

1 We thank all three reviewers for their comments and insightful suggestions. We have edited the manuscript to address  
2 them. We outline some of these changes here. One important addition is that we have now packaged our source code  
3 with trained model weights and data for open source release <sup>1</sup>. We believe that the availability of this code will aid in  
4 the reproduction and extension of our results by other researchers in the community.

5 *I think the author needs to argue why Avalon is a better agent for real-world hidden role scenarios than other*  
6 *games?* Among hidden role games, Avalon is one of the most popular and widely played games (according to  
7 boardgamegeek.com) including an active online community of player (proavalon.com). In the real world, there are  
8 subtle cues which can often be misinterpreted when others are acting under uncertainty. Avalon is not necessarily  
9 fundamentally better than other games like Saboteur, but its combination of discrete actions and natural language as  
10 well as its active online community make it a good candidate for study.

11 *For the methodology part, can you differentiate your method from existing methods such as MCTS + value network*  
12 *in alphaGO more explicitly?* Our approach uses CFR instead of MCTS. AlphaGo-like methods can be used when  
13 the board state alone is sufficient to determine the best move, but in imperfect information games it is necessary to  
14 consider how players acted to reach the current board state. We’ve added the following sentence: “Compared to  
15 MCTS-based methods like AlphaGo, CFR-based methods like DeepStack and DeepRole can soundly reason over  
16 hidden information.”

17 *Does the proposed method generalize to other games such as werewolf or saboteur? ... Do we actually want to a*  
18 *case-by-case AI or general intelligent agents?* We agree that the goal is to create generally intelligent agents rather  
19 than specific agents designed to play a particular game. The DeepRole algorithm extends the CFR algorithm which  
20 was developed for two-player competitive imperfect information games such as poker. We showed that CFR and value  
21 networks can be used to build a more general system capable of cooperation in addition to competition. In principle,  
22 DeepRole could be applied directly to Saboteur. However, for other hidden role games such as Werewolf or Mafia,  
23 understanding and producing natural language is a key missing component. We mention in the discussion: “In future  
24 work, we will investigate whether the interpretable belief state of DeepRole could also be used to ground language,  
25 enabling better coordination through communication.”

26 *Need ablation and analysis — we all know trained agents are vulnerable to adversarial human players — e.g. the*  
27 *online dota bots, who beat professionals, can be easily beaten by streamers after a few days. Are you claiming your*  
28 *bot is not 'hackable'?* We have carried out experiments with lesioned versions of DeepRole which showed that our  
29 novel neural network architecture and deductive reasoning component were key drivers of performance (Figure 5). We  
30 are not claiming that the bot is not exploitable. Because it is possible for agents to team up on others in flexible ways,  
31 exploitability is less straightforward to measure. Additionally, over the 1-2 week period of online play, the Avalon  
32 community (including streamers) did not collectively find a strategy that could consistently win against DeepRole.

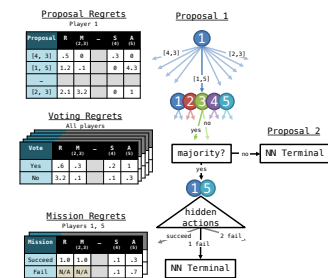
33 *Another interesting observation is the bot does not need conversation. Does this mean the game is not well-designed or*  
34 *a good strategy is to close your eyes and shut your mouth during playing?* We aren’t sure, but this is an interesting  
35 question. Algorithms like DeepRole could be used to probe these ideas and challenge human players to find a distinct  
36 role for language.

37 *Clarify unexplained notation and terminology.* We’ve added the following sentence to  
38 line 110: “...be the joint strategy of all  $p$  players. We write  $\sigma_{I \rightarrow a}$  to mean strategy  $\sigma$ ,  
39 modified so action  $a$  is always played at information set  $I$ .” We’ve added the following to  
40 the end of line 118: “..., where  $z[I]$  is the  $h \in I$  such that  $h \sqsubseteq z$ ”. We have also  
41 added additional text on the "replicator dynamics" used for evaluation: “The replicator  
42 dynamic gradient describes the direction a player playing meta-strategy  $\sigma$  can update  
43 their strategy for maximal gain, assuming other players are also playing  $\sigma$ . Both vector  
44 field sinks and points with zero gradient correspond to nash equilibria in the replicator  
45 dynamic gradient.”

46 *On a higher level, I would liked to see a brief explanation of CFR(+)...* We have  
47 enriched the supplement to more clearly point to the full algorithm and have added  
48 an additional figure which describes the complete architecture with all components in  
49 place.

50 *Provide the raw data for all tables and figures for reproducibility.* This data is included in the open source repository.

51 *Also, Table 1 does not explain what the +/- indicates, nor whether e.g. approximate Gaussianity holds.* We have  
52 changed the figure caption to read: “Confidence interval is the standard error of the mean calculated over a binary  
53 outcome”. This does not require the assumption of Gaussianity.



<sup>1</sup><https://bitbucket.org/Anonymou314159/deeprole/src/master/>