Hydrologic Model Assessment from Automated Spatial Pattern Comparison Techniques

S.R. Wealands, R.B. Grayson and J.P. Walker

Department of Civil and Environmental Engineering, The University of Melbourne, Australia,
srweal@civenv.unimelb.edu.au

Abstract: Spatial patterns of hydrological response are observed from field studies and remote sensing, and predicted by distributed hydrologic models. However, techniques for using spatial patterns to assess model performance are presently very simple. Rather than make assessments based on the spatial patterns, spatially explicit predictions are generally assessed using point observations or integrated responses. The problem of comparing two spatial patterns is not unique to hydrology, and has been researched in many other disciplines (e.g. pattern recognition, image processing). Hydrologists currently rely on either simple quantitative comparisons or the human brain to assess the spatial similarity of predicted and observed patterns. While visual methods work well at local scales and with simple patterns, larger scale and more complex patterns demand alternative techniques. Automation of new quantitative techniques should prove applicable to hydrological applications. This paper discusses various methods currently used for comparing spatial patterns in hydrology. A review of promising techniques from other disciplines with potential application to hydrologically relevant spatial patterns is then presented. Some preliminary assessment of these techniques is undertaken and future directions for this research are outlined.

Keywords: *Spatial pattern; Model assessment; Pattern comparison; Hydrology; Distributed models*

1. INTRODUCTION

The need for techniques to compare predicted and observed spatial patterns stems from the widespread use of spatial data in hydrological modelling. Moreover, there has been an extensive development of spatially explicit (otherwise known as distributed) hydrologic models in the past 15 years. There are two main reasons for this development: 1) the increased availability of geospatial data and computing power, and 2) a desire for natural resource managers to have spatial estimates of hydrological attributes (Grayson and Blöschl 2000, Grayson et al. 2002). While the models and data sources have improved considerably over the recent past, techniques for assessing how well the hydrological models predict spatial patterns have remained relatively dormant.

Spatially explicit predictions from hydrological models are used for planning, monitoring and resolving issues in the natural environment. The complexity of these models and the large number of parameters used results in uncertainty in the model results (Beven 2001). As such, calibration, and subsequently testing of hydrologic models is essential to reduce the uncertainty and assess the user confidence of the predictions. This is most commonly done using comparison of observed and predicted point values (e.g. soil moisture) or integrated responses (e.g. runoff). To provide a comprehensive assessment of spatially explicit models, it is important to make our appraisal using spatially explicit data. Although data collection is a major undertaking, the lack of current model testing using spatial data is unacceptable (Beven 2001). By including spatial pattern comparisons in model assessments, we will "improve the confidence with which we can claim our models do indeed represent the right processes and get the right answers for the right reasons" (Grayson and Blöschl 2000).

Geospatial data of hydrologically relevant spatial patterns suitable for model assessment are available from a wide range of data sources. Remote sensing offers numerous measurements at various scales that can prove useful for assessing model performance (Schultz and Engman 2000). Intensive field campaigns, in which sampling is done at a regular spacing, also produce detailed spatial patterns of hydrologic attributes (e.g. Western et al. 1999). Such data sources make comprehensive testing of the spatial patterns produced by hydrologic models feasible. However, due to limited research in this area, the current techniques for making comparisons with

spatially explicit hydrologic data are simplistic (Grayson and Blöschl 2000).

Outside the field of hydrology there are many applications in which two spatial patterns (or images) are compared and their similarity assessed. Pattern recognition, pattern matching, computer vision, content-based image retrieval (CBIR) and medical imaging all require identification of features within an image and then comparison against a template or database of other images. These fields are active areas of research and have been used for identifying appropriate techniques for assessing the spatial predictions from hydrologic models.

This paper begins by concisely explaining spatial patterns in hydrology and the methods currently used for working with spatial patterns. A broad review of techniques for identifying features or aspects of spatial patterns follows. The common approaches for automatic comparison of two spatial patterns are then discussed, including comments about their suitability for hydrologic application. The future directions for this research are then briefly outlined.

2. SPATIAL PATTERNS IN HYDROLOGY

In a hydrological sense, a spatial pattern is any image or surface showing the spatial distribution of an attribute, especially where there is a degree of organisation (as opposed to the spatial pattern being random) (Grayson and Blöschl 2000). Such spatial patterns include: a RADAR field showing rainfall intensity; a Landsat TM image showing a vegetation index; an elevation model of a study catchment; or a dense array of soil moisture samples. The raster data format is most common for representing spatial patterns in hydrology and is the format is assumed in this paper. To compare raster data, patterns should have the same pixel size and extents. Figure 1 shows three spatial patterns of soil moisture resulting from the work at Tarrawarra in Australia (Western et al. 1999) and will be used later in discussing methods of pattern identification and comparison.

Observed and predicted spatial patterns come from different sources. Hydrological models make spatially explicit predictions, while interpolation of sparse point samples produces a predicted pattern. Observed patterns are obtained directly from sensors or measurements and have a 'scale triplet' defining their spacing, extent and support (Blöschl and Sivapalan 1995). 'Directly measured' patterns are obtained from exhaustive ground sampling with a spatial reference at each sample (e.g. Western et al. 1999). It is important to collect sufficient samples to capture the pattern of interest (Grayson et al. 2002).

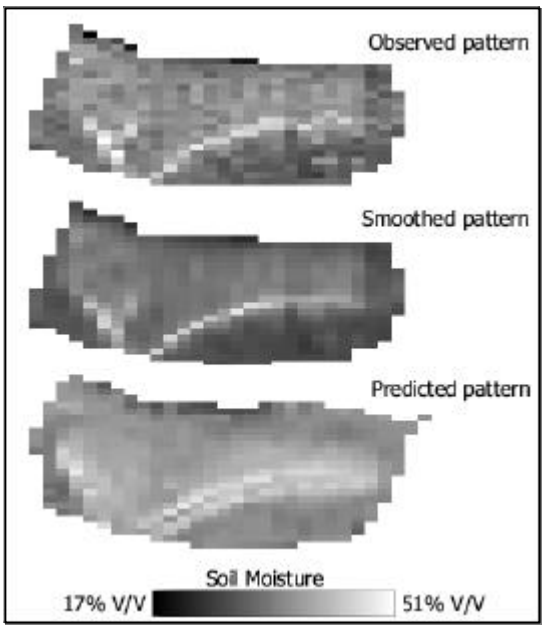

Figure 1. Spatial patterns of soil moisture from Tarrawarra (adapted from Grayson et al. 2002).

'Remotely sensed' patterns are observed from an aircraft or satellite. Observations of electromagnetic radiation and response from the Earth's surface produce a raster image covering a large area at an instant in time. Remote sensing measurements provide an integrated response from a regular area (the pixel size). Schmugge et al. (2002) gives examples of remotely sensed patterns used in hydrology. Remotely sensed patterns require significant processing through interpretation models before the measured values are useful, however use of the broader pattern may prove sufficient for assessing model performance.

A 'surrogate pattern' is one that shows a degree of correlation to the spatial pattern of interest, however it is not an actual measurement of that attribute. One of the most common surrogates is terrain (e.g. Western et al. 2001), which is often used due to the widespread availability of digital elevation models. Terrain has been used as a surrogate for solar radiation exposure, soil properties, vegetation distributions and many others (e.g. Wilson and Gallant 2000). All patterns are effectively surrogates, as there is rarely a direct measurement of the attribute of interest (Grayson et al. 2002). Surrogate patterns are useful when the attribute of interest is difficult to collect, as they provide a means of assessing spatial predictions (albeit with greater uncertainty).

With all patterns the 'level of measurement' can vary (Clarke 1999). A pattern of snow cover could be binary (presence/absence) or ordinal

(e.g. 0, 1 or 2 denoting amount of snow cover). It is common for a trade off to occur between this level of measurement (or information content) and the spatial detail of mapping (Grayson et al. 2002).

3. IDENTIFYING AND COMPARING SPATIAL PATTERNS IN HYDROLOGY

To utilize a spatial pattern for model assessment it is important to either identify the prominent features or enhance important aspects of the pattern. This can be done visually, with statistics or via image processing techniques. These three techniques have all been used in hydrology. However, only the first two are used in common practice. Comparing two spatial patterns to obtain a measure of similarity and an explanation of the differences has been attempted in a simplistic manner, however there are many weaknesses with current techniques. Visual comparisons are the most common, but some pixel-based techniques are widely used. Grayson and Blöschl (2000) provide details of many case studies in which spatial patterns are identified, enhanced and compared. The approaches used in various case studies are now reviewed.

3.1. Visual identification and comparison

The typical approach that hydrologists use when presented with an observed and predicted pattern is visual identification (a plausibility check), followed by comparison (a qualitative similarity assessment). The human brain intuitively identifies a hierarchy of similarities between the patterns, starting with global differences and then noticing the finer (or local) similarities (Power et al. 2001). By incorporating background knowledge and expertise (and sometimes personal biases) into the assessment of the patterns, humans are able to explain what is seen. Computer systems cannot replicate this analysis (Pavlidis 2003). However, they can provide unbiased, quantitative estimates of similarity. Bow (2002) recognises the main disadvantage with visual human interpretation is the training and labour required. The human vision system also struggles to process all the information existing in large, complex patterns, thus making similarity assessment difficult in those instances (Grayson and Blöschl 2000).

3.2. Pixel based comparisons

To compare observed and predicted spatial patterns, hydrologists commonly compute the difference between the coincident pixels on each pattern. Residual maps are used to show differences visually for detecting consistent over- or under-prediction in a particular location. The residuals are also used for computing the root mean square error (RMSE) as a measure of 'matching' between the patterns. Scatter plots showing observed against predicted pixel values are produced to look for consistent trends between the patterns, while coefficients of determination and efficiency are computed to quantify how well the patterns match one another. These techniques are commonly used to detect correlation between hydrologic attributes and surrogates (such as elevation, land use, aspect or soil type). Pixel based comparisons do not carry any spatial information, like connectivity, through the analysis and are sensitive to spatial shifts, which can lead to large 'false' residuals (Power et al. 2001). A similar problem exists in time series comparisons of hydrographs, where a 'phase shift' has a large effect on comparative statistics.

3.3. Statistical identification

The statistics of a spatial pattern provide a succinct, quantitative summary of the pattern. Many global statistics, such as the mean, variance, histogram and RMSE are summaries, yet they describe nothing of the spatial structure (arrangement of features) or distribution. Geostatistical methods can capture some features of the spatial distribution of a pattern, but they do not capture the arrangement of features directly.

Western et al. (2001) used variograms to describe soil moisture patterns and compare pattern characteristics over time. They then looked further into characterising spatial connectivity within patterns using connectivity functions. Zepeda-Arce et al. (2000) devised statistics to represent the patterns of precipitation fields for verification of quantitative precipitation forecasts. A statistical approach, called the frequency scaling ratio (FSR), is used to provide information about the consistency of features (in this case, pixels with positive or negative anomalies) between observed and interpolated soil moisture fields (Thattai and Islam 2000). These statistical approaches provide quantitative measures of patterns, which can be compared numerically with other patterns.

3.4. Alternative techniques

Image processing techniques, such as principal components analysis (PCA), enable spatial patterns to be enhanced to improve pattern identification. PCA was applied to remotely measured soil moisture patterns in Grayson and Blöschl (2000) to identify features not apparent in

the original patterns. Transforming the original patterns into their principal components separated out patterns due to terrain, land use and soil moisture.

Grayson et al. (2002) investigated two approaches for detecting shifts between observed and predicted patterns. Extracting pixel values along a transect in each pattern allowed shifts in one dimension to be observed. For two dimensional shifts, an implementation of optical flow methods (Barron et al. 1994), called optimal local alignment (OLA), was used. This technique uses the cross correlation between regions of each pattern to create a map of likely shifts. While OLA does not provide a quantitative assessment of similarity, it is useful for understanding where patches of the patterns contain lateral shifts.

4. SPATIAL PATTERN IDENTIFICATION

Identification of a spatial pattern involves both the enhancement of the pattern and the segmentation (or classification) of interesting regions within the pattern. Enhancement is done to make interesting parts of the pattern more apparent or to provide alternative representations of the pattern (e.g. PCA mentioned in Section 3). Segmentation is undertaken to separate regions of interest from the background pattern, so that they can then be compared. Segmentation produces a similar result to image classification in remote sensing, whereby pixels are assigned into homogenous groups. This section presents ways of enhancing and segmenting spatial patterns.

Within pattern recognition, the Fourier and wavelet transformations (Chen and Bui 1999) are often used to obtain an alternate description of the pattern. These transforms change the representation of the pattern and enable calculation of descriptors that do not change value, whether the pattern is scaled or rotated (which is particularly useful for applications such as text recognition). It is then possible to compare patterns rapidly using the descriptor. For hydrology, these approaches are unlikely to be suitable, as the spatial patterns contain more noise and complexity (as compared to characters or more basic images). However, the computational efficiency of these transformations makes them of interest.

Segmentation of a spatial pattern is done to identify homogenous regions within the pattern. Enhancement plays a role in improving segmentation by clarifying features within the spatial pattern. Some common enhancement techniques include low or high pass filtering, smoothing (as shown in Figure 1), resampling and normalising pixel values within a pattern (Bow 2002).

In pattern recognition, segmentation is undertaken to extract regions that are then subject to object recognition, whereas in hydrology we want to segment a pattern so we have features to compare. Some patterns, like binary snow cover patterns, are already segmented into object and background, making them ready for comparison (Di Gesu and Starovoitov 1999). Obtaining a good segmentation is a difficult task and many techniques exist, although no method has proven to be equally good for all types of patterns (Pal and Pal 1993). A basic approach is to segment the pattern using a threshold value, such that pixels in the pattern with values higher than the threshold become the features of interest, while the remainder becomes the background. By thresholding the patterns shown in Figure 1 at a value of 45% volumetric soil moisture, it is possible to segment the original pattern into wet regions and the background. Western et al. (2001) used this approach to identify pixels exhibiting hydrologically significant connectivity. Thresholding makes no use of any spatial relationship between pixels and the threshold value generally needs to be user defined, making this unsuitable for automated segmentation (although adaptive thresholding has been researched) (Pal and Pal 1993).

To extract interesting (or salient) regions from the original pattern, it is necessary to define what is meant by salient. These may be any regions containing values above the mean, or they may be regions that are distinctly different from their surroundings. There is extensive research that looks at obtaining segmentation via clustering (Jain et al. 1999), which is the grouping of well-connected pixels with similar values. Clustering is a common approach to unsupervised learning or classification, making it suitable for automated applications. Most approaches to clustering require a priori knowledge of the number of clusters expected. Nonparametric clustering (Pauwels and Frederix 1999) is an alternative to traditional methods (such as K-means) that creates many smaller clusters and merges them later. This approach relies on careful assessment of cluster validity to ensure ‘correct’ clusters are created. Pauwels and Frederix (1999) present an approach to pattern segmentation that is suitable for automated applications. It focuses on obtaining a perceptually salient segmentation, using clusters to detect obvious regions within the pattern, by looking at alternate characteristics like shape and texture.

Clustering is a common approach to detecting groups of interesting pixels, however spatial

coherence is more strictly enforced using 'region growing' to segment (Matas and Kittler 1995). Starting from an initial seed pixel, regions are grown to encompass neighbouring pixels that meet certain criteria. For automated applications, region growing requires an adaptive approach to select the starting pixels and the homogeneity criteria (e.g. Chang and Li 1994). Another approach to segmentation is to use 'edges' to define regions of interest. An edge occurs where there is a discontinuity in pixel value (or intensity) (Bow 2002). This is a subjective measure and many algorithms exist for detecting edges. The edges detected ideally surround regions of similar value, however, edges are very sensitive to noise. A small gap in an edge boundary can permit dissimilar regions to be merged (Sharma 2001), resulting in false segmentation. Pattern enhancement and filtering can reduce edge related problems, and edge detection is common practice in many facets of pattern matching. For natural patterns, edges are less prominent and subsequently less stable features for defining interesting regions.

The numerous techniques discussed above are not used in hydrology and can offer improvements for identifying interesting regions in a spatial pattern. Segmentation by clustering (e.g. Pauwels and Frederix 1999) is the most promising technique for hydrology, as it detects regions that are naturally and spatially grouped. Other techniques could also be suitable, however their use may be compromised by noise within the patterns.

5. SPATIAL PATTERN COMPARISON

Comparison techniques for spatial patterns vary with the information content of the pattern. Binary patterns are compared differently to those that contain gray level (single value) or colour (multiple value) information (Di Gesu and Starovoitov 1999). Comparison is undertaken to produce a measure of how similar two patterns are using some algorithm to compare the regions or pixel values in each pattern. The following section identifies techniques used for producing a measure of similarity.

Tompa et al. (2000) made a modification to the standard mean square error (MSE) statistic. After identifying perceptually important parts of the pattern (using techniques discussed in Section 4), the MSE is weighted accordingly. Differences in more interesting regions increase the standard MSE, while differences present in the background are less influential. This idea of 'perceptually weighted' pattern comparison is promising for hydrology and applicable to binary and gray level patterns, providing an improvement on most existing global techniques. Experimentation has shown that global similarities are often too crude an approach for matching natural patterns (Pauwels and Frederix 1999). Di Gesu and Starovoitov (1999) have developed an image distance function (IDF) that combines global intensity with some structural information (e.g. the geometric distribution of 'local' pixel values) about the pattern. This function does a direct comparison of non-binary patterns, without having to calculate regions within the image. The experimental results show this approach is suitable for coarse similarity assessment. Some patterns can be simplified to a representative set of points (commonly done with fingerprints) or regions (as done with natural pattern segmentation). Using these simplified representations, a similarity measure is used to assess the resemblance of one aspects of a pattern to another (Veltkamp and Hagedoorn 1999). Similarity measures come in many varieties. However, they all aim to produce an assessment of similarity, regardless of whether the object of interest within the source pattern is scaled, rotated or shifted within the matched pattern.

A widely used measure for point and region matching is the Hausdorff metric (Veltkamp and Hagedoorn 1999), which is basically the minimum distance fit between two point sets that represent the object of interest. Many variations of this metric exist, some of which make it more reliable for noisy patterns (as dealt with in natural patterns). This approach is widely used for rapid detection of a known object within images, like in content-based image retrieval applications (where an image database is queried for a particular image). Similarity measures have been successfully applied to natural spatial patterns in ecology. Maps of species abundance and presence/absence have been compared to help select successful model results (Fewster and Buckland 2001), showing the versatility of these metrics to varied applications. Comparison of two segmented patterns, where the pattern has been separated into related regions, is useful to hydrologists. For comparison of individual region objects, Veltkamp and Hagedoorn (1999) suggest the 'area of overlap' as a similarity measure. If the two regions are identical they will overlap 100%, however slight differences could result in only 80% overlap. Using such a measure also allows one region to be translated and rotated, with the maximum overlap defining the optimal shift.

The task of comparing two patterns containing segmented regions is similar to the comparison of land use maps (or any other type of classified image). Land use maps are commonly compared

to detect land use change over time. The traditional approach has been pixel-by-pixel comparison, however a fuzzy pattern matching approach has been investigated (Power et al. 2001). Fuzzy maps are more appropriate for representing complex spatial patterns, as they allow pixels located at the edge of regions to be members of both bordering regions. Power et al. (2001) uses a polygon-by-polygon overlapping area comparison approach to assess the similarity of two fuzzy land use maps. They suggest future research into the comparison technique for incorporating shape and complexity.

6. FUTURE RESEARCH

The techniques associated with spatial pattern identification and comparison suggests many avenues for future research in hydrology. While some techniques appear more suited to automated implementation and for model assessment, others offer ways of representing the subtle differences between patterns. The spatial patterns encountered in hydrology are most likely to contain regions of interest, as opposed to specific objects (which is the focus of much work in pattern recognition). This research will experiment with various techniques, applying them to hydrologically relevant spatial patterns. The investigation will help evaluate the methods that are useful. With further research, effective techniques for automatically assessing the similarity of observed and modelled spatial patterns will be developed.

7. REFERENCES

Barron, J.L., S.S. Beauchemin and D.J. Fleet, On optical flow, 6th Int. Conf. on Artificial Intelligence and Information-Control Systems of Robots, Bratislava, Slovakia, September 12-16, 1994.

Beven, K.J., How far can we go with distributed hydrological modelling?, Hydrology and Earth Science Systems, 5(1), 1-12, 2001.

Blöschl, G. and M. Sivapalan, Scale issues in hydrological modeling - a review, Hydrological Processes, 9(3-4), 251-290, 1995.

Bow, S., Pattern recognition and image preprocessing, Marcel Dekker, 698 pp., New York, 2002.

Chang, Y.L. and X. Li, Adaptive image region-growing, IEEE Transactions on Image Processing, 3(1057-7149), 868-872, 1994.

Chen, G.Y. and T.D. Bui, Invariant Fourier-wavelet descriptor for pattern recognition, Pattern Recognition, 32(7), 1083-1088, 1999.

Clarke, K.C., Getting started with geographic information systems, Prentice Hall, 338 pp., New Jersey, USA, 1999.

Di Gesu, V. and V. Starovoitov, Distance-based functions for image comparison, Pattern Recognition Letters, 20(2), 207-214, 1999.

Fewster, R.M. and S.T. Buckland, Similarity indices for spatial ecological data, Biometrics, 57(2), 495-501, 2001.

Grayson, R. and G. Blöschl, Spatial Patterns in Catchment Hydrology: Observations and Modelling, Cambridge University Press, 404 pp., Cambridge, 2000.

Grayson, R.B., G. Blöschl, A.W. Western and T.A. McMahon, Advances in the use of observed spatial patterns of catchment hydrological response, Advances in Water Resources, 25(8-12), 1313-1334, 2002.

Jain, A.K., M.N. Murty and P.J. Flynn, Data clustering: a review, ACM Computing Surveys, 31(3), 264-323, 1999.

Matas, J. and J. Kittler, Spatial and feature space clustering: Applications in image analysis, 6th International Conference on Computer Analysis of Images and Patterns, Prague, Czech Republic, Sept 6-8, 1995.

Pal, N.R. and S.K. Pal, A review on image segmentation techniques, Pattern Recognition, 26(9), 1277-1294, 1993.

Pauwels, E.J. and G. Frederix, Finding salient regions in images - Nonparametric clustering for image segmentation and grouping, Computer Vision and Image Understanding, 75(1-2), 73-85, 1999.

Pavlidis, T., 36 years on the pattern recognition front, Pattern Recognition Letters, 24(1-3), 1-7, 2003.

Power, C., A. Simms and R. White, Hierarchical fuzzy pattern matching for the regional comparison of land use maps, International Journal of Geographical Information Science, 15(1), 77-100, 2001.

Schmugge, T.J., W.P. Kustas, J.C. Ritchie, T.J. Jackson and A. Rango, Remote sensing in hydrology, Advances in Water Resources, 25(8-12), 1367-1385, 2002.

Schultz, G.A. and E.T. Engman, Remote sensing in hydrology and water management, Springer, 483 pp., New York, 2000.

Sharma, M., Performance evaluation of image segmentation and texture extraction methods in scene analysis, Masters Thesis, University of Exeter, United Kingdom, 227 pp., 2001.

Thattai, D. and S. Islam, Spatial analysis of remotely sensed soil moisture data, Journal of Hydrologic Engineering, 5(4), 386-392, 2000.

Tompa, D., J. Morton and E. Jernigan, Perceptually based image comparison, International Conference on Image Processing, 2000.

Veltkamp, R.C. and M. Hagedoorn, State-of-the-art in shape matching, Technical Report UU-CS-1999-27, Utrecht University, the Netherlands, 1999.

Western, A.W., G. Blöschl and R.B. Grayson, Toward capturing hydrologically significant connectivity in spatial patterns, Water Resources Research, 37(1), 83-97, 2001.

Western, A.W., R.B. Grayson and T.R. Green, The Tarrawarra project: high resolution spatial measurement, modelling and analysis of soil moisture and hydrological response, Hydrological Processes, 13(5), 633-652, 1999.

Wilson, J.P. and J.C. Gallant, Terrain analysis: principles and applications, J. Wiley, 512 pp., New York, 2000.

Zepeda-Arce, J., E. Foufoula-Georgiou and K.K. Droegemeier, Space-time rainfall organization and its role in validating quantitative precipitation forecasts, JGR Atmospheres, 105(D8), 10129-10146, 2000.