Optimizing Business Intelligence using the OODA Framework

Morten Middelfart *

* *TARGIT A/S, Aalborgvej 94, 9800 Hjørring, Denmark (e-mail: morton@targit.com).*

Abstract: Business Intelligence is applied to improve decision making in organizations, and as globalized trading and connectivity increase, the pace and unpredictability of the business environment increase as well. Therefore traditional methods of user-driven data discovery will be too slow, and traditional long-term forecasting methods will be unreliable since the environment is behaving chaotically. For business users and their organizations this means that the ability to act swiftly based upon changes in the environment is the key determining factor for success and failure. The ability to react fast to changes in the environment can be achieved by pursuing two objectives: First, the speed with which a user travels through the four phases of Observation, Orientation, Decision and Action (OODA) can be increased, and secondly the time-horizon for the warnings can be expanded by applying so-called sentinel rules. In this article a research project that pursues each of these two objectives is described in terms of motivation, research methods, and expected results. Upon completion the project is expected to be a valuable contribution of concrete solutions that can improve the "horizon" and speed of an OODA cycle for a user of a Business Intelligence system. In addition, it is expected to form a good research foundation that others can use in their research going forward.

Keywords: Business Intelligence, OLAP, Management systems, Man/machine interaction

1. INTRODUCTION

For the past few decades, a lot of focus has been put on the immediate data warehousing challenges of obtaining and storing the data needed for business purposes, in a manner so that it can be accessed with a reasonable performance (Zurawski, 2005), and although a number of applications have emerged in the On-Line Analytical Processing (OLAP) space, there has been little theorizing about the management challenge of turning stored data into meaningful and valuable decisions, with perhaps the exception of Kimball's "Analytical Cycle" that suggests a way to conduct a business analysis (Kimball, 1998).
The framework proposed as part of the Computer Aided Leadership & Management (CALM) thesis (Middelfart, 2005) can assist with a heuristic approach to applying the technologies available at any given time, to the challenges faced by leaders and managers; moreover this framework can assist in uncovering new areas for future research that will benefit organizations in the years to come.
The idea is to take the Observation-Orientation-Decision-Action (OODA) loop, which was originally pioneered by Top Gun fighter pilot John Boyd in the 1950s (Lind, 1985), and group and categorize the available Business Intelligence technologies according to their role in the four processes in the OODA loop, namely:

Observation - in this phase we use technology to look at the data with an expectation of what it should be, e.g. through dashboards, reports and agents.

Orientation - in this phase we look at the data in different ways depending on what it shows, e.g. through analysis, data mining and simulation. Typically this phase is initiated after something in the observation phase has proven to be different from what we expected.

Decision - this phase is currently undertaken primarily by human intelligence, but the results found in the data warehouse can be evaluated against other external or internal - typically unstructured - data.

Action - based on the decision made, we seek to implement the course of action chosen, either through the IT infrastructure as a business process or as communication such as e-mail, phone, or plain face-to-face communication.

If we succeed in allowing a user in an organization to move as fast as possible all the way from the observation phase to the action phase, we are basically using information technology to channel the core competency of that organization as effectively into its environment as possible. From a competitive perspective, success is not granted, but if maximum speed in the application of core competency is achieved, there is nothing more we can do for that organization. Therefore, the ultimate challenge for developers of information systems that cater leaders and managers in an organization is to ensure the speed from the observation to the action phase, and in some instances allow a deeper understanding of what the organization's core competency really is.

To meet these challenges there are obvious areas that need to be improved with contemporary technologies. Typically the different technologies within the field of Business Intelligence have been fragmented, meaning that different applications were needed in order to create and distribute reports (observation) or to create and conduct analysis (orientation) (Middelfart, 2002). Moreover, there are very few technologies that seek to go beyond the phases of observation and orientation to also cater the needs in the decision and action phases.
This article presents two directions in which research will be performed during the Ph.D. project. In general, the research will aim for technologies that assist users to act fast upon changes within their area of interest or responsibility. There are basically two directions where we can improve the ability to react fast to changes in the environment, first we can *improve the speed with which we go through the four phases in the OODA loop*, and secondly, we can seek to *expand the horizon* in the environment where we conduct our observation and orientation.
The rest of this article is structured as follows: section 2 presents the research methods that will be applied during the Ph.D. project. Section 3 presents preliminary results as well as expectations for the research in progress. Section 4 summarizes and concludes on the findings, concepts and ideas presented in this article.

2. METHODS

2.1 Improving the Speed in OODA Loops

In order to measure the speed with which a user is able to move through the phases in the OODA loop, we need to understand if such a measure exists. If successfully identified, such a measure will be able to provide a foundation for structured improvement of OODA speed in the future. In this project, we propose the number of interactions the user makes during the phases, and in contemporary Business Intelligence technology these interactions can be counted as the number of (mouse-) clicks.
This research will investigate the feasibility of few clicks being a valid measure for the effectiveness of the user traveling through the phases of the OODA loop. In practice we will set up a number of tasks that a number of test users will have to do in a number of applications. During these tasks we will observe the number of clicks as well as the success of the tasks, and in addition, we will ask the users own opinion of perceived usability. Using these data we will investigate if any correlations exist between the number of clicks, the number of successes/errors and the perceived usability. We will select a number of tasks and applications that are as closely related to the field of Business Intelligence as possible. We will also conduct research in the real world by observing the number of user clicks on a real Business Intelligence application to assess whether few clicks is a meaningful measure. In this context, we will look into the number of clicks, the time between the clicks as well as the complexity of the information presented to the user, where complexity is a measure based on the number of cells, dimensions and measures that are returned to the user. Through these experiments we hope to gain important insight into how we can better design Business Intelligence applications that move the user through the OODA loop faster, in particular, we expect to gain valuable insight from the discrete observation of users since this is not biased as opposed to the users' opinions gathered.
Based on the data collected, we will test the hypothesis that few clicks is indeed a meaningful measure to assess usefulness of a Business Intelligence application.

2.2 Expanding the Horizon with Sentinel Rules

Expanding the horizon in the environment can in theory be done using the sentinel concept as it is proposed in the CALM thesis. However, in the CALM thesis there is no actual process described of how to create the link between the sentinel and the Key Performance Indicator (KPI). The research in this project will seek to bridge this gap by proposing and testing methods to discover so-called *sentinel rules* for Business Intelligence in the real world.
Let us imagine a company selling a product on a global scale where we have discovered the following sentinel rule: "IF negative blogs go up THEN revenue goes down within two months AND IF negative blogs go down THEN revenue goes up within two months". Let us assume that every day Google is searched for negative blogs, and the number of negative blogs is stored along with all other measures in the company's OLAP database. In addition, the company's Business Intelligence solution has the capability of notifying users based on the sentinel rules when data is loaded. When a user receives an agent notification on his desktop or PDA which tells him that the number of negative blogs go up, he will know ahead of time that revenue will go down in two months with a certain confidence. Depending on the situation, the user might have a number of possible evasive actions such as: post positive blogs to sway the mood, or perhaps reduce cost in order to cope with an anticipated reduction in revenue. Whatever the course of action, the sentinel rule has contributed to awareness of a potential problem and thus reduced the time from observation to action.
One could say that sentinel rules, once identified, reduce the time from observation to action, and that the discovery of sentinel rules reduces the time in the orientation phase. One way to look at the sentinel concept is that it expands the "horizon" by allowing the user to see data from the external environment, and not only for the internal performance of the organization. The idea behind placing "sentinels" at the outskirts of the data available for an organization seeks to harness both these ways of improving reaction time and thus organizational competitiveness. Metaphorically, what we seek to do for navigation of organizations with sentinels, is what radars do for navigation of ships. A sentinel rule is a relationship between two measures, A and B, in an OLAP database where we know, that if measure A changes at one point of time, then measure B changes within a certain warning period, with a certain confidence. If such a relationship exists, we call measure A the source measure, and measure B the target measure. Usually, the target measure is a KPI or something critical to fulfilling a KPI. The source measure is ideally a measure that either represents the external environment, or a measure as close to the external environment as possible (from the perspective of the organization's OLAP database). Examples of source measures

for an organization could be: the number of negative blog entries (external), the number of positive articles in papers (external), or the number of complaints from customers (internal, yet as close to external as possible). Examples of target measures could be: revenue and contribution margin.
The scientific method applied in the sentinel research is a blend of analytical, constructive and experimental approaches during which a series of prototypes will be developed. These prototypes will be validated on synthetic data to assess functional correctness and on real-world data to assess them for usefulness. In addition, both types of validation will assess the usefulness of the algorithms from a performance point of view.

3. RESULTS

In November 2007 an article which introduced OODA as a heuristic approach to improving Business Intelligence was presented on the annual Data warehousing and OLAP (DOLAP) conference. By using the approach proposed, three new desired directions for Business Intelligence technologies were identified (Middelfart, 2007). Specifically: (1) artificial intelligence to reduce human interaction in the OODA loop, (2) "sentinels" that can give early warnings about a later influence on a business critical measure, and (3) the ability to analyze the speed and quality of an OODA loop in order to quantify and analyze organizational talent and core competencies. To this date, the hypothesis that the OODA approach can identify new technologies and areas of improvement in Business Intelligence has not been rejected.
It is expected that the future research into the optimization of OODA using the number of human interactions as a measure will further prove the value of this methodology while, at the same time, result in interesting findings on how to design better Human-Computer Interfaces (HCI) for decision making.
The first article about sentinel rules is currently in the process of conference submission. The article proposes the concept of sentinel rules for multi-dimensional data that warns users when data concerning the external environment changes. The first prototype for sentinel rule discovery has been developed and tested on both synthetic and real-world data. Based on both theoretic assessment and practical testing, we found that useful sentinel rules can indeed be found with a reasonable performance. In addition, the solution handles the fuzziness of real-world data by using a weighted elimination process, this means that a user has to deal with fewer, more general rules, which means that less time is consumed to interpret the output.
The future research in sentinel rules introduces a parameterized implementation of the sentinel algorithm that exploits the dimensions to discover sentinel rules that only appear at some levels of the dimensional hierarchies. The approach will expand the algorithm from the previous article to exploit the hierarchical data to identify clusters where certain sentinel rules apply. Most likely the algorithm will rely on a bottom-up approach where sentinel rules are grown from the lowest granularity of the data, and as we travel higher up in the hierarchy, some rules are retired if they do not apply on these levels; effectively we expect to end up with a set of sentinel rules that apply to a cluster bounded by the dimension levels over which no rules are applicable. For example one can think a scenario where the rule: "IF negative blogs go up THEN revenue goes down within two months AND IF negative blogs go down THEN revenue goes up within two months" applies to the state of California, this rule is then tested and found to apply to all states in the United States, however, it is not found to apply to any other countries in the world. In other words, we grew the rule from the lowest level, in this case California, and from there we went to the next hierarchical level, the United States, and found the rule to be true as well. At the global level we found the rule not to be true, and thus the largest cluster for which the rule is true is the United States. The use of this bottom-up approach is expected to give a very fast performing and useful addition to the algorithm for sentinel discovery.
In addition to adding multi-dimensionality to sentinel rules, some research will be dedicated to expose the prototypes developed to massive amounts of real-world data. We expect to find that the algorithm will find meaningful rules and while doing so give us interesting insight by identifying certain clusters where specific rules apply. Moreover, we expect to find that in many cases the entire discovery process can run in main memory which means that performance on real-world data is highly feasible and within reach for many organizations.

4. CONCLUSION

The research described in this article will be compiled into a Ph.D. thesis called *"Optimizing Business Intelligence using the OODA Framework"*. The six articles, whereof two have already been written, is expected to be a valuable, original contribution of concepts and solutions that can improve the "horizon" and speed of an OODA cycle for a user of a Business Intelligence system. In addition, it is expected to form a good research foundation that others can use in their research going forward.

REFERENCES

R. Kimball. *The Data Warehouse Lifecycle Toolkit: Expert Methods for Designing, Developing, and Deploying Data Warehouses.* Wiley, 1998.

W.S. Lind. *Maneuver Warfare Handbook.* Westview Press, 1985.

M. Middelfart. *Business Analytics - From Chosen Few to Common World View.* <www.morton.dk> current as of October 2nd 2008

M. Middelfart. *CALM: Computer Aided Leadership & Management - How Computers can Unleash the Full Potential of Individuals and Organizations in a World of Chaos and Confusion.* iUniverse, 2005.

M. Middelfart. Improving Business Intelligence Speed and Quality through the OODA Concept. *In Proc. of DOLAP*, pp. 97–98, 2007.

R. Zurawski (ed.), C.S. Jensen, and T.B. Pedersen. *The Industrial Information Technology Handbook*, Multidimensional Databases and OLAP. CRC Press, 2005.