Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, Oct. 7-10, 2007, Montréal.

Decision-Making Under Severe Uncertainty for Autonomous Mobile Robots

Daniel Berleant and Gary T. Anderson, *Senior Member, IEEE*

***Abstract*— The field of robotics is on a growth curve, with most of the growth expected in the areas of personal and service robots. As robots become more prevalent in chaotic home and industrial settings, they will be required to make increasingly independent decisions about how to accomplish their tasks. A key to accomplishing this is the development of techniques to allow robots to handle severe uncertainty. This paper introduces the use of Information Gap Theory as a way to enable robots to make robust decisions in the face of uncertainty, and illustrates this with an example problem.**

I. Introduction

Decision-making under severe environmental uncertainties is an important challenge for mobile robots. This paper introduces a new approach for robots to make robust decisions in the face of severe uncertainty. By "severe" we refer to uncertainty that contains model (often called epistemic) uncertainty in addition to the ordinary stochastic variety. The goal is to expand the range of practical applications for robot systems. To date, most successful applications of robotics either use human teleoperation (e.g. most surgical robots) or highly structured environments (e.g. industrial robots). However, many of the tasks that robots will be required to perform in the near future will require them to make decisions in relatively unstructured, changing settings.

Robots have two types of difficulties in making good decisions: dealing with lack of sensor information and choosing appropriate actions based on the information they do have. Even if the sensors available to a robot are perfect (i.e., error free), they do not always provide all of the information a robot needs to make a decision. For example, a camera can provide visual information about the type of terrain a sloped path has, but cannot directly measure how slippery it is. Faced with sparse and only partially relevant data, robots must still decide on what, if any, action to take. Considerable research on this problem seeks to enable robots to choose the action most likely to be correct based on the data available. Our work takes a complementary approach. We propose allowing a robot to decide if the information available supports a decision problem it needs to solve. If it does not, the robot attempts to obtain the data it needs to make an appropriate decision. Among the advantages of this method are that it enables robots to make easy or moderately difficult decisions quickly, and it lessens the need for complex sensor systems for robots.

The approach we propose is to integrate new advances from the rapidly developing fields of Information-Gap Decision Theory [1] and imprecise probabilities for representing severe uncertainty and supporting the need for robots to make decisions in highly uncertain circumstances. *Information-Gap Theory* is an approach to decision-making under epistemic uncertainty – uncertainty due to lack of complete knowledge about modeled relationships. Epistemic uncertainty involving techniques for manipulation of flexible representations of uncertainty. Probabilistic approaches to epistemic uncertainty have advantages over other methods in being consistent with widely accepted axiomatizations [2]. This is useful because it permits evidence for effectiveness based on theoretical rather than only empirical grounds. The resulting *imprecise probabilities* permit successful modeling of problems that may include probability distributions of uncertain form, intervals, or both [3]. An important need, and the goal of this research, is to enable robotic agents to determine when they lack sufficient information to make decisions and what information they would need. Once this objective is manageable, mobile robots will be able to effectively operate in chaotic or confusing environments that they would not normally be able to operate in. As an example, relatively simple robots will be able to traverse paths over terrain they could not otherwise handle, because they will be able to determine which portions are traversable, which are not, which might be and, for those, what information is needed to decide whether they are or not.

Many of the environments where robots would be useful are chaotic, changing, and require a high tolerance for ambiguity. This work proposes a theoretically sound method for autonomous mobile robots to handle such severe uncertainties. As such, it holds the promise of extending the range of robot applications in the home and workplace. Potential applications for such robots include 1) agricultural tasks; 2) cargo transportation in rough terrain; 3) navigation in 3-D environments; 4) navigation in dangerous environments; 5) autonomous construction of bridges or other structures; 6) autonomous operation of robots in chaotic work and home environments. We focus here on navigation as an important example of the more general technique.

G. T. Anderson is with the Department of Applied Science, University of Arkansas at Little Rock, Little Rock, AR 72204 USA (e-mail: gtanderson@ualr.edu).

D. Berleant is with the Department of Information Science, University of Arkansas at Little Rock, Little Rock, AR 72204 USA (e-mail: jdberleant@ualr.edu).

II. Related Work

While better characterization of the terrain will often lead to better decisions about whether a robot can travel over a given path, these measurements will generally involve uncertainty. Our project addresses this problem by determining whether a decision problem can be solved *given the available information.* As such, it could be applied in conjunction with many other techniques to make robot decision-making more robust.

A. Navigation

There has been a great deal of recent work in trying to assess the traversability of terrain, especially for space and military applications. Typically, the approach is to estimate selected terrain characteristics (e.g. slop, roughness, and hardness) from sensor information. Once this is done, pattern classification techniques such as neural networks are used to classify the terrain into one of several predefined types, such as sand, gravel, asphalt, etc. (e.g. [16] [17]). Stereo vision systems have been proposed as the main sensor type for determining many terrain parameters [18]. However, this requires distinct features to be present in an image. Moreover, this method suffers from sensitivities to lighting conditions and distance of the camera to the feature being measured. Additionally, cameras can detect only the topmost layer of terrain, while subsurface features may be important in determining traversability. An example of this would be grass covering a muddy field.

Because vision systems cannot extract all needed environmental information, researchers have looked at measuring terrain parameters at the immediate site of the robot. These include methods to measure specific properties that can be placed into terramechanics equations [17], and methods to map sensor readings to soil characteristics through pattern matching routines [16]. To this end, sensors have been used to measure parameters such as wheel sinkage, soil cohesion, soil friction, vertical load, torque, motor current, motor voltage, angular wheel speed, linear wheel speed, robot roll/pitch/tilt, robot vibrations and other features [19] [17] [16]. Sensors such as cameras, accelerometers, gyro inertial measurement units, ladar, microphones, ultrasound sensors and more have been used to determine these parameters. Once the traversability of different regions is found, a cost function or rule base is often used to guess the best path [18].

B. Decision-Making

Agents that have autonomy in navigation or other tasks must, in general, make decisions. This is why research on decision-making is critically important to robotics. This importance has motivated considerable progress to date. Yet despite the headway that has been made, the problem of robot decision-making is far from solved. Thus the significance, both intellectual and practical, of our objective in advancing this area is evident. We have identified an approach that incorporates two formal inferencing techniques whose integration promises significant advancement in the area.

1. An algebra of severely uncertain random variables. This has roots going back decades [4] and has continued to be pursued over the intervening years (e.g. [3]) with a more recent surge in interest [5] [6] [7] [23].
2. Information-Gap Decision Theory, a way to use the results of computations yielding highly uncertain outputs to make rational decisions. This technique forms a platform for a two-level decision strategy. In level one, severely uncertain inputs are tested to see if they support solving a desired decision problem. If not, meta-decisions are enabled about what and how much new information to acquire to solve the original decision problem.

Uncertainty is inherent in the operation of autonomous robot systems. There is uncertainty in sensing the environment and, given the estimated state of the environment, uncertainty as to what the best action might be. One way to deal with this problem is to improve knowledge of an environment by improving sensing capability. For example, [8] examines careful calibration of sensors to improve the accuracy of information. Another approach often used is to combine information from several sensors to reduce errors in individual measurements, as was done by [9] for robot localization. While these techniques can be valuable, it is impossible to achieve perfect knowledge of an environment, no matter how much effort is placed on sensing capabilities.

Another approach to handling uncertainty is to try to determine the "best" answer, given the information available. For example [10] describes using a fuzzy logic-based controller to enable a rover to navigate over challenging terrains. The final outcome is a recommendation with no uncertainty associated with it. However, *the conclusion that is most likely to be correct given the information available might in fact not be reliable enough*. For example, if the best decision has a probability in the range 40% - 90% of being correct, but robot mission specifications require at least an 85% chance of success, then the robot should gather more information before making a final decision so that in the worst case the 85% minimum is met.

Other ways to deal with uncertainty in robotics, especially in planning, use Bayesian inference [11], Markov Decision Processes [12], and the related Partially Observable Markov Decision Processes or POMDPs [13]. Because practical problems require a large number of possible states, Markovian techniques are mainly viable for small problems unless suitably modified [14]. Moreover, the availability of probability values for transitions between states is assumed, although these might actually be unavailable. One reason for unavailability is that the true probabilities can change with environmental conditions in realistically complex environments. This is a problem for Bayesian approaches as well. Ordinary Bayesian inference relying on Baye's Theorem relies on inputs that are probabilities or probability distributions. The problem with such representations of uncertainty is that real problems are often characterized by severe uncertainty best expressed as interval ranges for probabilities or probability distributions with error bounds. Similarly, in Bayesian networks, the inputs should permit error bounds, rather than being single numbers or single distributions. Credal networks [15] do permit such flexible, realistic representations, but have not yet been seriously investigated for robotics. Our approach shares with credal networks the ability to relax the often-unrealistic assumption of single, known probability distributions, while at the same time allowing computation of functions of these imprecise probabilities.

In our approach, uncertainty is represented as bounded families of probability distributions (Figs. 2 through 7 and 9 through 11). Such severe, epistemic (ignorance-based) uncertainties can arise, for example, from measurements containing a known noise distribution which, however, is parameterized by a mean whose deviation from the measured value depends on bounded environmental variations. Another example would be when two random variables with known distributions both contribute to a third, "derived" [20] random variable, but the dependency relationship between the two input r.v.'s is unknown. Then every conceivable dependency relationship predicts its own distribution for the derived r.v. and, properly, one must represent the derived r.v. as a family of distributions rather than as a single one.

Our approach includes an algorithm for numerically deriving families of distributions (e.g. [20]), coupled with a method for using such families to make decisions. This method, Info-Gap Theory [1], seeks decisions that are *satisficing* – that is, that meet minimal acceptable requirements – rather than optimal as in many other decision-making techniques. Models are parameterized with an externally provided value, called α, for the degree of ignorance (i.e. the epistemic uncertainty). In our technique, α represents the amount of space between the envelopes bounding a family of distributions. If an acceptable decision can be made despite the existing ignorance, well and good. If not, our technique mandates acquiring additional information in an effort to reduce ignorance, and hence α, to the point where a satisficing decision can either be made, or ruled out as unachievable.

III. Applying Information-Gap Theory

We have previously developed a technique to implement Information Gap Theory for problems of severe uncertainty, and publicly available software [21] for doing the necessary calculations on imprecise probabilities. We represent uncertain quantities as probability boxes, and perform arithmetic on the probability boxes [22]. Here, we apply this technique to a illustrative problem in robot navigation.

Consider an outdoor setting in which a robot needs to move a cart filled with cargo to a new location. It can take a safe but long path, or it could choose a shorter, but less safe path that involves pulling the cart up a hill. Depending on the environmental conditions, this could be a difficult decision for a human to make. The example shows how to apply Information-Gap Theory to the process of deciding whether or not to attempt to take the shorter path up the hill.

A. An Example

Consider a robot whose task is to move a cart containing cargo from point A to point B. The robot has intrinsic knowledge of its task, but not of the specifics of its environment. In other words, the robot knows how to find the cart, how to connect with it, where points A and B are, how to travel with the cart and how to disconnect from the cart. However, it lacks specific knowledge of the nature of the cargo or the details of its environment. Suppose the most direct path from point A to B contains an incline the robot must pull the cart up. The question the robot must answer is

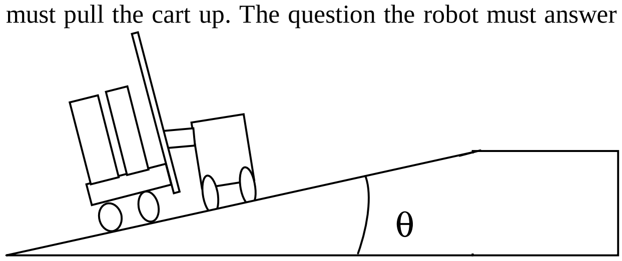

Fig. 1. A robot pulling a cargo-laden cart up an incline.

whether or not it is capable of doing this. If the robot is carrying a cart up an incline, the force of friction on the robot must be greater than the gravitational force pulling the cart down the incline (see Fig. 1). The force of gravity on the cart is $m_{cart}*g*\sin\theta$. The force of friction on the wheels of the robot is $\mu_{friction}*m_{robot}*g*\cos\theta$. In this simplified example, the weight of the cargo-carrying cart is assumed much higher than that of the robot. So, for the robot to successfully pull the cart up the incline, assuming an all-wheel drive robot with each drive matched to the weight on the wheel (so that all wheels slip or none do) we need:

$$\mu_{friction}*m_{robot}*g*\cos\theta > m_{cart}*g*\sin\theta.$$

In other words, it must be that $m_{robot} > (m_{cart}*\tan\theta)/\mu_{friction}$. The unknowns in this are $\mu_{friction}$, m_{cart} and θ. These can be roughly estimated visually by the robot: m_{cart} from the size of the cargo; $\mu_{friction}$ by the color and glossiness of the incline's surface; and θ from stereo vision estimates of its depth at the top and bottom.

The estimates will have large uncertainties associated with them. Sometimes this will not matter. For example, if the slope of the incline is shallow and the cargo small, the robot can be confident that it will be able to traverse the incline. If the uncertainties are such that the robot cannot decide whether or not it can pull the cargo up the incline, then we need to determine what information will reduce this uncertainty sufficiently. If the robot knows the cost of obtaining that information, it can calculate the net gain of obtaining it.

The above example can be applied to specific situations such as a robot moving cargo from several airplane drops to a central location. In that case, once the robot has found the cargo, it must decide how to navigate to the designated location. The most direct route may contain dangerous slopes and gulches, while safer routes may take the robot away from the most direct path. Each route will have difficulties associated with it, and our proposed method will enable the robot to make good decisions about what to do. Similar considerations apply to robots undertaking navigation in dangerous environments; cargo transportation in rough and/or dangerous terrain; navigation in 3-D environments; reconnaissance in dangerous situations; autonomous construction of bridges or other structures; navigation across rivers, streams, swamps, hills and trenches; and so on.

The Robot Problem: Is it Possible to Climb a Given Slope?

We model whether a robot can climb a slope as whether the following equation holds.

$$\mu_{friction}\, m_{robot} > m_{cart} \tan\theta \qquad (1)$$

Let m_{robot} be a known constant, $m_{vehicle}$ be known within a range, such as ±10% (Operand X in Fig. 2), and m_{cargo} be represented by a probability distribution with given mean (Operand Y, Fig. 2). Then $m_{cart} = m_{vehicle} + m_{cargo}$ looks like Result Z in Fig. 2. Result Z is represented by two cumulative probability functions (CDFs): an upper CDF (upper staircase curve in Result Z graph in Fig. 2) and a lower CDF (lower staircase curve in Result Z graph in Fig. 2). The actual value of Result Z will be in the area between these two CDFs. The algorithm for this is described in [3] and several more recent papers. Other algorithms also produce equivalent results on benchmark problems [23].

Next, import m_{cart} (from Result Z in Fig. 2) as Operand X (Fig. 3). Let tanθ be a distribution (Operand Y, Fig. 3). Then the RHS of inequality (1), m_{cart} tanθ, is Result Z in Fig. 3, assuming stochastic independence of m_{cart} and tanθ.

Now let us work out the LHS of inequality (1). Suppose $\mu_{friction}$ is a function of both the mileage on the tires and a "ground condition factor." Model the mileage as a distribution that is uniform over a range, meaning the tires are equally likely to have any of the mileage within that range. Now assume we have three different sources of information regarding the ground condition factor. Let information source A bound the ground condition factor within an interval $\boldsymbol{i}_A$, information source B bound it within an interval $\boldsymbol{i}_B$, and information source C bound it within an interval $\boldsymbol{i}_C$. For example, interval A might represent the ground condition factor of a paved road, interval B of a gravel road and interval C of a dirt road. To further explore the flexibility of the mathematics, suppose further that: (i) intervals $\boldsymbol{i}_A$, $\boldsymbol{i}_B$, and $\boldsymbol{i}_C$ are disjoint (no overlaps), (ii) the actual value might not be within $\boldsymbol{i}_A$, $\boldsymbol{i}_B$, or $\boldsymbol{i}_C$ (a fact modeled as an information source specifying a uniform distribution over a range of conceivable coefficients of friction), and (iii) source A is not as reliable as B or C and should have correspondingly less weight. This situation for the ground condition factor is modeled by the probability box of Fig. 4. (While this may sound like an intractable set of data at first glance, one might readily model it as a Dempster-Shafer structure, which in essence is what is occurring in this case).

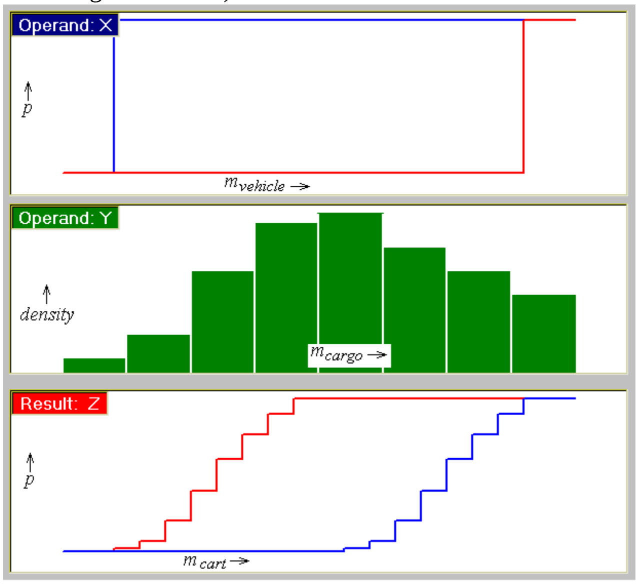

Fig. 2. An interval looks like two cumulative distribution functions (CDFs) forming a probability box as shown for Operand X. The vertical axis is a probability that varies from 0 to 1, while the horizontal axis is the value of the input (Operand X or Y) and the output (Result Z).

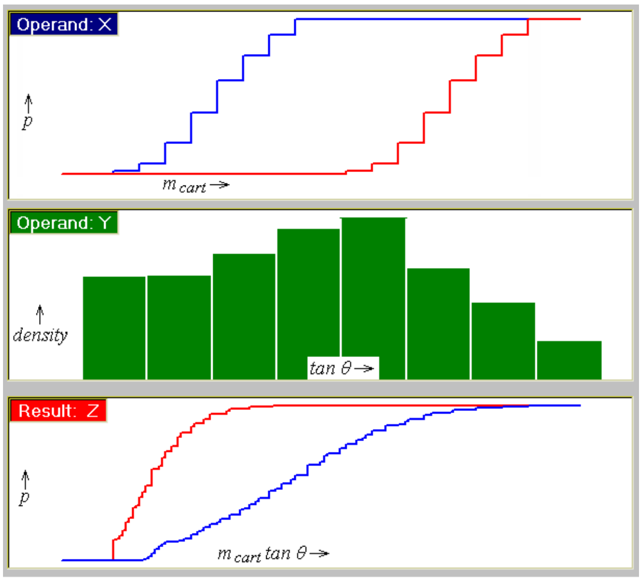

Fig. 3. Operand X is from Z in Fig. 2. Operand Y is histogram discretization of the PDF for tan θ, and Result Z is m_{cart} tan θ, the RHS of inequality (1).

Now that we have modeled mileage m and ground condition factor c, we need to model the binary function $f(m, c)$ to get the coefficient of friction, $\mu_{friction} = f(m, c)$. While the software we have been using in this example can compute a wide range of functions of two parameters, for expository simplicity we will just divide them. Then, if m and c have an unknown dependency relationship, $\mu_{friction}$ looks like Fig. 5. If instead we model mileage and ground condition as independent (perhaps a more realistic assumption), then $\mu_{friction}$ looks like Fig. 6.

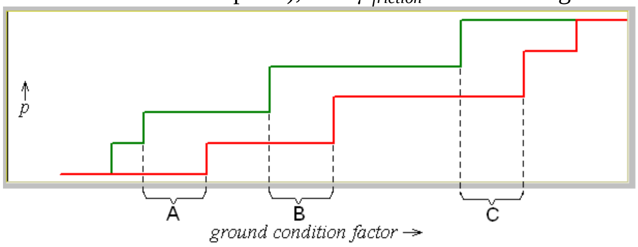

Fig. 4. A probability box for the combination of disjoint intervals comprising evidence sources A, B, and C. (See, e.g., [23]).

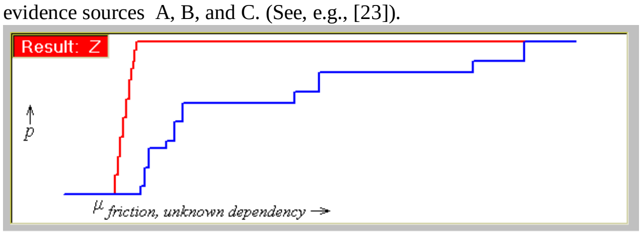

Fig. 5. $\mu_{friction}$ = *f*(*mileage*, *ground condition*), if no assumptions are made about the dependency relationship between *mileage* and *ground condition.*

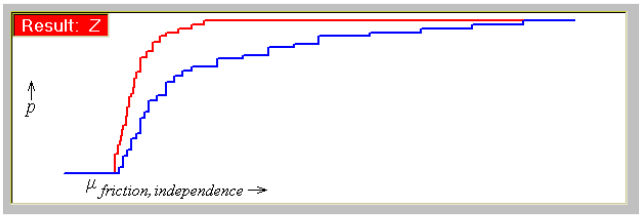

Fig. 6. $\mu_{friction}$ = *f*(*mileage*, *ground condition*), assuming *mileage* and *ground condition* are independent random variables.

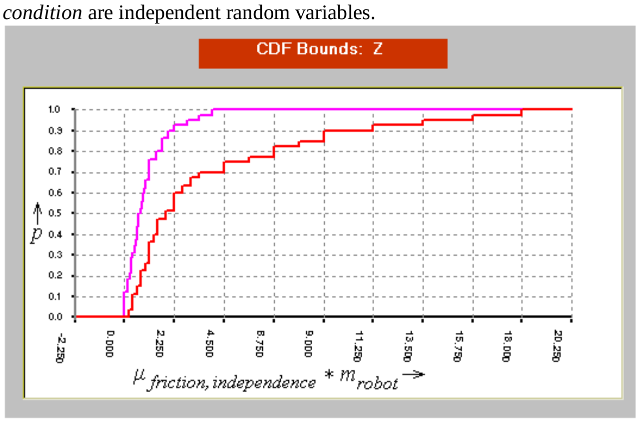

Fig. 7. Probability box bounding the CDF for the LHS of inequality (1).

To get the LHS of inequality (1) we next need to multiply $\mu_{friction}$ (Fig. 6) by m_{robot}. Assuming this to be a known constant, this simply stretches Fig. 6 horizontally. The result is shown in Fig. 7, which shows the *x*-axis quantitatively. The numbers are mathematically consistent across the various figures, though they are not intended to be physically realistic. That is why most of the figures are shown without numeric labelings. Now that we have computed the RHS and LHS of inequality (1) and graphed them (Figs. 3 & 7, respectively), let's compare them, hoping that LHS > RHS. This is the requirement given by (1) if the robot is to be able to traverse the slope. If the probability density functions (PDFs) of the LHS and RHS are such that the right tail of the RHS descends to zero before the left tail of the LHS rises from zero, as illustrated in Fig. 8, all is well. In CDF form, which is more suitable for expressing epistemic uncertainty, we wish for the configuration of Fig. 9. Often, however, we will obtain a situation like Fig. 10 in which the left envelope of the right probability box rises from the horizontal axis before the right envelope of the left p-box reaches its maximum at 1. Then (1) is not assured and it would be risky for the robot to attempt to climb the slope.

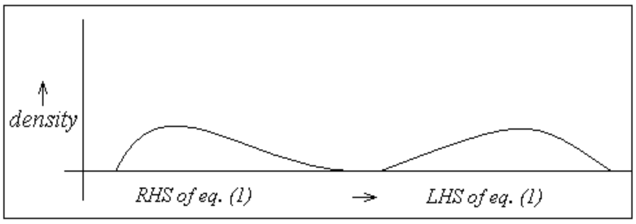

Fig 8. Two non-overlapping PDFs. It is guaranteed that the corresponding random variables satisfy an inequality.

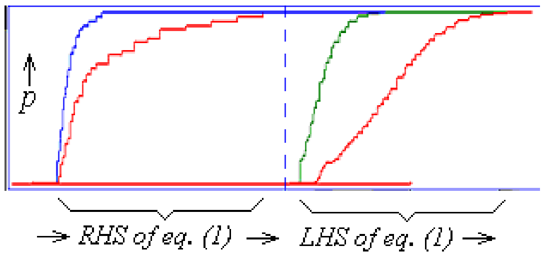

Fig. 9. The random variable whose Probability box is on the right is unambiguously greater than the one whose probability box is on the left.

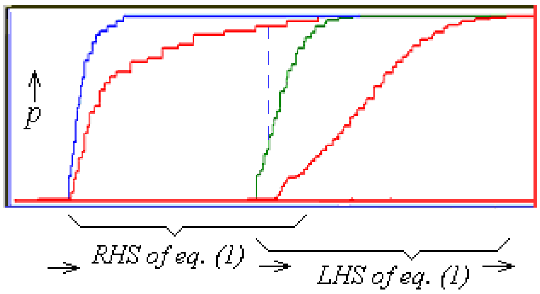

Fig. 10. Samples of the random variable whose probability box is on the right are probably above samples of the one whose probability box is on the left.

We can check the probability of inequality (1). Based on the probability boxes for the RHS and LHS of (1), imported into Operand X and Y of Fig. 11 from Figs. 3 & 7, and making no assumption about the dependency relationship between RHS & LHS, the Statool software produced the probability range [0.19, 1] portrayed in Result Z of Fig. 11 and in more detail in Fig. 12. We conclude that p(LHS > RHS)∈ [0.19, 1] if the dependency is unknown. Such a wide range of probability for success may be insufficient to support a decision either to climb the hill or not climb it. However if we assumed that the LHS and RHS of eq. (1) are independent, we could get a higher quality estimate of p(LHS > RHS). This can readily be done by performing the same sequence of arithmetic operations on probability boxes as was done earlier in the example of this paper, except under the assumption of independence. Whether or not the result of such an analysis sufficed to support the climb/no climb decision would support an important decision. This decision is not about whether to climb or not, but rather, whether or not to investigate if independence holds so that the climb/no climb decision can, finally, be made.

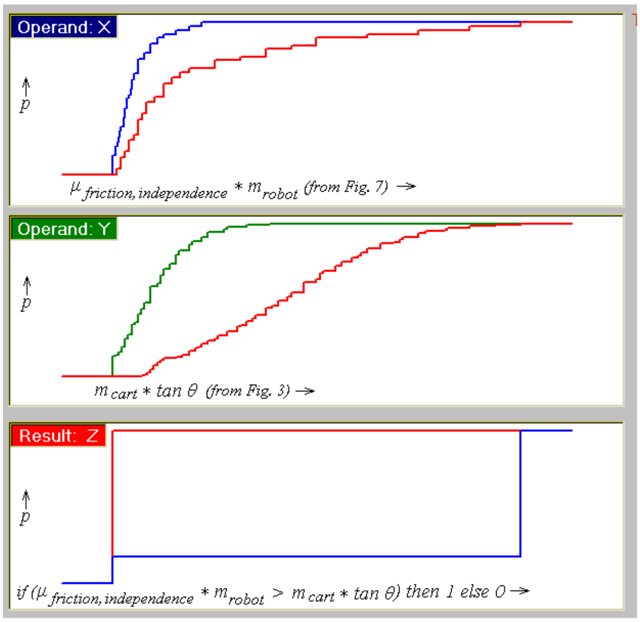

Fig. 11. Result Z shows the range of the probability that inequality (1) hold, based on the probability boxes for its RHS and LHS.

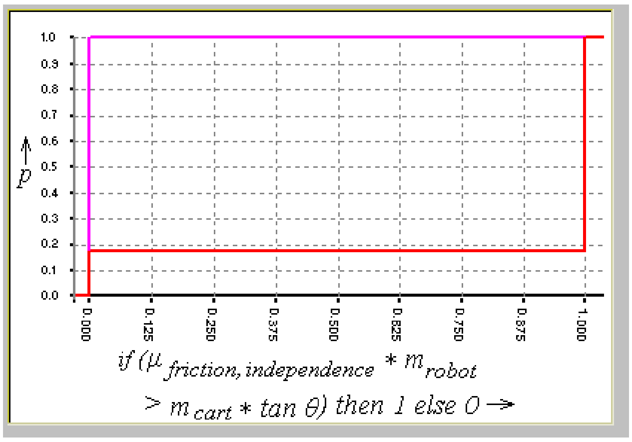

Figure 12. Numerical plot detailing Result Z of Figure 12. It shows that the probability of eq. (1) is in the range [0.19, 1].

IV. Summary

This paper deals with the problem of handling "severe" (epistemic, model) uncertainty for autonomous robots. We show how Information Gap Theory can be combined with the algebra of severely uncertain random variables to solve this problem. An example was given to illustrate how the approach can be applied to autonomous robot navigation.

References

[1] Y. Ben-Haim, *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edition. Academic Press, 2006. 1st edition, 2001.

[2] H. Regan, S. Ferson and D. Berleant, "Equivalence of methods for uncertainty propagation of real-valued random variables," *Internat. Journal of Approximate Reasoning*, vol. 36, pp. 1-30, 2004.

[3] D. Berleant, "Automatically verified reasoning with both intervals and probability density functions," *Interval Computations*, 1993 no. 2, pp. 48-70.

[4] M. Fréchet, "Généralisations du théorème des probabilités totales," *Fundamenta Mathematica*, vol. 25, pp. 379-387, 1935.

[5] Society for Imprecise Probabilities: Theory and Applications. http://www.sipta.org/. Founded in Feb., 2002.

[6] D. Berleant, guest editor, Special issue on intervals and probability, *Reliable Computing*, vol. 9, issue 6, 2003.

[7] D. Berleant and S. Nesterov, guest co-editors, Special issue II on intervals and probability, *Reliable Computing*, vol. 10, issue 2, 2004.

[8] P. Stone, M. Sridharan, D. Stronger, G. Kuhlmann, N. Kohl, P. Fidelman and N. Jong, "From pixels to multi-robot decision-making: a study in uncertainty," Special Issue on Planning Under Uncertainty in Robotics, *Robotics and Autonomous Systems Journal* (RAS), vol. 54, no. 11, pp. 933-943, 2006.

[9] S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 781-795, Oct. 2002.

[10] H. Seraji and A. Howard, "Behavior-based robot navigation on challenging terrain: a fuzzy logic approach," *IEEE Trans. Robotic Autom.* vol. 18, no. 3, pp. 308-321, June 2002.

[11] S. Thrun, W. Burgard and D. Fox, *Probabilistic Robotics*. Cambridge, MA: MIT Press, 2005.

[12] P. Fabiani and F. Teichteil-Konigsbuch, "Symbolic focused dynamic programming for planning under uncertainty," in *IJCAI Workshop on Reasoning with Uncertainty in Robotics*, Edinburgh, Scotland, July 30, 2005, pp. 70-77.

[13] J. Pineau, "Tractable planning under uncertainty: exploiting structure," Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, 2004.

[14] X. Li, W. K. Cheung and J. Liu, "Towards solving large-scale POMDP problems via spatio-temporal belief state clustering," in *IJCAI Workshop on Reasoning with Uncertainty in Robotics*, July 30, 2005.

[15] F. G. Cozman, "Credal networks," *Artificial Intelligence*, vol. 120, no. 2, pp. 199-233, July 2000.

[16] L. Ojeda, J. Borenstein, G. Witus and R. Karlsen, "Terrain characterization and classification with a mobile robot," *Journal of Field Robotics*, vol. 23, no. 2, pp. 103-122, 2006.

[17] K. Iagnemma, S. Kang; H. Shibly and S. Dubowsky, "Online terrain parameter estimation for wheeled mobile robots with application to planetary rovers," *IEEE Trans. on Robotics*, vol. 20, no. 5, pp. 921-927, Oct. 2004.

[18] A. Howard and H. Seraji, "Vision-based terrain characterization and traversability assessment," *J. of Robotic Systems*, vol. 18, no. 10, pp. 577-587, 2001.

[19] P. Schenker, P. Pirjanian, T. Huntsberger, H. Aghazarian, E. Baumgartner, K. Iagnemma, A. Rzepniewski and S. Dubowsky, "Reconfigurable robots for all-terrain exploration," in *Proceedings of the SPIE Symposium on Sensor Fusion and Decentralized Control in Robotic Systems III*, vol. 4196, Boston, MA, September 2000.

[20] M. Springer, *The Algebra of Random Variables*. Wiley and Sons, 1979.

[21] D. Berleant, L. Xie and J. Zhang, "Statool: a tool for Distribution Envelope Determination (DEnv), an interval-based algorithm for arithmetic on random variables," *Reliable Computing*, vol. 9, no. 2, pp. 91-108, 2003. (Tool available at http://ifsc.ualr.edu/jdberleant/statool/index.htm.)

[22] D. Berleant and J. Zhang, "Representation and problem solving with Distribution Envelope Determination (DEnv)," *Reliability Engineering and System Safety*, vol. 85, no. 1-3, pp. 153-168, 2004

[23] J. C. Helton and W. L. Oberkampf, eds., Special Volume on Alternative Representations of Epistemic Uncertainty, *Reliability Engineering and System Safety*, vol. 85, no. 1-3, pp. 1-369, July-Sept. 2004.