Explaining the Impact of Actions☆

Johan Schubert, Pontus Hörling
Department of Decision Support Systems
Division of Information and Aeronautical Systems
Swedish Defence Research Agency
SE-164 90 Stockholm, Sweden
schubert@foi.se
http://www.foi.se/fusion

***Abstract*—In this application oriented paper we develop information fusion explanation functions for simulation-based decision support for evaluation of military plans in expeditionary operations. The explanation function is based on a sensitivity analysis on the impact of different actions upon the success of the plan where we systematically vary the alternatives of each action of the plan, one action at a time, keeping all the other actions unchanged in a series of simulations. This sensitivity analysis shows the relative level of importance of making the correct selection of alternative for each action. Using the explanation function, a decision maker can find the most important actions of a plan and focus his attention on actions where successful decision making is crucial to the success of the entire plan.**

Keywords-decision support; simulation; information fusion; explanation functions, Effects-based Approach to Operations.

I. Introduction

In this application oriented paper we develop information fusion explanation functions for simulation-based decision support for evaluation of military plans in expeditionary operations. We have developed a simulation-based decision support methodology with which we can test operational plans as to their robustness [1][2][3]. Primarily, this methodology highlights the dangerous options in an operational plan, leaving the decision maker free to focus his attention on the set of remaining robust plans. By using a decision support tool, a decision maker is able to test a number of feasible plans against possible courses of events and decide which of those plans is capable of achieving a desired military end state. The purpose is to evaluate plans and understand their consequences through simulating the events and producing outcomes which result from making alternative decisions regarding actions. Each plan consists of many actions, where several actions can be performed in a number of alternative ways. When a good plan has been found by the system we may use the explanation functions derived in this paper to find the actions within the plan with the highest impact on the success of the plan.

Actors and actions are modeled using a scenario used by the Swedish Armed Forces in their Combined Joint Staff Exercises. The actions of the plan are simulated together with all actors and their reactions on our planned actions, and their possible follow-on interactions. As the actions may each have several different alternatives in which manner they can be carried out, together they span-up an action tree. The tree is searched by an A^*-algorithm [4][5] where each level in the tree is an action and each node in the tree is an alternative for an action. As the action tree is searched, each node (i.e., sequence of alternatives leading from the root of the three to this node) is evaluated by the simulator and results are stored and communicated to the decision support side of the system. By using an A^*-search to guide the tasks of the simulator we let the simulator work in a manner to achieve maximum information value gain. In addition, a simulation control interface lets the decision maker put constraints on the search, in order to simulate actions within his area of interest [1].

How we model a phenomenon depends on the purpose of the model and the questions we want to answer. Since our simulation system aims to support decision-making within an Effects-based Approach to Operations (EBAO) [1] [6][7] the modeling has to be based on EBAO and the concepts used within it, such as plan, action, effect, end state, etc.

The planning process we develop corresponds to the selection of a subset of actions which are chosen from a set of alternative actions. A chosen combination of alternative actions constitutes a plan. The number of plans can theoretically grow very large since each permutation of alternative actions will constitute a separate plan.

By systematically varying one action at a time keeping all the other actions unchanged in a series of simulations, we are able to perform a sensitivity analysis for each action in the plan based on the change in evaluation score of the plans. This sensitivity analysis shows the relative level of importance of making the correct selection of alternative for each action. Using the explanation function, a decision maker becomes informed as to which actions of the plan are crucial to its success.

In section II we describe the scenario of an expeditionary operation. In section III we describe the military plans under investigation and the A^*-search through simulation increments of alternative plan instances. We develop information fusion-based explanation functions for analyzing the impact of different actions of any plan on the ability of that plan to achieve the sought after military end state (section IV). In order to have full data input to the explanation function we perform

[1] EBAO is a military approach for the management and implementation of efforts at the operational level.

☆This work was supported by the FOI research project "Real-Time Simulation Supporting Effects-Based Planning", which is funded by the R&D programme of the Swedish Armed Forces.

additional simulations around the original plan instances where only one action at a time is changed compared to the original plan (section V). In section VI we analyze the impact of all actions using the derived explanation function for some of the best original plan instances found by the simulation. These sections IV–VI are the core of the paper. Finally, conclusions are drawn (section VII).

II. Scenario

We make use of the same scenario that has regularly been used by the Swedish Armed Forces in their Combined Joint Staff Exercises. The scenario comprises several fictitious countries. Background histories offer explanations to why and how sentiments, stances, identities, loyalties, economic dependencies and inequalities have evolved over time, occasionally resulting in shifts of power. Phenomena that are commonly found in conflict areas and post conflict areas have been embedded in scenario contexts that make the origins of the phenomena plausible.

In Bogaland, a newly industrialised country, a civil war broke out ten years ago when discontent within the minority ethnic-religious group had reached very high levels. The root cause was increasing social stratification caused by what members of the minority group perceived as unjust distribution of revenues from a natural resource located in an area populated by the minority group. The civil war put an end to the exploitation of the resource, in this case oil, and revenues dropped to very low levels. The country was split into two parts, roughly along ethnic lines, with each part having its own government. A post-war economy evolved over the next decade, and several irregulars and insurgents are now challenging the incumbent presidents.

The incumbent presidents have signed a peace-agreement, and an international force, BFOR, is present to support the implementation of the agreement. Irregular groups in Bogaland seek to preserve or increase their influence by undermining the efforts of BFOR, the governments or competing irregulars. Two of the neighboring countries have much at stake in the conflict, because of economic interests and shared identities with parties within Bogaland. Actors within these neighboring countries support irregulars in Bogaland.

This is the scenario that we use in simulation of alternative plans for the BFOR force. The scenario consists of participating actors, their initial states and probability distribution for different actions, environmental data, as well as the plan that is to be evaluated. The aim is to find a plan that drives the initial state of the scenario towards a predefined military end state. Specifically in this paper our focus is on finding the most important actions within a successful plan for this scenario by using explanation functions.

III. Event-based Simulation of Plans Using A^*

A. Plans

A *plan* as it is defined in the context of EBAO is a sequence of *actions* that together leads to a desired *end state* which is set by a military force.

A typical plan instance P_1 is

[1 2 41 61 6 9 23 12 4 50 51 52 16 63 55 56 72 76 67 85 79 80 94 95 96 105 31 57 46 47 5460.2 4448.2 1012.0]

where all but the last three numbers in this sequence is the number of the selected alternative for each action in this plan instance. For example, action number 3 (i.e., position 3 in the sequence) takes alternative number 41. Note that alternatives for different actions are numbered with running numbers in no particular order; they do not restart at 1 for each new action. A full plan instance is a path from the root of tree down to one particular leaf. Obviously, the depth of the tree is the length of the sequence minus three (i.e., not counting the f, g and h estimates). Plan instance P_1 above corresponds to a sequence of 30 specific simulations where the actions take the numbered alternative listed in the sequence as its input parameter [1].

The last three parameters are different evaluation measures called f, g and h ($f = g + h$). They are distance measures calculated from changes in the scenario state and used in the A^*-search algorithm.

Plan simulation is performed by the simulation engine. The engine basically contains an implementation of the A^*-search algorithm which uses the Monte Carlo principle for event based simulation.

B. Simulating action alternatives

The scenario consists of participating actors, their initial states and probability distribution for different actions, environmental data, as well as the plan that is to be evaluated. Furthermore the scenario contains an event list which consists of actions derived from the other actors' agendas, and spontaneous/natural events. The list is dynamic and changes during the course of the simulation.

Let's define the system state, S_n as the combination of all actors' state parameters and all environment parameters. Consider action A_n. It transforms system state S_n according to $S_n = f(S_{n-1}, A_n)$, in the time interval $[t_{n-1}, t_n]$. The implementation of A_n is rarely instantaneous. Instead, it is an interaction between our own action, other actors' agendas and response operations, and other external events. Hence, our function $f(S_{n-1}, A_n)$ is designed as an event-driven simulation model in order to manage the complex interactions in a transparent manner. The events in this case are: launching of actions (our own or any other actors' actions), an actor's observations of initiated actions, and occurrence of an external event.

During simulation an assessment is made of how well each action did perform. This is done by the function g, a function that measures the consequence of all performed actions as a distance from the initial state $S_{0,0}$ to the current simulated state S_{x,y_x} [1]. Function h is a heuristic estimate of the remaining distance from S_{x,y_x} to the end state.

We know that the goal of the simulation is to execute different plans and identify those plans that result in system states that are closest to our end state, i.e., has the shortest distance to it.

C. Searching among action alternatives

To find good combinations of alternatives for all actions of the plan we apply A^*-search. It means that, on the basis of a given system state, we simulate the effect of each alternative action in our plan, but only one step at the time. Doing so, for every alternative, we get a new system state whose distance to the desired end state is calculated. Given the alternative that is best, i.e., closest to our end state, we simulate possible subsequent alternative actions provided, but again only one step ahead in our action/event list. One of these alternatives leads to a condition that is closer than the others. However, it is possible that all the alternatives actually lead away from the target as seen by Figure 1.

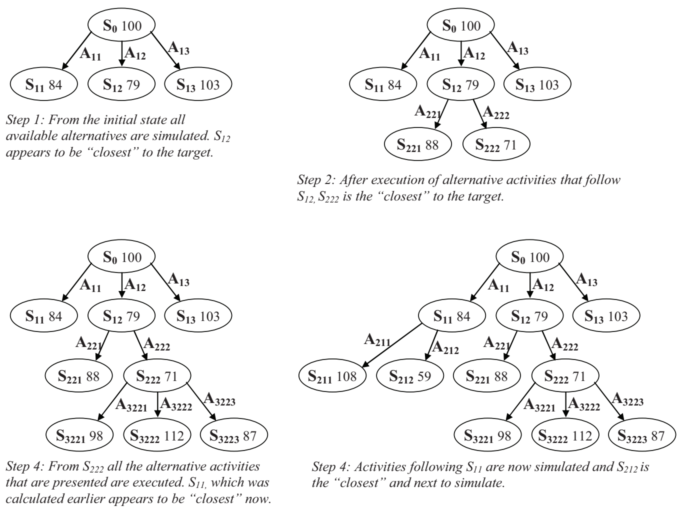

Figure 1. An example illustrating the four first steps in a simulation of a plan starting with initial system state S_0 with the *distance* of 100 to the desired end state. The available action alternatives A_x are executed successively in the currently most favourable plan option.

Therefore, we must also compare the new distance with the best of the distances that have been simulated and recorded in the previous simulation steps, but then had opted out in favor of a better sequence of alternative actions. The best sequence now becomes the basis for the next simulation step.

The distance from the initial state $S_{0,0}$ to a current state S_{x,y_x} is given by

$$g(y_x) = \sum_{i=0}^{x-1} \Delta\left(S_{i,y_i}, S_{i+1,y_{i+1}}\right). \tag{1}$$

We observe the difference in consequences between two plans. We compare the incremental changes of g called Δg as each plan P_i and P_j progresses down the sequence of additional actions A_k, where

$$\Delta g(P_i.A_k) = g(P_i.A_k) - g(P_i.A_{k-1}) \tag{2}$$

and i and j are indices for different plan instances and k is the index for action. Thus, $P_i.A_k$ is a variable referring to the kth action of the ith plan. It takes an integer as its value that is the number of the alternative for this action, e.g., $P_1.A_3 = 41$ imply that action number 3 of plan number 1 performs alternative number 41.

The estimated distance from the current state to the end state is given by

$$h(y_x) = \Delta\left(S_{x,y_x}, S_e\right). \tag{3}$$

With the total distance from the initial state to the end state via the current state is

$$f(y_x) = g(y_x) + h(y_x). \tag{4}$$

This is the distance function that is minimized by A^*.

IV. Explanation Function

Strat and Lowrance developed explanation functions [8][9] within Dempster-Shafer theory [10]. With these functions they where able to find the impact on the belief of any proposition from every contributing input factor through a sensitivity analysis based on numerical differentiation. They defined the differentiation of belief as

$$\left.\frac{\partial \mathrm{Bel}(A)}{\partial \alpha_i}\right|_{\alpha_i=1} \approx \frac{\left[\mathrm{Bel}^{\alpha_i}(A)\right]_{\alpha_i=1} - \left[\mathrm{Bel}^{\alpha_i}(A)\right]_{\alpha_i=1-\varepsilon}}{\varepsilon}. \tag{5}$$

The idea is that for each proposition A of interest a discounting of the belief in A with a small ε is performed. Then the belief in proposition A is calculated twice, once with the undiscounted m_i, and once with the discounted $m_i^{\alpha_i}$. From the difference, the impact of m_i on Bel(A) can be calculated. This may be repeated for all other propositions.

We extend this idea of numerical differentiation of the belief in propositions to the value domain that we work with in this application; i.e., f, g, h. As we work with plans consisting of several actions A_k we like to find the impact of each action on the evaluation $\{f_{ikl}(P_i.A_k = l)\}_l$ of plan P_i. This impact can be denoted $\partial f_{ikl}/\partial A_k$. Given that we have a discrete set of evaluations $\{f_{ikl}(P_i.A_k = l)\}_l$ we approximate the differentiation as a normalized difference between $\max_l f_{ikl}(P_i.A_k = l)$ and the average of all $\{f_{ikl}(P_i.A_k = l)\}_l$. We have,

$$\left.\frac{\partial f_{ikl}(P_i.A_k)}{\partial A_k}\right|_l = \frac{\frac{1}{n_{ik}}\sum_{j=1}^{n_{ik}} f_{ikj} - f_{ikl}}{\frac{1}{n_{ik}}\sum_{j=1}^{n_{ik}} f_{ikj}} \tag{6}$$

where $n_{ik} = \left|\{f_{ikl}(P_i.A_k = l)\}_l\right|$ is the number of alternatives for $P_i.A_k$, and l is the index of the alternative.

As the variance in the in this measure can be large between different plans P_i we may choose to study box plots for a small number of good plans for each action A_k. For example, in section VI we will study box plots for the five best plans over all alternatives for averages of all actions A_k,

$$\left\{ \left. \frac{\partial f_{ikl}(P_i.A_k)}{\partial A_k} \right|_l \right\}_{i=1}^{5} . \tag{7}$$

This may be the first presentation to analysts investigating which actions of the plan is most sensitive to accurate selection of alternatives on the overall success of the execution of the plan.

V. Generating Additional Simulations

As described in section III, the A^*-search algorithm is intended to continuously deliver the best plans it finds concerning the probability of success in reaching the end state, reflected in the distance g from start to end state; the lower, the better. Each of these plans consists of a set of actions where the actions have several alternative ways of execution and a plan must choose one alternative from each of these actions. Some actions in the simulation turn out to be more important than others for plan success. In order to find out how much a plan relies on a selection of a certain action alternative for its success, one might compare a good plan P_i found by the A^*-algorithm with plans that are structurally similar to it in some respect. One way to do this is to compare P_i with neighboring plans that only differs from P_i in the selection of alternatives for one single multi-alternative action, see Figure 2. This corresponds to all neighboring plans having a Hamming distance [11] to P_i of exactly 1. Using the formalism from section III for plans with n actions, we have

$$\sum_{k=1}^{n} \begin{cases} 0, P_i.A_k = P_j.A_k \\ 1, P_i.A_k \neq P_j.A_k \end{cases} = 1 . \tag{8}$$

We will in the next section perform a sensitivity analysis for the 150 best plans found by A^*.

Some of these plans might be good and may already have been simulated and delivered by A^*, and hence with a bit worse g. Furthermore, some of them might be only partly simulated, still residing in the A^* open set. However, since the space of potential plans can be extremely large (tens of million), and the A^* open set can reach a very large size, we do not try to find those partly simulated plans in the open set, and simulate them to completion, or try to find them among the already fully simulated plans. Instead we simulate all neighbors to each good plan P_i already found with a variation compared to P_i of exactly one action alternative a time. For each action A_k, we simulate P_i where the selected alternative for action A_k is replaced by another alternative to A_k in the additional simulations. This is the set P_{ik} consisting of $|P_i.A_k| - 1$ neighboring plans where

$$\forall P_i, P_i.A_k : P_{ik}(P_i, P_i.A_k) = \left\{ P_j \left| \begin{matrix} P_i.A_m = P_j.A_m, \forall m \neq k \\ P_j.A_k \neq P_i.A_k \end{matrix} \right. \right\}_{j=1}^{|P_i.A_k|-1} \tag{9}$$

After having worked through all actions with alternatives, changing only one action at a time, we get as many neighboring plans to P_i as the total number of additional alternative actions, excluding the alternatives that are part of P_i itself. For a set of n actions there are $|P_i.A_k| - 1$ alternatives to an action A_k in addition to the one in P_i. We have a total of

$$\sum_{k=1}^{n} \left(|P_i.A_k| - 1 \right) \tag{10}$$

neighboring plans to be compared with P_i. In our analysis, we use g as the quality measure of a plan and investigate how it is affected by systematic variations of each action of the plan.

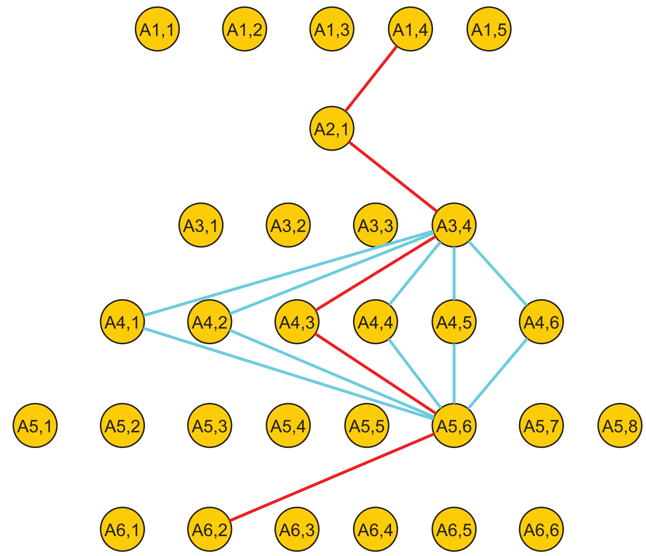

Figure 2. Conceptually, a plan P_i is a choice of alternatives for a sequence of actions, one for each consecutive action to be executed, like the red colored path. Each cyan colored path in this six-action planning problem corresponds to one neighboring plan out of $P_i.A_4 = \{1,2,4,5,6\}$ for alternative 3 of action 4.

VI. Analysing Results

In our example the planning problem is set up of 95 actions, similar to the 30 yellow nodes in Figure 2. Of these, 51 are grouped into 14 actions with 2, 2, 5, 4, 3, 2, 4, 2, 6, 8, 2, 7, 2, 2 alternatives per action, respectively. The remaining actions are all singleton actions without any alternatives that appear at some positions between the multi-alternative actions and always have to be executed. Singleton actions are not included in this analysis because of the lack of alternatives, and are hence not included in the figures in this section. A typical plan consists of one permutation of the alternatives of the 14 multi-alternative actions together with between 16 and 20 interleaved singleton actions. As described in the section V, the impact on plan quality when varying the alternatives for each multi-alternative action is the main outcome of this paper.

We perform simulations as described in section III and simulate 150 completed plans. "Completed" should be interpreted as the execution of all actions in a plan. Normally, we do not reach the military end state in these simulations, (corresponding to $h = 0$ in A^*) even though the plan is fully traversed. This corresponds to plans that do not have a combination of actions that, after completed simulation, are fully successful in obtaining the goal.

The evaluated plans have a g-distribution according to Figure 3.

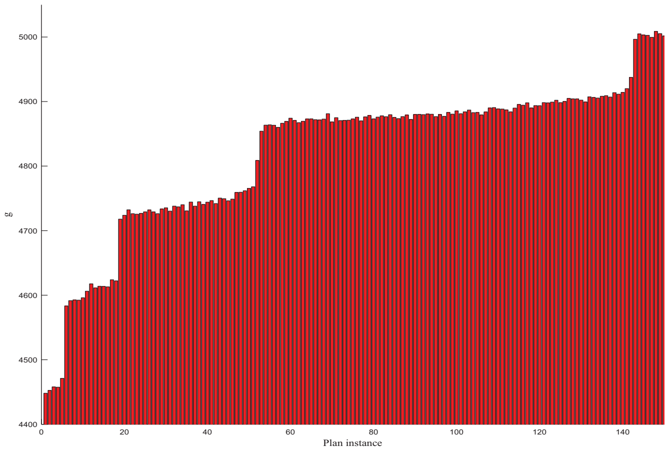

Figure 3. The g values of the 150 best plans found.

The steps in the g values are the result of additional singleton actions (one additional per step) being present in plans of lower goodness. The best found plan, delivered first by A^*, is represented by the leftmost, shortest, bar in the figure with $g = 4482.2$. Taking that plan as an example we can check how a variation of its multi-alternative actions affects g, see Figure 4. Each bar group in the figure contains the alternatives for each action not present in the original plan (red). For example, in Figure 4 we observe the third (cyan) bar group with alternatives numbered 3, 4, 5, and 6. These are four out of five alternatives for the third non-singleton action of the plan. Singleton actions are excluded in the figures. The fifth alternative for this action resides within the original plan (red) as the default for the third non-singleton action. Thus, Figure 4 shows the increase in g if one alternative for one action is changed from its default value in the original plan under investigation.

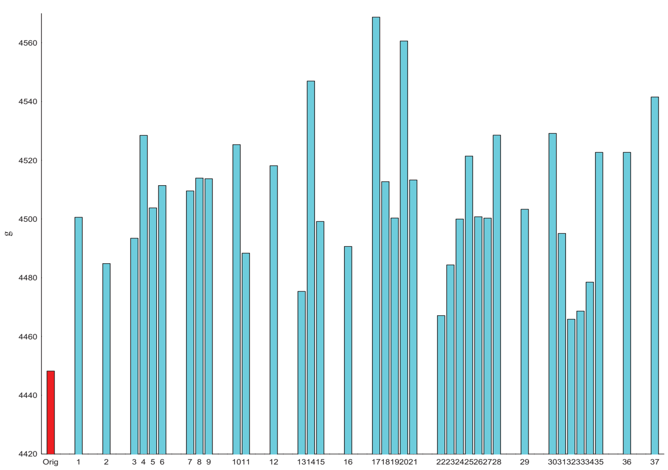

Figure 4. A bar chart showing the variation of g when varying the first plan (red) over the alternatives of each multi-alternative action (cyan). Only the variations of the first plan found by A^*-search are shown, the first plan itself is only shown once instead of being repeated in each bar group.

That analysis in Figure 4 was made for a single plan P_1. To get a hint of the variance of g values among different plans, we plot box charts for the five best plans, see Figure 5. They correspond to the five bars in the lowest step to the left in Figure 3.

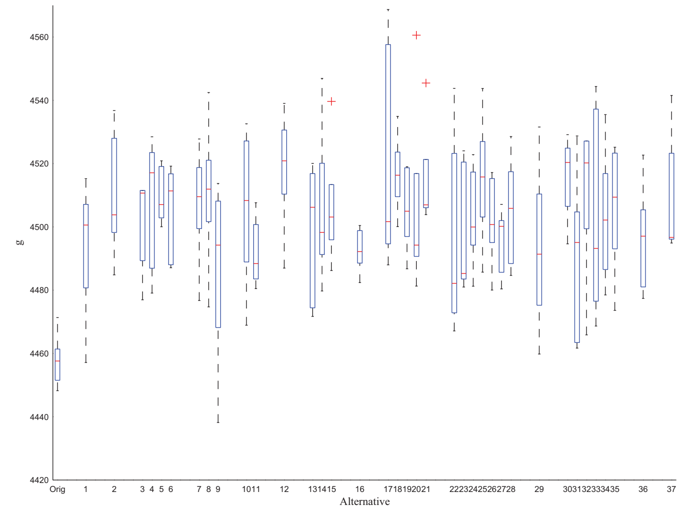

Figure 5. A box plot showing the variance of g when varying the alternatives of each multi-alternative action for the five best plans.

To get a feeling of the g-landscape for the best 150 plans from A^*-search together with their additional simulations varying one alternative at a time, a 3D bar plot is studied in Figure 6. The noise is mainly a result of the Monte Carlo simulation.

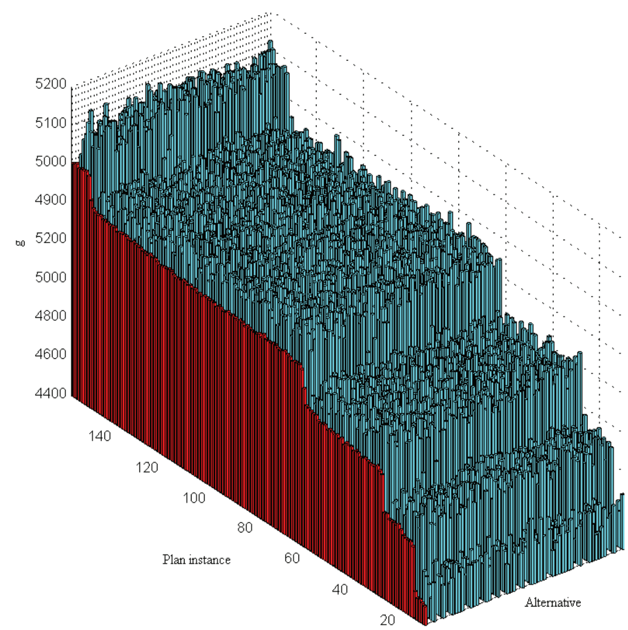

Figure 6. A 3D bar plot over the g values for 150 simulated plans. The red bars are the same as the bar chart in Figure 3 and the first row of bars along the "Alternative" axis corresponds to Figure 4.

However, the main interest of this study is the *sensitivity* of a plan as a function of the possible variations of its actions. The interest is two-fold; the performance of the A^*-search algorithm is checked since "shifting in" other alternatives to the actions found for good plans by A^*-search should give worse plans with higher g, that is a positive $\partial g/\partial A$. A high value of a certain $\partial g/\partial A$ also means that that the planning problem is more sensitive on the choice of alternative for that action which means that more importance should be payed to it in the planning process. When estimating P's sensitivity on the choice of alternative for a certain action, we use (6)(7), giving $\partial g/\partial A$ for an action. In the example, the variability of actions ranges between 1 and 8 alternatives. Observing the best plan P_1, the sensitivity (6) for different actions is plotted in Figure 7.

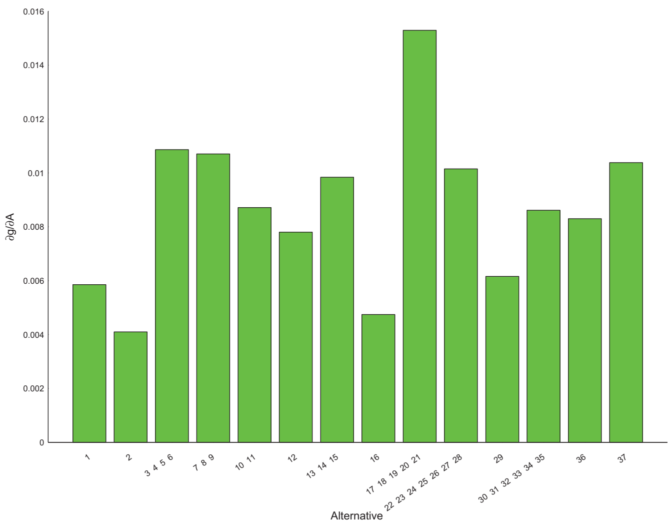

Figure 7. The sensitivity of the first plan P_1. Each bar represents the vaue of $\partial g/\partial A$ using (6) over all alternatives of each action.

As an example of the variance for action A_1 among different plans, the sensitivity is plotted for the best 150 plans $P_1.A_1$–$P_{150}.A_1$, Figure 8. For a few of the plans we observe that $\partial g/\partial A$ is negative which means that the neighboring plans on average gave a lower g value than the best plan found by A*-search. This is due to random effects of the Monte Carlo simulation of the subsequent additional simulation which introduces some noise.

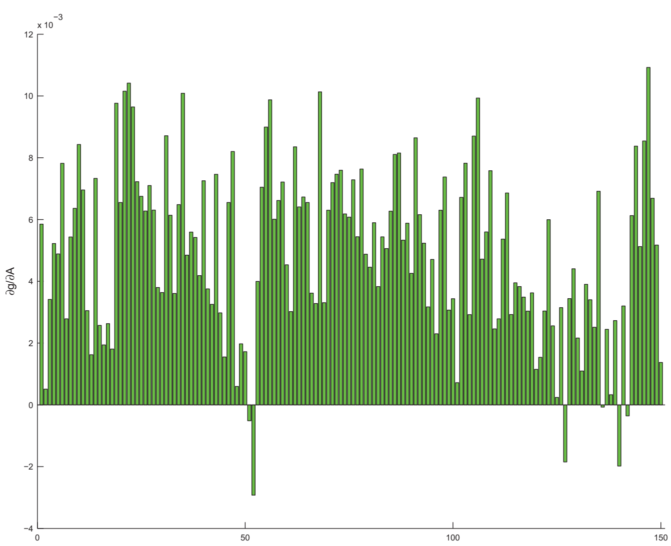

Figure 8. The sensitivity on the first action A_1 for each of the 150 plans (in all plans, the first action has only one alternative).

To introduce some additional robustness in the analysis of $\partial g/\partial A$ we produce box plots of (7) for the five best plans P_1–P_5, see Figure 9. We notice, for example, the low sensitivity and low variance of the eighth non-singleton action with alternative 16 and the rather high sensitivity of the ninth non-singleton action with alternatives 17–21. The latter is one of several actions that need to be in focus of attention of the decision maker.

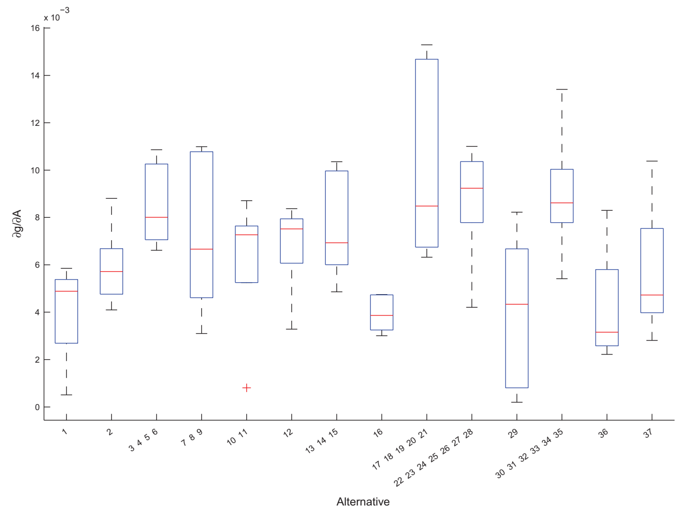

Figure 9. The variance in sensitivity of the five best plans.

Finally, to get an overview, for the best 150 plans and all possible alternatives, the sensitivity is plotted in Figure 10.

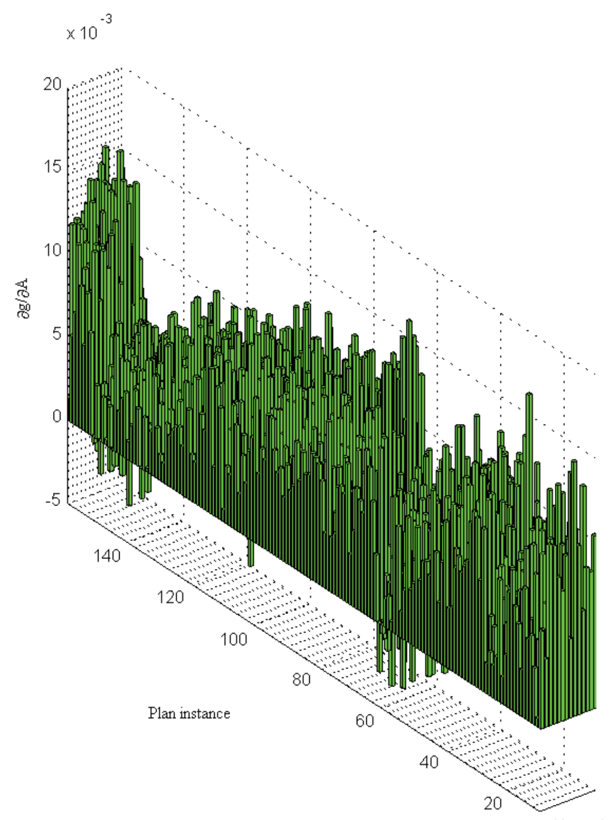

Figure 10. The sensitivity of all data; that is the 150 best plans found by A*-search together with all their additional simulations of neighboring plans. The first bar series facing the viewer is the same as in Figure 8 and the front bar series above the text "Alternative" corresponds to Figure 7.

VII. Conclusions

We have developed an explanation function that highlights the actions of a military plan where the impact of making the best selection of alternatives is most important.

It should be noted that this is not the same as which actions are most important for the overall success of the plan. There may for example be important actions where the impact on the success of the plan does not vary much between different alternatives, making a selection of alternative less important. There may even be singleton actions with only one way to perform it.

Thus, what we have developed is a method that highlights which actions of the military plan is crucial from a decision making point of view. This is where decision makers should focus their pre-execution attentions.

References

[1] J. Schubert, F. Moradi, H. Asadi, P. Hörling, and E. Sjöberg, “Simulation-based decision support for effects-based planning,” in Proceedings of the 2010 IEEE International Conference on Systems, Man and Cybernetics, October 2010. Piscataway, NJ: IEEE, 2010, pp. 636−645.

[2] F. Moradi and J. Schubert, “Modelling a simulation-based decision support system for effects-based planning,” in Proceedings of the NATO Symposium on Use of M&S in: Support to Operations, Irregular Warfare, Defence Against Terrorism and Coalition Tactical Force Integration (MSG-069), paper 11, pp. 1−14, October 2009.

[3] J. Schubert and A. Linderhed, “Learning boundaries on military operational plans from simulation data,” in Proceedings of the 2011 IEEE International Conference on Systems, Man and Cybernetics (SMC 2011), October 2011. IEEE, Piscataway, NJ, 2011, pp. 1325−1332.

[4] E. W. Dijkstra, “A note on two problems in connexion with graphs,” Numerische Mathematik, vol. 1, pp. 269–271, December 1959.

[5] P. E. Hart, N. J. Nilsson, and B. Raphael, “A formal basis for the heuristic determination of minimum cost paths,” IEEE Transactions on Systems Science and Cybernetics, vol. 4, pp. 100–107, July 1968.

[6] E. A. Smith, Complexity, Networking, and Effects-based Approaches to Operations. Washington, DC: Department of Defense CCRP, 2006.

[7] J. P. Hunerwadel, “The effects-based approach to operations: Questions and answers,” Air & Space Power Journal, vol. 20, pp. 53−62, Spring 2006.

[8] T. M. Strat and J. D. Lowrance, “Explaining evidential analyses,” International Journal of Approximate Reasoning, vol. 3, pp. 299–353, July 1989.

[9] H. Xu and P. Smets, “Some Strategies for Explanations in Evidential Reasoning,” IEEE Transactions on Systems, Man, and Cybernetics–Part A: Systems and Humans, vol. 26, pp. 599–607, September 1996.

[10] G. Shafer, A Mathematical Theory of Evidence. Princeton, NJ: Princeton University Press, 1976.

[11] R. W. Hamming, “Error detecting and error correcting codes,” The Bell Systems Technical Journal, vol. 29, pp. 147–160, April 1950.