

# Toward a Goal-oriented, Business Intelligence Decision-Making Framework

Alireza Pourshahid<sup>1</sup>, Gregory Richards<sup>2</sup>, Daniel Amyot<sup>1</sup>

<sup>1</sup> School of Information Technology and Engineering, University of Ottawa, Canada

<sup>2</sup> Telfer School of Management, University of Ottawa, Canada

alireza.pourshahid@gmail.com, richards@telfer.uottawa.ca, damyot@site.uottawa.ca

**Abstract.** Decision making is a crucial yet challenging task in enterprise management. In many organizations, decisions are still made based on experience and intuition rather than on facts and rigorous approaches, often because of lack of data, unknown relationships between data and goals, conflicting goals, and poorly understood risks. This paper presents a goal-oriented, iterative conceptual framework for decision making that allows enterprises to begin development of their decision model with limited data, discover required data to build their model, capture stakeholders goals, and model risks and their impact. Such models enable the aggregation of Key Performance Indicators and their integration to goal models that display good cognitive fit. Managers can monitor the impact of decisions on organization goals and improve decision models. The approach is illustrated through a retail business real-life example.

**Keywords:** Business Process Management, Business Intelligence, Decision Support Systems, Goal-oriented Modeling, Indicators

## 1 Introduction

Decision making is a crucial yet challenging task for many managers. Many challenges arise from the rapid growth of data within an operating environment of continuous change and increasing customer demands. Although many enterprises have applied different decision aids such as Business Intelligence (BI) tools in an attempt to improve decision-making capability, these approaches have not always met with success. We believe that one of the problems with the use of such tools is the lack of approaches that integrate goals, decision-making mechanisms and Key Performance Indicators (KPIs) into a single conceptual framework that can adapt to organizational changes and better fits manager's cognitive decision models. A secondary issue relates to the unavailability of sufficient data when performance models are first put in place.

The purpose of this paper is to describe the development of such a BI framework, and also the technical means to implement it in the enterprise. The paper first describes some of the issues related to decision making using BI tools. It then

describes extensions of the Goal-oriented Requirement Language (GRL) used as a modeling environment to support the aggregation of KPIs (whose values are either coming from external data sources through Web services or simulated as part of what-if strategies) and their integration to the rest of the goal model during formal analysis. A new formula-based goal evaluation algorithm is introduced that takes advantage of this aggregation of KPIs. In addition, the paper provides the implementation steps of the proposed BI-supported decision framework, which are applied iteratively to a retail business example where little data is available at first. Finally, lessons learned and conclusions are discussed.

## 2 Background Review

### 2.1 BI-Based Decision Making

Over the past 30 years, the growth of BI technology has helped managers make better decisions through improved organization of information, better data quality, and faster and more effective delivery of information. It has been estimated, however, that more than 50% of BI implementations fail to influence the decision-making process in any meaningful way [10]. Reasons for this include cultural resistance, lack of relevance, lack of alignment with business strategy, and lack of actionable and “institutionalized” decision support technologies [8]. Many of these problems could be attributed to approaches used for defining the data to be delivered by the BI tool.

Most data delivery schemes are based on dimensional models of the data. This approach often leads to a sound technical data model, but this view of the data might or might not fit with the user’s *decision model*. Indeed, Korhonen et al. [11] point out that one of the key challenges faced in institutionalizing decision aids is validation of decision models used by the *decision maker*. These authors argue that problems with model validation occur when relationships between the variables included in the decision model are not accurate and when the available data does not match the model’s specifications. Although the data model is often developed by first defining user needs in terms of the variables (i.e., data values) required, this approach does not necessarily illustrate *relationships* between the variables nor does it define variables in a cause-effect framework that matches the decision model used by decision makers. Thus the technical data model differs from the decision model.

Vessey [18] further suggests that more effective decision making results when the decision aid directly supports the decision task. The essence of this argument revolves around the notion of *cognitive fit*, which results when a good match exists between the problem representation (i.e., in this case the way data is presented by the BI tool) and the cognitive task (the way data is used) involved in making decisions.

The concept of cognitive fit is supported by research in the field of behavioral decision making which demonstrates that decision makers tend to make better use of information that is explicitly displayed. Moreover, they tend to use it in the form in which it is displayed. Slovic [17] for example points out that “information that is to be stored in memory, inferred from the explicit display, or transformed tends to be discounted or ignored.” Therefore, cognitive fit is enhanced when data is presented in

a form that fits well with the processes the decision maker uses to make decisions. This results in lower “cognitive load” (i.e., less manipulation of the data by the user), which facilitates the decision-making process.

In terms of the decision-making process itself, the key impact of a decision model in *goal-directed* systems is improving the probability of goal accomplishment. The “cause-effect” nature of such decisions is related to resource allocation. For example, should a manager invest more in advertising in order to improve revenues or would an investment in training have more of an impact? According to Vessey [18], these types of decisions call for an understanding of associations between variables (i.e., impact of advertising and training on revenue growth). The problem is that in most BI tools, such associations are not defined. Decision-makers need to process the data by estimating whether the cause effect model is correct then estimating the strength of the relationships. According to Popova and Sharpanskykh [13], even when relationships can be defined such as in the ARIS model [5], which allows users to define cause-and-effect relationships using Balanced Scorecards and connect KPIs to strategic goals, the analysis options are inadequate due to a lack of formal modeling foundations and proper representations. The more processing the decision-maker has to do, the higher the cognitive load and the less efficient the decision-making environment. The graphs versus tables literature [6,18] for example, argues that decisions can be improved (i.e., faster and more accurate decisions can be made) when cognitive load is reduced and when values for each of the variables in the model are displayed in their proper context.

This literature suggests that the failure of many BI tools to enhance decision making could be related to the lack of cognitive fit. The underlying data models used by multi-dimensional tools for example, provide data in tabular or graphical format, but these formats do not explicitly identify the variables important to the decision, the relationships between the variables, or the context for the decision itself. Therefore, it seems reasonable to assume that the use of business intelligence tools for decision making can be enhanced if the decision model (i.e., the cause effect relationships among variables relevant to the decision) is displayed by the tool, if the model is linked to the decision’s context (in this case, the desired strategic outcomes), and if the associations between the variables can be readily understood.

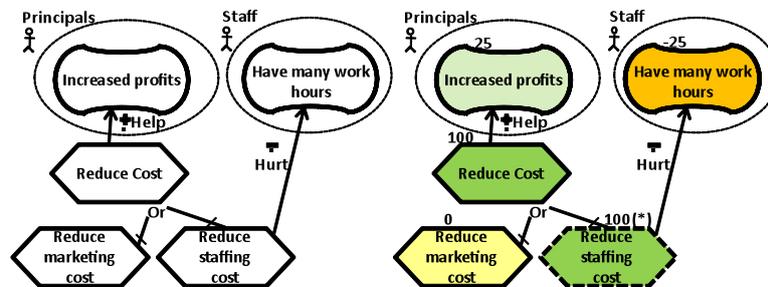
## **2.2 Goal-oriented Requirement Language**

Goals are high-level objectives of an enterprise, organization, or system. The requirements engineering community has acknowledged the importance of goal-oriented approaches to system development many years ago. Yu and Mylopoulos [20] observed that goals are important not just for requirements elicitation, but also to relate requirements, processes and solutions to organizational and business contexts, and to enable trade-off analysis and conflict resolution. Complete goal-driven development approaches now exist to support software development [19].

The Goal-oriented Requirement Language (GRL) is a graphical notation used to model and analyze goals. Although many goal-oriented languages exist, GRL is the first and currently only standardized one. GRL is part of the User Requirements Notation (URN), a standard of the International Telecommunication Union intended

for the elicitation, analysis, specification, and validation of requirements using a combination of goal-oriented and scenario-based modeling [9]. In URN, GRL is complemented by a scenario notation called Use Case Maps, which offers an operational or process-oriented view of a system.

GRL enables the modeling of stakeholders, business goals, qualities, alternatives, and rationales. Modeling goals of stakeholders with GRL makes it possible to understand stakeholder intentions as well as problems that ought to be solved. GRL enables business analysts to model strategic goals and concerns using various types of *intentional* elements and relationships, as well as their stakeholders called *actors* (○). Core intentional elements include *goals* (□), *softgoals* (◡) for qualities, and *tasks* (◡) for activities and alternative solutions. Intentional elements can also be linked by AND/OR *decompositions*. Elements of a goal model can influence each other through *contributions*, displayed as arrows. Qualitative positive (make, help, some positive) and negative (break, hurt, some negative) contribution levels exist, as well as quantitative contribution levels on a scale going from -100 to +100.



**Fig. 1.** Example of GRL model (left), with evaluation (right)

Fig. 1 (left) illustrates some of the above concepts with a toy retail store example, where principals (actor) want increased profits (softgoal) and the staff wants to have many work hours. Reducing costs (task), which can help satisfying the principals' objective, can be decomposed into two non-mutually exclusive options: reducing the marketing cost or reducing the staffing budget. The latter option however can hurt the staff's objective. As modelers get deeper knowledge of these relationships, they can move from a qualitative scale (e.g., *Help*) to a quantitative one (e.g., 35) for contributions and for satisfaction values. Such models can help capture stakeholder's objectives as well as their relationships in an explicit way in terms understandable by managers, and hence improve cognitive fit.

GRL *evaluation strategies* enable modelers to assign initial satisfaction values to some of the intentional elements (usually alternatives at the bottom of a goal graph) and propagate this information to the other elements through the decomposition and contribution links. Strategies act as *what-if* scenarios that can help assess the impact of alternative solutions on high-level goals of the involved stakeholders, evaluate trade-offs during conflicts, and document decision rationales. Different goal evaluation algorithms (using qualitative values, quantitative satisfaction values between -100 and +100, or mix of both types) for GRL are discussed in [1].

*jUCMNav* is an open source URN tool for the creation, analysis, and management of URN models [12]. It allows for the qualitative, quantitative, or hybrid evaluation of

GRL models based on strategies. To improve scalability and consistency, jUCMNav also supports the use of multiple diagrams that refer to the same model elements.

Fig. 1 (right) illustrates the result of a strategy for our example where the reduction of the staffing budget is selected, i.e., the satisfaction value of this task is initialized to 100. In jUCMNav, initialized elements are displayed with dashed contours. This strategy eventually leads to a weakly satisfied (+25) “Increased profits” softgoal and to a weakly denied (-25) satisfaction level for “Have many work hours”. The resulting satisfaction values in top-level goals and actors should be used to compare different strategies and find suitable trade-offs rather than be interpreted as some sort of satisfaction percentage. jUCMNav also uses a color-coding scheme to highlight satisfaction levels—red for denied, yellow for neutral, and green for satisfied (with various shades for values in between)—which again improve intuitive understanding.

### 2.3 GRL and KPI for Business Modeling

Although the primary application domains for URN target reactive systems and telecommunications systems, this language has also been applied successfully to the modeling and analysis of business goals and processes [21]. Goal models combined to process models have been used elsewhere to assess the risk and viability of business solutions [2] and model different concerns of interest to different stakeholders [4]. However, in order to better support business process monitoring and performance management, Pourshahid et al. [14] have extended standard GRL with the concept of *Key Performance Indicators* ( $\langle \text{KPI} \rangle$ ). KPIs can also be analyzed from various angles called *dimensions* ( $\langle \text{Dim} \rangle$ ), in a way similar to what is found in common BI systems. Dimensional data allows one to look at the data from different points of view and filter or aggregate the data based on the defined dimensions. For instance, in Fig. 2, staffing cost can be aggregated in all locations in all years of store operations or can be analysed for Store1, 2, 3 and the online store individually and in a specific month or year.

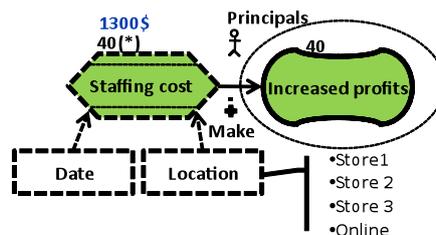


Fig. 2. Example of a KPI with dimensions and evaluation

KPIs include specifications of *worst*, *threshold*, and *target* values in a particular *unit*. For example, a Staffing cost KPI (see Fig. 2) could have a target value of \$1000, a threshold value of \$1,500, and a worst value of \$2,500. KPIs also contain a *current* value, which is either defined in a GRL evaluation strategy or provided by an external source of information such as a database, an Excel sheet, a BI tool, external sensors, or Web services. The KPI is a metrics of the system that normalizes the current value to a scale of -100 to 100, which enables it to be used like any other intentional

element in a GRL model. For instance, if the current Staffing Cost is \$1300, then the normalization function, which takes here  $|\text{threshold-current}| / |\text{threshold-target}| * 100$ , will result in a satisfaction level of 40. Furthermore, when the current value is between the threshold value and the worst value (e.g., 2500), then the normalization function becomes  $|\text{threshold-current}| / |\text{worst-threshold}| * (-100)$ , which results in a negative value (e.g., -100). If the result is higher than 100, then it becomes 100 (symmetrically, if it is lower than -100, then it becomes -100). Such an evaluation strategy was used in Fig. 2. Note also in this model that Staffing cost could be drilled down (e.g., explored) according to the Date and Location dimensions.

Although goal modeling and scorecards have been used in combination in the past [3, 15], we believe KPIs are also necessary because they act as an interface between conventional goal models and quantitative, external sources of information.

Furthermore, Pourshahid et al. [14] have introduced and implemented a service-oriented architecture enabling the use of underlying data and BI reports by the jUCMNav tool. jUCMNav is connected to BI systems via a Web service. All the information generated by the BI system, from raw data to very complex data warehouses, can hence be used as qualitative data to initialize the KPIs used in the GRL model, and against which goal satisfaction is evaluated.

Although several other goal modeling languages exist (e.g., *i\**, TROPOS, KAOS, and the Extended Enterprise Modeling Language), the combination of support for KPIs and performance modeling, the ability to combine process and goal models and perform analysis on both, existing tool support for using BI systems as sources of data, and the fact that URN is a standard modeling language, all together have convinced us that URN is the best language to be used in the context of our research.

### 3 New Formula-based Evaluation Algorithm

Although several GRL evaluation algorithms (qualitative, quantitative and hybrid) already exist [1], none of them provides the formula-based KPI aggregation required for the type of cause-effect analysis performed in our decision making context. As illustrated in the previous section, the current algorithms allow modelers to specify the contribution level of a KPI on another GRL intentional element and to calculate the satisfaction level of that target element [14]. However, these algorithms prevent one KPI from driving the computation of the *current value* of another KPI. Although the current evaluation methods allow computing the impact of one KPI on another KPI in terms of *satisfaction level*, when it comes to showing the impact of several KPIs on one KPI (e.g., their *aggregate* effect), the current evaluation methods quickly become a bottleneck and thus obstruct the cause-effect analysis.

Other modeling languages and enterprise modeling frameworks exist that can be used to model KPIs, however many have a limited computational power and do not allow one to define proper relationships between KPIs for advanced analysis [13]. In addition, there have been recent efforts in industry to use strategy maps and measurable objectives to help with decision making and process improvement [16]. However, influence of KPIs on one another has not been discussed.

In order to address this issue, we introduce further extensions to GRL and a novel evaluation algorithm that allow analysts and decision makers to define mathematical formulae describing relationships between the model elements. This method, which extends the bottom-up propagation algorithm defined in [1], enables the precise definition of accurate relationships between these elements. Analysts gain full control of the model and can change the impact of one element on another as desired.

The algorithm uses current/evaluation values of the source KPIs as inputs for the formula (described as metadata, see Fig. 3) and calculates the target KPI evaluation value using these inputs. Then, the satisfaction level of the KPI is calculated using the KPI's target/threshold/worst values as discussed previously. The impact of KPIs on other types of intentional elements (e.g., goals, softgoals and tasks) is computed using conventional GRL quantitative and qualitative algorithms. This unique combination allows one to have both quantifiable KPIs and strategic-level softgoals that are hard to quantify together in the same model and to show and monitor the impact of KPIs on the goals of the organization.

Fig. 3 shows a simple example where the current KPI values are displayed, with their units, above the usual satisfaction values. Note that the inputs can be of different units; the formula in the target KPI must take this into consideration. In this example, the current value of *Profit* is computed as  $Revenue - Costs - Stolen * 50$  (the first two are in dollars and the third is a number of items). Note also that the contributions have no weight; the satisfaction of the *Profit* KPI is based on the normalization of its computed current value (\$39,000) against its specified target, threshold and worse values. We have prototyped this new algorithm in the jUCMNav tool.

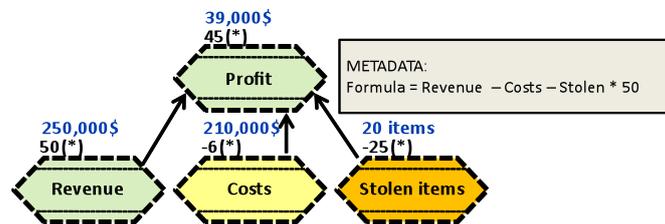
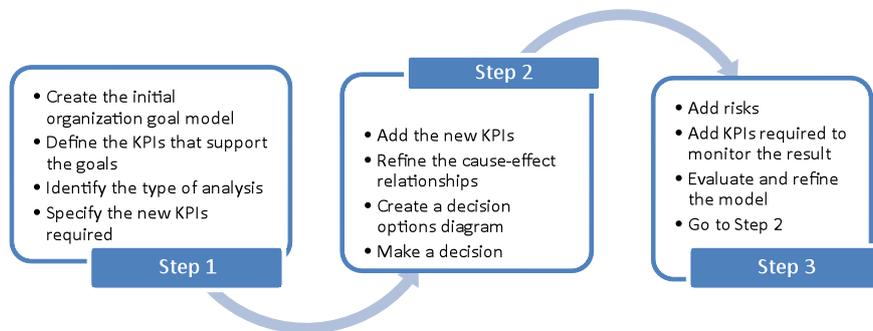


Fig. 3. New extension to KPI evaluation

Another benefit of this approach is the ability to account for risk. In organizations, cause-effect analysis and decision making usually involve an element of *risk*. Even though we could show risks as model elements in GRL diagrams (e.g., using softgoals stereotyped with «Risk»), it is very hard to quantify the impact of risks on the value of a KPI and consider it in the evaluation algorithm. In our new algorithm, we use risk as yet another input to the target KPI and connect it using a contribution link. The target KPI *threshold value* changes based on the level of contribution of the risk factor on the target element. This allows the modeler to vary the acceptable range of values for a KPI when there is expected risk involved.

## 4 Business Intelligence Decision-Making Framework

Based on the reasoning behind the notion of cognitive fit, the framework we are proposing defines the organizational goals and links these explicitly to a decision model and relevant Key Performance Indicators. The framework can be used by organizations at any level of maturity and readiness in terms of gathering and monitoring data for BI-based decision making. In particular, unlike many simulation approaches, it does not necessitate up front large quantities of data to be useful. We believe different organizations however have different needs and may be in different states when they decide to incorporate such a framework. Consequently, we are suggesting a spiral method consisting of three basic steps involving many iterations that build upon each other (Fig. 4).



**Fig. 4** Business Intelligence Decision-Making Framework

In the **first step**, an initial model of the organization's goals is created [7]. This model can be built based on interviews with executives and operational managers as we experimented with in our example. This goal model can consist of long term, short term, strategic and operational goals of the organization as well as contribution and decomposition relationships between them. Furthermore, in this step we define *the KPIs that support the goals* (e.g., financial KPIs) and add them to the model. This can be a challenging task and is very dependent on the level of maturity of the organization. For instance, in two cases we have studied as part of this research, one small organization had a very limited set of data and was using a spreadsheet to monitor the business while the other one had many indicators available and was using a sophisticated Business Intelligence system. Our discussions with both organizations however demonstrate that any organization at any point within this wide range of information management capabilities can benefit from applying this goal-based model. After defining the model, we *identify the type of analysis* we want to perform on the model and *specify the new KPIs required* to do so.

In the **second step**, we *add the new KPIs* and the new dimensions to the model. Note that not all the KPIs need to be dimensional and if the available data is not as granular as is required for a dimensional model, or if all the data is not available, a step-by-step approach can be used leading to a number of model iterations as additional data becomes available.

In addition, during this process we *refine the cause-effect relationships* between KPIs in the goal model (hence improving cognitive fit). These relationships create a BI-enabled decision framework which can be used to document the rationale for goal accomplishment, the decision context, and to analyze what-if scenarios. The framework also helps one to evaluate the impact of a decision on the enterprise's goals through the use of historical and trend data.

In cases where an organization does not have historical data and is in its early iterations of BI-based decision making, the initial formula used to define the decision framework can be based on industry standards. As the organization gathers more information, this historical data can be integrated to the model. As will be seen later in the retail example, a decision framework can be used to illustrate the expected impact of actions taken by managers. Furthermore, they can be continually adapted by saving the initial iteration as a “snap shot” and comparing it to actual results achieved by decisions. Gathering these snap shots will eventually create a “decision trail” that displays context, decisions taken and results of these decisions allowing managers to make better decisions in the future. In addition, decision trails allow organizations to refer back to the rationale they used for making successful or unsuccessful decisions.

We also add a *decision options diagram* (in the same GRL model) and connect these options to the goals and KPIs of the organization. A decision options diagram outlines the different options available to an organization to achieve a goal.

In the **third step**, we add the expected impact of the decision made in the second step to the model and *include risks* involved in the decision. In using GRL softgoals, we are able to show qualitatively the impact of risks on the rest of the model. The most challenging aspect of this step is modeling a qualitative risk factor that influences a quantitative KPI. In this case, we model the impact by increasing the range of acceptable values for a KPI. In other words, once we have estimated the inherent risk, we allow the acceptable range of the measured KPI to deviate accordingly from its target value.

In this step, we also add the *required KPIs* and dimensions to the model that allow one to better observe the impact of decisions. If we expect a decision to change anything in the organization, we will examine that hypothesis using the appropriate KPIs and GRL strategies. Finally, we *monitor* the impact of the decision and compare expected results against actual results. Based on this comparison, we *adjust* the decision framework as required and record the data.

In summary, the iterative cycle is based on creating an initial model which is then refined by expanding data sources, capturing decisions made and the results of those decisions, and building historical decision trails that informs future models.

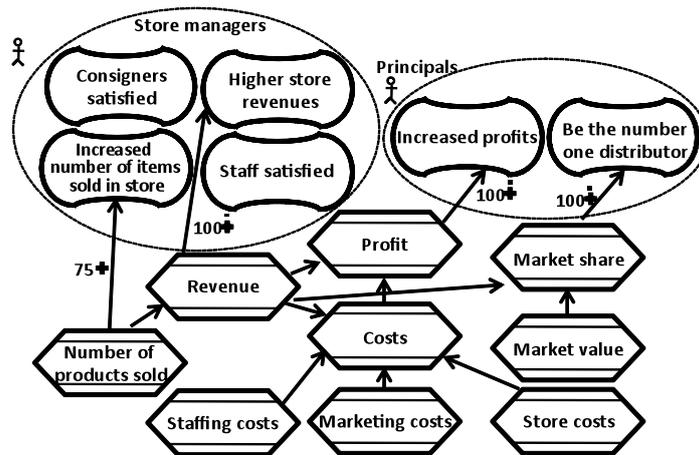
## 5 Retail Business Real-Life Example

In this example, we describe an initial application of the framework to a real, Ontario-based retail business. The retail business would be categorized as a small enterprise (revenues less than \$50 million) with 4 local stores and planned expansion nationally. The business has existed for over 15 years, establishing a strong foothold in one neighborhood. Three years prior to the study, the business was purchased by new

owners who set national growth as a key strategic objective. As part of the expansion plans, market and competitive studies were conducted. In addition, the owners had created a scorecard that tracked key operational indicators and provided the ability to conduct an assessment of business results. Some data however, for example, the flow of customers through each of the locations, was not yet available in the scorecard.

At the time of the study, most revenues were earned through consignment sales. The business however had started selling new items as well and was planning to invest in an online business. Revenue was driven by ensuring that stock was properly displayed which in turn depended on assuring that enough staff were available to sort, tag, and lay out the products. The supply side of the business depended on the number of consigners available, the amount of product they brought to each store, and the speed at which these products could be displayed. The demand side depended on local advertising and word of mouth that stimulated traffic flow. All stores were situated in prominent locations with good visibility that stimulated walk-in traffic.

To begin our investigation, we interviewed the CEO in order to identify the high-level goals as well as KPIs and drivers of organizational success. As depicted in Fig. 5, the goals of the principals were related to market growth: they wanted to be the number one distributor within their geographical market. Store managers were aware of the growth objective, but on the short term, they focused on increasing revenues and the number of items sold in their stores.



**Fig. 5.** First iteration model (aggregation formulas not shown here for simplicity)

As discussed above, the first iteration of the model provides an initial alignment of higher-level goals and KPIs. We started with a minimal decision model and limited set of data just to illustrate the business goals and financial targets and to identify the indicators and driver KPIs required to monitor the business and to make informed business decisions. Fig. 5 illustrates the first iteration of the model. At this stage, we also developed a rough dimensional model (Fig. 6) in order to ensure that the data needed for the decision model would be available. The dimensional model helps the store to analyze the impact of the KPIs on goals based on their different store

locations, in each period of time, by product type (e.g., clothing, electronics, etc.), and by product category (i.e., retail and consignment).

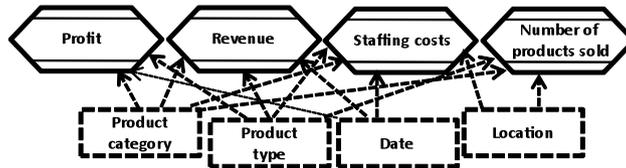


Fig. 6. First iteration dimensions

In Step 2 of the process, more KPIs were added to the model as drivers and then linked to the high-level financial KPIs and organizational goals. In addition, we added the new KPIs as well as a new dimension called “marketing type” (e.g., outreach, online advertising, etc.) to the dimensional model. This new dimension allows decision makers to analyze which marketing initiative has a more significant impact on the goals. In this step we also created a decision options diagram (Fig. 7) illustrating the specific actions managers can take to improve goal accomplishment. One of the decision options available to managers in this case, is to invest in an online business. We consider this option as the decision made by managers and update the models accordingly in step 3.

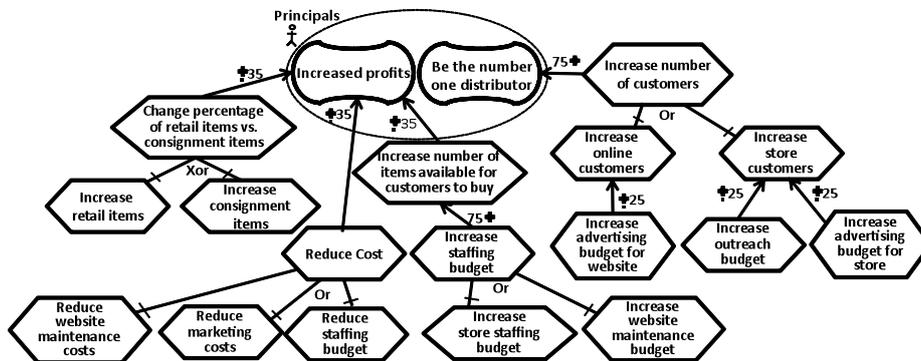


Fig. 7. Second iteration decision options diagram

Fig. 8 depicts the complete model, which defines the expected impact of the actions identified in Fig. 7 along with the KPIs and acceptable ranges for each of the relevant goals based on the risks associated with each goal. There is one new risk factor in the model that is associated with the investment in the online business. Furthermore, there is also a new KPI used to monitor the investment made in the online business and its impact on the costs. The GRL strategy used for the evaluation here focuses on the use of the online business investment (other GRL strategies were defined to evaluate different sets of options and find the most suitable trade-off). Fig. 9 depicts the final dimensional model (including its new “Sales method” dimension) used to ensure that

the relevant data can be delivered to decision makers. Note that, in jUCMNav, such figures can be split over many diagrams when they become complex.

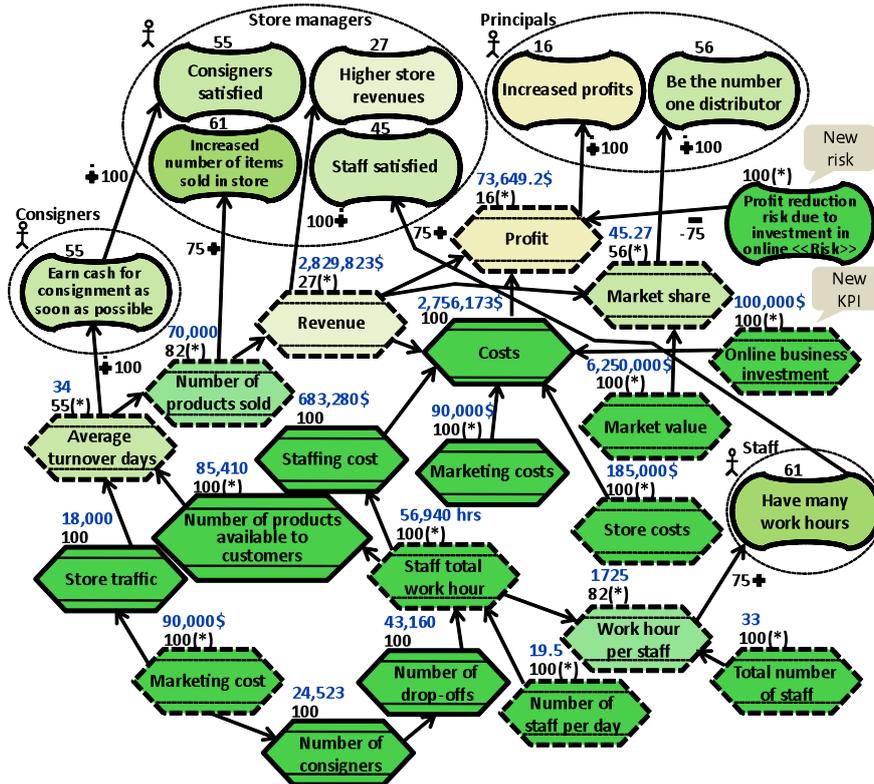


Fig. 8. Third iteration model – evaluated

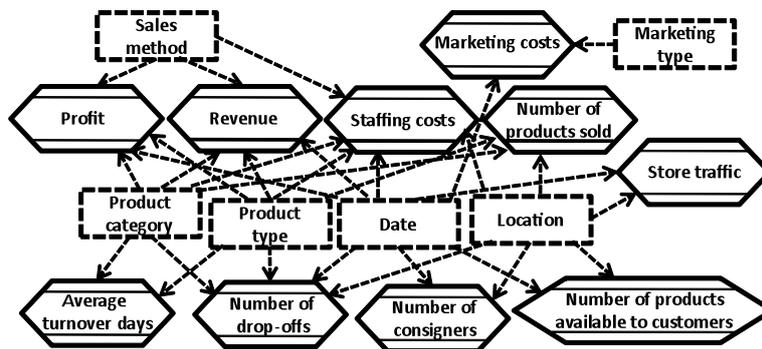


Fig. 9. Third iteration dimensions

## 6 Lessons Learned

The development of our framework and its application to the real world retail business data led to several lessons learned. From a business management perspective, we observe the following:

- Modeling goals and defining drivers and KPIs (i.e., creating the cause-effect decision model) not only helps to document the known aspects of the business but also helps to clarify unknown factors that might be driving goal accomplishment. Validation of the model through interviews with decision makers ensures that data and KPIs included are indeed relevant to the business. This can have a great impact especially for small businesses where initial goals might not have been clear.
- Even though modeling the indicators helps define the required information and the relationships between variables, we are still unsure about where to draw the line regarding the data we need to show in the model versus the data maintained in source systems (e.g., databases or BI reports). We still need to explore how to find the appropriate balance so we do not omit important information in the model for decision making while preventing the inclusion of too much data which can complicate the decision-making environment. We believe however that getting feedback on the right balance is facilitated by the use of graphical goal models with rapid evaluation feedback as provided by GRL strategies, which provide better cognitive fit than conventional BI reports. Note however that the goal-oriented view introduced here is complementary to what is found in BI tools, not a substitute.
- Defining relationships between the model elements without historical data is difficult. In some cases, managers themselves are not aware of the linkages because they have not had the historical data available to create cause-effect models. In this case, we first create the models using industry standards or “best guesses” and then use the different iterations of the framework to improve the cause-effect decision model.
- The ability to adjust the range of acceptable values for a KPI is useful for registering risk. For example, one might establish a wide range of acceptable values for an objective that carries a high level of risk, such as expected sales for a new product. On the other hand, objectives with lower profiles, such as sales of well-established products, might have a narrower range of acceptable values.

From a technical point of view, we have learned that:

- Our new extensions to GRL and the new formula-based algorithm provide a great deal of flexibility for model evaluation, especially as they are combined with standard goal satisfaction evaluation, hence offering the best of both worlds. However our new algorithm still has room for improvement, especially when it comes to using other intentional elements (e.g., goals) as contributors to KPIs. We have had limited experience with this idea by

considering risk as an input to KPIs, but this type of modeling may be useful in other situations that require further investigation.

- Creating different versions of a model in different iterations and keeping them consistent for comparison purposes can be painful with current tool support. Saving separate files for each version of the model quickly becomes a maintenance issue that requires a better technical solution.

## 7 Conclusions

There are critical issues related to the use of conventional Business Intelligence technology for decision making. The gap between the technical data model and the decision model creates a lack of cognitive fit, especially for supporting cause-effect decisions. This represents a challenge that will become even more important in the future given that organizations are nowadays gathering terabytes of information.

In this paper, we have provided several contributions toward a goal-oriented business intelligence decision framework, where we integrate in a novel way goal models, decision frameworks, action models, and risk, together with analysis capabilities. By integrating the decision framework into the BI system, we attempt to improve cognitive fit between the decision making task and the representation of information needed to complete the task. To do so, we extended a standard goal-oriented language, GRL, to better display relationships between Key Performance Indicators and objectives, enable formula-based evaluations of goal models, and integrate risk through the notion of acceptable ranges for KPIs. These extensions enable the combination of quantifiable KPIs with strategic-level softgoals in the same model, which in turn allows analysts to assess and monitor the impact of KPIs based on existing values and to explore what-if scenarios through GRL strategies. Tool support is provided as extensions to the open source jUCMNav tool.

From an implementation perspective, we also introduced a framework with iterative steps that support the construction of goal models (including KPIs and dimensions) even in situations where little information is available. Such models can be refined as more knowledge is gained about the organization and its context. Models can also be compared (as historical data) to validate newer models to assess the impact of past business decisions, leading to an ongoing system of record that permits continual adaptation. Our retail business example helped illustrate the framework in a real context, and suggests the feasibility of the approach. We also believe that such a graphical, goal-oriented approach, which delivers data values used to make decisions in context, supports the comprehension of important cause-effect relationships in a way that could complement existing current BI technologies, which often lack an appropriate goal view.

The framework is still evolving, and limitations and potential work items have been identified in our lessons learned. However, this framework brings new contributions and good value to the BI table, and we believe it has a promising future.

**Acknowledgments.** This research was supported by the Business Intelligence Network (BIN), a strategic network funded by the Natural Sciences and Engineering Research Council of Canada.

## References

1. Amyot, D., Ghanavati, S., Horkoff, J., Mussbacher, G., Peyton, L., Yu, E.: Evaluating Goal Models within the Goal-oriented Requirement Language. *International Journal of Intelligent Systems (IJIS)*, Vol. 25, Issue 8, August, pp. 841–877 (2010)
2. Ansar, Y., Giorgini, P., Ciancarini, P., Moretti, R., Sebastianis, M., Zannone, N.: Evaluation of Business Solutions in Manufacturing Enterprises. *Int. J. Bus. Intell. Data Min.* 3(3), December, pp. 305–329 (2008)
3. Babar, A., Zowghi, D., Chew, E.: Using Goals to Model Strategy Map for Business IT Alignment. In: Workshop on BUSITAL10, in conjunction with CAiSE'10, June. (2010)
4. Colombo, E., Mylopoulos, J., A Multi-perspective Framework for Organizational Patterns. In: *Conceptual Modeling, LNCS*, vol. 4215, pp. 451–467. Springer, Heidelberg (2006)
5. Davis, R., Brabänder, E. *ARIS Design Platform: Getting Started with BPM*, Springer (2007)
6. De Sanctis, G.: Computer graphics as decision aids: Directions for research. *Decision Sciences*, 15, pp. 463–487 (1984)
7. Feng, X., Richards, G., Raheemi, B.: The Road to Decision-Centric Business Intelligence. In: *2nd Int. Conf. on Business Intelligence and Financial Engineering*. Beijing (2009)
8. Hackathorn, R.: Making Business Intelligence Actionable. *DM Review*, 32 (2002)
9. International Telecommunication Union: Recommendation Z.151 (11/08), User Requirements Notation (URN) – Language definition (2008)  
<http://www.itu.int/rec/T-REC-Z.151/en>
10. Ko, I.S., Abdullaev, S.R.: A Study on the Aspects of Successful Business Intelligence System Development. In: *Int. Conf. on Computational Science (ICCS 2007)*. LNCS, vol. 4490, pp. 729–732. Springer, Heidelberg (2007)
11. Korhonen, P., Mano, H., Stenfors, S., Wallenius, J.: Inherent Biases in Decision Support Systems: The Influence of Optimistic and Pessimistic DSS on Choice, Affect, and Attitudes. *Journal of Behavioral Decision Making*, 21, pp. 45–58 (2008)
12. Mussbacher, G., Ghanavati, S., Amyot, D.: Modeling and Analysis of URN Goals and Scenarios with jUCMNav. In *17th IEEE Int. Requirements Engineering Conf., IEEE CS, USA*, pp. 383–384 (2009) <http://jucmnav.softwareengineering.ca/jucmnav/>
13. Popova, V., Sharpanskykh, A., Modeling organizational performance indicators. *Information Systems*, 35(4), Vocabularies, Ontologies and Rules for Enterprise and Business Process Modeling and Management, June 2010, pp. 505–527 (2009)
14. Pourshahid, A., Chen, P., Amyot, D., Forster, A.J., Ghanavati, S., Peyton, L., Weiss, M.: Business Process Management with the User Requirements Notation. *Electronic Commerce Research*, 9(4). Springer, pp. 269–316 (2009)
15. Siena, A., Bonetti, A., Giorgini, P.: Balanced Goalcards: Combining Balanced Scorecards and Goal Analysis. *Third Int. Conf. on Evaluation of Novel Approaches to Software Engineering (ENASE 2008)*. Funchal, Portugal (2008)
16. Silver, B. Collaborative Business Process Discovery and Improvement. *Industry Trend Reports*, Bruce Silver Associates, May (2010)
17. Slovic, P.: Psychological Study of Human Judgment: Implications for Investment Decision Making. *The Journal of Finance*, 27:4, pp. 779–799 (1972)
18. Vessey, I.: Cognitive Fit: A Theory-Based Analysis of the Graphs versus Tables Literature. *Decision Sciences*, 22 (2), pp. 219–240 (1991)
19. van Lamswerde, A.: *Requirements Engineering: From System Goals to UML Models to Software Specifications*. John Wiley & Sons (2009)
20. Yu, E., Mylopoulos, J.: Why Goal-Oriented Requirements Engineering. In *Fourth Intl. Workshop on Req. Eng.: Foundation for Software Quality (REFSQ'98)*, Pisa, Italy (1998)
21. Weiss, M., Amyot, D.: Business Process Modeling with URN. *International Journal of E-Business Research*, 1(3), pp. 63–90 (2005)