

# Multiple Changepoint Analysis of Noisy Nonlinear Data with an Application to Modeling Crack Growth in Additively Manufactured Titanium

Lucas Koepke<sup>\*†</sup>    Jolene Splett<sup>‡</sup>    Timothy Quinn<sup>‡</sup>    Nikolas Hrabe<sup>‡</sup>  
 Jake Benzing<sup>‡</sup>    Michael Frey<sup>‡</sup>

## Abstract

Noisy measurement data pose a challenge for changepoint analysis, especially in the presence of multiple changepoints and when the model is nonlinear. We explore various approaches to estimating changepoints and their standard errors under these conditions. We consider whether adding a monotonicity constraint improves the changepoint estimates and reduces their standard errors. We finish with a novel application to material science using crack growth data from additively manufactured titanium. As cyclic loading is applied to a test specimen, crack growth can be partitioned into three regimes: slow-growth, mid-growth, and high-growth. We improve estimates of the transition points between these regimes versus those made by experts in the field by adding confidence bounds to the changepoint locations, allowing for designed experiments to study treatment effects on changepoint location.

**Key Words:** changepoint, isotonic regression, nonlinear least squares, pool-adjacent-violators algorithm

## 1. Introduction

As a material is subjected to cyclic fatigue loading, cracks can form and grow over time even when the maximum applied force is well below the yield strength of the material. To measure the resistance of a material to fatigue crack propagation, a notched sample is cyclically loaded for many thousands of cycles. For each force cycle, the crack size  $a$  is estimated and the cyclic force range,  $\Delta P = P_{max} - P_{min}$  is recorded. For analysis, the stress intensity factor,  $\Delta K$ , is calculated from  $\Delta P$ ,  $a$ , and the geometry of the specimen, since during fatigue testing the force range is constant, while  $\Delta K$  increases as the crack propagates. The stress intensity can be calculated for any part where the forces and geometry are known.  $\Delta K$  is plotted against the change in crack length per cycle or fatigue crack rate,  $\frac{da}{dn}$  (see Figure 1).

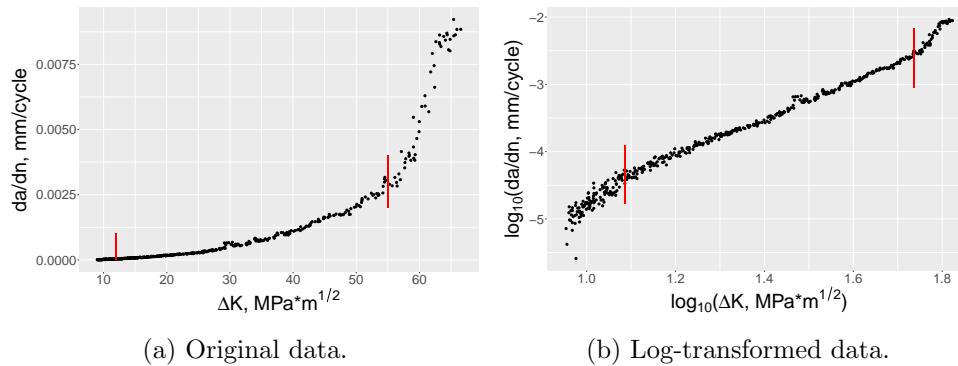
Figure 1a shows data on the linear scale. For lower values of  $\Delta K$ , the crack shows minimal growth. This transitions to an elbow region as  $\Delta K$  increases, where the crack growth rate can be predicted with a power law (Hertzberg, 1996, p. 614). This in turn leads to a short period of rapid crack growth. The transitions between these three regimes (slow-growth, mid-growth, and high-growth) can be used to characterize the fatigue properties of the material. Because the transitions are difficult to identify on the linear scale, both  $\frac{da}{dn}$  and  $\Delta K$  are typically transformed by taking the log (base 10) (Figure 1b). The transformation changes the elbow region into a linear region, with a lower changepoint occurring at the transition into the linear region and an upper changepoint at the transition out of the linear

---

<sup>\*</sup>Associate, National Institute of Standards and Technology, Boulder, CO 80305

<sup>†</sup>University of Colorado, Department of Physics, Boulder, CO 80309

<sup>‡</sup>National Institute of Standards and Technology, Boulder, CO 80305



**Figure 1:** Example of fatigue crack growth data from additively manufactured titanium on the linear scale (left) and after taking the log (base 10) of both the stress intensity factor  $\Delta K$  and the crack growth rate  $\frac{da}{dn}$  (right). Possible changepoint locations are shown by the vertical red lines.

region. These changepoints are typically determined through inspection by experts in the field. Our challenge is to estimate these two changepoints objectively.

We proceed as follows. Since we don't know the true changepoints for the experimental data, we start with a simulation study both to compare methods using data with a known changepoint and to explore how aspects of the data might bias the estimates. Additionally, we investigate whether it is possible to reduce noise in the data knowing that crack growth should always increase as  $\Delta K$  increases. We finish by estimating the changepoints on experimental data and comparing results with changepoints estimated using inspection by expert engineers. All computation was performed using R (R Core Team, 2018).

## 2. Methods

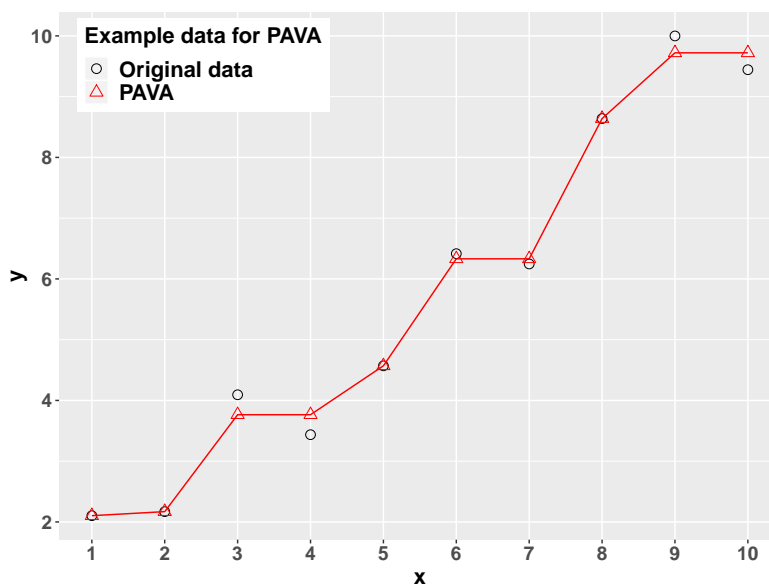
We assume that the model should be linear between the changepoints, so we concentrate on methods that can be made to accommodate this constraint. Specifically, we focus on fitting models using nonlinear least squares (NLS). We can parameterize the function that we fit to approximate different shapes below, between, and above the two changepoints. In each case, the function segments are constrained to meet at the changepoints, and the changepoint locations are estimated as parameters. The approach also has the advantage of providing standard errors on the parameter estimates and model fit diagnostics. For example, the model for estimating one changepoint with one linear and one quadratic segment is

$$\hat{y} = I_{x \leq \text{changepoint}} (a_1 + b_1 * x) + I_{x > \text{changepoint}} (b_2 * x + c_2 * x^2) \quad (1)$$

If  $c_2$  is set to zero this model reduces to one with two linear segments meeting at the changepoint.

### 2.1 Noise reduction

We know that  $\frac{da}{dn}$  should be monotonically increasing with  $\Delta K$  (cracks do not decrease in size in engineering materials under cyclic force), so it is natural to consider the benefit of isotonic regression (Barlow and Brunk, 1972; Dykstra, 1981; Robertson et al., 1988). Isotonic regression is unique in that it is guaranteed to reduce



**Figure 2:** Example of PAVA using simulated data. Black circles are the original data points, the red triangles show the result after PAVA. If the original data is already increasing, PAVA does nothing, but the pooling is evident whenever a  $y$  value decreases.

mean-squared error when estimating a monotonic mean function. We consider the isotonic regression calculated using the pool-adjacent-violators algorithm (PAVA) (Ayer et al., 1955). PAVA is implemented in R in the `isotone` package (Mair et al., 2009).

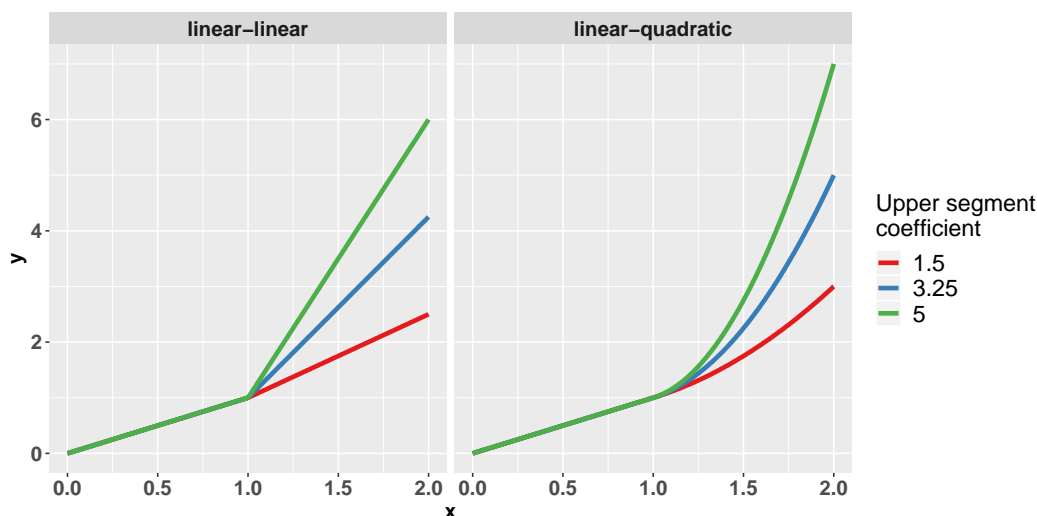
PAVA works as follows. Start with the data  $(x_i, y_i)$  for observations  $i \in 1, \dots, n$  ordered such that  $x_i \leq x_{i+1}$ . At this point the  $y_i$ 's have not been touched. We assume equal weights. Start at  $i = 1$  and check whether  $y_2$  is less than  $y_1$ . If this is the case, the monotonicity constraint is violated, so pool  $y_1$  and  $y_2$  by replacing both with their average  $\frac{y_1 + y_2}{2}$ . Then proceed as follows for  $i \in 2, \dots, n - 1$ :

1. Check for a monotonicity violation ( $y_{i+1}$  is less than  $y_i$ ). If that is the case, pool  $y_{i+1}$  and  $y_i$ .
2. If the new value of  $y_i$  is now less than  $y_{i-1}$ , pool all three ( $y_{i-1}$ ,  $y_i$ , and  $y_{i+1}$ ) by replacing the three values with their average. Then for  $j \in i - 1, \dots, 1$  repeat until  $j = 1$  or  $y_j < y_{j+1}$ .
3. Proceed to the next  $i$ .

After the PAVA computation has completed, the  $y$ 's are monotonically non-decreasing (Figure 2). Even though PAVA appears to naturally apply to fatigue crack growth data due to the monotonicity, it is not clear whether PAVA will bias the changepoint estimates, or to what extent the associated standard errors will be affected.

### 3. Simulation study

The simulation study is set up to explore different situations that might arise in the experimental data in a controlled way. We simulate data using (1) for the one

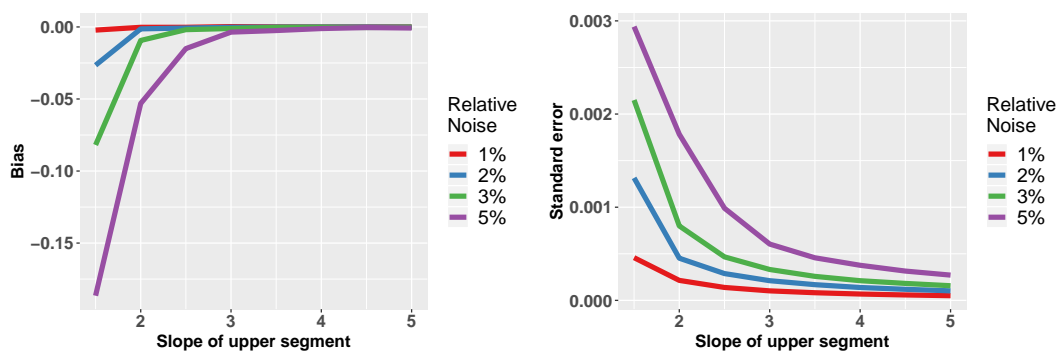


**Figure 3:** Data generation models shown without noise. Both have a linear first section with  $b_2 = 1$  up to the changepoint at  $x = 1$ . Multiple parameter values for the upper segments are shown, corresponding to the slope on the upper segment ( $b_2$ ) in the linear-linear model and the coefficient on the quadratic term ( $c_2$ ) for the linear-quadratic model.

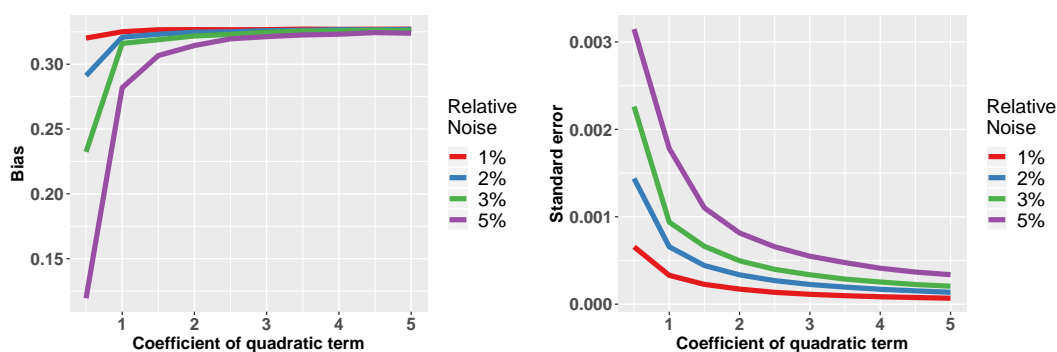
changepoint case since including a second changepoint is a straightforward extension of the NLS function. The first segment is a line with  $b_1=1$  for  $x$  between 0 and 1. The changepoint is at  $x = 1$ . For  $x$  between 1 and 2, we consider both a linear and a quadratic segment, the quadratic is chosen because it represents the simplest nonlinear case.

Two parameters are varied when simulating data. The first is the amount of additive noise. In the experimental data, the error relative to the range of the  $y$  measurements is about 1 %, so we consider a range of additive noise values including: 1 %, 2 %, 3 %, and 5 %. The other parameter we vary is the shape of the second segment after the changepoint. In the linear-linear model where  $c_2=0$ , this parameter is the slope of the second segment ( $b_2$ ), and will range between 1.5 and 5. For the linear-quadratic model, we fix the slope of the quadratic segment to be 1 at the changepoint ( $b_2=1$ ) so the function is continuous and smooth, but we vary the coefficient on the quadratic term ( $c_2$ ) between 0.5 and 5. Changing  $c_2$  is equivalent to changing the second derivative (differing by a factor of 2). Examples of these data generation models with different parameter values are shown in Figure 3.

For each model, parameter value, and relative noise level, we simulate 10,000 data sets, each with 200 points below the changepoint and 200 points above, and fit linear-linear and linear-quadratic models using NLS. The output of each iteration is an estimated changepoint location, and we use these to compute the bias and standard error of the changepoints. The second part of the simulation will treat noise reduction using PAVA. Although PAVA should reduce the noise in the data with no penalty in overall model fit, it is not clear whether PAVA will introduce bias in the specific parameter estimate for the changepoint. Thus for each simulated data set, we also fit the model to PAVA smoothed data, and compare the bias and standard error of the changepoint parameter estimates for unsmoothed and PAVA-smoothed data.



**Figure 4:** Bias (left) and standard error (right) for the estimated changepoint obtained by fitting the linear-linear model to linear-linear data versus the value of  $b_2$  used to simulate the data.



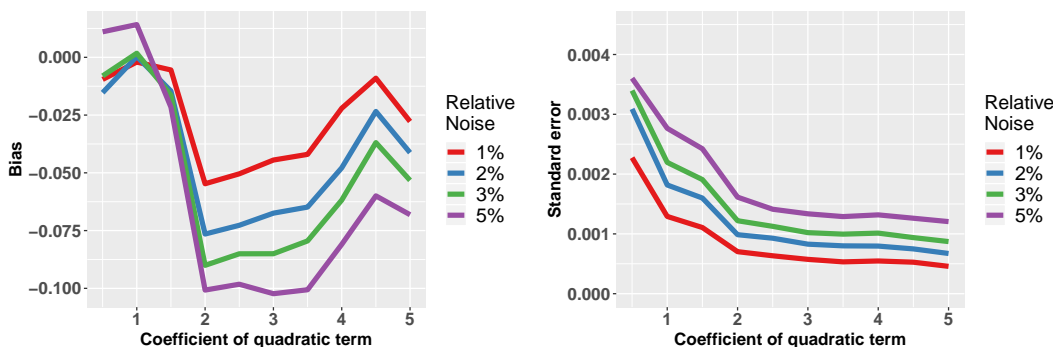
**Figure 5:** Bias (left) and standard error (right) for the estimated changepoint obtained by fitting the linear-linear model to linear-quadratic data versus the value of  $c_2$  used to simulate the data.

### 3.1 Bias and standard error

Figure 4 shows the bias (left) and standard error (right) of the estimated changepoint for the linear-linear model fit to the linear-linear data. On the  $x$ -axis is the slope of the upper segment, ranging between 1.5 and 5. In the noisiest scenario (5 % relative noise) with the least slope difference ( $b_2 = 1.5$ ), the bias is -0.18. When  $b_2 = 2$ , the bias drops to just -0.05 and keeps improving as the slope increases. The standard error follows a similar pattern.

Results for bias and standard error of the estimated changepoint for the linear-linear model fit to linear-quadratic data are shown in Figure 5. The shape of the curves for bias and standard error are similar to the previous case (fitting a linear-linear model to linear-linear data) but now instead of the bias trending towards zero it settles at roughly 0.32. The standard error decreases for all relative noise levels as the coefficient on the quadratic term increases.

Results for bias and standard error of the estimated changepoint for the linear-quadratic model fit to linear-quadratic data are shown in Figure 6. When  $c_2$  is between 0.5 and 1.5, the biases are near zero but the standard errors at the four relative noise levels are highest at the 0.5 value, decreasing as  $c_2$  increases. The bias is negative for coefficients 2 through 5. A negative bias here indicates that the transition to the quadratic segment, the changepoint, is estimated to be too low, so the quadratic segment consistently takes over part of the linear segment.



**Figure 6:** Bias (left) and standard error (right) for the estimated changepoint obtained by fitting the linear-quadratic model to linear-quadratic data versus the value of  $c_2$  used to simulate the data.

The results of the simulation show that fitting a linear-linear model to linear-linear data performs well. The recommendation is not as clear for the linear-quadratic data, since fitting a linear-quadratic model does not perform consistently. The parameter combinations with biases near zero have the largest standard errors, and there are erratic jumps in the bias for some coefficient values. Fitting a linear-linear model to the linear-quadratic data shows higher bias in the changepoint estimate, but the results are consistent over different coefficient values, and the standard errors are consistently lower compared with the linear-quadratic fit.

### 3.2 PAVA results

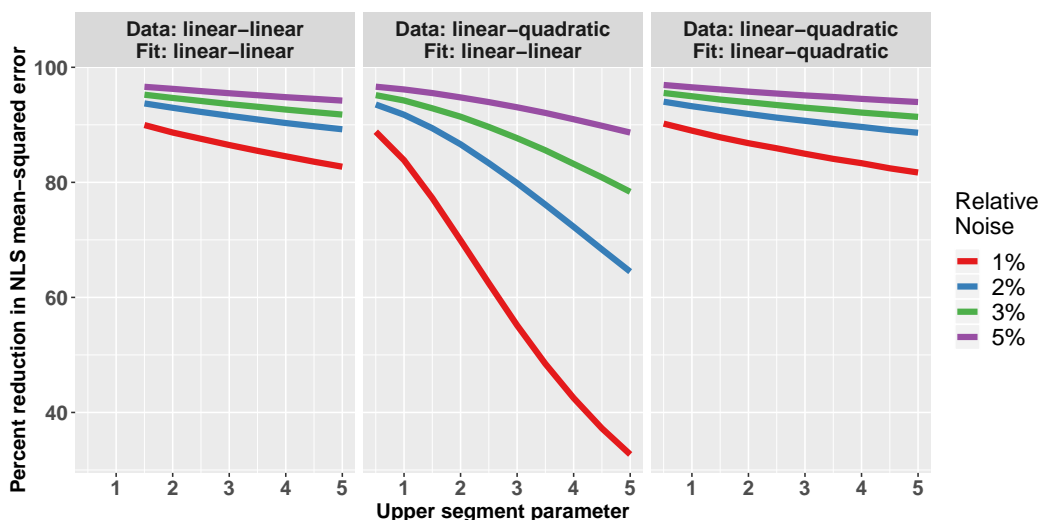
In the simulation study, we explore two ways that PAVA can influence the analysis. The first is to show the improvement in model fit, and the second is to explore the effect on the changepoint estimates themselves.

In terms of model fit, the reduction in mean-squared error from unsmoothed to PAVA-smoothed data is shown in Figure 7. The percent change was calculated going from unsmoothed to PAVA-smoothed data, relative to the unsmoothed data. For both cases where the model fit matches the data model, PAVA reduced the mean-squared error of the NLS fit by over 80 %. This reduction is larger at high relative noise levels, and decreases as the upper segment becomes steeper. The third case, fitting a linear-linear model to linear-quadratic data, shows similar improvement (over 80 %) when  $b_2$  is 0.5 or 1, but offers less of an improvement as  $b_2$  increases. When  $b_2 = 5$ , PAVA reduced the mean-squared error by less than 40 % at the lowest relative noise level.

The effect of PAVA on the bias of the changepoint estimates is shown in Figure 8. This plot shows the difference in absolute values,  $|\text{bias}_{\text{unsmoothed}}| - |\text{bias}_{\text{PAVA}}|$ . A positive value means that PAVA reduces the bias, while a negative value means the opposite. We chose to use this difference instead of calculating a percentage change since many values are close to zero, leading to unstable relative changes.

For a linear-linear model fit to linear-linear data, PAVA reduces the magnitude of the bias when  $b_2 = 1.5$  for the three highest relative noise levels (2 %, 3 %, and 5 %). When  $b_2 = 2$ , PAVA reduces the magnitude of the bias only for the 5 % relative noise case. PAVA increases the bias for the other combinations of slope and relative noise, but only by 0.01 or less.

PAVA does not improve the bias when fitting a linear-linear model to linear-



**Figure 7:** Percent change in mean-squared error of the NLS fit going from unsmoothed to PAVA smoothed data. (left) linear-linear model fit to linear-linear data, (center) linear-linear model fit to linear-quadratic data, and (right) linear-quadratic model fit to linear-quadratic data.

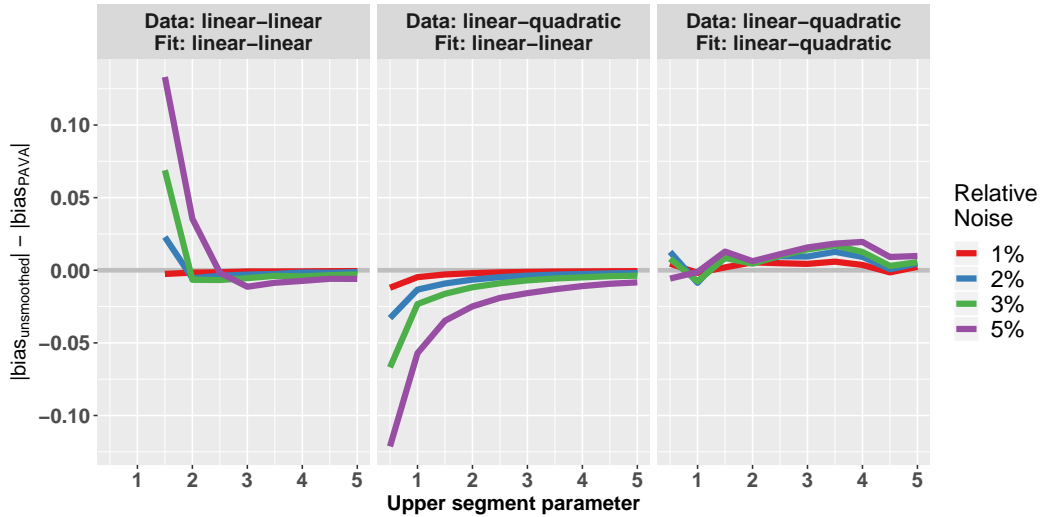
quadratic data. Bias is actually increased for all combinations of relative noise and the coefficient on the quadratic term ( $c_2$ ), although this increase is 0.02 or less for relative noise of 1%. The difference in bias decreases for all relative noise levels as  $c_2$  increases.

When fitting a linear-quadratic model to linear-quadratic data, PAVA reduces the bias at all noise levels for  $c_2$  values from 1.5 to 4, although by less than 0.025 in all cases. Bias increases for a quadratic coefficient of 1, but by less than 0.01.

While PAVA substantially reduces the overall error in the model fit, it does not offer a consistent reduction in bias for the changepoint estimate. However, in cases where PAVA increases the magnitude of the bias, the increase is less than 0.01 when the correct model is fit to the data. The bias increase is more than 0.02 when fitting the linear-linear model to linear-quadratic data when the  $c_2$  is less than 2. Although PAVA does not consistently reduce the bias of the changepoint estimates, the reduction in mean-squared error from the model fit means that PAVA could be a useful step in estimating changepoints with NLS under the right conditions (for example, the linear-linear model fit to linear-linear data with high noise and a small change in slope). PAVA should not be used for the case where a linear-linear model is fit to linear-quadratic data.

#### 4. Analysis of experimental data

We turn now to the problem of estimating the two changepoints on the experimental data. We extend the linear-linear and linear-quadratic models from the simulation study to now include two changepoints. For simplicity we call the linear-linear-linear model the “1-1-1” model and the quadratic-linear-quadratic model the “2-1-2” model, from the order of the polynomials in each segment. The equation that



**Figure 8:** Difference in the absolute value of the bias for unsmoothed data and the absolute value of bias for PAVA-smoothed data. A positive value indicates that the magnitude of the bias is smaller after PAVA, while a negative value indicates the opposite.

Parameter	1-1-1 model	2-1-2 model	Expert
Lower changepoint	1.087 (0.00579)	1.152 (8.735)	1.06 (0.0107)
Upper changepoint	1.735 (0.00469)	1.667 (0.0283)	1.75 (0.00974)
$b_0$	5.180 (0.162)	2.710 (161.622)	N/A
$c_0$	N/A	-9.268 (1.756)	N/A
$a_1$	-7.333 (0.0252)	-7.299 (0.0368)	N/A
$b_1$	2.735 (0.0179)	2.712 (0.0261)	N/A
$b_2$	6.893 (0.353)	1.865 (0.783)	N/A
$c_2$	N/A	21.821 (3.954)	N/A

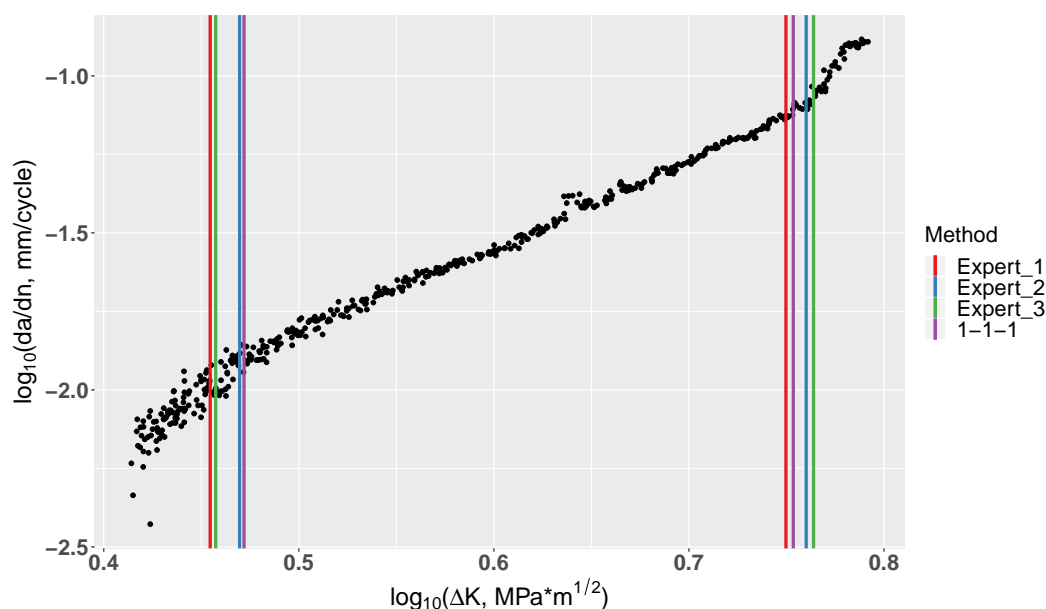
**Table 1:** Parameter estimates for the 1-1-1 and 2-1-2 models, and the mean values from the expert engineers. Standard errors for each point estimate are shown in parentheses. Units on the changepoint estimates are  $\log_{10}(\text{MPa} * \text{m}^{1/2})$ .

we fit using NLS is given by

$$\hat{y} = I_{x < \text{lower changepoint}}(b_0 * x + c_0 * x^2) + I_{\text{lower changepoint} \leq x \leq \text{upper changepoint}}(a_1 + b_1 * x) + I_{x > \text{upper changepoint}}(b_2 * x + c_2 * x^2) \tag{2}$$

where  $c_0$  and  $c_2$  are set at zero for the 1-1-1 model. There are no intercept terms estimated for the lower and upper segments because they are not free parameters. Since we do not know the true changepoints, we first had three expert engineers provide their best guess of the changepoint locations, independent of one another, for comparison with our results. Results for the 1-1-1 and 2-1-2 models are compared with the expert estimates in Table 1. This comparison illustrates why we use the 1-1-1 model for fitting even though the fatigue crack growth data appears to follow a 2-1-2 model. The 2-1-2 model fails to produce a lower changepoint with a reasonable standard error, and the 2-1-2 estimates are not as close to the expert changepoints as the 1-1-1 estimates. The experimental data may correspond to a high bias,





**Figure 9:** Experimental data with changepoints from each of the three experts and the 1-1-1 model. The 1-1-1 estimates agree almost exactly with at least one expert for both the upper and lower changepoints.

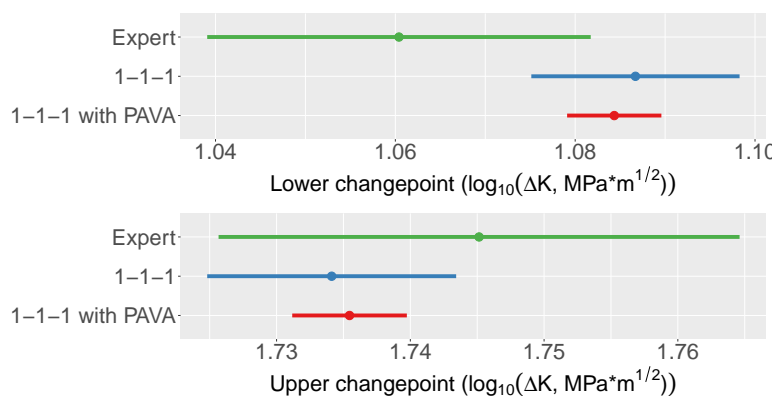
high standard error case in the simulation study where a linear-linear model is fit to linear-quadratic data. In addition, the 2-1-2 model may not be a good fit to the experimental data. The 1-1-1 estimates are much closer to the average values from the experts for both changepoints, and the estimates have reasonable standard errors. Note that the 1-1-1 and 2-1-2 models estimate almost identical values for the slope and intercept of the central linear segment.

Individual estimates from the three experts and the 1-1-1 model are shown in Figure 9. The expert estimates show variability in the changepoint locations, but the estimates from the 1-1-1 model are both in the right area and very close to at least one expert for both the upper and lower changepoints. Results for data from other fatigue tests are similar, so we conclude that the 1-1-1 method is at least as reliable as an expert engineer in estimating the two changepoints.

For the experimental data, Figure 10 shows the point estimates and 95 % confidence intervals for three methods: experts (mean  $\pm 2 \times$  standard error), the 1-1-1 model, and the 1-1-1 model with PAVA. While PAVA isn't recommended for fitting a linear-linear model to linear-quadratic data, we estimate changepoints for PAVA-transformed experimental data for completeness. Based on the simulation results for the case where a linear-linear model is fit to linear-quadratic data, the upper changepoint should be biased high; however, the 1-1-1 model produces an upper changepoint that is less than the average computed for the experts. Although the point estimates for the 1-1-1 method with and without PAVA are similar, PAVA does reduce the widths of the confidence intervals.

## 5. Conclusions

The simulation study provides insight into the behavior of changepoint estimation using NLS for a variety of scenarios. Fitting a linear-linear model to linear-linear data produces the most consistent changepoint estimates, with biases smaller than



**Figure 10:** Comparison of point estimates and 95 % confidence intervals for the lower changepoint (top) and the upper changepoint (bottom).

0.01 in most cases. When fitting a linear-quadratic model to linear-quadratic data, many biases were in the range of -0.025 to -0.1. Changepoint estimates are biased, converging to a bias of 0.32, for all cases considered when fitting a linear-linear model to linear-quadratic data. Additionally, PAVA always reduces the mean-squared error for the model fit, but did not always reduce the bias in the changepoint estimates and should be used with caution.

The application to fatigue crack growth data shows that the 1-1-1 model provides realistic changepoint estimates even though the model doesn't accurately represent the data in the lower and upper segments. This method performs quite well compared with changepoint estimates made by experts in the field.

Future work will focus on two areas. First, we will estimate changepoints in the context of a designed experiment to evaluate how material conditions affect the changepoint locations. The second area of work involves missing data. Because valid data can only be collected while the crack is below a certain percentage of the sample width (ASTM E647), a particular test may only provide data on one of the two changepoints. Utilizing data from these tests would be useful, since fatigue tests are expensive.

This work builds the foundation for a methodical, data-driven approach to the analysis of fatigue crack growth data. The objectivity provided by our statistical approach will be useful to the scientific community.

## References

- M. Ayer, H.D. Brunk, G.M. Ewing, W.T. Reid, and E. Silverman. An empirical distribution function for sampling with incomplete information. *The annals of mathematical statistics*, pages 641–647, 1955.
- R.E. Barlow and H.D. Brunk. The isotonic regression problem and its dual. *Journal of the American Statistical Association*, 67(337):140–147, 1972.
- R.L. Dykstra. An isotonic regression algorithm. *Journal of Statistical Planning and Inference*, 5(4):355–363, 1981.
- R.W. Hertzberg. *Deformation and fracture mechanics of engineering materials*. J. Wiley & Sons, 1996. ISBN 9780471012146.

- P. Mair, K. Hornik, and J. de Leeuw. Isotone optimization in r: pool-adjacent-violators algorithm (pava) and active set methods. *Journal of statistical software*, 32(5):1–24, 2009.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2018. URL <https://www.R-project.org/>.
- T. Robertson, F.T. Wright, and R. Dykstra. *Order Restricted Statistical Inference*. Probability and Statistics Series. Wiley, 1988. ISBN 9780471917878.