An Evaluation of Stochastic Model-Dependent and Model-Independent Glider Flight Management

Amor A. Menezes, *Senior Member, IEEE*, Dhaval D. Shah, and Ilya V. Kolmanovsky, *Fellow, IEEE*

***Abstract*—This paper develops and compares two stochastic strategies that manage glider flight in a dynamic environment, to tackle an open problem concerning the necessity, applicability, complexity, and validity of using an environment model when making flight management decisions. Strategy performance is compared on two surveillance and pursuit tasks that require optimal control: the maximization of expected range while maintaining altitude within given limits, and the maximization of expected range while following a moving ground vehicle within a prescribed distance and maintaining altitude within given limits. Both tasks involve flight management decisions of glider airspeed and the time to spend climbing in a randomly encountered thermal. The first strategy is stochastic drift counteraction optimal control (SDCOC), a dynamic programming method that relies on a model. Here, SDCOC uses a computationally-inexpensive yet accurate environment model that reflects transition probabilities between updrafts and downdrafts as well as thermal locations and strengths based on existing glider flight data. The second strategy is selective evolutionary generation, a model-free Markov chain Monte Carlo method that is not subject to the curse of dimensionality, efficiently searches for an optimal control in an online and tunable fashion, and easily adapts to a dynamic environment. However, this strategy requires a tolerance for learning and decision-space exploration. Both strategies perform satisfactorily, and showcase an environment/decision-space exploitation-exploration trade-off.**

***Index Terms*—Dynamic programming, dynamic uncertain environments, efficient search, glider unmanned aerial vehicle (UAV) flight, optimal control.**

I. Introduction

A. Background

FLIGHT management decisions to maximize gliding range are required during normal and competition gliding flight. For powered aircraft, gliding range maximization is necessary after engine failure [1] and to reduce the weight of military unmanned aerial vehicle (UAV) power plants [2]. Gliding range maximization is also important for glider UAVs, which are a noiseless means of conducting ground vehicle surveillance and pursuit for border patrol applications. Glider UAV feasibility has been experimentally demonstrated in [3].

The glider flight decision-making process is complicated by a dynamic and uncertain environment [4], due to a lack of *a priori* knowledge about the location and strength of updrafts, thermals, and downdrafts. Updrafts are caused by the upward deflection of the prevalent wind over raised ground. Thermals are columns of rising air at lower altitudes of the Earth's atmosphere that are caused by uneven heating of the ground surface. Downdrafts (or sink) occur on the leeward side of raised ground deflecting the prevalent wind, and when air flows to the base of a nearby thermal. Exploiting the lift produced by updrafts and thermals and avoiding or minimizing the time spent in downdrafts extends flight time and range. This extension involves decisions (e.g., reducing airspeed, or ignoring a weak thermal) that can be optimized with atmospheric models. However, modeling the environment to predict the locations and intensity of updrafts, downdrafts, and thermals—which may be time-varying because of changing wind speed and direction, time of day, and weather—is a complex task.

Hence, four concerns affect the formulation of glider flight management strategies: (1) the necessity of an environment model based on the desirability of flight management decision optimality; (2) the tailoring of a chosen strategy to an environment model or vice-versa; (3) the requisite fidelity, complexity, and added computational cost of an environment model, with implications on practicality for real-time deployment; and (4) the validity of a strategy that harnesses one comprehensive environment model when separate models exist for individual environment conditions. There is a wealth of literature (Section I-B) on each of these four concerns, but there are few works that provide comparative strategy insights by tackling all concerns simultaneously. In this paper, for two tasks of glider range maximization for surveillance and pursuit, a model-dependent and a model-independent stochastic flight management strategy are each developed and compared, to provide insight into what amounts to an environment and decision-space exploitation-exploration tradeoff for glider flight.

B. Related Literature

Soaring results exist for: dynamic soaring [5]–[13], where flight energy is provided by wind velocity gradients, orographic features, and wind gusts; static soaring [14]–[16], where flight energy is influenced by thermals, updrafts, and downdrafts; and trajectory generation and optimization [17]–[31]. Stochastic optimal control approaches to glider flight

Manuscript received February 1, 2016; revised December 19, 2016; accepted February 20, 2017. Manuscript received in final form March 8, 2017. The work of I. V. Kolmanovsky was supported by the National Science Foundation under award ECCS 1404814. Recommended by Associate Editor Y. Ebihara. *(Corresponding author: Amor A. Menezes.)*

A. A. Menezes is with the California Institute for Quantitative Biosciences, University of California, Berkeley, CA 94704-5230 USA (e-mail: amenezes@berkeley.edu).

D. D. Shah is with Powertrain Controls, FCA North America, Auburn Hills, MI 48326 USA (email: dhaval.d.shah@fcagroup.com).

I. V. Kolmanovsky is with the Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI 48109-2140 USA (e-mail: ilya@umich.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCST.2017.2692723

management have been described in [4] (using the theory of [32]) and in [33]. A model-independent approach to guidance that utilizes online measurements for optimizing performance in stochastic wind (albeit for a powered UAV) is presented in [34], while [35] balances exploitation and exploration by generating energy-gain paths via a model-free wind map that is maintained and improved with in-flight observations.

Flight environment modeling is reviewed in [36], and additional related literature is cited here for completeness. Traditional techniques to model updrafts and downdrafts treat a given region as a fixed 3-D volume and calculate the interacting fluid flow at each point in that volume subject to boundary conditions. These calculations involve solving the Navier-Stokes equations within the volume, which is a task that is computationally expensive and difficult to implement within a stochastic control scheme. Previous work on topographic flow can also be found in [4], [37]–[40]. These techniques solve a fluid flow problem and estimate the horizontal component of wind velocity over a hill. For the optimal control tasks of interest in this paper, the vertical component of wind velocity needs to be estimated. The present work utilizes a simple model from [41] that is validated and shown to predict this component with good accuracy.

Thermals are difficult to model due to their somewhat random appearance over various land features (such as roads, cities, and ploughed fields), and due to their variable intensities that are dependent on the time of day and weather conditions (e.g., thermals are more common on bright sunny days than overcast days). Traditionally, thermal intensity is modeled empirically using information obtained from flight sensors and environment conditions, or by observing birds [15], [42]–[46]. However, the thermals predicted by these models do not incorporate thermal intensity changes as a function of wind speed. Further, these models do not predict the presence of sink caused by thermals. In this paper, estimates of thermal and sink intensity are utilized that account for these factors.

C. Technical Approach

This paper develops two stochastic strategies to manage glider airspeed and thermalling time for two tasks: the maximization of expected range while maintaining altitude within given limits, and the maximization of expected range while following a moving ground vehicle within a prescribed distance and maintaining altitude within given limits. The first strategy is stochastic drift counteraction optimal control (SDCOC) [47], which maximizes a cost functional reflective of expected time-to-violate specified constraints or expected total yield before violation of specified constraints while in the presence of disturbances that cause system states to drift. Such an optimal control law may be viewed as providing drift counteraction; hence, it is referred to as SDCOC. This control framework is of a dynamic programming type [48], but differs from average cost or discounted factor approaches as indicated in Section III-A.

The SDCOC method directly utilizes aggregated information about the glider flight environment through a transition probability model. These transition probabilities predict the conditional probabilities of updraft, downdraft, and thermal strength in the next flight segment given measured updraft, downdraft, and thermal strength in the current flight segment. Such transition probabilities provide a simple approach that captures aggregate information about the flight environment for best on-average performance.

The requisite transition probabilities are constructed for a given region by adapting an approach from [41] to model updrafts and downdrafts based on a representation of topographic lift. The updraft or downdraft intensity due to the land topography is simulated by sampling the ground elevation of the terrain predominantly upwind of the current position of the aircraft. The method is computationally simple, and, as we establish by comparing topographic lift to measured updraft data, is remarkably accurate. This updraft model can then be used to determine the location and intensity of thermals from previously recorded flight data. The resultant transition probabilities for updrafts and thermals can be utilized by SDCOC or by any other stochastic controller for glider flight management in an uncertain environment in the given region.

The second strategy is selective evolutionary generation systems (SEGSs) [49], a biologically inspired [50] Markov chain Monte Carlo method that can be search-efficient [51]. The method uses online probabilistic optimization to cope with uncertainty and perturbations in an external environment, and is essentially memoryless and therefore cost-effective and flexible in the face of change while regulating an output.

A SEGS strategy for glider flight management eschews an internal model. Instead, it heavily utilizes sensed feedback from the glider flight environment, feedback that can be emulated in simulation by deploying a model similar to the one used for SDCOC. This model is external to the SEGS controller and is only used as a flight environment substitute in simulations. For instance, this model can be easily reconfigured, changed, or adapted during the simulated flight while retaining the same SEGS controller, which will evolve the best action in a stochastic environment over time.

This paper deploys SDCOC and SEGS to separately make control decisions during simulated glider flight in the same, simulated, stochastic flight environment. Thus, this paper is able to compare the range maximization performance of a model-dependent approach (SDCOC) and a model-independent approach (SEGS) for glider flight management.

D. Original Contributions

Accordingly, the original contributions of this paper are as follows:

1) a further development and extension of the approach in our preliminary conference paper [52] to account for the modeling results of our preliminary conference paper [53], which leads to more realistic SDCOC policies and their evaluation for glider flight management;
2) the implementation and evaluation of SEGS as a glider flight management strategy;
3) a performance comparison of the SDCOC and SEGS glider flight management strategies;
4) additional details, discussions, and interpretations not found in [52] and [53].

E. Paper Outline

Section II specifies the two tasks for strategy evaluation. Section III outlines applicable SDCOC and SEGS theory. Section IV highlights and validates a flight environment model. Section V discusses SDCOC and SEGS results obtained for an exemplar glider. Section VI presents conclusions.

The units in this paper follow standard aeronautical practice, and we provide equivalent SI units for any imperial units that we use. We cite altitude in feet and airspeed in (statute) miles per hour because such units are used in our exemplar glider's instruments [54] and polar curve [55].

II. Task Formulation

A. Range Maximization

We assume that the glider flight path is partitioned into segments, and we let Δs denote the distance of a single flight segment. In each flight segment, the glider can encounter a thermal with strength (altitude increase rate) w_2, and the glider can spend u_2 seconds climbing this thermal. We also assume that the glider flies each flight segment with longitudinal speed u_1 when outside the thermal. This results in the altitude change rate $(-f(u_1)+w_1)$, where $-f(u_1)$ is the polar curve for the sink rate of the glider in still air given glider longitudinal speed u_1, and w_1 denotes the time rate of change of the altitude due to the updraft (or downdraft, if w_1 is negative).

Equations that approximately model the glider's flight are

$$h^+ = h + (-f(u_1)+w_1)\frac{\Delta s}{u_1} + w_2 u_2, \tag{1}$$

$$t^+ = t + u_2 + \frac{\Delta s}{u_1}, \tag{2}$$

where h is the altitude at the start of a flight segment, t is the total time prior to the start of the flight segment, h^+ is the altitude at the end of the flight segment, and t^+ is the total time at the end of the flight segment. The variables u_1 and u_2 are control variables. The variables w_1 (updraft strength) and w_2 (thermal strength) are determined by flight and weather conditions, and are typically unknown *a priori*. We employ a Markov chain model to describe the evolution of thermal strength and updraft strength along the flight path. Specifically, we assume that a prediction of the probability distribution of thermal and updraft strength in the next flight segment can be made once a thermal and updraft of certain strength in a given flight segment is encountered (Sections II-C and IV).

The objective of the stochastic range maximization task is to determine a control law that maximizes the expected distance that the glider can travel within a given time, t_{max}, i.e., the expected distance that the glider can travel before the system states exit a prescribed set

$$G = \{(h,t):\ h_{min} \le h \le h_{max},\ t \le t_{max}\}. \tag{3}$$

The constraint $h \ge h_{min}$ is imposed because the glider has to execute a landing maneuver upon descending to the minimum altitude h_{min}. The altitude of the glider should also not exceed the maximum altitude, h_{max}, which is typically determined by oxygen requirement laws or airspace ceiling restrictions.

B. Ground Vehicle Surveillance

In this task, the glider must fly so that a moving ground vehicle target is kept under observation. The formulation of this task is similar to the formulation of the range maximization task for the glider. The update equations are

$$h^+ = h + (-f(u_1)+w_1)\frac{\Delta s}{u_1} + w_2 u_2, \tag{4}$$

$$d^+ = d + w_3\left(u_2 + \frac{\Delta s}{u_1}\right) - \Delta s, \tag{5}$$

where h is the altitude at the start of a flight segment, d is the relative distance to the ground vehicle at the start of the flight segment, h^+ is the altitude at the end of the flight segment, and d^+ is the relative distance to the ground vehicle at the end of the flight segment. The ground vehicle velocity, w_3, varies according to a Markov chain model.

The objective of the stochastic ground vehicle surveillance task is to determine a control law that maximizes the expected range before the system states exit a prescribed set

$$G = \{(h,d):\ h_{min} \le h \le h_{max},\ d_{min} \le d \le d_{max}\}, \tag{6}$$

where d_{min} and d_{max} are the minimum and maximum relative distances, respectively, at which surveillance can be conducted.

C. Markov Chain Uncertainty Modeling

A Markov Chain model is used to represent the evolution of w_1, w_2, and w_3. The transition probabilities are

$$p_{ij} = \Pr[w^+ \in W_j \mid w \in W_i], \tag{7}$$

with W_i, $i=1,\ldots,m$, defining a partitioning of the feasible range of the uncertainty. As an approximation, specific values of $w^i \in W_i$ may be associated with each W_i.

This Markov chain approach to stochastic uncertainty modeling permits a prediction of the distribution of updraft and thermal strength that may be encountered in the next flight segment given estimates of updraft and thermal strength in the current flight segment (Section IV). The approach also facilitates control laws that use stochastic dynamic programming.

Alternatively, a jump (Poisson process based) model can be employed to represent w_1 and w_2, while a Wiener process may be employed to represent w_3. Related SDCOC techniques from [56] can be applied to construct the optimal control law. The treatment of jump based models will be considered in future publications.

III. Control Strategies

A. Stochastic Drift Counteraction Optimal Control Law Construction

Employing the same notations as [56], we consider a class of control problems for discrete-time nonlinear systems with stochastic disturbances

$$x(t+1) = f(x(t), v(t), w(t)), \tag{8}$$

where $x(t)$ is a vector state, $v(t)$ is a vector control, and $w(t)$ is a vector disturbance, the value of which is known at time instant $t \in \mathbb{Z}^+$ while future values are unknown. The control input must satisfy control constraints $v(t) \in U$ where U is a specified set, and additional state constraints of the form $(x(t), w(t)) \in G$. Unlike conventional controllers that respond to setpoints, here the set G is specified so that if $(x(t), w(t)) \in G$, then the system performance is acceptable.

The disturbance $w(t)$ is modeled by a Markov chain with a finite number of states, $w(t) \in W = \{w^j,\ j \in J\}$. The transition probability from $w(t) = w^i \in W$ to $w(t+1) = w^j \in W$ is denoted by $p_{ij} = \Pr[w(t+1) = w^j \mid w(t) = w^i]$. A more general case of state-dependent transition probabilities is considered in [47].

If the drift caused by the disturbance is large, then there may exist a time when constraint violation is unavoidable regardless of the control law. In such a case, an SDCOC law can be constructed [47] to maximize the expected time that the system operates without violating the imposed constraints. That is, our objective is to determine a control function $u(x, w)$, such that with $v(t) = u(x(t), w(t))$, the cost functional

$$J^{x_0,w_0,u} = \mathrm{E}_{\,x_0,w_0}[\tau^{x_0,w_0,u}(G)] \tag{9}$$

is maximized. Here, $\tau^{x_0,w_0,u}(G) \in \mathbb{Z}^+$ represents the first-time instant when the trajectory of $x(t)$ and $w(t)$, which is denoted by $\{x^u, w^u\}$, resulting from the application of the control $v(t) = u(x(t), w(t))$ with values in the set U exits the prescribed compact set G. Note that $\{x^u, w^u\}$ is a random process, $\tau^{x_0,w_0,u}(G)$ is a random variable, and $\mathrm{E}_{\,x_0,w_0}[\cdot]$ denotes the expectation conditioned on initial values of x and w, i.e., on $x(0) = x_0$ and $w(0) = w_0$.

Let $V(x, w)$ be the values of state-disturbance pairs. Given a state vector, x^-, and disturbance vectors, $w^-, w^+ \in W$, let

$$\begin{aligned} L^u V(x^-, w^-) &\triangleq \mathrm{E}_{\,x^-,w^-}\left[V\left(f\left(x^-, u\left(x^-, w^-\right), w^-\right), w^+\right)\right] \\ &\quad - V(x^-, w^-), && (10) \\ &= \sum_{j\in J} V\left(f\left(x^-, u\left(x^-, w^-\right), w^-\right), w^j\right) \cdot p_{i,j} \\ &\quad - V(x^-, w^-). && (11) \end{aligned}$$

The SDCOC law, $u_*(x, w)$, and its value function, $V(x, w)$, yield the following set of sufficient conditions [47]:

$$\begin{aligned} L^{u_*} V(x, w) + 1 = 0, &\text{ if } (x, w) \in G, \\ L^u V(x, w) + 1 \le 0, &\text{ if } (x, w) \in G,\ u \ne u_*, \\ V(x, w) = 0, &\text{ if } (x, w) \notin G. \end{aligned} \tag{12}$$

The conditions (12) can be treated numerically using either a value iteration approach or a linear programming approach. This paper uses a value iteration approach, wherein a sequence of functions that converge to the value function is defined through the following iterative process:

$$\begin{aligned} &V_0 \equiv 0, \\ &V_n(x, w^i) = \max_{v\in U}\left\{\sum_{j\in J} V_{n-1}\left(f\left(x, v, w^i\right), w^j\right) p_{ij} + 1\right\}, \\ &\qquad \text{if } (x, w^i) \in G,\ n > 0. \end{aligned} \tag{13}$$

This sequence of functions $\{V_n\}$ is monotonically non-decreasing and converges pointwise to $V_*(x, w^i) = J^{x,w^i,u_*}$ [47]; this convergence is uniform if J^{x,w^i,u_*} is continuous.

Value iterations (13) produce a sequence of value function approximations, V_n, at grid-points $x \in \{x^k,\ k \in K\}$. A stopping criterion is $|V_n(x, w^i) - V_{n-1}(x, w^i)| \le \epsilon$ for all $x \in \{x^k,\ k \in K\}$ and $i \in J$, where $\epsilon > 0$ is sufficiently small. In each iteration, once the values of V_{n-1} have been determined, linear or cubic interpolation can approximate $V_{n-1}(f(x^k, v^m, w^i), w^j)$ on the right-hand side of (13), where $v \in \{v^m,\ m \in M\}$ is a specified grid for v. Formally,

$$\begin{aligned} &V_0(x^k, w^i) \equiv 0, \\ &V_n(x^k, w^i) = \max_{v^m,\ m\in M}\left\{\sum_{j\in J} F_{n-1}\left(f\left(x^k, v^m, w^i\right), w^j\right) p_{ij} + 1\right\}, \end{aligned}$$

where $F_{n-1}(x, w^i) = \text{Interpolant}[V_{n-1}](x, w^i)$ if $(x, w^i) \in G$, and $F_{n-1}(x, w^i) = 0$ if $(x, w^i) \notin G$. Once an approximation of the value function, V_*, is available, an optimal control law may be determined from the following relation:

$$u_*(x, w^i) \in \operatorname{argmax}_{v\in U}\left\{\sum_{j\in J} V_*\left(f\left(x, v, w^i\right), w^j\right) p_{ij}\right\}. \tag{14}$$

In lieu of a grid-based approach, functional approximations may be employed to represent V_n in each stage as in the conventional approximate dynamic programming [57]. This approach may scale better to larger dimensional problems but may also reduce the solution performance.

SDCOC pseudocode for Section II-A follows. The states x_1 and x_2 have respective grid sizes k_1 and k_2, the disturbances w_1 and w_2 have respective grid sizes j_1 and j_2, and the control inputs u_1 and u_2 have respective grid sizes m_1 and m_2.

Require: $x_1ValuesVector$, k_1, $x_2ValuesVector$, k_2, $w_1ValuesVector$, j_1, $w_2ValuesVector$, j_2, $u_1ValuesVector$, m_1, $u_2ValuesVector$, m_2, x_1Min, x_1Max, x_2Min, x_2Max, $w_1ProbTransMatrix$, $w_2ProbTransMatrix$, $MaxIterations$, ϵ

Initialize $ValIterErrorMax \leftarrow \infty$, $V \leftarrow 0$, $V_0 \leftarrow 0$.
// V is the current value function, an array of size $k_1 \times k_2 \times j_1 \times j_2$.
// V_0 is a similarly sized array.
while ($IterationCounter < MaxIterations$ **and** $ValIterErrorMax > \epsilon$) **do**
  $V_{suc} \leftarrow V_0$.
  // V_{suc} is an updated value function.
  $ValIterErrorMax \leftarrow 0$.
  for $ctr1 = 1$ **to** k_1 **do**
    for $ctr2 = 1$ **to** k_2 **do**
      for $ctr3 = 1$ **to** j_1 **do**
        for $ctr4 = 1$ **to** j_2 **do**
          $x_1 \leftarrow x_1ValuesVector(ctr1)$,
          $x_2 \leftarrow x_2ValuesVector(ctr2)$,
          $w_1 \leftarrow w_1ValuesVector(ctr3)$,
          $w_2 \leftarrow w_2ValuesVector(ctr4)$.
          // The above obtains the current settings; next, the control actions have to be evaluated.
          $uValMax \leftarrow 0$.
          $uBestVector \leftarrow 0$.
          for $ctr5 = 1$ **to** m_1 **do**
            for $ctr6 = 1$ **to** m_2 **do**
              $u_1 \leftarrow u_1ValuesVector(ctr5)$,
              $u_2 \leftarrow u_2ValuesVector(ctr6)$.

15: Update the states via the dynamical model.
16: **if** the updated states violate any of the respective constraints: x_1Min, x_1Max, x_2Min, x_2Max **then**
17: $CostToGo \leftarrow 0$.
18: **else**
19: **for** $ctr7 = 1$ **to** j_1 **do**
20: **for** $ctr8 = 1$ **to** j_2 **do**
21: $CostToGo \leftarrow CostToGo + w_1ProbTransMatrix(ctr3, ctr7) * w_2ProbTransMatrix(ctr4, ctr8) *$ Interpolant
// Interpolant is the aforementioned interpolation (2D here) from V.
22: **end for**
23: **end for**
24: **end if**
25: $uValuation = CostToGo + 1$.
26: **if** $uValMax < uValuation$ **then**
27: $uBestVector \leftarrow u$,
$uValMax \leftarrow uValuation$.
28: **end if**
29: **end for**
30: **end for**
// Next, store the best actions and values.
31: $U(ctr1, ctr2, ctr3, ctr4) \leftarrow uBestVector$.
32: $V_{suc}(ctr1, ctr2, ctr3, ctr4) \leftarrow uValMax$.
33: $ValIterError \leftarrow |V_{suc}(ctr1, ctr2, ctr3, ctr4) - V(ctr1, ctr2, ctr3, ctr4)|$.
34: **if** $ValIterError > ValIterErrorMax$ **then**
35: $ValIterErrorMax \leftarrow ValIterError$.
36: **end if**
37: **end for**
38: **end for**
39: **end for**
40: **end for**
41: $V \leftarrow V_{suc}$.
42: $ValIterErrorStore \leftarrow ValIterErrorMax$.
43: $IterationCounter \leftarrow IterationCounter + 1$.
44: **end while**

The above pseudocode makes clear that the SDCOC algorithm has polynomial time complexity. If q is the largest value of k_1, k_2, j_1, j_2, m_1, and m_2, and if s is the sum of the number of states with the number of control inputs and with twice the number of disturbances, then the time complexity is $O(q^s)$, or $O(q^8)$ here. Such severe complexity and computational expense quickly become worse with added states or control inputs, and doubly so for every added disturbance. Further, any decrease in grid spacing (an increase in q) also has detrimental computational effects. But the burden of the flight environment model is separate; once the probability transition matrices for Line 21 have been determined (Section IV), any associated computational cost in the above pseudocode is minor.

All of the above pseudocode can be processed offline, with the state updates of Line 15 following from (1) and (2), and the probability transition matrices of Line 21 following from Section IV, itself offline. Therefore, the substantial SDCOC computational burden can be mitigated during flight through the lookup of precomputed policies. Hence, the online control of two states x_1 and x_2 that have initial conditions x_{10} and x_{20}, respectively, using SDCOC is only the following pseudocode.

Require: x_{10}, x_{20}, x_1Min, x_1Max, x_2Min, x_2Max.
1: **while** ($x_1 \leq x_1Max$ **and** $x_1 \geq x_1Min$ **and** $x_2 \leq x_2Max$ **and** $x_2 \geq x_2Min$) **do**
2: Obtain (sense) the current values of disturbances w_1 and w_2.
3: Lookup the precomputed u_1 and u_2 for the current quadruple (x_1, x_2, w_1, w_2).
4: Update the states via the dynamical model.
// Glide using current control settings for one flight segment.
5: **end while**

B. Selective Evolutionary Generation Overview

SEGS is a model-free technique that iteratively improves control inputs based on the results of perturbing them, retaining perturbed control inputs depending on the achieved performance (fitness). Starting from a discretization of a search space U of possible controls u^i, $1 \leq i \leq n$, where the suitability of each control is indicated by its fitness, $F : U \to \mathbb{R}^+$, the SEGS algorithm constructs a Markov chain with stochastic dynamics that has a row vector of stationary probabilities, $\boldsymbol{\pi} = \begin{bmatrix}\pi^1 & \pi^2 & \ldots & \pi^n\end{bmatrix}$, given by

$$\pi^i = \frac{F\left(u^i\right)^N}{\sum\limits_{k=1}^{n} F\left(u^k\right)^N}, \ 1 \leq i \leq n, \tag{15}$$

with $N \in [0, \infty)$ called the level of selectivity. The fitness evaluation may depend on states and disturbances that are not explicitly included in the list of arguments of F.

Equation (15) represents a probabilistic version of the optimization of an objective function. The constructed Markov chain will choose the control of maximum fitness with the highest stationary probability, and, in the limit as N approaches ∞, this probability is 1. That is, N tunes the concentration of the stationary probability distribution around the control of maximum fitness, and in the limit as N approaches ∞, the problem and solution then revert to one of standard, offline optimization, i.e., a delta function at an element of U.

But there is a tradeoff [50] to choosing N approaching ∞, especially in dynamic environments that vary F, variation that can be caused by disturbances. The tradeoff concerns responsiveness, which is the sensitivity of the stationary distribution (15) to changes in fitness. Responsiveness is a transient ability to adapt the search for an optimal control in the aftermath of an environment fluctuation. This property is important in a very dynamic environment, where F can be constantly perturbed.

Defining the extrinsic resilience of u^i to changes in the fitness of u^j, $j \neq i$, as $\rho_{ij} = \frac{\partial \pi^i}{\partial F(u^j)}$, and the intrinsic resilience of u^i to changes in its own fitness as $\rho_{ii} = \frac{\partial \pi^i}{\partial F(u^i)}$, the Markov chain is responsive if $\rho_{ij} \neq 0$ for all i and j. Reference [50] proves that $\rho_{ij}\Big|_{\substack{N=0 \\ j\neq i}} = \rho_{ii}\Big|_{N=0} = 0$, and that $\lim\limits_{\substack{N\to\infty \\ j\neq i}} \rho_{ij} = \lim\limits_{N\to\infty} \rho_{ii} = 0$, i.e., purely random optimization ($N = 0$) and standard offline optimization ($N \to \infty$) are both unresponsive. Because the expected hitting time of the fittest element from U also decreases with increasing N [50], a tradeoff exists between this expected hitting time and responsiveness, with the tradeoff controlled by N.

Thus, in this paper, we use finite N of low value ($N < 10$) for responsiveness, and choose $N > 1$ to improve the expected hitting time to the optimal control. As in [49]–[51], we assume that N is constant over the entire Markov chain.

To construct this Markov chain that attains $\boldsymbol{\pi}$, we use a SEGS, which is a quintuple $\Gamma = (U, R, P, G, F)$ where

- U is the set in which to search, $U = \{u^1, u^2, \ldots, u^n\}$;
- R is a "resource" set whose elements can be utilized to transition between elements of U, $R = \{r^1, r^2, \ldots, r^m\}$;
- $P : R \to (0, 1]$ is a probability mass function on R, given by $P(r^i) = \Pr[\mathcal{R} = r^i] = p^i$, $\sum_{k=1}^{m} p^k = 1$;
- $G : U \times R \to U$ is a mapping, called a "generation function," from one element of U to another using a resource from R;
- $F : U \to \mathbb{R}^+$ is a positive function that evaluates the fitness of elements of U;
- U is reachable [58] through G and R; and
- the dynamics of the system are given by

$$\mathcal{U}(t+1) = Select(\mathcal{U}(t), G(\mathcal{U}(t), \mathcal{R}(t)), N), \tag{16}$$

where $Select : U \times U \times [0, \infty) \to U$ is a random function such that if $u^1 \in U$ and $u^2 \in U$ are any two elements, and $N \in [0, \infty)$ is the level of selectivity, then

$$Select(u^1, u^2, N) = \begin{cases} u^1 & \text{with prob. } \frac{F(u^1)^N}{F(u^1)^N + F(u^2)^N}, \\ u^2 & \text{with prob. } \frac{F(u^2)^N}{F(u^1)^N + F(u^2)^N}. \end{cases} \tag{17}$$

In (16), $\mathcal{U}(t)$ denotes the realization of a random element from U at time t, $\mathcal{R}(t)$ denotes the realization of a random resource at time t, $G(\mathcal{U}(t), \mathcal{R}(t))$ denotes the outcome mapped from the realized element from U at time t utilizing the resource at time t, and $\mathcal{U}(0)$ has a known probability mass function. Also in (16), the probability of realizing an element from U at some future time given the present realization of an element from U is conditionally independent of the past time history of U element realizations. Thus, the dynamics form a discrete-time homogeneous Markov chain.

To highlight the capacity of this Markov chain to search for an optimal control despite a lack of knowledge about the states being controlled, let $x : U \to X$ be an unknown, computable (i.e., numerically evaluatable), and possibly changing function that we are interested in. The set X is a metric space. Suppose that we are given a desired element x_{des} in the image of x, and we wish to find $u \in U$ such that $||x(u) - x_{des}||$ is small (i.e., $x(u) \approx x_{des}$). Formally, we want a probability mass function ϕ_U (where $\phi_U : U \to \mathbb{R}^+$) for choosing control actions that help achieve a known expected value $Y \geq 0$, i.e.,

$$\mathrm{E}_{\phi_U}[||x(u) - x_{des}||] = Y. \tag{18}$$

In (18), Y is effectively a tolerance, i.e., it is the acceptable mean distance between candidates in the image of x compared with the desired image value. An example x_{des} for glider flight is the altitude midpoint between h_{min} and h_{max} in (3).

An ideal scheme to find ϕ_U should be efficient in that it trades off prior information about U for search effort savings as quickly as possible. SEGS is a scheme with the property that $\phi_U = \boldsymbol{\pi}$ [50]. Reference [51] discusses how a fitness function like

$$F(u^i) = e^{-\left(\left(x(u^i) - x_{des}\right)^2\right)} \tag{19}$$

with SEGS guarantees efficient search-based optimization. This fitness function feeds back current state desirability to the SEGS algorithm, but the dependence on states is not explicit. Hence, the SEGS controller can evolve the best action toward (15) regardless of whether the mapping x in (19) changes or is disturbed. In fact, the entire right-hand side of (19) can change; in this case, if the new fitness function is not exponential, then only search efficiency is lost, not optimal control convergence.

Section V designs exponential fitness functions for the two tasks in this paper. The design goal is to increase gliding range over standardized actuator settings (i.e., no active control).

SEGS pseudocode for Section II-A follows, where flight environment conditions are online or simulated by a sample realization of the probability transition matrices of Section IV, and these conditions and the glider's performance in them are fed back to the controller through a designed fitness function. Each algorithm iteration examines the fitness of a current control input and the fitness when that input is perturbed. A choice is then made using (16) to either keep the new perturbed input or to revert to the original input. Therefore, each SEGS iteration produces a control action.

Let the control inputs be u_1 and u_2 with grid spacings μ_1 and μ_2, respectively. Let the initial states be x_{10} and x_{20}, and the initial control inputs be u_{10} and u_{20}. Because there are two inputs and two directions in which to perturb each input, we use four resources. The realization of a random resource here represents a random choice of positively or negatively perturbing an input by that input's grid spacing.

Require: N, x_{10}, x_{20}, u_{10}, u_{20}, μ_1, μ_2, x_1Min, x_1Max, x_2Min, x_2Max, u_1Min, u_1Max, u_2Min, u_2Max.

while ($x_1 \leq x_1Max$ **and** $x_1 \geq x_1Min$ **and** $x_2 \leq x_2Max$ **and** $x_2 \geq x_2Min$) **do**
  Update the states via the dynamical model.
  // Glide using current control settings for one flight segment.
  Evaluate fitness F_{curr} using the new states.
  // Next, pick a control input and a direction to perturb it in.
  Choose a random natural number between 1 and 4 as resource.
  if the resource is 1 **then**
    if $u_1 + \mu_1 \leq u_1Max$ **then**
      $u_1 \leftarrow u_1 + \mu_1$.
    end if
    // u_2 is unchanged.
  else if the resource is 2 **then**
    if $u_1 - \mu_1 \geq u_1Min$ **then**
      $u_1 \leftarrow u_1 - \mu_1$.
    end if
    // u_2 is unchanged.
  else if the resource is 3 **then**
    if $u_2 + \mu_2 \leq u_2Max$ **then**
      $u_2 \leftarrow u_2 + \mu_2$.
    end if
    // u_1 is unchanged.
  else if the resource is 4 **then**
    if $u_2 - \mu_2 \geq u_2Min$ **then**
      $u_2 \leftarrow u_2 - \mu_2$.
    end if
    // u_1 is unchanged.
  end if
  Update the states via the dynamical model.
  // Glide using new control settings for one flight segment.
  Evaluate fitness F_{new} using the updated states.
  // Next choose which control settings to retain.
  $ProbNewControlSettingsSuccessful \leftarrow \frac{F_{new}^N}{F_{curr}^N + F_{new}^N}$.

Obtain sample value (1 / 0) of a Bernoulli random variable of success probability $ProbNewControlSettingsSuccessful$.
if sample value is 0 **then**
  Revert the control input perturbation.
end if
end while

The above pseudocode makes clear that the SEGS algorithm has time complexity $O(1)$, thereby avoiding SDCOC's curse of dimensionality. Thus, SEGS can easily handle added states, control inputs, and disturbances, as well as decreased grid spacings. Like SDCOC, the burden of a flight environment model is separate; here, however, the disturbance probability transition matrices are not explicitly included in SEGS control computation. Instead, a model realization of the atmosphere is buffered by the fitness function, and this realization can be replaced by the actual flight environment. Sensors (e.g., an altimeter) then provide the necessary fitness function data.

The entire SEGS controller can be processed online. Each iteration is memoryless, and hence cost-effective and flexible to changes in environment (be it reality or a model).

IV. Flight Environment Modeling

A. *Updraft and Downdraft Modeling*

We utilize the flight simulation technique of [41], which continuously takes five samples of the ground elevation (one below, three upwind, and one downwind of the aircraft) and calculates slopes. The ground elevation profile can be obtained over any land region using Google Earth. The calculated slopes are then combined with wind strength and direction to produce an empirical lift factor to be applied to the glider. Since we consider only low and moderately high wind speeds (less than 50 knots), we assume that the wind flow is balanced (i.e., that vortexes do not exist on the leeward side of raised ground). Therefore, we neglect the effect of vortexes on the estimated updraft value. Empirical distances at which to sample as well as updated slope calculation methods and slope adjustment functions are available in [41]; we utilize the technique's original methodology here to baseline the results. The relevant expressions are also described in [53].

These expressions were implemented and tested for a region near Warner Springs, CA. The elevation profile and computed updraft values for a single flight path over the region are in Fig. 1(a). The wind velocity for this path is 20 knots from left to right, and the flight path is in the same direction.

The technique from [41] can compute updraft and downdraft profiles over 2-D regions while accounting for wind from various directions by utilizing multiple flight paths over a 2-D grid. For the considered Warner Springs area, a grid of 23 km $\times$ 61 km was divided into 14 paths. After resolving wind speed and direction into its components, the updraft or downdraft along each of the 14 paths was computed and thereafter interpolated within the grid to find the updraft profile over the entire region. Therefore, given any land terrain, wind speed and wind direction, estimates of the values of updraft and downdraft over the given region can be provided. Binning and counting the various updraft and downdraft values in the four cardinal directions yielded updraft and downdraft transition probabilities for the associated region.

Sample results for a 20 knot wind from two headings are presented in Fig. 1(b) and (c). The results are intuitive, since the windward side of every hill experiences lift and the leeward side experiences sink in proportion to the wind speed. The positions of updrafts and downdrafts are reversed with wind direction. The magnitudes of updrafts and downdrafts are proportional to the wind speed per the expressions in [41].

This approach has the following advantages: an updraft and downdraft value can be ascertained for every point on a given landscape; the updraft and downdraft profile is sensitive to local hillside contours; and the updraft and downdraft calculation is very simple and computationally inexpensive. This simple model is also quite accurate, as described next.

B. *Validation of the Updraft and Downdraft Model*

We validated the updraft and downdraft model described in the previous subsection using available flight data [59] because the original technique of [41] has not been formally validated in the published literature, even for single flight paths. The available flight data [59] consists of recorded glider flight path latitude and longitude coordinates, glider altitude at regular time instants along the flight path, and thermal locations and strength as experienced by numerous gliders and glider types in various regions. From this data, 10 different recorded flight paths in the Warner Springs region were chosen and a type of Vertical Speed Indicator (VSI) reading was generated for numerous points along each flight path. Each of these VSI readings incorporate only the combined effect of thermals and updrafts. The recorded thermal data was then used to eliminate thermal effects in these readings, revealing the sole VSI effects of updrafts and downdrafts over each flight path.

Four resultant flight path-specific VSI updraft profiles and model-generated updraft profiles are shown in Fig. 2. The model output (red dashed) is visibly consistent with actual flight data (black solid). The figure's table indicates the model output mean error and its standard deviation for the 10 paths. Using the table, the technique is, overall, 92.4% accurate, since the average mean absolute percentage range error is 7.6%.

C. *Modeling and Validation of Thermals*

By reversing the above updraft/downdraft validation process, it is possible to use the updraft/downdraft model with available flight data to determine the location and intensity of thermals in a given region. For Warner Springs, 20 flight paths (different from the 10 above for updraft/downdraft validation) were considered, and a type of VSI reading as previously described was obtained from the data. The updrafts over this region were then computed using the model, and the difference between the computed updrafts and VSI readings from the flight data provided an estimate of thermal intensities and locations.

Fig. 2 shows the resultant thermal intensity profile for the Warner Springs region. The 20 flight paths were taken at approximately the same time of day, between 1800–1930 Coordinated Universal Time (1100–1230 local time), and it was assumed that the prevailing wind direction was constant for these flights and that weather conditions were similar (i.e., suitable for competition glider flight). We note the presence of

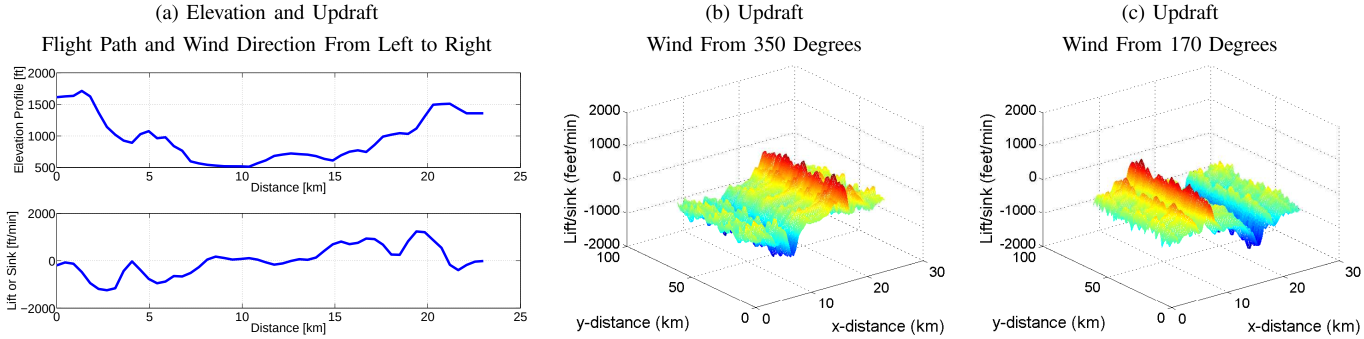

Fig. 1. Profiles over a region near Warner Springs, CA, with a wind speed of 20 knots and wind direction as indicated.

(a) Data-Based Non-Thermal VSI profiles and Modeled Updraft Profiles

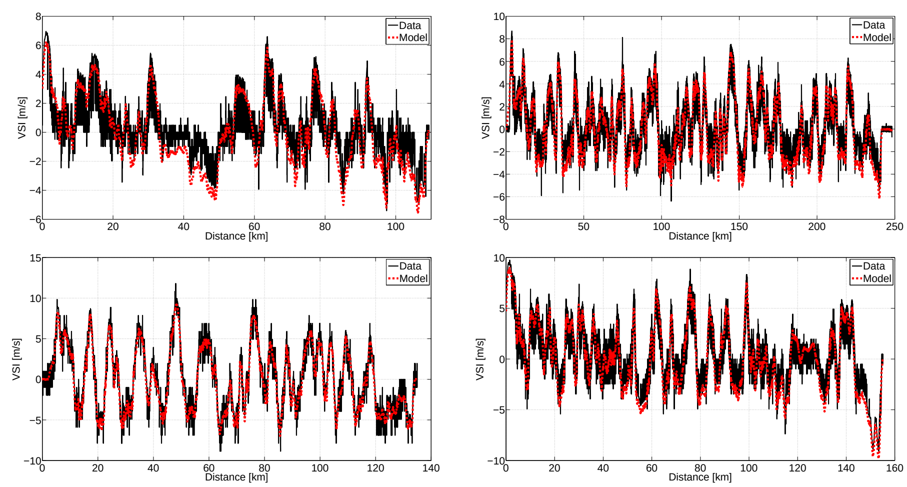

(b) Validation Results for 10 Flight Paths

Path	Mean Percentage Abs. Range Error	Std. Dev.
1	10.29	6.32
2	10.53	6.27
3	8.01	5.19
4	6.68	4.86
5	4.48	4.20
6	4.69	2.77
7	6.40	4.06
8	12.3	7.03
9	5.61	3.67
10	6.95	5.05

(c) Thermal Intensity Profile over the Warner Springs Region

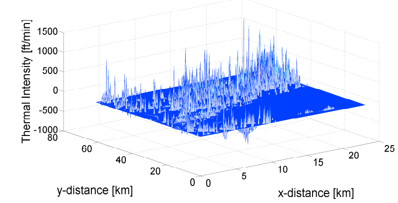

Fig. 2. Visual depiction of data-based non-thermal VSI profiles (black solid lines) and of modeled updraft profiles (red dashed lines), tabulated updraft validation results from 10 flight paths, and thermal intensity profile over the Warner Springs region.

"negative thermals," or locations of strong sink in the intensity profile. This is not unexpected, since in practice, the flight environment that is located near a thermal could experience sink due to air flow rushing to the base of the nearby thermal to replace rising air, or alternately, due to the cooling and descent of rising air after moving away from a heating source. It was observed that when the grid size was reduced to 100 m, the generated thermal profile was discrete enough that a thermal encountered during different flight paths at the same time of day on different days at roughly the same location had the same intensity. This is an intuitively expected outcome.

As before, binning and counting lift and sink values in the four cardinal directions yielded thermal transition probabilities for a region. These were calculated for Warner Springs.

Validation of the above thermal-determination approach was completed using a software package called SeeYou, which is used by soaring pilots to view their flights. This software provides the location and intensity of a thermal from flight data. The thermal profile generated for the Warner Springs region with our approach was compared with locations of thermals obtained from SeeYou. The software locations of thermals were found to match identically with locations in the model-generated profile. Moreover, the aforementioned computing of essentially the same thermal (encountered over multiple flight paths) as the same intensity for approximately the same conditions served as another validation mechanism.

Our above approach to computing thermal transition probabilities is contingent upon data availability for a specific region. However, it is possible to leverage our approach to produce a thermal transition probability model that depends only upon terrain condition, independent of terrain location. The thermal locations and intensities obtained with our approach or directly from SeeYou can be correlated to land cover classifications over the data's terrain condition (e.g., tree canopy, paved asphalt, dry land, a crop field), which will give a probabilistic estimate of thermal intensities over any land cover classification. Including these estimates in an extended model will eliminate a need for readily-available, region-specific flight data. Such an extended model is left to future work.

V. Application of Control Strategies

A. Glider Description and Model Parameters

We employ a glider model with parameters for Schweizer Aircraft's SGS 2-33A (Fig. 3), a common flight trainer in

North America that was manufactured during the late 1960s and throughout the 1970s [60]. In addition to the popularity of this two-seat glider, our choice is based on the relative inferiority of the glider's soaring performance (or glide ratio) when compared to modern sailplanes and competition gliders; hence, the optimal stochastic flight management decisions made with this trainer glider and the above theory may facilitate achieving comparable cross-country distances to those attained by skilled pilots in high-performance gliders with artificial aid. A third reason for this choice of aircraft is author familiarity with flight instruction on this glider, which enables a comparison of control strategy results with real-life practice.

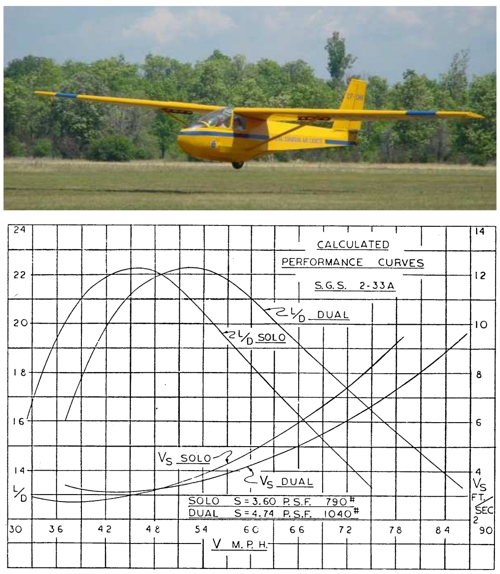

Fig. 3. The Schweizer SGS 2-33A trainer glider in the colors of the Air Cadet League of Canada, and the glider's polar curve [55].

The glider polar curve, depicted in Fig. 3 [55], determines the sink rate, $v_z = f(u_1)$, in still (no-wind) air versus the ground speed. In this plot, V_S designates the sink rate, v_z, and the airspeed V is taken to be u_1 (i.e., we assume a no-wind condition). The diagram also illustrates a best glide ratio (i.e., a best lift/drag ratio or L/D ratio) of 22.25:1, corresponding to 22.25 m of forward travel for every 1 m of descent, at V speeds of approximately 20.1 m/s or 45 mph solo and 23.2 m/s or 52 mph dual. A minimum sink V_S speed of 0.792 m/s or 2.6 ft/s is achieved at a V speed of 17.0 m/s or 38 mph solo, and V_S = 0.945 m/s or 3.1 ft/s at V = 18.8 m/s or 42 mph dual. Glider data is available in [54]. The polar curve for solo flight in the glider can be approximated by a quadratic function of the longitudinal speed,

$$v_z[\text{ft/s}] = 8.2582 - 0.2870u_1[\text{mph}] + 0.0038u_1[\text{mph}]^2. \quad (20)$$

For the range maximization task, the minimum and maximum altitudes for the glider are selected as h_{min} = 1,000 ft and h_{max} = 5,000 ft. For the value iterations, the altitude grid spacing is 100 ft, t_{min} = 0 s, t_{max} = 3,600 s, and time grid spacing is 30 s. The updraft grid is selected from -200 ft/min to 500 ft/min with a spacing of 100 ft/min. The thermal grid goes from -200 ft/min to 700 ft/min with a spacing of 100 ft/min. All values of updrafts and thermals are rounded to a grid point. Both these grids are treated independent of each other. The rounded transition probabilities for the thermals P_t and updrafts P_u are given in (21) and (22), respectively. Due to rounding, the rows may not sum to 1 exactly.

$$P_t = \begin{bmatrix} 0.21 & 0.12 & 0.66 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.03 & 0.29 & 0.66 & 0.02 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.98 & 0.01 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.01 & 0.67 & 0.26 & 0.03 & 0.01 & 0.01 & 0 & 0 & 0 \\ 0 & 0.01 & 0.74 & 0.1 & 0.08 & 0.03 & 0.01 & 0.01 & 0 & 0.01 \\ 0 & 0.01 & 0.78 & 0.05 & 0.05 & 0.06 & 0.02 & 0.01 & 0 & 0.01 \\ 0 & 0.01 & 0.8 & 0.06 & 0.03 & 0.04 & 0.03 & 0.01 & 0.01 & 0.01 \\ 0 & 0.01 & 0.81 & 0.04 & 0.05 & 0.02 & 0.02 & 0.02 & 0.01 & 0.02 \\ 0 & 0.01 & 0.79 & 0.06 & 0.03 & 0.02 & 0.02 & 0.01 & 0.02 & 0.03 \\ 0 & 0.01 & 0.82 & 0.04 & 0.02 & 0.02 & 0.01 & 0.02 & 0.02 & 0.04 \end{bmatrix}, \quad (21)$$

$$P_u = \begin{bmatrix} 0.95 & 0.05 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.07 & 0.82 & 0.1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.07 & 0.85 & 0.08 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.12 & 0.8 & 0.08 & 0 & 0 & 0 \\ 0 & 0 & 0.01 & 0.14 & 0.75 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & 0.01 & 0.16 & 0.72 & 0.11 & 0 \\ 0 & 0 & 0 & 0 & 0.02 & 0.17 & 0.65 & 0.16 \\ 0 & 0 & 0 & 0 & 0 & 0.01 & 0.06 & 0.93 \end{bmatrix}. \quad (22)$$

The thermal and updraft grid selections were based on author flying experience and flight data VSI readings. This grid selection also allowed comparisons between thermal and updraft profiles over different regions. The generated transition probabilities were independent of flight direction because the probability of transitioning from one lift value to another was calculated based on updrafts and thermals encountered in the four principal directions at each grid point. The grid for u_1 was selected as 36, 39, 42, 45, 48, 51, 54, 60, 66, 72, 78, 84, 90, 96 mph and the grid for u_2 was selected as 0, 15, 30, 45, 60, 120, 180, 240, 300, 420 s.

In the ground vehicle surveillance task, the ground vehicle velocities are 39, 54 and 66 mph, a grid size that reduces the curse of dimensionality. The probability transition matrix is:

$$P_3 = \begin{bmatrix} 0.5 & 0.25 & 0.25 \\ 0.25 & 0.5 & 0.25 \\ 0.25 & 0.25 & 0.5 \end{bmatrix}.$$

The vehicle's relative distance to the glider is limited to 20 km. The thermal strength model is the same as for range maximization, with updrafts considered to be zero for simplicity.

B. SDCOC Range Maximization Results

The solid blue line in Fig. 4 depicts the convergence of value iterations in the range maximization task after 300 iterations. Fig. 5 presents various cross-sections of the control policy.

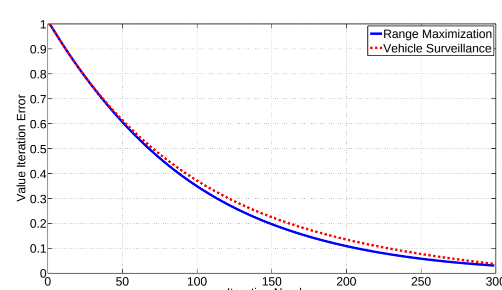

Fig. 4. Convergence of value iterations in the range maximization (blue solid) and ground vehicle surveillance (red dashed) tasks: $||V_n - V_{n-1}||$ versus n.

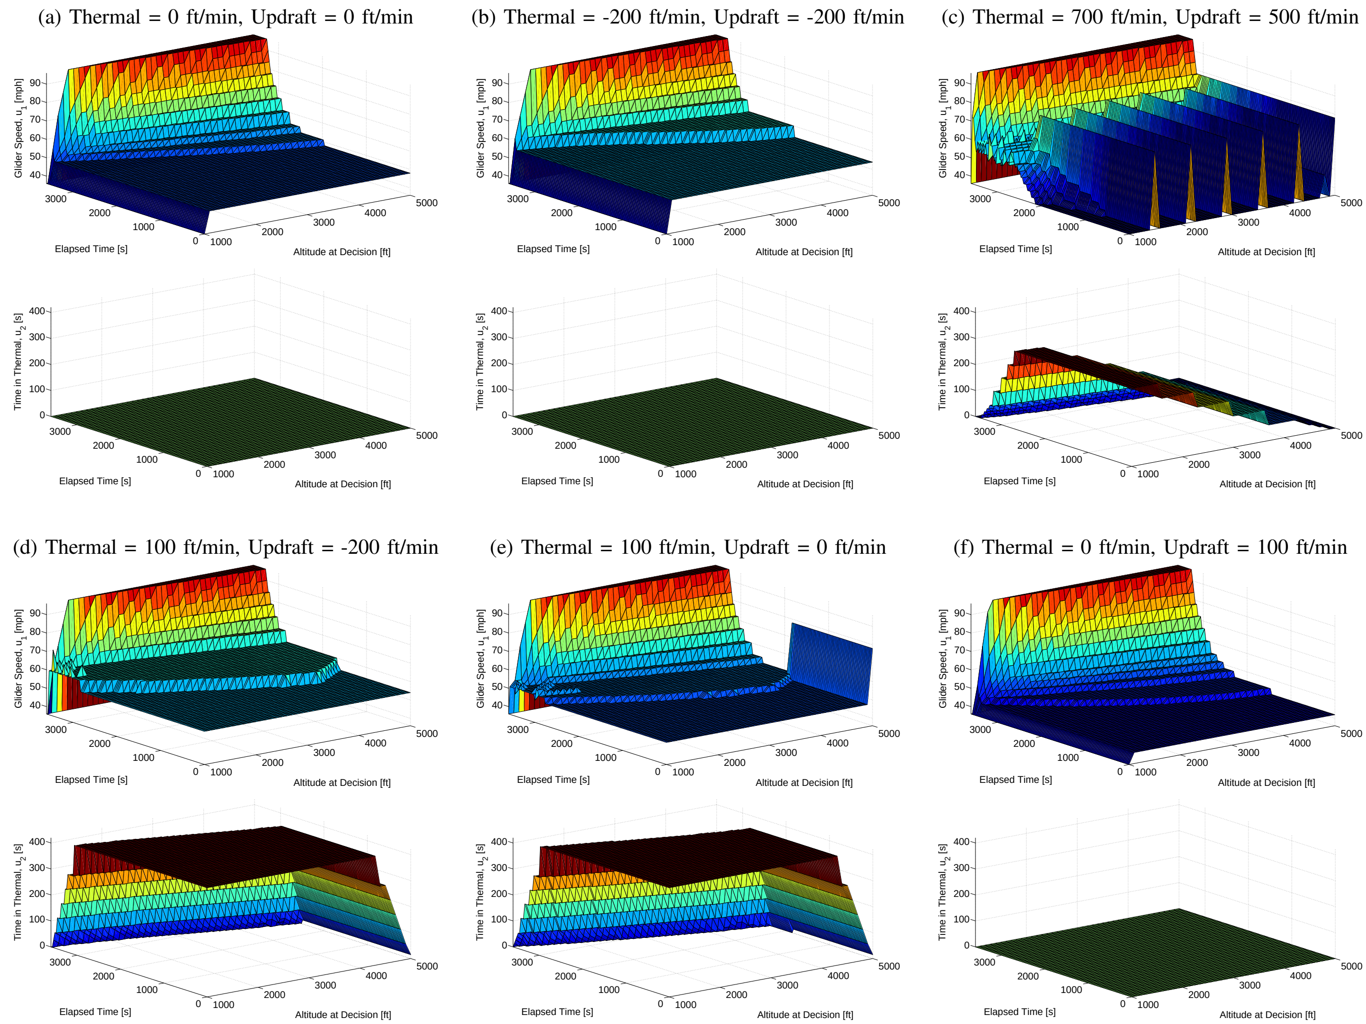

Fig. 5. Cross-sections of the SDCOC policy in the range maximization task.

Fig. 5 (a) considers gliding in still air, with neither thermals nor updrafts (or downdrafts). The computed control policy stipulates that flight takes place at the best L/D speed, the known range-maximization speed, except when the elapsed time is "close" to the end-time constraint. As t_{max} approaches, the computed control policy states that the flown glider airspeed must increase to the maximum, which effectively trades off altitude to augment traveled distance. When and how quickly glider airspeed is increased depends on altitude. Because encountered thermals do not boost altitude, the control policy does not require time to be spent in these thermals.

Fig. 5 (b) looks at gliding in adverse conditions, when encountered thermals are sink and flight occurs in downdraft. No time is to be spent thermalling in sinking air. The computed control policy desires glider airspeed to be slightly higher than the best L/D speed to quickly traverse through the downdraft, which is consistent with a rule-of-thumb that is taught to glider student pilots to maximize range. As in the still air condition, an impending t_{max} allows a trade off of altitude for range.

Fig. 5 (c) examines gliding in favorable conditions, when encountered thermals are strong and flight is within a very strong updraft. The optimal policy requires that thermals be used at low glider altitudes, with the thermalling duration decreasing with increasing glider altitude. This policy is mostly independent of elapsed flight time, unless very close to t_{max}. Interestingly, glider airspeed should be minimum sink with a substantial periodic increase to convert altitude into range before decreasing back. Again, as in the still air condition, a closeness to t_{max} affords a trade off of altitude for range.

Fig. 5 (d) pertains to gliding in downdrafts while encountering thermals of minimal strength. As before, an airspeed higher than the best L/D speed is suggested to quickly traverse the downdraft. Depending on how close t_{max} is to elapsed time, an airspeed increase can transfer altitude into distance. When thermals are encountered, particularly at low glider altitudes, the control policy calls for thermalling to increase altitude. If thermals are stronger (figure not shown), thermalling-initiation altitudes and time-to-spend cut-offs are lower.

Figs. 5 (e) and (f) investigate the difference in flight control policy when encountered thermals are of minimally positive strength with updrafts of neutral strength, and when encountered thermals are of neutral strength with updrafts of minimally positive strength. For the former scenario, gliding at the best L/D airspeed is optimal along with maximal thermalling at nearly all elapsed times and altitudes. For the latter scenario, gliding at closer to minimum sink airspeed is optimal. In both scenarios, the end time constraint allows for a conversion of excess altitude into distance traveled.

Thus, the following control policy trends are in Fig. 5. These intuitive trends are consistent with the rules-of-thumb that are taught to glider student pilots for range maximization.

1) Thermals should be taken advantage of, even if located within a broader region of downdraft or sink (assuming sufficient pilot skill or automation and sufficient thermal diameter), as long as the glider is not close to violating the maximum height constraint and there is enough feasible flying time available.
2) Since greater sink rates can be accommodated at higher altitudes, speeds that are slightly faster than best L/D are permissible at these altitudes to increase the gliding range. If the flight time is close to the maximum flight time constraint, then an increased speed is desirable to gain more gliding range. The time at which to begin this speed increase depends on the available altitude.
3) A greater airspeed is desirable in regions of sink to quickly traverse these areas. A slower airspeed (i.e., the minimum sink speed) is desirable in regions of updraft to gain altitude.

A simulation illustrates the above trends. Fig. 6 presents the flight environment. Fig. 7 (a) illustrates the time histories of the associated glider variables. The red dashed lines indicate constraints, and the blue dots indicate snapshots of the flight conditions during a flight segment. The glider utilizes the first thermals that it encounters.

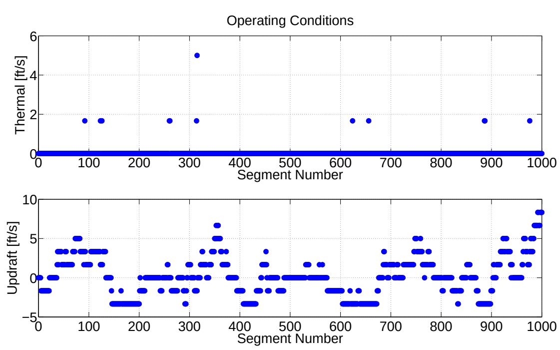

Fig. 6. Thermal strength and updraft versus flight segment number in the range maximization task.

For the same operating conditions and the polar curve:

$$\begin{aligned} v_z[\text{ft/s}] =& 7.9908241338 - 0.28352820676 u_1[\text{mph}] \\ &+ 0.0038324896852 u_1[\text{mph}]^2, \end{aligned} \quad (23)$$

a more accurate curve that, when used with the SDCOC law, represents the scenario of a perturbed and mismatched exploited model, gliding range is unaffected (Fig. 7 (b)). As a baseline, the glider covers 10.1 km in the approximately 501 seconds that it takes to descend from 2,500 ft to 1,000 ft while at the best L/D airspeed of 45 mph when thermals or updrafts are absent, a range value that was computed with the more accurate polar curve. Hence, SDCOC has more than tripled gliding range over baseline in the realized flight environment.

C. SDCOC Ground Vehicle Surveillance Results

The red dashed line in Fig. 4 depicts the convergence of value iterations in the vehicle surveillance task after 300 iterations. Fig. 8 presents cross-sections of the control policy.

Fig. 8 (a) considers gliding without thermals and a fast-moving ground vehicle. The intuitive computed result is to fly at the best L/D airspeed, with slight increases to compensate for greater vehicle distance, as altitude permits. No time needs to be spent in thermals that do not boost altitude.

On the other hand, if encountered thermals are the strongest possible, then these should be leveraged when the glider is at low altitudes and close to the vehicle, as expected. Depending on whether the ground vehicle is moving slow (Fig. 8 (b)) or fast (Fig. 8 (c)) and how far away it is, either a minimum sink airspeed or slightly higher is optimal. Similar small variations are apparent in the desirable thermalling time.

Fig. 9 illustrates the above control policy effects. The thermals are the same as in Fig. 6, but no updraft or downdraft is present. As before, the red dashed lines indicate constraints, and the blue dots indicate snapshots of the flight conditions during a flight segment. The glider tracks a ground vehicle for more than 15 km after its first identification at a relative distance of 10 km when at an altitude of 2,500 ft.

D. SEGS Range Maximization Results

We posit an exponential fitness function for efficient searching by a selective evolutionary generation system, a function that synthesizes available sensor feedback into a performance metric. One such function that uses sensor readings of glider altitude and already-traveled distance is a weighted sum of the fitness from the h^+ altimeter reading, F_A (i.e., future range fitness), and the fitness from the range reading, F_R (i.e., past range fitness). Here, F_A achieves maximum at the midpoint of permissible altitudes, 3,000 ft:

$$F_A = \exp\left(-\left(K_1\left(h^+ - 3000\right)\right)^2\right), \quad (24)$$

where K_1 is calculated so that altitude fitness at the low and high altitude extremes of 1,000 ft and 5,000 ft is 5%, i.e., $K_1 = \sqrt{\frac{\ln 0.05}{-(2000)^2}}$. Next, F_R is chosen to achieve maximum range fitness when the distance traveled by the glider $DistTraveled^+$ (equal to Δs multiplied by the number of flight segments traversed) approaches infinity:

$$F_R = \exp\left(-\left(K_2\left(\frac{1}{DistTraveled^+} - 0\right)\right)^2\right). \quad (25)$$

Here, K_2 is computed so that fitness past some arbitrary distance value exceeds 0.5. This arbitrary distance value represents a minimum desirable range, for example, 10 km, 15 km, etc. Solely for SDCOC comparative purposes in what follows, we use the glider polar curve to give some insight into what arbitrary distance value to pick, and we choose 13.44 km, which corresponds to the distance traveled by the glider at 45 mph when losing 2,000 ft of altitude in the absence of thermals or updrafts. That is, $K_2 = \sqrt{\frac{\ln 0.5}{-\left(\frac{1}{13.44}\right)^2}}$. Weighting F_A and F_R equally at 0.5 produces the fitness function that is plotted in Fig. 10 (a).

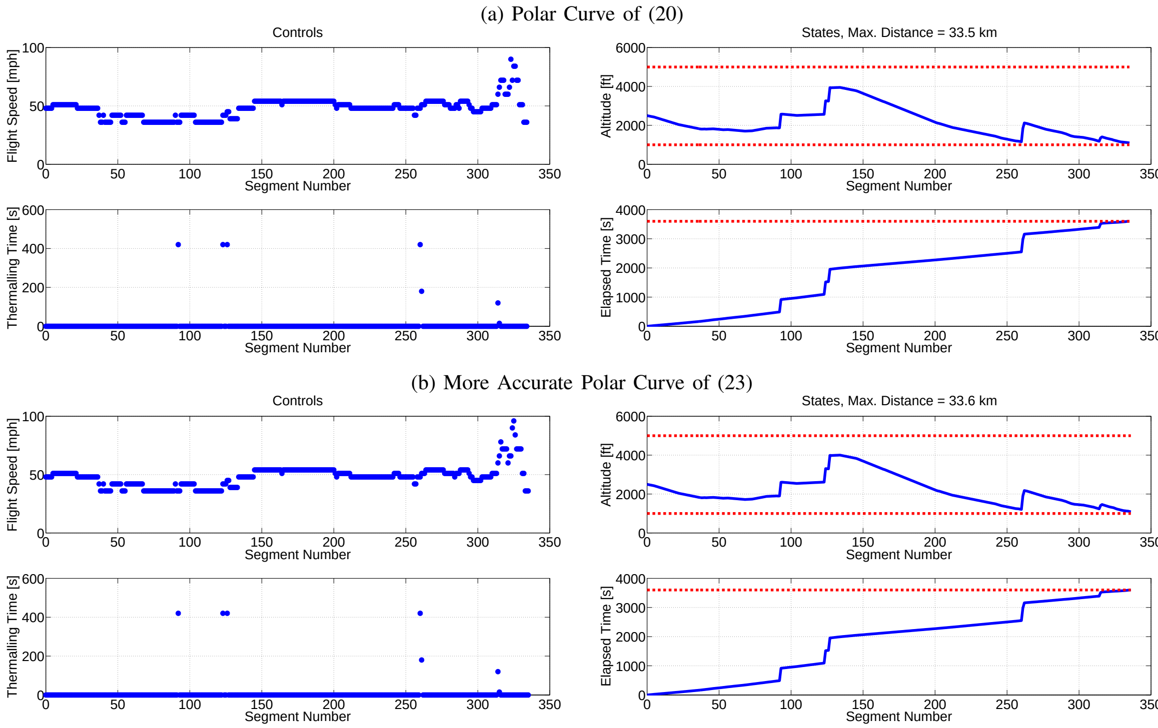

Fig. 7. Range maximization task with the SDCOC law.

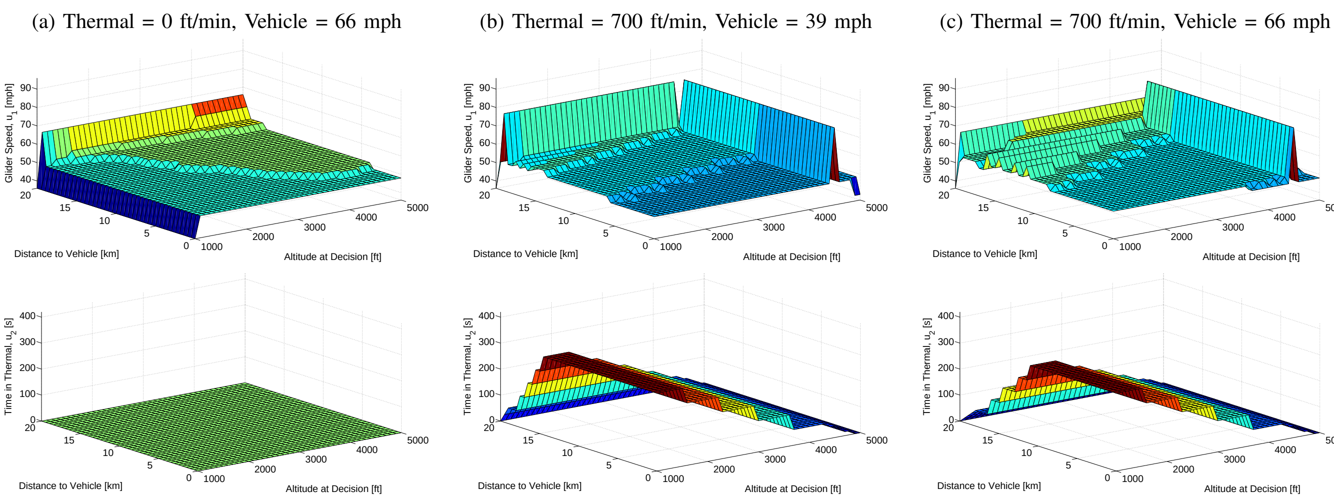

Fig. 8. Cross-sections of the SDCOC policy in the ground vehicle surveillance task.

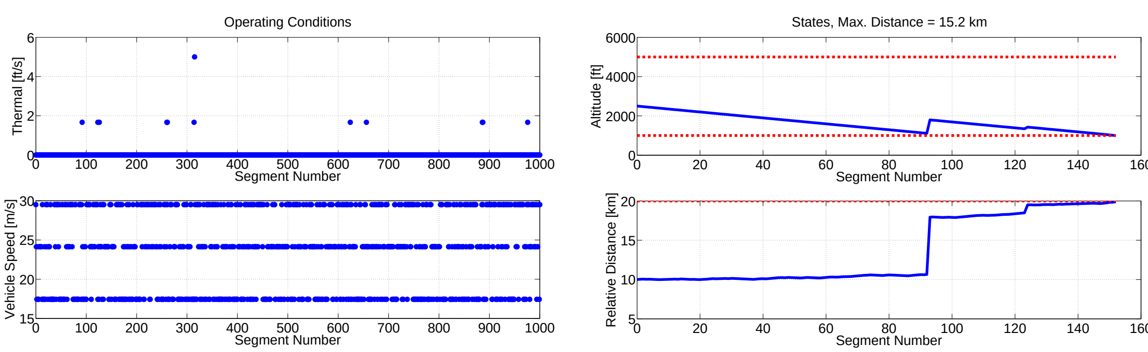

Fig. 9. Ground vehicle surveillance task with the SDCOC law.

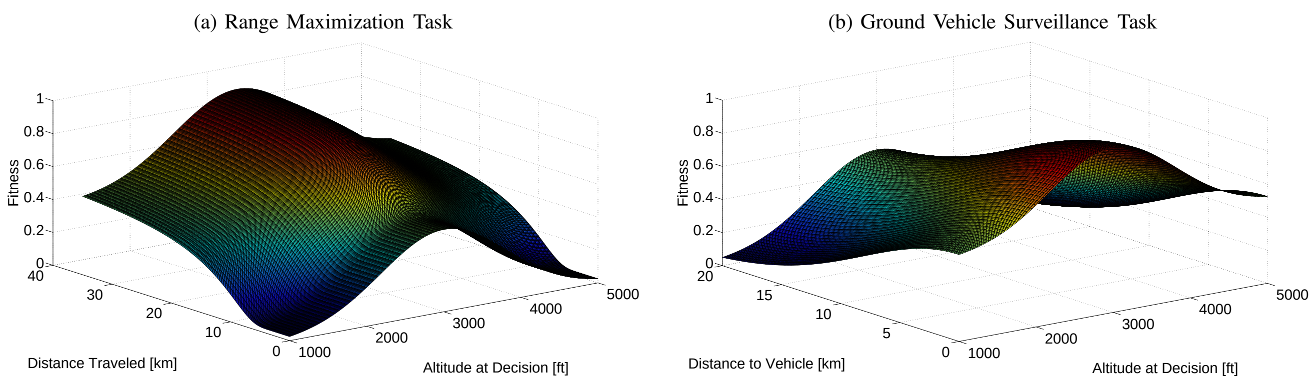

Fig. 10. Chosen SEGS fitness function.

We craft the following selective evolutionary generation system to maximize glider range online. The SEGS controller determines a control setting for airspeed and thermalling time prior to flight segment traversal by the glider, and the controller obtains the fitness of its setting based on sensor readings (altimeter and traveled distance) at the conclusion of that flight segment via the above fitness function. Thus, two consecutive flight segments have to be traversed to compare two control settings, as is done locally by SEGS. In our controller implementation, we also utilize another sensor reading to determine if the SEGS setting of thermalling time is indeed flown by the glider or if it is ignored; that is, if a thermal (positive or negative) is indeed present or not. This determination can be realized in practice very simply, through a sudden large positive or negative change in vertical speed that indicates the presence of a thermal (i.e., if the derivative of the VSI reading exceeds some threshold). We assume perfect thermal determination and valid application of the SEGS setting here.

The set in which to search, U, is the set of ordered pairs $(u_1(t), u_2(t))$, where

$$u_1(t) \in \{36, 39, 42, \ldots, 93, 96\}, \tag{26}$$

$$u_2(t) \in \{0, 15, 30, \ldots, 405, 420\}. \tag{27}$$

The set of resources, R, is the set $\{r_1, r_2, r_3, r_4\}$, with $r_i = \mathbf{e}_i,\ 1 \le i \le 4$ (where $\mathbf{e}_i$ are the standard basis vectors for $\mathbb{R}^4$). This set facilitates the perturbation of one of the two elements of U in either a positive or negative direction after a mapping by G. The probability mass function on R, P, is the discrete uniform distribution. This choice of probability mass function is a special case of SEGS theory validity [50]. The generation function, G, when applied to $u \in U$ using resource $r \in R$, yields $G((u_1(t), u_2(t)), r_i)$, $1 \le i \le 4$, the pair given by

$$\begin{cases} \begin{bmatrix} 3 & -3 & 0 & 0 \\ 0 & 0 & 15 & -15 \end{bmatrix} r_i + \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix}, & \text{if } 36 < u_1(t) < 96, \\ & 0 < u_2(t) < 420, \\ (u_1(t), u_2(t)), & \text{otherwise.} \end{cases} \tag{28}$$

The above is a modified version of a random walk over a discretized search space, where the modification involves selection dynamics described by the $Select$ function to produce a selective evolutionary generation system. By using an exponential fitness function, the dynamics are search-efficient.

A sample run of the evolution scheme when $N = 5$ is depicted in Fig. 11 (a). Thermal strength and updraft versus flight segment number are as per Fig. 6. SEGS control is computationally feasible for online implementation because it takes 4.97 s to traverse a 100 m flight segment at the best L/D airspeed of 45 mph, while benchmark SEGS computation time [49] averages 0.13 ms per iteration, which is four orders of magnitude less. The SEGS controller achieves satisfactory range maximization since it handily beats the 10.1 km range that is achievable at the best L/D airspeed of 45 mph when descending from 2,500 ft to 1,000 ft without thermals or updrafts. In fact, the SEGS control in Fig. 11 (a) realizes the optimal policy of a slower airspeed in an updraft and a faster airspeed in a downdraft. However, an insufficient number of thermals prevents the controller from exploring the effects of various thermalling time values; this is evident by changing initial $u_2(0)$ from 0 s to 420 s, which yields sample Fig. 11 (b). Such gliding range enhancement corroborates [49], where initial condition dependence is noted when few fit values exist in a search space. The augmented range is closer to the optimal SDCOC result (25.4 km vs. 33.5 km).

Deploying the more accurate polar curve to simulate sensor readings after each flight segment should not negatively impact the above results because of SEGS model-independence. As Fig. 11 (c) illustrates, a SEGS controller can improve performance with an improved state mapping.

E. SEGS Ground Vehicle Surveillance Results

A selective evolutionary generation system that is similar to the one described for range maximization can be utilized for ground vehicle surveillance. Let the total fitness function be the equally-weighted sum of F_A and F_D instead of the equally-weighted sum of F_A and F_R, where F_D is the fitness due to d^+, the relative distance between glider and ground vehicle after a flight segment, with greatest fitness at 0 km:

$$F_D = \exp\left(-\left(K_3\left(d^+ - 0\right)\right)^2\right), \tag{29}$$

and only 5% fitness at 20 km, setting $K_3 = \sqrt{\frac{\ln 0.05}{-(20)^2}}$. This fitness function is plotted in Fig. 10 (b).

A sample run of the evolution scheme when $N = 5$ and $u_2(0) = 420$ s is presented in Fig. 12. Again, the outcome is close to the optimal SDCOC result (approximately 12 km

(a) Polar Curve of (20), $u_2(0)$ = 0 s

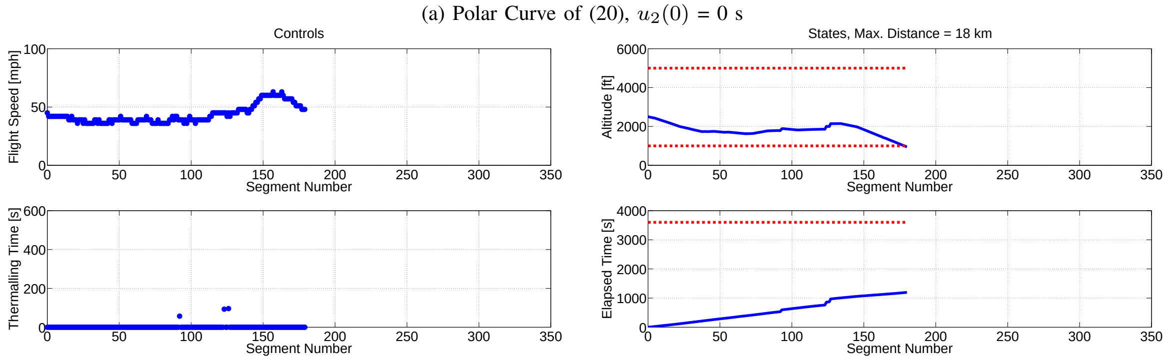

(b) Polar Curve of (20), $u_2(0)$ = 420 s

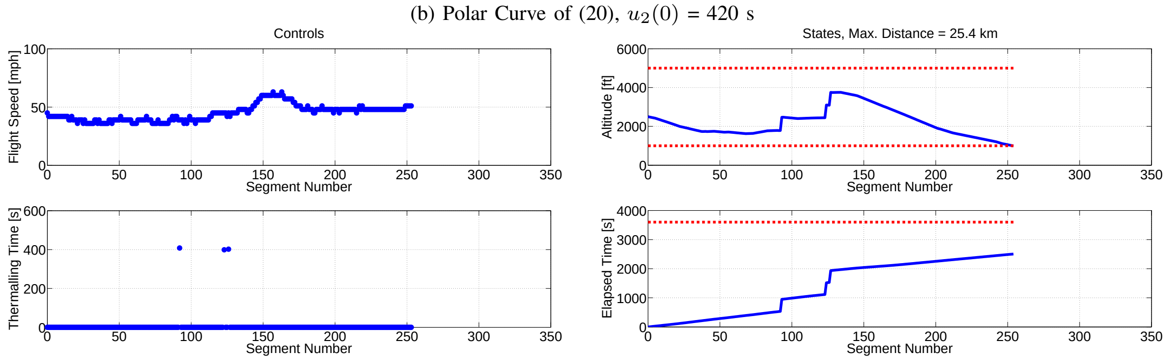

(c) More Accurate Polar Curve of (23), $u_2(0)$ = 420 s

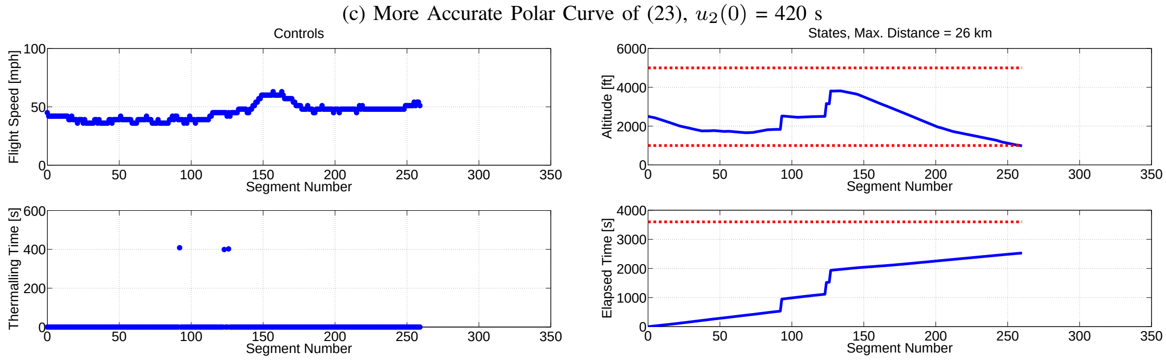

Fig. 11. Sample range maximization task with SEGS-based control evolution.

of vehicle tracking against about 15 km), and glider flight is mostly maintained at the best L/D airspeed to ensure continued tracking. As with the range maximization task, the SEGS controller does not experience enough thermals in the flight environment of Fig. 9 to learn the effects of various thermalling time values; it is while thermalling in the second encountered thermal that the system exits specified constraints.

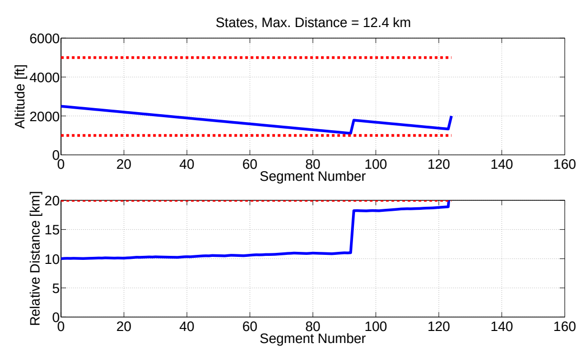

Fig. 12. Sample ground vehicle surveillance task with SEGS-based control evolution and $u_2(0)$ = 420 s.

F. SDCOC and SEGS Computation Results

SDCOC policy computation, which was accomplished in C on a personal Macintosh laptop, took weeks to complete for the range maximization task and several hours to complete for the ground vehicle surveillance task. The ensuing Matlab flight simulations that looked-up these precomputed SDCOC policies required several minutes to complete. In contrast, the combination of online SEGS gliding range or ground vehicle surveillance policy calculation with flight simulation in Matlab required only a few seconds to complete on the same laptop.

VI. SDCOC and SEGS Concluding Remarks

This paper has applied and contrasted two novel stochastic strategies to manage glider flight in a dynamic environment. Stochastic Drift Counteraction Optimal Control (SDCOC), a dynamic programming method, can suitably exploit a validated atmospheric environment model to produce intuitive control policies that effectively maximize gliding range and perform ground vehicle surveillance. Selective Evolutionary Generation Systems (SEGS), a model-independent Markov chain Monte Carlo method, can also produce satisfactory gliding range and ground vehicle surveillance policies, where satisfactory is defined as an increase in gliding and surveillance range over standardized actuator settings (i.e., no active control). However, SEGS policies yield results that fall short of results from SDCOC policies, because online learning and decision-space exploration is needed.

The simple and accurate flight environment model in this paper requires data of sufficient fidelity prior to exploitation by SDCOC. Such data may not be readily available, and the resultant region-specific environment model (based on region-specific data) may not be suitable for flight in another geographic region. SEGS needs no such model, and produces control actions much faster than SDCOC, even when the online version of the latter is policy lookup. But despite its computational expense, SDCOC's superior results may justify its use over SEGS.

This paper has provided quantitative insight into four glider flight management considerations: environment model use necessity, applicability, complexity, and validity. First, the presented SEGS performance confirms that an environment model is not necessary for satisfactory control. Second, by construction, SDCOC does not require an environment model to be tailored to its implementation, nor does SDCOC affect how an environment model generates disturbance transition probabilities. Likewise, an SDCOC implementation does not need to be tailored to an environment model since the optimality of its policies only depends on the values of the disturbance transition probabilities, not on how these probabilities were generated. Third, we show that a good environment model need not be complex, and that the computational burden of such a model can be separated from the computational burden of the control algorithm that relies on the model. And fourth, the optimal policies produced by SDCOC using a single environment model (that captures numerous flight environment features) are valid because they match the intuitive policies that are currently practiced by glider pilots.

The two stochastic strategies in this work help document the effects of a control-based, environment and decision-space exploitation-exploration trade-off for gliding range maximization and ground vehicle surveillance. It is conceivable that these two strategies can be combined to leverage their individual merits. Future work will include a comprehensive study of the sensitivity of both strategies to perturbation size.

VII. Acknowledgments

We thank Ian Forster-Lewis for flight environment modeling suggestions.

References

[1] P. F. A. Di Donato and E. M. Atkins, "An off-runway emergency landing aid for a small aircraft experiencing loss of thrust," in *Proceedings of the 53rd AIAA Aerospace Sciences Meeting and Exhibit*, no. AIAA 2015-1798, 5–9 January 2015, pp. 1–13.

[2] T. Coffey and J. A. Montgomery, "The emergence of mini UAVs for military applications," *Defense Horizons*, no. 22, p. 1, 2002.

[3] M. J. Allen, "Guidance and control of an autonomous soaring UAV," National Aeronautics and Space Administration Dryden Flight Research Center, Tech. Rep. NASA/TM-2007-214611, February 2007.

[4] J. H. Cochrane, "MacCready theory with uncertain lift and limited altitude," *Technical Soaring*, vol. 23, pp. 88–96, July 1999.

[5] F. Hendricks, "Dynamic soaring," Ph.D. dissertation, University of California, Los Angeles, 1972.

[6] G. Sachs and M. Mayrhofer, "Shear wind strength required for dynamic soaring at ridges," *Technical Soaring*, vol. 25, pp. 209–215, October 2001.

[7] T. Kiceniuk, "Dynamic soaring and sailplane energetics," *Technical Soaring*, vol. 25, pp. 221–227, October 2001.

[8] Y. J. Zhao, "Optimal patterns of glider dynamic soaring," *Optimal Control Applications and Methods*, vol. 25, pp. 67–89, March-April 2004.

[9] J. M. Wharington, "Heuristic control of dynamic soaring," in *Proceedings of the 5th Asian Control Conference*, vol. 2, 20–23 July 2004, pp. 714–722.

[10] C. K. Patel and I. Kroo, "Control law design for improving UAV performance using wind turbulence," in *Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit*, 9–12 January 2006, pp. 1–10.

[11] N. R. J. Lawrance and S. Sukkarieh, "A guidance and control strategy for dynamic soaring with a gliding UAV," in *Proceedings of the 2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 3632–3637.

[12] W. E. Shaw-Cortez and E. Frew, "Efficient trajectory development for small unmanned aircraft dynamic soaring applications," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 3, pp. 519–523, 2015.

[13] X.-Z. Gao, Z.-X. Hou, Z. Guo, and X.-Q. Chen, "Energy extraction from wind shear: Reviews of dynamic soaring," *Journal of Aerospace Engineering*, 2015.

[14] M. J. Allen, "Autonomous soaring for improved endurance of a small uninhabited air vehicle," in *Proceedings of the 43rd AIAA Aerospace Sciences Meeting and Exhibit*, January 2005.

[15] M. J. Allen and V. Lin, "Guidance and control of an autonomous soaring vehicle with flight test results," in *Proceedings of the 45th AIAA Aerospace Sciences Meeting and Exhibit*, January 2007.

[16] N. E. Kahveci, P. A. Ioannou, and M. D. Mirmirani, "Adaptive LQ control with anti-windup augmentation to optimize UAV performance in autonomous soaring applications," *IEEE Transactions on Control Systems Technology*, vol. 16, no. 4, pp. 691–707, 2008.

[17] R. Arho, "Optimal dolphin soaring as a variational problem," *OSTIV Publication XIII*, 1974.

[18] D. E. Metzger and J. K. Hedrick, "Optimal flight paths for soaring flight," *Journal of Aircraft*, vol. 12, pp. 867–871, 1975.

[19] H. Reichmann, *Cross-Country Soaring*. Thomson Publications, 1978.

[20] J. Sandauer, "Some problems of the dolphin-mode flight technique," *OSTIV Publication XV*, 1978.

[21] B. L. Pierson and J. L. D. Jong, "Cross-country sailplane flight as a dynamic optimization problem," *International Journal for Numerical Methods in Engineering*, vol. 12, pp. 1743–1759, 1978.

[22] G. Sander and F. X. Litt, "On global optimal sailplane flight strategy," in *Science and Technology of Low Speed and Motorless Flight*, P. W. Hanson, Ed. NASA, March 1979, no. NASA Conference Publication 2085, pp. 355–376.

[23] J. L. de Jong, "The convex combination approach: A geometric approach to the optimization of sailplane trajectories," *OSTIV Publication XVI*, pp. 182–201, 1981.

[24] J. W. Langelaan, "Long distance/duration trajectory optimization for small UAVs," in *Proceedings of the 2007 AIAA Guidance, Navigation and Control Conference and Exhibit*, no. AIAA 2007-6737, 16–19 August 2007.

[25] A. Chakrabarty and J. W. Langelaan, "Energy maps for long-range path planning for small- and micro-UAVs," in *Proceedings of the 2009 AIAA Guidance, Navigation and Control Conference and Exhibit*, vol. 2009, no. AIAA 2009-6113, 10–13 August 2009.

[26] J. L. Grenestedt and J. R. Spletzer, "Towards perpetual flight of a gliding unmanned aerial vehicle in the jet stream," in *Proceedings of the 49th IEEE Conference on Decision and Control*. IEEE, 2010, pp. 6343–6349.

[27] O. K. Ariff and T. H. Go, "Waypoint navigation of small-scale UAV incorporating dynamic soaring," *Journal of Navigation*, vol. 64, no. 01, pp. 29–44, January 2011.

[28] N. R. J. Lawrance and S. Sukkarieh, "Path planning for autonomous soaring flight in dynamic wind fields," in *Proceedings of the 2011 IEEE International Conference on Robotics and Automation*. IEEE, May 2011, pp. 2499–2505.

[29] D. Lee, S. Longo, and E. C. Kerrigan, "Predictive control for soaring of unpowered autonomous UAVs," in *Proceedings of the 4th IFAC Nonlinear Model Predictive Control Conference*, vol. 4, no. 1, 2012, pp. 194–199.

[30] S. M. Ross, R. G. Cobb, and W. P. Baker, "Stochastic real-time optimal control for bearing-only trajectory planning," *International Journal of Micro Air Vehicles*, vol. 6, no. 1, pp. 1–28, 2014.

[31] X.-Z. Gao, Z.-X. Hou, Z. Guo, R.-F. Fan, and X.-Q. Chen, "Analysis and design of guidance-strategy for dynamic soaring with UAVs," *Control Engineering Practice*, vol. 32, pp. 218–226, November 2014.

[32] P. B. MacCready, Jr., "Optimum airspeed selector," *Soaring*, p. 10, January-February 1958.

[33] R. Almgren and A. Tourin, "Optimal soaring via Hamilton-Jacobi-Bellman equations," *Optimal Control Applications and Methods*, 2014.

[34] K. Turkoglu, "Real-time guidance strategies for optimizing aircraft performance in stochastic wind conditions," in *Proceedings of the 2014 American Control Conference*. IEEE, 2014, pp. 2480–2485.

[35] N. R. J. Lawrance and S. Sukkarieh, "Autonomous exploration of a wind field with a gliding aircraft," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 3, pp. 719–733, 2011.

[36] R. Bencatel, J. T. de Sousa, and A. Girard, "Atmospheric flow field models applicable for aircraft endurance extension," *Progress in Aerospace Sciences*, vol. 61, pp. 1–25, August 2013.

[37] D. M. Deaves, "Wind over hills: A numerical approach," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 1, pp. 371–391, 1975.

[38] J. R. Pearse, "Wind flow over conical hills in a simulated atmospheric boundary layer," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 10, pp. 303–313, December 1982.

[39] M. H. Kobayashi, J. C. F. Pereira, and M. B. B. Siqueira, "Numerical study of the turbulent flow over and in a model forest on a 2D hill," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 53, pp. 357–374, 1994.

[40] L. M. S. Paiva, G. C. R. Bodstein, and W. F. Menezes, "Numerical simulation of atmospheric boundary layer flow over isolated and vegetated hills using RAMS," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 97, pp. 439–454, November 2009.

[41] I. Forster-Lewis, "Efficient simulation of topographic lift in Microsoft Flight Simulator (FSX)," December 2007. [Online]. Available: http://carrier.csi.cam.ac.uk/forsterlewis/soaring/sim/fsx/dev/sim_probe/sim_probe_paper.html

[42] F. Bernis, "Migration of falconiformes and *Ciconia* spp. through the straits of Gibraltar, part 2," *Ardeola*, vol. 21, pp. 489–580, 1980.

[43] S. Christensen, O. Lou, M. Muller, and H. Wohlmuth, "The spring migration of raptors in southern Israel and Sinai," *Sandgrouse*, vol. 3, pp. 1–42, 1981.

[44] L. S. Buurma, "Patterns of high bird migration over the North Sea area in October," *Limosa*, vol. 60, pp. 63–74, 1987.

[45] Y. Leshem and Y. YomTov, "The use of thermals by soaring migrants," *Ibis*, vol. 138, pp. 667–674, 1996.

[46] M. J. Allen, "Updraft model for development of autonomous soaring uninhabited air vehicles," in *Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit*, no. AIAA 2006-1510, January 2006.

[47] I. V. Kolmanovsky, L. Lezhnev, and T. L. Maizenberg, "Discrete-time drift counteraction stochastic optimal control: Theory and application-motivated examples," *Automatica*, vol. 44, pp. 177–184, January 2008.

[48] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Athena Scientific, 2007, vol. I and II.

[49] A. A. Menezes and P. T. Kabamba, "Efficient and resilient micro air vehicle flapping wing gait evolution for hover and trajectory control," *Engineering Applications of Artificial Intelligence*, vol. 54, pp. 1–16, September 2016.

[50] ——, "Efficient search and responsiveness trade-offs in a Markov chain model of evolution in dynamic environments," *Mathematical Biosciences*, vol. 276, pp. 44–58, June 2016.

[51] ——, "Optimal search efficiency of Barker's algorithm with an exponential fitness function," *Optimization Letters*, vol. 8, no. 2, pp. 691–703, February 2014.

[52] I. V. Kolmanovsky and A. A. Menezes, "A stochastic drift counteraction optimal control approach to glider flight management," in *Proceedings of the 2011 American Control Conference*, 29 June–11 July 2011, pp. 1009–1014.

[53] D. D. Shah, A. A. Menezes, and I. V. Kolmanovsky, "Glider flight environment modeling for optimal control," in *Proceedings of the 2012 American Control Conference*, June 27–29, 2012, pp. 926–931.

[54] Department of National Defense, Canada, *Air Cadet Gliding Program Manual*, 2010, no. A-CR-CCP-242/PT-005.

[55] *The 2-33 Sailplane Flight - Erection - Maintenance Manual*. Schweizer Aircraft Corporation.

[56] I. Kolmanovsky and T. Maizenberg, "Optimal containment control for a class of stochastic systems perturbed by Poisson and Wiener processes," *IEEE Transactions on Automatic Control*, vol. 47, pp. 2041–2046, 2002.

[57] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-Dynamic Programming*. Athena Scientific, 1996.

[58] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. MIT Press, 2009, pp. 1170.

[59] Online Contest (OLC). [Online]. Available: http://www.onlinecontest.org/olc-2.0/gliding/

[60] P. A. Schweizer, *Wings Like Eagles: The Story of Soaring in the United States*. Smithsonian Institution Press, 1988.

Amor A. Menezes (S'07–M'11–SM'17) received the B.A.Sc. degree in mechanical engineering from the University of Waterloo in 2005, and the M.S.E. and Ph.D. degrees in aerospace engineering from the University of Michigan in 2006 and 2010, respectively.

He is an Associate Project Scientist at the University of California, Berkeley. His research interests are in systems biology and synthetic biology.

Dr. Menezes is the Science Principal Investigator of the five-year, multi-university Center for the Utilization of Biological Engineering in Space (CUBES), a NASA Space Technology Research Institute in biomanufacturing for deep space exploration. He was a 2015 Emerging Leader in Biosecurity and a 2015 fellow of the Synthetic Biology Leadership Excellence Accelerator Program.

Dhaval D. Shah received the B.S. degree in electrical engineering from the University of Mumbai in 2010, and the M.S. degree in aerospace engineering from the University of Michigan in 2012.

He is an On Board Diagnostic Controls Engineer at FCA USA. His responsibilities include controls algorithm development for robust detection of EGR, EVAP, VVL, VVT, PCV, and braking system diagnostics.

Mr. Shah has two United States patents pending.

Ilya V. Kolmanovsky (F'08) received the M.S. and Ph.D. degrees in aerospace engineering, and the M.A. degree in mathematics, from the University of Michigan in 1993, 1995, and 1995, respectively.

He is presently a full professor in the Department of Aerospace Engineering at the University of Michigan, with research interests in control theory for systems with state and control constraints, and in control applications to aerospace and automotive systems.

Dr. Kolmanovsky is a Fellow of the IEEE and a past recipient of the Donald P. Eckman Award of the American Automatic Control Council. He is named as an inventor on 95 United States patents.