\gamma,d}(a^-)P_\gamma(a|a^-), [1 - P_{\gamma,d}(a^-)]P^k_\omega(a)\}. \tag{12}$$

Next, the birth move generates new events in the event window. These events are added to the hypothesis with $\Lambda^{ijk} = \zeta$ for each new event i. Subsequently, we repeat the following N times: one improve-arrival move for each arrival in the arrival window, and one improve-event move for each event in the event window. Finally, the death move kills some of the events, and we repeat one round of improve-arrival and improve-event moves. We describe these steps algorithmically next. The individual moves will be described in more detail later.

1. $t_0 = 0$; $t_1 = W$; $t_2 = W + M_T$.
2. Repeat while $t_0 <$ max time.
    - I. Give a default classification to arrivals in t_0 to t_2.
    - II. Add events from birth move ($t_0, t_1, \{a{:}t_0 \le a_t \le t_2\}$).
    - III. Repeat N times.
        - i. For each arrival a, such that $t_0 \le a_t \le t_2$, improve-arrival (a).
        - ii. For each event e^i, such that $t_0 \le e^i_t \le t_1$, improve-event (e^i).
    - IV. For all events e^i, death move (e^i).
    - V. For each arrival a, such that $t_0 \le a_t \le t_2$, improve-arrival (a).
    - VI. For each event e^i, such that $t_0 \le e^i_t \le t_1$, improve-event (e^i).
    - VII. $t_0 += S$, $t_1 += S$, $t_2 += S$.
    - VIII. Output e^i, Λ^{ijk} for all e^i such that $e^i_t < t_0 - M_T$.
3. Output any remaining e^i.

In order to simplify the computations needed to compare alternate hypotheses, we decompose the overall probability of equation (11) into the contribution from each event. We define the score S_e of an event as the probability ratio of two hypotheses: one in which the event exists, and another in which the event does not exist and all of its associated arrivals have the default classification (false or coda). If an event has score less than 1, an alternative hypothesis in which the event is deleted clearly has higher probability. Critically, this event score is unaffected by other events in the current hypothesis. From equations (7), (8), (9), (10), and (12) we have

$$S_e(e^i) = P_\theta(e^i)\prod_{j=1}^{J}\prod_{k=1}^{K}\Bigg[\mathbb{1}(\Lambda^{ijk} = \zeta)[1 - P^{jk}_{\phi,d}(e^i)] + \frac{\mathbb{1}(\Lambda^{ijk} \ne \zeta)P^{jk}_{\phi,d}(e^i)P^{jk}_\phi(\Lambda^{ijk}|e^i)}{\Upsilon^k(\Lambda^{ijk})}\Bigg].$$

Note that the final fraction in the previous equation is a likelihood ratio comparing interpretations of the same arrival as either the arrival of event i's jth phase at station k, or as a false arrival or a coda arrival. We can further decompose the score into scores S_d for each arrival. The score of Λ^{ijk}, defined when $\Lambda^{ijk} \ne \zeta$, is the ratio of the probabilities of the hypothesis, where the arrival is associated with phase j of event i at station k versus the default classification.

$$S^{jk}_d(\Lambda^{ijk}|e^i) = \frac{P^{jk}_{\phi,d}(e^i)}{1 - P^{jk}_{\phi,d}(e^i)}\frac{P^{jk}_\phi(\Lambda^{ijk}|e^i)}{\Upsilon^k(\Lambda^{ijk})}.$$

By definition, any arrival with a score less than 1 is more likely to be a false or coda arrival. Also, the score of an individual arrival is independent of other arrivals and events in the hypothesis. These scores play a key role in the following local search moves.

Birth Move

The birth move proposes events within a given time range, based on a list of arrivals. It starts off by inverting each of these arrivals to obtain an initial candidate list of event locations and times. The ability to invert an event follows from the fact that the slowness of an arrival is a monotonic function of distance (with fixed depth). If one assumes that an arrival is the P phase of a surface event, one can obtain a distance estimate from the slowness, which, combined with the arrival azimuth and time, gives an estimate for an event location and time. In Figure 20 we show the statistics of the distance between the inverted locations obtained from all arrivals in a one-week period and the corresponding ground-truth events during the same time period. The statistics in the figure show that for the great majority of events there is some arrival with an inverted location within 5°.

Having proposed a location and time by inverting one arrival, the birth move next attempts to construct the best possible event in a 5°-ball around the candidate location in steps of 2.5° using all available detections. Event depth