Modelling and Visualising Trends of Extreme Values in Acidifying Variables

Project Report

Leonardo Olivetti, Claudia von Brömssen

Swedish University of Agricultural Sciences, SLU
Department of Energy and Technology
Rapport 121(2021)

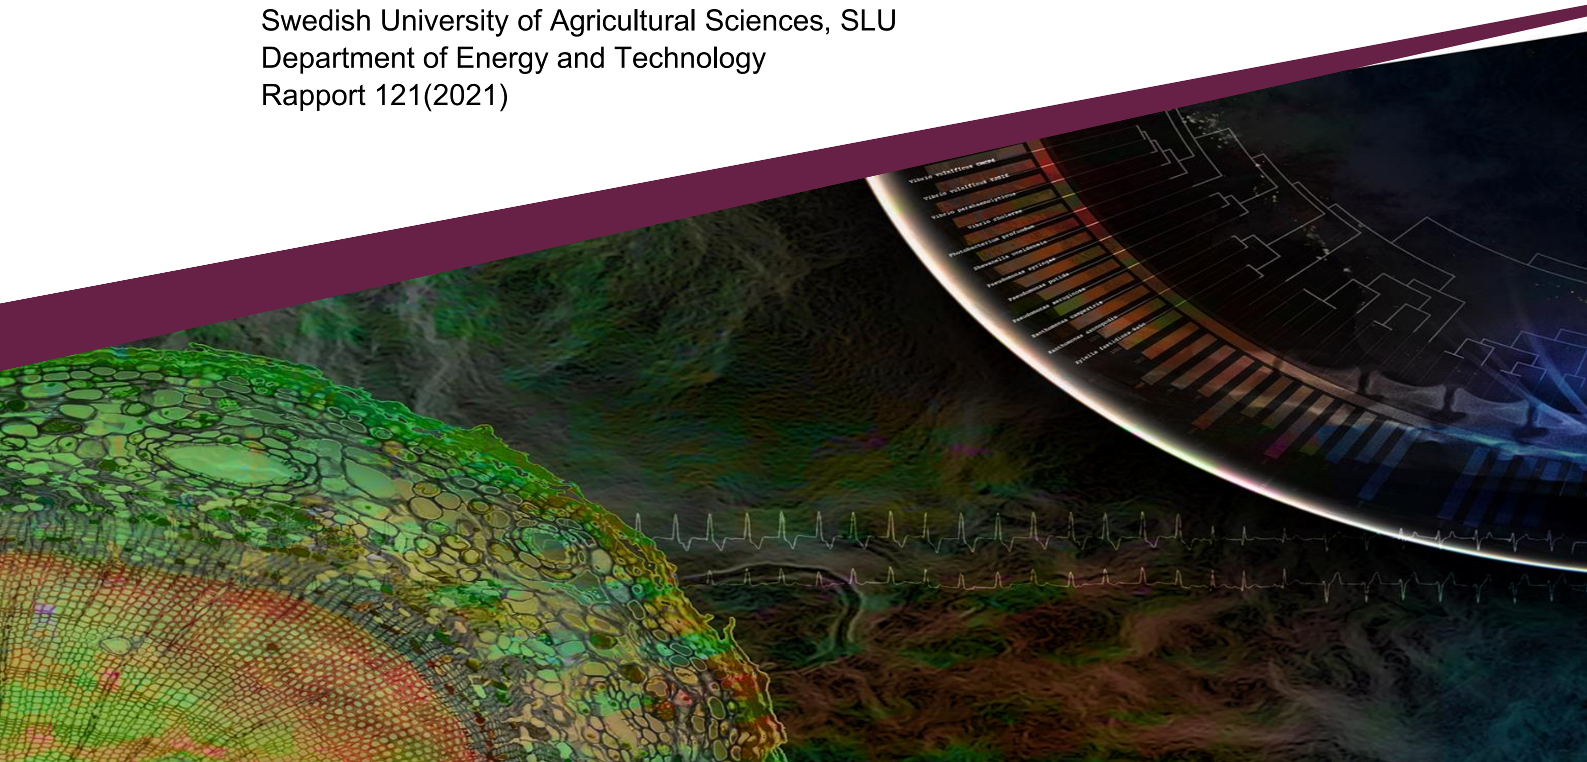

Preface

This report summarises the findings within the project “Modelling and Visualising Trends of Extreme Values in Acidifying Variables”, financed by the SLU Environmental Monitoring and Assessment program “Acidification”.

Acknowledgements

The authors are grateful to Jens Fölster, Marcus Wallin and Jesper Rydén for their valuable comments during the project work and on the report draft.

Table of Contents

1 Introduction

The southern regions of Sweden, Norway and Finland are among the areas in Europe most affected by surface water acidification (European Environment Agency, 2016). Since the 1980s, some steps have been taken to contrast this issue, which have contributed to a steady recovery of the water chemistry (Erlandsson et al., 2010). Among others, the Geneva Convention on Long Range Transboundary Air Pollution and its associated protocols have made an important contribution to the reduction of sulfate emissions (ibid.). In addition, monitoring programs have been put in place to detect patterns of acidification, and predict the long-term path to full recovery (European Environment Agency, 2016; Fölster et al., 2014).

A number of models have been developed in order to study the long-term trends in the levels of PH and other acidifying variables of relevance (Moldan et al., 2013; Wright and Cosby, 2003). Those models are useful to evaluate the overall trend in the levels of water acidity, and discuss whether the process of recovery is going in the right direction. However, a limitation of those models is that they might overlook the occurrence of episodic acidification, which can have severe impacts on local ecosystems (Baker et al., 1996; Heard et al., 1997; Laudon, 2008). Sudden changes in the levels of acidity can be deadly to a large number of fish species, and undermine long-term biodiversity (ibid.).

The purpose of this research is to contribute to fill this knowledge gap, by developing a framework to analyse, graphically and numerically, trends in the occurrence of episodic acidification. By combining the use of generalised additive models and quantile regression, models able to incorporate both seasonal and long-term time trends are developed. Patterns in episodic acidification are then illustrated with the help of visual tools first introduced by von Brömssen et al. (2021).

1.1 Aim

The overall objective of this project is to model and visualise trends in extreme values of acidifying variables. The reason for why this is interesting is that ecological effects of relevance are sometimes more closely related to the occurrence of episodic acidification than to the long-term average levels of acidity. In spite of this, many of the present models overlook episodic acidification, which is why more research on the patterns of episodic acidification is necessary.

1.2 Disposition

The structure of this project report is as follows:

Section 2 describes the dataset on which this research is based, and offers an empirical justification for why it is important to model the extreme values of the acidifying variables. Section 3 introduces the reader to the methodology employed in order to describe long-term trends in acidifying variables. Section 4 presents the main results of this research through examples of local and regional analyses. Section 5 summarises the results presented in Section 4, and discusses some possible strengths and limitations of the models employed.

1.3 Abbreviations

N: Northern Sweden

SV: Southern and Western Sweden

ÖM: Eastern and Central Sweden

GAMs: Generalised Additive Models

QGAMs: Quantile (Generalised) Additive Models

GAMMs: Generalised Additive Mixed Models

2 Data

The data available for this study were collected as a part of the Swedish national program for the monitoring of freshwater quality. The chemical analyses were performed by the Department of Aquatic Sciences and Assessment at the Swedish University of Agricultural Sciences (see also von Brömssen et al. 2021 and Fölster et al. 2014 for more details). The collection was performed on a regular basis (once a month in most cases) at over 200 different monitoring sites spread over the three macro-regions of Sweden (N, SV and ÖM). The data include information on several variables related to the acidity of freshwaters, such as pH, sulfates and ANC (Acid Neutralising Capacity). At least 10 years of data are available for all of the sites included in this paper. The methodology employed allows for time series of different length, and therefore makes use of all available data (see also Section 3 for additional details).

2.1 Long-term trends

This subsection briefly describes the long-term trends of acidifying variables in Swedish rivers. Mean, medians and key percentiles are presented for the three Swedish macro-regions (N, SV and ÖM).

Figure 1 shows how the mean levels of PH, ANC and sulfate concentration (mekv/l) have developed over the last 40 years. Important changes in the PH and concentration of sulfates have occurred before the year 2000. Since then, the mean value of all three acidifying variables has been rather stable.

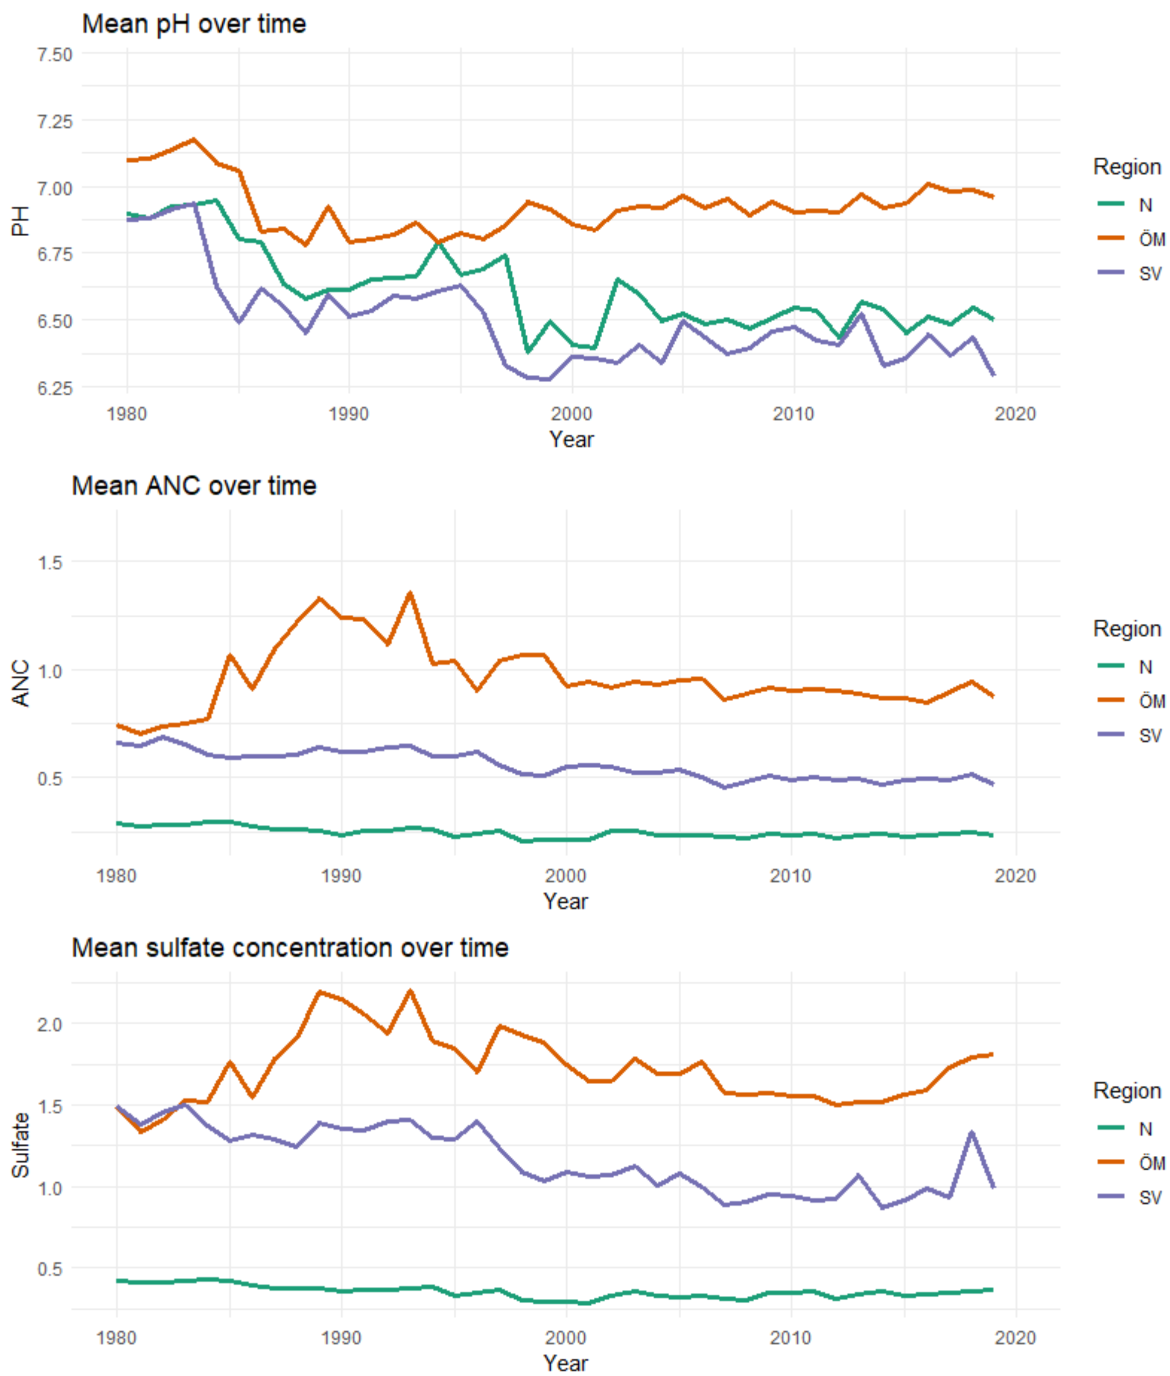

Figure 1: Mean PH, ANC and sulfate concentration over time

Figure 2 shows how the median levels of the three key acidifying variables have been evolving over time. A slightly different picture is painted by in this case: The overall trends are rather similar to those in Figure 1, but the differences among regions are smaller (Figure 2).

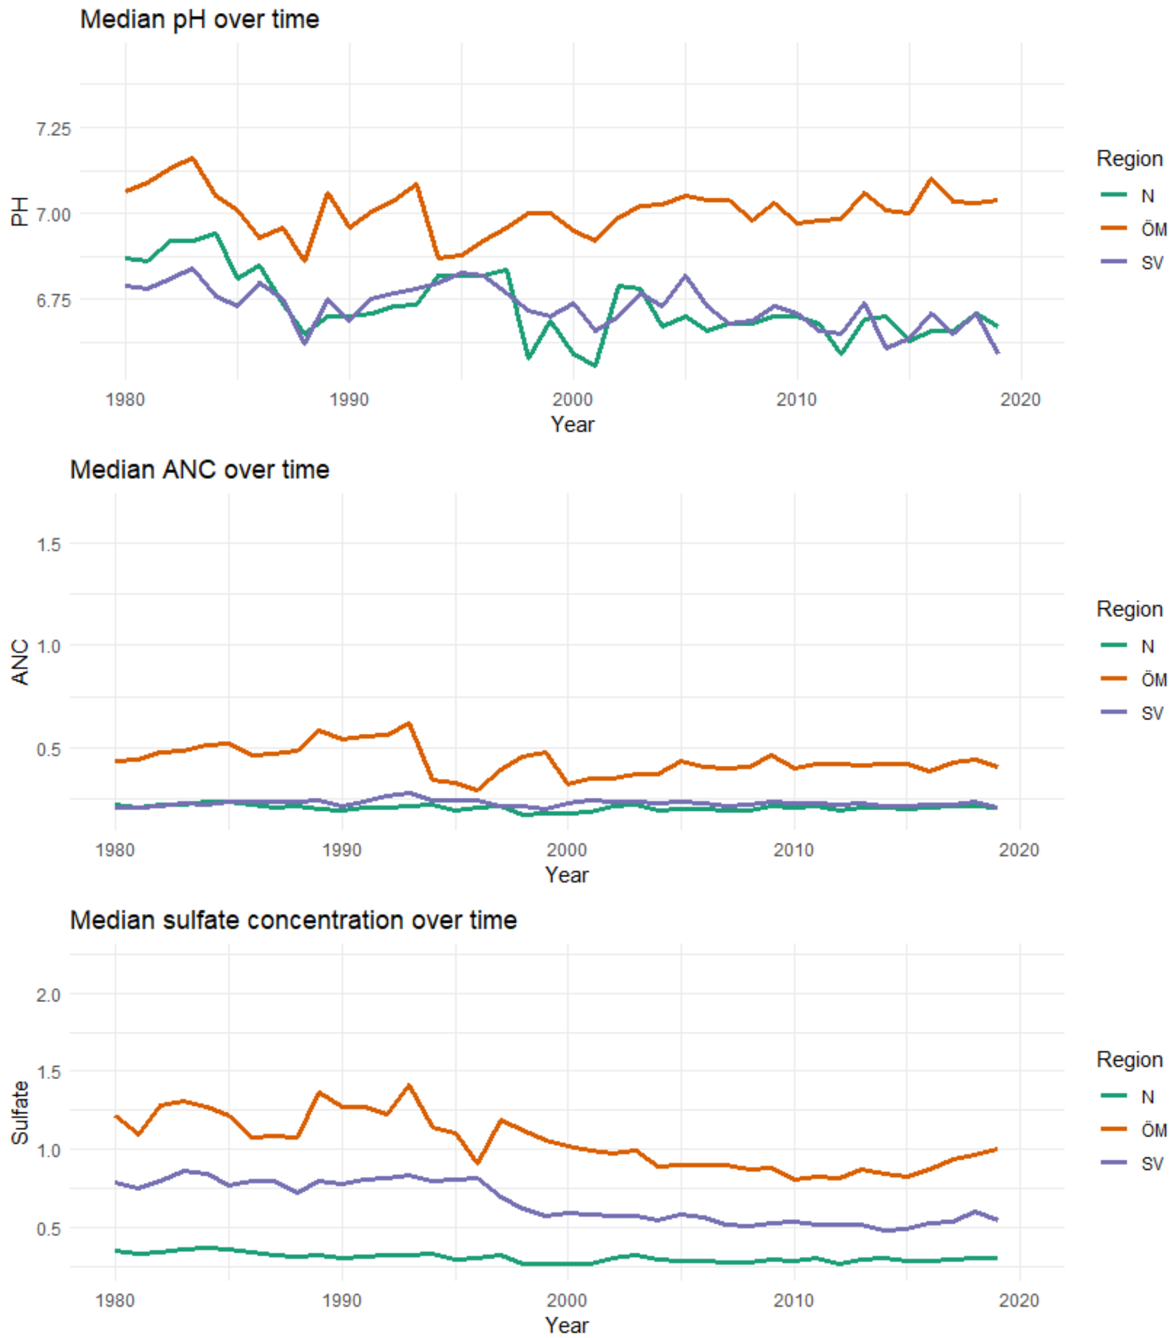

Figure 2: Median PH, ANC and sulfate concentration over time

Even greater differences can be observed when looking at the trends in the 10th percentile of PH and ANC, and in the 90th percentile of sulfate concentration (Figure 3). The 10th percentile of ANC has been stable over time, while PH and sulfate concentrations show some variation over time. In particular, it is interesting to note that the 90th percentile of sulfate concentration has been slowly creeping up in ÖM an SV in the last years, suggesting that the frequency of episodic acidification might have also been increasing.

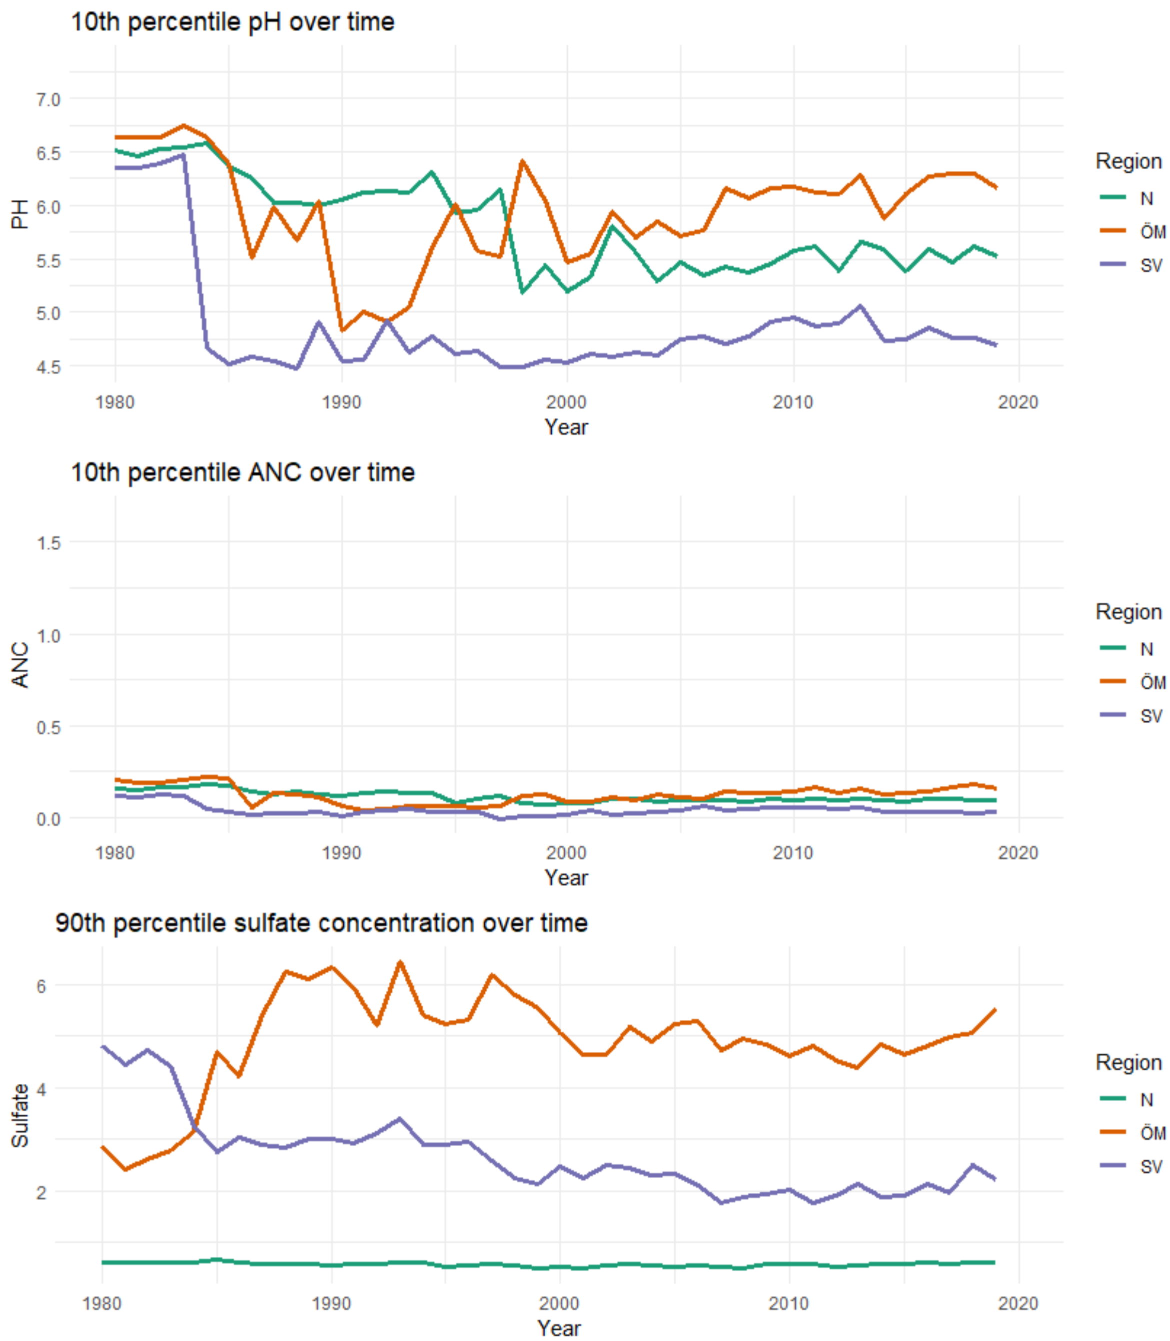

Figure 3: 10th percentile PH and ANC, and 90th percentile sulfate concentration over time

The obvious differences between the trends in average values (Figure 1 and Figure 2) and the trends in the values in the tales of the distribution (Figure 3) offer a clear picture of why an analysis of the trends in episodic acidification is necessary: The trends in the mean and median values provide us with little or no information on the frequency and severity of extreme events. A separate analysis of the trends in episodic acidification is necessary.

3 Methodology

Generalised Additive Models are a useful way of describing environmental trends (von Brömssen et al., 2021). The basic idea behind GAMs (Wood, 2006; Hastie and Tibshirani, 1986) is that a function of the dependent variable y is modelled as a series of splines of some explanatory variables x_i. The main advantage of GAMs over simpler models, such as linear regression, is that GAMs are able to capture both linear and non-linear trends, whereas linear regression assumes that the relationship between the predictors and the output follows a predetermined form.

GAMs are given by the following formula:

$$g(E(Y)) = \beta_0 + f_1(x_1) + f_2(x_2) + ... + f_i(x_i) \tag{1}$$

where $g(E(Y))$ is a function of the output variable, β_0 is an intercept, and $f_1(x_1) + f_2(x_2) + ... + f_i(x_i)$ are smooth functions of the predictors. The smooths can differ between predictors, and be estimated in a number of different ways. Most estimation methods penalise for excessive complexity.

3.1 Quantile Generalised Additive Models

An application of GAMs of particular interest for this research is quantile generalised additive models (Fasiolo et al., 2021), sometimes going under the name of QGAMs. QGAMs aim at modelling the effect of the predictors on a quantile of choice of the dependent variable, rather than on the more common choice of the mean value of the dependent variable. In mathematical terms, a QGAM model can be expressed in the following way:

$$Q_t(Y) = f_1(x_1) + f_2(x_2) + ... + f_i(x_i) \tag{2}$$

where $Q_t(Y)$ is a quantile of choice of the dependent variable. A quantile is, in less technical terms, a percentile out of all possible values of the variable of choice. For instance, if we model the first decile of the dependent variable, what we are actually modelling is the expected effect of the predictors on the 10th percentile of the dependent variable.

The main advantage of quantile regression and QGAMs over models looking at the expected effect of the predictors on the mean of the dependent variable is that quantile based models are more robust for outliers. Furthermore, those models allow for exploring the relationship between the predictors and the values of dependent variable in the tales of the distribution, which is what this research aims to do. By choosing a small or large enough quantile, we can identify the effect of the predictors on the extreme values of the acidifying variable of choice.

3.2 Regional visualisations of QGAMs

The regional plots included in the results section are based on von Brömssen et al. (2021). For a detailed explanation of the reasoning behind those visualisation tools we remind therefore to the aforementioned paper. All of the tools were carefully adapted to accommodate the use of QGAMs, and minor adjustments were made in order to enhance the visual experience in the new context.

3.3 A Generalised Additive Mixed Model for Count Data

An alternative way of identifying trends in extreme values is through a two steps procedure based on GAMMs (Generalised Additive Mixed Models)[1]. In the first step, a GAM model with time as explanatory variable and the acidification variable of choice as dependent variable is fit to each station (Expression 3).

$$E(Y) = f_1(date) + f_2(month) \tag{3}$$

In order to determine what values should be classified as extreme, a confidence interval for the values of the acidification variable is determined, either empirically or on the base of previous research. Many possible arguments may be employed to determine how large this interval should be. All observations falling outside of this interval are recorded as extreme. For the sake of simplicity, we use a wide confidence interval, given by ten times the size of the standard error estimated by the model described by Expression 3. In this way, only the most extreme values are identified as extreme.

In the second stage, we fit a GAMM model with a random intercept for each of the monitoring

[1]see also Pedersen et al. (2019) for a detailed treatment of GAMM models

sites included, in order to predict the number of extreme events likely to occur in a given year (Equation 4). Given the fact that the number of extreme values for a given site in a given year is discrete and usually small, we make use of a negative binomial GAMM model, which assumes that the number of extreme values for a given site in a given year follows a negative binomial distribution, and allows for over- and underdispersion. In order to account for unbalanced sampling over years and stations, we also include an offset representing the number of the available observations.

$$E(log(count\ of\ extremes)) = f(date) + \zeta_{site} \tag{4}$$

In expression 4, $E(log(count\ of\ extremes))$ is the number of extreme values for a given site in a given year, and ζ_{site} is the site-specific random intercept. By looking at the effect of time ($f(date)$) on the expected number of extreme events, site specific and overall trends can be obtained for any geographical areas of interest.

4 Preliminary results

The results in this section are to be seen as preliminary in the meaning that this research is still ongoing. The main goal of this section is to give a picture of what kind of questions can be answered with the help of QGAMs and specific visualisation tools adapted to those models. A number of relevant examples are presented based on the dataset first introduced in Section 2.

4.1 Quantile Generalised Additive Models

The quantile additive model presented in Section 3 can be used for the purpose of analysing the long-term trends in extreme values of acidifying variables for a given site.

We take as an example "Mälskarbäcken", a monitoring site in Northern Sweden. Figure 4 suggests an answer the question of whether extreme values in acidifying variables have become more or less frequent over time. Significant increases in the value of the quantile are marked in red, and significant decreases are marked in blue. Periods of time without significant changes are described by the colour yellow.

An increase in the value of a given quantile over a certain period of time suggests that the frequency of events above that value has increased, whereas a decrease in the value of that same quantile suggests that the number of events above that threshold has decreased. In other words, a red marking in the plots visualising changes in sulfate concentration over time can be interpreted as a sign of an increased frequency in the number of recordings of extreme values of sulfate concentration, whereas this same interpretation can instead be given to blue markings in the case of pH and ANC.

Figure 4 suggests that the frequency of extreme events at Mälskarbäcken has not increased markedly over time, and that the trend for the last years is rather stable.

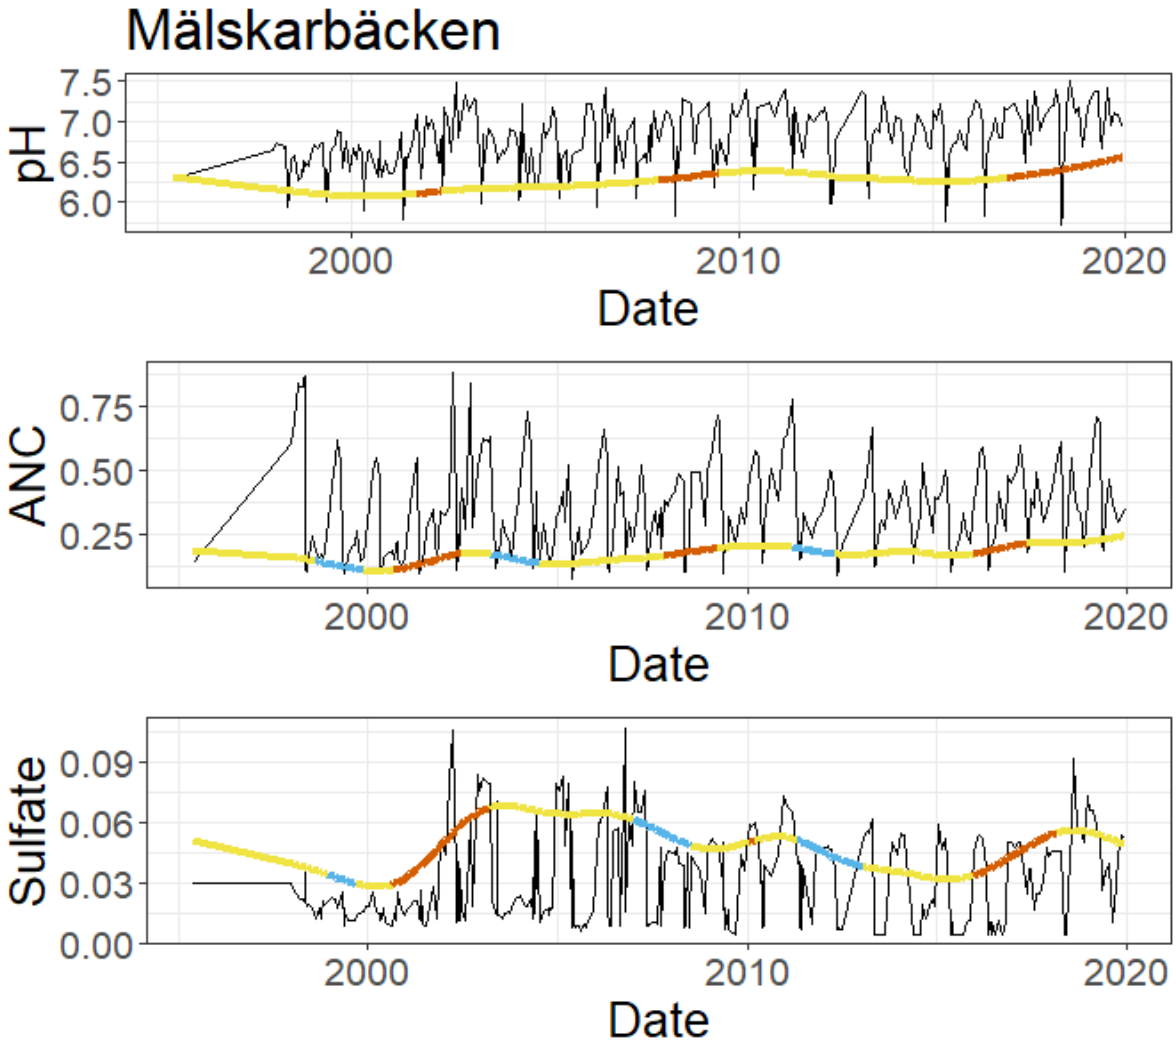

Figure 4: Mälskarbäcken, significant trend changes in pH (10th percentile), ANC (10th percentile) and sulfate (90th percentile) over time

It is important to note that the results we obtain depend on the significance level we choose to determine significant trend changes, and on how we define the extreme values. Figure 4 uses a significance level of 0.1, and considers values in the top (for sulfate) and bottom (for pH and ANC) decile respectively to be extreme. If we change some of those values, the results we obtain may be different.

Figure 5 provides an example of how different significance levels may affect the results: in this case, a lower significance level was used (0.05), and only the values in the top (sulfate) and bottom (pH and ANC) 5th percentile were identified as extreme.

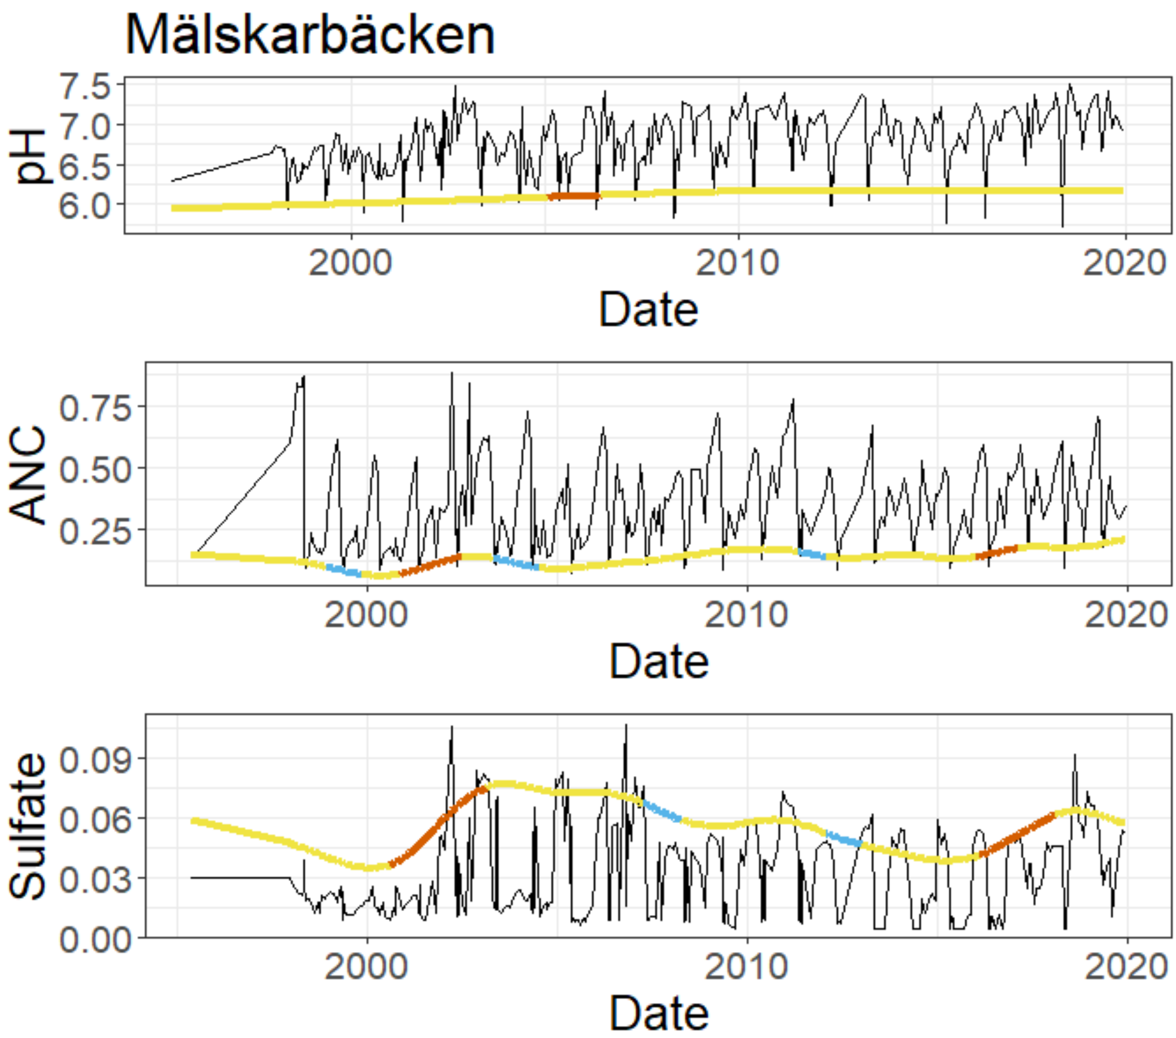

Figure 5: Mälskarbäcken, significant trend changes in pH (5th percentile), ANC (5th percentile) and sulfate (95th percentile) over time

The overall conclusion is the same: episodic acidification has not become more common in the last years. However, a number of differences in the trends of episodic acidification can be identified: For instance, whereas before it appeared that the number of extreme events in pH had become less common in the last five years, it now appears that they are approximately as common now as they were about five years ago. In addition, fewer significant changes occur.

A limit of Figure 4 and Figure 5 is that they concern a very limited geographical area, whereas we might be interested in the overall trends for multiple sites. Trends in episodic acidification might in fact be better understood when looking at several streams at once, as acidic deposition is likely to affect all sites in a given area in similar ways. The figures in the next section offer a remedy to this issue, by allowing us to look at multiple sites at once.

4.2 Regional visualisations of QGAMs

The quantile additive model presented in Section 3 can also be used for the purpose of analysing long-term trends in extreme values in acidifying variables for a number of sites at the same time. In this subsection, we offer some examples of how this can be done.

We begin by looking at five randomly chosen sites in Southern and Western Sweden (SV, Figure 6). To the left, we can see the coordinates of the individual sites according to the "SWEREF" system (latitude and longitude). The sites are ordered from North to South (largest latitude first). The colour scheme is the same as the one used in the plots in Subsection 4.1

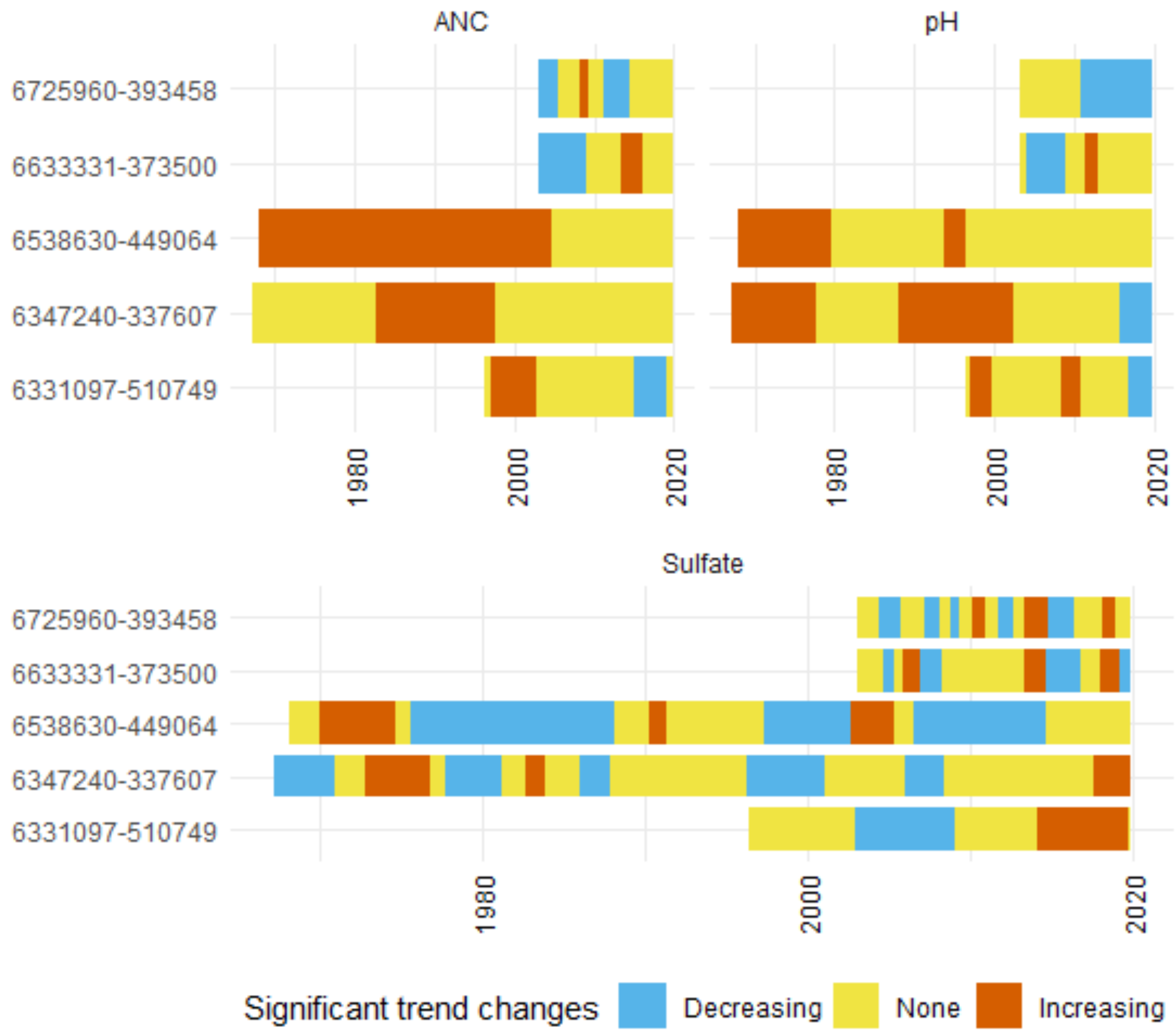

Figure 6: SV sample (five monitoring sites), significant trend changes (0.1 significance level) in ANC (10th percentile), pH (10th percentile), and sulfate concentration (90th percentile) over time

Figure 6 shows that three out of five streams have been experiencing more frequent extreme values of pH and sulfate concentration in the last years. However, this figure gives us only limited information on the overall trends in the region, as only five stations are included in this sample. Figure 7 and Figure 8 look at the trend for all monitoring sites in the region.

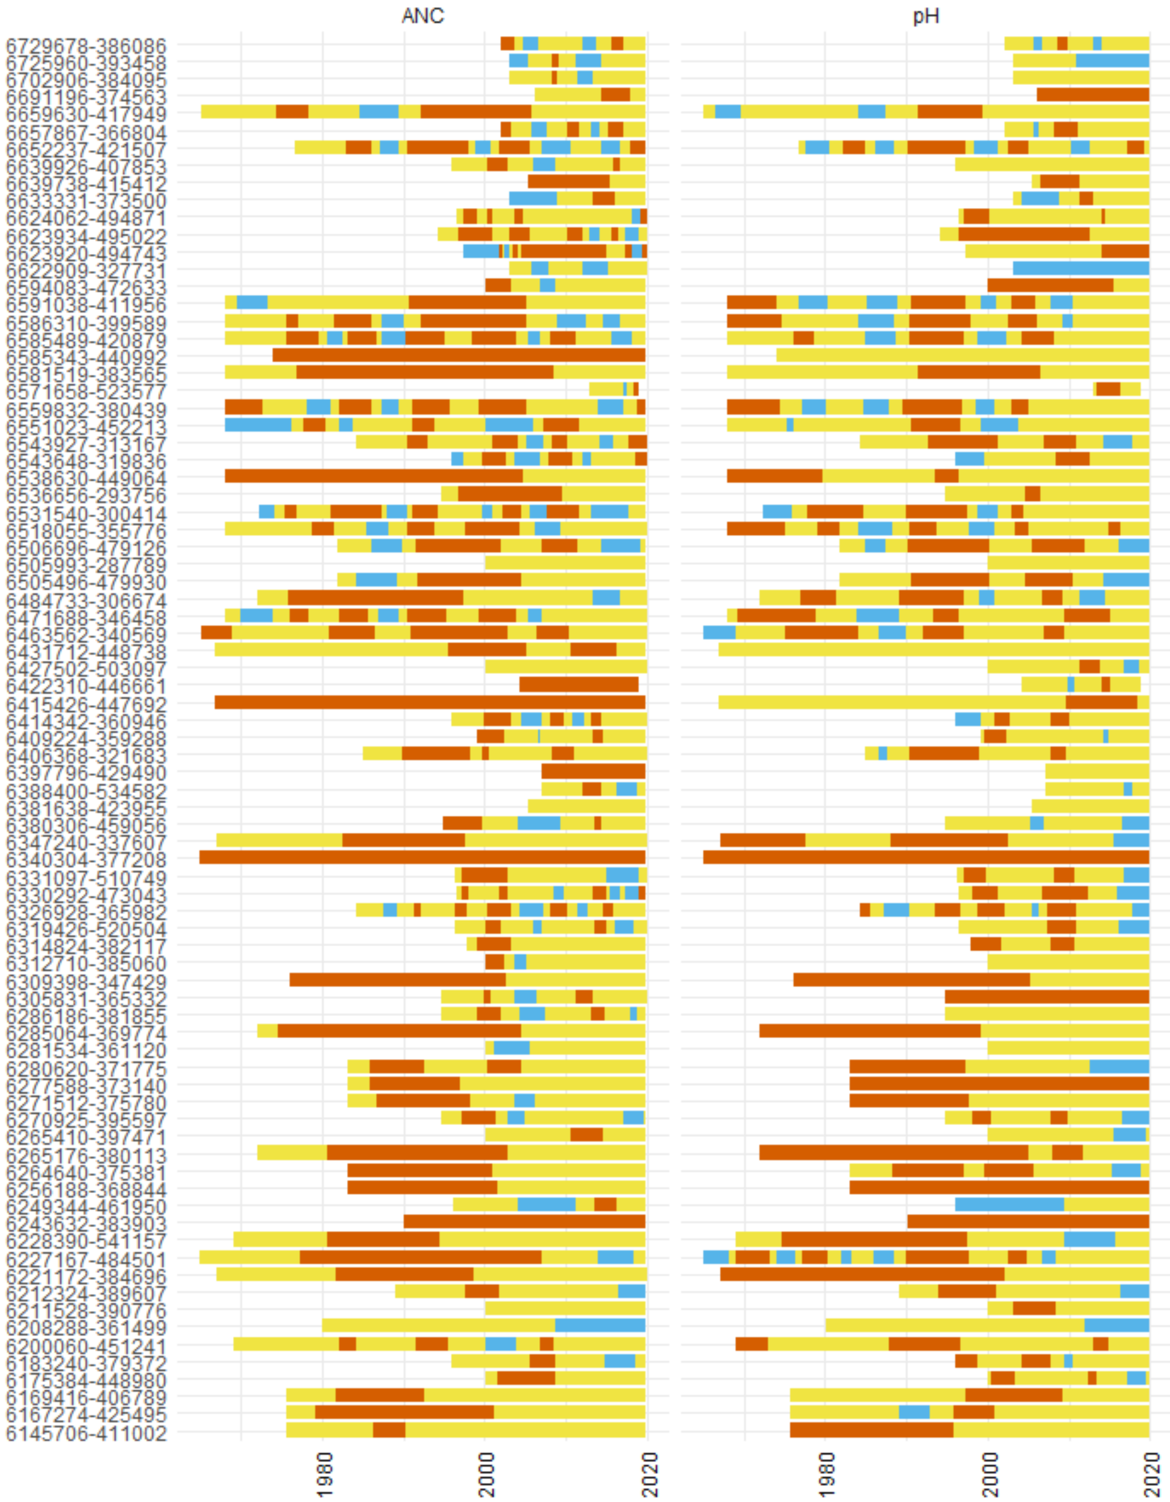

Figure 7: SV, significant trend changes (0.1 significance level) in ANC (10th percentile) and pH (10th percentile) over time

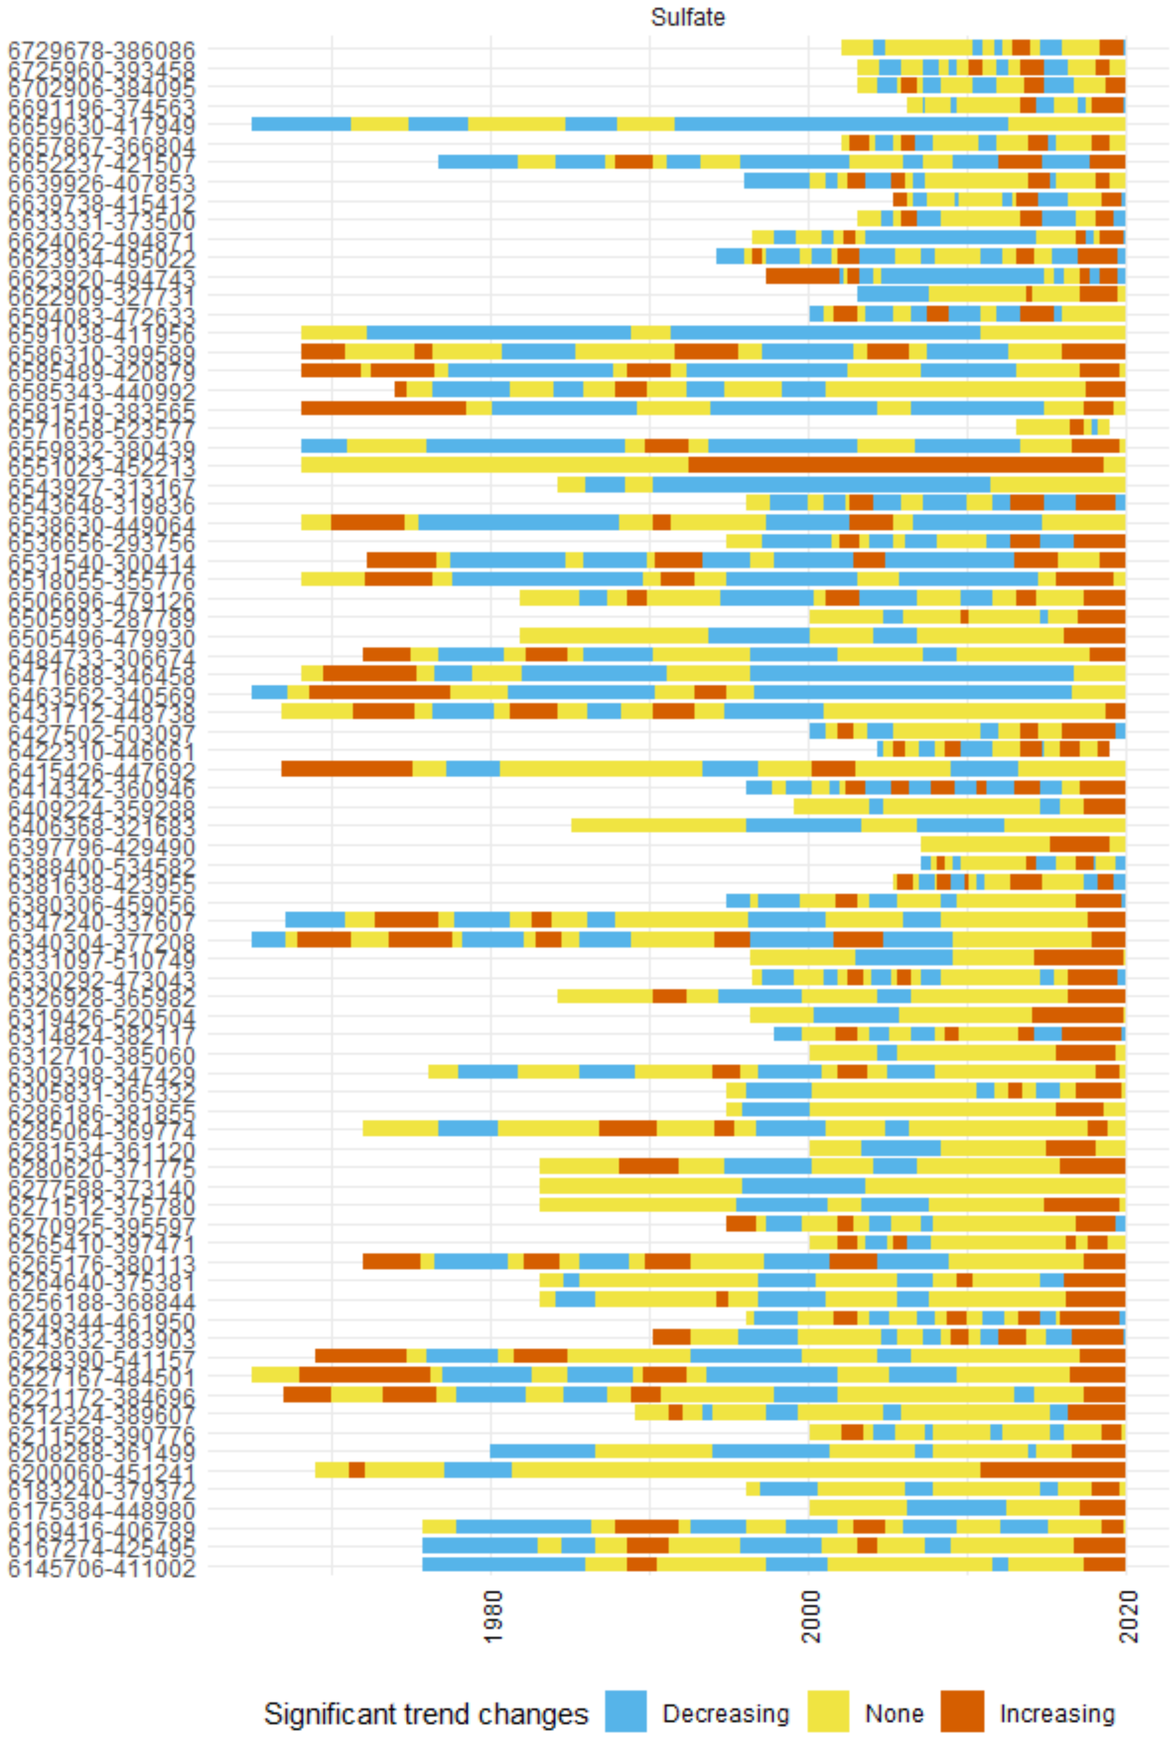

Figure 8: SV, significant trend changes (0.1 significance level) in sulfate concentration (90th percentile) over time

Figure 7 and Figure 8 offer a better picture of the overall trends in extreme values in the region. Figure 7 suggests that the frequency of extremely low ANC values has been rather stable over the last years, while the number of episodes of extremely low pH values has increased for some of the sites. Figure 8 reinforces the picture that episodic acidification has become more common, as more cases of extremely high values of sulfate concentration have been recorded in most sites.

A negative side of Figure 7 and Figure 8 is that they require much space in order to be easily readable. Furthermore, it might be hard to distinguish overall regional trends when trends from individual streams show contrasting results. For those reasons, an alternative plot, which better summarises overall trends, is also provided (Figure 9).

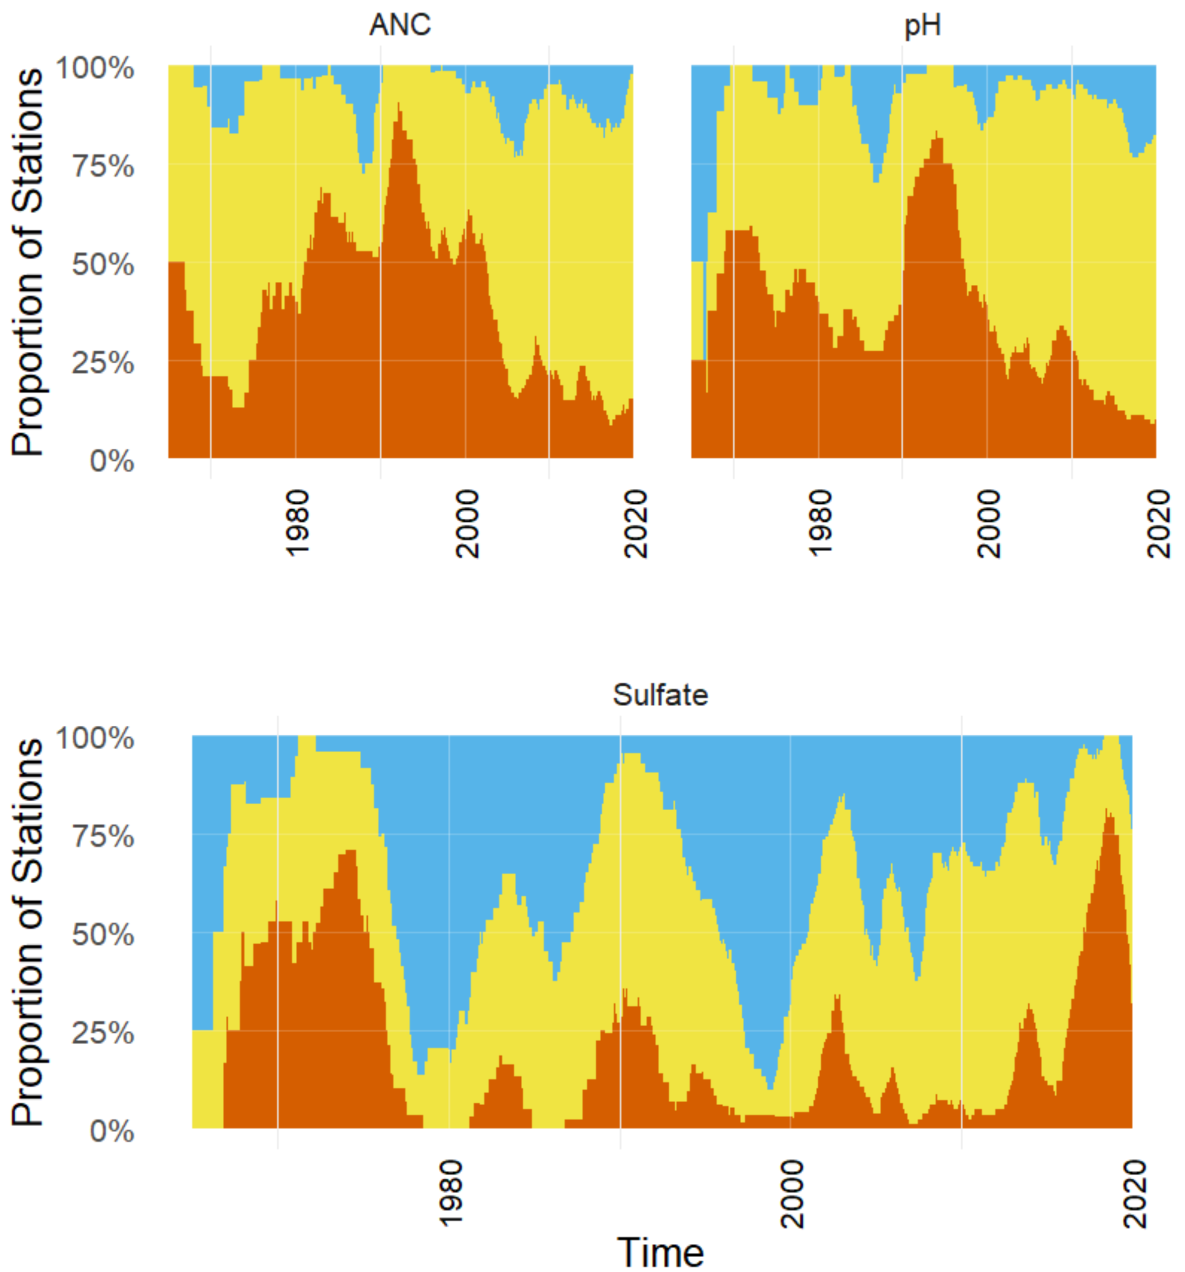

Figure 9: SV, percentage of sites per year experiencing significant trend changes (0.1 significance level) in ANC (10th percentile), PH (10th percentile) and sulfate (90th percentile)

Figure 9 gives a better overview of the overall trends in the region. It is clear that recordings of extremely high sulfate concentrations have become more common in most sites during the last years, while the frequency of extreme values in ANC and pH has remained more stable.

A possible limitation of all of the plots included in this section is that they might be excessively prone to show a large number of significant trend changes. The excess "wiggliness" in the trend estimate is due to the fact that the QGAMs we employed do not take into account the fact that the series might be autocorrelated. This might lead to an increased number of up and downs and the series, resulting in a large number of significant trend changes in the plots.

4.3 A Generalised Additive Mixed Model for Count Data

An alternative approach, first introduced in Section 3.2, is to use a two-stage GAMM model to estimate a long time trend in the number of extreme events per year.

An example is provided for three nearby rivers in the Gävleborg county: Ljusnan, Delångersån and Dalälven. Figure 10 shows the overall trend in episodic acidification for the three rivers, in black, and for the individual sites.

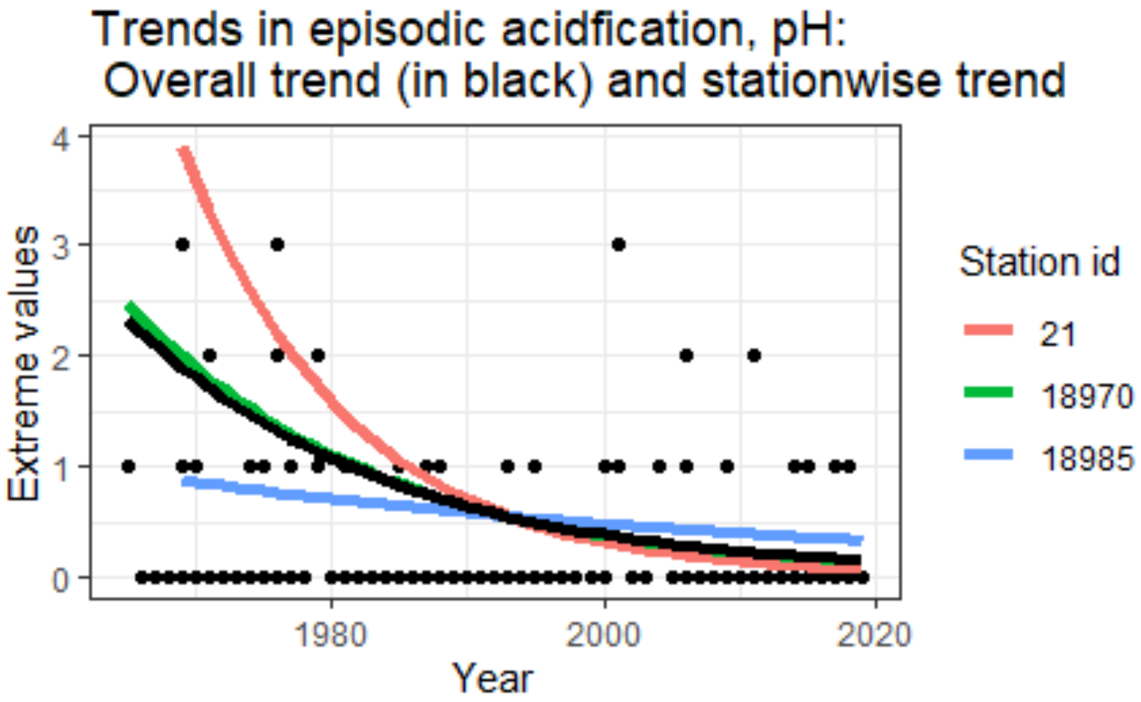

Figure 10: Trends in number of yearly events of episodic acidification (pH). Overall trend (in black) and trends for individual sites

Figure 10 suggests that the number of extreme events per year has decreased over time, both for the river as a whole and for the individual sites. We might, however, also want to try to quantify how sure we are about our conclusions, especially when it comes to the overall trend. An attempt to provide a solution to this question is made by Figure 11, which includes confidence intervals representing the levels of uncertainty surrounding our prediction for the overall trend.

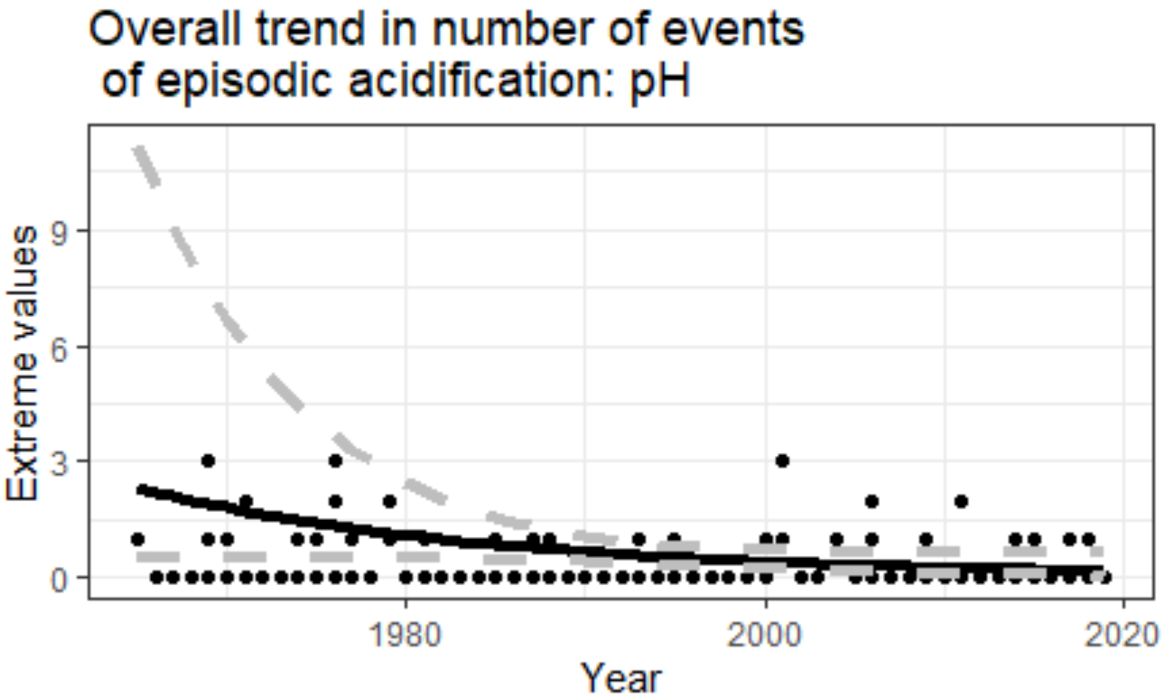

Figure 11: Trend in number of yearly events of episodic acidification (pH). Overall trend (in black) and confidence interval for the overall trend (dashed lines in grey)

Figure 11 is useful to get a first impression of how sure we are about the overall direction and size of the trend. However, two limitations affect the estimated confidence intervals: Firstly, the estimated intervals do not take into account the uncertainty introduced by the estimation process taking place in the first stage of the modelling process, described by Expression 3. Secondly, it is hard to correctly determine how wide the intervals should be in order to represent a confidence level of 95%: for this example, we made use of the traditional formula 1.96 x SE, but it is possible that we might be underestimating the overall level of uncertainty. A better way of computing the levels of uncertainty surrounding our prediction is among the future goals of this research project.

Similar plots can be obtained for any other acidifying variables of interest. For instance, Figure 12 and Figure 13 display, respectively, the overall trend in extreme values of ANC, and the trend in sulfate concentration for the river Ljusnan. Figure 13 represents an interesting example of why the estimated confidence intervals cannot always be relied on, especially if many of the observed values are zero: after the year 1990, the lower boundary of the interval overlaps with zero, while the upper boundary becomes progressively larger despite the fact that very few occurrences of episodic acidification could actually be observed.

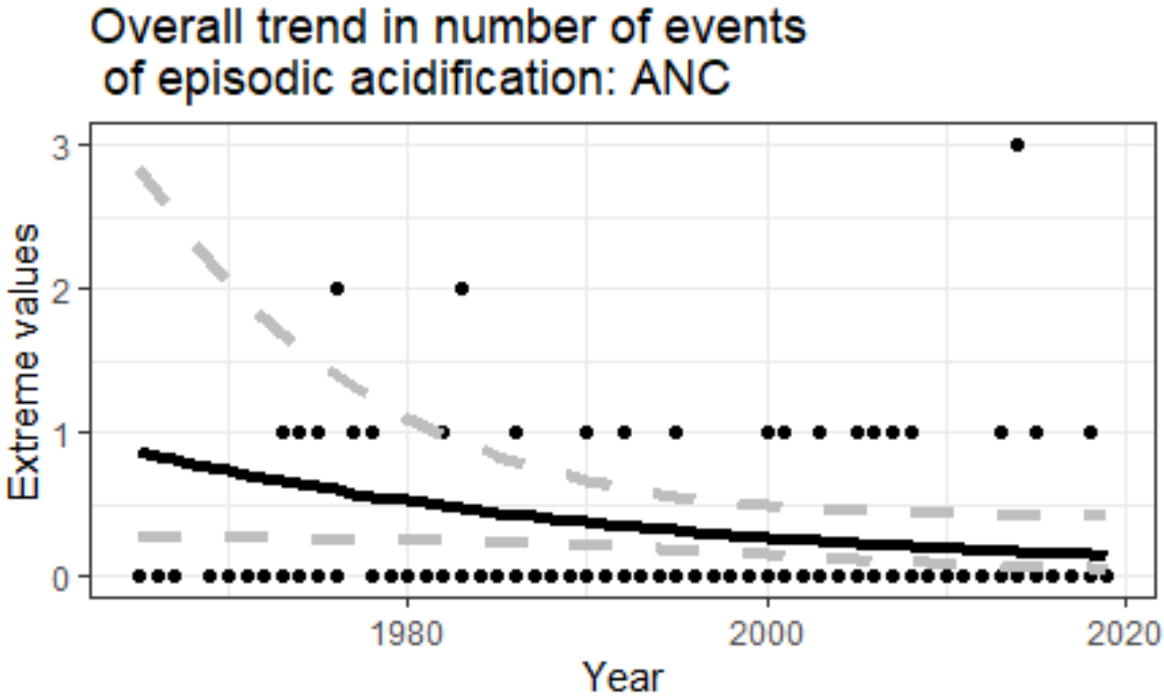

Figure 12: Trend in number of yearly events of episodic acidification (ANC). Overall trend (in black) and confidence interval for the overall trend (dashed lines in grey)

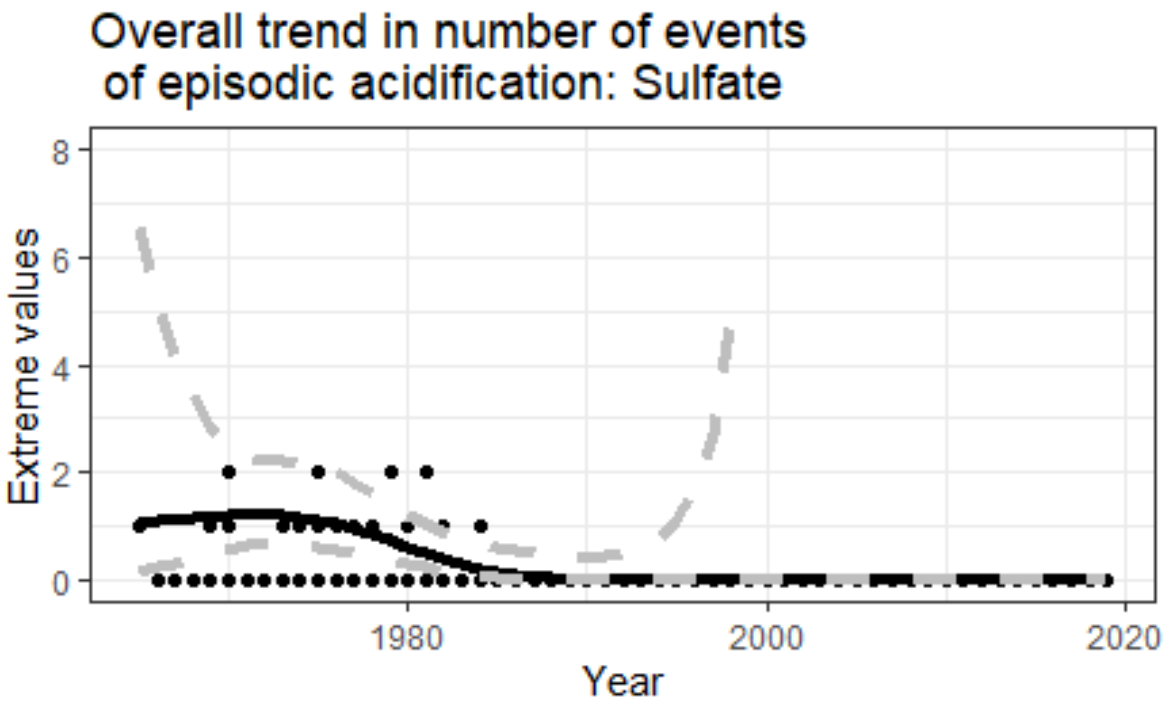

Figure 13: Trend in number of yearly events of episodic acidification (sulfate concentration). Overall trend (in black) and confidence interval for the overall trend (dashed lines in grey)

5 Conclusion

This report has been looking at the ways that trends of extreme values in acidifying variables can be identified with the help of QGAMs, Negative Binomial GAMMs, and visualisation tools adapted from von Brömssen et al. (2021). The need for those models is highlighted by Section 2, where the differences between long-term trends in average and extreme values of acidifying variables are introduced to the reader. A number of applications of the models and of the visualisation tools are presented in Section 4, where both local and regional data are analysed and visualised. Additional plots for the regional data can be found in the Appendix.

A strength of the models and visualisations tools presented in this paper is that they can be used equally effectively to identify patterns in extreme values for a single site and for a large number of monitoring sites at once. For instance, Figure 17, Figure 18 and Figure 19 offer an overview of how the number of extreme values has been changing in the whole of Southern and Western Sweden over time, whereas Figure 4 and Figure 5 look at just one of the monitoring sites in Northern Sweden. Another strength of these tools is that they are very flexible: they can be used to analyse the patterns of episodic acidification for any variable of interest.

One of the main limitations of QGAMs is that they might be harder to understand and apply than traditional regression and time-series models. Another important limitation is that their results can be affected by how the extreme values are defined and identified: More significant changes in the frequency of episodic acidification are likely to be detected when a larger number of events are recorded as extreme. The fact that QGAMs do not account for any autocorrelation in the series at the station level might also be a limitation.

Similar limitations also affect the GAMMs: how extreme should an event of episodic acidification be in order to be recorded as extreme? Another limitation has to do with the quantification of uncertainty surrounding the estimated trends: the confidence intervals presented in Subsection 4.3 are likely to underestimate the actual levels of uncertainty of the estimation, as they do not take into account the uncertainty introduced by the first step of the estimation process (Expression 3).

List of References

J. P. Baker, J. V. Sickle, C. J. Gagen, D. R. DeWalle, W. E. Sharpe, R. F. Carline, B. P. Baldigo, P. S. Murdoch, D. W. Bath, W. A. Krester, H. A. Simonin, and P. J. Wigington. Episodic Acidification of Small Streams in the Northeastern United States: Effects on Fish Populations. *Ecological Applications*, 6(2):422–437, 1996. ISSN 1939-5582. doi: 10.2307/2269380.

M. Erlandsson, H. Laudon, and J. Fölster. Spatiotemporal patterns of drivers of episodic acidification in Swedish streams and their relationships to hydrometeorological factors. *Science of The Total Environment*, 408(20):4633–4643, Sept. 2010. ISSN 0048-9697. doi: 10.1016/j.scitotenv.2010.06.010.

European Environment Agency. 4.4. Monitoring of surface water acidification, 2016. URL `https://www.eea.europa.eu/publications/92-9167-001-4`.

M. Fasiolo, S. N. Wood, M. Zaffran, R. Nedellec, and Y. Goude. qgam: Bayesian Nonparametric Quantile Regression Modeling in R. *Journal of Statistical Software*, 100:1–31, Nov. 2021. ISSN 1548-7660. doi: 10.18637/jss.v100.i09. URL `https://doi.org/10.18637/jss.v100.i09`.

J. Fölster, R. K. Johnson, M. N. Futter, and A. Wilander. The Swedish monitoring of surface waters: 50 years of adaptive monitoring. *AMBIO*, 43(1):3–18, Dec. 2014. ISSN 1654-7209. doi: 10.1007/s13280-014-0558-z. URL `https://doi.org/10.1007/s13280-014-0558-z`.

T. Hastie and R. Tibshirani. Generalized Additive Models. *Statistical Science*, 1(3):297–310, Aug. 1986. ISSN 0883-4237, 2168-8745. doi: 10.1214/ss/1177013604. Publisher: Institute of Mathematical Statistics.

R. M. Heard, W. E. Sharpe, R. F. Carline, and W. G. Kimmel. Episodic Acidification and Changes in Fish Diversity in Pennsylvania Headwater Streams. *Transactions of the American Fisheries Society*, 126(6):977–984, Nov. 1997. ISSN 0002-8487. Publisher: Taylor & Francis.

H. Laudon. Recovery from episodic acidification delayed by drought and high sea salt deposition. *Hydrology and Earth System Sciences*, 12(2):363–

370, Mar. 2008. ISSN 1027-5606. doi: 10.5194/hess-12-363-2008. URL `https://hess.copernicus.org/articles/12/363/2008/`. Publisher: Copernicus GmbH.

F. Moldan, B. J. Cosby, and R. F. Wright. Modeling Past and Future Acidification of Swedish Lakes. *AMBIO*, 42(5):577–586, Sept. 2013. ISSN 0044-7447, 1654-7209. doi: 10.1007/s13280-012-0360-8. URL `http://link.springer.com/10.1007/s13280-012-0360-8`.

E. J. Pedersen, D. L. Miller, G. L. Simpson, and N. Ross. Hierarchical generalized additive models in ecology: an introduction with mgcv. *PeerJ*, 7: e6876, May 2019. ISSN 2167-8359. doi: 10.7717/peerj.6876. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6542350/`.

C. von Brömssen, S. Betnér, J. Fölster, and K. Eklöf. A toolbox for visualizing trends in large-scale environmental data. *Environmental Modelling & Software*, 136:104949, Feb. 2021. ISSN 1364-8152. doi: 10.1016/j.envsoft.2020.104949.

S. Wood. *Generalized Additive Models: An Introduction With R*, volume 66. Jan. 2006. ISBN 978-1-315-37027-9.

R. F. Wright and B. J. Cosby. Future recovery of acidified lakes in southern Norway predicted by the MAGIC model. *Hydrology and Earth System Sciences*, 7:467–483, Jan. 2003. ISSN 1027-5606. doi: 10.5194/hess-7-467-2003.

Appendix

Regional visualisations QGAMs: Northern Sweden

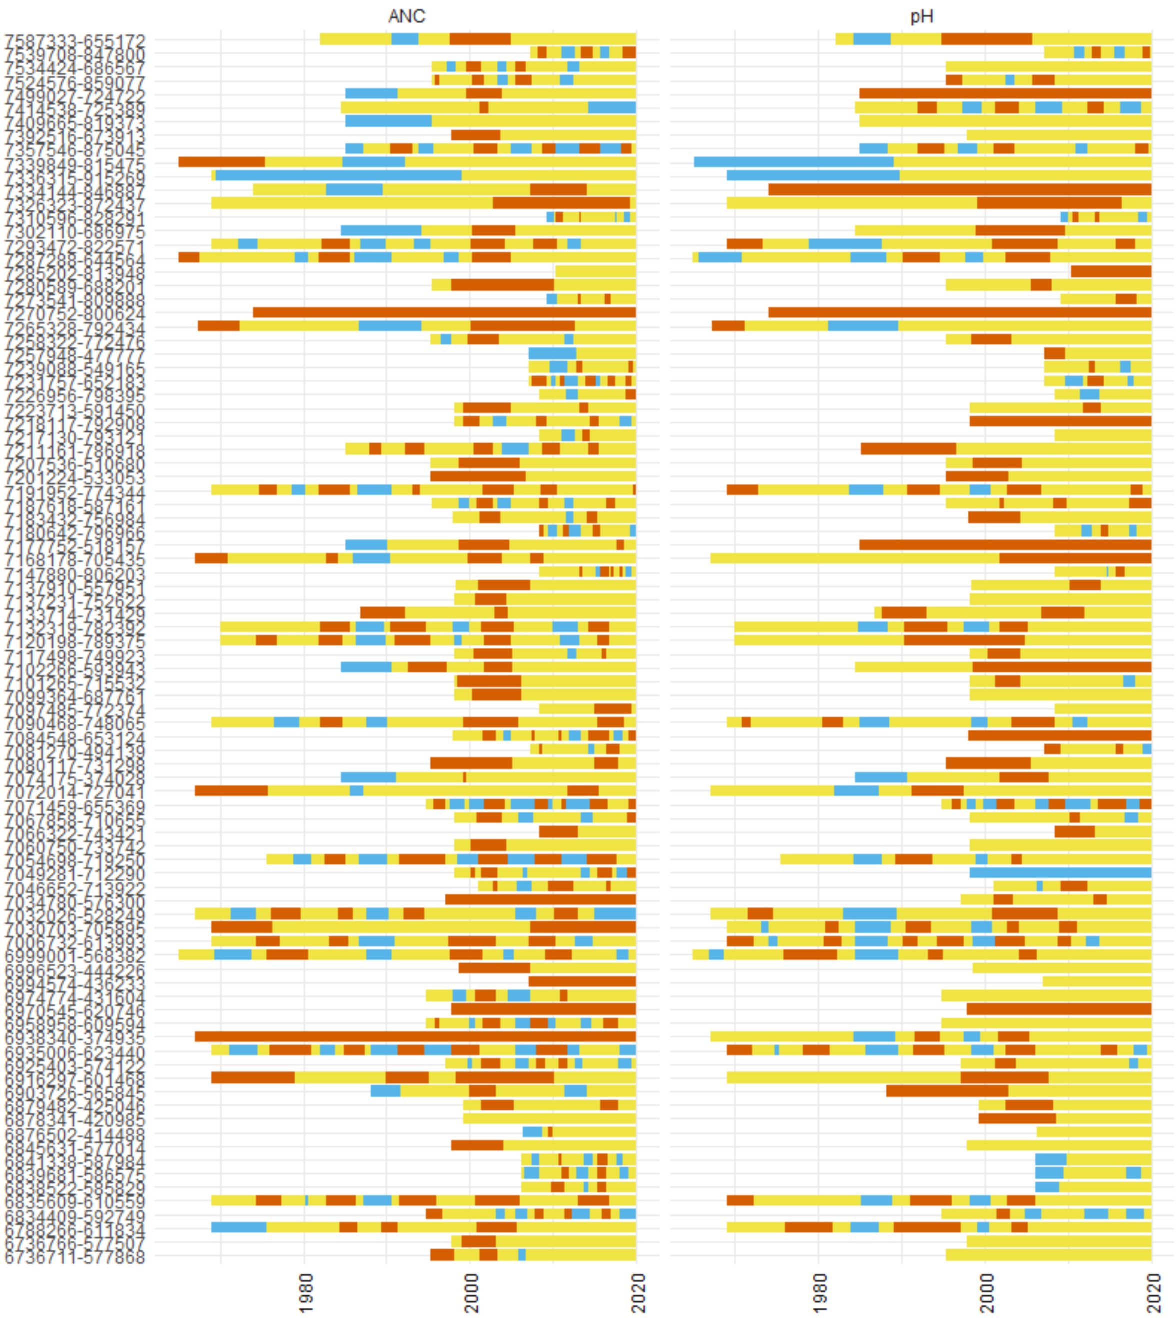

Figure 14: N, significant trend changes (0.1 significance level) in ANC (10th percentile) and pH (10th percentile) over time

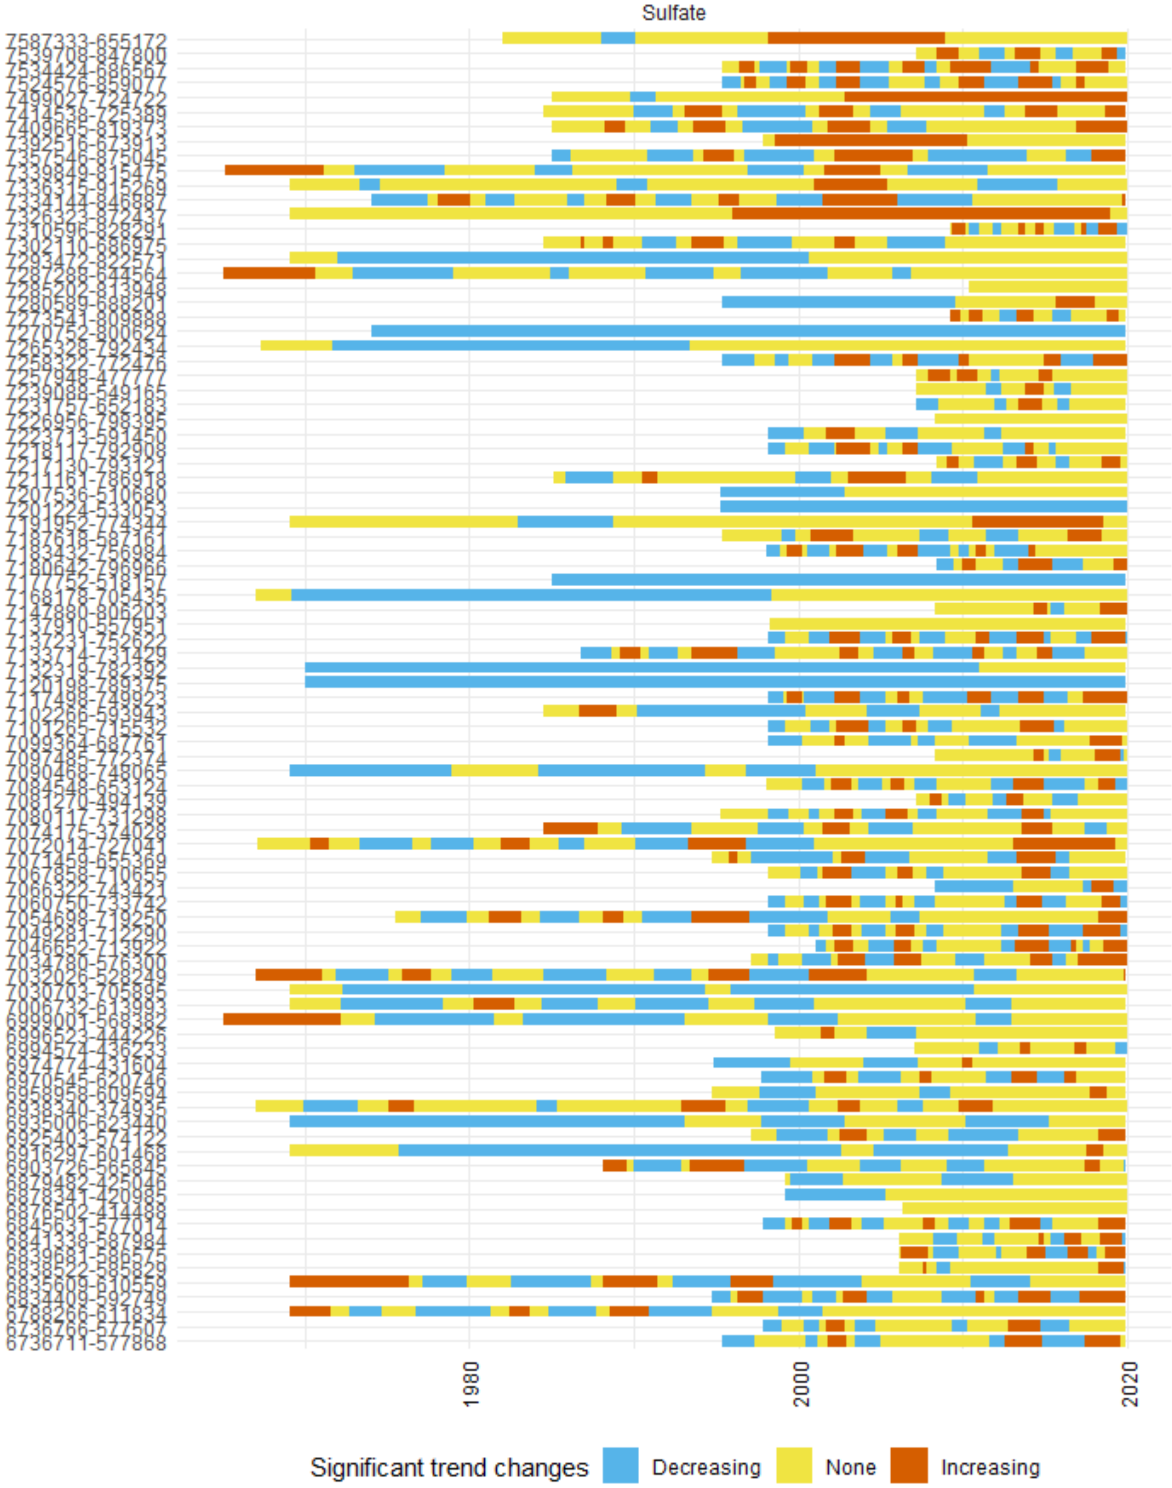

Figure 15: N, significant trend changes (0.1 significance level) in sulfate concentration (90th percentile) over time

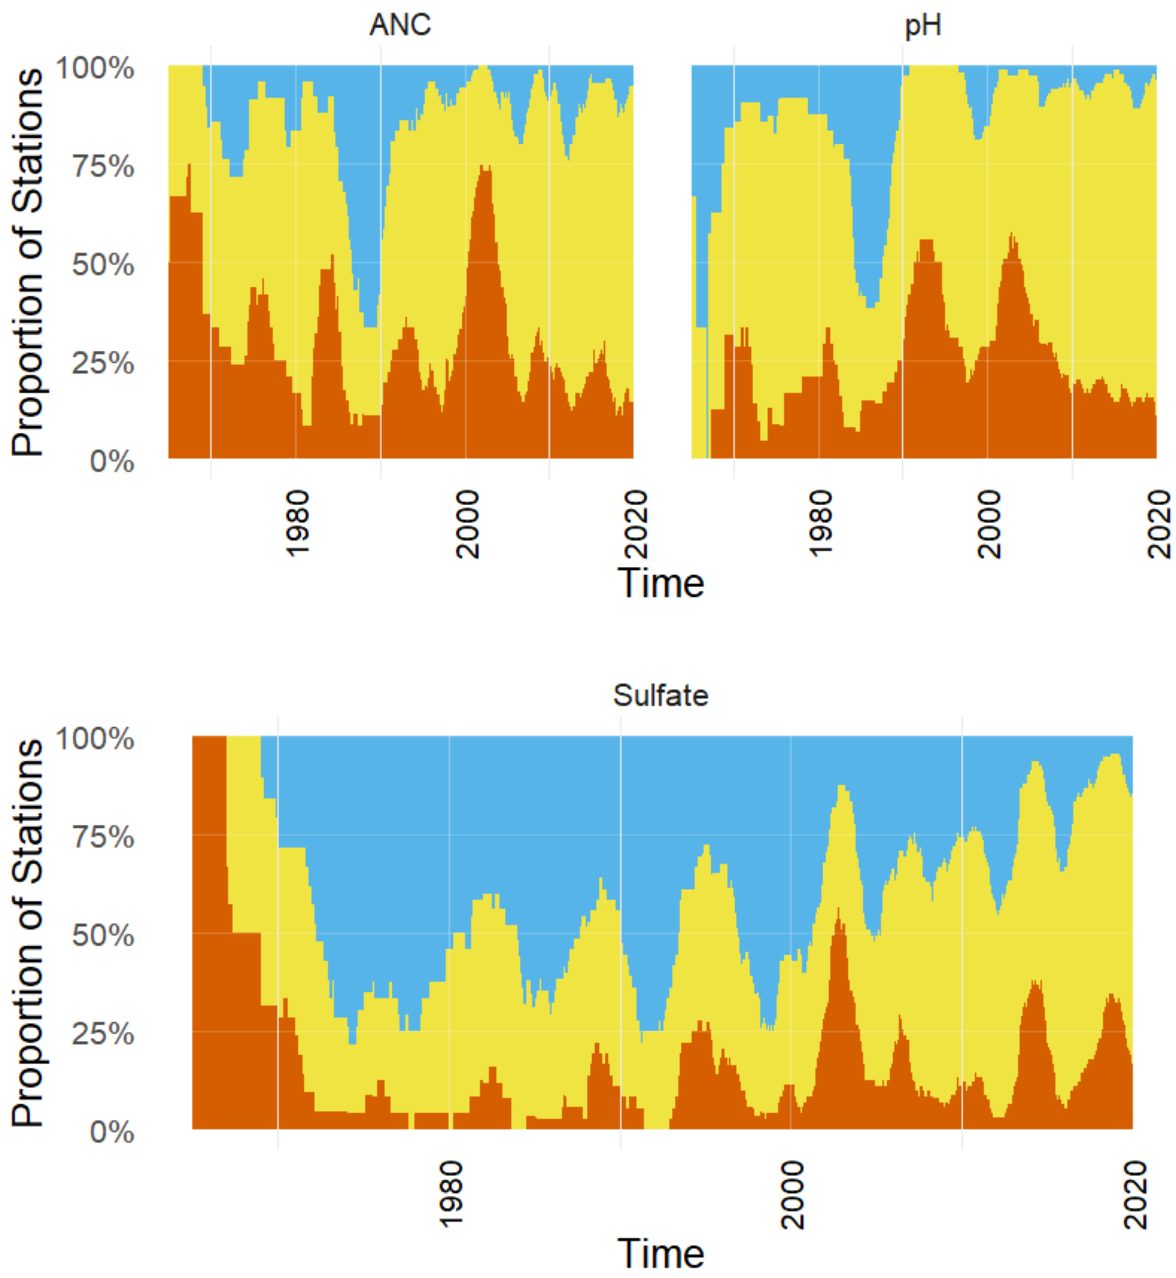

Figure 16: N, percentage of sites per year experiencing significant trend changes (0.1 significance level) in ANC (10th percentile), PH (10th percentile) and sulfate concentration (90th percentile)

Regional visualisations of QGAMs: Eastern and Middle Sweden

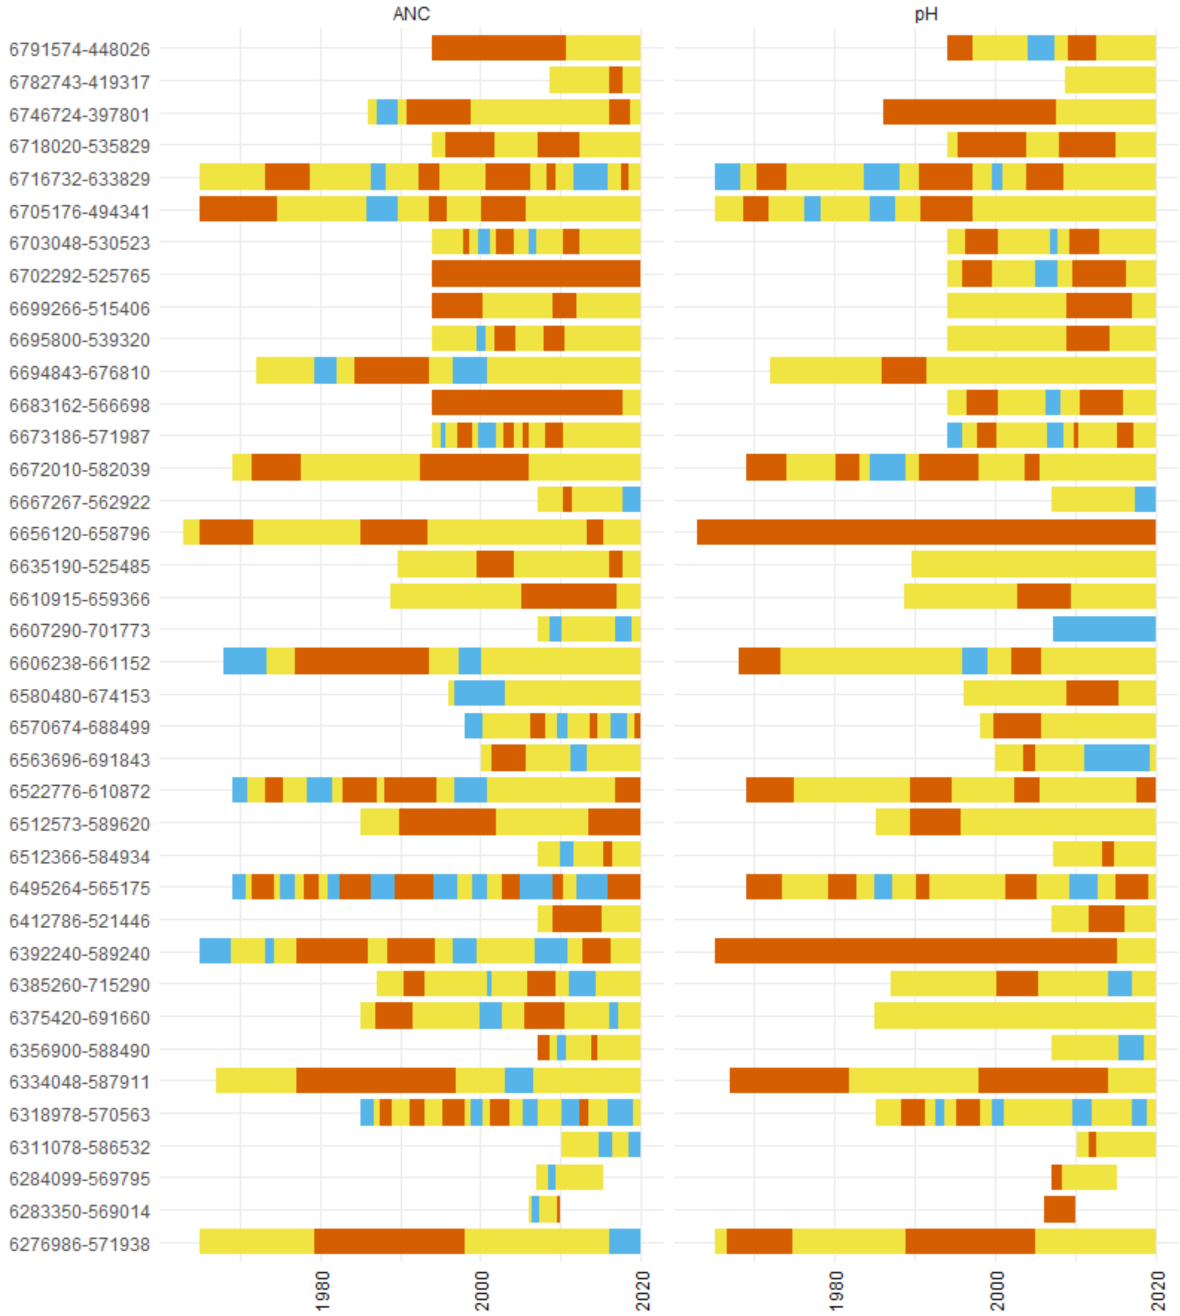

Figure 17: ÖM, significant trend changes (0.1 significance level) in ANC (10th percentile) and pH (10th percentile) over time

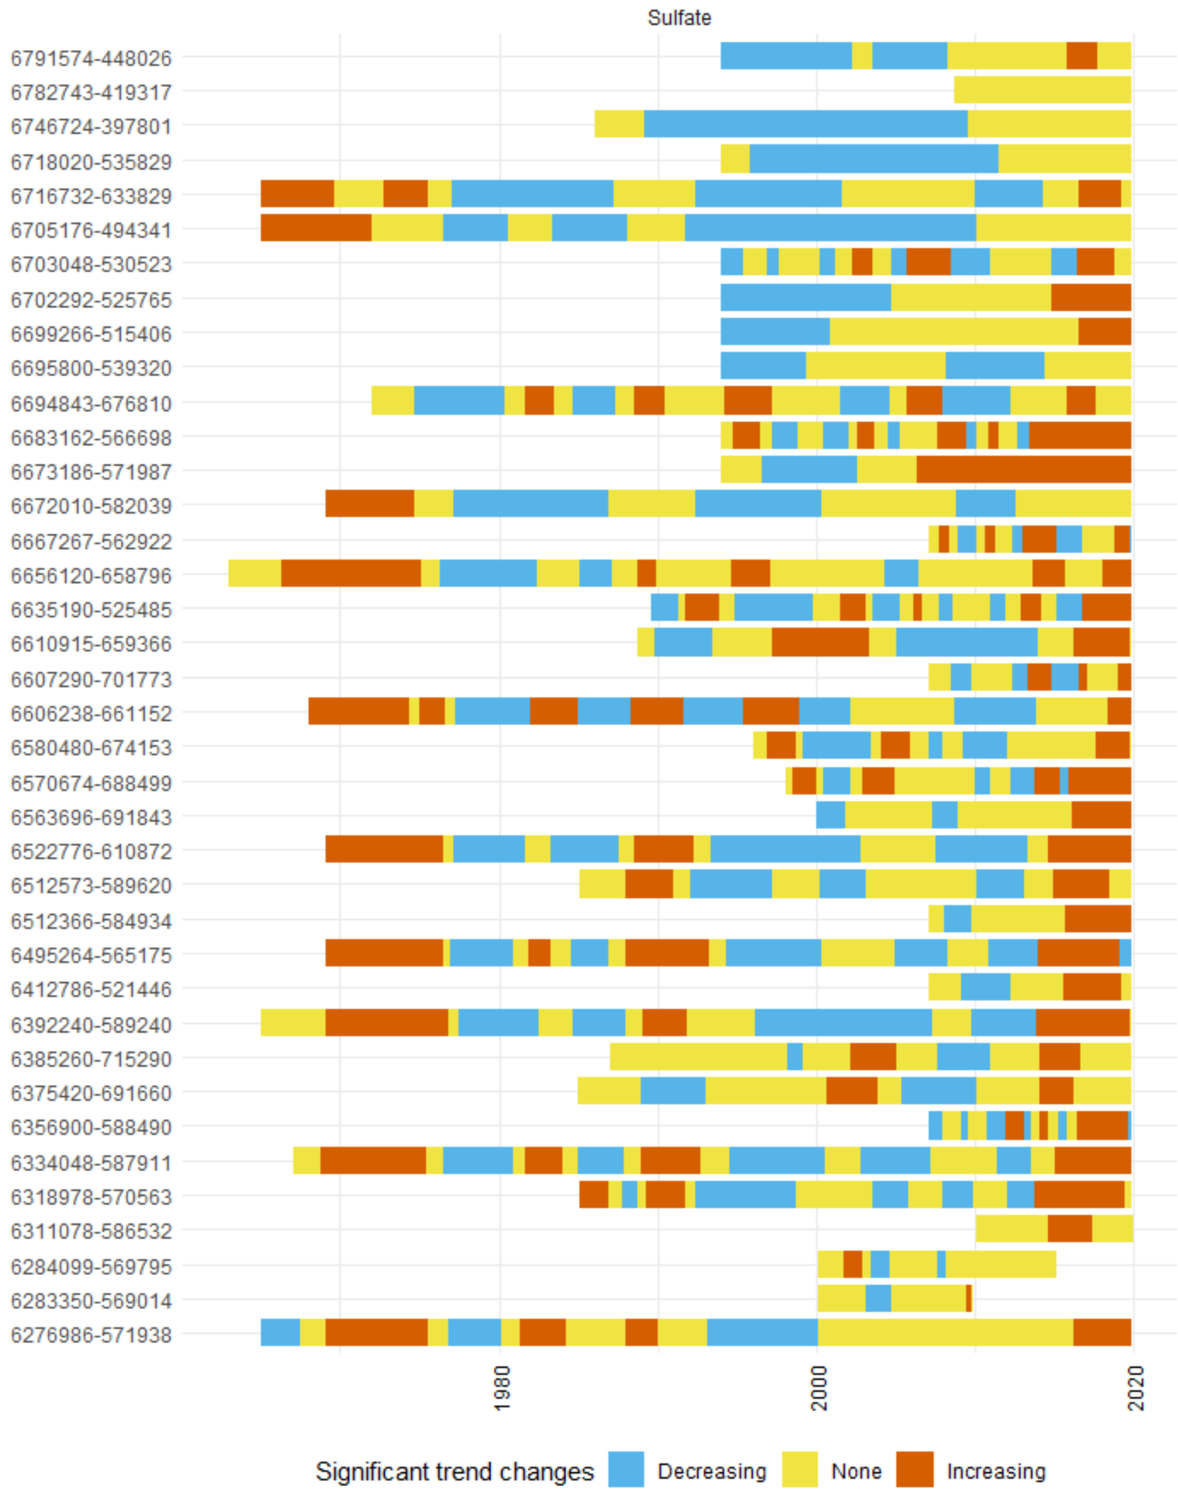

Figure 18: ÖM, significant trend changes (0.1 significance level) in sulfate concentration (90th percentile) over time

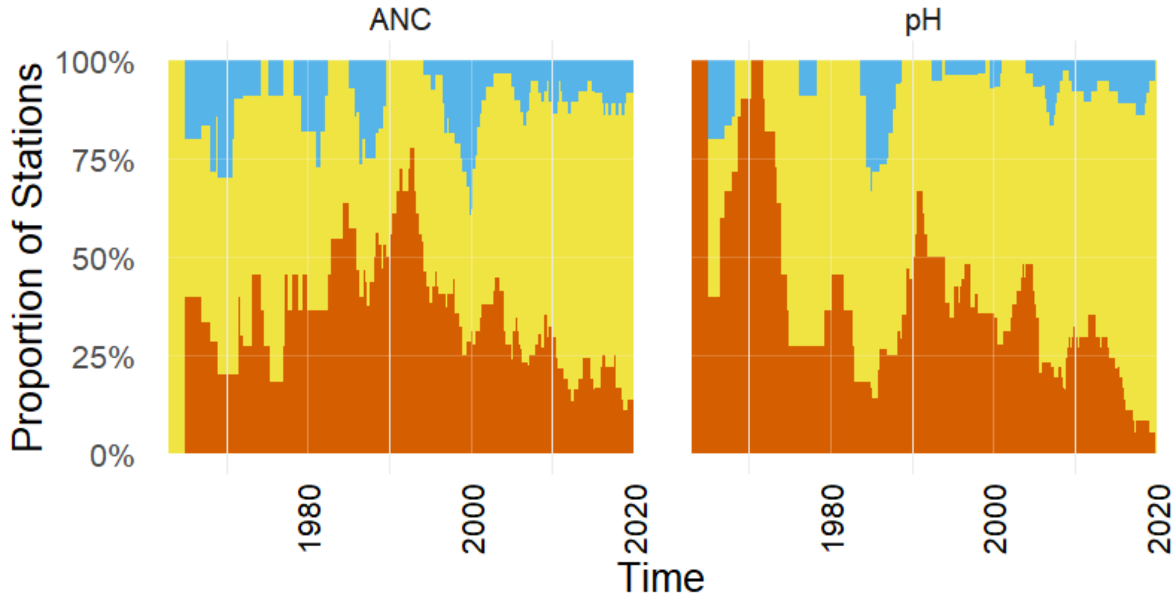

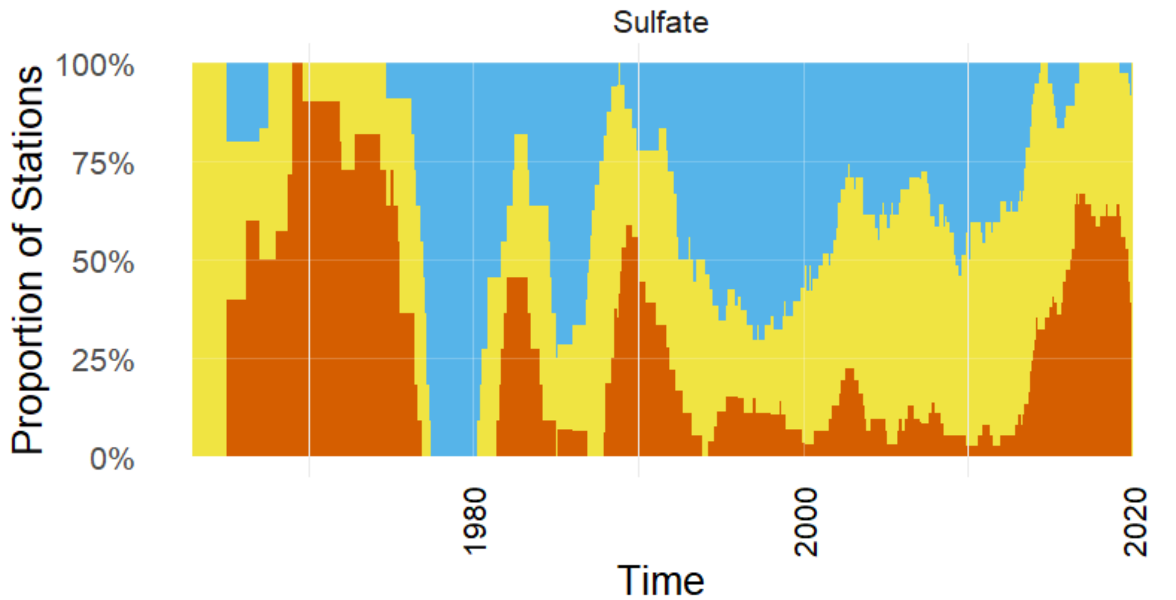

Figure 19: ÖM, percentage of sites per year experiencing significant trend changes (0.1 significance level) in ANC (10th percentile), PH (10th percentile) and sulfate concentration (90th percentile)