

# Dynamic warping of seismic images

Dave Hale<sup>1</sup>

## ABSTRACT

The problem of estimating relative time (or depth) shifts between two seismic images is ubiquitous in seismic data processing. This problem is especially difficult where shifts are large and vary rapidly with time and space, and where images are contaminated with noise or for other reasons are not shifted versions of one another. A new solution to this problem requires only simple extensions of a classic dynamic time warping algorithm for speech recognition. A key component of that classic algorithm is a nonlinear accumulation of alignment errors. By applying the same nonlinear accumulator repeatedly in all directions along all sampled axes of a multidimensional image, I obtain a new and effective method for dynamic image warping (DIW). In tests where known shifts vary rapidly, this new method is more accurate than methods based on crosscorrelations of windowed images. DIW also aligns seismic reflectors well in examples where shifts are unknown, for images with differences not limited to time shifts.

## INTRODUCTION

In seismic data processing we often must estimate relative shifts in time (or depth) between seismograms. Those shifts often vary with time and space coordinates. Examples cited by [Liner and Clapp \(2004\)](#) include alignment of synthetic and recorded seismograms, registration of P- and S-wave images, residual normal move-out correction, and alignment of images computed for different source-receiver offsets or propagation angles. They proposed a dynamic programming solution to this problem in the case where pairwise alignment of seismic traces is sufficient.

A different dynamic programming solution was developed by [Sakoe and Chiba \(1978\)](#) in the context of speech recognition, and is today widely known as dynamic time warping (DTW) (e.g., [Müller \[2007\]](#) chapter 4). Significantly, DTW imposes constraints

on the rate at which shifts may vary in time, and these constraints often enable DTW to accurately estimate shifts from sequences that are contaminated with noise, or that in some other ways are not warped versions of each other.

The use of DTW to estimate shifts in geophysical time series and other sequences is not new. Several applications of DTW to problems in geophysics were proposed by [Anderson and Gaby \(1983\)](#), who called this algorithm “dynamic waveform matching.”

Unfortunately, the most straightforward extension of DTW to the problem of estimating shifts in multidimensional images has been shown to have no computationally feasible solution ([Keysers and Unger, 2003](#)). Specifically, time required to compute an exact solution grows exponentially with image dimensions. Therefore, numerous authors have proposed practical solutions to problems that approximate this intractable problem. [Pishchulin \(2010\)](#) provides a recent summary.

In this paper, I propose an extension of an approximate solution developed by [Mottl et al. \(2002\)](#). Their solution and my extension for dynamic image warping (DIW) are especially simple, requiring very little software beyond what would already be available to implement DTW.

I first review the DTW algorithm, giving special attention to the so-called accumulation and backtracking parts of this algorithm. I then show how the accumulation part of DTW can be used to implement a nonlinear smoothing of alignment errors computed for two multidimensional images, and how this leads to a new method for DIW.

In tests with pairs of images related by shifts that are known, large, and rapidly varying, I demonstrate the accuracy with which DIW can estimate the known shifts. In further tests, I show that DIW can be more accurate than methods based on local crosscorrelations, especially when shifts vary rapidly in time or space. Crosscorrelation methods, such as those proposed by [Hall \(2006\)](#) and [Hale \(2009\)](#) to estimate shifts in time-lapse seismic images, are accurate only where shifts are more slowly varying.

I then illustrate DIW with applications to two problems, the estimation of displacements associated with geologic faults and the registration of PP and PS images. Although the shifts in these

Manuscript received by the Editor 15 August 2012; revised manuscript received 16 November 2012; published online 7 February 2013.

<sup>1</sup>Colorado School of Mines, Golden, Colorado, USA. E-mail: dhale@mines.edu.

© 2013 Society of Exploration Geophysicists. All rights reserved.

two examples are unknown, improved alignment after warping demonstrates the accuracy and robustness of DIW where differences between two images cannot be attributed to time shifts alone.

## DYNAMIC TIME WARPING

Consider the two synthetic seismograms  $f[i]$  and  $g[i]$  with length  $N = 500$  samples displayed in Figure 1. I computed the sequence  $f[i]$  by convolving a Ricker wavelet with a random reflectivity sequence. I then applied time-varying shifts  $s[i]$  to that reflectivity sequence and convolved again with the same wavelet to obtain the sequence  $g[i]$ . The two sequences are therefore approximately (but not exactly) related by  $f[i] \approx g[i + s[i]]$ .

In practice, these two sequences might be a recorded seismogram  $f[i]$  and a synthetic seismogram  $g[i]$  derived from well logs. Or they might be two sequences of sample values extracted from a seismic image on opposite sides of a fault. In any case, the practical problem considered here is estimation of the shifts  $s[i]$  given only the two sequences  $f[i]$  and  $g[i]$ .

In the example of Figure 1, the shifts  $s[i]$  are a simple sinusoidal function that is apparent in Figure 2a, which is an image of alignment errors defined by

$$e[i, l] \equiv (f[i] - g[i + l])^2. \quad (1)$$

Note that these alignment errors are nearly zero where the integer lag  $l$  approximately equals the shift  $s[i]$ . Also note the constant extrapolation of errors in the corners of  $e[i, l]$ , where  $i + l < 0$  or  $i + l \geq N$ .

In most discussions of DTW (e.g., Sakoe and Chiba, 1978), alignment errors are defined in a different way by  $e[i, j] \equiv (f[i] - g[j])^2$ , and are computed for all indices  $i$  and  $j$  for which samples  $f[i]$  and  $g[j]$  are available. I favor the alternative definition in equation 1 because in most applications to seismic data processing we can process the sequences  $f$  and  $g$  to be approximately aligned before using DTW. This implies that the number of lags  $L$  for which we must compute alignment errors via equation 1 is




Figure 1. Two synthetic seismograms  $f[i]$  (a) and  $g[i]$  (b) corresponding to misaligned reflection coefficients, used as inputs to the DTW algorithm. Reflections in the sequence  $f[i]$  appear to be squeezed toward the middle of the sequence  $g[i]$ .

much smaller than the number of time samples  $N$ , which leads to significant savings in memory and computation time.

The definition of alignment errors in equation 1 can be modified by changing the power 2 or by using some other nonnegative function of the differences  $f[i] - g[i + l]$ , without changing the DTW algorithm. For example, we might use the absolute values of those differences. I have squared the differences in all of the examples shown in this paper, as in equation 1.

## Constrained optimization

The simplest DTW computes a sequence  $u[0:N-1] \equiv (u[0], u[1], \dots, u[N-1])$  of integer shifts by solving the following optimization problem

$$u[0:N-1] = \underset{l[0:N-1]}{\operatorname{argmin}} D(l[0:N-1]), \quad (2)$$

where

$$D(l[0:N-1]) \equiv \sum_{i=0}^{N-1} e[i, l[i]], \quad (3)$$

subject to the constraint

$$|u[i] - u[i-1]| \leq 1. \quad (4)$$

As illustrated in Figure 2b, DTW yields a minimizing sequence of integer shifts  $u[0:N-1]$  that well approximates the (here, known) sequence of shifts  $s[0:N-1]$ .

The first and last shifts ( $s[0]$  and  $s[N-1]$ ) in this sequence need not be zero, although they are zero in this example. Zero-value boundary conditions are often assumed (e.g., Sakoe and Chiba, 1978) but this assumption is unnecessary. The DTW algorithm described below works as well to estimate a sequence of shifts that neither begins nor ends with zero.




Figure 2. Alignment errors  $e[i, l]$  (a) are small along the sinusoidal path corresponding to known shifts between reflections in the two sequences shown in Figure 1. DTW (b) yields (solid white) estimated integer shifts  $u[i]$  that approximate well the (dotted white) known shifts  $s[i]$  in reflection coefficients.

The function  $D$  defined by equation 3 is often referred to as distance, which makes sense if we think of the image  $e[i, l]$  in Figure 2a as representing topography. Large values in  $e[i, l]$  then correspond to tall hills (misalignments) that we wish to avoid as we choose a path from left to right, that is, from  $i = 0$  to  $N - 1$ . In this sense, DTW chooses a path  $u[0:N - 1]$  to minimize the total distance traveled, subject to the constraint (equation 4) that the shifts cannot change too rapidly from one sample to the next.

The constraint equation 4 is analogous to the simplest slope constraint of Sakoe and Chiba (1978). This constraint is slightly different here only because I define alignment errors by  $e[i, l] \equiv (f[i] - g[i + l])^2$  (equation 1), instead of by  $e[i, j] \equiv (f[i] - g[j])^2$ . Here, this constraint ensures that the argument  $i + u[i]$  in  $(f[i] - g[i + u[i]])^2$  neither decreases nor increases too rapidly with increasing sample index  $i$ .

This constraint is important. Where  $u[i] - u[i - 1] = 1$ , we stretch by 100%, such that two adjacent samples in the sequence  $f[i]$  correspond to two nonadjacent samples in the sequence  $g[i]$ . Where  $u[i] - u[i - 1] = -1$ , we squeeze by 100%, such that two adjacent samples in the sequence  $f[i]$  correspond to only one sample in the sequence  $g[i]$ . In many practical applications, 100% is an unreasonably large amount of strain, and we will see below how to reduce this upper bound.

It is significant that the sequence  $u[0:N - 1]$  computed by DTW minimizes exactly the distance  $D$  defined by equation 3, while satisfying the constraint equation 4. Differences in Figure 2b between the integer shifts  $u[0:N - 1]$  and the known shifts  $s[0:N - 1]$  are due entirely to the restriction of  $u[0:N - 1]$  to be integers and the approximation  $f[i] \approx g[i + s[i]]$ , not to any approximation in the optimization algorithm.

## Dynamic programming

As its name implies, DTW is a dynamic-programming algorithm (e.g., Cormen et al., 2001). The essential trait of this algorithm is decomposition of a problem into a sequence of nested and smaller subproblems.

Let  $u[0:m]$  denote the sequence of shifts  $l$  that minimizes the distance  $D$  defined by equation 3. To identify the sequence of smaller subproblems nested within this larger minimization problem, we consider a subpath  $u[0:m]$  of the minimizing path  $u[0:N - 1]$  and observe that

$$u[0:m] = \operatorname{argmin}_{l[0:m]} \sum_{i=0}^m e[i, l[i]]. \quad (5)$$

For if  $u[0:m]$  were not a minimizing subpath, then we could replace that part of  $u[0:N - 1]$  and thereby reduce the total distance  $D$ , which implies that  $u[0:N - 1]$  does not minimize  $D$ ; a contradiction.

This observation is important because it implies that we need not test all possible (roughly,  $3^N$ ) paths  $l[0:N - 1]$  that satisfy the constraint equation 4 in our search for the minimizing path  $u[0:N - 1]$ . Instead, we can find this minimizing path in two steps: accumulation and backtracking.

## Accumulation

In the first accumulation step of DTW, we recursively compute from the array of alignment errors  $e[i, l]$  an array of distances  $d[i, l]$  as follows

$$\begin{aligned} d[0, l] &= e[0, l], \\ d[i, l] &= e[i, l] + \min \begin{cases} d[i - 1, l - 1] \\ d[i - 1, l] \\ d[i - 1, l + 1] \end{cases}, \\ &\text{for } i = 1, 2, \dots, N - 1. \end{aligned} \quad (6)$$

For each index  $i$ , we cannot yet know in this first step whether or not the lag  $l$  lies on the minimizing path  $u[0:N - 1]$ , so that  $l = u[i]$ . Therefore, we must compute and store distances  $d[i, l]$  for all lags, assuming for the moment that lag  $l$  may lie on the minimizing path. Figure 2b shows distances  $d[i, l]$  computed in this way for the alignment errors shown in Figure 2a.

The constraint equation 4 implies that, when computing  $d[i, l]$  as in equation 6, we must consider only three previously computed distances  $d[i - 1, l - 1]$ ,  $d[i - 1, l]$ , and  $d[i - 1, l + 1]$ . In other words, if lag  $l$  lies on the minimizing path at sample index  $i$ , then either lag  $l - 1$ ,  $l$  or  $l + 1$  must lie on the minimizing path at sample index  $i - 1$ .

The computational complexity of this first step is  $O(N \times L)$ , where  $L$  is the number of lags  $l$  for which alignment errors  $e[i, l]$  and distances  $d[i, l]$  are computed.

I call this first step “accumulation” because the distances  $d[i, l]$  are running sums of alignment errors  $e[i, l]$ . At the end of this first step, we can loop over all lags  $l$  to find the minimum distance

$$D = \min_l d[N - 1, l]. \quad (7)$$

## Backtracking

The second step in DTW is to find the minimizing path, the sequence of shifts  $u[0:N - 1]$ , beginning with the last shift  $u[N - 1]$  and ending with the first shift  $u[0]$

$$\begin{aligned} u[N - 1] &= \operatorname{argmin}_l d[N - 1, l], \\ u[i - 1] &= \operatorname{argmin}_{l \in \{u[i] - 1, u[i], u[i] + 1\}} d[i - 1, l], \\ &\text{for } i = N - 1, N - 2, \dots, 1. \end{aligned} \quad (8)$$

This backtracking step begins with a simple loop over lags  $l$  to find the last shift  $u[N - 1]$  in the sequence of shifts  $u[0:N - 1]$ . Because this last shift must be on the minimizing path, it must equal the lag at which we found the minimum distance  $D$ . We then recursively find previous shifts  $u[i - 1]$  in this sequence, comparing the three distances  $d[i - 1, l - 1]$ ,  $d[i - 1, l]$  and  $d[i - 1, l + 1]$  to determine which of these was used in equation 6 to compute the minimum distance  $d[i, l]$ .

The computational complexity of backtracking is only  $O(N)$  because, for each sample index  $i$ , we compute a shift  $u[i]$  by comparing only three distances. Therefore, the complexity of DTW is the  $O(N \times L)$  complexity of the accumulation step, which is proportional to the number of samples in the array of alignment errors  $e[i, l]$ .

## REFINEMENTS

Figure 3 displays the same two synthetic seismograms  $f[i]$  and  $g[i]$  shown in Figure 1, after adding different sequences of bandlimited random noise to each of them. The rms signal-to-noise ratio (S/N) is 2:1. Although this level of noise obscures somewhat the sinusoidal warping path in the plot of alignment errors  $e[i, l]$  in Figure 4a, the shifts  $u[i]$  estimated using DTW roughly approximate the known shifts  $s[i]$ .

Again, it is important to remember that DTW solves exactly the constrained optimization problem of equations 2, 3, and 4. Differences in Figure 4 between estimated and known shifts are primarily due to errors in the approximation  $f[i] \approx g[i + s[i]]$  caused by the addition of random noise. The sequence  $g[i]$  in Figure 3b is not a warped version of the sequence  $f[i]$  in Figure 3a.

### Limiting strain

The robustness of DTW in the presence of random noise is due largely to the constraint equation 4. The number of shift sequences  $u[0:N-1]$  that satisfy this constraint ( $\approx 3^N$ ) is far less than the number that would be possible without it ( $\approx L^N$ ).

Of course, the constraint that strain (stretch or squeeze) be less than 100% is useful only when satisfied by the actual shifts  $s[i]$  that we wish to estimate. However, in many practical applications, this constraint is more than reasonable. Indeed, a strain as high as 100% may be unreasonably high, and we may be able to improve the accuracy of shifts estimated in DTW by reducing this upper bound on strain to a more reasonable value.

The simplest way to more tightly bound strain in DTW is to sample lags  $l$  more finely at some fraction of the time sampling interval. For example, if that fraction were  $\frac{1}{2}$ , then we would compute alignment errors  $e[i, l]$  for lags  $l = \dots, -1, -\frac{1}{2}, 0, \frac{1}{2}, 1, \dots$ . The maximum strain permitted would then be 50%, as the constraint equation 4 would become

$$|u[i] - u[i-1]| \leq \frac{1}{2}. \quad (9)$$




Figure 3. Same as Figure 1, except that bandlimited random noise sequences have been added to the synthetic seismograms  $f[i]$  (a) and  $g[i]$  (b). In this example the rms S/N is 2:1.

Although straightforward, this method for reducing the upper bound on strain requires a significant increase in computational cost. For a limit of 50%, computation time and memory will double if we compute errors  $e[i, l]$  and distances  $d[i, l]$  for twice as many lags. The increase in memory will be especially significant as we extend the DTW algorithm to the problem of multidimensional image warping.

A more efficient way to limit strain is to implement constraints much like the slope constraints proposed by Sakoe and Chiba (1978). As an example, for a limit of 50% strain, any change in shift for one sample should be preceded by no change in shift for the previous sample. To enforce this constraint for all samples (except those near the beginning), we need only modify the two branches of the accumulator in equation 6 that correspond to a change in shift

$$d[0, l] = e[0, l],$$

$$d[1, l] = e[1, l] + \min \begin{cases} d[0, l-1] \\ d[0, l] \\ d[0, l+1] \end{cases},$$

$$d[i, l] = e[i, l] + \min \begin{cases} d[i-2, l-1] + e[i-1, l-1] \\ d[i-1, l] \\ d[i-2, l+1] + e[i-1, l+1] \end{cases},$$

for  $i = 2, 3, \dots, N-1$ . (10)

A corresponding change is required in the backtracking step, in which we must now compute and compare the three expressions inside the min function of equation 10, to determine which of these was used to compute the distance  $d[i, l]$ . With backtracking implemented in this way, DTW is constrained to shift sequences in blocks of two or more samples. If any sample is shifted by the warping,




Figure 4. Known sinusoidal warping is obscured in alignment errors  $e[i, l]$  (a) computed for the noisy synthetic seismograms of Figure 3. DTW (b) yields (solid white) estimated integer shifts  $u[i]$  that roughly approximate the (dotted white) known shifts  $s[i]$  in reflection coefficients.



then at least one of the adjacent samples must be shifted by the same amount. We may uniformly distribute this shift over these two adjacent samples, thereby reducing strain to satisfy the constraint equation 9.

Modifications similar to equation 10 can be easily and efficiently implemented for any strain limit of the form  $1/b$ , where  $b$  is a positive integer. Figure 5 shows how strain limits implemented in this way can improve the accuracy of shifts estimated by DTW. Note, however, that the strain limit of  $\frac{1}{5}$  used to estimate shifts  $u[i]$  shown in Figure 5c is almost equal to the maximum strain in the known shifts  $s[i]$ . Any further reduction in the strain limit would yield poor shift estimates because strain for the correct shifts would exceed that limit. In practice, lacking any a priori limit on strain, we must take care to not reduce this limit so much that we prohibit estimates of the correct shifts.

### Smoothing alignment errors

To further improve the accuracy of estimated shifts  $u[i]$ , we might attempt to attenuate noise in the two sequences  $f[i]$  and  $g[i]$ , or we might instead try to attenuate noise in the alignment errors  $e[i, l]$ . Considering the second option, suppose that we apply some sort of smoothing filter to the alignment errors  $e[i, l]$ . Can we improve the accuracy of the estimated shifts  $u[i]$  by applying such a filter before DTW?

This question is suggested by the accumulation step in DTW defined by equation 6. Each distance  $d[i, l]$  computed in this step is a sum of alignment errors, which implies that the distances  $d[i, l]$  vary less rapidly with index  $i$  than do the alignment errors  $e[i, l]$ . In other words, the accumulation step is a smoothing filter.

This recursive smoothing filter is one-sided because each  $d[i, l]$  in equation 6 depends on only previous and present alignment errors, those with sample indices less than or equal to  $i$ . This filter is also nonlinear because of the min function in equation 6. In effect, this one-sided nonlinear smoothing filter already attenuates noise in alignment errors caused by noise in the two sequences to be aligned by warping.

One way we might improve this smoothing filter would be to make it two-sided and symmetric. We can implement such a two-sided symmetric smoothing filter by applying a one-sided filter in forward and reverse directions. Smoothing in the forward direction is the same as computing distances  $d[i, l]$  via equation 6

$$\begin{aligned} \tilde{e}_f[0, l] &= e[0, l], \\ \tilde{e}_f[i, l] &= e[i, l] + \min \begin{cases} \tilde{e}_f[i-1, l-1] \\ \tilde{e}_f[i-1, l] \\ \tilde{e}_f[i-1, l+1] \end{cases}, \\ &\text{for } i = 1, 2, \dots, N-1. \end{aligned} \quad (11)$$

Smoothing in the reverse direction is similar

$$\begin{aligned} \tilde{e}_r[N-1, l] &= e[N-1, l], \\ \tilde{e}_r[i, l] &= e[i, l] + \min \begin{cases} \tilde{e}_r[i+1, l-1] \\ \tilde{e}_r[i+1, l] \\ \tilde{e}_r[i+1, l+1] \end{cases}, \\ &\text{for } i = N-2, N-3, \dots, 0. \end{aligned} \quad (12)$$

Two-sided smoothing is then defined by

$$\tilde{e}[i, l] = \tilde{e}_f[i, l] + \tilde{e}_r[i, l] - e[i, l]. \quad (13)$$

Subtraction of  $e[i, l]$  in equation 13 ensures that this value is not counted twice as, for all  $i$ , it appears in  $\tilde{e}_f[i, l]$  and  $\tilde{e}_r[i, l]$ . In this way, each smoothed error  $\tilde{e}[i, l]$  is a sum of past, present, and future alignment errors.

Like the accumulator in DTW, this two-sided smoothing filter is nonlinear because it uses the min function in equations 11 and 12 to determine which errors to sum. Figure 6 displays smoothed alignment errors for the two noisy sequences shown in Figure 3. Observe that the known sinusoidal warping path is somewhat more apparent in these smoothed errors than in the unsmoothed alignment errors displayed in Figure 4a. We might therefore expect the shifts  $u[i]$  estimated by DTW from the smoothed errors  $\tilde{e}[i, l]$  would be more accurate than those estimated from the unsmoothed errors  $e[i, l]$ .

However, this sort of smoothing does not improve DTW. Although not shown here, the shifts  $u[i]$  computed by DTW for the smoothed alignment errors shown in Figure 6c are identical to those computed for the unsmoothed alignment errors in Figure 4a. The benefit of this two-sided smoothing lies in the extension of dynamic warping to multidimensional images.




Figure 5. Shifts  $u[i]$  estimated by DTW for different limits on strain, the rate at which shifts can change with sample index  $i$ . As we reduce the upper bound on this strain from 1 (a) to  $\frac{1}{2}$  (b) to  $\frac{1}{5}$  (c), the (solid white) estimated shifts  $u[i]$  better approximate the (dotted white) known shifts  $s[i]$ .

## DYNAMIC IMAGE WARPING

The simplest way to extend DTW for image processing is to think of an image as a collection of vertical columns and to estimate vertical shifts by applying DTW to each of those columns independently. We could likewise apply DTW to image rows to obtain estimates of horizontal shifts.

Figure 7 illustrates the application of this simple method for DIW for two seismic shot records, where the first record shown in Figure 7a has been warped to obtain the second record shown in Figure 7b. DTW applied to each corresponding pair of columns from these images yields the estimated shifts shown in Figure 7c. Except for small source-receiver offsets where seismograms are missing, these estimated shifts approximate well the known shifts shown in Figure 7d.

These shifts are large, up to eight times larger than the dominant period (about 40 ms) of most reflection events. Many of the events in the shot records (such as those for small offsets and late times) are ringy, almost periodic, which can make estimation of the shifts difficult. Nevertheless, DTW applied independently to each pair of seismograms in these shot records accurately recovers the correct shifts.

When applying dynamic warping to either sequences or images, upper and lower bounds for shifts must be specified, and in this example all shifts were assumed to lie in the range  $[-480, 480]$  ms. Although it is important that these bounds include




Figure 6. Alignment errors for the noisy sequences in Figure 3 after smoothing in the forward (a), reverse (b), and both directions (c). Smoothing in the forward direction is equivalent to the accumulation step in DTW, in which we compute distances  $d[i, l]$  via equation 6.

the range of actual shifts, which for this test is  $[0, 320]$  ms, the latter range need not be known precisely.

The success of the simple column-by-column method for DIW depends primarily on the fact that each pair of seismograms in these shot records satisfies exactly the DTW assumption that one sequence is a warped version of the other. When this assumption is not satisfied, this simple method for DIW can fail miserably. For example, if we add different bandlimited random noise images to each of the shot records before DIW, we obtain the results shown in Figure 8. In this example, the rms S/N is 1 : 1. For this noise level, it is difficult to estimate well the correct shifts from each pair of noisy seismograms in the two shot records. Therefore, the estimated shifts vary significantly for different offsets, and imply an unlimited amount of strain in the horizontal direction.

To improve these estimated shifts, we would like to limit strain in horizontal as well as vertical directions. In other words, we would like to minimize alignment errors as in equations 2 and 3 while satisfying constraints like those in equations 4 or 9 in horizontal and vertical directions. Unfortunately, this constrained optimization problem has been shown to have no computationally feasible exact solution (Keyser and Unger, 2003). We must therefore make




Figure 7. A recorded seismic shot record before (a) and after (b) warping with shifts that vary with time and are up to eight times larger than the dominant period of seismic reflections. Except for small offsets where data are missing, shifts (c) computed from seismograms in these two images by the DTW algorithm approximate well the known shifts (d) used to perform the actual warping.

approximations to this intractable problem, and methods for DIW differ in their approximations.

### Tree-sequential dynamic programming

One such approximation is that proposed by Mottl et al. (2002), and this approximation and its solution are today often referred to as “tree-sequential dynamic programming” (TSDP) (e.g., Pishchulin, 2010). Given software for DTW, implementation of the TSDP algorithm for DIW is almost trivial in the case considered here, where we seek to estimate only vertical shifts. The TSDP algorithm begins by computing alignment errors as for DTW. It then smooths those alignment errors in the vertical direction, by applying the nonlinear two-sided smoothing equations 11, 12, and 13 independently for each image column. TSDP ends by applying the DTW algorithm to the smoothed errors, but now in the horizontal direction, accumulating and backtracking for each image row, again independently. In this way, TSDP processes a multidimensional image with a cascade of 1D smoothing, accumulation, and backtracking.

Time shifts estimated by TSDP are shown in Figure 9a. For this example, I used strain limits of 25% in the vertical direction and 100% in the horizontal direction, and these values are close to the

maximum strains in the known shifts displayed in Figure 8d. Compared to the estimated shifts shown in Figure 8c, the shifts from TSDP shown in Figure 9a better approximate the known shifts.

The most obvious improvement is in reduced horizontal strain, greater continuity of shifts in the horizontal direction. This improvement is not surprising because TSDP as described above ends by applying DTW independently for each image row, and we know that DTW satisfies strain limits precisely. However, because TSDP ends by applying DTW independently for each row, we have no guarantee that vertical strain limits (here 25%) will be satisfied. Vertical discontinuities in shifts are in fact apparent in Figure 9a. Although vertical smoothing of alignment errors in TSDP reduces the likelihood that such discontinuities will occur, it does not entirely eliminate them.

Rather than smoothing vertically and then applying DTW horizontally, we might instead smooth horizontally and apply DTW vertically. Shifts estimated by this alternative implementation of TSDP are not shown here, but are significantly less accurate than those shown in Figure 9a. The reason to first smooth vertically is that, for each offset, a pair of image columns (seismograms) typically contains multiple events that will indicate a path of minimum alignment error like that apparent in Figure 2a, but the same is not




Figure 8. Same as Figure 7, except that different bandlimited random noise records have been added to the two shot records (a) and (b). Because these noisy seismograms are not related by time shifts, the shifts (c) estimated by the DTW algorithm vary wildly for different offsets, unlike the known shifts (d).




Figure 9. Time shifts estimated by DTW after vertical (a), vertical-horizontal (b), vertical-horizontal-vertical (c), and vertical-horizontal-vertical-horizontal (d) smoothings of alignment errors. Insignificant differences between shifts (c) and (d) indicate that this process of smoothing in alternating directions before DTW has converged.



true for each pair of image rows. By first smoothing alignment errors vertically, we extend these paths of minimum error to times for which little information about vertical alignment may be available, so that DTW applied horizontally can then more accurately estimate the shifts. Nevertheless, vertical discontinuities apparent in the estimated shifts shown in Figure 9a suggest that further improvement is possible.

### Improving TSDP

The key to improving TSDP lies in recognizing that it first smooths alignment errors in one direction before it applies DTW in another direction. Although I have not seen TSDP described in this way, the description is accurate. So why not first smooth in both vertical and horizontal directions?

Figure 9b shows the result of smoothing vertically and horizontally before applying DTW vertically to each column of smoothed alignment errors. As expected, vertical discontinuities in shifts are now eliminated, but a few horizontal discontinuities are apparent. However, if we apply more vertical and horizontal smoothings, this process quickly converges to the smooth shifts shown in Figure 9c and 9d.

Although we have no guarantee that this smoothing process will converge, I have not found a practical example in which more than

four smoothings (vertical-horizontal-vertical-horizontal) yielded any significant changes in shifts. The convergence shown in Figure 9 is typical. Therefore, in practice, the computational complexity of this method for DIW remains  $O(N \times L)$  where  $N$  is now the number of image pixels, and  $L$  is again the number of lags.

Still, even assuming that the smoothing process does converge, we cannot guarantee that estimated shifts will minimize alignment errors while satisfying vertical and horizontal strain limits, as we recall that this constrained optimization problem has no feasible solution (Keyser and Unger, 2003).

To my knowledge TSDP has never been described as nonlinear smoothing followed by DTW. However, the new DIW method proposed here is truly an extension of the TSDP method proposed by Mottl et al. (2002). Indeed, one way to view this new method is that it is TSDP with a larger tree, in which each vertical or horizontal smoothing before DTW represents a new level of branches.

### Dynamic warping and crosscorrelation

In tests of DIW discussed above, shifts are large (much larger than the dominant period of reflections) and vary rapidly with time and offset. Recalling that strain is the rate at which shift changes, the maximum strain in time is about 25%, and the maximum strain in offset is almost 100%. That is, time shifts change by as much as a quarter of one time sample from one sampled time to the next, and by almost one time sample from one sampled offset to the next.

Where shifts are not so rapidly varying, methods based on local crosscorrelation of images may be used instead to obtain accurate shift estimates. Figure 10 displays shifts estimated from noisy shot records like those in Figure 8a and 8b, using DIW and local crosscorrelations. The local crosscorrelation method used here is that described by Hale (2009), which finds shifts that maximize correlation coefficients computed for seamlessly overlapping windows of images. In these tests, those windows are Gaussian with half-widths equal to 320 ms in time and 240 m in offset.

Figure 10 illustrates how the success of this crosscorrelation method depends on whether or not shifts vary rapidly within the windows used to compute the correlation coefficients. The sinusoidal pattern of variation used for these tests is the same as that shown in Figure 8d, but the rates at which shifts change with time and offset (the strains) are smaller because the magnitudes of the shifts are smaller.

Where shifts vary slowly, as in Figure 10a and 10b (where the maximum time shift is less than four samples), DIW and local crosscorrelation yield estimated shifts that approximate well the known shifts. Shifts estimated using the crosscorrelation method show significant errors only for small times and large offsets where no reflections exist. Where shifts vary rapidly, as in Figure 10c and 10d (10 times more rapidly than for Figure 10a and 10b), shifts estimated using the local crosscorrelation method are unstable and inaccurate, whereas those obtained by DIW again approximate well the known shifts.

One way to stabilize shifts estimated in crosscorrelation methods is maximize a weighted sum of image correlation and shift smoothness. Hall (2006), for example, used such a regularization to improve the stability and accuracy of small shifts estimated from time-lapse seismic images. Regularization is, however, not the same as constraints in DTW, and it does not solve the fundamental problem in using crosscorrelation windows to estimate shifts that vary rapidly within those windows. For such shifts, correlation




Figure 10. Time shifts estimated from noisy shot records by DIW (a, c) and by local crosscorrelation (b, d). Local crosscorrelation without constraints yields reasonable estimates where shifts vary slowly (b), but not where shifts vary rapidly (d).



coefficients will be low for all lags because the two windowed images cannot be well aligned for any one of them. If we try to solve this problem by using smaller windows, then noise will degrade our shift estimates, and in any case we must use correlation windows that are at least as large as the shifts we seek to estimate. In summary, local correlation methods require that we choose a window size; for shifts that vary rapidly, as for Figure 10d, a suitable choice may not exist.

In contrast, dynamic time and image warping require no windows. We need not specify a window size when using the DIW algorithm proposed in this paper.

Another difference between crosscorrelation and dynamic warping methods is that correlation peaks are easily found with subsample precision, but dynamic warping yields only integer shifts. For this reason, I smooth the integer shifts estimated with dynamic warping using Gaussian filters with half-widths equal to image sampling intervals divided by maximum strains used to constrain the shifts.

## APPLICATIONS

In all tests discussed in the previous section, shot records  $f(t, x)$  and  $g(t, x)$  input to the DIW algorithm are related by known shifts  $u(t, x)$ , where  $t$  denotes time and  $x$  denotes offset. More precisely,  $f(t, x) \approx g(t + u(t, x), x)$ , and the approximation is necessary only because the records  $f$  and  $g$  include the addition of independent bandlimited random noise. These tests illustrate the accuracy of DIW in the presence of such noise.

In practice, however, the shifts  $u(t, x)$  are unknown, and differences between the two images  $f(t, x)$  and  $g(t, x)$  cannot be attributed entirely to time shifts and random noise. Here, I consider the fidelity of DIW with two examples of practical applications.

### Registration of PP and PS images

One example is the misalignment of seismic reflectors in PP and PS images caused by differences in P-wave velocities  $V_P$  and S-wave velocities  $V_S$ . These velocity differences cause differences in reflection amplitudes that one may wish to analyze in estimations of rock properties. To facilitate such an analysis of amplitudes, we often align reflectors in the PS image with corresponding reflectors in the PS image. In doing so, we assume that any subsurface reflector creates PP and PS reflections, but that amplitudes and phases of these reflections may differ, due to differences in reflection coefficients.

Several authors, including Gaiser (1996), Fomel et al. (2003, 2005), and Nickel and Sonneland (2004), have described methods for registration (alignment) of PP and PS images and the corresponding estimation of  $V_P/V_S$  ratios. Here, I illustrate briefly the application of DIW to only the image registration problem.

Figure 11a and 11b displays short one-second windows of PS and PP images, respectively. Both images are actually eight seconds long, but these short windows facilitate visual comparisons before and after warping. Amplitudes in these two images have been normalized so that rms amplitudes equal one within seamlessly overlapping 2D windows that decay smoothly to zero. The PS image has been converted from PS time  $T_{PS}$  to PP time  $T_{PP}$  using the ratio  $T_{PP}/T_{PS} = 2/3$ , which corresponds to a constant ratio  $V_P/V_S = 2$ .

Within the one-second window displayed here, average  $V_P/V_S$  ratios are greater than two, so that reflectors in the PS image of Figure 11a appear at PP times greater than those in the PP image

of Figure 11b. Letting  $f(t, x)$  denote the PP image and  $g(t, x)$  denote the PS image (both functions of PP time  $t = T_{PP}$ ), DIW yields the warped PS image  $g(t + u(t, x), x)$  shown in Figure 11c, for the estimated time shifts  $u(t, x)$  displayed in Figure 11d. These estimated time shifts imply an average  $V_P/V_S \approx 2.5$  within this short time window, with a slight decrease from left to right.

Significantly, reflectors in the warped PS image of Figure 11c are well aligned with those in the PP image of Figure 11b, despite differences in noise, spectral bandwidth, and reflection amplitudes




Figure 11. Seismic reflectors in PS (a) and PP (b) images remain misaligned after converting PS time  $T_{PS}$  to PP time  $T_{PP}$  for a constant ratio  $T_{PP}/T_{PS} = 2/3$ . Alignment with reflectors in the PP image (b) is improved in the warped PS image (c), using additional time shifts (d) estimated by DIW.

apparent in the two images. However, the estimated shifts in Figure 11d vary rather slowly, so that for this example we might expect correlation-based methods to perform as well.

### Estimation of fault throws

Figure 12 provides a similar sequence of images for a rather different problem, that of estimating fault throws. The 2D image in Figure 12a was extracted from a 3D seismic image along the hanging-wall side of a fault surface. The 2D image in Figure 12b




Figure 12. Reflectors in hanging-wall (a) and footwall (b) images extracted from a 3D seismic image alongside a curvi-planar normal fault are misaligned due to fault throw. Reflectors in the footwall image (b) are well aligned with those in the warped hanging-wall image (c), using shifts (d) estimated by DIW.

was extracted along the footwall side of the same surface. Neither 2D image is axis-aligned, and neither corresponds to a vertical slice of the 3D image because the fault surface is curvi-planar and nowhere vertical.

Because the fault in this example is a normal fault, reflectors in the hanging-wall image of Figure 12a are displaced downward relative to those in the footwall image of Figure 12b. Letting  $f(t, x)$  denote the hanging-wall image and  $g(t, x)$  denote the footwall image, DIW yields the warped hanging-wall image  $g(t + u(t, x), x)$  shown in Figure 12c, for the estimated time shifts  $u(t, x)$  displayed in Figure 12d.

Zeros near the bottom of the warped hanging-wall image in Figure 12c occur because the maximum time displayed equals the maximum time for which the 3D seismic image is sampled. As samples in the hanging-wall image are warped upward, no image samples exist to replace the samples near the bottom.

Estimated time shifts in Figure 12d are related to vertical components of fault throw vectors, which in this example are positive because the fault is a normal fault. Although generally increasing with time, the estimated shifts (and corresponding fault throws) vary significantly within this short time window. Strong reflectors in the warped hanging-wall image in Figure 12c are well aligned with those in the footwall image in Figure 12b, which confirms visually that the shifts estimated using DIW and displayed in Figure 12d are correct.

As others have observed (e.g., [Borgos et al., 2003](#); [Admasu, 2008](#)), the problem of estimating fault throws is difficult. Part of the difficulty lies in detecting fault surfaces alongside which we may extract footwall and hanging-wall images. Moreover, in my experience, even after finding fault surfaces, 2D images like those displayed in Figure 12a and 12b may be impossible to construct because one part of a fault surface may lie in front of another part of the same surface. This situation is not unusual, and where it occurs a fault surface has, in effect, not two but four (or more) sides.

Fortunately, the DIW algorithm proposed here can be modified to compute displacements alongside such fault surfaces. The modification is beyond the scope of this paper, but the key step is again nonlinear accumulation, up-down along fault dip and left-right along fault strike, of alignment errors computed from samples of the 3D seismic image adjacent to the fault surface.

### CONCLUSION

An appealing feature of DTW is that it solves exactly the constrained optimization problem of equations 2, 3, and 4. Although a practical and exact solution is unlikely to exist for the corresponding optimization problem for images, examples shown above indicate that the approximate dynamic warping solution proposed in this paper can produce accurate estimates of shifts from images, even where those shifts are large and rapidly varying.

The examples in this study are for 2D images, but DIW is easily applied to 3D images as well. We smooth alignment errors along all three image dimensions before eventually using DTW to estimate the shifts. For each dimension, this smoothing and DTW is especially efficient on computers with multiple processors, as each row or column of alignment errors can be processed in parallel.

As for 2D images, computational complexity for 3D images remains  $O(N \times L)$ , except that  $N$  is now the number of samples (voxels) in the 3D images. The amount of computer memory required is also  $O(N \times L)$  and, in practice, this requirement may exceed the amount of RAM available. In this case, we may either

use slower computer memory or process large 3D images in smaller overlapping blocks, alternatives that are widely used today in other methods for 3D seismic image processing.

In this paper, I have assumed that images can be aligned with only vertical warping. However, the DIW algorithm proposed here can also be used to estimate shift vectors with vertical and horizontal components. Estimating shift vectors requires computation of alignment errors for multiple components of lag, but this computation and corresponding modifications to the image warping algorithm are straightforward.

## ACKNOWLEDGMENTS

I am grateful to Luming Liang for emphasizing to me the shortcomings of crosscorrelation methods in estimating shifts that vary rapidly, and to Andreas Rueger for his careful review of this manuscript. I also thank Dengliang Gao, James Gaiser, and Brad Artman for corrections and suggestions that inspired significant revision. The seismic shot record used to test DIW is one of 40 shot records gathered and made freely available by Oz Yilmaz. The fault images were derived from the F3 data set made freely available by dGB Earth Sciences, and the PP and PS images were provided courtesy of Sinopec.

## REFERENCES

- Admasu, F., 2008, A stochastic method for automated matching of horizons across a fault in 3D seismic data: Ph.D. thesis, Otto-von-Guericke-University Magdeburg.
- Anderson, K., and J. 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).
- Borgos, H., T. Skov, T. Randen, and L. Sonneland, 2003, Automated geometry extraction from 3D seismic data: 73rd Annual International Meeting, SEG, Expanded Abstracts, 1541–1544.
- Cormen, T., C. Leiserson, R. Rivest, and C. Stein, 2001, *Introduction to algorithms*: The MIT Press.
- Fomel, S., M. Backus, M. DeAngelo, P. Murray, and B. Hardage, 2003, Multicomponent seismic data registration for subsurface characterization in the shallow Gulf of Mexico: OTC, Expanded Abstracts, 15117.
- Fomel, S., M. Backus, K. Fouad, B. Hardage, and G. Winters, 2005, A multistep approach to multicomponent seismic image registration with application to a West Texas carbonate reservoir study: 75th Annual International Meeting, SEG, Expanded Abstracts, 1018–1021.
- 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).
- Hale, D., 2009, A method for estimating apparent displacement vectors from time-lapse seismic images: *Geophysics*, **74**, no. 5, V99–V107, doi: [10.1190/1.3184015](https://doi.org/10.1190/1.3184015).
- Hall, S., 2006, A methodology for 7D warping and deformation monitoring using time-lapse seismic data: *Geophysics*, **71**, no. 4, O21–O31, doi: [10.1190/1.2212227](https://doi.org/10.1190/1.2212227).
- Keyser, D., and W. Unger, 2003, Elastic image matching is NP-complete: *Pattern Recognition Letters*, **24**, 445–453, doi: [10.1016/S0167-8655\(02\)00268-4](https://doi.org/10.1016/S0167-8655(02)00268-4).
- 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).
- Mottl, V., A. Kopylov, A. Kostin, A. Yermakov, and J. Kittler, 2002, Elastic transformation of the image pixel grid for similarity based face identification: *Proceedings of the 16th International Conference on Pattern Recognition*, IEEE, 549–552.
- Müller, M., 2007, *Information retrieval for music and motion*: Springer.
- Nickel, M., and L. Sonneland, 2004, Automated PS to PP event registration and estimation of a high-resolution  $V_p - V_s$  ratio volume: 74th Annual International Meeting, SEG, Expanded Abstracts, 869–872.
- Pishchulin, L., 2010, *Matching algorithms for image recognition*: M. S. thesis, Rheinisch-Westfälischen Technischen Hochschule Aachen.
- Sakoe, H., and S. Chiba, 1978, Dynamic programming algorithm optimization for spoken word recognition: *IEEE Transactions on Acoustics, Speech, and Signal Processing*, **26**, 43–49, doi: [10.1109/TASSP.1978.1163055](https://doi.org/10.1109/TASSP.1978.1163055).