Image-guided 3D interpolation of borehole data

Dave Hale
Center for Wave Phenomena, Colorado School of Mines, Golden CO 80401, USA

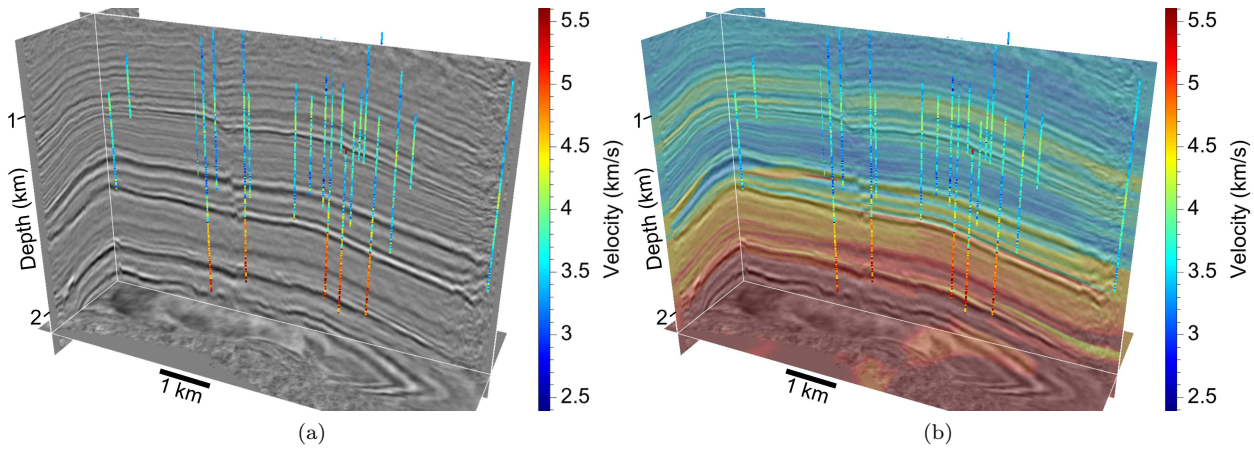

Figure 1. Slices of a 3D seismic image (a) with P-wave velocities measured in boreholes and an image-guided 3D interpolation (b) of those measured velocities.

ABSTRACT

A blended neighbor method for image-guided interpolation enables resampling of borehole data onto a uniform 3D sampling grid, without picking horizons and without flattening seismic images. Borehole measurements gridded in this way become new 3D images of subsurface properties. Property values conform to geologic layers and faults apparent in the seismic image that guided the interpolation.
The freely available Teapot Dome data set, which includes a 3D seismic image, horizons picked from that image, and numerous well logs, provides an ideal demonstration of image-guided interpolation of borehole data. In this example, seismic horizons picked by others coincide with thin layers apparent in the new 3D images of interpolated borehole data, even though the horizons were not used in the interpolation process.

Key words: seismic image well logs interpolation interpretation

1 INTRODUCTION

Seismic images are often used to guide the interpolation of subsurface properties that are measured more directly and (usually) more precisely in boreholes. Figure 1 provides an example for a 3D seismic image and sonic (P-wave velocity) logs from the Teapot Dome oilfield in Wyoming. These data are provided by the Rocky Mountain Oilfield Test Center, a facility of the U.S. Department of Energy (Anderson, 2009). Figure 1b shows interpolated velocities, displayed with translucent color so that the corresponding three slices of the 3D seismic image are visible as well. At depths where sonic logs are available, the interpolation of velocities is guided by the seismic image.

In a more conventional seismic interpretation, we might first pick horizons corresponding to coherent re-

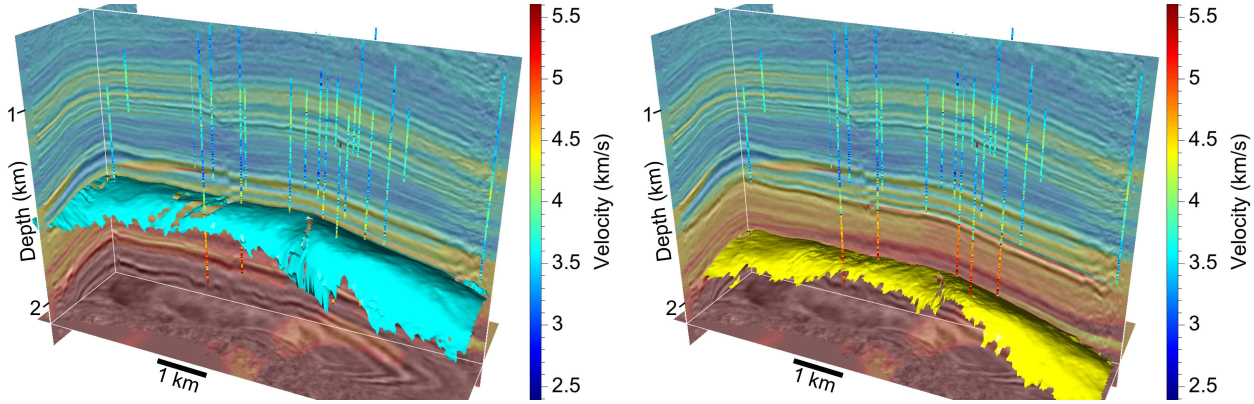

Figure 2. Two low-velocity layers in the 3D interpolated velocity image conform to the Crow Mountain (a) and Tensleep (b) horizons that were picked interactively (by others) from the 3D seismic image. Only the seismic image, not the horizons, was used to guide the 3D interpolation of the velocity logs.

flections in the seismic image. Two examples are shown in Figure 2. These two horizons correspond to the Crow Mountain and Tensleep formations, and are provided as part of the Teapot Dome data set. Typically, we would pick horizons like these interactively, with or without help from automatic event-tracking software. One reason we might construct horizon surfaces like these is to facilitate interpolation of properties measured in boreholes.

I interpolated the velocities shown in Figure 2 (and in Figure 1b) without using horizons. Instead, I used the seismic image to automatically and more directly guide 3D interpolation of the velocity logs. Although the horizons in Figure 2 were not used, they coincide with low-velocity layers apparent in the 3D interpolation shown in Figure 1b.

An obvious advantage of image-guided interpolation without horizons is that we save the time and effort of picking horizons. The savings may be significant, as seismic processing and interpretation have become interwoven parts of an iterative seismic imaging and inversion process. Another advantage in using a 3D seismic image directly is that we simultaneously interpolate at all locations between and on horizons that we might have picked. A third advantage is that our interpolation may be guided by images of geologic features, including unconformities and diapirs, that may be difficult to represent accurately and efficiently with picked surfaces.

Some of these advantages may be obtained by first flattening a seismic image (Stark, 2004; Lomask et al., 2006). By removing *structure* from a 3D image, flattening creates a stack of simpler 2D interpolation problems, like those we today solve routinely for 2D maps corresponding to picked horizons. However, automatic flattening as described by Stark (2004) and Lomask et al. (2006) is perfomed using vertical shifts that may distort distances measured within horizontal slices of a flattened image (Lee, 2001). Moreover, vertical shifts often cannot account for intrusions, such as overhanging salt diapirs; and flattened images are at best ambiguous in the presence of unconformities caused by erosion. Finally, flattening highlights *stratigraphic* features, such as channels, in 3D seismic images; and we may wish to use those features to guide the sequence of 2D interpolations. In other words, image-guided interpolation may be desirable even after flattening.

The purpose of this paper is to demonstrate image-guided interpolation of borehole data, without flattening and without picking horizons. I first review the blended neighbor interpolation method described by Hale (2009) using a 2D seismic image. I then describe the application of this method to the 3D seismic image and well logs from the Teapot Dome data set, and illustrate the method with several examples. Finally, I discuss current limitations and potential extensions of image-guided interpolation.

2 IMAGE-GUIDED INTERPOLATION

Let us assume that spatially scattered data to be interpolated are a set

$$\mathcal{F} = \{f_1, f_2, \ldots, f_K\} \tag{1}$$

of K known sample values $f_k \in \mathbb{R}$ that correspond to a set

$$\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_K\} \tag{2}$$

of K known sample points $\mathbf{x}_k \in \mathbb{R}^n$. Together these two sets comprise a set

$$\mathcal{K} = \{(f_1, \mathbf{x}_1), (f_2, \mathbf{x}_2), \ldots, (f_K, \mathbf{x}_K)\} \tag{3}$$

of K known samples. These samples may be scattered such that the n-dimensional sample points in the set

$\mathcal{X}$ may have no regular geometric structure. The classic interpolation problem is to use the known samples in $\mathcal{K}$ to construct a function $q(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}$, such that $q(\mathbf{x}_k) = f_k$.

As stated, this problem has no unique solution; there exist an infinite number of functions $q(\mathbf{x})$ that satisfy the interpolation conditions $q(\mathbf{x}_k) = f_k$. Additional criteria may include measures of smoothness, robustness, and efficiency. Because tradeoffs exist among such criteria, a variety of methods for interpolating scattered data are commonly used today.

In this paper I add the requirement that the interpolation should conform to features in a uniformly sampled image, as in Figures 1 and 2. That is, the interpolation must be *image-guided*.

2.1 Blended neighbor interpolation

The blended neighbor method (Hale, 2009) was developed specifically to facilitate image-guided interpolation. This process consists of two steps:

Step 1: solve the eikonal equation

$$\begin{aligned} \boldsymbol{\nabla}\, t(\mathbf{x}) \cdot \mathbf{D}(\mathbf{x})\, \boldsymbol{\nabla}\, t(\mathbf{x}) &= 1, \quad \mathbf{x} \notin \mathcal{X}; \\ t(\mathbf{x}_k) &= 0, \quad \mathbf{x}_k \in \mathcal{X} \end{aligned} \tag{4}$$

for

$t(\mathbf{x})$: the minimal time from $\mathbf{x}$ to the nearest known sample point $\mathbf{x}_k$, and

$p(\mathbf{x})$: the value f_k corresponding to the sample point $\mathbf{x}_k$ nearest to the point $\mathbf{x}$.

Step 2: solve the blending equation

$$q(\mathbf{x}) - \frac{1}{2}\, \boldsymbol{\nabla} \cdot t^2(\mathbf{x}) \mathbf{D}(\mathbf{x})\, \boldsymbol{\nabla}\, q(\mathbf{x}) = p(\mathbf{x}), \tag{5}$$

for the blended neighbor interpolant $q(\mathbf{x})$.

Here, *time* is simply a short word for *non-Euclidean distance*. By this measure of distance, a sample point $\mathbf{x}_k$ is *nearest* to a point $\mathbf{x}$ if the time $t(\mathbf{x})$ along some path to $\mathbf{x}_k$ is less than that for any other sample point. In step (1), I compute this minimal-time map $t(\mathbf{x})$ by solving the eikonal equation 4.

The metric tensor field $\mathbf{D}(\mathbf{x})$ provides the anisotropic and spatially varying coefficients of that eikonal equation. Intuitively, we must choose the tensor field $\mathbf{D}(\mathbf{x})$ so that, by our time measure of non-Euclidean distance, two points within the same geologic formation are near, while two points in different formations are much farther away. In this way, known sample values f_k for sample points x_k that are geologically nearby are given the most weight in any interpolated value $q(\mathbf{x})$.

In step (1), as I compute the time $t(\mathbf{x})$ from each point $\mathbf{x}$ to the location $\mathbf{x}_k$ of the nearest known sample, I also record the value $p(\mathbf{x}) = f_k$ of that nearest known sample. The function $p(\mathbf{x})$ is therefore a *nearest neighbor* interpolant.

In step (2), I compute the *blended neighbor* interpolant $q(\mathbf{x})$ by smoothing the nearest neighbor interpolant $p(\mathbf{x})$, and the extent of smoothing is controlled by the time map $t(\mathbf{x})$. At any known sample point $\mathbf{x}_k$, equation 4 states that $t(\mathbf{x}_k) = 0$, so that no smoothing is performed, and equation 5 becomes simply $q(\mathbf{x}_k) = p(\mathbf{x}_k) = f_k$. In other words, the function $q(\mathbf{x})$ interpolates exactly the known sample values.

Figure 3 illustrates the process of blended neighbor interpolation guided by a 2D seismic image. In this example I specified the 21 known samples illustated in Figure 3a. (For clarity, each sample is plotted with an opaque disk larger than the image pixel that represents the sample value.) The values $f_k \in [0, 1]$ in this example are arbitrary; I chose them to alternate vertically, while generally decreasing from left to right.

Figure 3b shows the time map $t(\mathbf{x})$ computed in step (1) for these known samples. The time map is displayed with translucent color on top of the seismic image displayed with shades of gray. Times are smallest near the known sample points $\mathbf{x}_k$ and largest in the corners that are farthest from any of those points. Contours of constant-time are not circular, because they are warped by the metric tensor field $\mathbf{D}(\mathbf{x})$. As described below, I computed this tensor field so that times would increase slowly in directions in which the seismic image is most coherent, while increasing rapidly near strong reflections and faults.

While computing the times $t(\mathbf{x})$ in step (1), I also computed the nearest neighbor interpolant $p(\mathbf{x})$ shown in Figure 3c. As expected, this interpolant conforms to structure in the seismic image, but it is discontinuous at locations $\mathbf{x}$ for which times $t(\mathbf{x})$ to two or more nearest known sample points $\mathbf{x}_k$ are equal.

These discontinuities are removed by solving the blending equation in step (2), which yields the continuous blended neighbor interpolant $q(\mathbf{x})$ shown in Figure 3d. Contours of constant color are well aligned with structures and faults in the seismic image, and interpolated values (colors) match the known sample values f_k at the known sample points $\mathbf{x}_k$.

As in the eikonal equation 4, the coefficients $\mathbf{D}(\mathbf{x})$ in the blending equation 5 are anisotropic and spatially varying. Therefore, the direction and extent of blending of the nearest neighbor values in step (2) depend on the metric tensor field $\mathbf{D}(\mathbf{x})$, in addition to the time map $t(\mathbf{x})$ computed in step (1).

2.2 Computing the metric tensor field

As discussed above, both steps (1) and (2) of blended neighbor interpolation are guided by a metric tensor field $\mathbf{D}(\mathbf{x})$. Blended neighbor interpolation becomes *image-guided* when we compute this tensor field from an image.

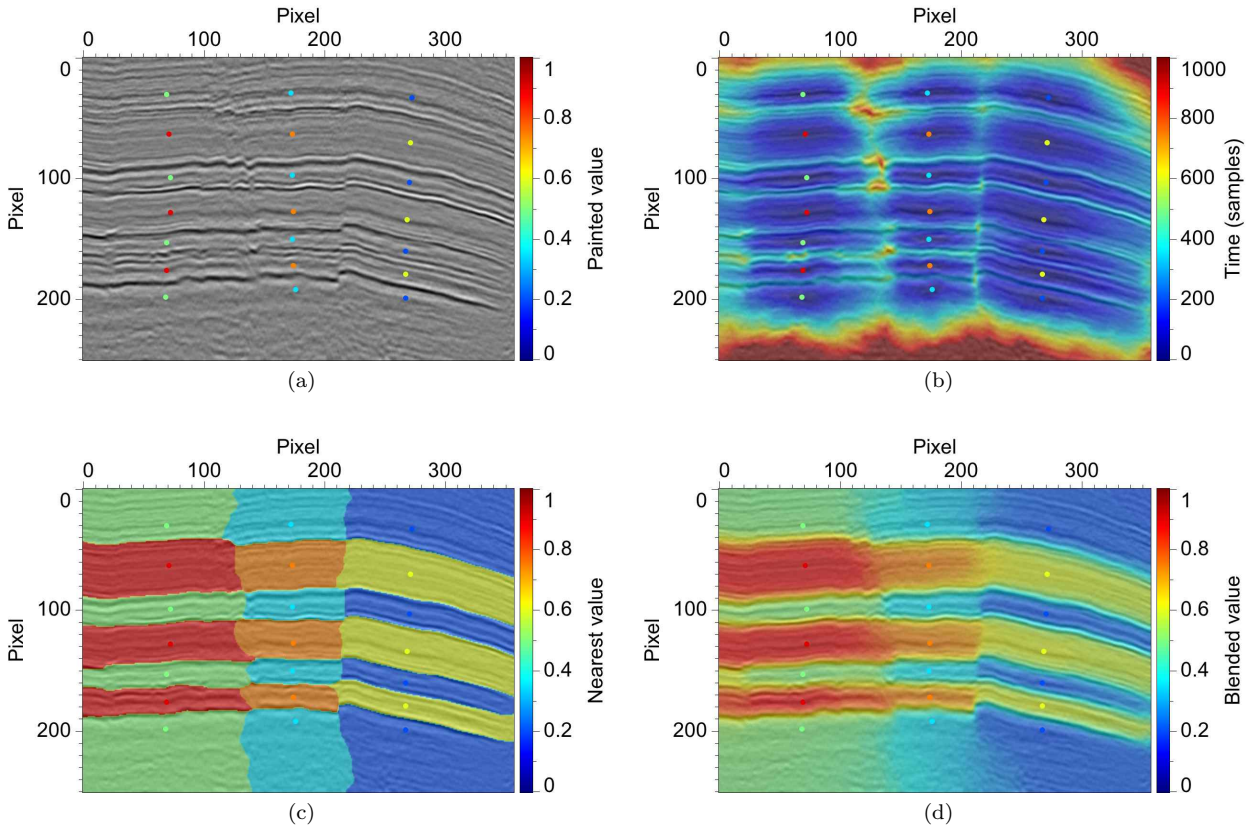

Figure 3. A simple example of image-guided 2D interpolation. For a set $\mathcal{K}$ of known (here, painted) samples (a), we first use equation 4 to compute the time map (b) and nearest neighbor interpolant (c), and then solve equation 5 for the blended neighbor interpolant (d).

I compute the metric tensor field $\mathbf{D}(\mathbf{x})$ from a seismic image by first computing structure tensors $\mathbf{S}(\mathbf{x})$. As described by van Vliet and Verbeek (1995) and Fehmers and Höcker (2003), these structure tensors are spatially smoothed outer products of image gradient vectors. In n dimensions, each structure tensor $\mathbf{S}$ is a symmetric positive-definite (SPD) $n \times n$ matrix, e.g., 2×2 for 2D images, and 3×3 for 3D images. I compute the eigenvectors of the metric tensors $\mathbf{D}(\mathbf{x})$ to be the same as those in the structure tensors $\mathbf{S}(\mathbf{x})$, but I modify the eigenvalues.

Equations 4 and 5 imply that the eigenvalues of $\mathbf{D}(\mathbf{x})$ have units of velocity squared. I scale the tensor field $\mathbf{D}(\mathbf{x})$ so that the maximum eigenvalue (maximum velocity squared) for any of these tensors is one. Eigenvalues less than one therefore imply slower velocities and larger times in directions of the corresponding eigenvectors. Times will be smaller in directions for which velocities (eigenvalues) are faster (larger).

In the 2D example of Figure 3, I computed the metric tensor field $\mathbf{D}(\mathbf{x})$ from a structure tensor field $\mathbf{S}(\mathbf{x})$ by

$$\mathbf{D}(\mathbf{x}) = s\frac{\mathbf{S}^{-1}(\mathbf{x})}{1 - c(\mathbf{x})}. \tag{6}$$

I computed the constant scale factor s so that the maximum eigenvalue in the metric tensor field $\mathbf{D}(\mathbf{x})$ is one. The function $c(\mathbf{x})$ is a measure of coherence or semblance (e.g., Bahorich and Farmer, 1995), computed for each image pixel along the slope of the most linear feature at that pixel. Any such measure in the range $0 \leq c(\mathbf{x}) < 1$ could be used. The effect of the divisor $1 - c(\mathbf{x})$ is to increase the eigenvalues of $\mathbf{D}(\mathbf{x})$, thereby decreasing times $t(\mathbf{x})$, between locations where features in images are most coherent.

Note that each matrix $\mathbf{D}$ in the metric tensor field $\mathbf{D}(\mathbf{x})$ is SPD, because the eigenvalues of each inverse matrix $\mathbf{S}^{-1}$ in equation 6 are reciprocals of the corresponding positive eigenvalues of an SPD matrix $\mathbf{S}$.

I chose the values and colors in Figure 3 to highlight the ability of image-guided interpolation to conform to structures and faults apparent in a seismic image. In practice we might interactively paint values that are

more realistic. Alternatively, the known samples might correspond to geophysical data, such as well logs.

3 TEAPOT DOME EXAMPLE

The freely available Teapot Dome data set, which includes a time-migrated 3D seismic image and hundreds of well logs (Anderson, 2009), enables a realistic demonstration of image-guided 3D interpolation of borehole data.

3.1 Seismic image

To reduce the large number of zero traces in the 3D seismic image, I rotated and trimmed the seismic survey coordinate rectangle, as shown in Figure 4. As illustrated by the constant-time slice shown there, roughly half of the traces in the original 3D seismic image are zero. After resampling to a spatial grid aligned with the solid (red) coordinate rectangle, a smaller fraction of traces are zero, and the spatial coordinate axes are more nearly aligned with the anticlinal structure apparent in the image. The original spatial sampling intervals are 110 ft in both inline and crossline directions. I used 2D sinc interpolation to interpolate traces on the resampled grid with spatial sampling intervals of 25 m in both directions.

In the original seismic survey coordinate system, the (east-west) axis is the inline direction, and the longer (north-south) axis is the crossline direction. Although the original and resampled coordinate grids are not aligned (because of the rotation in the coordinate transformation), I hereafter refer to the shorter (northeast-southwest) resampled coordinate axis as the *inline* direction and the longer (northwest-southeast) one as the *crossline* direction.

Figure 5 shows two sets of three orthogonal slices of the resampled seismic image, after conversion of the vertical axis from time to depth. (I discuss the time-to-depth conversion process in the next section.) In a typical 3D seismic image, many such slices are possible. I chose these slices for their intersections with structural features apparent in the image and with wells.

3.2 Well logs

Well logs in the Teapot Dome data set are provided in two groups. The numerous so-called "shallow" wells do not penetrate to the depths displayed in Figure 5. None of the "deeper" wells extends to the bottom depth (2.2 km) shown there, and less than twenty of them extend to the depth of the 1.5 km slice shown in Figure 5b. In the examples shown in this paper, I ignored all borehole data provided with the shallow wells. The velocity logs displayed in Figure 1 are those available for the deeper wells.

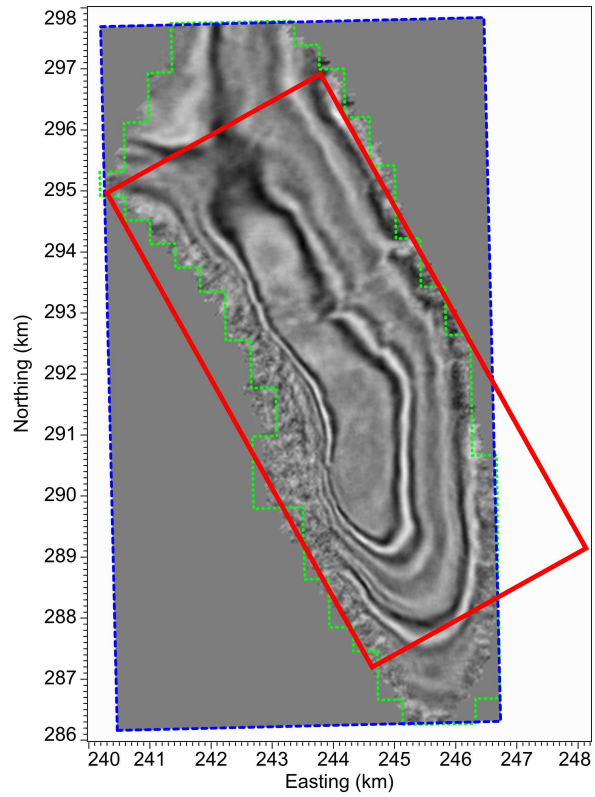

Figure 4. Original (dashed blue) and resampled (solid red) coordinate rectangles for the Teapot Dome data set. The (dotted green) polygon is the boundary of the Teapot Dome oilfield. The constant-time (0.95 s) slice shown here illustrates the anticlinal structure apparent in the 3D seismic image.

For the purpose of demonstrating image-guided interpolation, I selected four types of well logs: P-wave velocity, density, porosity and gamma ray.

Well logs are provided in LAS (Log ASCII Standard) format, and directional surveys are provided for boreholes that are not vertical. I performed only minimal pre-processing of the well logs. Specifically, I used elevations (of the kelly bushing, derrick floor, etc.) and directional survey data to convert distances measured along boreholes to inline, crossline and depth coordinates in the resampled seismic coordinate system. I discarded all logs with missing or clearly invalid elevation data, and all logs not entirely contained within the spatial boundaries of the resampled seismic image volume shown in Figure 5.

I also discarded entirely any well logs that contain clearly erroneous values: velocities outside the range $[0.2, 20]$ km/s, densities outside the range $[0.5, 10.0]$ gm/cc, porosities outside the range $[0, 0.8]$, and gamma ray radioactivies outside the range $[0, 300]$ API units. These bounds are broad and were chosen to exclude only those logs that contained data that are obviously

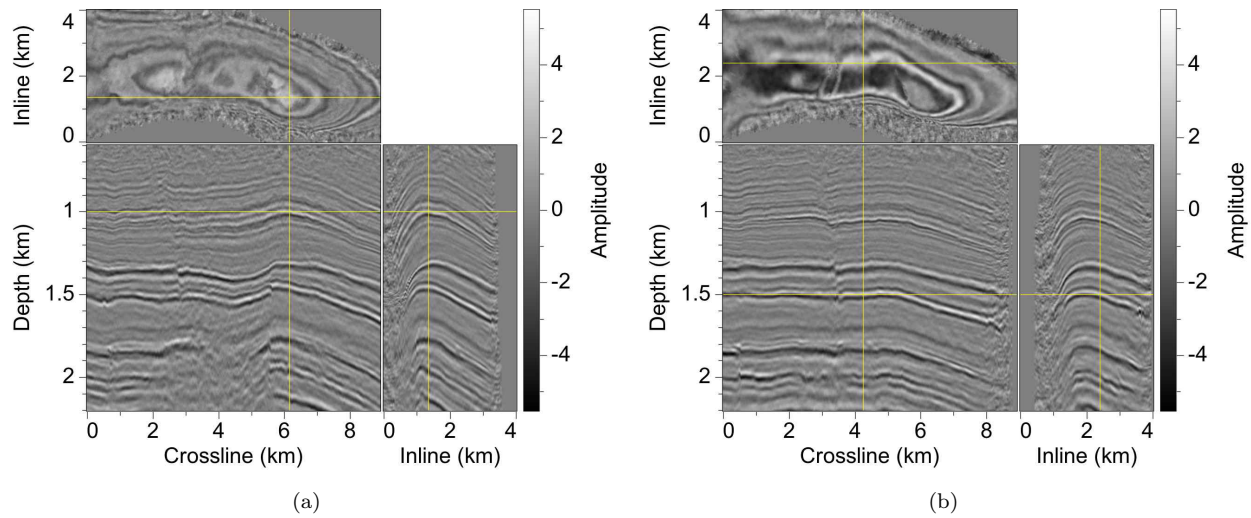

Figure 5. Two sets (a) and (b) of three orthogonal slices of the 3D seismic image used to guide interpolation of Teapot Dome borehole data. The horizontal constant-depth slice at 1 km (a) is intersected by many more wells than is the deeper slice at 1.5 km (b).

invalid. As shown below, the remaining well logs certainly contain measurements with significant errors, and those errors are especially apparent after image-guided 3D interpolation.

Whereas well logs are sampled every six inches along the boreholes, the depth sampling interval for the seismic image is 4 m. This difference in spatial sampling intervals (roughly a factor of 25) raises an important question. How finely should we sample the interpolated borehole data?

3.3 Initial gridding of well logs

Interpolation on a fine grid that would preserve all detail in the well logs would be about 25 times more costly than interpolation on the vertically coarser seismic grid. This high cost might be reduced by interpolating for only a subset of the seismic image and well data. For computational efficiency and convenience in this demonstration, I sampled interpolated values using the sampling intervals of the seismic image: 25 m in both inline and crossline directions, and 4 m in depth. Here *the interpolation grid* is that of the 3D seismic image.

For each type of log — velocity, density, porosity and gamma ray — I obtained the set of known samples defined by equation 3 with a simple binning and averaging procedure. First, I rounded the spatial coordinates of each well log sample to the coordinates of the nearest bin in the interpolation grid. Each known sample location $\mathbf{x}_k$ therefore corresponds to one such bin, and each known sample value f_k is the average of all well log samples for which $\mathbf{x}_k$ is the nearest bin. After this binning and averaging procedure, only those bins in the interpolation grid that are intersected by well logs of the appropriate type have values. Values for other bins in the grid are unknown and will be interpolated using the two-step process of equations 4 and 5.

3.4 Computing the tensor field

Before solving equations 4 and 5, we must first specify a metric tensor field $\mathbf{D}(\mathbf{x})$. As in the 2D example of Figure 3, I derived $\mathbf{D}(\mathbf{x})$ from structure tensors $\mathbf{S}(\mathbf{x})$ computed from the seismic image. For the 3D seismic image displayed in Figure 5, each structure tensor $\mathbf{S}$ is a 3×3 SPD matrix with eigen-decomposition

$$\mathbf{S} = \lambda_u \mathbf{u}\mathbf{u}^T + \lambda_v \mathbf{v}\mathbf{v}^T + \lambda_w \mathbf{w}\mathbf{w}^T, \tag{7}$$

where λ_u, λ_v and λ_w are the eigenvalues and $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{w}$ the corresponding eigenvectors of $\mathbf{S}$.

Let us label the eigenvalues and eigenvectors of $\mathbf{S}$ so that $\lambda_u \geq \lambda_v \geq \lambda_w \geq 0$. Then, eigenvectors $\mathbf{u}$, corresponding to the largest eigenvalues λ_u, indicate directions in which image gradients are highest, orthogonal to features that are locally linear or planar. The eigenvectors $\mathbf{w}$, corresponding to the smallest eigenvalues λ_w, will be aligned with locally linear features, such as channels and the intersections of geologic faults and layers, in seismic images. Both eigenvectors $\mathbf{v}$ and $\mathbf{w}$ lie within the planes of any locally planar features.

In other words, for each image sample, the orthonormal eigenvectors $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{w}$ specify the local orientation of the predominant image feature. The corresponding eigenvalues λ_u, λ_v and λ_w contain information about the shape of that feature. For example, locally linear features correspond to eigenvalues $\lambda_u \approx \lambda_v \gg \lambda_w$. For locally planar features, $\lambda_u \gg \lambda_v \approx \lambda_w$.

The eigenvalues λ_u, λ_v and λ_w of structure tensors

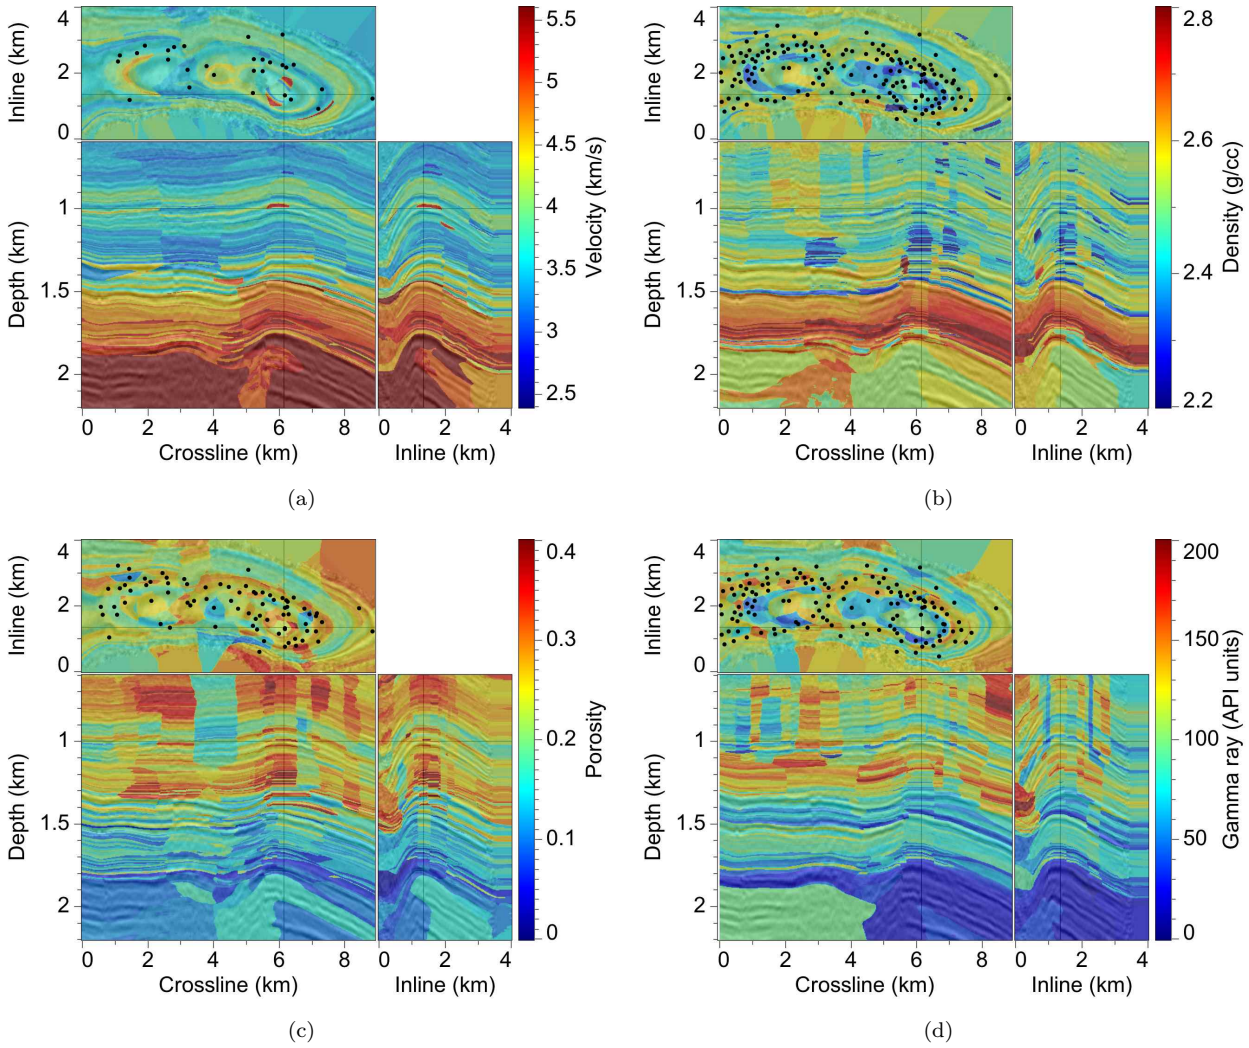

Figure 6. Image-guided nearest neighbor interpolation of velocity (a), density (b), porosity (c) and gamma ray (d) logs. Slices here correspond to those displayed for the seismic image in Figure 5a.

S are proportional to the magnitudes of image gradients squared, and therefore depend on the amplitudes of events in seismic images. Geologically, weak events may be as significant as strong ones; important geologic interfaces may or may not correspond to large contrasts in acoustic impedance.

Therefore, in image-guided interpolation, I discard the eigenvalues λ_u, λ_v and λ_w of the structure tensors **S** and use normalized local measures of semblance (coherence) to compute metric tensors

$$\mathbf{D} = \lambda_3 \mathbf{u}\mathbf{u}^T + \lambda_2 \mathbf{v}\mathbf{v}^T + \lambda_1 \mathbf{w}\mathbf{w}^T, \tag{8}$$

such that $0 \leq \lambda_3 \leq \lambda_2 \leq \lambda_1 \leq 1$.

Each eigenvalue λ_1 corresponds to a semblance of image samples measured along a 1D curvi-linear trajectory defined by eigenvectors **w**. Likewise, each eigenvalue λ_2 corresponds to a semblance of image samples along a 2D curvi-planar surface defined by the eigenvectors **v** and **w**. Finally, each eigenvalue λ_3 corresponds to a semblance of image samples within a local isotropic 3D window.

Recall that the eigenvalues of **D** have units of velocity squared. (See equation 4.) In equation 8 these eigenvalues are semblances, normalized measures of coherence in the range [0,1]. Therefore, at locations and in directions where semblance equals one, time in the map $t(\mathbf{x})$ is equivalent to Euclidean distance. Time exceeds Euclidean distance at locations and in directions where semblance is less than one, where image samples are less coherent.

In regions with no seismic image, the dead traces in Figure 5, I specified eigenvalues $\lambda_1 = \lambda_2 = 1$, $\lambda_3 = 0.01$, and eigenvectors **u**, **v** and **w** aligned with depth, crossline and inline coordinate axes, respectively.

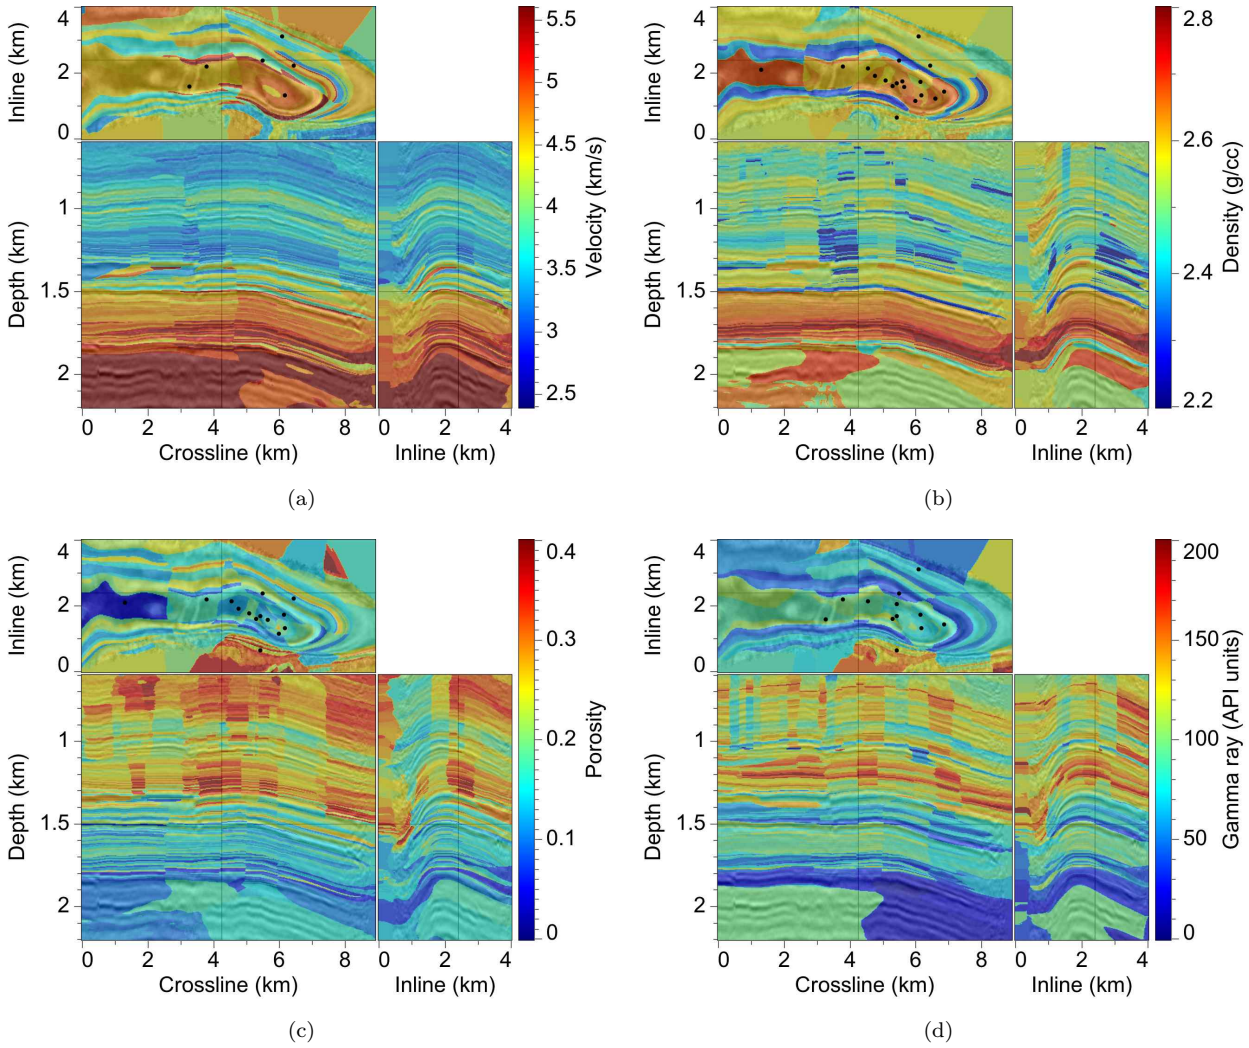

Figure 7. Image-guided nearest neighbor interpolation of velocity (a), density (b), porosity (c) and gamma ray (d) logs. Slices here correspond to those displayed for the seismic image in Figure 5b.

These default metric tensors $\mathbf{D}$ correspond to planar horizontal layering.

3.5 Nearest neighbor interpolation

The known samples $(f_k, \mathbf{x}_k)$ obtained by initial gridding of well log data and the tensor field $\mathbf{D}(\mathbf{x})$ computed from the image are the parameters required for step (1) of image-guided interpolation. In this step I simultaneously compute both the time map $t(\mathbf{x})$ and the nearest neighbor interpolant $p(\mathbf{x})$ by solving a finite-difference approximation of the eikonal equation 4.

Recall that "nearest" here implies *nearest in time*, based on a non-Euclidean distance that is defined by the metric tensor field $\mathbf{D}(\mathbf{x})$. Because the eigenvalues of $\mathbf{D}(\mathbf{x})$ are computed from semblances measured in local $(\mathbf{u}, \mathbf{v}, \mathbf{w})$ coordinate systems, times along paths of high image semblance (within imaged layers) are relatively small, while those along paths of low semblance (across imaged faults or layers) are relatively large.

Figures 6 and 7 display the nearest neighbor interpolants $p(\mathbf{x})$ for four different borehole measurements. Again, interpolated values are displayed with translucent color on top of the seismic image used to guide the interpolation.

For each log type, black dots in the horizontal constant-depth slices indicate the intersections of well logs with those slices. These dots represent only a tiny subset of the well log samples used to perform the 3D interpolation.

Figure 6 shows that a relatively small number of wells have velocity logs at a depth of 1 km, while a much larger number of wells have density, porosity and gamma ray logs at that depth. Well intersections plotted

at a depth of 1.5 km in Figure 7 indicate, for all four log types, that a much smaller number of well logs extend to this depth.

For example, Figure 7a indicates that only six velocity logs extend to a depth of 1.5 km. However, the velocity variation seen in the constant-depth slice at 1.5 km is not the result of interpolating only six logged velocities. The interpolation is three-dimensional, so that many logged velocities above and below this slice contribute to the interpolated velocities shown there.

As in the 2D example of Figure 3c, all of the nearest neighbor interpolants shown in Figures 6 and 7 exhibit discontinuities. I chose the slices displayed in Figure 6 specifically to highlight some of those discontinuities. Most of those discontinuities do not coincide with geologic faults. Rather, they reflect inconsistencies among properties measured within wells and those measured within their nearest neighbor wells.

For example, anomalously low (light blue) porosities are apparent in the upper middle part of the vertical crossline slice in Figure 6c. These low porosities are suspect because they are inconsistent with those measured in wells that are nearby with respect to the non-Euclidean metric tensor field $\mathbf{D}(\mathbf{x})$ computed from the seismic image. Image-guided nearest neighbor interpolation may lead us to look more closely at the porosity logs of nearby wells, to look for possible sources of error.

In the same way, we may use consistency with nearest neighbors as a measure of the fidelity of each well log sample. For example, the three slices of interpolated velocities shown in Figure 6 intersect an apparent high-velocity anomaly. These high velocities are likely caused by erroneous samples in the nearest velocity log. This hypothesis is supported by the fact that the anomaly coincides with the shallowest samples, which were acquired last, for that log. In the well logs provided with the Teapot Dome data set, the deepest (first recorded) and shallowest (last recorded) samples often exhibit anomalous values.

At depths greater than 1.9 km, large areas of constant interpolated values are apparent in Figures 6 and 7. Because no wells extend to these depths, all of the well log samples that lie in shallower geologic layers appear to be relatively far away, so that the nearest neighbor sample value is a poor interpolant.

3.6 Blended neighbor interpolation

Step (2) of image-guided interpolation is the solution of a finite-difference approximation of the blending equation 5. Parameters in this equation include the metric tensor field $\mathbf{D}(\mathbf{x})$, as well as the time map $t(\mathbf{x})$ and nearest neighbor interpolant $p(\mathbf{x})$. Figures 8 and 9 show slices of blended neighbor interpolants $q(\mathbf{x})$ corresponding to the nearest neighbor interpolants $p(\mathbf{x})$ shown in Figures 6 and 7.

As illustrated by these examples, the blending equation 5 smooths the nearest neighbor interpolants, and the extent of smoothing is controlled by the time map $t(\mathbf{x})$. Little smoothing is performed at locations $\mathbf{x}$ near the known well log samples, where times $t(\mathbf{x})$ are small; more smoothing is performed where those times are larger. In step (2) the metric tensor field $\mathbf{D}(\mathbf{x})$ causes this smoothing to be performed along seismically imaged geologic layers, but not across those layers or across faults. In this sense, the blending step (2) is an averaging of values from neighbors that are geologically nearby.

When solving the blending equation 5, I clipped all times in the time map $t(\mathbf{x})$ to be less than 10. Recall that, at locations and in directions where semblances are highest (that is, where eigenvalues of $\mathbf{D}$ equal one), one unit of time is equivalent to one spatial sample.

This time constraint limits the amount of smoothing performed. Where well log samples are dense, times are small anyway, and this limit has no effect on blended neighbor interpolants. In regions more sparsely sampled by well logs, this limit causes the blended neighbor interpolant to appear more like the nearest neighbor interpolant.

Setting an upper bound on times in the map $t(\mathbf{x})$ is analogous to setting an upper bound on the distance at which subsurface properties are correlated, as in kriging (e.g., Cressie, 1993). The difference here is that distance is defined by the metric tensor field $\mathbf{D}(\mathbf{x})$.

This upper bound also reduces the computational cost of solving the finite-difference approximation of the blended equation 5. For the conjugate-gradient solver that I use, that cost grows linearly with times in the map $t(\mathbf{x})$. For this example, the computation time required to solve the blending equation 5 is a few minutes on a modern workstation, roughly one tenth of the time required to solve the eikonal equation 4. If times had not been clipped, this cost would have been much higher.

Finally, by limiting the times in the map $t(\mathbf{x})$, we limit the range of influence of anomalous well-log values. After such values have been found and, if erroneous, corrected or discarded, we might increase the upper bound on times in $t(\mathbf{x})$, and thereby permit smoothing over greater non-Euclidean distances.

4 DISCUSSION

The Teapot Dome example demonstrates the process of image-guided 3D interpolation of borehole data. Instead of first picking horizons or flattening a seismic image, we may use the image to define a non-Euclidean metric tensor field that directly guides interpolation.

4.1 Two interpolants

In practice both the nearest neighbor and blended neighbor interpolants may be useful. The nearest neigh-

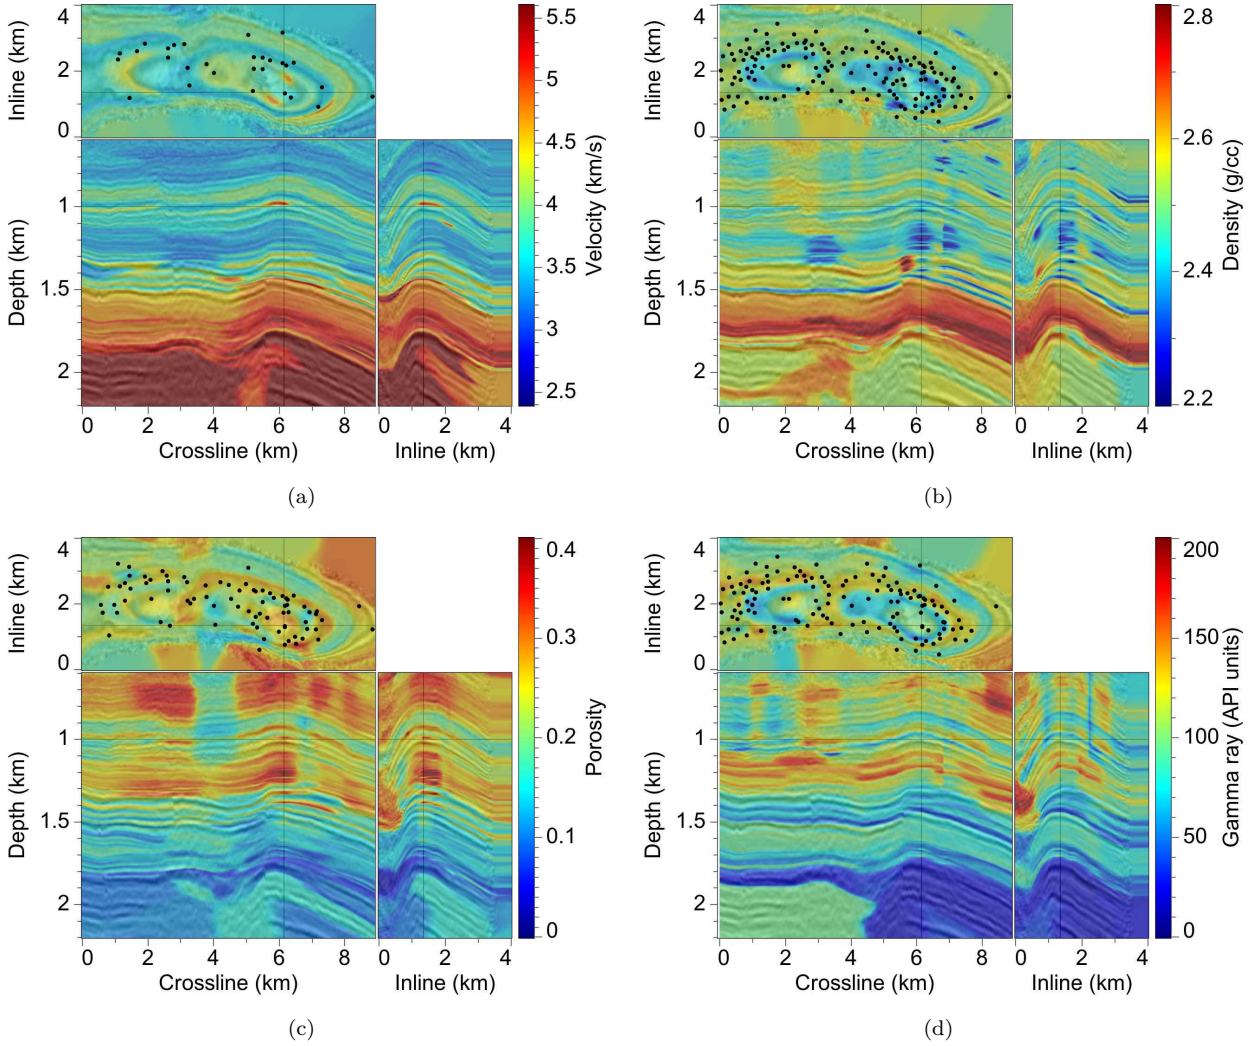

Figure 8. Image-guided blended-neighbor interpolation of velocity (a), density (b), porosity (c) and gamma ray (d) logs. Slices here correspond to those displayed for the seismic image in Figure 5a.

bor interpolant may be used to detect inconsistencies in borehole data acquired within the same seismically imaged geologic layers. Well log sample values that are inconsistent with those of geologically nearby log samples may be erroneous and perhaps should be discarded.

For example, we might compute, for each well log sample, the difference between the sample value and the mean of its nearest neighbor values. We might then discard log samples for which that difference exceeds some multiple of the standard deviation of the nearest neighbor values. We could also use more robust statistical measures in similar ways.

The nearest neighbor interpolant is also a useful first step toward computing the blended neighbor interpolant. Within seismically imaged layers, the blended neighbor interpolant is continuous and therefore geologically more reasonable than the discontinuous nearest neighbor interpolant.

The blended neighbor interpolants shown in Figures 8 and 9 are consistent with the borehole data and structures apparent in the corresponding seismic image. These interpolants are also consistent with expected trends. Velocities tend to increase with depth and porosities tend to decrease with depth. Also evident are some significant deviations from those trends.

For example, the strong reflector at a depth of about 1.5 km coincides with a significant change in both velocity and density, the factors of acoustic impedance. A thin layer at that depth with relatively low velocity, low density, high porosity, and low gamma ray radioactivity corresponds to the Crow Mountain sandstone formation marked by the light-blue horizon displayed in

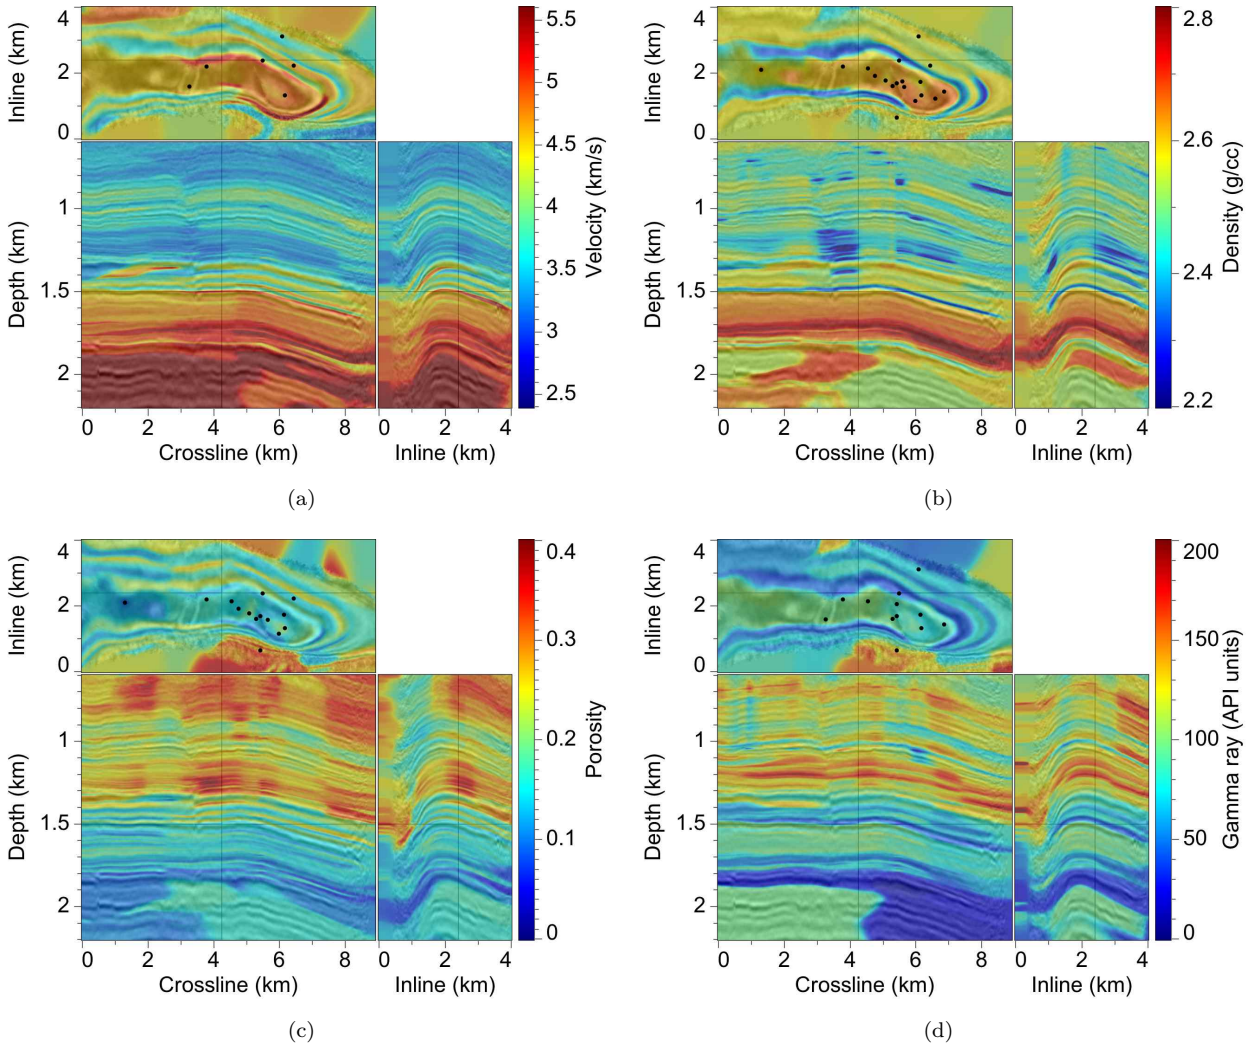

Figure 9. Image-guided blended-neighbor interpolation of velocity (a), density (b), porosity (c) and gamma ray (d) logs. Slices here correspond to those displayed for the seismic image in Figure 5b.

Figure 2a. The low (dark blue) density of this formation is especially visible in the slices of interpolated densities.

As another example, the Tensleep sandstone formation marked by the light-yellow horizon in Figure 2b corresponds to the low-velocity (light yellow) layer apparent at a depth of about 1.8 km/s in Figures 8 and 9. At the depths of both the Tensleep and Crow Mountain formations, the interpolated velocities shown here depend on velocities logged in only six wells.

After interpolating relevant borehole data onto a shared uniform 3D sampling grid, thereby creating 3D images of subsurface properties, we can easily combine them to create other images. For example, we might use the velocity and density images to compute a 3D image of acoustic impedance.

4.2 Time-to-depth conversion

Before seismic images can be used to guide interpolation of borehole data, the vertical axis of those images must be converted from vertical two-way time to depth. This conversion requires a uniformly sampled function $\tau(\mathbf{x}) = \tau(x, y, z)$ that specifies, for each point with horizontal coordinates x and y and depth coordinate z, the corresponding vertical two-way time τ. Given the uniformly sampled function $\tau(x, y, z)$, it is easy to convert a time-migrated seismic image $s_\tau(x, y, \tau)$ from time to depth using the mapping

$$s_z(x, y, z) = s_\tau[x, y, \tau(x, y, z)], \tag{9}$$

where $s_z(x, y, z)$ is the seismic image after time-to-depth conversion. The more difficult task is to construct the uniformly sampled function $\tau(x, y, z)$.

In the Teapot Dome example, that function was constructed in a typical manner, by correlating seismic horizons picked on time-migrated images with corresponding features in well logs (D. Witte, personnel communication, 2009). Specifically, for all well logs intersecting a seismic horizon, points with horizontal coordinates x and y and times τ were chosen from the horizon, and corresponding points with approximately the same x and y coordinates and depths z were chosen from the log. Depths z where then interpolated, first within each horizon for all x and y, using a minimum-curvature algorithm (Briggs, 1974), and then vertically between horizons for all times, using a simple linear interpolation, to obtain a uniformly sampled function $z(x, y, \tau)$. A simple inverse linear interpolation was then used to obtain the required function $\tau(x, y, z)$.

Image-guided interpolation suggests an alternative to this procedure that does not require picking seismic horizons. As in the typical procedure, we may first choose points (x, y, τ) from the seismic image and corresponding points (x, y, z) from the well logs. These points comprise scattered known samples of $z(x, y, \tau)$ that we may interpolate, using the time-migrated 3D seismic image $s_\tau(x, y, \tau)$ to guide our interpolation. Again, inverse linear interpolation would yield the required uniformly sampled function $\tau(x, y, z)$.

Other alternatives include direct interpolation of vertical traveltimes τ measured in checkshot surveys or vertical seismic profiles. In all of these alternatives, we replace three steps — (1) horizon picking, (2) interpolation within horizons, and (3) interpolation between horizons — with image-guided 3D interpolation.

4.3 Other interpolation methods

I developed the two-step blended neighbor method for image-guided interpolation to be both intuitive and computationally efficient. The method is intuitive because the blended neighbor interpolant is a smoothed version of the simplest nearest neighbor interpolant. The method is efficient primarily because it does not require the computation of times (non-Euclidean distances) from every interpolation grid point $\mathbf{x}$ to every known sample point $\mathbf{x}_k$.

Blended neighbor interpolation requires only the time $t(\mathbf{x})$ to the *nearest* (smallest in time) known sample point $\mathbf{x}_k$. Computation of the time map $t(\mathbf{x})$ displayed in Figure 3b does not require times from every interpolation grid point $\mathbf{x}$ to every known sample point $\mathbf{x}_k$.

Some well-known alternative methods, such as interpolation using Green's functions or radial basis functions (e.g., Wessel and Bercovici, 1998) or kriging (e.g., Cressie, 1993), require many more distance computations. For a constant metric tensor field $\mathbf{D}(\mathbf{x}) = \mathbf{D}$, the cost of computing these many distances is insignificant. However, the cost of computing non-Euclidean distances in a spatially varying metric tensor field $\mathbf{D}(\mathbf{x})$ is much higher, requiring numerical solution of the eikonal equation 4. This high cost makes many well-known alternative methods impractical for image-guided 3D interpolation of borehole data.

4.4 Limitations

When interpolating velocities and densities, properties that determine acoustic impedance, we should use seismic amplitudes to help estimate these properties between boreholes. However, as described here, image-guided 3D interpolation uses only estimates of image structure and semblance to guide interpolation of borehole data. It does not directly use the amplitudes of seismic reflections.

Seismic reflection amplitudes are especially useful in quantifying rapid vertical variations in velocities and densities. Those amplitudes are often less useful in quantifying long-wavelength vertical variations, because low frequencies are typically absent in recorded seismograms. One possible use of image-guided 3D interpolation would be to provide an a priori long-wavelength model for a more sophisticated joint inversion of seismic amplitudes and borehole data.

Another current limitation of image-guided interpolation is its reliance entirely on structure tensors $\mathbf{S}(\mathbf{x})$ computed from seismic images. While such automatic estimates of the orientations and shapes of subsurface structures are typically more reliable than reflection amplitudes, seismic interpreters routinely pick reflectors in noisy 3D seismic images for which automatic methods would fail. Moreover, not all subsurface properties conform to reflectors in seismic images. In practice a semi-automatic interpolation process, one guided by both seismic images and human interpreters, is likely to be optimal.

5 CONCLUSION

Notwithstanding its current limitations, image-guided interpolation provides an attractive new method for using a 3D seismic image to interpolate subsurface properties measured in boreholes. The method requires only a metric tensor field, which I compute automatically from the image, and the borehole data to be interpolated. In contrast to methods widely used today, image-guided interpolation does not require picking seismic horizons or faults; nor does it require image flattening.

The examples for the Teapot Dome data shown in this paper illustrate that the method produces interpolants consistent with seismic horizons picked by others. For depths where borehole data have been acquired, the most significant errors in the interpolants likely correspond to errors in well logs. The nearest neighbor interpolant naturally highlights such errors, as they cause

significant lateral discontinuities in interpolated subsurface properties at locations halfway (in time) between well logs. This observation suggests that we might use nearest neighbor interpolants to quantify the spatial consistency and, hence, the fidelity of well log samples.

Unlike nearest neighbor interpolants, blended neighbor interpolants are continuous, the latter being simply smoothed versions of the former. The extent of smoothing depends on times, non-Euclidean distances, to the nearest borehole measurements. By limiting these times to not exceed a specified maximum, we can reduce both the influence of erroneous measurements and the computational cost of image-guided interpolation.

ACKNOWLEDGMENTS

Thanks to the Rocky Mountain Oilfield Test Center, a facility of the U.S. Department of Energy, for providing the 3D seismic image, horizons, and well logs used in this study. Thanks also to Transform Software and Services, especially Amelia Webster, for providing their time-to-depth conversion of the 3D seismic image. I also thank Chris Engelsma for explaining to me the uses and significance of various types of well logs.

REFERENCES

Anderson, T., 2009, History of geologic investigations and oil operations at Teapot Dome, Wyoming: Presented at the 2009 AAPG Annual Convention.

Bahorich, M., and S. Farmer, 1995, The coherence cube: The Leading Edge, **14**, 1053–1058.

Briggs, I., 1974, Machine contouring using minimum curvature: Geophysics, **39**, 39–48.

Cressie, N., 1993, Statistics for spatial data: Wiley.

Fehmers, G., and C. Höcker, 2003, Fast structural interpretation with structure-oriented filtering: Geophysics, **68**, 1286–1293.

Hale, D., 2009, Image-guided blended neighbor interpolation of scattered data: 79th Annual International Meeting, SEG, Expanded Abstracts, 1127–1131.

Lee, R., 2001, Pitfalls in seismic data flattening: The Leading Edge, **20**, 160–164.

Lomask, J., A. Guitton, S. Fomel, J. Claerbout, and A. Valenciano, 2006, Flattening without picking: Geophysics, **71**, P13–P20.

Stark, T., 2004, Relative geologic time (age) volumes — Relating every seismic sample to a geologically reasonable horizon: The Leading Edge, **23**, 928–932.

van Vliet, L., and P. Verbeek, 1995, Estimators for orientation and anisotropy in digitized images: Proceedings of the first annual conference of the Advanced School for Computing and Imaging, 442–450.

Wessel, P., and D. Bercovici, 1998, Interpolation with splines in tension: a Green's function approach: Mathematical Geology, **30**, 77–93.