A Markov Chain Model of Tornadic Activity

Caren Marzban*
and
Peter Guttorp

National Research Center for Statistics and the Environment
University of Washington, Seattle, WA 98195

Abstract

Tornadic activity within the continental United States is modeled as a stochastic process. It is shown that the occurrence of a tornado on a given day is affected mostly by that of the previous day. In other words, the process appears to display no memory beyond one day. As such, the process is a Markov chain. The performance of the model in predicting tornadoes is assessed and is shown to be marginally superior to a model based on climatological forecasts.

---

*Permanent address: National Severe Storms Laboratory, Norman, OK 73069
http://www.nhn.ou.edu/~marzban

Figure 1: The number of matches to the phrase "Markov Chain" in the on-line journals of the American Meteorological Society, for every year from 1970 to 2000.

1 Introduction

Markov Chains have steadily been gaining popularity in meteorological circles. An exact phrase search for "Markov Chain" on the American Meterological Society's on-line journals suggests an increasing trend over the past 30 years (Figure 1). The interest, however, appears to be focused on precipitation, since of the 165 matches, 112 deal with precipitation, while the remaining few cover a somewhat wide range of applications. A few examples of the former are the works of Gates and Tong (1976), Hughes and Guttorp (1995), Katz (1997), Stern (1982), and Valdez and Young (1985). An example of the latter is the work of Lakshman (2001) in the context of data compression. It is interesting, though, that the number of matches drops to zero if the word "tornado" is included in the search.

The topic of Markov Chains is an instance of the larger topic of stochastic modelling or stochastic processes (Guttorp 1996, Kao 1996). A stochastic process is simply a collection of random variables. The closing price of a stock, or the amount of precipitation, at different times t are examples of stochastic processes. Given the generality of the definition of a stochastic process, it is not surprising that stochastic modeling has found wide-spread use. Although the price of a stock and the amount of precip-

itation are continuous variables, from a pedogogical view it is far simpler to consider a stochastic process whose states are discrete. In fact, most stochastic modeling text books begin their discussion with a process with binary states. This specialization does not hinder the application of the corresponding methodology, for many natural phenomena are inherently binary. The occurrence or non-occurrence of tornadoes on a given day is but one example.

In this article, the methodology of stochastic modeling will be reviewed through an application, and a model will be developed that can be employed to predict the occurrence of tornadoes.

2   Markov Chains

Here, the subject will be treated only in its most elementary form, sufficiently deep to introduce the general concepts, terminology, and notation. For additional details see (Guttorp 1995 - Ch. 2, Kao 1997). Further, generalities are avoided and the topic is considered within the specific context of tornadic activity.

Consider the binary variable, x_i, with $i = 1, 2, 3, ..., 366$, representing the occurrence (1) or nonoccurrence (0) of a tornado on day i of the year, somewhere within the continental United States. Given data (i.e., a time series for x_i) it is possible to estimate some important conditional probabilities; for example, $P(x_{i+1} = 1 | x_i = 1)$, which represents the probability of having a tornado on the $(i + 1)^{st}$ day, given that there is a tornado on the i^{th} day. This probability, together with $P(x_{i+1} = 0 | x_i = 0)$, $P(x_{i+1} = 0 | x_i = 1)$, and $P(x_{i+1} = 1 | x_i = 0)$, constitute the four elements of what is referred to as the 1-step transition matrix. Similarly, an element of the m-step transition matrix is $P(x_{i+m} = 1 | x_i = 1)$, the probability of a tornado m days after the i^{th} day, if the i^{th} is tornadic.

All of the above probabilities are said to be 1^{st}-order transition probabilities. A second order transition probability takes into account the state of x_i not just on one previous day, but on two past days. For example, one element of the second order, 1-step, transition matrix is $P(x_{i+2} = 1 | x_{i+1} = 1, x_i = 0)$ - the probability of a tornado

today, if there was a tornado yesterday, but there was no tornado on the day before yesterday. The generalization to higher orders and steps is straightforward.

Theoretically, the probability of a tornado on a given day could depend on the values of x_i for *all* past days. However, in practice such processes are rare. Most natural processes appear to have a "memory" that does not extend beyond one or two days (Guttorp 1995). In fact, the case with a 1-day memory is sufficiently ubiquitous to have been given a name: (Discrete-time) Markov chain. In other words, a process (i.e., a sequence of x_i's) is said to be a (discrete-time) Markov chain if the probability of x_i taking some (0/1) value, given the entire past, depends only on the value of x_{i-1}. If this probability is the same for all i, then the Markov chain is said to be (time) homogeneous.

A great many theoretical results have been developed for homogeneous Markov chains - adding to their popularity; however, it is important to check their assumptions prior to the application of the results. For the case at hand, the probability of a tornadic day clearly depends on the day of the year, for there is such a thing as a tornado season. As such, the process has nonhomogeneous transition probabilities, robbing one of the opportunity to apply the aforementioned results. And as to whether the process of tornadic activity is a Markov chain at all, in the next section it will be shown that it is (i.e., that the extent of the memory is one day).

3   Method

The data at hand span 49 years, and it can be represented as a matrix with 366 rows and 49 columns. Each element of the matrix, x_i^I, where $i = 1, 2, 3, ...366$, and $I = 1, 2, 3, ..., 49$, is a 1 or a 0, corresponding to the occurrence or nonoccurrence of some tornado somewhere in the continental United States. [1]

In order to determine the order of the process, several hypotheses must be tested against one another. One hypothesis is the one where the occurrence of a tornado on a given day is independent of that on any other day. For such a 0^{th}-order process one

---

[1] In this paper, the last day of a non-leap year is treated as a nontornadic day. This is not expected to adversely affect the results.

would have $P(x_{i+1} = 1|x_i = 0) = P(x_{i+1} = 1|x_i = 1) = P(x_i = 1)$ for every i. In other words, the process would have no memory. This hypothesis can be tested against the 1^{st} order process wherein the occurrence of a tornado on a given day depends on that of the previous day. Adopting the former as the null hypothesis, H_0, and the latter as the alternative hypothesis, H_1, if H_0 cannot be rejected, then there is no need to examine any other hypothesis, for the process can be said to be 0^{th}-order (or "random"). Otherwise, one must proceed to test H_1 against the alternative hypothesis that the process is second order (i.e., that the occurrence of a tornado on a given day depends on the occurrence/nonoccurrence of a tornado over the last two days). This procedure continues until a null hypothesis cannot be rejected, thereby establishing the most likely order of the process. A more sophisticated approach to estimating the order is to use an information criterion (Guttorp 1995, sec. 2.8).

Although the order of the process can be anticipated from a comparison of the empirical transition probabilities estimated from data, a formal starting point for the hypothesis testing is based on the likelihood (or the probability of observing the data):

$$L(\mathcal{P}) = \Pr(x_1^1, x_2^1, ..., x_{366}^1, x_1^2, x_2^2, ..., x_{366}^2, ..., x_1^{49}, x_2^{49}, ...x_{366}^{49}; \mathcal{P}),$$

where $\mathcal{P}$ is the vector of parameters (to be estimated) and each x_i^j is a 0 or a 1. The 49 years are assumed to be independent, which means that overall, large-scale trends are assumed nonexistent. Then,

$$L(\mathcal{P}) = \prod_{I=1}^{49} \Pr(x_1^I, x_2^I, ..., x_{366}^I; \mathcal{P}).$$

The likelihood for the 0-th order process is then

$$L_0 = \prod_{I=1}^{49} \prod_{i=1}^{366} \Pr(x_i^I; \mathcal{P}) = \prod_{i=1}^{366} [P_0(i)]^{n_0(i)} [P_1(i)]^{n_1(i)}, \tag{1}$$

where $n_0(i)$, and $n_1(i)$ are the number of years for which the i^{th} day is nontornadic and tornadic, respectively. Note that $n_0(i) + n_1(i) = 49$. $P_1(i) \equiv P(x_i = 1)$, and $P_0(i) \equiv P(x_i = 0) = 1 - P_1(i)$ are the marginal (climatological) probabilities. The maximum likelihood estimate for the parameter $P_1(i)$ is $n_1(i)/[n_0(i) + n_1(i)]$.

For a 1^{st} order process, the likelihood can be written as [2]

$$L_1 = \prod_{i=2}^{366} [P_{00}(i)]^{n_{00}(i)} [P_{01}(i)]^{n_{01}(i)} [P_{10}(i)]^{n_{10}(i)} [P_{11}(i)]^{n_{11}(i)} \times [P_0(i)]^{n_0(1)} [P_1(i)]^{n_1(1)}, \quad (2)$$

where $n_{00}(i)$ is the number of years for which the i^{th} day and the next day are both nontornadic. Similarly, $n_{01}(i)$ is the number of years for which the i^{th} day is nontornadic but the following day is tornadic. Etc. For simplicity, the following notation is introduced: $P_{kl}(i) = P(x_{i+1} = l | x_i = k)$. For example, $P_{01}(i)$ is the probability that a nontornadic day will be followed by a tornadic day; it can be estimated by $n_{01}(i)/n_0(i)$. Similarly, $P_{00}(i)$, $P_{10}(i)$, and $P_{11}(i)$ can be estimated by $n_{00}(i)/n_0(i)$, $n_{10}(i)/n_1(i)$, and $n_{11}(i)/n_1(i)$, respectively. Note that $P_{00}(i) + P_{01}(i) = 1$ and $P_{10}(i) + P_{11}(i) = 1$. As such, it is sufficient to examine only two of the four probabilities, say, $P_{01}(i)$ and $P_{11}(i)$, the probability of a tornado tomorrow, whether today is nontornadic or tornadic, respectively.

The likelihood for the 2^{nd} order process is

$$\begin{aligned} L_2 = \quad &\prod_{i=3}^{366} \quad [P_{000}(i)]^{n_{000}(i)} [P_{001}(i)]^{n_{001}(i)} [P_{010}(i)]^{n_{010}(i)} [P_{100}(i)]^{n_{100}(i)} \quad (3)\\ &\times \quad [P_{011}(i)]^{n_{011}(i)} [P_{101}(i)]^{n_{101}(i)} [P_{110}(i)]^{n_{110}(i)} [P_{111}(i)]^{n_{111}(i)} \\ &\times \quad [P_{00}(i)]^{n_{00}(2)} [P_{01}(i)]^{n_{01}(2)} [P_{10}(i)]^{n_{10}(2)} [P_{11}(i)]^{n_{11}(2)} \\ &\times \quad [P_0^(i)]^{n_0(1)} [P_1(i)]^{n_1(1)}, \end{aligned}$$

where, for instance, $P_{001}(i)$ is the probability that two nontornadic days will be followed by a tornadic day. The maximum likelihood estimate for $P_{001}(i)$ is $n_{001}(i)/[n_{000}(i) + n_{001}(i)]$. The other probabilities can be computed by a straightforward generalization. Since $P_{kl0}(i) + P_{kl1}(i) = 1$ only four of the eight probabilities are independent, and so, it is sufficient to consider only $P_{001}(i)$, $P_{011}(i)$, $P_{101}(i)$, and $P_{111}(i)$. The reason these four probabilities are selected is that the last day in all of them is tornadic.

Armed with the likelihood for each hypothesis, one can test for the most likely order of the process. To that end, one can compare L_1 with L_0, and if L_1 is significantly larger than L_0, then one can compare L_2 with L_1. The comparisons of the likelihoods is best made via a likelihood ratio test (Guttorp 1995, sec. 2.7). After the order of the

---

[2] This expression for the likelihood follows from repeated conditioning on the previous day; for example, $P(x_1, x_2, ..., x_{366}) = P(x_{366}|x_1, x_2, ..., x_{365}) P(x_1, x_2, ..., x_{365}) = P(x_{366}|x_{365}) P(x_1, x_2, ..., x_{365}) = P(x_{366}|x_{365}) P(x_{365}|x_1, x_2, ..., x_{364}) P(x_1, x_2, ..., x_{364}) = P(x_{366}|x_{365}) P(x_{365}|x_{364}) P(x_1, x_2, ..., x_{364})...$.

Figure 2: The probability of tornado tomorrow, if there is no tornado today (P_{01}), if there is a tornado today (P_{11}), and the climatological probability of tornado (P_1) for every day of the year.

process has been established, say M, then the corresponding transition probabilities can be employed for forecasting tornadic activity M or more days into the future.

4   Results

As mentioned previously, all of these probabilities (parameters) can be estimated from the data. Figure 2 shows P_{01}, P_{11}, and P_1 for every day of the year[3]. The most prominent feature in these plots is the bell shape of all three curves. This is simply a consequence of the increase in tornadic activity in the Spring and Summer Seasons, and implies that the process is not homogeneous.

The next important feature is the separation between the P_{01} and the P_{11} curves. It can be seen that the two probabilities are distinct throughout the year. If one had found $P_{00} = P_{01}$, then one would have concluded that the process underlying tornadic

---

[3]For purely visual purposes the curves have been smoothed by performing a running average with a window size of 21 days. The error bars are the standard errors representing the variation of the data within a window.

Figure 3: a) The 2-step and b) the 3-step transition probabilities.

activity is random in the sense that the occurrence of a tornado on any given day would be independent of tornadic activity on the previous day. However, as seen from Figure 2, this is not the case; whether or not there will be a tornado tomorrow depends on whether there is a tornado today (P_{11}) or not (P_{01}).

The independence of the process is slowly recovered as one examines higher-step transition probabilities. For example, if one computes the probability of tornado tomorrow, given the occurrence or nonoccurrence of a tornado yesterday, then one arrives at the 2-step probabilities plotted in Figure 3a. Figure 3b displays the 3-step transition probabilities. For the 2-step case, the difference between P_{01} and P_{11} is mostly within the error bars. This means that the 2-day forecasts are mostly climatological. Two exceptions are the intervals in the middle of the year and at the end of the year where the curves do display some separation beyond the error bars. For the 3-step case only the separation in the central range remains. Note that the displayed error bars are the standard errors. As such, even in the exceptional regions the difference between the curves is nonsignificant if one considers two (or higher) standard errors (i.e., approximately 95% confidence). In short, the memory of the process does not appear to extend beyond one day.

The comparison of P_{01} and P_{11} can also be made through a scatterplot of the

Figure 4: The scatterplot of P_{01} and P_{11} for the a) 1-step, b) 2-step, and c) 3-step transition probabilities.

two. Figures 4 show the scatterplots for the 1-step, 2-step, and 3-step models. The linear correlation coefficients are $r = 0.72$, $r = 0.82$, and $r = 0.82$, respectively. The scatterplots for larger steps are not illuminating; however, the linear correlation coefficient between P_{01} and P_{11} for as many as 30 steps is displayed in Figure 5. It can be seen that r reaches a constant value (within the error bars) at step 2. As such, one can expect that P_{01} and P_{11} are distinct only for the 1-step process, and they become identical (and equal to P_1) for $m \geq 2$.

In order to assess the 2-day memory of the process it is not sufficient to examine the 2-step process. One must examine the 2^{nd} order transition probabilities: P_{001}, P_{011}, P_{101}, and P_{111}. For example, a comparison of P_{001} and P_{011} can asses the dependence of tomorrow's tornadic activity on that of today's and yesterday's activity. Figure 6 displays all the pairwise comparisons. For clarity, the error bars are not shown.

The 2^{nd}-order transition probabilities generally confirm the results of 1^{st}-order, 1-step and 2-step analysis, namely that the memory of the process does not appear to extend beyond one day. Specifically, the difference between P_{001} and P_{011} is more pronounced than the difference between P_{001} and P_{101}. Similarly, the difference between P_{101} and P_{111} is more pronounced than the difference between P_{011} and P_{111}.

Figure 5: The linear correlation coefficient between P_{01} and P_{11} for the 1-step through the 30-step models.

Further qualitative features can be seen. For example (Figures 6a and b), the difference between P_{001} and P_{011} is almost negligible in the central region of the graph. As such one may conclude that the occurrence of a tornado during the most active time of the year is independent of tornadic activity on the previous day, if there is no tornado 2 days ago. By contrast, if there is a tornado 2 days ago, then the process does display a 1-day memory throughout the year.

The central region of Figure 6c suggests yet another interesting feature. During the most active time of the year P_{001} is larger than P_{101}. In other words, if there is no tornado today, it is more likely to have a tornado tomorrow if there was no tornado yesterday than if there was one. This phenomenon can be explained if one thinks of the occurrence of a tornado as an energy-releasing event that follows a phase of energy build-up. As such, the occurrence of a tornado on a given day might drain the system of the required energy to form another tornado on the next day. But after 2 days of energy build-up, there may be sufficient energy to release in the form of another tornado. This explanation, of course, is only a hypothesis that must be tested against alternative explanations.

10

Figure 6: A pairwise comparison of the four 2^{nd} order transition probabilities.

It is also interesting to note the overall right-shift of P_{111} with respect to P_{011} (Figure 6d). First, and not surprisingly, this suggests that for the latter 2/3 of the year it is more likely to have three consecutive tornadic days than a 3-day string that begins with a nontornadic day. However, for the earlier part of the year, it is more likely for a tornadic day to be followed by another tornadic day, if the previous day is nontornadic.

As mentioned previously, the analysis of the transition probabilities is somewhat qualitative, albeit illuminating. The likelihood ratio test can be employed to further quantify the findings. For instance, it can be shown that $2\log(L_1/L_0)$ and $2\log(L_2/L_1)$ have a chi-squared distribution with 1×365 and $2 \times 364 = 728$ degrees of freedom, respectively (Guttorp 1995) [4]. For the current data set, one finds $2\log(L_1/L_0) = 1106.07$ and $2\log(L_2/L_1) = 437.18$. With 365 and 728 degrees of freedom , respectively, the probability of having a chi-squared as large as 1106.07 and 437.18 is 0 and 1, respectively. Therefore, there is overwhelming evidence for rejecting H_0 in favor of H_1, while there is no evidence for rejecting H_1 in favor of H_2. This implies that the underlying process is most likely a 1^{st}-order process, namely a (nonhomogeneous) Markov chain.

5 Performance

One can assess the utility of the transition probabilities by computing and comparing the performance of several of the models mentioned above in predicting tornadoes. The question of performance (or forecast verification) is a complex one that will be more fully dealt with in a separate article. Here, only one measure of performance will be adopted to gauge the overall performance of the various models. The measure is the Receivers Operating Characteristic (ROC) curve (Masters 1993, Marzban 2000). Briefly, it is a parametric plot of the probability of detection (or hit rate) versus the false alarm rate (or false positive) as a threshold on a probabilistic forecast is varied

---

[4]The number of parameters in L_0, L_1, and L_2, as defined in eqs 1-3 is, respectively, 1×366, $2 \times (366 - 1) + 1 = 731$, and $4 \times (366 - 2) + 2 + 1 = 1459$. The pre-factors (1, 2, and 4) are the number of free parameters in P_i, P_{ij} and P_{ijk}, respectively. The degrees of freedom corresponding to the likelihood ratios are then given by the difference in the number of free parameters in the likelihoods, $731 - 366 = 365$, and $1459 - 731 == 728$.

Figure 7: The normalized frequency of tornado forecasts for nontornadic (0) and tornadic (1) days.

from zero to one.

First, note that P_{01}, and P_{11} together consitute a forecast for the occurrence of tornado one day into the future. It is then natural to ask if the forecasts discriminate between nontornadic and tornadic days at all. Figure 7 shows the normalized frequency of the forecasts for nontornadic (0) and tornadic (1) days. Evidently, the forecasts do discriminate between the two events. This figure provides a qualitative measure of the performance of the 1^{st}-order model in terms of its discriminatory capability.

Figure 8 shows the ROC curves for several models. The amount of "bowing" of an ROC curve is a measure of performance. A diagonal ROC curve represents random forecasts, and the more the ROC curve bows above the diagonal, the higher the performance. A scalar measure of performance based on an ROC curve is the area under the curve; however, the curve itself is a more faithful representation of performance due to its multi-dimensionality.

Note from Figure 8 that the climatological model (i.e., 0^{th}-order) has an ROC curve that overlaps that of the 1^{st}-order, 2-step model. This is consistent with what

Figure 8: The ROC curves (i.e., performance) of several models.

we found above - that $P_{01} = P_{11} = P_1$ for the 2-step process. The next best performance is obtained by the 1^{st}-order, 1-step process. It is important to emphasize that the curve lies above and beyond the one corresponding to climatological forecasts. The 2^{nd}-order, 1-step model has a performance that is comparable and only mildly superior to the 1^{st}-order, 1-step model. It is possible to quantify the statistical significance of the difference between the two; however, as shown in the previous section, the difference between the 2^{nd}-order and the 1^{st}-order model is not statistically significant.

6 Conclusion and Discussion

A time series spanning 49 years is employed to estimate the transition probabilities of the process underlying tornadic activity. It is found that the most likely process is a 1^{st}-order, 1-step process, with time-varying probabilities, namely a nonhomogeneous Markov chain. The corresponding transition probabilities generate a 1-day forecast for the occurrence of one or more tornadoes somewhere in the continental United States. The performance of the forecasts is found to be marginally (but statistically significantly) superior to climatological forecasts.

14

Figure 9: Scatterplot of the theoretical versus the empirical transition probabilties for a 2-step chain.

As mentioned previously, many theoretical results have been proven for Markov (i.e., 1^{st}-order) processes. Therefore, a consistency check on the order of the process is possible. For example, one well-known theoretical result is the Chapman-Kolmogorov Theorem (Guttorp 1995), according to which the matrix of m-step transition probabilities is equal to the m^{th} power of the matrix of 1-step transition probabilities. Figure 9 shows the scatterplot of the theoretical transition probabilities (i.e., as derived from the Chapman-Kolmogorov theorem) versus the empirical ones for a 2-step chain. The correlation coefficients for P_{01} and P_{11} are both 0.998. The same quantities for a 4-step process (not shown) show equally high correlation between empirical and theoretical estimates of the transition probabilities. Evidently, the theoretical transition probabilities are consistent with the empirical ones, providing another confirmation of the 1^{st}-order nature of the process.

The results found here are preliminary in that several issues must be examined further. These issues also point into future research. For instance, given that there is no physical reason why the occurrence of a tornado at two geographically distinct points should be correlated, one may ask how it is that the Markov Chain model developed

here can forecast tornadoes with any skill at all. The answer, of cousre, is that it cannot; the displayed skill shown above is most likley from the skill in forecasting tornadoes that *are* spatially correlated. In fact, it is likely that a geographic partitioning of the above analysis will lead to higher skill forecasts for each region. [5]

Another issue that has not been examined is the time-dependence of the performance of the model. In other words, in light of the exceptional regions in Figure 3, it is conceivable that the performance of the model depends on the time of the year. This would require producing an ROC-diagram for every period of interest (e.g., every day of the year, or every season). This type of time-partitioning may also lead to higher skill forecasts for certain times of the year.

7 Acknowledgements

8 References

Valdez, M. P., K. C. Young, 1985: Number Fluxes in Equilibrium Raindrop Populations: A Markov Chain Analysis. *Journal of the Atmospheric Sciences*, **42** pp. 1024-1036.

Stern, R. D., 1982: Computing a Probability Distribution for the Start of the Rains from a Markov Chain Model for Precipitation. *Journal of Applied Meteorology*,

---

[5]Some interesting geographic distributions of tornadoes for the years 1980-1994 are animated at http://www.nssl.noaa.gov/hazard/tanim8094/tornanim8094.html . Other climatological facts can be found at http://www.nssl.noaa.gov/hazard.

21, 420-422.

Gates, P., H. Tong, 1976: On Markov Chain Modeling to Some Weather Data. *Journal of Applied Meteorology*, **15**, 1145-1151.

Hughes, J. P., P. Guttorp, 1994: Incorporating Spatial Dependence and Atmospheric Data in a Model of Precipitation. *Journal of Applied Meteorology*, **33**, 1503-1515.

Katz, R. W., 1974: Computing Probabilities Associated with the Markov Chain Model for Precipitation. *Journal of Applied Meteorology*, **13**, 953-954.

Guttorp, P., 1995: *Stochastic Modeling of Scientific Data*. Chapman & Hall, pp. 372.

Kao, E. P. C., 1997: *An Introduction to Stochastic Processes*. Duxbury Press, pp. 438.

Lakshmanan, V., 2001: Lossless coding and compression of radar reflectivity data. Radar Conference, American Meteorological Society. In press.

Masters, T., 1993: *Practical Neural Network Recipes in C++*, Academic Press, pp. 493 .

Marzban, C., 2000: A neural network for tornado diagnosis. *Neural Computing and Applications*, **9** (2), 133-141.