

# Human Supervisory Control of Robotic Teams

## *Integrating Cognitive Modeling with Engineering Design*

Jeffrey R. Peters<sup>1</sup>, Vaibhav Srivastava<sup>2</sup>, Grant S. Taylor<sup>3</sup>, Amit Surana<sup>4</sup>, Miguel P. Eckstein<sup>5</sup>, and Francesco Bullo<sup>1</sup>

<sup>1</sup>Center for Control, Dynamical Systems, and Computation, UCSB

<sup>2</sup>Department of Mechanical and Aerospace Engineering, Princeton University

<sup>3</sup>US Army Aeroflightdynamics Directorate, Moffett Field

<sup>4</sup>United Technologies Research Center, Hartford

<sup>5</sup>Department of Psychology and Brain Sciences, UCSB

The emergence of sensor networks operating at different modalities, mobility, and coverage has opened the door to systems involving diverse data sources and analysis tools. These complex systems often contain both human and robotic elements, and, in many cases, it is the job of humans to process information generated by autonomous agents [1], [2]. The incredible amount of data generated by modern sensors makes these human operators susceptible to *information overload*, which can have detrimental effects on performance and may lead to dire consequences [3]. To alleviate this loss in performance, programs like the recent National Robotic Initiative [4] emphasize collaboration between humans and their robotic partners, and envision symbiotic mechanisms to facilitate interactions between diverse system components.

This article focuses on the design of systems in which a human operator is responsible for overseeing autonomous agents and providing feedback based on sensor data. In the control systems community, the term *human supervisory control* (or simply *supervisory control*) is often used as a shorthand reference for systems with this type of architecture [5]–[7]. In a typical human supervisory control application, the operator does not directly manipulate autonomous agents, but rather indirectly interacts with these components via a central data-processing station (see Figure 1). As such, system designers have the opportunity to easily incorporate automated functionalities to control how information is presented to the operator, and how the input provided by the operator is used by automated systems. The goal of these functionalities is to take advantage of the inherent robustness and adaptability of human operators, while mitigating adverse effects such as unpredictability and performance variability. In some contexts, to meet the goal of single-operator supervision of multiple automated sensor systems, such facilitating mechanisms are not only useful, but




Figure 1: A typical human supervisory control setup consisting of 3 main components: the human operator(s), the data-processing station, and the autonomous agents. Human operator(s) interact with the autonomous agents through the data-processing station. The degree to which human performance and input affects automation, as well as the method by which sensor data is presented to the operators is determined by the data-processing station and its internal functionalities.

necessary for practical use [8], [9].

A successful system design must carefully consider the goals of each part of the system as a whole, and seamlessly stitch components together using facilitating functionalities.

### Design Considerations

The design of any effective supervisory control system starts with a model of human cognitive processing [10]. This model, which forms the “backbone” of the human-centered system, must capture the operator’s underlying decision-making mechanisms, while still taking into account the variability that is inherent to human processing. Other factors, such as mental workload, memory, and fatigue, can significantly affect these driving mechanisms as well, and may also need to be incorporated into the model to achieve design goals.

Once an appropriate model has been constructed, the question becomes how to use the information that the model provides to manage data presentation and automated control schemes. For example, data collected by autonomous agents in supervisory

control applications is often of a visual nature, that is, photos or video. Given such visual imagery, can operator performance, imagery characteristics, and system parameters be used to decide which region of the image the operator should focus on? If multiple images are waiting to be processed, is it possible to determine how much time the operator should spend on each image? Can system parameters be adjusted to react to non-optimal user behaviors in real time? How should the autonomous agents take human responses into account?

It is apparent that an effective system design incorporates a broad range of theoretical and practical tools from many scientific disciplines, including control systems, human factors, and psychology. As such, practitioners face a series of diverse and complex choices when deriving models and strategies to govern system behavior. The goal for this article is to provide insight into some of these choices through examination of common theoretical tools relevant to each of the main components making up a supervisory control system; namely, the human operator, the autonomous agents, and the interface between them. In particular, the discussion is focused on those tools that have close ties to control and dynamical systems. Throughout the discussion, key challenges that arise both in practical implementation and in combining these tools for use in the overall system are highlighted. In some sense, this article can be thought of as a brief survey of work relevant to the design of human supervisory control systems; however, the article also serves to provide a proof of concept, by illustrating how basic, well-studied theory from various disciplines can work together for use in a broader, human-centered systems perspective.

The following discussion is not, by any means, intended to provide an exhaustive review of all relevant theory, but rather is intended to give control practitioners a flavor for the types of models that are being used and the unique issues that can arise in this type of application.

## State of the Art

*Automation* can be formally defined as the “execution by a machine agent of a function that was previously carried out by a human” [11]. In this broad context, the use of human operators to monitor the functionality of automated systems has arisen in widespread domains. Examples of current applications that incorporate human-centered automation systems include dynamic positioning systems in maritime applications [12], command and control systems for monitoring satellites and space assets [13], automated vehicle operation aids [14], aviation accident and emergency response systems [15], numerous military operations [16], [17], medical imaging systems [18], advanced traffic management and intelligent transportation systems [19], and many more.

As a consequence of this growing interest in human supervisory control, a large body of research has focused on the direct incorporation of human performance models into autonomous system design. Significant research efforts have gone into finding systematic ways of distributing operator cognitive resources. In some approaches, the human decision-making process is unregulated, but the automated system is tailored to the human operator’s cognitive requirements. The fundamental research questions under this approach include (i) optimal scheduling of the tasks to be processed by the operator [20]–[26]; (ii) enabling shorter operator reaction times by controlling the fraction of the total time during which the operator is busy [27], [28]; and (iii) efficient work-shift design to counter fatigue or interruption effects [29]. In other approaches, both the operator’s decision-making process and the autonomous agents are controlled. For example, the human operator is given a set time to spend on each task, and the operator’s decision is used to adaptively adjust overall automation schemes or parameters. The fundamental research questions under this approach include (i) determining optimal operator attention allocation both within and across tasks [30]–[32]; (ii) managing operator workload to enable better performance [33]; and (iii) controlling autonomous agents to collect the most useful information [33]–[35].

Many researchers have also studied adaptive strategies to human-centered system design, in which both physiological and performance measures are used to infer the operator’s cognitive state (such as mental workload and operator intentions), and automated functionalities are only triggered when a nonoptimal or undesirable state is detected [36], [37]. However, the majority of such adaptive systems to date have been experimental rather than practical due to difficulties in constructing accurate indicators of the user’s cognitive state [38]. Despite such difficulties, continually improving accuracy and affordability of physiological sensors, such as eye trackers and electroencephalography (EEG) devices, have led to a better understanding of objective measures that can give insight into operator cognitive behavior [39].

The remainder of this article discusses the three main components of a human supervisory control system as defined in Figure 1. Due to the vast amount of literature and theory that is available on each of these topics, an exhaustive survey is impractical for a single article. Therefore, the goal of providing an illustrative orientation to human supervisory control is accomplished by focusing the discussion on a subset of the available literature which (i) is representative of the state of the art approaches to human supervisory control system design, (ii) is accessible to readers unfamiliar with human-centered systems, (iii) is amenable to automated decision support and other facilitating functionality design, and (iv) effectively illustrates how familiar control theoretic tools can be used in this setting. Most of the following discussion is motivated by the control of mobile sensors that collect visual data (such as unmanned vehicles taking photos or video), although many of the concepts discussed readily extend to

## EXAMPLE: PERSISTENT SURVEILLANCE MISSION

### Problem Overview and Setup

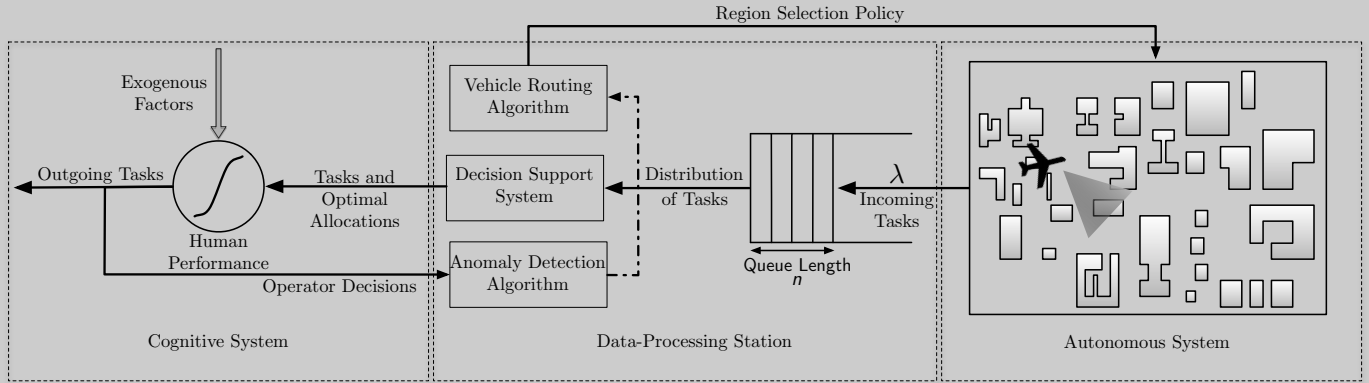


Figure S1: Setup for the example persistent surveillance mission.

### Overview

To further illustrate the design principles discussed, the sidebars in this article present an example human supervisory control problem that involves continuous search of target regions by a mobile sensor. This type of persistent surveillance using mobile sensors is applicable to a variety of real scenarios, including military applications, such as area reconnaissance and battlefield damage assessment, search and rescue operations, such as disaster assistance and target extraction, and environmental monitoring tasks, such as the control of forest fires and wildlife regulation.

An efficient persistent surveillance policy can have multiple objectives, including minimization of the time between subsequent visits to a region and minimization of the delay in detecting anomalous events, such as the appearance of an intruder or the onset of a fire. The fundamental tradeoff in persistent surveillance is between the amount of evidence collected from the visited region and the resulting delay in evidence collection from other regions. In this example, the objective is to address this tradeoff by designing an efficient surveillance policy that takes into account human responses to image analysis tasks, and subsequently collects evidence from regions that are highly likely to be anomalous. Human decisions regarding the collected evidence are considered in conjunction with a cognitive model to determine the likelihood of a region being anomalous. Finally, the integration of these tools is illustrated through the design of a simple decision support that determines how the operator should allocate time to multiple image-processing tasks.

### Setup

The primary objective in this example surveillance mission is to detect, within a prescribed accuracy, any anomaly in a discrete set of regions.

The mission setup is shown in Figure S1 and consists of three main components, consistent with the abstraction in Figure 1: (i) the autonomous system, (ii) the cognitive system, and (iii) the data-processing station.

The autonomous system is a single unmanned aerial vehicle (UAV) that surveys a set of regions according to a routing policy. The UAV is equipped with a camera, and during each visit to a region the UAV generates an image. The image is sent to the data-processing station, which, in turn, sends the image to the human operator (cognitive system).

The cognitive system is a single human operator who examines the image and decides whether an anomaly is present or absent in the associated region.

In this example, the data-processing station consists of three elements: (i) the decision support system, (ii) the anomaly detection algorithm, and (iii) the vehicle routing algorithm. The purpose of the decision support system is to use the performance of the operator to suggest the optimal amount of time that the operator should allocate to each perceptual task, that is, each image generated by the UAV. Decisions made by the human operator may be erroneous, and thus the anomaly detection algorithm is a sequential statistical algorithm that treats the operator's decision as a binary random variable and ascertains the desired accuracy of the anomaly detection. The anomaly detection algorithm also provides the likelihood of an anomaly at each region. The vehicle routing algorithm uses the likelihood of each region being anomalous to determine an efficient vehicle routing policy.

The goal of the overall system is to detect anomalies in the shortest time interval possible, subject to a false-alarm constraint. In subsequent sidebars, each problem component is examined in detail.

other related domains.

## HUMAN MODELING

### Approaches to Modeling Human Cognition

This section provides an overview of a few common strategies used to capture human cognitive behavior in tasks where a

person must choose between a set of alternative choices.

At a high level, the issue of producing a meaningful interpretation of human behavior that can be used in prediction and system design is a type of “black box” problem, similar to those encountered in the modeling of uncertain dynamical systems. That is, the aim is to model what occurs in a system whose precise internal behavior is unknown, given only information about inputs and outputs. If unknown system parameters can be estimated and output quantities can be isolated, then simple data fitting techniques can be used to produce a functional relationship between the nature of a stimulus, the state of the organism in question, and the output quantity of interest. This relationship can subsequently be used to predict future system behavior. Such fitting techniques have been studied and employed in psychological contexts for a number of years [40], [41]. These approaches are usually simple and straightforward to implement, making them an attractive option for applications that only require modeling on a coarse scale.

Black box approaches, however, are often not sufficient for capturing the relationship between physiological phenomena that occur on a fine scale, such as neuronal activity, and resulting behavior. In applications where these relationships are of particular importance, alternative psychological models that seek to explain cognitive behavior through direct links with detailed anatomy and physiology of the human’s contributive systems, such as the nervous and endocrine systems, may be more appropriate [42]. Many such models try to capture the decision process through the inherent dynamics of interconnected neurons (see, for example [43]).

At a slightly coarser level, some constructs, such as artificial neural networks, seek to explain behavior through massive parallel models that are composed of large numbers of simple and uniform interconnected processing elements. These constructs are called *connectionist models* or parallel distributed processing (PDP) networks [44], [45]. Another alternative modeling paradigm is the *symbolic* approach to cognition, which is inspired by logic and digital computing techniques, and sees reasoning as a process resulting from the structured manipulation of symbolic representations. The symbolic and connectionist approaches are complementary in the sense that the former is quite efficient at modeling knowledge representations, while the latter is more focused on capturing the learning process. This disparity has led to development of hybrid connectionist-symbolic models [46]. Examples of well-known cognitive architectures that fall into this category include ACT-R [47], Soar [48], EPIC [49], and CLARION [50].

Cognitive architectures that adopt intricate connectionist and/or symbolic components are usually very general, and have the ability to capture a wide variety of complex behavior [51]. However, these architectures are primarily used in

modeling sensory evidence retrieval and storage, rather than the dynamics of the evidence itself. As a result, cognitive architectures alone may lack information required to model low-level decision making behavior associated with particular tasks [52]. In the context of perceptual choice tasks, a variety of models can be used to more explicitly account for how behavioral performance improves over time as a result of the accumulation of sensory information. These models, called *accumulator* or *sequential sampling models*, provide significant insight into sensory evidence accumulation and behavioral correlates of decision-making [52]. Accumulator models have been used to predict human accuracy and reaction times, applied to allow interpretations on practical problems such as the effect of aging on performance, and also integrated with neurophysiological data to provide a framework to connect neuronal and behavioral measures [53]. Attempts to combine these dynamic approaches with cognitive architectures and develop a unified theory has been a subject of recent research [52].

In general, no single approach to modeling human cognition will be sufficient for all possible applications, and the plethora of models that have been developed over the past century all have merit in some domains. At a given level of granularity with respect to descriptions of stimuli and behavior, there will exist some model that describes relevant data as well as possible and more economically than is feasible at other levels [42]. For the purposes of designing automated systems to aid operators in human supervisory control, black box statistical models are often not detailed enough for the development of on-line control schemes, while intricate physiological models are usually too detailed or evolve at time scales that are too accelerated to be useful in macro-scale operations. Further, in supervisory applications, the dynamics of decision making are often of more immediate interest than the physiological mechanisms that drive sensory evidence retrieval and storage.

For these reasons, the remainder of the discussion on cognitive modeling is focused on accumulator models that seek to capture the sensory evidence accumulation process in forced choice tasks. In addition to the reasons already mentioned, accumulator models are relevant because they (i) have close ties to dynamical systems, (ii) are widely used in cognitive psychology, (iii) are relevant to visual perception which is a commonly encountered task in supervisory control, (iv) are abstractions of detailed physiological models [54], (v) have been proven to capture a large amount of relevant behavioral phenomena [55]–[57], and (vi) appropriately illustrate key challenges involved in modeling human behavior. This, of course, does not imply that other modeling techniques are not pertinent to control; only that accumulator models are powerful tools that are reflective of state of the art approaches to supervisory control, are conducive to the exploratory nature of this article, and give the unfamiliar reader a flavor of issues that arise in cognitive modeling.

## Two-alternative Forced Choice Tasks

A *two-alternative forced choice task* is one in which an operator must decide between two possible hypotheses. Models for two-alternative forced choice tasks within continuous sensory information acquisition scenarios rely on two assumptions: evidence is collected over time in favor of each alternative and a decision is made once a stopping criterion is met. A few simplistic models, such as the linear ballistic accumulator model, assume that evidence toward each alternative evolves in a linear and deterministic manner toward a decision threshold [58]. Such simplistic models allow for analytic solutions, which can be analyzed to infer changes in drift-rates or decision thresholds. Deterministic models are usually insufficient to adequately capture the complex nature of human cognition, and thus virtually all other accumulator models assume that the sensory evidence accumulation process has an element of randomness. In this stochastic context, several models for two-alternative forced choice tasks have been proposed [55]; however, almost all accumulator models are based on the drift diffusion model (DDM) [59]–[61]. The DDM is popular because: (i) it is simple and well characterized; (ii) it captures a significant amount of behavioral and neuro-scientific data; and (iii) many other models for two-alternative forced choice tasks reduce to the DDM under optimal parameter choices [55].

In the most basic version of the DDM, evidence toward an alternative is modeled as a variable  $x \in \mathbb{R}$  that evolves according to the stochastic differential equation

$$dx(t) = \mu dt + \sigma dW(t), \quad x(0) = x_0, \quad (1)$$

where  $\mu \in \mathbb{R}$  is the *drift rate*,  $\sigma \in \mathbb{R}_{>0}$  is the *diffusion rate*,  $W(\cdot)$  is the standard Wiener process, and  $x_0 \in \mathbb{R}$  is the initial evidence. For an unbiased operator, the initial evidence  $x_0 = 0$ , while for a biased operator  $x_0$  captures the odds or the prior probability of each hypothesis being true.

For the information aggregation model (1), human decision making is studied in two paradigms, namely, *free response* and *interrogation* [55] (see Figure 2). In the free response paradigm, the operator waits to make a decision until the evidence satisfies a pre-established criterion, while in the interrogation paradigm, the operator must make a decision within a pre-established time window. The free response paradigm is modeled via two thresholds (positive and negative) and the operator decides in favor of the first (second) alternative if the positive (negative) threshold is crossed from below (above). In contrast, the interrogation paradigm makes use of a single threshold, and the operator decides in favor of the first (second) alternative if the amount of accumulated evidence is above (below) the threshold at the end of the allotted time.

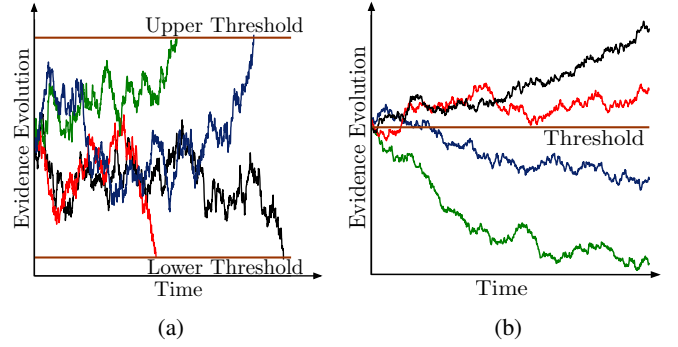


Figure 2: Curves illustrating decision-making paradigms for the DDM. Each colored line represents an independent decision-making task. (a) Free response paradigm for decision-making. Evidence evolves according to (1) and the operator makes a decision once a threshold is crossed. (b) Interrogation paradigm of decision-making. Evidence evolves according to (1) and the operator's decision depends on whether the evidence is above or below a threshold after a given amount of time.

### Free Response Paradigm

Typical evolutions of the DDM under the free response paradigm are shown in Figure 2a. For equally likely alternatives, the two decision thresholds are chosen symmetrically. If  $\pm\eta \in \mathbb{R}$  represents symmetrically chosen thresholds, the expected decision time ( $T_{\text{Decision}}$ ) under the free response paradigm is

$$T_{\text{Decision}} = \frac{\eta}{\mu} \tanh \frac{\mu\eta}{\sigma^2} + \frac{2\eta(1 - e^{-2x_0\mu/\sigma^2})}{\mu(e^{2\eta\mu/\sigma^2} - e^{-2\eta\mu/\sigma^2})} - \frac{x_0}{\mu}.$$

Reaction time on a task is  $T_{\text{Decision}} + T_{\text{SM}}$ , where  $T_{\text{SM}} \in \mathbb{R}_{>0}$  is the time taken by sensory and motor processes unrelated to the decision process. With proper choice of parameters, the DDM (1) can predict reaction times with some success (see Figure 3).




Figure 3: (a) Empirical reaction time data taken from [23]. (b) Decision times under the free response paradigm, as predicted by an appropriately chosen DDM (1).

The choice of the threshold is dictated by a tradeoff between

speed and accuracy. The two most common criteria to capture the speed-accuracy trade-off are (i) Bayes' risk (BR) and (ii) reward rate (RR) [55]. Bayes' risk is defined by  $BR = \xi_1 T_{\text{Decision}} + \xi_2 \mathbb{P}_{\text{Error}}$ , where  $\xi_1, \xi_2 \in \mathbb{R}$  are the cost per unit delay in decision and the cost of error, respectively, and  $\mathbb{P}_{\text{Error}}$  is the error rate. For the DDM, minimization of BR yields a transcendental equation for the threshold [55]

$$\frac{\xi_2}{\xi_1} \frac{2\mu^2}{\sigma^2} - \frac{4\mu\eta}{\sigma^2} + e^{-(2\mu\eta/\sigma^2)} - e^{(2\mu\eta/\sigma^2)} = 0.$$

Reward rate is generally defined by the proportion of correct trials divided by the average duration between decisions [62]

$$\begin{aligned} RR &= \frac{1 - \mathbb{P}_{\text{Error}}}{D_p \mathbb{P}_{\text{Error}} + T_{\text{Reaction}} + D} \\ &= \frac{1 - \mathbb{P}_{\text{Error}}}{D_p \mathbb{P}_{\text{Error}} + T_{\text{Decision}} + T_{\text{SM}} + D}, \end{aligned}$$

where  $D \in \mathbb{R}_{\geq 0}$  is the delay between a correct response and the next stimulus and  $D_p \in \mathbb{R}_{\geq 0}$  is a penalty delay introduced by an incorrect decision. Similar to the Bayes' risk, the optimal threshold for RR can be found by minimizing

$$\frac{\eta}{\mu} + D + T_{\text{SM}} + (D + T_{\text{SM}} + D_p - \frac{\eta}{\mu})e^{-(2\mu\eta/\sigma^2)} = \frac{1}{RR}.$$

These transcendental threshold equations can be solved numerically, and  $\xi_1$  and  $\xi_2$  can be estimated from empirical data [55]. The accuracy of these optimal threshold selection methods in predicting the speed-accuracy trade-off for human subjects has been investigated in [63]. The interested reader is also referred to [64] for a discussion on successes and shortcomings of such optimal methods.

The previous discussion pertains to a single two-alternative forced choice task. However, in most engineering applications, humans process a sequence of such tasks which imposes additional complexities. Researchers have found that, broadly speaking, human decision-making in a sequence of two-alternative forced choice tasks can often be effectively modeled by a sequence of drift-diffusion accumulators with sequentially updated initial conditions and thresholds [65]–[68].

### Interrogation Paradigm

Typical evolutions of the DDM under the interrogation paradigm are shown in Figure 2b. The interrogation paradigm relies upon a single threshold: for a given deadline  $T \in \mathbb{R}_{>0}$ , the operator decides in favor of the first (second) alternative if the evidence collected until time  $T$ , (that is,  $x(T)$ ) is greater (smaller) than a threshold. For equally likely alternatives, the threshold is chosen to be zero. From (1), the evidence collected until time  $T$  is a Gaussian random variable with mean  $\mu T + x_0$  and variance  $\sigma^2 T$ . Thus, if  $\nu \in \mathbb{R}$  represents the chosen threshold, the probability to decide in favor of the

first alternative under the interrogation paradigm can be written in closed-form as

$$\mathbb{P}(x(T) > \nu) = 1 - \mathbb{P}(x(T) < \nu) = 1 - \Phi\left(\frac{\nu - \mu T - x_0}{\sigma\sqrt{T}}\right), \quad (2)$$

where  $\Phi(\cdot)$  is the Gaussian cumulative distribution function.

### Generalizations of the DDM

A myriad of other accumulator models consider generalizations of the pure DDM. These variants often serve to capture additional behavioral characteristics that are not captured by the dynamics in (1). For example, the Ornstein-Uhlenbeck (O-U) model [69], [70] incorporates an additional linear term in the evidence accumulation equation

$$dx(t) = (\lambda x(t) + \mu)dt + \sigma dW(t), \quad x(0) = x_0, \quad (3)$$

where  $\lambda \in \mathbb{R}$ . The sign of  $\lambda$  determines whether evidence aggregation accelerates or decelerates with increasing evidence. The O-U model (3) has a fixed point at  $x = -\mu/\lambda$  and thus this model can represent situations where evidence accumulation asymptotes over time or, in other words, situations where the human is never perfectly accurate. This is a feature that the pure DDM does not have because (1) implies that in the absence of noise a human will always make a correct decision, given enough time. In the context of two-alternative forced choice tasks, if  $\lambda < 0$  then (3) is the reduction of what is sometimes called the *leaky competitive accumulator* (LCA) model [71]. The LCA model is characterized by leaky, stochastic, and competitive information accumulation in nonlinear decision units (one for each alternative), and has also been shown to capture neurally inspired properties, such as lateral inhibition and recurrent excitation [72].

Other generalizations, such as the *extended DDM* [60], [73] and the *full DDM* [74], incorporate additional parameters, including a noise parameter associated with the drift rate, a parameter characterizing initial latency, and a parameter capturing bias in the initial accumulation process. These variants have been shown to more accurately model the user response time distributions than the pure DDM [61], [75]. Further variants introduce the use of *collapsing thresholds*, where the decision-making threshold  $\eta \in \mathbb{R}$  is a function of the form  $\eta(t) = ce^{-rt}$ , where  $c$  and  $r$  are constants representing the initial threshold and the rate of convergence, respectively. These collapsing thresholds can be thought of as an “urgency signal” that prevents subjects from taking an excessive amount of time when drift rates are close to zero [74]. In some cases, collapsing threshold models can more accurately capture the higher reaction times that often occur in trials that result in an incorrect decision [76]. Optimality properties of collapsing threshold models have been explored in several contexts, including non-stationary environments [77], heterogeneous environments [78], and decision-making under deadlines [79].

The pure DDM is also related to classical hypothesis tests from probability theory. In the free response paradigm, the DDM (1) is the continuum limit of the *sequential probability ratio test* (SPRT) [80], a test that can be used when evidence is acquired sequentially at time steps  $\ell \in \mathbb{Z}_{\geq 0}$ . That is, the SPRT is equivalent to the DDM in the limit as the time between samples tends to zero [81]. Indeed, SPRTs utilize a statistic  $\Lambda_\ell$  that is incremented with each new observation  $y_\ell$ . A decision is made in favor of one of the alternatives once a threshold is reached. With symmetric thresholds  $\pm\Lambda_{\text{Thresh}}$ , unbiased initial evidence, and independent observations, a standard SPRT for deciding between hypotheses  $H^0, H^1$  is:

- 1: initialize  $\Lambda_0 := 0$ ;
- 2: at time  $\ell \in \mathbb{N}$ , collect observation  $y_\ell$ ;
- 3: integrate evidence  $\Lambda_\ell := \Lambda_{\ell-1} + \log \frac{\mathbb{P}(y_\ell|H^1)}{\mathbb{P}(y_\ell|H^0)}$ ;  
(decide only if the threshold  $\Lambda_{\text{Thresh}}$  is crossed)
- 4: **if**  $\Lambda_\ell < -\Lambda_{\text{Thresh}}$ , **then**  $H^0$  is true;
- 5: **else if**  $\Lambda_\ell > \Lambda_{\text{Thresh}}$ , **then**  $H^1$  is true;
- 6: **else** continue sampling (step 2);

The statistic  $\Lambda_\ell$  plays the role of evidence in favor each alternative. For sequentially accumulating data with known sampling likelihoods under each hypothesis, the SPRT is the optimal statistical test for two-alternative forced choice tasks, in the sense that it achieves a given error rate in minimum time [80], [82]. Despite this, some researchers argue that the standard SPRT does not capture reaction times observed in empirical data [73], and have turned to a variety of other variations in attempt to increase accuracy. In the interrogation paradigm, the DDM (1) is the continuum limit of the Neyman-Pearson hypothesis test [83], a test designed to decide among two hypotheses when the number of discrete data samples is fixed *a priori*. Given a set of observations  $\{y_1, y_2, \dots, y_n\}$ , the Neyman-Pearson test calculates the likelihood ratio

$$\Lambda(y_1, y_2, \dots, y_n) = \frac{\mathbb{P}(H^0 | y_1, y_2, \dots, y_n)}{\mathbb{P}(H^1 | y_1, y_2, \dots, y_n)},$$

and rejects the hypothesis  $H^0$  in favor of the hypothesis  $H^1$  if  $\Lambda$  is less than a threshold. Once again, the statistic  $\Lambda$  plays the role of evidence accumulated in favor of each alternative after a fixed amount of sampling. For a fixed number of data samples with known likelihoods, the Neyman-Pearson test is optimal in that it has the highest statistical power [84].

## Multi-alternative Forced Choice Tasks

A *multi-alternative forced choice task* is one in which an operator or observer must choose among multiple disjoint hypotheses. Broadly speaking, researchers have attempted to extend many of the same strategies for modeling two-alternative forced choice tasks for use with multiple alternatives through the use of *race models* [85], [86]. Race models

can be thought of as another variant of the pure DDM in which each alternative is assigned its own separate accumulator. That is, in a race model for an  $m$ -alternative forced choice task, there are  $m$  evidence accumulation variables  $x_1, x_2, \dots, x_m$ , representing evidence accumulated in favor of each respective alternative. Each of these variables then evolves according to a random process (such as the DDM), and a decision is made in favor of the alternative whose corresponding evidence  $x_i$  is the first to cross its respective evidence accumulation threshold (or has the largest value at the deadline in the case of the interrogation paradigm). The degree to which the accumulators interact varies depending on the problem setup.

Classical race models have exhibited some success in capturing behavioral phenomena. However, the multi-alternative scenario is inherently much more difficult to model than the two-alternative case [84]. For example, in discrete time the SPRT is the optimal test for achieving a given error rate in minimum time with sequentially accumulating data, but it has been shown that it is difficult to create an analogous constant-threshold multihypothesis sequential probability ratio test that is optimal in the same sense [87]. Despite this fact, some commonalities are generally accepted, such as *Hick's Law* [88], which states that in choosing between  $m$  alternatives, if accuracy is fixed at a high rate, then the mean reaction time increases at a rate proportional to  $\log(m)$  (although the exact form of this relation remains an open question). These commonalities are often used in attempts to validate models. For example, studies have shown that certain variations on the O-U model that are designed for multiple hypotheses [71], [84] outperform classic race models in some respects and capture the dynamics predicted by Hick's law.

More recent models have incorporated the use of modern technology. For visual stimuli, physiological sensing tools, namely eye-tracking, have been used in the context of race models as well. In this context, it is assumed that the relevant parameters that govern the dynamics of each accumulator are dependent upon the position of the observer's gaze. For example, it is assumed in [89] that the drift rate for a given accumulator is higher when the observer's attention is focused on the alternative in question. Other works, such as [90], have begun to use eye-tracking to explore the connection between visual characteristics such as saliency to the evidence accumulation process modeled via race model.

## Exogenous Factors

Human performance models discussed thus far only capture dynamics of evidence aggregation in decision-making. Exogenous factors, such as workload, fatigue, situational awareness, information retention, among others, also affect the decision-making process. For brevity, this article does not include an in-depth discussion of exogenous factors in evidence accumulation models. However, to give the reader a flavor for the types






Figure 5: Curves illustrating some key exogenous factors: (a) Performance as a function of operator stress, as described by the Yerkes-Dodson law; (b) Task effectiveness of an operator who wakes up at 6am after 6 hours of sleep, as predicted by the SAFTE model (5) using the default parameters in [91]; (c) Empirical memory retention data taken from [92], fitted with a curve that is the sum of two exponentials and a constant function.




Figure 4: Many factors can affect the decision-making process.

of models that exists, a few key factors are briefly mentioned. Note that these factors are closely linked with those discussed in “Key Challenges.”

#### Mental Workload, Stress, and the Yerkes-Dodson Law

Mental workload is the extent to which a task places demands on the operator’s cognitive resources [93], with a variety of models further reducing the construct into various sub-components [94]–[96]. Although operator mental workload is generally a subjective experience, many researchers attempt to capture this phenomena through more objective, quantifiable measures. For instance, operator workload is sometimes modeled as the utilization ratio (the fraction of recent history during which the operator was busy), with the utilization ratio  $u$  following the dynamics

$$\dot{u}(t) = \frac{b(t) - u(t)}{\tau}, \quad u(0) = u_0, \quad (4)$$

where  $b : \mathbb{R}_{>0} \rightarrow \{0, 1\}$  represents whether the operator is idle or busy,  $\tau \in \mathbb{R}_{>0}$  is the sensitivity of the operator, and  $u_0 \in [0, 1]$  is the initial utilization ratio [28]. Typically, system design focuses on methods of reducing workload to decrease the strain on the operator, but when taken too far this approach can result in performance degradation as well.

Closely related to mental workload is *operator stress*. The Yerkes-Dodson law [97], [98] is a classical model that captures the performance of an operator as a unimodal function of stress level. A typical representation of this relationship is shown in Figure 5a. The law demonstrates that there is a moderate level of stress, dependent on the task, that optimizes operator performance, while excessive stress (hyperstress) overwhelms the operator and too little stress (hypostress) leads to boredom and vigilance decrement [99]. Other work has expanded on this concept through more detailed models which differentiate between regions of psychological and physiological adaptability [100].

#### Fatigue, Sleep Cycle, and the SAFTE model

Fatigue is defined as the feeling of bodily discomfort after prolonged activity, and is known to have detrimental effects on operator performance [101]. Several models have been proposed to capture cognitive performance as a function of sleep deprivation [102]. One example is the *sleep activity fatigue task efficiency* (SAFTE) model [91], which assumes that a fully rested operator has a finite reservoir capacity  $R_c$  that depletes over time while the operator is awake, and replenishes when the operator sleeps. The SAFTE model



determines the task effectiveness (TE) as

$$\text{TE} = 100 \frac{R_c - 60KT_a}{R_c} + \left( a_1 + a_2 \frac{60KT_a}{R_c} \right) \left[ \cos \left( \frac{2\pi}{24} (T_d - p) \right) + \beta \cos \left( \frac{4\pi}{24} (T_d - p - p') \right) \right], \quad (5)$$

where  $T_a$  is the number of hours the operator has been awake,  $T_d$  is the time of the day in hours,  $K$  is reservoir drain rate due to wakefulness,  $a_1, a_2, \beta \in \mathbb{R}$  are constants,  $p$  is the time of the peak in the 24 hour circadian rhythm, and  $p'$  is the relative time of the 12 hour peak. Under this model, if the reaction time of a fully rested operator is  $T_{\text{Reaction}}$ , then the reaction time of the fatigued operator is  $T_{\text{Reaction}}/\text{TE}$ . An example TE curve generated using the SAFTE model is shown in Figure 5b.

### Information Retention and Situational Awareness

Information retention refers to the fraction of newly acquired information the operator remembers over time. Traditionally, the curve has been modeled as an exponential decay [103]. Some researchers [104] argue that the information retention curve should be modeled by a power-law function, while others [92] model the curve as a sum of two exponential functions and a constant function. An example of such a curve fitted to empirical data from [92] is shown in Figure 5c.

In many tasks, including supervisory tasks, the operator must not only perceive, process, and retain information, but also apply that knowledge to formulate an accurate mental image of their current situation. This leads to the notion of *situational awareness*, which can be defined as the sum of operator perception and comprehension of process information, and the subsequent ability to make projections of system states on this basis [105]. It has been argued that a lack of situational awareness results in poor performance by creating large waiting times, that is, the operator takes more time to start working on a task [106]. Situational awareness is critical as the operator is incapable of making timely and effective decisions without an accurate mental representation of the current and predicted future state of their operational environment, but can be difficult to moderate.

## COORDINATION OF AUTONOMOUS AGENTS

The design of coordination strategies for systems of autonomous agents is an issue that is at the heart of control theory and has generated a vast amount of research (for example, [107]–[111]). Here, the discussion is focused on *coverage problems* in the context of wireless sensor networks with a fixed number of nodes (agents), as this class of problem is applicable in many human supervisory control scenarios. Loosely speaking, the coverage problem is: given a compact area of interest  $\mathcal{Q} \subset \mathbb{R}^2$  and a team of agents

equipped with sensors capable of gathering information about their surroundings, determine a strategy to deploy and control the autonomous agents such that some coverage metric is maximized. In supervisory control, the agents generally can transmit data to a central location either by direct or multihop communication. A few of the most common coverage problems are discussed and some theoretical tools that can aid in solving them are highlighted below.

### Static Coverage

The most basic coverage problem is that of *static coverage*, that is, determining *a priori* a location where each of the agents will remain for some time. When the area  $\mathcal{Q}$  is relatively small, the static coverage problem often reduces to the problem of finding a location that maximizes the sensing footprint of the agents. The well-known *art gallery problem* [112] is an example of this type of coverage. The classic art gallery problem involves simple polygonal environments and visibility constraints; however, various extensions have been proposed to incorporate issues such as holes in the environment [113], additional coverage requirements [114], or sensor placement specifications [115].

In many cases, the environment of interest involves an element of stochasticity, that is, there is a probability density function  $\phi : \mathcal{Q} \rightarrow \mathbb{R}_{\geq 0}$  that encodes the likelihood of some event of interest occurring in any subregion. In this scenario, the goal is generally to place the sensors in a way that maximizes their ability react to events that may occur, proportional to the function  $\phi$ . Often the static sensor placement issue reduces to load-balancing. That is, each point in  $\mathcal{Q}$  is assigned to one agent, and the goal is to minimize a *multicenter function*  $\mathcal{H} : \mathcal{Q}^m \times (2^{\mathcal{Q}})^m \rightarrow \mathbb{R}_{\geq 0}$ ,

$$\mathcal{H}((c_1, \dots, c_m), (P_1, \dots, P_m)) = \sum_{i=1}^m \int_{P_i} d(k, c_i) \phi(k) dk,$$

where  $m \in \mathbb{N}$  is the number of sensors,  $(c_1, \dots, c_m)$  represents the location of the sensors,  $(P_1, \dots, P_m)$  is a partition of  $\mathcal{Q}$  (that is, satisfies  $\bigcup_{i=1}^m P_i = \mathcal{Q}$  and  $P_i \cap P_j = \emptyset$  for any  $i \neq j$ ), and  $d : \mathcal{Q} \times \mathcal{Q} \rightarrow \mathbb{R}_{\geq 0}$  is a distance metric. Finding global minimizers of the function  $\mathcal{H}$  is difficult in general; however, algorithms exist for finding high-quality approximate solutions in most typical cases. For example, if  $d$  is taken to be the square of the Euclidean distance, then the optimal sensor placement  $(c_1, \dots, c_m)$  and assignment  $(P_1, \dots, P_m)$  forms a centroidal Voronoi partition. Finding a (not necessarily globally optimal) centroidal Voronoi partition can easily be achieved through the well-known the Lloyd algorithm [116] and its variations. Additional partitioning schemes have also been proposed to incorporate constraints on coverage assignments [117], other cost functions [118], and varying communication protocols [119].

A thorough review of static coverage approaches in fixed

### EXAMPLE: PERSISTENT SURVEILLANCE MISSION

#### Human Performance Modeling

In the design of a human supervisory control system, the choice of the human model forms the basis for the cognitive system and supports virtually all other operations in the design strategy. This section focuses on the design of a performance function, which will drive the strategy for the rest of the system design.

This example uses the DDM (1) as the basis for constructing the human performance model. In general, human decision making will hinge upon a variety of factors not captured by the pure DDM. In this example, such exogenous factors are not explicitly considered. However, it should be noted that other decision-making models that do incorporate exogenous factors can also be used to construct a performance function in a similar manner.

The accuracy of decisions made by the operator is used as a measure of the operator's performance. Therefore, the probability of making the correct decision is selected to be the performance metric. The drift rates are assumed to be symmetric, that is, the drift rates are  $+\mu$  and  $-\mu$  when alternatives  $H^0$  and  $H^1$  are true, respectively. Recalling (2), the performance function when alternative  $H^0$  is true is  $f^0 : \mathbb{R}_{\geq 0} \times [0, 1] \rightarrow [0, 1]$ ,

$$f^0(t, \pi) = 1 - \Phi\left(\frac{\nu - \mu t - x_0}{\sigma\sqrt{t}}\right).$$

Similarly, the performance function when alternative  $H^1$  is true is  $f^1 : \mathbb{R}_{\geq 0} \times [0, 1] \rightarrow [0, 1]$ ,

$$f^1(t, \pi) = \Phi\left(\frac{\nu + \mu t - x_0}{\sigma\sqrt{t}}\right).$$

Given a prior probability  $\pi$  of the first alternative being true, the overall performance  $f : \mathbb{R}_{\geq 0} \times [0, 1] \rightarrow [0, 1]$  is

$$f(t, \pi) = \pi f^0(t, \pi) + (1 - \pi) f^1(t, \pi). \quad (6)$$

This performance function is a sigmoid function of time. Denote the  $k$ -th region by  $\mathcal{R}_k, k \in \{1, \dots, m\}$ . The surveillance mission is modeled as a sequence of two-alternative choice tasks and, accordingly, models the operator performance as in (6). The two alternatives  $H^0, H^1$  in this setting are the presence of an anomaly and the absence of an anomaly, respectively. The performance of the operator at region  $\mathcal{R}_k$  is denoted by  $f_k : \mathbb{R}_{\geq 0} \times [0, 1] \rightarrow [0, 1]$ . This presentation implicitly assumes that the evidence accumulated in the different regions is mutually independent.

sensor networks is provided in [120]

cases is briefly discussed below.

### Dynamic Coverage

*Dynamic coverage* typically refers to those problems in which a set of autonomous agents do not remain at fixed positions, but rather continually move throughout the environment to accomplish some task. Dynamic coverage is often used to accommodate certain performance goals or environmental characteristics that are not well suited to static coverage schemes. For instance, large and time varying environments (that is, those in which importance weights may change or the likelihood of events of interest is time-varying) may be better suited to dynamic coverage.

The dynamic coverage problem has several flavors, including random and non-uniform spatial-temporal fields [123], timevarying agent dynamics [124], dynamic vehicle routing problems [125], and informative path planning [126]–[134]. One particularly relevant class of dynamic coverage problems is persistent coverage or patrolling, where a set of vehicles is required to endlessly survey an environment. This type of coverage arises for applications such as the monitoring of oil spills [135], the detection of forest fires [136], the tracking of border changes [137], and general environmental monitoring [138]. In persistent coverage schemes, vehicles continuously visit regions in the environment according to some policy that is deterministic or stochastic. Each of these

### Deterministic Policies

For regions of interest that are represented as open subsets of Euclidean space, deterministic policies for persistent coverage include the construction of predetermined motion routines (such as lawnmower patterns), the adaptation the static coverage strategies [131], and the modeling of environments as random fields and subsequent design of optimal trajectories [139], [140].

In the context of discretized regions (that is, regions of interest that are represented as a graph), many deterministic policies rely on (i) computing a tour through the regions and (ii) requiring the vehicles to endlessly move along the tour (see for example [130], [134], [141], [142]). In several cases, the discrete case is closely related to network location, multiple traveling salesperson (TSP), graph exploration, or other classic vehicle-routing problems. Indeed, almost all traditional approaches to solving the discrete, deterministic, persistent coverage problem rely on state-space decomposition, and TSP tour computation [143]. However, recent works have looked at non-TSP tours [134], [140] as well as non-tour based policies [144].

Deterministic policies are often simple to implement, but are mostly periodic and predictable, which may be undesirable. If,

### EXAMPLE: PERSISTENT SURVEILLANCE MISSION Vehicle Routing and Anomaly Detection Algorithms

This sidebar focuses on the construction of a vehicle routing policy to govern motion of the UAV (autonomous agent). To this end, a simple routing policy that directs the UAV to a randomly chosen region during each visit is adopted. Recall that the goal is to detect anomalies in each of the regions of interest in the shortest amount of time, subject to a false-alarm constraint. To be consistent with this goal, it is desired that the probability of the UAV traveling to a given region should be proportional to the likelihood of that region being anomalous. However, since the decisions made by the human operator may be erroneous, his/her input cannot be accepted as a reliable indicator of the presence of an anomaly. Therefore, the routing strategy needs a tool to accurately determine the likelihood of an anomaly at each region.

The tool chosen for this example is a variation on the standard *cumulative sum* (CUSUM) algorithm [121] called the *ensemble CUSUM algorithm* [34], which is a statistical quickest-change-detection algorithm consisting of a set of  $m$  parallel CUSUM algorithms (one for each region). Accordingly, the binary decisions by the operator are treated as Bernoulli random variables whose distribution is dictated by the performance function. Subsequently, the ensemble CUSUM algorithm is run on these decisions to decide reliably on a region being anomalous. The standard CUSUM algorithm requires the observations from each region to be independent and identically distributed. However, the decisions made by the operator do not satisfy these requirements. Therefore, instead of the standard CUSUM algorithm, a CUSUM-like algorithm for dependent observations [122] is used instead. The ensemble CUSUM algorithm maintains a statistic  $\Lambda_k^\ell$  for each region  $\mathcal{R}_k$ ,  $k \in \{1, \dots, m\}$  and time step  $\ell$ . The statistic at region  $\mathcal{R}_k$  is updated using the binary decision of the operator whenever a task from region  $\mathcal{R}_k$  is processed. If the statistic associated with a region crosses a threshold  $\Lambda_{\text{thresh}}$ , then the region is declared to be anomalous. The choice of this threshold dictates the accuracy of the detection [121]. It is assumed that once an anomaly has been detected it is removed, and then, consequently, the operator's belief about the region being anomalous resets to the default value. Let  $k_\ell$  represent the region index of the  $\ell$ -th task, and let  $\pi_{k_\ell}^\ell$  represent the prior probability of an anomaly at region  $k$  after processing the  $\ell$ -th task. The ensemble CUSUM algorithm is

- 1: initialize  $\ell := 1$ ,  $\Lambda_k^0 := 0$ , for each  $k \in \{1, \dots, m\}$ ;
- 2: **if**  $\text{dec}_\ell == 1$  and  $t_\ell > 0$ , **then**

$$\Lambda_{k_\ell}^\ell = \max \left\{ 0, \Lambda_{k_\ell}^{\ell-1} + \log \frac{f_{k_\ell}^1(t_\ell, \pi_{k_\ell}^{\ell-1})}{1 - f_{k_\ell}^0(t_\ell, \pi_{k_\ell}^{\ell-1})} \right\};$$

- 3: **else if**  $\text{dec}_\ell == 0$  and  $t_\ell > 0$ , **then**

$$\Lambda_{k_\ell}^\ell = \max \left\{ 0, \Lambda_{k_\ell}^{\ell-1} + \log \frac{1 - f_{k_\ell}^1(t_\ell, \pi_{k_\ell}^{\ell-1})}{f_{k_\ell}^0(t_\ell, \pi_{k_\ell}^{\ell-1})} \right\};$$

(detect an anomaly if a threshold is crossed)

- 4: **if**  $\Lambda_{k_\ell}^\ell \geq \Lambda_{\text{thresh}}$ , **then**
- 5: declare an anomaly at region  $k_\ell$ ;
- 6:  $\Lambda_{k_\ell}^\ell = 0$ ;
- 7: set  $\ell = \ell + 1$ ; go to 2;

Having established this anomaly detection tool, a simple routing policy is employed that sends the UAV to each region with a probability proportional to the likelihood of that region being anomalous. In particular, the probability to visit region  $\mathcal{R}_k$  is initialized to  $q_k^0 = 1/m$  and after processing each task, the probability to visit region  $\mathcal{R}_k$  is chosen proportional to  $e^{\Lambda_k^\ell} / (1 + e^{\Lambda_k^\ell})$ . This simple strategy ensures that a region with a high likelihood of being anomalous is visited with a high probability. Moreover, it ensures that each region is visited with a nonzero probability at all times and consequently, an anomalous region is detected in finite time.

Note that such a simple vehicle routing algorithm only determines the probability with which the UAV should visit different regions and does not take into account factors such as the geographic location of regions, importance weights assigned to regions, vehicle travel times between regions, or the difficulty of detection at each region. These factors could be incorporated into the vehicle routing algorithm [34]; however, for simplicity of the presentation, such factors are not considered here. Vehicle travel time and importance weights are, however, taken into consideration in the design of the decision support system, which is presented in subsequent sidebars. Indeed, vehicle travel times are used to determine the rate at which the UAV generates imagery to send to the operator for analysis, and importance weights are used in deriving the reward function used to optimize time allocations.

for example, the goal is to detect the existence of an intruder, then the intruder may hide when a vehicle is nearby and thus, most deterministic policies will fail [145]. (Although there do exist a few deterministic strategies that partially address this issue, such as those in [146], [147], which use ergodic theory to produce vehicle trajectories that are largely unpredictable to an outside observer.)

### Stochastic Coverage Policies

In contrast to deterministic policies, stochastic coverage policies are often much less predictable. Although a few researchers have adopted elements of stochasticity into surveillance of regions that are represented by open subsets of Euclidean space (for example, [148]–[150]), the majority of existing policies assume discretized areas or discrete regions of interest (for example, [151]–[153]). In light of this fact, the remainder of the discussion is focused on discretized regions

of interest.

Stochastic coverage policies for discrete regions typically involve an ergodic Markov chain in which each region represents a state. Transition probabilities and stationary distributions are then designed according to an appropriate surveillance criterion. In general, the coverage criterion depends on the mission objective. For example, if the mission objective is the detection of anomalous regions, then the surveillance criterion may be chosen to minimize the average detection delay [34]. The minimization of the average detection delay inherently considers the difficulty of detection at each region, the travel times between the regions, and the likelihood of each region being anomalous.

For a single vehicle, there are two popular schemes to construct a Markov chain with a desired stationary distribution (surveillance criterion), namely, the *Metropolis-Hastings* algorithm and the *fastest-mixing Markov chain* (FMMC) method. The two schemes can be briefly described as follows. Consider a set of regions modeled by the graph  $\mathcal{G} = (V, \mathcal{E})$ , where  $V$  is the set of  $m$  nodes (each node corresponds to a region) and  $\mathcal{E}$  is the set of edges representing the connectivity of the regions. Let the surveillance criterion be  $\mathbf{q} = (q_1, \dots, q_m) \in \Delta_m$ . The Metropolis-Hastings algorithm [154] picks the transition matrix  $A$ , that is, the matrix of transition probabilities from each state to every other state of the Markov chain, as

$$A_{ij} = \begin{cases} 0, & \text{if } (i, j) \notin \mathcal{E}, \\ \min \left\{ \frac{1}{d_i}, \frac{q_j}{q_i d_j} \right\}, & \text{if } (i, j) \in \mathcal{E} \text{ and } i \neq j, \\ 1 - \sum_{k=1, k \neq i}^m A_{ik}, & \text{if } (i, j) \in \mathcal{E} \text{ and } i = j, \end{cases}$$

where  $d_i$  is the number of regions that can be visited from region  $\mathcal{R}_i$ . For the FMMC method, the transition matrix  $A \in \mathbb{R}^{m \times m}$  with a desired stationary distribution  $\mathbf{q} \in \Delta_m$  is determined by solving the semidefinite program [155]

$$\begin{aligned} & \text{minimize} \quad \|\sqrt{Q}A\sqrt{Q} - \mathbf{q}_{\text{root}}\mathbf{q}_{\text{root}}^\top\|_2 \\ & \text{subject to} \quad A\mathbf{1}_m = \mathbf{1}_m, \\ & \quad \quad \quad QA = A^\top Q, \\ & \quad \quad \quad A_{ij} \geq 0, \text{ for each } (i, j) \in \mathcal{E}, \\ & \quad \quad \quad A_{ij} = 0, \text{ for each } (i, j) \notin \mathcal{E}, \end{aligned}$$

where  $Q$  is a diagonal matrix with diagonal  $\mathbf{q}$ ,  $\mathbf{q}_{\text{root}} = (\sqrt{q_1}, \dots, \sqrt{q_m})$ , and  $\mathbf{1}_m$  is the vector of all ones. To achieve the coverage criterion at an accelerated rate, a time-varying Markov chain can also be constructed in the spirit of [152]. Variants on these algorithms exist that also seek to minimize additional heuristics related to the chain, such as the mean first-passage time (also known as the hitting time or Kemeny constant) [156].

For multiple vehicles, remarkably little is known about the design of cooperative surveillance based on “multiple Markov chains.” A naive stochastic policy that achieves the coverage criterion is to let each vehicle follow the single vehicle policy. A drawback of such a naive policy is that two or

more vehicles may survey the same region simultaneously, which may introduce a risk of collisions and nonoptimal coverage strategies. This drawback can be partially mitigated by constructing a Markov chain on a lifted space from which the undesired states are removed. Decentralized strategies, such as the message passing-based auction algorithm in [145], exist for constructing such policies.

## INTERFACING HUMANS AND AUTONOMOUS AGENTS

Once a model of human behavior has been established and the appropriate vehicle routing policy has been selected, the last step in the design of a human supervisory control system is the construction of the interface that links the two components. This step is essentially “closing the loop” by linking the autonomous agents with the human operator. As discussed, efficient designs must incorporate automated mechanisms to facilitate interactions between system components. Such mechanisms can take numerous forms, many of which can benefit directly from the incorporation of control-theoretic tools. This section provides illustrating examples of facilitating mechanisms known as *decision supports*, focusing on those that can be derived using control theory. The discussion concludes by highlighting some key challenges to effectively coupling humans and automated agents.

### Decision Supports

Researchers have established a simplified four-stage model of human cognition consisting of (i) information acquisition, (ii) information analysis, (iii) decision and action selection, and (iv) action implementation [161]. These abstract functions operate at various levels of granularity within a given task, and generally interact with each other in a continuous and complex fashion. Indeed, one cognitive process may be used to make decisions on low-level tasks, such as where to look next, while a different, simultaneous cognitive process may be working to make a dependent, higher-level decision, such as deciding if a target is present. As such, there is potential to improve system performance through incorporation of automated tools that focus specifically on aiding decision making, and consequently sharing the total cognitive load across system resources.

In this spirit, a *decision support* can be defined as any automated function that supports the *decision and action selection* stage of the cognitive process. In human supervisory control, this can mean directing operator attention, providing timing suggestions to the operator, preprocessing of tasks, and/or adjusting automation parameters, among many other possibilities. The specific form and potential for success of a given decision support system varies by application and system constraints. However, to illustrate how decision supports can be integrated at various levels of system operation, some

## EXAMPLE: PERSISTENT SURVEILLANCE MISSION Decision Support System

The design of a decision-support system for the example surveillance problem is now considered. Specifically, a support system is considered that uses operator performance to suggest the optimal amount of time to be spent on each task. Queuing theory has emerged as a popular paradigm to model supervisory control systems [20], [23], [24], [27]–[30]. Accordingly, it is assumed that the images collected by the UAV at the various regions over time are stacked in a queue, while they await operator analysis. The images arrive to the queue according to a stochastic process. A stability requirement is imposed on the queue length, namely, the queue length should remain finite for all time. The operator receives a reward for a correct decision on each task, and operator performance is quantified as the expected reward obtained after processing the task. The goal is to suggest time allocations to tasks such that the operator's overall reward per unit task is maximized. The system is designed under the following assumptions: (i) operator performance functions for a task originating from region  $\mathcal{R}_k$  in absence and presence of an anomaly are  $f_k^0 : \mathbb{R}_{\geq 0} \times [0, 1] \rightarrow [0, 1]$  and  $f_k^1 : \mathbb{R}_{\geq 0} \times [0, 1] \rightarrow [0, 1]$ , respectively; (ii) based on the importance of the region, a weight  $w_k \in \mathbb{R}_{>0}$  is assigned to each task collected from region  $\mathcal{R}_k$ ; (iii) tasks arriving to the queue while the  $\ell$ -th task is served are sampled from a probability distribution that assigns a probability  $q_k^\ell \in [0, 1]$  to region  $\mathcal{R}_k$ . Similar to (6), the average performance function  $f_k : \mathbb{R}_{\geq 0} \times [0, 1] \rightarrow [0, 1]$  at region  $\mathcal{R}_k$  is defined by  $f_k(t, \pi) = (1 - \pi)f_k^0(t, \pi) + \pi f_k^1(t, \pi)$ . Under the aforementioned assumptions, each task from region  $\mathcal{R}_k$  is characterized by the pair  $(f_k, w_k)$ .

For simplicity, the tasks in the queue are assumed to be processed by the operator on a first-come-first-serve basis. Let the  $\ell$ -th task in the queue be from region  $\mathcal{R}_{k_\ell}$ , and let the belief of the operator about region  $\mathcal{R}_k$  being anomalous before processing the  $\ell$ -th task be  $\pi_k^{\ell-1}$ . Initially the operator is unbiased about each region being anomalous, that is,  $\pi_k^0 = 0.5$ , for each  $k \in \{1, \dots, m\}$ . Given a time allocation  $t_\ell \in \mathbb{R}_{>0}$  to the  $\ell$ -th task in the queue, the operator's belief after processing the  $\ell$ -th task is estimated using the Bayes rule

$$\bar{\pi}_j^\ell = \begin{cases} \frac{\pi_j^{\ell-1} \mathbb{P}(\text{dec}_\ell | H_j^1, t_\ell)}{(1 - \pi_j^{\ell-1}) \mathbb{P}(\text{dec}_\ell | H_j^0, t_\ell) + \pi_j^{\ell-1} \mathbb{P}(\text{dec}_\ell | H_j^1, t_\ell)}, & \text{if } j = k_\ell, \\ \pi_j^{\ell-1}, & \text{otherwise,} \end{cases}$$

where  $H_k^0$  and  $H_k^1$  denote the hypothesis that region  $\mathcal{R}_k$  is nonanomalous and anomalous, respectively,  $\text{dec}_\ell \in \{0, 1\}$  is the operator's decision, and  $\mathbb{P}(\text{dec}_\ell | \cdot, t_\ell)$  is determined from the performance function of the operator,

$$\mathbb{P}(\text{dec}_\ell = 1 | H_{k_\ell}^1, t_\ell) = f_{k_\ell}^1(t_\ell, \pi_{k_\ell}^{\ell-1}).$$

The event that a region becomes anomalous corresponds to a change in the characteristic environment, which may happen at an arbitrary time. In a sequential change detection task, if the belief of the operator about a region being anomalous

is below a threshold, then the operator resets their belief to the threshold value [157]. The threshold is chosen to be 0.5. Consequently, the belief of the operator at region  $\mathcal{R}_k$  after processing the  $\ell$ -th task is  $\pi_k^\ell = \max\{0.5, \bar{\pi}_k^\ell\}$ .

The system should suggest to the operator the amount of time to spend on each task. To this end, support system is designed to maximize the infinite-horizon average reward, under the finite queue length constraint. The reward  $r : \mathbb{N} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$  obtained by allocating time  $t$  to the  $\ell$ -th task is

$$r(\ell, t) = w_{k_\ell} f_{k_\ell}(t, \pi_{k_\ell}^{\ell-1}),$$

where  $k_\ell$  is the index of the region that generated the  $\ell$ -th task. The objective of the decision-support system is to maximize the infinite-horizon average reward

$$V_{\text{avg}} = \liminf_{n \rightarrow +\infty} \frac{1}{n} \sum_{\ell=1}^n r(\ell, t_\ell), \quad (7)$$

while enforcing stability of the queue length. The solution to (7) is computationally intractable in general. However, a dynamic, approximate solution can be obtained under a certainty equivalent assumption, which approximates future uncertainties of the system by their expected values [158]–[160]. Specifically, the expected value of the operator's belief at a future time is equal to operator's current belief. Accordingly, the vehicle routing policy and the performance functions are stationary at all future decision times. The expected rate of arrival of tasks into the queue using current system parameters is  $\lambda_\ell = 1/(\mathbf{q}^\ell \top D \mathbf{q}^\ell + \mathbf{q}^\ell \top \mathbf{T})$ , where  $\mathbf{q}$  a vector of region visit frequencies (as determined by the vehicle routing policy),  $D$  is a  $m \times m$  matrix whose  $i, j$ -th entry represents the travel time between regions  $\mathcal{R}_i$  and  $\mathcal{R}_j$ , and  $\mathbf{T}$  is a vector whose entries represent image-generation times in the respective regions. Further, by the strong law of large numbers, the expected value of the function  $V_{\text{avg}}$  while the  $\ell$ -th task is processed is simply the expected average reward, calculated using current system parameters, that is,  $\bar{V}_{\text{avg}} = \sum_{k=1}^m q_k^\ell w_k f_k(t_k^{\text{reg}}, \pi_k^{\ell-1})$ , for some  $t_k^{\text{reg}}$  representing a stationary amount of time to be allotted to tasks originating from region  $\mathcal{R}_k$ . Therefore, under the certainty-equivalent assumption, the optimal time allocation to the  $\ell$ -th task can be approximated by solving at each time step,

$$\begin{aligned} & \underset{\hat{t}_1^{\text{reg}}, \hat{t}_1^{\text{reg}}, \dots, \hat{t}_m^{\text{reg}}}{\text{maximize}} && \sum_{k=1}^m q_k^\ell w_k f_k(\hat{t}_k^{\text{reg}}, \pi_k^{\ell-1}) \\ & \text{subject to} && \sum_{k=1}^m q_k^\ell \hat{t}_k^{\text{reg}} \leq \frac{1}{\lambda_\ell}, \\ & && \hat{t}_k^{\text{reg}} \geq 0, \text{ for each } k \in \{1, \dots, m\}, \end{aligned} \quad (8)$$

subsequently choosing  $t_\ell = \hat{t}_{k_\ell}^{\text{reg}}$ . Note that the first constraint enforces queue length stability [30], and that (8) is a knapsack problem with sigmoid utilities, whose solution can be approximated [30].




Figure 6: Distribution of foveal receptors as a function of radial eccentricity, taken with permission from [163].

examples that operate on different decision-making tasks and have the potential to drastically improve system performance are discussed in what follows.

#### Attention Allocation as an Optimization Problem

In supervisory tasks in which the operator has multiple simultaneous responsibilities, the question of where and how the operator should direct her/his attention becomes an important component to task success. In visual perception tasks, low-level decisions focus on where the operator should direct their attention within a given image or video.

It is well known in psychology literature that evidence accumulation in visual perception is highly dependent upon *radial eccentricity*, that is, the angular distance of a stimulus from the *foveal region*, the point in the visual field of highest resolution. The foveal region corresponds to the point on the stimulus at which the operator is directly looking. Indeed, due to the high density of foveal receptors when compared to the visual periphery (see Figure 6), evidence accumulation is generally much faster when a person is looking directly at a stimulus [162].

Suppose a model for a human operator’s accuracy in making a decision about a particular target as a function of time and radial eccentricity is given, and that access to a the operator’s fixation locations in real time is available (an assumption that is not unrealistic, given the increased availability and affordability of eye-tracking hardware [164]). Assume also that a given image is discretized into  $m$  disjoint, equally sized regions. In a static image, the amount of evidence about some target of interest is finite. Therefore, a differential equation of the form

$$\dot{x}_k = g(x_k, e)$$

can be associated to each region  $k$ , where  $x_k$  is the amount of evidence accumulated about the properties of some target in region  $k$ ,  $e$  is radial eccentricity, and  $g$  is a function that relates these two variables to the speed at which the operator accumulates evidence. Under this construction, the question of directing operator attention within a search task reduces to an optimization problem of the form

$$\begin{aligned} & \text{minimize}_{\{\text{Fix}_1, \dots, \text{Fix}_n, n\}} \quad \text{Dur}_{\text{Total}} \\ & \text{subject to} \quad \sum_{k=1}^m x_k = \alpha \sum_{k=1}^m \text{Evid}_k, \end{aligned}$$

where  $\text{Fix}_i = (\text{Loc}_i, \text{Dur}_i)$  is a tuple encoding the location and duration of the  $i$ -th fixation,  $n \in \mathbb{N}$  is the number of fixations,  $\text{Evid}_k$  is the total available information about the target in region  $k$ ,  $\alpha \in [0, 1)$  is a constant that captures the fraction of total information to be collected, and  $\text{Dur}_{\text{Total}} := \sum_{i=1}^n \text{Dur}_i$ . In other words, a sequence of fixations can be estimated that enables the operator to accumulate a specified fraction of available information about a particular image in the shortest amount of time (in general, it will not be possible to acquire all available sensory evidence in finite time). With this information, it may be possible to direct operator attention within image searches (assuming that it is possible to construct visual cues that the operator will respond to, an issue that will be discussed later).

In application, the function  $g$  could depend on many factors, such as visual clutter of the image [165] and task difficulty [166]. In addition, when the visual stimulus is a video, then the amount of evidence present evolves over time, and thus may not be finite. Further, a changing stimulus may alter the evidence accumulation process, and thus the model may need to be altered to take into account motion characteristics.

#### Timing as a Resource Allocation Problem

The goal for many supervisory control applications is to have a single operator who is capable of processing multiple tasks or data streams simultaneously [167]. In such tasks, the operator must not only decide how to allocate their attention within each task, but at a higher level must also decide how to allocate attentional resources across tasks. Assuming tasks to be processed are stacked in a queue, then the problem of deciding how much time the operator should spend on each task in the queue becomes a resource allocation problem. Indeed, time can be thought of as a resource that needs to be distributed among the necessary tasks. The inherent dynamics in persistent task analysis missions make these resource allocation problems a dynamic optimization. The solution to such optimization problems can be computed for small horizon lengths; however, for large horizons such computations are often not tractable. Despite these difficulties, high-quality suboptimal solutions can still be computed using tools from control theory.

In a general sense, the infinite-horizon optimization for a persistent task analysis mission can be formulated as follows. Consider a time-varying dynamical system of the form  $x_{\ell+1} = \text{evol}_\ell(x_\ell, t_\ell, d_\ell)$ , where  $\ell \in \mathbb{N}$  denotes the task index,  $x_\ell \in \mathcal{X}$  is the state variable,  $t_\ell \in \mathbb{R}_{\geq 0}$  is the time to be devoted to the  $\ell$ -th task,  $d_\ell \in \mathcal{D}$  is a disturbance, and  $\text{evol}_\ell : \mathcal{X} \times \mathbb{R}_{\geq 0} \times \mathcal{D} \rightarrow \mathcal{X}$  denotes the evolution map of the system. Given a stochastic model for  $\{d_\ell\}_{\ell \in \mathbb{N}}$ , the control objective is to solve the optimization

$$\begin{aligned} & \underset{\mathbf{t} := \{t_\ell\}_{\ell \in \mathbb{N}}}{\text{maximize}} && \liminf_{n \rightarrow +\infty} \frac{1}{n} \sum_{\ell=1}^n g(x_\ell, t_\ell, d_\ell), \\ & \text{subject to} && \mathbf{t} \in \mathcal{C}, \\ & && t_\ell \geq 0, \text{ for each } \ell \in \mathbb{N}, \end{aligned} \quad (9)$$

where  $x_1$  is the initial state,  $\mathbf{t}$  is the sequence of times devoted to each task,  $g : \mathcal{X} \times \mathbb{R}_{\geq 0} \times \mathcal{D} \rightarrow \mathbb{R}$  is the stage reward, and  $\mathcal{C}$  is a constraint set.

In general, the optimization (9) is hard; however, solutions can often be approximated in a tractable way using receding-horizon control [160]. *Receding-horizon control* approximates the solution to the infinite-horizon optimization (9) by solving a finite-horizon optimization problem at each iteration to sequentially determine the control input. In the presence of uncertainty, however, future parameters required for this finite horizon optimization may not be known. A common strategy for addressing this issue is to adopt a *certainty-equivalent* assumption, which replaces future uncertainties of the system by their expected values [158]–[160]. Specifically, the certainty-equivalent receding-horizon control scheme determines the control input at time  $\ell$  by solving the optimization problem

$$\begin{aligned} & \underset{\hat{t}_0, \hat{t}_1, \dots, \hat{t}_{n-1}}{\text{maximize}} && \frac{1}{n} \sum_{j=0}^{n-1} g(\bar{x}_{\ell+j}, \hat{t}_j, \bar{d}_{\ell+j}), \\ & \text{subject to} && (\hat{t}_0, \hat{t}_1, \dots, \hat{t}_{n-1}) \in \bar{\mathcal{C}}_\ell, \\ & && \hat{t}_j \geq 0, \text{ for each } j \in \{0, \dots, n-1\}, \end{aligned} \quad (10)$$

where  $\{\hat{t}_0, \dots, \hat{t}_{n-1}\}$  is time allocated to each task over the horizon,  $\bar{\mathcal{C}}_\ell$  is a modified constraint set, and  $\bar{x}_{\ell+j}$  is the certainty-equivalent evolution of the system, that is, the evolution of the system obtained by replacing the uncertainty in the evolution at each stage by its expected value,  $\bar{x}_\ell = x_\ell$ , and  $\bar{d}_{\ell+j}$  is the expected value of the uncertainty at stage  $\ell+j$ . The certainty-equivalent receding-horizon control scheme at iteration  $\ell$  solves optimization problem (10) and picks  $t_\ell = \hat{t}_0$ . If the deterministic dynamic program (10) can be solved efficiently, then certainty-equivalent receding-horizon control offers a computationally tractable suboptimal solution to problem (9).

In some cases, it may be tractable to solve the optimization (10) over very large or even infinite time horizons without resorting to additional approximations. For example, suppose that tasks are generated from  $k \in \{1, \dots, m\}$  different sources and stacked in the queue for analysis, and suppose that  $q_k^\ell \in [0, 1]$  represents the probability that the next task entering

the queue is generated by the  $k$ -th source at time step  $\ell$ . Further, suppose that under the certainty-equivalence assumption the utility functions for each source and probabilities  $q_k^\ell$ ,  $k \in \{1, \dots, m\}$  are stationary. In particular, for each  $k$  let  $f_k : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  be a continuous, monotonically increasing function representing the utility obtained by spending time  $t$  on a task generated by the  $k$ -th source. It is desired to maximize the infinite horizon reward while requiring a finite queue at all times. Invoking the strong law of large numbers makes solving (10) in the limit as  $n \rightarrow \infty$  reduce to solving a *knapsack problem*

$$\begin{aligned} & \underset{\hat{t}_1^{\text{reg}}, \hat{t}_1^{\text{reg}}, \dots, \hat{t}_m^{\text{reg}}}{\text{maximize}} && \sum_{k=1}^m q_k^\ell f_k(\hat{t}_k^{\text{reg}}), \\ & \text{subject to} && \sum_{k=1}^m q_k^\ell \hat{t}_k^{\text{reg}} \leq \frac{1}{\lambda}, \\ & && \hat{t}_k^{\text{reg}} \geq 0, \text{ for each } k \in \{1, \dots, m\}, \end{aligned} \quad (11)$$

where  $\lambda \geq 0$  is the rate of tasks entering the queue (the constraint function will ensure that the queue remains finite [30]). The control scheme would then choose  $t_\ell = \hat{t}_{k_\ell}^{\text{reg}}$ , where  $k_\ell$  is the region index of the  $\ell$ -th task. Although the problem (11) is simpler at first glance, it may still be difficult to solve exactly. For example, when the utility (reward) functions associated with knapsack problems are based on accumulator models of human decision making, they may take the form of sigmoid functions, that is, the utility functions  $f_k$  are sigmoid functions of time. The knapsack problem with sigmoid utilities is an NP-hard problem (although computationally tractable 2-factor solutions have been proposed [31]). If the problem (11) can be solved efficiently, an approximate solution to (9) can be obtained by solving an infinite horizon optimization under the certainty-equivalent assumption at each time step.

Simplifications, such as (11), must be used with caution since small modifications to the problem structure may cause such a formulation to lose validity. For instance, the addition of deadlines on tasks or latency penalties (penalties due to delay in processing a task) break this structure. In such scenarios, it may be necessary to go back to solving (10) over a short time horizon at each time step using standard dynamic programming techniques.

#### Adaptive Automation as a Feedback Control Problem

In addition to attention-allocation issues, at a higher level of granularity there is the issue of deciding what system functions should be left to the human operator and what functions should be automated. According to the 4-stage model in [161], each of the stages of the cognitive process may be automated to differing degrees within a single system. The choice of what aspects to automate may vary based on application, or based on the state of the operator.

In a supervisory role, most tasks in the action implementation



category (such as vehicle motion control) are automated at a high level; however, even in this case there are design decisions about how supervisory system should allow the operator to issue commands to automated subsystems [168]. For example, the authors of [169] show that a task-based control scheme in which the operator is only allowed to issue high-level commands, in many aspects outperformed a vehicle-based scheme, where the operator can control the motion of each automated agent individually.

Studies have shown that great care must be taken in choosing the right level of automation for a given system, since increased automation does not always lead to better performance. Indeed, increased automation can lead to increased operator complacency and bias [170]. As a result, researchers have begun to study an adaptive approach, in which the level of system automation is altered in an attempt to maintain the operator's cognitive state in some desired regime. In this simplified sense, the issue of adaptive autonomy reduces to a control problem. Indeed, if it can be verified that the level of autonomy of a given system has some effect on operator performance, then level of autonomy can be thought of as a control input that can be used to guide user performance.

For example, suppose the operator's workload is modeled via the utilization ratio presented in Equation (4). Then, operator workload is inversely related to the level of autonomy of a given system. The Yerkes-Dodson law suggests that moderate levels of operator workload (stress) lead to the highest levels of performance. It is a natural step, therefore, to design a feedback control law that adjusts automated functionalities to keep the utilization ratio, and thus the operator workload, within a moderate regime.

Of course, this type of control approach hinges on a myriad of assumptions about operator behavior and its relationship to performance, as well as the ability to accurately measure the complexities of human cognition and performance in real time. Some of these considerations are discussed in subsequent sections; however, the underlying concept of relating automation parameters to performance and using this easily adjustable parameter as a means of control is one that has received recent researcher attention and will remain an important topic.

## Key Challenges

Many of the formulations discussed assumed simplified models and optimization strategies that are loosely coupled across system components. Although this framework may suffice in some circumstances, in reality, a human supervisory control system and an associated decision-support system operate through a combination of driving factors that work together simultaneously. As such, the discussion is concluded by highlighting a few key challenges to improving design strategies

and effectively implementing them in practice.

### *Tightly Coupling System Components*

In human supervisory control, the design of a decision-support system and the design of coordination strategies for automated agents are often treated as completely decoupled or loosely coupled problems. In many instances there could be a tighter coupling in how the performance of decision supports and autonomous agents influence each other through the user. For example, in some operational contexts, autonomous agents may have to loiter until the operator can attend to them. Such operator-induced delays could lead to degradation in coverage performance. From a mathematical perspective, the problem of scheduling both the user and autonomous agents could be posed as a joint optimization problem to achieve overall system objectives. Hence, the design of decision support systems and control schemes for autonomous agents could be considered jointly. Further research is needed to develop appropriate formulations of such problems, incorporating different dynamic and performance characteristics, and developing tractable computational methods.

### *Assessing the Operator State*

The creation of effective decision supports often hinges upon the ability to accurately assess the operator's cognitive state, including situational awareness, perceived workload and fatigue. However, it is difficult to assess such a state with any degree of accuracy using current technology. Recent technological advances have made the use of physiological sensors, such as eye-trackers, electroencephalogram (EEG, measuring cortical electrical activity), and electrocardiogram (ECG, measuring heart beats), a viable option for providing real-time data in many applications [164]. As such, a large body of recent research has gone into finding correlations between cognitive activity and objective measures, such as pupil diameter [39], [171], blink rates [172], heart rate [173], [174], and EEG activity [175], [176].

Even though these studies have successfully found correlations in certain scenarios, it remains difficult to use such findings in practice. One reason is that it is difficult to control exogenous factors in real applications. For example, researchers have found correlations between pupil diameter variations and cognitive processing, but there are at least 23 factors that can affect pupil size [177]. Thus, it is hard to rely on pupil diameter alone as a reliable indicator of workload in a scenario where outside factors are not carefully controlled. Further complicating the issue, many physiological responses are highly task dependent or dependent upon the individual characteristics of the user. A more rigorous understanding of these physiological correlations and user cognitive states, as well as their sources of variation, are necessary.

### EXAMPLE: PERSISTENT SURVEILLANCE MISSION Numerical Simulations

The decision-support system designed in the previous sidebars is now illustrated through a numerical example. The sample surveillance mission involves four regions. The matrix of travel times (given in minutes) between the regions is

$$D = \begin{bmatrix} 0.0 & 22.1 & 34.5 & 9.0 \\ 22.1 & 0.0 & 19.3 & 14.6 \\ 34.5 & 19.3 & 0.0 & 25.6 \\ 9.0 & 14.6 & 25.6 & 0.0 \end{bmatrix}.$$

The time to collect information at each region is 10 minutes. It is assumed that the performance of the operator is the same at each region, and that the importance of each region is equal to that of all other regions. Let the drift rate in the DDM associated with the operator be  $\mu = -0.3$  for a nonanomalous region and  $\mu = +0.3$  for an anomalous region. Let the diffusion rate for the DDM associated with the operator be  $\sigma = 1$ . Suppose regions  $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ , and  $\mathcal{R}_4$  become anomalous at time instants 20, 80, 140, and 200 minutes, respectively.

The optimization problem (8) is solved before processing each task to determine the optimal time allocations for the human operator. A sample evolution of the system is shown Figure S2. For simplicity, it is assumed that the human operator allocates precisely the suggested amount of time to each task. The exact arrival time of each task is dictated by the region selection policy, information collection times, and UAV travel times. In this example, the average rate of arrival over all analysis tasks is 0.125 tasks per minute (1 task every 8 minutes). Note that the actual rate of arrival is non-uniform and thus varies over the course of the mission.

Note that the algorithm keeps the queue length close to unity. The queue length increases only if there is a high likelihood of anomaly at some region. Once an anomaly is detected, the allocation policy drops pending tasks in the queue until only one task remains. For this example, the threshold for the CUSUM algorithm is chosen equal to 4, and once an anomaly is detected the CUSUM statistic (shown in Figure S2c) resets to zero. Under the routing policy designed in the previous sidebars, with high probability, the UAV selects a region with a high probability of being anomalous, which is illustrated by the close correlation between the CUSUM statistics and the region selection probability (Figures S2c and S2d).

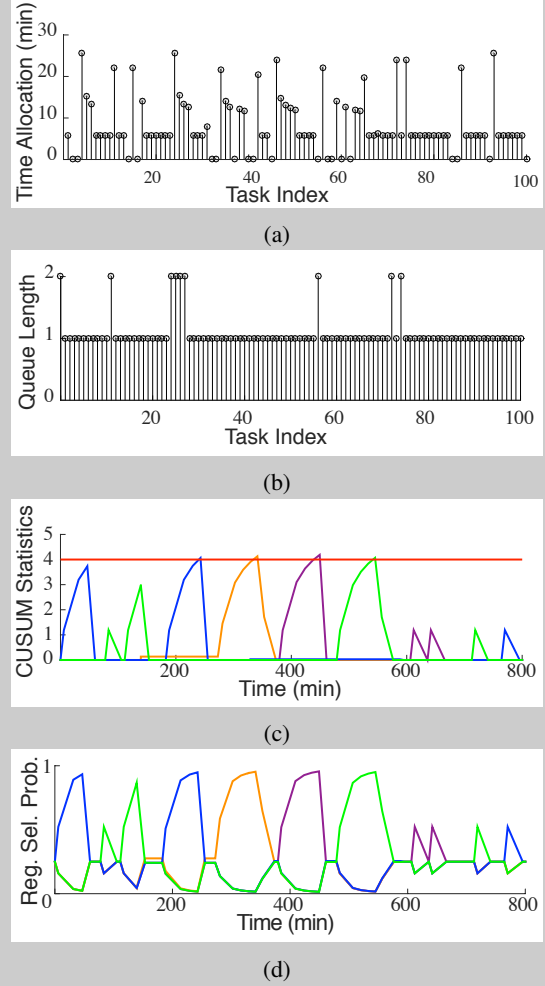


Figure S2: A typical evolution of the decision support system: (a) Suggested time allocations to each task as a function of the task index; (b) Queue length as a function of the task index; (c) Value of the CUSUM statistics associated with regions  $\mathcal{R}_1$  (blue),  $\mathcal{R}_2$  (orange),  $\mathcal{R}_3$  (purple), and  $\mathcal{R}_4$  (green) as a function of time, plotted along with the CUSUM decision threshold (red); (d) Region selection probabilities for regions  $\mathcal{R}_1$  (blue),  $\mathcal{R}_2$  (orange),  $\mathcal{R}_3$  (purple), and  $\mathcal{R}_4$  (green) as a function of time. Note that tasks arrive to the queue at an average rate of 1 task every 8 minutes.

### Graphical User Interface Design

Graphical User Interfaces (GUIs) can have a drastic impact on operator performance and the overall functionality of a system [178]. As such, the issue of designing effective user interfaces has been studied in a variety of contexts including computer science, marketing [179], human factors [180], psychology [181], and engineering [182]. Researchers have studied a wide range of aspects of the interface design problem and its impact on operator performance, including luminous-

ity [183], stimulus specifications [184], interactive window characteristics [185], and many many others.

The standard means of evaluating GUIs is through usability surveys. Some surveys, such as the *System Usability Scale* survey [186], [187], have been studied extensively, and provide benchmark statistics. Other researchers have turned to objective measures of usability that are more directly catered to the particular application under consideration [188]. In the context of supervisory control, GUIs play a key role in the

success of any system design, and must be carefully tested before being employed in application.

### *Automation Bias and Operator Trust*

Many automated systems that are designed to aid operator performance rely on the use of opportunistic cues or suggestions in attempt to guide operator behavior. Such forms of indirect control are of no use unless the operator responds to them in a meaningful way. Indeed, if the operator never takes into account any of the automated suggestions, then the whole purpose of providing them is defeated. On the other end of the spectrum, if the operator always heeds the automated suggestions without question, the operator can become complacent and lose situational awareness, resulting in performance degradation. This phenomena, sometimes referred to as *automation bias*, has been studied extensively in attempt to understand its effects and the conditions under which it occurs [170], [189], [190].

Both operator reluctance to follow automated prompts and automation bias are related to the issue of operator trust in automation. This complex phenomena is hard to model, due to its dynamic nature and its inherent situational and interpersonal dependencies [191], [192]. However, a successful human supervisory control system must employ tactics to maintain adequate operator trust, while mitigating automation bias.

### *Individual Differences and Statistical Uncertainty*

For the sake of simplicity, most system designers create a single, general-purpose model of human cognitive processing, with the intention of using this model for all potential operators. However, different operators may have varying responses to a given system design, and ignoring the vast differences between individuals can greatly reduce accuracy in predicting behavior. Indeed, factors such as personality traits [192], past experiences [191], and even the operator's current mood [193] can all affect performance. Studies of individual differences seek to resolve these shortcomings by identifying the ways in which people with distinct attributes react to the same situations in unique ways [194]. For example, research has shown that the personality trait *extraversion* moderates the relationship between stress and performance. Those with lower extraversion are more resilient to periods of hypostress, and those with higher extraversion are more capable of handling hyperstress [195].

Recent research has been somewhat successful in identifying specific traits that make a significant difference in operator interactions with autonomous systems [192]. For instance, in the context of supervisory control, high spatial ability, attentional

control, and video gaming experience have all been shown to lead to better performance in some aspects of a multiple agent supervisory control mission [196]. Despite these results, the relationship between individual operator differences and performance remains complex due to task dependencies and environmental sensitivity.

In addition to uncertainties caused by individual differences, there is also, in general, a large amount of statistical uncertainty involved in estimating model and system parameters. Indeed, even for a particular operator, it may be difficult (or impossible) to precisely determine the ideal parameters for describing behavior. Errors due to such statistical uncertainties can get propagated through a system design, causing undesirable results. Some models partially address this issue by directly incorporating auxiliary noise variables. For example, the extended DDM [60], [73] and the full DDM [74] introduce noise terms associated with the drift rate variable. Other approaches to mitigating the effects of uncertainty focus on optimal statistical data-fitting techniques associated with particular types of models [197].

Further complicating the issue, many decision support systems hinge on statistical models that make assumptions about some underlying process. Violations to these assumptions create another source of error. Some researchers have explored techniques for relaxing standard statistical assumptions with regard to tasks that are commonly encountered in supervisory control. For example, techniques for introducing spatial correlations in search tasks have been explored in the multi-arm bandit problem [198]. However, as in any modeling application, some assumptions will generally be unavoidable.

Although some work does exist, a thorough sensitivity analysis and characterization of model uncertainty in the context of human supervisory control is still largely an open problem. It is clear that a more thorough understanding of system robustness with respect to individual differences, as well as statistical and modeling uncertainties, is needed before successful implementation in practice.

## CONCLUSIONS

Human supervisory control of robotic teams is an area that has attracted a significant amount of research attention in recent years, and will only continue to grow as sensor and robotic technology becomes more advanced. The unique set of challenges that this application brings about spans many disciplines, including control systems, human factors, and psychology. In a broad sense the human supervisory control problem can be broken down into three components: the human, the autonomous agents, and the interface between them. In surveying each of these components and discussing examples of relevant theory for each, it becomes apparent

that well-studied tools from different scientific disciplines can work in conjunction with one another to create systems which have the potential to drastically increase productivity and efficiency in a given application. While many challenges still remain, continued collaboration among scientific disciplines will allow the maturation of future human supervisory control technology.

## ACKNOWLEDGMENTS

The authors would like to thank Casimir Ludwig from the University of Bristol, as well as P. Holmes and M. Shvartsman from Princeton University, for their insight and feedback regarding human modeling.

This work has been supported by the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

## REFERENCES

- [1] W. M. Bulkeley. Chicago's camera network is everywhere. *The Wall Street Journal*, November 17, 2009.
- [2] C. Drew. Military taps social networking skills. *The New York Times*, June 7, 2010.
- [3] T. Shanker and M. Richtel. In new military, data overload can be deadly. *The New York Times*, January 16, 2011.
- [4] E. Guizzo. Obama commanding robot revolution announces major robotics initiative. *IEEE Spectrum*, June 2011.
- [5] T. B. Sheridan. *Telerobotics, automation, and human supervisory control*. MIT press, 1992.
- [6] M. Lind. Plant modelling for human supervisory control. *Transactions of the Institute of Measurement and Control*, 21(4-5):171–180, 1999.
- [7] M. A. Goodrich and M. L. Cummings. Human factors perspective on next generation unmanned aerial systems. *Handbook of Unmanned Aerial Vehicles*, pages 2405–2423, 2015.
- [8] S. R. Dixon and C. D. Wickens. Automation reliability in unmanned aerial vehicle control: A reliance-compliance model of automation dependence in high workload. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 48(3):474–486, 2006.
- [9] R. Parasuraman, M. Barnes, K. Cosenzo, and S. Mulgund. Adaptive automation for human-robot teaming in future command and control systems. Technical report, Defense Technical Information Center (DTIC), 2007. Available at <http://www.dtic.mil/dtic>.
- [10] G. S. Taylor, L. E. Reinerman-Jones, J. L. Szalma, M. Mouloua, and P. A. Hancock. What to automate: Addressing the multidimensionality of cognitive resources through system design. *Journal of Cognitive Engineering and Decision Making*, 7(4):311–329, 2013.
- [11] R. Parasuraman and V. Riley. Humans and automation: Use, misuse, disuse, abuse. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 39(2):230–253, 1997.
- [12] L. J. Sorensen, K. I. Overgaard, and T. J. S. Martinsen. Understanding human decision making during critical incidents in dynamic positioning. In S. Sharples and S. Shorrock, editors, *Contemporary Ergonomics and Human Factors 2014*, page 359. CRC Press, 2014. Proceedings of the International Conference on Ergonomics & Human Factors, Southampton, UK, 7-10 April 2014.
- [13] D. C. Klein. *Using Adaptive Automation to Increase Operator Performance and Decrease Stress in a Satellite Operations Environment*. PhD thesis, Colorado Technical University, 2014.
- [14] V. A. Banks, N. A. Stanton, and C. Harvey. Sub-systems on the road to vehicle automation: Hands and feet free but not mind free driving. *Safety Science*, 62:505–514, 2014.
- [15] B. Terwilliger, D. Vincenzi, D. Ison, K. Witcher, D. Thirtyacre, and A. Khalid. Influencing factors for use of unmanned aerial systems in support of aviation accident and emergency response. *Journal of Automation and Control Engineering*, 3(3), 2015.
- [16] US Air Force. Report on technology horizons, a vision for Air Force Science and Technology during 2010–2030. Technical report, AF/ST-TR-10-01-PR, United States Air Force. Retrieved from <http://www.af.mil/shared/media/document/AFD-100727-053.pdf>, 2010.
- [17] U. E. Franke. Drones, drone strikes, and us policy: The politics of unmanned aerial vehicles. *Parameters*, 44(1):121, 2014.
- [18] S. M. Astley. Evaluation of computer-aided detection (CAD) prompting techniques for mammography. *The British Journal of Radiology*, 78(suppl\_1):S20–S25, 2014.
- [19] W. Barfield and T. A. Dingus. *Human Factors in Intelligent Transportation Systems*. Psychology Press, 2014.
- [20] C. Nehme, B. Mekdeci, J. W. Crandall, and M. L. Cummings. The impact of heterogeneity on operator performance in futuristic unmanned vehicle systems. *The International C2 Journal*, 2(2):1–30, 2008.
- [21] L. F. Bertuccelli, N. Pellegrino, and M. L. Cummings. Choice modeling of relook tasks for UAV search missions. In *American Control Conference*, pages 2410–2415, Baltimore, MD, USA, June 2010.
- [22] L. F. Bertuccelli, N. W. M. Beckers, and M. L. Cummings. Developing operator models for UAV search scheduling. In *AIAA Conf. on Guidance, Navigation and Control*, Toronto, Canada, August 2010.
- [23] K. Savla, C. Nehme, T. Temple, and E. Frazzoli. On efficient cooperative strategies between UAVs and humans in a dynamic environment. In *AIAA Conf. on Guidance, Navigation and Control*, Honolulu, HI, USA, 2008.
- [24] K. Savla, T. Temple, and E. Frazzoli. Human-in-the-loop vehicle routing policies for dynamic environments. In *IEEE Conf. on Decision and Control*, pages 1145–1150, Cancún, México, December 2008.
- [25] J. W. Crandall, M. L. Cummings, M. Della Penna, and P. M. A. de Jong. Computing the effects of operator attention allocation in human control of multiple robots. *IEEE Transactions on Systems, Man & Cybernetics. Part A: Systems & Humans*, 41(3):385–397, 2011.
- [26] C. E. Nehme. *Modeling Human Supervisory Control in Heterogeneous Unmanned*

- Vehicle Systems*. PhD thesis, Department of Aeronautics and Astronautics, MIT, February 2009.
- [27] K. Savla and E. Frazzoli. Maximally stabilizing task release control policy for a dynamical queue. *IEEE Transactions on Automatic Control*, 55(11):2655–2660, 2010.
  - [28] K. Savla and E. Frazzoli. A dynamical queue approach to intelligent task management for human operators. *Proceedings of the IEEE*, 100(3):672–686, 2012.
  - [29] N. D. Powel and K. A. Morgansen. Multiserver queueing for supervisory control of autonomous vehicles. In *American Control Conference*, pages 3179–3185, Montréal, Canada, June 2012.
  - [30] V. Srivastava, R. Carli, C. Langbort, and F. Bullo. Attention allocation for decision making queues. *Automatica*, 50(2):378–388, 2014.
  - [31] V. Srivastava and F. Bullo. Knapsack problems with sigmoid utility: Approximation algorithms via hybrid optimization. *European Journal of Operational Research*, 236(2):488–498, 2014.
  - [32] M. Majji and R. Rai. Autonomous task assignment of multiple operators for human robot interaction. In *American Control Conference*, pages 6454–6459, Washington, DC, June 2013.
  - [33] V. Srivastava, A. Surana, and F. Bullo. Adaptive attention allocation in human-robot systems. In *American Control Conference*, pages 2767–2774, Montréal, Canada, June 2012.
  - [34] V. Srivastava, F. Pasqualetti, and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. *International Journal of Robotics Research*, 32(12):1438–1458, 2013.
  - [35] V. Srivastava, K. Plarre, and F. Bullo. Randomized sensor selection in sequential hypothesis testing. *IEEE Transactions on Signal Processing*, 59(5):2342–2354, 2011.
  - [36] K. Cosenzo, R. Parasuraman, and E. De Visser. Automation strategies for facilitating human interaction with military unmanned vehicles. In M. Barnes and F. Jentsch, editors, *Human-robot Interactions in Future Military Operations*, pages 103–124. Ashgate Publishing, 2010.
  - [37] M. Scerbo. Adaptive automation. In R. Parasuraman and M. Rizzo, editors, *Neuroergonomics: The Brain At Work*, pages 239–252. Oxford University Press, 2001.
  - [38] K. M. Feigh, M. C. Dorneich, and C. C. Hayes. Toward a characterization of adaptive systems a framework for researchers and system designers. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 54(6):1008–1024, 2012.
  - [39] P. R. Murphy, J. Vandekerckhove, and S. Nieuwenhuis. Pupil-linked arousal determines variability in perceptual decision making. *PLOS Computational Biology*, 10(9):e1003854, 2014.
  - [40] R. C. Atkinson, G. H. Bower, and E. J. Crothers. *Introduction to Mathematical Learning Theory*. John Wiley & Sons, 1965.
  - [41] F. Restle and J. G. Greeno. *Introduction to Mathematical Psychology*. Addison-Wesley, 1970.
  - [42] J. T. Townsend and F. G. Ashby. *Stochastic Modeling of Elementary Psychological Processes*. Cambridge University Press, 1983.
  - [43] L. Trotta, E. Bullinger, and R. Sepulchre. Global analysis of dynamical decision-making models through local computation around the hidden saddle. *PLOS ONE*, 7(3):e33110, 2012.
  - [44] J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1):3–71, 1988.
  - [45] M. S. C. Thomas and J. L. McClelland. Connectionist models of cognition. In R. Sun, editor, *Cambridge Handbook of Computational Cognitive Modelling*, pages 23–58. Cambridge University Press, 2008.
  - [46] R. Sun. Artificial intelligence: Connectionist and symbolic approaches. In N. J. Smelser and P. B. Baltes, editors, *International Encyclopedia of the Social and Behavioral Sciences*, pages 783–789. Elsevier, 2001.
  - [47] J. R. Anderson. ACT: A simple theory of complex cognition. *American Psychologist*, 51(4):355, 1996.
  - [48] J. E. Laird. *The Soar Cognitive Architecture*. MIT Press, 2012.
  - [49] D. E. Kieras and D. E. Meyer. An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, 12(4):391–438, 1997.
  - [50] R. Sun. The CLARION cognitive architecture: Extending cognitive modeling to social simulation. In R. Sun, editor, *Cognition and Multi-Agent Interaction*, pages 79–99. Cambridge University Press, 2006.
  - [51] P. Langley, J. E. Laird, and S. Rogers. Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10(2):141–160, 2009.
  - [52] L. van Maanen, H. van Rijn, and N. Taatgen. RACE/A: An architectural account of the interactions between learning, task control, and retrieval dynamics. *Cognitive Science*, 36(1):62–101, 2012.
  - [53] R. Ratcliff and P. L. Smith. A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2):333, 2004.
  - [54] P. Holmes. Some joys and trials of mathematical neuroscience. *Journal of Nonlinear Science*, 24(2):201–242, 2014.
  - [55] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen. The physics of optimal decision making: A formal analysis of performance in two-alternative forced choice tasks. *Psychological Review*, 113(4):700–765, 2006.
  - [56] J. I. Gold and M. N. Shadlen. The neural basis of decision making. *Annual Reviews of Neuroscience*, 30:535–574, 2007.
  - [57] U. Basten, G. Biele, H. R. Heekeren, and C. J. Fiebach. How the brain integrates costs and benefits during decision making. *Proceedings of the National Academy of Sciences*, 107(50):21767–21772, 2010.
  - [58] S. D. Brown and A. Heathcote. The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3):153–178, 2008.
  - [59] R. Ratcliff. A theory of memory retrieval. *Psychological Review*, 85(2):59–108, 1978.
  - [60] R. Ratcliff and J. N. Rouder. Modeling response times for two-choice decisions. *Psychological Science*, 9(5):347–356, 1998.
  - [61] R. Ratcliff and G. McKoon. The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922, 2008.

- [62] J. I. Gold and M. N. Shadlen. Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36(2):299–308, 2002.
- [63] R. Bogacz, P. T. Hu, P. J. Holmes, and J. D. Cohen. Do humans produce the speed–accuracy trade-off that maximizes reward rate? *The Quarterly Journal of Experimental Psychology*, 63(5):863–891, 2010.
- [64] P. Holmes and J. D. Cohen. Optimality and some of its discontents: Successes and shortcomings of existing models for binary decisions. *Topics in Cognitive Science*, 6(2):258–278, 2014.
- [65] S. Goldfarb, K. Wong-Lin, M. Schwemmer, N. E. Leonard, and P. Holmes. Can post-error dynamics explain sequential reaction time patterns? *Frontiers in Psychology*, 3, 2012.
- [66] D. Laming. Choice reaction performance following an error. *Acta Psychologica*, 43(3):199–224, 1979.
- [67] K. Bicknell and R. Levy. A rational model of eye movement control in reading. In *Annual Meeting of the Association for Computational Linguistics*, pages 1168–1178, Uppsala, Sweden, July 2010.
- [68] R. L. Lewis, M. Shvartsman, and S. Singh. The adaptive nature of eye movements in linguistic tasks: How payoff and architecture shape speed-accuracy trade-offs. *Topics in Cognitive Science*, 5(3):581–610, 2013.
- [69] P. L. Smith and R. Ratcliff. Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3):161–168, 2004.
- [70] G. E. Uhlenbeck and L. S. Ornstein. On the theory of the Brownian motion. *Physical Review*, 36(5):823, 1930.
- [71] M. Usher and J. L. McClelland. The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550, 2001.
- [72] A. R. Teodorescu and M. Usher. Disentangling decision models: From independence to competition. *Psychological Review*, 120(1):1, 2013.
- [73] R. D. Luce. *Response Times*. Oxford University Press, 1986.
- [74] M. Milosavljevic, J. Malmaud, A. Huth, C. Koch, and A. Rangel. The drift diffusion model can account for value-based choice response times under high and low time pressure. *Judgment and Decision Making*, 5(6):437–449, 2010.
- [75] R. Ratcliff and P. L. Smith. A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2):333, 2004.
- [76] A. K. Churchland, R. Kiani, and M. N. Shadlen. Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6):693–702, 2008.
- [77] D. Thura, J. Beauregard-Racine, C. W. Fradet, and P. Cisek. Decision making by urgency gating: Theory and experimental support. *Journal of Neurophysiology*, 108(11):2912–2930, 2012.
- [78] R. Moran. Optimal decision making in heterogeneous and biased environments. *Psychonomic Bulletin & Review*, 22(1):38–53, 2014.
- [79] P. Frazier and A. J. Yu. Sequential hypothesis testing under stochastic deadlines. In *Advances in Neural Information Processing Systems*, pages 465–472, Vancouver, BC, Canada, December 2007.
- [80] A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 1945.
- [81] W. Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. John Wiley & Sons, 1968.
- [82] A. Wald and J. Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, pages 326–339, 1948.
- [83] J. Neyman and E. S. Pearson. On the use and interpretation of certain test criteria for purposes of statistical inference: Part I. *Biometrika*, pages 175–240, 1928.
- [84] T. McMillen and P. Holmes. The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, 50(1):30–57, 2006.
- [85] R. J. Audley and A. R. Pike. Some alternative stochastic models of choice. *British Journal of Mathematical and Statistical Psychology*, 18(2):207–225, 1965.
- [86] C. Bundesen. Visual attention: Race models for selection from multielement displays. *Psychological Research*, 49(2-3):113–121, 1987.
- [87] C. W. Baum and V. V. Veeravalli. A sequential procedure for multihypothesis testing. *IEEE Transactions on Information Theory*, 40(6):1994–2007, 1994.
- [88] W. E. Hick. On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, 4(1):11–26, 1952.
- [89] I. Krajbich and A. Rangel. Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33):13852–13857, 2011.
- [90] R. B. Towal, M. Mormann, and C. Koch. Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proceedings of the National Academy of Sciences*, 110(40):E3858–E3867, 2013.
- [91] S. R. Hursh, D. P. Redmond, M. L. Johnson, D. R. Thorne, G. Belenky, T. J. Balkin, W. F. Storm, J. C. Miller, and D. R. Eddy. Fatigue models for applied research in warfighting. *Aviation, Space, and Environmental Medicine*, 75(3, Suppl.):A44–A53, 2004.
- [92] D. C. Rubin, S. Hinton, and A. Wenzel. The precise time course of retention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(5):1161–1176, 1999.
- [93] D. Gopher and E. Donchin. Workload: An examination of the concept. In K. R. Boff, L. Kaufman, and J. P. Thomas, editors, *Handbook of Perception and Human Performance*, volume 2, pages 1–49. John Wiley & Sons, 1986.
- [94] N. Moray. *Mental Workload: Its Theory and Measurement*. Plenum Publishing Corporation, 1979.
- [95] S. G. Hart and L. E. Staveland. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Advances in Psychology*, 52:139–183, 1988.
- [96] C. D. Wickens. Multiple resources and mental workload. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(3):449–455, 2008.
- [97] R. M. Yerkes and J. D. Dodson. The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology*, 18(5):459–482, 1908.
- [98] C. D. Wickens and J. G. Hollands. *Engineering Psychology and Human Performance*. Prentice Hall, 3 edition, 2000.
- [99] J. S. Warm, W. N. Dember, and P. A. Hancock. Vigilance and workload in automated systems. In *Automation and Human Performance: Theory and*

- Applications (Human Factors in Transportation)*, volume 1, pages 183–200. Lawrence Erlbaum Associates, Inc, 1996.
- [100] P. A. Hancock. A dynamic model of stress and sustained attention. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 31(5):519–537, 1989.
  - [101] G. Matthews, D. R. Davies, S. J. Westerman, and R. B. Stammers. *Human Performance: Cognition, Stress, and Individual Differences*. Psychology Press, 2000.
  - [102] M. M. Mallis, S. Mejdal, T. T. Nguyen, and D. F. Dinges. Summary of the key features of seven biomathematical models of human fatigue and performance. *Aviation, Space, and Environmental Medicine*, 75(3, Suppl.):A4–A14, 2004.
  - [103] H. Ebbinghaus. *Memory: A Contribution to Experimental Psychology*. Teachers College, Columbia University, 1913.
  - [104] J. R. Anderson and L. J. Schooler. Reflections of the environment in memory. *Psychological Science*, 2(6):396–408, 1991.
  - [105] D. B. Kaber and M. R. Endsley. Team situation awareness for process control safety and performance. *Process Safety Progress*, 17(1):43–48, 1998.
  - [106] M. Cummings, C. Nehme, J. Crandall, and P. Mitchell. Predicting Operator Capacity for Supervisory Control of Multiple UAVs. In J. Chahl, L. Jain, A. Mizutani, and M. Sato-Ilic, editors, *Innovations in Intelligent Machines*, volume 70 of *Studies in Computational Intelligence*, pages 11–37. Springer, 2007.
  - [107] Y. Uny Cao, A. S. Fukunaga, and A. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4(1):7–27, 1997.
  - [108] L. E. Parker. Multiple mobile robot systems. In B. Siciliano and O. Khatib, editors, *Springer Handbook of Robotics*, pages 921–941. Springer, 2008.
  - [109] R. C. Arkin. *Behavior-Based Robotics*. MIT Press, 1998.
  - [110] T. Arai, E. Pagello, and L. E. Parker. Guest editorial: Advances in multirobot systems. *IEEE Transactions on Robotics and Automation*, 18(5):655–661, 2002.
  - [111] C. T. Abdallah and H. G. Tanner. Complex networked control systems: Introduction to the special section. *IEEE Control Systems Magazine*, 27(4):30–32, 2007.
  - [112] J. O’Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.
  - [113] K. J. Obermeyer, A. Ganguli, and F. Bullo. Multi-agent deployment for visibility coverage in polygonal environments with holes. *International Journal on Robust and Nonlinear Control*, 21(12):1467–1492, 2011.
  - [114] T. S. Michael and V. Pinciu. Art gallery theorems for guarded guards. *Computational Geometry: Theory and Applications*, 26(3):247–258, 2003.
  - [115] J. King and D. Kirkpatrick. Improved approximation for guarding simple galleries from the perimeter. *Discrete & Computational Geometry*, 46(2):252–269, 2011.
  - [116] S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982. Presented at the 1957 Institute for Mathematical Statistics Meeting.
  - [117] R. Patel, P. Frasca, and F. Bullo. Centroidal area-constrained partitioning for robotic networks. In *ASME Dynamic Systems and Control Conference*, Stanford, CA, USA, October 2013.
  - [118] J. G. Carlsson. Dividing a territory among several vehicles. *INFORMS Journal on Computing*, 24(4):565–577, 2012.
  - [119] J. W. Durham, R. Carli, P. Frasca, and F. Bullo. Discrete partitioning and coverage control for gossiping robots. *IEEE Transactions on Robotics*, 28(2):364–378, 2012.
  - [120] M. Younis and K. Akkaya. Strategies and techniques for node placement in wireless sensor networks: A survey. *Ad Hoc Networks*, 6(4):621–655, 2008.
  - [121] H. V. Poor and O. Hadjiladis. *Quickest Detection*. Cambridge University Press, 2008.
  - [122] B. Chen and P. Willett. Detection of hidden Markov model transient signals. *IEEE Transactions on Aerospace and Electronic Systems*, 36(4):1253–1268, 2000.
  - [123] N. Sydney and D. A. Paley. Multivehicle coverage control for a nonstationary spatiotemporal field. *Automatica*, 50(5):1381–1390, 2014.
  - [124] E. Bakolas and P. Tsiotras. The Zermelo-Voronoi diagram: A dynamic partition problem. *Automatica*, 46(12):2059–2067, 2010.
  - [125] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems. *Proceedings of the IEEE*, 99(9):1482–1504, 2011.
  - [126] R. Graham and J. Cortés. Adaptive information collection by robotic sensor networks for spatial estimation. *IEEE Transactions on Automatic Control*, 57(6):1404–1419, 2012.
  - [127] A. Singh, A. Krause, C. Guestrin, and W. J. Kaiser. Efficient informative sensing using multiple robots. *Journal of Artificial Intelligence Research*, 34(2):707, 2009.
  - [128] N. E. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. Davis. Collective motion, sensor networks and ocean sampling. *Proceedings of the IEEE*, 95(1):48–74, 2007.
  - [129] N. E. Leonard, D. A. Paley, R. E. Davis, D. M. Fratantoni, F. Lekien, and F. Zhang. Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay. *Journal of Field Robotics*, 27(6):718–740, 2010.
  - [130] S. L. Smith, M. Schwager, and D. Rus. Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*, 28(2):410–426, 2012.
  - [131] D. E. Soltero, M. Schwager, and D. Rus. Generating informative paths for persistent sensing in unknown environments. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 2172–2179, Vilamoura, Portugal, October 2012.
  - [132] C. G. Cassandras, X. Lin, and X. C. Ding. An optimal control approach to the multi-agent persistent monitoring problem. *IEEE Transactions on Automatic Control*, 58(4):947–961, 2013.
  - [133] R. N. Smith, M. Schwager, S. L. Smith, B. H. Jones, D. Rus, and G. S. Sukhatme. Persistent ocean monitoring with underwater gliders: Adapting sampling resolution. *Journal of Field Robotics*, 28(5):714–741, 2011.
  - [134] J. Yu, M. Schwager, and D. Rus. Correlated orienteering problem and its application to informative path planning for persistent monitoring tasks. *arXiv preprint arXiv:1402.1896*, February 2014.
  - [135] J. Clark and R. Fierro. Cooperative hybrid control of robotic sensors for perimeter detection and tracking. In *American Control Conference*, pages 3500–3505, Portland, OR, June 2005.



- [136] D. B. Kingston, R. S. Holt, R. W. Beard, T. W. McLain, and D. W. Casbeer. Decentralized perimeter surveillance using a team of UAVs. In *AIAA Conf. on Guidance, Navigation and Control*, San Francisco, CA, August 2005.
- [137] S. Susca, S. Martínez, and F. Bullo. Monitoring environmental boundaries with a robotic sensor network. *IEEE Transactions on Control Systems Technology*, 16(2):288–296, 2008.
- [138] Y. Elmaliach, A. Shiloni, and G. A. Kaminka. A realistic model of frequency-based multi-robot polyline patrolling. In *International Conference on Autonomous Agents*, pages 63–70, Estoril, Portugal, May 2008.
- [139] J. Le Ny and G. J. Pappas. On trajectory optimization for active sensing in Gaussian process models. In *IEEE Conf. on Decision and Control*, pages 6286–6292, Shanghai, China, December 2009.
- [140] X. Lan and M. Schwager. Planning periodic persistent monitoring trajectories for sensing robots in Gaussian random fields. In *IEEE Int. Conf. on Robotics and Automation*, pages 2415–2420, Karlsruhe, Germany, May 2013.
- [141] F. Pasqualetti, A. Franchi, and F. Bullo. On cooperative patrolling: Optimal trajectories, complexity analysis and approximation algorithms. *IEEE Transactions on Robotics*, 28(3):592–606, 2012.
- [142] F. Pasqualetti, J. W. Durham, and F. Bullo. Cooperative patrolling via weighted tours: Performance analysis and distributed algorithms. *IEEE Transactions on Robotics*, 28(5):1181–1188, 2012.
- [143] A. Almeida, G. Ramalho, H. Santana, P. Tedesco, T. Menezes, V. Corruble, and Y. Chevaleyre. Recent advances on multi-agent patrolling. In *Advances in Artificial Intelligence*, volume 3171 of *Lecture Notes in Computer Science*, pages 474–483. Springer, 2004.
- [144] V. Srivastava, P. Reverdy, and N. E. Leonard. Surveillance in an abruptly changing world via multiarmed bandits. In *IEEE Conf. on Decision and Control*, pages 692–697, Los Angeles, CA, December 2014.
- [145] K. Srivastava, D. M. Stipanović, and M. W. Spong. On a stochastic robotic surveillance problem. In *IEEE Conf. on Decision and Control*, pages 8567–8574, Shanghai, China, December 2009.
- [146] G. Mathew and I. Mezić. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240(4):432–442, 2011.
- [147] G. Mathew and I. Mezić. Spectral multiscale coverage: A uniform coverage algorithm for mobile sensor networks. In *IEEE Conf. on Decision and Control*, pages 7872–7877, Shanghai, China, December 2009.
- [148] G. Hoffmann and C. Tomlin. Mobile sensor network control using mutual information methods and particle filters. *IEEE Transactions on Automatic Control*, 55(1), 2010.
- [149] B. Grocholsky, J. F. Keller, V. Kumar, and G. J. Pappas. Cooperative air and ground surveillance. *IEEE Robotics and Automation Magazine*, 13(3):16–25, 2006.
- [150] C. M. Kreucher, A. O. Hero III, K. D. Kastella, and M. R. Morelande. An information-based approach to sensor management in large dynamic networks. *Proceedings of the IEEE*, 95(5):978–999, 2007.
- [151] J. P. Hespanha, H. J. Kim, and S. S. Sastry. Multiple-agent probabilistic pursuit-evasion games. Technical report, Electrical Engineering and Computer Science, University of California at Berkeley, August 1999. Available at <http://www.ece.ucsb.edu/hespanha/published>.
- [152] J. Grace and J. Baillieul. Stochastic strategies for autonomous robotic surveillance. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 2200–2205, Seville, Spain, December 2005.
- [153] T. Sak, J. Wainer, and S. Goldenstein. Probabilistic multiagent patrolling. In *Brazilian Symposium on Artificial Intelligence*, Advances in Artificial Intelligence, pages 124–133, Salvador, Brazil, 2008. Springer.
- [154] L. Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer, 2004.
- [155] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004.
- [156] R. Patel, P. Agharkar, and F. Bullo. Robotic surveillance and Markov chains with minimal weighted Kemeny constant. *IEEE Transactions on Automatic Control*, 2015. To appear.
- [157] R. Ratnam, J. Goense, and M. E. Nelson. Change-point detection in neuronal spike train activity. *Neurocomputing*, 52:849–855, 2003.
- [158] D. Bertsekas. Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11(4-5):310–334, 2005.
- [159] H. S. Chang and S. I. Marcus. Approximate receding horizon approach for Markov decision processes: Average reward case. *Journal of Mathematical Analysis and Applications*, 286(2):636–651, 2003.
- [160] J. Mattingley, Y. Wang, and S. Boyd. Receding horizon control: Automatic generation of high-speed solvers. *IEEE Control Systems Magazine*, 31(3):52–65, 2011.
- [161] R. Parasuraman, T. B. Sheridan, and C. D. Wickens. A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 30(3):286–297, 2000.
- [162] M. P. Eckstein. Visual search: A retrospective. *Journal of Vision*, 11(5):14, 2011.
- [163] D. Purves, G. J. Augustine, D. Fitzpatrick, L. C. Katz, A-S LaMantia, J. O. McNamara, and S. M. et. al. Williams. *Neuroscience*. Sinauer Associates, 2001.
- [164] K. Holmqvist, M. Nyström, R. Andersson, R. Dewhurst, H. Jarodzka, and J. Van de Weijer. *Eye tracking: A comprehensive guide to methods and measures*. Oxford University Press, 2011.
- [165] R. Rosenholtz, Y. Li, and L. Nakano. Measuring visual clutter. *Journal of Vision*, 7(2):17, 2007.
- [166] M. P. Eckstein, B. R. Beutter, and L. S. Stone. Quantifying the performance limits of human saccadic targeting during visual search. *Perception*, 30(11):1389–1402, 2001.
- [167] US Office of the Secretary of Defense. Unmanned aircraft systems (UAS) roadmap, 2005-2030, 2005.
- [168] L. Fern and R. J. Shively. A comparison of varying levels of automation on the supervisory control of multiple UASs. In *AUVSI's Unmanned Systems North America*, Washington, DC, USA, August 2009.
- [169] M. L. Cummings, L. F. Bertucelli, J. Macbeth, and A. Surana. Task versus vehicle-based control paradigms in multiple unmanned vehicle supervision by a single operator. *IEEE Transactions on Human-Machine Systems*, 44(3):353–361, 2014.
- [170] R. Parasuraman and D. H. Manzey. Complacency and bias in human use of

- automation: An attentional integration. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 52(3):381–410, 2010.
- [171] F. Vachon and S. Tremblay. What eye tracking can reveal about dynamic decision-making. *Advances in Cognitive Engineering and Neuroergonomics*, 11:157, 2014.
- [172] M. E. van Bochove, L. Van der Haegen, W. Notebaert, and T. Verguts. Blinking predicts enhanced cognitive control. *Cognitive, Affective, & Behavioral Neuroscience*, 13(2):346–354, 2013.
- [173] E. Haapalainen, S. Kim, J. F. Forlizzi, and A. K. Dey. Psycho-physiological measures for assessing cognitive load. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, pages 301–310. ACM, 2010.
- [174] G. Durantin, J.-F. Gagnon, S. Tremblay, and F. Dehais. Using near infrared spectroscopy and heart rate variability to detect mental overload. *Behavioural Brain Research*, 259:16–23, 2014.
- [175] F. G. Freeman, P. J. Mikulka, L. J. Prinzel, and M. W. Scerbo. Evaluation of an adaptive automation system using three eeg indices with a visual tracking task. *Biological Psychology*, 50(1):61–76, 1999.
- [176] G. Taylor, L. Reinerman-Jones, K. Cosenzo, and D. Nicholson. Comparison of multiple physiological sensors to classify operator state in adaptive automation systems. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, pages 195–199, San Francisco, CA, USA, 2010.
- [177] W. W. Tryon. Pupillometry: A survey of sources of variation. *Psychophysiology*, 12(1):90–93, 1975.
- [178] J. T. Hackos and J. Redish. *User and Task Analysis for Interface Design*. John Wiley & Sons, 1998.
- [179] S. Djasasbi, M. Siegel, and T. Tullis. Generation Y, web design, and eye tracking. *International Journal of Human-Computer Studies*, 68(5):307–323, 2010.
- [180] N. A. Stanton, P. M. Salmon, L. A. Rafferty, G. H. Walker, and C. Baber. *Human Factors Methods: A Practical Guide for Engineering and Design*. Ashgate Publishing, 2013.
- [181] R. Yan. Study of the user-involvement iterative computer game interface design method. *Review in Psychology Research*, 2(1), 2013.
- [182] J. Nielsen. *Usability Engineering*. Elsevier, 1994.
- [183] Y. Li, S. Bahn, C. S. Nam, and J. Lee. Effects of luminosity contrast and stimulus duration on user performance and preference in a P300-based brain-computer interface. *International Journal of Human-Computer Interaction*, 30(2):151–163, 2014.
- [184] J. Flach, P. Venero, T. Hughes, and A. J. Rowe. Supervisory control information management research (SCIMR). Technical report, Defense Technical Information Center (DTIC), June 2013. Available at <http://www.dtic.mil/dtic>.
- [185] J. Scarr, A. Cockburn, C. Gutwin, and S. Malacria. Testing the robustness and performance of spatially consistent interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3139–3148, 2013.
- [186] J. Brooke. SUS-A quick and dirty usability scale. *Usability Evaluation in Industry*, 189:194, 1996.
- [187] J. Sauro and J. R. Lewis. *Quantifying the User Experience: Practical Statistics for User Research*. Elsevier, 2012.
- [188] M. L. Cummings and B. Donmez. Metrics for supervisory control system evaluation. *The Oxford Handbook of Cognitive Engineering*, page 367, 2013.
- [189] M. L. Cummings. Automation bias in intelligent time critical decision support systems. In *AIAA 1st Intelligent Systems Technical Conference*, volume 2, pages 557–562, 2004.
- [190] K. Goddard, A. Roudsari, and J. C. Wyatt. Automation bias: Empirical results assessing influencing factors. *International Journal of Medical Informatics*, 83(5):368–375, 2014.
- [191] R. R. Hoffman, M. Johnson, J. M. Bradshaw, and A. Underbrink. Trust in automation. *IEEE Intelligent Systems*, 28(1):84–88, 2013.
- [192] J. L. Szalma and G. S. Taylor. Individual differences in response to automation: The five factor model of personality. *Journal of Experimental Psychology: Applied*, 17(2):71–96, 2011.
- [193] G. R. J. Hockey. Changes in operator efficiency as a function of environmental stress, fatigue, and circadian rhythms. In K. R. Boff, L. Kaufman, and J. P. Thomas, editors, *Handbook of Perception and Human Performance, Vol. 2: Cognitive Processes and Performance*, pages 1–49. John Wiley & Sons, 1986.
- [194] J. L. Szalma. Individual differences in human-technology interaction: Incorporating variation in human characteristics into human factors and ergonomics research and design. *Theoretical Issues in Ergonomics Science*, 10(5):381–397, 2009.
- [195] G. Matthews, D. R. Davies, and J. L. Lees. Arousal, extraversion, and individual differences in resource availability. *Journal of Personality and Social Psychology*, 59(1):150, 1990.
- [196] J. Y. C. Chen and M. J. Barnes. Supervisory control of multiple robots effects of imperfect automation and individual differences. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 54(2):157–174, 2012.
- [197] R. Ratcliff and F. Tuerlinckx. Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9(3):438–481, 2002.
- [198] P. B. Reverdy, V. Srivastava, and N. E. Leonard. Modeling human decision making in generalized Gaussian multiarmed bandits. *Proceedings of the IEEE*, 102(4):544–571, 2014.