Predicting the Likely Behaviors of Continuous Nonlinear Systems in Equilibrium

Alexander Yeh*

MIT Laboratory for Computer Science
545 Technology Sq., NE43-413
Cambridge, MA 02139
USA

This paper introduces a method for predicting the likely behaviors of continuous nonlinear systems in equilibrium in which the input values can vary. The method uses a parameterized equation model and a lower bound on the input joint density to bound the likelihood that some behavior will occur, such as a state variable being inside a given numeric range. Using a bound on the density instead of the density itself is desirable because often the input density's parameters and shape are not exactly known. The new method is called SAB after its basic operations: *split* the input value space into smaller regions, *and* then *bound* those regions' possible behaviors and the probability of being in them. SAB finds rough bounds at first, and then refines them as more time is given. In contrast to other researchers' methods, SAB can (1) find all the possible system behaviors, and indicate how likely they are, (2) does not approximate the distribution of possible outcomes without some measure of the error magnitude, (3) does not use discretized variable values, which limit the events one can find probability bounds for, (4) can handle density bounds, and (5) can handle such criteria as two state variables both being inside a numeric range.

1 Introduction

This paper introduces a method called SAB to predict the likely behaviors of a continuous nonlinear system in equilibrium in which the input values can vary. SAB uses a parameterized equation model and a lower bound on the input joint density to bound the likelihood that one or more state variables stay inside or outside of a given set of numeric ranges (the likelihood of meeting some criteria).

The reason for using one or more bounds on the probability density, and not the density itself, is that density parameters (means, etc.) and density shape are often not exactly known: One may only have confidence intervals or estimated ranges for some parameters. In fact, the actual parameters may vary (periodically) over time, such as the mean blood pressures in the chest which rise and fall due to pressure changes caused by the breathing cycle. Even if one has point estimates for all the parameters, they may be unusable due to slight inconsistencies. For example, the correlation estimates may

*This research was supported in part by the National Heart, Lung, and Blood Institute through grant R01-HL33041.

be slightly inconsistent, especially if some correlations are estimated using intuition or heuristics, as opposed to applying statistics over a large data-base. By relaxing correlation estimates to be ranges of values around the original estimates, one can examine densities with that type of correlation structure. For example, if some correlation coefficient estimate of 0.8 was found to be inconsistent with the other density parameters, one could still do preliminary analysis with a density where that correlation coefficient is high (near 0.8) by allowing that coefficient to be anywhere in the range of 0.7 to 0.9 (assuming some value within this range is consistent with all the other density parameters).

An alternative to using bounds on the probability density is to use a sample of possible probability densities. However, sampling is not as complete an examination as bounding. Some class of important behaviors may lie between all the samples and not be observed.

Other prediction-making methods have one or more of the following problems: not finding the likelihood of behaviors or only finding the likelihood of the variable values falling in certain ranges; producing approximate results without estimating the error and being unable to improve on an initial result's accuracy when given more computation time; not being able to handle density bounds, or handling them too slowly; needing all the set(s) of input values that satisfy the criteria to be explicitly mentioned.

Compared to the these other techniques, SAB produces analytic bounds, improves its answers as more samples or iterations are allowed, and deals with distributions of continuous variable values.

The next section of this paper gives a more detailed description of some other methods. Section 3 gives a simple example of using SAB. It is followed by three sections which give in order an overview of SAB, a demonstration of how it runs in the simple example, and SAB's details. The paper ends with a discussion section and an appendix on some bounds derivations.

2 Other Techniques

Other prediction-making methods fit into one of four categories. The first category of methods finds all the possible system behaviors (sometimes including impossible ones), but does not tell the likelihood of the behaviors. Such methods include systems either performing qualitative reasoning [1, 2], or providing numeric bounds [3].

Category two methods estimate the distributions of possible outcomes without giving some measure of each estimate's error and will not improve the accuracy of those estimates when given more computation time. One of these methods is using moments [4]. This method uses truncated Taylor series expansions of the model equations to find various moments (mean, variance, etc.) of the distributions of interest.

A third category is evidential reasoners [5, 6], which include most of the current work done on uncertainty in AI. These reasoners can only handle a variable value in terms of the possibility of it belonging to one or more regions in a preset discretization of the possible variable values. For example, blood pressure (bp) may be only thought of in terms of being low, normal, or high. This limitation is a problem because what is considered normal, desirable, etc. can change with each use of a model. For example, when trying to lower a patient's bp, an acceptable pressure depends on the patient's

former normal blood pressure and the patient's ability to withstand therapy side-effects.

Monte Carlo techniques [4, 7, 8], which fall into two general classes, constitute the fourth category. The first class simulates a system by generating samples according to some probability distribution. Most methods in this class cannot handle density bounds. The acceptance/rejection method can handle density bounds, but it is too slow due to the large number of potential samples it rejects.

The second class of Monte Carlo techniques integrates the density or density bound involved. These integration techniques include hit-or-miss and sample-mean Monte Carlo. Unfortunately, determining the interval(s) to be integrated over (the region(s) satisfying the criteria) is very hard. As an illustration of this, consider the PVR example to be given in Section 3. The interval of all possible input values is the region defined by $PAP \in [10, 80]$, $LAP \in [2, 45]$, and $CO \in [1, 30]$. To find a bound on $\Pr(PVR > 1.62)$, one needs to integrate over the density bound in all the regions of possible input values where the criterion of $1.62 < PVR = (PAP - LAP)/CO$ is satisfied. Two of these regions are

$$PAP \in [19, 80], LAP \in [2, 15], CO \in [1, 2]$$
and
$$PAP \in [19, 80], LAP \in [2, 11], CO \in [2, 4].$$

Finding all such regions in the interval of all possible input values is difficult. Also, as with all Monte Carlo techniques, every answer is inexact and has a standard deviation associated with it.

3 Simple Example Using PVR

A simple example of using SAB involves finding a patient's pulmonary vascular resistance (PVR) given the constraint

$$PVR = (PAP - LAP)/CO \tag{1}$$

and information on the patient's pulmonary arterial pressure (PAP), left atrial pressure (LAP) and cardiac output (CO). PVR is of interest because a high value indicates that the heart's right ventricle has to work very hard to keep the blood moving through the lungs [9, p. 234]. One threshold condition is $PVR \leq 1.62 mmHg/(l/min)$. Critically ill surgical patients with values above this are less likely to survive [10, p.54-59].[1] PAP, LAP, and CO have patient and time dependent values, and are not easy to measure accurately. Table 1 gives some statistics for the patient of interest, a heart attack victim. The question is, given information on PAP, LAP, and CO for the patient involved, is PVR at all likely to be above the threshold? If so, one ought to monitor PVR.

The numbers are close enough so that the answer is not obvious from looking at Table 1: For example, substituting the mean values into Equation 1 results in $PVR < 1.62$, but increasing PAP's value in the substitution by 3.38 (one standard deviation) while maintaining LAP and CO's values would result in $PVR > 1.62$. However, the latter is not that likely to happen because LAP tends to increase when PAP does (high positive correlation).

[1] Assume that patients have a body surface area of $1.74 m^2$, the average for humans.

NAME	MEAN	STD DEV	Correlation Coef. PAP	LAP	CO
PAP	23.94	3.38	1.0	.861	.096
LAP	15.29	3.08	.861	1.0	-0.044
CO	6.49	1.20	.096	-0.044	1.0

Table 1: PVR Example

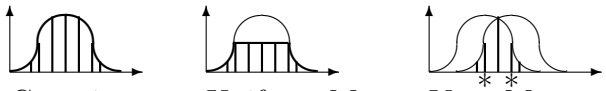

Figure 1: 3 Lower Density Bounds

So, one has to look at the joint density of PAP, LAP, and CO. Like most statistics, the ones in Table 1 are subject to sampling error, and in addition, the density shape is not exactly known. To get around this difficulty, one can hypothesize plausible bounds on the joint density and let SAB bound the probabilities of satisfying the criteria given each density bound. Ideally, the set of density bounds used will cover all the possible variations.

In this example, three lower density bounds are considered. They show the types of bounds that SAB can handle. One dimensional views of these are in Figure 1, where the areas under the density bounds are marked by vertical lines. As will be described later, the right-most bound covers all Gaussian densities where CO's mean is somewhere within a bounded interval and all the other parameters are as given in Table 1. The details on getting the results are given later on.

The first "bound" is a regular joint Gaussian density with the parameters listed in Table 1 and is shown in the left diagram of Figure 1. A 1000-sample Monte Carlo simulation with this bound (a normal probability density) indicates that $PVR > 1.62$ about 20% of the time. SAB analytically bounds this to be between 4% and 57%. This is consistent with the Monte Carlo simulation and with patient data, where 4 of 17 (23.5%) data points had $PVR > 1.62$.[2]

The second density bound is a

1. joint Gaussian density with the parameters listed in Table 1

2. in which the maximum value is limited to that of a jointly uniform density with the same means and standard deviations.

In other words, the density bound looks like a Gaussian far from the variables' means, but has the low flat top of a uniform density near the means. It is shown in the middle diagram of Figure 1. Integrating the bound indicates that it includes $\sim$ 70% of the probability mass. Using this bound SAB analytically bounds $\Pr(PVR > 1.62)$ to be between 4% and 79%. This is again consistent with the patient data.

The third density bound is the lower bound of a Gaussian density where CO's mean is allowed to be anywhere between 6.20 to 6.78.[3] This constraint might have been

[2]Here, the data could have been used by itself to answer the question of whether $PVR > 1.62$ is at all likely. SAB is meant to be used when such data is not available.

[3]The variances, covariances, and other means could also be allowed to vary.

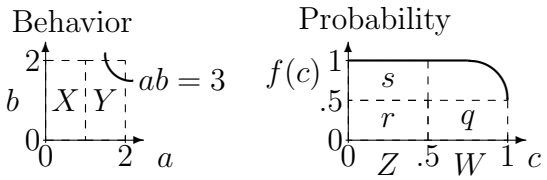

Figure 2: Examples of Splitting

determined by using information in some confidence interval for CO's mean. The right diagram of Figure 1 shows this bound: CO's mean can lie anywhere between the two *'s. The lower density bound is the intersection of the areas under all the densities possible due to allowable variations in CO's mean. Because Gaussian densities are unimodal, the lower density bound is the intersection of the areas under the two Gaussian density curves[4] shown. Integrating the bound indicates that it includes $\sim 65\%$ of the probability mass. Using this bound SAB analytically bounds $\Pr(PVR > 1.62)$ to be between 1% and 76%. This is also consistent with the patient data.

In all three input bound cases, $\Pr(PVR > 1.62) > 1\%$, so PVR should be monitored. Note that the results for each of the input bounds can be tightened. See the end of Section 5 for details.

4 SAB: Overview

SAB tightens the probability bound of achieving or failing some criteria by repeatedly selecting a region of possible input values, *splitting* that region into smaller regions α_i's, *and* then *bounding* both the possible behaviors within the α_i's and the probability of being in the α_i's (using the input probability density bound). SAB marks the α_i's whose possibilities always satisfy or fail the criteria.

Figure 2 shows two examples of splitting. In the one marked *Behavior*, the criterion is $ab < 3$, and the original region is $a, b \in [0, 2]$. In this region $ab \in [0, 4]$, so it sometimes passes and sometimes fails the criterion. Split the region along $a = 1$ into the two sub-regions X and Y. In X, $a \in [0, 1]$, so $ab \in [0, 2]$, and the criterion is always satisfied. Mark X. $a \in [1, 2] \rightarrow ab \in [0, 4]$ in Y, so Y is not marked. In the example marked *Probability*, α, the original region, is $c \in [0, 1]$, and $f(c)$ is a lower bound on probability density at c. SAB finds a lower bound on $\Pr(\alpha)$, the probability of being in α, of 0.5 (sum areas q and r) by multiplying 1, α's length, by 0.5, the lowest value of $f(c)$ in α.[5] Split the region at $c = 0.5$ into the two sub-regions Z and W. By a method similar to the one above, SAB finds a lower bound on $\Pr(Z)$ of 0.5 (sum areas r and s), and a lower bound on $\Pr(W)$ of 0.25 (area q). Sum the lower bounds of $\Pr(Z)$ and $\Pr(W)$ to get a new lower bound of 0.75 on $\Pr(\alpha)$.

As hinted by these two examples, *as long as* the bounding method used tends to reduce the range of possibilities as a region of input value space gets smaller, this continued splitting will mark more and more of the value space. And *as long as* the bounding method tends to reduce the gap between a density bound's upper and lower

[4] They are the ones with the extreme CO mean values.

[5] Better methods of bounding probabilities are described later.

bound[6] in a region as the region gets smaller, the bound on the probability of being in a marked region will improve.

To find a lower bound on Pr(satisfy criteria) sum the lower probability bounds of all the regions marked as satisfying the criteria. Similarly, one can find a lower bound on Pr(fail criteria). One minus the latter is an upper bound on Pr(satisfy criteria).

5 PVR Example Revisited

This section re-examines the introduction's PVR example when using the Gaussian density as a "bound" (first density bound). To bound $\Pr(PVR > 1.62)$, SAB looked at the space of inputs (given to SAB as one region):

$$PAP \in [1.0, 88.0], LAP \in [1.0, 88.0], CO \in [1.0, 100].$$

A lower bound on PVR, written $lb(PVR)$, is

$$\max(0, [lb(PAP) - ub(LAP)]/ub(CO)) = 0,$$

and an upper bound ($ub(PVR)$) is

$$[ub(PAP) - lb(LAP)]/lb(CO) = 87.0.$$

PVR can be either greater or less than 1.62, so SAB split the space in two along the CO dimension:

$$\begin{aligned} subspace1: &\quad PAP \in [1.0, 88.0], LAP \in [1.0, 88.0], CO \in [1.0, 50.5] \\ subspace2: &\quad PAP \in [1.0, 88.0], LAP \in [1.0, 88.0], CO \in [50.5, 100.0] \end{aligned}$$

SAB then checked and split as appropriate. Regions like

$$PAP \in [20.75, 25.47], LAP \in [15.95, 17.32], CO \in [6.41, 7.19], (PVR \in [0.756, 1.484])$$

where PVR is either always $>$, or ≤ 1.62, were marked. SAB found lower bounds on the probabilities of being in these marked regions (the one above has a probability ≥ 0.002).

As SAB recursively splits and checks regions, it tightens the probability bound for satisfying the criteria. When the bound is tight enough, or SAB runs out of time or another resource, it can be stopped. In this example, when SAB was stopped, it gave a lower bound of 0.438 on the probability of being in a passing region (one where $PVR \leq 1.62$), and 0.042 for a failing region ($PVR > 1.62$). If a tighter bound was desired, one could have restarted SAB with the then current set of regions. Since this input joint density bound includes all of the probability mass, SAB can, barring round-off error in the floating point math, get the bound to be arbitrarily tight if given enough computing time. In general, if an input joint density bound includes $n \times 100\%$ of the probability mass, SAB can, barring round-off error, get the bound to have a gap of $1.0 - n$ between the upper and lower figure. So if a density bound includes 70% of the probability mass, the tightest bound SAB could give on the chances of passing some criteria would have a gap of 0.3 between the lower and upper figures (such as a lower bound of 0.6 and an upper bound of 0.9).

[6] Yes, we are bounding a bound here.

6 SAB: Details

6.1 Main Loop

Perform the following cycle until told to stop:

1. Select the region α with the highest *rank* (see below). SAB can start with either one universal region (as in the example), or any number of predefined regions.

2. What type of region is it?

   (a) *Marked* for being known to always satisfy or fail the given criteria. An example is when a region's PVR range is 0.0 to 1.2 and the criterion is $PVR \leq 1.62$. Here, split the region into two, and using the given density bound, estimate and bound the greatest lower probability bound of being in each of the two sub-regions. Mark them for the same reason as the original region.

   (b) *Unsure.* The region can still either pass or fail the given criteria. An example is when a region's PVR range is 0.0 to 2.0 and the criterion is $PVR \leq 1.62$.

      i. If the possibilities of the region (PVR's range in the PVR example) have not been bounded yet, bound them (in the PVR example, use the given formulas for an upper and lower bound on PVR). *If* the region should be marked, do so and bound the probability of being in it.

      ii. If the possibilities have been bounded, split the region in two. Bound both sub-regions' possibilities, and estimate the greatest lower probability bound of being in each sub-region. *If* a sub-region should be marked, do so and bound that sub-region's probability.

The probability estimations made are just used to suggest the next best step for SAB by helping to rank the sub-regions. They are *not* used as part of any probability bound.

The only overlap allowed between regions is shared borders. No overlap is permitted if the probability density bound has impulse(s).[7]

6.2 Ranking Regions & Estimating Region Probabilities

A region's *rank* estimates how much splitting it will increase the known lower bound on the probability of either satisfying or failing the criteria. An "unsure" (unmarked) region's rank is the estimated greatest lower probability bound (using the given density bound) of being in that region. Estimate as follows:

1. Observe how many input and parameter sample points (out of a thousand picked using a "density" which resembles the given joint density bound) fall within the region. If > 10 samples (1%) fall inside, the fraction falling inside is the estimate.

[7] An impulse occurs when part of the bound becomes infinitely high and leads to a non-zero probability of the variables taking on a particular set of values. An example of such a set for the variables PAP and LAP is $(PAP = 45) \wedge (LAP = 30)$.

2. If ≤ 10 samples fall inside, estimate with a formula that quickly, but approximately integrates the density bound in the region. The PVR example uses formula C_n : 3-3 in [11, page 230].

These two parts compensate for each other's weaknesses:

1. The first part is bad for low probabilities because any region α will have large gaps between the sample points within it. So many sub-regions of α will have no sample points even though they may have high values for the lower probability density bound.

2. The second part is bad for high probabilities because the regions involved are either large or probably contain a complicatedly shaped part of the density bound.[8] The integration formulas only work well when a region's section of the density bound is easily approximated by a simple polynomial.

A marked region's rank is the gap between the estimated greatest lower probability bound of being in the region and the known lower bound on that probability. This works better than the gap between the upper and lower bounds on the greatest lower probability bound because SAB often finds very loose upper bounds, while the estimates are usually accurate.

6.3 Bounding Region Probabilities

The basic way SAB finds a lower bound on the probability of being in a region is to multiply the region's volume[9] by its minimum probability density lower bound value (found by the bounding mechanism described below). I derived the PVR example's first density bound expression (a Gaussian density) by taking the density parameters (Table 1) and substituting them into the general form for a Gaussian density. After some simplification, I got (numbers rounded-off):

$$0.01033 \exp(-0.01323(13.70P^2 - 26.09PL - 10.36PC + 16.42L^2 + 10.77LC + 28.08C^2))$$

where $P = (PAP - 23.94)$, $L = (LAP - 15.29)$, and $C = (CO - 6.487)$.

To help tighten this bound, SAB tries to use any monotonicity and/or convexity present in the region's part of the density bound in the following manner (derivations in Appendix A):

Let $f(x_1, \ldots, x_n)$ be the probability density and within a region α let x_i range between l_i and h_i. The probability of being in α is

$$F = \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} f(x_1, \ldots, x_n) dx_1 \ldots dx_n.$$

If $\partial f / \partial x_1$ is always > 0 in α, then

$$F \geq [\prod_{i=1}^{n} (h_i - l_i)][(\min_{*} f(l_1, x_2, \ldots, x_n)) + (\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n))(\frac{h_1 - l_1}{2})],$$

[8] Most of the common probability densities only have complicated shapes where the density values are high. I am assuming that this complication will be reflected in the corresponding part of the bound.

[9] For a region α, let its variables x_i $(i = 1 \ldots n)$ range between l_i and h_i. Then α's volume is $\prod_{i=1}^{n} (h_i - l_i)$. SAB only deals with n-dimensional rectangular regions.

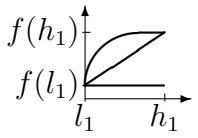

Figure 3: 1-D Convex Density and Lower Bound

where the minimization of f is over the x_2 through x_n values within α ($\min_*$ means that x_1 is NOT part of the minimization) and the minimization of $\partial f/\partial x_1$ is over the x_1 through x_n values within α. This bound is tighter than the basic lower bound:

$$[\prod_{i=1}^{n}(h_i - l_i)][\min f(x_1, \ldots, x_n)].$$

Similar expressions can be derived for the other variables and for when $\partial f/\partial x_i < 0$.

If $\partial^2 f/\partial x_1^2$ is always ≤ 0 in α (convex down), then

$$F \geq [\prod_{i=1}^{n}(h_i - l_i)][(\min_* f(l_1, x_2, \ldots, x_n)) + (\min_* f(h_1, x_2, \ldots, x_n))]/2,$$

where the minimizations of f are over the x_2 through x_n values within α. This bound is also tighter than the basic one. See Figure 3 for the one dimensional case: the $\cap$ curve is the density, the area under the diagonal line is F's new lower bound, and the area under the horizontal line is the original bound. Similar expressions can be derived for the other variables.

Several methods exist to integrate a region's probability density bound, including Monte Carlo [7] and quadrature (numeric integration) methods [11]. These cannot truly bound the integration error because they only take numeric samples at particular points.

6.4 Splitting Regions

SAB may split a selected region α in either step 2a or step 2(b)ii. In either, SAB picks a variable in α to split along and then bisects α. Select the variable as follows: in step 2(b)ii, find the one with the largest difference between its upper and lower bound within the region, normalized by its standard deviation. In step 2a, find the one with the largest apparent variation in the density's slope with respect to it.

6.5 Finding Numeric Bounds

Many of SAB's parts need to bound expressions. For algebraic expressions (the type in the models to be used), perfect bounding algorithms have not been built. The type of algorithm used here will find bounds that indicate the truly unachievable,[10] but may not be the tightest possible. Example: saying that $x < 7$, when in fact, $x < 3$. I have implemented an augmented version of bounds propagation [3]. It does the following interval arithmetic [12]:

- Bound an operation's result using bounds on the operation's operands. For example: $ub(a + b) \leq ub(a) + ub(b)$.

[10] In practice, the accuracy of this may be limited by round-off error.

- Bound an operand using bounds on an operation's result and the operation's other operands. For example: $ub(a) \leq ub(a+b) - lb(b)$.

The "bounder" examines expressions and updates bounds with these operations. It *iterates* over the expressions until every one that might produce a change has been examined at least once and all the recent bound changes are below a certain threshold.

7 Discussion

This paper introduces a method called SAB which uses a lower bound on the input joint density to analytically bound the likelihood of some possible behavior.

SAB entails much computation. When possible, first estimate the probability with some approximation method like moments [4] and then use SAB to insure that the probability is within certain bounds. Also, once regions are made in response to one set of criteria, they can be reused when examining other sets. This will cut down much of the computation for the remaining sets.

Future work on SAB itself includes testing how large a problem it can handle and expanding it to more quickly bound a variable's mean, variance, median, 90% confidence interval, etc. I will also explore splitting a region at the selected variable's median value (or some approximation) within the region. This can handle infinite intervals (bisection cannot), which permits an initial region where each variable is within the all-inclusive range of $[-\infty, \infty]$.

Another question about SAB is how important is its inability to handle upper density bounds. Some preliminary answers to this question and the question of SAB's speed can be found in [13], which is based on work done after the original version of this paper was written.

On matters other than algorithms, work needs to be done on finding the types of density bounds that are the most common, easiest to specify, and most useful. Candidates for easy-to-specify bounds are common densities with bounded parameters. An example is a Gaussian density with a mean between 0 and 1. One can generate such bounds by using information from parameter confidence intervals.

Despite uncertainty in input density shapes and parameter values, bounds on input densities have not really been utilized to bound the chances of events. This work describes a bounding method which has the features of being able to handle events beyond the ones in a pre-enumerated list, producing analytic probability bounds, and giving better answers as more iterations or samples are allowed.

A Some Region Probability Bounds Derivations

This appendix shows derivations for some of the expressions that bound the probability of being in a region α. Let $f(x_1, \ldots, x_n)$ be the probability density, and within α let x_i range between l_i and h_i. Then the probability of being in α is

$$F = \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} f(x_1, \ldots, x_n) dx_1 \ldots dx_n.$$

A.1 Basic Bound

This subsection derives the following lower bound on F:

$$[\prod_{i=1}^{n}(h_i - l_i)][\min f(x_1, \ldots, x_n)],$$

which is the 'volume' of the region multiplied by the lowest density value within it. The minimization of f is over the x_1 through x_n values within α.

$$\begin{aligned} F &\geq \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} [\min f(x_1, \ldots, x_n)] dx_1 \ldots dx_n \\ & [\min f(x_1, \ldots, x_n)] \int_{l_n}^{h_n} dx_n \cdots \int_{l_1}^{h_1} dx_1 \\ & [\min f(x_1, \ldots, x_n)] \prod_{i=1}^{n}(h_i - l_i) \end{aligned}$$

A.2 Bound Using Monotonicity

This subsection shows that if $\partial f/\partial x_1$ is always > 0 in α, then

$$F \geq [\prod_{i=1}^{n}(h_i - l_i)][(\min_{*} f(l_1, x_2, \ldots, x_n)) + (\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n))(\frac{h_1 - l_1}{2})],$$

where the minimization of f is over the x_2 through x_n values within α ($\min_*$ means that x_1 is NOT part of the minimization), and the minimization of $\partial f/\partial x_1$ is over the x_1 through x_n values within α.

$$\begin{aligned} F &= \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} f(x_1, \ldots, x_n) dx_1 \ldots dx_n \\ &= \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} [f(l_1, x_2, \ldots, x_n) + \int_{l_1}^{x_1} \frac{df}{dx_1}(x_1, \ldots, x_n) dx_1] dx_1 \ldots dx_n \\ &= \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} [f(l_1, x_2, \ldots, x_n) + \int_{l_1}^{x_1} [\sum \frac{\partial f}{\partial x_i}(x_1, \ldots, x_n) \frac{dx_i}{dx_1}] dx_1] dx_1 \ldots dx_n \end{aligned}$$

Since the x_i's are integrated independently of one another, dx_i/dx_1 is 0 for $i \neq 1$ and 1 for $i = 1$. So, the sum collapses down to the $\partial f/\partial x_1$ term:

$$\begin{aligned} F &= \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} [f(l_1, x_2, \ldots, x_n) + \int_{l_1}^{x_1} \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n) dx_1] dx_1 \ldots dx_n \\ &\geq \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} [(\min_{*} f(l_1, x_2, \ldots, x_n)) + \int_{l_1}^{x_1} (\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n)) dx_1] dx_1 \ldots dx_n \\ &\geq [\int_{l_1}^{h_1} [(\min_{*} f(l_1, x_2, \ldots, x_n)) + (\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n)) \int_{l_1}^{x_1} dx_1] dx_1] \times \\ & \quad [\int_{l_n}^{h_n} dx_n \cdots \int_{l_2}^{h_2} dx_2] \\ &\geq [\int_{l_1}^{h_1} [(\min_{*} f(l_1, x_2, \ldots, x_n)) + (\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n))(x_1 - l_1)] dx_1] \prod_{i=2}^{n}(h_i - l_i) \end{aligned}$$

$$
\begin{aligned}
&\geq [[(\min_{*} f(l_1, x_2, \ldots, x_n)) - l_1(\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n))] \int_{l_1}^{h_1} dx_1 \\
&\quad +(\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n)) \int_{l_1}^{h_1} x_1 dx_1)] \prod_{i=2}^{n} (h_i - l_i) \\
&\geq [(\min_{*} f(l_1, x_2, \ldots, x_n))(h_1 - l_1) \\
&\quad +(\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n))[-l_1(h_1 - l_1) + (\frac{h_1^2 - l_1^2}{2})]] \prod_{i=2}^{n} (h_i - l_i) \\
&\geq [(\min_{*} f(l_1, x_2, \ldots, x_n))(h_1 - l_1) + (\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n)) \frac{(h_1 - l_1)^2}{2}] \prod_{i=2}^{n} (h_i - l_i) \\
&\geq [(\min_{*} f(l_1, x_2, \ldots, x_n)) + (\min \frac{\partial f}{\partial x_1}(x_1, \ldots, x_n)) \frac{(h_1 - l_1)}{2}] \prod_{i=1}^{n} (h_i - l_i)
\end{aligned}
$$

A.3 Bound Using Convexity

This subsection shows that if $\partial^2 f / \partial x_1^2$ is always ≤ 0 in α (convex down), then

$$
F \geq [\prod_{i=1}^{n} (h_i - l_i)][(\min_{*} f(l_1, x_2, \ldots, x_n)) + (\min_{*} f(h_1, x_2, \ldots, x_n))]/2,
$$

where the minimization of f is over the x_2 through x_n values within α ($\min_*$ means that x_1 is NOT part of the minimization). Within α, f is convex down with respect to x_1, so $f(x_1, \ldots, x_n)$ is $\geq$ than the linear combination of

$$
q(x_1) f(l_1, x_2, \ldots, x_n) + (1 - q(x_1)) f(h_1, x_2, \ldots, x_n),
$$

where $q(x_1) = (x_1 - l_1)/(h_1 - l_1)$. So,

$$
\begin{aligned}
F &\geq \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} [q(x_1) f(l_1, x_2, \ldots, x_n) + (1 - q(x_1)) f(h_1, x_2, \ldots, x_n)] dx_1 \ldots dx_n \\
&\geq \int_{l_n}^{h_n} \cdots \int_{l_1}^{h_1} [q(x_1)(\min_{*} f(l_1, x_2, \ldots, x_n)) \\
&\quad +(1 - q(x_1))(\min_{*} f(h_1, x_2, \ldots, x_n))] dx_1 \ldots dx_n \\
&\geq [(\min_{*} f(l_1, x_2, \ldots, x_n)) \int_{l_1}^{h_1} q(x_1) dx_1 \\
&\quad +(\min_{*} f(h_1, x_2, \ldots, x_n)) \int_{l_1}^{h_1} (1 - q(x_1)) dx_1] \int_{l_n}^{h_n} dx_n \cdots \int_{l_2}^{h_2} dx_2 \\
&\geq [(\min_{*} f(l_1, x_2, \ldots, x_n))[\frac{x_1^2/2 - l_1 x_1}{h_1 - l_1}]_{l_1}^{h_1} \\
&\quad +(\min_{*} f(h_1, x_2, \ldots, x_n))[x_1 - \frac{x_1^2/2 - l_1 x_1}{h_1 - l_1}]_{l_1}^{h_1}] \prod_{i=2}^{n} (h_i - l_i) \\
&\geq [(\min_{*} f(l_1, x_2, \ldots, x_n))(\frac{h_1 - l_1}{2}) + (\min_{*} f(h_1, x_2, \ldots, x_n))(\frac{h_1 - l_1}{2})] \prod_{i=2}^{n} (h_i - l_i) \\
&\geq [(\min_{*} f(l_1, x_2, \ldots, x_n)) + (\min_{*} f(h_1, x_2, \ldots, x_n))][\prod_{i=1}^{n} (h_i - l_i)]/2
\end{aligned}
$$

Acknowledgments

The idea for SAB arose after a thought-provoking discussion with Mike Wellman, who also, along with Peter Szolovits, William Long, Alvin Drake and Peter Kempthorne have reviewed parts of this paper.

References

[1] Daniel G. Bobrow, editor. *Qualitative Reasoning about Physical Systems.* MIT Press, 1985. Reprinted from *Artificial Intelligence*, vol. 24, 1984.

[2] Brian C. Williams. MINIMA: A symbolic approach to qualitative algebraic reasoning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 264–269. American Association for Artificial Intelligence, August 1988.

[3] Elisha P. Sacks. Hierarchical reasoning about inequalities. In *Proceedings of the National Conference on Artificial Intelligence*, pages 649–654. American Association for Artificial Intelligence, 1987.

[4] G. Hahn and S. Shapiro. *Statistical Models in Engineering.* John Wiley & Sons, Inc., 1967.

[5] Eric J. Horvitz, David E. Heckerman, and Curtis P. Langlotz. A framework for comparing alternative formalisms for plausible reasoning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 210–214. American Association for Artificial Intelligence, 1986.

[6] Peter Szolovits and Stephen G. Pauker. Categorical and probabilistic reasoning in medical diagnosis. *Artificial Intelligence*, 11:115–144, 1978.

[7] J. M. Hammersley and D. C. Handscomb. *Monte Carlo Methods.* Barnes and Noble, Inc., 1965.

[8] Mark Johnson. *Multivariate Statistical Simulation.* John Wiley and Sons, New York, 1987.

[9] Dean T. Mason. *Congestive Heart Failure, Mechanisms, Evaluation, and Treatment.* Dun-Donnelley, New York, 1976.

[10] William C. Shoemaker. Physiology, monitoring and therapy of critically ill general surgical patients. In William C. Shoemaker and Edward Abraham, editors, *Diagnostic Methods in Critical Care*, chapter 4, pages 47–86. Marcel Dekker, Inc., New York, 1987.

[11] A. H. Stroud. *Approximate Calculation of Multiple Integrals.* Prentice-Hall, Inc., 1971.

[12] H. Ratschek and J. Rokne. *Computer Methods for the Range of Functions.* Halsted Press: a division of John Wiley and Sons, New York, 1984.

[13] Alexander Yeh. Finding the likely behaviors of static continuous nonlinear systems. *Annals of Mathematics and Artificial Intelligence.*, 1990. Revised version of the preliminary paper in the 2nd Intl. Workshop on AI and Statistics (1989). Accepted for a corresponding special issue of the journal.