Harrington, L. J., Gibson, P. B., Dean, S. M., Mitchell, D., Rosier, S. M., & Frame, D. J. (2016). Investigating event-specific drought attribution using self-organizing maps. *Journal of Geophysical Research: Atmospheres*, 121(21), 12766–12780.
<https://doi.org/10.1002/2016JD025602>

Publisher's PDF, also known as Version of record

Link to published version (if available):
[10.1002/2016JD025602](https://doi.org/10.1002/2016JD025602)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via AGU at <http://onlinelibrary.wiley.com/doi/10.1002/2016JD025602/abstract;jsessionid=49AD77052627E6A469366FD49B670509.f04t02> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

RESEARCH ARTICLE

10.1002/2016JD025602

Key Points:

- Evaluating human influence on specific drought events is challenging with present generation of climate models
- Method of event-specific drought attribution is proposed by using self-organizing maps
- Robust anthropogenic increase found in observing daily circulation features consistent with the 2013 NZ drought

Supporting Information:

- Supporting Information S1

Correspondence to:

L. J. Harrington,
luke.harrington@vuw.ac.nz

Citation:

Harrington, L. J., P. B. Gibson, S. M. Dean, D. Mitchell, S. M. Rosier, and D. J. Frame (2016), Investigating event-specific drought attribution using self-organizing maps, *J. Geophys. Res. Atmos.*, 121, 12,766–12,780, doi:10.1002/2016JD025602.

Received 30 JUN 2016

Accepted 15 OCT 2016

Accepted article online 18 OCT 2016

Published online 11 NOV 2016

Investigating event-specific drought attribution using self-organizing maps

Luke J. Harrington¹, Peter B. Gibson^{2,3}, Sam M. Dean⁴, Daniel Mitchell⁵, Suzanne M. Rosier⁴, and David J. Frame^{1,4}
¹New Zealand Climate Change Research Institute, School of Geography, Environment and Earth Sciences, Victoria University of Wellington, Wellington, New Zealand, ²Climate Change Research Centre, University of New South Wales, Sydney, New South Wales, Australia, ³ARC Centre of Excellence for Climate System Science, University of New South Wales, Sydney, New South Wales, Australia, ⁴National Institute of Water and Atmospheric Research, Wellington, New Zealand, ⁵Environmental Change Institute, University of Oxford, Oxford, UK

Abstract Previous studies evaluating anthropogenic influences on the meteorological drivers of drought have found mixed results owing to (1) the complex physical mechanisms which lead to the onset of drought, (2) differences in the characteristics and time scales of drought for different regions of the world, and (3) different approaches to the question of attribution. For a midlatitude, temperate climate like New Zealand, strongly modulated by oceanic influences, summer droughts last on the order of 3 months, and are less strongly linked to persistent temperature anomalies than continental climates. Here we demonstrate the utility of a novel approach for characterizing the meteorological conditions conducive to extreme drought over the North Island of New Zealand, using the January–March 2013 event as a case study. Specifically, we consider the use of self-organizing map techniques in a multimember coupled climate model ensemble to capture changes in daily circulation, between two 41 year periods (1861–1901 and 1993–2033). Comparisons are made with seasonal pressure and precipitation indices. Our results demonstrate robust (>99% confidence) increases in the likelihood of observing circulation patterns like those of the 2013 drought in the recent-climate simulations when compared with the early-climate simulations. Best guess estimates of the fraction of attributable risk range from 0.2 to 0.4, depending on the metric used and threshold considered. Contributions to uncertainty in these attribution statements are discussed.

1. Introduction

The growing area of event attribution research seeks to understand whether and to what extent the likelihood of an extreme weather event, often one which has been observed in the real world and resulted in damaging impacts, has increased in response to anthropogenic influences on the climate system [Stott *et al.*, 2016]. While the emergent evidence from previous studies suggests that a large proportion of recent heat-related extremes were made more likely by human influences on the climate system [Peterson *et al.*, 2013; Herring *et al.*, 2014, 2015; Fischer and Knutti, 2015], the anthropogenic influence on precipitation-related extremes, and particularly drought events, is less clear [National Academies of Sciences (NAS), 2016].

The mechanisms which contribute toward the onset of a drought event are complex and require the understanding of coupled land-atmosphere feedback processes in order to be accurately characterized [Dai, 2011b; Seneviratne, 2012]. The wide range of metrics employed in the peer-reviewed literature to quantify the severity of drought [Heim, 2002; Zargar *et al.*, 2011] is indicative of the difficulty faced by the research community in characterizing drought events. Given this difficulty in even quantifying the severity of an observed drought event, attempting to approach the question of attribution is exceptionally challenging.

1.1. Previous Research

Previous approaches toward drought attribution have found differing results for the same “event,” depending on which meteorological aspects of the observed drought event were considered for analysis. The majority of studies which have considered changes to the likelihood of extreme precipitation deficits (often linked to a subsequent drought event), for example, observed events in Australia [King *et al.*, 2014], Brazil [Otto *et al.*, 2015], the Horn of Africa [Marthews *et al.*, 2015], or the Central U.S. [Rupp *et al.*, 2015], did not detect any statistically significant differences between model simulations which included observed anthropogenic influences on the climate system and those which did not. Two independent studies did, however, find

human-induced increases in the likelihood of exceptionally low seasonal precipitation accumulations over the Middle East region [Bergaoui *et al.*, 2015; Kelley *et al.*, 2015], a result which can be explained as a unique regional response to the weakening of the Mediterranean storm tracks with climate warming, resulting in enhanced eastern Mediterranean drying through circulation changes [Hoerling *et al.*, 2012].

While precipitation deficits are a dominant factor in droughts [Williams *et al.*, 2015], there are nevertheless other critical nonlinearities and couplings in the system that are more difficult to characterize, including land-atmosphere feedback on small scales [Seneviratne *et al.*, 2010]. Some studies have looked at the co-occurrence of extreme temperatures and precipitation deficits: both King *et al.* [2014] and Rupp *et al.* [2015] consider this link in the context of whether increased dryness exacerbates the likelihood of concurrent hot years occurring but find no discernible change in the likelihood of annual precipitation deficits. Diffenbaugh *et al.* [2015] meanwhile notes that the record dryness over California for the 2012–2014 drought event coincides with record high-temperature anomalies and subsequently demonstrates that the increased chance of these extremes co-occurring has led to an increased probability of extreme drought over the region. Williams *et al.* [2015] goes one step further and decomposes the relative contributions of changing temperatures and precipitation patterns toward observing the exceptionally low recorded values of the Palmer Drought Severity Index, again with respect to the 2012–2014 California drought.

Other approaches have used multithousand member initial-condition ensembles to constrain an estimate of changes to simulated precipitation minus simulated evaporation ($P-E$) and found no attributable change in the likelihood of record low values for the 2014/2015 Brazil drought [Otto *et al.*, 2015]. Interestingly, the use of a coupled land surface model in an otherwise equivalent model ensemble approach by Rupp *et al.* [2015] actually showed a decrease in the likelihood of exceptional soil moisture deficits over the summer months for Texas, despite the large (at least tenfold) anthropogenic increase in the probability of record hot temperatures also occurring.

Simulating the causal chain of processes which leads to the occurrence of severe hydrological or agricultural drought is a nontrivial exercise [Seneviratne *et al.*, 2010; Orth *et al.*, 2015], and the ability of coupled climate models to realistically simulate the coupled mechanisms which contribute to a severe drought is a fundamentally limiting constraint [Burke and Brown, 2008; Dai, 2011b; Lorenz *et al.*, 2016]. While only looking at precipitation deficits as a key meteorological determinant of drought is unsatisfactory, simulating a more complex measure of drought severity will likely lead to mischaracterizing uncertainty when ascribing a change in event probability to anthropogenic influences. An alternative proposal suggested here is to primarily focus on the large-scale circulation patterns, which occur over the course of a given drought event. Without making prior assumptions about which specific circulation types are most conducive to drought, we assess how close modeled summers come to replicating the exact frequency and persistence of the circulation regimes which were observed during a specific drought event.

1.2. Event-Specific Attribution Using Self-Organizing Maps: Framing the Question

When considering probabilistic event attribution, there is a question of utility in how specifically tailored toward an observed event the analysis should be, in order for the result and any subsequent attribution statement to be suitably unbiased [NAS, 2016; Stott *et al.*, 2016]. The most common approach is for a specific event to actually be generalized as a class of event, whereby some metric is used to characterize the type of event in question (extreme daily temperatures, for example), and probabilistic estimates of the fraction of attributable risk (hereafter FAR [Allen, 2003]) are made by then considering occurrences or exceedances of a threshold which was observed for the actual event [Shepherd, 2016]. This way, attribution statements can have relevance for the occurrence of similar types of events in the future.

While explored less frequently, one can also consider the in-depth characteristics specifically related to the observed event only [Hannart *et al.*, 2016]—examples include using analogue approaches in an attempt to understand differences between dynamical and thermodynamic effects on a given event probability [Yiou *et al.*, 2007]. In this study, we propose to use self-organizing map (SOM) techniques to characterize the synoptic circulation of each day over the 3 month period of January to March 2013 (herein JFM13), during which the most severe drought on record occurred over the New Zealand region [Harrington *et al.*, 2014]. Each day is assigned 1 of 12 possible circulation regime types (hereafter SOM nodes), based on which pattern most closely resembles the modeled circulation for that day: these 12 SOM nodes are identified by using reanalysis

Figure 1. (a) MSLP anomalies over the New Zealand domain for January–March 2013. Anomalies are taken from the 20th Century Reanalysis Project v2c with respect to a 1981–2010 climatology. The domain size corresponds to that used for subsequent self-organizing map analyses. The blue and green rectangles show, respectively, the subregions used for the area-averaged MSLP and precipitation analysis in section 3.2.2. (b) Monthly self-calibrated PDSI values (see Dai [2011a] for details), averaged over the blue rectangle in Figure 1a, taken at the end of March for each calendar year. Negative values indicate soil moisture deficits. The red circle indicates value for 2013.

data over the period of 1979–2014 and are intended to characterize the range of typical summer circulation patterns. We will then look at how closely each individual summer in an ensemble of fully coupled climate model simulations from the Coupled Model Intercomparison Project Phase 5 (CMIP5 [Taylor *et al.*, 2012]) resembles the specific frequency and persistence of each SOM node which was observed for the JFM13 observed event. Hence, rather than evaluating the probabilistic change to some metric, which is intended to characterize drought over the North Island, this approach evaluates whether the specific sequence of synoptic circulation patterns observed over the summer of JFM13 was made more or less likely to occur, as a result of human influences on the climate system.

1.3. The 2013 North Island Drought Case Study: Event Description

Beginning in early January 2013, the region of the southwest Pacific over New Zealand experienced an unusual sequence of persistent, or “blocking,” high-pressure anticyclones [Harrington *et al.*, 2014]. Over the subsequent 3 month period to the end of March, the northern half of the country experienced the most extreme drought over 41 years of observations, both in terms of the severity of soil moisture deficit (when characterized by using a potential evapotranspiration deficit metric) and the areal extent over which record dryness was experienced [Harrington *et al.*, 2014] (Figure S1 in the supporting information). At the peak of the event, it was estimated that only 20 days of water supply remained for the capital city of Wellington (<http://bit.ly/1sQwD6u>; Accessed 26 May 2016). Subsequent economic analysis suggests that the drought reduced the country’s annual gross domestic product by approximately 0.6 percentage points, equating to at least US\$1.3 billion in lost revenue [New Zealand Treasury, 2013].

Monthly mean sea level pressure (MSLP) data, using ensemble-mean data from the 20th Century Reanalysis Project version 2c (hereafter 20CR [Compo *et al.*, 2011]), show anomalously high pressure situated over the New Zealand region for the summer of 2013 (Figure 1a), though not record breaking [Harrington *et al.*, 2014]. Despite this, the self-calibrated Palmer Drought Severity Index ($PDSI_{SC}$ [Dai, 2011a]) shows that the summer of 2013 exhibited record low levels of soil moisture over the North Island region (Figure 1b), consistent with the results using nationwide potential evapotranspiration deficit estimates. For subsequent analysis of MSLP and precipitation deficits in section 3.2, we use the regions marked as the green (37°S–45°S, 163°E–178°E) and blue (34°S–40°S, 172°E–178°E) boxes in Figure 1a, respectively—these correspond to the regions where anomalies of each corresponding variable were most extreme over the duration of the event (Figure S2).

2. Methods

2.1. Self-Organizing Maps

Self-organizing maps (SOMs) are a powerful technique to characterize the synoptic climatology of a given region and provide a method of quantifying how well climate models represent day-to-day variability in typical synoptic situations. In this study, a sequential SOM is used with a random initialization scheme, following *Hewitson and Crane [2002]*: a random initialization scheme was chosen over other methods (such as in *Reusch et al. [2005]*) to minimize computational memory requirements while having a negligible effect on the resultant SOM configurations [*Gibson et al., 2016a*].

The SOM training procedure involves the presentation of input data cases (in this case the daily MSLP data for the region), randomly chosen and presented individually, to iteratively update each node. The Euclidean distance measure between data cases and each node is then used to determine the “winning node,” consistent with previously published approaches [e.g., *Cassano et al., 2006; Alexander et al., 2010*]. Both the winning node and surrounding nodes are updated in this process by using a Gaussian neighborhood update function. The learning rate and radius parameter in the SOM were set to decrease linearly to zero and one, respectively. Sensitivity testing determined that the starting values of these parameters had negligible influence on the final maps produced. The reader is referred to *Hewitson and Crane [2002]* for further discussion and justification of methodological choices of the SOM procedure.

The end result of this process is a map of “self-organized” nodes where the order of these nodes on the two-dimensional array map relates to a degree of similarity between nodes (Figure 2), although we emphasize that the equidistant configuration of these nodes in grid form should not be interpreted as a quantitative measure of similarity between neighboring nodes. The MSLP patterns shown across the nodes are designed to preserve and span the typical range of daily spatial patterns found in the data. An attractive feature of this methodology is that node patterns trained from reanalysis data can then be used in the evaluation of climate models. For example, at each daily time step the MSLP field in a particular model is allocated to one of the SOM nodes trained through reanalysis.

We calculate SOMs over a spatial domain around New Zealand (25°S–55°S, 155°E–190°E) by using daily 20CR MSLP data over the months of January to March, for the period of 1979–2014: only training the SOM nodes with summertime MSLP data enables a more in-depth consideration of features specific to the 2013 drought event. We have chosen to use 20CR reanalysis data for this analysis, primarily because *Gibson et al. [2016b]* have demonstrated that 20CR is a reliable reanalysis product for studying large-scale circulation features over the New Zealand domain for the period of analysis considered. Comparisons with alternative reanalysis products are also briefly considered in section 2.3.

A configuration of 12 possible SOM nodes is selected for subsequent analysis, although alternative configurations of 20 and 30 nodes demonstrate very similar results (see supporting information). We then consider three metrics to represent the collective properties of synoptic circulation over each summer season: node frequency, average lifetime, and maximum lifetime. Node frequency is defined as the total sum of days when a given node occurs, expressed as a percentage; average lifetime is defined as the average duration (days) each node persists for; and maximum lifetime is defined as the longest continuous number of days over which the same node persists.

2.2. The 2013 Drought in the Context of Self-Organizing Maps

Figure 3 reveals how the SOM properties for the summer of 2013 compare from a climatological perspective. The box and whisker plots present statistics based on the aforementioned SOM metrics for each summer from 1979 to 2014: in each box, the center is the median, the boxes are the upper and lower quartiles, while the maximum whisker length is 1.5 times the interquartile range (following *Tukey [1977]*). The blue line in Figures 3a and 3b represent the SOM metrics when all summer days from all years are aggregated together. When compared with the climatological average, we see that the JFM13 (red circles) exhibited anomalously high frequencies of nodes a0, a1, c2, and c3, although none of these individual frequencies were the highest on record. In terms of persistence metrics, the average lifetime scores for 2013 were broadly consistent with the climatological average for all nodes—while node c3 represented an exception, there were still several other years which recorded exceptionally longer average lifetimes. The patterns which exhibited anomalously high (low) node frequencies also showed similar statistics for maximum node lifetime.

Figure 2. SOM nodes trained from 20CR daily MSLP for the months January to March over the period 1979–2014. In terms of referencing, node “a3” refers to the node in the lower left corner of the SOM plane. Inset numbers refer to the number of days each node was observed over the duration of the JFM13 drought.

2.3. Evaluating CMIP5 Models for Attribution Using SOMs

A more recent application of self-organizing maps has been to assess which models in the CMIP5 archive best simulate the observed synoptic climatology over a given region [e.g., Gibson *et al.*, 2016a]. Here we consider how this could be used to assess the climatology of summertime circulation over the New Zealand region for each model.

For 22 models from CMIP5, we compile daily MSLP data over the same 36 year period as the available reanalysis data, using “Historical” data from 1979 to 2005 concatenated with “Representative Concentration Pathway 8.5” (“RCP8.5”) simulations over the period of 2006–2014. For stability of the SOM procedure, a second-order conservative remapping scheme was employed to regrid all models and reanalyses to a common $1.5^\circ \times 1.5^\circ$ grid. By considering only the 3 months of JFM over which the SOM nodes were trained,

Figure 3. Box plots showing, for each SOM node, the observed (a) frequency, (b) average lifetime, and (c) maximum lifetime for all summers between 1979 and 2014. In each box, the center is the median, the boxes are the upper and lower quartiles, while the maximum whisker length is 1.5 times the interquartile range. All outliers beyond this 1.5*IQR range are shown as black crosses. The red filled circles show the observed statistics for JFM2013. The blue lines correspond to the climatological mean.

corresponding SOM time series are calculated for each model by identifying the closest resemblance to each SOM node for a given day. Based on these data, we consider (1) the total number of days each node is observed, comparing between each CMIP5 model and the reanalysis over the climatology period, as well as the (2) average and (3) maximum lifetime of each node, and calculate the mean absolute error (MAE) for each of these three metrics. MAE is defined as follows,

$$MAE = \frac{1}{n} \sum_{i=1}^n |f_i - y_i| \quad (1)$$

where n denotes the number of SOM nodes under consideration (default = 12), y_i refers to the frequency (or lifetime) of node i in the reanalysis, and f_i denotes the corresponding frequency (or lifetime) for the CMIP5 model of interest. An MAE of zero would indicate a perfect comparison in the frequency/lifetime of SOM nodes between a CMIP5 model and reanalysis. We have chosen to focus on mean absolute errors as the metric of choice for subsequent analysis, as this formula provides an absolute measure of model error and gives equal weight to errors from all SOM node types. While Pearson correlations and Spearman rank

Figure 4. (a) Spread in CMIP5 model performance based on the MAE scores, comparing the climatological node frequency with (a) average and (b) maximum persistence comparisons with 20CR data. A MAE of zero indicates a perfect match in the chosen metric between the model and reanalysis. Numbers refer to those CMIP5 models listed in Table S1 in the supporting information. The black circles correspond to comparison between 20CR and two reanalysis data sets: ERA-Interim and NCEP2.

correlations have also been considered (see supporting information), results for subsequent applications in section 3.1 were found to be less consistent across different SOM node configurations.

Figure 4 demonstrates that this approach used to compare the 20CR reanalysis with each CMIP5 model, as well as two alternative reanalysis products (ERA-Interim [Dee et al., 2011] and National Centers for Environmental Prediction Reanalysis 2 (NCEP2) [Kanamitsu et al., 2002]), across the full 36 year climatology. Based on this figure, it is apparent that some models do a better job than others at simulating climatological node characteristics of 20CR, although no model performs as well as the alternative reanalysis products. This latter point is expected, given that the CMIP5 simulations only have constrained boundary conditions (like external forcing rates) and are otherwise freely evolving, while the reanalysis data sets are, like 20CR, constrained by using actual observations over the entire period. It is also noted that the combination of (1) very low and (2) highly similar MAE scores found for the two alternative reanalysis products shows that the choice of reanalysis product has a negligible influence on the results of this SOM-based analysis.

An important step of any attribution study is to determine which models are a reasonable surrogate for reality and thus may be considered suitable for subsequent analysis. Based on the results of Figure 4, it might be argued that some models perform poorly relative to the model ensemble as a whole, and therefore may warrant exclusion from subsequent analysis of anthropogenic changes in circulation. However, further consideration reveals that the MAE scores of Figure 4, which use the full climatology period and using all nodes, do not necessarily translate to consistent model biases in the frequency distribution for each individual node (Figure S6). For this reason, and because of how we intend to utilize SOMs in an attribution context (section 2.4), we have chosen to include all models for subsequent analysis. However, our results are comparable even if the four models which are particularly “poor performers” based on Figure 4 (bold numbers) are excluded from the subsequent model ensemble (Figure S7).

2.4. Application of SOMs in the Context of Event-Specific Attribution

Evaluating the fidelity of a model’s synoptic climatology is only one application of self-organizing maps. Instead of aggregating all available years to compare climatological frequencies of each node type with the observed climatology of the same region, MAE scores can also be computed to compare each individual summer of a model with an observed summer in the real world. For example, in addition to the MAE score obtained by comparing the summertime climatology of 20CR over 1979–2014 with that of model “12” (Geophysical Fluid Dynamics Laboratory-Earth System Models version 2M (GFDL-ESM2M)), one can also calculate the MAE between the node frequency/lifetime of each summer in the model and the summer of 2013 in the reanalysis (filled circles shown in Figure 5). These individual MAE scores can be used to quantify the likelihood of a model simulating a summer with synoptic circulation features analogous to those of JFM13, and thus also be used to investigate changes in the probability of simulating summers like that of JFM13 in response to anthropogenic influences.

Figure 5. As in Figure 4 but showing the climatological MAE scores for a single model (number “12”, denoting GFDL-ESM2M) decomposed into the MAE scores (filled circles) between each summer of the model over 1979–2014 and the summer of 2013 in the reanalysis.

To consider changes to the probability of the summer of 2013 occurring in response to anthropogenic influences on the climate system, we now consider the output from the same CMIP5 simulations (under the same Historical and RCP8.5 experiments) for two time periods: 1861–1901 (hereafter “early-climate”) and 1993–2033 (hereafter “recent-climate”). The more recent time period, centered on the summer of 2013, is treated as a representative analogue of the present-day climate, while the earlier period is intended to represent a climate with a much smaller influence from anthropogenic factors. We assume that possible differences in natural forcings between the two periods will be negligible when smoothed over a 41 year time scale, and therefore consider differences between the two distributions as attributable to human influences on the climate system. The combination of using a 41 year aggregation period, and considering many different CMIP5 models, each of which are freely evolving, enables a further assumption that possible modes of decadal variability [Henley *et al.*, 2015] will have negligible influence on these results. Finally, it is important to note that the relative influences of ozone depletion versus greenhouse gas increases cannot be easily distinguished by using the model simulations available, and we therefore choose to consider “anthropogenic” changes to the climate as the aggregate signal of both effects.

For each model summer (JFM), we calculate the frequency, average lifetime, and maximum lifetime of each SOM node and compare with the corresponding characteristics for the observed summer in 2013. For each summer in each model, MAE scores are then computed for each of the three metrics, by comparing the value for each of the 12 nodes with the corresponding observed value in the summer of 2013. By calculating MAE scores to compare each modeled summer with JFM13 for the 22-model subset, we obtain a 41×22 element probability distribution (PDF) of MAE scores for each of the two periods considered.

To test the statistical significance of changes to each probability distribution between the two periods, we apply a two-sample Kolmogorov-Smirnov (K-S) test, and where appropriate, also calculate the estimated FAR for a given threshold as $\text{FAR} = (p_1 - p_0)/p_0$, where p_0 is the probability of exceeding a prescribed threshold in the early-period ensemble and p_1 denotes the probability of exceedance in the recent-period ensemble [Allen, 2003; Stott *et al.*, 2004]. We also repeat all calculations of FAR by resampling an equal number of data points from the distribution, with replacement, 10,000 times, and present the best guess (50th percentile) estimate, as well as the 10th and 90th percentiles of these bootstrapped distributions.

3. Results

3.1. Changes to the Likelihood of the JFM13 Drought From a Daily Circulation Perspective

Figure 6 shows the PDFs of aggregated model MAE scores for the early-climate and recent-climate ensembles, for each of the three SOM metrics, presented by using a kernel smoothing function. The distribution of frequency MAE scores appears to be near-Gaussian, while the distributions of average lifetime and maximum lifetime MAE exhibit heavier tails.

Figure 6. Probability distributions of the (left) frequency, (middle) average lifetime, and (right) maximum lifetime correlation scores for the 22-model ensemble. PDFs are estimated by using a kernel density smoother. The black lines correspond to the 1861–1901 ensemble; the red lines correspond to the 1993–2033 ensemble. Inset P values are found by using a two-sample K-S test.

The key question of this SOM-based approach to event-specific attribution is how these probability distributions have changed when comparing the early-climate model ensemble with the recent-climate ensemble centered on 2013. We find statistically significant ($p < 0.05$) shifts toward lower MAE scores for all three of the SOM metrics when comparing the recent-climate PDF (red) with the early-climate PDF (black). Similar results are also found when considering different SOM node configurations (see supporting information). This is a significant result and demonstrates that an anthropogenic signal can be found in the daily-scale properties of circulation related to a drought event. Specifically, these results suggest that the properties of circulation patterns coincident with the 2013 New Zealand drought were more likely to occur in response to anthropogenic influences on the climate system.

While the frequency MAE scores show a systematic shift in the entire distribution, the lifetime metrics instead show a smaller distribution tail and a more pronounced statistical mode. Understanding how these distribution changes can be interpreted in the context of a more traditional FAR-based framework is considered in more detail in section 4.

3.2. Comparison With Alternative Attribution Techniques

For both completeness, and to understand how these SOM-based results compare with more orthodox approaches to event attribution, here we present alternative assessments of human influence on the 2013 New Zealand drought, using methods that have been previously used in the peer-review literature. Specifically, we consider the frequency of only specific SOM node types, under the assumption that the notable prevalence of some node types during the drought event may suggest that they are more conducive to drying than others. We also compare the daily-scale circulation changes with changes in the distribution of an area-averaged, seasonal circulation metric and precipitation deficit metric.

3.2.1. Changes to the Frequency of Anomalous SOM Nodes

Some previous studies have investigated anthropogenic influences on changing circulation patterns by counting the number of occurrences of a particular circulation regime type, especially if this circulation pattern has been demonstrated as being anomalously prevalent during an extreme event of interest. For example, Schaller *et al.* [2016] considers the number of days in the UK wet winter of 2013/2014 when a south-westerly flow regime occurred, as record-breaking occurrences of this regime type coincided with record-breaking flooding over the region. Here we present an analogous approach for the 2013 drought, by counting the number of days which exhibited either SOM node “c2” or “c3.” As mentioned in section 2.2, the combined frequency of these two nodes was the highest over the 36 year reanalysis period in the summer of 2013.

Figure 7. (a) PDFs of the percentage of days in each summer which correspond to either SOM node c2 or c3. The black lines correspond to the 1861–1901 model ensemble; the red lines correspond to the 1993–2033 ensemble. The red circle corresponds to the 2013 value of 35.6%. (b) FAR estimates for thresholds corresponding to between the 80th and 98th percentiles of the 1861–1901 distribution, sampled at 2-percentile intervals. The green shading indicates an 80% confidence interval based on a 10,000-sample bootstrap. Inset P value corresponds to a two-sample K-S test between the early and late ensembles: bold font indicates $p < 0.05$.

Figure 7 shows a systematic shift in the recent-climate distribution to summers with more frequent occurrences of these node patterns, which were anomalously frequent in the summer of 2013 and represent anticyclonic MSLP patterns situated over the North Island (Figure 2). Corresponding FAR estimates for a series of arbitrary thresholds (chosen between the 80th and 98th percentiles of the early-climate distribution) show a bootstrapped 10th percentile of FAR of 0.3 for moderately extreme thresholds (>25 – 30% of days per summer); however, no discernible signal is apparent at the very tail of the distribution.

3.2.2. Comparison With 3 Month Area Average Approaches

For the two 41 year periods using the same 22-model ensemble, we extract model MSLP taken as the January–March mean value averaged over the region which exhibited a maximum anomaly for the 2013 drought event (green box in Figure 1). In addition, we calculate the total precipitation accumulation (mm) over the same 3 month period, averaged over the region which exhibited the most severe low-precipitation anomalies during the summer of 2013 (blue box in Figure 1). While it is noted that dry day frequencies have also been considered as a relevant metric to evaluate the 2013 New Zealand drought [Harrington *et al.*, 2014], we choose to focus on 3-monthly cumulative rainfall totals, as this is the most common approach in the peer-reviewed literature for assessing precipitation changes in a drought attribution context [Lott *et al.*, 2013; Barlow and Hoell, 2015; Bergaoui *et al.*, 2015; Diffenbaugh *et al.*, 2015; Funk *et al.*, 2015; Marthews *et al.*, 2015; Otto *et al.*, 2015; Rupp *et al.*, 2015; Swain, 2015].

Figure 8a reveals a weak negative correlation between the MSLP of a model summer and the corresponding 3-monthly total precipitation over the same period, suggesting that higher summertime MSLP is more likely than not to result in lower-than-normal precipitation over the North Island. There is also clear evidence for a systematic shift toward higher summertime MSLP in the more recent time period, consistent with our present understanding of changes to MSLP over the New Zealand region in response to human-induced climate change [Gibson *et al.*, 2016b]. An increase in variance of 3 month precipitation totals also translates to a small subsequent increase in the likelihood of exceptionally low precipitation totals. K-S tests show statistically significant differences ($p < 0.05$) exist between the early- and recent-climate PDFs for the area-average MSLP but not for precipitation.

Interrogating these results further, FAR estimates are considered for a series of arbitrary thresholds at the tail of the early-period distribution for each variable. Figure 8b confirms that a statistically robust increase in the likelihood of exceptionally high 3-monthly MSLP is found in the recent-climate model ensemble, with the 10th percentile of FAR estimates exceeding 0.3 for all MSLP thresholds considered, except at the very tail of the distribution. Meanwhile, Figure 8c shows a small but detectable anthropogenic increase in low

Figure 8. (a) Relationship between 3-monthly MSLP and 3-monthly total precipitation averaged over the regions shown in Figure 1. The black and red circles show each model summer in the early (1861–1901) and recent (1993–2033) time periods, respectively. Least squares linear regressions have been applied to each ensemble; r -values correspond to the Pearson's correlation score. (b and c) The best guess (median) fraction of attributable risk (FAR) calculated for (b) MSLP and (c) precipitation. FAR thresholds were considered between the 80th and 98th percentiles of the 1861–1901 distribution, sampled at 2-percentile intervals for MSLP (circles), while thresholds pertaining to the 20th to the 2nd percentiles were sampled for precipitation accumulations. The blue shading indicates an 80% confidence interval based on a 10,000-sample bootstrap. Inset are P values of a two-sample K-S test between the early and late ensembles of each variable: bold color indicates $p < 0.05$.

precipitation totals across all moderately extreme percentiles, with the bootstrapped ensemble of FARs showing a more distinct anthropogenic increase in the likelihood of 3 month totals below 450 mm. These results are consistent with similar analyses which found an increased likelihood of summers with exceptionally high dry day frequencies over the North Island due to anthropogenic influences [Harrington *et al.*, 2014].

4. Discussion

By considering probability distributions of MAE scores, the traditional determination of FAR estimates (by counting the exceedances of a metric threshold based on the observed event) is not possible, since by definition, the observed event corresponds to the lower bound of the probability distribution (an MAE score of zero). Further, it is not immediately clear how best to compare human-induced changes to the distribution

Figure 9. FAR estimates calculated at percentile thresholds of the early (1861–1901) time period comparable between the six metrics used to consider the 2013 New Zealand drought. The quantile thresholds actually correspond to the 20th to the 2nd percentiles for precipitation totals and the MAE scores. The bold lines with filled circles show the median estimate of FAR for each threshold (filled circles) based on a 10,000-sample bootstrap; the thinner lines correspond to the 10th percentile of FAR. Inset P values are the same as for Figures 6–8.

ally significant differences between the early-climate and recent-climate PDFs when based on K-S tests alone (inset P values). The bootstrapped 10th percentile estimate of FAR for the area-averaged MSLP index (thin red line) remains consistently above 0.33, suggesting at least a 50% increase in the likelihood of exceptionally high seasonal mean sea level pressures over the New Zealand domain between the early- and recent-climate model runs. Similar FAR values are also found for the node count-based metric (solid black lines), which is expected given that the two metrics represent related phenomena.

It is important to emphasize that these estimates of FAR for each metric consider fundamentally different characteristics of the 2013 drought, and the variety of estimates found in Figure 9 should therefore be expected. Moreover, metrics which consider variables over seasonal time scales, like that of the area-averaged pressure and precipitation, have inherently lower variance when compared with daily scale distributions and are thus more likely to experience an earlier emergence of a FAR signal [Christiansen, 2015; King *et al.*, 2016].

The estimates of FAR for thresholds between the 80th and 90th percentile are consistent for all six metrics—of most interest is that all three MAE-based metrics demonstrate a similar best guess FAR of approximately 0.2. However, as the FAR thresholds considered extend beyond the 90th percentile, this homogeneity breaks down, and the metrics with formerly robust signals show a lower-bounded FAR which overlaps with zero (suggesting the possibility of no change in likelihood), while the precipitation PDF begins to exhibit a robust increase in the likelihood of exceptionally dry summers. This divergence is indicative of the relatively small sample sizes associated with attribution by using a CMIP5 model ensemble—it is therefore difficult to make robust statements about the fraction of attributable risk for very extreme thresholds in the absence of a much larger model ensemble.

There are some further caveats that require consideration when interpreting the results presented in this study:

1. When several metrics on different time scales are being concurrently assessed, there are open questions as to the most appropriate methods of model validation for attribution analysis [Mitchell *et al.*, 2015; Mitchell, 2016]. In this study, we consider the climatological SOM statistics to determine which models

tails of the three MAE-based metrics with the area-averaged distributions of seasonal precipitation accumulations and MSLP, as well as the “c2 + c3” node counts. To circumvent these issues, we select values corresponding to every 2nd percentile between the 80th and 98th percentiles of the “early-period” distributions for each metric, thereby enabling a comparable sequence of thresholds to consider FAR estimates. It is noted that we actually consider quantile thresholds ranging from the 20th down to the 2nd percentile for the case of 3-monthly precipitation totals and the MAE scores, but label them to be the same as the other metrics. The FAR statements shown in Figure 9 thus demonstrate attributable probability increases in the tails of each distribution, rather than providing a single FAR estimate for the 2013 event.

It is evident that the best guess estimates of FAR for comparable thresholds along the tail of the distribution for each of the six metrics differ, but all are above zero—this is consistent with the fact that five out of six metrics show statistically

replicate the real-world frequency and persistence of summertime circulation patterns, but choose to include all models for subsequent attribution, as determining which models “performed poorly” depended on the validation method used. It is also important to note that the number of models which reliably simulate the observed distributions of seasonal precipitation and MSLP over the regions used for the area-average metrics may differ from those which meet a validation criteria based on daily-scale circulation characteristics [Masson and Knutti, 2011], and this must be considered when comparing subsequent FAR estimates.

2. As demonstrated in section 3.2.1, there is a robust anthropogenic increase in the likelihood of observing those SOM nodes which occurred frequently during the 2013 drought. However, it is also important to recognize that due to the complex topography of New Zealand, each of the circulation regimes characterized by the 12 SOM nodes in this analysis can enhance drying in some parts of the country, while potentially increasing rainfall in other regions (Figures S3 and S4). Moreover, the spatial patterns of drought for the 2013 event were not completely homogeneous, with record-breaking dryness observed over the western South Island and the majority of the North Island, while the more traditionally drought-prone eastern regions of the country experienced only moderate soil moisture deficits (Figure S1). Because of these complexities in the spatial evolution of the drought, assuming that one or two particular SOM nodes were the cause of the events’ severity, as is demonstrated in section 3.2.1, could lead to an overly simplistic characterization of the unique properties of the 2013 event. It is for this reason that we emphasize the use of MAE scores for our analysis, as this technique assigns equal weight to the difference in frequency (or persistence) of each node, when comparing model summers against JFM13, and thus enables a truly event-specific method of attribution.
3. The Kolmogorov-Smirnov test is a common approach used to establish whether two probability distributions can be interpreted as being statistically different from one another, or more accurately, whether they originate from the same continuous distribution [Steinskog *et al.*, 2007]. However, it is evident that changes between the early-climate and recent-climate distributions to the number of years with very high MAE scores will be less informative than evaluating changes to the very low end of the MAE distribution, especially when asking the question of whether a summer like that of 2013 was more likely to occur in response to anthropogenic climate change. Christidis and Stott [2015], for example, only subsampled model winters which obtained a Pearson correlation score greater than 0.6, when comparing the upper level circulation characteristics with the observed wet UK winter of 2013/2014. The SOM-based FAR estimates in Figure 9 are intended to be analogous to this fixed threshold approach, and thus provide supplementary insight to the full PDFs considered in Figure 6.

5. Summary and Conclusions

Any statement regarding the fraction of attributable risk for a specific extreme event needs to be carefully interpreted in the context of the question being asked [van Oldenborgh *et al.*, 2015]. In this study, we introduce a new approach to consider whether and to what extent the specific daily circulation characteristics observed over the New Zealand region during the record-breaking drought of early 2013 were more or less likely to occur, and break that down further to separately consider the frequency, average persistence, and maximum persistence of different circulation types. The introduction of SOMs in an attribution context represents a novel contribution toward an emerging focus on event-specific attribution techniques in the research community [Hannart *et al.*, 2016]. It also represents a method of circumventing difficulties in evaluating anthropogenic changes in drought likelihood, particularly for those locations where precipitation-temperature coupling mechanisms are less significant.

This SOM-based approach fills a gap in the toolset available for event-specific attribution. While different techniques remain more suited for extremes which occur on daily or multiday time scales, the approach suggested in this study will be useful for seasonal-to-annual scale extremes which cannot be sufficiently characterized by temporally averaging a single variable (such as MSLP) over the length of the event. While the analogue approach [Stott *et al.*, 2016] represents a complementary method, applying such a technique to studies of drought still requires the original assumption that a single, specific measure of drought (such as soil moisture) could be adequately simulated by climate models over time and also be representative of the underlying mechanisms which contribute to event severity. Our approach instead focuses only on changes to the properties of daily-scale circulation over the duration of the extreme event.

The final aim of this study was to illustrate how different approaches toward event attribution can yield a variety of results even when considering the same event. For the New Zealand drought case study, our analysis indicates that there was a small but discernible anthropogenic increase in the probability of observing daily circulation characteristics like those of the summer of 2013, with best guess FAR estimates of 0.2 being found for all three SOM-based metrics when using moderate percentile thresholds. There were also robust increases in the likelihood of exceptionally high seasonal MSLP patterns over the regions which were most extreme during the 2013 drought, with a best guess estimate of a 50% increase in likelihood. We emphasize that these results are sensitive to the threshold at which FAR was calculated, and discernible changes to likelihood in the extreme tail of the distribution could not be reliably identified by using the model ensemble available.

Future work is needed to explore the potential utility of self-organizing maps in the context of event attribution, as well as to determine how an event-specific framework could best complement other common attribution techniques.

Acknowledgments

We acknowledge the World Climate Research Programme's Working Group on Coupled Modelling, which is responsible for CMIP, and we thank the climate modeling groups for producing and making available their model output. For CMIP the U.S. Department of Energy's Program for Climate Model Diagnosis and Intercomparison provides coordinating support and led development of software infrastructure in partnership with the Global Organization for Earth System Science Portals. CMIP5 data are available via <http://pcmdi.llnl.gov>. The 20th Century Reanalysis data are provided by the NOAA/OAR/ESRL PSD, Boulder, Colorado, USA, and available at <http://www.esrl.noaa.gov/psd/>. Support for the Twentieth Century Reanalysis Project version 2c data set is provided by the U.S. Department of Energy, Office of Science Biological and Environmental Research, and by the National Oceanic and Atmospheric Administration Climate Program Office. The authors thank Andrew King for helpful comments on earlier versions of the manuscript and acknowledge support from Victoria University of Wellington (L.J.H. and D.J.F.) and the New Zealand Deep South National Science Challenge (L.J.H., S.M.R., and D.J.F.).

References

- Alexander, L. V., P. Uotila, N. Nicholls, and A. Lynch (2010), A new daily pressure dataset for Australia and its application to the assessment of changes in synoptic patterns during the last century, *J. Climate*, 23(5), 1111–1126, doi:10.1175/2009JCLI2972.1.
- Allen, M. (2003), Liability for climate change, *Nature*, 421(6926), 891–892, doi:10.1038/421891a.
- Barlow, M., and A. Hoell (2015), Drought in the Middle East and Central–Southwest Asia during winter 2013/14, *Bull. Am. Meteorol. Soc.*, 96(12), S71–S76, doi:10.1175/BAMS-D-15-00127.1.
- Bergaoui, K., D. Mitchell, F. Otto, M. Allen, R. Zaaaboul, and R. McDonnell (2015), The contribution of human-induced climate change to the drought of 2014 in the Southern Levant region, *Bull. Am. Meteorol. Soc.*, 96(12), S66–S70, doi:10.1175/BAMS-D-15-00129.1.
- Burke, E. J., and S. J. Brown (2008), Evaluating uncertainties in the projection of future drought, *J. Hydrometeorol.*, 9(2), 292–299, doi:10.1175/2007JHM929.1.
- Cassano, J. J., P. Uotila, and A. Lynch (2006), Changes in synoptic weather patterns in the polar regions in the twentieth and twenty-first centuries. Part 1: Arctic, *Int. J. Climatol.*, 26(8), 1027–1049, doi:10.1002/joc.1306.
- Christiansen, B. (2015), The role of the selection problem and non-Gaussianity in attribution of single events to climate change, *J. Climate*, 28(24), 9873–9891, doi:10.1175/JCLI-D-15-0318.1.
- Christidis, N., and P. A. Stott (2015), Extreme rainfall in the United Kingdom during winter 2013/14: The role of atmospheric circulation and climate change, *Bull. Am. Meteorol. Soc.*, 96(12), S46–S50, doi:10.1175/BAMS-D-15-00094.1.
- Compo, G. P., et al. (2011), The Twentieth Century Reanalysis Project, *Q. J. R. Meteorol. Soc.*, 137(654), 1–28, doi:10.1002/qj.776.
- Dai, A. (2011a), Characteristics and trends in various forms of the Palmer Drought Severity Index during 1900–2008, *J. Geophys. Res.*, 116, D12115, doi:10.1029/2010JD015541.
- Dai, A. (2011b), Drought under global warming: A review, *Wiley Interdiscip. Rev. Clim. Change*, 2(1), 45–65, doi:10.1002/wcc.81.
- Dee, D. P., et al. (2011), The ERA-Interim reanalysis: Configuration and performance of the data assimilation system, *Q. J. R. Meteorol. Soc.*, 137(656), 553–597, doi:10.1002/qj.828.
- Diffenbaugh, N. S., D. L. Swain, and D. Touma (2015), Anthropogenic warming has increased drought risk in California, *Proc. Natl. Acad. Sci.*, 112(13), 3931–3936, doi:10.1073/pnas.1422385112.
- Fischer, E. M., and R. Knutti (2015), Anthropogenic contribution to global occurrence of heavy-precipitation and high-temperature extremes, *Nat. Clim. Change*, 5(6), 560–564, doi:10.1038/nclimate2617.
- Funk, C., S. Shukla, A. Hoell, and B. Livneh (2015), Assessing the contributions of East African and West Pacific warming to the 2014 boreal spring East African drought, *Bull. Am. Meteorol. Soc.*, 96(12), S77–S82, doi:10.1175/BAMS-D-15-00106.1.
- Gibson, P. B., P. Uotila, S. E. Perkins-Kirkpatrick, L. V. Alexander, and A. J. Pitman (2016a), Evaluating synoptic systems in the CMIP5 climate models over the Australian region, *Clim. Dyn.*, 1–17, doi:10.1007/s00382-015-2961-y.
- Gibson, P. B., S. E. Perkins-Kirkpatrick, and J. A. Renwick (2016b), Projected changes in synoptic weather patterns over New Zealand examined through self-organizing maps, *Int. J. Climatol.*, 36, 3934–3948, doi:10.1002/joc.4604.
- Hannart, A., A. Carrassi, M. Bocquet, M. Ghil, P. Naveau, M. Pulido, J. Ruiz, and P. Tandeo (2016), DADA: Data assimilation for the detection and attribution of weather and climate-related events, *Clim. Change*, 136(2), 155–174, doi:10.1007/s10584-016-1595-3.
- Harrington, L. J., S. Rosier, S. M. Dean, S. Stuart, and A. Scahill (2014), The role of anthropogenic climate change in the 2013 drought over North Island, New Zealand [in “Explaining Extremes of 2013 from a Climate Perspective”], *Bull. Am. Meteorol. Soc.*, 95(9), S45–S48.
- Heim, R. R. (2002), A review of twentieth-century drought indices used in the United States, *Bull. Am. Meteorol. Soc.*, 83(8), 1149–1165, doi:10.1175/1520-0477(2002)083<1149:AROTDI>2.3.CO;2.
- Henley, B. J., J. Gergis, D. J. Karoly, S. Power, J. Kennedy, and C. K. Folland (2015), A tripole index for the Interdecadal Pacific Oscillation, *Clim. Dyn.*, 45(11–12), 3077–3090, doi:10.1007/s00382-015-2525-1.
- Herring, S. C., M. P. Hoerling, T. C. Peterson, and P. A. Stott (2014), Explaining extreme events of 2013 from a climate perspective, *Bull. Am. Meteorol. Soc.*, 95(9), S1–S104, doi:10.1175/1520-0477-95.9.S1.1.
- Herring, S. C., M. P. Hoerling, J. P. Kossin, T. C. Peterson, and P. A. Stott (2015), Explaining extreme events of 2014 from a climate perspective, *Bull. Am. Meteorol. Soc.*, 96(12), S1–S172, doi:10.1175/BAMS-ExplainingExtremeEvents2014.1.
- Hewitson, B. C., and R. G. Crane (2002), Self-organizing maps: Applications to synoptic climatology, *Clim. Res.*, 22(1), 13–26, doi:10.3354/cr022013.
- Hoerling, M., J. Eischeid, J. Perlwitz, X. Quan, T. Zhang, and P. Pegion (2012), On the increased frequency of Mediterranean drought, *J. Climate*, 25(6), 2146–2161, doi:10.1175/JCLI-D-11-00296.1.
- Kanamitsu, M., W. Ebisuzaki, J. Woollen, S.-K. Yang, J. J. Hnilo, M. Fiorino, and G. L. Potter (2002), NCEP–DOE AMIP-II Reanalysis (R-2), *Bull. Am. Meteorol. Soc.*, 83(11), 1631–1643, doi:10.1175/BAMS-83-11-1631.
- Kelley, C. P., S. Mohtadi, M. A. Cane, R. Seager, and Y. Kushnir (2015), Climate change in the Fertile Crescent and implications of the recent Syrian drought, *Proc. Natl. Acad. Sci.*, 112(11), 3241–3246, doi:10.1073/pnas.1421533112.

- King, A. D., D. J. Karoly, M. G. Donat, and L. V. Alexander (2014), Climate change turns Australia's 2013 big dry into a year of record-breaking heat [in "Explaining Extremes of 2013 from a Climate Perspective"], *Bull. Am. Meteorol. Soc.*, 95(9), S41–S44.
- King, A. D., M. T. Black, S.-K. Min, E. M. Fischer, D. M. Mitchell, L. J. Harrington, and S. E. Perkins-Kirkpatrick (2016), Emergence of heat extremes attributable to anthropogenic influences, *Geophys. Res. Lett.*, 43, 3438–3443, doi:10.1002/2015GL067448.
- Lorenz, R., et al. (2016), Influence of land-atmosphere feedbacks on temperature and precipitation extremes in the GLACE-CMIP5 ensemble, *J. Geophys. Res. Atmospheres*, 121, 607–623, doi:10.1002/2015JD024053.
- Lott, F. C., N. Christidis, and P. A. Stott (2013), Can the 2011 East African drought be attributed to human-induced climate change?, *Geophys. Res. Lett.*, 40, 1177–1181, doi:10.1002/grl.50235.
- Marthews, T. R., F. E. L. Otto, D. Mitchell, S. J. Dadson, and R. G. Jones (2015), The 2014 drought in the Horn of Africa: Attribution of meteorological drivers, *Bull. Am. Meteorol. Soc.*, 96(12), S83–S88, doi:10.1175/BAMS-D-15-00115.1.
- Masson, D., and R. Knutti (2011), Spatial-scale dependence of climate model performance in the CMIP3 ensemble, *J. Climate*, 24(11), 2680–2692, doi:10.1175/2011JCLI3513.1.
- Mitchell, D. M. (2016), Attributing the forced components of observed stratospheric temperature variability to external drivers, *Q. J. R. Meteorol. Soc.*, 142(695), 1041–1047, doi:10.1002/qj.2707.
- Mitchell, D. M., et al. (2015), Solar signals in CMIP-5 simulations: The stratospheric pathway, *Q. J. R. Meteorol. Soc.*, 141(691), 2390–2403, doi:10.1002/qj.2530.
- National Academies of Sciences, Engineering, and Medicine, Committee on Extreme Weather Events and Climate Change Attribution, Board on Atmospheric Sciences and Climate, and Division on Earth and Life Studies (2016), *Attribution of Extreme Weather Events in the Context of Climate Change*, Natl. Academies Press, Washington, D. C.
- New Zealand Treasury (2013), Economic impacts of the drought—Budget economic and fiscal update 2013—The Treasury - New Zealand.
- van Oldenborgh, G. J., F. E. L. Otto, K. Haustein, and H. Cullen (2015), Climate change increases the probability of heavy rains like those of storm Desmond in the UK—An event attribution study in near-real time, *Hydrol Earth Syst Sci Discuss*, 2015, 13,197–13,216, doi:10.5194/hessd-12-13197-2015.
- Orth, R., M. Staudinger, S. I. Seneviratne, J. Seibert, and M. Zappa (2015), Does model performance improve with complexity? A case study with three hydrological models, *J. Hydrol.*, 523, 147–159, doi:10.1016/j.jhydrol.2015.01.044.
- Otto, F. E. L., et al. (2015), Factors other than climate change, main drivers of 2014/15 water shortage in Southeast Brazil, *Bull. Am. Meteorol. Soc.*, 96(12), S35–S40, doi:10.1175/BAMS-D-15-00120.1.
- Peterson, T. C., P. A. Stott, and S. Herring (2013), Explaining extreme events of 2012 from a climate perspective, *Bull. Am. Meteorol. Soc.*, 94(9), S1–S74, doi:10.1175/BAMS-D-13-00085.1.
- Reusch, D. B., R. B. Alley, and B. C. Hewitson (2005), Relative performance of self-organizing maps and principal component analysis in pattern extraction from synthetic climatological data, *Polar Geogr.*, 29(3), 188–212, doi:10.1080/789610199.
- Rupp, D. E., S. Li, N. Massey, S. N. Sparrow, P. W. Mote, and M. Allen (2015), Anthropogenic influence on the changing likelihood of an exceptionally warm summer in Texas, 2011, *Geophys. Res. Lett.*, 42, 2392–2400, doi:10.1002/2014GL062683.
- Schaller, N., et al. (2016), Human influence on climate in the 2014 southern England winter floods and their impacts, *Nat. Clim. Change*, 6, 627–634, doi:10.1038/nclimate2927.
- Seneviratne, S. I. (2012), Climate science: Historical drought trends revisited, *Nature*, 491(7424), 338–339, doi:10.1038/491338a.
- Seneviratne, S. I., T. Corti, E. L. Davin, M. Hirschi, E. B. Jaeger, I. Lehner, B. Orlowsky, and A. J. Teuling (2010), Investigating soil moisture–climate interactions in a changing climate: A review, *Earth-Sci. Rev.*, 99(3–4), 125–161, doi:10.1016/j.earscirev.2010.02.004.
- Shepherd, T. G. (2016), A common framework for approaches to extreme event attribution, *Curr. Clim. Change Rep.*, 2(1), 28–38, doi:10.1007/s40641-016-0033-y.
- Steinskog, D. J., D. B. Tjøstheim, and N. G. Kvamstø (2007), A cautionary note on the use of the Kolmogorov–Smirnov test for normality, *Mon. Weather Rev.*, 135(3), 1151–1157, doi:10.1175/MWR3326.1.
- Stott, P. A., D. A. Stone, and M. R. Allen (2004), Human contribution to the European heatwave of 2003, *Nature*, 432(7017), 610–614, doi:10.1038/nature03089.
- Stott, P. A., et al. (2016), Attribution of extreme weather and climate-related events, *Wiley Interdiscip. Rev. Clim. Change*, 7(1), 23–41, doi:10.1002/wcc.380.
- Swain, D. L. (2015), A tale of two California droughts: Lessons amidst record warmth and dryness in a region of complex physical and human geography, *Geophys. Res. Lett.*, 42, 9999–10,003, doi:10.1002/2015GL066628.
- Taylor, K. E., R. J. Stouffer, and G. A. Meehl (2012), An overview of CMIP5 and the experiment design, *Bull. Am. Meteorol. Soc.*, 93(4), 485–498, doi:10.1175/BAMS-D-11-00094.1.
- Tukey, J. W. (1977), *Exploratory Data Analysis*, Addison-Wesley, Reading, Mass.
- Williams, A. P., R. Seager, J. T. Abatzoglou, B. I. Cook, J. E. Smerdon, and E. R. Cook (2015), Contribution of anthropogenic warming to California drought during 2012–2014, *Geophys. Res. Lett.*, 42, 6819–6828, doi:10.1002/2015GL064924.
- Yiou, P., R. Vautard, P. Naveau, and C. Cassou (2007), Inconsistency between atmospheric dynamics and temperatures during the exceptional 2006/2007 fall/winter and recent warming in Europe, *Geophys. Res. Lett.*, 34, L21808, doi:10.1029/2007GL031981.
- Zargar, A., R. Sadiq, B. Naser, and F. I. Khan (2011), A review of drought indices, *Environ. Rev.*, 19(NA), 333–349, doi:10.1139/a11-013.