Presented at the 2019 ICEAA Professional Development & Training Workshop

Adaptive Curve Fitting: An Algorithm in a Sea of Models

MICHAEL L. SCHIAVONI

Tecolote Research, Inc.

Adaptive Curve Fitting (ACF) is a novel technique that analyzes finite time series data such as monthly expenditures for government acquisition contracts. By intelligently fitting known resource phasing curve forms to an existing sequence of data, the automated algorithm generates a custom model that extrapolates all remaining future values. It performs joint forecasting of cost, schedule, and phasing for in-progress programs and can be updated on a monthly basis. In this regard, it can provide a cross-check to more traditional earned value methods, independent estimates, customer budgets, and contractor-reported plans. Furthermore, the total cost and duration of the forecast can optionally be bounded, thereby allowing the rapid exploration of excursions and "What if" scenarios. ACF has primarily been developed with Department of Defense satellite acquisitions in mind; however, it potentially has wide-reaching applications to other commodities and industries. ACF is based in theory and empirical research, shares similarities with existing techniques, and introduces several key innovations to the field. The algorithm is explained and validation results presented for a sample of completed programs with known actuals.

Introduction

A considerable amount of research has been done on the topic of the ideal or most likely resource phasing shape for large acquisition programs. Early influential works were published by Norden (1958, 1970), who first noted that the Rayleigh distribution approximates the staffing buildup and phase-out of hardware development, and Putnam (1978), who applied the Rayleigh form to software estimation. More recently, researchers also began assessing the more general Weibull distribution (of which Rayleigh is a specific case) and Beta distribution (Unger 2001), (Burgess 2006), (Brown et al. 2015).

Typically at the start of a large contract, a simple phasing model based on one of these forms is used in order to budget the upcoming years. The width of the curve is fixed to the duration estimate, and the total area under the curve is fixed to the cost estimate (or some other resource of interest, e.g. labor hours, staffing numbers, etc.). This model remains useful as long as the actual phasing curve continues to approximately follow it. However, consider the situation in which an in-progress program has strayed significantly from its original spending plan. Since the baseline has failed to accurately predict the actuals to date, it no longer makes sense to continue using it to forecast the future. Additionally, the failure to meet the phasing model calls into question the accuracy of the cost and duration estimates used to generate it.

This situation highlights the need for an objective technique that can optimally re-phase the remainder of the program while taking into consideration the performance to date. Proposed in this paper is an *algorithmic* solution to this problem that has applications extending beyond phasing alone. An algorithm is a defined set of instructions that takes an input, performs some sequence of calculations, manipulations, or decision-making processes, and produces an output. It can be thought of as a small computer program, and can often be visualized as a flowchart. An adaptive algorithm is one that changes its behavior based on the input information. Note that this is in stark contrast to traditional parametric models, which utilize historical training data to develop a static

equation that is applied the same way from that point forward. Adaptive algorithms are common in certain industries outside the cost/schedule/EVM world, for example radar detection (Finn & Johnson 1968), guidance/navigation/tracking (Hu et al. 2003), audio compression (Tsutsui et al. 1992), and noise-cancelling headphones (Kuo & Morgan 1999), to name just a few.

Adaptive Curve Fitting (ACF) is an automated procedure that analyzes the burn rate trend of an in-progress contract to generate updated forecasts for the phasing profile, total cost, and duration. Figure 1 illustrates the concept with simulated expenditure data that closely resembles a real acquisition program. The left portion of the plot with the turquoise background represents the past, and is therefore known; this is the data that is fed to ACF. The right portion with the white background represents the future, and is therefore unknown. The algorithm cannot see this data; it is provided here for comparison purposes only. The noisy data are monthly actuals from the time series of interest (it represents expenditures in this example but could alternatively represent another resource). The solid black points are smoothed monthly actuals that approximate the general trend of the burn rate. The dashed magenta lines are phasing curve forms that are fit to the smoothed data, and finally the solid magenta line is the resulting forecast.

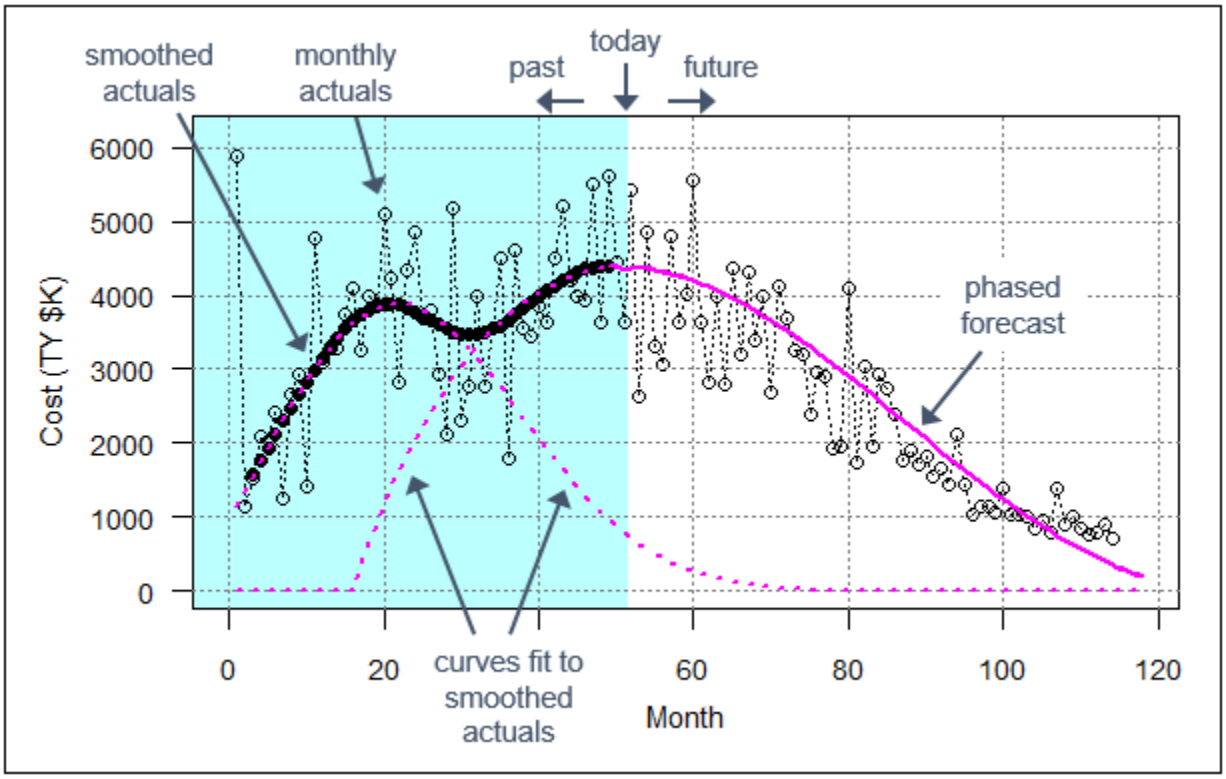

FIGURE 1 Adaptive Curve Fitting Example with Known Past Data and Unknown Future Data

Prior related techniques that take advantage of some of the same fundamental principles as Adaptive Curve Fitting include: Multiple Model Adaptive Estimation (Gallagher & Lee 1995), The Rayleigh Analyzer® (Dukovich et al. 1999), N-R Curve Generation Tool (Chelson et al. 2004), Executive Cost and Schedule Assessment tool (Davis 2008), and Weibull Analysis Method (Burgess et al. 2014).

Once ACF generates its monthly forecast, there are several avenues for analysis. Figure 2 shows a five-year annual phasing forecast created by binning the monthly values into years, which

could be useful for planning purposes. Additionally, if a user is interested in a group of programs, they could apply ACF to all of them and use the resulting outputs to optimally budget resources between programs within the portfolio.

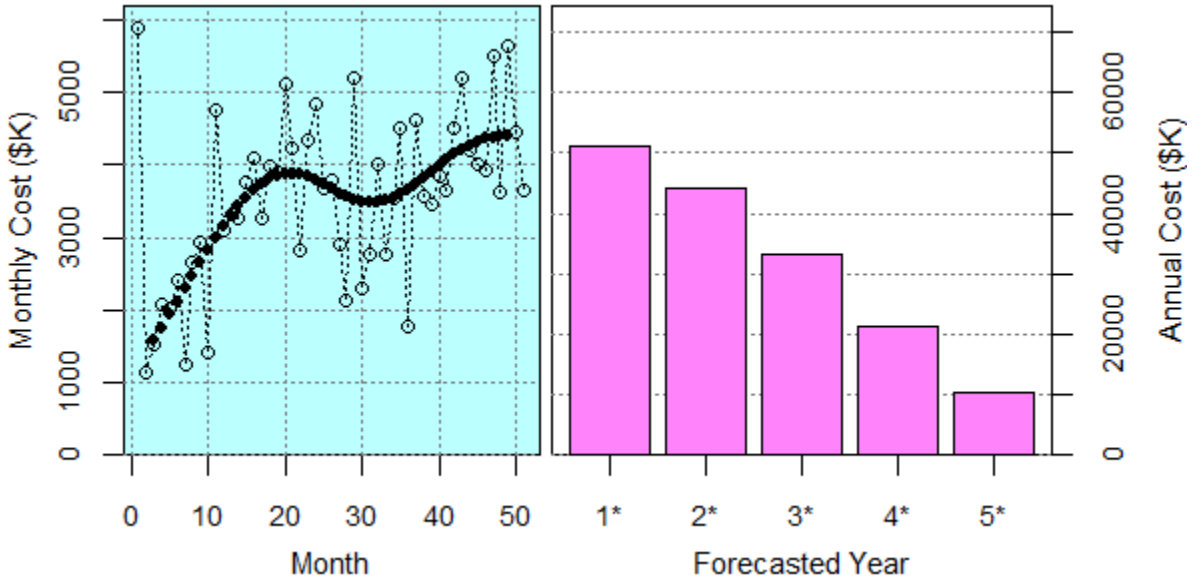

FIGURE 2 Annual Phasing Forecast Output

Figure 3 displays the predicted *cumulative* curve, which is simply the running total of the monthly forecast. The maximum value yields the ACF-derived estimate at completion (EAC). Additionally, if a particular milestone date is of interest (e.g. first launch for a satellite program), ACF applies a percent spent metric to the cumulative curve to predict the month of occurrence. The value chosen for the metric can be based on historical data, subject matter expert input, contractor plan, or some other source.

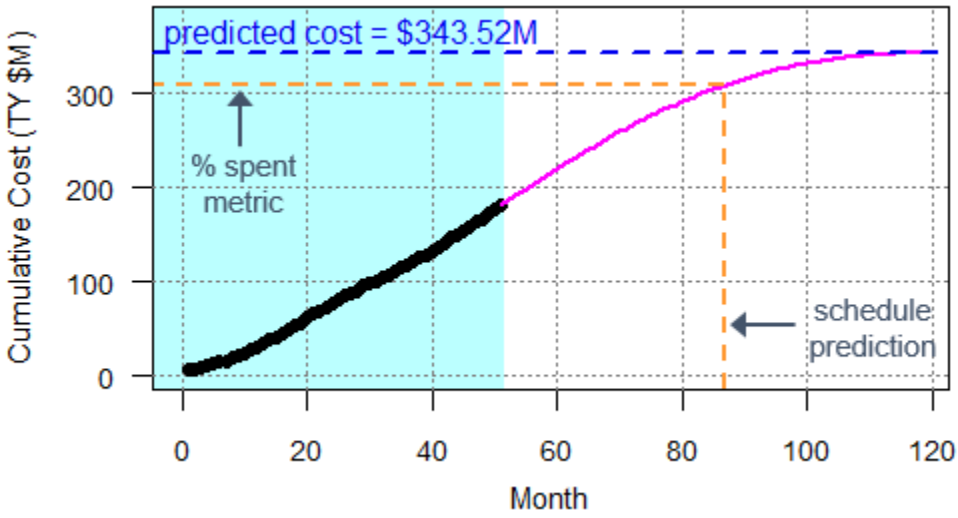

FIGURE 3 Cumulative Monthly Forecast Output with Cost and Schedule Predictions

One of ACF's features is that it can be applied continuously throughout a contract's lifecycle. Figure 4 shows how the schedule prediction has trended over time for the example dataset. The vertical axis represents the schedule prediction, and the horizontal axis represents the number of months of available data. This type of plot can assist the analyst in determining whether the forecast is stable or trending upward or downward.

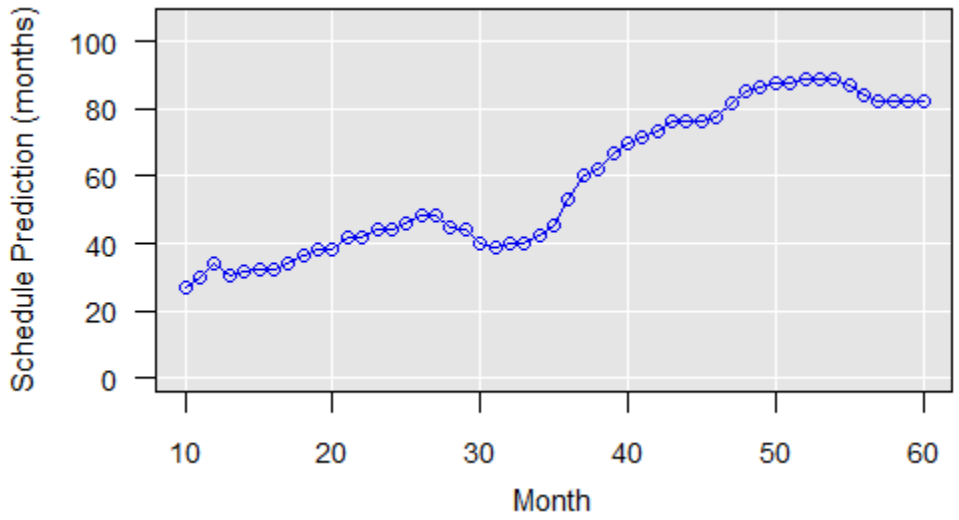

FIGURE 4 Schedule Prediction Time Trend Output

Another important feature of ACF is that it can constrain its forecast to be consistent with an existing duration estimate and/or EAC. In Figure 5 the dashed magenta line represents the initial model that was fit to the smoothed data, and the solid magenta line is the final model that has been expanded until certain conditions were met (e.g. the total area under the model meets some total cost, or the width of the model equals some duration). This allows for the generation of a phasing forecast that is consistent with the latest cost and schedule estimates. It also enables the exploration of "What if" scenarios, for example predicting the cost effect of a schedule delay, or the phasing effect of a cost overrun.

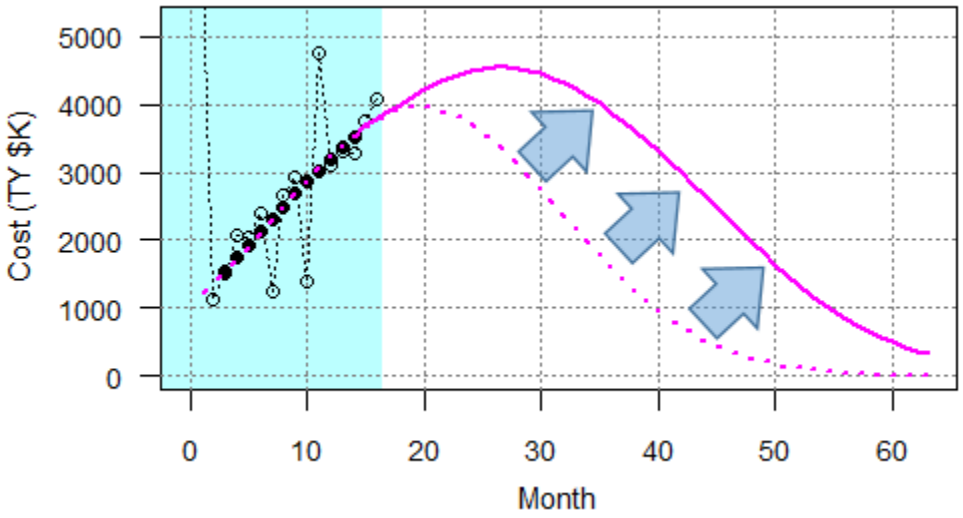

FIGURE 5 Constrained Forecast

Methodology

The following flowchart (Figure 6) summarizes the Adaptive Curve Fitting algorithm.

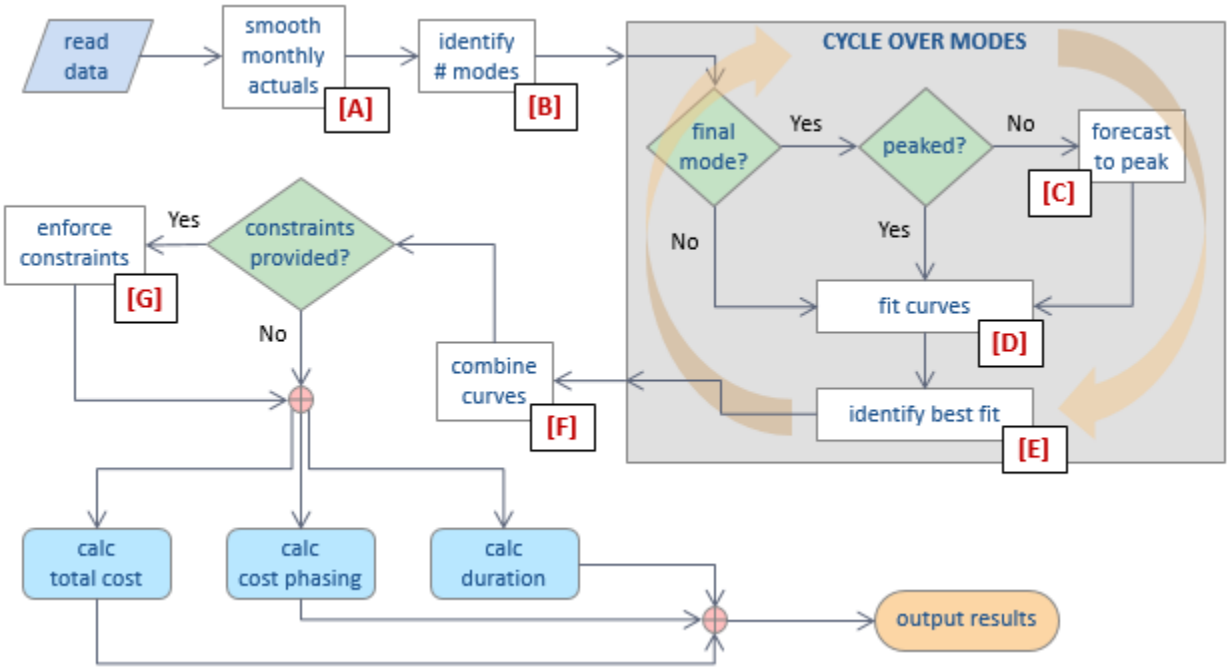

FIGURE 6 Algorithm Flowchart

The purpose of smoothing the data [A] is to extract the general trend of the burn rate. It is a multi-step process:

1) Replace extreme positive outliers (if applicable)
2) Reallocate negative values (if applicable)
3) LOESS compression
4) Iterative moving average filter
5) Prune initial and final values

Figure 7 illustrates the positive outlier replacement process with simulated expenditure data. The black time series represents raw data as reported by a contractor, and the red time series is the corrected data. In this case, the first six months of data is unavailable and all those dollars were dumped into month seven. Additionally, there is a significant spike at month 38, which likely represents a large one-time purchase of materials or a subcontractor fee payout. Since these months are in the upper 90% of data and are greater than three times the value of both of their nearest neighbors, they are flagged and replaced by the nearest neighbor. This rule was determined by analyzing historical data streams from real acquisition programs. If left in the data, these points could potentially exert undue influence on the rest of the smoothing process, thereby yielding an artificially perturbed trend. It is important to note that this process only adjusts the most severe outliers, and in practice, many datasets have no such extreme values.

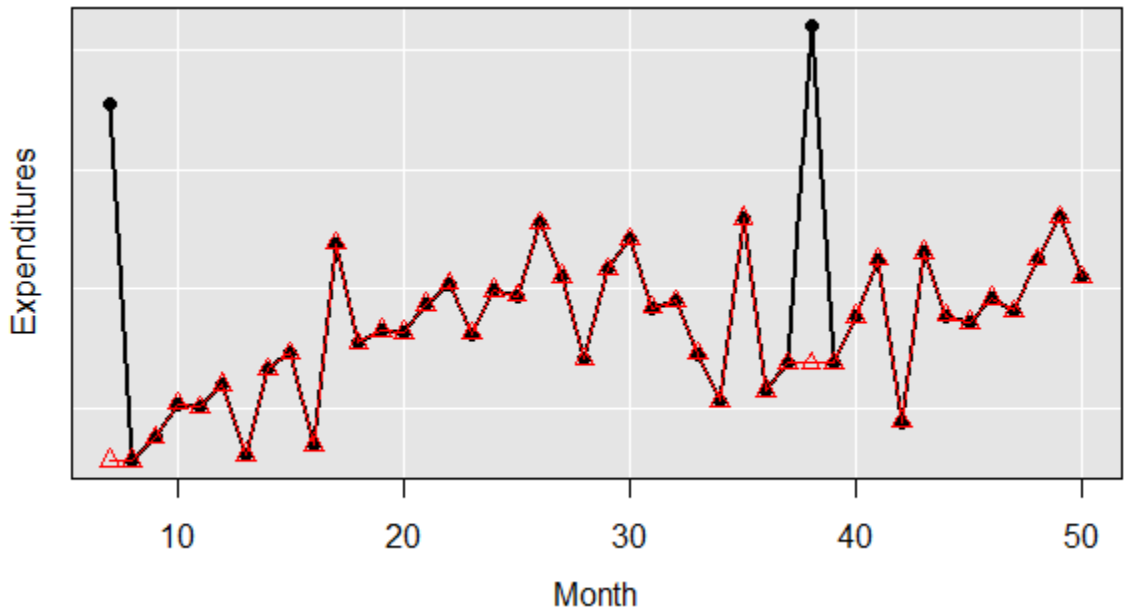

FIGURE 7 Positive Outlier Replacement: Input (black) and Output (red)

Figure 8 illustrates the negative value reallocation process. In this case, month 16 was reported as a negative expenditure. This is most likely the result of an accounting adjustment, e.g. correcting previous unallowable costs, or moving expenses from one WBS/CLIN to another. After all, it is unlikely that the contractor did no work this month and instead paid money to the customer. The algorithm replaces this negative value by the average of its two nearest neighbors in attempt to better approximate the true expenditure burn rate over time. Then, the amount of adjustment (in this case, $\boldsymbol{y}_{16}^{new} - \boldsymbol{y}_{16}^{old}$) is removed proportionally from the preceding months.

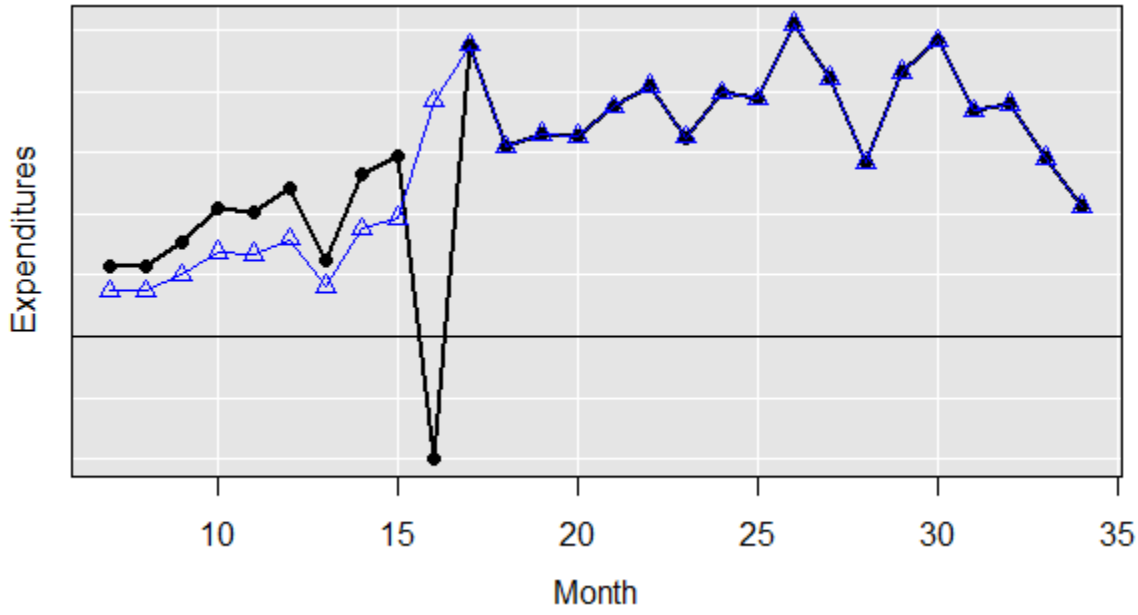

FIGURE 8 Negative Value Reallocation: Input (black) and Output (blue)

After extreme outliers and negative values have been processed, the data continues to the next step, which the author has named "LOESS Compression". It is conceptually related to dynamic range compression in audio signal processing, although the process and result are different. Figure 9 illustrates the technique. The black data is the input, and the magenta curve represents a quadratic locally weighted regression (LOESS) fit. The green time series is the output, which is calculated by proportionally compressing the input signal toward the LOESS curve. In this plot, each of the green points lies halfway between the input data and the magenta curve. This serves to temper the noise in the time series, since more extreme values will have a greater compensation than values that lie closer to the local trend.

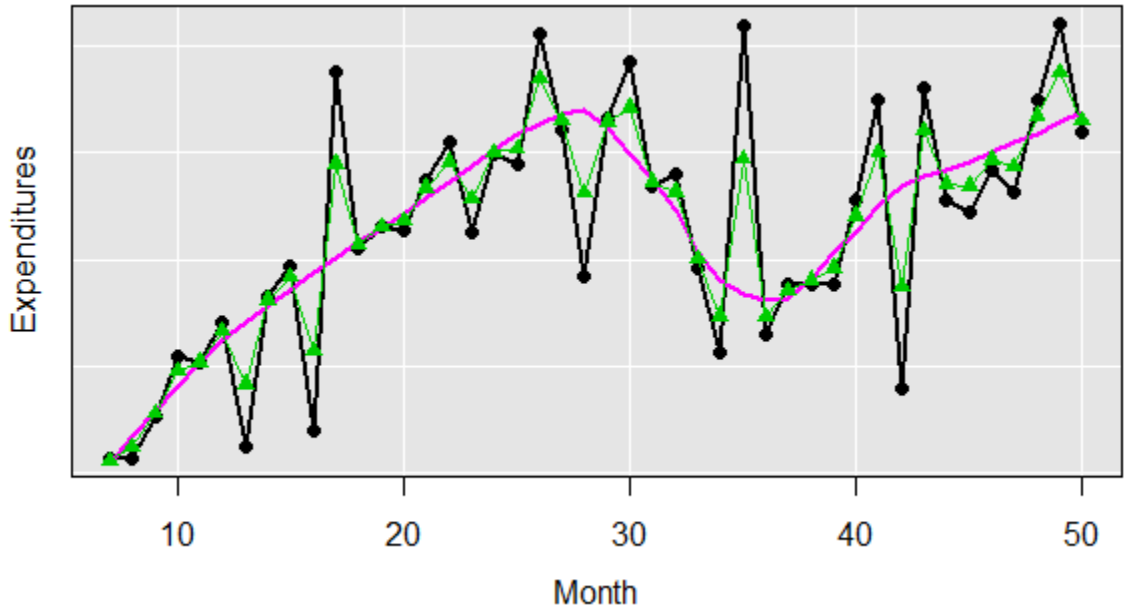

FIGURE 9 LOESS Compression: Input (black), Local Fit (magenta), Output (green)

Next, an iterative moving average filter performs the bulk of the smoothing. Figure 10 shows the input data in black, as well as the results of the first and third moving average iterations in red and blue, respectively. The final step is to prune the first two and last two data points from the smoothed trend. Endpoints are notoriously difficult to smooth (Müller 1991), (Vint & Hinrichs 1996), and simply removing them aids some of the subsequent procedures of the algorithm.

After the smoothing procedure, ACF scans the burn rate trend for any points at which the slope transitions from negative (downward) to positive (upward) [B]. This is known as a local minimum. A nominally executed program often has none, but the example dataset in Figure 1 displays one such point at month 31. This point is used as the dividing line to separate the smoothed trend into two discrete segments (the author refers to this as a "multi-modal" program, or more specifically in this case a "bi-modal" program). This phenomenon of a burn rate peaking and decreasing only to increase again has been observed in real government acquisition programs, and some apparent causes have been identified:

- an engineering change proposal (ECP) adds significant scope to the existing contract
- a serious technical issue is discovered, thereby requiring redesign of a major component
- the contractor is behind schedule and tries to catch up by ramping up resources

Additionally, annual budgeting constraints affect the phasing of resources, and therefore budget changes could conceivably be another cause of multi-modal phasing shapes.

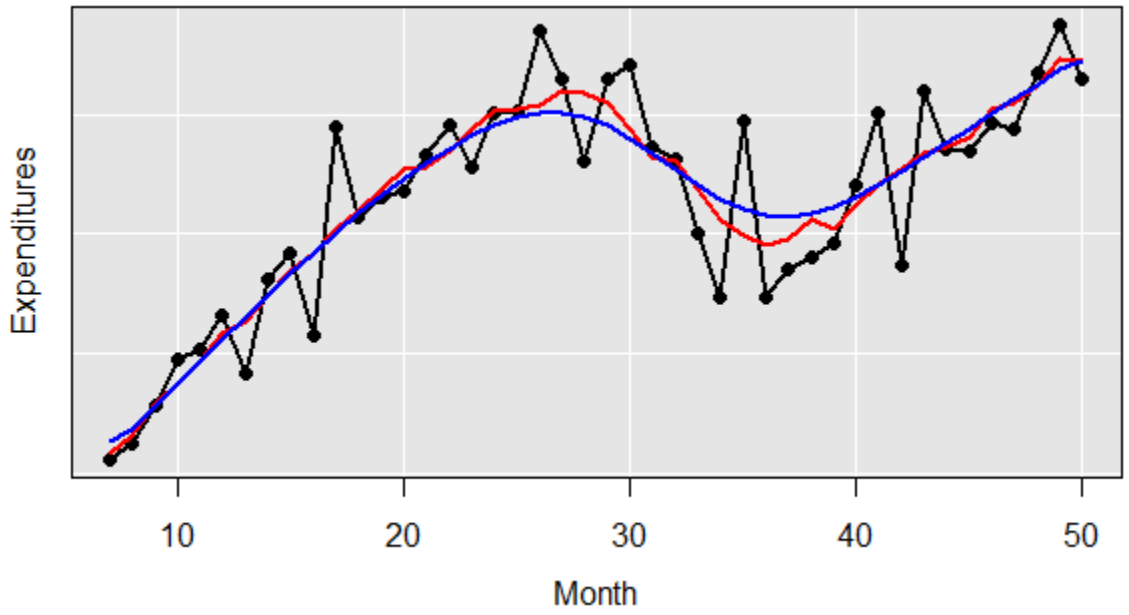

FIGURE 10 Iterative Moving Average: Input (black), 1st Iteration (red), 3rd Iteration (blue)

Next, a nonlinear least-squares solver fits the best curve to each smoothed data segment independently [D]. The three established phasing forms of Rayleigh, Weibull, and Beta are fit, in addition to the Normal distribution. The Normal form is included for two reasons: First, in analyzing real data, a need for overall symmetric phasing shapes was identified. Second, even for a more typical front-loaded time series, in the multi-modal case the individual segments can exhibit a symmetric behavior. Whichever of these four curve forms yields the minimum sum of squared errors for the smoothed trend is utilized for that segment [E], and the initial model is formed as the maximum value of the overlapping curve segments at each month (if there are more than one) [F]. Additionally, the tail of the model is truncated by enforcing a minimum burn rate value. The current rule-of-thumb being utilized is 5% of the peak model value, where "model" in this usage refers to the concatenation of the smoothed trend and the forecasted curve.

Note in the algorithm flowchart (Figure 6) in the scenario in which the final mode has not yet peaked (i.e. its smoothed trend is on the rise), it is forecasted to its likely peak before any curve fitting takes place [C]. This is depicted in Figure 11, and it is accomplished by analyzing the recent smoothed values to calculate the latest 1st, 2nd, and 3rd-order derivatives ($d^{(1)}$, $d^{(2)}$, and $d^{(3)}$, respectively). $d^{(1)}$ is the slope, $d^{(2)}$ represents the concavity, and $d^{(3)}$ equals the rate of change of the concavity. This information is used to treat the forecast like a projectile motion problem in which the future path is predicted based on the recent trajectory. In a typical introductory physics course, such problems assume a constant $d^{(2)}$, i.e. the acceleration due to gravity. ACF extends this principal to allow a varying $d^{(2)}$, which enables a trend that is increasing at an increasing rate (as is the case in Figure 11) to transition to curving downward, thus eliminating the possibility of a runaway forecast. For the example shown, the combined data of the smoothed trend plus the projection to the predicted peak is used for curve fitting. This process acknowledges the inherent "momentum" in acquisition programs, as it is difficult to change staff sizes up or down rapidly.

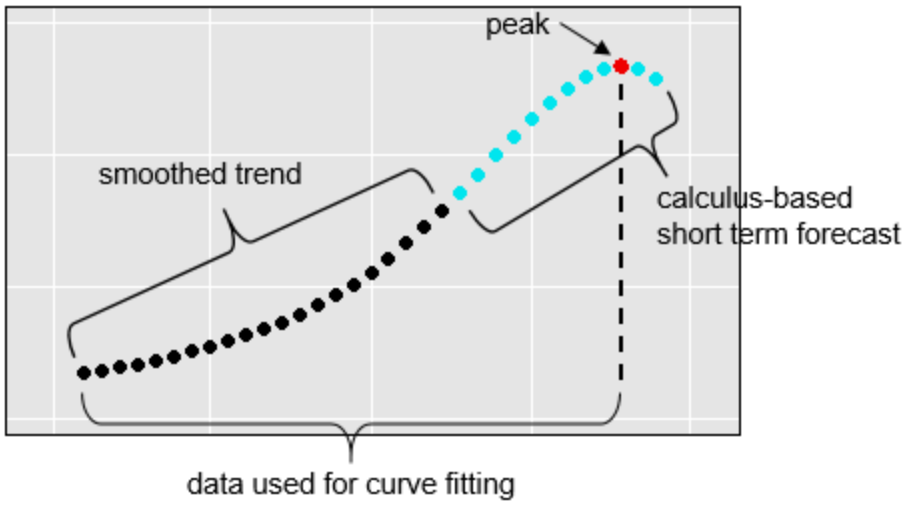

FIGURE 11 Curve Projection with Peak Detection for Increasing Curve Segment

ACF has the ability to enforce limits on the duration and/or total area of its forecasts [G]. Area constraining involves multiplying the forecast by a tapering function to reduce it, or a growth function to increase it. Figure 12 shows an unconstrained forecast on the left (dotted curve) along with four different constrained forecasts (thin solid curves). The constraints are applied by iteratively multiplying the initial model by the linear tapering function pictured in the right cell.

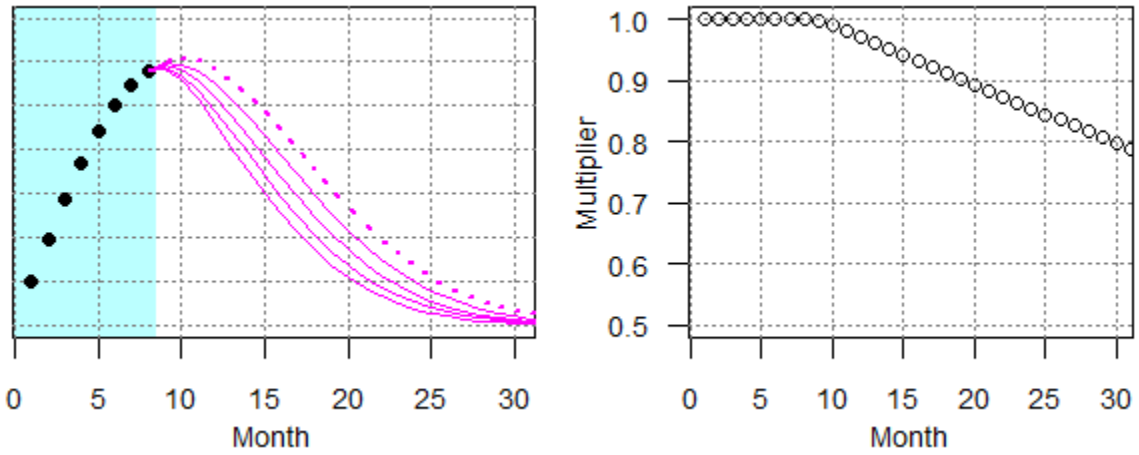

FIGURE 12 Area Constraining (Reduction, left) by Tapering Function (right)

Figure 13 shows how duration constraints are handled. The upper-left cell depicts an unconstrained forecast, which is the resource phasing form that best fits the smoothed data. This curve is sampled once per month, and the forecast has a 17 month to-go duration. The upper-right cell is the same forecasted curve, however it has been oversampled via cubic spline interpolation (there are now 27 points packed into the same 17-month span). The lower cell displays the result when the oversampled curve is evenly spaced out to one point per month, yielding an expanded forecast that remains consistent with the smoothed actuals. The amount of oversampling determines the amount of expansion, and the inverse process also works: *under*sampling yields a shortened forecast.

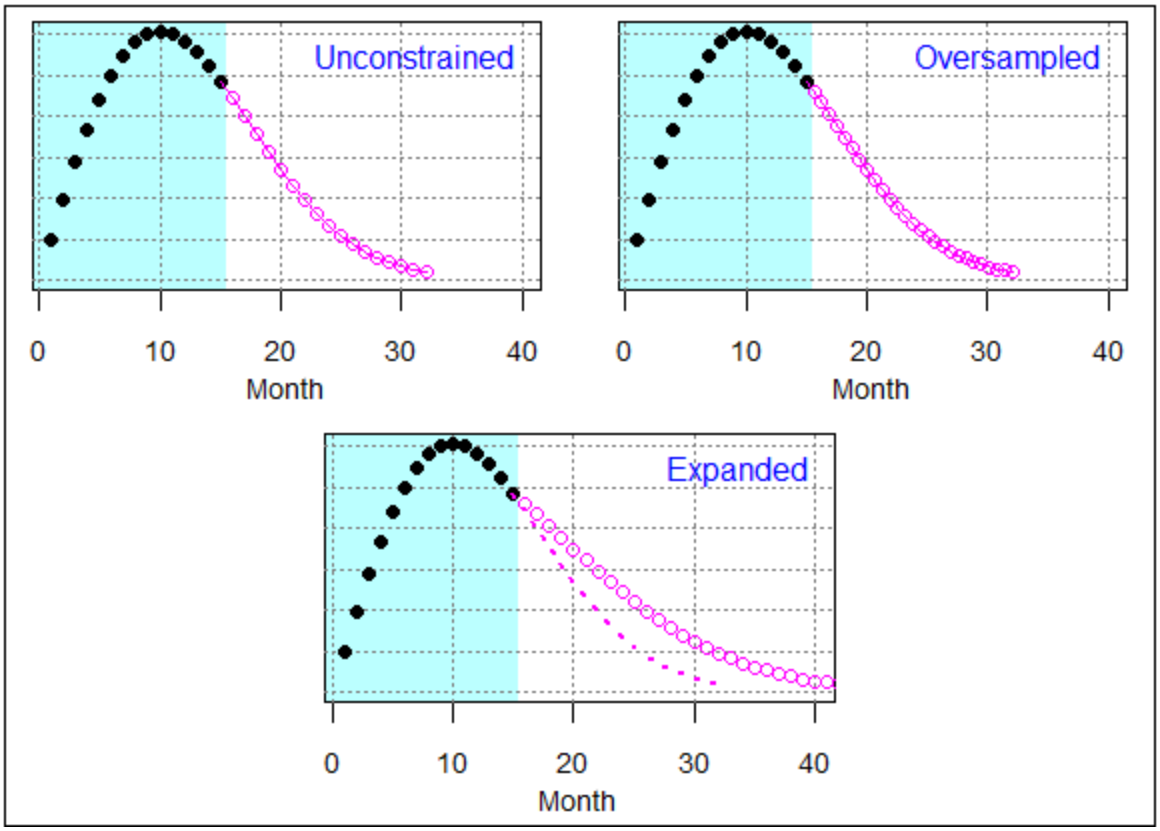

FIGURE 13 Duration Constraining (Expansion by Oversampling via Interpolation)

An important point about the Adaptive Curve Fitting algorithm is that it makes no assumptions about the units of the input data or the type of commodity from which it came. What it does assume is that the data is finite (i.e. it will have an endpoint) and that its shape can be modeled as one or more known probability density functions. Therefore, one should be able to apply it with little or no modification to any time series data that meets these two requirements. Some examples of other potential data streams include: earned value data, labor hours/heads, and number of concurrent schedule tasks. Conversely, data sources that are probably not amenable to this technique include stock market indices, sustainment contracts and other constant level of effort tasks, and production contracts with many units manufactured in an assembly line process.

Results

All errors presented in this section are "true" errors, i.e. $(prediction - truth)/truth$. Negative values represent under-predictions and positive values represent over-predictions. Efforts are ongoing to improve upon the following results and gather more data to increase the sample size.

To date, ACF has primarily been developed and tested on expenditure data for satellite acquisition programs. A diverse sample of 20 programs with known monthly actuals was assessed for accuracy of phasing, schedule, and cost predictions. It includes military and NASA programs, development and follow-on contracts, space and ground domains, and sensing and communications missions.

Figure 14 shows annual phasing errors for five years of forecasting. The blue points represent the sample mean, the error bars signify the +/– 1σ bounds, and the horizontal axis denotes how far

into the future is being forecasted. The unconstrained case depicts phasing forecast accuracy when no information is available about the final cost or duration; real-world performance improves when reasonable bounds are enforced on one or both of these variables. The constrained case shows the accuracy when the true cost and duration are known, which provides a theoretical upper limit for the accuracy of ACF's phasing forecasts. Furthermore, the approximate zero bias and reasonable uncertainty bounds indicate that constrained ACF is an appropriate technique for generating an updated annual phasing forecast that is consistent with the existing cost and duration estimates.

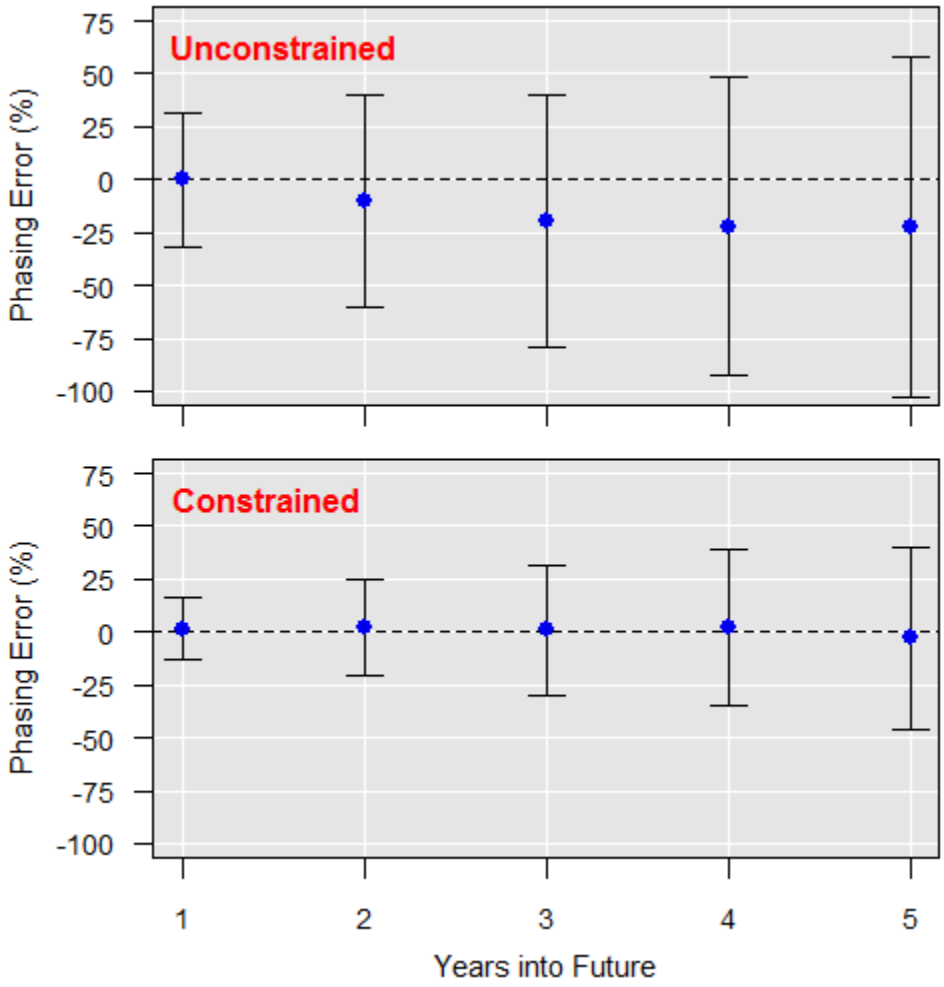

FIGURE 14 Annual Phasing Errors (mean +/– 1σ)

Figure 15 shows schedule prediction errors over time for first launch availability. The dashed blue line is the sample mean, and the shaded area is the +/– 1σ bounds. The horizontal axis denotes how far along the program is (e.g. 0.4 could represent 40 months into a 100-month program, or 24 months into a 60-month program). The upper cell depicts unconstrained results, meaning no information is available about total cost. Performance becomes good beginning around 40% schedule completion, as denoted by the low bias and variance. For reference, this tends to align with some time shortly after critical design review (CDR) for a typical space vehicle program. The middle cell illustrates reduction in bias and variance when relatively loose constraints are applied

(+/– 20%) to the total cost of the ACF model (i.e. the area under the curves). The lower cell is a theoretical upper limit, as the model is constrained to true final cost.

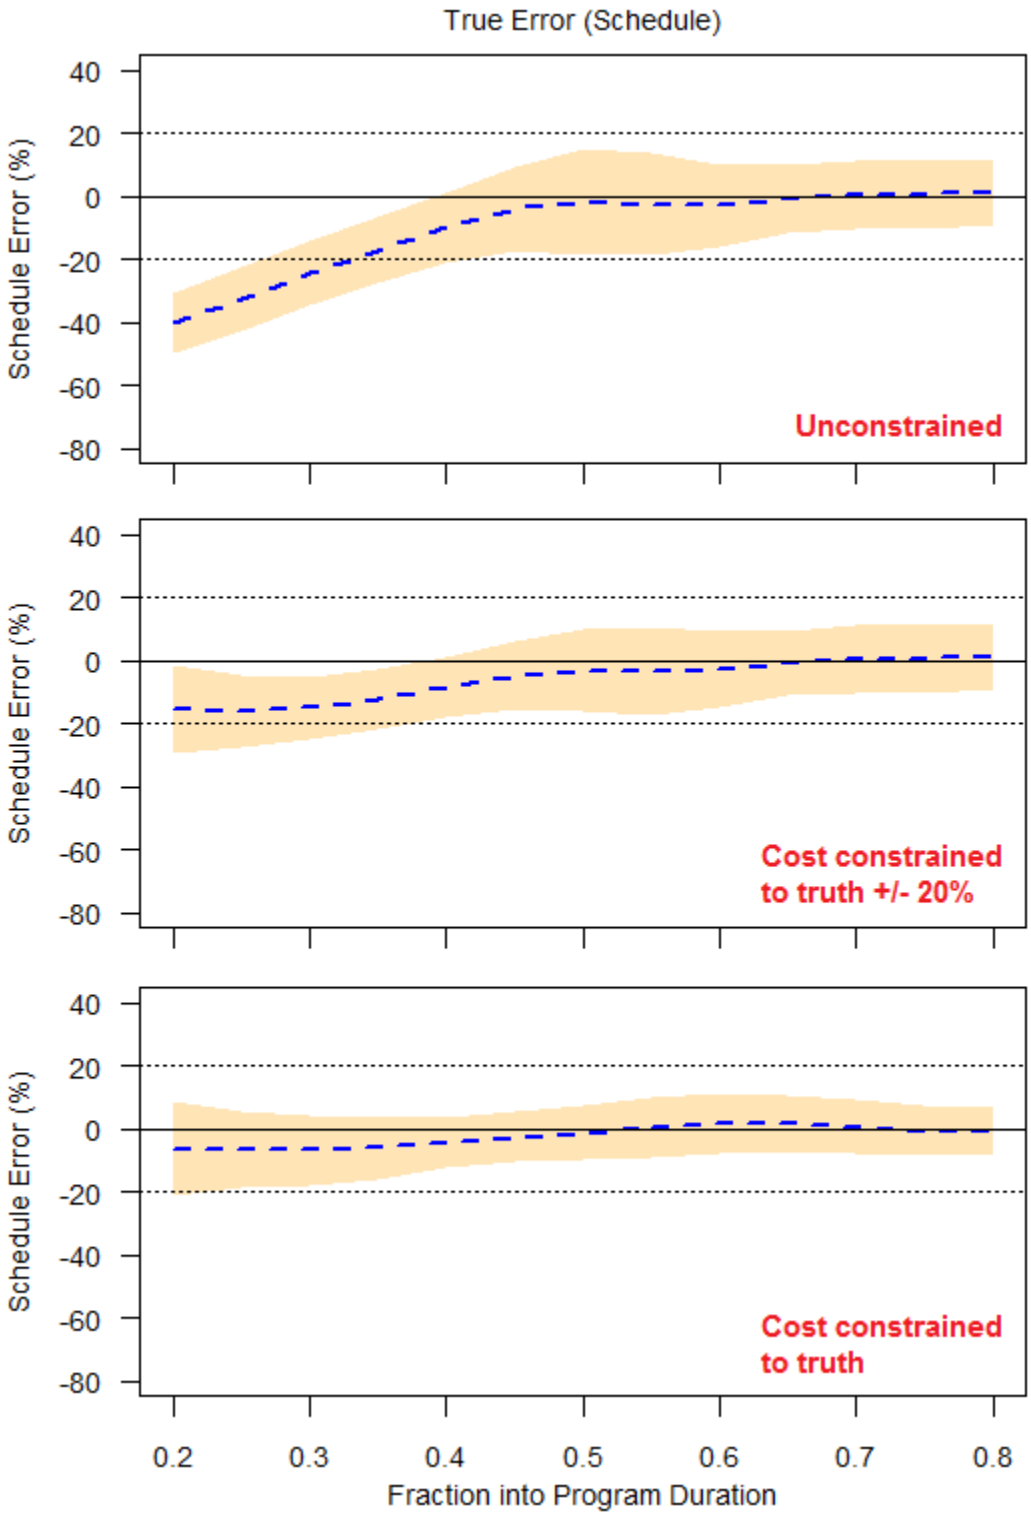

FIGURE 15 Percent Errors over Time for Predicting Month of First Launch

Figure 16 shows similar plots to Figure 15 except total cost at final launch is being predicted, and the effects of duration constraints are illustrated. One can draw similar conclusions from them.

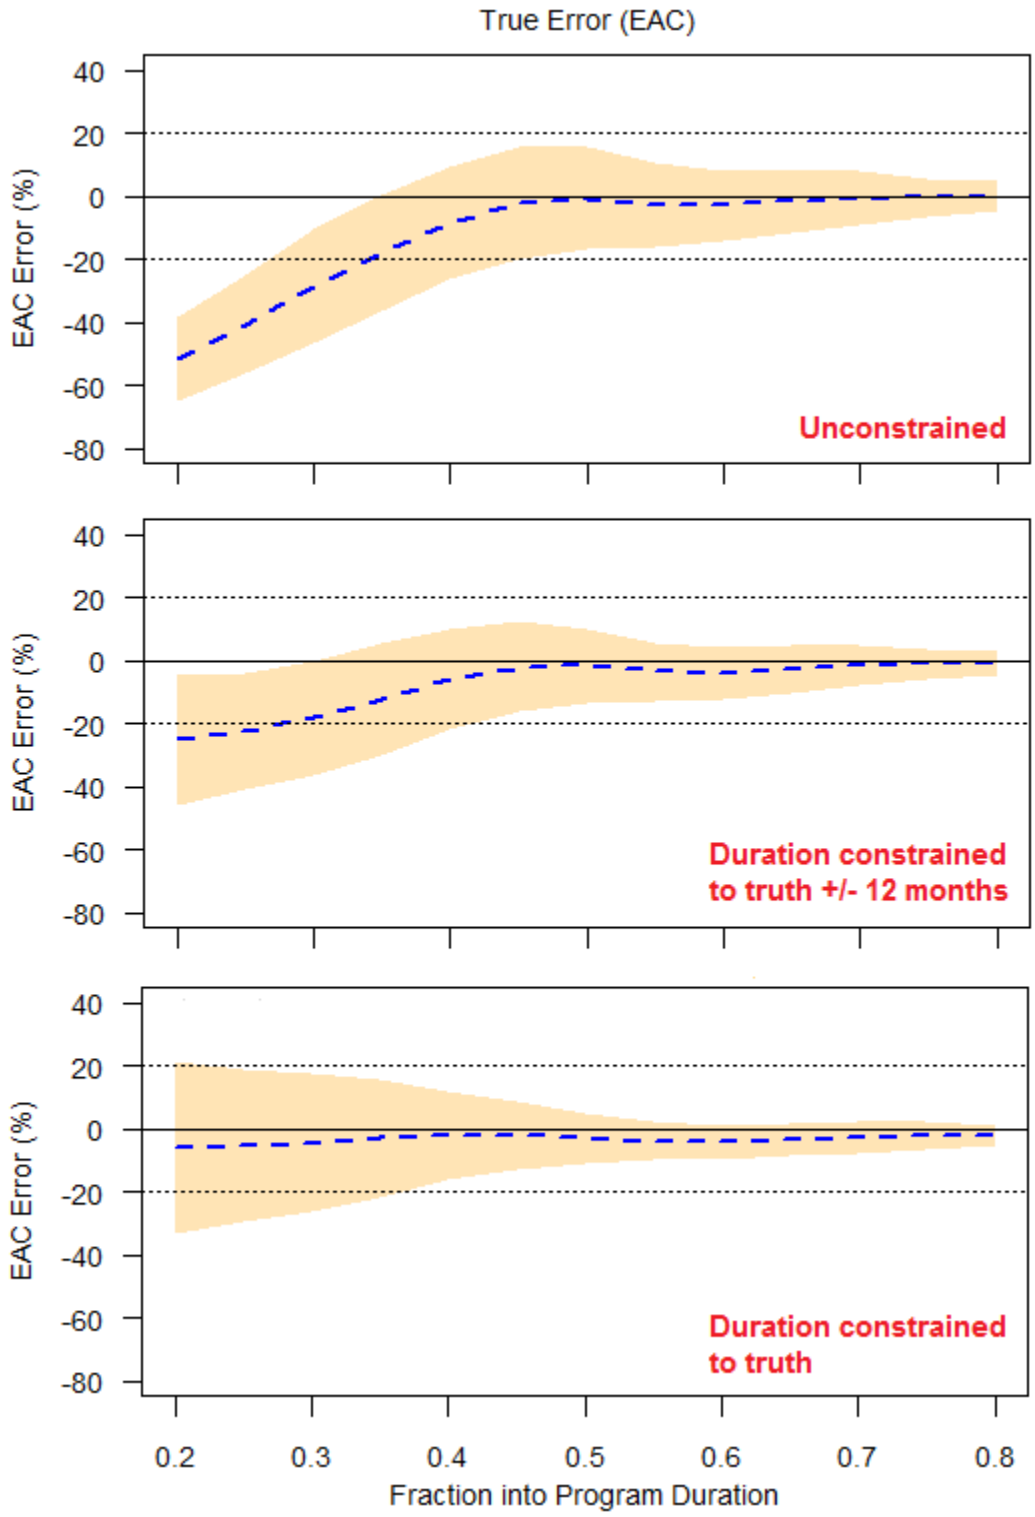

FIGURE 16 Percent Errors over Time for Predicting Cumulative Cost at Final Launch

What are not shown on the preceding error plots are comparisons with predictions from other contemporaneous sources such as contractor-submitted integrated master schedules (IMS) or latest revised cost estimates. When these comparisons are made, in many cases ACF has shown the ability to detect cost overruns and schedule slips earlier. Individual program-specific data is not included in this publication to avoid any potential data rights concerns.

Conclusions

The Adaptive Curve Fitting algorithm intelligently fits known resource phasing curve forms to monthly time series data in order to forecast the future while adjusting to current program performance. It extracts the burn rate trend through a multi-step smoothing procedure that is robust to the variations commonly found in data from monthly contractor submittals. Additionally, the method automatically detects compound phasing shapes, also known as multi-modal curves. Since it makes limited assumptions about the source of the input data, ACF is potentially applicable to a wide variety of data across commodities. It can be applied consistently across programs by different individuals: For the analyst, it offers the ease of generating independent crosschecks for existing cost and schedule estimates, while enabling rapid exploration of "What-if" scenarios. Program managers can use ACF on a monthly basis to monitor for potential schedule slips and cost overruns. Lastly, organizations responsible for a number of contracts can apply the technique during the budgeting process to assist in optimally distributing resources across the portfolio.

ACF provides objective analyses based in theory and empirical research, and validation efforts have shown that the algorithm works well on historical programs when utilized as intended. Space vehicle expenditure data has been most thoroughly assessed: unconstrained prediction of both schedule and total cost reaches low bias and variance after approximately 40% completion of program duration, and constraining the forecasts yields improved accuracy. As a rule of thumb, one should not perform analysis prior to CDR unless constraints are enforced or a bias compensation factor is applied to correct for the known tendency toward under-prediction during this early time period. For annual phasing forecasts, constraining is recommended to ensure consistency with existing cost and duration estimates.

Acknowledgments

This work was funded in part by Space and Missile Systems Center (SMC) and Air Force Cost Analysis Agency (AFCAA).

Acronyms

ACF	Adaptive Curve Fitting
CDR	Critical Design Review
CLIN	Contract Line Item Number
EAC	Estimate at Completion
ECP	Engineering Change Proposal
EVM	Earned Value Management
IMS	Integrated Master Schedule
LOESS	Locally Weighted Regression
NASA	National Aeronautics and Space Administration
WBS	Work Breakdown Structure

References

Brown, G.E., White, E.D., Ritschel, J.D., Seibel, M.J. (2015). *Time Phasing Aircraft R&D Using the Weibull and Beta Distributions*, Journal of Cost Analysis and Parametrics, vol. 8:3, pp. 150-164.

Burgess, E.L. (2006). *R&D Budget Profiles and Metrics*, Journal of Parametrics, vol. 25, pp. 11-29.

Burgess, E.L., Smirnoff, J., Wong, B. (2014). *Weibull Analysis Method*, ICEAA Professional Development & Training Workshop.

Chelson, H.F., Coleman, R.L., Summerville, J.R., Van Drew, S.L. (2004). *Rayleigh Curves – A Tutorial*, SCEA-ISPA Joint Annual Conference and Training Workshop.

Davis, D. (2008). *Early Warning Model for Acquisition Program Cost and Schedule Growth*, SCEA-ISPA Joint Annual Conference and Training Workshop.

Dukovich, J., Houser, S., Lee, D.A. (1999). *The Rayleigh Analyzer® Volume 1 – Theory and Applications*, Logistics Management Institute.

Finn, H.M., Johnson, R.S. (1968). *Adaptive Detection Mode with Threshold Control as a Function of Spacially Sampled Clutter-Level Estimates*, RCA Review, vol. 29, pp. 141-164.

Gallagher, M., Lee, D.A. (1995). *Final-Cost Estimates for Research & Development Programs Conditioned on Realized Costs*, Report from the Office of the Secretary of Defense.

Hu, C., Chen, W., Chen, Y., Liu, D. (2003). *Adaptive Kalman Filtering for Vehicle Navigation*, Journal of Global Positioning Systems, vol. 2, no. 1, pp. 42-47.

Kuo S.M., Morgan, D.R. (1999). *Active Noise Control: A Tutorial Review*, Proceedings of the IEEE, vol. 87, issue 6, pp. 943-973.

Müller, H-G. (1991). *Smooth Optimum Kernel Estimators Near Endpoints*, Biometrika, vol. 78, no. 3, pp. 521-530.

Norden, P.V. (1958). *Curve Fitting for a Model of Applied Research and Development Scheduling*, IBM Journal of Research and Development, vol. 3, no. 2, pp. 232-248.

Norden, P.V. (1970). *Useful Tools for Project Management*, Management of Production, pp. 71-101.

Putnam, L.H. (1978). *A General Empirical Solution to the Macro Software Sizing and Estimating Problem*, IEEE Transactions on Software Engineering, vol. SE-4, no. 4, pp. 345-361.

Tsutsui, K., Suzuki, H., Shimoyoshi, O., Sonohara, M., Akagiri, K., Heddle, R. (1992). *ATRAC: Adaptive Transform Acoustic Coding for MiniDisc*, 93rd Audio Engineering Society Convention, San Francisco.

Unger, E.J. (2001). *Relating Initial Budget to Program Growth with Rayleigh and Weibull Models*, Master's Thesis, Air Force Institute of Technology.

Vint, P.F., Hinrichs, R.N. (1996). *Endpoint Error in Smoothing and Differentiating Raw Kinematic Data: An Evaluation of Four Popular Methods*, Journal of Biomechanics, vol. 29, no. 12, pp. 1637-1642.