On the Impact of Distribution Information on the Effectiveness of Statistical Decision Criteria

Barbara Engelhardt, Kim Reece, Steve Chien
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
{firstname.lastname}@jpl.nasa.gov

Abstract

A statistical decision criterion is used to determine when a sufficient number of independent samples have been made of a random variable in order to statistically guarantee the proximity (within error ε) of *the current estimated expected value* to the *actual expected value*. When information is available regarding the distribution being estimated (e.g., parametric, normal, or bounds on the random variable) these assumptions enable use of specialized statistical decision criteria that may be more efficient. We examine four decision criteria that make varying assumptions about characteristics of the random variable. We evaluate these criteria on real data from three selection problems from NASA applications, quantifying the value of such information on the cost and accuracy of each decision criteria. We find that the algorithm presuming normality (e.g., Nádas) performs better than two-sided bound (e.g. Bernstein), and the two-sided bounds perform better than a one-sided bound (e.g. Chernoff) algorithm.

1 Introduction

Decision criteria enable a choice among a set of possible hypotheses (e.g., actions, parameters) when the consequence of the choice depends on the interactions of the hypothesis in some stochastic environment. We can determine the value of each choice from sample interactions at a cost to our system. Decision criteria can be applied to many problems, including learning statistical decision models [Maron & Moore, 199], optimizing function parameters [Dubrawski & Schneider, 1997], design optimization [Fukunaga et al, 1997], algorithm selection [Chien et al 1999], or optimizing planner control strategies [Engelhardt & Chien, 2000, Gratch & DeJong 1994]. The selection of an appropriate decision criterion to use for a particular data set depends heavily on known features of the data. This paper examines four decision criteria: a probability approximately correct (PAC) criteria [Valiant, 1984], Chernoff Bounds [Hagerup & Rub, 1990], Hoeffding's inequality [Hoeffding, 1963], and Bernstein's inequality [Bernstein, 1946]. Each of these criteria uses information about the data in order to converge faster than the costly brute-force alternatives.

Stopping rule	Assumption
Nádas (PAC)	Random variable distributed normally
Chernoff	Random variable has upper bound B
Hoeffding	Random variable is bounded by $[a, b]$
Bernstein	Maximum difference between random variable mean and single instance bounded by M

Table 1: List of the assumption made for each decision criteria

Parametric decision criteria assume that the random variable distribution is based on a distribution function given some parameter, and the goal is to estimate the parameter. Our parametric PAC algorithm, based on a stopping rule introduced in [Nádas, 1969], assumes that the random variable distribution is normal, and requires fewer samples. In a specific domain, if the distribution of the data cannot be estimated by a common distribution function, non-parametric (distribution-free) bounds (e.g., due to Chernoff, Hoeffding, and Bernstein) may be substituted for the normal-theory probabilities above. The complexity of these bounds are described by their rate of convergence: the probability that the estimated expected utility is not within error ε of the actual expected utility goes to 0 exponentially fast as the number of samples m increases [Hoeffding, 1963]. Because of the special nature of the utility estimates (i.e., sample means), these bounds typically give accurate

results due to the concentration of measure phenomenon [Talagrand, 1991]. Chernoff bounds require the upper bound of the data ("one-sided" bound) be specified. Hoeffding's inequality makes the assumption that the data fall within a given [*a, b*], and will converge at least as fast as Chernoff bounds because of the additional information ("two-sided" bound). Bernstein's inequality makes the assumption that the maximum variance of a single sample utility from the expected utility is bounded by a value *M* ("two-sided" bound), which further improves convergence properties of the criterion. Convergence does not guarantee accuracy when the random variable does not satisfy the imposed assumptions, and this discrepancy is illustrated in the accuracy measure.

2 The Statistical Decision Problem

Statistical decision criteria can be used to select a particular hypothesis from a set of hypotheses, based on some ordering scheme such as maximum expected value, when the actual ranking of the hypotheses is not known. The choice of hypothesis will depend on an estimate of the true ranking, and this estimate can be obtained through sampling the hypothesis in the real world. The outcome of sampling the hypotheses in the real world will depend on the current state of the world. Moreover, sampling will have a cost, such as a monetary or temporal cost.

More specifically, a statistical decision criterion is given a set of hypotheses H from which it must select a single "best" hypothesis. There is also a set of "experiments," E, which take a hypothesis and apply it to the state of the world Θ, which is unknown and unpredictable, and returns an outcome z in the space Z of possible outcomes. A utility evaluator assigns a utility to the outcome z based on the experiment e, the hypothesis h, and the state of the world θ [Raiffa & Schlaifer, 1961]. In this paper, we will use the notation $U(h_i)$ to mean the actual expected utility of hypothesis h_i, and $\hat{U}(h_i)$ to mean the estimated expected utility distribution of hypothesis h_i, based on the expected utility of the set of samples taken.

In order to determine how the sample outcome should impact the ranking, the utility of that sample is factored into the current estimate of the utility distribution. An assessment is made as to the likelihood of the state θ and the outcome z to scale the utility of each sample appropriately. When a sufficient number samples have been made in the current world to statistically guarantee the error for estimated expected utility relative to alternative hypotheses, the statistical decision criteria can estimate the actual ranking of the set of hypotheses.

Decision criteria do not guarantee absolute accuracy because a complete guarantee on accuracy is impossible without an infinite number of samples. Instead, they guarantee the upper limit of error in the ranking based on estimated expected utility will be some given error ε, which is the upper bound on the difference between the actual expected utility of a single non-selected hypothesis and the actual expected utility of the selected hypothesis.

3 Ranking Algorithm

We use a limited number of hypotheses, and sample sequentially using the decision criterion as a stopping rule in order to increase efficiency over fixed-sample techniques [Gratch *et al.*, 1994]. In the actual evaluation, a small number of initial samples (7 per hypothesis) are taken to generate a starting expected utility and estimated variance and to select the top hypothesis thus far (h_{sel}). The algorithm checks to see whether all of the pairwise comparisons between h_{sel} and the other hypotheses satisfy the particular decision criteria. If the comparisons satisfy the criterion, no more samples are taken. If the comparisons fail to satisfy the criterion, additional samples are taken until the criterion is satisfied for all of the comparisons. After each sample, the algorithm reevaluates the choice of the selected hypothesis. The choice of which pair of hypotheses to sample is based on rational selection for the Nádas criterion, described below, and estimates on the number of samples required to satisfy the criteria for the non-parametric criteria, described in the algorithm sections.

There are many possible decision criteria that can be plugged in to the evaluation function. Some simplify the problem by assuming that the stochastic data is normally distributed, while others do not. Sampled utility is averaged over m samples to calculate $\hat{U}(h)$. Due to the central limit theorem, the difference $U(h) - \hat{U}(h)$ can often be approximated by a normal distribution. If desired, the validity of the normal assumption can be verified using tests such as the quantile-quantile (Q-Q) [MathSoft, 1998].

3.1 Blocking

We use a statistical technique called "blocking", in which we analyze the distribution formed by the difference of two random variables directly in order to show the mean of one is greater than the mean of the second [Buringer, Martin, Schriever, 1980]. This is more efficient in cases where the distributions are positively correlated (as in our domains where they are variant attribute vectors or algorithms).

The blocked distribution is the pairwise comparison of two different random variable distributions, *A-B*, which enables the decision criteria to rank based on *dominance*, or decide that the distributions are indistinguishable (*indifference*). Dominance is the area of the distribution above 0, or the probability that random variable *A* will be greater than random variable *B* for a random sample. Indifference occurs when the mean of the blocked distribution falls within [-ε, ε], and the two hypotheses are so close that the criterion cannot distinguish between the two. Using blocking, we enable the decision criteria to guarantee that one random variable is expected to perform better than another, or whether the difference falls with the error allocated to the decision.

3.2 Rational Selection

We can use discrepancies based on the decision criteria to determine how many training examples to allocate to each comparison, given the error bound on the probability of a mistake, an estimate of the difference in expected utility, and an estimate of the variance of each hypothesis:

The decision requirement is only that the sum of the selection errors for the pairwise comparisons remains less than the given bound. If one pairwise comparison requires many more samples to achieve the same amount of accuracy as another pairwise comparison, then if the first comparison is allowed to have more error and the second is allowed less, the overall cost of achieving those local requirements might be reduced. In practice, this method significantly reduces the number of samples necessary to achieve the requirement for certain domains. In general, we cannot solve the allocation problem optimally since estimates for parameters required to compute optimal solutions will include sampling error. For more information regarding these techniques, see [Gratch *et al.*, 1994, Chien et al. 1995].

Only the Nádas decision criterion uses this estimate to determine how to rationally allocate the samples such that the overall selection error is bounded appropriately and the criterion is satisfied. This is because the concept of rationally is tied to distributing the required confidence over the n pairwise comparisons, and the non-parametric criteria are calculated individually for each comparison and do not include the concept of confidence as an input parameter.

3.3 Expected Utility

Expected utility is the average utility of hypothesis h_i in domain θ over an infinite number of runs. In reinforcement learning, it is interpreted as expected reward [Kaelbling, 1990]. As described earlier, utility can be determined by sampling hypothesis h_i on domain θ and calculating the utility of the outcome z, based on utility function U. Thus, the calculation of the expected utility generally takes into account the probability of the current state of the world.

$$U(h_i) = \sum_{j=0}^{N} (\Pr(h_i, \theta) \bullet u(z_{i,j})) \qquad (1)$$

In our simulations, all h_i in H and θ in Θ are equally likely, as we are searching a space of uniformly distributed random numbers, so $\Pr(H)$ is uniform across all h_i in H.

4 Statistical Decision Criteria

For each of the four decision criteria, we will describe the assumptions made by the criterion, the equation to be satisfied, and the formula for calculating the number of samples required for satisfaction. For the non-parametric criteria, the formula for approximating the number of samples required to satisfy the criterion is used to determine the difficulty of the comparison. The largest estimated value of m for a pairwise comparison indicates that the next samples should be for that comparison.

The bounds on the random variables required by the non-parametric decision criteria are often unknown before the sampling begins. For the three non-parametric criteria that require bounds, we created general functions to generate estimates of those bounds from the initial samples. We ran a version of the algorithm that recalculated the bounds after each subsequent sample, and did not find significant improvement from the preliminary analyses.

4.1 Probably Approximately Correct (Nádas)

Much previous work in the machine learning community has focused on Probably Approximately Correct (PAC) learning [Valiant, 1984]. With the PAC decision criterion, an algorithm makes decisions with a given confidence (probability 1-δ, for small δ) to select a good hypothesis (within error ε of the best hypothesis). Because any specific decision either satisfies or does not satisfy this requirement, the PAC criterion holds that over a large number of decisions that the accuracy rate must meet 1-δ.

We use the stopping rule proposed in [Nádas, 1969], which includes the normality assumption to calculate confidence based on the number of samples m. We are requiring the satisfaction of the decision criterion over a set of pairwise comparisons, so the selection error (δ) is distributed to each comparison using the rational allocation methods described above. The normality assumption and rational allocation, along with the sequential sampling algorithm, which sacrifices additional samples for average-case accuracy[1], make the Nádas PAC model extremely efficient in practice.

The requirement that a selected hypothesis be better than the remaining hypotheses is a conjunctive statement, and the conjunction of the difference in expected utility (we use summation for conjunction) is the total selection error.

$$\Pr\left[\bigwedge_{i=1}^{n} (\hat{U}(h_i) - \hat{U}(h_{sel}) > \varepsilon)\right] < \delta \qquad (2)$$

Each decision that the algorithm makes is a selection of a single hypothesis h_{sel} from a set of hypotheses $h_1 \ldots h_n$. This is represented by Equation (2) below, in which the probability that the selected hypothesis h_{sel} is better than all the rest within error ε must be at least 1-δ.

For fixed number of hypotheses n, a sufficient number of samples m may always be taken to satisfy this condition or show indifference.

Using the assumption of normality for the blocked distribution (where $\Phi(k)$ is the normal density function) and blocking, it is straightforward to compute the confidence (and correspondingly the selection error) in a pairwise

[1] In the Nádas stopping rule, the application of the triangle inequality, which allows us to substitute the difference between two estimated expected values for the difference between actual expected values and estimated expected values, theoretically is not valid, because of the distribution of error over all of the pairwise comparisons. In practice this estimation does not hurt the accuracy of the criterion. For more information, see [Turmon, 1995]

comparison as a function of the number of samples m used to produce the estimates, the accuracy parameter ε, the estimated utilities, and an estimated variance σ^2.

$$\alpha_i = \Phi\left(-\sqrt{m}\left[\hat{U}(h_{sel}) - \hat{U}(h_i)\right] \Big/ \sqrt{\sigma^2_{sel,i}}\right) \quad (3)$$

From this equation we can solve for the expected number of samples required to satisfy the Nádas criterion for each pairwise comparison using the normal cumulative distribution function.

$$m_{sel,i} = \frac{\sigma^2_{sel,i}}{(\hat{U}(h_{sel}) - \hat{U}(h_i))^2}\left[\Phi^{-1}(\alpha_i)\right]^2 \quad (4)$$

PAC learning requires the number of samples required to satisfy this criterion to be theoretically bounded by a polynomial function [Valiant, 1984], although in practice the Nádas criterion appears linearly bounded.

4.2 Chernoff Bounds

Chernoff bounds are an effective non-parametric criterion when the random variable is bounded above by constant B:

$$\Pr\left(U(h) - \hat{U}(h) > \varepsilon\right) \le \exp\left(\frac{-2m\varepsilon^2}{B^2}\right) \quad (5)$$

Solving equation 5 for m, we can determine the number of samples required to satisfy the inequality. The estimate for the number of samples required is calculated separately for each blocked distribution. We can use this estimate of m to compare the difficulties in satisfying the criterion across the pairwise comparisons.

$$m \le \frac{-B^2 \log\left(\Pr\left[U(h) - \hat{U}(h) > \varepsilon\right]\right)}{2\varepsilon^2} \quad (6)$$

The value B is estimated for each pairwise comparison from the initial samples (in our case, 7 samples per hypothesis) by finding the largest sample utility thus far, and multiplying this difference by a scalar value. Our analysis has shown that a sufficient generic scalar is 3, based on the blocked distribution functions for these (and other) domains.

4.3 Hoeffding's Inequality

In the case where the distribution of the random variable cannot be estimated by a normal distribution, but the upper and lower bounds of the data are known, Hoeffding's inequality can be applied. Hoeffding's inequality uses both the upper and lower bounds to evaluate the inequality:

$$\Pr\left(\hat{U}(h_i) - \hat{U}(h_{sel}) > \varepsilon\right) \le 2\exp\left(\frac{-2m^2\varepsilon^2}{(b-a)^2}\right) \quad (7)$$

The number of samples to convergence is estimated for each pairwise comparison as the inequality is applied to each blocked distribution separately. The estimates are based on the following equation (from equation 7):

$$m \le \frac{-(b-a) * \sqrt{\log\left(\frac{\Pr\left[\hat{U}(h_i) - \hat{U}(h_{sel}) > \varepsilon\right]}{2}\right)}}{e\sqrt{2}} \quad (8)$$

Hoeffding's inequality, as in Chernoff bounds, estimates a, b for each pairwise comparison from the initial samples (in our case, 7 samples per hypothesis) by finding the minimum and maximum sample utilities, and multiplying this by a scalar value (as in Chernoff, 3).

4.4 Bernstein's Inequality

Bernstein's inequality bounds the maximum difference between the expected utility and a single utility instance (M). In the case where a variance estimate is known, Bernstein's inequality may give somewhat better results, but at a minimum it will converge more quickly. Bernstein's inequality will tend to perform better when the variances of the random variables are small.

Given $|u_i(h_{sel,j}) - \hat{U}(h_{sel,j})| \le M$ for all i, j with probability 1, then

$$\Pr\left(\hat{U}(h_{sel}) - \hat{U}(h_i) > \varepsilon\right) \le \exp\left(\frac{-\varepsilon^2 m}{2\sigma^2 + 2M\varepsilon/3}\right) \quad (9)$$

The total number of satisfying samples for each blocked distribution is based on the following equation (solving for m in equation 9):

$$m \le \frac{-\left(2\sigma^2 + \frac{2M\varepsilon}{3}\right)\log\left(\Pr\left[\hat{U}(h_{sel}) - \hat{U}(h_i) > \varepsilon\right]\right)}{2\varepsilon^2} \quad (10)$$

The value M is estimated for each pairwise comparison from the initial samples by calculating the largest difference between the utilities of two hypotheses on a single domain problem, and multiplying this by a scalar value. To estimate the M from initial samples, the following generic equation was used for all three domains.

$$M = \max\left|\hat{U}(h_{i,sel}) - u_i(h_{i,sel})\right| * 3 \quad (11)$$

We can thus use such exponential inequalities to get accurate bounds on the probability of incorrect selection, retaining a probabilistic guarantee, while avoiding the normality assumption.

5 Empirical Results

The four decision criteria were run on three different domains described in Table 1. Unfortunately, at submission time, the Hoeffding inequality runs had not produced enough samples to be statistically significant (we hope to have these results shortly). In all cases, we used the quantile-quantile test described above to determine their degree of normality.

For each of the three models, the Nádas stopping rule is satisfied in the fewest number of samples, and scales linearly for all three domains (see Table 1, Figure 2). The other two criteria scale exponentially for all three domains, with Bernstein's criterion the lowest of the non-parametric criteria. Because of the non-normality of the penetrator domain, the scalar estimate of M was an over-estimate, and the resulting number of samples for Bernstein's criterion is increased because of this over-estimate.

Name	Description	Shape
Aeroshell	Design parameters for a lander aeroshell. Utility: appropriate entry velocity with minimum weight. Each Hypothesis: Floating-point dimensions for aeroshell forebody overlap, halfcone angle, bluntness ration, fillet radius, outer diameter, and tail geometry.	Normal
Image	Performance of lossless image compression methods on Galileo image data. Utility: relative size of compressed file. Hypotheses: Seven different compression programs.	Close to Normal
Penetrator	Determine dimensions for a penetrator (a small, robust probe). Utility: stability under surface impact at high velocity in extreme cold. Each Hypotheses: Floating-point dimensions for penetrator outer diameter, total length.	Highly Non-normal

Table 2: Descriptions of the three domains and distributions.

The confidence rate for the Nádas criterion was held at 95%. The error rates for all four criteria were determined using the ratio σ/ε, which was set at 2, 3, 4 for each domain, and as expected, the number of samples required for convergence increases as this ratio increases. Note that for the Chernoff and Bernstein inequalities, if the distribution bounds are accurate, as the number of trials approaches infinity, the accuracy should approach 1.0 (100%).

Aeroshell Domain	Samples for 3 Hypotheses		Samples for 5 Hypotheses		Samples for 10 Hypotheses	
Criterion	**Mean**	**Var**	**Mean**	**Var**	**Mean**	**Var**
Nádas	28.16	10.68	35.58	2.71	70.47	2.30
Bernstein	58.38	24.82	104.05	24.96	3864.00	467.01
Chernoff	302.37	35.62	833.47	56.92	4344.96	670.49

Table 3: Illustrative Aeroshell samples for three decision criteria, showing the linear scaling of the Nádas criterion and the exponential scaling of the other two criteria.

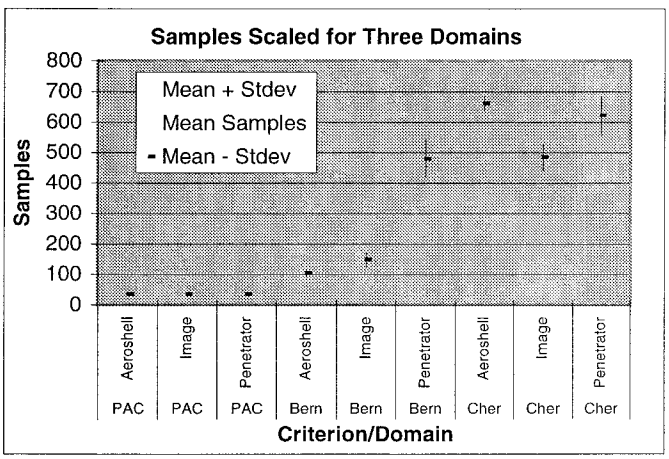

Figure 1: Scaling of each criterion across the three domains.

These empirical results confirm that the Nádas algorithm significantly outperforms the Bernstein, which in turn outperforms the Chernoff algorithm. Because in our experience a two-sided bound is not significantly harder to estimate than a one-sided bound, and Bernstein significantly outperforms Chernoff, we view Bernstein as more useful when normality cannot be assumed (such as in the penetrator domain).

Although the cost of convergence depends on the run time properties of the algorithm itself and less on the domain, the accuracy rate depends heavily whether the assumptions of the decision criteria are satisfied in the blocked utility distributions (Table 3). Accuracy is measured by running a deep sample of each hypothesis (2000 samples), and then comparing the rankings of the decision criteria with the actual rankings, including ε error.

Algorithm	σ/ε	Aeroshell	Image	Penetrator
Nádas	2	0.93	1.0	0.64
Nádas	3	1.0	1.0	0.73
Nádas	4	0.99	1.0	0.8
Bernstein	2	1.0	1.0	0.99
Bernstein	3	1.0	1.0	1.0
Bernstein	4	1.0	1.0	0.99
Chernoff	2	0.48	0.98	1.0
Chernoff	3	0.87	1.0	1.0
Chernoff	4	0.99	0.99	1.0

Table 4: Accuracy measurements for the four criteria, with different error ratios, for the three domains, on 3 hypotheses (equivalent accuracy as the number of hypotheses grew).

6 Related Work

Much of this work descends from the COMPOSER algorithm, developed by Gratch and DeJong, which uses the Nádas stopping rule, but COMPOSER uses a greedy allocation scheme [Gratch *et al.*, 1994]. The Probably Approximately Locally Optimal (PALO) algorithm uses a Chernoff-based stopping criterion and stops searching when it has identified a near-local maximum with a high probability, using hill-climbing [Greiner, 1992].

Maron and Moore propose using Hoeffding races, which quickly prune inferior models using a PAC-learning version of Hoeffding's inequality and focus computational efforts on selecting from the best models using brute-force methods [Maron & Moore, 1993]. The improvements are minimal using Hoeffding races unless there is a clear subset of winners within the complete set of models.

Bayesian learning models are closely related to the PAC models presented here. Bayesian models require the specification of prior knowledge related to the causality of random variables [Pearl, 1988]. Instead of relying on prior knowledge, the algorithm presented here uses a frequentist approach by continually adapting the statistical model to include the most recently sampled utility. This approach avoids using non-informative prior knowledge and in many cases is equivalent to bayesian models [Gratch *et al.*, 1994].

7 Conclusion and Further Work

Currently, we are running simulations of the four decision criteria on different domains with unique characteristics, including sets with bimodal and Poisson distributions, to

determine the accuracy of the decision criterion for alternative distributions. Prior probabilities can be included in the decision criteria in order to incorporate prior knowledge about the hypotheses into the current decision. One algorithmic extension is to include prior knowledge in the model and compare versions of the criteria with and without encoded prior knowledge (based on the accuracy of the knowledge). Another extension is to adjust expected utility to include variance in the scoring.

This paper describes four decision criteria, evaluated on three actual NASA selection domains. The results presented in the paper indicate that the Nádas algorithm assuming normality significantly outperforms the Bernstein algorithm, with its two-sided bound, which in turn outperforms the Chernoff algorithm, with a one-sided bound. In domains where two-sided bounds can be found almost as easily as the one-sided bound, we found Bernstein's criterion of more practical use. We also found that the accuracy rates of all the algorithms are significantly affected by how well the random variable satisfies the distribution knowledge encoded in the decision criteria.

8 Acknowledgements

This paper describes work performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract from the National Aeronautics and Space Administration.

References

[Bernstein, 1946] S. N. Bernstein (1946). *The Theory of Probabilities*, Gastehizdat Publishing House.

[Buringer, Martin, Schriever, 1980] H. Buringer, H. Martin, K. Schriever. *Nonparametric Sequential Selection Procedures, Boston*, MA: Birkhauser, 1980.

[Chien et al, 1995] S. Chien, J. Gratch, M. Burl 1995. "On the Efficient Allocation of Resources for Hypothesis Evaluation: A Statistical Approach." In *IEEE Trans. Pattern Analysis & Machine Int.* 17(7), p. 652-665.

[Chien et al 1999] S. Chien, A. Stechert, D. Mutz 1999. "Efficient Heuristic Hypothesis Ranking." *Journal of Artificial Intelligence Research* Vol 10: 375-397.

[Dubrawski & Schneider, 1997] A. Dubrawski, J. Schneider, (1997). "Memory Based Stochastic Optimiation for Validation and Tuning of Function Approximators" *6th Intl Workshop on Artificial Intelligence and Statistics*, Fort Lauderdale, Florida.

[Engelhardt & Chien, 2000] B. Engelhardt, S. Chien,, " An Empirical Evaluation of Local Search Methods for Adapting Planning Policies in a Stochastic Environment," *Local Search in Planning & Scheduling*, Springer-Verlag, *to appear.*

[Ferguson, 1967] T. S. Ferguson, 1967. *Mathematical statistics: A decision theoretic approach.* Academic.

[Fukunaga et al, 1997] Fukunaga, A. Stechert, A., Chien, S. 1997. "Towards a Self-Configuring Optimization System for Spacecraft Design," in *Proceedings of the 1997 International Symposium on Artificial Intelligence, Robotics, and Automation for Space*, Tokyo, Japan, July 1997.

[Gratch & Chien, 1996] J. M. Gratch and S. A. Chien, "Adaptive Problem-solving for Large Scale Scheduling Problems: A Case Study," *Journal of Artificial Intelligence Research* Vol. 4 (1996), pp. 365-396.

[Gratch *et al.*, 1994] J. Gratch. S. Chien, and G. DeJong (1994) "A Decision-theoretic Approach to Adaptive Problem Solving," *Proc. AAAI-94*, , pp. 36-45.

[Greiner, 1992] R. Greiner and W.W. Cohen, "Probabilistic Hill Climbing," *Proceedings of Computational Learning Theory and Natural Learning Systems,* 1992.

[Hagerup & Rub, 1990] T. Hagerup, C. Rub, (1990). A Guided Tour of Chernov Bounds, *Information Processing Letters* 33: 305-308.

[Haussler, 1992] D. Haussler (1992). "Decision Theoretic Generalization of the PAC Model for Neural Net and Other Learning Applications", *Information and Conputation*, Vol. 100, pp. 78-150.

[Hoeffding, 1963] W. Hoeffding (1963). Probability inequalities for sums of bounded random variables, *Journal American Statistical Association* 58(301):13-30

[Kaelbling, 1990] L. P. Kaelbling. *Learning in Embedded Systems*. Ph.D. Thesis;. TR-90-04, Stanford University, Department of Computer Science, June 1990.

[Maron & Moore, 1993] O. Maron, A. Moore, (1993) "Hoeffding Races: Accelerating Model Selection Search for Classification and Function Approximation," *Advances in Neural Information Processing Systems 6*, Morgan Kaufmann, 1993.

[Nádas, 1969] A. Nádas (1969), "An extension of a theorem of Chow and Robbins on sequential confidence intervals for the mean," The Annals of Mathematical Statistics 40:2, pp. 667-671.

[Pearl, 1988] J. Pearl (1988), *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufmann.

[Raiffa & Schlaifer, 1961] H Raiffa, R Schlaifer. *Applied Statistical Decision Theory*, Division of Research, Harvard Business School.

S-Plus 5 for UNIX Guide to Statistics, Data Analysis Products Division, Mathsoft, Seattle; 1998.

[Talagrand, 1991] "A new isoperimetric inequality and the concentration of measure phenomenon," Lecture Notes in Mathematics, 1469: 94–124, Springer, 1991.

[Turmon, 1995] M. Turmon (1995). "Assessing Generalization of Feedforward Neural Networks" Ph.D Thesis,. Cornell University.

[Valiant, 1984] L. Valiant, "A Theory of the Learnable," *Communications of the ACM*, 1984.