

20 May 2015

Dr. Harold Hawkins
ONR Code 341
Office of Naval Research
875 North Randolph SL
Arlington, VA 22203-1995

Reference: US Navy Contract N00014-12-C-0653: "The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation"
Charles River Analytics Contract No. C12186

Subject: Contractor's Quarterly Status Report #11
Reporting Period: 20-February-2015 to 19-May-2015

Dear Dr. Hawkins,

Please find enclosed 1 copy of the Quarterly Status Report for the referenced contract. Please feel free to contact me with any questions regarding this report or the status of the "The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation" effort.

Sincerely,



W. Scott Neal Reilly
Principal Investigator

cc: Cheryl Gonzales, DCMA
Annetta Burger, ONR
Whitney McCoy, Charles River Analytics

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 20 MAY 2015		2. REPORT TYPE		3. DATES COVERED 20-02-2015 to 19-05-2015	
4. TITLE AND SUBTITLE The Model Analyst's Toolkit:Scientific Model Development, Analysis, and Validation				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Charles River Analytics,,625 Mount Auburn Street,,Cambridge,,MA, 02138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 19	19a. NAME OF RESPONSIBLE PERSON
a REPORT unclassified	b ABSTRACT unclassified	c THIS PAGE unclassified			

Charles River Analytics

Monthly Technical Progress Report No. R12186-11

Reporting Period: February 20, 2015 to May 19, 2015

Government Contract No. N00014-12-C-0653

Charles River Analytics Contract No. C12186

The Model Analyst's Toolkit: Scientific Model Development, Analysis, and Validation Quarterly Status Report

Principal Investigator: Scott Neal Reilly

Charles River Analytics
625 Mount Auburn Street
Cambridge, MA 02138
617-491-3474

May 20, 2015

Distribution Statement A: Approved for public release distribution is unlimited.

The views, opinions, and findings contained in this report are those of the authors and should not be construed as an official Agency position, policy, or decision unless so designated by other official documentation.

1. Executive Summary

The proposed research effort builds on and extends the work of the previous ONR-funded “Validation Coverage Toolkit for HSCB Models” project. The overall objectives of the ongoing research program are:

- Help scientists create, analyze, refine, and validate rich scientific models
- Help computational scientists verify the correctness of their implementations of those models
- Help users of scientific models, including decision makers within the US Navy, to use those models correctly and with confidence
- Use a combination of human-driven data visualization and analysis, automated data analysis, and machine learning to leverage human expertise in model building with automated analyses of complex models against large datasets

Specific objectives for the current effort include:

- **Fluid temporal correlation analysis.** Our objective is to design a new method for performing temporally fluid correlation analysis for temporal sets of data and implement the method as a new prototype component within the Model Analyst’s Toolkit (MAT) software application.
- **Automated suggestions for model construction and refinement.** Our objective is to design and implement a prototype mechanism that learns from data how factors interact in non-trivial ways in scientific models.
- **Data validation and repair.** Our objective is to design and implement a prototype capability to identify likely errors in data based on anomalies relative to historic data and to use models of historic data to offer suggested repairs.
- **System prototyping.** Our objective is to incorporate all improvements into the MAT software application and make the resulting application available to the government and academic research community for use in scientific modeling projects.
- **Evaluation of applicability to multiple scientific domains.** Our objective is to ensure (and demonstrate) that MAT can be applied to a wide range of scientific domains by identifying and building at least one neurological and/or physiological model and analyze the associated data with MAT, making any extensions to the MAT tool that are needed to support the analysis of such a model.

2. Overview of Problem and Technical Approach

2.1. Summary of the Problem

One of the most powerful things scientists can do is to create models that describe the world around us. Models help scientists organize their theories and suggest additional experiments to run. Validated models also help others in more practical applications. For instance, in the hands of military decision makers, human social cultural behavior (HSCB) models can help predict instability and the socio-political effects of missions, whereas models of the human brain and

mind can help educators and trainers create curricula that more effectively improve the knowledge, skills, and abilities of their pupils.

While there are various software tools that are used by the scientific community to help them develop and analyze their models (e.g., Excel, R, Simulink, Matlab), they are largely so general in purpose (e.g., Excel, R) or so focused on computational models in particular (e.g., Simulink, Matlab), that they are not ideal for rapid model exploration or for use by non-computational scientists. They also largely ignore the problem of validating the models, especially when the models are positing causal claims as most interesting scientific models do. To address this gap, Charles River Analytics undertook the “Validation Coverage Toolkit for HSCB Models” project with ONR. Under this effort, we successfully designed, implemented, informally evaluated, and deployed a tool called the Model Analyst’s Toolkit (MAT), which focused on supporting social scientists to visualize and explore data, develop causal models, and validate those models against available data (Neal Reilly, 2010; Neal Reilly, Pfeffer, Barnett et al., 2011, 2010).

As part of the development of the MAT tool, we identified four important extensions to that research program that would further support the scientific modeling process:

- Correlation analyses are still the standard way of identifying relationships between factors in a model, but correlations are fundamentally flawed as a tool for analyzing potentially causal or predictive relationships as they assume instantaneous effects. Even performing correlation analyses with a temporal offsets between streams of data is insufficient as the temporal gap between the causal or predictive event and the following event may not be the same every time (either because of variability in the system being modeled or because of variability introduced by a fixed sampling rate). What we need is a novel way of evaluating the true predictive power across streams of data that can deal with fluid offsets between changes in one stream of data and follow events in the other stream of data.
- Modeling complex phenomena is a fundamentally difficult task. Human intuition and analysis is by far the most effective way of performing this task, but even humans can be overwhelmed by the complexity of modeling the systems they are studying (e.g., socio-political system, human neurophysiology). Automated tools, while not especially good at generating reasonable scientific hypotheses, *are* extremely good at processing large amounts of data. We believe there is an opportunity for computational systems to enhance human scientific inquiry. Under the “Validation Coverage Toolkit for HSCB Models” project, we demonstrated how automated tools could help human scientists to analyze and validate their models against data. We believe a similar approach can be used to help suggest modifications to the human-built models to make them better match the available data. To be useful, however, such automated analyses will need to be rich enough to suggest subtle data interactions that are most likely to be missed by the human scientist. For instance, correlations (especially correlations that take into account fluid temporal displacements) could be used to identify likely relationships between streams of data, but such an approach would miss complex, non-linear relationships between interrelated factors that cannot be effectively analyzed with

simple two-way correlations. For instance, if crime waves are associated with increases in unemployment or drops in the police presence, that would be hard to identify with a correlation analysis. We need richer automated data analysis techniques that can extract complex, non-linear, multi-variable relationships between data if we are to effectively suggest model improvements to human scientists.

- Even if a scientific model is sound, if the data sets provided as inputs to the model are unreliable, the results of the model are still suspect. And, unfortunately, data will often be wrong. For instance, HSCB surveys are notoriously unreliable and biased for a variety of reasons, and neurological and physiological data can be corrupted by broken or improperly used sensors. If it were possible to identify when data was unreliable and, ideally, even repair the data, then the models that are using the data could once again be effectively used.
- The MAT tool we developed under the “Validation Coverage Toolkit for HSCB Models” project was focused primarily on assisting social scientists in the analysis, refinement, and validation of HSCB models. In parallel with that effort, however, we also took an opportunity to apply MAT to evaluating neurological and physiological data under the DARPA-funded CRANIUM (Cognitive Readiness Agents for Neural Imaging and Understanding Models) program. We discovered the generality of the MAT tool makes it potentially applicable to a great number of different scientific domains. MAT proved to be a useful, but peripheral tool, in CRANIUM. We believe MAT could be applied to a broader suite of scientific modeling problems than it has been so far.

2.2. Summary of our Approach

To address these identified gaps and opportunities, we are extending MAT’s support for model development, analysis, refinement, and validation; enhancing MAT to analyze and repair data; and demonstrating MATs usefulness in additional scientific modeling domains. Our approach encompasses the following four areas, which correspond to the four gaps/opportunities identified in the previous section:

- **Temporally Fluid Correlation Analysis.** We are designing a new method to perform Temporally Fluid Correlational Analysis on temporal sets of data, and we are implementing the method as a new component within the MAT software application. The version of MAT at the beginning of the new effort supported correlation analysis for temporally offset data; it shifts the two data streams being compared by a fixed offset that is based on the sampling rate of the data (i.e., data that is sampled annually will be shifted by one year at a time), performs a standard correlation on the shifted data, plots the correlation value against the amount of the offset, and then repeats the process for the next offset amount. If two data streams are shifted by a fixed offset (e.g., changes in one stream are always followed by a comparable value in the other stream after a fixed time), then this method will find that offset. Under the current effort, we are expanding on this capability to support fluid temporal shifts within the data streams. That is, we are making it possible to identify when the temporal offset between the

change in the first data stream and its effect in the second stream is not a static amount of time.

- **Automated suggestions for model construction and refinement.** We are designing and implementing a mechanism to learn how factors interact in non-trivial ways in scientific models. In particular, we are developing a method for learning disjuncts, conjuncts, and negations. This mechanism starts with the model developed by the scientist user and make recommendations for possible adjustments to make it more complete by performing statistical data mining and machine learning.
- **Data validation and repair.** Recognizing that data contains errors is plausible once we understand the relationships between data sets. That is, if we are able to develop models of the correlations between sets of data, then we can build systems that notice when these correlations do not hold in new data, indicating possible errors in data. For instance, if we know that public sentiment tends to vary similarly between nearby towns, then when one town shows anomalous behavior, we can reasonably suspect problems with the data. There might be local issues that cause the anomaly, but it is, at least, worth noting and bringing to the attention of the user of the data and model. As MAT is designed to help analyze models and recognize inter-data relationships, it is primed to perform exactly this analysis. Existing methods perform similar types of analysis for environmental data (Dereszynski & Dietterich, 2007, 2011) . For instance, a broken thermometer can be identified and the data from it even estimated by looking at the temperature readings of nearby thermometers, which will generally be highly correlated.
- **Application to multiple scientific modeling domains.** To ensure (and demonstrate) that MAT can be applied to a wide range of scientific domains, we are identifying and building at least one neurological and/or physiological model and analyzing the associated data with MAT, making any extensions to the MAT tool that are needed to support the analysis of such a model. The initial MAT effort focused on HSCB models; by focusing this effort on harder-science models at much shorter time durations, we believe we can effectively evaluate an interesting range of applications of the MAT tool.

3. Current Activities and Status

During the current reporting period, we focused primarily on improving the causal analysis functionality of MAT. This included three subtasks: using the MAT tools for a real-world analysis to demonstrate and evaluate the various causal analysis methods, adding a new causal analysis method that we hope will provide better performance for certain kinds of analysis in the face of significant noise in the data, and create a centralized method for running and summarizing the results of the various analysis methods. These efforts are described in Sections 3.1, 3.2, and 3.3. We also made initial progress on data validation, which is described in Section 3.4. Finally, we made a number of other improvements to the usability and efficiency of the software that are summarized in Section 3.5. Our on-going transition and marketing efforts are described in Section 5.

3.1. Causality Analysis Use Case Development and Demonstration

During the current reporting period, we developed an in-depth, real-world demonstration use case of MAT as a causal-analysis and modeling tool, which we included in our AHFE paper and which we hope will provide a sound basis for ongoing evaluations and demonstrations. In the case study, we demonstrate a representative exploration of the causal/predictive relationship between poverty and conflict. A large body of literature exists that explores the “conflict trap”—the process whereby countries get stuck in a repeated pattern of violent conflict and economic underdevelopment (Collier et al., 2003). There have been several studies evaluating the causal/predictive link between these two features using standard statistical approaches, with some finding evidence for poverty driving societies into conflict (Collier & Hoeffler 2004, Braithwaite 2014), while others (Djankov 2008) indicate that civil conflict may be the cause of depressed economic growth. Using the methods described in the previous section, we can better untangle and characterize this relationship and gain insight into the processes that lead to the conflict trap.

The choice of data is itself a challenge for causal/predictive analyses, as the complex and abstract concepts of “poverty” and “conflict” are difficult to represent as measurable variables. To measure conflict we use the UCDP/PRIO dataset (Themnér & Wallensteen 2014), which tracks the incidence and intensity of global armed conflict between 1946 and 2013. To capture the notion of poverty, which is not merely a measure of income, but also of relative well-being, we use two variables from the World Bank World Development Indicators dataset (The World Bank 2013)—infant mortality rate, measured as the number of infants per thousand live births that die each year, and GDP, to measure the overall level of development. We consider conflict as both a categorical variable ranging from 0 to 3 indicating the intensity of a conflict in a given year, and as a numerical value with counts of the battle deaths due to conflict within a country. We focused on the timeframe from 1960-2013 as both data sets were more complete for this time period.

In our first experiment, we analyzed the relationship between poverty and conflict using Granger causality, varying the time lag between 1 and 10 years. Out of the 100 countries under study, we found strong evidence that conflict causes poverty in about 30% of the cases with a time lag of 1 year, as shown in Figure 1, with strength of the causal linkage degrading slightly as the time lag increased. Interestingly, there is also strong evidence of a causal relationship from poverty to conflict, but this actually consistently *increases* as we stretch out the time lag. This result may indicate the nature of conflict and poverty as persistent conditions with longer duration impacts, but may also be due to uneven time lags that cannot adequately be captured by Granger causality.




Figure 1. Results from Granger causality analysis with increasing time lag.

For our second experiment, we used convergent cross mapping (CCM) to further characterize the causal relationship between conflict and poverty. Social processes are often best described by complex dynamical systems, with multiple layers of feedback and interaction, and CCM can help identify these more complex causal interactions, particularly reciprocal or bidirectional causality. Because CCM examines the relationships between projections of the time series, we normalized the data to measure the percent change at each time point to account for the vastly different scales of conflict casualties, infant mortality, and GDP. We observed convergence in 80% of countries supporting the hypothesis that conflict causes poverty, and 12% for poverty leading to conflict. However, these results may be skewed by the perfect predictability of conflict in countries that experienced no conflict during the time period under study. Figure 2 shows an example of convergence to support the hypothesis that conflict causes poverty in Comoros.




Figure 2. Results from CCM analysis illustrating convergence indicative of a causal link from conflict to poverty.

While dynamic time warping and convergent cross-mapping can be useful analytic tools, the nature of our case study data is not ideal for these types of methods that look explicitly for point-by-point relationships across the time series. However, even though the relationships between the conflict and poverty data are difficult to quantify through these types of measurements, we found they can be reasonably described through qualitative featurization analysis. While conflict and poverty are linked to one another, this phenomenon does not manifest as similar patterns of proportional increases or decreases in values offset in time. Instead, across the countries studied, we saw that rapid increases in conflict or periods of recurring conflict are associated not with similar fluctuations in poverty, but by continually decreasing or statically depressed levels of economic activity and by statically high levels of infant mortality. Similarly, we found that when the conflict ended, we saw decreases in poverty follow. In essence, this illustrates the notion of the conflict and poverty traps, where violence is associated not with rapid declines into poverty, but with sustained levels of minimal development.

Figure 3 shows an example using the qualitative feature-based approach to analyze the data from Senegal. The top plot indicates GDP in current US\$ from 1989 to 2013, the middle plot shows the number of battle related deaths, and the bottom chart is the infant mortality rate. The human-guided qualitative featurization algorithm has divided these data series into important component pieces representing distinct features. From these features, it is evident that there was a period of violent conflict from 1989 to about 2004, with several spikes in the number of casualties. During this same time, infant mortality was consistently high, and GDP was

consistently low. However, after 2004, GDP and infant mortality both begin to steadily improve, while conflict remains very limited. Using these features to represent concepts such as “spikes in conflict” and “high infant mortality,” we can identify causal patterns between these more complex features that are not visible when doing a lower-level comparison of individual time points.




Figure 3. Fig. 1. Qualitative featurization showing the relationships between poverty and conflict.

3.2. Integrated Causal Analysis Report

We are developing an integrated causality analysis reporting feature. Previously users of MAT would have to use different individual analytical techniques to formulate their own specialized view of indicators which could suggest a causal relationship between one or more data series. The new functionality combines multiple causality-related analytical tool calculations into a single report. Here is a mock-up of the report that is currently being developed:




Figure 4. Mockup integrated causality analysis interface design.

To access the new report, the user selects two series in the MAT Data View and right clicks to access the normal context menu, which now has a new choice "Causality Analysis Report." Activating the report runs the analysis and pops up the report in a dialog box. The analysis automates various parameterizations which the user would normally have to perform manually. For example, the analysis examines different temporal offsets (e.g., for Granger Causality analysis) and determines which is most likely, then uses causality metrics at those particular offsets to determine the degree of causality.

3.3. New Causal Analysis Method: Beta Process Predictor Analytical Method

We have developed an additional causality analysis methodology provisionally called a *beta process predictor*. The idea behind this method is to model causal distributions as beta distributions. Most causal or statistical methods make a Gaussian assumption about the distribution of data, which is reasonable if each event is fully independent. In practice, however, real data tends to follow beta distributions because individual data events are not independent of each other.

The method works by examining for each effect-hypothesis data point every possible percent change in the data at different offsets leading up to that point. The best match is added as the preferred offset for that point. For example, if the effect shows a percent change of +15% and there are three causal points at offsets of 1 month previously, 2 months, and 3 months with values of +5%, +17% and +25%, then an offset of 2 months will be chosen (because the difference between 15 and 17 is the least among the three candidates). This results in an *offset distribution*. For example, if we repeat this process for 1000 points in the effect dataset, then we might have 43 points at -5 months, 117 points at -4 months, 175 points at -3 months, 511 points at -2 months, and 223 points at -1 months and smaller numbers beyond -6 months back. This offset distribution is fitted to a beta distribution which yields an alpha and beta parameter which fully characterize a beta distribution. We can then measure the peakedness, or kurtosis of the resulting distribution by the equation:

$$\frac{6[\alpha^3 + \alpha^2(1 - 2\beta) + \beta^2(1 + \beta) - 2\alpha\beta(2 + \beta)]}{\alpha\beta(\alpha + \beta + 2)(\alpha + \beta + 3)}$$

The sharper this peak, the more likely the causal relationship is taken to be and vice versa. If there is no causal relationship, we would expect the distribution to be flat (kurtosis = 0). As a final step we perform this analysis both ways, for $A \rightarrow B$ and $B \rightarrow A$. The ratio of the higher value over the lower is taken to be the degree of causality.

We have tested this metric against synthetic data by creating fully causal series and injecting differing amounts of random noise into data, then testing to see whether the metric can detect the causality. So far the results of these tests have been very encouraging, so we are developing the metric further and including it in our latest release.

3.4. Data Validation

We began work this period on data validation. Data validation will support data analysts who are concerned that they are developing models with unreliable data or who are developing models to help detect when a system breaks. For instance, if we can develop a model for how survey data results from various regions tend to correspond to each other, then we can develop tools that can flag data that appears to violate those historic patterns, and so might not be appropriate for use in decision making or model building.

The basic idea for our approach is similar to that described in Dereszynski and Dietterich (2007, 2011) in the domain of environmental science. That is, we create a graphical, probabilistic model from “good” data that models and learns how various data sources relate to each other. For MAT, we plan to have the use identify the structure (that is, which variable are likely to related) and to use machine learning techniques to learn the patterns (traditional multi-linear regression would be acceptable in many cases, but we also happen to have access to richer probabilistic machine learning tools in house that let us build more sophisticated models if needed). We use the resulting model to identify (and suggest repairs for) anomalies in a dataset that is not known ahead of time to be “good.” We currently assume that a domain expert user is able to recognize “good” data from which to learn.


We have developed three simple scenarios/data sets for exploring this concept. The first is learning to detect when a thermostat or temperature sensor are faulty by learning how the actual temperature and thermostat setting relate to each other and then detecting inconsistencies. The second is similar, but with a second thermostat (e.g., a second zone in the same house), where we expect there to be a link between the two thermostats/temperatures, but there is a less directly causal link. Third, we have developed a survey scenario with three villages where the survey taker in one starts filling in random data at some point.

During the current period, we used the Figaro open source probabilistic programming language to develop and learn a model for the first scenario with promising results. During the next period, we plan to incorporate this into MAT more directly and to focus on the second and third scenarios.


3.5. Software Improvements

3.5.1. Improvements in Handling Missing Values

In order to provide a clearer picture to the user, it is important to provide information about what data was used in an analysis. It is common for data to have missing values that could influence the results, so improvements were made to show how many data points were used in an analysis. Furthermore, some analyses have a temporal offset as input which can change the number of values used for a calculation. For example, when performing the correlation of two time series with a temporal offset, the following screenshot lets the user know that with no temporal offset that nine points were used in the correlation calculation and there was one missing value.



By increasing the temporal offset to one, we now have two missing values.



3.5.2. Improvements in Handling Large Datasets

In order to increase rendering speed so that large datasets can be displayed in MAT, the newest stable release of the graphics library used by MAT (Processing) was introduced. While exact

performance depends on the machine's speed, improvements allowed for better handling of datasets with millions of data points. This will help make MAT useful in a broader array of domains. Also, additional testing with large datasets has made MAT more robust by expanding the limits of what the application can handle.

4. Planned Activities

During the upcoming reporting period, we plan to focus on the following tasks:

- Completing our data validation component and integrating it into MAT.
- Completing work on the integrated causality report.
- Completing the port of Convergent Cross Mapping from Matlab to MAT.
- Completing our speed-up improvements to MAT.
- Improving the feature extraction algorithms to provide better, faster results.
- Completing all deliverables, including the MAT software and final report.
- Preparing for the briefing on MAT in Arlington in June.
- Transition and marketing efforts, including presenting MAT at the AHFE conference in July and submitting drafts of our book chapter.
- Continuing to support MAT users, both internal and external.

5. Evaluation and Transition

We continue to focus on making MAT available to the government and academic research communities and to look for opportunities to use MAT on a variety of ongoing research efforts.

During the current reporting period, we submitted our final version of "Tools for Validating Causal and Predictive Claims in Social Science Models," which we will be presenting at the 6th International Conference on Applied Human Factors and Ergonomics (AHFE 2015) in July.

We have also been invited to submit a chapter to a forthcoming book, "Modeling sociocultural influences on decision making," edited by Denise Nicholson, CDR Joseph Cohn, LT David Combs, and Sae Schatz. Our chapter will be on using MAT to help validate social science models and will include representative use cases of planning for the continuing growth and associated challenges of megacities.

As reported previously, we have also used the explorations into causal analysis and validation done under MAT as the basis for seedling pitches to DARPA (Steve Jameson) and IARPA (Steve Rieber), both of whom we have spoken to and have expressed initial interest in the MAT work and pursuing follow-on ideas. The DARPA effort has progressed to the point of contract negotiations and we hope to begin work on that effort in the next week of two. The IARPA effort is still in discussions. Dr. Rieber gave us some guidance and feedback on an initial white paper and we are reworking the white paper based on that feedback.

Table 1 summarizes our transition progress to date. We will continue to update this table as we make additional progress and will include it as a regular part of future status reports.

Table 1. MAT Transition and Use Progress

Program	Customer	Comments
On-going efforts		
Tourniquet Master Trainer (TMT) (Phase II SBIR)	US Army's Telemedicine & Advanced Technology Research Center (TATRC)	MAT is being used to visualize and analyze data from sensors on a medical manikin that indicate whether a number of novel medical devices used to combat junctional and inguinal hemorrhaging are being applied properly. This is an ongoing program.
Laparoscopic Surgery Training System (LASTS) (Phase II SBIR)	US Navy's Office of Naval Research (ONR)	Under lasts, Charles River and Caroline Cao at Wright State University are using MAT to analyze data collected from the location of the laproscopic surgery tools tools during an experiment. Surgical tools are instrumented with markers and 3D data is collected on their location as the person performs the task. This is a now-completed program.
Cognitive Readiness Agents for Neural Imaging and Understanding Models (CRANIUM) (Phase I SBIR)	US Navy's Office of Naval Research (ONR)	MAT was used to visualize and extract patterns of stress and workload from neuro-physiological data for training systems. This was a Phase I SBIR program that did not progress to Phase II.

Business Intelligence Visualization for Organizational Understanding, Analysis, and Collaboration (BIVOUAC) Phase II SBIR	US Navy's Space and Naval Warfare Systems Command (SPAWAR)	MAT is being evaluated as part of the BIVOUAC SBIR program, which provides data analysis and visualization for Enterprise Resource Planning (ERP) systems for the Navy. This is an ongoing Phase II SBIR program.
Adaptive toolkit for the Assessment and augmentation of Performance by Teams in Real time (ADAPTER) (Phase I SBIR)	US Air Force Research Lab Human Effectiveness Directorate (AFRL/RH)	MAT is being used to analyze neuro-physiological data from cyber operators to evaluate cognitive workload during team-based cyber operations. This is an ongoing Phase II SBIR program.
Anticipated Efforts		
A system for augmenting training by Monitoring, Extracting, and Decoding Indicators of Cognitive Load (MEDIC)	US Army's Telemedicine & Advanced Technology Research Center (TATRC)	We are evaluating the practicability of using MAT to analyze and visualize neuro-physiological data from combat medic trainees to identify periods of stress and cognitive overload. This is an active Phase II SBIR program where MAT is being considered for data analysis.
Causality Analysis Seedling	DARPA I2O through AFRL/RI	We are developing datasets for training and evaluating causal analysis methods and evaluating ensemble approaches to causal analysis of temporal data. This contract is currently being negotiated.

In addition we have provided copies of MAT to the following institutions based on their requests for the software: the University of Michigan, Arizona State University, Kansas State University, University of California at Los Angeles, the Naval Medical Research Unit at Wright Patterson Air Force Base, Concordia University (Montreal), the University of Wisconsin, the University of Maryland, and the Air Force Research Laboratory's Human Effectiveness Directorate, the Intelligence Advanced Research Projects Agency (IARPA), and the Joint Advanced Warfighting Division (JAWD).

6. Budget and Project Tracking

As of April 30, 2015, we have spent \$808,114, or 87% of our total budget of \$928,224, in 87% of the scheduled time. Our current funding is \$928,224, so we have been allocated 100% of our available funding.

Overall, we believe we are in good shape to complete the project on time and on budget.

7. References

- Braithwaite, Alex, Niheer Dasandi, and David Hudson (2014). "Does poverty cause conflict? Isolating the causal origins of the conflict trap." *Conflict Management and Peace Science*.
- Collier, Paul; Elliott, V. L.; Hegre, Håvard; Hoeffler, Anke; Reynal-Querol, Marta; Sambanis, Nicholas. (2003). *Breaking the Conflict Trap: Civil War and Development Policy*. Washington, DC: World Bank and Oxford University Press.
- Collier, P. and Hoeffler, A. (2004). "Greed and Grievances in Civil Wars." *Oxford Economic Papers* 56.
- Dereszynski, E. W.& Dietterich, T. G. (2011). Spatiotemporal models for data-anomaly detection in dynamic environmental monitoring campaigns. *ACM Transactions on Sensor Networks (TOSN)*, 8.
- Dereszynski, E.& Dietterich, T. (2007). Probabilistic models for anomaly detection in remote sensor data streams.
- Djankov, S. and Reynal-Querol, M. (2008). "Poverty and Civil War: Revisiting the evidence", CEPR No. 6980.
- Neal Reilly, S. (2010). *Validation Coverage Toolkit for HSCB Models*. Charles River Analytics Inc.
- Neal Reilly, W. S., Pfeffer, A., and Barnett, J. (2010). A Metamodel Description Language for HSCB Modeling. In *Advances in Cross-Cultural Decision Making*.
- Neal Reilly, W. S., Pfeffer, A., Barnett, J., Chamberlain, J., and Casstevens, R. (2011). A Computational Toolset for Socio-Cultural Data Exploration, Model Refinement, and Model Validation. In *Human Social Culture Behavior (HSCB) Focus*.
- The World Bank. *World Development Indicators*. 2013.
- Themnér, Lotta & Peter Wallensteen (2014) *Armed Conflict, 1946-2013*. *Journal of Peace Research* 51(4).

