The Thirty-First International Florida
Artificial Intelligence Research Society Conference (FLAIRS-31)

Machine Learning from Observation to Detect Abnormal Driving Behavior in Humans

**Josiah Wong,[1] Lauren Hastings,[1] Kevin Negy,[1] Avelino J. Gonzalez,[1]
Santiago Ontañón,[2] Yi-Ching Lee[3]**

[1] Department of Computer Science, Univ. of Central Florida – Orlando, FL, USA
[josiah.w.ucf@knights.ucf.edu, hastingsl@knights.ucf.edu, kevinnegy@gmail.com, gonzalez@ucf.edu]
[2] Drexel University – Philadelphia, PA, USA [so367@drexel.edu]
[3] George Mason University – Fairfax, VA, USA [ylee65@gmu.edu]

Abstract

Detection of abnormal behavior is the catalyst for many applications that seek to react to deviations from behavioral expectations. However, this is often difficult to do when direct communication with the performer is impractical. Therefore, we propose to create models of normal human performance and then compare their performance to a human's actual behavior. Any detected deviations can be then used to determine what condition(s) could possibly be influencing the deviant behavior. We build the models of human behavior through machine learning from observation; more specifically, we employ the Genetic Context Learning algorithm to create models of normal car driving behaviors of different humans with and without ADHD (Attention Deficit Hyperactivity Disorder). We use a car simulator for our studies to eliminate risk to our test subjects and to other drivers. Our results show that different driving situations have varying utility in abnormal behavior detection. Learning from Observation was successful in building models to be applied to abnormal behavior detection.

Introduction

Abnormal behavior is a deviation from a behavior expected from someone under certain assumed conditions. Humans detect it naturally in their daily lives, from being alert for suspicious persons on a dark street to assessing someone's honesty. Abnormal behavior detection has also played a crucial role in computational applications, from detecting deceptive writing styles in text (Afroz, Brennan, and Greenstadt 2012) to identifying impending simulated car crashes (Stanley et. al. 2005). Thus, many abnormal behavior detection applications can benefit from research in informing proper behavioral expectations without being explicitly told what proper behavior is, and in deciding when such expectations have been violated. This can be done with behavior models built by *Learning from Observation*.

Learning from Observation (LfO) is a machine learning approach wherein a computer system learns to emulate an observed performer's behavior through nonintrusive observation alone, without any direct communication with the performer. This approach allows implicitly unique traits of a particular performer's behavior to be captured. We hypothesize that when observing an unknown performer, LfO models can be used to determine the condition (physical, mental, or emotional) most likely affecting the performer, such as whether the performer is behaving normally or under the effect of an impairing influence. The results of this research could be used to detect dangerous impediments in performing humans, such as fatigue when operating complex machinery.

In this paper, we present an approach for detecting abnormal behavior in a simulated car driving domain. We use a modified version of *Genetic Context Learning* (GenCL) (Fernlund et. al. 2006), to learn models from observation of different human drivers. We use our models to discriminate between behaviors by different performers, and by one performer under different influences, specifically medicated and unmedicated Attention Deficit Hyperactivity Disorder (ADHD). By learning the driving traits of each human and each ADHD condition, we seek to detect abnormal driving by indicating when observed behavior doesn't match expectations (as expressed by the model).

This paper is organized as follows. First, we briefly review LfO and GenCL. Then, we describe our driving domain, data set, and approach to creating and validating driving models. Finally, we test our prototype system with the data set compiled and discuss the results of these. Lastly, we present conclusions and future research directions.

Learning from Observation

Learning from Observation is an intuitive approach to machine learning when the objective is learning to perform observable tasks. Fernlund et. al. (2006) define LfO as:

> *"The agent shall adopt the behavior of the observed entity solely from interpretation of data collected by means of observation*."

In LfO, the learning agent learns to emulate a behavior that it observes a human or computer-controlled performer exhibit, even if that behavior is suboptimal (or even incorrect). This is important for capturing behavioral nuances, flaws, and implicit knowledge specific to the performer. We hypothesize that LfO-built models can capture the specific aspects of an individual's performance and differentiate between different observed behaviors.

A full review of the field is outside the scope of this paper (the interested reader is referred to (Ontañón, Montaña, and Gonzalez 2014)); however, a brief discussion is presented here. One of the earliest LfO works was in program synthesis by observations of program instructions (Bauer 1979). LfO was later adapted to learning complex control mechanisms for flying a simulated airplane (Sammut et al. 1992) and driving a car (Sidani and Gonzalez 2000). Other early works include learning to play air hockey (Bentivegna and Atkeson 2001) and how to control a real car's steering by observing the road (Pomerleau 1989).

More recently, LfO has been extended by several researchers. Stensrud and Gonzalez (2008) learned high-level behaviors from observation in maze navigation and in poker. Stein and Gonzalez (2011) bootstrapped agent learning with LfO for tactical tasks such as driving a car and controlling a crane in simulation. Trinh and Gonzalez (2013) identified the contexts experienced by a performer, to be used as input to LfO. Johnson and Gonzalez (2014) learned team behaviors in bucket brigade and ship pursuit domains. Ontañón, Montaña, and Gonzalez (2014) developed a unified LfO framework and defined the types of behaviors that can be learned through LfO. Floyd (2013) developed a general framework for LfO that could learn without knowledge of its task or environment. Much LfO research has focused on creating agents that learn to replace the observed performer, but little work has been done on using LfO agents to detect abnormal behavior, one exception being (Fernlund et. al. 2009) who evaluated trainee behavior with LfO agents modeled after expert tank crews.

Genetic Context Learning

Our chosen LfO algorithm was Genetic Context Learning (GenCL). The interested reader is referred to (Fernlund et. al. 2006) for a detailed description of GenCL, but we briefly describe it here. GenCL uses genetic programming (GP) and Context-Based Reasoning (CxBR) to create tactical agents from observation. In CxBR (Gonzalez, Stensrud, and Barrett 2008), agent control is divided into several *contexts* or situations where the agent exhibits behaviors in conformity to implicit assumptions about its environment. At any time, the agent is controlled by exactly one *active* context. The active context contains *action knowledge*, which controls what the agent does within that context, and *transition knowledge*, which dictates to which other contexts (including a *default* context) the agent could transition when the active context becomes irrelevant to the agent's current situation. The active context can also call a *subcontext* to handle a more specialized situation. In GenCL, the action and transition knowledge are represented by segments of C code. These code segments are generated by GP using evolutionary algorithms (Koza 1992).

GenCL combines CxBR and GP by learning an action rule and transition rule separately for each context that comprises an agent. Segments of a performer's trace where the performer is behaving in a specific context are used in the fitness function. A GP individual is evaluated by running it in a *Micro Simulator*, a graphic-less model of the driving world. For each road segment used for training in a given context, the individual is presented with the same sensory inputs the human perceived at the start of that segment. Then, the individual emits an output to the control mechanism. The Micro Simulator uses this output to update the environment and present the individual with new sensory inputs. This process repeats for each time step in the segment. At regular intervals, the deviation of the individual from the human at a given time step in certain variables (such as speed in a car driving domain) is measured and the individual's final fitness value is the average deviation recorded (lower fitness values are better).

Car Driving Domain

GenCL was originally applied to a simple driving domain in an urban setting with traffic lights and turning at intersections (Fernlund et. al. 2006). In that study, agents observed a human driver's actions in a simulation and learned to control a car's throttle/brake. The sensory input to the learning agents included car speed, distance to nearest intersection or traffic light, traffic light color, and Boolean values for impending intersection turns and traffic lights. Fernlund et. al. (2006) predefined a context hierarchy governing agent behavior; its contexts were Traffic Light Driving (with Green Light and Red Light subcontexts), Intersection Turning, and Urban Driving, the default context.

We expanded the work by Fernlund et. al. (2006) in two ways. First, they omitted rural areas, "hazardous" situations (such as another car cutting in front of the driver, road maintenance, etc.), and other vehicles from their agents' training. In contrast, we include those scenarios. Second, Fernlund et al. (2006) limited the sensory input to five variables, but our sensory input uses 45 variables. This potentially makes our learning task harder because not all variables will be needed for every situation and GenCL may be required to do feature selection in addition to learning how to combine variables to form an agent's procedural knowledge. This may also make our resulting driving models more complex.

The increased complexity of our driving domain required us to modify the context hierarchy in (Fernlund et al. 2006) by adding stop sign and car following contexts (see Figure 1). Other driving scenarios of note include construction zones, hill driving, and curving turns, which are handled by the Default Driving context. For future work, we hope to use an automatic contextualization technique, like that in (Trinh and Gonzalez 2013), to automatically infer a context hierarchy for our driving agents.

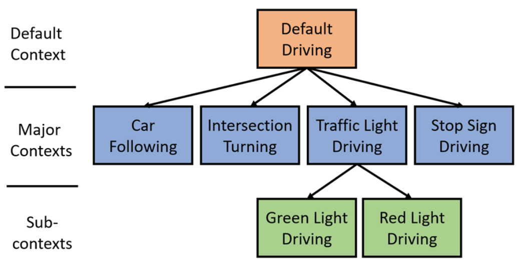

Figure 1: Driving domain CxBR context hierarchy

Data Set

As of the time of this publication, we have collected data from 22 ADHD and 16 non-ADHD teens via a realistic driving simulator. There were five driving scenarios, including a practice drive not used for training or assessment of driving agents (Drive 0). Drives 1 and 2 feature rural driving and scenarios such as curving turns and hills. Drives 3 and 4 feature urban driving and scenarios such as turning at intersections and traffic lights. Figure 2 shows a snapshot of the driving simulator used by the test subjects. These driving scenarios featured situations with high potential for driver error in beginner drivers; more details about the data are provided in (Ontañón et. al. 2017).

Our goal has been to use LfO-built models to detect abnormal behavior in human drivers by training agents that emulate human drivers under different conditions vis-a-vis ADHD. These agents control a car's throttle and brake by treating both as one variable (brake is negative throttle), as was done in (Fernlund et. al. 2006). The angle of the car's steering wheel was not controlled by the agent. Instead, the agent's next steer value was set to that of the human driver when he/she was closest to the agent's location.

Figure 2: A first-person view of the driving simulation the human participants experienced (Drive 1)

We created three driving agents from two humans. *Agent Baseline* (B) was modeled from a non-ADHD teen, who generated four traces, one per drive. *Agent Medicated* (M) was modeled from an ADHD teen after he/she had taken medication to control the effects of ADHD. *Agent Unmedicated* (U) was modeled after the same ADHD teen, but without medication. Thus, the ADHD subject generated two traces per drive, one with and one without medication. As part of our ongoing work, we will in the future train and analyze agents with data from all 38 participants.

The sensory input to our learning agents used 45 variables: 12 internal car variables (e.g., speed, acceleration, heading, XYZ position), 10 driving world variables (e.g., speed limit, distances to stop signs and impending turns, road slope, traffic light color), 20 variables for other cars (ID, XY position, relative distance/speed of closest car in front, behind, left, and right), cumulative distance traveled (2D and 3D), and simulation time. Sensory variable values are persistent; for example, speed limit is recorded for all time steps and not just when the speed limit sign is visible.

For a given agent, we manually selected the parts of the human performer's trace that corresponded to each context in the hierarchy of Figure 1. These segments were used to train the action and transition rules for each context. Segments from Drives 1, 2, and 4 were used for training. Drive 3, the only drive with both intersection turns and traffic lights, was reserved as a validation drive to assess the agent's generalization capabilities, except for three intersection turning segments from Drive 3, because no other drive had intersection turns. In future work, we plan to use an automated contextualization tool, (see (Trinh and Gonzalez 2013)), to assign trace segments to each context.

Assessment

This section describes how we evaluated the goodness of our agents after training and then how we used them to determine whether a driver's behavior was abnormal.

We measured the effectiveness of agent learning through speed RMS error. To see how well a trained agent learned its human counterpart's driving behavior, we simulated the agent over each of its training segments and compared its behavior to the trace used to train it. Then, we computed the speed RMS error (meters/second) for each training segment. Finally, the average of all speed RMS errors of segments for a given context was calculated for each context, both unweighted and weighted by segment time duration. Low speed RMS errors indicate successful learning.

We used our agents to detect abnormal driving in an observed driving segment, given an assumed condition (i.e. Baseline, Unmedicated, or Medicated), by 1) initializing our agents with the sensory input of the human at the start of the segment, then 2) generating a trace for each agent (B, U, and M – described in the "Data Set" section) for the duration of the human's operation in this segment, and finally, 3) labeling the influencing condition of the observed behavior as that of the agent with the lowest speed RMS error when its trace is compared to the observed behavior. Therefore, *abnormal behavior* in the scope of this work is defined as follows: "Given an observed behavior and an assumed condition of influence, the behavior is *abnormal* if it is more similar to the behavior of a driving model of a different condition than it is to that of the model with the assumed condition."

For example, let us initially assume that a driver with ADHD has taken his/her medication. During specific driving situations (e.g. approaching a car from behind), the Unmedicated and Medicated ADHD agents, trained a priori from this driver's behavior while under these conditions, will generate traces indicative of what he/she would do in this situation. After the situation has concluded, these agents' traces will be compared to what the human just did in this situation via the speed RMS error. If the Medicated agent has the lower speed RMS error, then the human's behavior is "normal". Otherwise, the human's behavior is deemed more similar to unmedicated driving, which is indicative of abnormal behavior in violation of our initial assumption of "medicated driving"; that could indicate that either the medication is wearing off or the human forgot to take it before driving. If the same diagnosis is made for several observed segments, then proper action can be taken (e.g., notifying the driver and/or his or her guardian).

Results

Table 1 contains the context-specific speed RMS errors for agents B, U, and M, for all contexts in Figure 1, except subcontexts. Across all agents, the Car Following context had relatively high RMS errors while the Traffic Light and Intersection Turning contexts had relatively low RMS errors. Overall, RMS errors seem low, but accumulating errors over time may hamper agent utility in abnormal behavior detection. Generally, agents adequately learned the high-level speed trends of their human trainer but tended to "smooth" out the low-level irregularities observed in their human trainer's behavior, as seen in Figure 3. In agent control applications, this smoothing out is desirable, but this loss of information may be detrimental to abnormal behavior detection when the abnormal actions are of short duration and smoothed over.

Table 1: Learning Effectiveness Results

Agent	Context	# Segs	Speed RMS Unweighted	Speed RMS Weighted
B	Car Foll.	14	1.655	1.800
	Intersect.	3	0.993	0.943
	StopSign	7	1.589	1.697
	Traffic	7	1.189	1.248
	Default	16	1.440	1.631
U	Car Foll.	15	1.545	1.643
	Intersect.	3	0.754	0.787
	StopSign	7	1.185	1.185
	Traffic	7	0.606	0.627
	Default	15	1.765	1.751
M	Car Foll.	11	1.266	1.303
	Intersect.	3	1.023	1.027
	StopSign	6	1.110	1.200
	Traffic	7	0.694	0.693
	Default	20	0.777	0.830

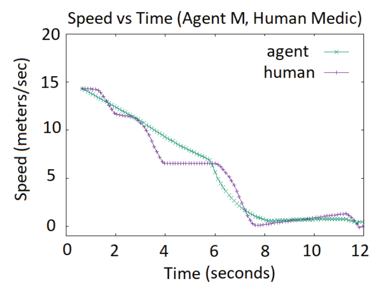

Figure 3: Car Following segment comparing speed over time for medicated ADHD human and agent M

To show how our agents detect abnormal behavior, we used our agents to classify the driving behavior of each training segment as either Baseline, Unmedicated, or Medicated driving. For each segment, the squared speed deviations used to compute the lowest and second lowest speed RMS errors by the agents were compared with a one-tailed t-test (α = 0.05) to determine statistical significance. A correct prediction was a "success" if it had statistical significance; otherwise, it was a "partial success". An incorrect prediction was a "failure", but if the second lowest speed RMS error by the agents was modeled after the actual behavior and the difference in speed deviations had no statistical significance, it was only labeled a "partial failure". Tables 2-4 show success/failure frequencies for Baseline, Unmedicated, and Medicated driving, respectively.

For Baseline driving (Table 2), Intersection Turning had the highest success rate (100%), followed by the Stop Sign context (71%). In contrast, Default Driving had the highest failure rate (19%), followed by Traffic Light (14%). Traffic Light also had the highest combined partial success/failure rate (44%), which is indicative of a high propensity for diagnostic uncertainty if the behavior is Baseline driving. Therefore, Stop Sign and Intersection Turning seem to be the best contexts for detecting Baseline driving.

Table 2: Classified Driving Accuracy, Baseline

Context	Success	Part. Suc.	Part. Fail.	Failure
Car Follow	9 (64%)	4 (29%)	0	1 (7%)
Intersect	3 (100%)	0	0	0
Stop Sign	5 (71%)	0	2 (29%)	0
Traff Light	3 (43%)	2 (29%)	1 (14%)	1 (14%)
Default	10 (63%)	2 (13%)	1 (6%)	3 (19%)
Overall	30 (64%)	8 (17%)	4 (9%)	5 (11%)

Table 3: Classified Driving Accuracy, Unmed. ADHD

Context	Success	Part Suc	Part Fail	Failure
Car Follow	9 (60%)	2 (13%)	2 (13%)	2 (13%)
Intersect	3 (100%)	0	0	0
Stop Sign	3 (43%)	2 (29%)	1 (14%)	1 (14%)
Traff Light	7 (100%)	0	0	0
Default	7 (47%)	0	0	8 (53%)
Overall	29 (62%)	4 (9%)	3 (6%)	11 (23%)

Table 4: Classified Driving Accuracy, Med. ADHD

Context	Success	Part Suc	Part Fail	Failure
Car Follow	6 (55%)	1 (9%)	3 (27%)	1 (9%)
Intersection	3 (100%)	0	0	0
Stop Sign	4 (67%)	1 (17%)	1 (17%)	0
Traff Light	6 (86%)	0	1 (14%)	0
Default	15 (75%)	0	0	5 (25%)
Overall	34 (72%)	2 (4%)	5 (11%)	6 (13%)

For Unmedicated ADHD driving (Table 3), Traffic Light and Intersection Turning had 100% success rates. Default Driving had the highest failure rate (53%), followed distantly by the Stop Sign context (14%), which also had the greatest combined partial rate (43%). Therefore, Intersection Turning and Traffic Light seem to be the best contexts for detecting Unmedicated driving.

For Medicated ADHD driving (Table 4), Intersection Turning again had the highest success rate (100%), followed by Traffic Light (86%) and Default Driving (75%). However, Default Driving also had the highest failure rate (25%), which indicates that Medicated driving is best detected by Traffic Light and Intersection Turning.

Overall, Intersection Turning seems to be the best predictive context across all behaviors. The Traffic Light and Stop Sign contexts also had moderate to high success. In contrast, our agents had greater difficulty properly classifying segments from the Car Following and Default Driving contexts. These two contexts also had the most training segments. This may indicate a need to divide these contexts into finer, more specialized contexts to detect subtler forms of abnormal driving.

To further our understanding of which contexts can assist abnormal behavior detection, we present the confusion matrices for behavior classification in Table 5 for all contexts except Intersection Turning, which had 100% success for all behaviors. Overall, Baseline driving was approximately equally confused with Unmedicated and Medicated driving. However, it is most often confused with Medicated driving for Default Driving segments and with Unmedicated driving for Traffic Light segments. Overall, Unmedicated driving was most often confused with Medicated driving because of two segments in the Stop Sign context; it had equal confusion with Baseline and Medicated driving in all other contexts. Overall, Medicated driving was most often confused with Unmedicated driving; Car Following had the most segments (four) that confused Medicated driving with Unmedicated driving.

Table 5: Context-Specific Confusion Matrices

Context	Actual Behavior	Predicted Behavior		
		Baseline	Unmedic	Medic
Car Following	Baseline	13	1	0
	Unmedic	2	11	2
	Medicat.	0	4	7
Stop Sign	Baseline	5	1	1
	Unmedic	0	5	2
	Medicat.	0	1	5
Traffic Light	Baseline	5	2	0
	Unmedic	0	7	0
	Medicat.	0	1	6
Default Driving	Baseline	12	1	3
	Unmedic	4	7	4
	Medicat.	3	2	15
Overall - All Contexts	Baseline	38	5	4
	Unmedic	6	33	8
	Medicat.	3	8	36

These results align with the intuition that comparing behaviors from different humans (Baseline versus Unmedicated or Medicated) will produce less confusion than comparing behaviors from the same person under different conditions (Unmedicated versus Medicated). In particular, Default Driving had the most confused segments, while Intersection Turning had the least. Some confusion will be inevitable in abnormal behavior detection, particularly in situations where only one reasonable action is possible (e.g. brake hard or crash). Thus, future work will focus on determining how many segments must be observed in which contexts before a confident diagnosis can be made.

Our results show that our chosen contexts have varying utility in abnormal behavior detection. In contexts like Intersection Turning, different behaviors are discernible and abnormal behavior can be detected relatively easily. In contexts like Default Driving, this task is harder and more specialized contextualization may be needed to detect subtler driving abnormalities. In the end, we have shown success in discerning different behaviors in abnormal behavior detection, specifically regarding ADHD, and we have begun to identify in which contexts abnormal driving is likely to occur. This indicates great promise in general detection of abnormal driving behavior in humans from observation of specific driving situations using context-based LfO.

Conclusion

Abnormal behavior detection is important in many computational applications, such as determining an unknown person's identity or whether he or she is under the influence of a specific condition. LfO is an intuitive approach for capturing implicit traits that inform behavioral expectations. We used Genetic Context Learning to train driving agents that emulate observed human drivers' behavior in a simulation. We modeled agents after baseline non-ADHD, medicated ADHD, and unmedicated ADHD drivers, and used them to detect abnormal driving behavior in humans.

Our results show that abnormal behavior detection can be used to distinguish between behaviors by different human drivers and by one human under medicated and unmedicated ADHD in the Intersection Turning context and to a lesser degree in the Traffic Light and Stop Sign contexts. Abnormal driving was relatively harder to detect in the Default and Car Following contexts, which indicates a need for finer contextualization to detect subtler abnormal driving in these situations. Overall, a context-based LfO approach shows great promise in detecting general abnormal driving in humans of varying internal influence.

For future work, we plan to investigate how to use agents enabled with transition knowledge to allow a comparison over an entire trace rather than context-by-context; automatically contextualize human traces for more specialized context hierarchy inputs as done in (Trinh and Gonzalez 2013) as well as to eliminate the need for the extensive manual data preparation required in this work so far; and allow agents to use memory of past events in abnormal behavior detection to better emulate human performance. We also plan to develop metrics that detect abnormal behavior in real-time and to investigate which traits of abnormal behavior are generalizable across different human drivers with the same internal influence, such as ADHD.

Acknowledgements This material is based on work supported by the National Science Foundation under Grant No. 1521972.

References

Afroz, S., Brennan, M., and Greenstadt, R. 2012. Detecting Hoaxes, Frauds, and Deception in Writing Style Online. In *Proceedings of IEEE Symposium on Security and Privacy*, 461–475.

Bauer, M. A. 1979. Programming by Examples. *Artificial Intelligence* 12(1): 1–21.

Bentivegna, D. C., and Atkeson, C. G. 2001. Learning from Observation Using Primitives. In *Proceedings of International Conference on Robotics and Automation (ICRA)*, 1988–1993.

Fernlund, H. K. G., Gonzalez, A. J., Georgiopoulos, M., and DeMara, R. F. 2006. Learning Tactical Human Behavior Through Observation of Human Performance. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 36(1): 128–140.

Fernlund, H., Gonzalez, A. J., Ekblad, J., and Rodriguez, A. 2009. Trainee Evaluation Through After-Action Review by Comparison. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology* 6(3): 135–150.

Floyd, M. W. 2013. A General-Purpose Framework for Learning by Observation. Ph.D. diss., Department of Systems and Computer Engineering, Ottawa-Carleton Institute of Electrical and Computer Engineering, Ottawa, Ontario, Canada.

Gonzalez, A. J., Stensrud, B. S., and Barrett, G. 2008. Formalizing Context-Based Reasoning: A Modeling Paradigm for Representing Tactical Human Behavior. *International Journal of Intelligent Systems* 23: 822–847.

Johnson, C. L., and Gonzalez, A. J. 2014. Learning Collaborative Team Behavior from Observation. *Expert Systems with Applications* 41: 2316–2328.

Koza, J. R. 1992. *Genetic Programming*. Cambridge, Mass.: MIT Press.

Ontañón, S., Lee, Y., Snodgrass, S., Winston, F. K., and Gonzalez, A. J. 2017. Learning to Predict Driver Behavior from Observation. In *Proceedings of the AAAI 2017 Spring Symposium on Learning from Observation of Humans*, 506–512. Palo Alto, Calif.: AAAI Press.

Ontañón, S., Montaña, J. L., and Gonzalez, A. J. 2014. A Dynamic-Bayesian Network Framework for Modeling and Evaluating Learning from Observation. *Expert Systems with Applications* 41: 5212–5226.

Pomerleau, D. 1989. Alvinn: An Autonomous Land Vehicle in a Neural Network. *Advances in Neural Information Processing Systems* 1: 305–313.

Sammut, C., Hurst, S., Kedzier, D., and Michie, D. 1992. Learning to Fly. In *Proceedings of the Ninth International Workshop on Machine Learning*, 385-393.

Sidani, T. A., and Gonzalez, A. J. 2000. A Framework for Learning Implicit Expert Knowledge through Observation. In *Transactions of the Society for Computer Simulation* 17: 54–72.

Stanley, K. O., Kohl, N., Sherony, R., and Miikkulainen, R. 2005. Neuroevolution of an Automobile Crash Warning System. In *Proceedings of the Seventh Conference on Genetic and Evolutionary Computation*, 1977–1984. ACM.

Stein, G., and Gonzalez, A. J. 2011. Building High-Performing Human-Like Tactical Agents through Observation and Experience. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 41: 792–804.

Stensrud, B. S., and Gonzalez, A. J. 2008. Discovery of High-Level Behavior from Observation of Human Performance in a Strategic Game. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 38: 855–874.

Trinh, V. C., and Gonzalez, A. J. 2013. Discovering Contexts from Observed Human Performance. *IEEE Transactions on Human-Machine Systems* 43: 359–370.