	NW	Error	ND	State	Initial Errors	Final Errors
Error+State-1	1238	14.0%	0.6%	85.4%	202	29
Error+State-3	493	0%	0.2%	99.8%	24	24
Human	2	50%	0%	50%	42	41
Human+Casebase	27	55.6%	3.7%	40.7%	42	27

Table 1: The analysis results and quality of each trace after trace analysis

ing with an internal state or non-deterministic behaviour as error and, when cleaning the noise, actually introduce more error into the case base. Additionally, noise reduction is only one aspect of our analysis approach and reducing error is not the only goal of the analysis.

6 Conclusions

This paper has described an approach to analyze and clean traces of an expert's behaviour. The analysis identifies when a single sensory input can, at different times, result in different actions being performed. The expert is made to replay the original trace in order to generate several new versions of the trace and those traces are used to determine if the expert reasoned with an internal state, performed non-deterministic behaviour or performed any errors. The trace can then be cleaned in order to remove any detected errors.

The major assumptions of this approach are that the expert is available to generate new traces and that the agent is able to present the inputs in a realistic manner. If the expert is not available, this approach can not be used since it relies on the generated traces. If the agent does not present inputs to the expert in a way that is similar to how the environment presents them, the expert may behave differently which can compromise the quality of the generated traces. Both of these issues result in a significant limitation of our approach.

Our experiments demonstrated the applicability of the analysis in an obstacle avoidance domain. The results showed that the analysis was able to correctly detect which of the three properties were present in the trace and cleaning was able to remove many of the errors. All of the threshold values used during analysis (τ, α and β) were selected intuitively so future work will look to examine the effects of changing these values. Also, future work will examine alternative approaches for trace analysis that do not require the generation of new traces but are still able to differentiate between errors, non-determinism and stateful behaviour.

References

Coates, A.; Abbeel, P.; and Ng, A. Y. 2008. Learning for control from multiple demonstrations. In *25th International Conference on Machine Learning*, 144–151.

Cummins, L., and Bridge, D. G. 2011. On dataset complexity for case base maintenance. In *19th International Conference on Case-Based Reasoning*, 47–61.

Delany, S. J., and Cunningham, P. 2004. An analysis of case-base editing in a spam filtering system. In *7th European Conference on Case-Based Reasoning*, 128–141.

Floyd, M. W., and Esfandiari, B. 2009. An active approach to automatic case generation. In *8th International Conference on Case-Based Reasoning*, 150–164.

Floyd, M. W., and Esfandiari, B. 2011. Learning state-based behaviour using temporally related cases. In *16th United Kingdom Workshop on Case-Based Reasoning*, 34–45.

Gillespie, K.; Karneeb, J.; Lee-Urban, S.; and Muñoz-Avila, H. 2010. Imitating inscrutable enemies: Learning from stochastic policy observation, retrieval and reuse. In *18th International Conference on Case-Based Reasoning*, 126–140.

Ontañón, S., and Ram, A. 2011. Case-based reasoning and user-generated AI for real-time strategy games. In *Artificial Intelligence for Computer Games*. Springer-Verlag. 103–124.

Romdhane, H., and Lamontagne, L. 2008. Forgetting reinforced cases. In *9th European Conference on Case-Based Reasoning*, 474–486.

Rubin, J., and Watson, I. 2010. Similarity-based retrieval and solution re-use policies in the game of Texas Hold'em. In *18th International Conference on Case-Based Reasoning*, 465–479.

Tomek, I. 1976. An experiment with the edited nearest-neighbor rule. *IEEE Transactions on Systems, Man, and Cybernetics* 6(6):448–452.

Wooldridge, M. 2002. *An introduction to multiagent systems*. John Wiley and Sons.