

Bayesian hierarchical modelling: incorporating spatial information in water resources assessment and accounting

G.S. Chiu^a, E.A. Lehmann^b

^aCSIRO Mathematics, Informatics and Statistics, GPO Box 664, Canberra, ACT 2601, Australia

^bCSIRO Mathematics, Informatics and Statistics, Private Bag 5, Wembley, WA 6913, Australia
Email: grace.chiu@csiro.au

Abstract: Water resources assessment and accounting research requires making the best use of multiple sources of data for producing reliable accounting predictions. For any given quantity of interest, e.g. soil moisture, available sources of data could be directly observed *in situ*, derived indirectly from a remotely-sensed surrogate (e.g. brightness temperature) using retrieval models, or output from deterministic hydrological models. Whether it be model-data fusion or the evaluation of remotely-sensed products, multiple data sources for individual quantities of interest must be assimilated to optimize the use of available information. Popular data assimilation techniques focus heavily on assimilating time series at fixed geographical locations. Yet, the resulting predictions typically show unrealistically large variability over space. This suggests the need to assimilate spatial maps at fixed time points, in addition to time series at fixed locations. Spatial statistical models, constructed in a Bayesian hierarchical framework, offer an intuitive and unifying approach for this purpose. It cohesively utilizes multiple data sources by addressing the mismatch in both spatial and measurement scales between ground-based and remotely-sensed products; it can also spatially interpolate missing data due to the remote sensor's "blind spots." In this paper, we use the evaluation of the remotely-sensed soil moisture product, AMSR-E, to illustrate this statistical framework. Specifically, a single statistical model can incorporate both point-level *in situ* data (benchmark) and pixel-based AMSR-E data (product), as well as other related variables such as precipitation.

Keywords: Statistical modelling, spatial statistics, conditional autoregressive model, Bayesian inference, Markov chain Monte Carlo

1 INTRODUCTION

The Water Information Research and Development Alliance (WIRADA) conducts large-scale water resources research. In a recent WIRADA project entitled *Australian Water Resources Assessment* (AWRA), one of the primary objectives is the development of model-data fusion methods to blend deterministic model outputs with observed data (van Dijk, 2010a). For any given quantity of interest to AWRA (soil moisture, catchment runoff, etc.), available sources of data could be directly observed *in situ*, derived indirectly from a remotely-sensed surrogate (brightness temperature, etc.) using retrieval models, or output from deterministic hydrological models (van Dijk, 2010b). For the blending to be meaningful, a thorough understanding of the error structure of each blended component is essential. Thus, a related research problem under AWRA is the evaluation of remotely-sensed products, such as Vrije Universiteit Amsterdam's soil moisture product based on the Advanced Microwave Scanning Radiometer – Earth Observing System (VUA AMSR-E) (Owe *et al.*, 2008). That is, we wish to assess how well these surrogate-based AMSR-E soil moisture data (henceforth, simply AMSR-E) reflect the true level of on-ground soil moisture.

In situ measurements of soil moisture recorded with ground probes are often referred to as “ground truth,” a misnomer from a statistical standpoint. This is because measurements contain inherent uncertainty, even if the ground probe technology were deemed most accurate and precise for measuring soil moisture. Thus, the term “benchmark” is more appropriate for *in situ* soil moisture data. A quantitative evaluation of the remotely-sensed AMSR-E requires, among other procedures, a sensible comparison of the product to the benchmark. Such a task is not necessarily straightforward due to the mismatch in both spatial and measurement scales between ground-based and remotely-sensed products. To complicate matters, the remote sensor's “blind spots” (due to a combination of heavy vegetation, unfavourable weather conditions, instrumentation failure, etc.) often lead to swaths of missing data.

Evaluation methods from the literature (e.g. Draper *et al.*, 2009) typically address missing data and the mismatch in scales of observed data in a semi-qualitative fashion. Data assimilation techniques based on the Kalman filter (KF) are also used (Crow and Zhan, 2007) and conventionally applied to each pixel, thus focusing heavily on assimilating time series at fixed geographical locations. Yet, the resulting assimilated values typically show unrealistically large variability among adjacent pixels (e.g. Fig. 3 in Crow and Zhan, 2007). The 3D-KF (de Lannoy *et al.*, 2010; Reichle and Koster, 2003; Reichle *et al.*, 2002) induces spatial smoothing without a parametric spatial correlation model, although basic understanding of the physical phenomenon of interest can suggest the parametric structure for both spatial and temporal dependence.

Thus, we consider complementary modeling frameworks that assimilate spatially (using data from spatial neighbours to update state estimates) at fixed time points in addition to temporally at fixed locations. Statistical spatio-temporal modelling (Cressie and Wikle, 2011) offers an intuitive and unifying approach for this purpose. This statistical inferential approach simultaneously models multiple spatial time series, each being a temporal sequence of maps; a single model can address the mismatch in both spatial and measurement scales between ground-based and remotely-sensed products, without relying on qualitative scale realignment or linearity/Gaussian assumptions of the KF. The approach would yield a single model that can address various AWRA research problems comprehensively and simultaneously, including product evaluation, data assimilation, and model-data fusion.

Focus of this paper. Constructing and fine-tuning a unifying model to address the wide range of AWRA objectives is a laborious process. As a start, we focus on evaluating AMSR-E via statistical spatial modelling (Banerjee *et al.*, 2004). Thus, we assimilate soil moisture maps from the same static time frame. We employ Bayesian hierarchical modelling to incorporate both point-level *in situ* data (typically sparse) and pixel-based remotely-sensed data (mostly contiguous), as well as another related variable, precipitation (of a third spatial scale). As a bonus, this framework regards interpolation of missing data as part of the one-step statistical inference. Model implementation and data imputation are via Markov chain Monte Carlo (MCMC) with Gibbs sampling (Chiu and Lehmann, 2011). Hierarchical modelling and various MCMC sampling methods are reviewed by Gelfand and Smith (1990).

2 SOURCES OF SOIL MOISTURE AND RELATED DATA

The Murrumbidgee River Catchment (MRC) in eastern Australia is one of the main geographical regions of focus by AWRA. For our hierarchical spatial model, we include the following MRC quantities, directly




Figure 1. Data on AMSR-E (v , left) and transformed AWAP moving average (p , right) for the MRC on 2007-01-20 (horizontal and vertical axes denote longitude and latitude). Note the different intensity scales. *In situ* OzNet stations are marked by “+”. AMSR-E pixels with missing data are outlined in dark blue; intensity shown for these pixels are imputed according to the hierarchical model in Section 3.

comparable to those used by Draper *et al.* (2009) and Crow and Zhan (2007) for their product evaluation. A detailed description of these data sources appears in Section 3 of Henderson *et al.* (2011).

1. *Point-level soil moisture data from 38 MRC sites in the OzNet Hydrological Monitoring Network* (<http://oznet.org.au>) (Fig. 1). Of the available variables, we consider ground soil moisture readings (% volume) at the shallowest depth (0–5 or 0–8cm, depending on the station) at time-of-day closest to 01:30 (01:20 or 01:30, depending on the station), so as to minimize incompatibility between *in situ* and descending pass remotely-sensed data. (Ascending pass data are not considered.) Availability of data varies between the years 2001 and 2010, depending on the station.
2. *The gridded AWAP precipitation product (Jones et al., 2009) for daily rainfall.* The AWAP time frame spans the latter part of 2002 to the end of 2009 (taken to coincide with AMSR-E data). AWAP pixels are 0.05° by 0.05° . We consider AWAP data for the MRC only.
3. *The remotely-sensed AMSR-E product for soil moisture* (Fig. 1). We consider (VUA) AMSR-E data over 0.25° by 0.25° pixels, available for the entire Australian mainland except for “blind spots.” Edges of these pixels align perfectly with AWAP pixel edges. Thus, each AMSR-E pixel contains exactly 25 AWAP pixels. For spatial modelling, we consider C-band data from the descending pass over the MRC only, comprising $M=312$ AMSR-E pixels.

3 A HIERARCHICAL STATISTICAL SPATIAL MODEL

For a preliminary model, we consider precipitation to be the sole driver of soil moisture. Inspired by McMillan *et al.* (2010), we construct our fixed-time spatial model to comprise the following statements.

Relating ground data and small-pixel truth. Let s_{mr} be the r th AWAP pixel inside the m th AMSR-E pixel for $r=1, \dots, 25$ and $m=1, \dots, M$; $\psi_{mr} \equiv \psi(s_{mr})$ the true (latent) soil moisture spatially averaged across s_{mr} ; and g_{mrk} the point-level OzNet soil moisture reading at the k th ground station inside s_{mr} for $k=1, \dots, K_{mr}$ (defined only for AWAP pixels with $K_{mr} \geq 1$). Then, a measurement model is

$$g_{mrk} = \psi_{mr} + \varepsilon_{mrk}, \quad \varepsilon_{mrk} \stackrel{\text{iid}}{\sim} \text{N}(0, \sigma_m^2) \quad (1)$$

where the ground instrument is unbiased (typically assumed for the benchmark), and its behaviour, invariant within AMSR-E pixels but is subject to variation across pixels. Spatial variation of the error structure is not assumed across the smaller AWAP pixels to ensure estimability of the within-pixel variance σ_m^2 .

Relating AMSR-E and big-pixel truth. Let B_m be the m th AMSR-E pixel; $v_m \equiv v(B_m)$ the AMSR-E value at B_m ; and $\bar{\psi}_m \equiv \psi(B_m) = (1/25) \sum_r \psi_{mr}$ the true soil moisture spatially averaged across B_m . Then,

$$v_m = \alpha_0 + \alpha_1 \bar{\psi}_m + \delta_m, \quad \delta_m \stackrel{\text{iid}}{\sim} \text{N}(0, \sigma_\delta^2). \quad (2)$$

Hence, (2) assumes that (i) AMSR-E data have been calibrated to represent a linear transformation of the underlying truth plus uncertainty with regression coefficient $\alpha = (\alpha_0, \alpha_1)$, and (ii) the behaviour of the AMSR-E device and the calibrated data is invariant over space with constant error variance σ_δ^2 .

Relating small-pixel truth and antecedent precipitation. We wish to model the true soil moisture as (i) being driven by antecedent precipitation (AP), and (ii) depending linearly on a transformation of this moving average, subject to additive error with constant variance σ_η^2 . Then,

$$\psi_{mr} = \beta_0 + \beta_1 p_{mr} + \eta_{mr}, \quad \eta_{mr} \stackrel{\text{iid}}{\sim} \text{N}(0, \sigma_\eta^2), \quad (3)$$

where $p_{mr} \equiv p(s_{mr})$ is the 4th-root transformation of a 20-day AWAP moving average at s_{mr} (Fig. 1). Here, the transformation and moving average bandwidth are based on exploratory analyses (see Henderson *et al.*, 2011, Section 3 for details).

Relating precipitation and space. Let $\mathbf{x}_{mr} = (x_{mr1}, x_{mr2})$ be a linear transformation of the longitude-latitude coordinates of the top-left corner of s_{mr} ; h_γ a polynomial of \mathbf{x}_{mr} with a pre-specified order and coefficient vector γ ; and $\phi_{mr} \equiv \phi(s_{mr})$ the spatial random error associated with s_{mr} . Then,

$$p_{mr} = h_\gamma(\mathbf{x}_{mr}) + \phi_{mr} + \zeta_{mr}, \quad \zeta_{mr} \stackrel{\text{iid}}{\sim} \text{N}(0, \sigma_\zeta^2). \quad (4)$$

We take the quadratic $h_\gamma(\mathbf{x}_{mr}) = \gamma_1 x_{mr1} + \gamma_2 x_{mr2} + \gamma_3 x_{mr1} x_{mr2} + \gamma_4 x_{mr1}^2 + \gamma_5 x_{mr2}^2$ as an initial model here.

Spatial patterns. We take a 5th order Gaussian conditional autoregressive (CAR) model for $\phi = \{\phi_{mr}\}$,

$$\phi \sim \text{CAR}(5; \tau^2) \quad (5)$$

with parameter τ^2 . A 5th order adjacency of AWAP pixels induces spatial dependence beyond one but under two AMSR-E pixels. See Chiu and Lehmann (2011) Appendix A for technical details of (5).

Prior distributions. Finally, we take values of a, a_1 and a_2 so that the following priors are reasonable (see Chiu and Lehmann, 2011, Appendix B):

$$\alpha_1, \beta_1 \stackrel{\text{iid}}{\sim} \text{N}(3a^{-1/2}, a^{-1}), \quad \beta_0, \gamma_q \stackrel{\text{iid}}{\sim} \text{N}(0, a^{-1}), \quad \sigma_m^2, \sigma_\delta^2, \sigma_\eta^2, \sigma_\zeta^2, \tau^2 \stackrel{\text{ind}}{\sim} \text{IG}(a_1, a_2) \quad (6)$$

for all $q=1, \dots, 5$ and all $m \in \mathcal{S} \equiv \{m: \sum_r K_{mr} \geq 1\}$. Here, $\text{IG}(a_1, a_2)$ is the inverse-Gamma distribution with mode $a_2/(a_1 + 1)$. Centring the slope priors at 3 times the prior standard deviation is to impose a large prior probability of positive association between v and ψ , and between ψ and p .

Technical remarks. (i) Statements (1)–(6) constitute our single hierarchical model. Marginally, response quantities in all of (1)–(5) exhibit spatial autocorrelation due to the model’s hierarchy. (ii) The hierarchical modelling framework does not require linearity/Gaussian assumptions; they are assumed in our preliminary research here to reduce complexity. (iii) In (3), the driver of true soil moisture is taken to be a manifestation of the true AP. At the expense of estimating extra nuisance parameters, (3) can be broken down further: AP being the regressor of each of ψ and p ; true AP would then replace p in (4). We have yet to pursue this extension. (iv) It may be unreasonable to assume either $\alpha_0=0$ given $\bar{\psi}_m$ in (2) or $\beta_0=0$ given p_{mr} in (3). Yet, having $\alpha_0, \beta_0 \neq 0$ makes these intercepts statistically unidentifiable/confounded. Since (3) is nested in (2), we preserve $\beta_0 \neq 0$ and set $\alpha_0=0$. The same reason applies to the absence of an intercept in (4). Thus, β_0 should be interpreted as the combined bias from all of (2)–(4) in the model hierarchy. (v) Gibbs sampling and data imputation details for (1)–(6) appear in Chiu and Lehmann (2011).

4 EVALUATING AMSR-E FOR THE MRC

As a preliminary exercise for this ongoing research, this paper focuses on the inference of daily ψ maps on selected days from January (summer) and June (winter) between 2007 and 2009. Then, the approach as described in Section 3.3 of Chiu (2011) is employed to provide a quantitative evaluation of AMSR-E. This approach relies on the ratio $R_m = \sigma_\delta^2 / \sigma_m^2$ as a model-based comparison of performance between the AMSR-E product and OzNet benchmark in $B_m, m \in \mathcal{S}$. Specifically, $R_m \leq 1$ would suggest that AMSR-E is no worse than OzNet for m . Unlike Chiu (2011), our paper does not model temporal dependence alongside spatial dependence. Thus, we make inference for R_m per selected day, where a graphical display of, say, 95% credible intervals (CIs) for R_m (over all 21 B_m s in MRC for which $m \in \mathcal{S}$) can be used to summarize the performance of AMSR-E against the benchmark over the MRC on that particular day. For example, consider Fig. 2. It shows that on 2007-01-20, a very small segment of each 95% CI for R_m is below 1; this suggests that, given the data and according to our model, it is unlikely for AMSR-E to perform better than OzNet anywhere in the MRC on that day. Similar patterns are seen earlier on 2007-01-08, except for the high credibility that AMSR-E performs better than benchmark in 7 of the 21 pixels (3 along the upper MRC perimeter, 1 in the Yanco study area, 2 and 1 respectively in the Kyeamba and Adelong sub-catchments). However, this day-to-day fluctuation of AMSR-E’s performance could well be due to temporal autocorrelation not being handled by the model of this paper, and is analogous to the pixel-to-pixel variation mentioned in Section 1 for temporally assimilated quantities.

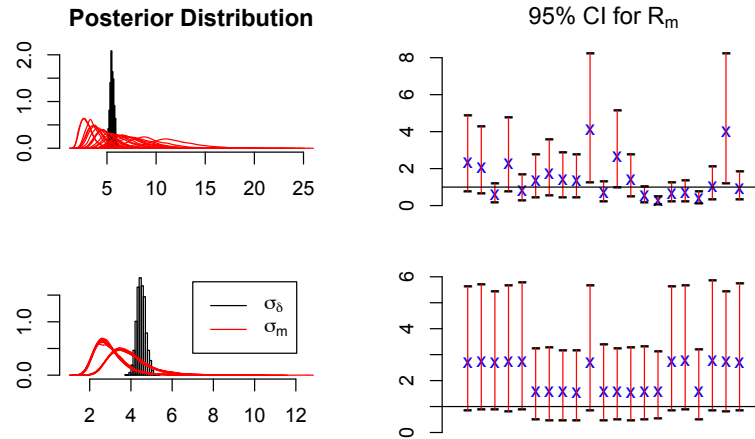


Figure 2. Posterior inference for σ_δ and the 21 values of σ_m s and R_m s on the 8th (top panels) and 20th (bottom panels) of January 2007. “Posterior” refers to the conditional behaviour of an unobserved quantity given the OzNet, AMSR-E, and AWAP observations. There is a 95% posterior chance that the true $R_m = \sigma_\delta^2 / \sigma_m^2$ lies inside the 95% CI in red; a blue “x” denotes the posterior mean for R_m .

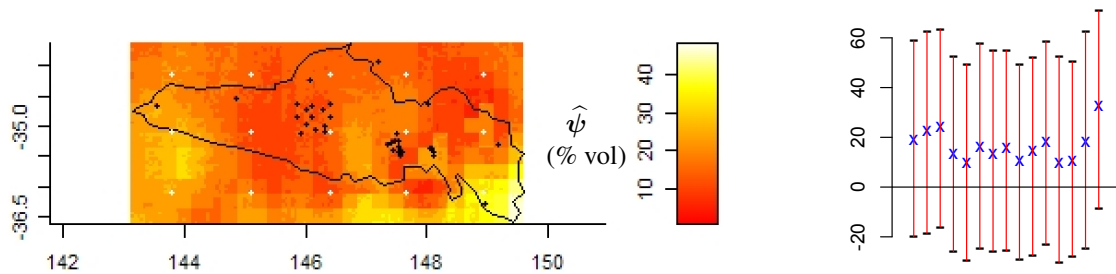


Figure 3. Left: Posterior mean, denoted by $\hat{\psi}$, for latent soil moisture ψ across the MRC on 2007-01-20. Right: 95% CI (red bar) and posterior mean (blue “x”) for ψ across selected AWAP pixels appearing as white “+” on the left panel, ordered from top to bottom then left to right.

5 LATENT SOIL MOISTURE AND OTHER UNOBSERVED QUANTITIES

While the evaluation of AMSR-E against benchmark is the primary goal of our model, a convenient bonus due to the model is our ability to define a new soil moisture product, $\hat{\psi}$ (Fig. 3), which estimates ψ according to the fixed-date spatial assimilation with OzNet, AMSR-E, and AWAP data. Given the sparsity of OzNet stations, one cannot expect g to have a heavy influence on the overall $\hat{\psi}$ map. In contrast, the resemblance on 2007-01-07 between v in Fig. 1 and $\hat{\psi}$ in Fig. 3 (despite their different spatial scales) suggests that the information about latent soil moisture is dominated by AMSR-E and less so by AWAP. This dominance is not consistent over time. Specifically, Fig. 4 shows that for 2009-06-22, $\hat{\psi}$ combines the flatness of the v surface and the spatial features of the p surface.

Date-specific posterior behaviour is also evident in dispersion parameters such as σ_m , as we have discussed for R_m which depends on σ_m . Moreover, despite only two different prior distributions (see Chiu and Lehmann, 2011, Appendix C) being used for all 21 values of σ_m , the assimilation for 2007-01-08 yields 21 distinguishable posterior distributions (Fig. 2, top left). In contrast, the assimilation for 2007-01-20 had little influence on the inference for σ_m so that the posterior and prior are virtually identical (Fig. 2, bottom left). On both days, the posterior of σ_δ greatly differs from the prior (not shown), suggesting the data’s heavy influence on the inference for this parameter.

Our inference also allows us to diagnose the adequacy of the model. Indeed, the linearity and Gaussian assumptions throughout the model hierarchy are somewhat inadequate, as suggested by negative soil moisture values covered by the 95% CIs for ψ (Fig. 3) and other quantities (not shown). The dilemma here

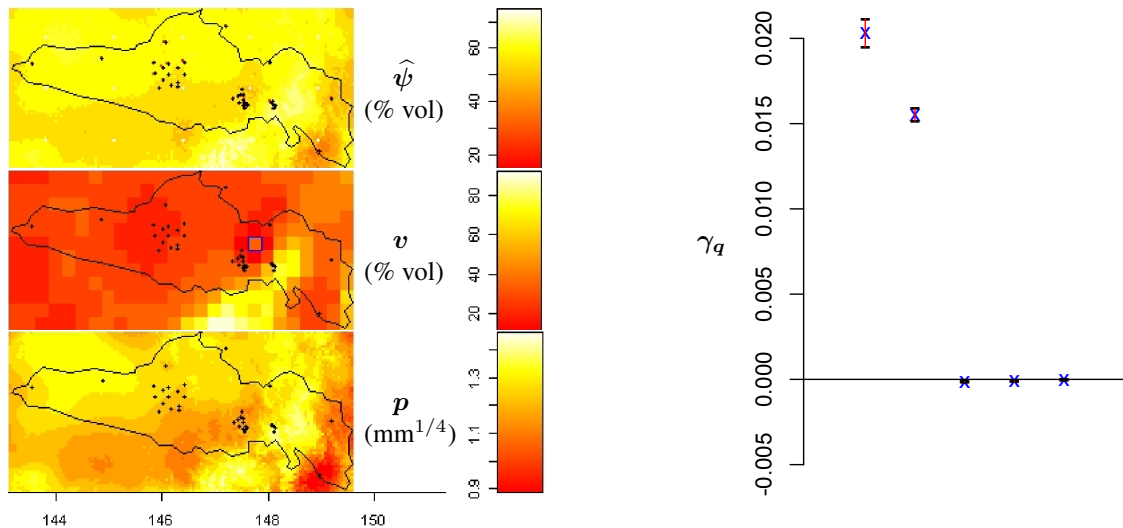


Figure 4. Left: Estimated latent soil moisture ($\hat{\psi}$), AMSR-E (v), and transformed AWAP moving average (p) on 2009-06-22 (latitude axis suppressed). Note the different intensity scales. Right: 95% CI (red bar) and posterior mean (blue “x”) for $\gamma_q, q=1, \dots, 5$ from left to right, on 2007-01-20.




Figure 5. Left: Posterior mean, denoted by $\hat{\phi}$, for spatial random effect ϕ across the MRC on 2007-01-20. Right: 95% CI (red bar) and posterior mean (blue “x”) for ϕ across selected AWAP pixels appearing as white “+” on the left panel.

is that these assumptions are standard for measurement errors such as those for a benchmark; yet, if latent soil moisture were 0, then unrealistically these assumptions would imply a 50% probability for a probe to give negative values. Currently, we are considering alternative assumptions to address this dilemma.

Another feature to diagnose is the form of the spatial polynomial h_γ in (4). The 95% CIs for γ_q in Fig. 4 suggest a high credibility for h_γ to not be a quadratic but a linear function of latitude and longitude with no interaction term. This is the case also for other selected dates we have studied (not shown).

A key assumption for taking the modelling approach of this paper is that spatial autocorrelation is substantial enough to warrant an entirely different framework from that for fixed-grid temporal assimilation. Fig. 5 shows the map of $\hat{\phi}$, the posterior mean for the spatial random effect ϕ that underlies the model hierarchy. Non-random spatial patterns (hence, spatial autocorrelation) are clearly visible.

A final set of unobserved quantities of interest is missing AMSR-E and OzNet data (AWAP data are complete). Figs. 1 and 4 show imputed AMSR-E values using the posterior mean (pixels outlined in dark blue), and similarly for OzNet values in Fig. 6. As is the case for any unobserved quantity in a Bayesian model, CIs can be easily produced for missing values of v and g (not shown).

6 CONCLUSIONS

Through a statistical spatial model for the evaluation of AMSR-E, we have laid out the unifying aspect of Bayesian hierarchical spatial modelling in water resources research involving complementary sources of data. Indeed, Bayesian hierarchical modelling is a powerful tool for direct statistical inference of any




Figure 6. OzNet soil moisture observations (grey) and imputed values (white) on 2007-01-20. “A” stations are in Adelong, “K” are in Kyeamba, “M” are along the MRC perimeter, and “Y” are in Yanco.

unobserved quantity in the model, whether it be missing *in situ* data, an abstract latent parameter such as true soil moisture on an infinitesimal time and spatial scale, and so on. Future attempts to improve model assumptions, incorporate additional drivers of latent soil moisture, and extend the modelling framework of this paper to the temporal domain will lead to a fully unified spatio-temporal approach.

ACKNOWLEDGEMENT

This ongoing research is part of the WIRADA-AWRA Project between CSIRO and the Bureau of Meteorology. We thank project teammates W. Jin (CSIRO Mathematics, Informatics and Statistics) for collating OzNet data and L. Renzullo (CSIRO Land and Water) for providing AMSR-E and AWAP data.

REFERENCES

- Banerjee, S., B.P. Carlin, and A.E. Gelfand (2004). *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC.
- Chiu, G.S. (2011). *Statistical Model-Based Alternatives to Integrate and Evaluate Soil Moisture Products*. Technical Report #CMIS2011-EP11722. CSIRO Mathematics, Informatics and Statistics.
- Chiu, G.S. and E.A. Lehmann (2011). *Appendices to Bayesian hierarchical modelling: incorporating spatial information in water resources assessment and accounting*. <ftp://ftp.csiro.au/MODSIM2011>
- Cressie, N., and C.K. Wikle (2011). *Statistics for Spatio-Temporal Data*. Wiley.
- Crow, W., and X. Zhan (2007). Continental-scale evaluation of remotely sensed soil moisture products. *IEEE Geoscience and Remote Sensing Letters*, 4, 451–455.
- de Lannoy, G.J.M., R.H. Reichle, P.R. Houser, K.R. Arsenault, N.E.C. Verhoest, and V.R.N. Pauwels (2010). Satellite-scale snow water equivalent assimilation into a high-resolution land surface model. *Journal of Hydrometeorology*, 11, 352–369.
- Draper, C.S., J.P. Walker, P.J. Steinle, R.A.M. de Jeu, and T.R.H. Holmes (2009). An evaluation of AMSR-E derived soil moisture over Australia. *Remote Sensing of Environment*, 113, 703–710.
- Gelfand, A.E. and A.F.M. Smith (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85, 398–409.
- Henderson, B., G. Chiu, W. Jin, E. Lehmann, and S. Barry (2011). *AWRA model-data fusion: Evaluation, error characterisation and blending*. Water for a Healthy Country Flagship. CSIRO.
- Jones, D.A., W. Wang, and R. Fawcett (2009). High-quality spatial climate data-sets for Australia. *Australian Meteorological and Oceanographic Journal*, 58, 233–248.
- McMillan, N.J., D.M. Holland, M. Morara, and J. Feng (2010). Combining numerical model output and particulate data using Bayesian space-time modeling. *Environmetrics*, 21, 48–65.
- Owe, M., R.A.M. de Jeu, and T.R.H. Holmes (2008). Multi-sensor historical climatology of satellite-derived global land surface moisture. *Journal of Geophysical Research*, 113, F01002.
- Reichle, R.H. and R.D. Koster (2003). Assessing the impact of horizontal error correlations in background fields on soil moisture estimation. *Journal of Hydrometeorology*, 4, 1229–1242.
- Reichle, R.H., D.B. McLaughlin, and D. Entekhabi (2002). Hydrologic data assimilation with the ensemble Kalman filter. *Monthly Weather Review*, 130, 103–114.
- van Dijk, A. (2010a). *Monitoring Water Resources from Catchment to Continent*. Water for a Healthy Country Flagship. CSIRO.
- van Dijk, A.I.J.M. (2010b). *The Australian Water Resources Assessment System: Technical Description*. WIRADA Technical Report No. 3. Water for a Healthy Country Flagship. CSIRO.