Artificial Agents Inspired by Human Motivation Psychology for Teamwork in Hazardous Environments

Anupama Arukgoda[1] , **Erandi Lakshika**[1] , **Michael Barlow**[1] and **Kasun Gunawardana**[2]
[1]University of New South Wales, Canberra, Australia
[2]University of Colombo School of Computing, Colombo, Sri Lanka
a.arukgoda@adfa.edu.au, e.henekankanamge@adfa.edu.au, m.barlow@adfa.edu.au, kgg@ucsc.cmb.ac.lk

Abstract

Multi-agent literature explores personifying artificial agents with personality, emotions or cognitive biases to produce "typical", believable agents. In this study, we demonstrate the potential of endowing artificial agents with a motivation, using human implicit motivation psychology theory that introduces 3 motive profiles - power, achievement and affiliation, to create diverse, risk-aware agents. We first devise a framework to model these motivated agents (or agents with any inherent behavior), that can activate different strategies depending on the circumstances. We conduct experiments on a firefighting task domain, evaluate how motivated teams perform, and draw conclusions on appropriate team compositions to be deployed in environments with different risk levels. We find that motivational diversity within teams is beneficial in dynamic collaborative environments, especially as the task risk level increases. Furthermore, we observed that the best team composition in terms of the performance metrics used to evaluate teams, does not remain the same as the collaboration level required to achieve goals changes. These results have implications for future designs of risk-aware autonomous teams and Human-AI teams, as they highlight the prospects of creating better artificial teammates and performance gains that could be achieved through anthropomorphized motivated agents.

1 Introduction

With the continuous advancements in Artificial Intelligence (AI), autonomous agents have become a ubiquitous part of the workforce, rising up from being mere subordinate tools to equal fellow teammates of humans [Schelble *et al.*, 2020; O'Neill *et al.*, 2020]. Recent advancements in research have been invested to identify what makes an autonomous agent, not just a viable teammate, but also a good teammate. This compendium of research has shown that a good autonomous teammate engenders factors beyond accuracy, such as trust [McNeese *et al.*, 2021], team cognition [Cooke *et al.*, 2013], situation-awareness [Gorman *et al.*, 2006] and most importantly, a sense of predictability [Bansal *et al.*, 2021; Siu *et al.*, 2021]. Siu *et al.* [2021] empirically demonstrate that humans prefer predictable agents with rule based systems, over learned agents even when both models have the same accuracy. Therefore, in the pursuit of building predictable autonomous agents, we explore the potential of human motivation psychology theory to personify autonomous agents to create effective teams.

Three Needs Theory, a popular motivation psychology theory defines three implicit motives that drive human decisions [McClelland and Mac Clelland, 1961]. These three motive profiles are affiliation, achievement and power. While humans possess these three motives in varying forms, they are involuntarily guided by one dominant motive when making decisions thus, leading to behavioral diversity [McClelland and Mac Clelland, 1961]. Based on the dominant motive, certain unique characteristics can be attributed to individuals forming three distinct archetypes. Power motivated individuals prefer taking high risks, targeting high incentive goals, and prefer working alone. They are highly concerned about status and recognition and attempt to influence others as positively as mentors or as negatively as oppressors [Magee and Langner, 2008; Merrick, 2016]. Achievement motivated individuals also prefer working alone. They like challenging yet achievable goals, hence they tend to take moderate risks. On the other hand, affiliation motivated individuals place more importance in social relationships, so they voluntarily seek opportunities to form affiliations and alliances. They are also characterized to prefer low risk goals and prefer collaboration while avoiding conflicts [Merrick, 2016].

Analyzing these three archetypes, we observe that they mainly distinguish from each other in two aspects; risk preference and social preference. Based on this observation, we conjecture that endowing agents with these three motives would help create social, risk-aware agents making them predictable, explainable autonomous agents. Risk-aware agents are vital for the exploration of extreme/hazardous environments that are too dangerous and uncertain for humans to explore [Hunt *et al.*, 2021; Vielfaure *et al.*, 2022]. Inspired by the behavior of the social spider, Hunt *et al.* [2020] have designed adaptive risk-taking swarms to explore similar settings. Therefore, by using computational models of implicit motives [Merrick and Shafi, 2011], we intend to add an innate behavior to agents making them predictable, risk-aware agents that can be effective teammates in Human-AI teams deployed in hazardous environments.

We experiment on an inherently risky task domain - a fire-fighting task. We first devise a framework that can model rule based agents with an inherent behavior but can also adopt other strategies appropriately based on the defined rules. We created different heterogeneous teams of agents with different motives and analyzed their goal selection behavior, team dynamics and performance using appropriate metrics. Furthermore, we explored the same statistics for homogeneous teams with new motive profiles created by perturbing the three original motives. The use of an established psychology theory paves the path for personality profiling when putting together organic teams to meet motivational needs and increase productivity.

In summary, the contributions of this work include:

1. Proposing a framework to model rule based agents with probabilistic goal-selection capabilities, with an innate behavior (implicit motivation) which can be overridden by other strategies/rules appropriately when required.
2. Conducting experiments in a hazardous domain that require different levels of collaboration, and empirically demonstrate performances of different team compositions on a range of task risk levels.
3. Demonstrating the suitability of a well-established human motivation psychology theory to create risk-aware agents with the ultimate vision of synergized human-AI teaming.

2 Computational Models of Motivation

The three motive profiles affiliation, achievement and power are characterised under incentive-based motivational psychology. Incentives are situational characteristics that drive behavior to satisfy a motive [Heckhausen and Heckhausen, 2008]. Incentives are inversely proportional to the probability of achieving a goal [Merrick, 2016]. This implies that a goal easily achievable with a high probability of success will be associated with a low incentive value. On the other hand, a goal that is much harder to achieve, meaning with a low probability of success, will have a high incentive value.

The three motive profiles are modeled based on their characteristic aspirations towards different incentive ranges. Individuals with high power, achievement and affiliation motives will prefer high, moderate and low incentive goals respectively. To further explain how these computational models could create diverse risk-aware agents, power motivated individuals prefer high-risk goals by nature. A high risk goal can be attributed to having a low probability of success of achieving. Therefore, power motivated agents are computationally modeled such that they will have a high tendency to select high incentive goals. In contrast, since affiliation motive prefer low risks, agents endowed with an affiliation motive will have a high tendency to pick low incentive goals; achievement motivated agents who prefer moderate risks will have a high tendency to prioritize and select goals with moderate incentive values.

These motivation tendencies driven by incentives ($T(I)$), are modeled as a sum of 3 inverted U-shaped curves [Merrick, 2016], where each curve represent the strength of each motive. Therefore, a motive profile can be represented as the sum of the three motives as provided in Equations 1 and 2.

$$T(I) = T_{ach}(I) + T_{aff}(I) + T_{pow}(I) \tag{1}$$

$$\begin{aligned} T(I) = &\frac{St_{ach}}{1+e^{-\rho^{+}_{ach}((1-I)-M^{+}_{ach})}} - \frac{St_{ach}}{1+e^{-\rho^{-}_{ach}((1-I)-M^{-}_{ach})}} \\ &+ \frac{St_{aff}}{1+e^{-\rho^{+}_{aff}(M^{+}_{aff}-I)}} - \frac{St_{aff}}{1+e^{-\rho^{-}_{aff}(M^{-}_{aff}-I)}} \\ &+ \frac{St_{pow}}{1+e^{-\rho^{+}_{pow}(I-M^{+}_{pow})}} - \frac{St_{pow}}{1+e^{-\rho^{-}_{pow}(I-M^{-}_{pow})}} \end{aligned} \tag{2}$$

I is the incentive, ρ^{+}_{aff}, ρ^{+}_{ach}, and ρ^{+}_{pow} control the gradient of approach of each motive while ρ^{-}_{aff}, ρ^{-}_{ach}, and ρ^{-}_{pow} control the gradient of avoidance of the 3 motives. St_{aff}, St_{ach}, St_{pow} are the relative strengths of each motive. M^{+} defines the turning point of success approach and M^{-} defines the turning point of failure avoidance for each motive which are represented by the turning points of the sigmoid curves.

Equation 3 models the 3 motive profiles using only the dominant motive component of Equation 2. In this equation, mot in the parameters St_{mot}, ρ_{mot}, M_{mot} represent motive profiles which could bear the values pow, ach or aff to model power, achievement or affiliation motive respectively. These equations are discussed at length by Merrick [2016].

$$T(I) = \frac{St_{mot}}{1+e^{-\rho^{+}_{mot}(I-M^{+}_{mot})}} - \frac{St_{mot}}{1+e^{-\rho^{-}_{mot}(I-M^{-}_{mot})}} \tag{3}$$

3 Methodology

3.1 Task Domain

The hazardous domain of our choice is a fire-fighting scenario. To create our task domain, we took inspiration from C3Fire microworld [Granlund, 2003], a popular experimental platform to analyze distributed team decision making. C3Fire defines certain roles, a hierarchy of individuals within a team, resource constraints etc., making the abstraction of real-world fire outbreak as realistic as possible. However, we need a team of agents that are homogeneous excluding their motivation, to observe their collective goal selection behavior. Hence, we adopted the terrain, its content, possible actions of an agent, resource types on the terrain, proposed in an example simulation of C3Fire microworld [1] and created a more simplified fire-fighting scenario including selectable goals for the motivated agents. Our task domain is a toroidal world in a 20x20, discrete 2D grid, with houses (property), fast-burning trees and slow-burning trees (vegetation) randomly distributed on the grid. A team of 12 fire-fighting agents with perfect vision that behave according to the proposed framework, are spawned in random locations in the grid, and their overall task is to save the village from fire. In the beginning, fire is spawned in two random cells, and based on the defined fire-spreading likelihood (FSL), fire spreads to its 4 neighbouring cells recursively. That is, if a cell c_{ij} is burning at

[1] http://c3learninglabs.com/w/index.php/Doc/Simulation

time step t, at $t+1$, its neighboring cells ($c_{i+1,j}$, $c_{i-1,j}$, $c_{i,j+1}$, $c_{i,j-1}$) catch fire, gradually engulfing the entire village. A cell with or without a resource is "unharmed" at the beginning. If an agent has removed inflammable fuel (by performing a "fire-break" action) from a goal proactively, the goal is in the state "protected". Once it it is caught on fire, a cell's state changes to "on-fire". If an ignited cell is not extinguished before its burn-out time, the cell is considered "burnt-out" as it has run out of fuel. Fast-burning trees burn out after 3 time steps, slow-burning trees after 5 and houses and bare cells burn out after 10 time steps. These values were selected through pilot experiments in order to keep the simulations simple yet with adequate fidelity. Agents move horizontally or vertically one cell at a time. Whenever a goal is picked by an agent, it takes a single step towards the goal in the dimension with the least distance to the goal [Barrett *et al.*, 2011]. The order in which agents move in a single simulation tick is determined randomly in our simulation implementation (equivalent to a central moderator). If two agents can achieve the same goal without the help of the other in the current time tick, the agent that is selected to move first achieves the goal and the other agent is made to take a step in a random direction. Agents can perform two operations on a goal. They can put out fire (reactive measures) or they can "fire-break" (proactive measures). As the fire spreads engulfing the village, fire-fighters are faced with 7 types of goals listed in Table 1. The proposed framework defines how the motivated agents prioritize and select goals, based on various goal-selection strategies. Simulation ends when all the burning cells/resources are put-out and unharmed resources are protected or the maximum number of time steps (max_time) per simulation (100) ends. If everything burns out at time t before reaching max_time (i.e $t < max\_time$), our design choice was to end the simulation and collect the evaluation metrics at time t. We designed two variations of this task.

1. **Collaborative Task** - In this task, all resources (property and vegetation) need at least two agents to surround the resource by occupying 2 neighbouring cells, to extinguish or fire break. However, an agent can put out bare cells that are on fire by simply moving on to that cell.

2. **Non-Collaborative Task** - In this task, any cell or resource can be extinguished or proactively protected by an agent without the support of the other agents by simply moving on to the cell with the target goal.

Time taken for a fire-breaking or extinguishing action is independent of the number of agents working on the task. In both tasks, an action can be performed in 1 time step.

3.2 Proposed Framework

The proposed framework defines how an agent with an inherent behavior (motive) can probabilistically select goals while accommodating other goal-selection strategies that can supersede the inherent behavior, based on the level of importance of a particular strategy, at appropriate circumstances. Figure 1 depicts this framework.

First, the goal selection criteria/strategies are identified (S^*, S_2, .. S_k). The strategy that defines the characteristic

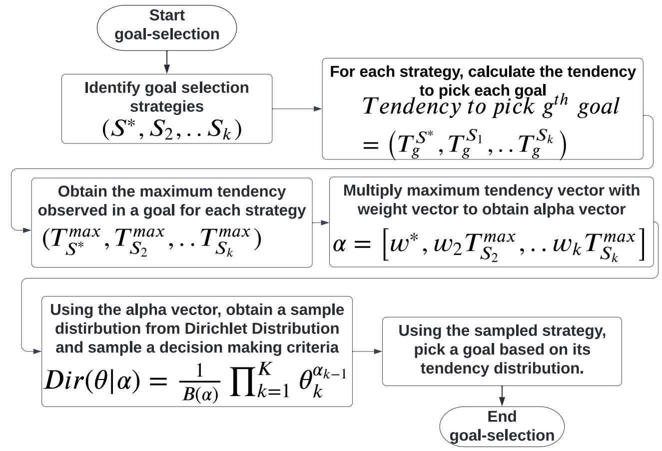

Figure 1: Goal selection process in the proposed framework.

behavior is S^*. Next, each available goal is scored by calculating the tendency of an agent to select that goal based on each strategy. As the third step, the highest tendency values obtained for each strategy is identified (highest tendency recorded by an agent to select a goal based on strategy S_s is $T^{max}_{S_s}$). The strategies are weighed based on their importance. Higher the priority, higher the weight assigned (w_1, w_2, .. w_k). In order to have a default behavior, the expected value of that default strategy (w^*) is denoted by its weight, while the expected value of other strategies is the product of highest tendencies and their weights [w^*, $w_2 T^{max}_{s_2}$, ..$w_n T^{max}_{s_k}$]. This vector of expected values constitute the alpha vector which is the *a priori* to the Dirichlet Distribution. Sampling a distribution from this, would give us which strategy should be used to pick a goal. Then we probabilistically pick the final goal using the selected strategy. As a design rule we recommend the use of soft-max functions to model the activation of a strategy other than S^*, such that values close to 1 are returned when the criteria for a strategy to be activated are met, and values closer to 0 are returned when those strategies should not be applied.

In our implementation, the agents are inherently motivated (S^*). They will always attempt to select goals that satisfy their motives. Assuming these agents all have similar selfless and self-interest levels, we model that they will also use 2 other goal selection strategies - greedy goal selection (S_2), selfless goal selection (S_3), by disregarding their implicit motivational needs for other situational needs. Following are the three goal selection strategies used in our experiments.

Goal Selection Based on Individual Motivation

This inherent strategy of goal selection based on individual motivation is represented by S^*. As the three motive profiles have three distinct risk preferences and social preferences, risk and an opportunity to affiliate, become the 2 situational characteristics that an agent considers when selecting a goal based on its motivation. To represent the risk level of each possible goal on the arena, we assigned a significance value sig (Table 1).

Protecting a fast-burning tree has a higher-incentive value than a slow-burning tree, as the damage caused if a fast-

Goal	Significance (sig)
Fire-break slow-burning tree	1
Fire-break fast-burning tree	2
Fire-break house	3
Extinguishing ignited cell	4
Extinguishing ignited slow-burning tree	5
Extinguishing ignited fast-burning tree	9
Extinguishing ignited house	13

Table 1: Significance values assigned to goals based on their risk and incentive in ascending order

burning tree is caught on fire, is higher. Furthermore, fire-fighters prioritize property over vegetation, hence protecting a house has a higher significance value. The risk is increased when the mentioned three types of goals are on fire, hence goals related to extinguishing fire are assigned higher incentive values as shown in Table 1. These significance values were chosen based on the above logic for demonstration purposes. Researchers adopting the framework should define them appropriately, based on their application domain.

To demonstrate how agents are incentivized by opportunities to collaborate, we calculate the number of agents in a defined vicinity of a goal which we represent by a. For our experiments, we consider agents that are within a range of 2 units to a goal. In their work [Hardhienata, 2015], the authors have proposed a generalized function to model the three motive profiles using the same two situational characteristics - risk and social preference (Equation 4).

$$I(sig, a) = c_1 + c_2 * e^{-(1-(sig/sig_{max}))} * e^{\alpha(1-(a/a_{max}))} \tag{4}$$

sig/sig_{max} is the relative significance value of a goal based on its risk and importance. sig_{max} is the highest significance value assigned to a goal. a/a_{max} represents the relative affinity of a goal. While a is the number of agents in the vicinity of a goal, it is normalized by the maximum number of agents that could occupy the defined vicinity of the goal a_{max}. α gives the required curvature to the graph.

Equation 4 proposed by Hardhienata [2015] is adopted so that, a power motivated agent will have a high tendency to pick goals with a high significance value and have 0 number of agents in the goal's vicinity ($a = 0$), achievement motivated agents would have a high tendency to pick goals that have a moderate significance value and have 1 or 2 agents in the vicinity, and affiliation motivated agents who characteristically prefer collaboration and low risk goals, will pick goals with a low incentive value and have more than 2 agents in the vicinity of a goal. The equation adjusted to model this behavior is visualized in Figure 2.

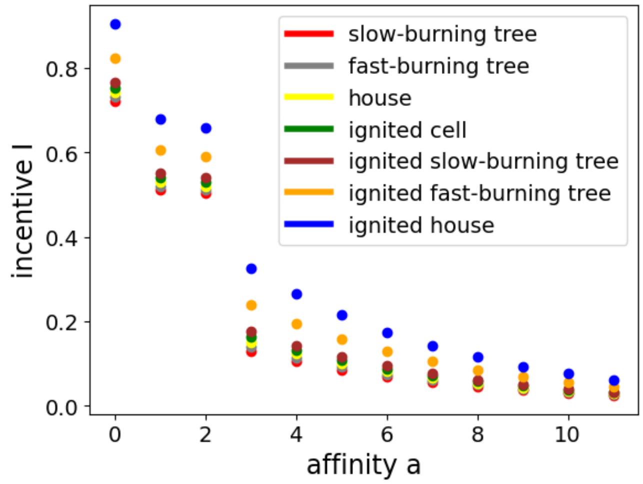

Figure 2: Individual incentive graph based on Equation 4

Tendency To Select Greedy Goal

In a dynamic fire-fighting domain, an agent moving towards a selected goal may encounter important, easily achievable goals that take precedence over its motivational needs. In our fire-fighting domain, if an agent observes a nearby goal on fire, it will prioritize and tend to choose that goal. This greedy behavior is the second goal-selection strategy S_2 employed, common to all agents regardless of their motivation. In our implementation, only the goals that were on fire and are in an immediate cell (Manhatten distance to the goal = 1) is prioritized by a greedy agent. When these conditions are met, the greedy strategy returns a high tendency value (1), and when a goal is not preferable based on the greedy criteria, we return a tendency value as low as 0.001, a value selected empirically.

Tendency To Select a Selfless Goal

We give agents a sense of selflessness where they pick goals that they could tackle by helping or with the help of another agent. This means, when an agent sees a house/tree that needs multiple agents to extinguish/protect, it will assist its fellow agents by disregarding a different goal that would have been preferable based on the other two strategies (S^* or S_2). In the task that requires collaboration, at least two agents are required to protect or extinguish a house. Therefore, when a fire-fighting agent sees a goal that has more than 1 and less than 3 agents in the goal's vicinity, this strategy S_3 is activated returning a high tendency value (1) to select this goal. If there are more than 2 agents or less than 1 agent in a goal's vicinity, S_3 returns a tendency value as low as 0.001. Agent counts 1 and 3 were identified empirically as appropriate criteria to activate this strategy.

Effective Tendency To Select a Goal in a Collaborative Dynamic Environment

The motivated agents modeled in this work are differently motivated, yet equally selfish and selfless. Hence, we define a logical priority ordering an agent assigns to the three strategies used in our experiments. We define the weights w_1, w_2 and w_3 as 1, 2, and 3 respectively to the three strategies so that majority of the time, an agent is driven to select a goal based on its motivation (S^*), but if the condition to select a greedy goal (S_2) is activated, greedy goal strategy is prioritized over motivation strategy. Moreover, if the right conditions are met and the selfless strategy is activated (S_3), S_3 will have a high probability for being the strategy to pick a goal. For the collaborative task we use all 3 strategies

Resource	Score
Cell (Land)	4
Slow-burning tree	8
Fast-burning tree	12
House	16

Table 2: Score assigned for each resource

while for the non-collaborative task only S^* and S_2 are used because the defined selfless strategy (S_3) is not required for non-collaborative goal achievement.

4 Experimental Settings and Results

In order to analyze the performance of a team composition, we considered 3 metrics. First being the *Total Steps*, which is the minimum of either the total step count to finish achieving all the goals or the maximum step count for a simulation. Secondly, we defined a metric that calculates the number of agents per step that pick goals based on a strategy other than motivation. This metric is called *Perceived Tension* and it tracks how often agents in a team have to suppress their inherent motivation behavior in order to pick goals that are deemed easier (S_2) or are helpful (S_3) for other agents. This metric resembles the frustration level of an agent, as in real-world, inhibition of one's true drives causes tension and frustration. Thirdly, we defined scores to each resource on the village to get a final valuation of the damage caused to the village at the end of the simulation. Table 2 defines the scores allocated for each resource type on the grid if they found in their original/unharmed state. If a resource is on-fire at the end of the simulation, 50% of its value is reduced and if its burnt-out by the end, 100% of its value is reduced. This valuation is represented by the metric *Saved Score*.

While both *Total Steps* and *Saved Score* represent a team's performance, we use the latter to identify which team composition is the best, as having a high saved score indicate teams with a better resource prioritization strategy with better damage control.

We conduct experiments with both heterogeneous teams and homogeneous teams. In these experiments, a heterogeneous team is a team with agents with different motives. In a heterogeneous team, each agent's motive is modeled using Equation 3 and the parameter values used by Hardhienata *et al.*, [2012]. We conducted experiments for 28 randomly selected team compositions of 12 agents, and plotted the three metrics on Ternary plots with 3 axes named Aff, Ach and Pow to represent the composition of the motives Affiliation, Achievement and Power in a team. These results were interpolated to obtain the metric values for other team compositions that were not explicitly experimented on, to show the distribution of the metric values across all possible team compositions that could be created with the 3 different motives. For each team composition, 20000 simulations were run to obtain representative averages of the three metrics. Furthermore, a non-parametric Kruskal-Wallis H test (one-way non-parametric ANOVA), followed by Conover's post-hoc test were conducted to identify the best performing teams based on each metric. In a heterogeneous team, an agent has a dominant motive of either affiliation, achievement or power. However, as provided in Equation 2, motivation can be represented as a sum of these three motives. Hence, we created homogeneous teams of 12 agents and created 28 different motive compositions perturbing the three motive strengths (of Equation 2) to sum up to 6 (St_{aff} + St_{ach} + St_{pow} = 6). Then, using the same Ternary plots, we represented the value distribution of 3 metrics for each motive-profile composition obtained by running simulations with the homogeneous teams.

We refer to Figures 3 and 4 for the results obtained for the two tasks with and without collaboration for heterogeneous teams.

When the task is collaborative, we observe that the cluster of heterogeneous teams with the lowest *Total Steps* changes as FSL increases (Figures 3a - 3e). When the task is relatively simple and less risky ($FSL < 0.3$), we observe that predominantly affiliation motivated teams complete the task faster (lower step count). This is because such teams prioritize low incentive goals and prefer cooperating. Therefore, they succeed in proactively protecting goals collaboratively while putting out the slowly spreading fire fast. However, when FSL gradually increases, we observed that teams that are heterogeneous and contain all three types of agents with a slight majority of affiliation motivated agents, complete the task faster. Hence, for collaborative tasks, the existence of motivational diversity is important. When the task requires no collaboration (Figures 4a - 4e) and FSL is increased gradually, predominantly power and achievement motivated teams complete the task faster than all-affiliation motivated teams. This is because as FSL increases, the number of high incentive goals increase. Teams with a majority of power and achievement agents prioritize such goals and contain the fire more strategically, making such teams beneficial for tasks with high severity but requires less collaboration. Although our observations are such, it was also observed that when FSL is high ($FSL \geq 0.4$), the simulation ends before reaching its maximum step count, as the fire has engulfed the entire village. Therefore, we cannot make concrete conclusions about the teams that finish the task efficiently at high FSLs, by relying on the *Total Steps* metric. This is why the best teams are ranked based on the *Saved Score* metric.

Analyzing the *Perceived Tension* plots for heterogeneous teams in Figures 3f - 3j and 4f - 4j, we observe that the cluster of teams with the least tension behave the same way for both collaborative and non-collaborative tasks. For tasks that are relatively less risky ($FSL < 0.3$), teams with a majority of achievement and affiliation motives have the least tension. This indicates that achievement and affiliation agents appropriately pick moderate incentive goals and low incentive goals as such goals are abundant in low to moderately risky environments. Hence, more opportunities for achievement and affiliation agents to pursue goals in line with their motivation exist, making such teams less tensed. As FSL increases, when the task is collaborative, results show that the cluster of teams with the highest tension move to the center of the plot making all-affiliation and all-power motivated teams relatively less tensed teams. This indicates that when the task

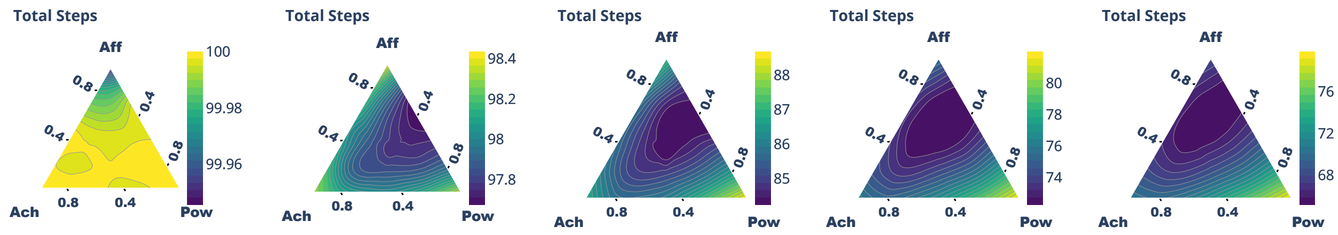

(a) FSL = 0.1 (b) FSL = 0.2 (c) FSL = 0.3 (d) FSL = 0.4 (e) FSL = 0.5

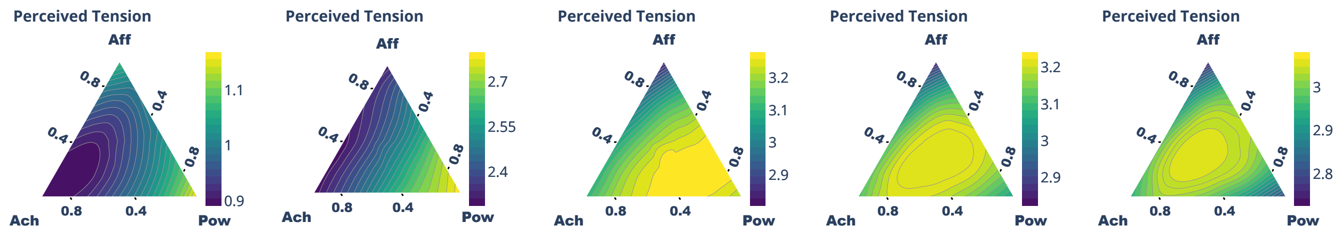

(f) FSL = 0.1 (g) FSL = 0.2 (h) FSL = 0.3 (i) FSL = 0.4 (j) FSL = 0.5

Figure 3: Ternary plots with the distribution of scores obtained for metrics *Total Steps* and *Perceived Tension* of each team composition for the **collaborative task**, when FSL (fire-spreading likelihood) vary between 0.1 - 0.5.

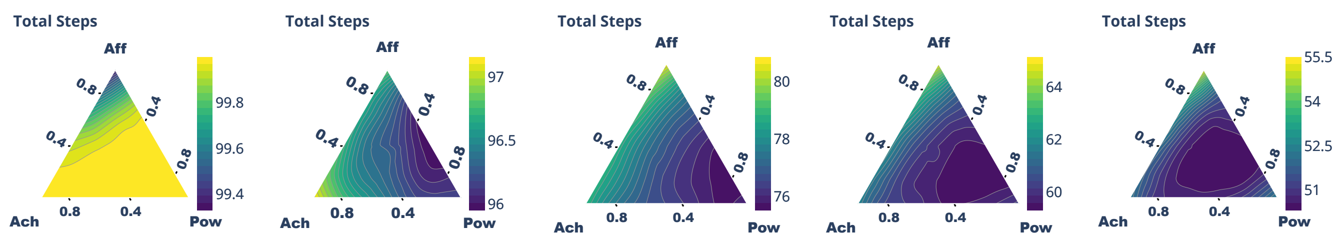

(a) FSL = 0.1 (b) FSL = 0.2 (c) FSL = 0.3 (d) FSL = 0.4 (e) FSL = 0.5

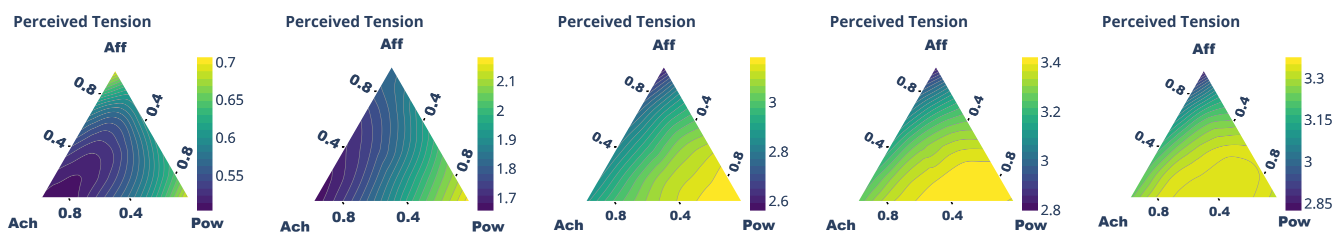

(f) FSL = 0.1 (g) FSL = 0.2 (h) FSL = 0.3 (i) FSL = 0.4 (j) FSL = 0.5

Figure 4: Ternary plots with the distribution of scores obtained for metrics *Total Steps* and *Perceived Tension* of each team composition for the **non-collaborative task**, when FSL (fire-spreading likelihood) vary between 0.1 - 0.5.

risk-level is high, teams that have similar motives experience less tension rather than teams that have diverse, conflicting motives.

In the non-collaborative task (Figures 4f - 4j), we hypothesized that power motivated agents will be the least tensed, especially in high FSLs, considering their characteristic affinity towards high incentive goals. However, this was not the case and they reported high motivation inhibition. This is introduced from the design of our incentive function. While the incentive function drives power motivated agents to pick high incentive goals, it also drives agents to explore goals that are isolated. When the fire is spreading fast, new goals are dynamically generated in close proximity. Hence, agents who are exploring goals away from other agents suffer and they resort to other goal picking strategies such as greedy/selfless strategies, causing high tension/inhibition.

Next, we analyze how the *Saved Score* metric increases with FSL in experiments with heterogeneous teams and homogeneous teams. Figures 5a and 5b summarize how the center of gravity of the cluster of best team compositions, with

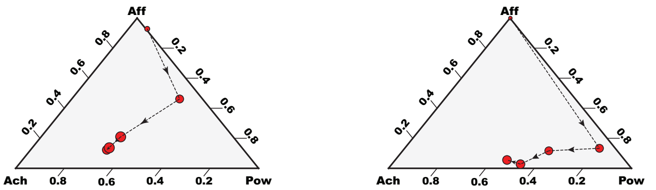

(a) *Saved Score* for **collaborative task**.

(b) *Saved Score* for **non-collaborative task**.

Figure 5: Team compositions with the highest *Saved Score* shift as FSL increases in experiments with heterogeneous teams

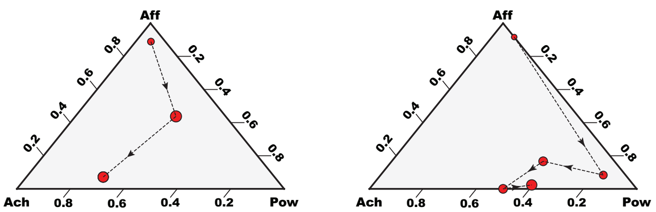

(a) *Saved Score* for **collaborative** task

(b) *Saved Score* for **non-collaborative** task

Figure 6: Best motive profile compositions based on the *Saved Score* shift as FSL increases in experiments with homogeneous teams

respect to *Saved Score* metric, shift as FSL increases from 0.1 to 0.5 in heterogeneous teams. The size of the circles are proportionate to the number of team compositions that are statistically identified as equally performing at a given FSL. Arrows guide how the affiliation, achievement, power ratios of agents in the teams with the least damage changes when FSL increases. From these plots it is evident that all-affiliation teams have the best strategy when FSL is low. A higher number of power agents and an even higher number of achievement agents are required for a team to perform more risky tasks in a collaborative setting. In contrast, we observe that in a non-collaborative task, best team compositions contain more power motivated agents, a moderate number of achievement agents and a lower number of affiliation agents. However, as FSL increases even further, teams with a combination of achievement and power becomes more robust. These observations align with the innate risk and social collaboration preferences of three motive archetypes, validating the proposed framework's ability to model agents with an inherent behavior while facilitating other strategies.

Figure 6 visualizes how the best motive profile for *Saved Score* changes as FSL increases in homogeneous teams in the two tasks (collaborative, non-collaborative). Arrows guide how affiliation (St_{aff}), achievement (St_{ach}) and power (St_{pow}) strengths of motive profiles in the best teams change as FSL increases. Although the plots pertaining to *Total Steps* and *Perceived Tension* for experiments with homogeneous teams are not depicted, it was observed that they follow a similar trend to the results obtained for the same metrics with heterogeneous teams. Analyzing the plots for *Saved Score* in Figures 6a and 6b, we observe how teams with high St_{aff} perform the best in both tasks, making them robust in low risk levels regardless of the level of collaboration the task needs. However, as FSL increases, the motive profiles that perform well require high St_{ach} and St_{pow} strengths. When the task is collaborative, having a high St_{ach} and a moderately high St_{pow} is more suitable. In non-collaborative tasks, teams with high St_{pow} motives excel by adopting effective strategies and minimizing harm to the village. High St_{pow} indicates a focus on high incentive goals. These teams thrive in non-collaborative tasks since they have a preference for working independently. Consequently, teams with high St_{pow} demonstrate superior performance in non-collaborative tasks compared to collaborative tasks.

5 Conclusion and Future Work

This paper applies a well-established human motivation psychology theory to create predictable artificial agents making risk-aware decisions. We present a framework for modeling agents with inherent predictability while allowing for various behavior strategies. Two tasks, one collaborative and one non-collaborative, were designed to analyze team performance and motive profile compositions across different risk levels. The results were evaluated using three metrics. Our experiments have shown that teams thrive when they constitute a portfolio of diverse motive profiles, especially in collaborative tasks, establishing the importance of functional heterogeneity in teams; an observation aligning with the work of O'Shea-Wheller, Hunt, and Sasaki [2021]. It also highlights the importance of adjusting team composition based on task risk levels to enhance performance. We found that a combination of affiliation and achievement agents is optimal for low-risk tasks, while a combination of power and achievement agents is more suitable for high-risk tasks. This observation is consistent with the findings of Di Pietrantonio *et al.* [2019], validating our framework's ability to model innate team behavior through incentive functions. Consequently, teams with the most effective strategies for different risk levels align with the characteristic preferences of their motives, resulting in predictable agents and laying the groundwork for enhancing Human-AI teams. In their research, [Noeldeke *et al.*, 2022] compare different decision-making models of humans using agents, ranging from Rational Choice Theory to Random choice models. Our proposed model aligns with the Theory of Planned Behavior, accommodating motivation, greediness, and selflessness as antecedents of behavioral intentions. Meanwhile, [Azaria, 2022] propose 7 solution concepts for modelling agents that can work with humans, from non-adaptive rule-based agents to an ensemble of models. Our work offers a new solution concept, allowing agents to select a strategy from a range of strategies and adapt to the situation, making it more mathematically versatile than previous work. Creating complementary agents with diverse risk and collaboration preferences improves predictability, perception, and team performance. Our findings on perceived tension inform the creation of organic teams with minimal motive inhibition. Initial experiments involve rule-based probabilistic goal-selection agents, but future research aims to extend these motives to learning agents, enabling adaptive risk-aware teams as the environment evolves.

References

[Azaria, 2022] Amos Azaria. Irrational, but adaptive and goal oriented: Humans interacting with autonomous agents. IJCAI, 2022.

[Bansal *et al.*, 2021] Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S Weld. Is the most accurate ai the best teammate? optimizing ai for teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11405–11414, 2021.

[Barrett *et al.*, 2011] Samuel Barrett, Peter Stone, and Sarit Kraus. Empirical evaluation of ad hoc teamwork in the pursuit domain. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 567–574, 2011.

[Cooke *et al.*, 2013] Nancy J Cooke, Jamie C Gorman, Christopher W Myers, and Jasmine L Duran. Interactive team cognition. *Cognitive science*, 37(2):255–285, 2013.

[Di Pietrantonio *et al.*, 2019] Josef Di Pietrantonio, Rachael Miller Neilan, and James B Schreiber. Assessing the impact of motivation and ability on team-based productivity using an agent-based model. *Computational and Mathematical Organization Theory*, 25(4):499–520, 2019.

[Gorman *et al.*, 2006] Jamie C Gorman, Nancy J Cooke, Harry K Pedersen, Jennifer Winner, Dee Andrews, and Polemnia G Amazeen. Changes in team composition after a break: building adapative command-and-control teams. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 50, pages 487–491. SAGE Publications Sage CA: Los Angeles, CA, 2006.

[Granlund, 2003] Rego Granlund. Monitoring experiences from command and control research with the c3fire microworld. *Cognition, Technology & Work*, 5(3):183–190, 2003.

[Hardhienata *et al.*, 2012] Medria K. D. Hardhienata, Kathryn E. Merrick, and Valery A. Ugrinovskii. Task allocation in multi-agent systems using models of motivation and leadership. *2012 IEEE Congress on Evolutionary Computation*, pages 1–8, 2012.

[Hardhienata, 2015] Medria KD Hardhienata. *Models of Motivation for Particle Swarm Optimization with Application to Task Allocation in Multi-Agent Systems.* PhD thesis, University of New South Wales, Canberra, Australia, 2015.

[Heckhausen and Heckhausen, 2008] Jutta Ed Heckhausen and Heinz Ed Heckhausen. *Motivation and action*. Cambridge University Press, 2008.

[Hunt *et al.*, 2020] Edmund R Hunt, George Jenkinson, Michael Wilsher, Carl P Dettmann, and Sabine Hauert. Spider: a bioinspired swarm algorithm for adaptive risk-taking. In *ALIFE 2020: The 2020 Conference on Artificial Life*, pages 44–51. MIT Press, 2020.

[Hunt *et al.*, 2021] Edmund R Hunt, Conor B Cullen, and Sabine Hauert. Value at risk strategies for robot swarms in hazardous environments. In *Unmanned Systems Technology XXIII*, volume 11758, pages 158–177. SPIE, 2021.

[Magee and Langner, 2008] Joe C Magee and Carrie A Langner. How personalized and socialized power motivation facilitate antisocial and prosocial decision-making. *Journal of Research in Personality*, 42(6):1547–1559, 2008.

[McClelland and Mac Clelland, 1961] David C McClelland and David Clarence Mac Clelland. *Achieving society*, volume 92051. Simon and Schuster, 1961.

[McNeese *et al.*, 2021] Nathan J McNeese, Mustafa Demir, Erin K Chiou, and Nancy J Cooke. Trust and team performance in human–autonomy teaming. *International Journal of Electronic Commerce*, 25(1):51–72, 2021.

[Merrick and Shafi, 2011] Kathryn E Merrick and Kamran Shafi. Achievement, affiliation, and power: Motive profiles for artificial agents. *Adaptive Behavior*, 19(1):40–62, 2011.

[Merrick, 2016] Kathryn E. Merrick. *Computational Models of Motivation for Game-Playing Agents*. Springer, 2016.

[Noeldeke *et al.*, 2022] Beatrice Noeldeke, Etti Winter, and Elisée Bahati Ntawuhiganayo. Representing human decision-making in agent-based simulation models: Agroforestry adoption in rural rwanda. *Ecological Economics*, 200:107529, 2022.

[O'Neill *et al.*, 2020] Thomas O'Neill, Nathan McNeese, Amy Barron, and Beau Schelble. Human–autonomy teaming: A review and analysis of the empirical literature. *Human Factors*, page 0018720820960865, 2020.

[O'Shea-Wheller *et al.*, 2021] Thomas A O'Shea-Wheller, Edmund R Hunt, and Takao Sasaki. Functional heterogeneity in superorganisms: Emerging trends and concepts. *Annals of the Entomological Society of America*, 114(5):562–574, 2021.

[Schelble *et al.*, 2020] Beau G Schelble, Christopher Flathmann, and Nathan McNeese. Towards meaningfully integrating human-autonomy teaming in applied settings. In *Proceedings of the 8th International Conference on Human-Agent Interaction*, pages 149–156, 2020.

[Siu *et al.*, 2021] Ho Chit Siu, Jaime Peña, Edenna Chen, Yutai Zhou, Victor Lopez, Kyle Palko, Kimberlee Chang, and Ross Allen. Evaluation of human-ai teams for learned and rule-based agents in hanabi. *Advances in Neural Information Processing Systems*, 34:16183–16195, 2021.

[Vielfaure *et al.*, 2022] David Vielfaure, Samuel Arseneault, Pierre-Yves Lajoie, and Giovanni Beltrame. Dora: Distributed online risk-aware explorer. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 6919–6926. IEEE, 2022.