



Guidance note on best statistical practices for TOAR analyses

Kai-Lan Chang¹, Martin G. Schultz², Gerbrand Koren³, and Niklas Selke²

¹Cooperative Institute for Research in Environmental Sciences, University of Colorado Boulder and NOAA Chemical Sciences Laboratory, Boulder, CO, USA

²Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany

³Copernicus Institute of Sustainable Development, Utrecht University, Utrecht, the Netherlands

This document has been produced based on discussions in the statistics focus working group of the TOAR-II initiative, and approved by the TOAR-II Steering Committee
April 26, 2023

The aim of this guidance note is to provide recommendations on best statistical practices and to ensure consistent communication of statistical analysis and associated uncertainty across TOAR publications. The scope includes approaches for reporting trends, a discussion of strengths and weaknesses of commonly used techniques, and calibrated language for the communication of uncertainty. The focus of this guidance note is placed on trend analysis, which is expected to be the main statistical topic of interest across many TOAR-II focus working groups, but some of the recommendations and principles provided below are also valid for other applications. Recommendations are highlighted and numbered from R1 to R9.

1 Definition of trend analysis

The purpose of trend analysis can be defined as

detecting and attributing the change and its uncertainty of the statistical properties in a time series of a predefined variable.

where the statistical properties represent how we define and express the trend (e.g. mean, median or extreme percentile changes). Each component of this definition will be discussed through relevant topics and recommendations as follows: Section 2 clarifies the distinct roles of linear and nonlinear techniques in trend quantification. Section 3 compares the commonly adopted statistical properties in trend assessment. Section 4 discusses data preparation for time series analysis. Section 5 identifies the fundamental factors affecting uncertainty quantification associated with the trend estimate. Section 6 demonstrates how the reliability of trend assessment can be determined through the relationship between quantified trend and uncertainty. Section 7 discusses the change point detection of trends. Section 8 provides additional resources.

This guidance note aims to provide recommendations of tropospheric ozone trend analysis from a practical aspect, without discussing technical details, although we do provide some R and Python code for quantile regression in the Annex E. Practitioners (referring to authors of publications written in the context of the TOAR activity) who are interested in the fundamental theory of trend analysis (e.g. the statistical relationship between the sample size, magnitude of trend and uncertainty) are referred to various textbooks (Chandler and Scott, 2011; Von Storch and Zwiers, 2001; Wilks, 2011) and articles (Chang et al., 2021; Weatherhead et al., 1998).

2 Quantitative versus qualitative summary of the trend

R1: Practitioners are recommended to provide quantitative trends and should avoid drawing conclusions when the magnitude of trends and associated uncertainties are not explicitly quantified.

Trend detection techniques can be generally classified as linear or nonlinear methods. By using a large number of parameters and numerical optimization, nonlinear methods offer a great deal of flexibility in modeling the nonlinear pattern of a time series. However, the result is often too complex mathematically to quantitatively provide a simple trend value and an uncertainty value. Unless there is a clear understanding of (non-linear) processes driving the trend, quantitative non-linear trend analysis poses the risk of overfitting.

Under this circumstance, nonlinear trends can be visualized and summarized qualitatively. We provide a further critical discussion of nonlinear trend methods with an example in Annex A for interested readers. Apparent changes in the magnitude of the trend, which are often the motivation to apply non-linear methods, can usually be handled well through piecewise linear trend analysis in combination with change point detection (see Section 7).

Due to mathematical challenges, we have limited options to quantitatively summarize the trends. The common statistics for reporting linear trends include the mean (e.g. from the least squares method), median (e.g. from the least absolute deviations method or Sen-Theil method), and a range of percentiles (e.g. from the quantile regression method). These statistics represent different data characteristics of trends and will be discussed more thoroughly in the next section.

To alleviate the disconnect between the issues of linearity and general nonlinearity, previous studies have chosen to conduct trend analysis by using both nonlinear and linear trend techniques (e.g., [Cooper et al., 2020](#); [Oltmans et al., 2006](#)); the former is used to visually inspect the general nonlinearity, and investigate the appropriateness of the linear assumption, and the latter aims to provide a quantitative statement and acknowledge the uncertainty. Therefore, the application of linear trends should not imply our belief in a strictly linear assumption, rather, we have chosen this approach in order to communicate with and convey messages to a more general audience ([Crimmins, 2020](#)).

It is possible to model nonlinearity in a linear additive regression model. For short-term nonlinearity, we recommend incorporating meteorological or relevant covariates into regression-based methods to attribute the data variability (see Section 4 in [Chang et al. \(2021\)](#) for a demonstration). For long-term nonlinear change, employing piecewise linear trends for change point analysis is preferable ([Banerjee et al., 2020](#); [Weatherhead et al., 1998](#)), since the trends and uncertainties before and after the change point can be explicitly quantified (see Section 7 and Annex A). The advantage of this attribution and detection approach is that the result is fully and quantitatively interpretable. Nevertheless, this exercise requires an iterative fitting process, because the best set of covariates and the location of the change point(s) need to be identified by the practitioners (see Annex B for further discussions), and the best fitted model often varies for different time series.

R2: Visualize data series and/or try non-linear trend detection methods to assess the validity of a linear trend model before reporting linear trends quantitatively.

R3: When time series exhibit clearly visible changes in the magnitude of a trend over time, we recommend the application of piecewise linear trends in combination with a change point detection algorithm.

3 Recommendation for trend analysis technique and comparison of statistical properties across common trend techniques

R4: In TOAR-II, we recommend quantile regression (QR) for trend analysis.

Briefly, the reasons are: 1) QR is a well-suited technique for detecting heterogeneous distributional changes ([Koenker and Hallock, 2001](#)), which is often the case for free tropospheric and surface ozone as they typically show diverse percentile trends; and 2) QR is a regression-based method that allows us to incorporate covariates for attributing data variability and piecewise trends for change point analysis. The main limitation to QR is that a larger sample size is required for a valid estimation of extreme variations, which means that datasets with small or insufficient sample sizes should be limited to analysis of the median trend.

More broadly, this section will also provide the basic context regarding the commonly applied linear trend techniques, so researchers can be reminded of the implicit assumptions behind the methods and what challenges may be encountered during the manuscript peer-review process. Here we select three broadly used techniques, including generalized least squares (GLS), the Sen-Theil method (which was adopted in TOAR-I, but abandoned for TOAR-II; as a candidate of median trend estimates, it can be replaced by QR which has additional advantages discussed below), and QR (which has been incorporated into the TOAR-II database, and is recommended for use). Table 1 gives a brief comparison of statistical capabilities from different techniques, which can be addressed in two aspects:

Table 1: A brief comparison of statistical capabilities of three commonly applied techniques.

	GLS	Sen-Theil	Quantile regression ¹
Inference statistics of a trend	Mean	Median	Percentile/Quantile
Robust to small sample size		✓	*2
Robust to outliers		✓	✓
Makes use of all available information	✓		✓
Inference with non-normal distribution on residuals		✓	✓
Unique solution	✓		
Adjustment for non-IID residuals (e.g. Incorporation of autocorrelation)	✓		✓
Incorporation of covariates	✓		✓
Extend to piecewise trends	✓		✓
Extend to adaptive nonlinear trends	*3		*3
Extend to extreme percentile estimate			✓ ²

¹ Least absolute deviations method (median regression) can be considered to be a special case of quantile regression, so it is not listed independently.

² For quantile regression, the median estimate is robust to the small sample size problem, but a larger sample size is required for extreme percentile estimates (depending on applications, [Das et al., 2019](#)). See Table 2 for TOAR-specific recommendations.

³ These extensions can be made via the class of standard or quantile generalized additive models.

- *The distinctions between means, medians and other percentiles:* The relevant principles are listed in the purple section of Table 1. We can see these principles are closely related to the concept of sample medians (e.g. robust to outliers, small sample size, and non-normal error distribution) or sample means (e.g. having a unique solution, and making use of all data points). It is worth emphasizing that there is a contradiction between making use of all information and resisting the influence of outliers. For example, the Sen-Theil estimator was shown to tolerate arbitrary corruption of up to 29% of the data points without deterioration of its estimation ([Wilcox, 2011](#)). However, this property also means that up to 29% of all data are automatically ignored and treated the same as “non well-behaved” data, even if they are valid samples. Being robust to outliers is not always superior when we have proper data quality control (see Annex C for further discussions), because there is no reason to waste the information about the extreme values when those data are deemed to be valid. Nevertheless, under this circumstance quantile regression can still account for all extreme events by providing a range of percentile trend estimates, instead of merely focusing on the mean or median trends.
- *The distinctions between regression-based methods and classical non-parametric methods:* The advantages of regression-based approaches over the classical methods (e.g. Sen-Theil) are listed in the orange section of Table 1. All of these features are considered to be necessary or desirable for trend analysis, i.e., accounting for autocorrelation, incorporating covariates to attribute data variability, or integrating adaptive components to study nonlinearity.

Table 1 provides a limited comparison and ignores certain unorthodox extensions. For example, 1) one can combine bootstrapping techniques with the GLS method so the results are more robust to outliers and small sample size problems, and 2) the idea of Sen-Theil estimator can be used to solve multiple regression models using an iterative process ([Wilcox, 2011](#)), but these are beyond our current scope. By focusing on the big picture, the items in Table 1 represent the basic methodology, and can be used as a reference for the strengths and weaknesses of each method. Based on this comparison the researchers can also better understand the hidden context behind the trend detection techniques. Indeed, for the practitioners who applied GLS for their trend analysis, the least squares method is frequently questioned for its lack of robustness to outliers.

Table 2: Recommendation on default quantiles to be reported with respect to different sample sizes.

No of samples	Recommended Quantiles
< 50	50%ile (median)
50 - 90	25%ile, 50%ile, 75%ile
91 – 120	10, 25, 50, 75, 90
> 120	5, 10, 25, 50, 75, 90, 95

Analyses of extreme quantiles (minima, 1%, 2%, 98%, 99%, maxima) require special care and detailed evaluation of the statistical model and diagnostic. In fact, the generalized extreme value (GEV) or threshold exceedance models are designed for the purpose of studying the very extreme percentiles and should be preferred (Gilleland, 2020b; Krock et al., 2022).

However, the Sen-Theil method is rarely criticized for its under-appreciation of data records, since it ignores a large portion of extreme values without identifying true outliers from other valid data (a "true" outlier is a high leverage/influence data point that should be identified by specialized knowledge, instead of solely by analytical tools). Familiarity with those principles will help the practitioner make stronger arguments and avoid inappropriate or erroneous statistical applications.

To facilitate comparisons of trends and uncertainties across the TOAR analyses, it is desirable to define the default quantiles to be estimated. Table 2 provides rule-of-thumb recommendations on default percentile trends to be reported according to various ranges of sample size (i.e. the number of data points fitted in the trend model, not raw data). To avoid over-interpretation, we do not recommend reporting percentile trends in intervals narrower than 5%.

R5: Choose the specified quantiles in QR regression depending on the number of samples.

4 Data preparation for trend analysis

R6: Seasonal adjustments are required for time series analysis.

Before a demonstration of QR is made, it is desirable to briefly discuss data preparation for trend analysis. Conventional regression analysis (the least squares method) requires that the error term has a constant variance. Therefore analyses of atmospheric composition data have been primarily carried out using monthly, seasonal, or annual aggregated values. The rationale behind such aggregations is simply to reduce heterogeneous data variability and meet the homoscedasticity assumption (but monthly values are considered to be less heterogeneous than daily values, only if the size of aggregated samples is large enough or the sample frequency is appropriate). Nevertheless, from a modern statistical perspective, some important information might be lost after data aggregation, which limits our understanding regarding the causes of the data variability (see Annex Figure A6 and Chang et al. (2023) for further discussions). We thus recommend preserving as much information as possible by utilizing QR to detect heterogeneous changes.

Ozone and many other atmospheric variables have distinct seasonal and/or diurnal cycles. These known variations should be included in the statistical model that is used to derive trend estimates (see Annex C), or the data should be deseasonalized prior to the trend analysis. Failure to do so will lead to unreasonably large uncertainty estimates and may even adversely impact on the trend estimator (for example, if a time series starts at a seasonal minimum and ends at a seasonal maximum). The requirement for seasonal adjustments was originally proposed in the field of economics (see a review by Hannan et al., 1970), since then seasonal adjustments have become a necessary step in time series analysis.

Demonstration: a comparison of trend analyses for ozone data at Mace Head, Ireland:

The code and technical details for deseasonalization are provided in Annexes E and F, respectively. Here the time series is deseasonalized in advance. Now we demonstrate the applicability of QR (For comparisons with other methods, we use the monthly time series here, but QR can be properly applicable to daily values too, see Annex F): Figure 1 compares the percentile trend estimates derived by QR to a range of broadly used mean/median trend methods (Chang et al., 2021; Yu and Yao, 2017), based on the monthly surface ozone anomalies measured at Mace Head, Ireland (Cooper et al., 2020). We can clearly observe the unique feature




Figure 1: A demonstration of the difference between a range of mean/median trend methods (upper panel) and percentile trends derived from QR (lower panel), based on surface ozone anomalies measured at Mace Head, Ireland.

Loess (locally estimated scatterplot smoothing, gray curve) fit is added to show nonlinearity of the sample means. Dashed lines in the lower right panel represent the 95% confidence intervals, which are computed through the Wilcoxon test for the Sen-Theil estimator (Wilcox, 2001) and a Huber sandwich estimate for QR (Koenker and Machado, 1999). Note that LAD (least absolute deviations) is a special case of QR when only the median trend is considered.

of QR for deriving trends in addition to the central tendency. It should be noted that even though the different methods may show similar trend values, the uncertainty estimators can vary greatly (as shown in the comparison of uncertainties from GLS, Sen-Theil and QR). The main reason for an unrealistically narrow uncertainty from the Sen-Theil estimator is its assumption that residuals are independent and identically distributed (IID, see Section 5 for further discussions). Applying other reasonable techniques for estimating the mean or median change is not forbidden or even discouraged, but we recommend QR be used as the standard for trend analyses in TOAR, because it is a well-designed technique for studying heterogeneous ozone change and quantifying uncertainty across a wide range of percentiles.

5 Derivation of trend uncertainty

R7: Practitioners should provide associated uncertainty for any statistical estimation, typically the 95% confidence interval.

One of the most critical components of statistical analysis is to acknowledge the uncertainty. Every estimation must be accompanied by a quantification of the associated uncertainty (or error bar), which is used to assess the reliability of the (trend) estimate and is considered to be as equally important as the estimate (Chang et al., 2021; Cressie, 2021; Tong, 2019). In terms of linear trend analysis, the uncertainty is referred to as the standard error (SE) of the trend value. The higher the SE, the larger the uncertainty for the trend estimate. Most programming languages provide a summary of fitted results (see Figure 2 for examples of output from Python and R), which at a minimum consists of estimate, uncertainty, t -value and p -value for each parameter. The t -value is the ratio between the estimate and uncertainty (also known as signal-to-noise ratio or SNR), which is used to determine the p -value. The interpretation of the reliability of a trend based on p -value or SNR is provided in Section 6.

There are many factors affecting the uncertainty value. In general, 1) the larger sample size, 2) the smaller data variability, and/or 3) the weaker autocorrelation and heteroskedasticity, result in a smaller uncertainty of a trend. Whereas the first two factors do not directly involve the model assumptions or formulation, IID is generally too strong an assumption because autocorrelation/heteroskedasticity is common

```

>>> import statsmodels.formula.api as smf
>>> print(smf.quantreg("y ~ x", data).fit(q=0.5).summary().tables[1])
=====
              coef      std err          t      P>|t|      [0.025      0.975]
-----
Intercept    -3.2765      1.069      -3.066      0.002     -5.378     -1.175
x              0.0059      0.002       3.114      0.002      0.002      0.010
=====

> library(quantreg)
> summary(rq(y~x, tau=0.5),se="nid")

Call: rq(formula = y ~ x, tau = 0.5)

tau: [1] 0.5

Coefficients:
              Value      Std. Error t value Pr(>|t|)
(Intercept) -3.27647    1.00916   -3.24673  0.00128
x              0.00588    0.00183    3.22137  0.00140

```

Figure 2: Output of summary statistics from quantile/median regression in Python (upper) and R (lower), based on ozone anomalies (y) measured at Mace Head, Ireland (as shown in Figure 1). The estimate and standard error (SE) for linear trend (x) is in units of ppbv/month.

Note that the default method for quantifying SE in Python is based on the theoretical asymptotic distribution approximated through a kernel estimator (Hall and Sheather, 1988). R does not provide a default method for quantifying SE and it has to be specified by the user (here `se="nid"` indicates a non-IID SE computed through a Huber sandwich estimate (Koenker and Machado, 1999)). Both methods account for heteroskedasticity in the residuals, so the resulting uncertainties are similar. More options for estimating SE are available for Python and R, but these options are designed for either IID or heteroskedasticity cases, and currently none of the options explicitly address autocorrelation.

Block bootstrap inference to account for autocorrelation in the trends uncertainty

<p>Moving block bootstrap algorithm Data: a time series $y_1 \dots y_N$ Input: block length B; number of iterations J Output: an uncertainty of the trend estimate</p> <ol style="list-style-type: none"> 1. for $1 \leq j \leq J$ do 2. generate N/B random blocks from data 3. fit a trend model to the random samples 4. extract the sampled trend value b_j 5. end for 6. return the standard deviation of $\{b_j\}$ 	<p>linear trends and 95% CIs for surface ozone at Mace Head, Ireland [in units of ppb/decade]:</p> <p>OLS: 1.06 [±0.37] → [±0.50] GLS-AR1: 1.09 [±0.53] WLS: 0.75 [±0.35] → [±0.46] LTS: 0.75 [±0.34] → [±0.54] Lasso: 1.05 [±0.20] → [±0.52] Ridge: 1.03 [±0.21] → [±0.53] Sen-Theil: 0.85 [±0.05] → [±0.53] Siegel: 0.70 [±0.16] → [±0.61] QR-50th: 0.71 [±0.44] → [±0.57] M-estimator: 0.88 [±0.20] → [±0.50]</p>
--	---

Figure 3: Code outline for the moving block bootstrap algorithm for deriving the trend uncertainty from autocorrelated time series, and its impact on the trend uncertainty for different trend techniques.

Results are based on a block length $B = N^{0.25}$ (i.e. each selected block contains B consecutive values), number of samples $J=1000$, and ozone data at Mace Head, Ireland. Block bootstrapping was not applied to GLS-AR1 as it already accounts for autocorrelation. Note that QR-50th is equivalent to LAD (least absolute deviations).

in atmospheric composition time series, so the practitioners are recommended to take non-IID residuals into account. The common approaches include incorporation of an autoregressive model (Weatherhead et al., 1998), adoption of the heteroskedasticity and autocorrelation consistent SE (Kiefer and Vogelsang, 2002), prewhitening (a process to remove autocorrelation from data, Yue and Wang, 2002), or inference with block bootstrapping methods (Politis and White, 2004), otherwise the resulting uncertainty is likely to be inappropriate or underestimated. Due to mathematical challenges, among these approaches, only block bootstrap resampling and prewhitening methods are currently applicable to all trend techniques (Here we only focus on moving or circular block bootstrap. While prewhitening is generally applicable, it was shown to drastically distort data structure in some cases (Razavi and Vogel, 2018)). The usual resampling methods are unlikely to preserve autocorrelation structure, since they do not sample a series of consecutive data in the time series. Moving block bootstrap fixes this issue by making inference from several blocks of sampled data, instead of from individual random samples (Fitzenberger, 1998; Lahiri et al., 2007). Circular (moving) block bootstrap further extends the methodology by wrapping around the start and end points of the time series, so the data near the beginning and end of the record are not systematically under-sampled (Gilleland, 2020a).

In general, this bootstrapping approach assumes that the uncertainty based on statistical theory can be

Table 3: Trend reliability scale

p -value	SNR (signal-to-noise) value	Term
$p \leq 0.01$	$\text{SNR} \geq 3$	very high certainty
$0.05 \geq p > 0.01$	$2 \leq \text{SNR} < 3$	high certainty
$0.10 \geq p > 0.05$	$1.65 \leq \text{SNR} < 2$	medium certainty
$0.33 \geq p > 0.10$	$1 \leq \text{SNR} < 1.65$	low certainty
$p > 0.33^1$	$\text{SNR} < 1$	very low certainty or no evidence

¹ This boundary is meant to be fuzzy around 1/3 (Mastrandrea et al., 2010).

replaced by empirical uncertainty (meaning that this approach is not limited by some items listed in Table 1, as long as the data are representative). Figure 3 shows the code outline for the moving block bootstrap algorithm, tailored to the purpose of taking autocorrelation into account in the trend uncertainty. The circular block bootstrap algorithm can be adopted by simply wrapping the time series. The optimal block length is often difficult to determine, here we adopt $B = N^{0.25}$ (the power of the sample size to a quarter) as suggested by Hall et al. (1995). We then apply this algorithm to adjust the trend uncertainty for the ozone anomaly time series measured at Mace Head, Ireland (only GLS considers autocorrelation in Figure 1, with the AR1 coefficient of 0.34, and therefore moving block bootstrap is not applied). We can see the uncertainty increased for all techniques after autocorrelation is accounted for, and all methods now have not only similar trend values, but also similar (appropriate) uncertainties.

In terms of nonlinear trend techniques, since they do not provide an overall trend estimate, the 95% confidence intervals apply to the entire fitted curve, which represents the uncertainty range of the sample means (and should not be interpreted as the 95% range of the data variability, see Figure 1 for a demonstration).

6 Calibrated language for communicating uncertainty

R8: Gradated calibrated language should be used as a replacement for dichotomous “statistical significance”.

In cases when the trend and its uncertainty need to be reported for a time series, and/or when the trends need to be compared from multiple sites or multiple products, it is desirable to use an uncertainty scale for assessing the reliability and likelihood of the estimated trend, instead of merely dichotomizing the outcome as significant or insignificant (Wasserstein and Lazar, 2016). A comprehensive guidance note for communicating uncertainties in terms of confidence, likelihood and probability was previously developed for the IPCC Fifth Assessment Report (Mastrandrea et al., 2010); here we adopt a few simplifications (only placing the focus on the level of reliability for trend assessment and with fewer likelihood scales) to allow for consistent and precise communication regarding trend reliability across TOAR analyses (Table 3). Trend reliability can be expressed with the p -value (Wasserstein et al., 2019) or the signal-to-noise (SNR) ratio (Chang et al., 2021).

This scale is similar to the scheme used in the first phase of TOAR (Fleming et al., 2018; Gaudel et al., 2018; Mills et al., 2018), but adds another gradation for “very high certainty”, approximately corresponding to a p -value ≤ 0.01 or $\text{SNR} \geq 3$, which allows us to investigate “the world beyond $p < 0.05$ ” (Wasserstein et al., 2019) and isolate the results with greater statistical confidence for further discussions. For the outcomes with even greater likelihood, the authors can consider directly using the SNR values as a reference. Importantly, this uncertainty scale allows us to avoid using the expression “statistically significant”, which is now widely recognized as being problematic and not useful (Wasserstein et al., 2019).

7 Change point detection

R9: Apply change point detection to validate any suspicious changes of data structure.

Change point analysis is highly relevant to the detection of trend change and attribution of intervention (e.g. evaluation of effectiveness of a measure to reduce emissions). A change point does not have a unique

meaning in the literature, it can be a break or turning point connecting two data series that are considered to have different averages (a constant shift, [Weatherhead et al., 2017](#)), variances (distinct variabilities, [Beaulieu et al., 2012](#)) and/or trends ([Chen et al., 2011](#)). Although these structural changes can be identified through numerical algorithms ([Killick and Eckley, 2014](#); [Muggeo, 2008](#)), it is up to practitioners to ensure that the change point of trends is not induced by inconsistent or incompatible data (e.g. a mean shift caused by new instruments, [Weatherhead et al., 2017](#)). A visual inspection of time series anomaly plots can help us to identify different scenarios.

Factors affecting the detection of trends (e.g. seasonality and autocorrelation) are relevant to the detection of change points, so deseasonalization and incorporation of necessary components in the regression model are suggested. The change point algorithms are generally not robust to small perturbations in the data (meaning that similar change points can be repeatedly detected due to small scale meteorological variations). Given ozone's high temporal and spatial variability, practitioners are recommended to consider the change point analysis of trends only when data are sufficiently long (e.g. no more than one meaningful change point of trends within a decade).

Examples of the change point analysis of a time series with a turnaround of trends is provided in Annex A, and with an abrupt change in Annex B. Based on a scientific point of view, when a change point of data averages or trends is identified and attributable, such as new air pollution control policies taking effect or changes of site environment (e.g. shutdown of a factory), it is not due to natural variability and therefore should be addressed by the trend model (some studies refer to this as the intervention analysis, see [Weatherhead et al. \(1998\)](#)). If the cause of a change point cannot be associated with any specific factors, practitioners are recommended to investigate the time series at nearby stations (see Annex C for further discussions) and/or compare to chemistry-climate model simulations.

8 Additional Resources

A step-by-step demonstration of trend analysis using quantile regression (with Python and R code) can be found in Annex E. A recording of an introductory lecture describing quantile regression can be found in [Chang \(2022\)](#).

Acknowledgments

KLC was supported in part by NOAA cooperative agreements NA17OAR4320101 and NA22OAR4320151. NS and MGS acknowledge the European Research Council and the European Union to have funded the ERC Advanced Grant 787576 (IntelliAQ). We would like to thank Owen R. Cooper, Erika von Schneidmesser, Eric Gilleland, Daan Hubert, Arno Keppens and Davide Putero for helpful and thought-provoking suggestions which improved the content of this work.

Annex A Discussion on the differences between the ways to capture nonlinearity and the ways to summarize trends

There are many statistical methods that can be used to study the nonlinearity of a trend, such as moving averages, polynomials, Loess (locally estimated scatterplot smoothing), smoothing splines, dynamical linear modeling, vector autoregressive modeling, and empirical mode decomposition. With the exception of polynomials, these methods are considered to be the adaptive approach, i.e. no assumptions of the shape/structure of trends have to be made in advance (a series of references on this topic is provided by [Chang et al., 2021](#)). However, these techniques are often too complex mathematically (e.g. involving too many hyperparameters) to quantitatively provide a simple trend value and an uncertainty value.

Based on the above discussions, there is a large disconnect between the ways to study nonlinearity and the ways to summarize the trends. We demonstrate this disconnect in Figure A1, using monthly surface ozone anomalies measured from the mountaintop site of Zugspitze, Germany (1978-2020) as an example ([Cooper et al., 2020](#)). The major feature of this time series is that from a visual inspection a clear overall increase of ozone can be observed between 1978 and the late 1990s, and the increasing trend appears to diminish afterward. In this figure three different methods are used to investigate the trends; even though these methods are able to capture the diminishment of the trends, their interpretations can be quite different:

- *Adaptive learning methods*: Here the Loess locally weighted regression is selected to identify the multi-year fluctuations in the time series, but it does not provide a quantitative estimate of the trend. Most adaptive methods have the capability to be tuned to capture the local nonlinearity (and their local 95% uncertainty range can be estimated), so the resulting trends can be highly irregular/fluctuating (as shown by the blue Loess curve). The rationale of this smoothness/roughness tuning is to remove as much short-term (unstructured) variation as possible, while maintaining the long-term tendency (in simple terms, the goal is to find a balance between data fidelity and model complexity). Nevertheless, as noted previously, the nonlinear trends derived from the adaptive methods are often too complex to condense into a simple equation, so the results can only be summarized qualitatively, instead of quantitatively. For example, the blue curve in Figure A1 should be interpreted as follows: *the long-term change can be observed as an increasing trend between 1978 to the late 1990s, and a relatively flat trend from 1998 to 2020, with some short-term variations attached.*
- *Piecewise linear segments*: In this approach each line segment represents a linear trend tailored to a specific period (preferably for a sufficiently long term), so any nonlinearity of the resulting trends will only occur when the practitioner deems it necessary. In contrast to adaptive learning methods, the specifications of this approach need to be customized by the practitioner (e.g. identifying any possible change point(s), determining any data discontinuity due to data instability issues, etc). Nevertheless, since the practitioner fully controls the specifications, they are able to explicitly and quantitatively summarize the trend value and its uncertainty for each linear segment. For example, using the Loess smoother as guidance, we select 1998 as the change point of the trends, the red curve in Figure A1 represents *a positive slope of $0.049 [\pm 0.008]$ ppbv/month (or $5.91 [\pm 1.04]$ ppbv/decade, $p \leq 0.01$, very high certainty) over 1978-1998 and a negative slope of $-0.007 [\pm 0.007]$ ppbv/month (or $-0.81 [\pm 0.88]$ ppbv/decade, $p = 0.07$, medium certainty) over 1998-2020. It should also be noted that even though we only specify two linear segments to represent the long-term change, the short-term variations can be accounted for by incorporating relevant covariates, such as meteorological variables into a multiple linear regression ([Chang et al., 2021](#); [Wells et al., 2021](#)).*
- *Polynomials*: While the use of polynomials may be appealing due to their simplicity of calculation (any spreadsheet can fit a polynomial to a time series), except for a specific scenario when the practitioner aims to study the acceleration/deceleration rate of the changes ([Paolo et al., 2015](#)), they are not recommended for the general purpose of ozone trend detection. This approach is an intermediate compromise between adaptive methods and piecewise linear trends. As can be seen in Figure A1 (orange curve), the disadvantages of an in-between method can be obvious: 1) Polynomials are not as flexible as adaptive methods for properly capturing shorter term interannual variations, and higher order polynomials can result in a very poor fit due to a lack of measures to prevent numerical issues at




Figure A1: Comparison of trend fitting techniques from Loess locally weighted regression (with 95% confidence interval), piecewise linear segments and second order polynomials, based on the monthly nighttime ozone anomalies measured from Zugspitze, Germany (1978-2020), where t_1 and t_2 in the equation represent the segment before and after the change point (thus the coefficients are the trends in units of ppbv/month), respectively, and t is the collection of t_1 and t_2 representing the temporal index of the time series.

boundary conditions (see the supplement to Cooper et al. (2020) for further details and examples); 2) To facilitate interpretation of the results, the applications of polynomial trends are mostly limited to the orders of 2 or 3, out of a desire to keep the equation simple and avoid overfitting. However, the polynomial equations do not directly provide any interpretable trend and associated uncertainty values (in contrast to the piecewise method). While the polynomial fit in Figure A1 minimizes the distance between the data points and the curve, it does not provide a realistic representation of the long-term data variability and trends. The Loess smoother shows a peak of ozone in 1998 and a relative minimum several years later in 2008; in contrast the polynomial implies an increase since 1978 that reaches a peak in 2007. The reason for the incorrect peak location is due to the fact that a simple 2nd order polynomial has to bend symmetrically about its vertex, regardless of the heterogeneous changes of the data variability; this symmetrical bend also gives the false impression that ozone has steadily decreased since 2007, even though the Loess smoother shows a small increase. Another problem common to polynomials is that as the time series extends into the future, the fit to the early part of the time series and the peak of the trend typically shift and vary in intensity, an unrealistic scenario that does not occur when using a stable curve fitting technique such as the Loess smoother (see the supplement to Cooper et al. (2020) for further details and examples).

Annex B Advice on change point analysis of trends

Searching for the location of change points: While Figure A1 adopted piecewise linear segments to carry out change point analysis, it may leave the readers unaware of how to select the change point in an objective way (when the change point is difficult to determine visually). The common approach is to find the optimal location of a change point, such that either the fitted residuals are minimized, or the difference in magnitude of trends or SNRs before/after the change point is maximized. It should be noted that the exact change point is not critical, as long as the rough locations can be identified (e.g. the year is known, but the month is unknown, or unimportant). Once the location of the change point is found, we can evaluate if the model with the change point outperforms the model without the change point, based on various model performance




Figure A2: A demonstration of change point analysis with a structural break in the time series, based on surface ozone anomalies measured at Centennial, Wyoming. The upper panel adopted a single linear trend (brown) and a single mean shift model after the structural break (orange, without considering the trends). The lower panel adopted the piecewise linear trends without (red) and with (purple) an offset in the structural break, respectively. Note that the Loess fit (blue) assumes smooth and continuous transition between data points, so it cannot handle an abrupt discontinuity properly.

metrics (Reeves et al., 2007). However, it might be infeasible to implement such a complex algorithm across a large database such as TOAR (e.g. over ten thousand time series). To simplify the methodology, we may conclude that a change point is valid if the difference in magnitude of SNRs before and after the change point is maximized and exceeds a certain threshold (e.g. $dSNR = |SNR(\text{before}) - SNR(\text{after})| \geq 2$).

Although various learning methods were proposed to automatically identify multiple change points (Baranowski et al., 2019; Verbesselt et al., 2010), one should be cautious when a change point is found near the beginning or end of the data record, and when multiple change points are detected within a short period, since the time frame might be too short to conclude that the long-term trend has changed.

Discussion on the change point of trends and an abrupt change: For a demonstrative purpose, let us consider a phenomenon known as Simpson’s paradox: When we combine data series from two consecutive periods that both show a positive trend, is it possible to produce an overall negative trend? This result is shown in Figure A2, based on surface ozone anomalies measured at Centennial, Wyoming (1990-2021, US EPA Clean Air Status and Trends Network or CASTNET). We can clearly observe an abrupt change in spring of 2004. If a linear trend is fitted to the entire period, a negative slope is found ($-0.46 [\pm 0.69]$ ppbv/decade, $p = 0.18$, low certainty). Nevertheless, as indicated by a mean shift model (i.e. only two constant terms are fitted and no trend is considered), a clear shift of data averages is observed between 1990-2004 and 2004-2021 ($-1.97 [\pm 1.15]$ ppbv, $p \leq 0.01$). So this change point can be identified as a mean shift, which suggests that the overall negative trend is induced by the abrupt change.

We further fit piecewise segments with a turning point (without an offset) in April 2004. This model finds that there is indeed a change point of trends, with a positive trend of $1.71 [\pm 1.64]$ ppbv/decade ($p = 0.04$, high certainty) over 1990-2004 and a negative trend of $-2.05 [\pm 1.27]$ ppbv/decade ($p \leq 0.01$, very high certainty) over 2004-2021 ($dSNR = 5.3$ and the 95% confidence intervals of pre-trend and post-trend do not overlap). However, this model is still inadequate because the positive trend before the change point is obviously underestimated and the negative trend after the change point is clearly induced by the abrupt change. Under the circumstance when a change point is considered to be both a mean shift and a trend change, a piecewise trend with an offset at the change point should be adopted. This result can be summarized as a positive trend of $4.98 [\pm 1.60]$ ppbv/decade ($p \leq 0.01$, very high certainty) over 1990-2004 and $0.01 [\pm 1.15]$ ppbv/decade

($p = 0.99$, no evidence) over 2004-2021 (dSNR = 6.2 and no overlap of the 95% confidence intervals), with a shift of $-5.58 [\pm 0.88]$ ppbv ($p \leq 0.01$) at April 2004. This example demonstrates how different model assumptions can lead to diverse conclusions, thus model diagnostics and inspection are crucial for the change point analysis.

Annex C Prior data quality control

Data quality control (QC) is essential for ensuring that any analysis results are credible. Since ozone and its precursors are highly variable, automated QC methods based on purely numerical criteria might not be fully reliable. Practitioners are recommended to, 1) inspect the time series and data histograms across multiple years when checking the data consistency (e.g. to validate the trends or identify the abrupt change), and 2) compare the results from a collection of stations (e.g. geographically nearby or similar conditions), in order to avoid over-interpreting and over-generalizing the result from one specific site.

Annex D Overview of studying trends from multiple time series

Whereas this guidance note places the focus mainly on trend detection of a single time series, it is often desired or necessary to summarize the overall trends from a collection of multiple time series (e.g. over an extensive region). The complexity of merging multiple time series is dependent on if data are observed on a regular grid (this case is straightforward and will not be discussed further) or an irregular grid. In general, a direct aggregation of all time series observed at irregularly distributed locations should be avoided (Chang et al., 2021; Diggle et al., 2010). According to the data availability (i.e. the number and distribution of sites), and assuming that the region is relatively homogeneous (e.g. similar terrain and climate), the analysis of multiple time series observed at unevenly distributed locations can be classified according to the following scenarios.

- *Sites are dense enough and have a good coverage across the region:* In this case it is adequate to simply aggregate the data onto a regular designated grid, then select (predefined) gridded locations to derive the trends.
- *Sites are abundant but heavily clustered within certain sub-regions:* Compared to the previous scenario, the result will be highly irregular even if the data are re-gridded. Under this circumstance the spatial heterogeneity needs to be explicitly accounted for, by using geostatistical modeling approaches (Cressie, 1990; Stein, 1999) or machine learning techniques (e.g. random forests and neural networks, see Berrocal et al. (2020) and Betancourt et al. (2022)), which prevents the regional trend from being dominated by certain subregions.
- *Sites are too few to determine the spatial pattern:* This scenario is often the case for free tropospheric measurements (profile data are sparse in space and infrequent in time). Due to the information being limited, deriving the overall trends in this scenario is challenging and subject to a larger uncertainty. Since observational sites are too few to establish a reasonable spatial model, alternative approaches are 1) to fill the spatial gaps by using validated model simulations (DeLang et al., 2021; Malashock et al., 2022); and 2) to consider the common variability between different data sources to be representative of the trend. Signal decomposition or mixed-effects modeling techniques can be used to identify such common patterns (Chang et al., 2022; Pedersen et al., 2019). The greater the number of data sources, the easier the overall signal can be identified.

It should be emphasized that the above discussion is focused on how to properly combine multiple time series, but not meant to explicitly address the issue of data compatibility. For example, time series observed at nearby locations could have very different trends and uncertainties (due to localized variability and spatial heterogeneity). To address this issue, we should no longer consider this merely a time series problem, but within the scope of a longitudinal study (referring to a research design that involves repeated observations of a group of individuals over a period of time). This type of study aims to study group behavior, but individual differences are not unexpected (individuals respond varyingly, but the overall tendency can be

quantified). The class of mixed-effects models in statistics is designed for the purpose of longitudinal data analysis, by including a component to quantify the group behavior (known as the fixed effect), and another component to address the individual behavior (known as the random effect). Figure A3 demonstrates the analysis result from the linear mixed effects modeling approach, by showing how to use random effects to adjust the variability from individual data sources, in addition to the overall trend. More in-depth discussions are provided by Diggle et al. (2002) or introductory articles (Fitzmaurice and Ravichandran, 2008; Harrison et al., 2018).

In terms of the TOAR analyses, if practitioners have doubts when combining potentially heterogeneous time series, instead of aggregating data (we cannot be sure what information will be lost after data aggregation), we recommend applying a mixed-effects modeling approach to account for all available data sources, and to allow for flexibility when considering individual data sources or local variability (see examples from Chang et al., 2017, 2022).




Figure A3: A demonstration of the linear mixed-effects models: Six annual time series (averaged maximum daily 8-h averages limited to the warm season) are shown in different colors. Black line represents the overall intercept and slope (the fixed effect). Two different random effects are specified: (a) each site shares the same slope, but a site-specific intercept is adjusted to each site, and (b) a site-specific intercept/slope is adjusted to each site, so each site has its flexibility to reveal different trends. Note that the overall intercept and slope are considered to be representative, because the random intercepts and slopes from all sites are expected to have a zero mean and follow a normal distribution. Data are selected for the purpose of demonstration only (see Section 6 of Chang et al. (2021) for utilizing all available sites).

Annex E Workflow process for trend analysis using quantile regression and moving block bootstrap

This section provides a step-by-step demonstration of trend analysis using quantile regression (with Python and R code). For a data presentation, Fig A4 shows the monthly nighttime ozone time series measured at Mauna Loa, Hawaii (data and code can be downloaded at https://github.com/Kai-LanChang/statistical_guidelines).

For the first step, we demonstrate how to deseasonalize the data. The most widely applied approach is to account for all components in a single model fit (i.e. seasonality is derived based on all available data). However, it is often desirable to use a common baseline period when we compare trends at multiple sites (or one might have greater confidence in data quality and coverage over a particular period). Therefore, in this example we demonstrate the deseasonalization based on the seasonality derived from the period 2000-2020. The seasonality is modeled as a sine-cosine combination with periodicities of 12 and 6 months.

```
>>> print(mlo)
   year  month    x     y
0  1974     1  205  32.243400
1  1974     2  206  41.480401
2  1974     3  207  47.719027
3  1974     4  208  49.147040
4  1974     5  209  46.732616
...
559 2021     8  776  37.181008
560 2021     9  777  38.806028
561 2021    10  778  38.103139
562 2021    11  779  38.768326
563 2021    12  780  35.386503
```

Figure A4: Monthly nighttime ozone time series (y) measured at Mauna Loa, Hawaii.

Python code	R code
<pre>import numpy as np import pandas as pd import scipy.stats import statsmodels.formula.api as smf #read data mlo = pd.read_csv("mlo.csv") mlo.head() #derive the seasonal cycle base = mlo[(mlo["year"] >= 2000) & (mlo["year"] <= 2020)] seasonality = smf.ols("y~np.sin(2*np.pi*month/12)+np.cos(2* np.pi*month/12)+np.sin(2*np.pi*month/6)+np.c os(2*np.pi*month/6)", base).fit(method="qr").predict(pd.DataFrame({" month": range(1, 13)})) #calculate anomalies mlo = pd.merge(mlo, pd.DataFrame({"month": range(1, 13), "yd": seasonality}), on="month") mlo["yd"] = mlo["y"] - mlo["yd"] mlo = mlo.sort_values("x")</pre>	<pre>library(quantreg) #read data mlo=read.csv("mlo.csv") head(mlo) #derive the seasonal cycle base=mlo[mlo\$year>=2000 & mlo\$year<=2020,] seasonality=predict(lm(y~sin(2*pi*month/12)+co s(2*pi*month/12)+sin(2*pi*month/6)+cos(2*pi* month/6), data=base), newdata=data.frame(month=1:12)) #calculate anomalies mlo=merge(mlo, data.frame(month=1:12, yd=seasonality), by="month") mlo\$yd=mlo\$y - mlo\$yd mlo=mlo[order(mlo\$x),]</pre>

- Four harmonic terms are used to capture the seasonality, these terms are subject to adjustment depending on data complexity.
- Data needs to be sorted by the temporal index (x), so we can sample the data in a correct order. In this data set the monthly temporal index is referenced to January 1957 (it can be defined by users).

```
>>> print(mlo)
   month  year    x     y     yd
0      1  1974  205  32.243400 -8.007737
1      2  1974  206  41.480401 -1.034788
2      3  1974  207  47.719027  1.696708
3      4  1974  208  49.147040  0.779685
4      5  1974  209  46.732616 -0.438169
...
559    8  2021  776  37.181008  3.242525
560    9  2021  777  38.806028  4.371030
561   10  2021  778  38.103139  1.228904
562   11  2021  779  38.768326 -0.083110
563   12  2021  780  35.386503 -4.164578
```

Figure A5: Monthly nighttime ozone time series (y) and anomalies (yd) at Mauna Loa, Hawaii.

The second step is to perform the quantile regression and apply a moving block bootstrap algorithm, so that the uncertainty can be properly estimated. The resulting trend value and its 1-sigma uncertainty can be found in ‘fit’ and ‘fit_se’ (in units of ppbv/year). Note that typically the result is reported by “trend [± 2 -sigma uncertainty]”.

Python code	R code
<pre> def mbfun(formula, data, tau): n = len(data) b = np.ceil(n**0.25).astype(int) nblocks = np.ceil(n/b).astype(int) blocks = np.array([list(range(i, i+b)) for i in range(n-b+1)]) bn = np.random.choice(np.arange(len(blocks)), nblocks, replace=True) samp_data = data.iloc[blocks[bn].flatten()] mod = smf.quantreg(formula, data=samp_data).fit(q=tau) return mod.params #median intercept and slope fit = smf.quantreg("yd~x", data=mlo).fit(q=0.5).params*12 #MBB standard error for intercept and slope np.random.seed(2013) op = [mbfun(formula="yd~x", data=mlo, tau=0.5) for _ in range(1000)] fit_se = np.nanstd(op, axis=0)*12 #p value for intercept and slope fit_pv = 2*scipy.stats.t.sf(x=abs(fit/fit_se), df=len(mlo)-2) </pre>	<pre> mbfun=function(formula,data,tau){ n=nrow(data) b=ceiling(n^0.25) nblocks=ceiling(n/b) blocks=lapply(seq_len(n-b+1), function(i) seq(i, i+b-1)) bn=sample(1:length(blocks),nblocks,replace=T) samp_data=data[unlist(blocks[bn]),] mod=rq(formula, data=samp_data, tau=tau) coef(mod) } #median intercept and slope fit=coef(rq(yd~x, data=mlo, tau=0.5))*12 #MBB standard error for intercept and slope set.seed(2013) op=t(replicate(1000, mbfun(formula=yd~x,data=mlo,tau=0.5))) fit_se=apply(op, 2, sd, na.rm=TRUE)*12 #p value for intercept and slope fit_pv=2*pt(q=abs(fit/fit_se), df=nrow(mlo)-2, lower.tail=FALSE) </pre>

1. The optimal block length suggested by Hall et al. (1995) is $b=n^k$, where $k=1/3$, $1/4$ or $1/5$. An alternative approach to determine the optimal block length based on autocorrelation and sample size is proposed by Politis and White (2004).
2. Different trend techniques can be adopted by this algorithm by replacing `quantreg()`/`rq()` with other valid models.
3. q/τ is ranged between (0,1) representing the quantile to be estimated, here 0.5 (median) is used.
4. Multiple quantiles can be estimated in R at once by `tau=seq(0.1,0.9,by=0.1)`.
5. `random.seed()` in Python and `set.seed()` in R create random numbers which can be reproduced (otherwise the bootstrapped results will not be the same every time the bootstrap is implemented). Note that even if the same seed number is set, random numbers are generated differently between Python and R.
6. The sigma value is estimated based on 1000 sampled trends.
7. The final trend and sigma values are converted from the units of ppbv/month to ppbv/year.
8. p -value for Student's t distribution is obtained by providing signal-to-noise ratio and degrees of freedom.

It is worth emphasizing that this code was written for a demonstrative purpose, and that data aggregation is not a necessary step for QR. Figure A6 shows the daily and monthly ozone time series at MLO, along with the monthly 95th and 5th percentiles derived from daily values. This figure clearly demonstrates how the (potentially valuable) information is lost after monthly data aggregations. Since the extreme ozone events are of great interest to the research community, we recommend preserving those extreme variabilities, and then utilize QR for detecting heterogeneous ozone changes.




Figure A6: Daily and monthly surface ozone variability at MLO, along with the monthly 95th and 5th percentiles derived from daily values.

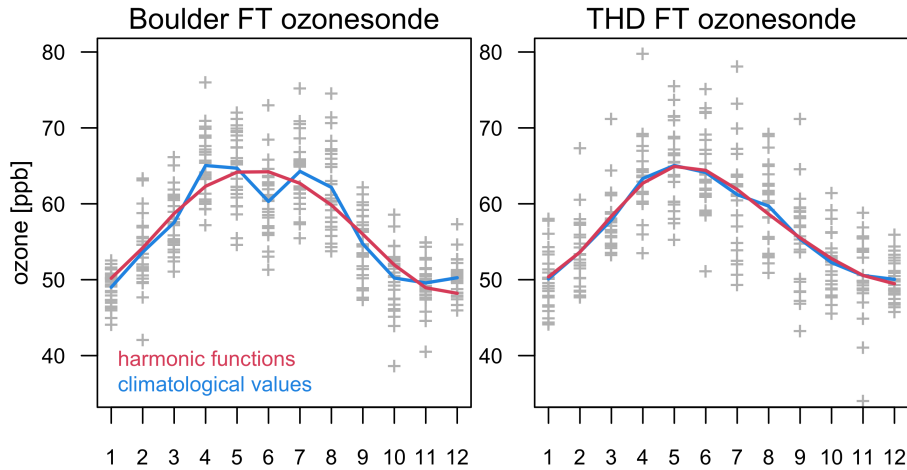


Figure A7: A demonstration of different approaches to derive seasonal cycles, using harmonic functions (red) and long-term mean of each individual month (blue). Data are taken from the Boulder (Colorado, 1995-2021) and Trinidad Head (California, 1997-2021) ozonesonde records in the free troposphere (700-300 hPa).

Annex F Practical issues

Different approaches to derive seasonal cycles and anomaly series

In general, different approaches to derive the regular seasonal cycle have a minor impact on the estimation of trends (Chang et al., 2023; Weatherhead et al., 1998). In Section E we use the harmonic functions to approximate the seasonality, which is a typical approach for linear regression models. Another common approach for deriving seasonality is calculated by the long-term mean of each individual month (i.e. the climatology). Then the anomaly is yielded by taking the difference between the long-term mean and observed monthly value. Those two approaches might or might not yield similar seasonal cycles. For example, Figure A7 shows the seasonal cycles derived by harmonic functions and climatological values from the free tropospheric ozonesonde data measured at Boulder, Colorado and Trinidad Head (THD), California. We can see the results are nearly identical for the THD data, but some discrepancies are observed above Boulder (two modes v. a single mode). Even though the trend estimate should not be adversely affected, should the purpose be to investigate the seasonality, practitioners need to evaluate if those discrepancies are real or an artifact due to an irregular sampling scheme. Two modes of variation above Boulder are likely to be real because the result is based on a long-term record (1995-2021) with a stable sampling frequency. Multimodality is difficult to capture by lower order harmonic functions (model fits in linear regression are typically assumed to be smooth with less wiggles), but it can be adopted by using the penalized spline functions from more sophisticated generalized additive models (Chang et al., 2017). Nevertheless, in terms of trend analysis both approaches should be practical.

An incorporation of covariates into multiple linear quantile regressions

Similar to conventional multiple linear regression analysis, an incorporation of covariates into QR models is possible. In this section we use several meteorological variables (dewpoint, relative humidity, wind direction, and wind speed) to fit daily surface ozone variability measured at MLO (1977-2021). The fit can be specified through the following R code:

```
formula=ozone~ind+dewpoint+relative_humidity+wind_direction+wind_speed+
  sin(2*pi*month/12)+cos(2*pi*month/12)+sin(2*pi*month/6)+cos(2*pi*month/6)
mod=rq(formula, data, tau=c(0.05,0.5,0.95))
```

The fitted results are shown in Table A1. We find that meteorological variables play different roles at different percentiles: dewpoint and wind speed show a strong correlation at all three percentiles (but wind speed has stronger influence on the low percentile than high percentile), a strong correlation with relative humidity only presents for the median ozone values, and a strong correlation with wind direction only presents at the

Table A1: Regression coefficients for meteorological variables (purple) and trends (orange, in units of ppbv/decade) at different ozone percentiles at MLO (1977-2021). Each meteorological variable is standardized so the coefficients are comparable.

	5th	50 th	95th
dewpoint	-6.89, p<0.01	-7.70, p<0.01	-6.47, p<0.01
relative humidity	-0.13, p=0.68	0.69, p<0.01	-0.12, p=0.76
wind direction	-0.06, p=0.64	-0.01, p=0.87	-0.46, p<0.01
wind speed	-1.24, p<0.01	-0.49, p<0.01	0.30, p=0.02
Trend (no met. variables)	0.14 [±0.69], p=0.68	1.06 [±0.58], p<0.01	1.33 [±0.65], p<0.01
Trend (with met. variables)	1.10 [±0.34], p<0.01	1.33 [±0.33], p<0.01	1.25 [±0.50], p<0.01

95th percentile. Therefore, we should expect the meteorological influence to vary at different percentiles. In terms of trend estimates, the magnitudes at the 50th and 95th percentiles are similar, but the magnitudes at the 5th percentile become stronger after the meteorological variables are accounted for (mainly due to dewpoint, not shown). Confidence intervals become narrower at all three percentiles, indicating that trend uncertainties can be attributed by the meteorological influence.

To visualize the roles of different meteorological variables, Figure A8 shows an incorporation of dewpoint and wind speed into the fit of the 95th percentile, respectively, in addition to the basic model fit with only a trend and seasonal components. Even though both dewpoint and wind speed show a strong correlation at the 95th percentile (Table A1), we can clearly see that dewpoint is more important to explain ozone variability.




Figure A8: Daily ozone time series (red) and QR fits with different model components (blue) at the 95th percentiles at MLO. Purple in the lower panels indicates observed and fitted values are overlapped. For a demonstrative purpose, only the results after 2011 are shown.

Interpolation of irregularly distributed observations onto a regular grid

Although there are many methods for fitting a spatial interpolation model, the rationale is the same: fitting a spatial model to irregularly distributed observations for the purpose of predicting interpolated values on a regular grid. For example, Figure A9 shows the locations of monitoring stations over the southwestern USA,

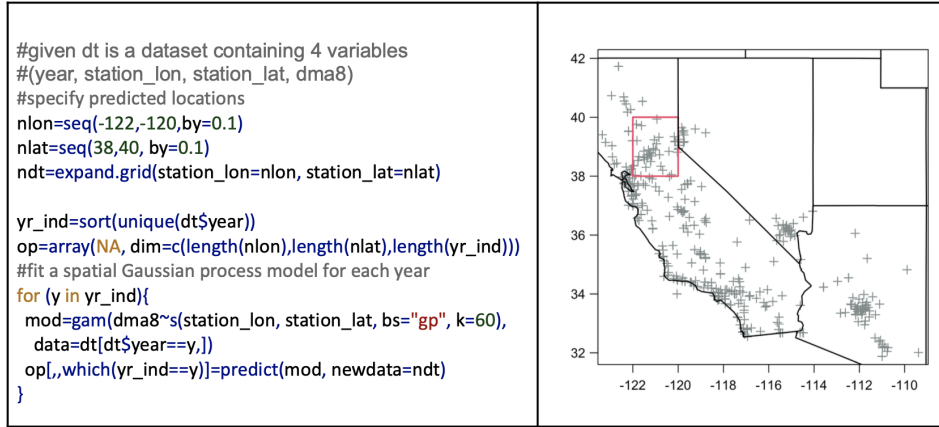


Figure A9: Locations of monitoring stations (crosses) and gridded area to be predicted (red rectangle).

and we aim to conduct the trend analysis for a grid box 38-40°N and 122-124°W. The example code first defines the targeted grid box, and sets a loop to perform spatial interpolation for each year. The method adopts a spatial Gaussian process (gp) model fitted through a generalized additive model (gam, Wood, 2001). Trend analysis can then be carried out on the resulting output.

Impact of lagged ENSO (El Niño-Southern Oscillation) correlations on ozone trends

The correlation between tropospheric ozone and ENSO is well recognized, especially across the Pacific Ocean. Its degree of correlation is varying not only with different locations, but also with different time lags. This section aims to investigate the difference between the peak and zero-lag ENSO correlations and quantify its impact on trends, by using the monthly gridded tropospheric column ozone observations measured from the OMI/MLS satellite instrument (Ziemke et al. (2006, 2019), data can be found at https://acd-ext.gsfc.nasa.gov/Data_services/cloud_slice/; the ENSO index used here is taken from <https://psl.noaa.gov/enso/mei/>). To facilitate discussion, we use ‘optimized lag’ to refer to the lag where the correlation peaks. The lag correlations can readily be produced by the cross-correlation function (ccf) in R and Python. For a demonstration, Figure A10 shows the correlation between ENSO and ozone at different monthly lags at a grid cell 57.5°E and 2.5°S. For this particular grid cell the optimized lag is located at -4, indicating that a strong ENSO signal might be highly correlated with ozone measured 4 months later.

We proceed by extracting zero-lag correlation and the peak correlation between ± 6 monthly lags at each cell (a longer correlation range is not considered), and the resulting maps are shown in Figure A11. For the regions with the strongest ENSO correlations (i.e. tropical Pacific Ocean and Southeastern Asia), we find that the correlation patterns are highly similar between two scenarios. Nevertheless, the grid cells with a very weak zero-lag correlation in the Northern Hemisphere and Tropics generally show a stronger positive correlation after the optimized lag is found, with some exceptions in the eastern North Atlantic Ocean where negative correlation is present. The median trend estimates are shown in Figure A12, three scenarios are considered here: 1) no ENSO is taken into account, 2) zero-lag ENSO correlation is incorporated, and 3) the optimized lag ENSO correlation is incorporated. Even if some minor discrepancies can be identified, highly consistent patterns can be observed between these three scenarios.

To explicitly quantify the ENSO impact, we define the relative change of a value A versus a value B as $(A - B)/B$. We calculate the relative change in trends and uncertainties for zero-lag v no ENSO correlations (case 1), and the peak v zero-lag ENSO correlations (case 2), respectively, and the result is shown in Figure A13. The only systematic bias is found from the relative change in trend uncertainties in case 1. After the ENSO correlation is accounted for, the uncertainties become smaller over the tropical Pacific Ocean (where the strongest ENSO correlation is located), indicating that certain trend uncertainty can be attributed by the ENSO signal. For the relative change in trends, 95.8% cells in case 1 and 97.0% cells in case 2 fall within $\pm 0.5\%$ bias, and 98.8% cells in case 1 and 98.4% cells in case 2 fall within $\pm 1\%$ bias. The relative change in uncertainties are all limited within 0.7% bias for both cases. In summary, in this case study we may conclude that an incorporation of ENSO in trend analysis could reduce trend uncertainty and avoid potential systematic bias at certain regions, but its impact on trend estimate is expected to be weak. Even though

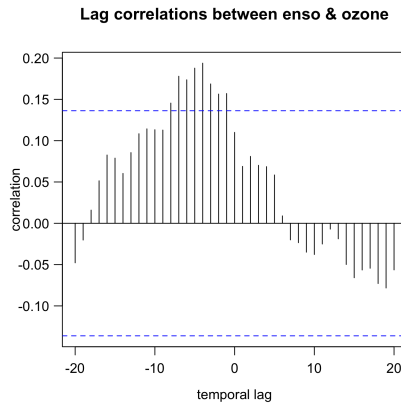


Figure A10: Correlation between ENSO and ozone at different monthly lags (OMI/MLS data, 57.5°E and 2.5°S). The peak (positive) correlation is found at -4 lag.

the peak correlation between ENSO and ozone might be temporally lagged, it does not necessarily have an impact on trends and uncertainties. The complete R code to produce all figures in this section is provided at https://github.com/Kai-LanChang/statistical_guidelines.

References

- Banerjee, A., Fyfe, J. C., Polvani, L. M., et al. (2020). A pause in southern hemisphere circulation trends due to the Montreal Protocol. *Nature*, 579(7800):544–548.
- Baranowski, R., Chen, Y., and Fryzlewicz, P. (2019). Narrowest-over-threshold detection of multiple change points and change-point-like features. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 81(3):649–672.
- Beaulieu, C., Chen, J., and Sarmiento, J. L. (2012). Change-point analysis as a tool to detect abrupt climate variations. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1962):1228–1249.
- Berrocal, V. J., Guan, Y., Muyskens, A., et al. (2020). A comparison of statistical and machine learning methods for creating national daily maps of ambient PM_{2.5} concentration. *Atmospheric Environment*, 222:117130.
- Betancourt, C., Stomberg, T. T., Edrich, A.-K., et al. (2022). Global, high-resolution mapping of tropospheric ozone—explainable machine learning and impact of uncertainties. *Geoscientific Model Development*, 15(11):4331–4354.
- Chandler, R. and Scott, M. (2011). *Statistical methods for trend detection and analysis in the environmental sciences*. John Wiley & Sons.
- Chang, K.-L. (2022). Lecture on quantile regression. Available at <https://b2share.fz-juelich.de/records/f52a5d792fbb42c3b1f1b92333d4d86e>.
- Chang, K.-L., Cooper, O. R., Gaudel, A., et al. (2022). Impact of the COVID-19 economic downturn on tropospheric ozone trends: an uncertainty weighted data synthesis for quantifying regional anomalies above western North America and Europe. *AGU Advances*, e2021AV000542.
- Chang, K.-L., Cooper, O. R., Rodriguez, G., et al. (2023). Diverging ozone trends above western North America: Boundary layer decreases versus free tropospheric increases. *Journal of Geophysical Research: Atmospheres*, e2022JD038090.
- Chang, K.-L., Petropavlovskikh, I., Cooper, O. R., et al. (2017). Regional trend analysis of surface ozone observations from monitoring networks in eastern North America, Europe and East Asia. *Elementa: Science of the Anthropocene*, 5:50.
- Chang, K.-L., Schultz, M. G., Lan, X., et al. (2021). Trend detection of atmospheric time series: Incorporating appropriate uncertainty estimates and handling extreme events. *Elementa: Science of the Anthropocene*, 9(1):00035.
- Chen, C. W. S., Chan, J. S. K., Gerlach, R., and Hsieh, W. Y. L. (2011). A comparison of estimators for regression models with change points. *Statistics and Computing*, 21:395–414.
- Cooper, O. R., Schultz, M. G., Schröder, S., et al. (2020). Multi-decadal surface ozone trends at globally distributed remote locations. *Elementa: Science of the Anthropocene*, 8(23).
- Cressie, N. (1990). The origins of kriging. *Mathematical Geology*, 22:239–252.
- Cressie, N. (2021). A few statistical principles for data science. *Australian & New Zealand Journal of Statistics*, 63(1):182–200.
- Crimmins, A. (2020). Improving the use of calibrated language in US climate assessments. *Earth’s Future*, 8(11):e2020EF001817.
- Das, K., Krzywinski, M., and Altman, N. (2019). Quantile regression. *Nature Methods*, 16(6):0–9.

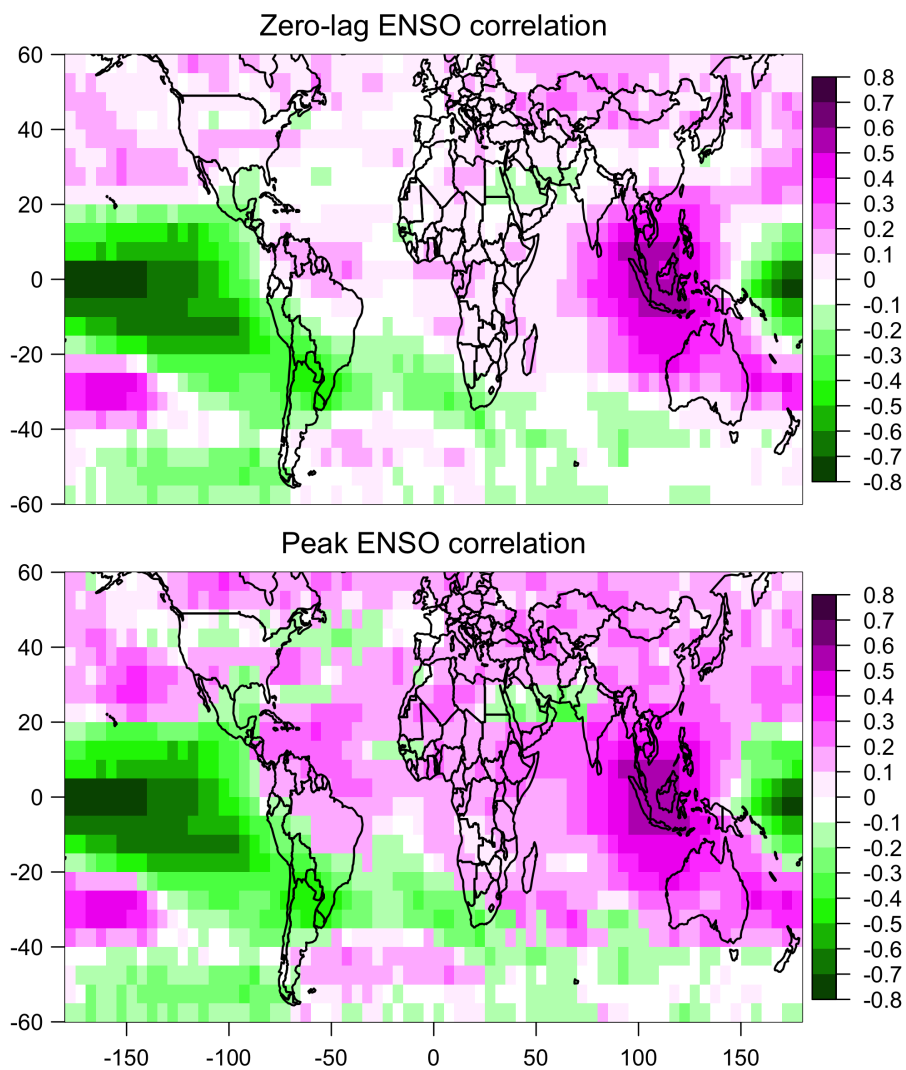


Figure A11: Correlation coefficients between ENSO and ozone from the OMI/MLS dataset (2004-2021): the upper panel shows zero-lag ENSO correlation, and the lower panel shows the peak ENSO correlation between ± 6 monthly lags.

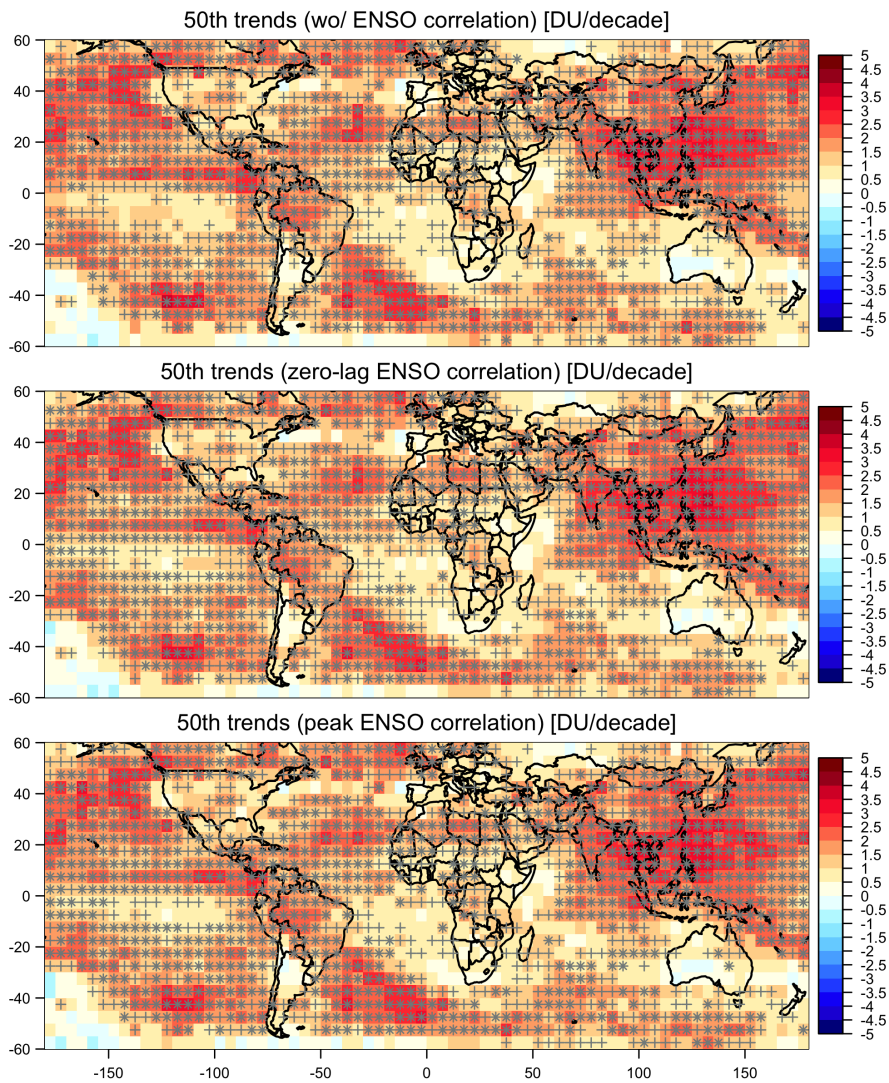


Figure A12: Median trend estimates [DU/decade] from the OMI/MLS dataset (2004-2021): the upper panel shows the estimation without incorporating ENSO correlation, the middle panel is based on the zero-lag ENSO correlation, and the lower panel is based on the peak ENSO correlation between ± 6 monthly lags. Gray crosses indicate grid cells with a p -value between 0.01 and 0.05, and stars indicate grid cells with a p -value less than or equal to 0.01.

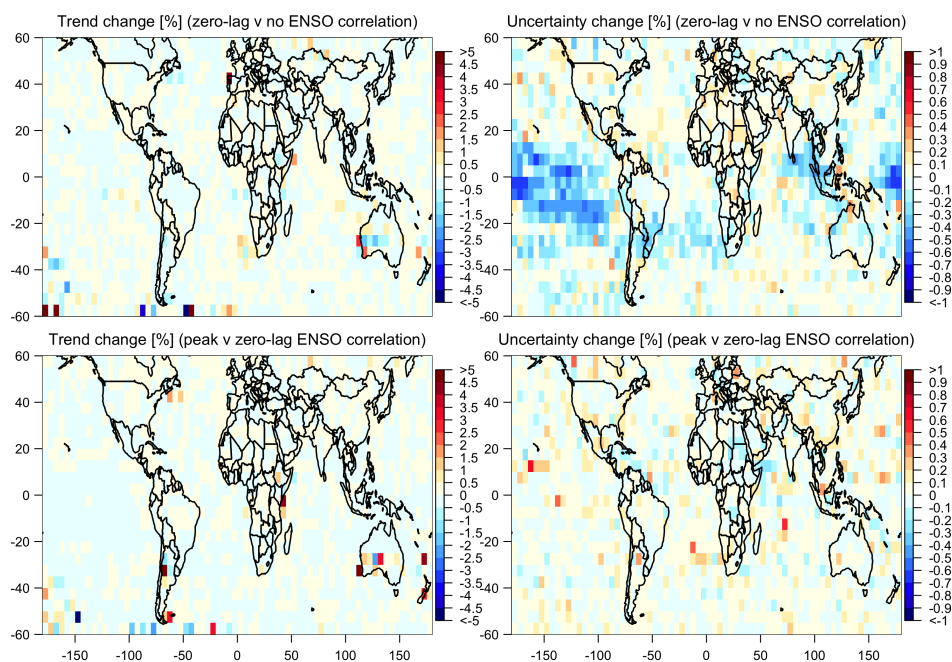


Figure A13: Relative change [%] in median trends (left) and uncertainties (right) for zero-lag v no ENSO correlations (top), and the peak v zero-lag ENSO correlations (bottom), respectively.

- DeLang, M. N., Becker, J. S., Chang, K.-L., et al. (2021). Mapping yearly fine resolution global surface ozone through the bayesian maximum entropy data fusion of observations and model output for 1990–2017. *Environmental Science & Technology*, 55(8):4389–4398.
- Diggle, P. J., Heagerty, P., Liang, K.-Y., et al. (2002). *Analysis of longitudinal data*. Oxford University Press.
- Diggle, P. J., Menezes, R., and Su, T.-I. (2010). Geostatistical inference under preferential sampling. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 59(2):191–232.
- Fitzenberger, B. (1998). The moving blocks bootstrap and robust inference for linear least squares and quantile regressions. *Journal of Econometrics*, 82(2):235–287.
- Fitzmaurice, G. M. and Ravichandran, C. (2008). A primer in longitudinal data analysis. *Circulation*, 118(19):2005–2010.
- Fleming, Z. L., Doherty, R. M., von Schneidmesser, E., et al. (2018). Tropospheric Ozone Assessment Report: Present-day ozone distribution and trends relevant to human health. *Elementa: Science of the Anthropocene*, 6:12.
- Gaudel, A., Cooper, O. R., Ancellet, G., et al. (2018). Tropospheric Ozone Assessment Report: Present-day distribution and trends of tropospheric ozone relevant to climate and global atmospheric chemistry model evaluation. *Elementa: Science of the Anthropocene*, 6.
- Gilleland, E. (2020a). Bootstrap methods for statistical inference. Part I: Comparative forecast verification for continuous variables. *Journal of Atmospheric and Oceanic Technology*, 37(11):2117–2134.
- Gilleland, E. (2020b). Bootstrap methods for statistical inference. Part II: Extreme-value analysis. *Journal of Atmospheric and Oceanic Technology*, 37(11):2135–2144.
- Hall, P., Horowitz, J. L., and Jing, B.-Y. (1995). On blocking rules for the bootstrap with dependent data. *Biometrika*, 82(3):561–574.
- Hall, P. and Sheather, S. J. (1988). On the distribution of a studentized quantile. *Journal of the Royal Statistical Society: Series B (Methodological)*, 50(3):381–391.
- Hannan, E. J., Terrell, R. D., and Tuckwell, N. E. (1970). The seasonal adjustment of economic time series. *International Economic Review*, 11(1):24–52.
- Harrison, X. A., Donaldson, L., Correa-Cano, M. E., et al. (2018). A brief introduction to mixed effects modelling and multi-model inference in ecology. *PeerJ*, 6:e4794.
- Kiefer, N. M. and Vogelsang, T. J. (2002). Heteroskedasticity–autocorrelation robust standard errors using the Bartlett kernel without truncation. *Econometrica*, 70(5):2093–2095.
- Killick, R. and Eckley, I. (2014). changepoint: An R package for changepoint analysis. *Journal of Statistical Software*, 58(3):1–19.
- Koenker, R. and Hallock, K. F. (2001). Quantile regression. *Journal of Economic Perspectives*, 15(4):143–156.
- Koenker, R. and Machado, J. A. (1999). Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association*, 94(448):1296–1310.
- Krock, M., Bessac, J., Stein, M. L., and Monahan, A. H. (2022). Nonstationary seasonal model for daily mean temperature distribution bridging bulk and tails. *Weather and Climate Extremes*, 36:100438.

- Lahiri, S. N., Furukawa, K., and Lee, Y.-D. (2007). A nonparametric plug-in rule for selecting optimal block lengths for block bootstrap methods. *Statistical Methodology*, 4(3):292–321.
- Malashock, D. A., Delang, M. N., Becker, J. S., et al. (2022). Global trends in ozone concentration and attributable mortality for urban, peri-urban, and rural areas between 2000 and 2019: a modelling study. *The Lancet Planetary Health*, 6(12):e958–e967.
- Mastrandrea, M. D., Field, C. B., Stocker, T. F., et al. (2010). *Guidance note for lead authors of the IPCC fifth assessment report on consistent treatment of uncertainties*. Intergovernmental Panel on Climate Change. Available at https://www.ipcc.ch/site/assets/uploads/2017/08/AR5_Uncertainty_Guidance_Note.pdf.
- Mills, G., Pleijel, H., Malley, C. S., et al. (2018). Tropospheric Ozone Assessment Report: Present-day tropospheric ozone distribution and trends relevant to vegetation. *Elementa: Science of the Anthropocene*, 6:47.
- Muggeo, V. M. (2008). Segmented: an R package to fit regression models with broken-line relationships. *R news*, 8(1):20–25.
- Oltmans, S. J., Lefohn, A. S., Harris, J., et al. (2006). Long-term changes in tropospheric ozone. *Atmospheric Environment*, 40(17):3156–3173.
- Paolo, F. S., Fricker, H. A., and Padman, L. (2015). Volume loss from Antarctic ice shelves is accelerating. *Science*, 348(6232):327–331.
- Pedersen, E. J., Miller, D. L., Simpson, G. L., and Ross, N. (2019). Hierarchical generalized additive models in ecology: an introduction with mgcv. *PeerJ*, 7:e6876.
- Politis, D. N. and White, H. (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, 23(1):53–70.
- Razavi, S. and Vogel, R. (2018). Prewhitening of hydroclimatic time series? Implications for inferred change and variability across time scales. *Journal of Hydrology*, 557:109–115.
- Reeves, J., Chen, J., Wang, X. L., et al. (2007). A review and comparison of changepoint detection techniques for climate data. *Journal of Applied Meteorology and Climatology*, 46(6):900–915.
- Stein, M. L. (1999). *Interpolation of spatial data: some theory for kriging*. Springer.
- Tong, C. (2019). Statistical inference enables bad science; statistical thinking enables good science. *The American Statistician*, 73(sup1):246–261.
- Verbesselt, J., Hyndman, R., Zeileis, A., and Culvenor, D. (2010). Phenological change detection while accounting for abrupt and gradual trends in satellite image time series. *Remote Sensing of Environment*, 114(12):2970–2980.
- Von Storch, H. and Zwiers, F. W. (2001). *Statistical analysis in climate research*. Cambridge University Press.
- Wasserstein, R. L. and Lazar, N. A. (2016). The ASA’s statement on p -values: context, process, and purpose. *The American Statistician*, 70(2):129–133.
- Wasserstein, R. L., Schirm, A. L., and Lazar, N. A. (2019). Moving to a world beyond “ $p < 0.05$ ”. *The American Statistician*, 73(1):1–19.
- Weatherhead, E. C., Harder, J., Araujo-Pradere, E. A., et al. (2017). How long do satellites need to overlap? Evaluation of climate data stability from overlapping satellite records. *Atmospheric Chemistry and Physics*, 17(24):15069–15093.
- Weatherhead, E. C., Reinsel, G. C., Tiao, G. C., et al. (1998). Factors affecting the detection of trends: Statistical considerations and applications to environmental data. *Journal of Geophysical Research: Atmospheres*, 103(D14):17149–17161.
- Wells, B., Dolwick, P., Eder, B., et al. (2021). Improved estimation of trends in US ozone concentrations adjusted for interannual variability in meteorological conditions. *Atmospheric Environment*, 248:118234.
- Wilcox, R. R. (2001). *Fundamentals of modern statistical methods: Substantially improving power and accuracy*. Springer.
- Wilcox, R. R. (2011). *Introduction to robust estimation and hypothesis testing*. Academic press.
- Wilks, D. S. (2011). *Statistical methods in the atmospheric sciences (3rd ed.)*. Academic press.
- Wood, S. N. (2001). mgcv: GAMs and generalized ridge regression for R. *R news*, 1(2):20–25.
- Yu, C. and Yao, W. (2017). Robust linear regression: A review and comparison. *Communications in Statistics-Simulation and Computation*, 46(8):6261–6282.
- Yue, S. and Wang, C. Y. (2002). Applicability of prewhitening to eliminate the influence of serial correlation on the Mann-Kendall test. *Water Resources Research*, 38(6):4–1.
- Ziemke, J. R., Chandra, S., Duncan, B. N., et al. (2006). Tropospheric ozone determined from Aura OMI and MLS: Evaluation of measurements and comparison with the Global Modeling Initiative’s Chemical Transport Model. *Journal of Geophysical Research: Atmospheres*, 111(D19).
- Ziemke, J. R., Oman, L. D., Strode, S. A., et al. (2019). Trends in global tropospheric ozone inferred from a composite record of TOMS/OMI/MLS/OMPS satellite measurements and the MERRA-2 GMI simulation. *Atmospheric Chemistry and Physics*, 19(5):3257–3269.