



UvA-DARE (Digital Academic Repository)

RLupus: Cooperation through emergent communication in The Werewolf social deduction game

Brandizzi, N.; Grossi, D.; Iocchi, L.

DOI

[10.3233/IA-210081](https://doi.org/10.3233/IA-210081)

Publication date

2021

Document Version

Final published version

Published in

Intelligenza Artificiale

License

Article 25fa Dutch Copyright Act (<https://www.openaccess.nl/en/policies/open-access-in-dutch-copyright-law-taverne-amendment>)

[Link to publication](#)

Citation for published version (APA):

Brandizzi, N., Grossi, D., & Iocchi, L. (2021). RLupus: Cooperation through emergent communication in *The Werewolf* social deduction game. *Intelligenza Artificiale*, 15(2), 55-70. <https://doi.org/10.3233/IA-210081>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

RLupus: Cooperation through emergent communication in *The Werewolf* social deduction game

Nicolo' Brandizzi^a, Davide Grossi^{b,c,*} and Luca Iocchi^a

^a*Dipartimento di Ingegneria Informatica, Automatica e Gestionale, Sapienza University of Rome, Italy*

^b*Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, The Netherlands*

^c*Amsterdam Center for Law and Economics and Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands*

Abstract. This paper focuses on the emergence of communication to support cooperation in environments modeled as social deduction games (SDG), that are games where players communicate freely to deduce each others' hidden intentions. We first state the problem by giving a general formalization of SDG and a possible solution framework based on reinforcement learning. Next, we focus on a specific SDG, known as *The Werewolf*, and study if and how various forms of communication influence the outcome of the game. Experimental results show that introducing a communication signal greatly increases the winning chances of a class of players. We also study the effect of the signal's length and range on the overall performance showing a non-linear relationship.

Keywords: Multi-agent systems, social deduction games, deep reinforcement learning, emergent communication

1. Introduction

Social deduction games (SGDs) are games characterized by the confrontation between two or more parties, one of which is usually seen as an evil faction. The other parties must deduce the real intention of the latter, seeing through lies and deceptions. While the details of these games may change, free communication—i.e., players can communicate with each other with no limitations—is a common aspect for them all.

In artificial settings, communication would lead to increased complexity in the environment, both on the user side, where engineers are tasked to design an

expressive and robust syntax [15, 17, 37], and on the artificial players who have to learn the syntax and the meaning of the available words. For shallow players, this task quickly becomes unfeasible. A common solution is to define a game-specific language containing communication semantics. This language is devised by usually injecting some expert knowledge into the system. The language becomes then part of the game, providing new available communication actions augmenting the agents' choices in the decision-making process. Instead of developing a language for the agents, in this article, we propose another approach enabling players to use a free communication mechanism.

*Corresponding author: Davide Grossi, Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen, The Netherlands. Amsterdam Center for Law and Economics and Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands. E-mail: d.grossi@rug.nl.

Context of the paper. Games arising from social interaction have been extensively studied within the multi-agent reinforcement learning (MARL) literature. Indeed, MARL systems have been successfully used to model a wide variety of social systems found

in nature and modern society, from food-collective ants systems to complex sharing of information in social networks [7]. In these settings, the agents' behavior can be cooperative, competitive or a mix of the two. In [7] a study of these different kinds of settings and the algorithms associated with them is pursued while [29, 46] focuses on the difference between competitive and cooperative agents.

Although MARL has been used for a wide variety of applications, the reinforcement learning paradigm may not scale up easily in complex multi-agent systems. For this reason, RL has been integrated with deep neural networks (DeepRL, DRL) [38]. This allows RL to scale to problems that were previously intractable, such as playing video games using pixels as input [21]. Shortly after [45], this approach has been applied to multi-agent systems to study the complex emergent behavior of multiple agents interacting with each other to reach a goal. Although DRL is a well-established field of study, agent interaction via actions or communication remains a challenging problem.

This research area is strictly tied to the study of complex behavior arising from the interaction of simple agents; these aspects are mainly studied through the usage of structured game systems. In their work, Baker et al. [3] showed how a few simple rules from the *Hide'n Seek* game can generate complicated behaviors to the point of exploiting environmental errors to their advantages. Along the same line, Leibo et al. [27] carried out extensive analysis on the problem of autotricula and non-stationary learning in multi-agent deep reinforcement learning (MADRL); they pointed out how the interaction of competitive agents can culminate in an endless cycle of counter-strategies due to the non-stationarity of the environment.

In particular, one such complex behavior consists in the emergence of communication between cooperative agents. Recent works investigate this aspect in various environments varying from joined image captioning [18], to negotiation [8] and simulated pointing games [35]. Our paper is a contribution to this line of research, focusing on communication in social deduction games.

Paper contribution and outline. The paper makes two main contributions:

- a formal description of social deduction games coupled with a general reinforcement learning solution framework, which allows for free

communication among agents without requiring to provide game-specific knowledge;

- an analysis of the performance obtained through learned communication behaviors in an instance of the above class of games: *The Werewolf*¹ social deduction game.

The paper is organized as follows. An overview of the related work is provided in Section 2. In Section 3, the problem statement is formalized together with the general RL framework. Section 5 describes the *Werewolf* game instance in detail, defining both the game logic and the actual implementation. This game provides a fit ground to study language emergence since its whole system is based on communication, indeed it is the subject of an annual *AiWolf* context in Japan. The experimental settings and results are reported and commented in Section 6 together with the comparison between our work and [23]. Finally, a discussion is provided in Section 7.

The code is available at https://github.com/nicofirst1/rl_werewolf.

2. Related work

Our work relies on the findings coming from the social interaction field of psychology, coupled with the multi-agent systems and reinforcement learning.

Social Deduction Games. Social deduction games have been studied in the broader context of social interactions [13]. In particular, they have been used to study the role of rationality in inter-personal interaction [10], analyze the different forms of social mechanics [11], and research the role of communal topology [1], however, the deduction part of these games has been neglected.

Indeed, in their work [9], the authors give a mathematical formulation for a general social game in order to simplify the way to design such games; however, no specific formulation is given for deduction games.

On the other hand, [54] study the most influential information source in social deduction games and concludes that the interactions that occurred prior to the game are regarded as most important to the player. Although this approach is reasonable in the context of acquaintances, no result is given for games in which the playing parties do not know each other.

In all these works, the goal is centered around the social interaction between players. In our work, we

¹Also known as *Lupus in Fabula*.

shift the focus to finding an optimal policy to improve the performances of a party.

Multi-agent Deep RL. Applying RL paradigms to find an optimal policy for multi-agent systems is a well-established line of work [6, 41, 46]. In recent years deep neural networks (DNN) have been used to solve complex tasks such as playing Atari games [21], cooperating in Hide-and-seek [3] and competing with humans on strategic games [42, 49].

A common paradigm for training RL agents in a deep setting is policy gradient methods, [43, 44] where the gradient of a parametrized policy is used to guide the agent in the direction of maximal expected reward. In our work, we leverage a particular instance of these methods called Proximal Policy Optimization [40].

In all such cases, the complexity of the environment coupled with a DNN generated unexpected behaviors that are usually counter-intuitive for humans but achieve greater performances in the task at hand. In our work, we leverage this aspect and further study how the coordination between agents varies under different communication instances.

Emergent Communication.. This phenomenon has been exploited in the newly born field of emergent communication [51], where agents are given the choice to use a communication channel to achieve a common goal.

The workshop on Emergent Communication (Emecom) includes many publications in the field of natural language processing [28] strictly tied to MADRL [16, 26, 29] and social deduction games as an environment. Standard games for this line of research are the Task & Talk [25], which is centered around dialogue, on the other hand, the Pointing Game [35] grounds the communication into natural image processing.

Another instance of these settings is *The Werewolf* game, where the players find themselves split into two opposite groups in a partially known environment. This game has gained increased popularity in the field of cooperation through emergent communication, especially in Japan, where the annual AiWolf contest [5, 20, 24, 36] sees artificial agents competing with and against human players to win the game with fixed language syntax. In particular [52], set up a 5-player game with additional roles and use a Deep Q-Network to determine who to trust or kill.

In our work, we choose *The Werewolf* as an instance of the general SDG framework. However, our implementation differs from the ones in the

AiWolf contest, the closest being [23], where the authors use Q-learning to study the winning chances of the villagers in a game with 16 players, divided into 14 villagers and 2 werewolves. Indeed we drop the hand-coded syntax and let the players develop their own communication by defining some general attributes of the channel.

3. Problem statement

Social deduction games (SDG) are characterized by the presence of a number of opposing parties

$$P^{(1)}, P^{(2)}, \dots, P^{(m)}$$

(typically $m = 2$), each containing a finite number $n^{(k)}$ of players, which may differ per party. For the sake of conciseness, we drop the k dependency and denote the number of players in a party as simply n :

$$P^{(k)} = \{p_1^{(k)}, p_2^{(k)}, \dots, p_n^{(k)}\}$$

with $1 \leq k \leq m$. We denote by $N = \bigcup_{1 \leq k \leq m} P^{(k)}$ the set of all players.

The game evolves as a sequence of actions performed by the players of the parties, typically in turns. The effect of these actions contributes to the definition of the game score.

The goal of each party is to prevail over the others by performing suitable social behaviors, including, for example, leveraging other players by means of bluffs and lies. These deceitful methods are the base for any SDG and force every player to perform a deductive analysis on the member of the other parties. In general we can identify two types of goals in SDGs:

- An agent-based *micro-goal*, which is the main factor steering the agent’s behavior, either in isolation (in competitive environments) or together with other agents’ goals (in cooperative environments).
- A party-based *macro-goal*, expressing the aligned interests of the members of the same party that, combined, make up the party goal.

During the execution of the game, agents can communicate among them, either implicitly or explicitly. We define explicit communication as the act of sharing information for the sole purpose of affecting other agents’ mental states. On the other hand, implicit communication regards all those actions that carry

more than one meaning, e.g., guiding someone to a goal. Decisions about how, when, and what to communicate are critical choices for the game’s success.

We thus distinguish two categories of actions performed by the players: 1) *game actions*, that are actions affecting the evolution of the game, 2) *communication actions* that are actions affecting only the mental state (i.e., the knowledge state) of the players. In this formalization, we consider only forms of explicit communication, while studying forms of implicit communication is left as future work.

With the previous assumptions, in this paper, we consider the formalization of an SGD with the following elements:

1. the action set $\mathcal{A} = \mathcal{G} \cup \mathcal{C}$ is made of two separate components:
 - a finite set of possible game actions \mathcal{G} the elements of which we will denote with g .
 - a set of unidirectional communication actions $C_{i,j}(b)$ intended to convey some information $b \in \mathcal{B}^2$ between two players:

$$C_{i,j}(b) : p_j \rightarrow p_i \quad \forall p_j, p_i \quad j \neq i \in N$$

2. the state set $\mathcal{O} = \mathcal{E}^N \times \mathcal{W}^N \times \mathcal{V}$, with N being the total number of players, built out of three elements³:
 - a set of agent’s features \mathcal{E} representing the game situation of each agent (typically visible to all other agents),
 - a set of agent’s internal states \mathcal{W} (e.g., representations of beliefs not visible to other agents).
 - a set of environment states \mathcal{V} that are common for all the agents (i.e., independent from the agent states).
3. an environment \mathcal{S} implementing the game logic. \mathcal{S} can be seen as a function taking as input the agents actions and yielding a new state obtained as the result of execution of such actions

$$\mathcal{S} : \mathcal{O} \times \mathcal{A}^N \rightarrow \mathcal{O} \quad (1)$$

Notice that, while it may be relatively easy to formalize the specifications of game actions, for example, in terms of pre-conditions and post-conditions using action representation formalisms, it is less clear how to formalize communication actions since it would require an explicit model of agents’

knowledge. For modeling this kind of communication actions, the use of typical action formalisms is not straightforward. For example, they may need to be extended with epistemic operators [14].

4. General solution framework

This article studies social deduction games that can be formalized as a multi-agent (deep) reinforcement learning (RL) scenario. In such scenarios, each party has to choose an optimal strategy or policy⁴ (i.e., an optimal assignment from states to actions) to maximize the game score. As already mentioned, a particular feature of SDGs is the presence of communication actions and the need to choose optimal communications among the players within a party.

The problem of learning optimal policies in multi-agent games is indeed well known, and many solutions are available. However, when communication actions are involved, using artificial intelligence techniques to make optimal decisions about how, what, and when to communicate is still a challenging problem under investigation, and fewer research works are available.

The advantage in defining a solution based on RL is that it does not require an explicit model of the transition function for communication actions. In other words, optimal behaviors can be computed without associating a semantic meaning to the communication actions. While this feature can be considered not desirable for some kinds of applications (e.g., for mixed human-AI teams), it is very convenient for AI teams based on RL that can learn their own communication language to win the game. Indeed AI agents can effectively learn a communication language without making the semantics of communication explicit. The explainability of learned communication actions is left as future work.

Action Policy. We consider a single game turn t as a sequence of $k + 1$ steps, where the first k steps are associated to only communication actions $c \in \mathcal{C}$ while the last step is a game action $g \in \mathcal{G}$. We denote each time step as $\mathcal{O}_t^j \quad \forall j \in [0, k + 1]$.

Now, we can define a policy π that uses the information from \mathcal{O}_t^j in order to choose an action as

²A possibly infinite set of all possible signals.

³We consider a setup where all the agents N choose an action simultaneously.

⁴Although the coordinated behavior favors the parties’ macro-goal, the actual policy works on the players’ micro-goal level.

follows:

$$\begin{aligned}\pi(\mathcal{O}_t^j) &\in \mathcal{C} \quad \forall j \leq k \\ \pi(\mathcal{O}_t^{k+1}) &\in \mathcal{G}\end{aligned}$$

Note that in some cases it can be useful to compute two distinct policies π_C ; π_G , one for communication and one for the game actions.

5. R-Lupus framework

In this section, we present the *RLupus* framework for the Werewolf social deduction game in which we apply the above formalization. In this specific context, we aim to study if and how various forms of communication can influence the outcome of the game, in which only one party is able to learn while the other one has a fixed, hand-coded policy.

Under the same circumstances, the hypothesis is that the agents who can communicate will perform much better than those not allowed to exchange information. Moreover, we speculate that different communication settings will have diverse influences on the amount of coordination among the agents and the outcome of the game.

In the following sub-sections, we give a brief introduction on *The Werewolf* game logic, a description of the RL environment with all its components and the policies used for the players in the game.

5.1. The werewolf game

Werewolf is a social deduction game modeling conflicts between two groups in a partially known environment. In its easiest version, the game sees two groups ($M = 2$), villagers $P^{(v)}$ and werewolves $P^{(w)}$ where $P^{(v)} > P^{(w)} + 1$. The wolves know exactly the identity of each player, while the villagers are certain exclusively about their role and the number of werewolves. In an open setup, an additional moderator is needed to coordinate the players. The game is divided into two phases: night and day, interleaving each other.

The game ends either when the villagers execute the last werewolf, or there is an even number of both roles. The latter case implies the wolves winning since the execution phase can be stalled, thus taking away the only possibility for villagers to kill the wolves.

5.2. Reinforcement learning

Before describing the *RLupus* environment, a brief description of the Reinforcement Learning paradigm must be introduced.

In the canonical RL environment, the problem satisfies the Markov property⁵ so it can be formulated as a Markov decision process (MDP). An MDP is defined by a tuple of five elements $(\mathcal{O}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$ defined as follows:

- A finite set of observations called the observation space \mathcal{O}
- A finite set of actions, the action space \mathcal{A}
- A transition model $\mathcal{T}_a(m_t, m_{t+1})$ defining the probability to transition from state m_t to a new state m_{t+1} given action a .
- $\mathcal{R}(m_t, m_{t+1})$ the reward associated to the previous transition.

At each timestep t , an RL agents receives a state m_t from a dynamic environment \mathcal{O} . The agent then selects an action a_t from the action space \mathcal{A} following a policy $\pi(a_t|m_t)$. The action is processed by \mathcal{S} , equation 1, which transitions to the next state m_{t+1} and yields a reward r_t to the agent. This feedback loop continues until the agent reaches a terminal state and restarts.

The agent aims to maximize the expected discounted reward given by:

$$R_t = \sum_{i=0}^{\infty} \gamma^i r_{t+i}$$

Where $\gamma \in (0, 1]$ is the discount factor governing the importance of future rewards.

5.3. *RLupus* environment

Dealing with RL implies the presence of an action and state set; the latter are referred to as action space and observation space, which are presented in the following section⁶.

Action Space. As mentioned in Section 5.4, the Werewolf uses a policy $\pi(\mathcal{O}_t^j)$ in order to choose both the communication and game actions. In a way, the agent can be seen as performing a game g_t and communication c_t action simultaneously.

⁵The future depends only on the current state and action.

⁶For the sake of conciseness some formulations are omitted from the body of the paper. The interested reader can refer to the Appendix A.

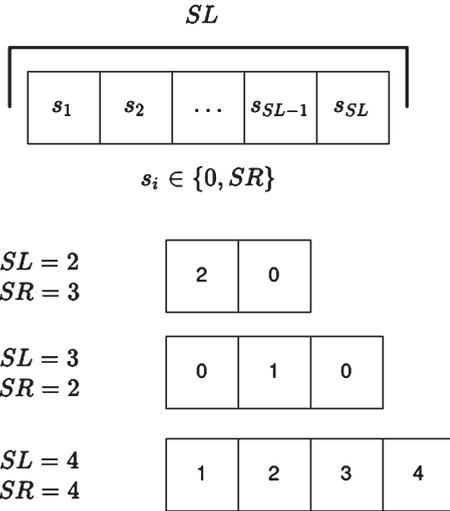


Fig. 1. Signal vector with general formulation as well as specific values for SL and SR .

Indeed the action space is divided into two parts:

- *Target*: The target g_t is an integer in range $g_t \in [0, N - 1]$, where N is the number of players. Its intended usage is to allow players to vote for other players during the game. The range of possible values never changes during the execution; instead, illegal actions, such as voting for dead players, are filtered out later in the model.
- *Signal*: The signal vector c_t (see Figure 1) is defined by two integer values: its length $SL \in [0, \infty[$ which can be any value from zero (no communication) to an arbitrary large integer; its range $SR \in [2, N]$ that defines the number of possible values it can have. SR bounded below by 2, since a signal with only one possible value would be considered a static vector carrying no information; the upper bound N comes from the necessity for the signal to be embedded with the target. Both are used to define the valid space for communication before training.

Observation Space. The observation space characterizes what the agents perceive in the environment. This space includes both other agents' actions and information about the environment⁷.

Transition Model. The transition model fully depends on the target action; that is, the environment

is independent of the communication between agents. Indeed \mathcal{S} switches between night and day at every timestep t and removes dead players from the environment.

Rewards. The rewards, or penalties, are the core of the environment and determine how the players interact, learn, and develop new strategies; an agent's main goal is to take actions that will maximize the expected reward.

Following our formalization, the environment is responsible for delivering a reward to each player⁸.

Metrics. To measure the changes in the agent behavior, the following normalized metrics are logged:

- *Suicide*: the number of times an agent votes for itself during an execution phase.
- *Wins*: the villagers' wins are plotted in the normalized range of values.
- *Average days*: average number of days before a match ends.
- *Accord*: This value represents, on average, the percentage of agents that vote for the same target during the two execution phases.

5.4. Policies

An agent's policy defines the behavior of a player during the game. In this environment, there are two kinds of policies: *trainable* policies use custom algorithms to collect experience and learn to maximize the reward; *static* policies are hard-coded behaviors that are used to guarantee a fixed baseline trough out the evaluation.

In this work, we assign static policies to opponent players (werewolves) and training policies to AI agents (villagers) learning how to win the game.

Static policies for werewolves. Static policies are reserved for the werewolf agents; their aim is to allow a baseline evaluation of the villager learning. Since wolves are more likely to win in a completely random environment, applying such policies is enough to prove the development of new strategies for the villagers if the winning rates are to change significantly. Three policy are implemented:

⁷The complete observation space is described in Appendix A.1.1.

⁸The complete set of rewards is described in Appendix A.1.2.

- *Random Target*: chooses a random non-dead player among the villagers during the execution phase.
- *Random Target Unite*: this policy targets the same player for every wolf, both during day and night execution; this allows the werewolf to dominate the day execution phase with random villagers.
- *Revenge Target*: with this policy, the wolves will either vote randomly or target a villager who previously voted for a wolf.

Trainable policies for villagers.. In our work, we chose Proximal Policy Optimization (PPO) as a training algorithm for its widespread success in multi-agent environment [19, 53], but future research could also focus on other algorithms such as TRPO [39] or MADDPG [31]. PPO uses a surrogate loss function to keep the difference between the old and the new policy within a safe range⁹. The model is a simple fully connected network with an LSTM cell¹⁰.

Learning Werewolves.. Up until now, only villager agents have been able to learn from experience, while the werewolf behave according to a static policy. Here, we address the possibility of an environment where both villagers and werewolves are able to learn. Having multiple agents learning simultaneously void the stationary assumption that is necessary for optimality in RL. The presence of multiple villagers already invalidates this assumption, but the coordination between them alleviates the problem. Introducing an adversarial set of learning agents would cause an increased complexity that would cloud the goal of this paper, which is to study the emergence of a language between agents.

6. Results

Following, an analysis of the results for both nine (Section 6.1) and twenty-one (Section 6.2) players is given. In both cases, a baseline setting with no communication is compared with multi-channel communication to show the increased performance reported using the metrics in Section 5.3.

⁹More information on the policy loss are available in Appendix A.1.3.

¹⁰Given the partial observability of the problem, we found the LSTM as a suitable element to approximate the state space.

Table 1

RLupus: Multi-channel metrics. The first column reports the type of **comm**(unication)channel regarding the **SignalLength** and the **SignalRange**. The next four show the metrics values for villagers winning rate, suicide rare, number of days elapsed and accordance rate

Comm	Win Vil	Suicide	Days	Accord
0SL	0.044	0.086	1.55	0.47
1SL-2SR	0.19	0.078	1.58	0.47
1SL-9SR	0.21	0.078	1.58	0.47
9SL-2SR	0.45	0.067	1.9	0.47
9SL-9SR	0.19	0.077	1.58	0.46

Moreover, for the nine players instance, results for the additional *revenge* and *unite* policies are reported against the *random* one.

Finally, in Section 6.3, a summary for a setting with 16 players is given for both the *AiWolf* environment [23] and the *RLupus* one.

6.1. Nine players

A game with nine players is relatively short but not trivial for the villagers. Indeed, in a completely random environment, they have a probability of 3.12% to win¹¹.

Since the random policy is believed to hold the least expert knowledge about the game environment, we present the results for this policy only in the next section. On the other hand, Section 6.1.2 also shows the results for the *revenge* and *unite* policies.

6.1.1. Multi-channel communication

Table 1 presents the results (columns) obtained with different forms of communication (rows), with different communication settings (SL=signal length, SR = signal range). From the table, it is clear that any form of communication improves over the non-communication setting (0SL). Moreover, it shows how the bit communication (2SR) performs much better than the full ranged one (9SR). Indeed, with the addition of just one bit of communication, the villagers can perform as well as the full extended communication (9SL-9SR). In fact, for the settings where $SR = 9$, increasing the channel length SL has the only effect of speeding up the convergence by a factor of circa 25% for each increase.

On the other hand, the bit communication, independently from the channel length, provides generally better results. This leads to the conclusion that the two parameters dictating the change in the

¹¹See Appendix A.2 for the complete estimation.

Table 2

RLupus: single/no channel metrics. The Design Choice part of the table shows which kind of policy Random, Unite or Revenge has been used in relation to the communication, while the Results half present the metric' values

Design Choice				Results			
Comm	Rnd	Unt	Rvg	Vil Win	Accord	Suicide	Days
OSL	X			0.044	0.478	0.086	1.55
OSL		X		0.03	0.695	0.059	1.5
OSL			X	0.12	0.482	0.078	1.64
1SL-2SR	X			0.19	0.47	0.078	1.58
1SL-2SR		X		0.08	0.685	0.055	1.52
1SL-2SR			X	0.4	0.479	0.065	1.9

communication channel are not equally influential when it comes to the agent's learning.

On a final note, the average time to train the agents was 6 hours on a single GPU¹² machine.

6.1.2. Policies comparison

Table 2 reports results also related to the werewolves policies (Rnd = Random, Unt = Unite, Rvg = Revenge).

When no communication is available (OSL), the accord value is almost identical for both the revenge and the random policy and much higher for the unite one, as previously anticipated.

The number of days reports a similar trend being smaller in the unite settings where the game ends sooner. Finally, the revenge winning rate reaches a much greater value than the other policies, 12%. The motivation being the simplicity of the policy itself, i.e., the wolves can not hide behind purely random action anymore and are not strong enough to drive the majority of the votes toward a villager.

Following the considerations for the three policies and bit communication ($SL = 1$ $SR = 2$) referred to Table 2:

- *Random*: the villager's winning rate reaches 20% ; 4.5 times more than the previous condition and 6.5 times the theoretical winning rate.
- *Unite*: as in the previous case, the coordination is much greater. Indeed the villagers are able to increase their winning rate by a factor of $\times 3$. This result alone proves that introducing a limited communication channel can greatly favor the outcome even in the most disadvantageous setting.
- *Revenge*: again, the revenge policy is the easiest to spot for a trained agent reaching a winning rate of 40%.

Table 3

Mapping between outcomes and probabilities

Outcome	Prob. %	Leaves	Total %
0-12	0.029	1	
0-10	0.143	4	
0-8	0.447	10	
0-6	1.162	20	
0-4	2.819	35	
0-2	7.02	56	11.62
1-1	21.06	56	
2-2	27.965	21	
3-3	26.017	6	
4-4	26.017	1	88.38
Total	1.0	210	1.0

Table 4

RLupus: 21 Player metrics

Comm	Win	Suicide	Days	Accord
OSL	0.42	0.072	7.84	0.56
1SL-2SR	0.25	0.075	7.74	0.57
9SL-2SR	0.98	0.04	7.3	0.56
21SL-2SR	0.94	0.05	7.6	0.56
1SL-21SR	0.72	0.062	8	0.56
21SL-21SR	0.61	0.066	7.9	0.56

6.2. Twenty-one players

Unlike what was studied in the previous section, the twenty-one players (21P) environment has more room for the villagers to win in a completely random setting. Indeed, by building a tabular representation of the expanded tree (shown in Table 3), the reader can see how the total villagers' winning possibilities increased to 11.62% in a completely random environment.

6.2.1. Multi channel communication

A comparison among all the multi-channel settings is reported in Table 4.

According to the previous section's findings, the bit communication (9SL – 2SR) seems to perform better than the full-ranged one (*SL – 21SR) both in terms of winning ratio and suicides.

On the other hand, the 1SL – 2SR instance performs worse than the no communication one; this anomaly can be attributed to an incorrect environment exploration. As mentioned in Table 3, the expanded tree size is greater than the instance with nine players; thus, the algorithm can get more easily stuck in a local minimum. However, every other setting with a communication channel has a higher winning rate than the one without; thus, one could confidently say

¹²Nvidia Geforce GTX 1080 Ti.

Table 5
Winning rate for [23] and our method in both baseline and active learning

Design		Win Vil	
		AiWolf	RLupus
No Learning	AiWolf baseline	38.6	/
	RLupus baseline	/	56.9
Learning	DQN / hand-coded	52.9	/
	PPO / OSL	/	90.2
	PPO / 1SL-2SR	/	98.4

that the communication signal is indeed helping the agents exploring the possible branches better than in other cases. For these reasons, we believe this outcome shows how the shape of the communication signal can improve the agents’ capacity to explore the environment and thus should be further investigated in future work.

6.3. Sixteen players

As previously mentioned, Kajiwara et al. [23] presented the AiWolf framework, implementing a well-defined semantic¹³ constrained to be both easy to use for the artificial players and understandable by humans. They extended the agents’ adaptive capabilities by adopting a Deep Q-network architecture and reported an increase in the villagers’ winning rate when they are the only ones capable of learning. Although their methodology is similar to ours, their work is not directly comparable to RLupus because the two frameworks represent the game in different ways. The different set-ups lead to related but orthogonal findings, which we discuss in this section.

We trained an environment with 16 players (2 werewolves and 14 villagers) and reported our results in Table 5, together with the results of the AiWolf framework.

It is interesting to observe the increase in performance due to the learning process. Indeed, the table shows baseline results without learning in the two frameworks and results obtained with RL.

In comparison with AiWolf, it can be seen how the OSL setting in RLupus already improves the winning chances by 33% reaching 90.2% and adding a single bit communication results in almost perfect victory on the villagers’ side (98.4%).

¹³More details about the AiWolf framework are available in Appendix A.3.

7. Conclusions and future work

In this paper, we have defined a formalism for social deduction games, in which communication is an essential part of the game together with the allowed actions. On top of that, a general resolution framework based on reinforcement learning was defined and applied to the *The Werewolf* SDG by studying how various forms of communication influenced the outcomes of a match.

As shown by the experimental results, the introduction of different forms of communication greatly increases the agents’ performance (villagers). In particular, we observed that a Boolean signal range is preferred to an integer one. The reason for this is unclear. We speculate it may lay in the duality of the roles of the game.

Moreover, we found how much of the villagers’ winning rate is determined by the agents not voting for themselves while keeping the accord value to a maximum. This translates into better coordination, which is possible only when there is a sufficient amount of communication present in the environment.

Also, we noticed that there is no linear map between the amount of communication permitted, i.e., *SL* and *SR*, and the overall performance. Indeed there seems to be an optimal combination of the two, which depends mainly on the signal length.

We conclude by identifying three directions in which our work could be extended.

Model of SDGs. One possible line of research consists of studying an instance of an SDG where the communication cannot be intrinsically tied to the action space, i.e., multi-step communication games such as Task & Talk [25]. Alternatively, one could decide to depart from the deep part of the RL resolution framework and choose an SDG instance whose environment can be optically solved with standard reinforcement methods, e.g., Pointing game [35].

The Werewolf. On the other hand, various possibilities arise from the study of The Werewolf game. One such could be the analysis of the language used for communication to highlight potential patterns, given that performance alone should not be considered a valid metric for the study of emergent communication [30]. Moreover, one could extend the RLupus environment either by adding other roles, for example, the medium and the witch, or using normalized

continuous vectors for the communication channel, which allow for the backpropagation of errors directly through the communication channel [16]¹⁴. Finally, a deeper analysis could be applied to understand the impact of the communication parameters (SR and SL) on the metrics of the system and the overall performance.

Human-machine coordination. Coordination among artificial agents is a key engineering challenge: from task allocation [12, 50], knowledge management [55], distributed constraint optimization problems [34] to multi-robot SLAM [47, 56] and language models [4, 48]. And even more challenging is coordination between artificial and human agents. The study of emergent communication in SDGs could provide useful novel insights for the development of more efficient coordination among artificial agents and between artificial and human agents.

References

- [1] G. Abramson and M. Kuperman, Social games in a social network, *Physical Review E* **63**(3) (2001), 030901.
- [2] Z. Ahmed, N.L. Roux, M. Norouzi and D. Schuurmans, Understanding the impact of entropy on policy optimization. In K. Chaudhuri and R. Salakhut-dinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA, volume 97 of Proceedings of Machine Learning Research*, pages 151–160. PMLR, 2019.
- [3] B. Baker, I. Kanitscheider, T.M. Markov, Y. Wu, G. Powell, B. McGrew and I. Mordatch, Emergent tool use from multi-agent autotutorials, In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, (2020).
- [4] Y. Bengio, R. Ducharme, P. Vincent and C. Janvin, A neural probabilistic language model, *The Journal of Machine Learning Research* **3** (2003), 1137–1155.
- [5] X. Bi and T. Tanaka, Human-side strategies in the werewolf game against the stealth werewolf strategy, In *International Conference on Computers and Games*, pages 93–102. Springer, (2016).
- [6] L. Busoniu, R. Babuska and B. De Schutter, A comprehensive survey of multiagent reinforcement learning, *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, **38**(2) (2008), 156–172.
- [7] L. Buşoniu, R. Babuška and B. De Schutter, Multiagent reinforcement learning: An overview, In *Innovations in multi-agent systems and applications-1*, pages 183–221. Springer, (2010).
- [8] K. Cao, A. Lazaridou, M. Lanctot, J.Z. Leibo, K. Tuyls and S. Clark, Emergent communication through negotiation, In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, (2018).
- [9] K.T. Chan, I. King and M.-C. Yuen, Mathematical modeling of social games, In *2009 International Conference on Computational Science and Engineering*, volume 4, pages 1205–1210. IEEE, (2009).
- [10] A.M. Colman, Cooperation, psychological game theory, and limitations of rationality in social interaction, *Behavioral and Brain Sciences* **26**(2) (2003), 139–153.
- [11] M. Consalvo, Using your friends: Social mechanics in social games, In *Proceedings of the 6th International Conference on Foundations of Digital Games*, pages 188–195, (2011).
- [12] M.M. de Weerd, Y. Zhang and T. Klos, Multiagent task allocation in social networks, *Autonomous Agents and Multi-Agent Systems* **25**(1) (2012), 46–86.
- [13] M. Eger and C. Martens, Keeping the story straight: A comparison of commitment strategies for a social deduction game, In *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, (2018).
- [14] R. Fagin, J.Y. Halpern, Y. Moses and M.Y. Vardi, *Reasoning About Knowledge*. MIT Press, (1995).
- [15] T. Finin, R. Fritzson, D. McKay and R. McEntire, Kqml as an agent communication language, In *Proceedings of the third international conference on Information and knowledge management*, pages 456–463, (1994).
- [16] J. Foerster, I.A. Assael, N. De Freitas and S. Whiteson, Learning to communicate with deep multi-agent reinforcement learning, In *Advances in neural information processing systems*, pages 2137–2145, (2016).
- [17] M.R. Genesereth, R.E. Fikes, et al., Knowledge interchange format-version 3.0: reference manual. (1992).
- [18] L. Graesser, K. Cho and D. Kiela, Emergent linguistic phenomena in multi-agent communication games, In K. Inui, J. Jiang, V. Ng, and X. Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3698–3708. Association for Computational Linguistics, (2019).
- [19] Y. Guan, Y. Ren, S.E. Li, Q. Sun, L. Luo and K. Li, Centralized cooperation for connected and automated vehicles at intersections by proximal policy optimization, *IEEE Transactions on Vehicular Technology* **69**(11) (2020), 12597–12608.
- [20] Y. Hirata, M. Inaba, K. Takahashi, F. Toriumi, H. Osawa, D. Katagami and K. Shinoda, Werewolf game modeling using action probabilities based on play log analysis, In *International Conference on Computers and Games*, pages 103–114. Springer, (2016).
- [21] I. Hosu and T. Rebedea, Playing atari games with deep reinforcement learning and human checkpoint replay, *CoRR*, *abs/1607.05077*, 2016.
- [22] E. Jang, S. Gu and B. Poole, Categorical reparameterization with gumbel-softmax, In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, (2017).
- [23] K. Kajiwara, F. Toriumi, H. Ohashi, H. Osawa, D. Katagami, M. Inaba, K. Shinoda, J. Nishino, et al., Extraction of optimal strategies in human wolf using reinforcement learning, *Proceedings of the 76th National Convention* **2014**(1) (2014), 597–598.
- [24] D. Katagami, S. Takaku, M. Inaba, H. Osawa, K. Shinoda, J. Nishino and F. Toriumi, Investigation of the effects of non-verbal information on werewolf, In *2014 IEEE International*

¹⁴On the other hand, [22] and [33] use a Gumbel approximation to backpropagate the error through a discrete distribution.

- Conference on Fuzzy Systems (FUZZ-IEEE)*, pages 982–987. IEEE, (2014).
- [25] S. Kottur, J.M.F. Moura, S. Lee and D. Batra, Natural language does not emerge ‘naturally’ in multi-agent dialog, In M. Palmer, R. Hwa, and S. Riedel, editors, *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2962–2967. Association for Computational Linguistics, (2017).
- [26] A. Lazaridou and M. Baroni, Emergent multi-agent communication in the deep learning era, *CoRR*, abs/2006.02419, (2020).
- [27] J. Z. Leibo, E. Hughes, M. Lanctot and T. Graepel, Autocurricula and the emergence of innovation from social interaction: A manifesto for multi-agent intelligence research, *CoRR*, abs/1903.00742, (2019).
- [28] Y. Li, E.M. Ponti, I. Vulic and A. Korhonen, Emergent communication pretraining for few-shot machine translation. In D. Scott, N. Bel, and C. Zong, editors, *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 4716–4731. International Committee on Computational Linguistics, (2020).
- [29] P.P. Liang, J. Chen, R. Salakhutdinov, L. Morency and S. Kottur, On emergent communication in competitive multi-agent teams, In A. E. F. Seghrouchni, G. Sukthankar, B. An, and N. Yorke-Smith, editors, *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, pages 735–743. International Foundation for Autonomous Agents and Multiagent Systems, (2020).
- [30] R. Lowe, J.N. Foerster, Y. Boureau, J. Pineau and Y.N. Dauphin, On the pitfalls of measuring emergent communication, In E. Elkind, M. Veloso, N. Agmon, and M. E. Taylor, editors, *Proceedings of the 18th International Conference on Autonomous Agents and Multi Agent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, pages 693–701. International Foundation for Autonomous Agents and Multiagent Systems, (2019).
- [31] R. Lowe, Y. Wu, A. Tamar, J. Harb, P. Abbeel and I. Mordatch, Multi-agent actor-critic for mixed cooperative-competitive environments, *arXiv preprint arXiv:1706.02275*, (2017).
- [32] J. Luo, S. Green, P. Feghali, G. Legrady and Ç.K. Koç, Visual diagnostics for deep reinforcement learning policy development, *CoRR*, abs/1809.06781, (2018).
- [33] C.J. Maddison, A. Mnih and Y.W. Teh, The concrete distribution: A continuous relaxation of discrete random variables, In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, (2017).
- [34] P.J. Modi, W.-M. Shen, M. Tambe and M. Yokoo, Adopt: Asynchronous distributed constraint optimization with quality guarantees, *Artificial Intelligence* **161**(1-2) (2005), 149–180.
- [35] I. Mordatch and P. Abbeel, Emergence of grounded compositional language in multi-agent populations, In *Thirty-Second AAAI Conference on Artificial Intelligence*, (2018).
- [36] N. Nakamura, M. Inaba, K. Takahashi, F. Toriumi, H. Osawa, D. Katagami and K. Shinoda, Constructing a human-like agent for the werewolf game using a psychological model based multiple perspectives, In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–8. IEEE, (2016).
- [37] P.D. O’Brien and R.C. Nicol, Fipa—towards a standard for software agents, *BT Technology Journal* **16**(3) (1998), 51–59.
- [38] J. Schmidhuber, Deep learning in neural networks: An overview, *Neural Networks* **61** (2015), 85–117.
- [39] J. Schulman, S. Levine, P. Abbeel, M. Jordan and P. Moritz, Trust region policy optimization, In *International conference on machine learning*, (2015), 1889–1897.
- [40] J. Schulman, F. Wolski, P. Dhariwal, A. Radford and O. Klimov, Proximal policy optimization algorithms, *CoRR*, abs/1707.06347, (2017).
- [41] Y. Shoham, R. Powers and T. Grenager, Multiagent reinforcement learning: a critical survey, *Web manuscript*, **2** (2003).
- [42] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., Mastering the game of go with deep neural networks and tree search, *Nature* **529**(7587) (2016), 484–489.
- [43] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra and M. Riedmiller, Deterministic policy gradient algorithms. (2014).
- [44] R.S. Sutton, D. McAllester, S. Singh and Y. Mansour, Policy gradient methods for reinforcement learning with function approximation, *Advances in Neural Information Processing Systems* **12** (1999), 1057–1063.
- [45] A. Tampuu, T. Matiisen, D. Kodelja, I. Kuzovkin, K. Korjus, J. Aru, J. Aru and R. Vicente, Multiagent cooperation and competition with deep reinforcement learning, *PloS one* **12**(4) (2017), e0172395.
- [46] M. Tan, Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, (1993), 330–337.
- [47] S. Thrun and Y. Liu, Multi-robot slam with sparse extended information filters, In *Robotics Research, The Eleventh International Symposium*, pages 254–266. Springer, (2005).
- [48] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, L. Kaiser and I. Polosukhin, Attention is all you need. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, (2017).
- [49] O. Vinyals, I. Babuschkin, W.M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D.H. Choi, R. Powell, T. Ewalds, P. Georgiev, et al., Grandmaster level in starcraft ii using multi-agent reinforcement learning, *Nature* **575**(7782) (2019), 350–354.
- [50] P. Vytelingum, T.D. Voice, S.D. Ramchurn, A. Rogers and N.R. Jennings, Agent-based micro-storage management for the smart grid. (2010).
- [51] K. Wagner, J.A. Reggia, J. Uriagera and G.S. Wilkinson, Progress in the simulation of emergent communication and language, *Adaptive Behavior* **11**(1) (2003), 37–69.
- [52] T. Wang and T. Kaneko, Application of deep reinforcement learning in werewolf game agents, In *2018 Conference on Technologies and Applications of Artificial Intelligence (TAAI)*, pages 28–33, (2018).
- [53] H. Wei, X. Liu, L. Mashayekhy and K. Decker, Mixed-autonomy traffic control with proximal policy optimization, In *2019 IEEE Vehicular Networking Conference (VNC)*, pages 1–8. IEEE, (2019).

- [54] S. Wiseman and K. Lewis, What data do players rely on in social deduction games? In *Extended Abstracts of the Annual Symposium on Computer-Human Interaction in Play Companion Extended Abstracts*, (2019), 781–787.
- [55] D.-J. Wu, Software agents for knowledge management: coordination in multi-agent supply chains and auctions, *Expert Systems with Applications* **20**(1) (2001), 51–64.
- [56] X.S. Zhou and S.I. Roumeliotis, Multi-robot slam with unknown initial correspondence: The robot rendezvous case, In *2006 IEEE/RSJ international conference on intelligent robots and systems*, pages 1785–1792. IEEE, (2006).

Appendix

A. Appendix

Following, some specific aspects for the *RLupus* environment are discussed. We decided to provide them in the Appendix to avoid interrupting the normal flow of the paper. Indeed what comes in the next paragraphs are implementation details which are not required to understand the contribution of this article.

A.1. Environment

A.1.1. Observation Space

The observation space characterizes what the agents perceive in the environment. Multiple elements make up this space in the environment; each of them is later packed into a common dictionary and fed to the players:

- *Phase*: $p \in \{1, 2, 3, 4\}$, to represent the four different phases of the game.
- *Day*: $d \in \mathcal{T}^+$, to represent the day. An upper bound for the possible maximum day is set to 10 which reduces the one-hot encoder representation length.
- *Status Map*: $d \in \mathcal{B}$, this array maps a boolean value to an agent index, effectively informing the players if agent i is alive = 1 or dead = 0.
- *Own id*: $oi \in [0, N]$, at the start of each game the agents' ids are shuffled. To keep an agent informed about its new position this integer is necessary.
- *Targets*: $t \in [-1, N]$, to group all the targets.
- *Signal*: $s \in [-1, SR - 1]$, to group all the signals when there are any.

A.1.2. Rewards

In the environment there are five conditions that determine if an agent is punished or rewarded:

- *Day*: At the end of each day every agent is penalized by a small factor (-1). The purpose is to train the agents to quickly win the game, trying to develop new policies to win faster.
- *Death*: Each time a player dies, it takes a -5 penalty. While dying can be a strategy in some matches, this behavior needs to be controlled to avoid a high number of suicides.
- *Target Accord*: To keep the voting system become another form of communication, the agents are penalized when they voted for someone that later on are not killed. Since execution depends on the coordination and communication of the required group of agents, this rewards is aimed to encourage a cooperative behavior
- *Win/Lost*: Until now the rewards were assigned per agent. At the end of a match each group, either werewolves or villagers, is rewarded or penalized by a factor ± 25 .

A.1.3. Trainable policy

Since the neural network architecture is positioned in between the policy and the value function, the loss function must take into account the contribution given by the latter, thus an additional error term must be included:

$$L^{CLIP+VF} = E_t \left[L_t^{CLIP}(\theta) - c_1 L_t^{VF}(\theta) \right]$$

Where $L_t^{VF}(\theta) = (V_\theta(s_t) - V_t^{arg})^2$ is the squared-error loss between the value function V_θ at state s_t and the target function, c_1 is its coefficient. This addition is of fundamental importance to estimate the critic loss, that is the how well the model is able to predict the value of each state.

An additional term is introduced to regularize the total loss:

$$L^{CLIP+VF+S} = E_t \left[L_t^{CLIP+VS}(\theta) + c_2 S[\pi_\theta](s_t) \right]$$

The entropy coefficient is maximized when all the policies have equal probability to be chosen, that is when the agent is acting at random. Adding such value to the loss function incentives the training algorithm to minimize

the entropy value, thus avoiding the premature convergence of one action probability dominating the policy and preventing exploration¹⁵. The term c_2 scales the value of the entropy.

A.2. Baselines

Understanding how AIs learn is a hard problem in the field of machine learning. In computer vision and specifically with convolutional neural networks (CNN), visualizing the activation layers together with the filters is a common practice to better understand the structure of the network itself.

On the other hand, for RL environments, there has been some research in the area of visual diagnostic for RL systems using CNNs to elaborate visual stimulus [32]; for emergent language, the preferred way to quantitatively evaluate the collaboration between agents has been accomplished by measuring the performances of the agents themselves.

It is trivial to see that, whichever method is preferred, there must be some kind of evaluation method for the communication system. In the following, such evaluation is estimated firstly by drawing some baselines with complete random agents, and then by comparing the latter with the results yield by the training.

A game with nine players is relatively short.

These agents determine a simple set of actions that can be easily mapped in a table; for this reason, the following section will focus only on nine players' settings (9P). Before showing any results the study of the complete behavior of a 9P setting is performed by inspecting its tree of possibilities.

An expanded tree, or tree of possibilities, shows every outcome of a particular process. At each possible intersection, a new branch is created associated with a probability p ; in the context of the werewolf game, each branch represents an outcome where, during the execution phase, the player has killed either a villager or a werewolf.

It is worth noticing that in a random 9P environment, the villagers are most likely to lose. As can be seen in the expanded tree for the game, 2, the villagers either execute a werewolf during the first day, with probability $p = 3/8$, or they lose. Moreover the only chance to win is achievable with probability:

$$p = \frac{3}{8} \cdot \frac{1}{3} \cdot \frac{1}{4} = 0.03125$$

Hence they have a probability of 3.12% to win with a complete random policy.

A.3. *AiWolf*

This section aims to provide some further details concerning the different implementation between the *RLupus* framework and the *AiWolf* one [23].

A.3.1. Syntax protocol

On their website they define a specific protocol for the structure of sentences which is the combination of 6 words, i.e. unit of meaning:

1. *Subject*: an agent identifier which has a clear mapping with the *Own Id* parameter specified in the previous Section A.1.1
2. *Target*: an agent identifier, again the parallelism is the same with the *RLupus* setting.
3. *Role*: one of the 6 implemented roles. This kind of specification is not needed in our setting since only two roles are available.
4. *Species*: either human or werewolf.
5. *Verb*: a set of 15 valid verbs.
6. *Talk number*: a unique id for each sentence.

These atomic information units can define 13 different type of sentences which can be broadly divided into 5 categories:

¹⁵For a complete understanding of the influence entropy has on policy optimizations, the reader is referred to [2].

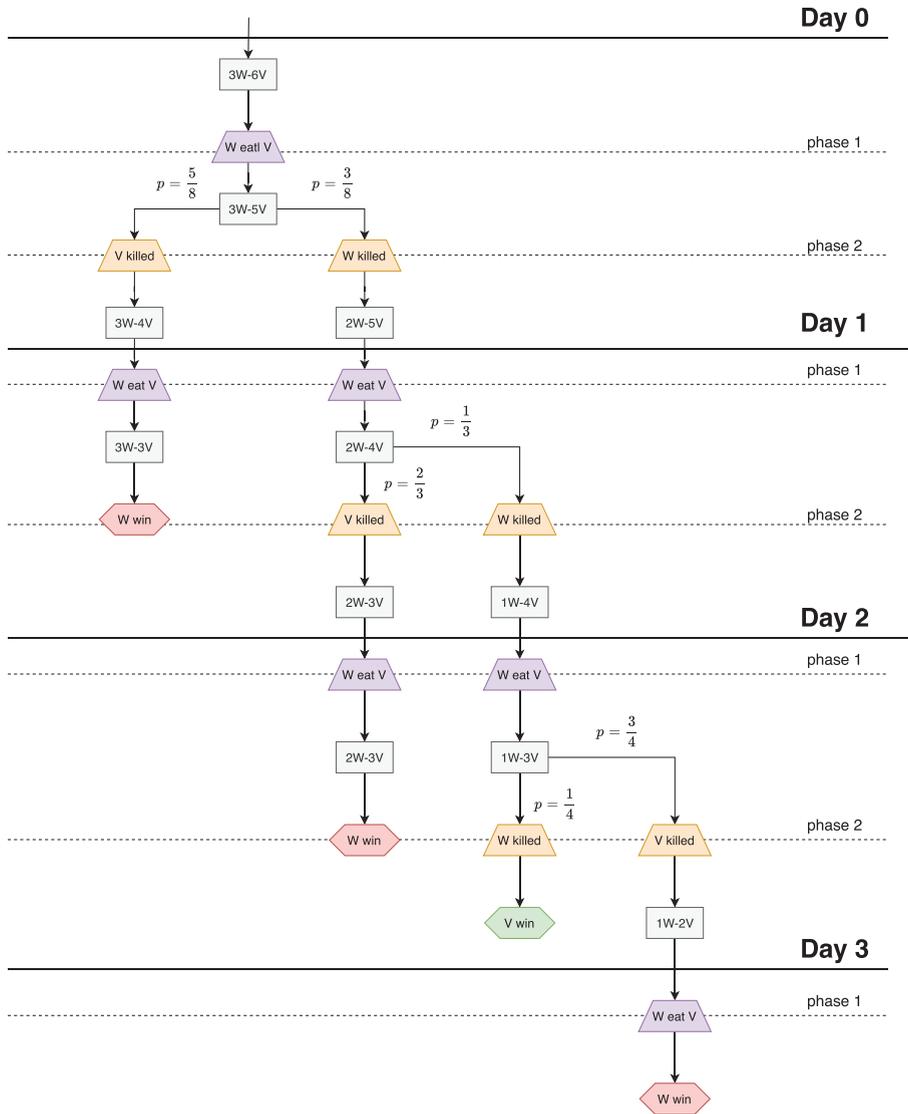


Fig. 2. Expanded tree for a 9 players game with 3 wolves and 6 villagers.

1. Sentences expressing knowledge or intent.
2. Sentences about ongoing actions.
3. Sentences about past actions and their results.
4. Sentences that express disagreement.
5. Sentences referring to the flow of the conversation.

Finally, there are 8 types of operators which are used to frame sentences and express their relationship:

1. Request operators which are directed towards the acquisition of new knowledge.
2. Reasoning operators: e.g., because.
3. Time indication operators.
4. Logic operators such as: and, not, or, xor.

A.3.2. Comparison

An initial parallelism can be traced between the number of words and the signal parameters¹⁶. Still, the main difference is that *AiWolf* forces a prior meaning on the available words while *RLupus* leaves them free to be used by the agents. On one hand, defining a grammar and a syntax beforehand allows for a highly interpretable language, but on the other, this constrains the agents to adapt to a syntax that could not be optimal.

Moreover, since a sentence must be grammatically correct, some constraint enforcing this aspect must be applied to the agent through policy-shaping or on their output in a later phase. This further increases the complexity of the *AiWolf* environment and define a solid discrepancy with the freedom expressed in the *RLupus* one.

Finally, in the *RLupus* framework, no direct relationship between sentences can be express by the agents since no unique id is defined for a sentence in the observation space.

¹⁶SL, SR in Section A.1.1