Towards Modeling the Behavior of Autonomous Systems and Humans for Trusted Operations

Weiqing Gu,[1] Ranjeev Mittu,[1] Julie Marble,[2] Gavin Taylor,[3] Ciara Sibley,[1] Joseph Coyne,[1] and W.F. Lawless[4]

[1]Naval Research Laboratory, 4555 Overlook Avenue SW, Washington, DC, first.last@nrl.navy.mil
[2]Office of Naval Research, 875 N Randolph St, Arlington, VA, julie.marble@nrl.navy.mil
[3]United States Naval Academy, 121 Blake Rd, Annapolis, MD, taylor@usna.edu
[4]Paine College, 1235 15th Street, Augusta, GA, wlawless@paine.edu

Abstract

Greater unmanned system autonomy will lead to improvements in mission outcomes, survivability and safety. However, an increase in platform autonomy increases system complexity. For example, flexible autonomous platforms deployed in a range of environments place a burden on humans to understand evolving behaviors. More importantly, when problems arise within complex systems, they need to be managed without increasing operator workload. A supervisory control paradigm can reduce workload and allow a single human to manage multiple autonomous platforms. However, this requires consideration of the human as an integrated part of the overall system, not just as a central controller. This paradigm can benefit from novel and intuitive techniques that isolate and predict anomalous situations or state trajectories within complex autonomous systems in terms of mission context to allow efficient management of aberrant behavior. This information will provide the user with improved feedback about system behavior, which will in turn lead to more relevant and effective prescriptions for interaction, particularly during emergency procedures. This, in turn, will enable proper trust calibration. We also argue that by understanding the context of the user's decisions or system's actions (seamless integration of the human), the autonomous platform can provide more appropriate information to the user.

Introduction

Unmanned systems will perform an increasing number of missions in the future, reducing the risk to humans, while increasing their capabilities. The direction for these systems is clear, as a number of Department of Defense roadmaps call for increasing levels of autonomy to invert the current ratio of multiple operators to a single system (U.S. Department of Defense 2011). This shift will require a substantial increase in unmanned system autonomy and will transform the operator's role from actively controlling elements of a single platform to supervising multiple complex autonomous systems. This future vision will also require the autonomous system to monitor the human operator's performance and intentions under different tasking and operational contexts, in order to understand how she is influencing the overall mission performance.

Successful collaboration with autonomy will necessitate that humans properly calibrate their trust and reliance on systems. Correctly determining reliability of a system will be critical in this future vision since automation bias, or overreliance on a system, can lead to complacency which in turn can cause errors of omission and commission (Cummings 2004). On the other hand, miscalibrated alert thresholds and criterion response settings can cause frequent alerts and interruptions (high false alarm rates), which can cause humans to lose trust and underutilize a system (i.e., ignore system alerts) (Parasuraman and Riley 1997). Hence, it is imperative that not only does the human have a model of normal system behavior in different contexts, but that the system has a model of the capabilities and limitations of the human. The autonomy should not only fail transparently so that the human knows when to assist, but autonomy should also predict when the human is likely to fail and be able to provide assistance. The addition of more unmanned assets with multi-mission capabilities will increase operator demands and may challenge the operator's workload just to maintain situation awareness. Autonomy that monitors system (including human) behavior and alerts users to anomalies, however, should decrease the task load on the human and support them in the role of supervisor.

Noninvasive techniques to monitor a supervisor's state and workload (Fong et al. 2011; Sibley, Coyne, and Baldwin 2011) would provide the autonomous systems with information about the user's capabilities and limitations in a given context, which could provide better prescriptions for how to interact with the user. However, many approaches to workload issues have been based on engineering new forms of autonomy assuming that the role of the human will be minimized. For the foreseeable future, however, the human will have at least a supervisory role within the system; rather than minimizing the actions of the human and automating those actions the human can already do well, it would be more efficient to develop a supervisory control paradigm that embraces the human as an agent within the system and leverages on her capabilities and minimizes the impact of her limitations.

In order to best develop techniques for identifying anomalous behaviors associated with the complex human-autonomous system, models of normal behaviors must be developed. For the purpose of this paper, an anomaly is not just a statistical outlier, but rather a deviation that prevents mission goals from being met, dependent on the context. Such system models may be based on, for example, mission outcome measures such as objective measures of successful mission outcomes with the corresponding behaviors of the system. Normalcy models can be used to detect whether events or state variables are anomalous, i.e., probability of a mission outcome measure that does not meet a key performance parameter or other metric.

The anomalous behavior of complex autonomous systems may be composed of internal states and relationships that are defined by platform kinematics, health and status, cyber phenomena and the effects caused by human interaction and control. Once the occurrence and relationships between abnormal behaviors in a given context can be established and predicted, our hypothesis is that the operational bounds of the system can be better understood. This enhanced understanding will provide transparency about the system performance to the user to enable trust to be properly calibrated with the system, making the prescriptions for human interaction that follow to become more relevant and effective during emergency procedures.

A key aspect of using normalcy models for detecting abnormal behaviors is the notion of context; and behaviors should be understood in the context in which they occur. In order to limit the false alarms, effectively integrating context is a critical first step. Normalcy models must be developed for each context of a mission, and used to identify potential deviations to determine whether such deviations are anomalous (i.e., impact mission success). Proper trust calibration would be assisted through the development of technology that provides the user with transparency about system behavior. This technology will provide the user with information about how the system is likely to behave in different contexts and how the user should best respond.

We present an approach for modeling anomalies in complex system behavior; we do not address modeling human limitations and capabilities in this paper, but recognize that this is equally important in the development of trust in collaborative human-automation systems.

Understanding the Value of Context

The role of context is not only important when dealing with the behavior of autonomous systems, but also quite important in other areas of command and control. Today's warfighters operate in a highly dynamic world with a high degree of uncertainty, compounded by competing demands. Timely and effective decision making in this environment is challenging. The phrase "too much data – not enough information" is a common complaint in most Naval operational domains. Finding and integrating decision-relevant information (vice simply data) is difficult. Mission and task context is often absent (at least in computable and accessible forms), or sparsely/poorly represented in most information systems. This limitation requires decision makers to mentally reconstruct or infer contextually relevant information through laborious and error-prone internal processes as they attempt to comprehend and act on data. Furthermore, decision makers may need to multi-task among competing and often conflicting mission objectives, further complicating the management of information and decision making.

Clearly, there is a need for advanced mechanisms for the timely extraction and presentation of data that has value and relevance to decisions for a given context. To put the issue of context in perspective, consider that nearly all national defense missions involve Decision Support Systems (DSS) – systems that aim to decrease the cycle time from the gathering of data to operational decisions. However, the proliferation of sensors and large data sets are overwhelming DSSs, as they lack the tools to efficiently process, store, analyze, and retrieve vast amounts of data. Additionally, these systems are relatively immature in helping users recognize and understand important contextual data or cues.

Context and the Complexity of Anomaly Detection

Understanding anomalous behaviors within the complex human-autonomous system requires an understanding of the context in which the behavior is occurring. Ultimately, when considering complex, autonomous systems comprised of multiple entities, the question is not what is wrong with a single element, but whether that anomaly affects performance of the team and whether it is possible to achieve the mission goals in spite of that problem. For example, platform instability during high winds may be normal, whereas the same degree of instability during calm winds may be abnormal. Furthermore, what may appear as an explainable deviation may actually be a critical problem if that event causes the system to enter future states that prevent the satisfaction of a given objective function. The key distinction is that in certain settings, it may be appropriate to consider anomalies as those situations that effect outcomes, rather than just statistical outliers. In terms of the team, the question becomes which element should have to address the problem (the human or the autonomy).

The ability to identify and monitor anomalies in the complex human-autonomous system is a challenge, particularly as increasing levels of autonomy increase system complexity and, fundamentally, human interactions inject significant complexity via unpredictability into the overall system. Furthermore, anomaly detection within complex autonomous systems cannot ignore the dependencies between communication networks, kinematic behavior, and platform health and status.

Threats from adversaries, the environment, and even benign intent will need to be detected within the communications infrastructure, in order to understand its impact to the broader platform kinematics, health and status. Possible future scenarios might include cyber threats that take control of a platform in order to conduct malicious activity, which may cause unusual behavior in the other dimensions and corresponding states. The dependency on cyber networks means that a network provides unique and complete insight into mission operations. The existence of passive, active, and adversarial activities creates an ecosystem where "normal" or “abnormal” is dynamic, flexible, and evolving. The intrinsic nature of these activities results in challenges to anomaly detection methods that apply signatures or rules that have a high number of false positives. Furthermore, anomaly detection is difficult in large, multi-dimensional datasets and is affected by the “curse of dimensionality.” Compounding this problem is the fact that human operators have limited time to deal with complex (cause and effect) and / or subtle (“slow and low”) anomalies, while monitoring the information from sensors, and concurrently conducting mission planning tasks. The reality is that in future military environments, fewer operators due to reduced manning may make matters worse, particularly if the system is reliant on the human to resolve all anomalies!

Below we describe research efforts underway in the area of anomaly detection via manifolds and reinforcement learning.

Manifolds for Anomaly Detection

A fundamental challenge in anomaly detection is the need for appropriate metrics to distinguish between normal and abnormal behaviors. This is especially true when one deals with nonlinear dynamic systems where the data generated contains highly nonlinear relationships for which Euclidean metrics aren’t appropriate. One approach is to employ a nonlinear “space” called a manifold to capture the data, and then use the natural nonlinear metric on the manifold, in particular the Riemannian metric, to define distances among different behaviors.

We view the path of an unmanned system as a continuous trajectory on the manifold and recognize any deviations due to human inputs, environmental impacts, etc. Mathematically, we transform the different data types into a common manifold-valued data so that comparisons can be made with regard to behaviors.

For example, a manifold for an unmanned system could be 12 dimensional, composed of position, pitch, roll, yaw, velocities of the position coordinates, and angular velocities of the pitch, roll and yaw. This 12- dimensional model captures any platform (in fact any moving rigid object’s) trajectories under all possible environment conditions or behaviors. This manifold is the tangent bundle, TM of $SO(3) \times \Re^3$. Here $SO(3)$ denotes the set of all possible rotations of the unmanned system which is a Lie group, and $\Re^3$ the set of all translations of the platform. Since rotations and translations do not commute, this is not a direct product of $SO(3)$ with $\Re^3$. The product between $SO(3)$ and $\Re^3$ is a “Semi-Product” $\ltimes$. Non-linear key geometric, dynamical and kinematic characteristics are represented using TM. This manifold model is able to encapsulate the unique structure of the environment, effects of human behaviors, etc. through continuous parameterizations and coherent relationships.

Once we have this manifold model and its Riemannian metric, it is possible to define concepts of geodesic neighborhood and other appropriate measurements and map those to mission cost. Such a mapping is done by designing a weighted cost function with dynamical neighborhoods around a trajectory of the platform. For example, if the weather is good in the morning, the neighborhood is smaller than it would be with bad weather. This innovative manifold method could be used to dynamically identify normal or abnormal behaviors occurring during a mission, taking into consideration whether a mission could be successfully achieved under a given cost constraint. We also have the freedom to adjust normal neighborhoods if a mission suddenly changes while en-route. Our model is robust and captures complicated dynamics of unmanned systems and is able to encapsulate very high dimensional data using only a 12 dimensional configuration space.

The algorithms use continuous parameterizations and coherent relationships and are scalable. Our manifold-based methods provide new techniques to combine qualitative (platform mechanics) and quantitative (measured data) methods and are able to handle large, nonlinear dynamic data sets.

Reinforcement Learning for Anomaly Detection

Another angle of approach for the problem is through reinforcement learning. We view the path of the platform as a trajectory through a discrete-time Markov Decision Process (MDP): $M = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$. Given a state $s_i \in \mathcal{S}$, the probability of a transition to a state s_j as a result of action $a \in \mathcal{A}$ is given by $P(s_j|s_i, a)$ and results in an expected reward of $R(s_i)$.

In this application, the state comprises the status of the platform at that moment in time, as collected via telemetry. For example, for a UAV, state may contain factors such as altitude, roll, pitch, yaw, ground velocity, wind speed, warning light status, fuel status, etc. The reward function may consider physical damage to the platform, stress on the control surfaces, or simply maintenance costs from another second of operation; these costs would be represented as negative rewards, while success in reaching a waypoint would receive a positive reward.

The optimal value function is the expected, discounted sum of future costs given the current state of the platform and an optimal pilot. Mathematically, this is represented with the Bellman equation,

$$V^*(s) = R(s) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, \pi^*(s))V^*(s'),$$

where s and s' are states in the state space $\mathcal{S}$ and π^* is the optimal policy, which maps states to the action maximizing the expected sum of future rewards.

For large or continuous state spaces, such as the ones of interest to us, it is impossible to solve for the optimal value function exactly, necessitating some form of approximation. There are a wide range of approaches, but the problem is well suited to L_1-Regularized Approximate Linear Programming (RALP) (Petrik et al. 2010; Taylor and Parr 2012). Given a feature matrix Φ of size $n \times k$, where n is the number of sampled states, and k is the number of features, RALP produces a linear approximation of the optimal value function; this is to say, RALP provides a weight vector w such that $\Phi w \approx V^*$.

RALP is based on preceding linear programming approaches to value function approximation presented by d'Epenoux (1963), Schweitzer and Seidmann (1985), and de Farias and Van Roy (2003). However, RALP adds L_1 regularization. Therefore, given a set of samples Σ where each sample consists of a state s, a reward received r, and a next state s', the linear program solved by RALP is as follows:

$$\begin{aligned} \min_{w} \quad & \rho^T \Phi w \\ s.t. \quad & r + \gamma \Phi(s') w \leq \Phi(s) w \;\; \forall (s, r, s') \in \Sigma \\ & \|w\|_{1,e} \leq \psi. \end{aligned}$$

ρ is a distribution over states, $\|w\|_{1,e} = \sum_i |e(i)w(i)|$, where $e(i)$ is 0 if i corresponds to the bias feature, and 0 otherwise, and ψ is the tunable regularization parameter.

RALP is unique in several useful ways. First, RALP approximates the optimal value function V^*, not the value function of a given intermediate sub-optimal policy. Second, it can do this with samples drawn from multiple sampling policies, which is convenient when collecting telemetry data from a range of missions. Third, the error in the value function is bounded, and has been experimentally demonstrated to be close to the true optimal value function. Finally, due to the well-known sparsity effect of L_1 regularization (see, for example, Tibshirani, 1996), nearly all elements of the vector w will be zero. This means RALP performs automated feature selection, and can do so from an extremely large, over-complete feature set. Only the most important features for value function approximation given the data set receive a non-zero weight, providing an objective way of identifying features that are most helpful for predicting the future success or failure of a mission. As a value function approximation, this is useful in its own right, but this feature set may be particularly powerful as an informative complement to other approaches mentioned in this paper.

In addition, an accurate value function approximation may prove to be a helpful tool in identifying anomalies. Sudden large-magnitude drops in value may indicate the occurrence of an adverse event, whether it be due to mechanical failure, worsening environmental factors, or pilot error. Additionally, because the value function predicts expected rewards, a sequence of states with values which underperform expectations may indicate a "slow and low" anomaly, such as a consistent headwind, which may endanger timely mission success.

Predictive and Prescriptive Analytics

It is clear the DoD and the U.S. Navy are increasingly reliant on autonomy, machines and robotics whose behavior is increasing in complexity. Most research indicates operational improvements with autonomy, but autonomy may introduce errors (Manzey, Reichenbach, and Onnasch 2012) that impact performance. These errors result from many factors, including faulty design assumptions especially in data fusion aids, stochasticity with sensor/observational data, and the quality of the information sources fed into fusion algorithms. Furthermore, additional factors may include greater sophistication and complexity and the subsequent inability of humans to fully comprehend the reasons for decisions made by the automated system (i.e., a lack of transparency). In spite of these known faults, some users rely on autonomy more than is appropriate, known as autonomy "misuse" (Parasuraman and Riley 1997). Another bias associated with human-automation collaboration is disuse, where users underutilize autonomy to the detriment of task performance.

We conjecture that in order to help overcome issues associated with misuse/disuse, the next generation of integrated human-autonomous systems must build upon the descriptive and predictive analytics paradigm of understanding and predicting, with a certain degree of confidence, what the complex autonomous system has done and what it will do next based on what is considered normal for that system in a given context. Once this is achievable, it will enable the development of models that proactively recommend what the user should do in response in order to achieve a prescriptive model for user interactions (Figure 1).

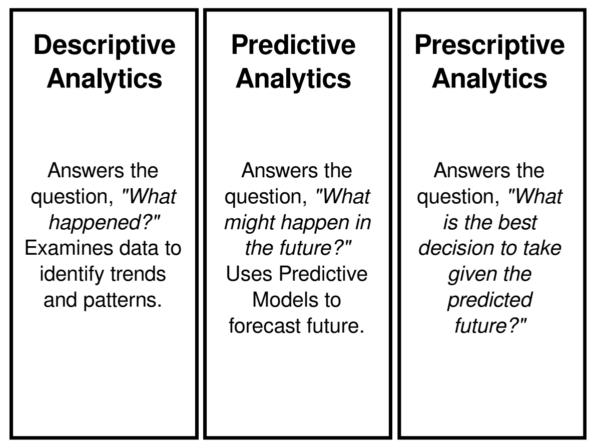

Figure 1: Different forms of Analytics

By properly presenting current and future system functioning to the user, and capturing user interactions in response to such states, we believe more effective human-automation collaboration and trust calibration can be established. The key question is "how are the best user interactions captured?"

Capturing User Interactions and Inference

Transparency in how a system behaves should enable the user to calibrate their level of trust in the system. However, there are still significant challenges that remain with regard to capturing and understanding the human dimensions of supervisory control in order to provide prescriptions for interaction. We envision several longer term challenges related to the notion of prescriptive analytics, specifically how best to understand and model the information interaction behaviors of the user. These information seeking behaviors may be in reference to the potential anomalies in the system, in relation to what is provided by the on-board sensors, etc. and may require the development of the following capabilities:

- Adequately capturing users' information interaction patterns (and subsequently user information biases)

- Reasoning about information interaction patterns in order to infer decision making context; for example, the work being done by researchers within the Contextualized Attention Metadata community and the Universal Interaction Context Ontology (Rath, Devaurs, and Lindstaedt 2009) might serve as a foundation
- Instantiating formal models of decision making based on information interaction behaviors (potentially using cognitive architectures)
- Leveraging research from the AI community in plan recognition to infer which decision context (model) is active, and which decision model should be active
- Recognizing decision shift based on work that has been done in the Machine Learning community with "concept drift," and assessing how well this approach adapts to noisy data and learns over time
- Incorporating uncertainty and confidence metrics when fusing information and estimating information value in relation to decision utility
- Using models of cognition and decision making (and task performance) to drive behavior development and interface development

Lastly, research is needed to address how the autonomous platform should adapt to user behaviors in order to balance both mission requirements as well as servicing the needs of the human supervisors.

Challenges and Opportunities

Elaborating on our ideas, longer-term research should be focused on the following: decision models for goal-directed behavior, information extraction and valuation, decision assessment and human systems integration.

With regard to decision models for goal-directed behavior, the key research question may include how to instantiate prescriptive models for decision making, which integrate information recommendation engines that are context-aware. Furthermore, what are the best techniques that can broker across, generalize, or aggregate individual decision models in order to enable application in broader mission contexts? Supporting areas of research may include the development of similarity metrics that enable the selection of the appropriate decision model for a given situation, and intuitive decision model visualizations.

The notion of information extraction and valuation would involve locating, assessing, and enabling, through utility-based exploitation, the integration of high-value information within decision models, particularly in the big data realm. This is a particular research challenge due to heterogeneous data environments when dealing with unmanned systems. In addition, techniques that can effectively stage relevant information along the decision trajectory (while representing, reducing and/or conveying information uncertainty) would enable a wealth of organic data to be maximally harvested.

In reference to decision assessment, research needs to address what are the most effective techniques for modeling decision "normalcy," in order to identify decision trajectories that might be considered outliers and detrimental to achieving successful outcomes in a given mission context. Furthermore, techniques that proactively induce the correct decision trajectory to achieve mission success are also necessary. Metrics for quantifying decision normalcy in a given context can be used to propose alternate sequences of decisions or induce the exact sequence of decisions. This would require pre-staging the appropriate information needed to support the evaluation of decisions, potentially improving the speed and accuracy of decision making.

Lastly, with regard to human systems integration, the key challenges are in understanding, modeling and integrating the human state (workload, fatigue, experience) as well as the human decision making component as an integral part of the aforementioned areas. Specific topics include: representing human decision-making behavior computationally; accounting for individual differences in ability and preferences; assessing human state and performance in real-time (during a mission) in order to facilitate adaptive automation; mathematically capturing the human assessment of information value, risk, uncertainty, prioritization, projection and insight; and computationally representing human foresight and intent.

Summary

The development of robust, resilient, and intelligent systems requires the calibration of trust by humans when working with autonomous platforms. We contend that this can be enabled through a capability which allows system operators to understand anomalous states within the system of systems, which may lead to failures and hence impact system reliability. Likewise, the autonomy should understand the decision making capabilities and other limitations of the humans in order to proactively provide the most relevant information given the user's task or mission context.

This position paper has discussed the need for anomaly detection in complex systems promote a human supervisor's understanding of system reliability. This is a challenging problem due to the increasing sophistication and growing number of sensor feeds in such systems which creates challenges for conducting big data analytics. Technical approaches that enable dimensionality reduction and feature selection should improve anomaly detection capabilities. Furthermore, building models that account for the context of each situation should improve the understanding of what is considered an anomaly. Additionally, we argue that anomalies are more than just statistical outliers, but should also be based upon whether they hinder the ability of the system to achieve some target end state. Understanding anomalies, we believe, should inform, and make more effective, the user's interaction with system. The interaction may include learning more about the anomaly through some form of query, command and control of the situation, entering into some emergency control procedure, etc.

Numerous research questions remain about the most effective interactions between human and autonomy. We believe the following research areas require further exploration in order to build more robust and intelligent systems. First, researchers should seek to capture users' interaction patterns (and subsequently user information biases) and reasoning about interaction patterns in order to infer decision making context. The work being done by researchers within the Contextualized Attention Metadata community and the Universal Interaction Context Ontology (Rath, Devaurs, and Lindstaedt 2009) might serve as a foundation for this approach. Second, instantiating formal models of human decision making based on interaction behaviors would lead to autonomous recognition of human capabilities and habits. Third, leveraging research from the AI community in plan recognition would allow for the inference of active decision contexts (model), and decision model selection. Fourth, adapting work that has been done in the Machine Learning community with concept drift, to recognize decision shifts and assess how well this approach adapts to noisy data. Finally, it is necessary to incorporate uncertainty and confidence metrics when fusing information and estimating information value in relation to decision utility.

In order to build trusted systems which include a human component performing supervisory control functions, it is vital to understand the behaviors of the autonomy as well as the human (and his/her interaction with the autonomy). This should provide a holistic approach to building effective collaborative human-automation systems, which can operate with some level of expectation and predictability.

Acknowledgements

The collaboration between the Naval Research Lab and the US Naval Academy was supported by the Office of Naval Research, grant number N001613WX20992.

References

Cummings, M. 2004. Automation Bias in Intelligent Time Critical Decision Support Systems. In *AIAA 3rd Intelligent Systems Conference*.

de Farias, D. P., and Van Roy, B. 2003. The Linear Programming Approach to Approximate Dynamic Programming. *Operations Research*.

d'Epenoux, F. 1963. A Probabilistic Production and Inventory Problem. *Management Science*.

Fong, A.; Sibley, C.; Coyne, J.; and Baldwin, C. 2011. Method for Characterizing and Identifying Task Evoked Pupillary Responses During Varying Workload Levels. In *Proceedings of the Annual Meeting of the Human Factors and Ergonomics Society*.

Manzey, D.; Reichenbach, J.; and Onnasch, L. 2012. Human Performance Consequences of Automated Decision Aids. *Journal of Cognitive Engineering and Decision Making* 6:57–87.

Parasuraman, R., and Riley, V. 1997. Humans and Automation: Use, Misuse, Disuse, Abuse. *Human Factors* 39:230–253.

Petrik, M.; Taylor, G.; Parr, R.; and Zilberstein, S. 2010. Feature Selection Using Regularization in Approximate Linear Programs for Markov Decision Processes. In *Proceedings of the 27th International Conference on Machine Learning*.

Rath, A.; Devaurs, D.; and Lindstaedt, S. 2009. UICO: an Ontology-Based User Interaction Context Model for Automatic Task Detection on the Computer Desktop. In *CIAO 09: Proceedings of the 1st Workshop on Context, Information and Ontologies*.

Schweitzer, P. J., and Seidmann, A. 1985. Generalized Polynomial Approximations in Markovian Decision Processes. *Journal of mathematical analysis and applications* 110(6):568–582.

Sibley, C.; Coyne, J.; and Baldwin, C. 2011. Pupil Dilation as an Index of Learning. In *Proceedings of the Annual Meeting of the Human Factors and Ergonomics Society*.

Taylor, G., and Parr, R. 2012. Value Function Approximation in Noisy Environments Using Locally Smoothed Regularized Approximate Linear Programs. In *Conference on Uncertainty in Artificial Intelligence*.

Tibshirani, R. 1996. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*.

U.S. Department of Defense. 2011. Unmanned Systems Integrated Roadmap FY2011-2036.