Reconciling Simultaneous Evolution of Ground Vehicle Capabilities and Operator Preferences

Christopher Slon[1] Vijitashwa Pandey[1]
David Gorsich[2] Paramsothy Jayakumar[2]
[1]Industrial and Systems Engineering Department, Oakland University Rochester MI 48309
[2]US Army CCDC Ground Vehicles Systems Center, 6501 E. Eleven Mile Rd., Warren MI 48397

ABSTRACT
An objective evaluation of ground vehicle performance is a challenging task. This is further exacerbated by the increasing level of autonomy, dynamically changing the roles and capabilities of these vehicles. In the context of decision making involving these vehicles, as the capabilities of the vehicles improve, there is a concurrent change in the preferences of the decision makers operating the vehicles that must be accounted for. Decision based methods are a natural choice when multiple conflicting attributes are present, however, most of the literature focuses on static preferences. In this paper, we provide a sequential Bayesian framework to accommodate time varying preferences. The utility function is considered a stochastic function with the shape parameters themselves being random variables. In the proposed approach, initially the shape parameters model either uncertain preferences or variation in the preferences because of the presence of multiple decision makers. We consider this utility distribution as the prior and update it to a posterior with feedback that can be acquired from actual system use. The framework improves the utility function and thereby the decisions made for the next generation systems, allowing continuous improvement. We present our approach on a ground vehicle selection problem.

1. Introduction

Formal decision-based approaches offer some of the best tools for developing product design specifications that are optimized to satisfy customer preferences. In engineering design, many times the final attributes of the product and the tradeoff between them is an afterthought. As a result, many products disappoint in practice because they do not take into account the preferences of the end user or customer directly. Additionally, techniques that do integrate end-user preferences into product design, work with static preferences on a static design problem. While offering an advantage over existing *bottom-up* engineering methods, such approaches neither take into account the fact that customer preferences evolve with time, nor that many times the technology itself is not fully developed. This leads to two challenges: products that are optimal today may not be so in the near future even with fixed preferences, and constant technological developments limit the trueness of the assessed preferences. Many times, it is argued, that the preference structure is not even fully formed in the customers' mind. In this paper, we present a Bayesian approach to accommodating this interplay between the evolution of customers' preferences and changes in technological capabilities.

The application area where the aforementioned challenges are expected to manifest themselves is the advent of autonomy in ground vehicles. An objective evaluation of ground vehicle performance is typically a challenging task, and off-road operation further complicates this issue. Different types of approaches have been identified in the literature for modeling vehicle terrain interactions for off road mobility such as empirical, semi-empirical, analytical and finite and discrete element methods, with their own concomitant challenges (Ahlvin and Haley, 1992, Taheri et al., 2015). The empirical models have lower computational cost and tend to be less accurate, while analytical and Finite Element or Discrete Element type methods that utilize the physics of the problem, generally are computationally very expensive. Another aspect of modeling and evaluating mobility is the difficulty in coming up with a metric that takes into account all the performance attributes of the vehicle that captures the operator's preferences. The above challenges are only exacerbated by the increasing level of autonomy, dynamically changing the roles and capabilities of these vehicles. For example, terrains considered untrafficable with traditional vehicles may become trafficable with assistive methods such as traction control and antilock braking systems. These capabilities keep enhancing and appear to be culminating in initially teleoperability and finally semi and full autonomy of these vehicles. The missions these vehicles are going to be expected to take part in, will only get increasingly complex (Department of Defense, 2012). How does one make acquisition and operational decisions regarding these vehicles when there is so much uncertainty in their future capabilities and concomitant expectations from them?

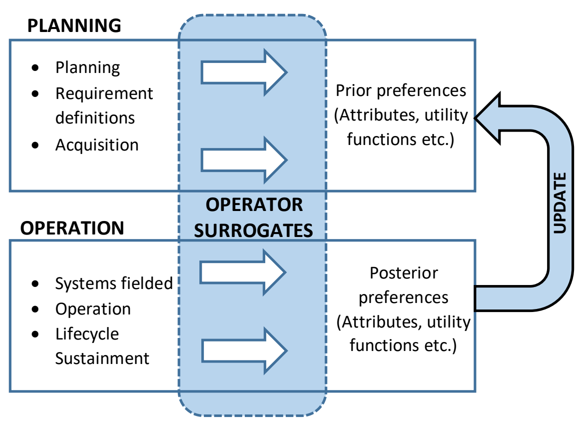

Figure 1. Prior utility function in planning and posterior utility function in operation.

Another challenge with the decision making described above is the presence of multiple decision makers. In cases where a single decision maker is present with static preferences, finding the optimal decision is a relatively straightforward process. But most decisions involve multiple attributes, and involve a wide range of decision makers across an organization (Chen et al, 2012). Additionally, the same decision-maker exhibits variability in his or her preferences from one time to another. This could be due to the change in their frame of mind as described above, variability introduced by the method of preference elicitation, or any number of other uncontrolled or uncontrollable factors (Becker, et al., 1963). Consider the following events in a decision analysis driven planning cycle depicted in Figure 1 (top rectangle). In the initial phases, experts will determine the acceptable levels of the attributes of the vehicles to evaluate the vehicle offerings. These attributes could include vehicle's average speed, the fraction of the terrain that the vehicle can safely negotiate (trafficable percent area, TPA), among others. Through commonly used lottery questions, a multi-attribute utility function that characterizes the preference behavior of the stakeholders with regard to these attributes can be acquired (Pandey, 2013; Barseghyan et al, 2018; Farquhar, 1984). The steps for system realization, whether in-house or at a supplier or contractor will include, definition of system architecture, design of parts with specifications of features, dimensions and tolerances, assembly processes all targeted towards maximizing the expected utility of the product. Common practices of associating attributes with design specifications include the House of Quality (Chowdhury, 2002), Design for Six Sigma (Babu and Asha, 2014), and system mapping (Hoffenson and Söderberg, 2015).

Once acquired and fielded, a system is evaluated by engineering managers and operators. It is clearly a natural progression to want to use this data to update the prior belief of utility based on a design and acquisition perspective, to a posterior belief of utility based on actual system and its performance. This is clearly the case when a system is expected to evolve in capabilities. This leads us to utilizing a Bayesian approach, which can help answer the following questions:

1. As usage data becomes available, how do we evaluate the utility of the fielded vehicles and compare it to the utility of the vehicles when ordered, given that the requirements (preferences) have changed?
2. From a group decision making standpoint, will a vehicle satisfy the group as a whole, or the largest number of end users, or both?
3. Does the data inform decisions about future design updates or adjustments needed?

Our proposed approach can be summarized as follows. By looking at DM preferences and their resulting utility functions governed by random variable shape parameters, we accommodate uncertainty in preferences. We take a Bayesian perspective and use preference data from vehicle operation to update the distribution of the utility function parameters used in the acquisition phase. This *posterior* stochastic utility function becomes the *prior* stochastic utility function for the next generation of the product.

The paper is organized as follows. The next section, discusses the evolution of DBD and identifies the gaps in the research relevant to this work. In section 3, the mathematical model for vehicle performance calculation is presented. Section 4 presents our proposed approach while section 5 presents the results of applying our approach to the vehicle selection problem. Section 6 concludes with a discussion of the results, strengths and weaknesses of the approach, as well as suggestions for future work.

2. Decision making background

The foundation of this work is in decision-based design (DBD), which posits that design essentially a decision-making process (Smith et al, 2015). Formal decision analysis (DA) provides a framework for making good design decisions by maximizing the expected utility (Pandey, 2013). Bentham was one of the first to define utility and even used it as a measure of social welfare (Bentham, 1879). John von Neumann and Morgenstern proposed an axiomatic framework for ordering of alternatives, incorporating decision maker's preferences and attitude toward risk (von Neumann and Morgenstern, 1947). It is agreed that the best choice in a decision problem has the highest (expected) utility (Friedman and Savage, 1948). Most researcher agree and are comfortable with the conditions outlined by von Neumann and Morgenstern to arrive at a cardinal measure of preferences when there is a single decision maker making decisions involving a single attribute. Some dispute and concomitant challenges arise when multiple attributes and multiple decision makers are involved.

There have been many practical approaches to multi-attribute decision problems—for example, Edwards (1977) discusses an additive model; Klein et al (1985) arrive at conditional utility functions, Thurston (1991) formalizes the design steps using a multiattribute utility function. More recently, Malak et al. (2009) use set-based design in the early design phase; while Abbas (2009) introduces multiattribute utility copulas. A dissenting line of research asserts (Wassenaar and Chen, 2001; Abbas and Cadenbach, 2018) that, when it comes to design for the market, its merit should be modeled by one attribute, typically in monetary units. It is unclear if modeling all design attributes in terms of dollar amounts will not be just as arduous as establishing a multi-attribute utility function. Therefore, in this paper, we use the well-established multi-attribute utility framework (Thurston, 2001; Keeney, 1972; Keeney and Raiffa, 1976). With regard to the presence of multiple decision makers, we acknowledge the challenge posed by Arrow's Impossibility Theorem, which states that it is impossible to aggregate a group of preference behaviors in a way that obeys some reasonable axioms (Arrow, 1950). Researchers since Arrow have shown that in practical decision making, at least some of the axioms can be reasonably relaxed. Goodman and Markowitz (1952) show that relaxing Arrow's requirement on irrelevant alternatives provides multiple avenues for aggregating preferences. Similarly, Keeney (1976) shows that given a cardinal utility function, u_i, from n decision makers, a grouped utility, u_G, is possible. Scott and Antonsson (1999) describe how design decision problems are not entirely constrained by Arrow's axioms.

Whether errors or inconsistencies appear during assessments of preferences of a single decision maker, or because of the existence of multiple decision makers, a probabilistic model of these variations can look similar (Debreu, 1958; Cyert and DeGroot, 1975; Manski, 1977; Karni and Safra, 2016). The idea of stochasticity in preferences has seen some attention in the literature (Blavatskyy, 2006, 2008). Most researchers describe the utility of an attribute at value x, $U(x)$, as a random variable with some probability distribution (Becker et al, 1963; Barseghyan et al, 2018; Loomes and Sugden, 1995; Chajewska et al, 2000). There are also techniques such as multi-dimensional scaling (Young, 2013), demand model estimation (Hauser and Rao, 2004), and conjoint analysis (Rao, 2014) that address variants of this problem.

Keeney showed, under some assumptions, that a group utility function over uncertain outcomes is possible if and only if $u(u_1, u_2, \ldots u_n\ ) = \sum_{i=1}^{n} \eta_i u_i$ (Keeney, 1976). We regard the scaling factor η_i for each utility function as the relative occurrence over a population of decision makers directly achievable if we make it proportional to the value of the probability distribution of the underlying random variable (parameters). This approach achieves two outcomes simultaneously: addressing Hazelrigg's objection to the existence of customer satisfaction in the aggregate (Hazelrigg, 1998) and extending the theoretical basis of *expected* expected utility Boutilier (2003). The work also opens the doors to a Bayesian approach which accommodates the co-evolution of end-user preference from prior to posterior as generations of systems go through cycles of acquisition and operation.

3. Mobility model used

In this section we briefly describe the mobility model which captures the input-output relationships for the ground vehicles exhibiting semi to full autonomy. The data needed for the development of a model of this type typically comes from physical tests, simulations or a combination of these. While many relationships have been derived for one or two aspects vehicle mobility, a comprehensive model is difficult to acquire. Physical tests, even high-fidelity simulations, tend to be expensive and time consuming. When considering multiple vehicle parameters, running an exhaustive combination of inputs takes an enormous amount of time and resources. To demonstrate our approach, in this paper, we combine some existing relationships generally seen in the literature with parameterized functions derived from experience and generally expected trends. If further information becomes available, it can be readily accommodated. Table 1 shows the main elements of the model and the related optimization problem, additional relationships are presented in the appendix. The objective is to maximize the expectation of the multiobjective utility function over four attributes in the vector $\mathbf{y}$, which are cost per mission, c, vehicle speed, s, trafficable area, a, and the normalized lane keeping error, ξ. The attributes in the vector $\mathbf{y}$ are functions of the inputs $\mathbf{x}$, and the vehicle characteristics in $\mathbf{v}$. The inputs considered are terrain roughness, δ_s and the soil strength s_s, therefore, $\mathbf{x} = (\delta_s, s_s)^T$. We acknowledge that many additional methods for terrain material characterization exist, such as soil moisture content, void ratio, density among others, however, there appears to be little consensus on which of these should be utilized (Shoop, 1993; Pinto, 2012; Umsrithong, 2010). The vehicle characteristics include inputs such as teleoperation fraction, t_t and vehicle intelligence level, ρ, among others. The single attributes utility functions are exponential, normalized between 0 and 1 as shown in equations. To calculate vehicle attributes, we consider operation in two modes, teleoperation and autonomous. The variable $t_t \in [0,1]$ is 1 for a teleoperated vehicle and 0 for an autonomous vehicle. A number strictly between 0 and 1 spans the continuum between teleoperation and full autonomy. We here outline how the attribute of speed is calculated, other attributes are calculated in a similar fashion. The speed of a vehicle, s, is the maximum speed of the vehicle, s_{max}, under the given conditions scaled by the speed ratio $s_r < 1$ as shown in equation A.1. This scaling factor takes into account that a vehicle need not be operated at its maximum speed. The maximum speed s_{max} is the weighted average of the speed during teleoperation, and that during autonomy as given in equation A.2. In the equation, s_t and s_a are the baseline speeds under the two modes and the fractions $\eta_t^s(\mathbf{x})$ and $\eta_a^s(\mathbf{x})$ model the cumulative effect of the other variables. Attributes of trafficable area and error are also similarly calculated, as shown in the appendix. Existing relationships are used from the literature, wherever available. For example, the speed-latency and speed-error relationships in equations A.7 and A.21 match Gorsich et al. (2018). Similarly, the TPA-speed characteristics are qualitatively similar to Lessem et al. (1996) and roughness-speed relationship is qualitatively similar to Vong et al (1999). If updated relationships are available for a vehicle under consideration, let us say those from physics based models for a specific vehicle, these equations can be updated without affecting the applicability of the remainder of the model.

Table 1: Formulation of the mobility model. Additional relationships described in the appendix.

Optimization Problem

$$\underset{\mathbf{v}}{\text{Maximize}}\ \mathbb{E}[U(\mathbf{y}(\mathbf{x},\mathbf{v}))] \tag{1}$$

Where:

$$\mathbf{y} = (c(\mathbf{x},\mathbf{v}), s(\mathbf{x},\mathbf{v}), a(\mathbf{x},\mathbf{v}), t(\mathbf{x},\mathbf{v}))^T \tag{2}$$

$$\mathbf{x} = (\delta_s, s_s)^T \tag{3}$$

$$\mathbf{v} = (c, t_t, \rho, \rho_{terrain}, \tau, a_0, s_t, s_a, s_r)^T \tag{4}$$

Utility functions

$$U(c,s,a,\xi) = \frac{1}{K}\left[(Kk_c U_c(c)+1)(Kk_s U_s(s)+1)(Kk_a U_a(a)+1)\left(Kk_\xi U_\xi(\xi)+1\right)-1\right] \tag{5}$$

$$U_c(c) = \frac{1}{n_c}\left(1-e^{-\frac{c_{max}-c}{R_c}}\right) \quad (6); \qquad U_s(s) = \frac{1}{n_s}\left(1-e^{-\frac{s-s_{min}}{R_s}}\right) \quad (7)$$

$$U_a(a) = \frac{1}{n_a}\left(1-e^{-\frac{a-a_{min}}{R_a}}\right) \quad (8); \qquad U_\xi(\xi) = \frac{1}{n_\xi}\left(1-e^{-\frac{\xi_{max}-\xi}{R_\xi}}\right) \quad (9)$$

4. Approach

Through a set of lottery questions one assesses the single attribute utility functions in equations 6-9. Further assessment questions lead to the multiattribute utility function of equation 5, which serves as an adequate way to measure the worth of multiattribute alternatives under uncertainty. In single attribute utility function assessments, for a sequence of values of an attribute in question, one gets the corresponding values for the utility which are then fitted to a smooth curve. Expectedly, fitting this curve by selecting the values of utility function parameters usually ignores the ramifications of representing many different risk behaviors with one curve (Ambrus et al, 2015), something we intend to address in this work. For the exponential functions shown in the equations 6-9, it is easy to show that a change in the value of R affects a decision substantially, both in the valuation of alternatives and also in preferences over risk. Given the probability distribution of $\mathbf{R}$, the vector of risk tolerances, the expected utility of equation 1 is now calculated over the distribution of both the outcome and the preferences as:

$$EU = \int_{D_R} \mathbb{E}_{\mathbf{x},\mathbf{v}}[U(\mathbf{y}(\mathbf{x},\mathbf{v}),\mathbf{r})]\, f_{\mathbf{R}}(\mathbf{r})d\mathbf{r} \qquad (10)$$

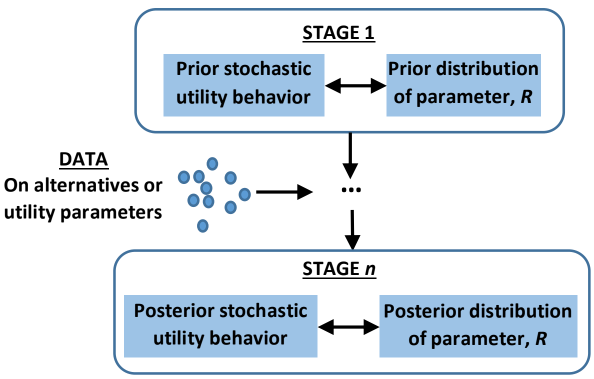

Figure 2. Summary of the proposed approach. Posterior from a stage acts as the prior for the next cycle of updating.

To model its probability distribution, we define a prior distribution on the risk tolerance R for each attribute individually assuming independence. This is a reasonable assumption since when using the multilinear utility function, the individual attribute utility functions are assessed independently. If there is dependence between attribute risk tolerances, the updating process presented later will automatically include it. The degree of our ignorance inversely corresponds to the degree of significance we place on the prior utility and this serves as a *starting point* for Bayesian updating. Figure 2 describes this approach. Notice that at each step, the assessment questions can be posed to the decision maker(s) to get data on R, or their assessment of existing designs can be directly noted, hence the bidirectional arrows in each stage. The latter option provides us the distribution of utilities which must then be converted to data on the distribution of R. This is straightforward as it can be shown for a single valued monotonic function $U = q(R)$ we can relate the pdf of U and R:

$$g_U(u) = f_R\left(q^{-1}(u)\right)\left|\frac{dq^{-1}(u)}{du}\right| \qquad (11)$$

Where $q^{-1}(u)$ is the inverse of the utility function in terms of R. Note that the attribute value is also an argument to the utility function but has been dropped for clarity.

Given data on U for s values, or the new values for R directly, we can get data for an attribute as $D = (r_1 \dots r_s)$. We can then determine the posterior distribution of R based on Bayes Theorem (Holicky, 2013):

$$f_R^{post}(r|D) \propto L(r|D) f_R^{prior}(r) \qquad (12)$$

Given the posterior distribution of the utility shape parameter R, we can now update the utility function.

Prior selection

Under conjugacy, which depends on certain assumptions for the types of prior and the likelihood functions, the posterior can be calculated without expensive calculations. Assuming normally distributed risk tolerance:

$$f_{\mathbf{R}}(\boldsymbol{r}) = \frac{1}{(2\pi)^{p/2}|\boldsymbol{\Sigma}|^{1/2}}\, exp\left[-\frac{(\boldsymbol{r}-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\boldsymbol{r}-\boldsymbol{\mu})}{2}\right] \qquad (13)$$

With priors defined as follows:

$$\boldsymbol{r}|\boldsymbol{\mu},\boldsymbol{\Sigma} \sim N(\boldsymbol{\mu},\boldsymbol{\Sigma}) \qquad (14)$$

$$\boldsymbol{\mu}|\boldsymbol{\Sigma} \sim N(\boldsymbol{\mu_0}, \frac{1}{n_0}\boldsymbol{\Sigma}) \qquad (15)$$

$$\boldsymbol{\Sigma} \sim Wi^{-1}(\alpha, \boldsymbol{\Psi}) \qquad (16)$$

Where

- $\boldsymbol{\mu_0}$ is the prior mean of $\boldsymbol{\mu}$.
- n_0 reflects the prior confidence and can be selected manually.
- α is the degrees of freedom of the inverse Wishart distribution where $\alpha > p - 1$; since the mean of the inverse Wishart is $\frac{\Sigma}{\alpha-p-1}$, $\alpha = p + 2$ will give a mean value of $\boldsymbol{\Sigma}$.
- $\boldsymbol{\Psi}$ is the conjugate prior of $\boldsymbol{\Sigma}$. It must be non-singular and symmetric.

The posterior is then given directly by the following expressions (Gill, 2014):

$$\boldsymbol{\mu}|\boldsymbol{\Sigma} \sim N\left(\frac{n_0\boldsymbol{\mu_0}+n\bar{\boldsymbol{r}}}{n_0+n}, \frac{1}{n_0+n}\boldsymbol{\Sigma}\right) \qquad (17)$$

$$\boldsymbol{\Sigma} \sim Wi^{-1}\left(\alpha + n, \boldsymbol{\Psi} + \bar{\boldsymbol{\Sigma}} + \frac{n_0 n}{n_0+n}(\bar{\boldsymbol{r}} - \boldsymbol{\mu_0})(\bar{\boldsymbol{r}} - \boldsymbol{\mu_0})^T\right) \qquad (18)$$

5. Vehicle selection decision problem

Let us consider the decision problem involving selection between eight vehicle offerings given below in Table 2. The general trends observable in the table are that the vehicle autonomy level increases from vehicle 1 to 8, and concurrently, so does the cost of operating them per mission. Vehicles 1 through 3 are teleoperated vehicles while the rest can both be teleoperated or operated in full autonomy. The teleoperation fraction lists the fraction of the time each vehicle is teleoperated. The rest of the parameters list vehicle characteristics such as base speed under teleoperation and that under autonomy. All vehicles are assumed to have some level of basic assistive features such as traction control and ABS included under the variable $\rho_{terrain}$.

The terrain input variables of roughness, Δ_s, and soil strength, S_s are assumed to be normally distributed with parameters given by: $\mu_{\delta_s} = 0.8$ inch rms, $\mu_{s_s} = 80$ Cone Index and $\boldsymbol{\Sigma} = \begin{bmatrix} 0.04 & 0 \\ 0 & 900 \end{bmatrix}$. The value of the soil strength is within the typical range specified for most types of soils (Priddy, 1995).

Table 2. Vehicle parameters for the eight types considered.

Vehicles	Cost ($)	Tele. Fraction	Latency	Auto. Level	Base speed	Base speed (autonomy)	$\rho_{terrain}$	Base area	Speed ratio
1	4000	1	0.4	0	55	0	0.9	90	0.8
2	5000	1	0.4	0	55	0	0.9	95	0.7
3	7000	1	0.4	0	60	0	0.9	95	0.8
4	7500	0.6	0.4	0.2	65	70	0.95	95	0.8
5	9500	0.4	0.4	0.4	65	75	0.95	95	0.8
6	12000	0.3	0.5	0.6	65	70	0.95	95	0.7
7	13000	0.3	0.5	0.6	65	75	0.95	95	0.7
8	18000	0.2	0.4	0.8	70	80	1	95	0.8

We start with the prior distributions on the mean and standard deviations for the risk tolerance for the attributes. As mentioned earlier, these represent the variation because of the assessment procedure(s) used or because of the multiplicity of the decision makers involved. In our work, these are sampled from first using an inverse Wishart prior for the covariance matrix, and using that information to then generate realizations of the mean. This provides us the parameters to generate the realizations of the risk tolerances. For demonstration purposes, we consider the risk tolerance for cost, speed and TPA as random while the risk tolerance for error is fixed at 5. The histograms of the three risk tolerances are given in Figure 3. The mean vector is given by $\boldsymbol{\mu_R} = \left(\mu_{R_c}, \mu_{R_s}, \mu_{R_a}\right)^T = (14000, 50, 70)^T$and the standard deviations are given by $\sigma_{R_c} = 1000$; $\sigma_{R_s} = 20$; and $\sigma_{R_a} = 20$. Notice that it is initially assumed that that risk tolerances are independent. This is a reasonable assumption since when using the multilinear utility function, the individual attribute utility functions are assessed independently.

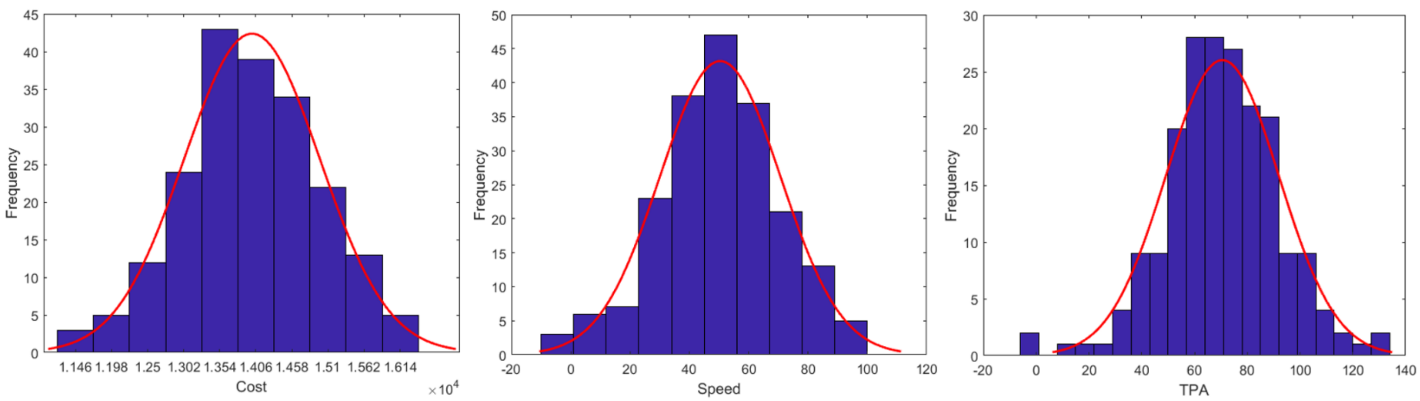

Figure 3. Histograms of risk tolerances associated with the three attributes.

Table 3 below shows the baseline expected-*expected* utilities (EU) that were calculated for the eight vehicles. The EU is calculated numerically using Matlab by averaging over many realizations of the vehicle input random variables and then by averaging over the realizations of the risk-tolerances. We see that vehicle 1 has the highest utility and is preferred over the others. The ranking of the vehicles suggests that the low cost nature of vehicles 1, 2, and 3 make them attractive for the given set of decision makers, with the given set of preferences.

Table 3. Prior utilities of the eight vehicles considered.

Vehicle	Utility
1	**0.5726**
2	0.566
3	0.569
4	0.5346
5	0.5637
6	0.5646
7	0.5594
8	0.5663

Let us now assume that a number of vehicle 1 have been acquired and fielded, and based upon observing their performance the decision makers have updated their preferences. These could be in the form of ranking of all the fielded vehicles, assessments of new utility functions or any combination of these. This data can then be incorporated into the Bayesian setup described earlier. The plot of the prior distribution, simulated example data, and the posterior are shown pairwise as two dimensional slices in Figures 4(a), (b) and (c). The resulting posterior distribution of risk tolerances is found to be $\boldsymbol{\mu}_{\boldsymbol{R}}^{\boldsymbol{post}} = \left(\mu_{R_c}^{post}, \mu_{R_s}^{post}, \mu_{R_a}^{post}\right)^T = (9977, 60.1, 81.3)^T$ and standard deviations are given by $\sigma_{R_c}^{post} = 776$; $\sigma_{R_s}^{post} = 14$; and $\sigma_{R_a}^{post} = 12.2$. This indicates that the mean risk tolerance decreased for cost, while it increased for the other two variables. Additionally, there is a decrease in the standard deviation for all the variables, as expected. There still is a significant amount of variation, especially in figure 4(c) because only 50 data points were used for updating. It was also observed that the variables now showed some dependence because of the mean shift term of equation 18, however, the correlation values were found to be mostly negligible. It may be interesting to investigate the induced dependence between attribute risk tolerances as more information about preferences is acquired, in a future work.

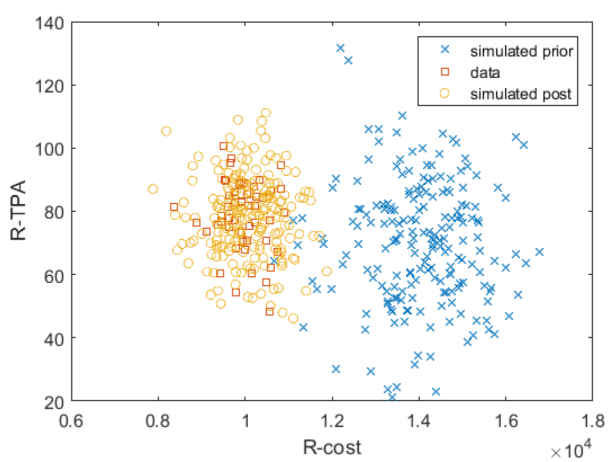

(a)

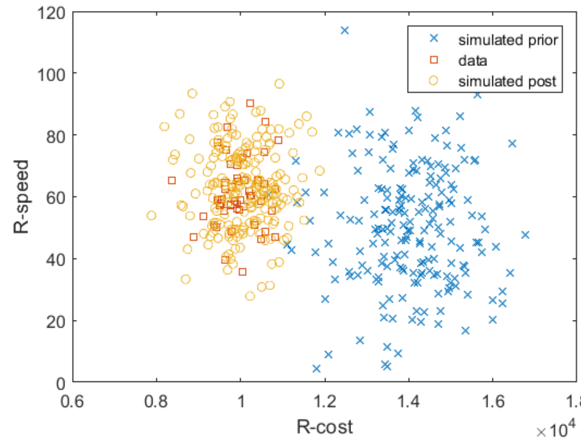

(b)

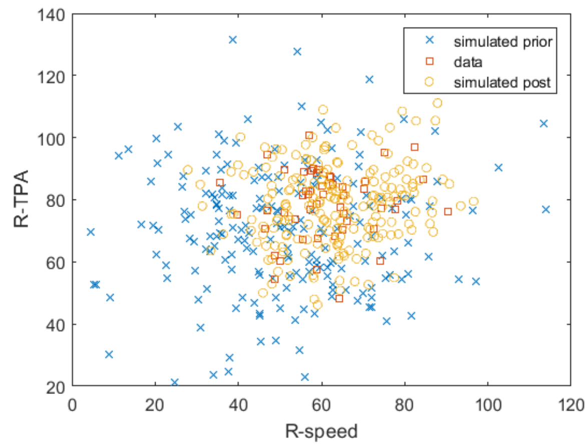

(c)

Figure 4. Scatter plots of the prior and posterior distributions and the incoming data for the three random variables.

We now plot the single attribute utility functions based on prior and posterior values of risk tolerance. The plots for the three attributes are shown in Figure 5, clearly showing the difference in preferences and the possible variations in the decisions made. The posterior cost utility functions, figure 5(a), show that the decision makers became more risk averse (increased concavity). This implies that they are more likely accept an expensive vehicle in this case, as even a small reduction in cost from the maximum acceptable $20,000 is enough to increase their utility functions substantially. Looking at the speed utility function, figure 5(b) we see that the decision makers in general have become less risk averse. This translates to them requiring better performance from the vehicles on the speed attribute. Similarly, for the trafficable area attribute, figure 5(c), we see that the decision makers have become less risk averse. All together, the updated preferences imply that the decision makers are more inclined to pay more for improved performance of the vehicles. It is interesting to note that some of the decision makers in the prior assessment actually exhibited risk-seeking behavior as evidenced by a few realizations of convex utility functions shown. This is expected in a group of decision makers with differing preferences and corresponds to a negative value of risk tolerance.

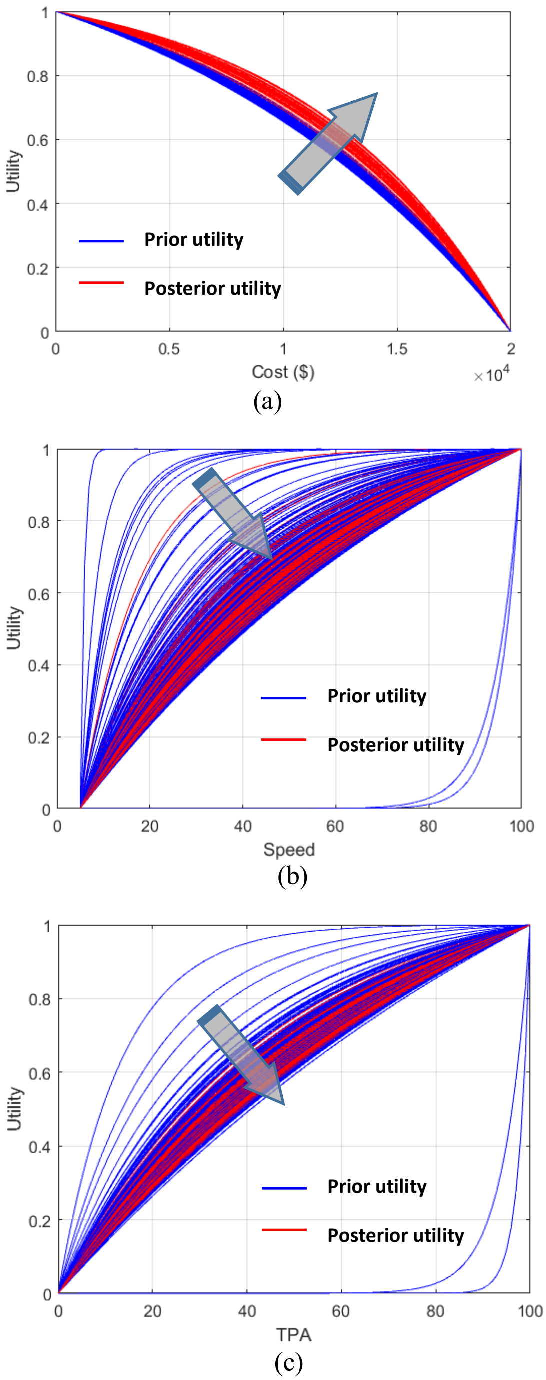

Figure 5. Plots of the prior and posterior utility functions for the three attributes.

The posterior utilities for the vehicles are shown below in Table 4. We see that in this case vehicle 6 has the highest utility and should be preferred if only one vehicle is to be selected by the organization. A close second is vehicle 8, while vehicle 1 underperforms compared to before. It is important to realize that utility values from tables 3 and 4 cannot be directly compared, only the relative rankings based on the utilities in each case can be compared. The change in the rankings is in line with the type of change exhibited by the decision makers that tend to now prefer a higher performing vehicle and are willing to accept an increased cost. A look at the average attribute values for vehicles 1, 6 and 8, shown in Figure 6, highlights this choice. Notice that some attributes have been scaled to better represent them on the chart. We see that both vehicles 6 and 8 outperform vehicle 1 on all attributes except cost. Vehicle 6, underperforms vehicle 8 on speed and error, but performs almost as well on TPA. The cost advantage ensures that the overall utility of vehicle 6 is higher.

Table 4. Posterior utilities of the eight vehicles considered.

Vehicle	Utility
1	0.5525
2	0.549
3	0.5531
4	0.5203
5	0.5515
6	**0.5549**
7	0.5499
8	0.5546

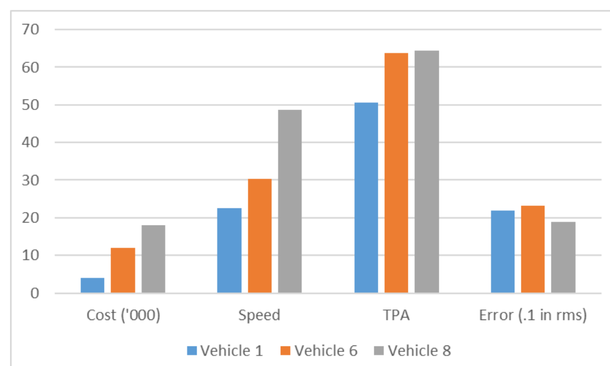

Figure 6. Relative attribute levels achieved by vehicles 1, 6 and 8.

Another way to look at how vehicles fare is to investigate the utilities of individual decision makers and based on it, the vehicles they would individually prefer. Recall that earlier we had averaged the decision maker utilities (equation 10). From the simulation of 200 decision makers based on posterior preferences, it is seen that about 67% of the decision makers prefer vehicle 6 over vehicle 1 *in a pairwise comparison*. Recall that vehicle 1 was preferred using prior preference data. This implies that if the decision were not changed based on the current distribution of risk tolerances, a majority of decision makers will have a suboptimal vehicle. Figure 7 below shows percentage wins for all the vehicles in the two cases. We see that between the prior and the posterior preferences, the choices are dominated by vehicles 1, 6 and 8. This information can be useful in that one can choose to remove 2, 3, 4 and 7 from further consideration. Another interesting observation is that while vehicle 8 does not win in either the prior or posterior case, it seems robust to changes in preferences. It is important to note that, while calculating the percentage of decision makers that prefer a vehicle gives insights to individual decisions, it may not always be a useful approach for two reasons. Firstly, in group decision making, it may be necessary to select just one solution as a group. Secondly, it is possible for decision makers that do not prefer, say vehicle 8, to have strong preferences against it. Therefore, using expected utility is the recommended method for making decisions in these cases (tables 1 and 4).

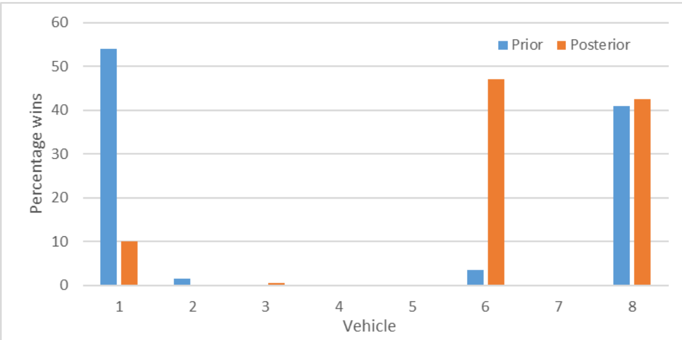

Figure 7. Percentage of decision makers preferring each vehicle.

6. Discussion and conclusions

In this paper, we addressed the issue of uncertain decision maker preferences in acquisition decisions focused on autonomous ground vehicles. Autonomous ground vehicles keep improving in their capabilities, challenging the current methods of evaluating vehicle performance. This is partly because, given the uncertainty about their future capabilities and the types of missions they may become capable of, it is unclear what an operator's preferences regarding them would be. In our approach, we accommodate uncertainty in DM preferences by modeling their utility function shape parameters as random variables. We then take a Bayesian perspective and use preference data that may be acquired from vehicle operation, to update the probability distribution of the parameters of the utility function used in the acquisition phase. This *posterior* stochastic utility function can then become the *prior* stochastic utility function for the next generation of the product. We applied our approach on selection decisions involving vehicles exhibiting partial to full autonomy. Our approach showed how one can accommodate the variation in the preferences of the decision makers to make a decision. Furthermore, if more information becomes available, how one can update the distribution of variables modeling these preferences. We showed how this resulted in a changed decision and also showed the effect of not updating the decision as new data became available.

We restricted the application to the parameters of single attribute utility functions. In the future, we intend to also consider the variation in the scaling constants for the multiattribute utility function. This way we can account for variation in the *tradeoff behavior* between different decision makers. Along the same lines, the assumption that all decision makers have the same worst and best case of each attribute may not always be correct and we will attempt to relax this assumption. We intend to also extend our approach to non-normal distributions as well.

The strength of a Bayesian approach is the choice of the prior, and contrastingly, it also is its biggest weakness. In our application, we selected distribution parameters that would best differentiate between the prior and the posterior. We believe exploration of the sensitivity to the distribution form and the parameter values is worth pursuing. Finally, we believe this is a powerful tool for continuous improvement of the decision making process, especially for problems as dynamic as encountered in vehicle autonomy. In large organizations, this approach and its future extensions should provide a feedback loop that translates anticipated design needs with experience from the field in the continuous cycle of design development.

ACKNOWLEDGEMENT

This work was supported by the Automotive Research Center (ARC) in accordance with Cooperative Agreement (W56HZV-14-2-0001) U.S. Army Ground Vehicle Systems Center, Warren, MI.

REFERENCES

1. Abbas, A.E. and Cadenbach, A.H., 2018. On the use of utility theory in engineering design. IEEE Systems Journal, 12(2), pp.1129-1138.
2. Abbas, A.E., 2009. Multiattribute utility copulas. Operations Research, 57(6), pp.1367-1383.
3. Ahlvin, R.B. and Haley, P.W., 1992. NATO reference mobility model: Edition II. NRMM user's guide. US Army Engineer Waterways Experiment Station.
4. Ambrus, A., Greiner, B. and Pathak, P.A., 2015. How individual preferences are aggregated in groups: An experimental study. Journal of Public Economics, 129, pp.1-13.
5. Arrow, K.J., 1950. A difficulty in the concept of social welfare. Journal of Political Economy, 58(4), pp.328-346.
6. Babu, J.R. and Asha, A., 2014. Tolerance modelling in selective assembly for minimizing linear assembly tolerance variation and assembly cost by using Taguchi and AIS algorithm. The International Journal of Advanced Manufacturing Technology, 75(5-8), pp.869-881.
7. Barseghyan, L., Molinari, F., O'Donoghue, T. and Teitelbaum, J.C., 2018. Estimating risk preferences in the field. Journal of Economic Literature, 56(2), pp.501-64.
8. Becker, G.M., DeGroot, M.H. and Marschak, J., 1963. Stochastic models of choice behavior. Behavioral Science, 8(1), pp.41-55.
9. Bentham, J., 1879. An introduction to the principles of morals and legislation. Clarendon Press.
10. Blavatskyy, P.R., 2006. Stochastic choice under risk. Institute for Empirical Research in Economics, Working Paper, (272).
11. Blavatskyy, P.R., 2008. Stochastic utility theorem. Journal of Mathematical Economics, 44(11), pp.1049-1056.
12. Boutilier, C., 2003, August. On the foundations of expected expected utility. In IJCAI (Vol. 3, pp. 285-290).
13. Chajewska, U., Koller, D. and Parr, R., 2000, July. Making rational decisions using adaptive utility elicitation. In AAAI/IAAI (pp. 363-369).
14. Chen, W., Hoyle, C. and Wassenaar, H.J., 2012. Decision-based design: Integrating consumer preferences into engineering design. Springer Science & Business Media.
15. Chowdhury, S., Design for Six Sigma: The Revolutionary process for achieving extraordinary profits 2002. Chicago, IL: Dearborn Trade Pub. xvii.
16. Cyert, R.M. and DeGroot, M.H., 1975. Adaptive utility. In Adaptive Economic Models (pp. 223-246). Academic Press.
17. Debreu, G., 1958. Stochastic choice and cardinal utility. Econometrica: Journal of the Econometric Society, pp.440-444.

18. Department of Defense, Defense Science Board, Task Force Report, 2012: "The Role of Autonomy in DoD Systems", Accessed from https://fas.org/irp/agency/dod/dsb/autonomy.pdf. Date Accessed 10/20/2019.
19. Edwards, W., 1977. How to use multiattribute utility measurement for social decision making. IEEE Transactions on Systems, Man, and Cybernetics, 7(5), pp.326-340.
20. Farquhar, P.H., 1984. State of the art—utility assessment methods. Management Science, 30(11), pp.1283-1300.
21. Friedman, M. and Savage, L.J., 1948. The utility analysis of choices involving risk. Journal of Political Economy, 56(4), pp.279-304.
22. Gill, J., 2014. Bayesian methods: A social and behavioral sciences approach. Chapman and Hall/CRC.
23. Goodman, L.A. and Markowitz, H., 1952. Social welfare functions based on individual rankings. American Journal of Sociology, 58(3), pp.257-262.
24. Gorsich, D.J., Jayakumar, P., Cole, M.P., Crean, C.M., Jain, A. and Ersal, T., 2018. Evaluating mobility vs. latency in unmanned ground vehicles. Journal of Terramechanics, 80, pp.11-19.
25. Hauser, J.R. and Rao, V.R., 2004. Conjoint analysis, related modeling, and applications. In Marketing Research and Modeling: Progress and Prospects (pp. 141-168). Springer, Boston, MA.
26. Hazelrigg, G.A., 1998. A framework for decision-based engineering design. ASME Journal of mechanical design, 120(4), pp.653-658.
27. Hoffenson, S. and Söderberg, R., 2015. Systems thinking in tolerance and quality-related design decision-making. Procedia CIRP, 27, pp.59-64.
28. Holický, M., 2013. Introduction to probability and statistics for engineers. Springer Science & Business Media.
29. J. von Neumann and O. Morgenstein, Theory of Games and Economic Behavior. Princeton, NJ: Princeton Univ. Press, 1947
30. Karni, E. and Safra, Z., 2016. A theory of stochastic choice under uncertainty. Journal of Mathematical Economics, 63, pp.164-173.
31. Keeney, R.L. and Raiffa, H., 1976. Decision analysis with multiple conflicting objectives. Wiley & Sons, New York.
32. Keeney, R.L., 1972. Utility functions for multiattributed consequences. Management Science, 18(5-part-1), pp.276-287.
33. Keeney, R.L., 1976. A group preference axiomatization with cardinal utility. Management Science, 23(2), pp.140-145.
34. Klein, G., Moskowitz, H., Mahesh, S. and Ravindran, A., 1985. Simplified Assessment of Single and Multi-Attribute Utility Functions, Working Paper 85-904.
35. Lessem, A., Mason, G. and Ahlvin, R., 1996. Stochastic vehicle mobility forecasts using the NATO reference mobility model. Journal of Terramechanics, 33(6), pp.273-280.
36. Loomes, G. and Sugden, R., 1995. Incorporating a stochastic element into decision theories. European Economic Review, 39(3-4), pp.641-648.
37. Malak Jr, R.J., Aughenbaugh, J.M. and Paredis, C.J., 2009. Multi-attribute utility analysis in set-based conceptual design. Computer-Aided Design, 41(3), pp.214-227.
38. Manski, C.F., 1977. The structure of random utility models. Theory and decision, 8(3), pp.229-254.
39. Pandey, V., 2013. Decision based design. CRC Press, pp 284.
40. Pinto, E.J., 2012. A Three Dimensional Discretized Tire Model for Soft Soil Applications (Doctoral dissertation, Virginia Tech).
41. Priddy, J.D., 1995. Stochastic Vehicle Mobility Forecasts Using the NATO Reference Mobility Model. Report 3. Database Development for Statistical Analysis of the NRMM II Cross-Country Traction Empirical Relationships (No. WES/TR/GL-95-8).
42. Rao, V.R., 2014. Applied conjoint analysis (p. 389). New York, NY: Springer.
43. Scott, M.J. and Antonsson, E.K., 1999. Arrow's theorem and engineering design decision making. Research in Engineering Design, 11(4), pp.218-228.
44. Shoop, S.A., 1993. Terrain characterization for trafficability (No. CRREL-93-6). Cold Regions Research and Engineering Lab Hanover NH.
45. Smith, W.F., Milisavljevic, J., Sabeghi, M., Allen, J.K. and Mistree, F., 2015, August. The realization of engineered systems with considerations of complexity. In ASME 2015 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference (pp. V007T06A019-V007T06A019).
46. Taheri, S., Sandu, C., Taheri, S., Pinto, E. and Gorsich, D., 2015. A technical survey on Terramechanics models for tire–terrain interaction used in modeling and simulation of wheeled vehicles. Journal of Terramechanics, 57, pp.1-22.
47. Thurston, D.L., 1991. A formal method for subjective design evaluation with multiple attributes. Research in Engineering Design, 3(2), pp.105-122.
48. Thurston, D.L., 2001. Real and misconceived limitations to decision based design with utility analysis. Journal of Mechanical Design, 123(2), pp.176-182.
49. Umsrithong, A. and Sandu, C., 2010. A 3D semi-empirical on-road transient tire model. SAE International Journal of Commercial Vehicles, 3(2010-01-1916), pp.42-59.
50. Vong, T.T., Haas, G.A. and Henry, C.L., 1999. Nato reference mobility model (NRMM) modeling of the demo III experimental unmanned ground vehicle (XUV) (No. ARL-MR-435).
51. Wassenaar, H.J. and Chen, W., 2001. An approach to decision-based design. ASME Paper No. DETC01/DTM-21683.
52. Young, F.W., 2013. Multidimensional scaling: History, theory, and applications. Psychology Press.

APPENDIX

Attribute calculations

$$s(\mathbf{x},\mathbf{v}) = s_r s_{max}(\mathbf{x},\mathbf{v}) \quad \text{A.1}$$
$$s_{max}(\mathbf{x},\mathbf{v}) = t_t s_t \eta_t^s(\mathbf{x},\mathbf{v}) + (1-t_t) s_a \eta_a^s(\mathbf{x},\mathbf{v}). \quad \text{A.2}$$
$$\eta_t^s(\mathbf{x},\mathbf{v}) = \eta_{\delta_s}(\delta_s)\eta_{s_s}(s_s)\eta_\tau(\tau) \quad \text{A.3}$$

$$\eta_a^s(\mathbf{x},\mathbf{v}) = \eta_{\delta_s}(\delta_s)\eta_{s_s}(s_s)\eta_\rho(\rho) \quad \text{A.4}$$
$$\eta_{\delta_s}(\delta_s) = e^{-\left(\frac{\delta_s}{1.1}\right)^2} \quad \text{A.5}$$
$$\eta_{s_s}(s_s) = 1 - e^{-\frac{s_s}{20}} \quad \text{A.6}$$
$$\eta_\tau(\tau) = -0.07e^{1.93\tau} + 1.07 \quad \text{A.7}$$
$$\eta_\rho(\rho) = 2.7182(1 - e^{-\rho}) \quad \text{A.8}$$
$$a(\mathbf{x},\mathbf{v}) = t_t a_0 \eta_t^a(\mathbf{x},\mathbf{v}) + (1 - t_t)a_0\eta_a^a(\mathbf{x},\mathbf{v}) \quad \text{A.9}$$
$$\eta_t^a(\mathbf{x},\mathbf{v}) = \eta_a^{s_t}(s_t)\eta_a^{\delta_s}(\delta_s)\eta_{t,a}^{int}(\rho_{terrain}) \quad \text{A.10}$$
$$\eta_a^a(\mathbf{x},\mathbf{v}) = \eta_a^{s_t}(s_t)\eta_a^{\delta_s}(\delta_s)\eta_{a,a}^{int}(\rho, \rho_{terrain}) \quad \text{A.11}$$
$$\eta_a^{s_t}(s_t) = 1 - \exp\left(-\frac{s_s}{20}\right) \quad \text{A.12}$$
$$\eta_a^{\delta_s}(\delta_s) = \exp\left(-\frac{\delta_s}{10}\right) \quad \text{A.13}$$
$$\eta_{t,a}^{int}(\rho_{terrain}) = exp\left(-\frac{s_r}{2\rho_{terrain}}\right) \quad \text{A.14}$$
$$\eta_{a,a}^{int}(\rho, \rho_{terrain}) = exp\left(-\frac{s_r}{2(\rho + \rho_{terrain})}\right) \quad \text{A.15}$$
$$\xi(\mathbf{x},\mathbf{v}) = t_t\eta_t^\xi(\mathbf{x},\mathbf{v}) + (1 - t_t)\eta_a^\xi(\mathbf{x},\mathbf{v}) \quad \text{A.16}$$
$$\eta_t^\xi(\mathbf{x},\mathbf{v}) = \eta_\xi^{s_s}(s_s)\eta_\xi^{\delta_s}(\delta_s)\eta_\xi^\tau(\tau) \quad \text{A.17}$$
$$\eta_a^\xi(\mathbf{x},\mathbf{v}) = \eta_\xi^{s_s}(s_s)\eta_\xi^{\delta_s}(\delta_s)\eta_\xi^a(\rho) \quad \text{A.18}$$
$$\eta_\xi^{s_s}(s_s) = 1 + \exp\left(-\frac{s_s}{50}\right) \quad \text{A.19}$$
$$\eta_\xi^{\delta_s}(\delta_s) = 2 - \exp\left(-\frac{\delta_s}{10}\right) \quad \text{A.20}$$
$$\eta_\xi^\tau(\tau) = 0.29e^{2.93\tau} + 0.71 \quad \text{A.21}$$
$$\eta_\xi^a(\rho) = 3e^{-\rho} \quad \text{A.22}$$