

# Volume Rendering Multivariate Data to Visualize Meteorological Simulations: A Case Study

Joe Kniss and Charles Hansen

Scientific Computing and Imaging Institute, School of Computing,  
University of Utah, Salt Lake City, UT, USA

Michel Grenier and Tom Robinson

Canada Meteorological Centre, Meteorological Service of Canada  
Montreal, Canada

---

## Abstract

*High resolution computational weather models are becoming increasingly complex. However, the analysis of these models has not benefited from recent advancements in volume visualization. This case study applies the ideas and techniques from multi-dimensional transfer function based volume rendering to the multivariate weather simulations. The specific goal of identifying frontal zones is addressed. By combining temperature and humidity as a multivariate field, the frontal zones are more readily identified thereby assisting the meteorologists in their analysis tasks.*

---

## 1. Introduction

High resolution computational weather models are widely used throughout the world<sup>13, 18, 19</sup>. Recent advancements in volume visualization can benefit the analysis of these increasingly complex and higher fidelity datasets. Meteorologists employ similar methods to analyze weather models as they would to track sensed weather data. These methods are effective at providing forecasts but lacks the detailed information to accurately model such features as baroclinic zones, more commonly known as fronts.

Baroclinic zones, or fronts, form the boundary between airmasses with different temperature and humidity characteristics<sup>14</sup>. It is precisely these differences in thermal characteristics which provide the energy by which, under the right circumstances, a wave-like perturbation along the front may amplify to form a mid-latitude storm. Dynamics along cold fronts are also responsible for much of the summer severe weather experienced in mid-latitudes. Hence, fronts are very important for meteorologists.

There exists a “classical” frontal theory, in which fronts consist of a zone of strong thermal gradient, both at the surface and extending in the vertical, potentially to the

tropopause<sup>†</sup>. The fronts slope in the vertical toward the colder airmass, with cold fronts having a sharper slope than warm fronts. Classic warm frontal weather consists of continuous rain, with possibly some embedded convection, while cold fronts tend to produce convection, including showers and thunderstorms, the latter potentially resulting in hail, damaging winds, tornadoes, *etc.*

In the real world, things are often not “classical”. The vertical extent of fronts are difficult to analyse explicitly due to the poor resolution of data in the free atmosphere, so meteorologists tend to concentrate on the surface manifestation. Surface features however can be masked by various phenomena including terrain, radiative effects, low-level moisture sources and sinks, cloud cover, *etc.*

### 1.1. Current Techniques

The standard joke is that if you put twelve meteorologists in a room, they will come up with a dozen frontal analyses for

---

<sup>†</sup> The tropopause is a term referring to the top of the troposphere, which is the layer of the atmosphere extending outward 7 to 10 miles from the earth’s surface.

the same situation, illustrating that there are many different ways to attack the problem. The critical surface fields are temperature and humidity<sup>14</sup>. Also important are the surface pressure, pressure tendency (the three-hourly change in pressure), and wind fields, since fronts tend to be found in pressure troughs<sup>5</sup>. There are characteristic signatures of fronts to be found in satellite imagery as well.

In terms of the upper structure of fronts, operational meteorologists at the Canadian Meteorological Centre use objectively analysed or forecast 1000-500 hPa<sup>‡</sup> thicknesses, 850 hPa thetaw<sup>§</sup> and 250 hPa jet stream chart in addition to other data sources such as radiosonde, satellite, and aircraft data. The thicknesses are proportional to the average temperature in that layer of the atmosphere, giving a good indication of where there are strong thermal gradients. The 850 hPa thetaw also gives a good indication of the thermal gradients at a level that is low enough in the atmosphere where most fronts will appear, while high enough to be separate from most terrain effects. The jet stream is an artifact of thermal gradients, through the thermal wind equation, which states that the vector difference in the wind between two levels in the atmosphere is proportional to the average horizontal temperature gradient in that layer<sup>12</sup>.

Since the summer of 2000, meteorologists at the Canadian Meteorological Centre (CMC) have been using a graphical editor (called Edigraf) to do their frontal analyses on-screen. This application allows the meteorologist to overlay any number of model generated meteorological fields, data and satellite imagery, enabling a comprehensive view of the atmosphere. This application takes advantage of color mapped scalar data, glyphs, and contours.

There has been a wealth of other visualization systems used in weather forecasting and analysis. The most widely used is VisSD<sup>3,4</sup>. This package uses scalar and vector visualization methods such as isosurfacing, vector glyphs, and volume rendering. However, all scalar render modalities are univariate. Other packages add various analysis capabilities but still are limited to univariate fields<sup>16,17,19</sup>.

## 2. Background

Volume rendering is an important technique for scalar field visualization. Most volume rendering systems set color and opacity based on a single scalar value. Levoy's seminal work<sup>10,11</sup> in volume rendering utilized 2D transfer functions of data value and gradient magnitude. The use of multi-dimensional transfer functions has all but disappeared from the literature and commercial volume rendering tools, with a few exceptions<sup>6,7</sup>. One reason for this is that the process

of selecting a good 1D transfer function is often trial and error. The complexity of this process is compounded by the additional degrees of freedom introduced by adding dimensions to the transfer function. Another reason stems from the popularity of graphics hardware volume rendering techniques. Until recently, multi-dimensional table lookups were not possible, even on very high-end graphics hardware.

Simulation and medical applications which require visualization often produce multiple values per sample point, or voxel. Clearly, multi-dimensional transfer functions could be beneficial in the analysis of such data. Analysis of multivariate fields in the context of color images has led to some interesting methods for both segmentation and rendering.

2D color image segmentation is an example of multivariate classification. A detailed description of this process can be found in<sup>15</sup> and the references contained therein. Sapiro's approach uses multivariate derivative measures to guide implicit surface evolution.

Ebert *et al.*<sup>2</sup> investigate the mapping of multi-valued color data to a scalar density or opacity. They have developed different techniques for managing the multi-value mapping, while directly rendering volumes from photographic data. They use a two step method which maps RGB into the CIE  $L = U = V$  color space. This allows them to explore photographic images without committing to an *a priori* segmentation. While this approach allows users to classify the data based on the behavior of a gradient magnitude measure, it lacks a mechanism to specify an arbitrary transfer function based on the color data directly.

Current volume rendering approaches for visualizing multi-valued datasets rely on separate transfer functions for each modality, and then combine them<sup>20</sup>. The composite visualization can simply overlay the classified modalities or use portions of the different transfer functions, for instance color from one modality and opacity from another. Others combine visualization techniques such as isosurfaces and glyphs to convey the relationship of values from each of the modalities<sup>3,4,18</sup>.

Laidlaw<sup>9</sup> describes a framework for contrast enhancement and pulse sequence optimization for spin-echo MRI data acquisition using multi-valued fields. This approach also demonstrates the utility of using multiple pulse sequences to better classify materials for visualization using 2D transfer functions.

Recent advancements in graphics hardware have enabled the use of multi-dimensional transfer functions within an interactive framework<sup>7</sup>. A key feature of this approach is the ability to set transfer functions that emphasize regions of interest using a straight forward point and click interface, a process known as dual-domain interaction. This work explores multi-dimensional transfer functions restricted to the scalar value, the magnitude of the gradient of the field, and

<sup>‡</sup> Pa is the official metric unit of atmospheric pressure. hPa means hecto-pascal, it is equivalent to mb or millibar.

<sup>§</sup> Thetaw or theta-w is the wet-bulb potential temperature.

the signed second derivative magnitude in the gradient direction of the field.

### 3. Multi-Dimensional Transfer Functions for Multi-Fields

The goal of this case study is to explore the utility of multi-dimensional transfer functions for the visualization of multi-variate fields. It is often the case that numerical simulations and medical image acquisition techniques produce a number of data values per sample point. Together, these data values may describe a phenomena or feature of interest better than any one of them in isolation.

Our system allows the user to specify a fully general multi-dimensional transfer function with up to three axes. For instance, the first two axes can map data values and the third can map the gradient magnitude of the multi-value field. For multi-dimensional transfer functions with greater than three axes, the additional axes are treated as separable portions of the transfer function. The use of multi-dimensional transfer functions is advantageous for volume visualization because a feature of interest may not be localized in the data space of any single variable. Creating a higher dimensional data space by adding independent or derived variables to the transfer function increases the likelihood that a feature can be classified. We have found this to be true in our attempts to visualize airmasses in atmospheric simulations. In this case the feature of interest is the frontal zones of airmasses. This feature is not adequately described by either temperature or humidity alone, nor does there exist a mathematical formulation capable of classifying it *a priori*.

Given the difficulty of specifying a multi-dimensional transfer function, we find dual-domain interaction to be an important interaction modality for dataset exploration. Dual-domain interaction is the process of setting a transfer function based on values queried, or pointed at, in the spatial domain of the volume rendering<sup>7</sup>. This is especially important for multi-dimensional transfer functions, given the complex relationship between a feature of interest and the ranges of data values that describe it.

It has been shown that gradient magnitude is a useful measure for data classification in volume rendering applications<sup>6, 7, 10, 11</sup>. Gradient measurement is a well defined operation for scalar datasets. Gradient estimation for multi-valued fields, however, is not well defined. For this reason we choose a gradient measure with demonstrated effectiveness for color image edge detection. Our formulation is taken from<sup>1, 15</sup>, but is repeated here for convenience with notation for 3D fields.

Let  $\Phi(u_1, u_2, u_3) : \mathfrak{R}^3 \rightarrow \mathfrak{R}^m$ . This defines a multi-valued 3D field with  $m$  elements at each sample, *ie.* each sample is a vector in  $\mathfrak{R}^m$ . The difference between two points,  $P = (u_1^0, u_2^0, u_3^0)$  and  $Q = (u_1^1, u_2^1, u_3^1)$ , is a vector in  $\mathfrak{R}^m$  given by:

$$\Delta\Phi = \Phi(P) - \Phi(Q)$$

When the distance between  $P$  and  $Q$  tends to zero the difference becomes the vector in  $\mathfrak{R}^m$

$$d\Phi = \sum_{i=1}^3 \frac{\delta\Phi}{\delta u_i} du_i$$

and its squared norm, a scalar, is given by the quadratic form

$$d\Phi^2 = \sum_{i=1}^3 \sum_{j=1}^3 \frac{\delta\Phi}{\delta u_i} \cdot \frac{\delta\Phi}{\delta u_j} du_i du_j$$

Making a substitution for the dot product

$$g_{ij} = \frac{\delta\Phi}{\delta u_i} \cdot \frac{\delta\Phi}{\delta u_j}$$

we have

$$d\Phi^2 = \sum_{i=1}^3 \sum_{j=1}^3 g_{ij} du_i du_j = \begin{bmatrix} du_1 \\ du_2 \\ du_3 \end{bmatrix}^T \begin{bmatrix} g_{11} & g_{12} & g_{13} \\ g_{21} & g_{22} & g_{23} \\ g_{31} & g_{32} & g_{33} \end{bmatrix} \begin{bmatrix} du_1 \\ du_2 \\ du_3 \end{bmatrix}$$

The use of this measure for color image edge detection is simplified by the fact that each of the color channels have the same dynamic range. This is most often not the case for simulation data. We handle this by multiplying each data channel's derivative by a weight constant. Our heuristic for determining the weight for each channel is to take the reciprocal of that channel's maximum data value minus its minimum data value. This prevents a channel with a very high dynamic range from dominating the derivative calculation. This heuristic assumes that each channel should have equal influence in derivative estimation relative to the other channels and there are no outliers. This may not always be desirable. Domain specific knowledge may be required to adjust these weights so that the feature or phenomena will be accurately represented in the visualization. We can express this component-wise multiply as a  $m \times m$  scale matrix  $W$  with diagonal elements equal to the  $m$  per-channel weight constants. We now have

$$d\Phi = W \left( \sum_{i=1}^3 \frac{\delta\Phi}{\delta u_i} du_i \right), \quad W := \begin{bmatrix} w_1 & 0 & 0 & \cdots \\ 0 & w_2 & 0 & \cdots \\ & & \ddots & \\ 0 & 0 & \cdots & w_m \end{bmatrix}$$

In the above equations,  $[g_{ij}]$  is a metric tensor which describes the rate of change in all directions. One way to access the total change at a location is to take the L2 norm of this matrix. Clearly there are other ways to extract meaningful information from this tensor, such as the magnitude of the largest eigenvalue or the sum of the eigenvalues (the trace). In any case, this value can then be used as an additional axis in the transfer function. Given the limitations of modern graphics hardware, we apply the gradient portion of the transfer function separably for fields with three or more values per sample.

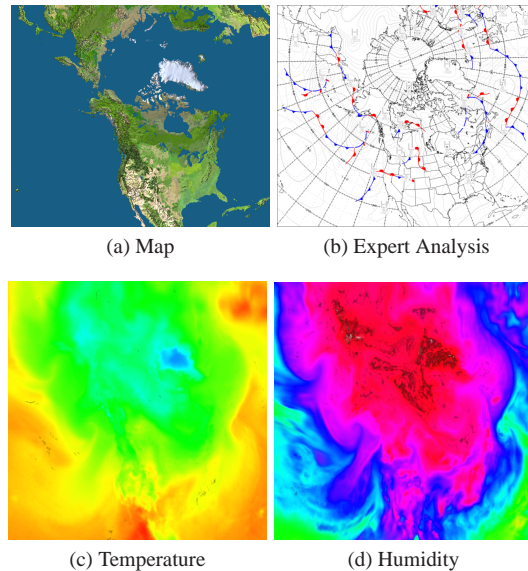
Lighting provides the human visual system with important cues about the shape and curvature of surfaces. Lighting for scalar volumes is often done by pre-computing the normalized gradient for each sample in the volume. This vector is then used as the surface normal for a graphics lighting model. This is appropriate for volume visualizations of features near regions with high gradient magnitude. The approach for generating multi-field gradients discussed above provides the orientation but not the absolute direction. The choice of gradient is often the eigenvector corresponding to the largest eigenvalue of the metric tensor. Thus, lighting with these gradients does not provide robust results since the gradient can flip direction in local neighborhoods, and the choice of eigenvector may not be clear when two or more eigenvalues have the same or similar values. Furthermore, this orientation may not even correspond to the surface of the classified region. Our approach uses an atmospheric lighting model which does not rely on a surface normal<sup>8</sup>, see Figures 3 and 4. This lighting model simply attenuates light through the volume and modulates the voxel's initial color by the summed attenuation, providing simple but robust lighting. It is the attenuated color which is composited to form the volume rendering.

#### 4. Results

A comparison with satellite imagery suggests that additional information can be gained from the volume rendered multi-fields.


Figure 1(a) shows a map of the globe underlying the simulation data. The data was generated using a forcing function from atmospheric measurements. Figure 1(b) shows the results of analysis by an expert with access to all relevant variables of the simulation including temperature (shown in Figure 1(c)), humidity (shown in Figure 1(d)), and pressure, as well as derived data such as dew point and wet-bulb temperature. We have investigated the use of multi-dimensional transfer functions with various combinations of these values and found temperature and humidity to be the most appropriate for this task.

Figures 2(a) and (b) show the results of a composite volume rendering that combines the attributes from two separate transfer functions, one for temperature and one for humidity. In this case color is taken from one transfer function



**Figure 1:** (a) shows a map of the dataset extent. (b) shows the expert analysis using Edigraph. (c) and (d) show slices of temperature and humidity, respectively, passed through a spectral color map.


and opacity is taken from the other. While a user could learn to read this type of visualization, we feel that this approach would not adequately capture the complex relationship between the data channels being visualized. However, if one desires this effect, it can easily be duplicated using a multi-dimensional transfer function.



**Figure 2:** This is an example of volume rendering using properties from each of the data channels. In (a), color varies only with temperature and opacity varies only with humidity. (b) shows the reverse of this, color with humidity, opacity with temperature.

Figures 3 and 4 show the results of a 3D transfer function specified using dual-domain interaction. These illustrations represent different timesteps in the simulation. The horizontal axis of the transfer function maps temperature, the vertical axis maps humidity, and a third axis, which is not explicitly shown, maps the multi-gradient magnitude described

in Section 3. The opacity assigned to low gradient magnitudes can be restricted by manipulating the sliders located at the top of the transfer function widget. These transfer functions could be set by hand in the transfer function domain. However, the additional feedback from the transfer function being updated as the queried location changes allows a user to identify features of interest more intuitively. A more detailed explanation of dual-domain interaction can be found in 7. This process is important given the fact that small changes to the transfer function may result in large unintuitive changes to the classified regions in the volume rendering. The interface between warm and cold airmasses is made apparent using dual-domain interaction; when the queried position passes through the interface between airmasses, the classified regions make a dramatic shift from one side of the boundary to the other. This effect can be seen in the second half of the accompanying animation. Adding the ability to animate the data over several 6-hourly analysis periods allows the meteorologist to see the evolution of frontal zones and airmasses; an example of this can be seen in the first part of the accompanying animation.



**Figure 3:** This figure demonstrates the full expressivity of a multi-dimensional transfer function. The transfer function was created using dual-domain interaction. The sliders on the top of the transfer function widget allow us to restrict the opacity applied to samples with low gradient magnitudes. Blue regions indicate cold airmasses, red regions indicate warm airmasses.

By combining temperature and humidity, frontal zones are very clearly delineated (see Figures 3 and 4), and additional structure in mid-latitude systems is evident as well. Using combinations of atmospheric data values in a higher di-



**Figure 4:** This image shows a transfer function similar to the one in Figure 3 applied to a different timestep.

mensional transfer function appears to be advantageous over other univariate methods.

## 5. Conclusions and Future Work

This case study applied the ideas and techniques from multi-dimensional transfer function based volume rendering to multivariate weather simulations. The specific goal of identifying frontal zones was addressed and proved to be useful by meteorologists. The combination of temperature and humidity as a multivariate field aided the identification of the frontal zones.

While this work has proven useful, additional investigation is warranted in several areas. The non-orthogonal three-dimensional volumetric representations of the data, as seen at the end of accompanying animation, have not proven as useful to the meteorologists as the plane level data shown in the figures and the beginning of the animation. This is an area for future research. Utilizing the 3D non-orthogonal representations may be more applicable to severe weather phenomena such as hurricanes. Another useful feature would be a heuristic based default transfer function well suited for this type of data. This transfer function would make visible regions which are likely to be of interest. Our current heuristic assumes that regions of change tend to be regions of interest. In this case we give higher opacity to higher multi-gradient magnitudes. Better heuristics, however, exist for identifying weather fronts.

In addition to frontal analysis, the visualization techniques

presented in this case study may be useful for the evaluation of numerical model objective analyses by applying the techniques to the trial fields used to create these analyses. Trial fields are usually forecast fields (e.g. 6-hour forecasts) valid at analysis time, which is adjusted according to the data. In this way an objective analysis is created, and is used to initialize the numerical model. The techniques may also be useful for creating subjective forecast products based on objective guidance.

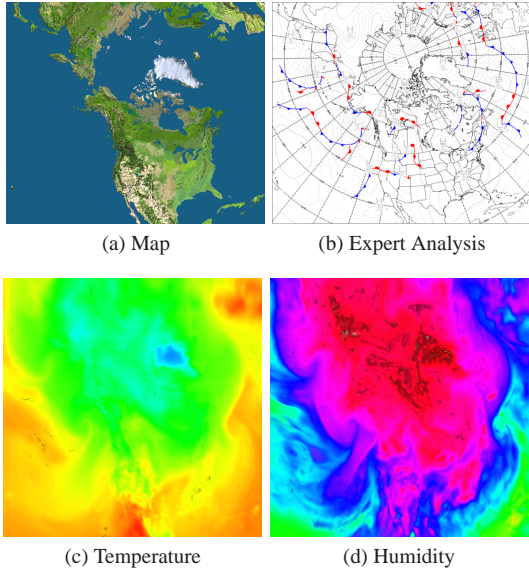
While our system has the potential to handle transfer functions with arbitrary dimensions, the issue of keeping the user interface convenient and intuitive becomes a problem for transfer functions with greater than three dimensions. We feel that there is still a great deal of work to be done generalizing multi-dimensional transfer functions for the visualization of numerical simulations and medical imaging. We also believe that the visualization techniques presented in this case study can be directly applied to other types of simulation data as well as multi-modal medical imaging. We intend to continue our investigation of this approach in these application areas.

## 6. Acknowledgments

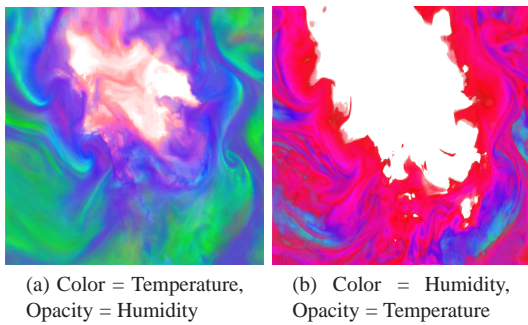
We would like to extend a special thanks to Gordon Kindlmann for helping us with the math for multi-gradient estimation.

## References


1. Silvano Di Zenzo. A Note on the Gradient of a Multi-Image. *Computer Vision, Graphics, and Image Processing*, 33(1):116–125, Jan 1986. 3
2. David Ebert, Christopher Morris, Penny Rheingans, and Terry Yoo. Designing Effective Transfer Functions for Volume Rendering from Photographic Volumes. *IEEE TVCG*, page to appear, 2002. 2
3. W. Hibbard and D. Sante. 4-d display of meteorological data. In *Proceedings of 1986 Workshop on Interactive 3D Graphics*, pages 23–36, 1986. 2
4. W. Hibbard and D. Sante. The vis-5d system for easy interactive visualization. In *Proceedings of IEEE Visualization '90*, pages 129–134, October 1990. 2
5. James R. Holton. *An Introduction to Meteorology*. Academic Press, 2nd edition, 1979. 2
6. Gordon Kindlmann and James Durkin. Semi-automatic generation of transfer functions for direct volume rendering. In *Proceedings of IEEE Symposium on Volume Visualization*, pages 79–86, October 1998. 2, 3
7. Joe Kniss, Gordon Kindlmann, and Charles Hansen. Interactive volume rendering using multi-dimensional transfer functions and direct manipulation widgets. In *Proceedings of Visualization 2001*, pages 255–262, October 2001. 2, 3, 5
8. Joe Kniss, Gordon Kindlmann, and Charles Hansen. Multi-Dimensional Transfer Functions for Interactive Volume Rendering. *TVCG*, 2002 to appear. 4
9. David H. Laidlaw. *Geometric Model Extraction from Magnetic Resonance Volume Data*. PhD thesis, California Institute of Technology, May 1995. 2
10. Marc Levoy. Efficient ray tracing of volume data. *ACM Transactions on Graphics*, 9(3):245–261, July 1990. 2, 3
11. Mark Levoy. Display of surfaces from volume data. *IEEE Computer Graphics & Applications*, 8(3):29–37, 1988. 2, 3
12. D. W. McCann and J. P. Whistler. Problems and Solutions for Drawing Fronts Objectively. *Meteorological Applications*, 8(2):195, 2001. 2
13. F. Molteni, R. Buizza, T. Palmer, and T. Petroliagis. The ECMWF Ensemble Prediction System: methodology and validation. *Quart. J. Roy. Meteor. Soc. (122)*, pages 73–119, 1996. 1
14. F. Sanders and C. A. Doswell III. A Case for Detailed Surface Analysis. *Bulletin of the American Meteorological Society*, 76:505–521, 1995. 1, 2
15. Guillermo Sapiro. Color Snakes. *CVIU* 68(2), pages 247–253, 1997. 2, 3
16. Lloyd A. Treinish. Task specific visualization design: A case study in operational weather forecasting. In *Proceedings of IEEE Visualization '98*, pages 405–409, October 1998. 2
17. Lloyd A. Treinish. Multi-resolution visualization techniques for nested weather models. In *Proceedings of IEEE Visualization 2000*, pages 513–516, October 2000. 2
18. Lloyd A. Treinish. Visual data fusion for applications of high-resolution numerical weather. In *Proceedings of IEEE Visualization 2000*, pages 477–480, October 2000. 1, 2
19. Yingcai Xiao, John Ziebarth, Chuck Woodbury, Eric Bayer, Bruce Rundell, and Jeroen van der Zijp. The challenges of visualizing and modeling environmental data. In *Proceedings of IEEE Visualization '96*, pages 413–416, October 1996. 1, 2
20. Karel J. Zuiderveld and Max A. Viergever. Multi-modal volume visualization using object-oriented methods. In *1994 Volume Visualization Symposium*, pages 59–66. IEEE, October 1994. ISBN 0-89791-741-3. 2



**Figure 1:** (a) shows a map of the dataset extent. (b) shows the expert analysis using Edigraph. (c) and (d) show slices of temperature and humidity, respectively, passed through a spectral color map.



**Figure 2:** This is an example of volume rendering using properties from each of the data channels. In (a), color varies only with temperature and opacity varies only with humidity. (b) shows the reverse of this, color with humidity, opacity with temperature.



**Figure 3:** This figure demonstrates the full expressivity of a multi-dimensional transfer function. The transfer function was created using dual-domain interaction. The sliders on the top of the transfer function widget allow us to restrict the opacity applied to samples with low gradient magnitudes. Blue regions indicate cold airmasses, red regions indicate warm airmasses.



**Figure 4:** This image shows a transfer function similar to the one in Figure 3 applied to a different timestep.