

Analyses of the connection between water body chemistry and wild rice.

Douglas M Hawkins,
Emeritus Professor and former Chair of Applied Statistics,
University of Minnesota, Twin Cities
dhawkins@douglashawkins.com

Executive Summary

My comments relate to Chapter 1, C, D, E and F of the *Final Technical Support Document Refinements to Minnesota's Sulfate Water Quality Standard to Protect Wild Rice*. The data set used in the study (the "Class B" data set) comprises 108 water bodies. Some analyses use a 96-body subset of the Class B data set – those water bodies with transparency greater than 30 cm, which is thought to make them more suitable for wild rice. Relevant variables include four analytes – porewater sulfide, porewater iron, porewater total organic carbon, and surface water sulfate – abbreviated in the discussion below to "sulfide, Fe, TOC and SO₄." Also measured was the presence or absence of wild rice, and (where wild rice was present) its stem density. Analyses were carried out in R Version 3.3.0. Weisberg (2014) and Fox and Weisberg (2011) are general references to the statistical methodology and its execution in R.

The analyses led to the conclusions:

1. The waterbody-specific sulfate standard proposed by MPCA does not differentiate waterbodies hosting wild rice from water bodies that do not.
2. More generally I have been unable to find any function of SO₄, TOC and Fe that can differentiate water bodies hosting wild rice from water bodies that do not.
3. Sulfide is a statistically significant but weak predictor of wild rice presence.
4. The MPCA assessment of the proposed sulfate rule's performance is questionable.
5. All four analytes vary substantially from time to time within the same water body.
6. SO₄, TOC and Fe are statistically significant but imprecise predictors of sulfide.
7. The proposed sulfide cutoff of 120 µg/L is not well supported and would lead to many false alarms.
8. A different approach using sulfide in a linear discriminant analysis incorporates explicit recognition of the implications of false positive and false negatives, and further motivates higher sulfide cutoffs.

In summary, the data presented give little reason to believe that changes in the sulfate standard will have any effect on the occurrence or health of wild rice, or indeed that a sulfate standard itself is required. A standard focused directly on sulfide would incur substantial numbers of false positives (water bodies with high sulfide but abundant wild rice) and false negatives (water bodies with low sulfide but no wild rice). More detailed study of these water bodies would be required to diagnose their specific properties and actions needed to enhance wild rice.

Analyses

1. *The waterbody-specific sulfate standard proposed by MPCA does not differentiate waterbodies hosting wild rice from water bodies that do not. This refers to Chapter 1 E Development of an equation to calculate a numeric sulfate standard for each wild rice water*

The Class B data set discussed in the MPCA document and analyzed here comprises 108 wild rice water bodies. As there is reason to believe that opaque water is inhospitable to wild rice, some of MPCA's analyses are restricted to 96 water bodies whose transparency exceeds 30 cm. This thinning can be justified by the observation that 11 of the 12 water bodies excluded did not have wild rice and only one did.

Directly cross-tabulating all 108 water bodies by the presence or absence of wild rice, and whether their SO4 is above or below the MPCA's Chapter 1E water-body-specific sulfate limit gives the table:

MPCA limit all water bodies			
	SO4 high	SO4 low	total
Rice absent	24	17	41
present	48	19	67
Total	72	36	108

A formal test of the association between SO4 and wild rice presence is given by Pearson's chi-squared test:

Pearson's Chi-squared test with Yates' continuity correction
 X-squared = 1.4203, df = 1, p-value = 0.2334

The P value of this test falls far short of statistical significance, confirming the visual impression that the proposed SO4 limit has no connection to the presence or absence of wild rice in the water body. Another indication of this is the total concordance – the proportion of water bodies correctly classified as wild-rice-hospitable or not by whether their SO4 is above or below the limit:

Concordance and CI 39.8% 31.1% 49.2%

In other words, 60% of the water bodies – a majority – would be misdiagnosed by the proposed standard.

Restricting the analysis to the 96 water bodies with suitable transparency gives the same conclusions:

MPCA limit 96 water bodies			
	SO4 high	SO4 low	total
Rice absent	18	12	30
present	47	19	66
Total	65	31	96

Pearson's Chi-squared test with Yates' continuity correction
 X-squared = 0.7285, df = 1, p-value = 0.3934

Concordance and CI 38.5% 29.4% 48.5%

FINDING 1: In both the broader and the narrower data sets, there is no association between the presence or absence of wild rice and whether the SO4 is above or below the waterbody-specific sulfate limit.

The performance of the proposed sulfate standard for identifying wild rice sites is akin to throwing a die and declaring the water body good if the die shows a 1 or 2, and bad if the die shows a 3, 4, 5 or 6.

2. *More generally I have been unable to find any function of SO₄, TOC and Fe that can differentiate water bodies hosting wild rice from water bodies that do not.*

Concentrating on the 96-water-body data set of sites where the water is transparent enough to be thought amenable to wild rice, presence or absence can be modeled directly from SO₄, TOC and Fe with a logistic regression:

```
glm(formula = Presence ~ logSO4 + logTOC + logFe, family = binomial)
```

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	4.3819	2.4913	1.759	0.0786 .
logSO4	-0.5050	0.3309	-1.526	0.1269
logTOC	-0.2771	0.4968	-0.558	0.5770
logFe	-0.7979	0.6716	-1.188	0.2348

Null deviance: 119.25 on 95 degrees of freedom
Residual deviance: 115.34 on 92 degrees of freedom
Explained 3.91 on 3 degrees of freedom

The overall model has an explained deviance of 3.91 with 3 degrees of freedom for a P value of 0.2713.

Thus neither the overall model, nor any of the terms in it, is statistically significant.

The same conclusion comes from a logistic regression using all 108 water bodies – neither the overall logistic regression, nor any of its terms, is statistically significant.

Another view of the data set is given by Hotelling's multivariate T squared test, which tests whether there is any difference in the triad log(SO₄, TOC, Fe) between the water bodies that do and that do not harbor wild rice. In the 96-water-body set, this test gives:

F-statistic: 1.283 on 3 and 92 DF, p-value: 0.285

confirming the lack of significant difference in these three concentrations between the water bodies that do and do not host wild rice.

The conclusion then is that these three predictors are not informative about the presence or absence of wild rice. Any model using them to predict presence or absence of wild rice can be no better than random guessing.

It is however conceivable that, even though these predictors cannot predict presence or absence of wild rice, they might nevertheless be able to differentiate water bodies with healthier wild rice. To explore this possibility, a linear regression of the stem density was fitted with the following results

Call:

```
lm(formula = Stems ~ logSO4 + logFe + logTOC, data = DDt)
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)
--	----------	------------	---------	----------

(Intercept)	33.675	47.627	0.707	0.481
logSO4	-7.481	6.400	-1.169	0.245
logFe	3.840	13.143	0.292	0.771
logTOC	-8.925	9.395	-0.950	0.345

Residual standard error: 42.82 on 92 degrees of freedom
Multiple R-squared: 0.01834, Adjusted R-squared: -0.01367
F-statistic: 0.573 on 3 and 92 DF, p-value: 0.6342

Neither the overall regression nor any of the terms in it is statistically significant, showing that SO4, Fe and TOC are not relevant in this context either.

FINDING 2: Whether for wild rice presence, or for the abundance of the wild rice, SO4, TOC and Fe do not show any predictive information in the field data.

- Sulfide is a statistically significant but weak predictor of wild rice presence. This refers to Chapter 1C. Identification of 120 µg/L as the protective sulfide concentration.*

External evidence cited in the document shows the potential for harm to wild rice from sufficiently high concentrations of sulfide, and sulfide is described in the MPCA document as a primary determinant of the presence of wild rice. The predictive power for sulfide can be quantified by a logistic regression of wild rice presence or absence on log sulfide within the full data set. This gives

```
glm(formula = Presence ~ log10(Sulfide), family = "binomial")
Coefficients:
              Estimate Std. Error z value Pr(>|z|)
(Intercept)   -0.7089     0.4567  -1.552  0.12059
log10(Sulfide) -1.3373     0.4522  -2.957  0.00311 **
```

Null deviance: 143.40 on 107 degrees of freedom
Residual deviance: 132.87 on 106 degrees of freedom

On the one hand, log sulfide is indeed a highly statistically significant predictor (P=0.00311).

On the other hand, however, sulfide explains only 10.53, or 7%, of the total deviance in wild rice presence, leaving the remaining 132.87, or 93% unexplained.

Performing the same calculation on the 96 water body data set gives the same substantive conclusions. The P value for sulfide in the regression is a 0.0114, still significant though not quite as strong as the full data set. However in this data set, sulfide explains an even-smaller 6% of the total deviances, leaving 94% unexplained.

This means that while porewater sulfide is a statistically significant part of the picture of wild rice presence or absence, it is only a modest part of it. Its contribution pales next to that of other characteristics and variables.

Exploring this further, receiver operating characteristic (ROC) curves are a standard methodology for exploring the ability of a predictor X to classify cases into a “good” and a “bad” class. The sensitivity associated with any cutoff C is the proportion of bad cases whose predictor X exceeds C and which are therefore correctly classified at the cutoff C. The false positivity FP associated with C is the proportion of

good cases whose X exceeds C and so are wrongly classified. The specificity is $1 - FP$: the proportion of good cases whose X does not exceed C and which are therefore correctly classified at the cutpoint C . Sensitivity, specificity and FP are often expressed as percentages.

The ROC is a plot of sensitivity versus FP, generated by varying C across the whole range of the data. A widely-used guidance (CLSI EP24 A2) discusses the use of ROC curves.

The ROC curve of a good classifier rises steeply from the origin before turning over and going to the point (1,1). The ROC of a worthless predictor would be a straight line from the origin to the point (1,1).

The area under the ROC curve, the AUC, is a summary measure of the ability of X to distinguish good from bad cases. The AUC is 1 for a classifier that separates good cases from bad perfectly, and is 0.5 for a worthless classifier. The AUC has a direct interpretation. In our context, if you take one random water body with wild rice and one without, the AUC is the probability that the one without wild rice has higher sulfide than the one with wild rice.

There is a formal statistical significance test, the Wilcoxon test, for whether the AUC is significantly better than 0.5, that is, whether X does better than blind guessing.

The ROC curve for sulfide and wild rice presence in the full data set is shown as Figure 1. Its AUC is 0.653. The Wilcoxon test gives $P = 0.0069$ showing that using the sulfide level perform significantly better than blind guessing. But the actual AUC of 0.653, though statistically significant, is much closer to the 0.5 you get by flipping a coin than it is to 1. While, at 65.3%, a water body without wild rice has a better than 50% chance of a higher sulfide levels than a water body with wild rice, its odds are not much better.

Like the proportion of deviance explained, the AUC paints a picture of sulfide as one fairly small part of the picture: statistically significant but far short of determinative.

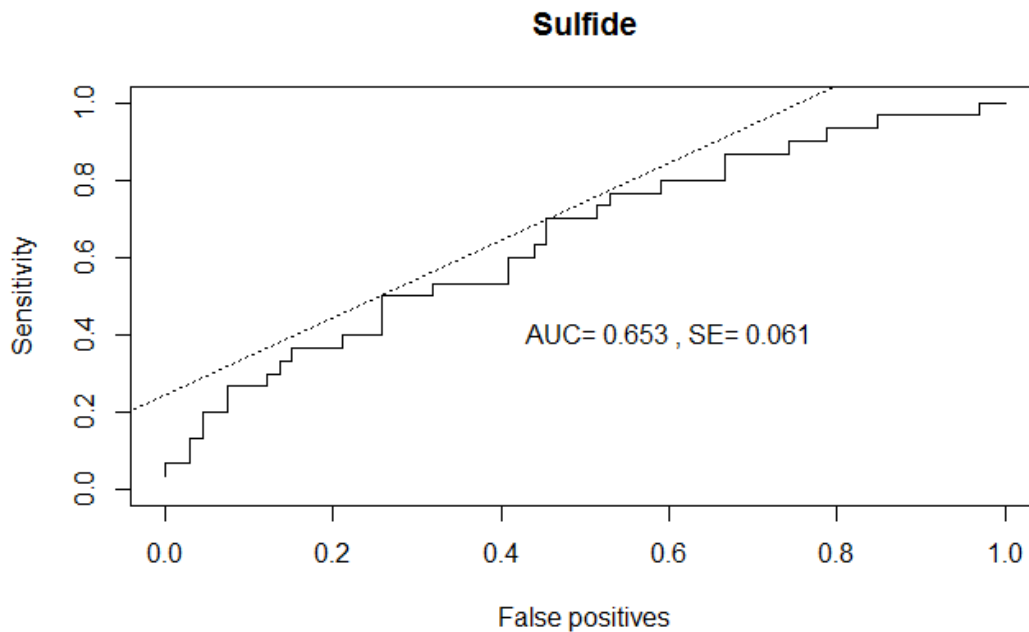


Figure 1. ROC curve for Class B waterbodies.

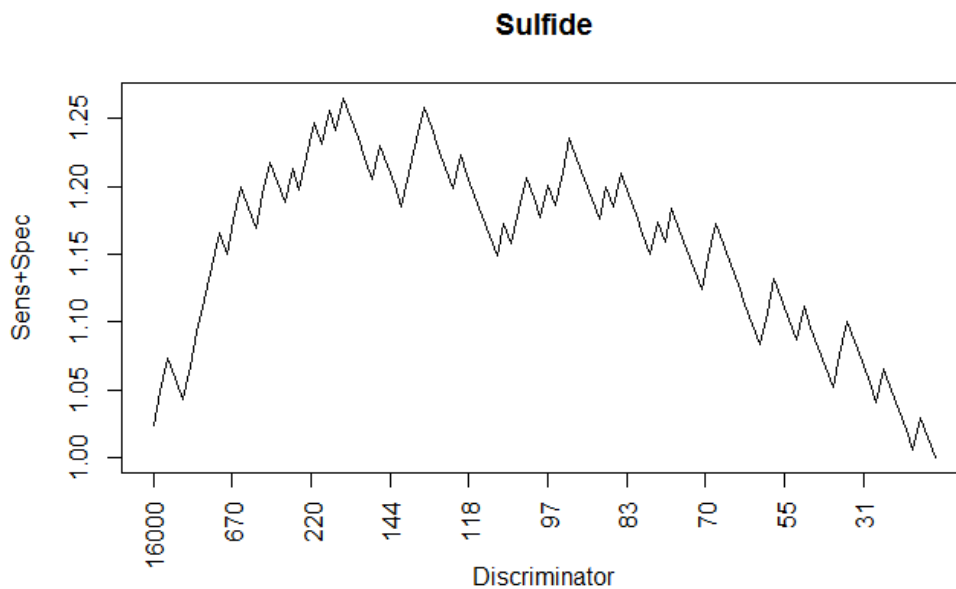


Figure 2. Youden index for different cutpoints.

The ROC curve also provides an objective way to determine a cutpoint. In an ideal situation, the ROC rises steeply to an “elbow” high up on the left of the graph, before leveling off and completing its path to the point (1,1). Such an elbow, when one exists, represents a natural cutpoint. The Youden index, defined as “sensitivity + specificity”, is an overall measure of the desirability of the associated cutpoint. A conventional way of selecting a cutpoint is to pick the value maximizing the Youden index, this being,

arguably, the point “highest on the left.” Figure 2 is an aid to this; it shows the Youden index as a function of the sulfide cutpoint.

The maximum Youden index is 1.265, given by a sulfide cutpoint of **181 µg/L**

But figure 2 also shows that the Youden index exceeds 1.2 for most cutoff values between 118 and 305, indicating that within this range, sensitivity and specificity are essentially trading off on a one for one basis and implying that a case could be made for any value within this range.

Repeating the ROC analysis on the subset of 96 water bodies with transparent water gave a slightly smaller AUC of 0.653 with less significance, $P=0.0172$, and leading to a Youden index of 1.245 at the cutpoint 0.093, somewhat below the values indicated for the full data set.

Another perspective on possible cutpoints comes from a changepoint analysis. The methodology of Hawkins (2001) was applied to the 96-body data set to find the cutpoint that optimally distinguishes the water bodies with wild rice from those without on the basis of their porewater sulfide. In this analysis, the optimal cutpoint was **274 µg/L**.

FINDING 3.1: Sulfide has a statistically significant separation between water bodies with and without wild rice, but is not particularly effective in differentiating between the two.

FINDING 3.2: The ROC curve does not identify a clear choice for a cutpoint on sulfide.

4. *The MPCA assessment of performance is questionable. This refers to Chapter 1 F Comparison of an equation-based standard to fixed standards: Error rates and concerns.*

Contrary to my conclusion that SO₄ has no perceptible connection to wild rice, the MPCA document reports quite favorable performance for the proposed water-body-specific sulfate standard. However this performance is against a surrogate endpoint – sulfide being below 120 µg/L – and not the actual endpoint of interest – the presence or absence of wild rice.

Surrogate endpoints are acceptable in some circumstances, notably

- when the surrogate is more easily available, or available sooner, than the primary endpoint; and
- the surrogate endpoint is closely related to the primary endpoint.

Neither of these circumstances motivating surrogate endpoints appears relevant in this problem. It is implausible that measuring the chemistry of a water body is faster, cheaper or more convenient than a visual assessment of its vegetation. On the second requirement, sulfide is a quite imperfect predictor of wild rice presence and health, a deficiency that the MPCA report itself notes.

FINDING 4: Thus the use of this surrogate endpoint seems questionable, as do the resulting conclusions.

5. *All four analytes vary substantially from time to time within the same water body. This relates to Chapter 1D Assumption that SO₄, TOC, iron and sulfide are in a steady state at field sites.*

Some water bodies were sampled more than once, but where a body had more than one measurement, MPCA's primary analyses used only one. Their analysis invoked the "steady state" concept that the water body chemistry does not change much over time, and it is appropriate to check on this.

The data set "MPCA_Field_Survey_Data_with_calculated_protective_sulfate_concentration" contained 267 records covering 165 waterbodies. Of these, 53 bodies provided more than one record of some or all of the key variables sulfide, SO₄, Fe and TOC. The repeat measurements at the same water bodies were taken at different dates. The standard deviation of the sampling date within a waterbody was 210 days, or some 7 months.

The multiple readings of these measures within the same water body were analyzed by a random effects analysis of variance to separate out the variability within and between water bodies. Calculations used the "lmer" command from the R package "lme4". To correct for any major seasonal effects, the model included sine and cosine terms with period one year and six months.

All four of the concentrations were transformed to common logs.

This results of this analysis follow.

Sulfide




Figure 3. Sulfide repeat sampling

Figure 3 is a comparative box and whisker plot of the log-transformed sulfide values measured at different times broken down by waterbody. Visually, the plot shows variability within a water body comparable in scale to that between water bodies. This visual impression is quantified by the analysis of variance. In this, the term “keeplake” corresponds to variation from one water body to another; “Residual” refers to variation over time within a water body.

```
REML criterion at convergence: 129.3
Random effects:
  Groups   Name                Variance Std.Dev.
keeplake (Intercept)  0.10763  0.3281
Residual                                0.08246  0.2872
```

The log-transformed sulfide level varies from one water body to another with a standard deviation of 0.33. However within the same water body, it varies from one time to another with a similar standard deviation of 0.29. A standard deviation of 0.29 on the log10 scale corresponds to a coefficient of variation of about 70% on the original scale. At this level, two sulfide readings on the same water body have a 1 in 3 chance of differing by more than 100%, a proportion supported by the actual successive sulfide readings.

In other words, the sulfide level of a water body is an elusive, moving target.

Fe




Figure 4. Fe repeat sampling

Figure 4 shows the same picture of the log-transformed iron level. This too is visually highly variable within a water body. The analysis of variance gives

```
REML criterion at convergence: 257.8558
Random effects:
Groups   Name             Std.Dev.
keelake (Intercept) 0.3376
Residual                0.4776
```

The standard deviation within a water body is considerably higher than that between water bodies.

As the four concentrations have been log transformed, they are dimensionless, and it is legitimate to compare the standard deviations of the different analytes. Thus one can note that the variability in Fe from one water body to another is comparable to that of sulfide, but within a water body, Fe is considerably more variable than sulfide.

TOC

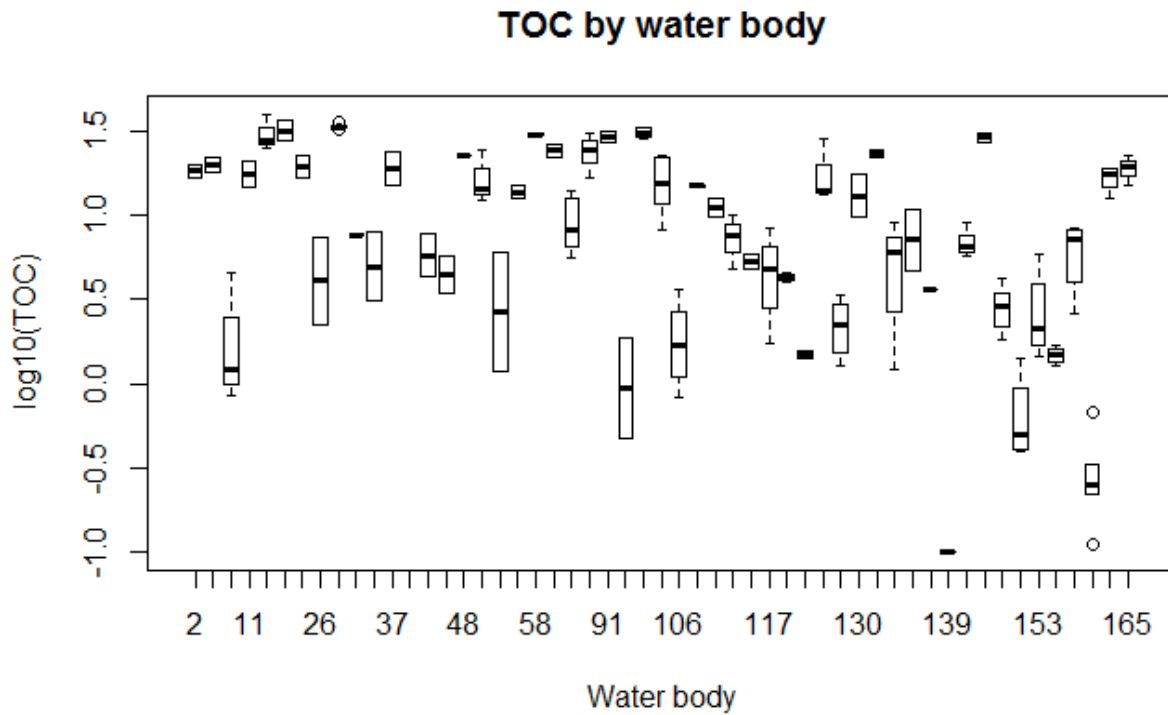


Figure 5. TOC repeat sampling

Figure 5 gives the box and whisker plot for log-transformed total organic carbon. The analysis of variance gives

```
REML criterion at convergence: 107.494
Random effects:
Groups   Name             Std.Dev.
keelake  (Intercept)      0.5401
Residual                               0.2013
```

TOC within a water body is much more stable than Fe or sulfide, but it varies more from one water body to another.

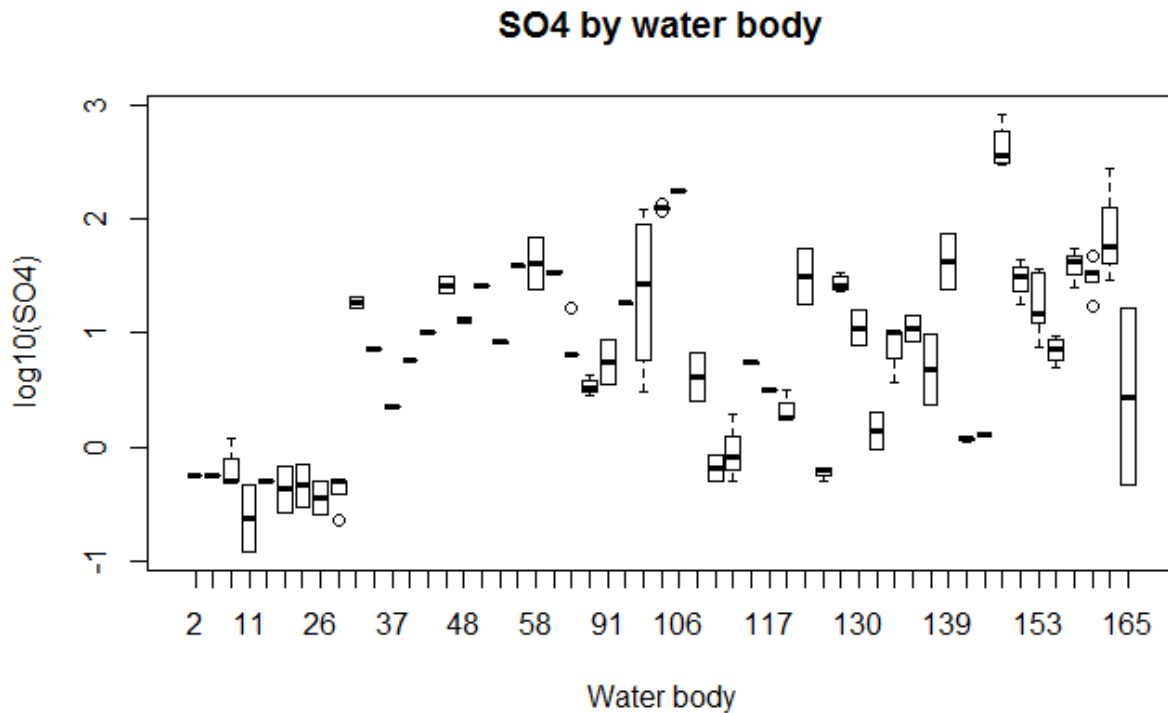


Figure 6. SO4 repeat sampling

Finally, Figure 6 shows log-transformed SO4. The analysis of variance gives

```
REML criterion at convergence: 185.077
Random effects:
Groups   Name             Std.Dev.
keplake  (Intercept)      0.7556
Residual                               0.2517
```

SO4 varies much more between waterbodies than do sulfide, Fe and TOC, and its variability within a water body is comparable with that of sulfide and TOC.

FINDING 5: In summary, all four analytes show substantial variability over time within the same water body. A snapshot of the chemistry at a given time may produce substantially different values than one made at another time. The steady state assumption is therefore not validated particularly well.

6. *SO₄, TOC and Fe are statistically significant but imprecise predictors of sulfide. This refers to Chapter 1 C. Relationship between surface water sulfate and porewater sulfide.*

The regression model connecting sulfide to SO₄, TOC and Fe in the full data set is:

```
lm(formula = logsulfide ~ logSO4 + logTOC + logFe
```

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	0.97145	0.38938	2.495	0.0142	*
logSO4	0.40241	0.05368	7.497	2.27e-11	***
logTOC	0.45564	0.07832	5.818	6.65e-08	***
logFe	-0.69130	0.10748	-6.432	3.91e-09	***

Residual standard error: 0.3751 on 104 degrees of freedom
 Multiple R-squared: 0.491, Adjusted R-squared: 0.4763
 F-statistic: 33.44 on 3 and 104 DF, p-value: 3.253e-15

As expected, all three terms in the model are highly statistically significant, as is the overall regression. However, while significant, the regression explains less than half the variability ($R^2 = 0.491$), implying that other factors and random variability are responsible for most of the sulfide variability.

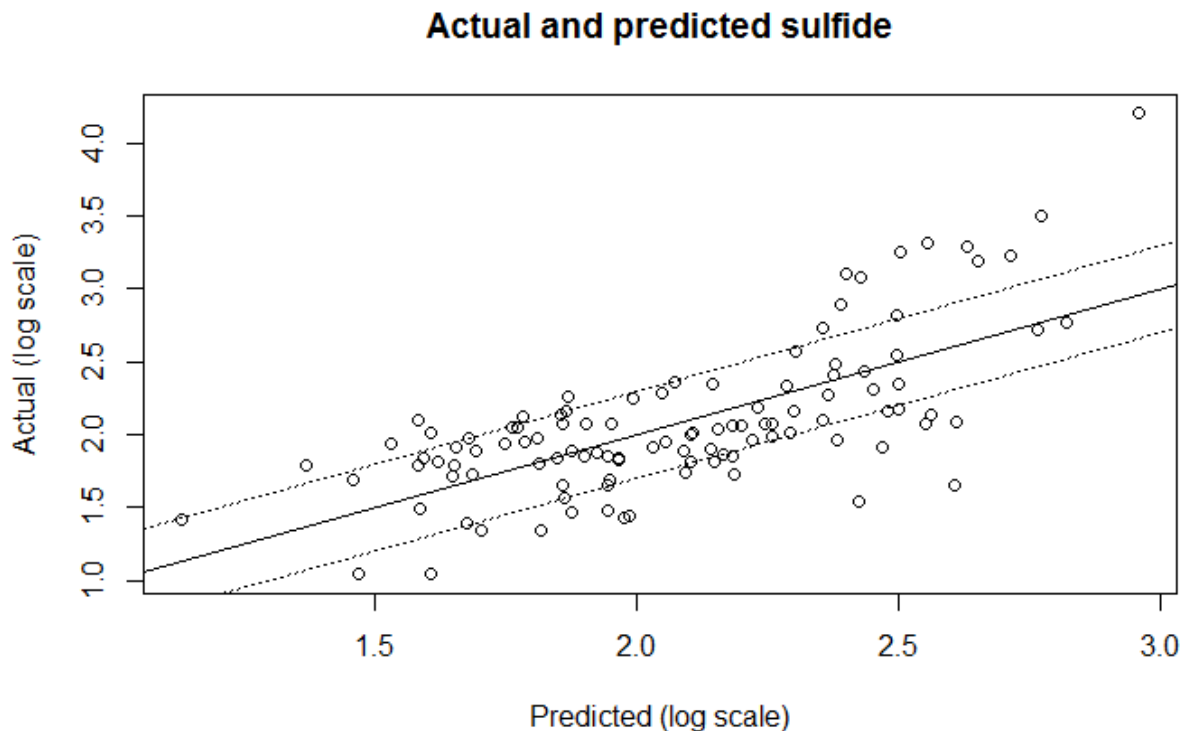


Figure 7. Predicting sulfide from SO₄, FE and TOC

Figure 7 shows this graphically. It is a plot of the actual sulfide values against the value predicted by the regression on SO₄, Fe and TOC. The plot is on a double log scale. The solid line is the line of identity.

The two dotted lines mark where the actual sulfide differs from the model prediction by a factor of 2. As the graph makes clear, the actual sulfide level quite commonly differs from its prediction by this factor of 2 or more. This is shown by the points lying outside the dotted lines.

FINDING 6.1 *In other words, the highly significant regression nevertheless makes sulfide predictions that are commonly wide of the mark.*

FINDING 6.2: *Putting various pieces of the puzzle together,*

- *There is a statistically significant but imprecise relationship of SO₄, Fe and TOC to sulfide, and*
- *There is a statistically significant but modest relationship of sulfide to wild rice.*
- *This chain of relationships falls apart when the intermediate of sulfide is removed and one attempts to predict wild rice directly from SO₄, Fe and TOC. Then the unexplained variability in the two relationships overwhelms the modest associations, leading to the lack of significant association between SO₄, Fe and TOC and the presence or absence of wild rice.*

7. *The proposed sulfide cutoff of 120 µg/L is not well supported and would lead to many false alarms.*

Continuing with the possibility of using the sulfide level as a classifier and taking a closer look at some proposed cutoffs, the usual measures of performance at the MPCA's 120 µg/L applied to the 96 water bodies with acceptable transparency are

	Cutoff 120		
	Sulfide high	Sulfide low	total
Rice absent	15	15	30
present	18	48	66
Total	33	63	96
Sensitivity and CI	50.0%	33.2%	66.8%
Specificity and CI	72.7%	61.0%	82.0%
Sens + Spec and CI	122.7%	101.9%	143.6%
Concordance and CI	65.6%	55.7%	74.4%
PPV and CI	45.5%	29.8%	62.0%
NPV and CI	76.2%	64.4%	85.0%

The sensitivity, specificity and Youden index (sensitivity + specificity) have been mentioned. Looking beyond them to the outcomes of testing, the positive predictive value, PPV, is the probability that a high sulfide truly corresponds to lack of wild rice. This highly relevant as it tells you what fraction of followup after a signal of high sulfide will be productive in identifying genuine problems.

At the 120 µg/L cutoff, the PPV is less than 50%. The majority of high sulfides will therefore be false alarms and so most of the effort involved in following up high sulfide values will be wasted.

The negative predictive value, NPV, is the mirror image of this – the probability that a water body with sulfide below the cutoff does indeed host wild rice. A high NPV would imply that the water bodies that are categorized as good on the basis of low sulfide most likely are good and do not need much attention.

The NPV of 76.2% is fair to good. It does however mean that a quarter of the water bodies with this low sulfide level nevertheless do not have wild rice, leading one to wonder whether some simple intervention might bring wild rice to these water bodies,

The concordance, $65/96 = 66\%$, is the proportion of water bodies identified correctly.

The corresponding figures for the 181 $\mu\text{g/L}$ that optimizes the Youden index, and the 274 $\mu\text{g/L}$ that optimizes the changepoint test are:

Cutoff 181			
	Sulfide high	Sulfide low	total
Rice absent	11	19	30
present	10	56	66
Total	21	75	96
Sensitivity and CI	36.7%	21.9%	54.5%
Specificity and CI	84.8%	74.3%	91.6%
Sens + Spec and CI	121.5%	102.2%	140.8%
Concordance and CI	69.8%	60.0%	78.1%
PPV and CI	52.4%	32.4%	71.7%
NPV and CI	74.7%	63.8%	83.1%

Cutoff 274			
	Sulfide high	Sulfide low	total
Rice absent	8	22	30
present	6	60	66
Total	14	82	96
Sensitivity and CI	26.7%	14.2%	44.4%
Specificity and CI	90.9%	81.6%	95.8%
Sens + Spec and CI	117.6%	100.3%	134.9%
Concordance and CI	70.8%	61.1%	79.0%
PPV and CI	57.1%	32.6%	78.6%
NPV and CI	73.2%	62.7%	81.6%

The higher cutoffs give progressively better concordance, going from 66% to 71% as the cutoff goes from the MPCA's suggested 120 $\mu\text{g/L}$ to the Youden optimum of 181 $\mu\text{g/L}$ and the changepoint optimum of 274 $\mu\text{g/L}$. The PPV increases substantially, going from 45.5% to 57.1% indicating that effort spent in diagnosing high sulfide values is spent more productively.

The NPV decreases slightly, from 76% to 73%. This means that the clean bill of health coming from a sulfide below the cutoff becomes less clean as the cutoff increases. However the small change – from 76% to 73% -- shows that the reduction is not substantial.

FINDING 7: In summary, going from a sulfide cutpoint of 120 to 274 $\mu\text{g/L}$ produces many fewer alarms, and those alarms that are produced are much more likely to indicate real problems with the wild rice.

8. Discriminant Analysis Approach using Sulfide

The problem faced in monitoring water body chemistry is a decision – on the basis of the current chemistry, deciding whether or not to flag the water body as suspicious of being inhospitable to wild rice and requiring closer investigation. The conventional statistical model for this problem is a linear

discriminant analysis (LDA, Anderson 2003, Chapter 6). Unlike the other calculations made so far, LDA pays explicit attention to the tradeoff between the consequences of dealing with false positives and with false negatives.

It was applied to the 96 water bodies transparent enough to be thought hospitable for wild rice.

Using X , the common log of the sulfide as the discriminator, these water bodies fall in two populations:

1. Those in which wild rice is absent, of which there are 30,
2. Those in which wild rice is present, of which there are 66.

The sulfide data from the 96 water bodies of interest give summary statistics

Mean of X in population 1 = 2.26

Mean of X in population 2 = 1.96

Pooled variance = 0.22

Following Anderson section 6.5.1, the optimal classification rule is to classify a water body as suspect if its sulfide value satisfies

$$\frac{X - 0.5(2.26 + 1.96)}{0.22} (2.26 - 1.96) > \log_e k$$

The constant k is defined as

$$k = \frac{q_2}{q_1} R, \quad \text{where } R = \frac{L(1|2)}{L(2|1)}$$

The constants q_1 and q_2 are the proportions of water bodies that do not, and do harbor wild rice. In the 96-body data set, it is reasonable to estimate the ratio q_2/q_1 by 66/30, the proportions in the data set.

The constants L reflect the “loss” incurred by the two types of potential misclassification. $L(1|2)$ is the loss when you declare a water body suspect when in fact it can harbor wild rice, and $L(2|1)$ is the loss when you declare a water body acceptable when in fact it can not harbor wild rice. Note that only the ratio of these numbers, and not their actual values, is relevant.

Solving the optimal classifying equation for X using these values for q_1 and q_2 classifies the water as suspect if

$$X > 2.11 + \frac{0.22}{0.30} (0.79 + \log_e R) = 2.69 + 0.73 \log_e R$$

Values for the ratio R could be found by considering the follow-up steps needed to determine that the classification was wrong and evaluating how onerous they are, but a sensitivity calculation is illustrative.

Consider the values $R = 0.5, 1$ and 2 , ranging from the two types of error being equally severe to one being twice as bad as the other. These values of R lead to the sulfide cutoffs

R	Cutoff (µg/L)
0.5	153
1	490
2	1570

FINDING 8. Even the lowest of these numbers is above the 120 µg/L proposed in the MPCA document. These numbers provide further evidence that, if sulfide is used as an indicator of suitability for wild rice, a higher sulfide cutoff should provide a better use of resources for followup.

References

Anderson, T. W., (2003), *An Introduction to Multivariate Analysis 3rd Edition*, Wiley New York.

CLSI. *Assessment of the Diagnostic Accuracy of Laboratory Tests Using Receiver Operating Characteristic Curves. Approved Guideline – Second Edition*. CLSI document EP24-A2, Wayne, PA. Clinical and Laboratory Standards Institute, 2011.

Fox, J., and Weisberg, S. (2011), *An R Companion to Applied Regression 2nd Edition*, Sage, Los Angeles.

Hawkins, D. M., (2001) 'Fitting Multiple Change-Points to Data', *Computational Statistics and Data Analysis*, 37, 323-341.

Weisberg, S. (2014) *Applied Linear Regression, 4th Edition*, Wiley New York.