

High-quality spatial climate data-sets for Australia

David A. Jones, William Wang and Robert Fawcett
National Climate Centre, Australian Bureau of Meteorology

(Manuscript received March 2009; revised September 2009)

In this paper, we describe a new high-quality set of historical and ongoing real-time climate analyses for Australia. These analyses have been developed for improving the definition of past climate variability and change over Australia and to improve on estimates of recent climate. The climate analyses cover the variables of rainfall, temperature (maximum and minimum) as well as vapour pressure at daily and monthly timescales and are complemented by remotely sensed and model-derived data described elsewhere.

New robust topography-resolving analysis methods have been developed and applied to *in situ* observations of rainfall, temperature and vapour pressure to produce analyses at a resolution of $0.05^\circ \times 0.05^\circ$ (approximately 5 km \times 5 km). The new methodologies are similar to those applied internationally, but in applying them to Australia we found it necessary and desirable to introduce a number of innovations. The resulting analyses represent substantial improvements on operational analyses currently produced by the Australian Bureau of Meteorology, and have a number of advantages over other similar data-sets currently available.

Careful attention has been paid to developing systems and data-sets which are robust and useful for the monitoring of both climate variability and climate change. These systems are now running in real time and are expected to form the basis for the ongoing monitoring of Australia's surface climate variability and change by the Australian Bureau of Meteorology. The underlying data and associated error surfaces (grids and station data) are updated in real time and are all available free of charge through the Bureau's climate website (www.bom.gov.au/climate).

Introduction

Recent years have seen much of Australia suffer from severe meteorological drought and a series of climatic extremes (e.g. Bureau of Meteorology 2008a, 2008b). As a result, water resources have widely fallen to record lows, and agricultural production in southern and eastern parts of Australia has been poor with a series of crop failures (e.g. Murray-Darling Basin Commission 2007; Australian Bureau of Agriculture and Resource Economics 2007). High-quality national climate information is clearly needed to place these climate-

driven events in a proper historical perspective and to provide a context for understanding the associated impacts on humans and the environment.

A key to better management of Australia's physical resources is ensuring that expectation and demand match the long-term supply. Matching a demand to the available resource is clearly required in the case of water, where the idea of sustainable yields is fairly well developed (e.g. Chiew et al. 2008). However, it is also clear that other climate variables such as temperature can be considered as a resource, with agricultural productivity (for example) being closely tied to temperature in much the same way as it is to rainfall (Cline 2007). A better characterisation of Australia's climate and associated variability should lead to better risk management and improved decision-making processes.

Corresponding author address: David A. Jones, National Climate Centre, Bureau of Meteorology, GPO Box 1289, Melbourne, Vic. 3001, Australia.
Email: d.jones@bom.gov.au

In this paper, we describe a new set of daily and monthly spatial climate analyses for Australia, covering the variables of rainfall, temperature and vapour pressure. These analyses extend through the twentieth century and are updated in real time. The analyses are a component of a new effort to measure Australian climate, combining *in situ* surface observations, satellite data and a water-carbon model (Raupach et al. 2008), in a project called the Australian Water Availability Project (AWAP) (2004).

There is a very extensive literature on spatial analysis techniques for analysing *in situ* observations dating back over half a century (e.g. Cressman 1956; Daley 1993). These range from relatively simple inverse distance-weighting schemes (Cressman 1956; Koch et al. 1983), through to highly complex systems using dynamical model frameworks in which surface parameters are directly assimilated or otherwise estimated (Kalnay et al. 1996; Compo et al. 2006; Ebert et al. 2007). More recently, remotely sensed data have been applied directly to the estimation of surface climate parameters, though in the case of rainfall with varying degrees of success (Ebert et al. 2007; Renzullo et al. 2008).

The most significant challenge to the consistent spatial analysis of historical climate data over Australia is the general decline in data availability which occurs as one goes back in time (e.g. Jones and Trewin 2002) and the non-availability of modern data types (e.g. data from radar and satellites) for much of the period of the instrumental record. When analysing historical data we require that the analyses are not only accurate as measured by some error metric but in addition:

- can be compared to and are consistent with the long-term averages (climatology);
- provide a more accurate representation of short-term climate than does the long-term average; and
- contain values which are limited to a physically realistic range.

The last of these points, while seemingly trivial, is important as it is possible for interpolated surfaces to become unrealistic in data voids when meteorological gradients are strong. This is particularly important back in time with networks becoming rather sparse during the first half of the twentieth century.

Generating daily analyses of meteorological data which are consistent with monthly and longer-term analyses is far from straightforward (e.g. Rayner et al. 2004). Temporal averaging dampens the small-scale variability, and also reduces random observational errors (Jones and Trewin 2002). This means that the characteristic spatial scales tend to increase with the temporal scale. In addition, the relationship with topography is significantly stronger on longer timescales, though often not well-resolved by networks.

The favoured methods for analysing historical climate observations are geostatistical techniques, typically applied in a univariate framework (e.g. Jones and Trewin 2000a; Jeffrey et al. 2001). For the most part these methods behave similarly, and often show relatively similar accuracies (e.g. Jones and Trewin 2000a), despite rather different levels of sophistication. It is common practice for these techniques to be di-

vided into three broad classes; empirical interpolation, statistical interpolation and function fitting (Jones and Trewin 2000a). In this work we describe a hybrid technique which combines empirical interpolation with function fitting, in a system which is similar to those described by Hunter and Meentemeyer (2005) and Xie et al. (2007).

Climate data

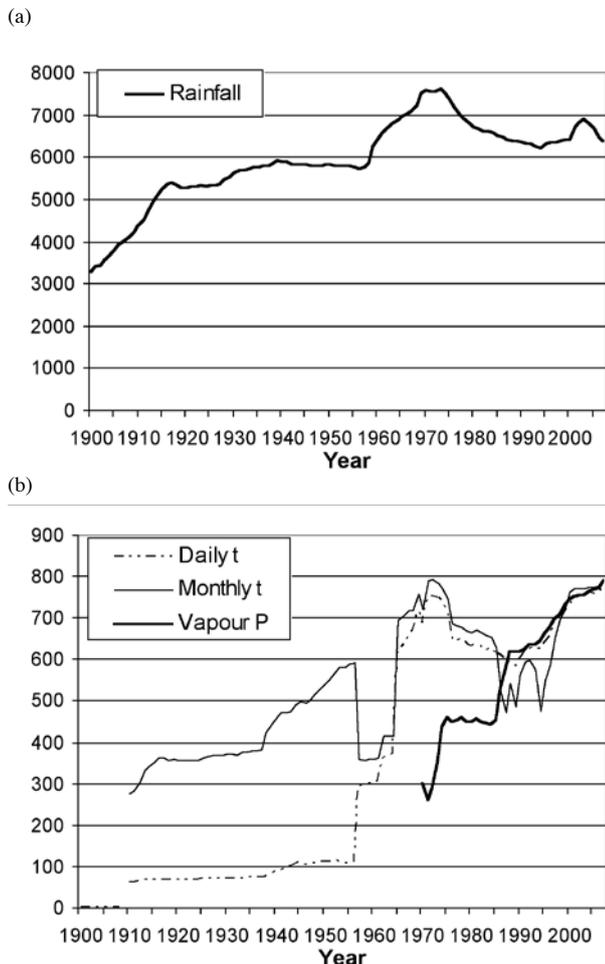
In constructing these new climate data we have deliberately used only *in situ* data managed by the Bureau of Meteorology. We recognise that rather better analyses are possible for recent years, drawing on remotely sensed information from satellites and radar. However, these newer data types typically do not go back far in time, and even long-run data such as infrared data from geostationary satellites show large and systematic shifts in resolution and quality arising from multiple changes in satellites, orbits and instrumentation.

The meteorological data used in this study are taken from the national climate databank of the Bureau of Meteorology, called the Australian Data Archive for Meteorology (ADAM). The climate analyses are generated using both daily and monthly data contained in this database. ADAM is updated in real time, with significant non-real-time inputs for some meteorological variables (in particular, rainfall), meaning that the analyses can (and need to) be updated over time. For rainfall about one-third of the current network in ADAM reports in real time, while the remaining data arrive via mail on paper reports with most records added within three months of the end of the month. The ADAM database is an evolving resource, with new historical data added from time to time (in addition to the continual input of newly observed data) and ongoing quality control, meaning that improvements to the station data are frequently made and subsequently available for analysis.

Figure 1 shows the total number of stations used for the analysis of rainfall, temperature and vapour pressure by year for the periods considered in this paper. For rainfall there is a nearly monotonic increase in the number of reports from around 3000 at the start of the twentieth century to more than 7000 in the 1970s. Recent decades have seen an overall slight decline in the network, with some variability. The decline reflects a slow loss of manual observations taken by volunteer observers, in particular.

The available temperature data at both daily and monthly timescales vary rather more erratically in response to an increasing and evolving network and also the fact that a large volume of paper reports remains undigitised (Jones and Trewin 2002). In recent years the vapour pressure (dew-point temperature) network is very similar to the daily temperature network, but prior to about 1980 there were rather fewer dew-point observations available, with a sharp decline in the network apparent in the 1970s. The station network for rainfall is substantially larger than those for temperature and dew-point, as rainfall is easier to record and the majority of rainfall observations come from volunteer observers.

Fig. 1 The number of stations contributing to the (a) rainfall and (b) temperature (daily and monthly) and vapour pressure analyses (twice daily) by year.



The earliest temperature data in ADAM are from 1844, while rainfall data extend back to 1832 and dew-point data (the basis for vapour pressure) back to 1864. However, during the earliest years observation practices were often very different from modern-day practices and the networks sparse. For Australian temperature, Stevenson Screens became widespread around 1910 (Torok and Nicholls 1996) and for this reason the AWAP temperature analyses commence in 1911. The earlier recordings were made in a variety of instrument enclosures (Torok 1996), which can lead to misleading conclusions about the nature of past climate variations (e.g. Trewin 1997).

The Australian rainfall network is particularly sparse prior to 1900, with very little data in central Australia and large data voids in parts of the south and west. For this reason, the AWAP rainfall analyses commence in 1900, as has been the case with the existing operational system (Jones and Weymouth 1997). The available vapour pressure network is sparse prior to 1971 and accordingly the analyses start in this year.

There is the potential to extend the national temperature and vapour pressure analyses back some years earlier with further digitisation of historical data and suitable quality control. In addition there is the potential for extending regional rainfall analyses into the nineteenth century in some better-observed regions such as Victoria. This will be a focus of future work.

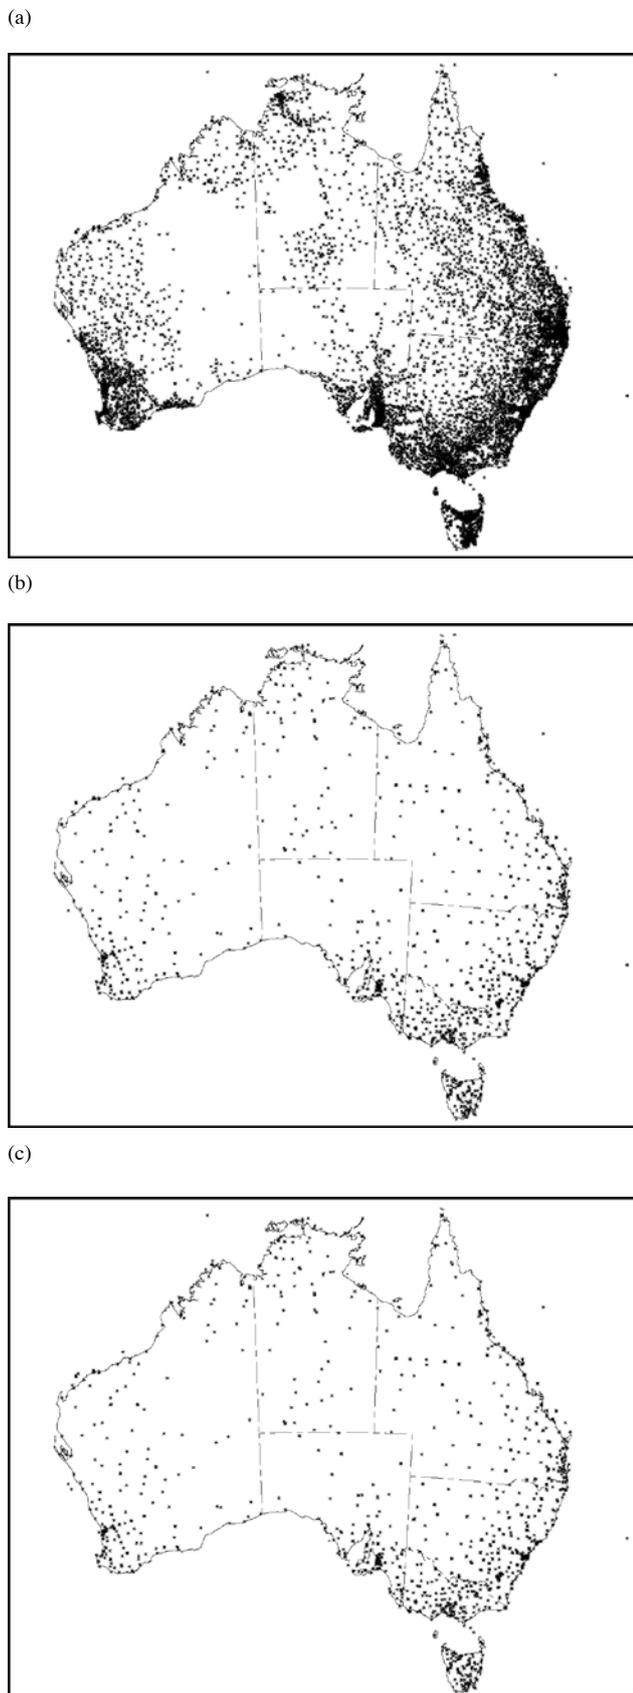
For rainfall, the base daily data and the spatial analyses represent the total precipitation (including rain, snow, hail and dew) accumulated in the 24-hour period to 0900 local time (0800 in the case of pre-Federation Queensland, which is effectively the years 1900–1907 as far as the analyses are concerned). The maximum (minimum) temperatures are the highest (lowest) temperature for the 24-hour period starting (ending) 0900 local time. This convention means that the minimum and maximum temperatures will have usually occurred on the same calendar day (in the morning and afternoon, respectively). The vapour pressure has been calculated at stations using observations of dew-point temperature, following Murray (1967). The vapour pressure is that observed at 0900 and 1500 (local time). These two times have the best data coverage for Australia, and are chosen for this reason. We note that the vapour pressure has a rather weak diurnal cycle through the day (Jeffrey et al. 2001; Jones et al. 2007). We also note that accumulated maximum/minimum temperature and rainfall values over more than one day have been omitted from the analyses where these are identified in the ADAM database.

The meteorological variables and analyses are summarised in Table 1. We note that the early data come entirely from manual observations, while more recently an increasing fraction of the data comes from automatic weather stations. Figure 2 shows the network of available stations for recent years (1980 to the end of 2007) showing a good coverage except in the arid interior.

Table 1. Definition of the meteorological data and associated analyses.

Variable	Source	Temporal resolution	Spatial resolution
Precipitation	Analysis of rain gauge data	Daily and monthly total	0.05°×0.05°
Daily maximum temperature	Analysis of thermometer data	Daily and monthly average	0.05°×0.05°
Daily minimum temperature	Analysis of thermometer data	Daily and monthly average	0.05°×0.05°
Vapour pressure 'humidity'	Analysis of vapour pressure data (calculated from temperature and dew-point)	Daily and monthly average at 9 am and 3 pm	0.05°×0.05°

Fig. 2 The networks of (a) rainfall, (b) temperature and (c) dew-point temperature (vapour pressure) stations contributing to the analyses from 1980 to 2007.



We recognise that there are a number of important climate variables which are not considered in this paper for reasons of data quality and/or differences in analysis methodology. Specifically, we do not analyse wind data because these data are extremely problematic and show large and artificial drifts over time due to changes in instrumentation (Rayner 2007). Estimates of solar radiation form part of the AWAP product set, and are derived from high-resolution visible imagery from the Geostationary Meteorological Satellites, while evaporation estimates are generated via a water balance model (see Weymouth and Le Marshall 2001; Jones et al. 2006; Raupach et al. 2008).

Generating the spatial analyses

Seaman and Hutchinson (1985) and Jones and Trewin (2000a) describe a number of the common approaches to the generation of climate analyses for Australia using geostatistical techniques, of which there are many.

We have used an anomaly-based approach, in part because this means that the full set of analyses is largely consistent with the long-term climatology. A further motivation is that anomalies tend to be spatially rather smoother than the raw data and climatology provides information beyond that contained in individual observations and data for a single day or month. Our anomaly method is similar to that described by Hunter and Meentemeyer (2005) and Xie et al. (2007). It uses a decomposition of the meteorological variable being analysed (such as rainfall) into its long-term average and an associated 'anomaly'.

The analysis methodology represents an extension of the background and increment method which has been popular in meteorological analysis (e.g., Koch et al. 1983; Daley 1993; Jones and Trewin 2000a; Rayner et al. 2004). It is important that analyses which are to be used for operational climate monitoring are robust and hence the similarity with these systems is seen as an advantage.

An important factor is that anomalies tend to be weakly related to altitude, in part because of the tendency for atmospheric anomalies to be nearly barotropic. This means that they can be more adequately analysed with a two-dimensional analysis procedure, giving efficiency savings. In addition, this approach reduces the impact of network changes back in time, which sees a sharp decline in the number of stations generally in earlier years, and more specifically for those at higher elevations.

For both temperature and vapour pressure, the anomaly is defined as a simple difference from the climatology, leading to $T(t) = \bar{T} + T'(t)$ for the meteorological variable T for some day or month t . The over-bar denotes a long-term average (always monthly) for the meteorological variable. For rainfall, we have defined the anomalies using division rather than subtraction, leading to the representation $R(t) = \bar{R} \times R'(t)$.

For all variables the representation given above is applicable for both station observations and gridded analyses, and for both daily and monthly data. In each case the station average or grid average is that for the relevant calendar month.

For rainfall, the use of ratios in interpolation can become unstable when the monthly mean is very low. This is most commonly a problem in northern Australia during winter when large areas have an average monthly rainfall of less than 10 mm but very occasionally experience heavy rainfall (an example is Wyndham in northwest Australia which averages just 7.5 mm of rain for the whole of winter but which has experienced daily totals in excess of 50 mm during this season). To address this we have applied a floor of 5 mm to the denominator in the ratio. This heuristic fix was applied after testing on a series of case examples.

We have used the Barnes successive-correction method for the analysis of the daily and monthly anomalies (Koch et al. 1983; Seaman 1989; Jones and Weymouth 1997) and three-dimensional smoothing splines for the analysis of monthly climatological averages of 0900 and 1500 vapour pressure, maximum and minimum temperature and rainfall (Hutchinson 1995). These two techniques have been widely used for climate analyses and have been found to be robust for the types of analyses for which we are using them (e.g. Hutchinson 1995; Jones and Trewin 2000a; Bureau of Meteorology 2000; Jeffrey et al. 2001; Rayner et al. 2004).

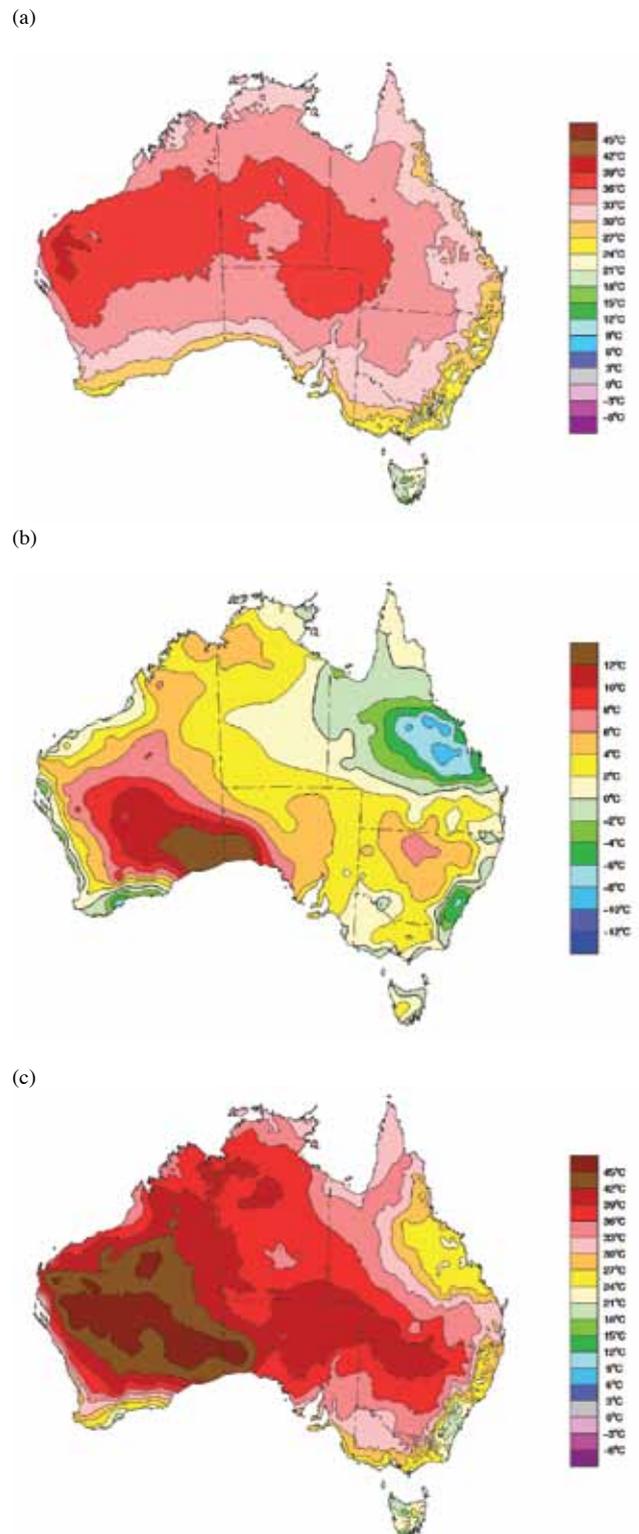
The anomaly analysis is generated using an optimal two-dimensional Barnes successive correction analysis procedure, as described by Jones and Trewin (2000a) analysed to grid locations. The weighting function is obtained using the iterative Barnes algorithm described by Jones and Weymouth (1997) with the analysis parameters generated via exhaustive cross-validation (Seaman 1989) on pre-2000 data. The Barnes analysis technique has a number of advantages, including being efficient, robust (coping with strong gradients and data voids) and highly tunable. Jones and Trewin (2000a) have previously shown that the accuracy of this method is similar to that of more sophisticated techniques, and avoids some of the weaknesses such as extrapolation of unrealistic values into data voids and the dampening of variance.

The smoothing spline approach is particularly suited to analysing smooth climatological relationships between meteorological variables and three-dimensional position (i.e. latitude, longitude and altitude) but not well suited to noisy or sparse data. The climatological surfaces from the spline have been fitted by minimising the generalised cross-validation error (Wahba and Wendelberger 1980), as is common practice (e.g. Hutchinson 1995; Jeffrey et al. 2001).

The Barnes and spline methods are both 'statistically optimal', in that the analysis fields have the smallest error subject to constraints on the smoothness and spectrum of the final field. The analysis methods have been implemented in a modular fashion, which allows for the use of alternative methods for generating the climate (background) and anomaly (increment) fields in future as improved techniques are developed. These might include the use of weather-model forecasts or satellite estimates, for example.

The final analysis is a simple sum (or multiplication in the case of rainfall) of the climatological and the anomaly analyses. A graphical example of this procedure is shown in Fig. 3

Fig. 3 Climatological (monthly) (a) average for February for 1971-2000, (b) daily anomaly analysis and (c) summation, for maximum temperature on 1 February 2007.



for daily temperature. The rainfall, temperature and vapour pressure climatologies for January and July are provided in Jones et al. (2007).

The monthly and daily AWAP analyses are not constrained to be absolutely consistent, and indeed they are not for a number of reasons. Particularly, for historical data (pre-1957) and in real time there is a difference in the quantity of data available for the daily versus monthly analyses. Consistency can be achieved in a straightforward way by a rescaling of the daily data using the monthly analyses. We have applied a rescaling to daily rainfall to produce an additional daily data set for historical data which is consistent with the monthly analyses, although this requires a delay in real time. We do not consider these rescaled data further, but note their existence.

The analysis scheme as applied requires station latitude, longitude and altitude. Prior to the mid 1970s, a significant number of (now closed) rainfall stations lack a station altitude in ADAM: for example, nearly 1400 stations with rainfall data in ADAM in 1920 do not have a station altitude, representing about 20 per cent of the total network for that year.

We have generated estimated station altitudes for a subset of these stations from a high resolution 0.0025° latitude/longitude grid subject to constraints on the smoothness of the local topography to minimise estimation errors. Specifically, we required that the altitude vary by less than 75 m in an 8 × 8 grid-point lattice (representing an area of approximately 2 km × 2 km) around the station. We have chosen not to augment the network for temperature and vapour pressure in the same manner owing to the strong control altitude places on these variables.

Defining the station and gridded climate normals

The climate normal fields for the analysis are produced for each calendar month for each variable (rainfall, maximum and minimum temperature and 0900 and 1500 vapour pressure) for the 1911-1940, 1941-1970 and 1971-2000 base periods. The climatologies are for 30-year periods following World Meteorological Organization convention (WMO 1989). Note that the terms 'climate averages', 'climate means' or 'climate normals' are all interchangeable. They refer to arithmetic calculations based on observed climate values for a location (or through space) over a specified time period and are used to describe the climatic characteristics of that location.

The climate normals are used to form the final analysis through the addition of (or, for rainfall, multiplication by) the anomaly analysis fields. The climate base period used for each daily or monthly analysis depends on the date of the analysis, with the anomaly calculated with respect to the monthly average for the nearest base period. For example, all daily and monthly rainfall analyses for 1900-1940 use the 1911-1940 base period. In practice, many stations have incomplete data in each of the 30-year periods owing to missed observations and station openings and closings. Following extensive testing, we found that stations with twelve or more complete monthly observations for the same calendar month provided useful information for the spatial climate analyses for that month (for example, twelve valid January monthly observations for the period 1971-2000). These are subsequently referred to as 'qualified' stations.

The use of incomplete station records

The analysis procedure requires the calculation of station climate normals (for 1911-1940, 1941-1970 and 1971-2000) as well as climate normal grids based on station data. As described previously, the station normals used for generating the climate normal grids are limited to those 'qualified' stations with twelve or more years of records (or four in the case of high altitude stations as described below).

Using data for stations with short records (less than twelve complete years) for calculating the individual monthly and daily anomaly grids requires us to form an estimate of the local 30-year average at the station. We have achieved this using a trade-off between the station's temporal average (using all available observations in the 30-year period) and an estimate of the temporal average calculated by interpolation to the station location of the smoothing spline analysis applied to those stations with twelve or more observations (for each of the 1911-1940, 1941-1970 and 1971-2000 periods).

A linear combination of these two independent estimates at the stations with incomplete records was used, with the weights chosen through an optimisation process using cross-validation with the same values applied across the whole of Australia. This optimisation revealed that the final analysis accuracy was only weakly dependent on the exact form of the combination at stations. This weighted combination estimate for incomplete stations is internal to the analysis procedure and not used for the climatological analyses (Fig. 3(a)).

Ensuring a consistency between adjacent climate normal grids

The Australian station network shows substantial changes over time. For the most part, there is a tendency for improved coverage (e.g. Jones and Weymouth 1997; Jones and Trewin 2000a), but in some areas climatologically unique stations have closed, an effect which is most apparent in semi-arid coastal locations such as around Shark Bay in Western Australia.

The network changes can cause locally large and spurious differences between the climate normal fields for adjacent periods (e.g. between 1911-1940 and 1941-1970, and 1941-1970 and 1971-2000). These then flow into the final analysis, as the anomaly analysis grids are added to (or multiplied by) the climate normal grid. The issue of network changes is more particularly a problem going backwards in time; the 1911-1940 period has less data than the 1941-1970 period, which in turn has less data than the 1971-2000 period.

We have preserved continuity in the climatological analyses across the three normal periods by using consecutive climate normals at suitable stations (i.e. those stations which have enough data to calculate normals in at least one of the adjacent 30-year periods such as 1941-1970 and 1971-2000) perturbed by the difference between the normal periods. Obviously, for stations which are present in both normal periods (the 'qualified' stations) this adds no additional information to the analysis process, but for those stations which

are ‘qualified’ for only one of the normal periods this introduces substantially improved consistency between adjacent climate normals. Most importantly, this method allows the denser recent network to augment the climatology for the earlier periods on the assumption that the differences in the 30-year averages are spatially smooth.

The difference (ratio for rainfall) at a station has been calculated by an interpolation of a monthly difference grid using those stations which are qualified in consecutive normal periods (e.g. 1911-1940 and 1941-1970). The difference grids are analysed using the same Barnes analysis method as for the daily and monthly anomaly analyses. The overall impact on analysis accuracy is relatively small (a very slight improvement). The reason for using this technique is to ensure the analyses more faithfully capture both the variability and secular change across the full period (see below).

High altitude stations

High altitude stations play a key role in defining the vertical gradient of climate parameters for the climate normal fields and hence for the final analyses. Unfortunately, it has only been in recent years, with the installation of automatic weather stations (AWS), that data have become available from the more isolated high altitude locations such as the Grampians in western Victoria, the central plateau and western ranges of Tasmania and parts of the Victorian and New South Wales Alps (see Fig. 4).

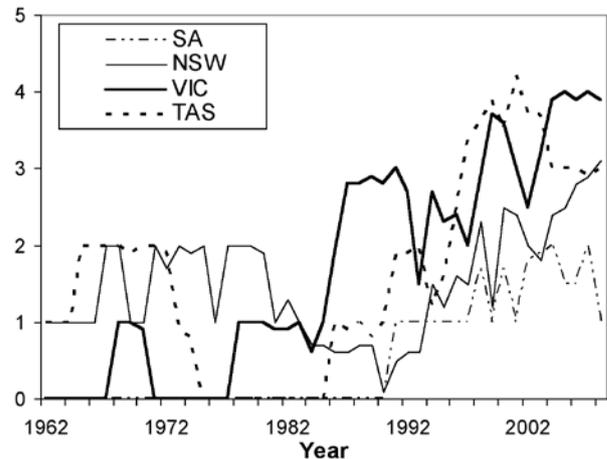
We have found it necessary to apply a heuristic modification to the analysis procedure to incorporate information from high altitude locations which are under-represented in the Australian station network. To capture these short-lived stations we have relaxed the rule for the inclusion of stations for the calculation of the climate normal fields to four (down from twelve) or more complete monthly observations where the station altitude is 1000 m or higher. The inclusion of additional high altitude stations was found to make a modest but important improvement in representing high altitude climate. We note that W. Wright (personal communication) found it useful to do something similar when constructing the *Climatic Atlas of Australia* (Bureau of Meteorology 2000).

Measuring the accuracy of the analyses

The accuracy of the spatial analyses has been determined for the full analysis period using verification against station observations. Fully cross-validated estimates have been generated for the seven years 2001-2007, following the broad methodology described by Jones and Trewin (2002). It is possible to fully cross-validate both the anomaly analyses and also the station climatologies for this period, noting that the climatologies use data only up to and including 2000.

Additional cross-validated verification statistics have been produced for the earlier periods 1911-2000 for temperature, 1900-2000 for rainfall and 1971-2000 for vapour pressure. We note that the earlier verification statistics are not absolutely cross-validated because dependent station data have been

Fig. 4 The average number of high-elevation stations operating in January of the listed year. High-elevation stations are defined as those above 1500 metres in NSW and Victoria, above 1000 metres in Tasmania and above 700 metres in South Australia.



used for the calculation of the climate normal grids. It is not practical to fully cross-validate the climate normal grids. In practice, however, the impact of not cross-validating the climatology is extremely slight as most of the analysis error for individual months and days is associated with the anomaly interpolation rather than the gridded climatology. In addition, the climatological grids are only slightly affected by individual station observations owing to the underlying smoothness of these. This can be seen by the similarity of the analysis errors in years before and after the year 2000 (see below).

Cross-validation has been achieved by randomly deleting five per cent of the stations in the network, performing an analysis using the remaining 95 per cent of station observations and then calculating the analysis errors for the omitted stations. This process was repeated twenty times for each month/day, providing independent verification statistics at stations. Every single monthly grid was cross-validated in this process, while the daily cross-validation was applied to ten days in each calendar month providing 120 analyses per year. We note that the errors are subsequently accumulated across stations which are not distributed evenly in space.

Jones et al. (2006, 2007) describe a range of issues with the method of cross-validation. Importantly, cross-validation will tend to give somewhat inflated analysis errors, as the method involves a modest degrading of the data network compared to reality (e.g. Jones and Trewin 2000a; Jeffrey et al. 2001).

In addition, calculating analysis errors by independent cross-validation against station observations introduces a bias due to observation ‘error’ (see Daley 1993; Jones and Trewin 2002). Consider a cross-validated estimate of a station value T at station k (at three-dimensional location \mathbf{r}_k) and time t , denoted by $\hat{T}(\mathbf{r}_k, t)$. This is calculated using the 95 per cent of the network which is retained in the cross-validation

step. The cross-validated analysis error is given by

$$E_k(t) = \hat{T}(\mathbf{r}_k, t) - T_k(t) \quad \dots 1$$

Aggregating across time, we can calculate a station root mean square analysis error (RMSE);

$$RMSE_k = \sqrt{\frac{1}{N} \sum_{t=1}^N [E_k(t)]^2} = \sqrt{\frac{1}{N} \sum_{t=1}^N [\hat{T}(\mathbf{r}_k, t) - T_k(t)]^2} \quad \dots 2$$

Note that N will vary from station to station and according to whether the analysis is for daily or monthly data. The observation $T_k(t)$ can be divided into a ‘true’ component, and an ‘observational error’ component, $e_k(t)$. The true component is what would be measured if the observation at station k was completely accurate, while the error component is the error introduced due to factors such as instrument miscalibration, misreading by the observer, errors in spatial representativeness arising from specific factors at the observation site and so on. Hence, we have

$$RMSE_k = \sqrt{\frac{1}{N} \sum_{t=1}^N [(\hat{T}(\mathbf{r}_k, t) - T_k^{True}(t)) - e_k(t)]^2} \quad \dots 3$$

Clearly, even a perfect analysis will have a non-zero cross-validated error because all observations have some level of error. To obtain a zero cross-validated error, the observations also need to be ‘perfect’. While it is common practice for the cross-validated differences between independent observations and analyses to be treated as ‘analysis errors’, it is important to keep in mind that they also contain an observation error component. Daley (1993) and Jones and Trewin (2000a) describe how the observational errors can be estimated statistically.

In defining the analysis errors averaged across time and stations we have used the additional measures of bias and mean absolute error (MAE). These are both defined in the usual way (e.g. Jones and Weymouth 1997);

$$BIAS_k = \frac{1}{N} \sum_{t=1}^N [E_k(t)] \quad \dots 4$$

and

$$MAE_k = \frac{1}{N} \sum_{t=1}^N |E_k(t)| \quad \dots 5$$

Quality of the analyses

Cross-validated statistics for the fully independent 2001 to 2007 period and for the earlier years are provided in Tables 2 to 4, with maps of RMSE for 2001 to 2007 in Figs 5, 6, 8 and 10. For reference, we also provide national average statistics for the current operational Barnes analysis system used at the Bureau of Meteorology (Jones and Weymouth 1997; Jones and Trewin 2000a). This operational system uses a relatively simple two-dimensional analysis with an isotropic weighting function, something very similar to that used in the new system for the direct analysis of the rainfall and temperature ‘anomalies’. There are no ‘operational’ climate analyses for vapour pressure, so a comparison is not possible

Table 2. Verification statistics for (a) monthly and (b) daily maximum and minimum temperatures. The units are °C.

(a)					
	<i>Mean</i> 2001- 2007	<i>RMSE</i> 2001- 2007	<i>Bias</i> 2001- 2007	<i>RMSE</i> 1910- 2000	<i>Bias</i> 1910- 2000
Monthly maximum temperature					
AWAP	24.9	0.7	0.0	0.7	0.0
Operational		1.6	0.0	1.7	0.0
Monthly minimum temperature					
AWAP	12.7	1.0	0.0	0.9	0.0
Operational		1.5	0.0	1.7	-0.1
(b)					
	<i>Mean</i> 2001- 2007	<i>RMSE</i> 2001- 2007	<i>Bias</i> 2001- 2007	<i>RMSE</i> 1910- 2000	<i>Bias</i> 1910- 2000
Daily maximum temperature					
AWAP	24.9	1.2	0.0	1.7	0.0
Operational		1.9	0.0	3.6	-0.1
Daily minimum temperature					
AWAP	12.8	1.7	0.0	2.0	0.0
Operational		2.1	-0.1	3.1	-0.1

Table 3. Verification statistics for (a) monthly and (b) daily rainfall. The units are mm and % for MAE/Mean.

(a)					
	<i>Mean</i>	<i>Bias</i>	<i>RMSE</i>	<i>MAE</i>	<i>MAE/Mean</i>
Monthly rainfall 2001-2007					
AWAP	54.3	0.5	21.2	11.5	21.0
Operational		0.1	24.4	12.8	23.5
Monthly rainfall 2001-2007					
AWAP	61.8	0.0	19.6	11.2	18.1
Operational		-0.1	24.7	13.1	21.5
(b)					
	<i>Mean</i>	<i>Bias</i>	<i>RMSE</i>	<i>MAE</i>	<i>MAE/Mean</i>
Daily rainfall 2001-2007					
AWAP	1.8	0.0	3.1	0.9	50.0
Operational		0.0	3.8	1.1	61.1
Daily rainfall 1900-2000					
AWAP	2.0	0.0	3.7	1.2	59.6
Operational		0.0	3.9	1.3	63.4

for this variable. We note that these verification results are not directly comparable to those provided by Jones and Weymouth (1997) (nor, for example, Jeffrey et al. 2001), because of slightly different cross-validation procedures and different verification periods. Beesley and Frost (2009) provide a more direct comparison of these different data-sets.

Table 4. Verification statistics for (a) monthly and (b) daily vapour pressure. The units are hPa.

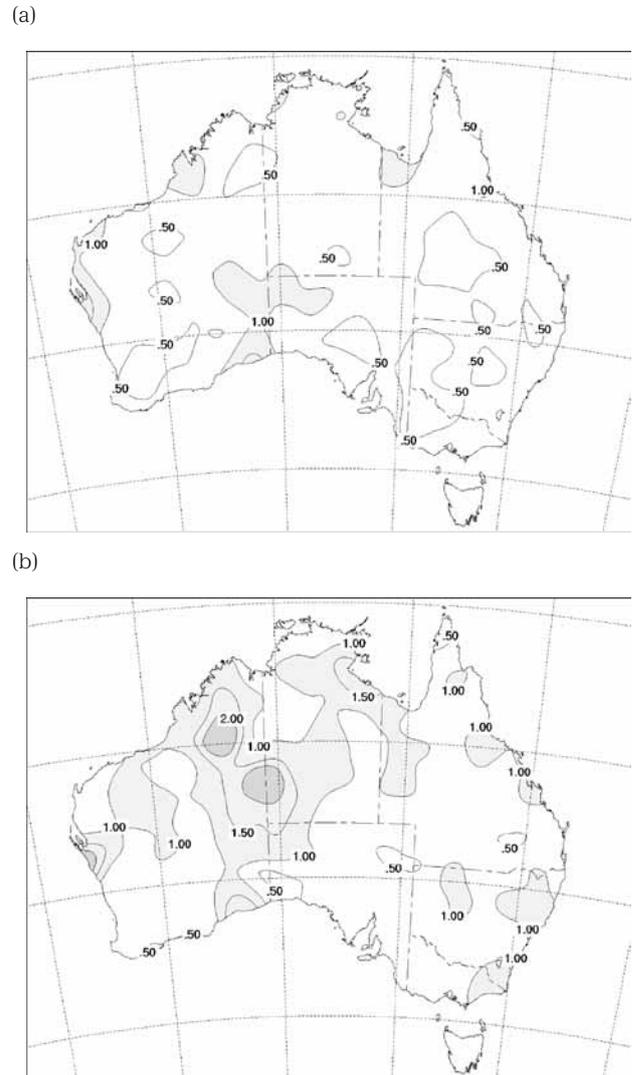
(a)				
	<i>Mean</i>	<i>Bias</i>	<i>RMSE</i>	<i>MAE</i>
Monthly vapour pressure 2001-2007				
AWAP 9 am	13.7	0.0	1.1	0.8
AWAP 3 pm	13.1	-0.1	1.7	1.1
Monthly vapour pressure 1971-2000				
AWAP 9 am	13.9	0.0	1.0	0.7
AWAP 3 pm	13.5	0.0	1.3	0.9
(b)				
	<i>Mean</i>	<i>Bias</i>	<i>RMSE</i>	<i>MAE</i>
Daily vapour pressure 2001-2007				
AWAP 9 am	13.7	0.0	1.8	1.2
AWAP 3 pm	13.1	-0.1	2.5	1.6
Daily vapour pressure 1971-2000				
AWAP 9 am	13.9	0.0	1.8	1.2
AWAP 3 pm	13.5	0.0	2.5	1.6

Maximum and minimum temperatures

The RMSEs for monthly maximum and minimum temperatures are typically between 0.5 and 1°C, while those for daily temperatures are a little larger reflecting shorter length scales and larger station error variances (Jones and Trewin 2000a). There is clearly a strong correspondence between station density (Fig. 2) and the analysis errors (Figs 5 and 6), with the largest errors tending to occur in the poorly observed western interior. There is also a tendency for large errors to occur where the variance of temperature is greater (Jones and Trewin 2000b).

The new analyses are a substantial improvement on the current Bureau of Meteorology practice for maximum and minimum temperatures at both the monthly and daily timescales. For maximum temperatures, the RMSE is reduced by around 40 per cent for daily and nearly 60 per cent for monthly analyses. The percentage improvement for minimum temperatures is smaller but still substantial, being around 0.5°C as measured by the RMSE. We note that the MAEs (not shown) tend to be a little smaller than the RMSEs, indicating a modest (positive) skewness in the analysis errors.

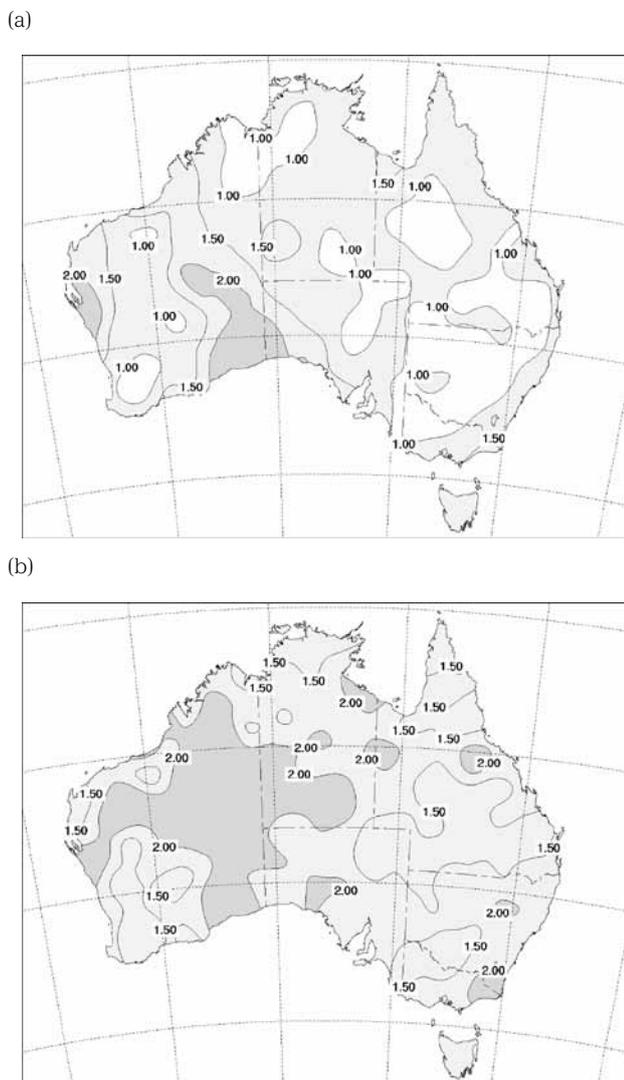
The improvement in the temperature analyses compared to current Bureau practice is quite general across Australia. The most substantial improvements are in regions of significant topography. For example, near the Victorian Alps and Snowy Mountains the RMSE is reduced from more than 2°C (not shown) to around 0.6°C for daily data. Similar improvements have been reported for California (USA) by Hunter and Meentemeyer (2005), who used topography-resolving techniques like ours.

Fig. 5 Cross-validated root mean square error for monthly (a) maximum and (b) minimum temperatures for the seven years 2001-2007. The units are °C.

The spatial maps of analysis error for maximum and minimum temperatures (Figs 5 and 6) show little evidence of increased values near significant topography. This confirms that the anomalies at both the daily and monthly timescale tend to be nearly barotropic, supporting the two-step approach which we have adopted.

Spatially the RMSE represents a trade-off between the temperature variance, the network density and the local difficulty of analysis (due to errors of representativeness and variations in correlation length scales). Locally large RMSEs highlight regions where spatial analysis is particularly difficult or the network insufficiently dense. There is some evidence that analysis errors for maximum temperature are larger near the coast around northwest Australia and about the Nullarbor Plain, with the areas near Shark Bay and Eucla standing out in particular. These two coastal regions often experience

Fig. 6 Cross-validated root mean square error for daily (a) maximum and (b) minimum temperatures for the seven years 2001-2007. The units are °C.



very strong gradients in maximum temperatures between the coast and inland deserts, and are difficult to analyse with a relatively sparse network. It is quite possible that much of the coast of Western Australia and parts of the Northern Territory experience similar analysis issues during the warmer part of the year, but the historical network is not sufficiently dense to show this.

Analysis errors for minimum temperature are greater than those for maximum temperatures. This is because minimum temperatures tend to have larger errors of representativeness and shorter length scales (e.g. Jones and Trewin 2000a). In addition minimum temperatures often show complex and variable relationships with topography (e.g. Trewin 2005). The weaker link between altitude and minimum temperature also largely explains the lesser improvement over

the current operational practice. These factors combined imply that a denser network is required for minimum temperature to achieve the same analysis accuracy as that for maximum temperature. This will be clearly important for analysing individual events such as frosts where a difference of 1°C to 2°C may be very significant in terms of impact.

We note that the RMSEs for monthly maximum and minimum temperatures are now not much larger than the theoretical lower bounds calculated by Jones and Trewin (2000a, 2002) in parts of inland eastern Australia where the station network is most dense. This suggests that in these regions future improvements will require the use of very different analysis procedures and/or new data-sets (particularly with lower representativeness errors), such as those obtained by remote sensing or dynamical models.

Figure 7 shows the historical variation of analysis RMSE from 1911 through to 2007 for both the daily and monthly data. Most obviously there is a general improvement in the analysis accuracy through time as the station density improves (see Fig. 1). The improvement of daily maximum and minimum temperature errors is relatively monotonic up until about 1970; after that time subsequent changes are slight, consistent with the findings of Jeffrey et al. (2001). Interestingly, the abrupt increase in available daily data in the 1950s has only a modest improvement in the overall analysis accuracy, confirming the results of Jones and Trewin (2002) that daily temperature analyses are only slightly improved when going beyond a network of about 100 stations.

A somewhat surprising property of the analyses is a tendency for the monthly analyses to show larger RMSE errors in recent years. There are a number of reasons for the increase in the errors. A partial explanation is the decline in the amount of available monthly data from a 1970s peak, which has only recently been reversed. This means that in recent years a significant number of stations lack a robust estimate of the station climate normal.

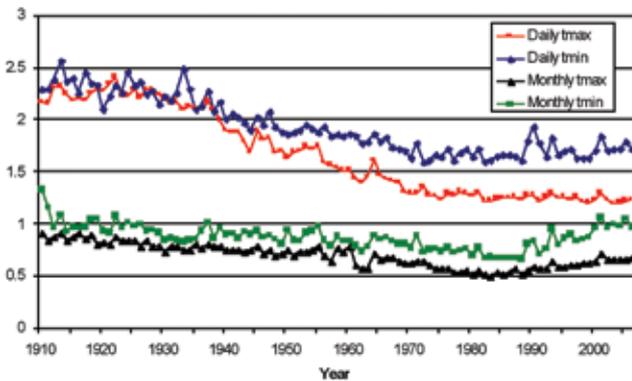
The interaction between spatial analysis and the strong warming observed over Australia (Fig. 12) is also contributing to the evolution of the RMSE and the increase in values towards the end of the series. A warming trend means that the variance of temperature will tend to increase away from the centre of each of the 30-year average periods (see Fawcett and Jones 2007). This will tend to lead to slightly increased RMSEs.

An additional factor is the spread of stations into remote coastal and alpine locations as automatic weather stations have come to dominate the station network. These more remote locations tend to have considerably larger analysis errors as a result of unique and often complex microclimates.

Monthly and daily rainfall

Table 3 and Fig. 8 provide verification statistics for the monthly and daily rainfall analyses, including a comparison with current operational practice. The monthly rainfall analyses show a modest but significant improvement over current Bureau practice, and the RMSEs are below those which

Fig. 7 Cross-validated root mean square error for maximum (tmax) and minimum (tmin) temperature for the full analysis period. The units are °C.



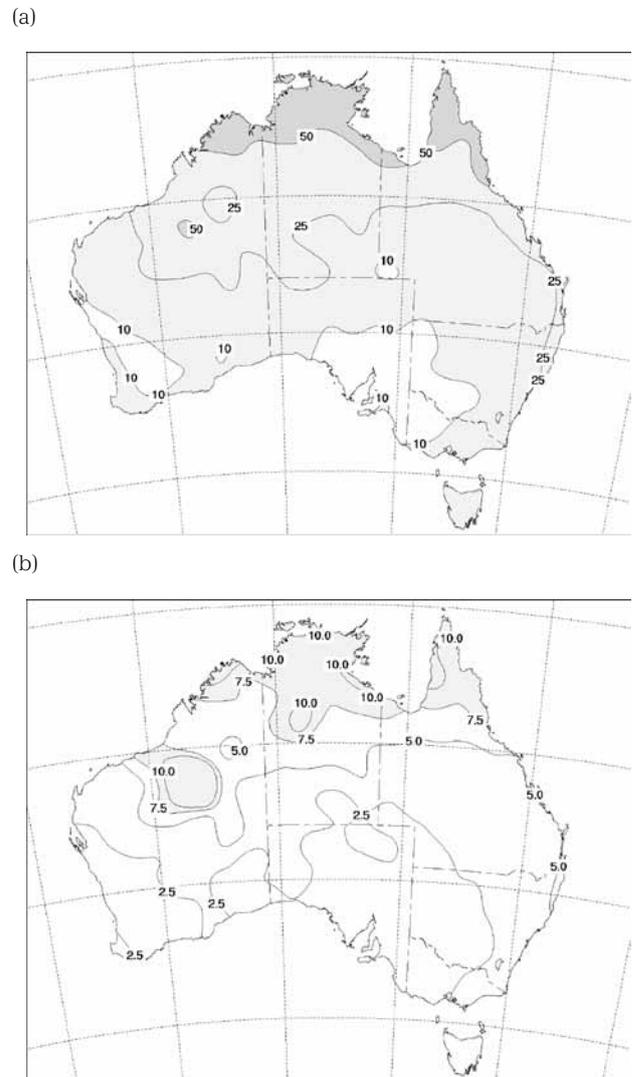
have been reported previously for Australia (Jones and Weymouth 1997; Jeffrey et al. 2001). The analysis improvement is most marked in southern Australia where the climatological signals captured in the climate normals are most robust from month to month. We note that the RMSE is substantially larger than the MAE, while the bias is small. This is because of a significant skewness in the distribution of rainfall errors, with a relatively small number of large errors, particularly in the tropical parts of Australia. This skewness provides some caution against taking a non-linear transformation of the rainfall data prior to spatial analysis (e.g. taking the square root or cube root). Such a transformation may tend to de-emphasise large rainfall events and emphasise smaller ones, leading to less accurate analyses overall.

There is a substantial north-south gradient in the RMSE for rainfall across Australia on both monthly and daily timescales (Fig. 8). In part this reflects the higher rainfall in the tropical regions which will lead to larger analysis errors for a given data smoothness and station network (e.g. Daley 1993). This pattern has been noted previously by Mills et al. (1997), Jones and Weymouth (1997) and Jeffrey et al. (2001). This pattern is further amplified by the tendency for rainfall to be highly convective in tropical parts and hence to have shorter and more variable length scales (e.g. Mills et al. 1997; Ebert et al. 2007).

The RMSEs for the historical period are a little better than those for the Bureau's current operational system and those reported previously for Australia (Mills et al. 1997; Jeffrey et al. 2001; Beesley and Frost 2009). The insensitivity of the errors to the analysis method for the earlier years is somewhat surprising, given that the underlying analysis systems are very different.

For the most recent years (2001-2007), the new analyses are rather better than those for the operational system, supporting the observations of Beesley and Frost (2009). A year-by-year comparison of errors (not shown) shows that there

Fig. 8 Cross-validated root mean square error for (a) monthly and (b) daily rainfall for the seven years 2001-2007. The units are mm.



is rather little difference between the quality of the AWA and operational analyses prior to the 1950s, but thereafter the AWA analyses show an improvement which grows over time. This improvement coincides with a significant expansion of the rainfall network and an increase in observations at higher elevations. Hunter and Meentemeyer (2005) found modest impact from including climate-topography relationships in daily rainfall analyses in California.

While the improvement is positive we note that the analysis errors remain large with the MAE being 50 per cent of the average daily rainfall amount. A possible way of improving the analyses might be to develop rainfall-altitude relationships (climatologies) which are conditional on weather type, such as light wind convective situations versus strong upslope flow situations.

We also note that the analysis errors for daily rainfall are only weakly dependent on the Barnes parameters obtained through the optimisation process described by Seaman (1989). Mills et al. (1997), Weymouth et al. (1999) and Jeffrey et al. (2001) found similar insensitivities in accuracy for their analyses of daily rainfall. We interpret this as indicating that the length scales for rainfall vary markedly from day to day (and also spatially), and hence are not well-approximated by a single parameter set. It is clear that further substantial improvements in daily rainfall analyses will require either much denser networks or the use of remotely sensed and/or model-derived data (e.g. Ebert et al. 2007), thereby introducing a substantial inhomogeneity with respect to historical analyses.

The annual mean RMSE for monthly rainfall by year together with the all-Australian annual mean rainfall is shown in Fig. 9. The time series for daily rainfall (not shown) shows a relatively monotonic decline from near 4 mm to near 3 mm, overlaid on substantial interannual variability.

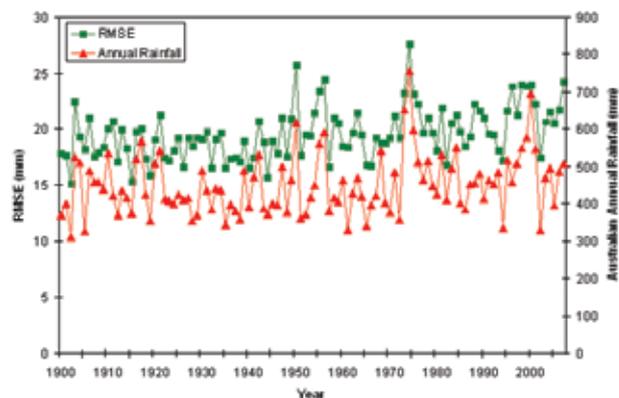
The RMSEs by year are dominated by a tendency for errors in the monthly and daily data (not shown) to be larger in wet years than in dry ones, as noted previously by Jeffrey et al. (2001). There is some evidence for a slight reduction in analysis error as the available rainfall network expands through the first 20 years of the twentieth century, with relatively less change thereafter. There is also some evidence for a recent slight increase in errors, which is likely to be associated with a slight reduction in the station network. In addition, recent years have been characterised by very wet conditions in large parts of northern Australia compensated for by low rainfall in southern parts (Bureau of Meteorology 2008b). It would be expected that such a pattern of rainfall will lead to increased analysis errors overall as analysis is more difficult in the north of Australia than in the south.

Monthly and daily vapour pressure

Figure 10 shows the distribution of analysis errors for monthly and daily vapour pressure (at 0900). The errors for 1500 are similar (see Table 4) though they do tend to be somewhat larger. Following the climatology, the vapour pressure analysis errors increase towards the north where average values are substantially higher (Jones et al. 2007). There is also evidence of somewhat increased errors close to the coast, where gradients often tend to be large between moist maritime air and drier continental air, in agreement with Jeffrey et al. (2001). The lowest analysis errors are found in the well-sampled southeast and southwest parts of Australia.

These vapour pressure analyses are the first of their type to be produced by the Bureau of Meteorology, and consequently they cannot be directly compared to existing analyses. Comparison with Jeffrey et al. (2001) suggests these analyses provide a slightly lower RMSE and MAE overall (1.4 hPa versus around 1.5 hPa). We note that direct comparison is not possible given different analysis periods. An

Fig. 9 Cross-validated root mean square error for monthly rainfall together with the Australian annual mean rainfall. The units are mm.



important property is the absence of inflated errors near topography. This suggests that the vapour pressure/altitude relationship is rather robust and amenable to the two-step anomaly analysis method we have developed.

Application of the data to climate change

The data which we described have been developed for the accurate description of daily and monthly climate over Australia in a way which is consistent with long-term climatology. An obvious application of these data is in documenting long-term trends and climate change. In this section we compare our new AWAP analyses with the most widely-used Bureau of Meteorology climate change data-sets for rainfall and temperature. We focus on national annual averages of the data as well as trend maps as described by Jones et al. (2004).

The data-set described by Jones and Weymouth (1997) is currently used by the Bureau of Meteorology for describing the historical variation in local spatially averaged rainfall, while trend maps for rainfall use the station data of Lavery et al. (1992, 1997) with updates (see Jones et al. 2004). Updates to these two sets are made as data become available in the climate data bank (ADAM).

The annual average rainfall (taken from the monthly grids) across the whole of Australia in the 0.05° AWAP data and the data of Jones and Weymouth (1997) has a correlation of 0.995, while the trends as calculated by least-squares regression for 1900 to 2007 are +0.66 mm/year and +0.75 mm/year, respectively. This reveals very good agreement between these nationally averaged data on both annual and longer timescales, with a modest wetting trend in both.

Figure 11 shows the spatial map of the linear trend in annual rainfall over the 1900-2007 period in units of mm/decade for the Jones et al. (2004) and AWAP data. The Jones et

Fig. 10 Cross-validated root mean square error for (a) monthly and (b) daily 9 am vapour pressure for the seven years 2001-2007. The units are hPa.

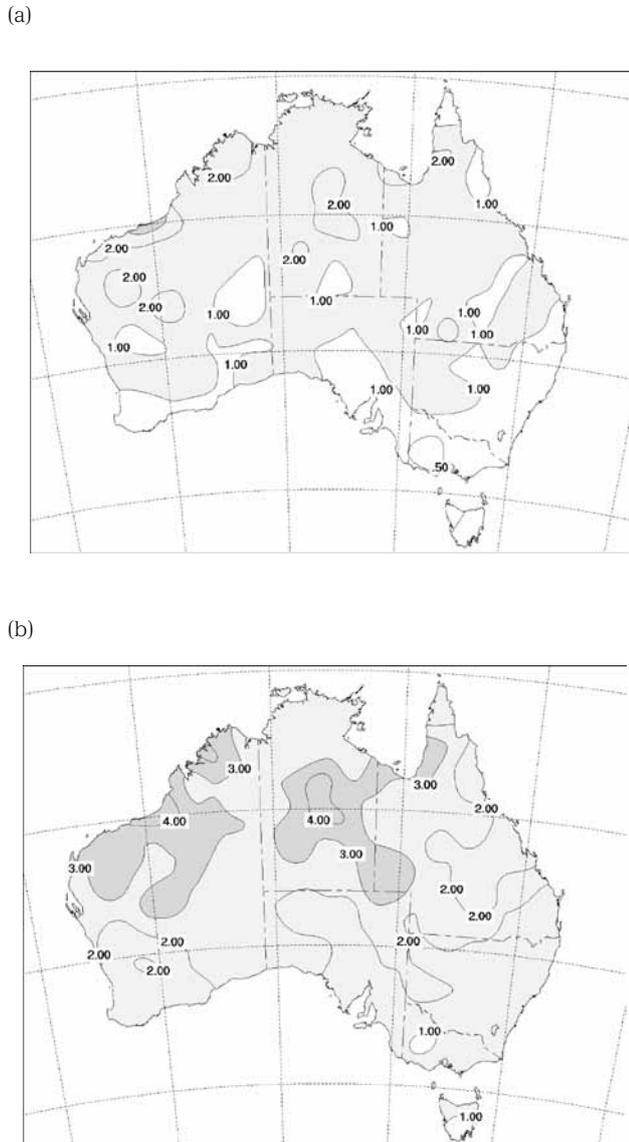
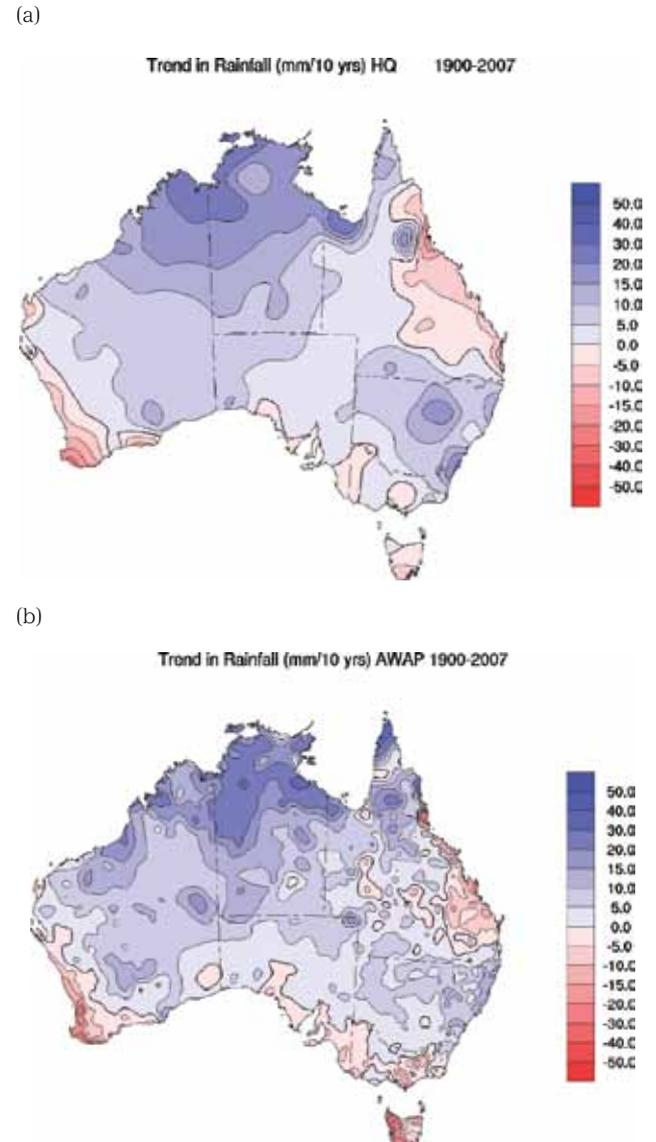


Fig. 11 The linear trend in annually summed rainfall based on (a) the 0.25° Jones et al. (2004) data and (b) the monthly AWAP data. Units are mm/decade.



al. (2004) data include 184 stations with records sufficiently long to provide trend estimates, and as a result yield far smoother spatial trend maps. Overall, there is reasonably good agreement and in many places the AWAP data appear to be more physically consistent; for example they show marked and consistent drying in those parts of southern Australia where rainfall tends to occur from mid-latitude westerlies (including southwest Western Australia, western Tasmania, and the highlands of Victoria and New South Wales). In contrast, the Jones et al. (2004) analyses show trends through these regions which are less well linked to topography and geography.

A notable difference in the analyses is the strong drying in western Tasmania in the AWAP data. Jones and Beard (1998) have previously described the difficulty in producing historical analyses in this region due to the lack of stations in the first half of the twentieth century which can lead to an artificial wetting trend. A comparison with the available somewhat fragmented station data in this region (not shown) does suggest that this area may have experienced substantial drying over the course of the last century as suggested by the AWAP data.

The slightly enhanced wetting trend in the AWAP data for inland Western Australia appears to be associated with net-

work changes and, in particular, the opening of the remote meteorological station at Giles in 1956. There are no data in this broad region prior to the 1950s, and so an extrapolation of spatially remote information for the first half of the trend period is required. These earlier interpolated data appear to slightly underestimate the actual rainfall in this region. This local trend highlights that it will never be possible to fully remove the effects of network change on analyses of historical climate change.

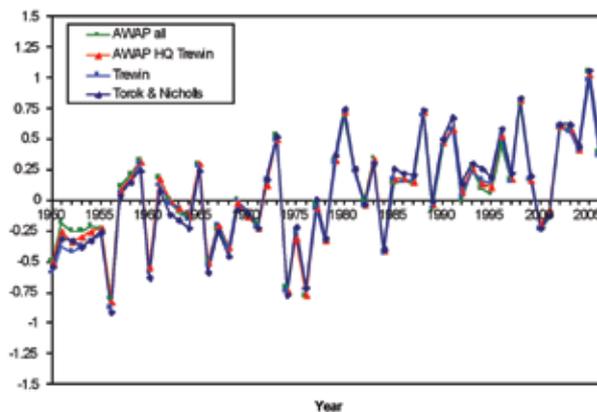
The use of raw temperature data for the monitoring of climate change is somewhat problematic owing to a range of non-climatic effects. These include local changes in station environments (e.g. urbanisation), changes to observation practice (e.g. the introduction of daylight saving) and changes to instrument exposure. There currently exist two widely used homogenised temperature data-sets for Australia, one being homogenised at the annual timescale (Torok and Nicholls 1996; Della-Marta et al. 2004) and the other at the daily timescale (Trewin 2001; Jones and Trewin 2002). The first of these currently yields annually resolved gridded data back to 1910, while the second yields monthly resolved gridded data back to 1950. The lack of pre-1950s data in the second set relates to the relative lack of digitised daily temperature data for use in local data homogenisation (Fig. 1). Both data-sets are available in gridded form on a 0.25° national grid covering Australia (Jones et al. 2004).

For climate change applications, we have developed a version of the AWAP monthly maximum and minimum temperature grids by drawing only on the homogenised data of Trewin (2001) with updates (a total of 109 stations) analysed onto a 0.25° grid. The stations used for this grid-set have long and relatively homogeneous records, meaning that their use minimises the impact of network changes and artificial inhomogeneities.

In Fig. 12 we show the Australian annual mean temperature (as a departure from the 1961-1990 average) using the station data of Trewin (2001) and Torok and Nicholls (1996) together with the new 0.05° AWAP data and the 0.25° resolution AWAP data (termed here the 'high-quality' set) using the station data from Trewin (2001) to form the anomaly grids. While a national average clearly smooths out much detail, these four data-sets show good agreement over this 58-year period.

There is a tendency for the full AWAP data-set to show slightly less warming than the high-quality data-sets. The lesser warming is in part due to smaller rates of warming over inland Western Australia as a result of network changes since 1950 which introduce a slight cool bias as the stations are relatively high and cool. The difference is somewhat ameliorated when our new methodology is applied to the station data of Trewin (2001). The use of the climate normal grids in the AWAP data-set, however, means that there is some lingering dependency on network changes in this region. These results suggest that the all-Australian temperature is a very robust statistic which is insensitive to the use of lower quality temperature data and urban sites.

Fig. 12 The all-Australian annual mean temperature (anomaly from 1961-1990 average) for the AWAP data, the AWAP data using the Trewin (2001) station set, and the Trewin (2001) and Torok and Nicholls (1996) data (from Jones et al. 2004).



Elsewhere the trends (not shown) are similar though the full AWAP temperature data tend to contain substantially more detail as a result of the larger number of stations. At the large scale at least, these results suggest that the new data are robust for defining trends, except in regions of large network changes, and data voids.

Summary and conclusions

In this paper we have provided a detailed description of a series of new meteorological analysis products developed by the Australian Bureau of Meteorology as a contribution to the Australian Water Availability Project. Careful attention has been paid to developing systems and data-sets which are robust and useful for the monitoring of both climate variability and climate change. These systems are now running in real time and are expected to form the basis for ongoing monitoring and mapping of Australia's climate by the Australian Bureau of Meteorology.

The analyses make use of a new two-step analysis system which partitions the analysis field into a climatological component and an anomaly component. This approach has been found to be robust, to preserve the background climatology in the long term and to be computationally efficient. These systems produce a substantial improvement on existing Bureau practice as measured by error statistics.

It is acknowledged that the accuracy of analyses will be limited in regions where the station network is insufficient to resolve detail, particularly on the daily timescale. Maps of the RMSE do reveal some tendency for errors to increase where climate gradients are tight, such as in coastal areas where the network may not be sufficient to resolve maritime effects.

In designing these systems we have deliberately confined the analysis to the use of *in situ* data, as the introduction of modern data-sets such as from satellites may lead to homo-

geneity issues. Clearly, better analyses are possible for the more recent period making use of these new data types, particularly under a multivariate analysis paradigm. A focus for future development will be in developing these systems.

There are ongoing issues which have emerged through this study and which will be the focus of future development and work. Foremost, there is a need to improve the daily rainfall analyses, for which all currently available Australian analyses have rather poor accuracy. The evidence is that this will require either very different analysis techniques which make use of data not currently used (such as from remote sensing and numerical weather prediction) or a substantial improvement in the national rain-gauge network. There is also a clear need to maintain consistent high-quality data networks across Australia, as the present analysis shows that even temporary declines in networks can have substantial impacts on our ability to monitor climate.

Acknowledgments

This work was undertaken as a contribution to the Australian Water Availability Project. This project was a partnership between the Bureau of Rural Sciences, CSIRO, and the Bureau of Meteorology.

References

- Australian Bureau of Agriculture and Resource Economics 2007. *Australian Crop and Livestock Report*, 30 October 2007 Drought Update. Available from Australian Bureau of Agricultural and Resource Economics, Canberra, ACT, Australia, 11 pp.
- Australian Water Availability Project 2004. *Work plan 2004-2006*. Bureau of Rural Sciences, 12 pp.
- Beesley, C. and Frost, A. 2009. *A comparison of the BAWAP and SILO spatially interpolated daily rainfall datasets*. Bureau of Meteorology, Melbourne, Australia.
- Bureau of Meteorology 2000. *Climatic Atlas of Australia - Rainfall*. Available from Australian Bureau of Meteorology, Melbourne, Victoria, Australia, 25 pp.
- Bureau of Meteorology 2008a. *Annual Climate Summary 2007*. Available from Australian Bureau of Meteorology, Melbourne, Victoria, Australia, 16 pp., http://www.bom.gov.au/climate/current/statement_archives.shtml.
- Bureau of Meteorology 2008b. *Special Climate Statement 16: Long-term rainfall deficiencies continue in southern Australia while wet conditions dominate the north*. Available from Australian Bureau of Meteorology, Melbourne, Victoria, Australia, 9 pp., <http://www.bom.gov.au/climate/current/>.
- Chiew, F., Vaze, J., Viney, N., Jordon, P., Perraud, J.-M., Zang, L., Teng, J., Arancibia, J.P., Morden, R., Freebairn, A., Austin, J., Hill, P., Wisemfeld, C. and Murphy, R. 2008. *Rainfall-runoff modelling across the Murray-Darling Basin: A report to the Australian Government from the CSIRO Murray-Darling Basin Sustainable Yields Project*. CSIRO, Australia, 70 pp.
- Cline, W.R. 2007. *Global Warming and Agriculture: Impact Estimates by Country*. Washington: Center for Global Development and Peterson Institute for International Economics, 178 pp.
- Compo, G.P., Whitaker, J.S. and Sardeshmukh, P.D. 2006. Feasibility of a 100-Year Reanalysis Using Only Surface Pressure Data. *Bull. Am. Met. Soc.*, *87*, 175-90.
- Cressman, G.P. 1956. An operational objective analysis system. *Mon. Weath. Rev.*, *87*, 367-74.
- Daley, R. 1993. *Atmospheric data analysis*. Cambridge University Press, New York, USA, 457 pp.
- Della-Marta, P.M., Collins, D.A. and Braganza, K. 2004. Updating Australia's high-quality annual temperature dataset. *Aust. Met. Mag.*, *53*, 75-93.
- Ebert, E., Janowiak, J.E. and Kidd, C. 2007. Comparison of near-real-time precipitation estimates from satellite observations and numerical models. *Bull. Am. Met. Soc.*, *88*, 1-18.
- Fawcett, R. and Jones, D. 2007. When does the global warming signal come to dominate interannual temperature variability? *Bull. Aust. Met. Oceanogr. Soc.*, *20*, 120-6.
- Hunter, R.D. and Meentemeyer, R.K. 2005. Climatologically aided mapping of daily precipitation and temperature. *Jnl Appl. Met.*, *44*, 1501-10.
- Hutchinson, M.F. 1995. Interpolating mean rainfall using thin plate smoothing splines. *Int. J. Geog. Inf. Sys.*, *9*, 385-403.
- Jeffrey, S.J., Carter, J.O., Moodie, K.B. and Beswick, A.R. 2001. Using spatial interpolation to construct a comprehensive archive of Australian climate data. *Environ. Model. Software*, *66*, 309-30.
- Jones, D.A. and Beard, G.S. 1998. Verification of Australian monthly district rainfall totals using high resolution gridded analyses. *Aust. Met. Mag.*, *47*, 41-54.
- Jones, D.A. and Trewin, B.C. 2000a. On the description of monthly temperature anomalies over Australia. *Aust. Met. Mag.*, *49*, 261-76.
- Jones, D.A. and Trewin, B.C. 2000b. On the relationships between the El Niño-Southern Oscillation and Australian land surface temperature. *Int. J. Climatol.*, *20*, 697-719.
- Jones, D.A. and Trewin, B.C. 2002. On the adequacy of historical Australian daily temperature data for climate monitoring. *Aust. Met. Mag.*, *51*, 237-50.
- Jones, D., Collins, D., Nicholls, N., Phan, J. and Della-Marta, P. 2004. A new tool for tracking Australia's climate variability and change. *Bull. Aust. Met. Oceanogr. Soc.*, *17*, 65-9.
- Jones, D.A., Wang, W., Fawcett, R. and Grant, I. 2006. The generation and delivery of Level-1 historical climate data sets. *Australian Water Availability Project Milestone Report*. Bureau of Meteorology, Melbourne, Australia, 32 pp.
- Jones, D.A., Wang, W., Fawcett, R. and Grant I. 2007. Climate Data for the Australian Water Availability Project. *Australian Water Availability Project Milestone Report*. Bur. Met., Australia, 37 pp.
- Jones, D.A. and Weymouth, G. 1997. An Australian monthly rainfall data set. *Technical Report No. 70*, Bur. Met., Australia, 19 pp.
- Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., Zhu, Y., Chelliah, M., Ebisuzaki, W., Higgins, W., Janowiak, J., Mo, K.C., Ropelewski, C., Wang, J., Leetmaa, A., Reynolds, R., Jenne, R. and Joseph, D. 1996. The NCEP/NCAR 40-year reanalysis project. *Bull. Am. Met. Soc.*, *77*, 437-71.
- Koch, S.E., DesJardins, M. and Kochin, P.J. 1983. An interactive Barnes objective map analysis scheme for use with satellite and conventional data. *Jnl Clim. Appl. Met.*, *22*, 1487-503.
- Lavery, B., Joung, G. and Nicholls, N. 1997. An extended high-quality historical rainfall dataset for Australia. *Aust. Met. Mag.*, *46*, 27-38
- Lavery, B., Kariko, A. and Nicholls, N. 1992. A historical rainfall data set for Australia. *Aust. Met. Mag.*, *40*, 33-9.
- Mills, G.A., Weymouth, G., Jones, D., Ebert, E.E., Manton, M., Lorkin, J. and Kelly, J. 1997. A national objective daily rainfall analysis system. *Bureau of Meteorology Development Report No. 1*, Bur. Met., Australia, 30 pp.
- Murray-Darling Basin Commission 2007. *River Murray System - Drought Update No. 7*, 8 pp. Available from http://www.mdbc.gov.au/rmw/drought_updates.
- Murray, F.W. 1967. On the computation of saturation vapour pressure. *Jnl Appl. Met.*, *6*, 203-4.
- Rayner, D., Moodie, K., Beswick, A., Clarkson, N. and Hutchinson, R. 2004. *New Australian daily historical climate surfaces using CLIMARC*. Available from Department of Natural Resources, Mines and Energy, Queensland, Brisbane, Queensland, Australia, 76 pp.
- Rayner, D. 2007. Wind run changes: the dominant factor affecting pan evaporation trends in Australia. *Jnl Climate*, *20*, 3379-94.
- Raupach, M.R., Briggs, P.R., Haverd, V., King, E.A., Paget, M. and Trudinger, C.M. 2008. *Australian Water Availability Project*. CSIRO Marine and Atmospheric Research Component. Final Report for Phase 3. Available from CSIRO, Canberra, Australia, 67 pp.

- Renzullo, L.J., Toscas, P., Ahmed, R. and Sutton, G. 2008. *Precipitation and actual evapotranspiration products for water resources applications* (gridded precipitation data intercomparison), CSIRO. Available from CSIRO, Canberra, Australia, 48 pp.
- Seaman, R.S. 1989. Tuning the Barnes objective analysis parameters by statistical interpolation theory. *J. Atmos. Ocean. Tech.*, 6, 993-1000.
- Seaman, R.S. and Hutchinson, M.F. 1985. Comparative real data tests of some objective analysis methods by withholding observations. *Aust. Met. Mag.*, 33, 37-46.
- Torok, S.J. 1996. The development of a high quality historical temperature data base for Australia. PhD Thesis, School of Earth Sciences, Faculty of Science, University of Melbourne, Australia.
- Torok, S.J. and Nicholls, N. 1996. An historical annual temperature dataset for Australia. *Aust. Met. Mag.*, 45, 251-60.
- Trewin, B. 1997. Another look at Australia's record high temperature. *Aust. Met. Mag.*, 46, 251-6.
- Trewin, B.C. 2001. Extreme temperature events in Australia. PhD Thesis, School of Earth Sciences, Faculty of Science, University of Melbourne, Australia.
- Trewin, B.C. 2005. A notable frost hollow at Coonabarabran, New South Wales. *Aust. Met. Mag.*, 54, 15-21.
- Wahba, G. and Wendelberger, J. 1980. Some new mathematical methods for variational objective analysis using splines and cross validation. *Mon. Weath. Rev.*, 108, 1122-43.
- Weymouth, G., Mills, G.A., Jones, D.A., Ebert, E.E. and Manton, M.J. 1999. A continental-scale daily rainfall analysis system. *Aust. Met. Mag.*, 48, 169-79.
- Weymouth, G.T. and Le Marshall, J.F. 2001. An operational system to estimate global solar exposure over the Australian region from satellite observations 1. Method and the initial climatology. *Aust. Met. Mag.*, 48, 181-95.
- World Meteorological Organization 1989. Calculation of monthly and annual 30-year standard normals. *WCDP-No. 10, WMO-TD/No. 341*. Washington, D.C., U.S.A.
- Xie, P., Yatagai, A., Chen, M., Hayasaka, T., Fukushima, Y., Liu, C. and Yang, S. 2007. A Gauge-Based Analysis of Daily Precipitation over East Asia. *J. Hydrometeor.*, 8, 607-26.