# Neuro-Symbolic Modelling for Operational Decision Support

**Jeroen Voogd**

**Patrick Hanckmann**

**Paolo de Heer**

**Jeroen van Lith**

TNO
Oude Waaldorperweg 63
2597 AK The Hague
THE NETHERLANDS

{jeroen.voogd, patrick.hanckmann, paolo.deheer, jeroen.vanlith}@tno.nl

## *ABSTRACT*

*During military operations much information (physical, human, information) must be integrated and interpreted to have situational understanding (SU) and make the right decisions while under time pressure in potentially life-threatening circumstances.*

*Simulation can play a role in decision making, for example by doing what-if simulations or determining key factors via data farming. Such a simulation should reflect all important aspects of the operations area; thus the models and their underlying data must continuously be updated. Based on an understanding of the current situation, the different models (geographical, behavioural, etc.) are composed into one simulation model, and configured with the current data. To obtain current and relevant understanding, building the simulation model and executing it must be as quick as possible. Automation of (parts of) these steps is therefore necessary.*

*In this work we discuss an automated support for gaining SU where multiple sources of information are integrated. We combine sub-symbolic information (e.g. images and sensor data) and technologies (e.g. machine learning with neural networks) with symbolic knowledge (e.g. from military experts) and technologies (expert systems). The case study is military terrain analysis, and more precisely identifying buildings and their function. On the one hand, the sub-symbolic part takes images of buildings as input and classifies them, giving a probability to be of a certain type/function. At the same time, expert and symbolic information (e.g. schools are commonly situated in the city centre and well connected by roads) is used to guide, steer and refine the classification.*

## 1.0 INTRODUCTION

Military decisions are made in a complex environment. To obtain situational understanding, information from different sources need to be integrated. Information from sensors, which may include human observations, and military expert knowledge, e.g. obtained in previous operations, are important sources to build an understanding of the situation.

The data coming from sensors is often unstructured, ambiguous, and sometimes purposefully false, while military knowledge is spread over different experts in different locations. In modern operations, the amount of sensor data can be overwhelming, and analysts have a hard time keeping focussed on the most important pieces of data. During analyses of possible interventions and their expected effects, it is important for the decision maker to have insight into uncertainties and possible second or third order (side) effects.

To support military decision makers, TNO uses various analysis tools such as Marvel [1]. Using these tools, a model is created by hand based on military expert knowledge and available data. This model is then used to analyse the effects of interventions. Building these models, however, is labour intensive and takes considerable time. Adaptation or extension during the operation requires diverse military expertise, and good knowledge of the model and the tooling, which may not be at hand. The consequence is that crucial factors may not be included in the automated part of the analysis. Another issue is that during use it is not possible for the decision maker to engage in a question and answer type of interaction, which might reveal 'hidden' assumptions or omissions.

Hence, these tools may be perfectly suitable for higher command layers. On a strategic level where planning horizon is typically measured in months or years, there is time to consult (model) experts. On a lower level, where planning concerns days or even hours, this is not possible.

It is expected that future military operations will be more and more complex (e.g. mega cities, see section 3.0). In these environments all three landscapes (physical, human, and information) need to be integrated for full situational understanding. There are, however, several factors that limit understanding. There may be too little information on some aspects, there may be unreliable, or intentionally false, information, and there may even be too much information, leading to a needle in the haystack problems. In that situation human analysts and decision makers may not be well suited for sifting through the data to find trustworthy and useful information and being able to integrate it in the decision-making process. Especially given the time pressure.

In this work we seek to aid the analysts and decision makers with automated tooling. We focus on two types of tools: Artificial Intelligence (AI) and simulation.

The AI part of this work is used to make sense of the available information in the area of operations. This can be data from all kinds of sensor systems, knowledge of military experts, lessons learned from previous operations and theories on e.g. how societies or economies work. Analysts should be able to interact with this part of the system to enter information, correct errors, and to obtain new insights based on the whole collection of information. It should be possible to elicit all factors that must be considered during the decision-making process based on the current situation and the commander's objective.

Simulation provides the means to calculate many different scenarios using complex models, and produce an unbiased overview of what (higher order) effects a commander might expect when executing an action. This may require many simulation runs. Commanders will want to interact with this part of the system to enter their goals but also question the answers. Explainability of the results is important to show how the results were obtained. In this way commanders can gain trust in the answers or indicate what is wrong or missing, to have this corrected or added.

In this paper we first sketch the envisioned support system, and indicate some requirements and which technologies are to be developed (Chapter 2). In Chapter 3 a scenario is presented that forms the backdrop of our work, which is described in chapter 4. In that chapter we search for a combination of AI techniques for processing sub-symbolic data, as well as AI techniques for symbolic reasoning, that together are able to form the AI part of the envisioned system. Finally, in chapter 5 some conclusions are drawn.

## 2.0 PROPOSED ANALYSIS AND DECISION SUPPORT SYSTEM

In this chapter we present an overview of the system we envision that is capable of aiding analysts in their effort to obtain situational understanding, as well as helping decision makers in their planning process.

## 2.1 System Overview

Figure 2-1 presents an overview of the envisioned system, see also . On the left, sensors feed data into the system. Sensor can e.g. be cameras, microphones, wire taps, and various social media platforms. When analysed via machine learning techniques they contribute to the concept network (the left part in the grey box in the Figure 2-1). This network consists of a hierarchy of concepts (important aspects of the world, typically expressed as nouns) with relations between them. The stream of concepts from various sources may have all kinds of relations between them that are not explicit. Machine learning techniques are employed to find these relations. Military experts can enter their knowledge and lessons learned into the network as a series of relations between concepts. These are called 'rules'. The analysts and decision makers can use the network of concepts and relations to gain insights in what is or isn't important in the operational area. They may also choose to (automatically) build a simulation model (the right-hand part in the grey box in the Figure 2-1) that covers all important aspects and can be used to help them build Courses of Action (CoA). The simulation results can be used to determine the broad spectrum of what might happen, what the crucial factors are and help construct the best CoA. The constructed CoA may even be fed back into the system such that it can search for positive and negative signs of how well its execution progresses.

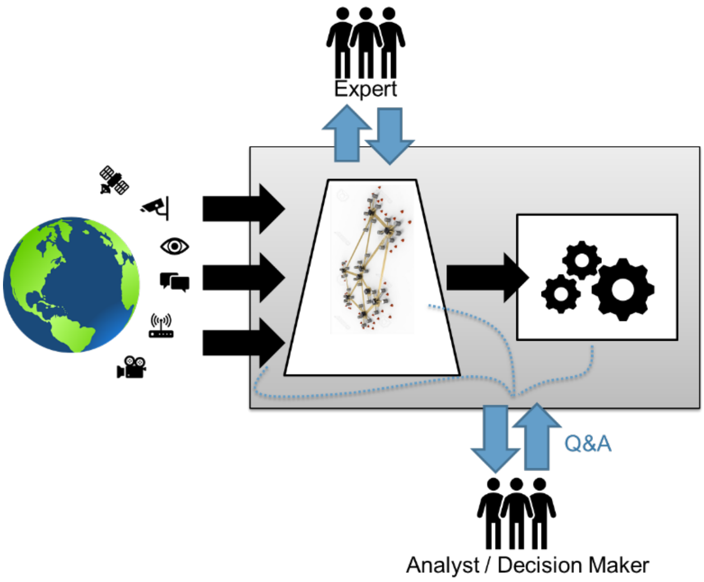


**Figure 2-1: System overview**

The network of concepts and relations will at some point have a structure comparable to the network in Figure 2-2. There the concepts and relations coming from the analysed sensor data, and expert knowledge are integrated. The sensor data is of a sub-symbolic nature, i.e. the individual data items, for example the pixels in an image, do not have an explicit meaning. Using machine learning, these can be brought up to the symbolic level. The knowledge of experts is already expressed on the symbolic level.

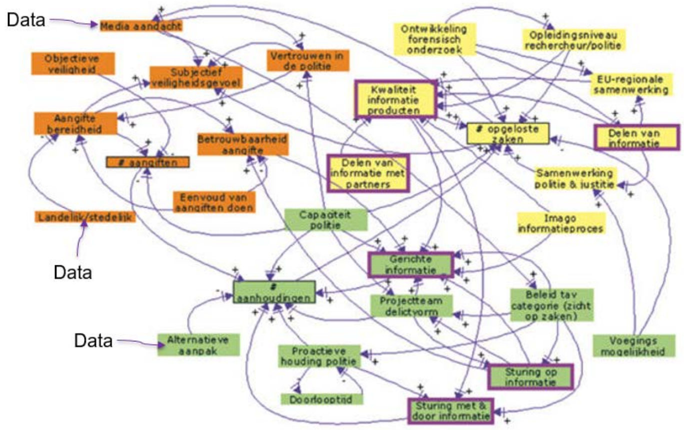


**Figure 2-2: Example of a network of concepts with their relations**

## 2.2 System Properties

The above sketched system will need to have specific properties to be useful for operational decision support. See e.g. also [5] for additional desired properties.

On the military planning levels where results need to be available in hours to days, time is of the essence. There is no time for broad studies or consulting many experts. The cycle from collecting and analyzing data, via simulating courses of action with many random variations, to analyzing the results should be almost a continuous process.

The basis of the knowledge available in the concept network can be formed by lessons learned from similar operations. Transferability of knowledge is an important aspect to make the aggregated knowledge applicable to other operations.

The concept network is updated by incoming and analyzing sensor data as well as military expert knowledge. The expert can input relations they know or expect to hold in the current theatre. The interfaces for entering, but also using, knowledge must be user friendly.

The concept network may change when additional information becomes available. It is important to keep the knowledge up to date for each use in the analysis or planning process. This also means that the best possible CoA that was determined yesterday, may no longer be the best today.

The results shown to the analysts and decision makers must be explainable. These results may be based on incomplete knowledge, and machine learning algorithms may have learned to use unimportant aspects of the data. By analyzing the results, users may be able to correct errors by adding lacking knowledge (concepts and/or relations) or making relations between existing concepts more, or less, important.

Even though the concept network is (currently) thought of as being local to the operational area, the models used in the simulation do not need to be local. The models can be a large collection of small models that live in the cloud and can be composed into one consistent model that contains exactly what is needed to simulate the current situation for the user's purpose (e.g. determine the best CoA).

It is inevitable that uncertainty must be dealt with. Uncertainty comes from all sources of knowledge. Not only can input be uncertain, it can even contain purposefully false information. Uncertainty must be monitored and where possible mitigated. When performing simulations, the uncertainty can be translated into ranges of input data for the simulation parameters as well as how many runs are required.

In the theater of operations, it is unlikely that sufficient computing power is available for performing thousands, or even millions, of simulations. Therefore, a connection to 'the cloud' is required that can offer the needed computing power. This connection needs to be secure and the data to be send and received must fit the available bandwidth, which may determine what data needs to be local or in the cloud.

## 2.3 Relevant technologies developed within NATO context

To build a system as described above, may different technologies will need to work together. Some of those technologies already exist, others are possibly still far off. In this section we briefly describe some relevant technologies that are actively developed within an NATO context. Besides NATO, many other researchers are - of course - also working on related topics. However, we feel that the topic of connecting low-level knowledge to high-level knowledge is still very immature, and that is why we focus on that topic in our research, of which some examples are given in chapter 4.0.

*Cross-panel: Big Data and Artificial Intelligence for Military Decision Making*

In this recently defined cross panel topic [6], initiated by the IST-160 meeting, many technologies were discussed that are relevant to the AI part of the proposed decision support system. Besides presentations and discussions on technology, this cross-panel topic is also meant to structure the work on these topics across panels. For this work, NATO efforts on topics such as (un)supervised learning and semantic integration of heterogeneous data are very interesting.

*Modelling and Simulation as a Service (MSaaS): NMSG-136*

Recent technical development in cloud computing technology and service-oriented architecture (SOA) offers opportunities to utilize M&S capabilities better to satisfy NATO critical needs [7]. A new concept that includes service orientation and the provision of M&S applications via the as-a-service model of cloud computing may enable composable simulation environments that can be deployed rapidly and on-demand. This new concept is known as M&S as a Service (MSaaS). NATO MSG-136 MSaaS – Rapid deployment of interoperable and credible simulation environments investigates MSaaS with the aim of providing the technical and organizational foundations for a future permanent service-based M&S ecosystem within NATO and partner nations, called the Allied Framework for MSaaS.

As described above, for a simulation-based decision support system in a military operational environment, this type of technology is very important in order to obtain the necessary models and computing resources.

*Datafarming: NMSG-155*

The methods and processes of Data Farming have been developed in the six areas of model development, rapid prototyping of scenarios, design of experiments, high performance computing, analysis and visualization of large simulation data output, and collaborative processes. These six domains of Data Farming have been documented as part of the work of the MSG-088 Task Group that codified the data farming concept. In addition, the follow-on MSG-124 Task Group turned the concept into actionable data farming decision support in part by developing a cyber simulation model and a decision support tool. These activities have proven Data Farming ready for application and implementation. Now Data Farming can be made accessible and usable by NATO through MSG-155 efforts to develop the groundwork for analysis and simulation-based decision support.

The results of this effort should largely fill in the simulation part of the envisioned system as described in chapter 2.0, especially when available as a service as worked upon in NMSG-136.

C2-Simulation interoperability: NMSG-145

Simulation based decision support in operations requires that the system can understand the current situation as well as the user's purpose. Therefore, information exchange between command and control devices used by commanders, and simulations need to be in place. NMSG-145 develops a standard for C2-Simulation interoperability: the C2SIM standard. This will offer fast and consistent information exchange among systems.

In this work the C2SIM standard should facility the communication between decision makers (i.e. the commander and his staff) and the simulation-based decision support system.

*Human factors*

Due to the uncertain circumstances and ill-defined problems, AI cannot yet do this autonomously. Instead, deriving decisions from predictions and analysed data should be organized as an interactive human-technology activity, in which both parties become aware of one another's strengths, limitations, and objectives. Within NATO the Human Factors and Medicine Panel studies several topics related to decision support and simulation are studied, as well as the cooperation between humans and decision support systems, see e.g. [8].

## 3.0 EXAMPLE SCENARIO: FUTURE URBAN WARFARE

The world is rapidly changing, and the military domain – in all three landscapes – will change with it. In 20 years, 70% of the population will live in cities [2], many of which will be megacities located in coastal areas [3]. In such environments, military operations become much more complex, for example due to the sheer density and interconnectedness of people and buildings, and factions with different goals and values. Along with that, a large part of warfare will have shifted from the physical landscape to the information landscape, indirectly and subtly influencing the social landscape and creating a grey zone where entities employ hybrid strategies, and it might not be clear how or even if an area is at war, see Figure 3-2.

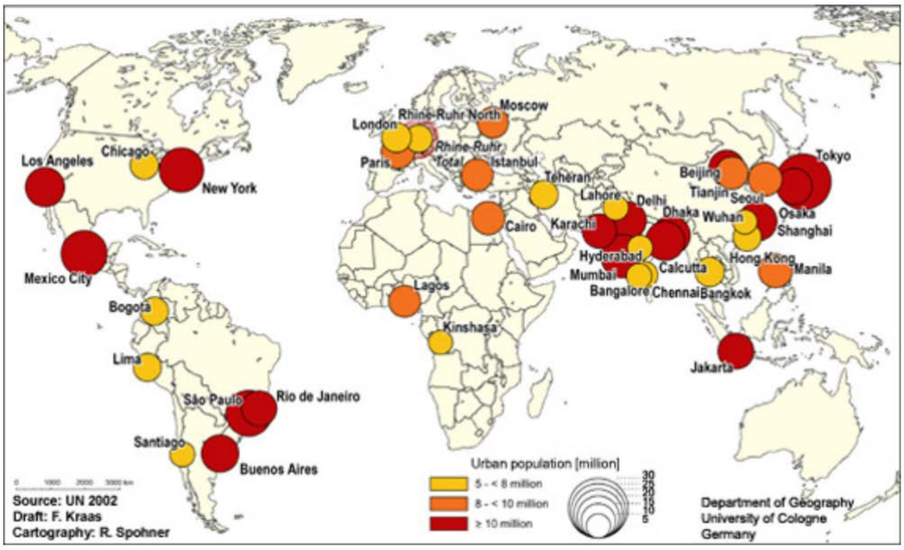


Figure 3-1: Global growth of megacities [4].

As the military domain shifts to these more complex situations and a larger presence in the information landscape, it will become paramount to be able to make sense of the vast amounts of data collected. Already today, analysts are having difficulty to keep up with the data collected and making sense of the operational situation. As time passes, more ways to collect data are introduced. Interpreting the data will no longer be feasible without using AI algorithms.

This urban, littoral domain of the future is the scenario chosen for this research. In this scenario, typical critical information requirements for commanders can be: the seats of power in the area, the location and outreach of key media, local patterns of life and which notable groups or splinter cells might be active.

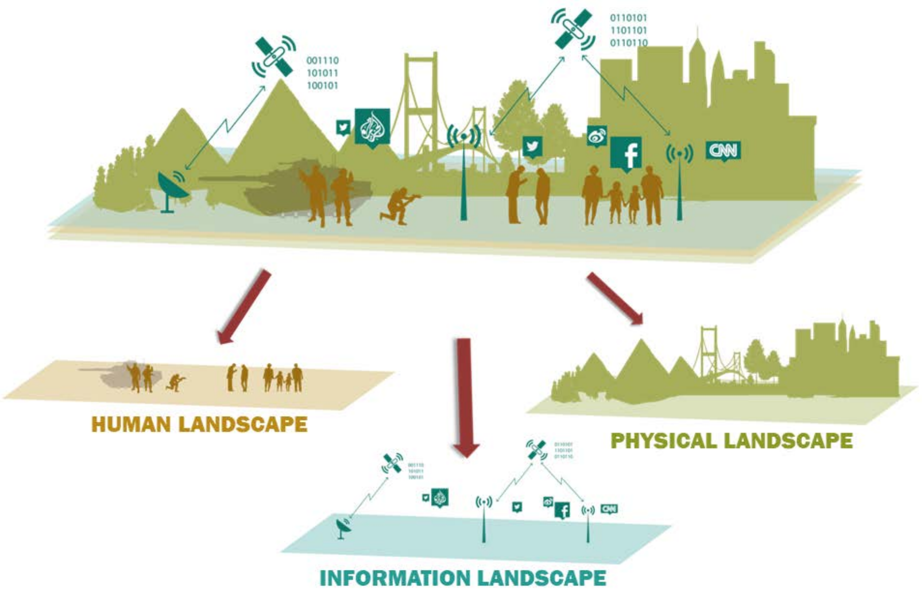


**Figure 3-2: The three interconnected landscapes of operation.**

# 4.0 EXAMPLES OF SUBSYMBOLIC-SYMBOLIC MODELLING

In this chapter some projects the authors have been or are working on are discussed and some (preliminary) results are shown.

## 4.1 Pattern Of Life Identification

In military operations detectors may be used which produce low-level sub-symbolic data. Because this data is hard to interpret for a decision maker, it is explored how this sub-symbolic data can be transformed into higher-level concepts. Higher-level concepts are obtained from the raw data detectors by finding correlations between their various input data items and represent them as logical relationships. One experiment shows that with relatively little, low-level, input data, higher level information can automatically be obtained.

In the experiment, data is obtained from e.g. a UAV that can hoover above a part of a city for an extended period. Note that we used a simulation of a population in a city for this experiment. From an image analysis module (which was not needed in the simulation) data is obtained about people entering and leaving a specific building, here a school building. The enter and leave data items are timestamped. From this data alone, it is tried to obtain part of the pattern of life of people living in the neighbourhood of that school.

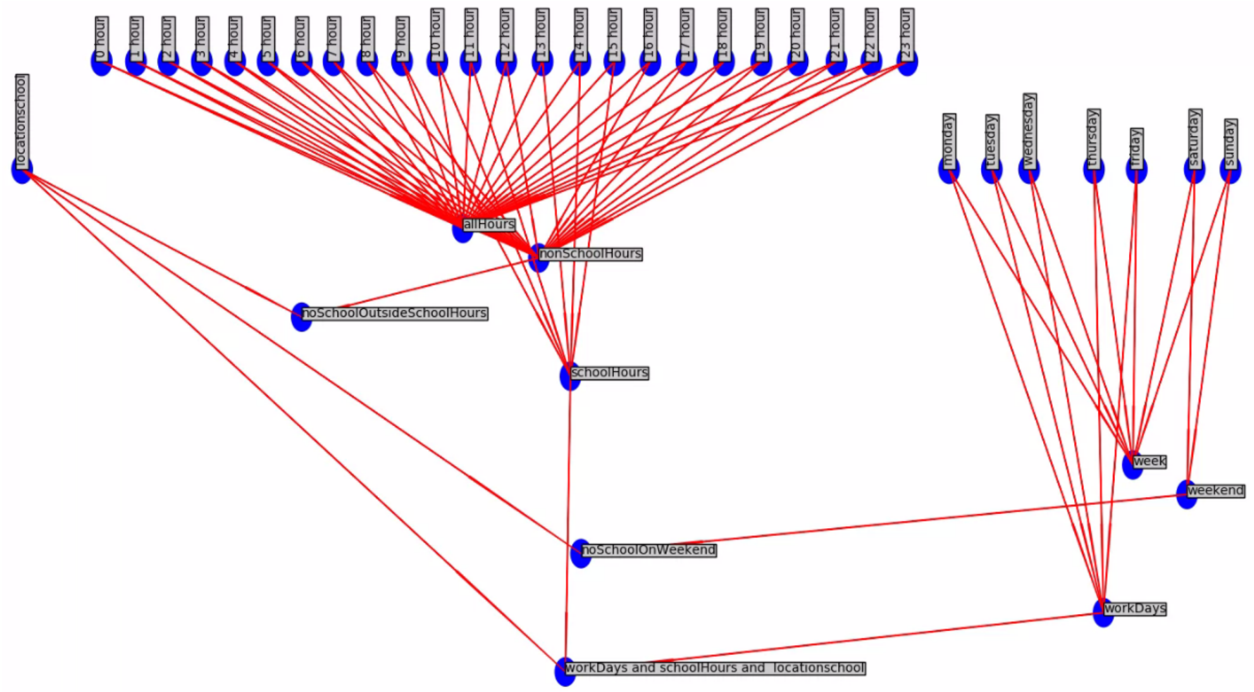


**Figure 4-1: Finding correlations in low-level data to obtain higher level concepts (top = low-level, bottom = high-level).**

In Figure 4-1 at the top there are two sets of input: the time and data represented by the 24 hours and seven days of the week, and, at the left, a signal that indicates whether people are in the school. By analysis of the data recorded over a period of weeks, and trying to find various combinations of AND, OR and XOR correlations, higher level concepts (towards the bottom in Figure 4-1) are found automatically. In the experiment shown in Figure 4-1 the highest-level node (at the bottom) the automatic process provides a high-level signal that indicates that "during working days, at school hours, people will be present in the school building". The automated analysis finds correlations and defines new nodes, the labels in the nodes (e.g. week, weekend, workdays, school hours, etc.) currently must be provided by an analyst. We are studying ways to compare the topology of nodes and relations with existing knowledge networks to propose labels of nodes to a user.

This approach offers an approach to self-learning concept networks that can translate large quantities of low-level data into meaningful data with the aid of an expert. It could indicate possibly unexpected correlations in the pattern of life that may be of influence on decision making. This experiment was conducted with clean data from one type of sensor, further experimentation is needed to see how resilient the approach is to noisy real-world data and how well it can scale with the inclusion of more sensors.

### 4.2 Combining machine learning with symbolic knowledge

The project Neuro Symbolic Modelling aims to combine the strength of sub-symbolic methods, here Neural Network modelling, with the strength of Symbolic modelling, see Table 4-1. From the desired system properties (section **2.2**) it can be derived that weaknesses of machine learning as well as symbolic reasoning may be overcome by the combination of both.

**Table 4-1: Strengths and weaknesses of various approaches (- undesirable, 0 is neutral, + is desired)**

| Property | Neural Network Modelling | Symbolic Modelling | NSM |
|---|---|---|---|
| Perceptual tasks | ++ | -- | + |
| Conceptual tasks | - | ++ | + |
| Explainable model | -- | ++ | ++ |
| Required dataset size | -- | ++ | 0 |
| Train time | -- | ++ | 0 |
| Performance | ++ | 0 | + |

The method we use is called Relational Concept Network (RCN) Reasoning. An RCN is a network of nodes with relations, which is designed by the military expert. Each node represents a (symbolic) concept and provides a value (belief) of the concept being true. The relations between concepts may be, but are not required to, be explicit (e.g. node A is related to node B, as opposed to: node A is a part of node B). Relations between nodes represent that the two related concepts influence each other (e.g. the presence of smoke and heat provides evidence for the presence of fire, thus when observing smoke and heat the belief of the presence of fire will increase). The relations will often be bidirectional (e.g. taking the previous example; when observing fire, the belief for the presence of smoke and heat will increase).

In a first experiment we augmented the outcome of a purely neural part with a RCN based symbolic part. The symbolic part provides the explainability. The results indicate that even when the neural part is trained with relatively little data points, it is difficult for human experts to improve on the end results. The reason for this is that humans can be expected to express their knowledge in a fuzzy form as "the probability of concept A being true increases if the value of concept B is high, and diminishes if concept C is low", but not as "concept A += 0.5703 * high(concept B) - 0.234 * low(concept C)". With the latter form it is possible to improve on the results, but generally not with the former.

Based on these results it is foreseen that to use the strength of both modelling approaches, these methods should be combined in a tightly interwoven manner. Expert knowledge (or semantic/otologic networks) is used to indicate which concepts are related without specifying how exactly they relate. Each concept has a small sub-symbolic part (e.g. a neural network) that can be trained to determine the weights with which concepts influence each other. See Figure 4-2 for a simple example. If neural networks are used, the whole network may actually be seen as one large neural network where concepts are neurons with a name.

Directly related concepts in an RCN have a shallow view, meaning that they only perceive the output belief of a node as an input value without having knowledge of how this believe has been constructed. The concept, in turn, can use the inputs in various ways: e.g. for regression or classification using machine learning methods (neural networks) or symbolic methods (a rule-based system): if an expert does have exact knowledge of how inputs are related to the concept, this can be used instead of the sub-symbolic approach.

Input can come from external sensors (which may or may not require complex data analysis) or manual input by e.g. an analyst. Each concept is “activated” when its input(s) are provided, and the internal process has produced its results (e.g. the beliefs for smoke and heat are provided to a symbolic rule which outputs its believe for the presence of fire). In turn this concept can be connected to another concept, which can produce a result (if the required inputs are provided). This way, available information can trigger a cascade of concepts being activated.

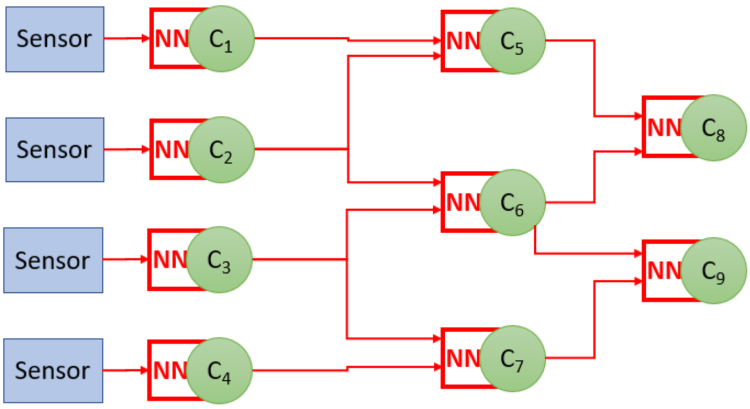


**Figure 4-2: Simplistic example of symbolic concepts (Cx) mixed with neural networks (NN).**

Even though Figure 4-2 shows a network where data flows from left to right, in practise connections may go in any direction, forming an arbitrary directed graph.

In a standard neural network approach, the whole network is trained using a full set of data. An RCN separates the train-step into multiple separate train steps, which compared to a neural network will require less training data, which in turn results in faster training times.

Using the RCN does not require all inputs to be provided. Only the available information is provided and those concepts which use that information as an input produce results, e.g. in Figure 4-2, providing values for C3 and C4 will also produce C7; if also data for C2 is provided C6 and C9 can also produce output. Hence, with even little information the RCN can provide new information and/or clues to what to expect as well (e.g. an audible fire alarm is detected, smoke is detected, and a loud noise has been detected. The network could provide feedback on that there is a high believe for the presence of fire, and that there is a high believe that firefighters will be present).

The performance of the model in terms of accuracy in the classification or regression task at hand, is likely to perform less well than a more traditional fully connected neural network. This is because a RCN cannot learn relations between concepts which are not expected by the designer.

Initial results indicate that it is possible to train this system piece by piece and obtain accurate results, although not as good as those for a fully connected network or decision tree. This network can indeed provide explainability by indicating which concepts triggered another concept, as well as indicate about which concepts information is available given a subset of input data. The RCN can indicate whether a rule provided by an expert performs well or not.

### 4.3 Towards Applying Decision Support using Neuro-Symbolic Modelling

The results of a neuro-symbolic model can give insights to an analyst or commander in the field. An example of this is a commander inputting the available expert knowledge of the situation into a pretrained, portable neuro-symbolic modelling system, that automatically collects the relevant sensor data. The system can then provide high-level insights such as expected positioning or behaviour of enemies, and identify impactful locations, such as gathering places, market squares and municipal buildings.

In addition, the (results from the) neuro-symbolic model can be used as an input component for other intelligence augmentation systems. At TNO, several initiatives exist to enable the use of this data for decision support. One example is a system that supports a commander and staff in better analysis of the mission environments and the effects of prospective battle plans. The neuro-symbolic model provides (part of) a conceptual model, which needs to be transformed into a computational model. Simulation uses this computational model to calculate a spectrum of expected results of actions, which after analysis can be used to select the best performing CoA.

In another project at TNO, predictive warfare methods are being developed to generate and evaluate CoAs for blue and red forces for dynamic situations using AI algorithms such as Deep Reinforcement Learning. Specifically, the project will be able to determine the most dangerous enemy CoA, the most likely enemy CoA, and generate different friendly CoA's to determine the chances of success, depending on available data about the current operational situation. The system will be able to convey the possible strategies, weaknesses and chances of success in a high-level to the commander and analysts. A neuro-symbol model enables the use of conceptual data and helps translate found strategies into high-level descriptions.

## 5.0 CONCLUSIONS

The envisioned simulation-based decision support system should have certain properties to be truly useful and trustworthy. One of the important aspects is that situational understanding can be translated into a simulation model. In this work we focus on the approach to obtain situational understanding using AI tools. We found that a combination of sub-symbolic and symbolic methods is necessary. However, combining sub-symbolic and symbolic approaches is not a straightforward, only certain integration strategies can potentially give the properties that are desirable for military decision support systems: explainability as well as data-efficiency, while utilizing the performance of sub-symbolic algorithms.

Because this approach is still in its infancy, further research is needed. Some identified areas of further research include the following: The amount, and correctness, of expert knowledge required to have a functioning system. Gradually training and growing the RCN while maintaining a stable behaviour. Robustness to small errors in knowledge from experts and observed data, which may include intentionally false information. Transferability of knowledge build-up in the RCN from one operation to another.

We expect these properties to (at least in part) hold for our framework thanks to the specific integration strategy of sub-symbolic parts per symbolic concept in a large network of concepts.

This Neuro-symbolic modelling system can be used in conjunction with other high-level decision support tools to provide AI-aided intelligence augmentation for military operations. Several projects within TNO are identified that can benefit from this approach.

NATO is working on several key technologies needed for the described decision support system, we hope to benefit from that work in the future. We propose that NATO renews its interest in this topic (there has been a NATO conference on neurocomputing in 1990!). Possibly via the cross-panel initiative on Big Data, AI and military decision making [6].

## 6.0 REFERENCES


[1] Veldhuis et al, "Collaborative problem structuring using MARVEL", EURO Journal on Decision Processes (2015) 3:249–273

[2] United Nations Department of Economic and Social Affairs https://www.un.org/development/desa/en/news/population/2018-revision-of-world-urbanization-prospects.html, 16 May 2018

[3] Glasow et al, "Megacities and Large Urban Agglomerations in the Coastal Zone", AMBIO February 2013, Volume 42, Issue 1, pp 13–28

[4] University Cologne, Megacity Commission of the International geographical Union, http://www.megacities.uni-koeln.de/_frame.htm?http://www.megacities.uni-koeln.de/documentation/start.htm

[5] Pham et al, "Prevailing in a Complex World: ARL's Essential Research Area on AI & ML", STO-MP-IST-160, 2018

[6] NATO STO, "Big Data and Artificial Intelligence for Military Decision Making", Meeting proceedings, STO-MP-IST-160, 2018

[7] Hannay, van den Berg, " The NATO MSG-136 Reference Architecture for M&S as a Service", STO-Meeting proceedings, MP-MSG-149-03, 2017

[8] van den Bosch, Bronkhorst, "Human-AI Cooperation to Benefit Military Decision Making", STO-MP-IST-160-S3-1, 2018

[9] Kerbusch et al, "Roles of AI and Simulation for Military Decision Making", STO-MP-IST-160 PT-4, 2018