A Hybrid Approach to Improving Rainfall Forecasts

High-resolution rainfall forecasting has important benefits, such as enabling flood prediction, yet little progress has been made toward developing an effective strategy. The hybrid approach presented here combines weather physics, statistics, and artificial neural networks.

High-resolution rainfall forecasting—the ability to predict rainfall intensities at enhanced resolution in space and time—is one of the most difficult, but also most useful, problems in applied meteorology and hydrology. Such forecasting could help scientists predict floods and manage water resource operations. Despite its well-known benefits,[1] however, scientists have made little progress over the years.[2,3] Although better understanding of atmospheric physics and faster computing speeds have resulted in improved numerical models of the weather,[4] the weather system's chaotic nature and the lack of knowledge of rainfall physics at higher resolutions[2] limit the predictability of rainfall. Statistical extrapolations of radar measurements at enhanced resolutions are useful for short-term (about one- to two-hour lead time) distributed forecasts but have limited use at longer lead times.[3]

The complexity of the problem and the potential benefits have led researchers to explore statistical and mathematical models, parameterized physics combined with Kalman filters, Gaussian mixtures, and even multifractal models. Investigative studies have demonstrated that neither available physics nor traditional statistical models, nor even data-mining tools such as neural networks, can improve forecasts by themselves.[1] Previous results using neural networks have been mixed.[5–7]

In this article, I present a hybrid strategy for high-resolution rainfall forecasting combining weather physics, traditional statistics, and a neural network-based approach.[1] The strategy is able to draw on all available information, account for and use aspects of the domain physics that are better understood, and exploit the strengths of the available data-dictated tools.

1521-9615/02/$17.00 © 2002 IEEE

Auroop R. Ganguly
Massachusetts Institute of Technology

Problem setting

The best archived rainfall information is available as six-hour, 48-km Numerical Weather Prediction (NWP) model (the *Eta model*) outputs and hourly Nexrad radar at 4×4 km.[2,4] At six hours or more, we achieve the best rainfall forecasts with numerical-model outputs at lower spatial resolutions, with or without statistical postprocessing. For one-hour or less lead time, extrapolation gives the best forecasts at the higher radar resolutions. Several applications require high-resolution forecasts at one to six hours;[3,8] both radar extrapolation and numerical model outputs can be useful here.[8] This article focuses on improving rainfall forecasts at high resolutions (4×4 km and hourly) with one- to six-hour lead times.

Existing rainfall forecast tools

Previous researchers have attempted to use different sets of tools to improve high-resolution rainfall forecasts, with limited success.

Process physics

Researchers can develop mathematical models for the weather system's basic physics, so they can forecast atmospheric variables such as temperature, pressure, or instability indices with reasonable accuracy.[1] Rainfall physics, on the other hand, operates at several scales in space and time, making rainfall forecasts less accurate and uniform. Several factors contribute to this variability:

- Large-scale air mass movements, rainfall cell translation and evolution, small-scale convection, topographic effects, and even raindrop evaporation can all affect observed rainfall at higher resolutions.
- Rainfall is not a state variable of the weather system, but needs to be derived. However, the generation process suffers from thresholds and intermittence.
- Cloud formation and convection depend on aerosol concentrations, which are difficult to measure, and cloud microphysics, which is not well understood.

Higher-resolution rainfall processes are superimposed on the results of lower-resolution and larger-scale effects. For example, small-scale structural changes are often superimposed on the translation of larger scale rainfall cells.[3] This translation, or *advection*, in turn occurs on top of the larger-scale movements and weather pattern changes. Some researchers, such as Mircea Grecu and Witold Krajewski, believe that advection is the only physical process that improves rainfall forecasts at higher resolutions, as the true physics of the evolving distributed structure is not well-known and cannot be accurately modeled from data.[7] Others have tried to model structural changes using mathematical or statistical models, data-mining techniques such as neural networks, and conceptual methods that use simplistic parameters to capture the process physics.

Numerical weather models capture lower-resolution rainfall physics and generate grid-averaged forecasts of rainfall. Data assimilation techniques can improve model initializations or parameters using higher-resolution radar observations, but alone they cannot capture higher-resolution physical processes. Statistical post-processing improves numerical model outputs at lower resolutions by comparing them with observations. For example, using radar observations aggregated in space and time and an autoregressive model, I statistically correct for the errors in the lower-resolution grid-averaged rainfall forecasts produced by a numerical model.[1]

As Grecu and Krajewski point out, rainfall predictability decays at enhanced resolutions, complicating high-resolution forecast generation.[7] We could potentially disaggregate error-corrected rainfall data at lower resolutions to higher resolutions. Rainfall variability and localized processes such as convection can complicate the disaggregation process, however. Multifractal models can achieve this disaggregation for simulation purposes but have limited application in forecasting.[9]

Using model outputs that correspond to atmospheric variables often improves disaggregation results.[9] The US National Weather Service obtains rainfall forecasts at specific locations within a grid by linearly regressing with grid average outputs (for rainfall and other atmospheric variables) from numerical weather models. NWS studies indicate that grid-averaged rainfall is not the most important predictor for rainfall at a specific point,[2] underscoring rainfall's spatial variability as well as the significance of atmospheric model output information. Hydrologists have developed a class of physically based models that could estimate parameters using either observed quantities or numerical model outputs (or both). The physics is overly parameterized, however, applicable only in very specific situations, and the parameters are difficult to estimate.[3]

Statistical models

Statistical techniques for rainfall forecasting continue to evolve. One approach uses observations to model, analyze, and forecast time series and space–time phenomena. Another estimates and updates the parameters of physically based formulations or mathematical cell evolution and decay models in real time with new observations, for example, with discrete Kalman filter formulations. Statistical time series methods, such as the Multivariate Autoregressive Integrated Moving Average (Marima) and state space formulations, improve forecasts for point- or basin-averaged rainfall.[1] Complex statistical or mathematical models, however, are suboptimal for high-resolution rainfall forecasts in space and time.[3] Some researchers have suggested combining information from models or measurement sources using Bayesian statistics and human–computer interaction.

As the UK Meteorological Office's recently operational Nimrod system demonstrates, statistical combinations of radar extrapolation and numerical weather model outputs yield the best results for one- to six-hour forecasts.[8]

Artificial neural networks

Researchers have attempted to improve rainfall forecasts using artificial neural networks, which could use information from remote sensors like radar and satellites.[1,7] Artificial neural networks are complex data-dictated tools that act as universal function approximators and converge faster than traditional approximators. Rather than replace numerical weather models, which have detailed physics, the researchers postprocess model results. Researchers have often combined information from satellites, numerical weather models, and other sources using neural networks to get average rainfall forecasts over larger areas or to directly obtain the discharge at river basin outlets.

Well-designed neural networks complement and outperform traditional forecasting methods.

One technique, which outperformed NWS-generated operational forecasts, uses the results from a numerical weather model as input to a neural network that generates the average and maximum rainfall at a Texas river basin.[1] In addition to neural networks, researchers have used Marima, state space, and Bayesian or likelihood-based formulations to improve rainfall forecasts at aggregate scales or at specific locations. In a comparison of techniques (neural network, *k*-nearest neighbor, and autoregressive models) for basin-averaged rainfall forecasts, Elena Toth and colleagues found that the neural-network-based method outperformed the rest.[6] Elsewhere I attempt to quantify the applicability of neural networks to rainfall forecasting at aggregate and distributed scales.[1]

Techniques based on neural networks have shown some promise for estimating and forecasting in recent years. In the Santa Fe Time Series Prediction and Analysis Competition, some of the best forecasts were based on neural network methods.[10] A need for caution exists, however: neural networks are not applicable to all situations. Chris Chatfield illustrates how easy it is to misuse these tools,[11] and the Sante Fe competition showed that poor neural network techniques can give misleading results.

Despite some success in forecasting rainfall, neural networks have had minimal success modeling high-resolution rainfall structure. Keith Smith and Geoff Austin indicate that complex methods do not improve high-resolution rainfall forecasts and often perform worse than simpler techniques.[3] Some researchers use the rainfall intensities at all pixels simultaneously as the outputs of a large neural network, and the corresponding intensities at lagged time-steps as the inputs.[1] Such methods increase the dimensionality of the underlying function approximation. Another approach is to forecast a few specific rainfall field characteristics. Grecu and Krajewski, for example, used spatial correlation to estimate rainfall velocity scales and a neural network-based time series to forecast them.[7] They then used the velocity forecasts to translate, or advect, the rainfall maps. Neural networks did not significantly improve the forecasts in that research effort, however.

Elsewhere I describe several experiments that modeled the change in rainfall intensity over time (that is, the rainfall evolution) at any pixel as a function of neighboring pixel intensities and the large-scale atmospheric state (as represented by the numerical weather model outputs for atmospheric variables).[1]

Benefits of neural networks

While researchers recognize the value of neural networks for pattern recognition (classification and regression) problems, they sometimes debate their applicability in forecasting applications,[10,11] and the neural network community has been moving toward traditional statistics for theoretical rigor and for validation of results.[12] However, well-designed neural networks complement and outperform traditional forecasting methods.[5,10] In particular, the statistical community recognizes multilayer perceptrons, or MLPs as "especially valuable because ... the complexity of the model [can be changed] from a simple parametric model to a highly flexible, nonparametric model."[12] Jerome Connor and colleagues demonstrate how neural networks can model the nonlinear form of the traditional autoregressive moving average (ARMA) or nonlinear ARMA (NARMA) model.[13]

Neural networks add value to rainfall forecasting in three areas:

- Neural networks utilize the information content of numerical model outputs by regressing against the observed rainfall. In the winter, the variability explained by the linear and neural network-based strategies is comparable, indicating a linear or quasilinear relationship. In addition, the model outputs corresponding to atmospheric variables do not result in improvements. In the summer, neural networks improve over linear methods, and atmospheric outputs contain information over and above the grid average rainfall forecast.
- Using numerical model atmospheric outputs and neighboring pixel radar measurements as neural network inputs can significantly improve forecasts for short (three- to four-hour) lead times, in the summer, and in terms of distributed structures. For longer lead times and in terms of aggregate measures of accuracy, however, these components do not improve the overall model performance.
- Combining a simple disaggregation method with a complex neural network-based strategy on a pixel-by-pixel basis, and on spatially aggregate scales, outperforms both methods alone, as well as current forecasting technologies.

New hybrid model

Earlier investigative studies confirm that we can best model and forecast certain complex phenomena such as precipitation by decomposing the problem into component processes and using weather and rainfall physics and data-dictated tools such as neural networks or traditional statistics where they are best suited.[1] In other words, the studies suggest a hybrid-model approach using the most applicable tool for each component.

Hybrid models have been used in several other areas related to forecasting and optimization. The meteorological community has used hybrid models that statistically combine low-resolution rainfall forecasts from numerical weather models with high-resolution radar extrapolation.[3] Brian Golding's Nimrod model forecasts rainfall at high resolutions for one- to six-hour lead times.[8] It essentially combines radar advection with rainfall forecasts from numerical weather models using relative weights, and improves forecasts over both.

There are two primary directions in which we can improve Golding's approach:

- We can disaggregate in space and time the low-resolution rainfall forecasts from the numerical weather model using higher-resolution radar extrapolation.
- Using process physics or data-dictated techniques, we can model localized processes (such as convection) that are known to dominate in the summer and significantly influence the high-resolution rainfall structure.

We can best model complex phenomena by decomposing the problem into component processes.

I selectively use neural network-based techniques and traditional statistical methods—a design based on insights from statistical data analysis and an understanding of the domain physics. My hybrid model consists of four components, which must be performed in sequence.[1]

Radar extrapolation

Radar extrapolation uses remotely sensed radar measurements and produces high-resolution forecasts through advection. I obtain advection velocity scales at each time step through 2D correlation in space between the current and the lagged rainfall maps. A single exponential smoothing-based time-series formulation then obtains velocity scale forecasts, which I use to translate the rainfall maps. This extrapolation strategy does not consider any low-resolution physics or changes in high-resolution structure.

Large-scale physics

Large-scale physics disaggregates the outputs for rainfall obtained from the (large-scale) physics-based numerical weather models. I obtain disaggregation in time (six hours to one hour) by linear interpolation and error-correct the grid-averaged rainfall by comparing it with aggregated radar observations using an autoregressive time-series formulation. I then disaggregate the grid-averaged and error-corrected rainfall in space by scaling the high-resolution extrapolation forecasts to reflect the grid-averaged mean rainfall.

Overall, this disaggregation strategy blends the high-resolution radar information, the corresponding advection component, and the large-scale physics from the numerical weather model. The strategy improves upon Golding's.[8] If the disaggregation strategy used all available information and domain knowledge, the overall approach could terminate here. However, this strategy does not consider rainfall physics at higher resolutions, including localized processes like convection, and does not use the numerical model output information that corresponds to atmospheric variables. Convection effects should dominate in the summer, and they tend to be more transient (compared to large-scale physics) and more nonlinear. On the other hand, nonlinear convective processes are difficult to model from data or process physics. Thus, depending on the specific forecast scenario, this simple disaggregation strategy and more complex methods optimized to handle localized processes could outperform each other.

Localized evolution

Localized evolution attempts to approximate rainfall physics at high resolutions. Data analysis and domain knowledge indicate that the evolution of rainfall intensity at each pixel is a function of the neighboring pixels' current intensity and the large-scale atmospheric state (represented by the lower-resolution numerical model outputs). I performed this function approximation using neural networks, specifically MLP, under the implicit assumption that a functional form dictating this evolution exists in space and time. I used each pixel in space over all training times to calibrate the MLP, and then used MLP to approximate the time-invariant evolution function. After randomly dividing the pixels into training and cross-validation, roughly in a 2:1 ratio, I combined Bayesian and ensemble techniques with neural networks to obtain confidence bounds with the forecasts at each point in space and for each lead time. Details on this strategy are published elsewhere.[1]

The localized evolution component should improve forecasts as long as the implicit assumptions are valid—for example, convective processes dominate in the summer, last for a few hours, and account for localized effects. Thus I expected the neural network-based strategy to produce better forecasts in the summer for shorter lead times, and to be more predominant in terms of forecast skills at higher resolutions. At longer lead times and in the winter, the neural network-based model's validity might not hold, and these complex models could perform worse than the disaggregation strategy alone. The performance of the simple disaggregation strategy could be better in terms of the aggregate skills, as the neural network-based strategy would focus on the localized structure. The balance between the low-resolution aggregate skills and the high-resolution distributed skills (as quantified by the conditional bias and the normalized root mean squared errors, for example) is similar to the bias-variance tradeoff in calibration problems. In fact, there seemed to be distinct domains where the simple disaggregation and the complex neural network-based strategies were applicable. This motivated the residual structures component.

Residual structures

The residual structures component combines the results from the large-scale physics and localized evolution components using neural networks at high resolutions (radar pixels) and aggregate errors at low resolutions (NWP grids). The inputs to the neural network were the results of the two individual components (less the results of extrapolation alone); the output was the observed rainfall less the results of extrapolation alone. At aggregate scales, I achieved the combination using a weighing scheme, where the aggregate error statistics at each lead time determined the weights. I scaled the spatially distributed results of the pixel-by-pixel combination with the results of the aggregate combination to attain the overall combined rainfall forecast at high resolutions in space and time.

I present detailed descriptions of the overall hybrid strategy, the individual components, and the statistical and neural network algorithms elsewhere.[1]

Model skills and performance analyses

Because skills for space–time forecasts should be viewed from a variety of perspectives, I use different qualitative and quantitative measures. I use results from persistence (where the forecast is assumed to be exactly equal to the current observation at each radar pixel) as a baseline (for comparison purposes). I apply the four individual components of my hybrid strategy in sequence:

1. Radar extrapolation is the starting point of the proposed strategy.
2. Large-scale physics combines low-resolu-

tion rainfall from NWP with radar extrapolation through a simple disaggregation strategy.
3. Localized evolution is conditioned on radar extrapolation and large-scale physics, and uses MLP for function approximation.
4. Residual structures combine the results of simple disaggregation and localized evolution.

Current models for rainfall forecasting at high resolutions include persistence, radar extrapolation, and combinations of low-resolution rainfall forecasts from numerical weather prediction models with radar extrapolation. Thus, the results shown here not only demonstrate the relative performance of the individual components of the proposed hybrid strategy but also show improvement over current techniques. Earlier work presents detailed results and performance analysis, both in terms of the most likely forecasts and confidence bounds.[1]

Table 1 shows performance in terms of the average distributed (high-resolution) skills for three rainfall events (two in summer, one in winter) in and around Topeka, Kansas, quantified by the inverse of the normalized root mean squared errors (1/NMSE). I performed normalization by dividing the RMSE with the standard deviation of the verification data. Radar pixels with nonzero rainfall were considered in the calculations. A zero NMSE indicates a perfect forecast. For a stationary process, an NMSE of unity indicates that the forecast is no better than the process mean; however, rainfall is not stationary.[1] Radar extrapolation marginally improves over persistence. Large-scale physics also results in marginal improvement. The neural network-based localized evolution component results in significant improvement up to three- to four-hour lead times but makes the skill go down thereafter. This is not unexpected, as neural network-based techniques tend to perform better than simpler tools when the complex functional forms they approximate remain valid, but worse in other situations. The residual structures strategy, which combines the individual component results, improves over any of these components alone.

Table 2 shows performance in terms of the average aggregate skills for six rainfall events (three in summer, three in winter) in southwest Oklahoma, quantified by the bias, or error, in the mean forecasts aggregated to 48 km and six hours. Radar extrapolation fails to improve over persistence. Large-scale physics results in significant improvement. The neural network-based localized evolution component, however, results in less accurate forecasts. There are two reasons for this:

Table 1. Average performance of the hybrid strategy's components for three storms in Topeka, Kansas. Larger values indicate better forecast skills.

	Lead time (in hours)				
	1	2	3	4	5
Baseline persistence	0.820	0.820	0.700	0.690	0.605
Component					
Radar extrapolation	0.845	0.842	0.745	0.710	0.645
Large-scale physics	0.865	0.843	0.685	0.700	0.690
Localized evolution	1.170	0.950	0.875	0.715	0.540
Residual structures	1.170	0.970	0.900	0.825	0.735

Table 2. Aggregate errors (in mm, for six hours and 48 km) of the hybrid strategy's components for six Oklahoma storms. Larger values indicate lower forecast skills.

	Errors (mm)
Baseline persistence	3.12
Component	
Radar extrapolation	3.30
Large-scale physics	2.64
Localized evolution	4.68
Residual structures	2.05

- The neural network-based model's performance decays significantly at longer lead times (more than three to four hours).
- The neural network attempts to model the details of the localized structure, thus improving the NMSE at shorter lead times. This improvement often hurts performance in terms of the overall bias, however. As mentioned earlier, this is analogous to the bias-variance tradeoff in model calibration or training.

The residual-structures component improves over all the individual components in terms of skill measures. Combined, Tables 1 and 2 demonstrate that the proposed hybrid approach improves overall forecasts.

Figure 1 shows the relative performance of the hybrid modeling strategy components in terms of their empirical cumulative distribution function (CDF) and the corresponding Kolmogorov-Smirnov (K-S) statistic. These are shown for a typical storm event in Oklahoma for all pixels,

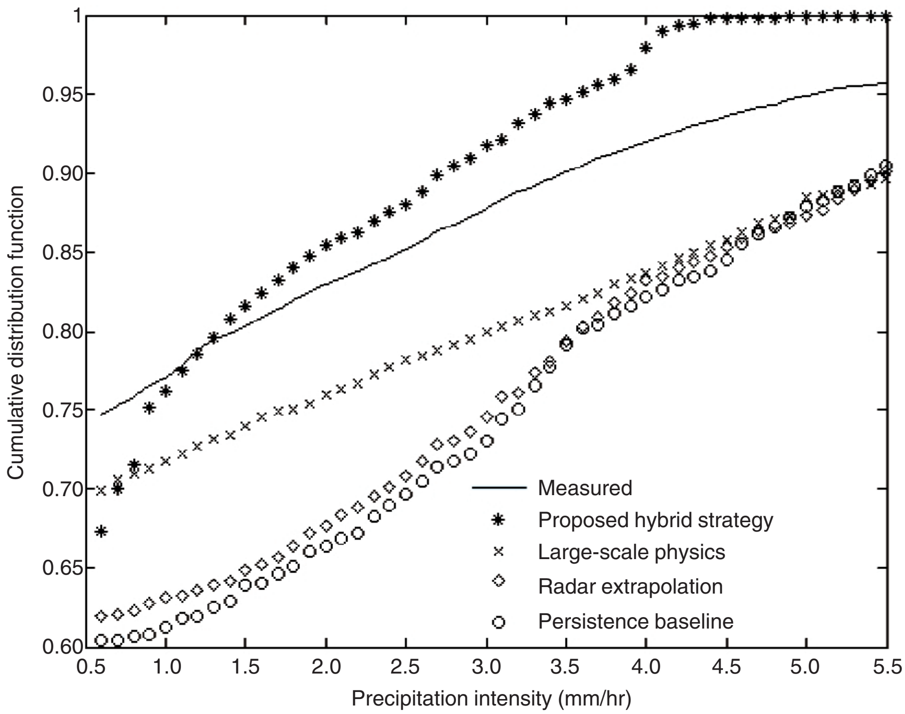

Figure 1. Empirical cumulative distribution function for measured rainfall (radar measurements corrected with ground-based measurements) in Oklahoma and the corresponding empirical CDF from the various rainfall forecasting strategies. The radar pixels for all (one- to six-hour) forecast lead times were used for the computations. The K-S statistics (highest absolute deviation between measured and forecast CDF) are persistence, 0.1693; radar extrapolation, 0.1574; large-scale physics, 0.0851; and proposed hybrid model, 0.0726. The hybrid model performs best both qualitatively (visual comparison of CDF) and quantitatively (K-S statistic).

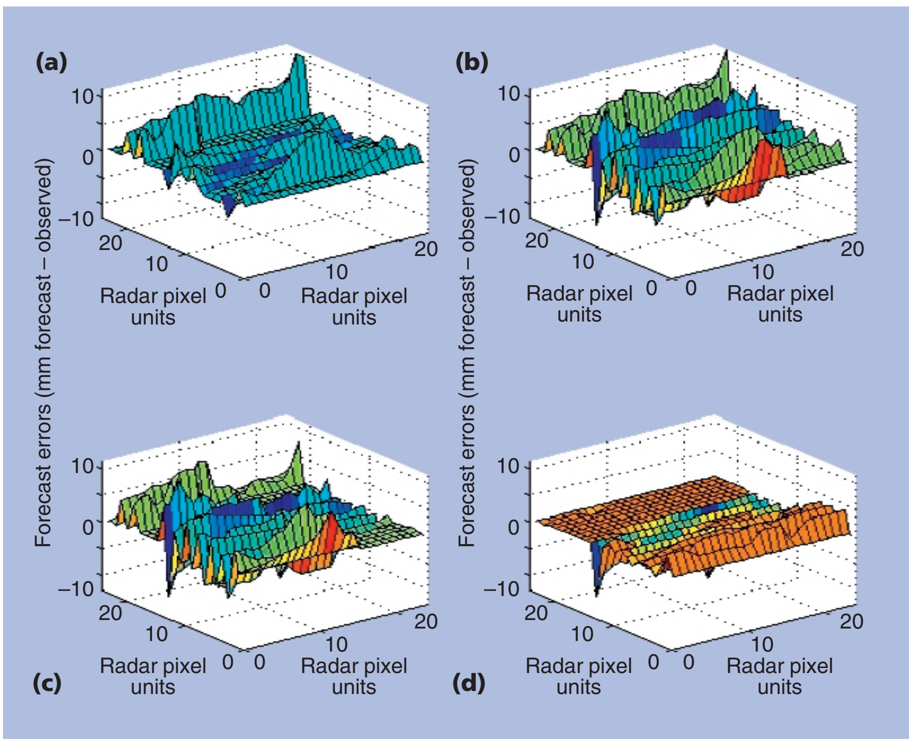

Figure 2. Precipitation forecast error (mm/hour) surface plots for an Oklahoma storm (*X* and *Y* axes in pixel units of 4 × 4 km; lead time is one hour). (a) Persistence retains significant error structures. (b) One-hour advection, or radar extrapolation, does not visually improve over persistence, but the error structure changes. (c) Combining radar extrapolation with error-corrected rainfall forecasts from the numerical weather model seems to improve the spatial aggregate skills marginally, but does not improve distributed skills. (d) The proposed hybrid strategy significantly improves over the existing methods, both in terms of the distributed and the aggregate errors structures.

and for one- to six-hour forecast lead times. I compared four strategies: persistence, radar extrapolation, a combination of extrapolation and numerical weather model outputs (that is, large-scale physics), and the proposed hybrid model. These measures indicate how well the strategies predict rainfall ranges and thresholds (considering all pixels and all lead times at once).

Radar extrapolation performs marginally better than persistence, while large-scale physics results in some improvement. The empirical CDF qualitatively demonstrates that the proposed hybrid strategy most closely matches the observations. The K-S statistic confirms this by quantifying the largest absolute deviations between the measured and the forecast CDF.

Figure 2 shows the forecast errors using surface plots for an Oklahoma storm event at a one-hour lead time. Once again, I compare four strategies: persistence, radar extrapolation, large-scale physics, and the proposed hybrid model. Lack of structures in the error surface plots indicate better forecasts. Visually, the improvement from radar extrapolation is not too apparent; the improvement from large-scale physics is marginal; and the improvement from the hybrid strategy is significant.

There are other ways to combine traditional time-series and forecasting concepts with neural networks. For example, you can isolate and model longer term trends and certain short-term or time-varying components using standard statistical and time-series techniques, and model nonlinear time-invariant components using neural networks.

Future research must explore efficient algorithms for data-mining and hybrid solutions in operational settings, as well as the possibility of exploiting more information sources. One emerging area is the use of satellite information to classify storms. Other applications of forecasting in space and time—for example, in such areas as the earth sciences, fluid dynamics, and astrophysics—could also be explored.

Acknowledgments

I thank my thesis supervisor, Rafael L. Bras, and Dennis McLaughlin, Elfatih Eltahir, Dara Entekhabi, Amar Gupta, Shafiqul Islam, and Ana Barros for their helpful suggestions. The anonymous reviewers provided valuable comments.

References

1. A.R. Ganguly, "Distributed Quantitative Precipitation Forecasts Combining Information from Radar and Numerical Weather Prediction Model Outputs," doctoral dissertation, Civil and Environmental Eng. Dept., Mass. Inst. of Technology, Cambridge, Mass., 2002.
2. M.S. Antolik, "An Overview of the National Weather Service's Centralized Statistical Quantitative Precipitation Forecasts," *J. Hydrology*, vol. 239 (1–4), Dec. 2000, pp. 306–337.
3. K.T. Smith and G.L. Austin, "Nowcasting Precipitation—a Proposal for a Way Forward," *J. Hydrology,* vol. 239 (1–4), Dec. 2000, pp. 34–45.
4. T.L. Black, "The Use of NMC Mesoscale Eta Model: Description and Forecast Examples," *Weather and Forecasting*, vol. 9, June 1994, pp. 265–278.
5. D.W. McCann, "A Neural-Network Based Short-Term Forecast of Significant Thunderstorms," *Weather and Forecasting,* vol. 7, Sept. 1992, pp. 525–532.
6. E. Toth, A. Brath, and A. Montanari, "Comparison of Short-Term Rainfall Prediction Models for Real-Time Flood Forecasting," *J. Hydrology,* vol. 239 (1–4), Dec. 2000, pp. 132–147.
7. M. Grecu and W.F. Krajewski, "A Large-Sample Investigation of Statistical Procedures for Short-Term Quantitative Precipitation Forecasting," *J. Hydrology*, vol. 239 (1–4), Dec. 2000, pp. 69–84.
8. B.W. Golding, "Quantitative Precipitation Forecasting in the UK," *J. Hydrology,* vol. 239 (1–4), Dec. 2000, pp. 286–305.
9. S. Perica and E. Foufoula-Georgiou, "A Model for Multiscale Disaggregation of Spatial Rainfall Based on Coupling Meteorological and Scaling Descriptions," *J.Geophysical Research*, vol. 101, no. D21, 1996, pp. 26, 347, 361.
10. A. Weigend and N. Gershenfeld, *Time Series Prediction: Forecasting the Future and Understanding the Past*, Addison-Wesley, Reading, Mass., 1994.
11. C. Chatfield, "Neural Networks: Forecasting Breakthrough or Passing Fad?" *Int'l J. Forecasting,* vol. 9, no. 1, 1993, pp. 1–3.
12. A. Flexer, "Connectionists and Statisticians, Friends or Foes?" *From Natural to Artificial Neural Computation, Proc. Int'l Workshop Artificial Neural Networks*, J. Mira and F. Sandoval, eds., Lecture Notes in Computer Science 930, Springer-Verlag, New York, 1995, pp. 454–461.
13. J.T. Connor, R.D. Martin, and L.E. Atlas, "Recurrent Neural Networks and Robust Time Series Prediction," *IEEE Trans. Neural Networks,* vol. 5, no. 2, Mar. 1994, pp. 240–254.

Auroop R. Ganguly is a research associate in both the Civil and Environmental Engineering Department and the Sloan School of Management at Massachusetts Institute of Technology. At MIT-CEE, his research focuses on rainfall and flood forecasting, modeling complex systems, and developing prediction strategies that combine information from multiple sources and blend process physics with data-dictated tools. At Sloan, he works in the areas of business forecasting, time series analysis, demand planning, inventory optimization, and the use of statistical and IT tools such as OLAP and data mining for business applications. He is also the product manager for Oracle's Demand Planning in the Analytic Solutions (for e-business applications) product development group. He received a BTech (Hons.) in civil engineering from the Indian Institute of Technology, Kharagpur, and a PhD in civil and environmental engineering from MIT. He is a member of the IEEE, Sigma XI, the American Geophysical Union, and the American Meteorological Society. Contact him at 16 Royal Crest Dr., #4; Nashua, NH 03060; auroop@ ieee.org.