Geospatial analysis for Machine Learning in Tactical Decision Support

Nico M. de Reus, Philip J.M. Kerbusch, Maarten P.D. Schadd
The Netherlands Organisation for Applied Scientific Research TNO
THE NETHERLANDS
nico.dereus@tno.nl, philip.kerbusch@tno.nl, maarten.schadd@tno.nl

Maj. Ab de Vos
Ministry of Defence JIVC/KIXS
THE NETHERLANDS
a.d.vos@mindef.nl

ABSTRACT

Tactical military land operations heavily depend on the terrain, thus the terrain is always taken into account in the military decision making process. Terrain related (geospatial) tactical information products, such as optimal routes or avenues of approach are usually determined by terrain analysts in the Intelligence cell, however automated generation is possible as well. These products can be used in decision support tools to support the planning process. When machine learning is used in these decision support tools, the products can also be of benefit for modelling the behaviour of military units that is required for finding well-performing courses of action by machine learning. This work presents an overview of geospatial products and classifies them into a tier-based architecture in which products are based on products of underlying tiers. We furthermore formalize the steps of creating tactical terrain models and tactical mission models that are required for machine learning. Based on two practical examples we demonstrate how geospatial products can be generated in the proposed architecture, how these products can be used in machine learning for tactical planning, and how the learned courses of actions and intelligence products can be supplied to the planner in support of decision making.

1.0 INTRODUCTION

Military planning takes place at levels ranging from operational/joint planning down to tactical planning. The military decision making process at the tactical level calls for commanders and staff to analyse the mission and intent two levels up. Therefore Battalions or Brigades can copy and paste their higher headquarters mission statement into their own mission analysis. Due to the nature of tactical land combat operations, the guidance from these higher levels of command is mainly terrain related and the analysis of terrain in land-combat-operations decision making therefore is of great importance. For instance [1] shows the importance of terrain analysis by providing historical examples of battles where the impact of terrain on the battle was large. The concept of terrain analysis is part of the so-called "Intelligence Preparation of the Battlefield (IPB)"[1]. It is defined in [2] as "the systematic process of analysing the mission variables of enemy, terrain, weather, and civil considerations in an area of interest to determine their effect on operations". IPB is a collaborative staff effort that is led by the intelligence staff (S-2/G2).

The terrain analysis part of the IPB delivers so-called "geospatial products". These, usually hand-crafted, products are central to the decision making process taking place in the other staff sections and staff officers use them in their manually performed processes.

This paper focuses on the idea that these geospatial products can have a wider application in the automation of the planning process because the abstracted terrain view that these products provide can serve as a world model that can be leveraged in the automation or at least support of military decision making. The reason is

[1] IPB is sometimes also called IPE (Intelligence Preparation of the Environment) in order to include non-combat missions.

that currently emerging artificial intelligence (AI) applications for military decision making need such world models to perform their optimization processes. Therefore the products that result from the IPB can be used as a crucial step towards the support of decision making using AI technologies.

On the one hand efforts are ongoing to automate the generation of geospatial products and on the other hand efforts are going on to perform automation in support of decision making steps. Therefore combining these efforts can result in a future where most of the process steps, from the initial terrain description to a Course-Of-Action (COA) advice can be performed automatically. This is the vision we present in this paper.

Section 2 describes which geospatial products are used in the planning process and also how automated production of these products will support the military planner. After that in Section 3 the concept of world modelling that is required in machine-learning-based concepts for plan optimization is discussed and finally in Section 4 two examples are shown. Section 5 concludes by describing future possibilities.

2.0 GEOSPATIAL ANALYSIS AND THE MILITARY PLANNING PROCESS

To get an overview of the available geo-spatial products, we first look closely at the IPB analysis in Section 2.1. Section 2.2 classifies these products into several tiers, following ideas from the US Battlefield Terrain Reasoning and Awareness program. Each tier depends on its underlying tiers as input. The most basic tier (Tier-0), which does not depend on any other tier, is discussed as well.

2.1 Intelligence Preparation of the Battlefield (IPB)

The Intelligence Preparation of the Battlefield (IPB) is defined in [1] as *"the systematic process of analysing the mission variables of enemy, terrain, weather, and civil considerations in an area of interest to determine their effect on operations … it results in intelligence products that are used during the military decision-making process (MDMP) to assist in developing friendly courses of action (COAs)."* It is performed in the following four steps:

1. Define the Operational Environment.
2. Describe environmental effects on operations.
3. Evaluate the threat.
4. Determine threat Courses Of Action (COAs).

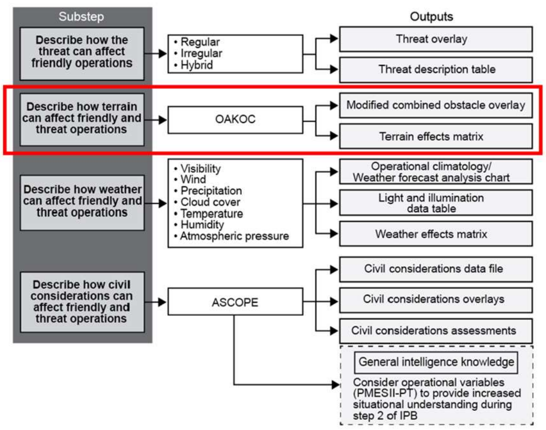

Figure 1 Sub-steps of IPB step-2, in this paper we only concentrate on the OAKOC part (the red rectangle) (figure taken from [1])

Terrain analysis is the key element of step-2 of the IPB which is worked out in detail in [3]. This step consists of several sub-steps which are shown in Figure 1. The second sub-step analyses how terrain can affect operations through a so-called OAKOC analysis (Obstacles, Avenues of Approach, Key terrain, Observation and fields of fire, Cover and concealment). Two main products that are determined after the OAKOC analysis are the so-called Modified Combined Obstacle Overlay (MCOO) and the terrain effects matrix. An MCOO enables the decision maker to see where mobility corridors are, such that the most relevant ones for the mission can be chosen and the impact on the mission be discussed. The terrain-effects matrix describes OAKOC factor effects on operations.

2.2 Leveraging terrain analysis products from the IPB for tactical decision making

Several nations have worked on automation of the terrain analysis part of the IPB and on leveraging the terrain analysis tools and products that are normally the domain of the terrain analysts for tactical decision making. Bringing terrain analysis products into the tactical domain of the field commander was for instance the aim of the (2002) US Battlefield Terrain Reasoning and Awareness (BTRA) research program [4]. The (2006) US Geospatial Battle Management Language (geoBML) project [5] focused on providing a semantic and syntactic bridge between the domain of terrain reasoning and analysis and the domain of the operational commander's tactical decision-making process by extending the lexicon of the Battle Management Language (BML)[2] [6, 7]. The Common Ground project [8] took this a step further by also supporting the military decision maker with digital orders and so-called engineered knowledge (a database of tasks and capabilities by unit, echelon, service, and nation).

Tactical Spatial Objects (TSOs) are the key products that provide the link between the world of the terrain analysts and the world of tactical decision making. They are defined as objects that are developed with geographic information systems that directly support the planning of tactical operations and that may contain the relationships to specific military operations, missions and tasks as well as various types of military organizations. They are categorized in layers, called tiers and can typically be calculated using BTRA's Commercial Joint Mapping Toolkit (CJMTK) which is the replacement of the Joint Mapping Toolkit, previously used by the US Department of Defence. The CJMTK uses commercial of-the-shelve components to calculate these products based on foundational terrain data which comprises an integrated description of terrain characteristics that are sufficient for generating higher level products. The way these features are characterized is related to the foreseen tactical use of the terrain. For example, it is not sufficient to state that a certain area is a forest because, in order to travel through forest, the tree density must be known and therefore this must be part of the foundational data. The purpose of TSOs is not to replace humans with automation in regard to the geospatial dimension of mission command, but rather to allow commanders to evaluate geospatial variables more quickly.

Inspired by the tier structure used in the BTRA project and extending this structure with foundational data, which we have called Tier-0, the TSOs can be divided into the following tiers:

- **Tier-0 TSOs** are the foundational products, comprising an integrated description of terrain characteristics.
- **Tier-1 TSOs** are based only upon the terrain and can be pre-computed without being informed by the other factors of METT-TC[3]. These TSOs in general concern large data sets, are generic/reusable and primarily static, i.e. not dependent on the dynamics of the mission environment such as weather.
- **Tier-2 TSOs** can be derived from foundational terrain data and Tier-1 TSOs. They are mission specific as they depend on the tasks that a unit needs to perform in that mission. Therefore these TSOs cannot be precomputed before mission information becomes available. These TSOs in general are fine grained,

[2] BML was the language for connecting C2 systems and simulations that was under development at that time, currently the BML standard has been replaced with the C2SIM standard [21]

[3] METT-TC: Mission, Enemy, Terrain and Weather, Troops and Support Available, Time Available and Civil Considerations

concern information rather than data, are generic/reusable for different COAs but not across missions, and are more dynamic, i.e. more dependent on the dynamics of the mission environment such as weather.

- **Tier-3 TSOs** are specific objects that have been selected to support a specific COA and are associated with a plan or order. In many cases they have been chosen from the Tier-2 candidate TSOs and been further refined based upon METT-TC. They may also include graphic control measures and other items that are often associated with or influence the perception of terrain.

Tier-1 products either are solely computed from Tier-0 products or partly dependent on other Tier-1 products. Tier-2 products are computed based on Tier-1 products and Tier-3 products can either be selected from Tier-2 or Tier-1 products. All these characteristics are visualized in Figure 2.

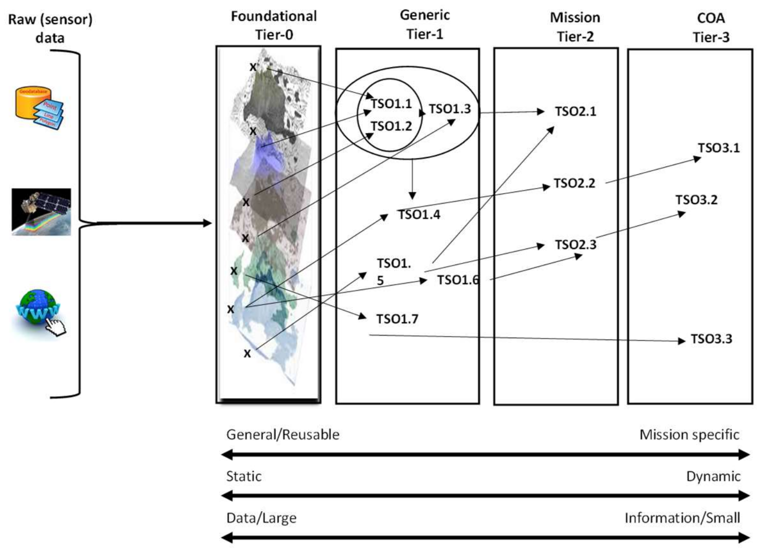

Figure 2 TSOs, Tier structure and tier characteristics

Examples of TSOs are:

Tier-0 Digital Terrain and Elevation Data (DTED) and other terrain features such as infrastructure, hydrology, vegetation, soil types, surface water, roads. Tier-0 products can be generated in a sequence of steps starting from raw sensor data, which we omit from this article. This data can be collected using (geo) data sources [9], satellites [10], drones [11], from open sources on the internet [12], amongst others.

Tier-1 Cross Country Mobility; Combined Obstacle Overlay; Cover and Concealment, Manoeuvre Networks.

Tier-2 Assembly Areas, Attack by Fire Positions, Avenues of Approach and Indirect Fire Positions.

Tier-3 These are selected and possibly refined Tier-1 or Tier-2 products for use in a COA. Examples are Phase lines, lines of contact, engagement priorities and criteria, checkpoints, section boundaries.

Many examples of Tier-1 and Tier-2 products are given in [3]. The following summarizes a few to give an impression (illustrated in Figure 3):

- **Cross Country Mobility (CCM)** demonstrates the off-road speed for a vehicle as determined by the terrain (soil, slope, and vegetation) and vehicle performance capabilities; however, it does not consider the effects of roads and obstacles.
- **Combined Obstacle Overlay (COO)** integrates obstacles to vehicular movement (built-up areas, slope, soils, vegetation, hydrology).

- **Mobility Corridors** are a combination of cross-country mobility, transportation, and linear obstacle overlays to show mobility corridors that are based on the restrictiveness of the terrain, vehicle capabilities, and preferred movement formations.

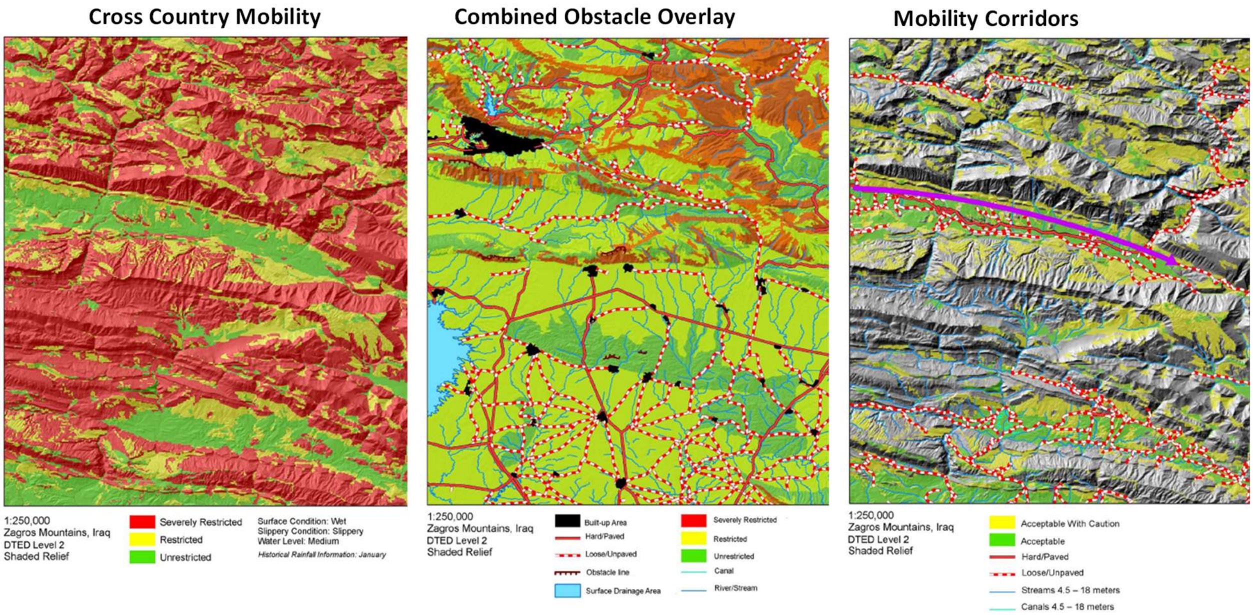

Figure 3 Three TSO examples (CCM, COO, MC) (figures taken from [3])

Another example Tier-2 TSO which combines many underlying TSOs is the so-called Modified Combined Obstacle Overlay (MCOO). The MCOO generation depends on the following Tier-1 products:

- Cross-Country Mobility.
- Avenues of Approach / mobility corridors.
- Counter-mobility obstacle systems.
- Combined Obstacle Overlay
- Defensible terrain.
- Engagement areas.
- Key terrain.

An example of an MCOO is visualised in Figure 4. Here, the Avenues of Approach (which depend on the size of the unit) are clearly visible which is an important TSO for the example application that will be discussed in Section 4.

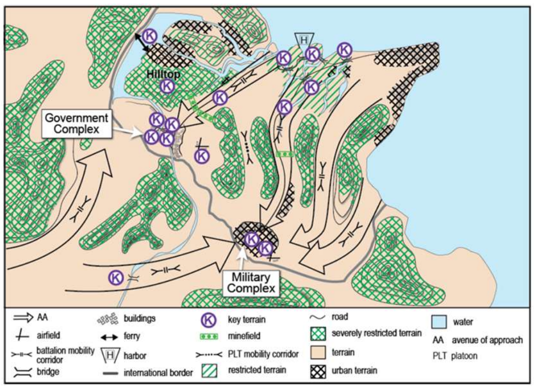

Figure 4 Example MCOO (figure taken from [1])

3.0 OPTIMIZATION OF MILITARY COURSES OF ACTION

This section investigates how, in the planning of military operations, military COAs are optimized and how the optimization steps, that currently are primarily performed manually, can be automated. The section starts with describing the steps after which the automation possibilities are discussed. Since, for automation of COA optimization, simulation models are essential, this is discussed separately.

3.1 Military COA planning

The following figure visualises the steps that are currently used, whether implicitly or explicitly, in the optimization process of COAs where in general, COAs consist of many elements that can be represented by Tier-3 products. Also visualized is the type of information that needs to be provided in the intermediate steps. For the generation of Tier-2 TSOs this is mission specific METT-TC data, but for the optimization process, more information is required because the optimal solutions usually depend on objectives (both from own troops as well as enemy), unit capabilities and rules of engagement.

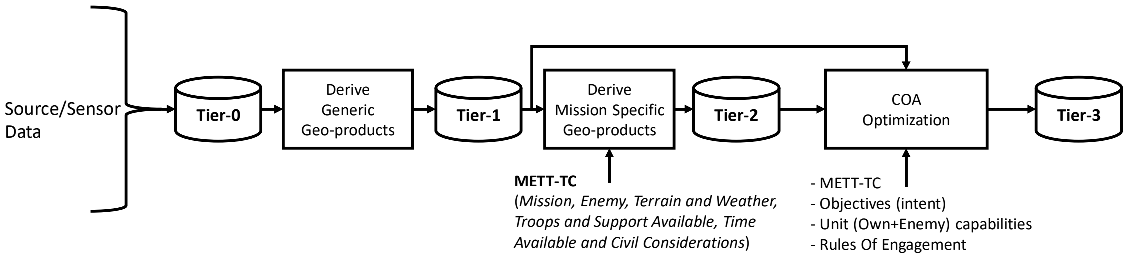

Figure 5 Description of the process steps for finding optimal COA elements (Tier-3 products) using the tiered approach.

Currently the planners in the plans cell (S5) take the information that is provided by the Intel cell (S2) and use that in combination with the mission objective to find best Courses Of actions (COAs). Translating this into TSO terms and referring to Figure 5 this means that the Intel cell provides Tier-1 and Tier-2 products which are used by the planners to find optimal COA elements (Tier-3 products).

Generating generic geo-products (Tier-1) and Mission specific geo-products (Tier-2) products as done by the Intel cell currently is a mainly manual process, although pieces of this process are currently automated by military GIS applications.

Generating Tier-3 products currently in practice is a fully manual process. Since humans can only hold a limited and small amount of data in their mind at any point in time, in order to increase the understanding of any problem, humans look for patterns and abstractions that help in reducing the need for memorizing data and create useful information. For these abstractions, map overlays are created that abstract a single aspects of the underlying, highly complex terrain. These overlays are used by the planner, who, given the military objective, foreseen enemy intent and doctrine and available units and their capabilities, uses an implicit model to envision a best COA. An example of a tactical important aspect is the presence of water in the form of rivers and lakes. The created overlay highlights this terrain feature with blue shapes, and only significantly large water features are added to the overlay to reduce mental load. Once enough of these overlays have been created, the entire tactical reasoning will be performed using the overlays and not the underlying data.
Another example is the case of finding optimal casualty collection point locations and routes. The route and collection points need to meet certain conditions, such as being close to one another. The overlays are used to identify several suitable collection point locations and several attractive routes. The most combinations that are likely to work well are compared to one another and one chosen.

3.2 COA planning automation

This section describes the possibilities for automating the process steps described in the previous section.

Currently the generating Tier-1, Tier-2 products in the military world currently is only done partially. A gradual shift towards more automation can be seen within NATO armies with the ultimate future vision of full automation.

For the COA optimization, that is, the generation of Tier-3 products based on Tier-1 and Tier-2 input, this is a different matter. In the manual casualty collection points optimization process described in the previous section the example is a simple case in which a human only needs to find a solution to a problem with a few degrees of freedom. In reality the potential solution space can be enormous. Take for instance that there exist 20 suitable collection points and 30 likely routes. A human cannot consider all 600 combinations and will prune combinations based on (possibly fallible) heuristics. With more variables, such as fire support, extraction points, unknown hostile movement, the probability of good solutions being discarded by the planner increases.

Automation of optimization of COA elements can be done using Artificial Intelligence techniques, which basically are search techniques for finding optimal solutions in large solution spaces. In order to define what the ‘optimal’ solution entails, an evaluation function has to be designed that describes the attractiveness of each COA solution based on its properties. In a military setting this is usually called the Measure Of Effectiveness (MOE). For a military planner the COA solution may consist of various (configuration) settings of a given COA, such as what actions to execute and in what ordering. In many situations, the function between the COA’s input features and evaluation score cannot be determined directly because, especially in manoeuvre warfare, time/space considerations are required and actual simulations need to be performed to find the MOE of a certain solution. Other difficulties can arise from the fact that the function from COA input space to MOEs is hyperdimensional and the solutions space is not a convex surface, making it hard to find the optimum. In such occasions machine learning techniques can be employed to find well-performing values for the input features. Example search techniques that aim at finding good solutions are Genetic Algorithms and Reinforcement Learning, see e.g. [13, 14]. In order to be able to use such techniques, simulation models are required. The next section discussed these.
Future vision: We envision a future where the process steps as depicted in Figure 5 can be performed automatically to a high degree. When also the steps to generate Tier-0 products could be automated, this vision can be presented as the process *“From satellite image to plans.”*.

3.3 Simulation modelling for COA optimization

In order to perform the automation as described in the previous section, (simulation) models are required that can be executed in order to evaluate a chosen COA and determine the corresponding MOE. A requirement for such simulations in a decision support context is that it can be computed fast enough because (1) in general, not much time is available to set-up a simulation environment and (2) machine learning technologies in general require many simulations to be performed (depending on the technology used this can range from thousands to millions of simulation runs) and (3) the context of the mission can change and may require a quick response of the system to deliver decision support in an evolving mission. Therefore simulation models should be as simple as possible but still have a level of detail that captures the essence of the mission.

Simulation models currently in use for the military domain have been mainly based on simulating “from the dirt level” as described in [15], which means that the terrain is usually modelled in sufficient detail to enable a realistic movement of units through terrain which in general is not required for COA optimization. However, these types of models in general are too detailed and too slow for optimization purposes because

(1) setting up these simulations requires extensive configuration and (2) executing simulation runs is too slow to perform many simulations.

Therefore, for plan optimization, the idea that is presented in this paper is to integrate M&S with GIS information and thus enable simulations to be performed at the tactical level rather than at the "dirt" level and use sematic information rather than basic terrain representation as used in many simulators. Next to this environment modelling, also unit behaviour needs to be modelled. Therefore, in order to facilitate plan optimization, we recognize that the we need (abstract) tactical *mission* models. Such mission models need to contain aspects of terrain as well as aspects of the entities involved in the mission. Therefore we differentiate between *mission* models and *terrain*[4] models. Tactical *terrain* models and tactical *mission* models are described below in more detail.

Tactical Terrain Model (TTM): This model is the set of Tier-0, 1, 2 TSO products abstracted in such a manner that they can be used to simulate units quickly, yet sufficiently realistic. Instead of units moving through detailed terrain they move along a grid, nodes, edges, cells or another form of an "abstract terrain". The abstracted terrain has to capture the essence of the "dirt" level as far as this is relevant for the entities involved. These models are mission dependent because they depend on Tier-2 products and also because terrain characteristics that relate to the mission are part of the model. For instance, in a delay operation, aspects such as maximum reachable delay as a function of geographic location and unit to be delayed can be part of a TTM.

Tactical Mission Model (TMM): Units and their behaviour are modelled in such a way as to enable them to operate in the tactical terrain model. The abstraction of the environment imposes an abstraction of the actions. For example when the terrain is captured in nodes and edges, units can only be placed on nodes, and move to other nodes via an edge. It is not possible to leave the graph, although in the real world moving to locations outside the graph edges would be perfectly feasible. When military units interact (e.g., own versus enemy) the interactions are modelled via "game" rules such as who can fire at whom, with what speed can be moved, what are detection and firing ranges, what are the relative strengths when units fire at each other, etc. Capturing the dynamic and complex reality in such game rules can be a challenge on itself, but is essential for optimization. Figure 6 visualises TTM and TMM

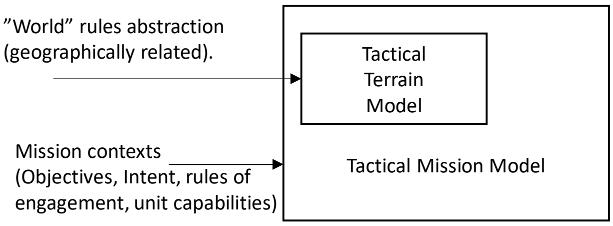

Figure 6 Tactical Terrain Models (TTM) and Tactical Mission Models (TMM), their interrelation and dependency on

The basics of the idea of tactical terrain modelling for different purposes than COA optimization have already been investigated in [15], [16], [17]for training simulators. In [18] such types of models are also described but in this case for yet another purpose, namely to abstract the context of a detailed simulator to enable the simulation of intelligent agents within this simulator. It also describes ideas for using this information for decision makers in the operational process.

[4] Note that we currently call this: terrain model while in the future, when other aspects of the mission, such as cyber, are also taken into account, the more generic term "environment model" may be used.

4.0 IMPLEMENTATIONS OF AUTOMATED COA OPTIMIZATION

We describe two example cases, namely (1) the automated “engineering effect” plan and (2) the small-unit plan optimization, both using the concepts described above, namely the use of environment and tactical modelling in order to facilitate the use of AI optimization techniques for optimal planning.

4.1 Creating an effect plan based on an Avenues of Approach overlay

The Netherlands Combat Operations Support with Modelling & Simulation (COSMOS) project works on AI based decision support for land operations. The following use-case was part of this project.

The military engineer is able to place effects in the physical landscape to hinder the approach of an opponent. For instance in a defensive operation, delaying the opposing forces may be the main objective. The military engineer then is tasked to study the terrain and find the best placements for effects. The engineer takes knowledge of the opponent, the terrain and of course own capabilities and limitations into account. The number of combinations (what effect is placed where) is too large to comprehend for a human, and uncertainty and ambiguity in observations make it even harder to think through all possible solutions. Thus, decision support by an AI seems a valuable addition.

The possible effects that can be placed are **block, fix, disrupt and turn**. With a **block effect**, the opposing force is hindered in such a way that they cannot break through the blockade (during the relevant mission time). With a **fix effect** the opponent is fixed on the current position for a certain amount of time, but is able to continue its journey later. A **disrupt effect** temporarily disrupts the approach of the opponent, but the opponent is able to continue the approach quickly. With a **turn effect** the opponent is being suggested to take a desired direction by blockading the other directions in such a manner that the time investment to take those directions is too large.

The above described effects cost resources to be achievable. Examples include mines, concertina wire, time, fire support. The amount of resources needed are heavily dependent on the terrain characteristics of each location. For instance can a block effect be achieved much easier when natural obstacles such as a lake or river form natural blockades, and only a single road or bridge needs to be blocked by the engineers. Blocking an enemy in an open field is possible, but very costly to achieve, as then bulldozers need to dig long trenches so that the opponent cannot simply travel cross country to circumvent the obstacles.

A military engineer needs to take all these factors into account and at the same time be robust against the most-likely, most-dangerous and all other enemy courses of actions that might occur. An AI-based decision support system may be able to support the engineer by optimizing the obstacle plan to achieve the maximal delay of the opponent while adhering to the constraints of its own capabilities in the given terrain. The next steps describe a system where information products from Tier-0, Tier-1 and Tier-2 are used to support the engineer’s decision making. An overview of such a system is shown in Figure 7. A discussion of the steps is given below the figure.

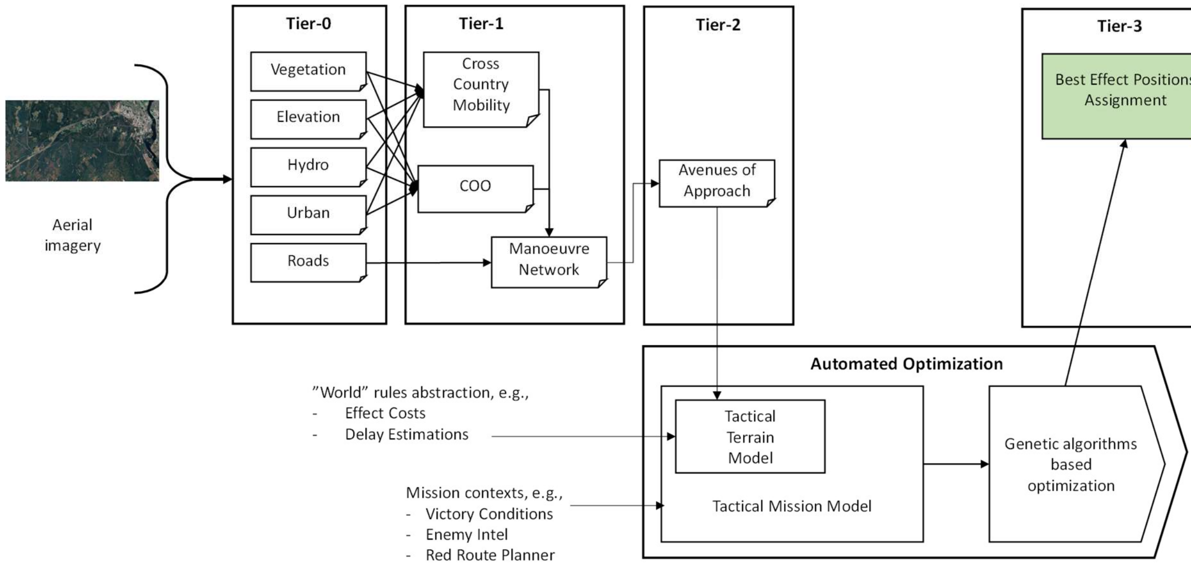

Figure 7: Overview of the products involved in the COSMOS prototype. The green Tier-3 TSO product was generated by the Automated Optimization process, in this case using a genetic algorithm.

Tier-0:
Based on for example satellite images a foundational model of the terrain is created. This model includes information such as roads, soil composition, vegetation, building locations. This model can be created weeks before the mission, as the real world is not changing that quickly.

Tier-1:
The foundational model (Tier-0) is analysed to generate a Combined Obstacle Overlay (COO) based on natural and man-made obstacles for the current mission context. Also a manoeuvre network is generated that takes into account both the Cross-Country Mobility (CCM) network as well as the road network. Such a manoeuvre network can be precomputed for different kinds of vehicles and different weather conditions and when the units are known requires only retrieving the characteristics for the specific unit types. For instance, with loose soil, traveling cross country with wheeled vehicles may not be feasible when it has been raining the past weeks, while travelling with a tracked vehicle this may be possible. With the estimated composition (Orbat) of the opponent as determined by the intelligence cell in mind, the COO can be annotated on the map as an overlay. An example of such obstacles is shown in Figure 8. In this terrain, dense forests and densely populated cities are marked as unpassable (hatched areas). Also a river and a bridge is annotated.

Figure 8: Natural obstacles based on terrain: dense forests, cities, river.

This task is commonly performed by the military engineer manually. This takes time however, and many maps, satellite images and sometimes field studies are performed to have accurate results. Here automated Tier-1 analysis tooling can in the future create a first calculation of such obstacles. For this study, we created this manually with the support of a military engineer.

Tier-2:
With the manoeuvre network and COO in hand, all possible manners of movement in the current terrain and mission context are known. Given this and the mission characteristics (objective area of opponent) and opponent's characteristics (unit size), the so-called Avenues Of Approach (AA) Tier-2 TSO can be generated. In our case this was still done manually, but we foresee that this will be possible to be automated in the future. Such AAs indicate the main routes that the opponent may travel on, given its size and objective area. Therefore, in order to delay the enemy, these avenues should be targeted with effects to achieve the desired delay.

Figure 9: Manoeuvre Network (left) and Avenues of Approach (Right).

With these AAs as input, suitable positions for fire support can be determined. Such fire support positions should be able to cover the avenues of approach in order to increase the effectives of the effects that are placed on the avenues of approach. In our use case we assumed that adequate firing positions can be found. Although we did not optimize firing positions, we want to mention fire support here as it is an essential part of making an effect for delaying the opponent more efficient.

Tactical Terrain Model and Tactical Mission Model:
The tactical terrain model consists of two main components. (1) For each location, the costs required to achieve an effect on a certain position is needed. Given the terrain, how many mines, man-hours, wires, explosive charges are needed to create a certain effect on a certain location. The terrain properties play a large role in the effect costs. Next to the costs, (2) the achieved delay that a chosen effect on a location creates has to be determined. In case of turn effects, rather than delaying an opposing force, the movement directions are being limited. Information from Tier-1 and Tier-0 products may be analysed to set appropriate values for these parameters.

In the tactical mission model we define the victory conditions. In our case the victory condition for Red is achieved when the second unit arrives at an important bridge. We choose the second unit as in our scenario one unit has not enough momentum to conquer the bridge, and only when two red units work together is the bridge lost. Another part of the mission model is the enemy intel. In different settings we experimented with the enemy have various reconnaissance capabilities that enabled red to spot obstacles beforehand and choose a circumventing route. The route preferences of the red force, according to a route planner, describes the likely choices of the opponent in the manoeuvre network.

Tier-3:
Given the Tier-0, Tier-1 and Tier-2 products (the possible avenues of approach, where effects can be placed and what the cost of these effects are in terms of resources for each location). Given these models, the following constraint satisfaction problem can be formulated “Given that the average delay of the opponent should be maximized and given the constraints in terms of resources, what are the most effective effect location and types?” To answer this question, we employ a genetic algorithm that iteratively improves a (set of) solutions.

Such an optimized solution can be used by the military engineer as inspiration, or confirm the effect position that the engineer him/herself envisioned. Additional functionality may include features such as what-if analysis and explanations why certain effect setups do not work well against a (set of) enemy courses of action.

4.2 The ZebraSword experiment

The Defence research programme AI for Military Simulation performs R&D on AI based decision support for military operations. The following use-case was part of this project.
ZebraSword was a Royal Netherlands Army 2020 training and experimentation exercise for robotic platoon operations. One of the scenarios during the exercise covers a *remote controlled ambush*, wherein several technical and procedural innovations were evaluated such as remote standoff engagements. In the scenario the opponent platoon arrives from the north and wants to capture a village in the south. The friendly platoon’s task is to prevent the enemy from capturing the village, making optimal use of terrain features and unit capabilities. Inspired by this scenario we developed a prototype decision support tool based on the automated COA optimization approach described in Section 3.2 with the primary objective to aid the commander in selecting the best ambush positions for his/her forces (for two infantry groups and a single remote controlled heavy weapons platform). A selection of the TSO products generated by this decision support tool to support the platoon commander in mission analysis and planning, can be seen in Figure 10 below. These include a heatmap of likely enemy contact positions, a map of routes to follow towards ambush positions, and what most likely enemy routes towards their target location would be.

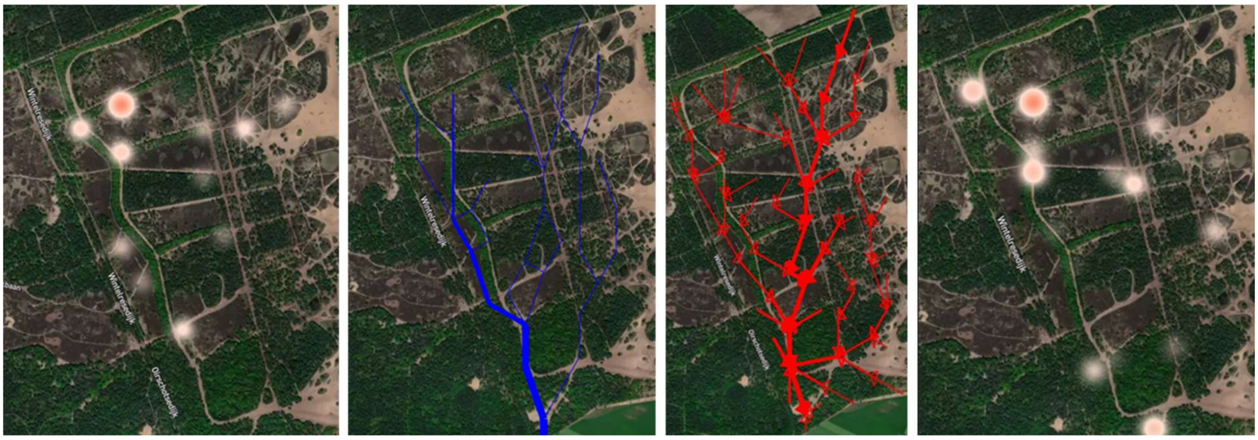

Figure 10: Decision Support for the Platoon Commander. From left to right: heatmap of likely enemy contact position, preferred routes towards ambush positions, for each position the most likely enemy travel direction towards their target, and the top ambush positions.

The next steps describe the prototype system in terms of the automated COA optimisation approach presented in Section 3.2, of which an overview is given in Figure 11. With this approach products from Tier-0 and Tier-1 are used to create a Tactical Terrain Model and Tactical Mission Model that can be optimised using a combination of *reinforcement learning* techniques to generate the desired Tier-2 and Tier-3 TSO products for decision support.

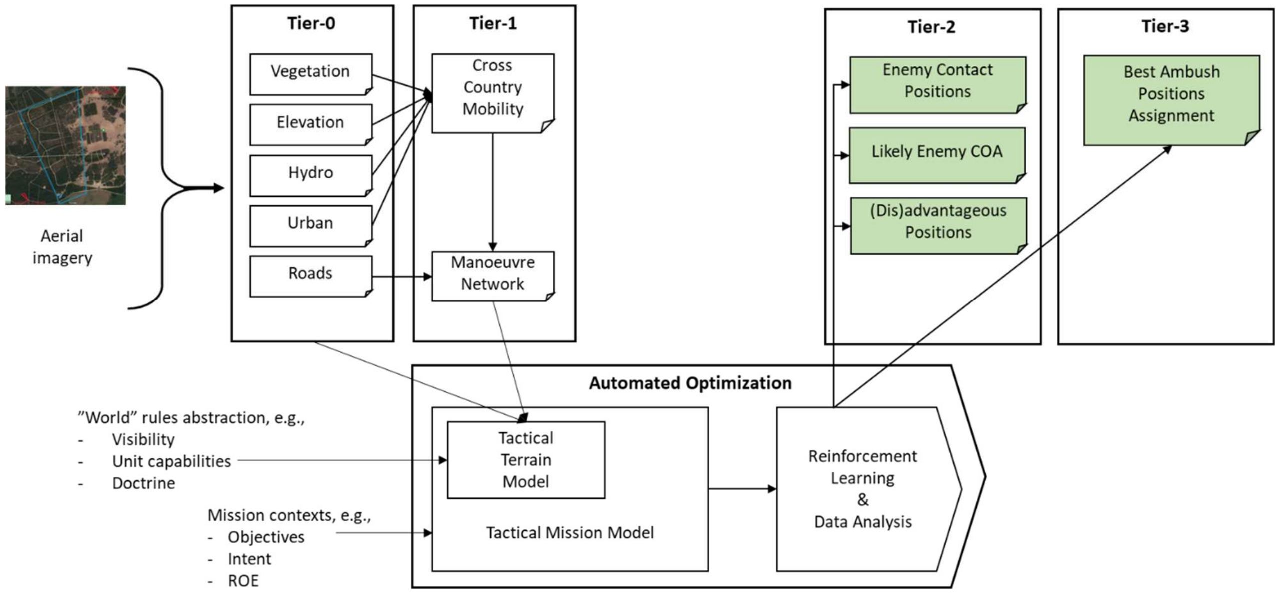

Figure 11: Overview of the products involved in the ZebraSword prototype. Green Tier-2 and Tier-3 TSO products are generated by the Automated COA Optimization process using Reinforcement Learning.

Tier 0:

A foundational model of the terrain was not created explicitly for the prototype. Nevertheless, Tier-0 foundational data that was used to build other TSO's is manually derived from aerial imagery and includes elevation, vegetation, hydro, urban and road maps.

Tier-1:

Based on manual analysis of the Tier-0 products, several Tier-1 products can be created. These include a Cross Country Mobility and Manoeuvre Network map.

Tactical Terrain Model and Tactical Mission Model:

From these products and the Tier-0 products, the Tactical Terrain Model is created. In this case the world is manually abstracted into hexagons with diameter of 100 meters, which corresponds roughly to the area a single group within a platoon can cover. Each hexagon holds properties such as the terrain type (plains, forest, water, urban), traverse cost, whether it can provide cover, and whether it blocks line of sight. Overlayed on the hexagon map a sparse manoeuvre graph is placed based on the Manoeuvre Network, in such a way that units can use the terrain to their advantage, for instance units in ambushes positions in forest edges have an advantage because such a unit is able to see any units on the nearby planes, while itself remaining invisible. In order to facilitate that, the nodes in forest edges were given the characteristics to enable these concealment calculations.

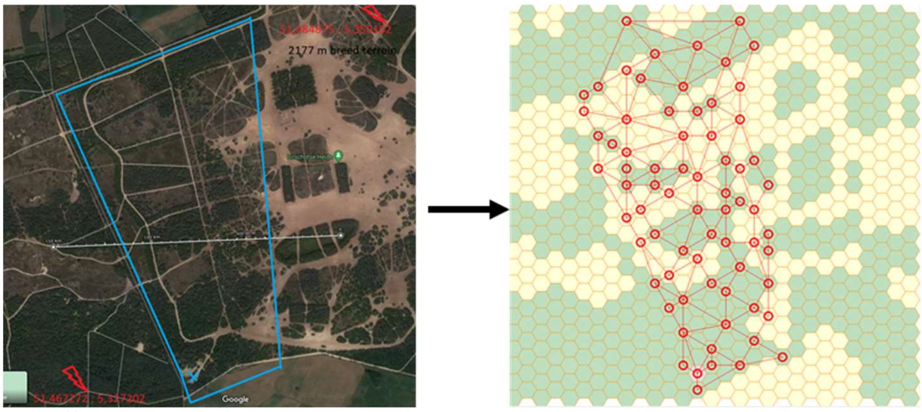

Figure 12: Transformation of the raw data to an abstracted Tactical Terrain Model using hexagons and a sparse manoeuvre graph.

For the tactical mission model, we modelled infantry units that have properties such as health, traveling speed on different terrains, sight range (varying depending on whether the unit is in open area, in a forest or in the edge of a forest), weapon range, munition, amongst others. Units move in a turn-based fashion, but are able to react first to opponents when lying in an ambush. The red force has as goal to reach the village in the south, and the blue force has as goal to neutralize red. Although this is a very rough abstraction of the real world, we can argue that similar tactics have to be employed to win the abstract game as well as in the real world. We tested this during a demonstration with platoon commanders which found the results of the analysis logical.

Tier-2:
Desired (unofficial) Tier-2 products the prototype decision support tool should provide include a map of engagement positions where enemy contact is to be expected, several likely enemy courses of action (routes), and (dis)advantageous positions. These products depend on terrain, but, as well as the primary Tier-3 product, are also heavily dependent on the mission context (objective, doctrine, units) of both friendly and (expected) enemy forces. Therefore, these products lend themselves well to be generated using an automated COA optimisation approach. Rather than manually modelling this behaviour to predefined rules, we choose to have *reinforcement learning* algorithms optimise behaviour of both forces, similar to our previous work described in [19]. This *co-learned* optimal behaviour is then recorded, from which the required products can be created and visualised for the commander.

Tier-3:
The primary required product is an (advisory) assignment of all the commanders forces to their respective ambush positions. Using the automated COA optimisation approach described in Section 3.2 these can be easily obtained from the learned optimal behaviours. Interestingly, it is also possible to do several *what-if* analyses. For example, during optimisation the intent of the enemy was supposed to be unknown and therefore varied (i.e., *be fast* vs. *be safe*). It is thus possible to get ambush position assignments, or any Tier-2 product, given assumptions on both enemy intents.

Decision support:
Visualizing optimized behaviour is not straightforward, as the behaviour is internally presented by many parameters instead of a human-readable explanation. It is also behaviour that is robust against a wide set of opponent tactics, and it is hard to understand generic tactics in cases when it fails against a specific tactic. We choose to create explainable results by letting the behaviour of both parties play against each other many times, and provide customizable views and statistics to the platoon commander, as seen in Figure 10.

At this point we want to stress that the behaviour learned by an algorithm based on abstracted models may not resemble the human opponent. Humans follow doctrines, ideology and intuition, while machine learning algorithms optimise a goal function in a simulated environment. Such an algorithm is prone to overestimating its own performance and it may stumble upon bugs in the simulator or find a loophole in the Tactical Terrain or Mission Model and exploit it with unrealistic outcomes as a result. Nevertheless, while AI-technologies are being continuously improved, an application of decision support as shown can already prove useful now when the learned behaviour is realistic enough to challenge implicit assumptions of the platoon commander and widen his/her view during analysis or planning.

5.0 CONCLUDING REMARKS AND FUTURE WORK

Supporting the tactical military decision making process with (machine learning) optimization techniques requires tactical mission modelling. The currently available simulation models in general take too much detail into account and more abstract models are required. We have recognized that products that are developed during geospatial engineering as performed in the intelligence preparation of the battlefield process can support the more abstract modelling. This work presented an overview of geospatial products and classifies them into a tier-based architecture in which products are based on products of underlying tiers; inspired by [4]. We furthermore formalize the steps of creating tactical terrain models and tactical mission models that abstract the detailed models from lower tiers with the goal of speed for optimization. We present how these abstract models can be used to optimize elements of COAs. Two examples are presented of this approach that are currently in a so-called Concept Development & Experimentation (CD&E) process which means that iteratively the results are presented to stakeholders/users in an experimentation and based on the feedback the concepts are refined.

However promising, many challenges still exist such as how to deal with uncertainties in the tactical terrain model and tactical mission model, how to define good measures of effectiveness (reward functions) and how to prevent the learned results to be biased because of discrepancies between the real world and the modelled world. We want to share some thoughts on how in the future the creation of tactical terrain models and tactical mission models in a (semi-)automated manner can be achieved. The general problem formulation is how to abstract the world in such manner that calculations can be done in a fraction of a time without significantly losing accuracy so that machine learning applications can use these models to create large amounts of data. We currently manually make choices on which details are relevant and which ones are not relevant to arrive at a mission specific model that is reasonably quick and accurate. Several approaches come to mind that can significantly improve on this method.
(1) Current simulators, such as VR forces, are able to simulate any encounter on any terrain. This is usually done for individual units, but can also be done in an aggregated manner [20]. It should be possible to create a meta-model of such simulations for a specific mission when collecting statistics of running this simulation in many variations. Once these statistics are available, running the underlying simulator is not needed any more (e.g., blue inf squad A against red tank squad B on location C results in some average amount of losses of A and B with a certain standard deviation). Such results should be fairly accurate and the machine learning algorithm can sample from the resulting distribution rather than using the simulator itself.

(2) The design of the goal function needs to be performed carefully, as machine learning algorithms search for strategies and loopholes to increase the score of the chosen function [21]. The goal function has to be closely related to the commanders intent. As each mission has its own intent and it has to be interpreted within the current mission context, the translation is not trivial. At TNO we are currently investigating possibilities to analyse the free text of a commanders intent and translate it to computable goal function components that are being briefed back to the commander. The goal is that the military commander will be able to design the goal function in natural language based on the intent.

(3) With the current successes of deep neural networks, we can envision a multitude of applications that may be possible in the future that enable shifts from traditional models to a computation by a neural network. Rather than defining game rules and behaviour, neural networks may learn create realistic simulations, or even only their outcomes, based on provided examples. As data is generally sparse in the military domain, the use of simulators seems promising. Such networks may be able to learn intrinsic relations within mission concepts and create fast and accurate meta-modelling and predictions.

The idea is by using the CD&E approach we will be able to iteratively improve the concepts in order to bring tactical decision support to the war fighter.

REFERENCES

[1] US Army, "Intelligence Preparation of the Battlefield ATP 2-01.3 (Army Techniques Publication)," US Army, Washington DC, 2019.

[2] J. Roskin, "The Role of Terrain and Terrain Analysis on Military Operations in the Late Twentieth to Early Twenty-First Century: A Case Study of Selected IDF Battles," in *Military Geoscience, Bridging History to Current Operations*, Springer, 2020, pp. 145-160.

[3] US Army, "Geospatial Engineering (Army Techniques Publication ATP 3-34.80)," US Army, Washington DC; USA, 2017.

[4] D. L. Visione, "Battlespace Terrain Reasoning and Awareness," in *Esri International USer Conference Proceedings*, 2005.

[5] M. R. Hieb, M. W. Powers, J. M. Pullen and M. Kleiner, "A Geospatial Battle Management Language (geoBML) for Terrain Reasoning (I-110)," in *11th International Command and Control Research and Technology Symposium*, 2006.

[6] Wikipedia, "Battle Management Language," [Online]. Available: https://en.wikipedia.org/wiki/Battle_management_language. [Accessed August 2021].

[7] SISO, "Standard for Coalition Battle Management Language (C-BML) Phase-1 (SISO STD-011-2014)," SISO, 2014.

[8] M. C. I. Eastburg, "Common Ground: Advanced Geospatial Analytics," Army Engineer School,Engineer Professional Bulletin,464 MANSCEN, Fort Leonard Wood,MO,65473, 2011.

[9] M. Breunig and S. Zlatanova, "3D geo-database research: Retrospective and future directions," *Computers & Geosciences,* vol. 37, no. 7, pp. 791-803, 2011.

[10] F. Kuijper, R. v. Son, F. v. Meurs, R. Smelik and K. J. d. Kraker, "Techniques for automatic creation of terrain databases for training and mission preparation," in *IMAGE 2010 Conference*, Scottsdale, AZ, USA, 2010.

[11] R. McAlinden, E. Suma, T. Grechkin and M. Enloe, "Procedural Reconstruction of Simulation Terrain Using Drones," in *Proceedings of Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, Arlington, VA, 2015.

[12] T. Esch, J. Zeidler, D. Palacios-Lopez, M. Marconcini, A. Roth, M. Mönks, B. Leutner, E. Brzoska, A. Metz-Marconcini, S. L. F. Bachofer and S. Dech, "Towards a large-scale 3D modeling of the built environment—joint analysis of TanDEM-X, Sentinel-2 and open street map data," *Remote Sensing,* vol. 12, no. 15, p. 2391, 2020.

[13] K. J. Maroon, "Genetic Algorithms in the Battlespace Terrain Reasoning and Awareness - Battle Command (BTRA-BC) Battle Engine," Naval Postgraduate school, Monterrey - California, 2009.

[14] P. Schwartz, D. O"Neil, M. Bentz, A. Brown and B. Doyle, "AI-enabled wargaming in the Military Decision Making Process," in *SPIE Defense + Commercial Sensing*, on-line only, California US, 2020.

[15] T. Stanzione and K. Johnson, "Geo-Enabled Modelling and Simulation," in *BML conference*, Fairfax, Virginia, 2007.

[16] T. Stanzione and K. Johnson, "Tools for the Creation of Semantic Information for Modeling and Simulation," in *ICCRTS*, Cambridge, MA, 2007.

[17] NATO MSG-156, "STO-TR-MSG-156; Correlated Dynamic Synthetic Environments for Distributed Simulation (MSG-156) (pre-release)," NATO MSG, Paris, 2021.

[18] A. Bitoun, Y. Prudent and A. Hubervic, "Smart Simulation for Decision Support at Headquarters," in *I/ITSEC*, Orlando, 2018.

[19] P. de Heer, N. de Reus, L. Tealdi and P. Kerbusch, "Intelligence augmentation for urban warfare operation planning using deep reinforcement learning," in *SPIE Defense + Commercial Sensing, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, Baltimore, MD, 2019.

[20] VT MAK, "What is Aggregate-Level Simulation Anyway?," *What's Up MÄK,* vol. 17, no. 5, pp. 2-2, 2015.

[21] B. Baker, I. Kanitscheider, T. Markov, Y. Wu, G. Powell, B. McGrew and M. Igor, *Emergent Tool Use From Multi-Agent Autocurricula,* arXiv:1909.07528, 2019.

[22] SISO, "Standard for Command and Control Systems - Simulation Systems Interoperation (SISO STD-019-2020)," SISO, 2020.

[23] A. Grogan, "Creating a spatial analysis model for composite cost surfaces to depict cross country mobility in natural terrain," in *ASPRS/MAPPS 2009 Fall Conference*, San Antonio, Texas, 2011.