Incorporating expert knowledge in decision support for logistics

Balancing costs and customer relations

TRAIL Research School, Delft, November 2006

Authors

Tamas Mahr MSc, Dr. Mathijs de Weerdt

Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology

F. Jordan Srour MSc, Dr. Rob Zuidwijjk

Rotterdam School of Management, Department of Decision and Information Sciences, Erasmus University

Contents

Abstract

Decision support within transport companies should not only use traditional objective functions, but also reason about qualitative effects on all involved actors. We propose a fuzzy-logic rule base that can be used in addition to traditional operations research tools to calculate not just optimal solutions, but solutions that are optimal with respect to knowledge about the preferences of and long-term effects on customers, employees, and the environment.

We propose a fuzzy logic based judge module that is capable of evaluating logistical performance considering all parties involved in the act transporting a container. It is based on measurements of selected key performance indicators that are fuzzified and combined into satisfaction scores of customers, employees and society.

Our proposed method not only enables the continuity of the quality of planning by storing and maintaining valuable expert knowledge, but can also explain decisions based on this knowledge.

Keywords

Decision support systems, vehicle routing, logistics, fuzzy logic

1 Introduction

The current best practice of planning in a logistics company consists of human planners interacting with an order-assignment database and a tool for calculating optimal routes. The most advanced logistics tools focus on balancing well-defined goals such as minimizing the number of used vehicles, the distance covered, and the number of empty kilometers driven. The task of the human planners is to make choices, often by imposing additional constraints on the optimization tool, such that all involved parties are more or less satisfied with the schedule. This process is time-consuming, but at some point execution has to start, and thus decisions have to be made. Consequently, the resulting schedule is generally believed to be sub-optimal, and, given more time, better schedules could have been found.

In larger companies, it takes up to three years to train human planners to make the right choices. They need to learn the preferences of all drivers, customers, and of the management of the company, they need to weigh all these against each other, and they need to understand the long-term effect of their choices. Losing one of these planners often means losing valuable information, and starting up another intensive training process.

To summarize, current best practice has two problems: it is time-consuming, resulting in potentially sub-optimal decisions and it heavily depends on the experience of the planning experts, making the quality of the core process of logistics companies very vulnerable to changes in personnel.

In this paper we propose a method for storing the knowledge of these experts. This knowledge can then be leveraged by automated reasoning to expedite and improve the quality of the daily planning tasks. The task for the human planners then shifts towards giving feedback to the system, by selecting options, and maintaining the knowledge database. This paper, focused on capturing and leveraging knowledge for logistics planning, is an outgrowth of earlier work focused on capturing performance measurement in terms of stakeholder satisfaction (Srour *et al*, 2007).

There has been quite some work done already on modeling such logistic knowledge. In the next section we will briefly summarize this work, and discuss how we can use these results. In Section 3 we then argue why we chose fuzzy logic to store such knowledge, and show how this can be done. Subsequently, in Section 4, this method is applied to a specific logistic case.

2 Related work

There are two separate, but related, streams of research that support this work. The first is that of performance evaluation. This stream explores questions of how to extract and measure metrics important to logistics performance. The second stream is that of how to convert hard metrics into human readable control mechanisms. This second stream (the topic of Section 3) focuses on elements of artificial intelligence; specifically, fuzzy logic.

The human planners in a logistic company have to consider many financial and non-financial results. A balanced scorecard (see, for example, the work by Brewer (2000)) is a formalized mechanism to achieve this. The development of the balanced scorecard must, however, be undertaken by managers and is thus capturing only a snapshot of management's perceived critical measures.

Kaplan and Norton (1996) extend the balanced scorecard concept to capture the relationship between performance measures and core activities. They propagate the balanced scorecard as a means to translate a strategy into a comprehensive and useful overview of business performance. They consider a strategy to be a set of hypotheses about cause and effect, such as: "In case we improve our logistics services, our customers will be more satisfied." Ultimate effects usually are in terms of financial performance, such as cash flow. A well constructed balanced scorecard should incorporate those performance measures that represent critical cause and effect relationships.

However, just a list of such hypotheses is not enough. Krauth *et al*. (2005) argue that there is a need to represent multiple points of view, because of conflicting desires by stakeholders such as the owners, the customers, the employees and the society. A basic example is that of a service provider who prefers to charge high prices and deliver a low-cost service in contrast to a customer who desires a lower price and a high-quality service.

Our model of dispatching expertise is designed using these various recommendations from the literature. Therefore, the first step is to identify all of the stakeholders. To facilitate this step we utilize the stakeholder categories described in Krauth *et al*. (2005). These categories are *management*, *employees*, *customers*, and *society*. Within these large classifications there may be one or more specific stakeholder groups that should be specifically designated. For example, in the category of employees there may be multiple stakeholder groups such as drivers, maintenance personnel, dispatchers, etc.

Next, a proper understanding of the cause and effects between hard metrics and satisfaction measures is critical. For example, consider the cause and effect relationship: "If my trucks arrive on time at the customer's, my customer will be satisfied". This cause and effect relationship proposes a positive relationship between timeliness of trucks at the customer site and customer satisfaction. Most people will agree with the proposal, but as such it does not provide a managerial lever to improve customer satisfaction through improving truck timeliness. There are several ambiguities present in the cause and effect relationship as indicated.

First, the customer could have measured the timeliness of trucks in terms of average amount of minutes too late, where minutes too early are neglected. The customer may just as well base her satisfaction on an extreme case or on the delivery last week. In other words, the customers' appreciation may be associated with a derived aggregate measure from the set of measurements of all deliveries, explaining her level of satisfaction. Secondly, the level of satisfaction itself is a "soft" measure that may be expressed in an ordinal scale. Each individual may respond in a different way, not only based on the actual state (e.g. level of satisfaction), but also on the understanding of the measure (e.g. what "good" stands for).

In the next section we show how fuzzy logic can help handle these ambiguities.

3 Model

The framework we propose is designed to capture the knowledge of logistical planners about their business partners, colleagues and environment. In their struggle to balance the requirements of all parties, they can use physical measurement data describing the current situation (length of routes, arrival times, etc). Our model also uses hard metrics and from these extracts soft measures representing the satisfaction of each party involved in the business process (the length of the routes is good, the number of late deliveries is too high, etc). These satisfaction measures are described via a satisfaction evaluation component built for each of the three classes of drivers (more generally this could be employees), customers, and society. Note that "managers" are treated as a special case given their specific business rules. Generally, it is difficult to model satisfaction of all these parties properly, due to vague (verbal) boundaries of evaluation classes (e.g. "good", "ok", "bad" service). To tackle such issues, we propose the use of fuzzy logic for the design of these satisfaction measures.

3.1 Fuzzy logic

Fuzzy logic applies fuzzy set theory in the design of control systems. Originally it was proposed by Lotfi Zadeh (Zadeh, 1965) for data processing and it gained wide recognition through successful applications only later. Admittedly, early uses of fuzzy logic were limited. However, with the advent of modern computing, fuzzy logic has seen an increase in attention by the research community. One example of an early application of fuzzy logic may be seen in the 1977 paper by Bass and Kwakenaak, in which they apply fuzzy logic to the problem of choosing between multiple alternatives. There are several papers or books that introduce fuzzy logic in general (e.g. Bernardinis, 1993, Cox, 1992, Nguyen and Walker, 1997). Here we introduce only the core concepts that are necessary to understand our contribution.

In fuzzy set theory variables can be partial members of multiple sets. A variable, depending on its value, can be a member of a set to a certain degree between zero and one. This partial membership is described by a fuzzy variable (also known as a linguistic variable [Zadeh, 1973]); the value of the fuzzy variable expresses how much the variable is a member of the given set. For example the variable 'number of late deliveries' with a value of 5 may be considered to be a member of the set 'good' to a degree of 0.7. Thus the linguistic or fuzzy variable 'number of late deliveries is good' has a value of 0.7. A variable can be a member of multiple sets but the sum of its membership values does not have to sum up to one (in contrast to probability theory). For example, the variable 'number of late deliveries' with the same value is also member of the set 'bad', but only in a degree of 0.2. Thus the fuzzy variables 'number of late deliveries is good' and 'number of late deliveries is bad' have values 0.7 and 0.2 respectively.

Fuzzy variables can be combined by fuzzy rules to yield other fuzzy variables. The rules are defined in terms of linguistic variables and logical operators (AND, OR, NOT) in the form of IF ... THEN. Whenever a linguistic variable changes its value it impacts all the rules it is part of. The result of these rules in turn changes the values of other variables (notably those in the THEN clause), which in turn may impact other rules. For example, we can construct the following rule: IF ‘number of late deliveries’ IS good AND ‘number of early deliveries IS ‘good’ THEN customer is happy. In general, it is possible to express arbitrary complex inter-relations of variables at the cost of lengthy evaluation iterations.

It is possible to convert fuzzy variables to *crisp* (non-fuzzy) variables and back. The conversion from crisp to fuzzy (as shown above) is called *fuzzification* and from fuzzy to crisp is called *defuzzification*. Defuzzification, as opposed to fuzzification, takes a couple of linguistic variables and converts them into a single crisp value. These processes provide interfaces to connect the fuzzy rules into the embedding system.

This triumvirate of fuzzification, rules and defuzzification constitutes the three steps of fuzzy logic that aims to process some input variables in order to produce some output values. In our case we fuzzify key performance indicators of logistical performance (over one day of commercial transportation). We compute their membership value in the sets “good”, “bad” and sometimes “ok”. These fuzzy variables, obtained from the key performance indicators, are then combined by fuzzy rules to yield satisfaction measures for customers, drivers (employees), and society. These satisfaction variables are defuzzified in the end to provide the evaluation scores of logistical performance.

To summarize, fuzzy logic offers several unique features that make it a particularly good choice to capture planners’ knowledge.

1. It is inherently robust since it does not require precise, noise-free inputs. The output is a smooth function despite a wide range of input variations.
2. Since it processes user-defined rules, it can be modified and tweaked easily to improve or drastically alter system performance. New inputs can easily be incorporated into the system simply by generating appropriate governing rules.
3. Because of the rule-based operation, any reasonable number of inputs can be processed and numerous outputs generated, although defining the rule base quickly becomes complex if too many inputs and outputs are chosen since rules defining their interrelations must also be defined.
4. Fuzzy logic can express nonlinear systems that would be difficult or impossible to model mathematically.

The appeal of fuzzy logic for us is the ability to express verbal description of business rules, and the inherent non-linearity that mimics human reasoning. In the following we describe the three main components of our fuzzy system for logistics performance evaluation. We start (as noted in Section 2) by identifying all the stakeholders and the performance indicators they care about. We then define fuzzy sets and membership functions to obtain the soft measures from the performance measures (section 3.2). In section 3.3, we show how fuzzy rules can be set up to combine the soft measures into satisfaction variables. Finally, in section 3.4 the defuzzification process in described.

3.2 Fuzzification

As described in section 2 we classify the stakeholders into three groups: customers, employees (drivers) and society. Once all stakeholder groups have been identified, a list of key performance indicators is constructed for each stakeholder group. These performance indicators represent key factors on which stakeholders are judging their satisfaction with the vehicle routing or logistics performance. The KPIs are such that they can be quantifiably measured (extracted or derived) from the output of the company's operations. Note that each stakeholder group has one or more KPI – if no logical KPI could be defined for a stakeholder group then that stakeholder group was removed from the set of stakeholders. Examples of KPIs include the number of late deliveries for customers, number of places visited that were not listed as preferred locations for drivers, and overall truck utilization for society.

To fuzzify the identified KPIs, fuzzy sets (or corresponding fuzzy or linguistic variables) need to be defined. Fuzzy sets in effect allow the translation of a quantified metric into a verbal description of performance or satisfaction – i.e. "good", "ok", or "bad". Making this translation allows the meaning of a hard metric to be captured in a unique manner for each stakeholder. For example, considering the case of schedule deviation, in the context of order delivery, some customers may rate their satisfaction as "good" if an order arrives two minutes late. However, if the order is ten minutes late they may consider their satisfaction to be "bad"; but in the case of five minutes late – this may fall into the verbal "grey area" somewhere between "good" and "ok". Thus, fuzzy sets must be carefully constructed for each stakeholder and KPI combination. The identification of KPIs and the definition of fuzzy sets for each selected KPI require expert knowledge and it is far from being an automated process. Nevertheless, it partly encodes the planners' knowledge about their business environment, and as such is progress towards achieving the main goal of this study.

3.3 Fuzzy Rules

After fuzzy sets have been constructed, fuzzy rules must be defined to merge all "fuzzified" metrics into fuzzy measures of satisfaction per stakeholder group. A fuzzy rule has the form of IF 'linguistic variable' AND/OR 'linguistic variable' AND/OR ... THEN 'linguistic variable'. In our model, the variables in the IF clause are the fuzzy variables derived from the KPIs. They refer to a certain computed performance indicator value to be 'good', 'bad', or 'ok'. The variables in the THEN clause express the following six satisfaction measures: 'customers happy' or 'customers unhappy', 'drivers happy', 'drivers unhappy' and 'society happy', 'society unhappy'.

In theory it is possible to consider all combinations of variables, but it leads to an exponential explosion of the number of rules. In our model the set of rules that bridge the fuzzy variables based on the KPI measures and the satisfaction measures are defined by expert knowledge. Logistical planners can assert basic business rules they consider during planning and these assertions can be translated into linguistic variables and fuzzy rules in a straightforward manner. This property is the main reason why we promote the usage of fuzzy logic in modeling experts' knowledge. The rule base can easily be verified, modified or extended by human interaction, which makes the system implementing this model flexible and human friendly.

3.4 Defuzzification

In order to evaluate logistical performance by the six satisfaction variables described in section 3.3, they have to be defuzzified. Defuzzification converts fuzzy values to crisp values; in practice there are several methods by which to do this. In our model we suggest using of a method that result in a continuous value (center of area, center of gravity), so that the results can be linearly combined.

The satisfaction variables are defuzzified in pairs. ‘Customer happiness’ and ‘customer unhappiness’ is converted to a customer satisfaction score, ‘driver happiness’ and ‘driver unhappiness’ are defuzzified into employee satisfaction and ‘society happiness’ and ‘society unhappiness’ yield a society satisfaction score. The three satisfaction scores are continuous real values between zero and ten, which can be linearly combined with similarly scaled managerial scores (cost, profit, etc scores).

4 Application

Presently, researchers from RSM Erasmus University, TU Delft, Free University Amsterdam, the Centre for Applied Mathematics and Computer Science (CWI) and Almende BV are working together with industrial partners Post-Kogeko, Vos Logistics, and CarrierWeb on the application of agent-based technologies to the vehicle routing problem. Specifically, decision support systems are being developed to support the transport of containers over the road by the logistics service provider (LSP), Post-Kogeko. Post-Kogeko is a mid-size LSP active in several sectors, one being the transport of import and export containers, generally of the merchant haulage type. Post-Kogeko has a fleet of around 40 trucks active in this sector, handling around 100 customer orders each day.

The process of executing an order starts with the reception of an order, generally one day before required execution. An order is the request from a customer to Post-Kogeko to pickup a container at a container terminal (in case of an import container) and transport it to the customer, with delivery within a certain time window. Arriving at the customer requested location, the container is then unloaded, and the empty container is brought back to the same or another container terminal or empty depot – depending on the contract the customer has with the ocean carrier or shipping agent. This concludes the order, and the truck is ready for its next order. The process is similar for export containers, except that an empty container is picked up at the terminal, it is filled at the customer site (instead of emptying it) and the full container is returned to the terminal. What complicates matters is that not all containers are available at the start of operations early in the morning: either they have not physically left the ship yet, or they are delayed for administrative reasons – often due to an unsettled payment or customs. Post-Kogeko can only transport containers that have been released, and are allowed to leave the container terminal. For this reason it is hard to optimize the system in a traditional sense, since not all information is known beforehand, and will only become available sometime during the day. Large variation in the daily work load, (i.e., the number of orders per day and the distance to travel per order), complicates the planning process.

One of the first steps in developing a decision support system for this problem is the development of an evaluation module that can judge the performance of one day’s execution. Our judge module implements the fuzzy based model outlined in section 3 to capture experts’ knowledge to mimic planners’ evaluation. It requires a database of orders and execution records for a given day. The database should contain information on the orders that were due on the given day and data on exactly how the execution was carried out (when were the containers were delivered, by which truck, etc). Using this database the judge module computes a list of KPIs (section 4.1) fuzzifies them by the given membership functions (section 4.2) applies the rules (section 4.3) and finally defuzzify them (section 4.4). The result of all this process is three satisfaction score for customers, drivers and society.

4.1 Selected Key Performance Indicators

According to our model, key performance indicators reflecting the preferences of four stakeholder groups are defined. Managers have a set of well-defined measures that relate to cost and profit:

1. Empty distance traveled
2. Profit per delivery
3. Profit per kilometers

Drivers care about how long of a rest they can have between two jobs (KPI 4), how many times they are interrupted during their ongoing activity to be sent somewhere else (KPI 5), and if they are sent to addresses, cities or countries they do not like (KPI 6):

4. Driver idle time
5. Number of plan deviations
6. Geographical range

Customers are mainly interested in the quality of service they receive. They usually do not like early or late deliveries (KPI 7), sometimes they do not like specific drivers (KPI 8), and they definitely dislike when their orders are rejected (KPI 9) by the transportation company:

7. Schedule deviation
   a) Number of early deliveries
   b) Maximum span of consecutively early deliveries (the maximum number of early deliveries that happened in a series at the customer's site)
   c) Total minutes of earliness attributable to the maximum span of consecutively early deliveries
   d) Number of late deliveries
   e) Maximum span of consecutively late deliveries
   f) Total minutes of lateness attributable to the maximum span of consecutively late deliveries
   g) Maximum span of consecutively on-time deliveries
8. Driver serving each customer
9. Number of jobs rejected

Finally, society is expected to be interested in environmental issues and traffic conditions. All these point to the direction of as few trucks on the road as possible, while maintaining standard of living. We define this as:

10. Capacity utilization

4.2 Membership Functions

To fuzzify the KPIs, we define linguistic variables and membership functions for each one of them. Due to their special business rules, we do not fuzzify the managers' performance indicators, only those of the drivers, customers and the society. At least two linguistic variables are defined for all KPIs. These are 'good' and 'bad'. Additionally, for the driver idle time (4) and the number of plan deviations (5), as well as for the capacity utilization (10) a third linguistic variable is defined, which is called ‘ok’. Note that all KPI values are normalized to have a value between 0 and one. This makes it easier to define the membership functions. Membership functions to fuzzify KPI 7.a, 7.b, 7.c, 7.d, 7.e, 7.f and 9 are depicted in ***Figure 1***.

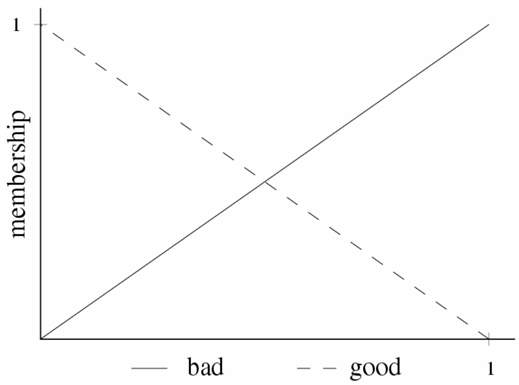

Figure 1: Membership functions with two linguistic variable

A KPI value of 0.6, for instance, defines a membership value of 0.6 for the ‘bad’ linguistic variable and a membership value of 0.4 for the ‘good’ linguistic variable. Membership functions for KPI 6, 7.g and 8 are very similar to those in ***Figure 1***, except that the linguistic variables are assigned to the functions the other way around.

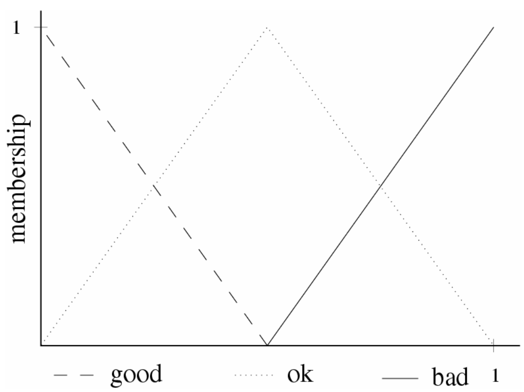

Figure 2: Membership functions with three linguistic variable

KPIs 5 and 6 are fuzzified by the membership functions on ***Figure 2***. Here a given value of a KPI yields in three membership values for the ‘bad’, ‘ok’ and ‘good’ linguistic variables. To fuzzify KPI 4, the same functions are used, but the ‘good’ and ‘bad’ variables are exchanged.

4.3 Fuzzy Rules

Our goal is to convert the fuzzy variables obtained by fuzzifying the KPIs to other fuzzy variables describing satisfaction of customers, drivers and society. For every stakeholder group we define two fuzzy variables: ‘happy’ and ‘unhappy’ and for each of the six satisfaction variables there is a rule constructed. A rule combines linguistic variables (fuzzified KPIs) by AND and OR operations. AND operations take the minimum of the two arguments, while OR operations take the maximum. The results of the rules are the ‘happy’ and ‘unhappy’ fuzzy variables for every stakeholder group. Following are the applied rules in the notation of the *Fuzzy Control Language* (FCL) as defined in IEC 1131-7 CD1.

Driver satisfaction:

1. **IF** (driver idle time **IS** good **AND** plan deviations **IS** good **AND** geographic range **IS** good) **OR** (driver idle time **IS** ok **AND** plan deviations **IS** bad **AND** geographic range **IS** good) **OR** (plan deviations **IS** ok **AND** geographic range **IS** good) **THEN** driver happy.
2. **IF** (driver idle time **IS** bad **OR** plan deviations **IS** bad **OR** geographic range **IS** bad) or (plan deviations **IS** ok **AND** geographic range **IS** bad) **THEN** driver unhappy.

Customer satisfaction:

3. **IF** (maximum span of consecutively on-time deliveries **IS** good **OR** [number of early deliveries **IS** good **AND** maximum span of consecutively early deliveries **IS** good] **OR** total minutes attributable to maximum span of consecutively early deliveries **IS** good) **AND** (maximum span of consecutively on-time deliveries **IS** good **OR** [number of late orders **IS** good **AND** maximum span of consecutively late deliveries **IS** good] **OR** total minutes attributable to maximum span of consecutively late deliveries **IS** good) **AND** drivers serving **IS** good **AND** jobs rejected **IS** good **THEN** customer happy.
4. **IF** number of rejected jobs **IS** bad **OR** drivers serving **IS** bad **OR** (drivers serving **IS** good **AND** [number of late deliveries **IS** bad **OR** maximum span of consecutively late deliveries **IS** bad **OR** maximum span of consecutively on-time deliveries **IS** bad **OR** total minutes attributable to maximum span of consecutively late deliveries **IS** bad]) **OR** (number of rejected jobs **IS** good **AND** [number of late deliveries **IS** bad **OR** maximum span of consecutively late deliveries **IS** bad **OR** maximum span of consecutively on-time deliveries **IS** bad **OR** total minutes attributable to maximum span of consecutively late deliveries **IS** bad]) **OR** (number of early deliveries **IS** bad **OR** maximum span of consecutively early deliveries **IS** bad **OR** maximum span of consecutively on-time deliveries **IS** bad **OR** total minutes attributable to maximum span of consecutively early deliveries **IS** bad) **THEN** customer unhappy.

Society Satisfaction:

5. **IF** capacity utilization **IS** good **OR** ok **THEN** society happy.
6. **IF** capacity utilization **IS** bad **THEN** society unhappy.

4.4 Defuzzification

Pairs of satisfaction variables are converted to crisp satisfaction scores. Linguistic variables ‘driver happy’ and ‘driver unhappy’ are converted to a driver satisfaction score, ‘customer happy’ and ‘customer unhappy’ to a customer satisfaction score and ‘society happy’ and ‘society unhappy’ to a society satisfaction score. All scores are continuous real variables in the interval [0, 10]. To achieve a continuous domain the center-of-area defuzzification method is used with the membership functions depicted on ***Figure 3***.

Using this method, first the triangles representing the fuzzy values are scaled down (in height) proportional to the actual values of Happy and Unhappy. The point that divides the combined area of the two discounted triangles equally is returned as the result. Note that the satisfaction value will always be a real value between 0 and 10. In the event that either Happiness or Unhappiness is zero, the satisfaction score will be either 0 or 10 respectively - regardless of the value of the other variable.

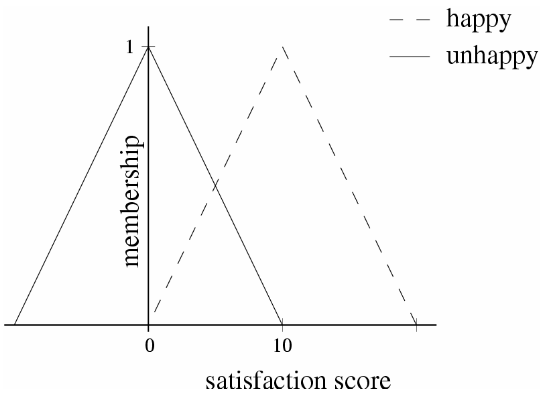

Figure 3: Membership functions to use with the center-of-area defuzzification method

5 Conclusions

This paper demonstrates the need to move beyond traditional vehicle routing approaches by including expert knowledge in the process. We recommend a mechanism, premised on fuzzy logic, to capture this expert knowledge. We further demonstrate the exact specification of this model as applied to a real-world case in container logistics.

Given the capability of this approach to capture the satisfaction of various stakeholders within the logistics industry we are confident that this approach can enhance vehicle routing practices. In addition to ensuring that the routing is efficient and cost effective, this methodology for measuring satisfaction can also ensure that the routing is also promoting the logistic company’s business position with their stakeholders.

Furthermore, in using multi-agent systems to find solutions to the vehicle routing problem (as is underway at the universities involved in this study) the quality of the solution is highly dependent on the ability of the agents to learn from past performance. At the crux of this work is the need to accurately evaluate performance. As indicated in the beginning of this work, we believe that while logistics performance is highly dependent on traditional measures (such as empty distance traveled), it is also subject to measures not commonly studied (such as satisfaction).

Utilizing fuzzy logic as a means to derive stakeholder satisfaction yields a highly tractable mechanism by which to measure performance and the factors influencing performance. Furthermore, the control mechanisms afforded by fuzzy logic is a promising technique for machine learning and control in multi-agent systems for vehicle routing. The next steps for this research fall along both of these streams. First, we wish to use the framework described here as a means by which to compare the quality of the solutions emerging from the different multi-agent systems under development at RSM Erasmus University, TU Delft, Free University Amsterdam, the Centre for Applied Mathematics and Computer Science (CWI) and Almende BV. Second, we wish to embed the fuzzy logic framework as a control mechanism within the vehicle routing agent systems.

Acknowledgements

This work is funded by Novem via the DEAL project and also partly by Transport Research Center Delft. The research by Mathijs de Weerdt is supported by the Technology Foundation STW, applied science division of NWO, and the technology program of the Ministry of Economic Affairs. The research is conducted at Almende BV, at RSM Erasmus University, and at TU Delft.

References

Bass, S.M., and H. Kwakenaak (1977) Rating and ranking of multiple aspect alternatives using fuzzy sets, in *Automatica*, Vol. 13, pp. 47–58.

Bernardinis, L.A. (1993) Clear Thinking on Fuzzy Logic, *Machine Design*, April 23, 1993.

Brewer, P.C. and Th.W. Speh (2000) Using the Balanced Score-card to Measure Supply Chain Performance, *Journal of Business Logistics*, Vol. 21, No. 1, pp. 75–93.

Cox, E. (1992) Fuzzy Fundamentals, *IEEE Spectrum*, pp. 58-61.Kaplan, R.S. and D.P. Norton (1996) Linking the Balanced Scorecard to Strategy, *California Management Review*, Vol. 39, No. 1, pp. 53–79.

Krauth, E., H. Moonen, V. Popova and M. Schut (2005) Performance Measurement and Control in Logistics Service Providing, *The ICFAIAN Journal of Management Research*, Vol. 4, No. 7. August 2005.

Nguyen, H.T. and E.A. Walker (1997) *A First Course in Fuzzy Logic*, CRC Press, NY.

Srour, F. J., T. Mahr, M. de Weerdt, R. Zuidwijk, H. Moonen (forthcoming 2007) Performance Evaluation of a Networked Enterprise: Balancing Local Objectives and Network Relations, to be published in, *Title pending*, Springer books, Berlin.

Zadeh, L.A. (1965) Fuzzy sets, in *Information Control*, Vol. 8, pp. 338–353.

Zadeh, L. A. (1973) Outline of a new approach to the analysis of complex systems and decision processes, *IEEE Transactions on Systems, Man, and Cybernetics*. Vol. SMC-3, pp. 28-44.