

A semiautomatic method to tie well logs to seismic data

Roberto H. Herrera¹ and Mirko van der Baan¹

ABSTRACT

We evaluated a semiautomatic method for well-to-seismic tying to improve correlation results and reproducibility of the procedure. In the manual procedure, the interpreter first creates a synthetic trace from edited well logs, determines the most appropriate bulk time shift and polarity, and then applies a minimum amount of stretching and squeezing to best match the observed data. The last step resembles a visual pattern recognition task, which often requires some experience. We replaced the last step with a constrained dynamic time-warping technique, to help guide the interpreter. The method automatically determined the appropriate amount of local stretching and squeezing to produce the highest correlation between the original data and the created synthetic trace. The constraint ensured that stretching and squeezing were kept within reasonable bounds, as determined by the interpreter. Results compared well with the manual method, leading to ties along the entire trace length in contrast to the shorter analysis window in the conventional method. Yet, we advise against unsupervised applications because the method is intended as a guide instead of a fully automated blind approach.

INTRODUCTION

Tying well logs to seismic traces is a crucial step in seismic interpretation to correlate subsurface geology to observed seismic data. The ease and quality of the tying procedure depend on the availability of high-quality logs and the estimation of a suitable wavelet. White and Simm (2003) and Newrick (2012) describe best practices and general recipes for the tying procedure. An estimated wavelet is convolved with the reflectivity series calculated from the well logs (sonic log and bulk density log), to generate the first synthetic trace (Anderson and Newrick, 2008). The wavelet can be estimated directly from the amplitude spectrum of the seismic data,

giving a zero-phase wavelet, which is phase rotated and time shifted during the matching process (e.g., White and Simm, 2003) or using a coherency matching technique (White, 1980; Walden and White, 1998), giving a frequency-dependent wavelet. The next step is to match the generated synthetic with the seismic data. The target seismic trace intuitively should be the trace at the well location, but time-migration effects could move the best match location away from the well. White and Simm (2003) propose to find the best match time location by scanning several traces in the vicinity of the well. Alternatively, as a compensation for dispersion at the well location, a composite trace can be generated from the neighboring traces (e.g., Hampson-Russell, 1999).

Well logs are commonly used as ground truth to correlate the seismic signal with the earth's stratigraphy (White and Simm, 2003), but several key steps render the procedure challenging and rarely fully reproducible. Density logs are greatly affected by borehole conditions, e.g., washouts or invasion, which negatively impact the measurement of the bulk's rock density (Anderson and Newrick, 2008). On the other hand, sonic logs are less sensitive to borehole conditions but they are measured at different length scales than the seismic data, and thus calibration plays an important role via dispersion and check shot corrections. This obviously assumes that high-quality velocity and density logs are available; if not, pseudocurves have to be estimated by empirical relationships (Newrick, 2012).

Only after creation of a first, correctly positioned, synthetic trace following advocated best practices (White and Simm, 2003) should the actual matching process start. The goal is not to force perfectly matching waveforms between the synthetic and selected observed trace, but to find distinctive reflectors on the seismic section that are also observable on the seismic trace (Newrick, 2012). These reflectors are rarely perfectly aligned, for instance, due to dispersion effects, necessitating some corrections in timing. This is done by changing the sonic velocities by a small amount, i.e., by stretching or squeezing the synthetic seismogram. The acceptable amount of imposed interval velocity changes is left to the interpreter's discretion. This particular step is clearly prone to pitfalls due to subjectivities in interpretation and procedures. In addition, even the

Manuscript received by the Editor 5 July 2013; revised manuscript received 27 October 2013; published online 12 March 2014.

¹University of Alberta, Department of Physics, Edmonton, Alberta, Canada. E-mail: rherrer@ualberta.ca; mirko.vanderbaan@ualberta.ca.

© 2014 Society of Exploration Geophysicists. All rights reserved.

selection of major matching reflections on the synthetic and observed traces can be challenging in practice.

Furthermore, the quality of the tie between the synthetic and the seismic trace is based on the correlation coefficient, which is limited to a constant time-shift similarity. The time-variant nature of the actual propagating wavelet, due to dispersion effects, adds nonlinearities to the seismic trace that cannot be easily followed by a linear metric. Likewise, nonconsistent dispersion effects in the sonic logs and seismic data affect the estimated reflectivity series, also adding time-varying nonlinear variations. We propose to use dynamic time warping (DTW) to match these time series, where the quality of the fit is not limited to the selected correlation window but the entire trace. DTW, as a correlation-based technique, measures the goodness of fit but not the reliability of the end results. For that purpose, we use a quality control (QC) step, which is done by watching the relative velocity changes produced by the tying.

DTW has been proven to handle nonstationarities in signals. Nonstationarity in time series due to physical processes is common in many areas such as speech processing (Rabiner and Juang, 1993), medicine, industry, and finance (Keogh, 2002). DTW is a robust tool to match time series even if the shifts vary with time (Keogh, 2002).

In this work, the application of DTW is limited to the specific step of stretching and squeezing. This is a controversial (White and Simm, 2003), but necessary, part in the tying procedure. Our procedure substitutes the manual stretching and squeezing step by an optimization algorithm, which is still supervised by the interpreter. This improves the repeatability of the tying, while the critical and often-abused stretch and squeeze (Newrick, 2012) is still under control. We assume that all the best practices have been followed (White and Simm, 2003; Anderson and Newrick, 2008).

Related works using DTW in seismic applications originate with attempts to automate well-to-well log correlation (Lineman et al., 1987; Zoraster et al., 2004). Well logs from different wells are correlated to infer common earth features. Zoraster et al. (2004) find that the crosscorrelation was unable to follow local distortions such as stretching or shrinking of stratigraphic intervals, typical of logs collected even from closely spaced wells. Essentially, these methods aim to correlate common features in various logs (Anderson and Gaby, 1983). An early approach to DTW is presented by Martinson et al. (1982) and Martinson and Hopper (1992). They develop a mapping function able to track stretching and squeezing in a time series, based on a correlation technique to establish a point-for-point correlation between traces. Liner and Clapp (2004) propose a global optimization method to create pairwise alignments between seismic traces. More recently, Hale (2013) uses DTW in seismic image processing, which aligns seismic images with differences not limited to time shifts. Our work follows the constrained optimization of Sakoe and Chiba (1978) with DTW, and it is not pairwise, but allows for one-to-many connections in both directions between two traces. We (Herrera and Van der Baan, 2012a) already present an early stage application, and here we develop a refined and more complete method.

In this paper, we first explain the theory behind DTW, and then we introduce the QC parameters. Next, we focus on the practical implementation showing a real example. Finally, we discuss our results with practical recommendations for application.

THEORY

The conventional method of well ties depends on the quality of the matching between the synthetic seismograms with the actual seismic traces. The measure of this matching has been extensively studied (e.g., White and Simm, 2003), but it remains based mainly on linear techniques. We propose the use of dynamic programming techniques commonly used in data mining (Keogh and Kasetty, 2003; Keogh and Ratanamahatana, 2004) and in speech processing (Rabiner and Juang, 1993).

The correlation coefficient is often used to measure the quality of the seismic-to-well tie (Hampson-Russell, 1999). Comparing two (time-dependent) sequences $S = [s_1, s_2, \dots, s_n]$ and $T = [t_1, t_2, \dots, t_n]$, both of length n , gives a correlation coefficient (γ_{ST}) at the time lag τ :

$$\gamma_{ST}(\tau) = \frac{\sum_{i=1}^n [s_i - \mu_S][t_{i+\tau} - \mu_T]}{(\sum_{i=1}^n [s_i - \mu_S]^2 \sum_{i=1}^n [t_i - \mu_T]^2)^{1/2}}, \quad (1)$$

where μ_S and μ_T are the means of S and T , respectively, and the denominator supplies the energy normalization term.

The optimal time lag τ is generally set at the maximum correlation coefficient. This measure works well if a constant time shift τ characterizes both signals. When this time alignment is constant, the problem is reduced to the correction of the time lag by crosscorrelation. But this measure fails to find the best match in nonstationary cases.

Most geophysical applications have nonstationary time-alignment problems (Anderson and Gaby, 1983). An alternative to the crosscorrelation is to find the Euclidean distance (L_2 -norm) between the two time series (Keogh and Kasetty, 2003):

$$D(S, T) = \sqrt{\sum_{i=1}^n (s_i - t_i)^2}, \quad (2)$$

where $D(S, T)$ is the one-to-one distance between the synthetic S and the trace T .

The Euclidean distance (L_2 -norm) is the most widely used distance measure. It is trivial to implement but also is very sensitive to small distortions in the time axis (Berndt and Clifford, 1994; Keogh and Kasetty, 2003). Taking the advantages of the Euclidean distance and adapting it for nonstationary matching, Berndt and Clifford (1994) propose DTW.

DTW distance can accommodate stretching and squeezing in the time series by linear programming. It uses the Euclidean distance as the initial metric but allows for the one-to-many alignment. The warping distance is represented as the minimum path in a grid representation of both sequences. In Figure 1, the warping path, $W = w_1, w_2, \dots, w_k$, aligns the elements of S and T in such a way that the distance between them is minimized.

In this matrix, the squared distance δ in the elements (i, j) is calculated by

$$\delta(s_i, t_j) = (s_i - t_j)^2. \quad (3)$$

To find the best alignment between these two sequences, we have to retrieve the warping path W through the distance matrix that minimizes the total cumulative distance for each path (Berndt and Clifford, 1994; Keogh, 2002) as illustrated in Figure 1. The

optimal path that minimizes the total warping cost (Berndt and Clifford, 1994) is

$$W(S, T) = \min_W \sum_{k=1}^p \delta(w_k), \quad (4)$$

where each w_k corresponds to a point $(i, j)_k$ and each grid point (i, j) corresponds to an alignment between s_i and t_j . From Figure 1, we can extract the first samples of the new stretched sequences as $S_k = s1, s2, s3, s3, s4, s5, s5, s5, s6, s7$ for the synthetic and $T_k = t1, t2, t2, t3, t3, t4, t5, t6, t7, t7$ for the stretched trace.

The dynamic programming approach uses the following recurrence to find the warping path (Berndt and Clifford, 1994):

$$\gamma(i, j) = \delta(s_i, t_j) + \min[\gamma(i-1, j), \gamma(i-1, j-1), \gamma(i, j-1)], \quad (5)$$

where $\delta(s_i, t_j)$ is the distance defined in equation 3 and the cumulative distance $\gamma(i, j)$ is the sum of the distance between the current elements and the minimum cumulative distance of the three neighboring cells.

Where the cumulative distance $\gamma(i, j)$ gets a minimum point of all the allowed points in the distance matrix, we store this value in the matrix W , which contains the optimum values of warping paths w_k . Thus, w_k is a two-column array with points $(i, j)_k$ that is used to construct the warped signals by mapping the original points (i, j) to $(i, j)_k$. This mapping process produces stretched versions of the original signals with length p . Only if the two signals are fully aligned from the start, do we have $i = j = k$.

For the final warping process, we compute a new argument $\hat{i}$ by extracting the indices of the intersection of the two sets $\{i_k\}$ and $\{j_k\}$:

$$\hat{i} = i_k(j_p), j_p \subseteq \text{position of } \{j_k\} \text{ in } \{i_k\} \cap \{j_k\}. \quad (6)$$

In our implementation, we use the MATLAB function `intersect`, which generates the indices of the intersection of two vectors. Following the hypothetical example shown in Figure 1, with the two stretched signals S_k and T_k the new argument is $\hat{i} = 1, 3, 4, 5, 5, 5, 7$. With the new argument $\hat{i}$, we can now get the warped signal $S_{\hat{i}}$. This signal is the best approximation of T_j following the optimum warping path ($S_{\hat{i}} \approx T_j$).

In time-domain signals, this monotonic transformation of the initial time interval into itself with different axis distribution is called “curve registration” (Ramsay and Silverman, 2005). The argument $\hat{i}$ accelerates or decelerates the synthetic signal along the time axis to match the seismic trace, such that well tops can be matched to corresponding reflectors.

Global constraints

There are cases when the two signals are not correlated by a stationary time lag; i.e., the warping path is not a diagonal in Figure 1. The warping path could then become quite nonlinear, indicating strong nonstationarity.

To prevent the occurrence of nonphysical alignments between both signals, we use a global distance constraint r to limit the maximum allowed amount of stretching and squeezing (Sakoe and

Chiba, 1978). This is illustrated by the two diagonal dotted lines in Figure 1. Elements of the warping matrix are restricted by the warping window $|i_k - j_k| < r$, where r is the window width. The warping window prevents pathological warping (Keogh, 2002). In our case, it prevents unrealistic velocity changes due to stretching and squeezing of the synthetic trace. This method will be called “constrained DTW.”

Proper specification of the permitted warping window guarantees meaningful results in the tying process. A large constraint r allows more stretching, which is sometimes undesirable, but also provides better matching. Its choice allows for a trade-off between algorithmic performance and physical meaning.

Oversampling

The connection of two points in the warping path could lead to unrealistic time shifts. If we add more samples, the warping path gains resolution because the connection is done at a shorter step. In this way, nonphysical connections appear more clearly on the warped signals. Likewise, similar signal parts are better represented. A modest oversampling of two to three times improves the algorithm performance without a high increment in the computational load. With the oversampled signals, we run the DTW algorithm and then decimate to the original time rate fitting the warping path using a smoothing spline filter.

Quality control: The relative velocity change

QC is done by computing the effects of the stretching and squeezing process on the resultant velocity curve. It is similar to the time-depth correction applied to sonic logs when check-shot points are available. In our method, the corrected signal produces a new time-depth curve, which comprises all the automatic stretching and squeezing represented by the warping path.




Figure 1. Warping path at the minimum distance of two sequences. The stretched versions of the original signals following the warping path are $S_k = s1, s2, s3, s3, s4, s5, s5, s5, s6, s7$ and $T_k = t1, t2, t2, t3, t3, t4, t5, t6, t7, t7$. The dashed lines represent the constraint lines for a warping window of width r .

The ratio of interval (or local) velocity change at depth z_i is given by the derivative of the smoothed warping path $w'(z_i)$ for the mapped variable $\hat{i}$ from equation 6. The local change in velocity at depth z_i is then $V_{p,new}(z_i) = V_{p,old}(z_i)/w'(z_i)$.

The new time depth curve is then

$$t_{d,new}(z_i) = 2 \sum_{k=1}^{\hat{i}} \frac{\Delta z_k}{V_{p,new}(z_k)}. \quad (7)$$




Figure 2. Seismic section with a well at CDP 13. The inset shows the initial correlation for that well. In red is the generated synthetic, and in blue is the composite seismic trace. The initial correlation for a synthetic created with a zero-phase statistical wavelet is 0.60, in the time window 850–1100 ms.




Figure 3. Unconstrained DTW: Warping path showing large stretching and squeezing at sample 350 of the synthetic (horizontal trace). The unconstrained algorithm is able to match the synthetic trace with the seismic trace (vertical trace) even if this may lead to unrealistic interval velocity changes.

RESULTS

We apply the DTW algorithm to obtain a well tie alignment between observed seismic data and synthetic traces created from well logs. The data set used in our experiments consist of a 3D poststack seismic profile, with 13 wells and their corresponding logs. The data set is provided as benchmark seismic data within a commercial package (Hampson-Russell, 1999). Figure 2 shows a seismic section at crossline 42 with one well at inline 13. The seismic-to-well tie for this well is shown in the inset figure with a crosscorrelation of

0.60. This synthetic trace (red) and its corresponding composite trace (blue) have been exported for postprocessing using the DTW method. The composite trace was built as the average of the neighboring traces centered at the well location with radius 4. The sampling rate is 500 Hz, and both signals have the same length; i.e., the seismic trace has been shortened to the well log length. Also, both signals have been standardized in amplitude. The initial synthetic is created as the convolution of the computed reflectivity from well logs with a zero-phase statistical wavelet. This zero-phase wavelet is calculated from the seismic amplitude spectrum — It is estimated as the inverse Fourier transform of the square root of the autocorrelation spectrum (Hampson-Russell, 1999).

Following the theoretical setup in Figure 1, the original synthetic trace is the bottom signal in Figure 3. The seismic trace is the left vertical signal. The main plot shows the cumulative distance matrix with the warping path. Note that the vertical portion of the warping path at sample 350 in the synthetic trace matches this point with 100 samples (from 240 to 340) from the seismic trace.

The unconstrained DTW algorithm seeks the minimum cost path even for unrealistic alignments as shown in Figure 3. In the resulting well tie, such large stretches and squeezes between the synthetic and the observed trace would represent unrealistic changes in the interval velocities, which are highly nonphysical.

To avoid such unrealistic connections, the original observed data and the synthetic are subjected to the constrained DTW approach. We use a global constraint based on the Sakoe-Chiba band (Sakoe and Chiba, 1978), which constrains the alignment process to a limited window. This reduces the freedom of the warping path to align events within a limited time shift. The estimated warping path is shown in Figure 4, where we used $r = 10$ samples to limit the maximum amount of permitted point-to-point shifting. With a sampling rate of 2 ms, two peaks can then still be aligned even if they are 20 ms apart.

Figure 5 shows how the synthetic and seismic traces match after constrained DTW. The originals (Figure 5a) and stretched versions (Figure 5b) are shown, with the same convention:

red for the synthetic trace and blue for the seismic trace. The DTW technique allows for one-to-many sample connections in both directions. Note that the original signals have 435 samples and their stretched versions have 597 samples. This new length comes from mapping the initial time (length n) onto the warping path (length k). The gray lines connecting individual time samples are useful for QC because they are indicative of time shifts between traces (forward- or backward-slanting lines) and changes in the relative velocity profiles (diverging or converging neighboring lines).

Visual inspection reveals that automatic matching leads to physically acceptable results for points 200–400 in the original series (Figure 5a), but likely excessive stretching and squeezing has occurred in the first 200 points. This point will be revisited after comparison with a manual tie.

Figure 6 shows the matched traces after final warping, where both traces now have the same number of samples as the original observed seismic trace. The new argument ($\hat{i}$) was estimated by applying equation 6 to the constrained warping path shown in Figure 4, such that the synthetic trace samples in the new argument approximate the seismic trace (synthetic ($\hat{i}$) $\approx$ seismic (i)).

For comparison, we perform a manual tie, following recommended practices in seismic-to-well tying (Hampson-Russell, 1999; White and Simm, 2003). We first select a high-quality region of interest for the synthetic and seismic traces where the correlation window is placed. This region of interest comprises the main stratigraphic features and is situated where both signals are more similar. In this case, the correlation window is between 850 and 1100 ms; only a constant time shift is needed in this region to reach the best match. The correlation of the manual seismic-to-well tie is 0.80 inside this window as is shown in Figure 7 (top). The left portion in the manual tie shows little similarity and would require significant stretching and squeezing before an acceptable fit is obtained. Because excessive stretching and squeezing are never recommended, we exclude this portion here from the manual tie. We will consider any automated matching in this portion as suspect for this reason as well.

We use two warping windows to test the reliability of the automated tying procedure and the effect of constraining the warping window on the resultant velocity curves. The automated method with a warping window of $r = 20$ (bottom of Figure 7) is able to completely reproduce the tying inside the manual correlation window and shows a total correlation coefficient along the entire traces of 0.92. Naively, this could be considered a great result based on the naked-eye observation in matching time series and the value

of the correlation coefficient, but the velocity changes produced are unrealistic with local perturbations in excess of 40% as is shown in Figure 8 (Figure 8c; the dotted line corresponds to $r = 20$). To prevent such velocity changes, we need to constrain the warping path, allowing less stretching and squeezing. With $r = 10$, the correlation coefficient is 0.89, slightly smaller than the previous one, but the velocity changes in the right part of the signal is minimum; i.e.,




Figure 4. Constrained DTW. The warping path represented by the center line is limited by the warping window $r = 10$ samples. Vertical and horizontal signals are the seismic trace and the synthetic one, respectively.




Figure 5. Illustration of the warping process. In panel (a), the seismic trace (blue) is connected to the synthetic trace (red) through gray lines representing connected time samples. Note the one-to-many connections in both directions. (b) The stretched version of both signals. They are highly correlated, but the number of samples increases compared with the original signals.

the local velocity change is close to one (Figure 8b). Thus, the semi-automated method corrected the simple time shift observed in the analysis window for the manual tie and produced the best possible match in the left part with mostly reasonable velocity changes (less than 10%) except possibly between 400 and 550 ms. Inspection of the relative slowness variations as well as the updated local velocities will reveal if and where the automated tying procedure has produced unrealistic stretching and squeezing. Such areas can then be discarded, or the warping window is to be limited further.

DISCUSSION

The DTW method aims to guide the interpreter in the well tying process by suggesting the optimum solution to create highly correlated traces. However, verification of the results using the suggested QC measures is strongly advised. In particular, nonphysical local

velocity changes or relative slowness perturbations are highly suspect, even if the final traces look perfectly matched.

The unconstrained DTW produces an optimal tying irrespective of physical implications; hence, the importance of properly bounding the warping paths. A large window may produce unrealistic connections; conversely, a very short one may lead to little to no improvement. The use of the global constraint leads to a robust matching process. The warping window parameter should be carefully adjusted to satisfy the trade-off between good matching and realistic alignment of events. This is similar to controlling the amount of stretching and squeezing in the manual method. This critical parameter keeps the velocity curve limited to realistic values.

The constrained DTW method was able to match similar events along the full seismic trace providing better correlation coefficients than the manual method. What validates the proposed approach is

that a good agreement is found between the manual and automatic well tie inside the manual correlation window. Thus, the matching outside this window actually represents the best possible tying, but it is the velocity variation that helps identify whether these changes are acceptable or not.

As a general tying procedure, we recommend the following:

- 1) First, an appropriate wavelet is to be estimated.
- 2) Next, the correct wavelet polarity is to be chosen.
- 3) Then, a global bulk shift is to be applied, for instance, using a simple correlation. This avoids requiring large warping windows.
- 4) Finally, the length of the warping window should be slowly increased until the QC procedure indicates unrealistic stretching and squeezing results.

Estimation of an appropriate wavelet is a crucial step because it can greatly affect the ease and quality of manual and semiautomated well tying. The wavelet power spectrum is generally best estimated using spectral averaging. However, this leads to a zero-phase wavelet, which is likely nonoptimal. For the appropriate wavelet phase, we recommend the use of kurtosis-based methods for constant-phase wavelets (Van der Baan, 2008) or short-time homomorphic wavelet estimation (Herrera and Van der Baan, 2012b) for frequency-dependent wavelet phases. These methods use statistical means to obtain the wavelet estimates directly from the data, and they have been shown to be reliable even for relatively low signal-to-noise ratios (Edgar and Van der Baan, 2011; Herrera and Van der Baan, 2012b).

There is an implicit assumption that the wavelet is stationary within the chosen correlation window. This is generally true for shorter windows, say, less than 0.5 s long, but for longer ones, seismic attenuation and dispersion make




Figure 6. Warped synthetic at the bottom producing a tie of 0.89 after the automated process with a warping window of $r = 10$. Note that the sample length is corrected to match the size of the seismic trace at the top. Major reflections are aligned.




Figure 7. Comparison between the automated and manual seismic-to-well ties. (a) Manual well tie with correlation of 0.80 in the window 850–1100 ms (bold vertical lines). (b) Semiautomatic tie with a warping window $r = 10$. Using the full length of the traces, the computed correlation is 0.89. The bottom pair (c) corresponds to a warping window of $r = 20$ with a correlation of 0.92 over the entire trace length. In all pairs, the top signal is the seismic trace and the bottom one is the synthetic trace.




Figure 8. QC: (a) The depth time curves — original curve and the resultant changes with two different warping windows r . The local velocity change is plotted in (b) for both warping windows with more stability for $r = 10$. (c) The original and updated local velocities due to the warping process, with reasonable velocity changes for $r = 10$.

it less likely. If attenuation and dispersion corrections have been successfully applied to the data prior to the seismic-to-well tie, then the full correlation window can be used with a single wavelet. If not, three options exist:

- 1) Split the correlation window into several parts with individually estimated wavelets.
- 2) Estimate the time-varying wavelet directly from the data using statistical means.
- 3) Predict the wavelet's nature using deterministic methods, such as inverse Q -filtering.

Van der Baan (2012) gives a discussion on the advantages and inconveniences of corrections based on statistical estimates such as time-varying Wiener deconvolution and deterministic ones based on inverse Q -filtering.

Polarity is another crucial parameter in well ties. Polarity reversal can significantly improve the quality of the well tie (Gratwick and Finn, 2005). Our recommendation is to compute a crosscorrelation between the synthetic and observed trace first to check if there is a significant difference in absolute value between the most positive and negative correlation coefficient. If not, we suggest to run the algorithm with the normal and reverse polarity and compare both outputs taking into account the effects on the velocity changes. The best tie connects reflectors with the same polarity and produces meaningful velocity changes.

Finally, other applications of DTW are envisioned for seismic data. These include alignment of baseline and monitor surveys in 4D seismic data (Fomel and Jin, 2009; Hale, 2013), PP and PS wavefield registration for 3C data (Gaiser, 1996), seismic offset balancing (Ross and Beale, 1994), and depth conversion and VSP calibration (Hackert and Parra, 2002). It could also be used for log-to-log correlations (e.g., Bois et al., 1972; Anderson and Gaby, 1983), and to correlate logs and extracted cores (e.g., Price et al., 2008).

CONCLUSIONS

Manually tying the generated synthetic to the seismic trace is a labor-intensive task, subject to interpreter bias. DTW improves the reproducibility of the tying process and creates a more guided, less subjective procedure. It also generates superior matches over conventional manual ties, yet possibly at the expense of unrealistic velocity changes. Unsupervised applications are therefore not advisable, and examination of the associated velocity changes is highly recommended.

Neither the visual impression nor the correlation coefficient are objective metrics of trust to evaluate the quality of a well tie. By integrating the robustness of the semiautomatic method into the traditional procedure, one generates a hands-free tying, which by no means implies unsupervised.

ACKNOWLEDGMENTS

The authors thank Hampson-Russell for software licensing and the sponsors of the Blind Identification of Seismic Signals (BLISS) project for their financial support. We also thank the three reviewers: M. Hall for the excellent review and positive comments that helped to improve this manuscript, R. White for his criticism and great questions, and a third anonymous reviewer whose comments shaped this manuscript with more geophysical insights. Finally, we thank the associate editor, A. Baumstein, for his valuable recommendations.

REFERENCES

- Anderson, K. R., and J. E. Gaby, 1983, Dynamic waveform matching: Information Sciences, **31**, 221–242, doi: [10.1016/0020-0255\(83\)90054-3](https://doi.org/10.1016/0020-0255(83)90054-3).
- Anderson, P., and R. Newrick, 2008, Strange but true stories of synthetic seismograms: CSEG Recorder, **12**, no. 12, 51–56.
- Berndt, D. J., and J. Clifford, 1994, Using dynamic time warping to find patterns in time series: Presented at KDD-94: AAAI-94 Workshop on Knowledge Discovery in Databases, 359–370.
- Bois, P., M. L. Porte, M. Lavergne, and G. Thomas, 1972, Well-to-well seismic measurements: Geophysics, **37**, 471–480, doi: [10.1190/1.1440273](https://doi.org/10.1190/1.1440273).
- Edgar, J. A., and M. Van der Baan, 2011, How reliable is statistical wavelet estimation?: Geophysics, **76**, no. 4, V59–V68, doi: [10.1190/1.3587220](https://doi.org/10.1190/1.3587220).
- Fomel, S., and L. Jin, 2009, Time-lapse image registration using the local similarity attribute: Geophysics, **74**, no. 2, A7–A11, doi: [10.1190/1.3054136](https://doi.org/10.1190/1.3054136).
- Gaiser, J., 1996, Multicomponent V_p/V_s correlation analysis: Geophysics, **61**, 1137–1149, doi: [10.1190/1.1444034](https://doi.org/10.1190/1.1444034).
- Gratwick, D., and C. Finn, 2005, What's important in making far-stack well-to-seismic ties in West Africa?: The Leading Edge, **24**, 739–745, doi: [10.1190/1.1993270](https://doi.org/10.1190/1.1993270).
- Hackert, C. L., and J. O. Parra, 2002, Calibrating well logs to VSP attributes: interval velocity and amplitude: The Leading Edge, **21**, 52–57, doi: [10.1190/1.1445848](https://doi.org/10.1190/1.1445848).
- Hale, D., 2013, Dynamic warping of seismic images: Geophysics, **78**, no. 2, S105–S115, doi: [10.1190/geo2012-0327.1](https://doi.org/10.1190/geo2012-0327.1).
- Hampson-Russell, 1999, Theory of the Strata program: Technical report May 1999, CGGVeritas Hampson-Russell.
- Herrera, R. H., and M. Van der Baan, 2012a, Guided seismic-to-well tying based on dynamic time warping: 82nd Annual International Meeting, SEG, Expanded Abstracts, doi: [10.1190/segam2012-0712.1](https://doi.org/10.1190/segam2012-0712.1).
- Herrera, R. H., and M. Van der Baan, 2012b, Short-time homomorphic wavelet estimation: Journal of Geophysics and Engineering, **9**, 674–680, doi: [10.1088/1742-2132/9/6/674](https://doi.org/10.1088/1742-2132/9/6/674).

- Keogh, E., 2002, Exact indexing of dynamic time warping, in P. A. Bernstein, Y. E. Ioannidis, R. Ramakrishnan, and D. Papadias, eds., Proceedings of the 28th International Conference on Very Large Data Bases, VLDB Endowment, Morgan Kaufman, 406–417.
- Keogh, E., and S. Kasetty, 2003, On the need for time series data mining benchmarks: A survey and empirical demonstration: Data Mining and Knowledge Discovery, **7**, 349–371, doi: [10.1023/A:1024988512476](https://doi.org/10.1023/A:1024988512476).
- Keogh, E., and C. A. Ratanamahatana, 2004, Exact indexing of dynamic time warping: Knowledge and Information Systems, **7**, 358–386, doi: [10.1007/s10115-004-0154-9](https://doi.org/10.1007/s10115-004-0154-9).
- Lineman, D., J. Mendelson, and M. Toksoz, 1987, Well to well log correlation using knowledge-based systems and dynamic depth warping: Presented at SPWLA 28th Annual Logging Symposium.
- Liner, C., and R. Clapp, 2004, Nonlinear pairwise alignment of seismic traces: Geophysics, **69**, 1552–1559, doi: [10.1190/1.1836828](https://doi.org/10.1190/1.1836828).
- Martinson, D., and J. Hopper, 1992, Nonlinear seismic trace interpolation: Geophysics, **57**, 136–145, doi: [10.1190/1.1443177](https://doi.org/10.1190/1.1443177).
- Martinson, D., W. Menke, and P. Stoffa, 1982, An inverse approach to signal correlation: Journal of Geophysical Research, **87**, 4807–4818, doi: [10.1029/JB087iB06p04807](https://doi.org/10.1029/JB087iB06p04807).
- Newrick, R., 2012, Well ties basics — Well tie perfection, in M. Hall, and E. Bianco, eds., Fifty-two things you should know about geophysics: vol. 10, Agile Libre, 104–107.
- Price, D., A. Curtis, and R. Wood, 2008, Statistical correlation between geophysical logs and extracted core: Geophysics, **73**, no. 3, E97–E106, doi: [10.1190/1.2890409](https://doi.org/10.1190/1.2890409).
- Rabiner, L., and B. H. Juang, 1993, Fundamentals of speech recognition: Prentice Hall.
- Ramsay, J., and B. W. Silverman, 2005, Functional data analysis, 2nd ed.: Springer, Springer Series in Statistics.
- Ross, C. P., and P. L. Beale, 1994, Seismic offset balancing: Geophysics, **59**, 93–101, doi: [10.1190/1.1443538](https://doi.org/10.1190/1.1443538).
- Sakoe, H., and S. Chiba, 1978, Dynamic programming algorithm optimization for spoken word recognition: IEEE Transactions on Acoustics, Speech, and Signal Processing, **26**, no. 1, 43–49, doi: [10.1109/TASSP.1978.1163055](https://doi.org/10.1109/TASSP.1978.1163055).
- Van der Baan, M., 2008, Time-varying wavelet estimation and deconvolution by kurtosis maximization: Geophysics, **73**, no. 2, V11–V18, doi: [10.1190/1.2831936](https://doi.org/10.1190/1.2831936).
- Van der Baan, M., 2012, Bandwidth enhancement: Inverse Q filtering or time-varying Wiener deconvolution?: Geophysics, **77**, no. 4, V133–V142, doi: [10.1190/geo2011-0500.1](https://doi.org/10.1190/geo2011-0500.1).
- Walden, A., and R. White, 1998, Seismic wavelet estimation: A frequency domain solution to a geophysical noisy input-output problem: IEEE Transactions on Geoscience and Remote Sensing, **36**, 287–297, doi: [10.1109/36.655337](https://doi.org/10.1109/36.655337).
- White, R. E., 1980, Partial coherence matching of synthetic seismograms with seismic traces: Geophysical Prospecting, **28**, 333–358, doi: [10.1111/j.1365-2478.1980.tb01230.x](https://doi.org/10.1111/j.1365-2478.1980.tb01230.x).
- White, R. E., and R. Simm, 2003, Tutorial: Good practice in well ties: First Break, **21**, 75–83.
- Zoraster, S., R. Paruchuri, and S. Darby, 2004, Curve alignment for well-to-well log correlation: Presented at SPE Annual Technical Conference and Exhibition.