

Measuring Event Probabilities in Uncertain Scalar Datasets using Gaussian Processes

Steven Schlegel
University of Leipzig
schlegel@informatik.uni-leipzig.de

Sebastian Volke
University of Leipzig

Gerik Scheuermann
University of Leipzig

ABSTRACT

In this paper, we show how the concept of Gaussian process regression can be used to determine potential events in scalar data sets. As a showcase, we will investigate climate data sets in order to identify potential extrem weather events by deriving the probabilities of their appearances. The method is implemented directly on the GPU to ensure interactive frame rates and pixel precise visualizations. We will see, that this approach is especially well suited for sparse sampled data because of its reconstruction properties.

Keywords

Gaussian Process Regression, OpenCL Programming, Climate Data

1 INTRODUCTION

The visualization of tensor data that is given on discrete positions typically requires interpolation of the data values in between the sample positions. Usually, this task is solved by using linear interpolation. However, in the case that the given data is uncertain, this method is not feasible [BUR96]. In [SCH12], Schlegel et al. proposed the use of Gaussian process regression to overcome the aforementioned problems.

We will show how this method can be used to calculate occurrence probabilities of certain events. Therefore, we present an approach that combines fast computation on the GPU as well as visualizations of arbitrarily dense samplings. This is achieved by using the reconstruction properties of Gaussian Process Regression. We use datasets from the climate research domain to demonstrate our method. One of the major tasks in climate research is the prediction and the understanding of extreme weather events. Events like heat waves, heavy precipitation (resulting in floods) or hurricanes have a large impact on society and politics. Decision makers rely on climate simulation results as accurate as possible. Those results should hold characteristics of extreme events like location, frequency and intensity. There is a lot of research that points out that climate and extreme events undergo a change especially in frequency and intensity. Emanuel [EMA05], for instance, pointed out that the destructiveness of hurricanes immensely increased of the past 30 years. A quick overview for other observed changes can be viewed here: http://www.ipcc.ch/publications_and_data/ar4/wg2/en/ch10s10-2-3.html#table-10-3.

2 RELATED WORK

As pointed out by [PAN96] it is important to keep in mind that uncertainty of scalar data can reside in the data value or in the position of the data point or in both. In this paper, we deal with the uncertainty of the data value itself. Pöthkow et al. [POE11a] also targeted this issue and presented an uncertain counterpart for isocontours [LOR87]. Therefore, they calculate the so called level-crossing probability (LCP). In a given interval, the probability is computed that a certain threshold is crossed within. Therefore, they interpolated the expected values and the roots of the central moments to interpolate the probability density function. Schlegel et al. [SCH12] proposed the method of Kriging to interpolate the mean and the variance in an uncertain Gaussian field. They also applied their method to compute LCP. Several acceleration methods were employed in [SCH15] to enable a fast computation of Kriging in 3D scalar fields. They created interactive 3D visualizations of the mean field and showed the confidence of the computation by depicting areas of high uncertainty. Therefore, they computed an upper boundary for the posterior variance. In contrast, we aim to provide probabilities for the data to exceed (or fall below) certain thresholds using the exact posterior variance. [ATH13] analyzed the effects of uncertainty to linear interpolation and iso-surface extraction. The extension of [POE11a] to correlated data was done in [POE11b]. To reduce the heavy computation time (mainly caused by Monte Carlo Sampling), two methods called maximum edge crossing probability and linked-pairs to approximate the level-crossing probabilities were introduced in [POE13b]. To overcome the restrictions of predefined probability distributions [POE13a] introduced nonparametric models (empirical distributions, histograms, and kernel density estimators) to compute the probability of features in an

uncertain field. Based on [POE11a], Pfaffelmoser et al. [PFA11] developed an algorithm to compute the so called isosurface first crossing probability. It is an algorithm that incrementally uses a front-to-back volume ray casting to visualize that probability. The rendering is enriched by additionally depicting surfaces of the stochastic distance function (SDF-surfaces). Additional work to compute the gradient of the probability density function of uncertain 3D scalar fields was done in [PFA12]. Kniss et al. [KNI05] try to perform classification of medical volume data under uncertainty. They base their transfer function on what they call the decision boundary distance that is computed for every class, which is a maximal log-odds ratio of all the other classes. Roughly speaking, it is a measurement of the risk of being wrong to assume that the current class is the correct one. A more thorough overview of visualization of uncertain data can be read in [BRO12].

3 GAUSSIAN PROCESS REGRESSION

As pointed out earlier, we can not rely on standard techniques like linear interpolation, when the need for interpolation of uncertain data arises. We need to regard the uncertainty of the sampled data points, as well as their respective correlation. If these samples are Gaussian-distributed random variables, it is suitable to use the concept of *Gaussian processes*. Interpolating random variables in a Gaussian process is also known as *Kriging* [KRI51] or *Gaussian process regression* [RAS06]. The basic approach is to assume a prior Gaussian distribution for any (continuous) position in the data set. By considering the given samples and a defined covariance between the data points, this prior distribution is turned into a posterior Gaussian distribution that matches the given data more precisely in the sense of reducing the variance of the distribution. For details, on how to define the prior distribution, we refer to [SCH12].

A Gaussian process given on a domain S defines a Gaussian distributed random variable at any position $s \in S$. It is defined by a mean function at every position s and a covariance function between any two positions s and s' , e.g., see [ADL11]:

$$\begin{aligned} \mu : S &\mapsto \mathbb{R} & \mu(s) &= \mathbb{E}[f(s)], \\ k : S \times S &\rightarrow \mathbb{R} & k(s, s') &= \mathbb{E}[(f(s) - \mu(s))(f(s') - \mu(s'))], \end{aligned} \quad (1)$$

where the mean function is assumed to be constant. Our choice for the covariance function throughout this paper is the squared exponential. This covariance function models an exponential drop of the covariance with increasing distance of the data points. It is often used in the field of Geostatistics and is given by

$$k(s, s') = \sigma_p^2 \exp\left(-\frac{1}{2l^2}|s - s'|^2\right), \quad (\sigma_p^2, l > 0). \quad (2)$$

The parameters σ_p^2 (prior variance) and l (length scale) are hyperparameters. Throughout this paper our choice for l will be 1. The choice for σ_p^2 will be discussed in section 4.1. Gaussian processes, as well as the optimization of the hyperparameters, are discussed in detail in [RAS06].

Let S be sampled with N Gaussian distributed variables at positions s_i , $i = 1, \dots, N$, with $X(s_i) = X_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$ and the covariance function $k(s, s')$. Then one can calculate the covariances between those sample points and generate the covariance matrix

$$K = \begin{pmatrix} k(s_1, s_1) + \sigma_1^2 & \dots & k(s_1, s_N) \\ \dots & \dots & \dots \\ k(s_N, s_1) & \dots & k(s_N, s_N) + \sigma_n^2 \end{pmatrix}, \quad (3)$$

The posterior distribution at position s is then defined as

$$X(s) \sim \mathcal{N}\left(k(\vec{s})^T K^{-1} \vec{\mu}_i, \quad k(s, s) - k(\vec{s})^T K^{-1} k(\vec{s})\right), \quad (4)$$

with $\vec{\mu}_i$ being the vector of the means of the sampled X_i and $k(\vec{s}) = (k(s, s_1), \dots, k(s, s_N))^T$. It can be shown that the variance of the posterior distribution is minimized, when estimating $X(s)$ in that way.

Additionally, by defining the basis functions (see [SCH15])

$$\phi_i(s) = \begin{cases} -1, & \text{if } i = 0 \\ \sum_{j=1}^N (K^{-1})_{ij} k(s_j, s), & \text{otherwise} \end{cases} \quad (5)$$

we can write the computation of $X(s)$ as

$$X(s) = \sum_{i=0}^N X_i \phi_i(s). \quad (6)$$

with

$$\begin{aligned} X(s) &\sim \mathcal{N}\left(\mu(s) = \sum_{i=1}^N \phi_i(s) \mu_i, \right. \\ &\left. \sigma^2(s) = k(s, s) - \sum_{i=1}^N \phi_i(s) k(s_i, s)\right). \end{aligned} \quad (7)$$

Using Eq. 6, it is possible to compute the derivative of $X(s)$ by differentiating the basis functions:

$$\begin{aligned} \frac{\delta \mu(s)}{\delta s^{(n)}} &= \mathbb{E}\left(\sum_{i=1}^N \frac{\delta \phi_i(s)}{\delta s^{(n)}} X_i\right), \\ \frac{\delta \sigma^2(s)}{\delta s^{(n)}} &= \text{Var}\left(\sum_{i=1}^N \frac{\delta \phi_i(s)}{\delta s^{(n)}} X_i\right) \end{aligned} \quad (8)$$

This form of Kriging is called *simple* Kriging. Other forms of Kriging differ mainly in the form of the assumption of the prior (e.g. see [GOV97]).

4 METHOD

In this section, we show how to work with Gaussian regression on climate research datasets. Our goal is to calculate the probability that a certain threshold will (not) be exceeded. Thus, we construct uncertainty variables on the basis of the original climate data. We use time dependent data from a global climate simulation given on a rectilinear 2D grid. This could be interpreted as a time series of data at each grid position.

4.1 Modeling the Gaussian Process

First we normalize each time series by removing the seasonal component. Normalized time series have a Gaussian distribution [LAR12]. So the normalization enables us to compute the variability of the data and, of course, to apply our method. Therefore, we replace the value at each time step $v(ts)$ with the average of these values from the annual cycle. For example, if we have monthly means, we replace $v(ts)$ with the average of the period $[ts - 6; ts + 6]$. Furthermore, we can remove a linear trend in the dataset by replacing each time step with its forward difference, i.e. $v(ts) = v(ts + 1) - v(ts)$. Although the trend in temperature time series is not necessarily linear, we can use this simplification for relatively short time periods (i.e. 30 years). This step is optional and best suited for datasets where linear trends may disturb the normalization (e.g. temperature data). In order to estimate the variability of the simulated data at that position, we derive the empirical variance for every grid point based on the normalized time series.

As described in section 3, the basic principle of Gaussian process regression is to turn a prior Gaussian distribution, which we observe on the data, into a posterior Gaussian distribution by taking the given samples and the covariance into account. The prior distribution describes the uncertainty in the data acquisition method. The posterior distribution on the other hand is in general a better estimator for the uncertainty in the dataset in the sense of having less variance. To model the prior distribution, we need its mean and its variance. The mean is constantly zero. This can be accomplished by subtracting the empirical mean of the data from the samples. When we display the results, we simply add the posterior mean back on each sample point. This applies to an error model, where the observed value is the sum of the true (unknown) value and a zero mean Gaussian error. The prior distribution variance is the maximum of all the variances which we extracted at the grid points. The maximum variance in the dataset is an obvious choice for the prior variance, because the variance of the data acquisition method is at least as big as the maximum variance residing in the dataset. The prior variance (or signal variance) is the factor σ_p^2 in the covariance function, see eq. 2, which results in $k(s, s) = \sigma_p^2$ (see eq. 4).

Algorithm 1 Creating The Cell Cache on the CPU

```

·  $l$  := length scale
·  $d$  := cell diameter
·  $n$  := number of Cells
· CellCache[ $n$ ]
for  $i = 0$  to  $n$  do
    ·  $b$  := barycenter(Cell)
    ·  $P$  := all sample points in radius  $[b - (3l + d), b + (3l + d)]$ 
    · create and invert covariance matrix using all points in  $P$ 
    · CellCache[ $i$ ] := inverted covariance matrix, its positions, and its samples
end for
· send CellCache to graphics card
for all pixels do
    · calculateColor()
end for
    
```

Algorithm 2 "calculateColor()" - GPU Colormap Algorithm.

```

·  $t$  := threshold
·  $idx$  := cell index of Pixel
·  $pos$  := position in (2D) world space of Pixel
·  $n$  := number of sample points in CellCache[ $idx$ ]
for  $i = 0$  to  $n$  do
    · compute basisfunction  $\phi_i$  using  $p$  and CellCache[ $idx$ ]
end for
· compute distribution using the basisfunctions and Eq. 7
· calculate probability  $p$  that the value at  $pos$  falls below  $t$ 
· color pixel according to  $p$  and given color map
    
```

4.2 Implementation

Gaussian process regression performs poorly on many datasets. The reason is the storage and the inversion of the covariance matrix. A method to reduce those requirements, is the use of many small Gaussian processes (and thus covariance matrices) instead of one large process. For regular sampled datasets, it is feasible to create a small Gaussian process for each grid cell composed of the data points lying in a $3l + d$ radius of the bary center of the cell. Where l is the length scale of the covariance function and d is the diameter of the cell. This approach is described in more detail in [SCH12]. The result is, that we have to invert one relatively small covariance matrix for every cell instead of one large matrix, which can also be done in parallel for another speed up.

When an inverted covariance matrix (see Eq. 3) for each cell of the dataset is computed, we send those matrices to the GPU. We also store the dataset itself as well as the indices of the data points that belong

to each of those local Gaussian processes on the GPU. The next step is to compute the probability distribution according to eq. 4 for every pixel position which lies inside our dataset domain. Now, we are able to compute probabilities that values at that pixel fall below or exceed certain thresholds (which in our case are indices for extreme weather events). Furthermore, we are able to accumulate those probabilities over several time steps in order to compute probabilities that the values fall below or exceed the threshold over a given period of time. Given the fact, that we compute everything on the GPU, we finally use the given probabilities for each pixel to create a pixel precise color map which can be rendered immediately by writing the color into the frame buffer. The main advantage of this approach is, that we send the required data (inverted covariance matrix and the point indices) to the GPU once, which will process both tasks, namely computing and rendering. There is no need to send the data back to the CPU. Our implementation uses the OpenCL framework. The algorithm's pseudocode for processing the data on the CPU to send it to the GPU is given in algorithm 1. The pseudocode for the GPU implementation of the calculateColor() function is depicted in algorithm 2.

5 RESULTS

The data we use is a temperature data set from a global climate simulation of IPCC scenario A1B with the coupled atmosphere ocean general circulation model (AOGCM) ECHAM5-MPIOM, which was carried out as a contribution to the International Panel on Climate Change Assessment Report 4 (IPCC AR4) [SOL07]. It is given on a 192 x 96 rectilinear grid. At each grid point, there is a temperature data time series of monthly means from the year 1860 to the year 2100. We used the data from 1860 to 1890 in order to calculate the variances for every grid position.

After the variances are calculated, we assigned them to temperature data (simulated by the same model) given on a 6 hour basis to calculate the probability that the temperature of 273.15°K (0°C) was not exceeded in the whole month of January 2001 (124 time steps). The prior distribution is calculated as described in section 4.1. The result is given in Fig. 1. Additionally, we zoomed into one area containing probability transitions to demonstrate that this kind of interpolation in fact enables rendering using arbitrary zoom factors and still providing smooth results.

This kind of application is also interesting with respect to regional climate changes. Therefore, we used as a second example a data set from a simulation with the regional climate model CLM [HOL08]. This is a community model for the German climate research, originally based on the LM forecast model of the German Weather Service (DWD). The CLM simulation was



Figure 1: Colormap of the probability that the temperature of 273.15°K (0°C) is not exceeded in the whole month of January 2001.




Figure 2: Colormap of the probability that the surface runoff exceeds a threshold of 60 kg/m² at least five consecutive days in the summer of 1961.

forced with results of the IPCC scenario A1B simulation with ECHAM5 / MPI-OM. The particular dataset we used is the surface runoff. The surface runoff is the amount of water that cannot be absorbed by the soil. It is an accumulated quantity mainly based on precipitation, snow melting, and the water content of the soil surface and is an indicator for floods. If large volumes of surface runoff flow into a river in a short period of time, the likeliness of a flooding increases. In climate research, one typically counts how often the data exceeds or falls below a threshold within a certain interval to identify weather extremes; see [SIL] and references. As in the example above, we can normalize the time series data at each grid point and calculate the probability that a certain threshold of the surface runoff is exceeded. With our method, we are able to interpolate the data, incorporate the uncertainty of the simulations into the interpolation and compute the probability pixel by pixel. A second step would be to assign the results to the corresponding river catchment basin and accumulate the probabilities over this area to derive potential risks for people living near those rivers, see [SCH13]. Unfortunately this is beyond the scope of this paper.





Figure 3: Images showing the regression error for $k = 2$ (Fig.3(a)) and $k = 3$ (Fig.3(b)).


The data we used is a cutout of 65×50 grid points of daily CLM data for Europe centered on Germany. The grid is regular (data stream 3) and has a spacing of 0.2° (approximately 20km). It is a simulation run for the 20th century (20C) from 1961-1990. In Fig. 2, we depicted the probability that the surface runoff exceeds the amount of 60 kg/m^2 in at least five consecutive days in the summer (June, July, and August) of 1961 as a showcase. We can judge from the image that within Germany especially regions inside the catchment basin of the river Rhine have a high probability of exceeding the threshold. This method can be a valuable tool when performing research on larger time scales to evaluate the development of such quantities in order to draw conclusions on climate change.

5.1 Error Analysis

In Section 4.2, we showed that cutting off the exponential covariance function provides smaller cell caches and thus a faster computation of the regression result. On the other hand, this technique introduces errors, which we will analyze using a sample 2D climate data set (sea level pressure) on a 192×96 grid. Therefore, we first calculate the inverted covariance matrix for all the grid points, i.e. we do not cut off the covariance function. Then, we do an regression of the samples with this covariance matrix and use this as a ground truth. The interpolated field again is regularly sampled at 2880×1440 positions. In the next step, the field is interpolated and resampled the same way but using short-



(a)



(b)

Figure 4: Diagrams showing the exponential drop of the error with increasing covariance influence radius (Fig.4(a)), as well as the exponential gain of computation time (Fig. 4(b)) for different k .

ened covariance functions, each with a different length: $kl + d$, $k = 1, \dots, 10$. Finally, we compare those fields with the ground truth and calculate the absolute average error.

An error colormap for $k = 2$ and $k = 3$ can be seen in Fig. 3. We can conclude from those images, that the error in fact decreases, when the covariance function gains a larger influence radius. But this comes at a comparably high computational cost. The development of error and computational cost to create the cell cache is depicted in Fig. 4. We can conclude an exponential drop of the error as well as an exponential increasing computation time with increasing k . The average relative error for this particular field ranges from 0,0057% to 4,27%.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we showed how climate data can be interpolated in an arbitrarily dense matter. Therefore, we use the framework presented in [SCH12] as a basis and extend it by implementing it on the GPU. This enables interactive visualizations with arbitrarily dense samplings. We normalized the time series data to extract the simulation variance. Then we were able to calculate probabilities for certain occurrences like the appearance of extreme events while considering the variance. The resulting color maps can be computed in any desired resolution. Especially when the underlying data is sparse (like in Fig. 2), we nevertheless are able to provide visualizations with smooth transitions.

We want to point out that the variance we used in this paper is computed from the given data. Of course there

are cases, when the variance may known a priori, for example when the simulation model has a known error. The proposed method also works with that kind of uncertainty as long as it is Gaussian distributed.

In general, Gaussian process regression works on any type of scalar data regardless of the underlying topological structure. The only precondition is, that there has to be covariance defined between any of the data points in the given domain. This covariance is often modeled with a covariance function, which then serves as the interpolation kernel. Since we use Gaussian process regression for the means of data interpolation, it is suitable to use a distance based covariance function. With these prerequisites, Gaussian process regression resembles inverse distance based interpolation methods (e.g. Shepard interpolation).

A suitable application of this paper is to study the probabilities of extreme weather events. As mentioned before in Sec. 5, this work can therefore be extended by assigning the calculated probabilities to certain areas of interest to draw conclusions on the danger for flooding, droughts et cetera. We consider this as future work.

7 REFERENCES

- [ADL11] Adler, R.J. and Taylor, J.E. Topological complexity of smooth random functions: École d'Été de Probabilités de Saint-Flour XXXIX - 2009. Lecture Notes in Mathematics. Springer, 2011.
- [ATH13] Athawale, T. and Entezari, A. Uncertainty quantification in linear interpolation for isosurface extraction. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2723–2732, Dec 2013.
- [BRO12] Brodlié, K. and Osorio, R.A. and Lopes, A. A review of uncertainty in data visualization. In *Expanding the Frontiers of Visual Analytics and Visualization*, pages 81–109. Springer, 2012.
- [BUR96] Bursal, F.H. On interpolating between probability distributions. *Applied Mathematics and Computation*, 77(2-3):213 – 244, 1996.
- [EMA05] Emanuel, K. Increasing destructiveness of tropical cyclones over the past 30 years. *Nature*, 436(7051):686–688, 2005.
- [GOV97] Goovaerts, P. *Geostatistics for Natural Resources Evaluation*. Applied geostatistics series. Oxford University Press, 1997.
- [HOL08] Hollweg, H.-D., Böhm, U., Fast, I., Hennemuth, B., Keuler, K., Keup-Thiel, E., Lautenschlager, M., Legutke, S., Radtke, K., Rockel, B., Schubert, M., Will, A., Woldt, M., and Wunram, C. Ensemble simulations over europe with the regional climate model clm forced with ipcc ar4 global scenarios. Technical Report, 2008.
- [SOL07] Intergovernmental. *Climate Change 2007 - The Physical Science Basis: Working Group I Contribution to the Fourth Assessment Report of the IPCC*. Cambridge University Press, Cambridge, UK and New York, NY, USA, September 2007.
- [KNI05] Kniss, J.M., Van Uitert, R., Stephens, A., Li, G. S., Tasdizen, T., and Hansen, C. Statistically quantitative volume visualization. pages 287–294, October 2005.
- [KRI51] Krige, D G . A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Chemical Metallurgical and Mining Society of South Africa*, 52(6):119–139, 1951.
- [LAR12] Larsen, R.J., and Marx, M. *An introduction to mathematical statistics and its applications*; 5th ed. Prentice Hall, Boston, MA, 2012. The book can be consulted by contacting: IT-ES-DNG: Abler, Daniel.
- [LOR87] Lorensen, W.E., and Cline, H.E. Marching cubes: A high resolution 3d surface construction algorithm. *COMPUTER GRAPHICS*, 21(4):163–169, 1987.
- [PAN96] Pang, A.T., Wittenbrink, C.M., and Lodh, S.K. Approaches to uncertainty visualization. *The Visual Computer*, 13:370–390, 1996.
- [PFA11] Pfaffelmoser, T., Reitingger, M., and Westermann, R. Visualizing the positional and geometrical variability of isosurfaces in uncertain scalar fields. In *Computer Graphics Forum*, volume 30, pages 951–960. Wiley Online Library, 2011.
- [PFA12] Pfaffelmoser, T., Mihai, M., and Westermann, R. Probability distributions for gradient orientations in uncertain 3d scalar fields. Technical report, Technische Universität München, 2012.
- [POE11a] Pöthkow, K., and Hege, H.-C. Positional uncertainty of isocontours: Condition analysis and probabilistic measures. *IEEE Transactions on Visualization and Computer Graphics*, 17(10):1393 – 1406, 2011.
- [POE11b] Pöthkow, K., Weber, B., and Hege, H.-C. Probabilistic marching cubes. *Computer Graphics Forum*, 30(3):931 – 940, 2011.
- [POE13a] Pöthkow, K., and Hege, H.-C. Nonparametric models for uncertainty visualization. *Computer Graphics Forum*, 32(3):131 – 140, 2013.
- [POE13b] Pöthkow, K., Petz, C., and Hege, H.-C. Approximate level-crossing probabilities for interactive visualization of uncertain isocontours. *International Journal for Uncertainty Quantification*, 3(2):101 – 117, 2013.
- [RAS06] Rasmussen, C.E., and Williams, C. *Gaus-*

- sian Processes for Machine Learning. MIT Press, 2006.
- [SCH12] Schlegel, S, Korn, N., and Scheuermann, G. On the interpolation of data with normally distributed uncertainty for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2305–2314, 2012.
- [SCH13] Schlegel, S, Böttinger, M., Hlawitschka, M, and Scheuermann, G. Determining and visualizing potential sources of floods. In *EuroVis Workshop on Visualisation in Environmental Sciences*, Leipzig, 2013.
- [SCH15] Schlegel, S., Goldau, M., and Scheuermann, G. Interactive GPU-based visualization of scalar data with gaussian distributed uncertainty. In David Bommes, Tobias Ritschel, and Thomas Schultz, editors, *VMV 2015 - Vision, Modeling and Visualization*. Eurographics Association, 2015.
- [SIL] Sillmann, J., and Roeckner, E. Indices for extreme events in projections of anthropogenic climate change. *Climatic Change*, 86(1-2):83–104, 2008.