

Visualizing Time-Specific Hurricane Predictions, with Uncertainty, from Storm Path Ensembles

L. Liu¹, M. Mirzangar², R.M. Kirby², R. Whitaker², and D. H. House¹,

¹Clemson University, ²University of Utah

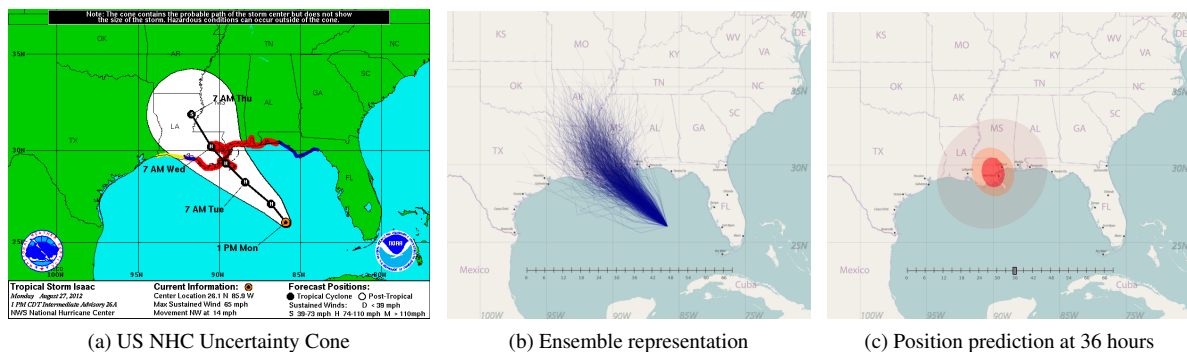


Figure 1: Construction of a time-specific hurricane prediction, *NHC Advisory: Hurricane Isaac, 1 PM CDT, Aug. 27, 2012.*

Abstract

The U.S. National Hurricane Center (NHC) issues advisories every six hours during the life of a hurricane. These advisories describe the current state of the storm, and its predicted path, size, and wind speed over the next five days. However, from these data alone, the question “What is the likelihood that the storm will hit Houston with hurricane strength winds between 12:00 and 14:00 on Saturday?” cannot be directly answered. To address this issue, the NHC has recently begun making an ensemble of potential storm paths available as part of each storm advisory. Since each path is parameterized by time, predicted values such as wind speed associated with the path can be inferred for a specific time period by analyzing the statistics of the ensemble. This paper proposes an approach for generating smooth scalar fields from such a predicted storm path ensemble, allowing the user to examine the predicted state of the storm at any chosen time. As a demonstration task, we show how our approach can be used to support a visualization tool, allowing the user to display predicted storm position – including its uncertainty – at any time in the forecast. In our approach, we estimate the likelihood of hurricane risk for a fixed time at any geospatial location by interpolating simplicial depth values in the path ensemble. Adaptively-sized radial basis functions are used to carry out the interpolation. Finally, geometric fitting is used to produce a simple graphical visualization of this likelihood. We also employ a non-linear filter, in time, to assure frame-to-frame coherency in the visualization as the prediction time is advanced. We explain the underlying algorithm and definitions, and give a number of examples of how our algorithm performs for several different storm predictions, and for two different sources of predicted path ensembles.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Viewing algorithms, information visualization, uncertainty, ensembles, hurricane prediction

1. Introduction

The US National Hurricane Center (NHC) begins posting advisories when a tropical storm, in either the Atlantic or Eastern Pacific region, develops into a cyclone, meaning an “organized system of clouds and thunderstorms that originates over tropical or subtropical waters and has a closed low-level circulation.” [NOA14d] Advisories take the form of several text documents, including the Forecast Advisory, which, along with other information, includes the storm center’s predicted latitude and longitude, wind intensity, and storm size for 12, 24, 36, 48, and 72 hours from the time of the advisory. Advisories are issued every six hours at 04:00, 10:00, 16:00, and 22:00 US Eastern Standard Time. They are downloadable from [NOA14b], and easily parsed to extract prediction information.

Besides the text documents with each advisory, the NHC produces several visualizations to assist in interpreting the information in the advisory. The most well-known of these is officially named the *Track Forecast Cone*, but is most often referred to as the *uncertainty cone* or *cone of uncertainty*. An example is shown in Figure 1a. According to the NHC website [NOA14a],

The cone represents the probable track of the center of a tropical cyclone, and is formed by enclosing the area swept out by a set of circles (not shown) along the forecast track (at 12, 24, 36 hours, etc). The size of each circle is set so that two-thirds of historical official forecast errors over a 5-year sample fall within the circle.

Thus, the width of the cone is an estimate of the uncertainty in the prediction, based on the NHC’s own performance in the recent past.

While this visualization gives an overall view of the path of the hurricane and its associated uncertainty, it does not facilitate important time and location-specific queries such as “What is the likelihood that the storm will hit my area, with hurricane strength winds, by 8:00 a.m. on Friday?” Emergency managers, responsible for planning in advance of an oncoming hurricane, are anxious to have such time and location-specific information readily available (personal communication: Matthew Green, Federal Emergency Management Agency representative at the NHC, March 2014). In addition, moving away from path-based predictions to time and location-specific predictions would facilitate the superposition of multiple storm variables, such as wind speed and storm size, on the display.

In a step towards providing more time and location-specific information, the NHC recently began augmenting an advisory with an ensemble of potential paths generated using Monte Carlo methods that follow an advisory’s path prediction, while accounting for its uncertainty (personal communication: Mark DeMaria, Technology and Science Branch Chief, NHC, March 2014). Figure 1b is an example of such

an ensemble, containing 1000 paths, for the same advisory as the uncertainty cone. Since the paths in an ensemble are sampled in time, they can carry with them time-based predicted storm characteristics such as storm size and wind speed, and since the paths are projected geospatially, they can be used to produce time and position based visuals. For example, the NHC uses them to produce “heat maps” of wind speed probabilities across the region predicted to be affected by the hurricane [NOA14c]. While such heat maps can be used to provide useful information, because they are spatially sampled on a grid they are coarse grained and subject to artifacts due to undersampling.

The primary contribution of this paper is to demonstrate an approach to generating and smoothly interpolating robust statistics from path ensembles, including outlying paths, to produce time-specific visualizations that inherently include uncertainty. As a demonstration piece, we outline the development of a visualization encoding three levels of positional storm-strike risk, for a specific point in time. An example of this visualization is shown in Figure 1c. Beyond strike position, the methods of the paper should be applicable to the visualization of other predicted variables such as storm speed, wind strength, storm size, and flood risk. The approaches used that will be of interest to the visualization community include:

- sampling each path from the ensemble at a specific time, to create an ensemble of points fixed in time,
- applying the concept of simplicial depth to provide a centrality ordering of time samples,
- developing an adaptive radial basis function interpolation technique that smoothly interpolates simplicial depth,
- designing a geospatial visualization, incorporating the concept of risk, based on the simplicial depth field.

2. Background and related work

2.1. Ensembles as an alternate visualization to the uncertainty cone

Although the uncertainty cone, shown in Figure 1a, is well known, and reasonably easy to explain, it has several apparent drawbacks. Broad et al. [BLWS07] have pointed out that the probabilistic concepts underlying the uncertainty cone can be easily misinterpreted. For instance, instead of reading the cone as the 66% likelihood region through which the storm center will pass, it is very easily misread as indicating an increasing storm size. Indeed, the NHC has begun placing a notice to this effect at the top of their most recent displays. In addition, the cone is a binary representation, possibly leading one to a false sense of security outside of the cone, or an exaggerated sense of certainty inside the cone.

In order to overcome some of the problems with the uncertainty cone visualization, Cox et al. [CHL13] proposed an alternative ensemble path visualization. Figure 2a is an example of their method, showing a prediction for Hurricane

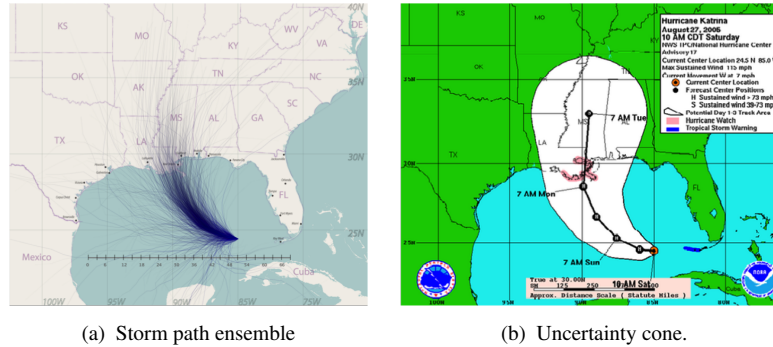


Figure 2: Cox et al.'s ensemble display vs. the uncertainty cone, NHC advisory 10 AM CDT, August 27, 2005.

Katrina, compared with Figure 2b, showing the NHC uncertainty cone for the same advisory. Using a Markov Random Field approach, their technique uses a combination of historical hurricane tracks and the current NHC advisory to continuously generate and draw possible hurricane tracks in such a way that the statistical distribution of the resulting ensemble closely matches that of the distribution implied by the cone of uncertainty. A user experiment demonstrated that this method results in estimates of hurricane direction that are as accurate as those made when viewing the uncertainty cone, but with some improvement in the user's estimation of the possibility of strikes outside of the cone.

None of the hurricane path prediction visualization methods, either proposed or in use, provide an integrated visualization of the storm, at a specific time and place, including representations of both uncertainty and storm characteristics.

2.2. Uncertainty visualization from ensembles

Uncertainty visualization has received much recent attention. The state of the art in the field has been carefully reviewed by Pang et al. [PWL96, Pan08], and more recently by Potter et al. [PRJ12]. Here, we concentrate on providing an overview of the techniques most relevant to dealing with path ensembles.

Approaches to gleaning statistical information from ensembles fall into two main categories: parametric and non-parametric. Parametric methods require an *a priori* assumption of the model describing the data distribution and focus on estimating the parameters (e.g. mean and variance for a Gaussian distribution) best matching the data. Non-parametric methods attempt to describe the data distribution without any assumption of a model. Since we have no basis on which to assume a given model, non-parametric methods seem to be the most attractive choice for our work.

Liu [LY90] developed the notion of *simplicial depth*, which is a powerful non-parametric approach for describing robust statistical summaries of an ensemble of samples. Simplicial depth defines the centrality of an individual point within an ensemble of points, and may be used to compute

a center outward ordering of the data. A sample point with larger simplicial depth is considered to be closer to the center of the ensemble, and thus more representative of the set of points. A sample point with smaller simplicial depth is considered to be less representative. Once the simplicial depth of each point in an ensemble has been determined, the points can be sorted based on their depth, with the indices of the sorted samples providing the structure of a cumulative distribution. We divide these indices by the number of samples to produce a normalized ranking of the points.

Simplicial depth is defined as follows. Let $V = \{v_0, v_1, v_2, \dots, v_{n-1}\}$ be the positions of an ensemble of n 2D points and let v_i, v_j , and v_k be three arbitrarily selected members of V . Let $\Delta_{i,j,k}$ denote the triangle formed by these points. Thus, the number of triangles is $N_\Delta = \binom{n}{3}$. The simplicial depth of a point in the ensemble is simply the number of such triangles containing the point. A straight-forward implementation of the simplicial depth calculation in the two-dimensional case takes $O(n^3)$ computational time. A more efficient algorithm taking $O(n \log n)$ time has been proposed by Rousseeuw and Ruts [RR96].

Whitaker et al. [WMK13] extended the idea of data depth to contour band depth, enabling statistical analysis of isocontours extracted from scalar fields. They used this contour band depth to estimate median, order statistics, and outliers for drawing what they call contour boxplots. Later, Mirzargar et al. [MWK14] built on their ideas to derive statistical characteristics from ensembles of multivariate curves extracted from flow fields, allowing them to draw curve boxplots. As one potential application, they demonstrated how ensembles of hurricane forecast tracks can be summarized using their method. Since these methods apply to paths they are not directly applicable to ensembles of points.

As an extension of their ensemble path visualization, Cox and House [CH13] began to explore the idea of interactive visualization of path ensembles at fixed points in time. Rather than rendering complete paths, showing a prediction over an entire forecast period, they implemented a time slider to fix a time within the prediction and rendered the corresponding point on each path in the ensemble. Fig. 3 shows

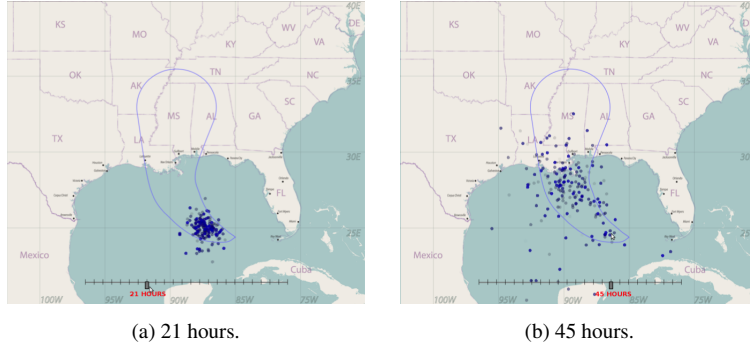


Figure 3: Path ensembles sampled at two different times, NHC advisory 10 AM CDT, August 27, 2005.

examples of their visualization for Hurricane Katrina, with the time set at 21 and 45 hours from the start of the advisory, compared with the corresponding uncertainty cone, shown in outline. While this method highlights the uncertainty in the forecast in both space and time, the naive rendering of predicted positions has clear drawbacks. Near the start time of the advisory, data points are tightly clustered, resulting in many overlapping points, making it impossible to take advantage of the point-based display to visually encode other variables using glyphs. At later times in the advisory, there is more spread in the points, but the visual clutter of the display, and the overemphasis of outliers makes the position distribution difficult to estimate visually. In the work being reported here, our goal is to build a coherent display of the distribution of the data, at a specific point in time, starting with similar time-specific point ensembles.

Going from a set of spatially-distributed points to a continuous representation of geospatial uncertainty requires the ability to derive a continuous scalar field over the spatial region covered by the data samples. In our work, we use simplicial depth to associate a scalar value with each of the data samples. We then build a continuous scalar field over this spatial region using *radial basis function* interpolation.

2.3. Radial basis function interpolation

Radial basis functions [BL88] have important applications in several fields requiring scattered data interpolation, most notably in machine learning [Ort96] and in computer graphics [CBC*01].

Radial basis function (RBF) interpolation builds a continuous function from a set of samples using radially symmetric kernel functions of position x , of the form $f(x) = \phi(\|x - x_0\|)$, where x_0 denotes the kernel center. A number of functions can be used as the radial basis kernel, with one of the most popular being the Gaussian kernel, which we use in the work reported here.

We associate with each data point i a location v_i , a weight w_i , and a kernel function ϕ_i . Then the RBF interpolation at a

given point x , is

$$f(x) = \sum_{i=0}^{n-1} w_i \phi_i(\|x - v_i\|). \quad (1)$$

If f_i is a scalar value known at each data sample, and we impose the condition that $f(x)$ interpolates the available data, for each data sample i we have the linear combination

$$f_i = f(v_i) = \sum_{j=0}^{n-1} w_j \phi_j(\|v_i - v_j\|).$$

Letting $\phi_{i,j} = \phi_j(\|v_i - v_j\|)$ yields the linear system

$$\begin{bmatrix} \phi_{0,0} & \phi_{0,1} & \dots & \phi_{0,n-1} \\ \phi_{1,0} & \phi_{1,1} & \dots & \phi_{1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{n-1,0} & \phi_{n-1,1} & \dots & \phi_{n-1,n-1} \end{bmatrix} \begin{bmatrix} w_0 \\ w_1 \\ \vdots \\ w_{n-1} \end{bmatrix} = \begin{bmatrix} f_0 \\ f_1 \\ \vdots \\ f_{n-1} \end{bmatrix},$$

whose unknowns are the weights w_i .

Since the gaussian kernel has infinite support, the solution matrix tends to be densely filled. We use this kernel because of the very broad spread of our data points, but kernels of finite support might be used to advantage in speeding computation by exploiting matrix sparsity.

A technique, often used with radial basis functions, is that of matrix regularization. Briefly, what is done is to add a small value to each element of the diagonal of the matrix Φ . This allows the solution to closely approximate the data at the sample points, rather than forcing a strict interpolation [ALP14]. In all of our RBF work reported here, we are using a regularization constant of 10^{-4} .

Equation 1, together with the above method of finding a set of weights, forms the basis of our method for turning an ensemble of predicted storm centers into the smooth continuous function of simplicial depth.

3. Methodology

Our approach to creating time-specific visualizations from predicted storm path ensembles begins by sampling the

paths at specific times. Instead of an ensemble of paths, this gives us ensembles of locations fixed in time, from which we can construct visualizations. How to do this in a compelling way is the primary question addressed by this paper. As already demonstrated in Figure 3, due to spatial under-sampling, a simple scatter plot of predicted locations tends to create a confusing display as the prediction time increases. One possible improvement would be to portray the underlying spatial density distribution implied by the hurricane predictions as a “heat map”. However, our early attempts to produce heat maps, by laying down a spatial grid and counting data points, led to displays that were too coarse where data points were tightly clustered, and too incoherent where they were widely spread.

3.1. Visualizing simplicial depth

Following our early, unsuccessful experiments using data density, we turned to the concept of simplicial depth. Simplicial depth is a measure of the centrality of data elements within a data set, giving a clean measurement associated directly with a data sample, and not dependent upon the local sampling density. Once simplicial depth is assigned to each sample, interpolation methods can be used to create a continuous simplicial depth scalar field from the available samples. We approached this task in two steps.

The first step is to compute the simplicial depth values for all sample points. The sample points are the predicted locations from a path ensemble, generated to correspond with a storm advisory. We then compute the simplicial depth of each sample point, using the fast algorithm of Rousseeuw and Ruts [RR96], and sort the sample points in ascending order by their depth. If n is the number of samples, a point’s sorted array index, divided by $n - 1$, is its normalized rank. The set of points contained within a ranking interval can be visualized by drawing its convex hull. Figure 4a is a scatter plot of the rainbow mapped simplicial depth values, and the red line is the convex hull of the [0%, 67%] rank interval. Although there are good reasons not to use a rainbow color map for displaying levels [BT07], we began by following the NHC’s convention for drawing heat maps, moving to a better designed system later in the study.

The second step is to interpolate across the evaluated simplicial depth values in order to provide a smoothly varying continuous representation. One interpolation method is to splat each point into the map, which produces results like those in Figure 4b, where a transparency is applied to each splat proportional to its depth. Splatting leaves many uncolored regions, and depends on sampling to a spatial grid. To overcome these problems, radial basis function (RBF) interpolation can be used, which produces a depth value anywhere in space, and can be used to produce very smooth visualizations, as illustrated in Figure 4c. While this is a big improvement over splatting, with the central region filled smoothly, the outer region is highly serrated. This is because

we were using a constant RBF kernel spread parameter, with data samples that are very unevenly spread.

3.2. Varying the radial basis function kernel spread

We found that dynamically adjusting the kernel radius, used in Eq. 1, solves the RBF interpolation problems caused by a highly nonuniform data density. We do this by selecting the kernel spread parameter based on the prediction density distribution – dense regions are interpolated with a narrow spread, while sparse regions are interpolated with wide spread.

To assign a density value to each sample point, we first create a density field, again using RBF interpolation. Do this by constructing a uniform rectangular grid over the regional map and counting the number of the sample points that fall into each cell. Since the grid cells are evenly distributed, the required density field can be obtained by using RBF interpolation with a constant kernel spread parameter. The Gaussian kernel centered at x_0 is given by $\phi(x) = \exp(-\frac{\|x-x_0\|^2}{2c^2})$. We associate the kernel spread parameter c with a bounding box, of major dimension w , containing all predicted locations, and let

$$c = sw, \quad (2)$$

where $0 \leq s \leq 1$ is a user defined fractional scale factor. Using this kernel spread parameter, we treat all of the grid cells containing any samples as sample points for building a set of weights to apply in Equation 1 for interpolating density.

Now, for each sample point v_i , in addition to simplicial depth d_i , we can ask for a density value ρ_i from the density field, which we use to determine an appropriate kernel spread parameter c_i for that point. Given the width of an individual grid cell δw , we choose the kernel spread to be inversely proportional, with constant of proportionality λ , to the number of data points per unit linear dimension,

$$c_i = \lambda \frac{\delta w}{\sqrt{\rho_i}}. \quad (3)$$

This gives us a spread parameter that adapts to the density of sample points in the neighborhood of each sample point, and provides smooth interpolation across all of the samples. Example visualizations using this interpolation approach are shown in Figure 5, with the NHC uncertainty cone shown in blue for reference. These examples also move away from the rainbow color map, using a color encoding meant to clearly show three nested risk regions.

3.3. Visualization design

The GIS and the mobile device communities have adopted the convention of presenting a geolocation containing uncertainty by a pale (i.e. transparent) blue dot, with radius conforming to some (e.g. 95%) confidence interval. Often this blue dot is augmented by a marker indicating the center

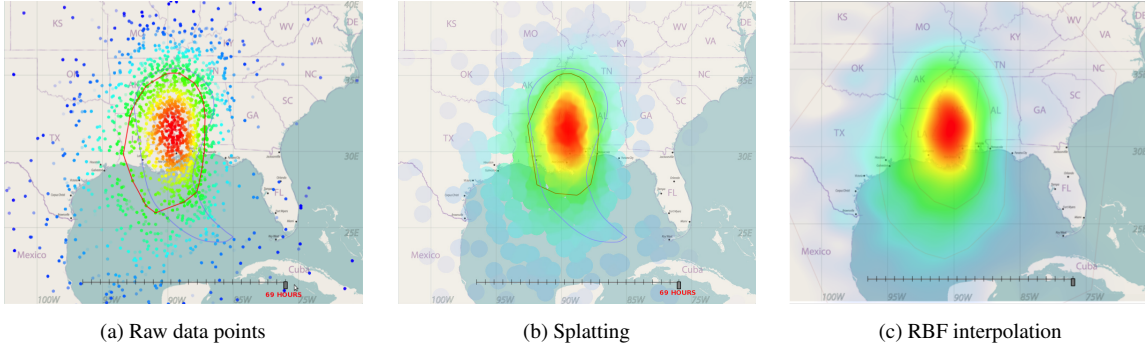


Figure 4: Simplicial depth visualization. NHC advisory 10 AM CDT, August 27, 2005.

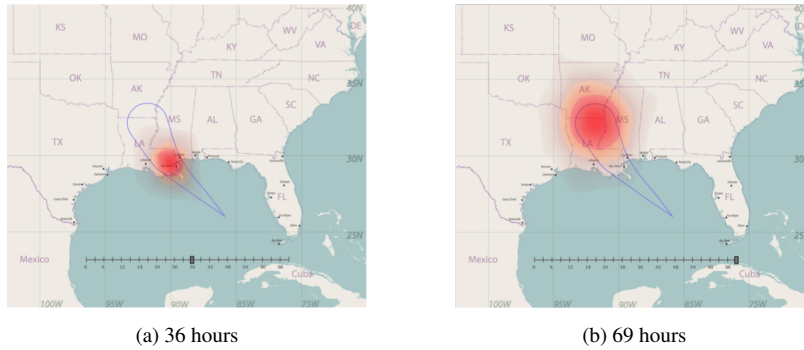


Figure 5: RBF interpolation with dynamically adjustable kernel radius. NHC Advisory: Hurricane Isaac, 1 pm CDT, Aug. 27, 2012.

of the dot, an outline, and sometimes by a transparency fade indicating the probability distribution. A recent series of experiments [BHM14] provides strong evidence that the pale blue dot, without border or center marking, and (contrary to intuition) without a transparency fade, provides visual cues most helpful in aiding experimental subjects to make correct spatial judgements incorporating uncertainty.

In our visualization design, having a strong feel for the uncertainty in a prediction is of paramount importance. Therefore, we elected to present the storm position as three overlapping confidence intervals: 33%, 66% and 99%. These intervals are unembellished except that each is of a different color, and each is of a different transparency. The 33% region is most opaque, the 66% region less opaque, and the 99% region is highly transparent.

Our color choices started with the color coding common in emergency systems, e.g. the U.S. Homeland Security Advisory System, which employ red, orange and yellow to present the top three levels of warning. However, yellow is a poor choice for our application, since highly transparent yellow over a white background is almost invisible. Thus, in our design we use red to indicate the region of highest risk, orange to indicate the medium risk region, and maroon to indicate the cautionary region. Given a depth interval and an associated color, the opacity is given by

$$\alpha = \alpha_0 + d_{\min}\beta, \quad (4)$$

where d_{\min} is the minimum normalized data depth of this interval, $0 \leq \alpha_0 \leq 1$ is the minimum desired opacity, and $0 \leq \beta \leq 1$ is a user supplied gain. For all of the relevant figures in this paper we have set $\alpha = 0.02$ and $\beta = 0.6$.

While the images shown in Figure 5 are close to what we envisioned, the colored depth intervals are somewhat irregular shapes, unlike the standard blue dots. The irregularity is induced by the Monte Carlo ensemble generation process and does not carry any useful information. Since the irregular risk regions are already nearly elliptical, we decided to replace them by minimum enclosing ellipses, rather than circular dots.

Minimum enclosing ellipses have the property that they preserve the aspect ratio of a region along two orthogonal axes. This orthogonality corresponds to the two sources of uncertainty in the prediction: hurricane bearing, and speed. Although their effects are not entirely independent, speed uncertainty tends to manifest in elongation of the risk region along the predicted path, while bearing uncertainty tends to broaden the region orthogonal to the path.

Figure 6 shows three snapshots of a Hurricane Isaac advisory, with the risk regions presented in this way. To determine the center, lengths of minor and major axes, and the rotation angle of the ellipses, we use an image moments-based algorithm proposed by [RVC02].

One of our eventual goals is to develop an interactive ap-

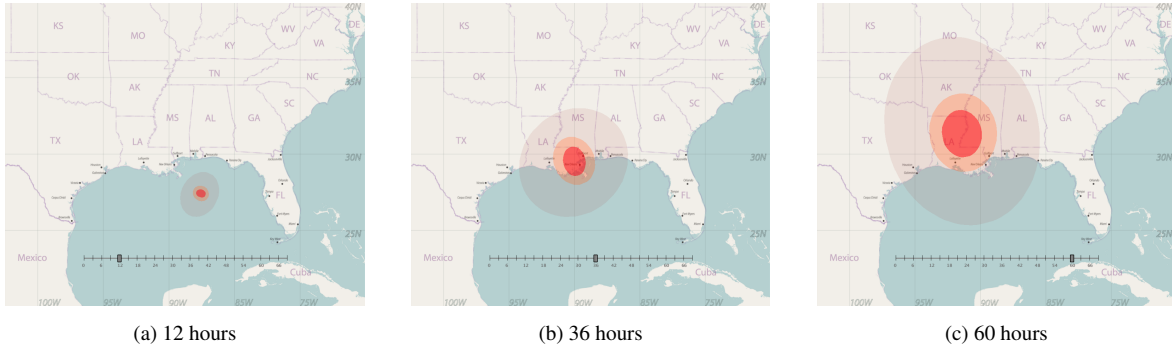


Figure 6: Minimum enclosing ellipses of depth intervals. NHC Advisory: Hurricane Isaac, 1 pm CDT, Aug. 27, 2012.

plication that embeds this approach to storm position visualization, allowing the user to “scrub” through time. Interpolation in time is easily achieved via any one of a number of interpolation methods across the known data points in a path in the ensemble. These are every hour in the NHC ensembles, and every three hours in the method by Cox and House. In our current work we are using simple linear interpolation, after determining that using a higher order method produced visually indiscernible results.

A problem with the ellipse representation became apparent while producing an animation to simulate scrubbing through time. When the ellipses become nearly circular, the choice of minor and major axis is not stable, leading to rapid 90° flips of ellipse orientation, which results in disturbing jitter in the animation. Figure 7 plots major axis angle over a series of 138 animation frames for the Hurricane Isaac example. This instability is very apparent in the top curve of the figure, showing several of these axis flips.

To eliminate the visual noise resulting from this instability, we developed a non-linear smoothing filter designed to ignore small changes in angle across time steps but suppress large changes. Detecting the potential for axis flips could be done by eigenvalue analysis, but our filter works well and fits more naturally into the signal analysis pipeline. We first compute $\Delta\theta = \frac{\theta^{[i]} - \theta^{[i-1]}}{\pi}$, the normalized difference between the ellipse orientation angle in time step i and the previous time step $i - 1$. We then compute two weights

$$w_1 = \frac{1}{1 + (q\Delta\theta)^2}, \quad w_2 = \frac{(q\Delta\theta)^2}{1 + (q\Delta\theta)^2},$$

where q is a parameter controlling the gain of the filter. The filtered angle at the current time is given by

$$\theta^{[i]'} = w_1\theta^{[i]} + w_2\theta^{[i-1]}. \quad (5)$$

If $\Delta\theta$ is small, w_1 dominates, selecting the current angle, while if $\Delta\theta$ is large, w_2 dominates, selecting the previous angle. In our experiments, setting $q = 14$ gave the best results. The middle curve in Figure 7 shows the result after applying this non-linear filter.

While the large angular jumps observed in the original

curve have been successfully removed, there are still small perturbations that interfere with frame-to-frame visual coherency. To filter out these bumps, we utilize a Gaussian filter, with kernel width 5, centered on the current time. This gives filtered results like those shown in the bottom curve in Figure 7, and provides smooth frame-to-frame transitions.

A possible criticism of this approach is that the filtered result may not be faithful to the data. The combination of the two filters is applied only to the orientation angle of the ellipse, not to the radii of the major and minor axes. The non-linear filter is not strongly sensitive to small angle changes, and is thus only removing large orientation flips in the sequence of visualizations. The smoothing filter is only removing small perturbations from the data, thus removing jitter. These have a negligible effect on the overall orientation of the ellipse angle, as can be clearly seen by comparing the curves in Figure 7.

4. Results

In this section, we show experimental results we have obtained, demonstrating the utility of the proposed visualization technique to explore time-specific predictions both from ensembles produced by the NHC, and generated by method of Cox et al. [CHL13] We also suggest settings for all user-defined parameters required in our approach.

Recall that we employ two RBF interpolations in our approach. We utilize a RBF interpolation with constant kernel radius to obtain a density field. Fractional parameter s adjusts the kernel spread parameter c based on the longest dimension of the sample bounding box, as given in Equation 2. To interpolate a simplicial depth field, we use another RBF interpolation with adaptive kernel radii, computed using a constant of proportionality λ as given in Equation 3. For rendering, we control opacities of the risk zones using Equation 4, using parameters to set a minimum opacity, and to scale opacity by simplicial depth, but these are fixed based on visual preference, as given in Section 3.3, and our ellipse angle filter parameters, are also fixed based on experimental results, as also explained in that section.




Figure 7: Rotation angles as a function of frame number.

The results of applying our visualization technique on ensembles produced by the National Hurricane Center have already been shown, for three different times in the prediction, for a Hurricane Isaac advisory in Figure 6. The parameters used were selected experimentally, and were $s = 0.35$, and $\lambda = 0.3$.

We also applied our method to advisories for several different storms, using path ensembles generated using the method of Cox et al. The results are shown in Figure 8. The images are for advisories for hurricanes Katrina, Rita, and Ida. In the top row we show the NHC uncertainty cone for each hurricane, and in the next three rows we show the predictions for 12 hours, 36 hours and 60 hours. The system parameters used to generate these results were $s = 0.35$, and $\lambda = 0.15$.

Importantly, none of the user-settable parameters in our approach needed to be adjusted across a variety of different storm advisories. The only parameter needing adjustment across ensemble generation methods was λ , controlling the adaptive kernel spread parameter in Equation 3. Because the method of Cox et al. generates a broader spread of hurricane paths than the method used by the NHC, the density in the denominator of Equation 3 tends to be low, so the value of λ must be decreased to compensate. Thus, our method appears to be robust, requiring only the tuning of one parameter, and this only if there is a change in ensemble generation method.

A criticism of our approach is that, as time progresses into a prediction, the sizes of the risk regions increase, leading to the very strong perception that the storm itself is increasing in size. This problem is also inherent in the NHC uncertainty cone, and indeed in any geospatial display that attempts to track dispersion in a prediction using a summary display. This is not a soluble problem, as long as spatial extent is being used as an uncertainty measure. As yet unpublished studies, underway in our research group, are producing strong evidence that ensemble displays do not induce this same perceptual anomaly. Therefore, our research plan is to build on the work reported here, resampling of the simplicial depth field in a well-distributed way to produce a set of exemplar storm positions that can be displayed as points, but without the visual clutter and confusion of the early work of Cox and House shown in Figure 3.

5. Conclusions

We have presented a visualization technique to provide exploration of time-specific predictions from an ensemble of potential hurricane paths. These paths are sampled in time, to create a set of points for each time period, which are assigned a scalar value associated with their simplicial depth. We then create a scalar field over the region covered by the samples using radial basis function interpolation. Using this field, we determine risk regions based on simplicial depth, and render them using best-fit elliptical approximations. The approach has been shown to be robust across a number of storm predictions, and across two different Monte Carlo path ensemble generation approaches.

This work provides a simple geospatially located visualization, incorporating uncertainty, and keyed to a particular point in time. Our intent is that this will form the basis for future research leading to a set of interactive tools for exploring a hurricane prediction in both time and space. Given the structure for spatial interpolation that we have developed, it should be possible to interpolate storm parameters other than strike risk, such as storm speed, bearing, wind speed, size, and flood risk. This has the potential to enable development of an integrated hurricane prediction visualization application to be used in the field by emergency managers.

One impediment to developing such an interactive application is the speed of the current algorithm. While most stages of the computation can be easily accelerated to interactive rates, the solution for RBF interpolation weights involves solving an $N \times N$ linear system in the number of sample points. This is prohibitively slow for a typical system of 1000 or more samples. Our plans include investigating fast algorithms for getting a good approximate solution to this system, with special attention to choosing a subset of the samples that minimizes approximation error.

While this study was intended to support future work on visualization of time-specific predictions from time-parameterized path ensembles, it also stands alone as a project to develop a new visualization tool for evaluating hurricane risk. The next natural step in this side of the work will be to conduct a study comparing how users perform on time and place specific risk evaluation tasks using this visualization versus other proposed alternatives, including the uncertainty cone itself, and a scattered point approach displaying color-coded path samples.

6. Acknowledgements

The authors would like to thank the anonymous EuroVis 2015 reviewers for their suggestions, which resulted in several improvements to the paper. This material is based upon work supported by the US National Science Foundation under Grant Nos. IIS-1212501 and IIS-1212806.

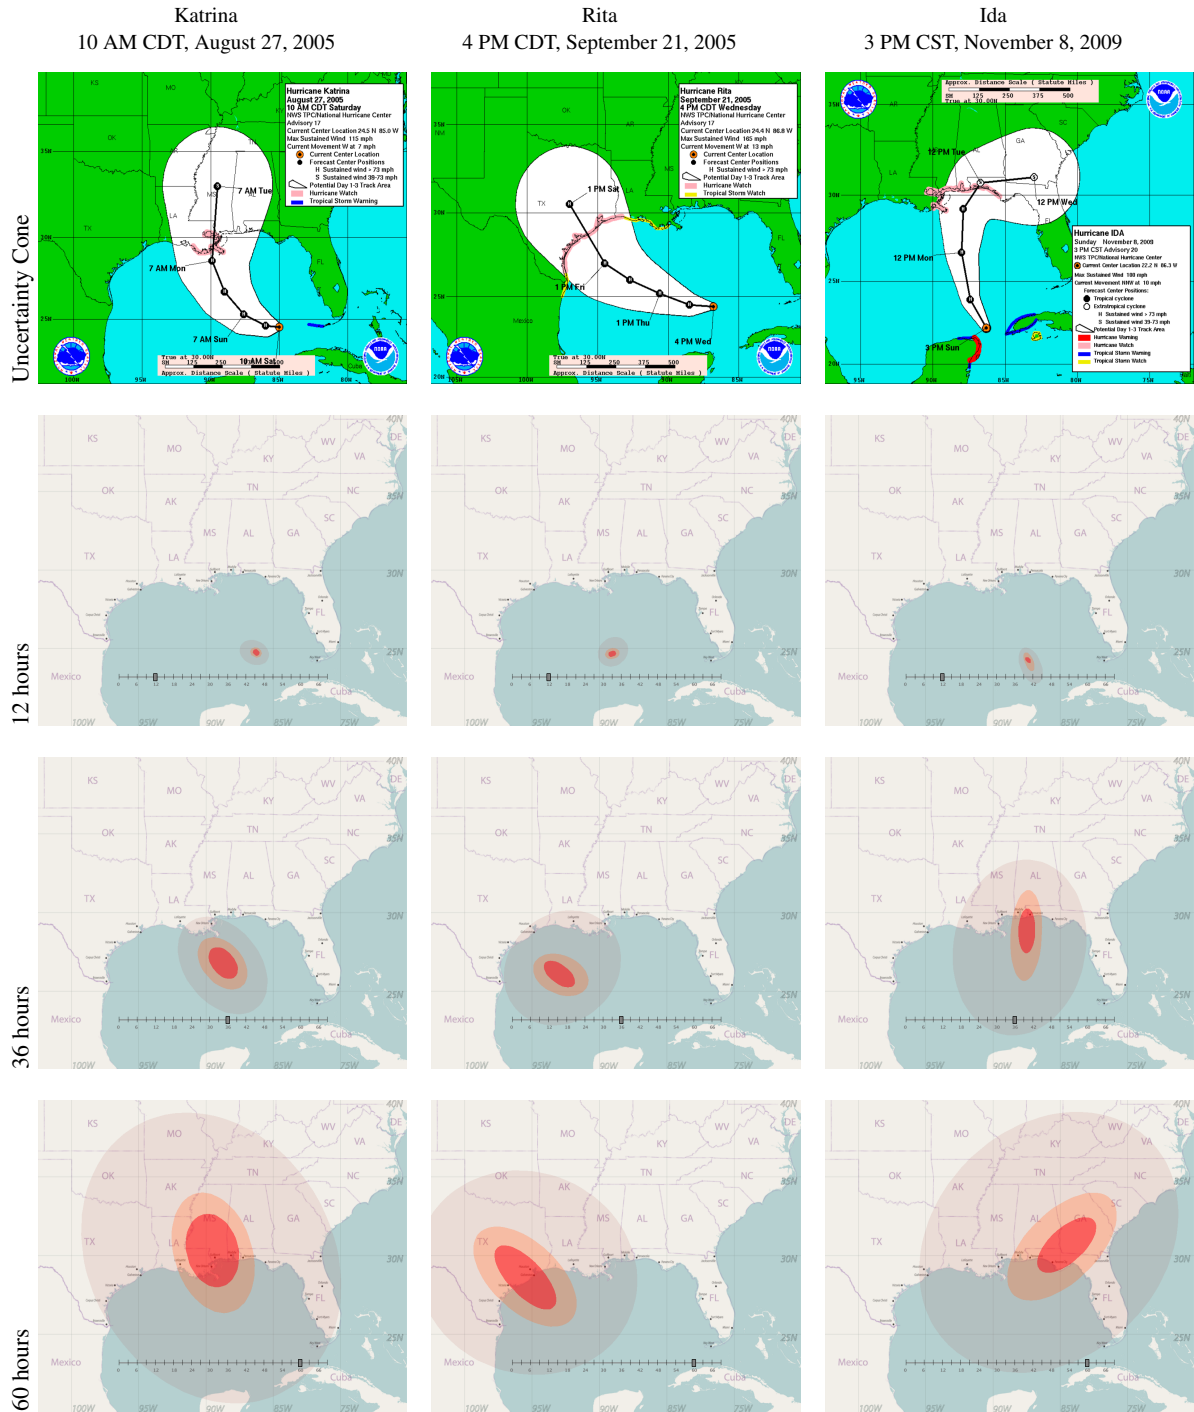


Figure 8: Time-specific visualizations of risk regions for four different hurricane advisories.

References

- [ALP14] ANJYO K., LEWIS J., PIGHIN F.: Scattered data interpolation for computer graphics. In *SIGGRAPH Course Notes*, 2014. 2014, pp. 28–31. 4
- [BHM14] BARRETT T., HEGARTY M., MCKENZIE G., GOODCHILD M.: Am I really there? Evaluating visualizations of geospatial uncertainty. *Spatial Cognition poster*, Sept. 2014. 6
- [BL88] BROOMHEAD D. S., LOWE D.: Multivariable functional interpolation and adaptive networks. *Complex Systems* 2 (1988), 321–355. 4
- [BLWS07] BROAD K., LEISEROWITZ A., WEINKLE J., STEKETE M.: Misinterpretations of the “cone of uncertainty” in Florida during the 2004 hurricane season. *Bulletin of the American Meteorological Society* 88, 5 (2014/11/11 2007), 651–667. 2
- [BT07] BORLAND D., TAYLOR R.: Rainbow color map (still) considered harmful. *Computer Graphics and Applications, IEEE* 27, 2 (March 2007), 14–17. 5
- [CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (2001), SIGGRAPH '01, pp. 67–76. 4
- [CH13] COX J., HOUSE D.: Visualizing uncertainty as an interactive ensemble. *COSIT Workshop on Visually-Supported Reasoning with Uncertainty*, Sept. 2013. 3
- [CHL13] COX J., HOUSE D., LINDELL M.: Visualizing uncertainty in predicted hurricane tracks. *International Journal of Uncertainty Quantification* 3, 2 (2013), 143–156. 2, 7
- [LY90] LIU, Y R.: On a notion of data depth based on random simplices. *The Annals of Statistics* 18, 1 (1990), 405–414. 3
- [MWK14] MIRZARGAR M., WHITAKER R., KIRBY R.: Curve boxplot: Generalization of boxplot for ensembles of curves. *Visualization and Computer Graphics, IEEE Transactions on* 20, 12 (Dec 2014), 2654–2663. 3
- [NOA14a] NOAA N. H. C.: Definition of the NHC track forecast cone, Dec. 2014. URL: <http://www.nhc.noaa.gov/aboutcone.shtml>. 2
- [NOA14b] NOAA N. H. C.: NHC active tropical cyclones, Dec. 2014. URL: <http://www.nhc.noaa.gov/cyclones/>. 2
- [NOA14c] NOAA N. H. C.: Surface wind speed probabilities (120 hours), Dec. 2014. URL: http://www.nhc.noaa.gov/archive/2012/ISAAC_graphics.shtml. 2
- [NOA14d] NOAA N. H. C.: Tropical cyclone climatology, Dec. 2014. URL: <http://www.nhc.noaa.gov/climo/>. 2
- [Orr96] ORR M. J. L.: Introduction to radial basis function networks, 1996. URL: <http://www.anc.ed.ac.uk/rbf/papers/intro.ps>. 4
- [Pan08] PANG A.: Visualizing uncertainty in natural hazards. In *Risk Assessment, Modeling and Decision Support*, vol. 14 of *Risk, Governance and Society*. 2008, pp. 261–294. 3
- [PRJ12] POTTER K., ROSEN P., JOHNSON C. R.: From quantification to visualization: A taxonomy of uncertainty visualization approaches. *IFIP Advances in Information and Communication Technology Series* (2012), 226–249. (Invited Paper). 3
- [PWL96] PANG A. T., WITTENBRINK C. M., LODH S. K.: Approaches to uncertainty visualization. *The Visual Computer* 13 (1996), 370–390. 3
- [RR96] ROUSSEEUW P. J., RUTS I.: Algorithm AS 307: Bivariate location depth. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 45, 4 (1996), pp. 516–526. 3, 5
- [RVC02] ROCHA L., VELHO L., CARVALHO P.: Image moments-based structuring and tracking of objects. In *Computer Graphics and Image Processing, 2002. Proceedings. XV Brazilian Symposium on* (2002), pp. 99–105. 6
- [WMK13] WHITAKER R., MIRZARGAR M., KIRBY R.: Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. *Visualization and Computer Graphics, IEEE Transactions on* 19, 12 (Dec 2013), 2713–2722. 3