Computational Creativity for Intelligence Analysis

Robert Forsgren, Peter Hammar and Magnus Jändel

Swedish Defence Research Agency
164 90 Stockholm, Sweden
{robert.forsgren, peter.hammar, magnus.jaendel}@foi.se

Abstract. We describe a decision support system for hypothesis assessment in which exploration is supported by computational creativity. A software tool for morphological hypothesis analysis and evidence handling is extended with a creative assistant that on demand suggests hypotheses that the analyst should consider. Suggested hypotheses are chosen so that they are far from hypotheses that the analyst previously has paid attention to but nevertheless are supported by evidence in an interesting way. For the purpose of providing thought-provoking suggestions, the creative assistant employs ensembles of novelty and value assessment methods and proposes hypotheses that stand out in this multi-ensemble analysis. Preliminary experiments investigate the system's potential for infusing novel and valid ideas into the decision making process.

Keywords: Decision support system, hypothesis assessment, situation assessment, computational creativity

1 Introduction

Situation assessment is the core of intelligence analysis. Once the analyst understands what really is going on it is comparatively easy to come up with reasonable actions and evaluate costs, risks and likely outcomes. Classical intelligence analysis means that the space of all relevant hypotheses is defined and that the analyst relates each piece of evidence to each hypothesis as for example in the *Analysis of Competing Hypotheses* (ACH) method [1]. Realizing that such classical methods are impractical in face of the profusion of hypotheses and evidence provided by present day information handling systems, Gustavi et al. [2] pioneered that evidence should be connected to hypothesis attributes rather than to each hypothesis per se thus greatly reducing the number of evidence links and hence simplifying the analysis process.

The *Multi-Hypothesis Management and Analysis* (MHMA) method of Gustavi et al. [2] extends legacy *Morphological Analysis* tools in use by intelligence analysts. Morphological analysis is a three-phase qualitative analysis methodology in which a first phase maps out the hypothesis space, a second phase identifies consistency constraints and a third phase applies judgment to analyse and select preferred hypotheses. The method was pioneered by the astrophysicist and polymath Fritz Zwicky [3] and has been applied in diverse fields including future studies, policy analysis, law and

adfa, p. 1, 2011.

technology. There is a vast literature on morphological analysis[1] for decision support that we will not attempt to review here (for a recent monograph see [4]).

Situation assessment is, however, not just about mechanically evaluating how the weight of pre-existing evidence is distributed over some hypothesis space. Based on what hypotheses that are found to be interesting, analysts will look for more evidence and re-evaluate existing evidence thereby following what essentially is a scientific research methodology. The creativity of this process is as important as the formal representations and information processing tools. Human creativity is, however, a fickle resource at best and succumbs easily to group-think and prejudice. In this paper we investigate how computational creativity can contribute to the analysis process.

Computational creativity is a new and burgeoning branch of computer science. For a recent review see [5]. In [6] we outline a research program for computational creativity in decision making where six different options for incorporating computational creativity in decision processes are described. One of these options pertains to using computational creativity for situation assessment i.e. coming up with the main hypothesis about what kind of situation that is as hand. The present paper describes a system and initial experiments that address this issue and provides hence a first step towards realizing the outlined research program. [6] provides also a brief review of the literature on computational creativity in decision support. As a harbinger of nascent interest in the decision analysis community, note the recent article [7].

In this paper we take a first step towards computational creativity for supporting hypothesis assessment in intelligence analysis by exploring the design space and design principles as well as experimentally comparing the simplest possible implementation with a few selected more complex designs. Section 2 introduces the intelligence analysis methods from the user's point of view and describes how computational creativity is integrated in the user interface. Section 3 defines and motivates the computational creativity algorithms whereas section 4 describes the implemented system that is explored experimentally in section 5.

2 The User Perspective

We introduce evidence supported morphological analysis according to [2] by an example using the same scenario as used in our experiments in section 4. The scenario is about the 2001 anthrax attacks [8] in which, shortly after the 9/11 events, letters containing lethal anthrax spores were mailed killing several infected victims. Both al-Qaeda and Iraq were suspected but B.E. Ivins, a U.S. biodefense scientist was declared to be the sole perpetrator although this conclusion is disputed [8]. Suppose now that we want to make a situation assessment in this scenario.

The hypothesis space is set up by morphological analysis. We define the *parameters* or conceptual dimensions of the problem and a domain of discrete *values* for each parameter. To facilitate presentation in readable figures and tables we use a small

[1] Note that "morphological analysis" has different meanings in decision support and linguistics. In this paper we use the term only in the former meaning.

morphological chart with only three parameters and a handful of values per parameter. *Culprit*, *Motive* and *Source* (of the anthrax spores) are chosen as parameters. The set of values for each parameter is shown in Fig. 1. This morphological chart was developed in [2] where a detailed discussion of how it connects to the scenario is provided. A hypothesis is formed by picking a specific value (cell) from each of the parameters (columns), the hypothesis space being the set of all such combinations.

Culprit	**Motive**	**Source**
Al-Qaeda	9/11 inspired	Gov. special lab
Iraq	Domestic politics	Simpler lab U.S.
B.E. Ivins	Personal gain	Simpler lab Afg.
S.J. Hatfill	Target recipients	Iraqi special lab
People from U.S. gov.		

Fig. 1. Morphological chart with evidence weight represented by shade of colour. The cells with darker colour have twice as much supporting evidence as the lightly coloured cells. White cells are not related to any evidence.

Using the MHMA tool, analysts enter evidence and connect evidence as having a positive, neutral or negative impact on each value cell. Hypotheses are scored based on the aggregated evidence weight on all the values that compose the hypothesis. The essential simplification of MHMA compared to ACH is that evidence is linked to cells rather than hypotheses thus avoiding the combinatorial explosion of evidence links in a large hypothesis space. To illustrate the method of coupling evidence, we have in Fig. 1 entered two pieces of evidence supporting each of the value cells *al-Qaeda*, *9/11 inspired* and *Simpler Afghan lab*. One piece of evidence supports each of the cells *B.E. Ivins*, *Target recipients* and *Gov. special lab*.

Assume now that a forensics team is strongly biased in favour of the al-Qaeda hypothesis maybe to the degree that it is considered disloyal to explore any alternatives to the current orthodoxy. A good computational creative assistant could help overcoming the ingrained bias of the group by proposing that analysts should look into the Ivins hypothesis. Given that the creative assistant knows that only the al-Qaeda hypothesis has been considered by the analysts, it would look for other hypotheses that both has evidential support and are different from most from the well-known hypotheses. The Ivins hypothesis fulfils both these criteria although some variations of the dominating hypothesis actually are better supported by the evidence.

The user interface of our tool consists of four fields as shown in Fig. 2. The *morphological chart* shows the hypothesis space. Analysts can enter evidence in *the evidence list* and connect each piece of evidence by positive or negative links to selected value cells of the morphological chart. *The hypotheses and evidence list* collates the hypotheses according to evidential support and allows the analyst to sort hypotheses according to evidence weight. By selecting a specific hypothesis the analyst can explore all evidence related to the hypothesis. The fourth field is the only new user interface component in our *Computational Creativity Supported Multi-Hypothesis Management and Analysis* tool (CC-MHMA). If the user clicks on the light bulb icon, the system suggests a creative hypothesis.

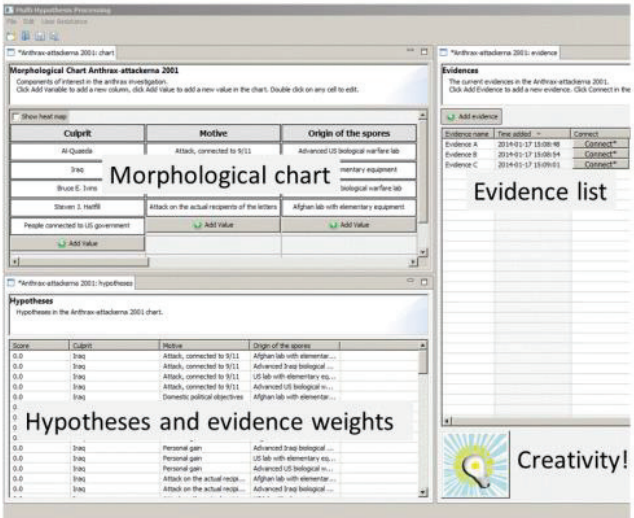

Fig. 2. Annotated user interface of the CC-MHMA tool consisting of the MHMA interface enhanced with the light bulb icon. Legacy MHMA features are the morphological chart, the hypothesis and evidence weight list and the evidence list.

The purpose of the computational creativity function is to provide useful creative input to the analysts. Creative means in this context that the suggestions should be novel from the point of view of the analysts and germane to human goals and the situation at hand. The creativity of a suggestion can hence only be evaluated by human experts. One of the main problems is to provide the computational creativity agent with sufficient knowledge for producing novel and valuable suggestions. Novelty requires insight into what the analyst has considered before and might be aware of. Value calls for knowledge about the situation, the objectives of the analyst and general background information about the application domain. It is often much more difficult to compile the required knowledge in a computer readable form than to come up with suitable processing algorithms. The main hurdle for practical applications of computational creativity is that users cannot be bothered with extra work or cumbersome sensors in order to provide a computational creativity agent with the necessary knowledge. Therefore, we have in this work designed the system to harvest all inputs to the computational creativity from the pre-existing user interface and to include the light bulb as the single minimalistic and unintrusive extension of the user interface. The underlying computational creativity algorithm strives to suggest a hypothesis that has not yet been considered by the analyst and that is supported by the evidence in an interesting way. The next section will explain exactly what this means and how it is done. From the user perspective, the suggestion is shown by a colour-coded selection of values in the morphological chart, highlighting of the proposed hypothesis in the hypothesis list and highlighting of the evidence that impacts on the proposed hypothesis. The data provided by the tool to the computational creativity agent is the morphological chart, the evidence, the links between evidence and parameter values and the history of hypotheses selected for examination by the user. Note that any hypothesis that is suggested by the computational creativity agent is put into this list of known hypotheses to the effect that the computational creativity agent will not suggest something similar again.

3 The Algorithms

The first subsection describes the formal representation of the computational creativity agent while the following subsections flesh out the details. We discuss a broad range of possible algorithms leaving the definition of what has been explored in our preliminary tests to the next section.

3.1 Hypothesis Space and Formal Creativity Model

The hypothesis space H is the set of all hypotheses combined with a distance measure. The distance measure $d(h_1, h_2)$ is a real-valued function of the two hypotheses h_1 and h_2. The distance matrix D is formed by indexing the hypotheses according to an arbitrary order and defining matrix elements according to $D_{ij} = d(h_i, h_j)$.

A variable that is associated with a specific hypothesis and furthermore is defined for all hypotheses in the morphological chart forms a scalar field in hypothesis space. Any such scalar field will in the following be called a *charge field* and be denoted by a bold face symbol.

The next hypothesis to be recommended by the creative assistant is computed from by some multi-objective optimization algorithm Ω,

$$h_c = \Omega(\mathbf{n}, \mathbf{v}), \tag{1}$$

where **n** and **v** are the novelty and value[2] charge fields respectively. Multi-objective optimization algorithms (see [9] for a review) typically includes (1) finding the set of (Pareto) efficient hypotheses such that no other hypothesis have better novelty charge without having a worse value charge or vice versa and (2) selecting a solution among this set.

The *novelty charge* **n** of a hypothesis is a predictor for how novel the hypothesis appears to the user. It depends on how aware the analyst is of the hypothesis, the distance of the hypothesis to, and the level of awareness for, other hypotheses that the analyst is aware of. Formally we express this as,

$$\mathbf{n} = \Psi(\mathbf{a}, D), \tag{2}$$

in which **a** is the awareness charge field and D is the distance matrix. The function Ψ should be designed to allocate high novelty to hypotheses with low awareness charge

[2] We suffer from a terminology collision regarding the term *value* between the fields of morphological analysis and computational creativity. Where the distinction is not obvious from the context *a cell in the morphological chart* is referred to as *value cell* whereas *usefulness* is called *value charge*.

that are far from any hypotheses with high awareness charge. The awareness charge represents how conscious the user is of the hypothesis.

The value charge **v** of a hypothesis is a predictor for how well the hypothesis describes the real-world situation that the analyst is interested in. Formally, the value charge field depends on the evidence E according to some function,

$$\mathbf{v} = \Xi(E) . \tag{3}$$

Both the novelty charge and the value charge are real-valued variables in the domain [0, 1]. Ascending numerical value means rising novelty and value respectively.

3.2 Hypothesis Distance Measures

If the creative assistant has no insight in the semantics of the parameters and values of the morphological chart, the most obvious choice of distance measure $d(h_1, h_2)$ seems to be the Hamming distance [10] i.e. count the number of value cell substitutions that is required for transforming one of the hypotheses to the other and use the result as our distance measure.

Many possible enhancements of the distance measure depend on that the user can be coaxed to input more information. An example would be that some parameters in the table could be given more importance than others, creating a weighted Hamming distance where value substitutions would contribute to the distance in proportion to the weight of the parameter. In certain domains values could be ordinal, creating additional structure for a distance measure. Furthermore, we could ask the user to express the relations between the values in the domain of a given parameter as a graph. The contribution to the distance from each parameter could be counted as the number of edges that connect the value cells in each hypothesis. Also such a graph could be enhanced by weighing the edges. Other options for improving the distance measure with additional user input can also be envisaged.

3.3 Modelling Value Charge

We model the value charge of a hypothesis with a sigmoid form,

$$v_h = (1 + \exp(-\sum_{i=1}^{N_e} w_{ih}))^{-1} , \tag{4}$$

where the evidence weight w_{ih} describes how the piece of evidence with index i supports or contests the hypothesis h whereas N_e is the number of evidences. We take w_{ih} to be real-valued with positive values denoting degrees of support.

The sigmoid in (4) is symmetric, and saturates as the amount of (positive or negative) evidence increases, i.e. for a large collection of evidence the value v_h will change very little by the addition of one more piece of evidence. This captures the assumption that the belief (or lack thereof) in a hypothesis that is already well supported (or refuted) by a massive amount of evidence will not change much if another piece of evidence is added.

The MHMA software invites the analyst to define a negative, neutral or positive relation between each value cell in the morphological chart and each piece of evidence. By this device, MHMA implicitly connects each piece of evidence to many hypotheses in one fell swoop. To simplify, we will in the following only consider the resultant connections between evidence and hypotheses. The user input describing the relation between hypothesis *h* and the evidence with index *i*, is henceforth called the evidence impact factor and is denoted e_{ih}. In general, e_{ih} is real-valued with positive values representing levels of support.

The evidence weighting model relates the evidence weights to the evidence impact factors,

$$w_{ih} = \begin{cases} \gamma e_{ih} & if\ e_{ih} \geq 0 \\ \gamma \alpha_{-} e_{ih} & if\ e_{ih} < 0 \end{cases}, \quad (5)$$

in which γ is used for defining how much evidence that is needed for near certainty. The sigmoid function in (4) saturates when the absolute value of the sum of weights is about three which means that γ should be set so that this happens when the analyst would judge that the total weight of evidence is such that any further corroboration makes little difference for the conclusion. The parameter α_{-} is used for controlling the importance of negative evidence. The MHMA tool has primed the analyst to assume that $\alpha_{-} = 1$. The creative assistant could for example explore $\alpha_{-} = \infty$ according to which hypotheses that are disfavoured by any negative evidence have zero value or $\alpha_{-} = 0$ where the impact of negative evidence is disregarded. The creative assistant can furthermore use different perspectives on the relation between the evidence and the hypotheses by using several different evidence weighting models each represented by a choice of values of γ and α_{-}.

3.4 Capturing Awareness

To provide novelty, the computational creativity agent needs to know what hypotheses that the analyst already have considered. User interface actions, the timing of user interface actions and biometry are the three main means for learning about this. We represent the analyst's awareness of hypothesis *h* by a single real-valued variable a_h with values in the domain [0, 1] where $a_h = 0$ means no knowledge of *h* and $a_h = 1$ means maximum awareness of *h*.

Relevant user interface actions include selecting a hypothesis, selecting a group of hypotheses and examining evidence related to selected hypotheses. For each such action *A* and each hypothesis *h* that is associated with the action, the system should update the awareness charge according to some function $a_h = F(A, h, a_h)$ which operates on the present value of a_h and outputs an updated value of a_h. We can for example initialize all awareness charges to zero, allocate a baseline charge $a_h = 0.5$ the first time h is selected by the user and further increase a_h if the user scrutinizes related evidence or selects the hypothesis again.

The timing of user interface actions related to a hypothesis could also be taken into account which means that a_h is updated at regular time intervals $\{t_1, t_2, \ldots\}$ by an updating function that incorporates the history of user interactions according to,

$$a_h(t_n) = F(\mathbf{A}, h, a_h(t_{n-1}), t_n)\,, \tag{6}$$

where t_n is the current time, t_{n-1} is the time of the previous update and $\mathbf{A} = \{A(t_1), A(t_2), \ldots\}$ is the history of user actions. In (6) we could for example model the forgetfulness of analysts by letting a_h decay over time according to $a_h(t_n) = c(t_n - t_h)^{-\mu}$ where c, $t_h < t_n$ and μ are parameters. Psychological research indicates power laws for forgetting [11].

Timing analysis of user interactions is complicated by the lack of knowledge about what the user is doing while a given hypothesis is selected. The analyst may be vigorously ruminating over the hypothesis or alternatively be on a coffee break. Biometric methods, including for example video analysis or eye tracking, could provide crucial information about user behaviour to be encoded as special types of actions in (6).

3.5 Modelling Novelty Charge

Although there are many possible algorithms for computing the novelty charge field according to (2), we shall presently only define a simple baseline method and a somewhat more generic parameterized model.

According to the baseline method, the novelty charge of a hypothesis is zero if the awareness charge of the hypothesis is positive i.e. the user is considered to be aware of it. Otherwise, the novelty charge is proportional to the distance to the closest other hypothesis with positive awareness charge. Finally, we normalize so that novelty charges fall in the domain [0, 1].

The parameterized model is based on computing a *familiarity potential*,

$$\varphi(h) = \sum_{h' \in H} \frac{a_{h'}}{d^{\kappa}(h, h')}\,, \tag{7}$$

where the sum is taken over all hypotheses and κ is a positive integer. To handle the self-potential issue, we define $d(h, h) = \delta$ where δ is another parameter. In this equation, awareness charge is similar to electric charge and φ is analogous to electric potential. Consequently, the familiarity potential decreases as the distance to known hypotheses increases. The novelty charge is a sign-reversed normalized version of φ according to,

$$n_h = \frac{\varphi_{\max} - \varphi(h)}{\varphi_{\max} - \varphi_{\min}} \tag{8}$$

in which $\varphi_{\max}$ and $\varphi_{\min}$ are the maximum and minimum values of the familiarity potential respectively.

3.6 Putting It All Together: Creative Suggestions

The main assumption driving the design of our first prototype system is that users may be able to second-guess deterministic algorithms and will then feel that the creative clout wanes as they gain increasing experience of the tool. To avoid this, the creative assistant should choose randomly from an ensemble of different algorithms. In future experiments we intend to test if this assumption is valid. A generic creative suggestion method consists of, (A) a distance measure $d_i(h_1, h_2)$, (B) an algorithm for computing the novelty charge field Ψ_j (see (2)), (C) an algorithm for computing the value charge field Ξ_k (see (3)), (D) a multi-objective optimization algorithm $\Omega_m(\mathbf{n}, \mathbf{v})$ (see (1)). Each of these are selected from a corresponding ensemble of distance measures, novelty or value charge fields or optimization algorithms, respectively. The method for capturing the user awareness is considered to be fixed. At each new round of suggestion production, the creative assistant will select a creative suggestion method comprised of a randomly collection of components from these four aspects.

4 The Implementation

This section describes the experimental setups used in our initial explorative experiments. The computational creativity agent is implemented as an extension to the MHMA tool described in [2].

4.1 Baseline Implementation

The baseline implementation of the creative assistant is intended to investigate computational creativity in the simplest possible setting thereby providing a reference point for more complex implementations. The user awareness charge is initiated to zero for all hypotheses and is changed to one the first time that the user selects a hypothesis for examination. There is no time decay, and no other relations between awareness charge and user actions. Hence there are two distinct set of hypotheses: known hypotheses with $a_h = 1$ and unknown hypothesis with $a_h = 0$. This simplistic model of the user state was selected because it only uses information that is available in the MHMA tool.

The Hamming distance is the only distance measure employed by the baseline implementation. There is just one novelty charge algorithm according to which the novelty charge (n_h) of a known hypothesis is zero and the novelty charge of an unknown hypothesis is proportional to the distance to the closest known hypothesis. Furthermore, we employ one single value charge algorithm according to (4) and (5) with $\gamma = 1$ and $\alpha_{-} = 1$.

The multi-objective optimization algorithm applies the utility function,

$$f(v_h, n_h) = \beta \frac{n_h - \hat{n}}{\sigma_n} + (1-\beta) \frac{v_h - \hat{v}}{\sigma_v}, \qquad (9)$$

for selecting the hypothesis to suggest. In (9), $\hat{n}$ and σ_n are the average and standard deviation of the novelty charge with the corresponding notation for the value charge. Only unknown hypotheses are considered as candidates for selection and for calculating the averages and standard deviations. The parameter β balances the influences of novelty and value and is in our initial experiments ad hoc selected to 0.75. Since the optimum of this particular utility function always is Pareto efficient there is no need to explicitly compute the efficient set.

5 The Experiments

Experiments are performed on the morphological chart presented in the introduction. Three models are tested, representing a baseline implementation, an enhanced value model and an enhanced novelty model.

5.1 Testing the Baseline Implementation

The experiments use a collection of evidence comprising 35 pieces of information as further described in [2] and summarized in Fig. 3. Note that this set of evidence has been selected only for the purpose of testing the decision support tool and is not claimed to accurately represent the factual circumstances. A brief look at Fig. 3 suggest that there are much evidential support of Ivins being involved as well as domestic politics and the 9/11 events as motives. There is more evidence rejecting than supporting an Afghan lab as the source of the anthrax spores, whereas an Iraqi lab have an equal amount of positive and negative evidence.

Culprit	Motive	Source
Al-Qaeda	9/11 inspired	Gov. special lab
Iraq	Domestic politics	Simpler lab U.S.
B.E. Ivins	Personal gain	Simpler lab Afg.
S.J. Hatfill	Target recipients	Iraqi special lab
People from U.S. gov.		

Fig. 3. Morphological chart showing the total weight of the evidence considered in our experiments. Positive evidence weight is indicated by solid green shading with darker tone representing more evidential support. Negative evidence weight is indicated by striped cells with the level of red tone indicating the absolute value of the evidence weight.

We will first consider a situation in which the user just keeps pressing the light bulb icon and does not select any hypotheses other than those suggested by the creative assistant. This is not the normal mode of usage but serves to illustrate how the creative assistant works. The first ten hypotheses suggested by the creative assistant are shown in Table 1. We use abbreviated versions of the value cell names in the tables.

The first suggestion in Table 1 is the hypothesis with highest total value charge. For the next three suggestions, the assistant selects the hypothesis with the highest value charge that does not include any value cells from hypotheses already selected (i.e. they are maximally novel). The fifth hypothesis to be suggested is the one with highest evidence support but not sharing more than one value cell with any of the already known hypotheses.

Table 1. Hypotheses in the order suggested by the baseline creative assistant assuming that the user is aware only of previously suggested hypotheses. The Evidence column shows the number of evidences supporting and refuting the hypothesis as well as the summed evidence weight.

Culprit	Motive	Source	Evidence	
Ivins	Domestic	Simple U.S.	+21, -7 =	+14
U.S. gov.	9/11	Gov. special	+11, -2 =	+9
Al-Qaeda	Recipients	Iraqi special	+9, -3 =	+6
Hatfill	Personal	Simple Afg.	+5, -1 =	+4
Ivins	Recipients	Gov. special	+18, -7 =	+11
Hatfill	9/11	Simple U.S.	+12, -2 =	+10
Hatfill	Domestic	Gov. special	+12, -2 =	+10
Ivins	9/11	Iraqi special	+18, -8 =	+10
U.S. gov.	Recipients	Simple U.S.	+11, -2 =	+9
U.S. gov.	Domestic	Iraqi special	+11, -3 =	+8

Note that the wide scope of the hypothesis space combined with the simple MHMA evidence handling model occasionally give high value to intuitively rather unlikely hypotheses such as al-Qaeda targeting the individual mail recipients or Ivins using spores from an Iraqi lab. The user could, however, use even unrealistic suggestions as creative stepping stones rather than literally as candidates for the most likely solution. Although Ivins may not have had access to spores from Iraqi labs, investigators could be inspired to consider the possibility of a U.S. perpetrator producing a strain that appears to come from an Iraqi lab.

As we keep generating suggestions, the creative assistant eventually runs out of creativity at a point when all not known hypotheses have the same distance to the closest known hypothesis and thus the same novelty. The value charge is then the only decisive factor in (9) which makes any further suggestions trivial since the analysts have other tools for sorting hypotheses according to evidential support. Note that this effect is a consequence of using the Hamming distance as novelty measure and that analysts in practice will reach this state only in scenarios with quite small hypothesis spaces.

The fourth suggestion in Table 1 has a rather low value; only three out of 80 hypotheses have a lower value. Its high position in the list is caused by the combination of using the coarse-grained Hamming distance as novelty measure combined with the strong bias towards generating novel hypotheses engendered by the choice of $\beta = 0.75$

in (9). We tried setting $\beta = 0.25$ with the result that the fourth suggestion in Table 1 was removed from the top ten suggestions while all other hypotheses remained in the same order with a new hypothesis {*al-Qaeda*, *Personal gain*, *Simpler lab U.S.*} with evidence sum +8 appearing at the bottom of the list.

5.2 Enhanced Value Model

The sigmoid function in (4) saturates for most of the hypotheses in the baseline implementation. This means that differences in the creative utility function (9) is dominated by differences in novelty. Reducing the parameter γ should extend the range of value charge explored by (4). By running a series of experiments with different γ and β we found that the combination $\gamma = 0.3$ and $\beta = 0.3$ appears to give a reasonable balance of novelty and value. Table 2 provides an example of the output indicating a higher preference for value compared to Table 1.

Table 2. Hypotheses in the order suggested by creative assistant using the baseline implementation enhanced with the value model $\gamma = 0.3$, $\alpha_{-} = 1$ and using utility function parameter $\beta = 0.3$.

Culprit	Motive	Source	Evidence
Ivins	Domestic	Simple U.S.	+21, -7 = +14
U.S. gov.	9/11	Gov. special	+11, -2 = +9
Ivins	Recipients	Gov. special	+18, -7 = +11
Hatfill	9/11	Simple U.S.	+12, -2 = +10
Hatfill	Domestic	Gov. special	+12, -2 = +10
Ivins	9/11	Iraqi special	+18, -8 = +10
Al-Qaeda	Recipients	Simple U.S.	+11, -2 = +9
Ivins	Domestic	Gov. special	+20, -7 = +13
Ivins	9/11	Simple U.S.	+20, -7 = +13
U.S. gov.	Domestic	Iraqi special	+11, -3 = +8

As suggested in the discussion of modelling value charge, different value models can be obtained by varying, α_{-}. We have briefly investigated the case of $\alpha_{-} = 0$ which requires additional adjustment of γ in order to obtain results substantially differing from the baseline implementation. Setting $\gamma = 0.1$ and $\beta = 0.5$ produces a new and different set of creative suggestions.

5.3 Enhanced novelty model

The enhanced novelty model according to (7) and (8) is controlled by parameters δ and κ. The self-potential (δ) has no effect in the present implementation since only hypotheses with zero awareness charge are candidates for creative suggestions. Using the baseline implementation with an enhanced novelty model with either $\kappa = 1$ or $\kappa = 2$ is, however found to make a significant difference in the output of the creative assistant. Table 3 shows the resulting hypotheses. Comparing this enhanced novelty model with the baseline implementation shows that the first four suggestions are the same but that the following suggestions are quite different reflecting that the potential model will differ more from the Hamming distance as the inventory of known hypotheses accumulates.

Table 3. Hypotheses in the order suggested by creative assistant using the baseline implementation enhanced with the novelty model of (7) and (8) in which both $\kappa = 1$ and $\kappa = 2$ give the same output.

Culprit	Motive	Source	Evidence
Ivins	Domestic	Simple U.S.	+21, -7 = +14
U.S. gov.	9/11	Gov. special	+11, -2 = +9
Al-Qaeda	Recipients	Iraqi special	+9, -3 = +6
Hatfill	Personal	Simple Afg.	+5, -1 = +4
Iraq	Domestic	Gov. special	+10, -2 = +8
Hatfill	9/11	Simple U.S.	+12, -2 = +10
Ivins	Recipients	Simple Afg.	+14, -6 = +8
U.S. gov.	Personal	Iraqi special	+8, -3 = +5
Al-Qaeda	Personal	Simple U.S.	+10, -2 = +8
Iraq	9/11	Iraqi special	+8, -3 = +5

6 Discussion and Conclusions

Our initial experiments have focused on exploring key aspects of the generative algorithm and in particular the effect of varying selected components. The general impression from the experiments is that the creative assistant shows some promise and should be properly evaluated. This means that we must use much larger scenarios than in the present experiments. We cannot expect users to get a genuine eureka experience from creative suggestions unless the hypothesis space is so large that humans find it impossible to systematically consider all alternatives. Furthermore, we need a rich and complex evidence situation. In order to judge the creativity of the hypotheses suggested, and the value of the creative assistant, an experiment would use subject matter experts working on a realistic case for which they do not know the solution beforehand. The evaluation should be done in the process as well as post mortem.

We primarily regard the computational creativity as a tool for inspiring human analyst to explore a wider range of ideas and consider more alternatives rather than as generator of optional solutions per se. Furthermore, the infusion of creativity can stimulate users to expand the analysis model with new data, such as adding new value cells or looking for further evidence.

We recognize that morphological analysis may not be the ideal platform for building computational creativity. One could argue that a morphological chart spans a limited and static hypothesis space so that the only scope for creativity is to explore a predefined domain. However, we hold forth that the contribution of the computational creativity can only be evaluated by real users and that access to a user community of a perhaps less than ideal tool is a better starting point for researching computational creativity than to build a perhaps theoretically better tool with no practical opportunity for real-life evaluation. Likewise, we understand that a much better model of the user state can be built with state of the art behavioural research methods but we also know that more intrusive probes would deter professional analysts. By unobtrusively extending an existing, actively used baseline system for evidence supported morphological analysis we will in future more comprehensive experiments benefit from the crucial resource of the existing user community.

Although some of the suggestions produced by the creative assistant may seem to be obvious to a detached viewer, such ideas may still be useful in real-life decision making. Human free-thinkers may find it hard to get attention in a team that is locked into group-think. Originators of dissident ideas could be accused of disloyalty or having ulterior motives. That a divergent suggestion originates from a supposedly impartial and objective machine may help to make analysts consider it seriously and perhaps take it as a stepping stone for further exploration. Application of computational creativity to situation assessment in intelligence analysis, as described here, has not been explored before. Despite the limitations of the present implementation and the experimental scenario we feel that the novelty of the approach makes it interesting as a basis for further investigations and in particular full-fledged user trials.

Acknowledgements. Christian Mårtensson suggested that we should consider how to apply computational creativity in morphological analysis. We thank Tove Gustavi, Maja Karasalo and Christian Mårtensson for making the MHMA tool available for our research. This research is commissioned by FMV, the Swedish Defence Materiel Administration.

References

1. Heuer, R.J.: Psychology of Intelligence Analysis. US Gov. Printing Office (1999)
2. Gustavi, T., Karasalo, M., Mårtenson, C.: A Tool for Generating, Structuring, and Analyzing Multiple Hypotheses in Intelligence Work. In Intelligence and Security Informatics Conference (EISIC), 23–30 (2013)
3. Zwicky, F.: Discovery, Invention, Research - Through the Morphological Approach. The Macmillian Company, Toronto (1969)
4. Ritchey, T.: Wicked Problems – Social Messes: Decision Support Modelling with Morphological Analysis. Springer (2011)
5. Colton, S., Wiggins, G.A.: Computational Creativity: The Final Frontier?. Proc. 20th European Conference on Artificial Intelligence (2012)
6. Jändel, M.: Computational Creativity in Naturalistic Decision-Making. In Proceedings of the Fourth International Conference on Computational Creativity, 118–122 (2013)
7. Varshney, L.R.: Surprise in Computational Creativity and Machine Science. Decision Analysis Today 32: 25–28 (2013)
8. Cole, L.A.: The Anthrax letters. Skyhorse publishing (2009)
9. Ehrgott, M.: Multiobjective Optimization. AI Magazine 29 (4): 47-57 (2009)
10. Hamming, R.W.: Error Detecting and Error Correcting Codes. Bell System Technical Journal 29: 147–160 (1950)
11. Wixted, J.T., Ebbesen, E.B.: Genuine Power Curves in Forgetting: A Quantitative Analysis of Individual Subject Forgetting Functions. Memory & Cognition 25 (5): 731–739 (1997)