UC Davis
IDAV Publications

Title
Interactive Visualization of Function Fields by Range-Space Segmentation

Permalink
https://escholarship.org/uc/item/2p38115s

Journal
Computer Graphics Forum, 28

Authors
Anderson, John C.
Gosink, Luke
Duchaineau, Mark A.
et al.

Publication Date
2009

Peer reviewed

Interactive Visualization of Function Fields by Range-Space Segmentation

John C. Anderson[1], Luke J. Gosink[1], Mark A. Duchaineau[2], and Kenneth I. Joy[1]

[1]Institute for Data Analysis and Visualization, Department of Computer Science, University of California, Davis
e-mail: {janderson, ljgosink, kijoy}@ucdavis.edu
[2]Center for Applied Scientific Computing, Lawrence Livermore National Laboratory
e-mail: duchaine@llnl.gov

Abstract

We present a dimension reduction and feature extraction method for the visualization and analysis of function field data. Function fields are a class of high-dimensional, multi-variate data in which data samples are one-dimensional scalar functions. Our approach focuses upon the creation of high-dimensional range-space segmentations, from which we can generate meaningful visualizations and extract separating surfaces between features. We demonstrate our approach on high-dimensional spectral imagery, and particulate pollution data from air quality simulations.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.6]: Methodology and Techniques

1. Introduction

With increasing computing power and the ability to gather more and more data via increasingly powerful imaging and sensor technology, we can generate data sets of ever increasing complexity. Datasets that represent physical phenomena now contain billions of multi-valued, multi-dimensional, time-varying elements, and are difficult (or impossible) to analyze by the classical scalar- and vector-field algorithms commonly used in the visualization community [LC87, BW01].

In this paper, we address the visualization and analysis of *function fields*, a class of high-dimensional, multi-variate data. Function fields directly arise in applications where an entire spectrum of values is simulated/collected at each data point. From hyperspectral imagery to ground cover distributions, ocean, weather, and air quality simulations, we find data in which samples do not correspond to collections of disjoint scalar values, but rather one-dimensional scalar functions:

$$F : p \in \mathbb{R}^n \rightarrow f_p \in \mathcal{F}_I,$$

where $\mathcal{F}_I$ is the set of functions over a closed interval I. Consider hyperspectral imagery, the structure of which is depicted in Figure 1; here sophisticated sensors produce images in which individual pixels correspond to sampled functions of the spectral intensity of visible and infrared light. Functions are typically represented by a discrete set of m samples over the functional domain. Furthermore, m is often large, leading to tens or hundreds of samples per data point.
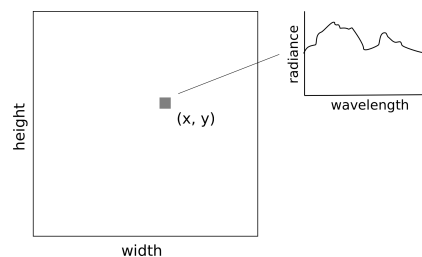
Figure 1: *Hyperspectral images are spatially two-dimensional, with pixels that are sampled functions of radiance (or reflectance) versus wavelength.*

We approach the visualization and analysis of function fields by creating range-space segmentations of the function field data. To begin, we define a similarity metric over the space of functions. Next, a list of *function samples* within the data domain is created – guided by application-specific knowledge, data statistics, or by directly manipulating a spatial probe. These samples are used to compute a range-space segmentation of the data. From such a segmentation, we are able to generate meaningful visualizations, and also extract separating surfaces between features.

We visualize these range-space segmentations by defining a set of transfer functions that operate over each segment. Modifications of the function samples can be used to interactively modify the segmentation of the data, while interactions with transfer functions can be used to interactively generate meaningful visualizations of the data. These interaction techniques provide users with the ability to quickly and directly resolve collisions created by dimension reduction (i.e., when dissimilar high-dimensional values map to similar low-dimensional values). We exhibit a system where feature segmentation does not rely upon fragile high-dimensional queries or clustering, and within which users have great flexibility in exploring complex function fields.

2. Related Work

In addition to domain-specific techniques, dimension reduction, clustering, and query-driven approaches have been used for the visualization and analysis of function field data. Dimension reduction methods project high-dimensional data to fewer dimensions so that traditional visualization techniques can be applied; clustering assigns labels to data based upon some criteria; and queries explicitly segment the data by evaluating constraints upon the original, high-dimensional space.

A common approach for visualizing function fields involves casting them as scalar fields, either directly, by treating the interval *I* over which a field's functions are defined as an extra space or time dimension [ESG97, HAF*96], or through local operations. Kao et al. [KLDP02, KKL*05] and Luo et al. [LKDP03] use parametric statistics and shape descriptors to describe functions using scalar values. For example, a two-dimensional hyperspectral image might be viewed as a three-dimensional image cube, or as a two-dimensional scalar field of averaged radiance.

Principal Component Analysis (PCA) [Jol02] is an ubiquitous dimension reduction technique. For a set of vectors in *m*-dimensional space, PCA identifies a set of ordered, orthonormal basis vectors. Transforming the data vectors into a space spanned by the first $k < m$ of these basis vectors yields a dimension reduction that maximally preserves variance. PCA has been used to display hyperspectral imagery by associating components with color channels to produce color images [TKDO03, JG05].

Multidimensional scaling (MDS) [CC00] may be used to embed high-dimensional data samples in a low-dimensional metric space, such that similar samples are close and dissimilar samples are distant. Once MDS has been performed, the low-dimensional space may be visualized (e.g., by using software such as Voromap [PdOMA06], or as by Fang et al. [FMHC07]) to study the similarity structure of the original data. Spatial datasets, such as function fields, are ill-suited to MDS visualization, however, since the original spatial layout of the data is lost.

Queries have been used to extract and visualize features within function fields. The general idea of query-driven visualization is to isolate and analyze spatial regions that satisfy Boolean range constraints [SSWB05]. Anderson et al. [AGDJ07] have demonstrated that certain function field features can be extracted by constructing queries over function space. Query-driven approaches can be hard to use, however, due to the experimentation required to successfully create a query that extracts the desired feature. Clustering techniques such as *k*-Means and Vector Quantization [AKCM90, Mac67] may also be applied to segment and visualize function field data, however high dimensionality can lead to poor clustering results [JMF99].

The extra dimension inherent to function fields can often be eliminated through domain-specific specialization. In hyperspectral imagery, for example, each pixel may be colored by integrating the radiance versus wavelength functions with color matching functions, such as CIE XYZ, which models the wavelength-dependent response of the human eye [WS00], or the spectrally weighted envelopes of Jacobson and Gupta [JG05]. Furthermore, hyperspectral imagery can be processed using linear spectral unmixing [SD93] to estimate the ratios of material within each pixel. Information theoretic approaches have also been presented to optimize band selection in spectral images [ABS*05, SPS07].

Recent work has focused on using distance or similarity measures to perform dimension reduction of function fields for visualization. Anderson et al. [AGDJ07] derive scalar fields from function fields by computing function-space distance to a "probe" within the data. Fang et al. [FMHC07] present a similar approach for visualizing time-varying data from medical imaging sensors using both function-space and geometric distance measures.

This paper develops a dimension reduction and feature extraction approach based upon range-space segmentations of the original, high-dimensional function field space. We utilize a function-space metric to define a segmentation of the range space, coupled with a visualization approach based upon individual per-segment transfer functions. This approach provides an intuitive dimension reduction for these complex datasets, and allows fast interaction methods to be developed that modify segmentations and visualizations to better analyze the data.

3. Range-Space Segmentation

Consider a function-space distance metric $||\cdot||$ such that $||f-g||$ represents the "similarity" of function f to g. Using such a metric it is possible to project a function field to an scalar field by comparing each of the dataset's functions against a known, exemplar function f. The scalar value at point p with corresponding function f_p is defined to be the distance in function-space between f_p and the exemplar function:

$$S^f : p \in \mathbb{R}^n \to ||f - f_p||. \tag{1}$$

We can extend this approach to produce a *range-space segmentation* of a function field. Consider an ordered set of m function samples $\mathbf{M} = (f_1, f_2, \ldots, f_m)$. From such a set, we can construct *multiple* scalar fields S^{f_i}, one for each function in $\mathbf{M}$. These fields describe the function space distance from the function at p to each of the functions in $\mathbf{M}$. Range-space segmentations are formed by keeping two pieces of information per point p: first, a *classification field* value with the index i of the function f_i in $\mathbf{M}$ that is closest to f_p, and second, a *multi-function scalar distance field* value that stores the distance from f_i to f_p.

Thus, in the multi-function scalar distance field S^*, the scalar value at point p becomes the minimal function-space distance from f_p to any function in $\mathbf{M}$:

$$S^* : p \in \mathbb{R}^n \to \min_{i \in (1,\ldots,m)} ||f_i - f_p||. \tag{2}$$

S^* can be calculated from a set of function samples $\mathbf{M}$ either by computing each scalar distance field S^{f_i} and then their minimum, or by computing the minimum for each point p sequentially.

We also generate an integer-valued classification field L that specifies the *index* of the function in $\mathbf{M}$ used to minimize the value at point p in S^* (rather than the minimum value itself):

$$L : p \in \mathbb{R}^n \to \underset{i \in (1,\ldots,m)}{\arg\min} ||f_i - f_p||. \tag{3}$$

For example, if the nth probe is used to minimize Equation (2) at point p, then $L(p) = n$. Under a Euclidean distance metric, one way to interpret the value at a point p in L is as the label of the cell to which f_p belongs in the function-space Voronoi tessellation [Aur91] created by the function samples.

These two scalar fields, S^* and L, represent the range-space segmentation of the function field formed by the set of function samples $\mathbf{M}$. Before turning to direct visualization and feature segmentation, however, we must discuss two important aspects of the segmentation construction process: what similarity metric to use, and how to choose function samples.

3.1. Similarity Measures

Our approach is very flexible with respect to the distance metric used to create the range-space segmentation. To obtain results, the distance metric simply needs to reflect a measure of similarity between two function-space samples.

An example of a general, function-space metric is the weighted Euclidean metric. Given a function f defined over a closed interval I, the weighted Euclidean metric is defined as:

$$||f|| = \left(\int_I w(x) f(x)^2 dx \right)^{\frac{1}{2}}, \tag{4}$$

where $w(x)$ is a weight function. If f is defined discretely – i.e., represented by a sequence of m points $(f_1, f_2, \ldots, f_m)$, the metric is defined as:

$$||f|| = \left(\sum_{i=1}^{m} w_i f_i^2 \right)^{\frac{1}{2}}, \tag{5}$$

where $(w_1, w_2, \ldots, w_m)$ are a set of weights. We measure the distance between two functions f and g by calculating $||f - g||$. An interesting note is that if the weight function (or vector) is constant, then the segmentation produced by multiple function-space samples under this metric corresponds to a Voronoi tessellation [Aur91] of the range-space.

We have applied the above metric over hyperspectral imagery and particulate pollution data with good results (Section 4). However, in these and other application domains the choice of distance metric will lead to different range-space segmentation results. In addition to weighted Euclidean, other commonly used metrics in the context of sampled functions include Earth Mover's Distance [RTG98], and Chang's spectral distance metrics [Cha00]. Cox and Cox [CC00] also suggest a number of additional metrics.

3.2. Function Samples

In order to create a range-space segmentation, the user must specify a set of function samples $\mathbf{M}$. In some situations, users might have meaningful exemplar functions *a priori* in the form of "test sets." An example is in the domain of hyperspectral imagery, where it is likely that analysts already have a list of reflectance functions corresponding to known materials (i.e., a spectral library). Other domains are also likely to have their own "known" function signatures, and our approach fully supports this type of foreknowledge.

In our software implementation, we provide the user with flexible controls to specify the function samples, including:

- functions from test sets,
- analytic functions,
- hand-drawn functions, and
- functions derived from the data under various distribution statistics.

In addition, we allow users to specify function samples through an interactive spatial probing process. A *probe* is a user-specified point in the data domain $p \in \mathbb{R}^n$, controlled by a *full space cursor* [NDRO87]. The function sample associated with the probe is the function f_p at the point p in the function field.

3.3. Visualization

Range-space segmentations are effective vehicles for producing direct visualizations of function field data. A segmentation is the combination of a distance field and an integer-valued classification field, both of which are scalar fields in $\mathbb{R}^n$. We can apply traditional scalar field rendering techniques, largely unchanged, upon a range-space segmentation to produce images of two-dimensional function fields and volume renderings of three-dimensional fields.

In the case where the segmentation is created by a single function sample – i.e., $m = 1$, the user can directly associate colors with scalar values in the distance field as done by Anderson et al. [AGDJ07]. For volume rendering, where a color's opacity is important, users are able to modify an opacity function as part of the transfer function. We use 1D transfer functions during volume rendering, but two-dimensional [Lev88], multidimensional [KKH05], and local [LLY06] transfer functions may be applied to emphasize local structures.

Most often, however, the range-space segmentation will be derived from multiple function samples. To visualize non-trivial segmentations we turn to the classification field. We associate a transfer function with each of the m function samples to create an ordered list of transfer functions $\mathbf{T} = (t_1, \ldots, t_m)$. During rendering, the classification field value at p determines the transfer function $t_{L(p)}$ used to color the scalar value at p in S^*:

$$\text{color}(p) = t_{L(p)}(S^*(p)).$$

Geometrically, we associate a different transfer function within each "cell" defined by the range-space segmentation. Figure 2 illustrates this rendering approach, in which a point p is shaded using the transfer function associated with the nearest function sample to f_p.

3.4. Feature Segmentation

Range-space segmentations facilitate the construction of segmenting surfaces between feature regions in two and three dimensions. The key to producing segmenting surfaces in our framework is to perform surface extraction over the classification field L. Recall that integer values in the classification field encode the range-space segmentation "cell" membership for each point, as defined by the similarity metric and the current function samples (Equation 3). Thus, surfaces that partition the classification field into homogeneously labeled regions correspond to boundaries between function-space features.
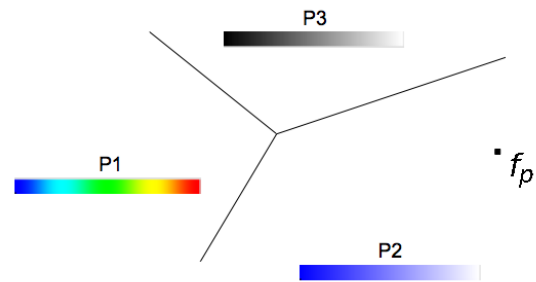
Figure 2: *Our approach creates a Voronoi-like tessellation of the function field range space (i.e., space of functions). Each "cell" of the tessellation is assigned its own colormap for visualization. The color of a point p is determined by the location of its corresponding function f_p within the range-space segmentation.*

For a range-space segmentation constructed from two function samples, the classification field will be a binary labeling of the function field domain. We can extract the segmenting surface(s) between features by performing isosurfacing with an isovalue of 0.5. Algorithms such as Marching Cubes [LC87] can be used to extract a surface representing the set of points I with a constant isovalue v through a the classification field – i.e., $I : \{x | L(x) = v\}$.

In more complex cases, where the classification represents a segmentation derived from three or more function samples, segmenting surfaces can be extracted using one of various multi-label segmentation algorithms. Examples include multi-label Marching Cubes methods [HSSZ97, WJMS03, BL03], Dual Contouring [JLSW02], or the method of Nielson and Franke [NF97] on an implicit tetrahedrization of the rectilinear domain.

4. Results

Function fields arise in many application domains. In this section we discuss the results of our method across multiple datasets: hyperspectral imagery from the domain of remote sensing, and simulated particulate pollution data.

4.1. Hyperspectral Imagery

Hyperspectral imaging systems are used in remote sensing for a broad range of applications, including environmental studies and military preparation. Each pixel in a hyperspectral image contains data for multiple spectral channels (instead of only grayscale or RGB), thus allowing more in-depth image analysis. The Airborne Visible InfraRed Imaging Spectrometer (AVIRIS) [VGC*93] is aircraft-mounted and acquires calibrated 614x512 images of up-welling spectral radiance. In AVIRIS images each pixel consists of 224

PCA

Incremental construction of range-space segmentations.   VQ Clustering
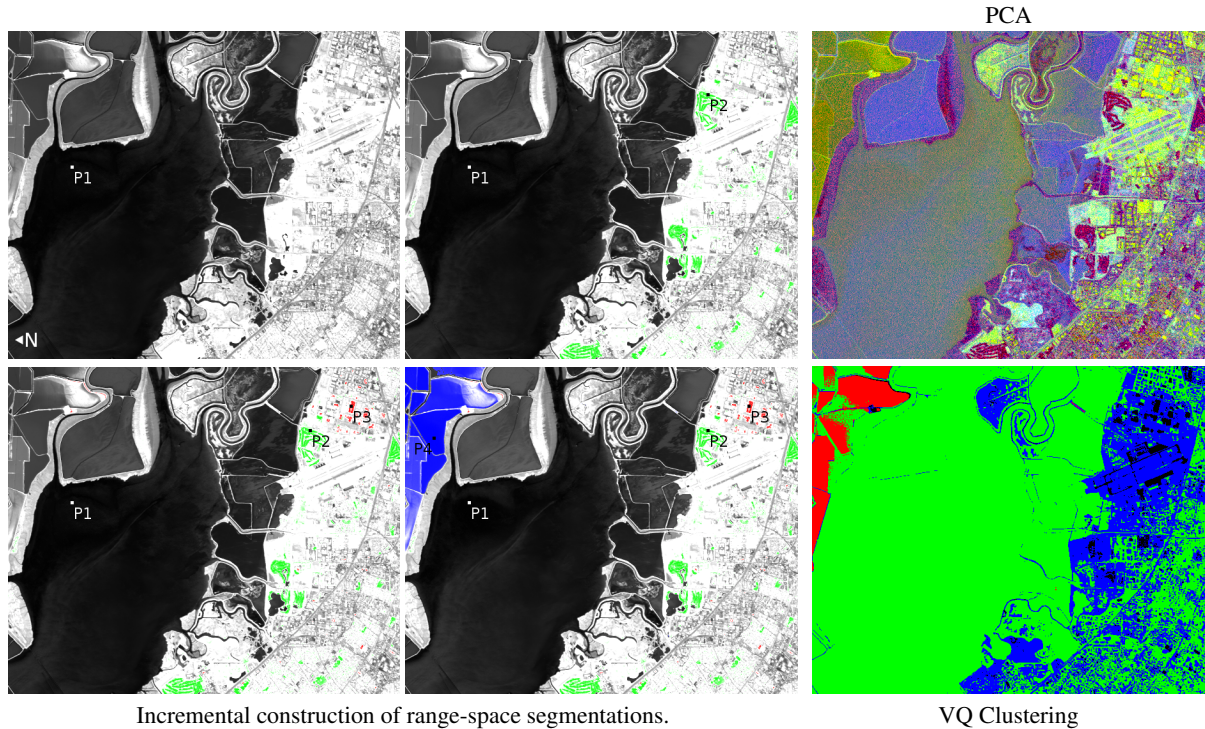
Figure 3: *Visualizations of a hyperspectral image of Moffett Field and the San Francisco Bay. The leftmost set of images shows the construction of a range-space segmentation with 1, 2, 3, and 4 probes. On the right are images generated by mapping PCA components to RGB (top), and by Vector Quantization (VQ) clustering (bottom).*

radiance (or reflectance) samples over visible and short-wave infrared wavelengths, yielding an image size of approximately 270 megabytes.

The leftmost images of Figure 3 show the incremental construction of a range-space segmentation for a hyperspectral image of Moffett Field and the San Francisco Bay. The probes, and their associated transfer functions, were interactively added in the following order: 1) over water with function ⌐⌐⌐⌐⌐⌐ (black-to-white), 2) on a golf course with function ⌐⌐⌐⌐⌐⌐ (green), 3) on a building with function ⌐⌐⌐⌐⌐⌐ (red), and 4) over evaporation ponds containing brine shrimp with function ⌐⌐⌐⌐⌐⌐ (blue). Because segmentation is done in the original spatial and functional domains, users can track particular features of interest while still "managing" unknown features with tentative function samples and transfer functions. It is often the case that the context provided by the initial distance field visualization helps the user identify and segment addition features. Furthermore, spatial coherency in the function field helps our visualizations to remain relatively stable when adding and changing function samples by probing.

In the upper right image of Figure 3 we show the result of applying PCA over the hyperspectral image for visualization [TKDO03, JG05]. Here, individual dimensions are mapped to RGB color channels after the PCA transform; we

have used the mapping $(P_4, P_5, P_6) \rightarrow (R, G, B)$. Unlike our method, PCA requires an expensive preprocess of the data and is a "static" dimension reduction. The only choice in PCA visualization is the set of principle components to consider. This can be difficult: we found (P_4, P_5, P_6) to be the first set of components that produce an image not dominated by noise, and cycling through additional sets of components leads to little additional insight into the data.

We also compare our approach to clustering. The bottom right image of Figure 3 shows the results of applying Vector Quantization (VQ) clustering [AKCM90]. We have clustered the first 10 components from a PCA transform of the hyperspectral image into four clusters; clustering over all 224 dimensions in either the original or PCA transformed data produces an *extremely* noisy clustering. With the correct settings and preprocessing, clustering is able to capture similar features to our method (e.g., the brine shrimp ponds in the image), because both are based upon a segmentation of function range-space.

4.2. Particulate Pollution

Time-varying, three-dimensional particulate pollution datasets used for air quality research often take the form of function fields. In the datasets we consider, each cell contains a sampled function of particle concentration
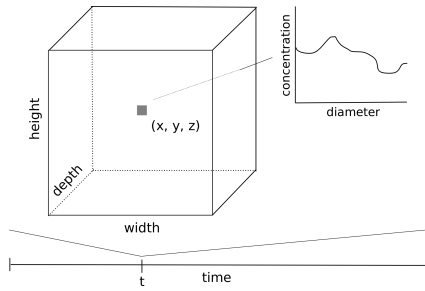
Figure 4: *Data layout of three-dimensional, time-varying function field from particulate pollution simulations. Each cell contains a sampled function of aerosol particulate concentration versus diameter – an important factor in toxicity.*

Table 1: *Timings for range-space segmentation creation.*

Dataset	S	S^*	**Total** (ms)
Hyperspectral Imagery	80	16	336
H_2O Aerosol (National)	20	11	51
SO_4 Aerosol (CRPAQS)	27	23	104

versus diameter as depicted in Figure 4. This functional dependence is crucial, because both concentration and size are crucial factors in the toxicity of aerosol particles. The first dataset (National) is a 148x112x19 grid of particulate H_2O concentration over the continental United States. The second dataset, from the California Regional Particulate Air Quality Study (CRPAQS), is a 185x185x15 grid of particulate SO_4 concentration throughout the San Joaquin Valley, California, U.S.A. Each dataset contains cell-centered, 9-sampled functions of particle concentration versus diameter, over 25 timesteps. The CRPAQS dataset is approximately 450 megabytes, and the National dataset is approximately 260 megabytes (much larger than scalar fields with similar spatial extents).

Important to our method is that by using multiple probes, and simple transfer functions, users are able to create renderings that meaningfully highlight different aspects of the same dataset. Figure 5 shows H_2O concentration from the National particulate dataset rendered over multiple timesteps using two probes. In both (a) and (b), the first probe, with a black-to-white transfer function, is located over central Mexico, and corresponds to low H_2O concentration. In (a), the second probe, with a rainbow transfer function, is located in an area of low to moderate moisture in the United States mid-west. In (b), the second probes is placed in an localized area of functions with high total moisture content.

Our method also provides flexibility in the visualization and segmentation of time-varying function fields. In time-varying fields, probes become points in $\mathbb{R}^n \times \mathbb{T}$. In Figure 6(a), we show a direct visualization of the range-space segmentation created by three function samples in different timesteps for the CRPAQS dataset. The first probe is located outside of the central San Joaquin valley and has low total SO_4 concentration. The second and third probes, however, are located at the same spatial position, but at different points in time. The second probe with a red transfer function is in timestep 0 and corresponds to function of high total SO_4 concentration. The third probe with a blue transfer function

is in timestep 18 and corresponds to moderate SO_4 concentration. Figure 6(b) shows a closeup of the direct visualization produced by the range-space segmentation. In 6(c) we highlight feature segmentation: multi-material surface extraction as described in Section 3.4 is used to extract boundaries between spatial regions with functions having high, medium, and low total SO_4 concentration.

4.3. Performance

The techniques presented herein are best utilized in an interactive setting, where operations such as changing function samples, creating segmentations, and deriving new visualizations are rapidly realized. We have performed testing on an Apple MacBook Pro notebook computer (2.33 GHz Intel Core 2 Duo processor, 2 GB memory, and an ATI Radeon X1600 graphics card). For the datasets considered our method is interactive.

Table 1 shows timings in milliseconds for the generation of the segmentations used in Figures 3, 5, and 6(b) (four, two, and three probes, respectively). The column S lists the time required to generate one single-function distance field (Equation 1). The S^* column lists the time required to combine all distance fields into a range-space segmentation (Equations 2 and 3). The totals listed reflect the time required to *fully* generate a new segmentation: i.e., to generate each single-probe field and combine them into a multi-probe field. Often, however, end users will experience far less latency. When the user modifies a function sample (for example, by repositioning a probe), they only modify one of the m distance fields, which the remaining $m - 1$ fields remain unchanged. Thus, the latency experience by users when changing a single function sample will be $S + S^*$ from Table 1.

5. Conclusion

In this paper, we have presented a range-space segmentation framework for the visualization and analysis of function field data: one of the myriad of possible data types that can populate the variables in a multi-dimensional, multi-variate dataset. The presented methods increase our capacity to visualize these complex fields, and help us gain new insight about the data. Future work will focus upon generalizing and extending our approach to other types of multi-variate data.

Acknowledgements

References

[ABS*05] AIAZZI B., BARONTI S., SANTURRI L., SELVA M., ALPARONE L.: Information-theoretic assessment of multi-dimensional signals. *Signal Processing 85*, 5 (2005), 903–916.

[AGDJ07] ANDERSON J. C., GOSINK L., DUCHAINEAU M. A., JOY K. I.: Feature identification and extraction in function fields. In *Proc. of EuroVis* (May 2007).

[AKCM90] AHALT S. C., KRISHNAMURTHY A. K., CHEN P., MELTON D. E.: Competitive learning algorithms for vector quantization. *Neural Networks 3*, 3 (1990), 277–290.

[Aur91] AURENHAMMER F.: Voronoi diagrams – survey of a fundamental geometric data structure. *ACM Computing Surveys 23*, 3 (1991), 345–405.

[BL03] BANKS D. C., LINTON S.: Counting cases in Marching Cubes: Toward a generic algorithm for producing substitopes. In *Proc. of IEEE Visualization* (Oct. 2003), pp. 51–58.

[BW01] BRODLIE K., WOOD J.: Recent advances in volume visualization. *Computer Graphics Forum 20*, 2 (2001), 125–148.

[CC00] COX T. F., COX M. A. A.: *Multidimensional Scaling*, second ed. Chapman & Hall/CRC, Sept. 2000.

[Cha00] CHANG C.-I.: An information-theoretic approach to spectral variability, similarity, and discrimination for hyperspectral image analysis. *IEEE Trans. on Information Theory 46*, 5 (2000), 1927–1932.

[ESG97] EHLSCHLAEGER C. R., SHORTRIDGE A. M., GOODCHILD M. F.: Visualizing spatial data uncertainty using animation. *Computational Geosciences 23*, 4 (1997), 387–395.

[FMHC07] FANG Z., MÖLLER T., HAMARNEH G., CELLER A.: Visualization and exploration of time-varying medical image data sets. In *Proc. of Graphics Interface* (2007), pp. 281–288.

[HAF*96] HIBBARD W. L., ANDERSON J., FOSTER I., PAUL B. E., JACOB R., SCHAFER C., TYREE M. K.: Exploring coupled atmosphere-ocean models using Vis5D. *International J. of Supercomputer Applications and High Performance Computing 10*, 2/3 (Summer/Fall 1996), 211–222.

[HSSZ97] HEGE H.-C., SEEBASS M., STALLING D., ZÖCKLER M.: A generalized Marching Cubes algorithm based on nonbinary classifications. ZIB Preprint SC-97-05, 1997.

[JG05] JACOBSON N., GUPTA M.: Design goals and solutions for display of hyperspectral images. In *IEEE Image Processing* (Sept. 2005), vol. 2, pp. 622–625.

[JLSW02] JU T., LOSASSO F., SCHAEFER S., WARREN J.: Dual contouring of Hermite data. *ACM Trans. on Graphics 21*, 3 (2002), 339–346.

[JMF99] JAIN A. K., MURTY M. N., FLYNN P. J.: Data clustering: a review. *ACM Comput. Surv. 31*, 3 (1999), 264–323.

[Jol02] JOLLIFFE I. T.: *Principal Component Analysis*, second ed. Springer, Oct. 2002.

[KKH05] KNISS J., KINDLMANN G., HANSEN C.: *Multidimensional Transfer Functions for Volume Rendering*. Elsevier, 2005, ch. 9, pp. 189–210.

[KKL*05] KAO D., KRAMER M., LOVE A., DUNGAN J., PANG A.: Visualizing distributions from multi-return lidar data to understand forest structure. *Cartographic Journal, Special Issue on GeoVisualization 42*, 1 (June 2005), 1–14.

[KLDP02] KAO D., LUO A., DUNGAN J. L., PANG A.: Visualizing spatially varying distribution data. *Proc. of Information Visualization* (2002), 219–226.

[LC87] LORENSEN W. E., CLINE H. E.: Marching Cubes: A high resolution 3D surface construction algorithm. In *Proc. of SIGGRAPH* (1987), pp. 163–169.

[Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications 08*, 3 (1988), 29–37.

[LKDP03] LUO A., KAO D., DUNGAN J., PANG A.: Visualizing spatial distribution data sets. In *Proc. of the Symposium on Data Visualisation* (2003), pp. 29–38.

[LLY06] LUNDSTROM C., LJUNG P., YNNERMAN A.: Local histograms for design of transfer functions in direct volume rendering. *IEEE Trans. on Visualization and Computer Graphics 12*, 6 (2006), 1570–1579.

[Mac67] MACQUEEN J.: Some methods for classification and analysis of multivariate observations. pp. 281–297.

[NDRO87] NIELSON G. M., DAN R. OLSEN J.: Direct manipulation techniques for 3D objects using 2D locator devices. In *Proc. of the Workshop on Interactive 3D Graphics* (1987), pp. 175–182.

[NF97] NIELSON G. M., FRANKE R.: Computing the separating surface for segmented data. In *Proc. of IEEE Visualization* (Oct. 1997), pp. 229–233.

[PdOMA06] PINHO R., DE OLIVEIRA M. C. F., MINGHIM R., ANDRADE M. G.: Voromap: A Voronoi-based tool for visual exploration of multi-dimensional data. In *Proc. of Information Visualization* (2006), pp. 39–44.

[RTG98] RUBNER Y., TOMASI C., GUIBAS L. J.: A metric for distributions with applications to image databases. In *Proc. of the International Conference on Computer Vision* (1998), pp. 59–66.

[SD93] SETTLE J., DRAKE N.: Linear mixing and the estimation of ground cover proportions. *International J. of Remote Sensing 14*, 6 (1993), 1159–1177.

[SPS07] SOTOCA J. M., PLA F., SNCHEZ J. S.: Band selection in multispectral images by minimization of dependent information. *IEEE Trans. Systems, Man and Cybernetics 37*, 2 (Mar. 2007), 258–267.

[SSWB05] STOCKINGER K., SHALF J., WU K., BETHEL E. W.: Query-driven visualization of large data sets. In *Proc. of IEEE Visualization* (Oct. 2005), pp. 167–174.

[TKDO03] TYO J. S., KONSOLAKIS A., DIERSEN D. I., OLSEN R. C.: Principal-components-based display strategy for spectral imagery. *IEEE Trans. on Geoscience and Remote Sensing 41*, 3 (Mar. 2003), 708–718.

[VGC*93] VANE G., GREEN R., CHRIEN T., ENMARK H., HANSEN E., PORTER W.: The airborne visible infrared imaging spectrometer. In *Remote Sensing Environment* (1993), vol. 44, pp. 127–143.

[WJMS03] WU Z., JOHN M. SULLIVAN J.: Mutliple material Marching Cubes algorithm. *International J. for Numerical Methods in Engineering 58*, 2 (July 2003), 189–207.

[WS00] WYSZECKI G., STILES W. S.: *Color Science: Concepts and Methods, Quantitative Data and Formulae*, second ed. Wiley, 2000.
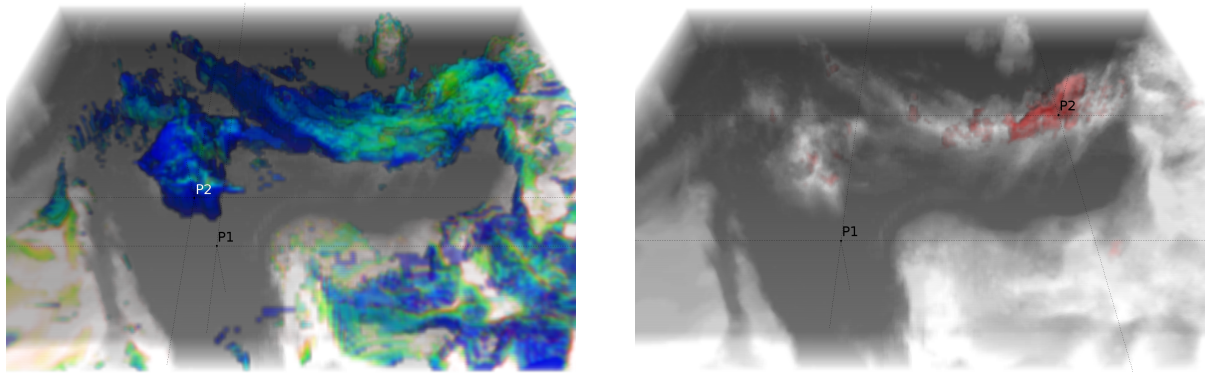
(a) Broad region of moderate moisture functions

(b) Localized region of high moisture functions

Figure 5: *Volume renderings produced from range-space segmentations of the National H_2O particulate concentration dataset. By using multiple probes, and simple transfer functions, users are able to create renderings that meaningfully highlight different aspects of the same dataset.*
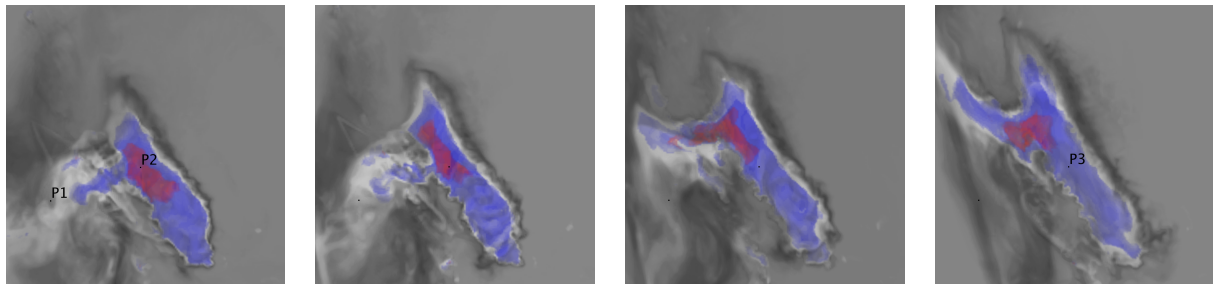
(a) Time Probes
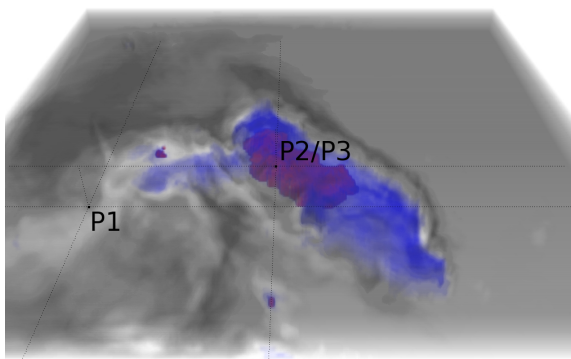
(b) Volume Rendering
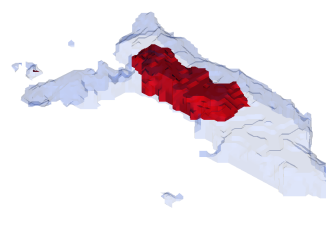
(c) Feature Segmentation

Figure 6: *Range-space segmentation of the CRPAQS dataset using multiple probes in different timesteps: the first two probes are in timestep 0, while the third probe is in the last timestep. In (a) we use these probes to visualize the movement of high (red) and moderate (blue) SO_4 concentration features over time through the San Joaquin Valley. In (b), we show a closeup of the first timestep, while (c) shows the segmentation of high, moderate, and low concentration regions.*