Practical Polygraph: Seven Things to Know About Feature Extraction with Electrodermal and Cardio Data

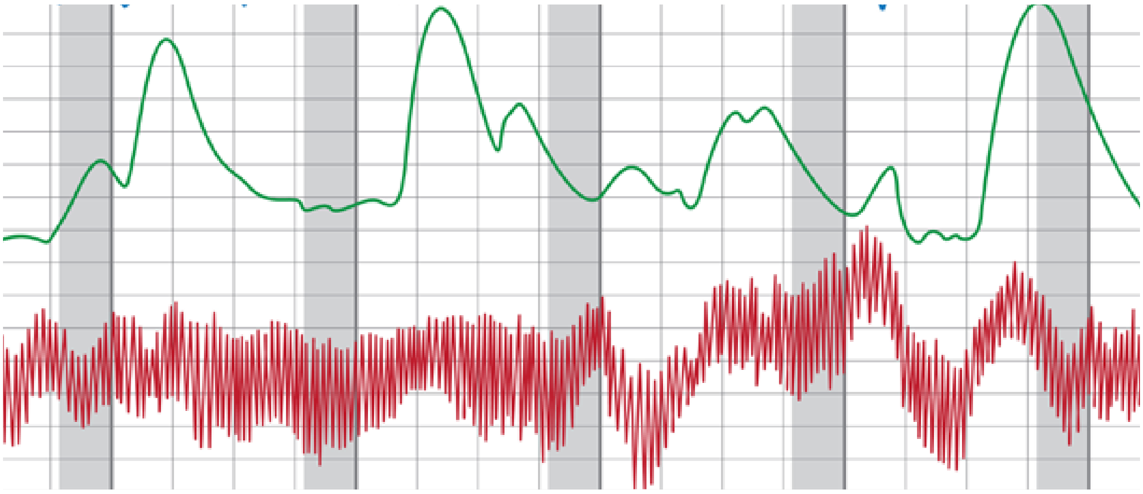

by Raymond Nelson and Mark Handler

This document describes seven different parameters that can affect the feature extraction with polygraphic electrodermal (EDA) and cardio activity data. Those parameters include: the response latency requirement, the response onset window and whether to interpret the response onset window in a strict or permissive manner, the evaluation window and whether to interpret the evaluation window in a strict or permissive manner, the imputing of a response onset when the data are already ascending at stimulus onset, and the interpretation of descending segments during the measurement period. Each of these parameters is an opportunity for inconsistency, differences of opinion, subjectivity and unreliability in field practice.

All scientific tests can be thought of as consisting of a common set of operations. These operations can include: feature extraction; numerical transformations and data reduction; a likelihood function in the form an empirical reference distribution, theoretical reference distribution, or other device for the computation of a statistical or probabilistic value for the observed data; and structured rules that determine the interpretation of the statistical and categorical test result. Bayesian analytic methods will also include a prior probability that will be computed together with the data and likelihood function. It is possible that some of these operations can occur together; for example, when dealing with big data (a term that refers to the use of large datasets) some data reduction may be completed prior to feature extraction in order to reduce redundancy of information within the data. Or, decision rules may specify which values, after numerical transformation and data reduction, are to be computed with the normative reference data or likelihood function.

Regardless of the exact design and organization of an analysis method, all test data analysis begins with feature extraction. Feature extraction is the process of identifying the useful information or signal of interest within the data, so that numerical scores can be obtained for analysis. For polygraphic EDA and cardio data, response amplitude is the primary signal of interest. Response amplitude has been described as a function of positive slope[1] activity in response to the test stimuli (Bell, Raskin, Honts & Kircher, 1999; Boucsein, 2012; Harris, Horner & McQuarrie, 2000; Kircher & Raskin, 1988; Kircher, Kristjannoson, Gardner & Webb, 2005; Podlesny & Truslow, 1993). Positive slope activity is easily observed visually by human experts, though there are a number of issues of ambiguity and potential inconsistency. Descriptions of feature extraction research and computer algorithm development can illustrate the complexity of the actual logic and process that experienced human experts can execute intuitively and virtually automatically, with little executive attention.

Boucsein (2012), Kircher and Raskin (1988) and Podlesny and Truslow (1993) described response onsets as the beginning of positive slope segment. They described response peaks as the highest point between onsets, and EDA and cardio response amplitude as the maximum difference between a low point and subsequent high point in the data within the evaluation window (Bell et.al., 1999; Kircher & Raskin, 1988). Similarly, Harris, Horner and McQuarrie (2000) also described response amplitude as a function of the difference between a peak of positive slope activity within the evaluation window and an onset of positive slope activity that occurred prior to the response peak and within a specified period for an expected response onset. In another similar description, Kircher et.al. (2005) described response onsets as changes from negative or zero slope to positive slope and response peaks as changes from positive slope to zero or nega-

1 The slope of a time-series data segment is said to be positive when it is moving upward, because the difference between each successive data point will be a positive number. A time-series data plot is said to be negative when moving in a downward direction because the difference between each data point will be a negative number. Time-series data that are moving horizontally – neither upward nor downward – can be said to have zero slope.

tive slope, with response amplitude defined as the maximum observed difference between a response onset and each subsequent response peak.

Although succinct, this definition has proven unsatisfactory when the slope of the EDA or cardio data is already positive at stimulus onset, when there is no change from negative or zero slope to positive slope, and when there are multiple changes in slope within the data segment of interest. With experience, many of the subtleties and nuances of feature extraction can be executed quickly and easily, almost automatically and with potentially little executive attention – giving rise to the possibility that some professionals may not be prepared to discuss the exact details of how they extract a particular score. Others may execute the feature extraction tasks inconsistently and may therefor be vulnerable to a variety of secondary influences. For these and other reasons, polygraphic test data analysis is a complex skill with numerous areas for potential ambiguity and subjectivity and inconsistency. The potential magnitude of inconsistency can be illustrated mathematically; although several of these parameters can have numerous possible solutions, with only two options for each parameter the result is that over 128 different feature extraction solutions exist (2^7 = 128) for the primary feature of interest when evaluating the EDA and cardio data. For this reason, understanding these various decision parameters – and their potential to foment errors and disagreement – should be of great interest to polygraph experts.

Response latency requirement

Response latency generally refers to the period from stimulus onset to response onset. The response latency period refers to the short period of time immediately following a stimulus onset during which a change in physiological activity is not interpreted. The rationale for a response latency requirement is that the speed of activity within the nervous system makes it unlikely that an immediate change in physiology is due to the test stimulus.

Boucsein (2012) describes a response latency of one to two seconds for EDA responses. Edelberg (1972) suggested a 1.2 to 4 second EDA response latency. Levinson and Edelberg (1985) listed all EDA latencies published in Psychophysiology between 1977 and 1982, and showed that 1 to 4 seconds and 1 to 5 seconds were the most often described.

In contrast to research psychophysiologists, field polygraph practitioners have typically used a shorter EDA response latency requirement. Bell et.al. (1999) described a response latency requirement of one-half second following the stimulus onset. Dutton (2000) also described the use of a one-half second response latency requirement for electrodermal reactions. Krapohl and Shaw (2015), citing Kircher and Raskin (1988), and also described the use of a one-half second response onset latency requirement. Figure 1 shows a change in EDA immediately upon stimulus onset followed by another change a few seconds later. Because manual scoring has continued to depend almost completely on visual feature extraction methods, it is likely that some field examiners are more attentive than others in their adherence to the response latency requirement. Also, no published description exists for a minimum required latency period for cardiovascular data.

Figure 1. Response latency for EDA data.

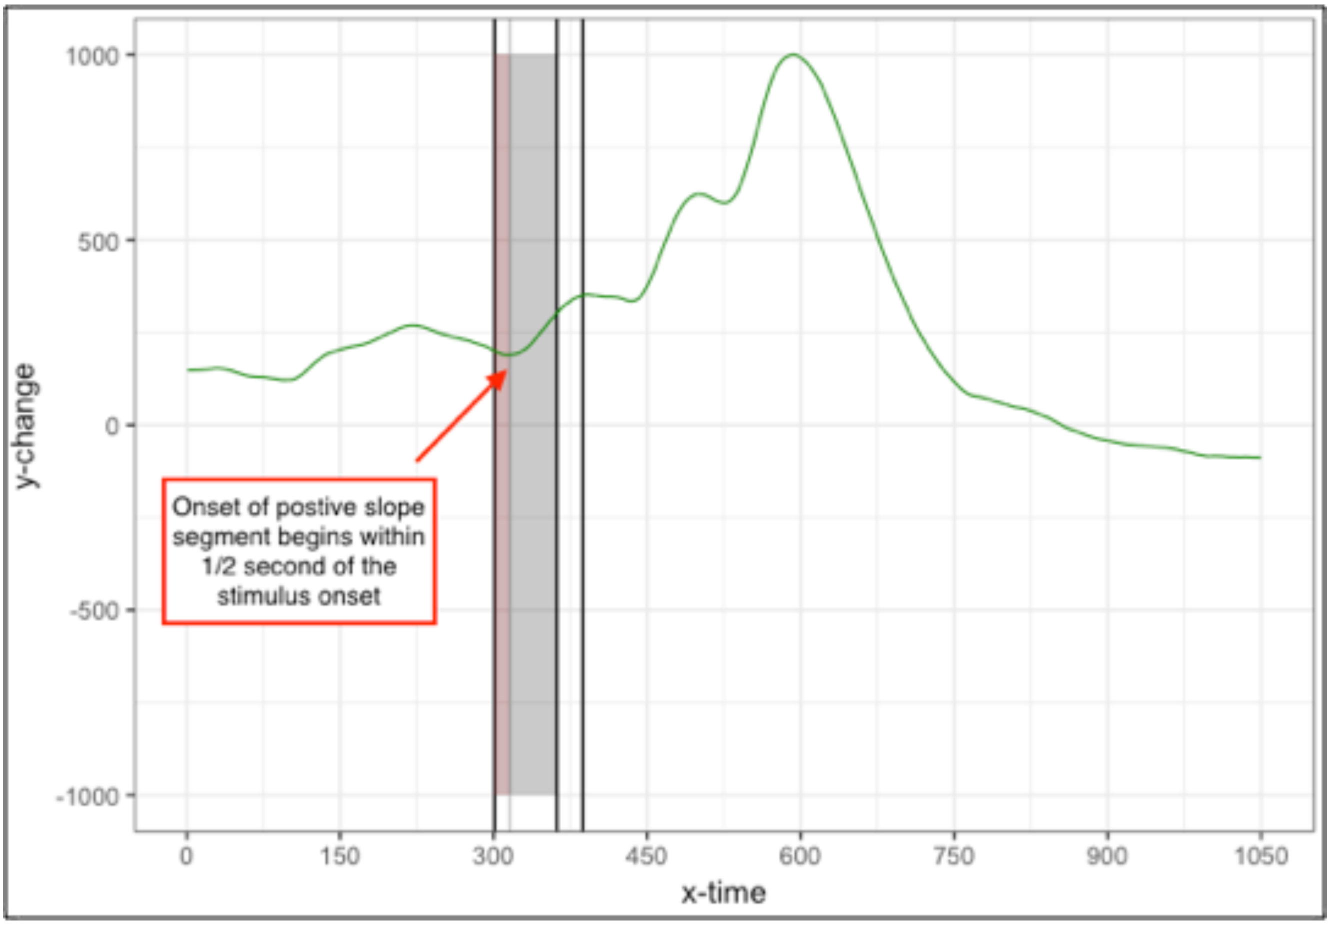

Response onset window

The response onset window and response latency are related. Changes in physiological activity can be attributed to the test stimuli when two requirements are satisfied: 1) the change in physiology is timely with the test stimulus, and 2) there is an absence of any other observable cause for the change in activity. Obviously, if there is any other noted possible cause then change in activity cannot be reliably attributed to the test stimulus. Changes in physiological activity can be interpreted as timely with the test stimulus if they begin within a defined response onset window. The response onset window should be consistent with reasonable assumptions about the ability of normal functioning persons to maintain undistracted attention.

Bell et.al., (1999) described the response onset window as existing from stimulus onset, following the minimum response latency requirement, until five seconds after the verbal answer. Figure 2 shows a cardio data segment with the response onset window shaded for five seconds after the verbal answer. Others have published different descriptions of the response onset window. Dutton (2000) described the response onset window as the period from stimulus onset to about 8 seconds. A slightly different solution was offered by the Department of Defense (2006) which defined the response onset window as existing from the stimulus onset to the verbal answer under normal circumstances. Harris, Horner and McQuarrie (2005) also described the exclusion of electrodermal and cardio response onsets that began more than one second after the verbal answer. Regardless of the details, the response onset window should not be confused with the evaluation window.

Figure 2. Cardio data with the response onset window shaded for five seconds after the verbal answer.

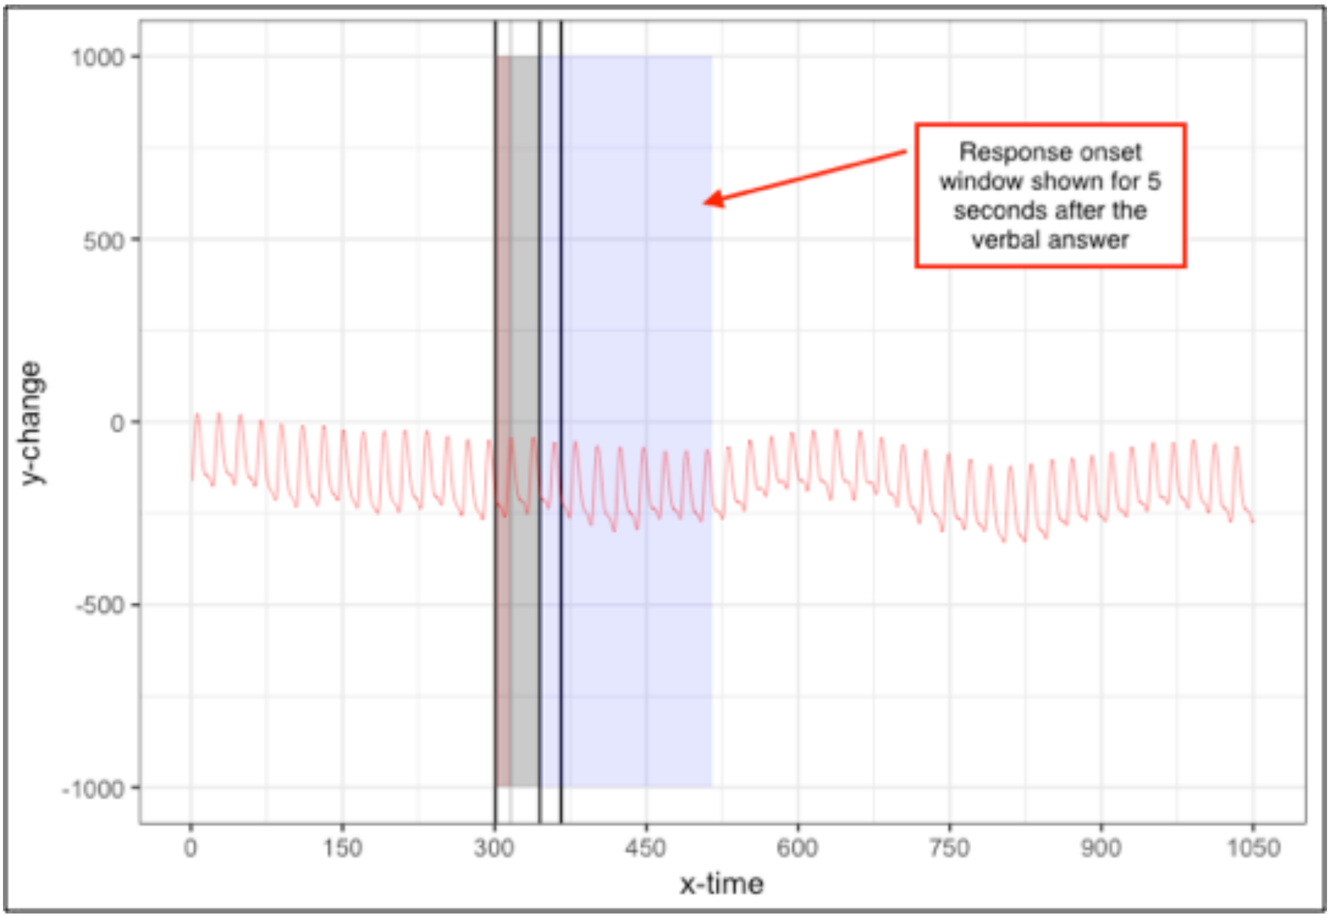

Evaluation window

The purpose of the evaluation window is similar to that of the response onset window – to ensure that changes in physiological activity can be reliably attributed to the test stimuli. For this reason, the length of time for the evaluation window should conform to reasonable knowledge about the ability of normal functioning persons to maintain undistracted attention and concentration. Different descriptions have been published regarding the length of the evaluation window.

Kircher and Raskin (1988) used a fixed evaluation window of 20 seconds beginning after stimulus onset and including a short latency period. Kircher et.al., (2005) also described the use of a 20 second evaluation window. The OSS-3 algorithm (Nelson, Krapohl & Handler, 2008) was developed using a 15 second evaluation window. Figure 3 shows a segment of EDA and cardio data with a 15 second evaluation window. Harris, Horner and McQuarrie (2005) described the selection of an EDA response peak in a 13 second window of data, and the selection of a cardio response peak in a 10 second window beginning at stimulus onset.

Figure 3. Evaluation window of 15 seconds.

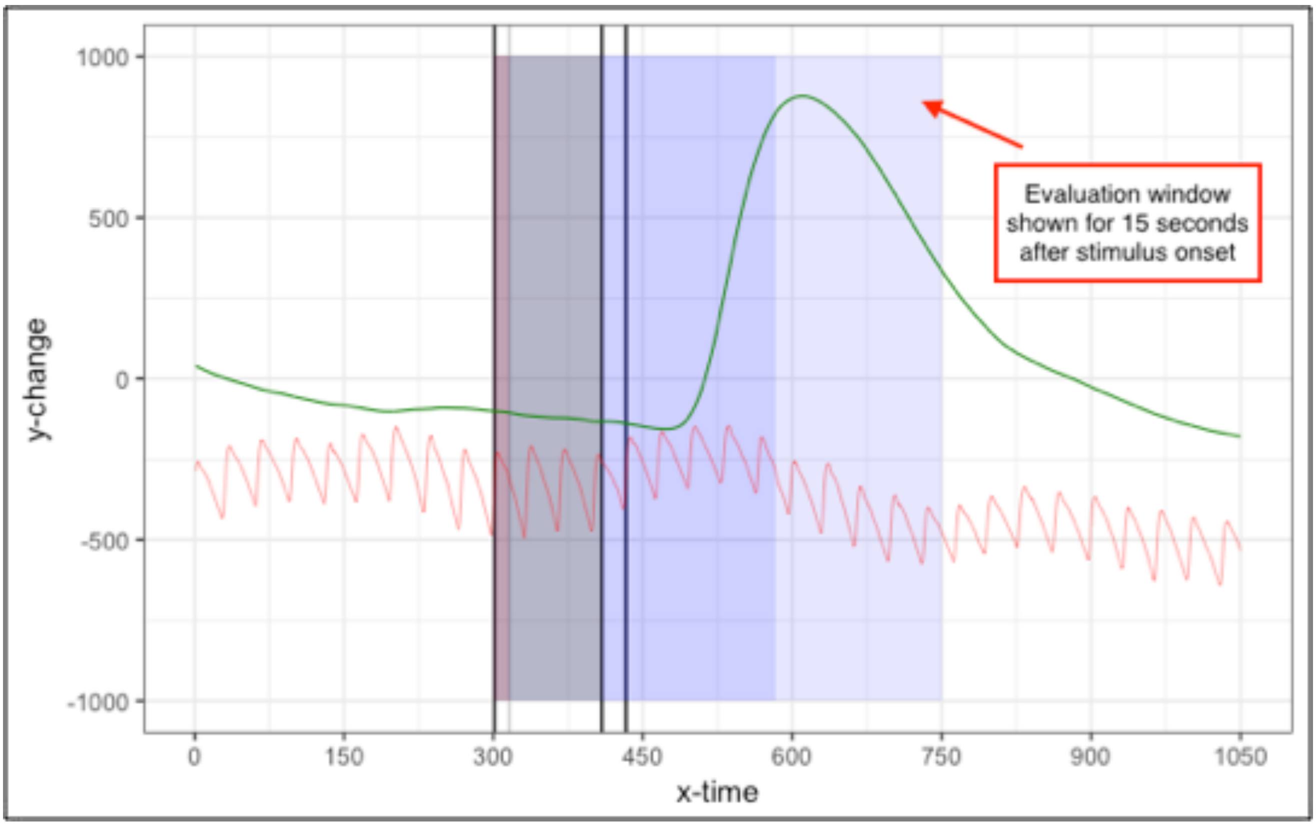

In contrast to a fixed length evaluation window, neither the ESS procedural reference (Nelson, et.al., 2011) nor the Federal Polygraph Examiners Handbook (Department of Defense, 2006) describe the use of a fixed evaluation window for EDA feature extraction – and instead indicate that changes in physiological activity are evaluated until the peak of response. Krapohl and Shaw (2015) reported the evaluation window as beginning after stimulus onset, and following a short latency period, and ending at the onset of the next stimulus.

When the length of the evaluation window is determined by the onset of the next stimulus event the effect can be that there is no fixed length for the evaluation window. This is because experienced field practitioners will generally not present each test stimulus at fixed intervals but will instead introduce each question while observing both the examinee and recorded test data to ensure readiness. Computerized scoring algorithms have commonly used a fixed-length evaluation window.

Strict or permissive interpretation of the response onset window

Use of a fixed-length evaluation window can introduce additional ambiguity to the feature extraction process when dealing with complex reactions. Simple reactions include only one positive slope segment. Simple reactions are also characterized by exactly two changes in slope: the onset of a positive slope segment following a negative or zero slope segment, and the peak of response or end of a positive slope segment. Complex reactions are those that consist of two or more positive slope segments, or three or more changes in slope.

A strict interpretation of the response onset window would extract information from only those positive slope segments that begin within the response onset window, excluding any positive slope segments from analysis if they begin after the response onset window. Figure 4 shows an EDA segment for which a second positive slope segment begins within the 15 second evaluation window yet outside the response onset window. A permissive interpretation of the response onset window would allow the extraction of information from all positive slope segments within the evaluation window as long as the onset of the first positive slope segment has occurred within the response onset window. A permissive interpretation of the response onset window may be reasonable when using a fixed-length evaluation window but may become problematic when using an evaluation window of undefined length – such as when the length of the evaluation window is determined by the onset of the next stimulus question.

Figure 4. EDA segment with a second positive slope change outside the response onset window.

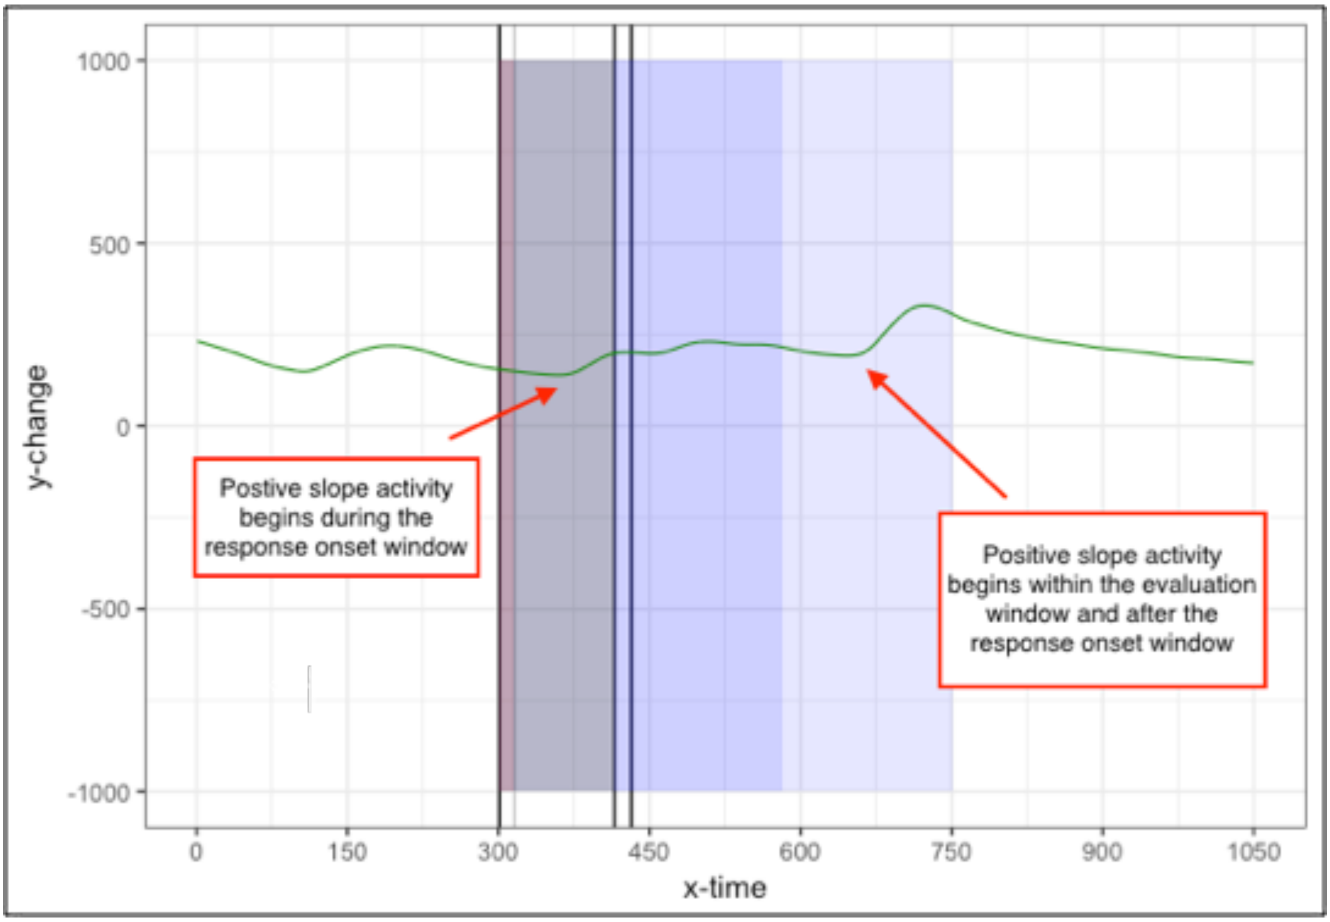

Strict or permissive interpretation of the evaluation window

Use of a fixed-length evaluation window can also introduce another option or ambiguity to the feature extraction process when a positive slope segment continues past the end of the fixed-length evaluation window. A strict interpretation of a fixed-length evaluation window would require the termination of feature extraction at the end of the fixed-length evaluation window. Figure 5 shows a cardio data segment that continues outside of the fix-length evaluation window. A permissive interpretation of the fixed-length evaluation window would permit the extraction of information until the peak of response – even if the response peak occurs after the end of the fixed-length evaluation window.

Figure 5. Cardio reaction that continues outside a fixed-length evaluation window.

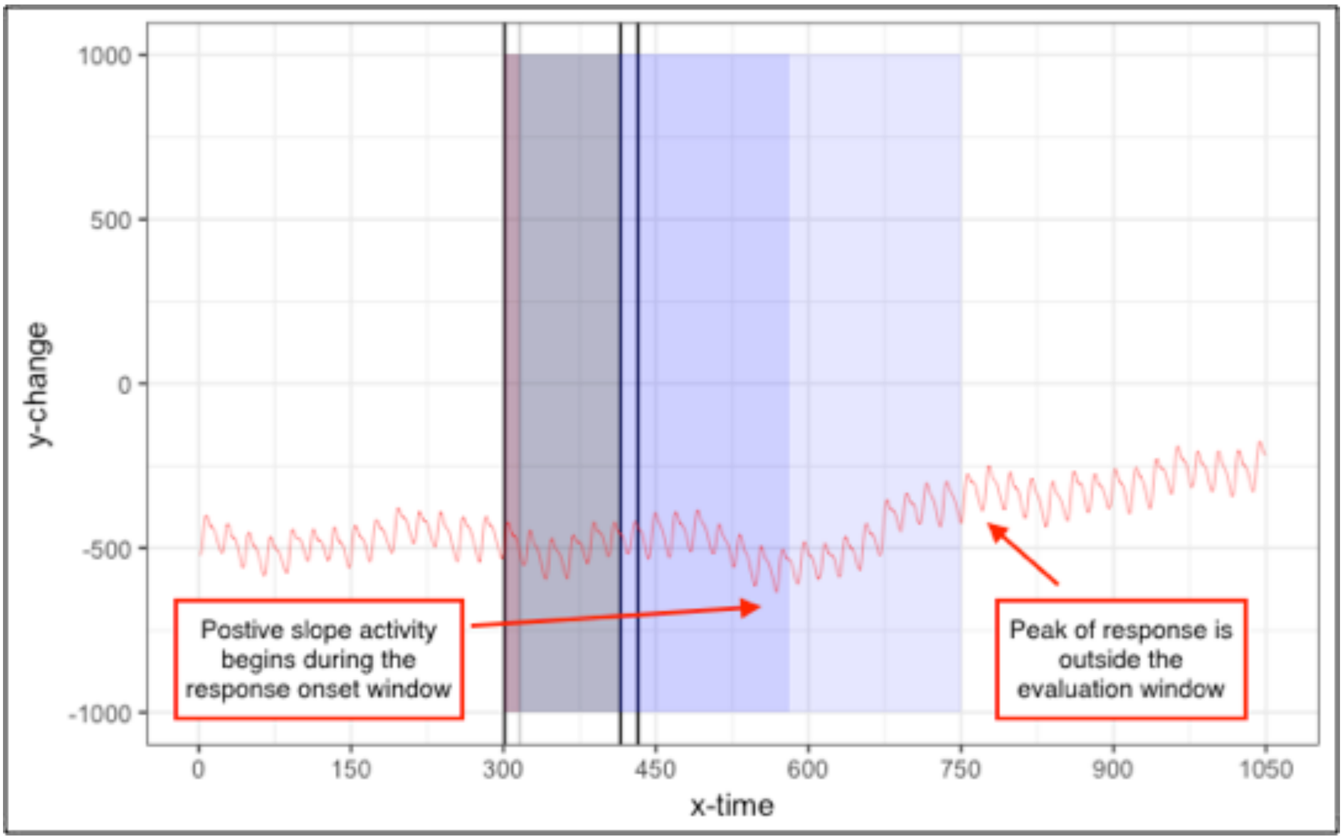

This matter of ambiguity could be eliminated by the use of an evaluation window that is determined by the onset of the next test stimulus question. However, an evaluation of window of undetermined length would impose a strict interpretation of the response onset window, such that some positive slope segments that begin after the response onset window yet within a fixed-length evaluation window might not be scored.

Interpretation of a response onset as a function of a change in positive slope angle

Positive slope segments that occur prior to stimulus onset are referred to as non-specific physiological responses[2] (Boucsein, 2102). They are often changes in physiological activity that cannot be attributed to the test stimuli. A positive slope prior to stimulus onset may also occur as a result of a positive slope tonic trend in the EDA or cardio data. If the onset of response is defined as the onset of a positive slope segment, then no feature extraction can occur when there is no positive slope segment that begins within the response onset window. Figure 6 shows an EDA data segment for which the slope is positive prior to stimulus onset. Automated feature extraction algorithms, because they are structured and procedural with no real knowledge or intuition about the data and context, have sometimes been unable to extract a response under these conditions.

Human experts (and perhaps some machine learning algorithms with "artificial intelligence") may have contextual knowledge about this the data and this potential condition and might therefor apply some creative intuition to these situations. In field practice, when the slope of the EDA and cardio data are already positive at stimulus onset many examiners will look for a change in positive slope angle within the response onset window and impute the onset of response at the point of change. Bell et.al., (1999) described this practice, and it can be often observed in use among field examiners.

Automated or computerized feature extraction algorithms can also be imbued with information or "knowledge" that can enable the feature extraction algorithm to impute the onset of a response in a manner similar to human experts. Harris, Horner and McQuarrie (2000) described the identification of a

2 Some have attempted to refer to these as "anticipated" or "early" reactions but these terms can be problematic because they can be misinterpreted as attributing the cause to the examinee's thinking about the forthcoming stimulus question. The cause of these reactions cannot in fact be known and it is possible their cause is unrelated to the test stimuli. It is known only that they have begun before the stimulus. The occurrence of numerous non-specific physiological reactions may be an indicator of problems with attention or cooperation during testing.

response onset within a positive slope segment as a function of a slope tolerance and the proportion of change relative to the peak of response with the goal of discriminating between positive slope activity that is tonic vs phasic. Other automated solutions also exist for this situation using statistical learning theory or machine learning methods.

When visually imputing a response onset within a positive slope data segment there may always exist some subjective differences in tolerance among human experts as to how much of a change in positive slope activity is usable or un-useable as a response onset. For automated feature extraction algorithms, the question of tolerance can be resolved in three ways: through the selection of arbitrary parameters, through heuristic observation and experimentation, or through statistical or machine optimization of the parameters that maximize a stated goal in terms of test accuracy or error rates.

Figure 6. EDA data with positive slope prior to stimulus onset.

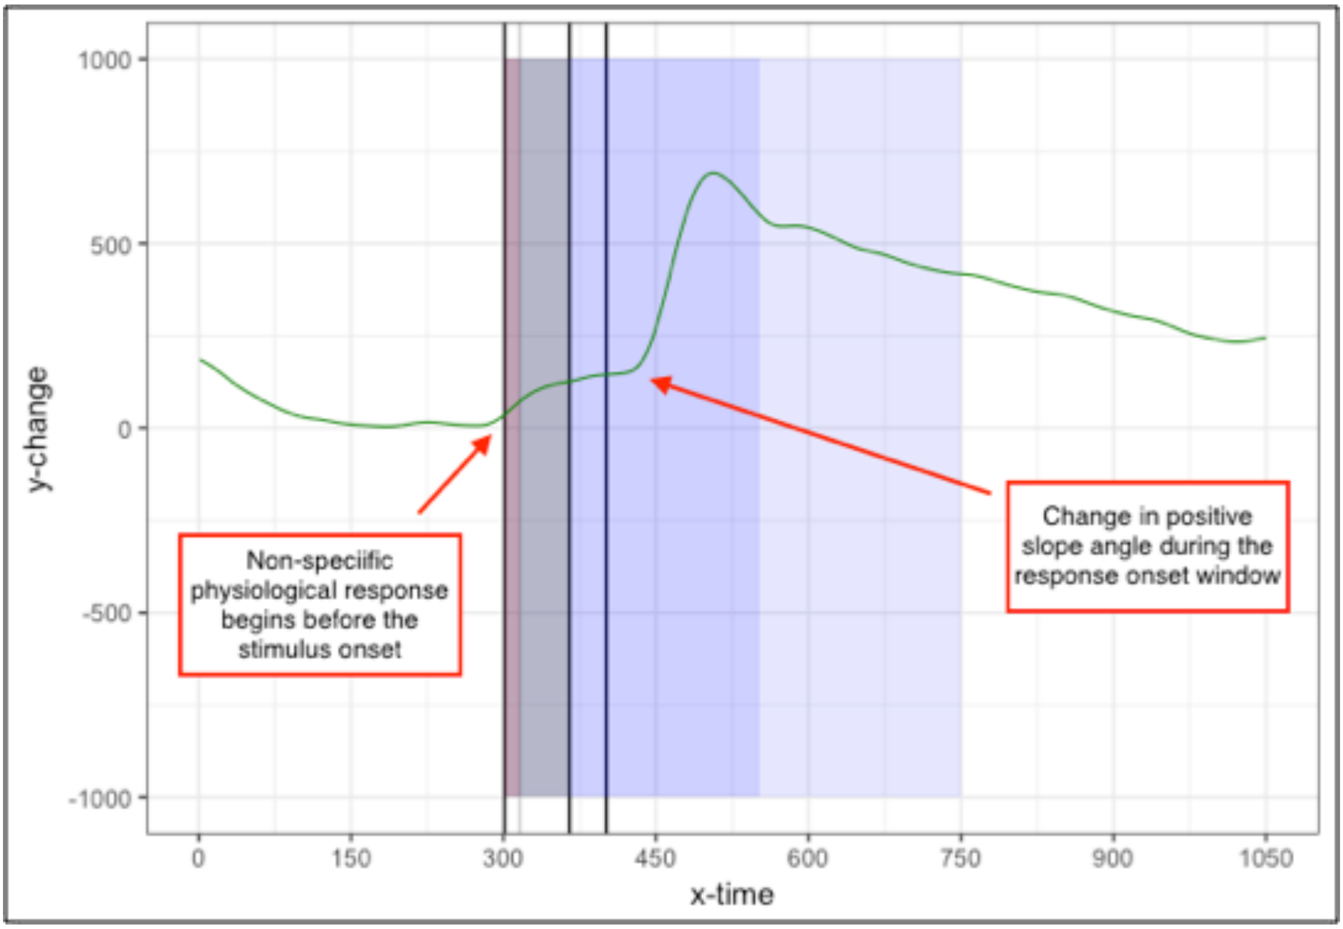

Interpretation of descending segments of complex reactions

A remaining issue of ambiguity in EDA and cardio feature extraction involves the interpretation of complex responses. As described earlier, complex responses consist of two or more positive slope segments or three or more changes in slope. Complex reactions are those that include at least one negative slope segment within two or more positive slope segments during the evaluation window.

Kircher and Raskin (1988), using the term "electrodermal burst frequency," reported a negative correlation (r = -.05) for response complexity and the criterion of deception and truth-telling. Harris, Horner and McQuarry (2000) reported that response complexity used alone was of little practical value, and further described that counting the number of peaks within the evaluation window was no better than the simple observation of multiple response peaks within the evaluation window. Kircher, et.al. (2005) also reported a weak negative correlation between response complexity and the criterion of deception and truth-telling (r = -.11).

Regardless of its weak diagnostic contribution, Bell, et.al., (1999) included response complexity as a scoring feature for EDA responses, as did the Department of Defense (2006). However, Krapohl and McManus (1999) did not include EDA complexity when scoring the Objective Scoring System (OSS), nor did and Krapohl (2000) with the OSS2. Similarly, Nelson Krapohl & Handler (2008) used only EDA and cardio amplitude when scoring the OSS-3. Nelson et.al, (2011) described the use of only EDA and cardio amplitude for the Empirical Scoring System. Both Harris, Horner and McQuarrie (2000) and Kircher et.al., (2005) reported that use of primary response features alone would produce information with the same diagnostic value as the traditional use of both primary and secondary features.

Although the evidence is weak and inconsistent regarding the diagnostic value of response complexity itself, response complexity has been described as a consideration when scoring response amplitude. According to the Department of Defense (2006), descending segments within the evaluation window are meaningful when they descend below the level of the response onset, under which condition the segment is not interpreted as a complex response even though there are multiple positive slope changes subsequent positive slope segment within the evaluation window is not interpreted as a complex reaction. In this case, a negative slope segment that descends below the initial response onset level is referred to as a "trough." Figure 7 shows a segment of EDA data that includes two positive slope sections for which the intervening negative slope section descends below the onset level of the first positive slope section.

Figure 7. EDA segment with negative slope that descends below the response onset level.

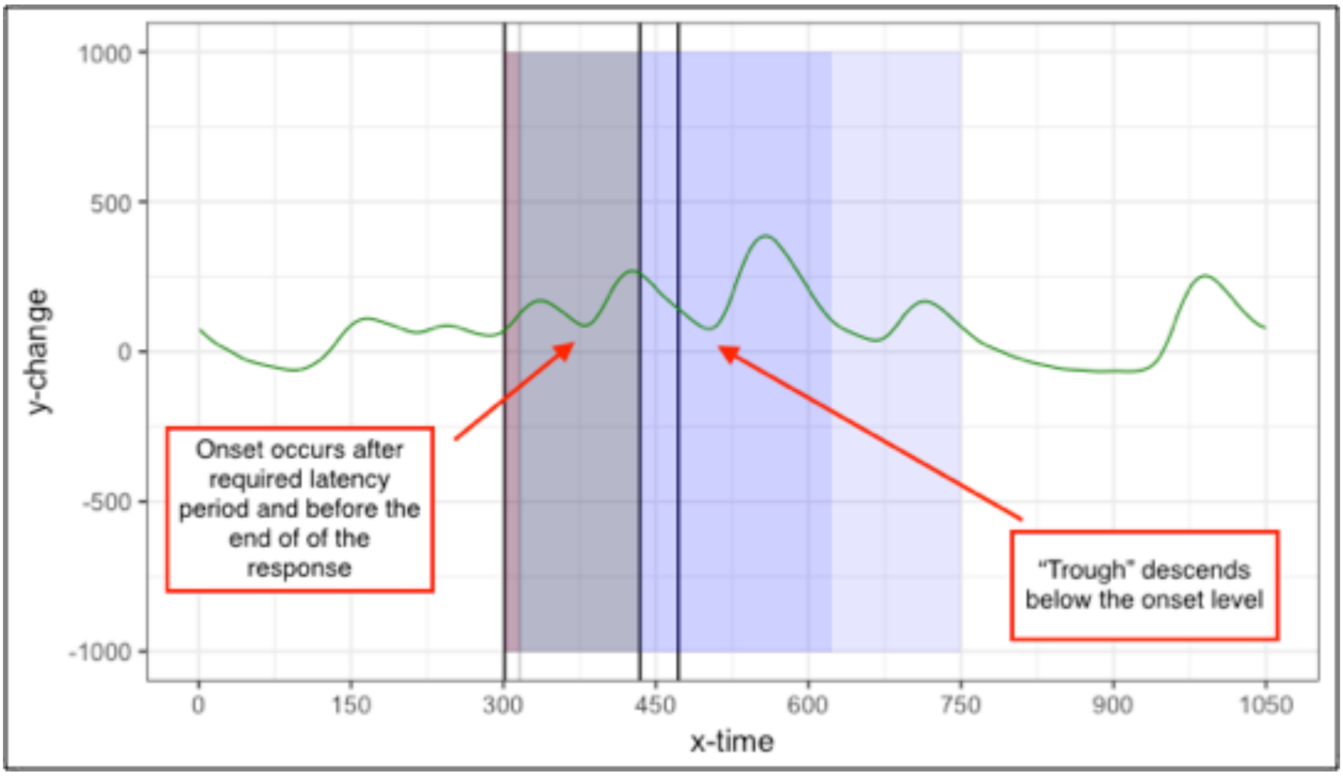

It is possible that different examiners can interpret the EDA segment in Figure 7 in different ways. If the "trough," is interpreted as indicating that the initial response is complete – with the corollary that the subsequent positive slope segment is not assumed to be caused by the test stimulus – then some will score only the first positive slope segment. Others might score the second positive slope segment – in the absence of any observable artifact or outside stimulus – because it begins within five seconds of the verbal answer and is larger than the first positive slope segment.

Another variant of this rule can be observed in discussion among field practitioners wherein a segment that would normally be interpreted as a complex reaction is viewed as two distinct reactions when the negative slope segment had descended more than 50% of the distance from peak of the initial positive slope segment and the response onset level (Boucsein, 2012). Figure 8 shows an EDA segment with two positive slope segments within the evaluation window for which the intervening negative slope segment descends about halfway from the peak to the onset level before the subsequent positive slope segment.

Figure 8. EDA segment with negative slope that does not descend below the response onset level.

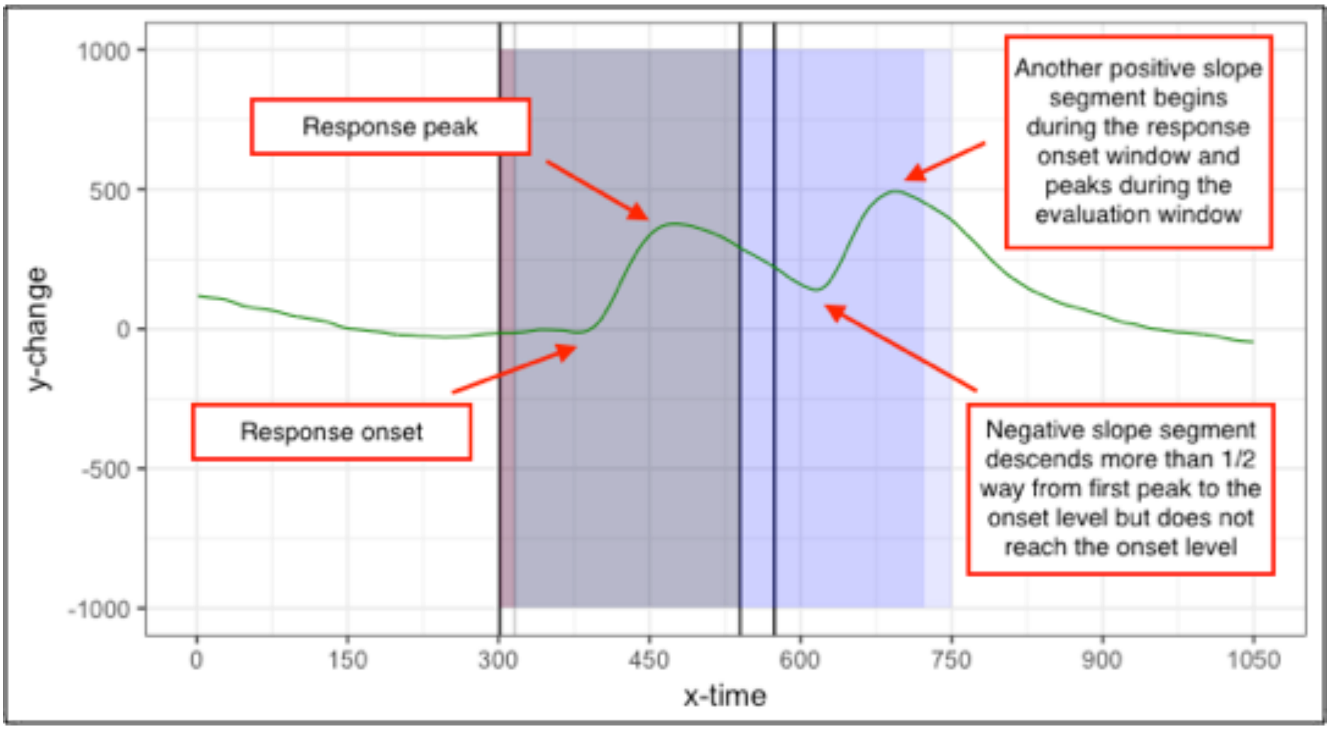

The variant descending rule shown in Figure 8, wherein the EDA data descend more than half-way from the peak to onset value, ignores that factors other than the examinee may influence these negative slope segments. Those factors can including environmental factors such as temperature, humidity and convection, which may affect evaporation and hydration at the surface of the skin, and other factors such as differing design characteristics of the Auto EDA filter for different polygraph instruments. Also, this variant rule is not extant in the authoritative publications.

References

Bell, B. G., Raskin, D. C., Honts, C. R. & Kircher, J.C. (1999). The Utah numerical scoring system. Polygraph, 28(1), 1-9.

Boucsein, W. (2012). Electrodermal activity(2nd ed.). New York: Springer.

Department of Defense (2006b). Psychophysiological Detection of Deception Analysis II -- Course #503. Test data analysis: DoDPI numerical evaluation scoring system. (Retrieved from http://www.antipolygraph.org/documents/federal-polygraph-handbook-02-10-2006.pdf on 3-31-2007).

Dutton, D. (2000). Guide for performing the objective scoring system. Polygraph, 29, 177-184.

Edelberg, R. (1972). Electrodermal activity of the skin: Its measurement and uses in psychophysiology. In N.S. Greenfield & A. Sternbach (Eds.), Handbook of Psychophysiology, 9, 512-520.

Harris, J., Horner, A. & McQuarrie, D. (2000). An evaluation of the criteria taught by the department of defense polygraph institute for interpreting polygraph examinations. Johns Hopkins University, Applied Physics Laboratory. SSD-POR-POR-00-7272.

Kircher, J. C. & Raskin, D.C. (1988). Human versus computerized evaluations of polygraph data in a laboratory setting. Journal of Applied Psychology, 73, 291-302.

Kircher, J.C., Kristjansson, S.D., Gardner, M.K. & Webb, A. (2005). Human and computer decision-making in the psychophysiological detection of deception.

University of Utah.

Krapohl, D. J. (2002). Short report: Update for the objective scoring system. Polygraph, 31, 298-302.

Krapohl, D. J. & Shaw, P, K. (2015). Fundamentals of Polygraph Practice. Elsevier.

Krapohl, D. & McManus, B. (1999). An objective method for manually scoring polygraph data. Polygraph, 28, 209-222.

Levinson, D.F. & Edelberg, R. (1985). Scoring criteria for response latency and habituation in electrodermal research: A critique. Psychophysiology, 22, 417-426.

Nelson, R., Handler, M., Shaw, P., Gougler, M., Blalock, B., Russell, C., Cushman, B. & Oelrich, M. (2011). Using the Empirical Scoring System. Polygraph, 40, 67-78.

Nelson, R., Krapohl, D. & Handler, M. (2008). Brute force comparison: A Monte Carlo study of the Objective Scoring System version 3 (OSS-3) and human polygraph scorers. Polygraph, 37, 185-215.

Podlesny, J. A. & Truslow, C.M. (1993). Validity of an expanded-issue (modified general question) polygraph technique in a simulated distributed-crime-roles context. Journal of Applied Psychology, 78, 788-797.

Raskin, D., Kircher, J. C., Honts, C. R. & Horowitz, S.W. (1988). A study of the validity of polygraph examinations in criminal investigations. Final Report, National Institute of Justice, Grant No. 85-IJ-CX-0040.