

# Computer-Generated Forces Team Behavior within Air Combat Simulations: Concept and Agent Structure

**Fabian Reinisch, Philippe Ruther, Luca Winkler, Michael Strohal, Peter Stütz**  
University of the Bundeswehr  
Munich, Germany  
fabian.reinisch-, philippe.ruther-, luca.winkler-, michael.strohal-,  
peter.stuetz@unibw.de

## ABSTRACT

The growing interest in situationally adaptive, tactically intelligent Computer-Generated Forces (CGF) for military training underscores the necessity for a conceptual framework using interpretable Artificial Intelligence (AI) methods to generate such behavior. Prioritizing system interpretability enables effective debriefing of training exercises. Our concept to generate tactically intelligent CGF team behavior involves three essential processes: Situation Analysis, which brings information into context to generate higher-value information to estimate the opposing team's tactics; Risk-/Opportunity-Assessment, evaluating the information; as well as Decision-Making and Planning, utilizing the assessment for tactical decision-making, which involves selecting and sending actions to the CGF. A critical challenge in simulating these processes is formulating comprehensible risks and opportunities from predicted data in a team-vs-team scenario. Within our concept, first, a situation is assessed qualitatively and semi-quantitatively based on the relations between all entities. Second, a full quantification of the identified risks and opportunities, considering team, group and entity perspectives is performed. Finally, the results are evaluated to formulate interpretable decisions on the CGF tactics. This approach allows us to adjust and simulate the experience levels as well as risk tolerances of the CGF, while also providing explainability by recording the decisions and the relevant parameters they are based on. In addition to describing our conceptual framework, we introduce a model for selecting team-oriented tactics based on an evaluated situation and inferring the necessary maneuvers of individual entities. Here, we focus on creating a solution space for each level of decision-making based on the assigned task and the identified risks and opportunities. The primary challenge resides in minimizing the re-planning of higher order tactics while allowing individual entities to adapt to changing environments.

## ABOUT THE AUTHORS

**Fabian Reinisch** graduated from the Technical University of Munich, Germany in Informatics (B.Sc., M.Sc.). Presently, he holds a position as a research associate and Ph.D. candidate at the Institute of Flight Systems at the University of the Bundeswehr Munich, Germany.

**Philippe Ruther** received a B.Sc. and an M.Sc. in Aerospace Engineering from University of Stuttgart, Germany. Currently, he works as a research associate and Ph.D. candidate at the Institute of Flight Systems at University of the Bundeswehr Munich, Germany.

**Luca Winkler** holds a B.Sc and M.Sc in Physics from the University of Augsburg. Currently, he is a research associate and Ph.D. candidate at the Institute of Flight Systems at the University of the Bundeswehr Munich.

**Michael Strohal** is a graduate of the University of the Bundeswehr Munich, Germany, with a doctorate in Aerospace Engineering. He has 10 years of experience as head of department in the field of weapon system simulation and more than 30 years of experience in aerospace vehicle simulation including helicopters, missiles and airplanes.

**Peter Stütz** is the professor of Aeronautical Engineering at the University of the Bundeswehr Munich and co-director of the Institute of Flight Systems since 2008. He is a graduate of TU Munich and received his doctorate in cockpit human factors from the University of the Bundeswehr Munich. His scientific focus lies on various aspects of airborne, sensor-based environmental perception as well as aircraft platforms, systems and scenario simulation techniques.

# Computer-Generated Forces Team Behavior within Air Combat Simulations: Concept and Agent Structure

**Fabian Reinisch, Philippe Ruther, Luca Winkler, Michael Strohal, Peter Stütz**  
**University of the Bundeswehr**  
**Munich, Germany**  
**fabian.reinisch-, philippe.ruther-, luca.winkler-, michael.strohal-,**  
**peter.stuetz@unibw.de**

## INTRODUCTION

In the domain of military Modeling and Simulation (M&S), there is a growing demand for situationally adaptive, tactically intelligent Computer-Generated-Forces (CGF) that serve as opponents (red forces) in training simulations and tactical evaluations (Roessingh et al., 2011; Toubman et al., 2016). The primary goal of these forces is to enhance realism and immerse trainees (blue forces) in a realistic training scenario. Most military scenarios involve two competing teams whose actions directly influence each other, as both strive to achieve their overarching mission goal. To effectively oppose trainees or validate blue tactics in tactical evaluations, red CGF must exhibit adaptive and tactically intelligent team behavior (van Oijen et al., 2019).

To enable a situationally adaptive and tactically intelligent team of red CGF, a behavior model must be designed which meets various requirements. First, the behavior model must be able to coordinate a team of CGF, taking into consideration the team's risks and opportunities. The basic requirement for this is to understand the tactical capabilities of both teams, red and blue, potentially incorporating intel data in this process. The model must operate on multiple levels of decision-making to maintain tactically intelligent behavior, regardless of team size. An overarching team tactic must be developed on an informed foundation, with necessary tasks delegated to individual groups. The model needs to manage these groups independently and develop their tactics in accordance with their individual tasks. This process should be designed iteratively to accommodate any number of sub-groups. Additionally, the state evaluation of each subdivision must also occur independently, in order to avoid unnecessary re-planning of higher-order tactics. Third, the use in training simulators or tactic validation demands the interpretability of the red CGF team behavior. During simulation debriefings, the system must provide the information on which decisions were based to make the behavior comprehensible and trace the decision-making process of the CGF team. Lastly, the model must be capable of anticipating the future course of a situation, a requirement we consider essential for creating tactically intelligent behavior.

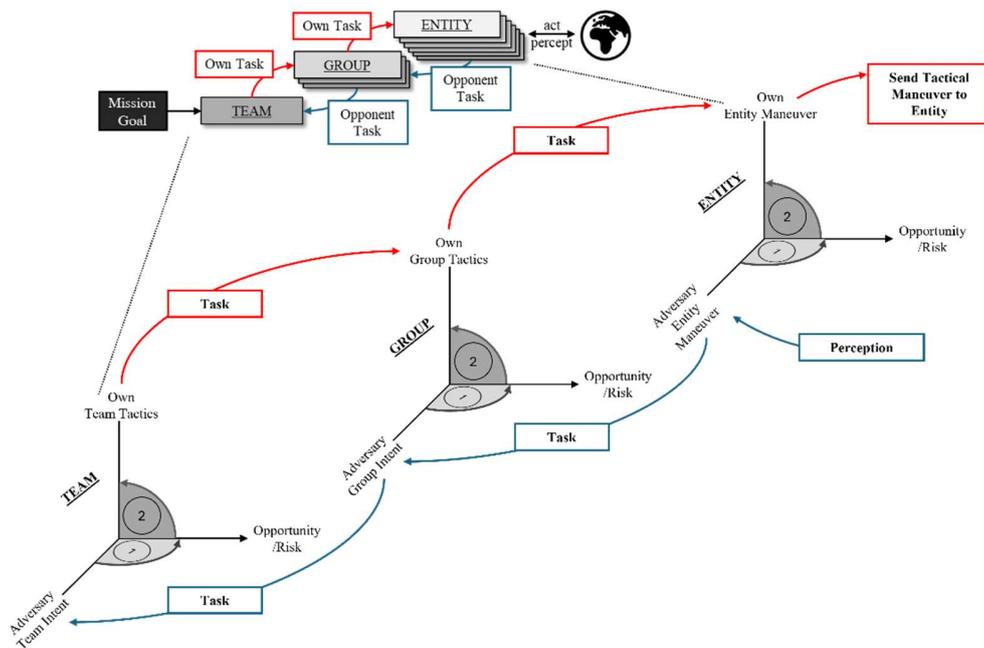
There are already several concepts and modeling approaches to generate adaptive and intelligent behavior of CGF. However, existing papers mostly focus on a small subset of engagements such as target detection and tracking (V. Chandrakanth et al., 2022) or engagement support (Joao P. A. Dantas et al., 2021), however research taking the full workflow of team behavior modeling into consideration is still a mostly unexplored field. When generating CGF behavior, most research assumes that complete information regarding the adversary side is available (Dongyuan Hu et al., 2021; Jing Xianyong et al., 2022). However, in real-world scenarios (and therefore in training simulations), multiple factors, such as sensor limitations, can lead to incorrect information about the opponents. Some studies have tried incorporating these limitations and conducted research on sensor errors (Jung Ho Bae et al., 2023), however, while they are considering some sensor limitations, they have not considered effects such as sensor resolution. When it comes to behavior generation, one study (R.E Smith et al., 2000) concentrated on modeling tactics with a learning classifier system, while another (Zhixiao Sun et al., 2021) used a multi-agent hierarchical policy gradient in order to discover and generate air combat tactics via self-play. More recently, researchers have been actively applying Reinforcement Learning (RL) techniques to generate air combat maneuvering (Dongyuan Hu et al., 2021; Xinwei Wang et al., 2024). Although RL-based simulations have shown promising results, their practical implementation will require extensive verification, since interpreting the maneuvering principles and situational judgment learned by RL-based systems may present challenges for human operators and limit the comprehension of the learned decisions.

Our approach to fulfilling the mentioned requirements and achieving situationally adaptive and tactically intelligent team behavior of CGF involves a centralized agent function based on a multi-agent system utilizing expert knowledge

(Russel & Norvig, 2021). By employing this approach combined with the usage of explainable AI methods, we are able to model centrally controlled adaptive and intelligent team behavior, while also ensuring interpretability of the behavior for users.

**CONCEPT**

To model situation-adaptive CGF that behave tactically intelligent as a team, we are guided by a hierarchical structure that emerges within a team operating in a competitive environment against another team. In this section, we first describe our general concept, and then illustrate it with a domain-specific example later in this paper. As shown in the upper part of Figure 1, a team can be divided into multiple levels. Initially, a team leader has a mission objective, which must be achieved through the tactically intelligent use of available resources (team members). To achieve the mission goal, an initial tactic is pursued at the team level. A team can consist of several groups, which are assigned specific tasks in accordance with the overall team tactic. Each group, in turn, can consist of several entities that collaborate to accomplish the group task. Each entity is given a specific task, which it executes through a series of defined tactical maneuvers. After establishing initial tactics at the team level, it is necessary to assess its viability. In a competitive environment, the two team behaviors are directly correlated (Myerson, 2019). Therefore, it is essential to identify opposing entities and their intentions. Each opposing entity performs a specific task and interacts with other entities within a group, which, in turn, collaborates with other opposing groups to fulfill the team tactic derived from the opposing mission objective. A clear parallelism between the hierarchical structures of both teams can be observed.



**Figure 1. Schematic representation of the team behavior model**

The depicted hierarchical levels, from the entity level to the entire team, can be decomposed into three essential components. Initially, it is necessary to identify the adversary’s actions and intentions in order to recognize the risks and opportunities that arise at each respective level (as indicated by the first process step between Intent and Risk axis in Figure 1). Based on this evaluation, decisions for subsequent actions are made and further planning is conducted. In the Situation Analysis (SA) process, all perceptions of the own entities are collected and used to infer the tactical maneuvers of the opposing entities. By identifying interacting opposing entities and their respective tactical maneuvers, it is possible to determine the intentions of the identified groups. Observing these groups over a longer period allows the interpretation of the opposing team's tactics. A Risk-/Opportunity Assessment (ROA) evaluates the risks and opportunities for the own team, group, or entity. This assessment represents a measurable metric. The upper limit of the axis represents the maximum risk that an entity, group, or team is allowed to take. The lower limit describes the minimal opportunities at the respective level. For instance, if a group exceeds the predefined maximum risk, this

necessitates re-planning. Conversely, if the behavior model identifies no opportunities for the group to fulfill its task, independent of the chosen group tactic, re-planning is also required. The third process component of the behavior model describes the Decision-Making and Planning (DMP), where the most appropriate tactic or tactical maneuver is chosen for the current situation. From the selected team tactic, tasks for the groups are derived, which must be fulfilled to achieve the team tactic. At the group level, based on the identified risks and opportunities, a group tactic is decided to accomplish the team tactic. This results in specific tasks for the entities within the respective group, for which sequential tactical maneuvers must be planned to execute the group tactic. A more detailed explanation is provided later in this paper, where this concept is applied to a specific domain.

These three processes in each hierarchical level allow the model to make proactive conclusions and assessments about its own behavior. Moreover, this concept is essential for providing tactical responses to maneuvers executed by the opposing team, ensuring effective adaptation to dynamic situations. However, this requires a knowledge base where potential tactics and maneuvers of the opponent can be stored, as well as tactics and tactical maneuvers of the own team. The storage of past identified maneuvers also plays a crucial role, as the identification of the opposing team's tactics can only be achieved through observations over an extended period. By continuously evaluating the current tactic, it is possible to reintegrate an entity that has been manually overwritten by the user into the team tactic.

## AGENT STRUCTURE

After identifying the three essential processes for tactically intelligent behavior in team-based conflicts, the next step is to establish a framework to implement this behavior model within a simulated environment. A utility-based agent is highly suitable for implementing centralized CGF team behavior due to its capability to make decisions and perform actions with the overarching goal of achieving specific objectives. It consists of an agent function and an agent architecture that interact with their environment (Dilger). The agent function perceives the environment and issues specific actions to the agent architecture, which contains the CGF dynamic models to execute the received actions. A utility-based agent operates based on a utility function that evaluates multiple action alternatives and selects the one anticipated to yield the highest utility or overall success (Dilger; Russel & Norvig, 2021). Furthermore, an agent structure is well-suited to operate in uncertain, highly complex, and dynamic environments. To achieve this, the agent requires a comprehensive knowledge base that includes both static information available in advance (such as information provided by an ATO), as well as dynamic information generated during the mission (Dilger; J J M Roessingh & GG Verhaaf; J N Roux & J H Van Vuuren, 2007). Uncertainties in static information do not directly impact the agent, as this information is assumed to be accurate due to the absence of alternative data sources. In contrast, uncertainties in dynamic information, such as sensor data from each CGF entity, must be considered in further evaluation. Information that cannot be precisely defined is referred to as fuzzy information (Moral et al., 2021). Therefore, the methods applied in the agent function are chosen to be capable of processing fuzzy information.

A mentioned requirement of situational adaptive behavior is predicting future developments. This requires incorporating past indicators that have shaped the current situation into the agent's knowledge base. To achieve this, the knowledge base can be structured with two distinct data layers. The Situation Representation (SR) stores data from individual frames, capturing the most current information. Past SRs are archived in the Agent's Memory (AM), enabling the extrapolation of data to predict how a situation might evolve (Endsley, 2000). Integrating the AM into the agent architecture can be achieved by using a database to store all generated SRs.

We do not believe that constant anticipation of new future courses leads to tactically intelligent behavior. Instead, the last anticipated situation, which caused the current tactics to be selected, is continually evaluated (Floyd et al.; Floyd et al., 2017; Klenk et al., 2013; Sado et al., 2023). When the current tactics are no longer deemed effective, new tactics must be generated, necessitating a fresh anticipation of future developments. This requires the collaboration of all three identified subsystems, as schematically illustrated in Figure 2. Our concept is divided into two distinct phases: the agent function operates in a loop executed in each frame (solid arrows) and a loop that is only invoked upon request from DMP (dashed arrows). The process begins with the integration of perceived data from the synthetic environment into the SA. The result is stored as the current SR (indicated by the yellow portion of the AM). This, along with historical information from the AM, is forwarded to the ROA. The evaluated data is then passed to DMP. At this point, the final decision is made whether re-planning is required. If DMP identifies the need for a new maneuver or tactic, a collection of tactically intelligent alternatives for the own team is sent to the SA (Own Maneuver Possibilities: small cyan disk). The SA conducts a precise trajectory prediction based on the proposed maneuvers, assuming the opposing entity maintains its current intent and speed (Predicted Trajectories: small yellow disk). The results are then evaluated

by the ROA, which identifies the opportunities, risks, and the critical point of the maneuver alternatives (Critical Point Assessment: small light red disk). The so-called critical point indicates the last moment in which a new decision will have to be made, as an entity or group will no longer be able to execute its assigned task while maintaining an acceptable level of risk. The outcome is forwarded to the DMP process. At this point, a choice is made from the proposed maneuver alternatives. These are sent to the individual entities for execution, while simultaneously recording all information generated in this process from the three subsystems in the SR. Finally, the entire SR is added to the AM, becoming part of the historical data for future simulation steps. This modular structure offers multiple benefits. The collection of SRs enhances the accuracy of the inferred knowledge generated by SA of the current situation and improves the precision of trajectory extrapolations necessary for evaluating proposed plans. In the following chapter, the described procedure will be further explained by applying it to a specific domain.

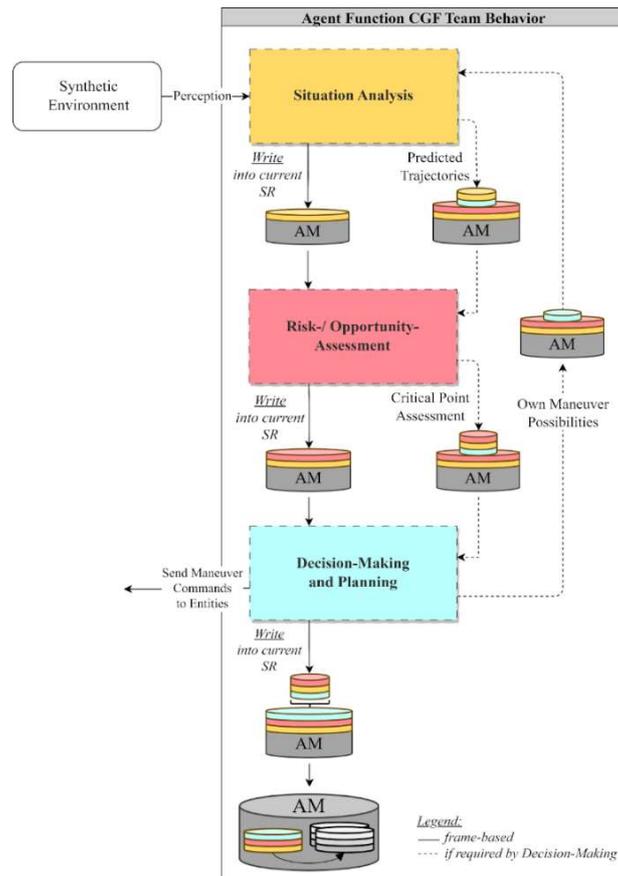


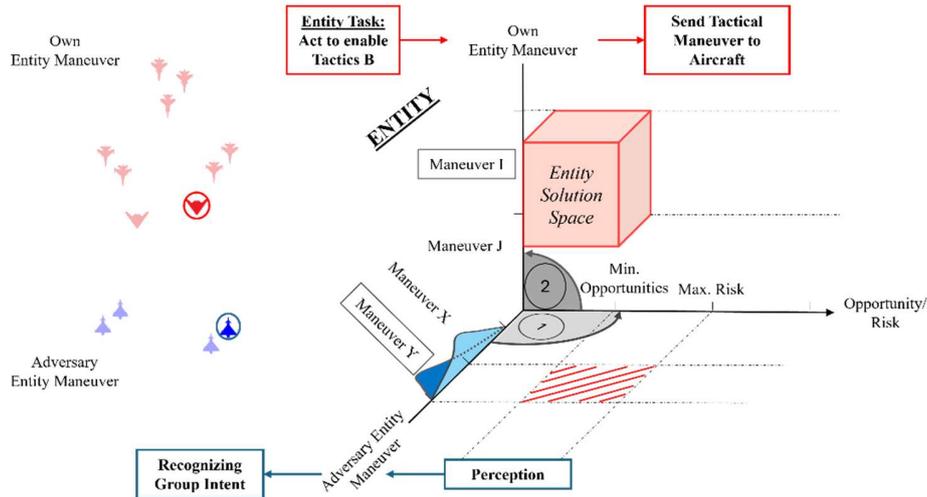
Figure 2. Interaction between the three subsystems to anticipate future events

### APPLYING THE CONCEPT TO THE BVR-AIR COMBAT DOMAIN

After presenting the concept for mapping an agent function to model CGF team behavior in the previous section, in the following section we will apply this concept to a Beyond Visual Range (BVR) air combat scenario. Figure 3 illustrates a scenario on the left side, as it might occur in a military training simulation. As we aim to model the behavior of simulated red forces, we will refer to their behavior as the “own” behavior and that of the blue forces as “adversary” behavior. As mentioned in the description of Figure 1, the initial state of the CGF behavior model involves selecting a team tactic based on the underlying mission objective. This results in initial group tasks, which must be executed at the group level to implement the team tactic. From these group tasks, specific tasks are defined for each entity within the group, which influence the adversary (or environment) through the execution of specific tactical maneuvers. Since situational adaptability is a central aspect of behavior modeling, it must be assessed at each level whether the task set by the higher level can be executed or if re-planning at the higher level is necessary. This principle is detailed in the following example of a red CGF fighter team in BVR air combat.

**Entity Level**

Once initial tactics within the team have been established and the resulting tasks defined, a continuous perception of the adversary (blue) forces takes place. This perception is achieved through sensor information collected by each individual entity and combined in a centralized air picture. By perceiving the adversary entities and utilizing the available static information (see agent structure) on known adversary maneuvers, the currently executed tactical maneuvers can be identified in relation to the own entities (see Maneuver Y in Figure 3). Uncertainties play a significant role, as the actions of adversary entities can never be definitively identified (see the probability distribution of adversary entity maneuvers axis). Therefore, each adversarial maneuver is represented as a range. This comparison is performed for every pairwise configuration between adversary and own entity, making it computationally intensive.



**Figure 3. Illustration of an entity solution space instance**

Once the corresponding tactical maneuver of an adversary entity is identified, a ROA is conducted for the respective own (red) entity (see Figure 3, process step 1). By defining a maximum risk level as part of the static information, the acceptable risk for the entity is limited. This allows for individual risk affinity of each entity. For instance, an Unmanned Aerial Vehicle (UAV) entity can have a higher maximum risk tolerance compared to a manned entity. Additionally, a lower bound on the axis describes the minimal opportunities or conditions that must be met for the entity to fulfill its assigned task while maintaining an acceptable level of risk. Thus, the entity's current state falls within the shaded area of the two described axes. If the entity's state is outside this defined area—due to identifying a different adversary maneuver, exceeding the maximum risk level, or the opportunity assessment not providing a feasible way to fulfill the entity's task—this signals the need for re-planning at the entity level through the DMP process (third axis). A decision for a new tactical maneuver must be made by estimating its impact on the future course, corresponding to the outer loop depicted in Figure 2. The DMP process sends possible pre-selected tactical maneuvers to the SA, which predicts the adversary's intentions under these conditions. The ROA evaluates the results, which then informs the DMP process to decide on a tactical maneuver. If this process reveals that none of the own tactical maneuvers, known through the knowledge base, enable the entity to accomplish the task commanded by the group level, re-planning at the group level is required. If this is not the case, the DMP process plans further maneuvers considering the overarching group task. This tactical maneuver is executed in relation to an adversary entity but may vary in precise execution (e.g., speed, altitude, etc.). Due to these uncertainties, each tactical maneuver is represented as a range. Subsequently, the tactical maneuver is issued to the entity. Adding the third axis to the shaded area results in a three-dimensional state space (Entity Solution Space). A state outside this space necessitates re-planning.

**Group Level**

The principle described at the entity level also extends to the group and team levels. By identifying interacting adversary entities, it is possible to infer group structures (see Figure 4, blue side, e.g., Sub-Tactics C). Considering the previously identified tactical maneuvers of each entity, the tactical intent of the identified group can be deduced. This assessment is used to conduct a ROA for each own (red) group. If the current risk level affecting the group is outside the shaded area, re-planning within the DMP process at the group level is necessary (see Figure 4, example Tactic B

Version X). A decision for a new group tactic must be made by estimating the impacts and future developments of the scenario, corresponding to the outer loop described in Figure 2. In this process, the DMP sends tactical maneuvers of individual entities to the SA based on potential group tactics, which predicts the adversary's intentions under the perceived conditions. The ROA evaluates the result, upon which the DMP process decides on a tactical maneuver. If this process determines that none of the known tactics from the knowledge base enable the group to accomplish the group task commanded by the team level, re-planning at the team level is required.

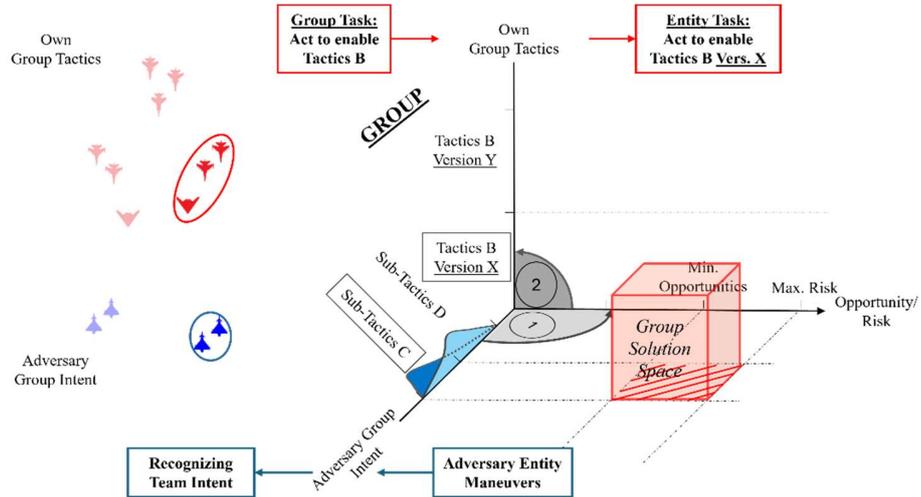


Figure 4. Illustration of a group solution space instance

**Team Level**

From the identified group intentions, the overall adversary tactic can be inferred (see Figure 5, Adversary Team Intent, Example Tactics A). Based on the recognized tactic, a ROA is conducted for the entire team. The maximum risk for the team is predetermined in the Air Tasking Order (ATO) as part of the static information. The DMP process decides on a team tactic considering the mission objective (see Figure 5, Tactics Band commands individual group tasks (see Figure 4, Tactics B, Version X or Y) to the respective subgroups. If it is determined at the team level that the mission risk is exceeded or no opportunity of achieving the mission goal could be identified, the offensive red team aborts its attack. This corresponds to the successful achievement of the mission goal by the defensive blue team (trainee pilots).

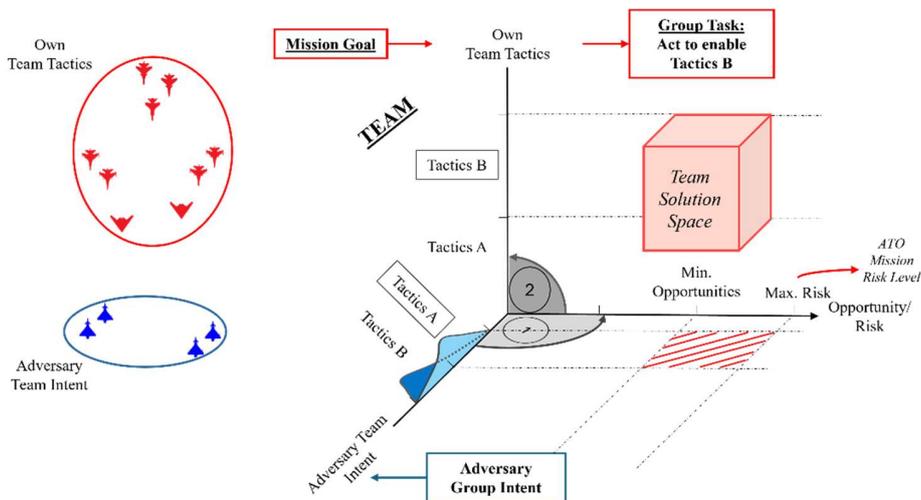


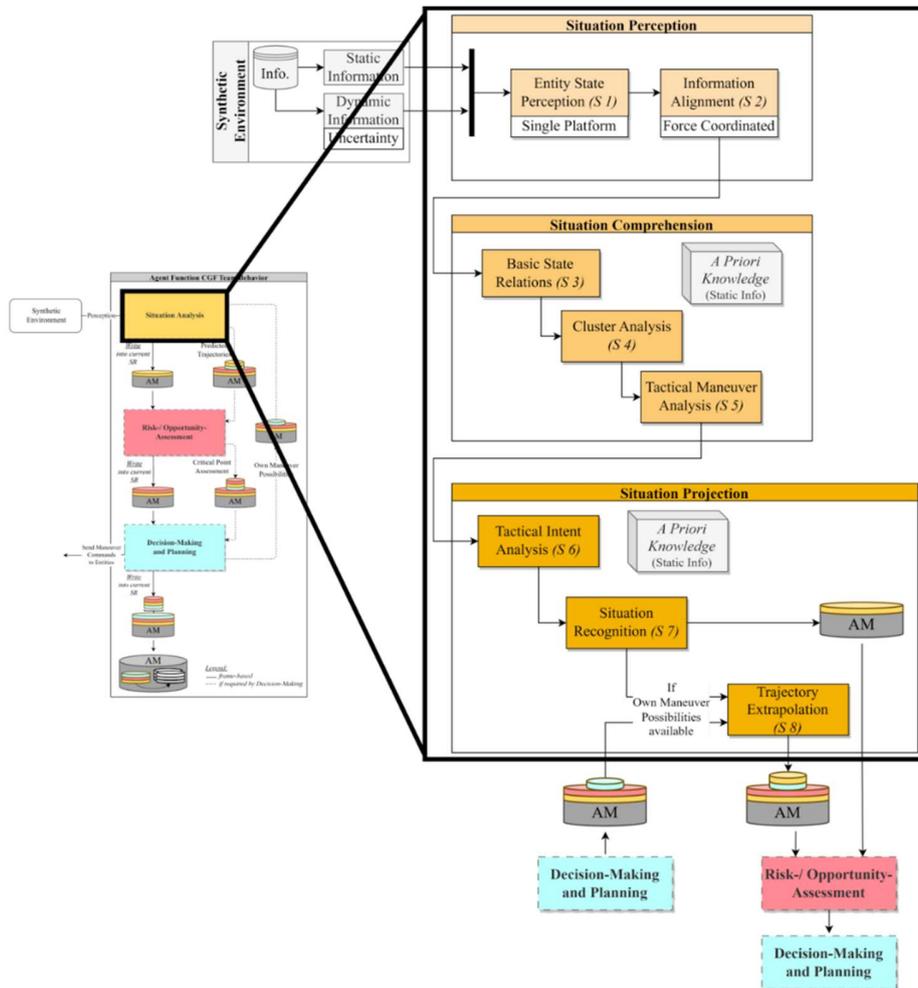
Figure 5. Illustration of the team solution space instance

**PROCESS IN DETAIL**

To generate situationally adaptive, tactically intelligent CGF team behavior, it is essential to model the three major processes within the utility-based agent. The following section will provide an explanation of the specific processes.

**Situation Analysis**

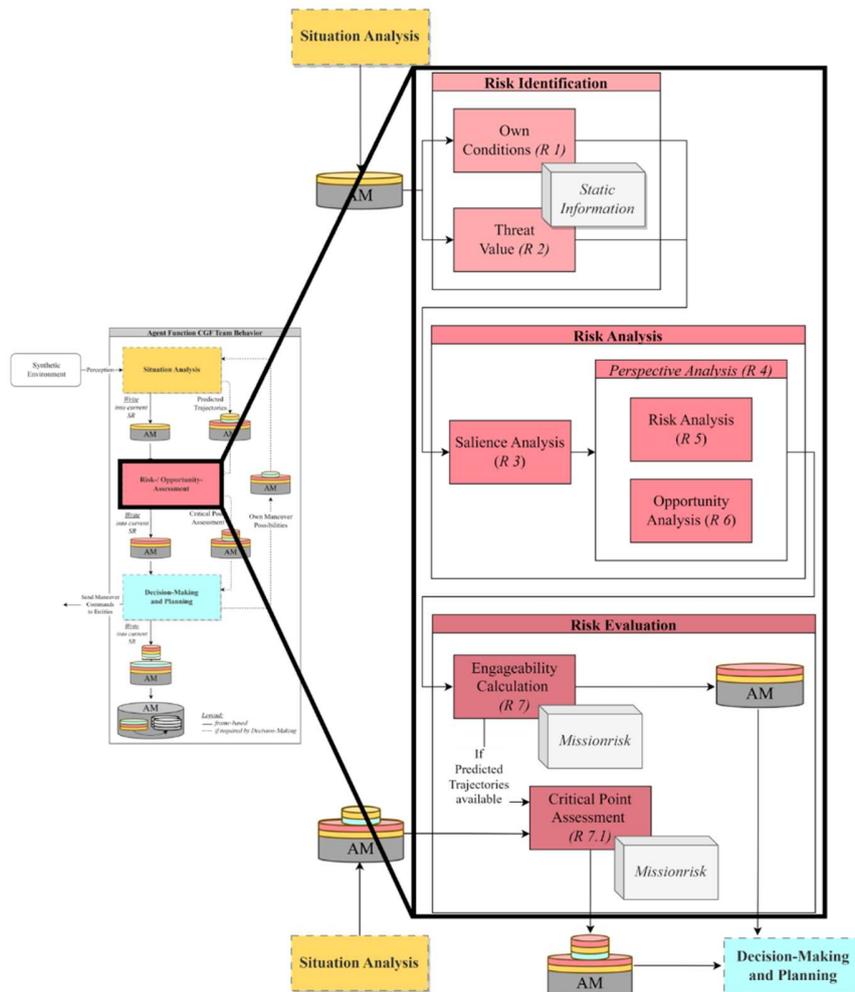
The SA process, visualized in Figure 6, examines a situation and writes its output into the Situation Representation (SR). A situation is defined as a specific combination of observations or states at a given time. This information describes entities existing in a situational environment (e.g., detected radar tracks) as well as their currently executed maneuvers (actions). The information written to the SR depends on the first level within the SA: Situation Perception, through which the agent becomes aware of its surroundings. For example, if an adversary entity is not detected by a red CGF radar, this information is not written to the SR. Here, the above mentioned pairwise comparison of several entities is performed and information is combined. While perception involves awareness of information, Situation Comprehension involves linking perceptions to obtain higher-level information. A utility-based agent being capable of comprehending a situation can derive operationally relevant information through logical inference. For example, identifying enemy entities which collaborate as a group (see Cluster Analysis in Figure 6), the tactical maneuvers of these groups and their tactical intent can be classified. A group can consist of any number of entities. The ability to anticipate intentions and future flight paths based on current events is known as Situation Projection (Jean, 2001), enabling timely ROA and the creation of a successful tactical plan through subsequent processes.



**Figure 6. Illustration of the Situation Analysis composed of the three parts Situation Perception, Situation Comprehension and Situation Projection**

**Risk-/Opportunity Assessment**

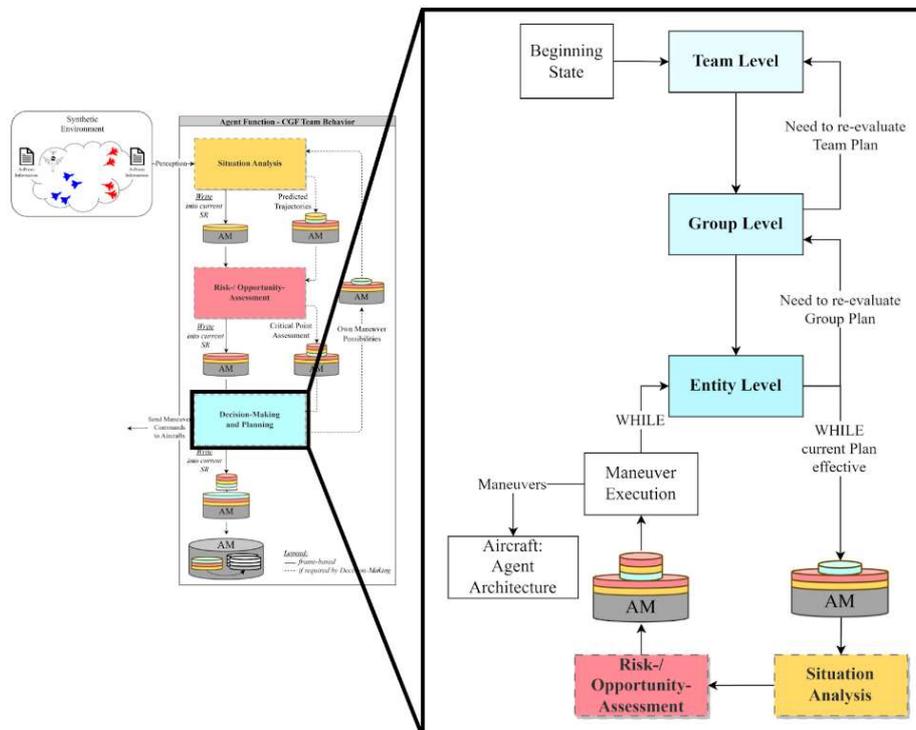
The ROA process calculates the quantifications of the risks and opportunities faced by the own team in any given situation. We structured it to consist of three steps: risk identification, risk analysis, and risk evaluation, as illustrated in Figure 7. The process of risk identification involves systematically detecting and classifying potential parameters or events that could impact the goals of the team, a group or an individual entity. Initially, both teams are evaluated independently. Then, a semi-quantitative investigation of the possible opportunities and risks between the entities is conducted, based on their relative states to each other. During the risk analysis, a comprehensive examination of all identified risks and opportunities is carried out, considering their extent and likelihood of occurrence. In this process, the likelihood is understood as a quantitative measure of the probability that a specific event will occur within a defined period. This probability depends on several parameters, which are incorporated into the probability calculation with varying weights. Considered parameters include both the inherent uncertainties of the situation and the results of the risk identification. The extent is represented in levels, indicating the potential drawback or benefit to the relevant level of decision making if an identified event occurs. The combination of extent and likelihood for each identified event results in a matrix that can be evaluated in a subsequent step. This procedure must be performed for each level of decision making, with the results of individual entities informing the calculations at group and team level. During the risk evaluation, the previous results are compared against predefined thresholds. The comparison aims to cumulate the results of the previously established matrices and provide a meaningful and comprehensible assessment of the situation.



**Figure 7. Illustration of the Risks-/Opportunity Assessment composed of the three parts Risk Identification, Risk Analysis and Risk Evaluation.**

**Decision-Making and Planning**

The DMP process is the core of the centralized behavior model, as it generates behavior affecting the entire CGF team, as well as decisions made at the level of CGF groups or individuals using the given solution spaces from the previous chapter. Additionally, we distinguish between two temporal phases: (1) Task-based, where the relevant decision-making switches to the next task after the current one was completed or aborted, and (2) Time-based, where the current plan is continuously evaluated throughout the scenario. As illustrated in Figure 1, the DMP operates on three levels: (1) Team Level, (2) Group Level, and (3) Entity Level, which are being visualized in Figure 8 together with their dependencies. The Team Level marks the beginning of the DMP process. Utilizing the adversary team intent, primarily collective tasks such as selecting and evaluating the overarching team tactics or re-grouping the own forces in order to switch tactics, are performed on this level. The Group Level is reached when CGF are assigned into different Groups and follow the previously selected tactics. Here, using the recognized opposing group intent, it is first decided whether the current tactical plan for the different groups is still effective or needs to be changed. If no changes are necessary, the tasks for individual CGF are determined in the Entity Level based on the recognized target entity's maneuver and the identified opportunities and risks. Finally, the maneuvers that the CGF should perform in real-time are selected. This process involves three steps. First, potential tactically intelligent maneuvers are selected. Then, using SA and ROA, the possible resulting risks and opportunities of these maneuvers are assessed. Based on this assessment, the maneuvers with the best risk/opportunity ratio are selected and sent to the CGF (Agent Architecture).



**Figure 8. Illustration of the Decision-Making and Planning Process composed of Team-, Group-, and Entity Levels as well as the Maneuver Execution.**

**CONCLUSION AND FUTURE WORK**

In conclusion, we introduced a novel approach for modeling CGF team behavior established by three distinct processes: Situation Analysis, Risk-/Opportunity Assessment, as well as Decision-Making and Planning. The first stage involves assessing the environmental context and the states of all entities, estimating the current intent of the opposing forces. Following, potential risks and opportunities are identified, and the impact of various risk factors evaluated, providing a quantification of the maximum tolerable risk and a minimum required opportunity to execute the assigned task. Finally, the Decision-Making and Planning stage involves invoking methods capable of selecting and executing predefined team tactics, by assigning tasks to the necessary team subdivisions and generating

maneuvers based on the preceding analyses. Furthermore, defining acceptable solution spaces for each team-based subdivision allows us to minimize the re-planning of higher order tactics while allowing individual entities to adapt to changing environments. We have illustrated the applicability of this method in detail using a BVR scenario. Furthermore, we introduced an approach to integrate this model into an agent function, showcasing its potential for enhancing coordinated team behavior in competitive environments. By implementing behavior trees at the specified levels, we have already developed a demonstration framework where CGFs can compete against pilots.

Going forward, we aim to focus on combining behavior trees with symbolic and sub-symbolic methods for the identified subprocesses. To meet the requirement for explainability, the selected methods must be evaluated in the context of verification and validation. Finally, we plan to conduct trials with subject matter experts serving as reviewers to validate the generated behavior of the CGF.

## REFERENCES

- Dilger, W. *Einführung in die Künstliche Intelligenz*.
- Dongyuan Hu, Rennong Yang, Jialiang Zuo, Ze Zhang, Jun Wu, & Ying Wang (2021). Application of Deep Reinforcement Learning in Maneuver Planning of Beyond-Visual-Range Air Combat. *IEEE Access*, 9, 32282–32297. <https://doi.org/10.1109/ACCESS.2021.3060426>
- Endsley, M. R. (2000). Situation Models: An Avenue to the Modeling of Mental Models. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 44(1), 61–64. <https://doi.org/10.1177/154193120004400117>
- Floyd, M. W., Karneeb, J., Moore, P., & Aha, D. W. (2017). A Goal Reasoning Agent for Controlling UAVs in Beyond-Visual-Range Air Combat. *IJCAI*(Vol. 17), 4714–4721.
- J J M Roessingh, & GG Verhaaf. Training Effectiveness of Embedded Training in a (Multi-) Fighter Environment - A Discussion Paper. Advance online publication. <https://doi.org/10.13140/RG.2.2.31697.53606>
- J N Roux, & J H Van Vuuren (2007). Threat evaluation and weapon assignment decision support: A review of the state of the art. *ORiON*, 23(2), 151–187. <https://www.ajol.info/index.php/orion/article/view/34267>
- Jean, R. (Ed.) (2001). *From data fusion to situation analysis*.: Vol. 2.
- Jing Xianyong, Manyi Hou, Gaolong Wu, Zongcheng Ma, & Zhongxiang Tao (2022). Research on Maneuvering Decision Algorithm Based on Improved Deep Deterministic Policy Gradient. *IEEE Access*, 10, 92426–92445. <https://doi.org/10.1109/ACCESS.2022.3202918>
- Joao P. A. Dantas, Andre N. Costa, Diego Geraldo, Marcos R. O. A. Máximo, & Takashi Yoneyama (2021). Engagement Decision Support for Beyond Visual Range Air Combat. *CoRR*, abs/2111.03059.
- Jung Ho Bae, Hoseong Jung, Seogbong Kim, Sungho Kim, & Yong-Duk Kim (2023). Deep Reinforcement Learning-Based Air-to-Air Combat Maneuver Generation in a Realistic Environment. *IEEE Access*, 11, 26427–26440. <https://doi.org/10.1109/ACCESS.2023.3257849>
- Klenk, M., Molineaux, M., & Aha, D. W. (2013). Goal-driven autonomy for responding to unexpected events in strategy simulations. *Computational Intelligence*, 29(2), 187–206. <https://doi.org/10.1111/j.1467-8640.2012.00445.x>
- Moral, Castiello, Magdalena, & Mencar. (2021). *Explainable fuzzy systems*. <https://link.springer.com/content/pdf/10.1007/978-3-030-71098-9.pdf>
- Myerson, R. B. (2019). *Game theory: Analysis of conflict* (First Harvard University Press paperback edition). <https://permalink.obvsg.at/>
- R.E Smith, B.A Dike, R.K Mehra, B Ravichandran, & A El-Fallah (2000). Classifier systems in combat: two-sided learning of maneuvers for advanced fighter aircraft. *Computer Methods in Applied Mechanics and Engineering*, 186(2), 421–437. [https://doi.org/10.1016/S0045-7825\(99\)00395-3](https://doi.org/10.1016/S0045-7825(99)00395-3)
- Roessingh, Rijken, Merk, Meiland, Huibers, Lue, & Montijn (2011). Modelling CGFs for tactical air-to-air combat training motivation-based behaviour and machine learning in a common architecture. <https://reports.nlr.nl/bitstream/10921/465/1/tp-2011-540.pdf>
- Russel, S., & Norvig, P. (2021). *Artificial Intelligence: A Modern Approach* (4th ed.). Pearson.
- Sado, F., Loo, C. K., Liew, W. S., Kerzel, M., & Wermtter, S. (2023). Explainable Goal-driven Agents and Robots - A Comprehensive Review. *ACM Computing Surveys*, 55(10), 1–41. <https://doi.org/10.1145/3564240>
- Toubman, A., Roessingh, J. J., van Oijen, J., Lovlid, R. A., Ming Hou, Meyer, C., Luotsinen, L., Rijken, R., Harris, J., & Turcanik, M. (2016). Modeling behavior of Computer Generated Forces with Machine

- Learning Techniques, the NATO Task Group approach. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE. <https://doi.org/10.1109/smc.2016.7844517>
- V. Chandrakanth, V. S. N. Murthy, & Sumohana S. Channappayya (2022). UAV-based autonomous detection and tracking of beyond visual range (BVR) non-stationary targets using deep learning. *J. Real Time Image Process.*, *19*(2), 345–361. <https://doi.org/10.1007/s11554-021-01185-w>
- van Oijen, J., Toubman, A., & Poppinga (2019). Effective behaviour modelling for computer generated forces. *Interservice/Industry Training, Simulation and Education Conference*. [https://www.researchgate.net/profile/joost-van-oijen/publication/337797509\\_effective\\_behaviour\\_modelling\\_for\\_computer\\_generated\\_forces/links/5dea70f592851c83646862bd/effective-behaviour-modelling-for-computer-generated-forces.pdf](https://www.researchgate.net/profile/joost-van-oijen/publication/337797509_effective_behaviour_modelling_for_computer_generated_forces/links/5dea70f592851c83646862bd/effective-behaviour-modelling-for-computer-generated-forces.pdf)
- Xinwei Wang, Yihui Wang, Xichao Su, Lei Wang, Chen Lu, Haijun Peng, & Jie Liu (2024). Deep reinforcement learning-based air combat maneuver decision-making: literature review, implementation tutorial and future direction. *Artif. Intell. Rev.*, *57*(1), 1. <https://doi.org/10.1007/S10462-023-10620-2>
- Zhixiao Sun, Haiyin Piao, Zhen Yang, Yiyang Zhao, Guang Zhan, Deyun Zhou, Guanglei Meng, Hechang Chen, Xing Chen, Bohao Qu, & Yuanjie Lu (2021). Multi-agent hierarchical policy gradient for Air Combat Tactics emergence via self-play. *Eng. Appl. Artif. Intell.*, *98*, 104112. <https://doi.org/10.1016/j.engappai.2020.104112>.