Safeguarding a Lunar Rover with Wald's Sequential Probability Ratio Test

P. Michael Furlong, Michael Dille, Uland Wong, Ara Nefian

***Abstract*— The virtual bumper is a safeguarding mechanism for autonomous and remotely operated robots to conservatively avoid collisions even in the face of sensor uncertainty. In this paper we take a new approach to the virtual bumper system by applying a powerful but rarely examined statistical test. By using a modified version of Wald's sequential probability ratio test, we demonstrate that we can reduce the number of false positives reported by the virtual bumper, thereby saving valuable mission time. We further use the concept of sequential probability ratio to control vehicle speed in the presence of possible obstacles in order to increase certainty about whether or not obstacles are present. With this principled obstacle certainty measure, our new algorithm reduces the chances of collision by approximately 98% relative to traditional virtual bumper safeguarding without speed control.**

I. INTRODUCTION

It is not uncommon for robots to employ a "virtual bumper" as a safeguarding mechanism in a wide variety of scenarios. [1]–[7] Fairly basic in concept, a virtual bumper is a procedure that checks a range sensor and, according to some decision rule, determines whether a collision is imminent. If a collision is considered likely, the vehicle will stop. Stopping the vehicle will likely trigger interventions by remote operators to regain situational awareness.

Remote interventions consume execution time, the most precious quantity in short lunar missions that are not designed to survive lunar night. As observed by Shah *et al.*, [8] the mean time between interventions has a greater impact on the time to complete a mission than the time it takes complete an intervention. For that reason we have developed an algorithm for increasing the mean time between interventions when using a virtual bumper.

Constraints on hardware that may be used in space missions eliminate more reliable approaches to obstacle detection. LIDAR, while a preferred tool of many roboticists – authors included – has not yet been deployed on a planetary mission. The GESTALT obstacle avoidance system on MER takes approximately 70 seconds to operate, [9] and while this is acceptable on Mars, it does not enable the high-tempo operations necessary for lunar operations constrained to one lunar day. By necessity faster, and therefore less reliable, obstacle detection must be used, motivating a principled technique to reduce the number of false readings from the virtual bumper.

A virtual bumper can be viewed as a classification of terrain into the classes of "obstacle" and "clear." Any classifier will have false positives and negatives which need to be removed – called denoising. Traditional approaches to denoising a virtual bumper often include straightforward temporal filtering such as averaging the last k reported values. These do not take into account a characterization of the virtual bumper's false positive and false negative rates. Probabilistic sensor models have been used to improve readings from the underlying sensor, [10] and we propose a similar approach to modelling the decision made with such a sensor.

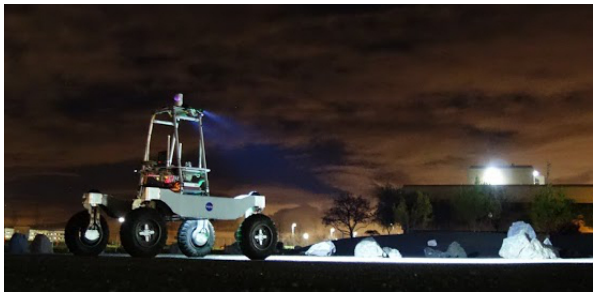

Fig. 1: NASA Ames' autonomous wheeled rover, KRex2.

The readings from a virtual bumper can be viewed as a series of sequential experiments attempting to confirm one of two hypotheses, $\mathcal{H}_0$ – no obstacle is present – or $\mathcal{H}_1$ – an obstacle is present. In this paper we describe a modification to Wald's sequential probability ratio test (SPRT) that takes readings from the virtual bumper along with a probabilistic characterization of the virtual bumper and produces an estimate of hypothesis that is more credible.

Our approach to filtering the virtual bumper performs at least as well as colloquial approaches, specifically averaging the last k reported values, but has the added advantages of being tunable for false positive rates. We have expanded the algorithm to control vehicle speed in response to uncertainty about the presence of an obstacle. We demonstrate a substantial improvement in rover performance over window averaging. By merely slowing down the vehicle instead of requesting a complete halt, we increase the mean time between interventions and thereby increase mission effectiveness.

We test the denoising algorithm first on simulated data drawn from the empirically-derived Shoemaker-Morris lunar rock distribution. [11] Then, using a laser-based structured light system installed on the NASA Ames Krex2 rover shown in Figure 1 and under evaluation as a potential virtual bumper for lunar exploration, we evaluate it on actual rover data collected during live operation on outdoor terrain. Finally, speed adaptation experiments were conducted in simulation

Intelligent Robotics Group, SGT / NASA Ames Research Center, Moffett Field, CA 94035, USA Contact: padraig.m.furlong@nasa.gov

on obstacles drawn from this same realistic distribution.

II. BACKGROUND

Virtual bumpers are a mechanism of detecting nearby navigation hazards faster than processing of other on-board sensors or a remote operator can act. [2] They have been deployed on robotic platforms as a mechanism of safeguarding vehicles in the absence of full situational awareness. [5], [12]

Wald in 1945 [13] produced a method to determine when to stop collecting samples called the sequential probability ratio test. The problem setup was to determine whether one of two hypotheses were true, and the action the decision maker can take is to either collect another sample or to stop collecting samples altogether. With the tth sample that is collected the experimenter computes

$$\Lambda_t = \sum_{i=1}^{t} \log\left(\frac{f_1(x_i)}{f_0(x_i)}\right), \tag{1}$$

where $f_1(x)$ is the probability density function given $\mathcal{H}_1$, and $f_0(x)$ is the probability density function given $\mathcal{H}_0$. If $\Lambda_t \geq b$ then the experimenter accepts $\mathcal{H}_1$, and if $\Lambda_t \leq a$ then the experimenter accepts $\mathcal{H}_0$, where a and b are preselected to control for type I (false positive) and type II (false negative) errors. The experimenter determines an acceptable false positive and false negative rates, α and β, respectively, $a = \log(\beta/(1-\alpha))$ and $b = \log((1-\beta)/\alpha)$. While Λ_t remains bounded by a and b, the experimenter continues to collect samples. The Maximized Sequential Probability Ratio Test (MaxSPRT) [14] is an improvement on Wald's SPRT in that it allows for composite alternative hypotheses. In our setting the hypothesis is simple – obstacle or clear – so the additional benefits of the MaxSPRT are not immediately relevant.

Meanwhile, planetary exploration applications are ideal scenarios for a virtual bumper. Great conservatism in operation is always desirable due to the easily-fatal and immensely costly consequences of striking an obstacle with a vehicle far from human assistance. Indeed, it is quite typical for spacecraft in general to be designed to fall back into a so-called quiescent "safe mode" in the presence of any possible anomaly so that ground-based operators can study reported telemetry before continuing the mission. [15] Planetary rovers have been no exception, with stereo imagery [9] and laser-based [16] virtual bumper systems used on previous missions to safely halt an autonomously executed trajectory and wait for operator direction should a potential obstacle be encountered rather than risk a collision. As mentioned however, avoiding mission downtime during unnecessary remote operator intervention inevitably reduces scarce science data collection time, reduces overall traversal distances, and is particularly incompatible with aggressively scheduled lunar missions as here considered.

Forms of the sequential probability ratio test have even seen limited use in non-terrestrial scenarios. Carpenter and Markley [17] employ a modification of the sequential probability ratio test to choose between actions in conducting orbital manoeuvres. This approach plans uncertain trajectories in known environments, and because our environment is not known ahead of time – lunar orbital imagery resolution is too coarse – it is not applicable. Should a rover have an on-board map and planner, then such an approach should be considered.

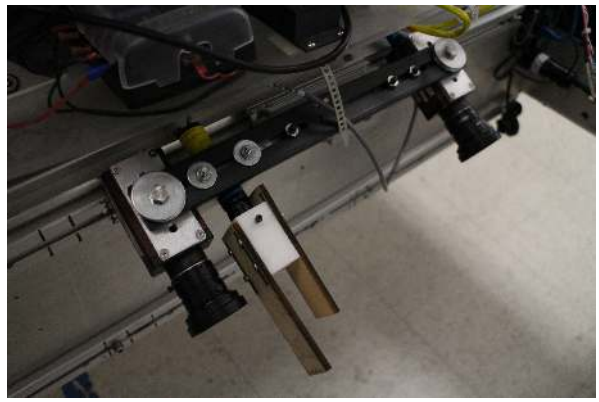

Fig. 2: The virtual bumper apparatus. The laser projector is collinear with the camera. Two cameras are here shown as the virtual bumper system is shared with a stereo-based navigation system. Only one is used in these experiments, as they are performed under varying lighting conditions in which stereo reconstruction is not always possible.

III. METHOD

Test data was generated from KRex2's virtual bumper system, shown in Figure 2, which consists of a solid-state laser emitting highly monochromatic light through a diffraction grating to project a series of dots onto a scene, examples of which may be found in Figure 3. Based on the location of each dot in the image of an adjacent camera, the 3D location of a point on the terrain may be determined by triangulation, similarly to the system employed on Mars Pathfinder. [16] As such, this may be thought of as form of a sparse stereo reconstruction that is able to operate under both sunlit (given sufficient laser power) and dark conditions likely to be found in potential missions to the lunar poles.

Due to the difficulty of distinguishing projected dots from nearby terrain under adverse conditions and the unfortunate geometry resulting from attempting to detect obstacles along a horizontal line in front of the vehicle, reliable sparse 3D reconstruction is difficult when performed in the straightforward manner necessary on an extremely low-speed radiation-tolerant processor. Instead, for the purposes of evaluating operation of the proposed method on a primitive candidate obstacle detector, the overall motion of the projected dots is considered. The geometry of virtual bumper apparatus guarantees that the laser dots will always be in the same image row, and by tracking the lateral motion of the dots, obstacles can be detected. Because the focus of this paper is not on the virtual bumper algorithm itself, but to explicate

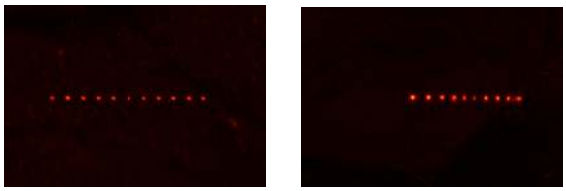

(a) Laser dots on benign terrain.

(b) Laser dots on an obstacle.

Fig. 3: The view from the virtual bumper when the scene is clear (3a) and when an obstacle is laterally displacing the laser dots (3b). The laser dots moved to the right when incident on an obstacle that threatens the vehicle with collision. By detecting deviations from the expected projection, the virtual bumper can detect obstacles.

the denoising and speed control processes, we employ this simple detection algorithm merely as an illustrative example.

Pathologically thin obstacles, ones that are narrow in the coronal plane of the rover, can avoid the virtual bumper by lying between the projected laser dots. However, for our analysis we will ignore those obstacles and are highly unlikely to encounter them in physical terrain. We consider that the visibility of obstacles depends purely on their height and their depth (extent in rover sagittal plane).

From one trial run of the KRex2 robot we collected a number of images that represent "clear" and "obstacle" scenes. We then computed two images from ten arbitrarily chosen clear and obstacle images. From these two sets of ten images we computed two pixel-wise mean images, μ_c and μ_o, respectively. For a candidate image we compute the L^2 error between the candidate image and the two mean images, $\delta_x = ||img - \mu_x||_2 \quad x \in \{c, o\}$. An image is classified as clear if $\delta_c < \delta_o$ otherwise it is considered an obstacle. It is important to stress that this particular algorithm is not necessarily the best one for detecting obstacles with this particular setup, but it is an exemplar for illustrating the role of the Wald test in denoising a sensor. Any algorithm could be used in place of $detector$ in Algorithm 1.

The image label, obstacle or clear, is passed to the Wald test (see Algorithm 2) as part of a window of previous image classifications. The standard Wald test is conducted on all collected samples. We modified the test to prevent spatially distant readings from having an effect on currently sensed terrain. Using a limited window size allows the algorithm to be more responsive to changes in terrain as the vehicle moves.

The Wald test has two regimes, one where $\mathcal{H}_0$ (clear) is definitely accepted and another regime where $\mathcal{H}_1$ (obstacle) is definitively accepted. In between these two regimes the experimenter should continue to collect samples. If $\mathcal{H}_0$ is accepted then the rover should continue collecting samples, i.e. driving. Likewise, when the choice between the two hypothesis is still being resolved, $a < \Lambda_t < b$, the rover will also continue to drive in order to collect more samples. If $\Lambda_t \geq b$, the rover will stop driving.

Algorithm 1

function VIRTUALBUMPER($windowSize, InputImages$)
  $window \leftarrow Queue\,(\emptyset)$
  for all $img \in InputImages$ **do**
    $obstacleDetected \leftarrow 0$
    $detectorValue \leftarrow detector(img)$
    if $detectorValue < obstacleThreshold$ **then**
      $obstacleDetected \leftarrow 1$
    end if
    $window.push_front(obstacleDetected)$
    if $window.size() > windowSize$ **then**
      $window.pop_back()$
    end if
    $requestStop \leftarrow waldTest\,(window, a, b)$
    if $requestStop = 1$ **then**
      return $requestStop$
    end if
  end for
end function

Algorithm 2 The modified Wald test tests the window previously reported by the virtual bumper. Since the reaction to being certain that there is no obstacle present ($\Lambda < a$) is to continue driving, that half of the Wald test is ignored.

function WALDTEST($window, a, b$)
  $\Lambda = \sum_{z \in window} \log\left(\frac{\mathbb{P}(Z=z|X=0)}{\mathbb{P}(Z=z|X=1)}\right)$
  return $\Lambda > b$
end function

A. Modifying Vehicle Speed to Increase Safety

Initially we assume the rover moved at a fixed speed, $0.1m/s$, with the virtual bumper operating with a fixed frame rate of 2 frames per second (fps). Because the current value of the probability test, Λ_t, and the threshold that triggers a stop request, b, are known, it is possible to adjust the speed of the vehicle depending on the difference $b - \Lambda_t$.

We construct a simple recurrence relationship to determine the expected number of frames necessary to wait until the agent has become certain that there is an obstacle present. We choose to slow down when $\Lambda_t > (a+b)/2$, above the halfway

Algorithm 3 The Wald test is used to adjust vehicle speed such that the expected number of frames to be certain of $\mathcal{H}_1$ in the presence of an obstacle is obtained. The function returns a new speed that should make sure that $extraFrames$ many frames are collected before the vehicle reaches the limits of what the virtual bumper can see.

function ADAPTSPEED($\Lambda_t, a, f_1\,(\cdot)\,, f_0\,(\cdot)$)
  $\mathbb{E}[\Lambda_1] = f_1(1)\log\left(\frac{f_1(1)}{f_0(1)}\right) + f_1(0)\log\left(\frac{f_1(0)}{f_0(0)}\right)$
  $extraFrames \leftarrow \lceil (b - \Lambda_t)/\mathbb{E}[\Lambda_1] \rceil$
  $distance \leftarrow speed/fps$
  $time \leftarrow (windowSize + extraFrames)\,/fps$
  $newSpeed \leftarrow distance/time$
  return $newSpeed$
end function

point to confirming $\mathcal{H}_1$. If $\alpha = \beta$, then $(a+b)/2 = 0$.

$$
\begin{aligned}
f_0(z) &= \mathbb{P}\left(Z = z | X = 0\right) \\
f_1(z) &= \mathbb{P}\left(Z = z | X = 1\right) \\
\implies 1 &= f_x(0) + f_x(1) \quad x \in \{0, 1\} \\
\mathbb{E}[\Lambda_1] &= f_1(1) \log\left(\frac{f_1(1)}{f_0(1)}\right) + f_1(0) \log\left(\frac{f_1(0)}{f_0(0)}\right) \\
\mathbb{E}[\Lambda_t] &= f_1(1)\left(\log\left(\frac{f_1(1)}{f_0(1)}\right) + \mathbb{E}[\Lambda_{t-1}]\right) \\
&\quad + f_1(0)\left(\log\left(\frac{f_1(0)}{f_0(0)}\right) + \mathbb{E}[\Lambda_{t-1}]\right) \\
\mathbb{E}[\Lambda_t] &= \mathbb{E}[\Lambda_{t-1}] \\
&\quad + f_1(1) \log\left(\frac{f_1(1)}{f_0(1)}\right) + f_1(0) \log\left(\frac{f_1(0)}{f_0(0)}\right) \\
\mathbb{E}[\Lambda_t] &= \mathbb{E}[\Lambda_{t-1}] + \mathbb{E}[\Lambda_1] \\
\implies \mathbb{E}[\Lambda_t] &= t\mathbb{E}[\Lambda_1] \\
&\quad \text{where} \quad t \in \mathbb{N}^+, t > 1
\end{aligned}
$$

where $\mathbb{E}[\Lambda_1]$ is the expected change in the sequential log probability ratio sum, Λ_t, due to sensing in the presence of an obstacle. In order to become certain that an obstacle is present in the scene the rover needs to collect on average $\lceil (b - \Lambda_t) / \mathbb{E}[\Lambda_1] \rceil$ more frames.

With a specified frame rate and the expected number of frames the appropriate speed of the vehicle is trivial to compute, and given in Algorithm 3. Should the value of $\mathbb{E}[\Lambda_1]$ ever be non-positive then the expected number of frames to yield confidence in the belief in an obstacle is undefined. In this way $\mathbb{E}[\Lambda_1] > 0$ is a criterion for whether or not the sensor can be used to confidently sense the environment.

IV. EXPERIMENTS

We present three experiments. The first is used to characterize the instantaneous classification error of the virtual bumper we defined. The second experiment determines the probability of identifying an obstacle based on the width of the obstacle in collected virtual bumper frames, assuming a fixed frame rate and vehicle speed. The capability of the SPRT algorithm is then demonstrated using actual data collected on the KRex2 rover at NASA Ames' Roverscape planetary analogue testing grounds. The third experiment shows the ability to control the speed of the vehicle in response to uncertainty. In all experiments $\beta = 0.005$.

A. Sensor Characterization

In this experiment we take a collection of N "clear" frames and N "obstacle" frames from one data set of virtual bumper images, where $N = 10$. We extract from the raw images the three rows of pixels that contain the laser dots. The average images μ_o and μ_c are computed from the respective N obstacle and clear frames.

(a) Average clear image, μ_c. The individual laser dots are clearly visible.

(b) Average obstacle image, μ_o. The laser dots are smeared over the scene.

Fig. 4: The average images produced for clear and obstacle frames where the laser dots would be incident on a flat plane. For each, $N = 10$ frames were averaged to produce the image.

We apply the virtual bumper algorithm on the remaining data using the L^2 norm as a similarity metric between μ_o, μ_c, and the current image frame. We compute the true and false positive and true and false negative classification rates on this hold out data.

The mean images that are used for all experiments, μ_c and μ_o, are seen in Figure 4. The statistics of the virtual bumper algorithm on the hold out data are given in Table I. We can see that our illustrative example detector does substantially better than chance at detecting obstacles, however with 10% false negative and 16% false positive rates it is not necessarily suitable for high-reliability lunar operations. At rover velocities of $0.1m/s$ and a frame rate of $2fps$, for every meter of travel in obstacle free terrain there would be approximately three false positive incidents. With this characterization, however, we can apply the SPRT algorithm to denoise the virtual bumper.

		Obstacle present (X)	
		F	T
Obstacle	F	0.90	0.10
reported (Z)	T	0.16	0.84

TABLE I

B. Obstacle Detection Probability

In simulation we determined the probability that the Wald test would identify an obstacle as a function of the number of frames an obstacle would be visible for. That is to say given that the virtual bumper sensor is reporting for each image frame either "clear" or "obstacle" we treat the sensor like a Bernoulli random variable with the characteristics determined by the experiment defined in Section IV-A.

Height and width as determined by a cm-scale histogram determines visibility at a given time, while speed and frame rate determine the duration of visibility. Modelling rocks as hemi-spheres, the number of frames that an obstacle is visible for is determined by the radius of the hemisphere, the speed of the vehicle, the frame rate, and the height threshold for obstacle-ness, in our case $0.1m$. In this experiment the sensor is considered to be flying over the scene, so it can drive over obstacles. This is done to determine the instantaneous classification performance of the competing algorithms.

To supplement the simulation results, we collected real-world data using the NASA Ames' KRex2 platform (Figure 1). The rover was driven around the NASA Ames' Roverscape towards obstacles. The rover would drive up to

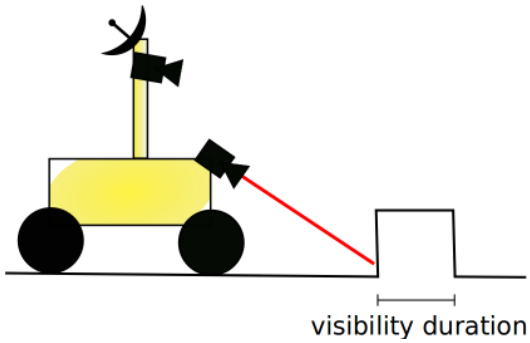

Fig. 5: The line projecting out of the lower camera represents the laser projector projecting dots onto the terrain. The height of the obstacle in this diagram is 1, representing any obstacle or portion of an obstacle with a height $\geq 0.1m$. The number of frames the obstacle is visible for is determined by the speed of the vehicle and the sensor frame rate.

obstacles, then reverse direction and continue driving. We ran Algorithm 1 on all the datasets to determine its success rate, and compared it to the control algorithm of averaging the last k reported virtual bumper classifications. The frames in the scene were manually labelled as either "clear" or "obstacle."

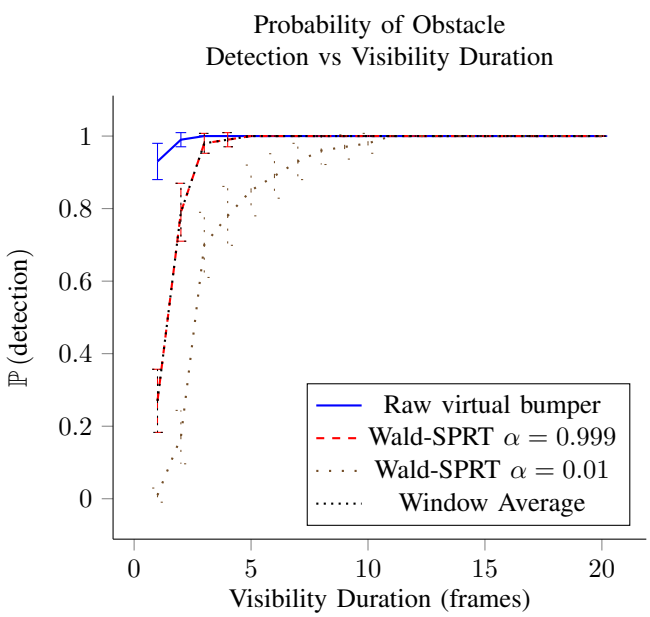

Fig. 6: The probability of detecting an obstacle given the width of the obstacle in virtual bumper frames on 100 randomly generated transects. The probability of detecting an obstacle given the number of frames it appears in the virtual bumper field of view. The Window Averaging algorithm had identical performance compared to the SPRT algorithm where $\alpha = 0.999$. When $\alpha = 0.01$ the SPRT had worse performance on shorter visibility durations, but ultimately converged to the same performance of the other algorithms as bump size increased. In all cases $\beta = 0.005$. The error bars represent a 95% confidence interval.

Figure 6 shows that the performance of the SPRT algorithm is statistically indistinguishable to the performance of the window averaging algorithm. As α decreases, it takes larger obstacles before there is 100% detection rate. While reduced true positives is not ideal, the benefit of the tunable algorithm is the reduction in false positives. Depending on the risk tolerance of the mission, which may vary over time, the SPRT can be modified to give the exact same performance of the window-averaging algorithm or be tuned to substantially reduce the false positive rate. Figure 7 shows that α can be used to control the false positive rate independent of the obstacle visibility duration.

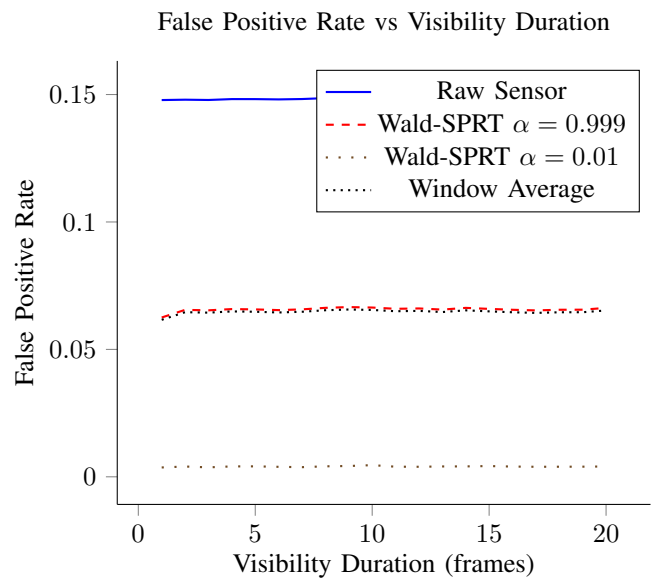

Fig. 7: The instantaneous false positive rate varying with the size of the bump obstacle. The obstacle width is the number of frames that an obstacle appears in the sensors field of view. The size of the obstacle does not have a substantial effect on the false positive rate, however α does have an effect on false positive rate.

The first trial on rover data demonstrates an improvement in the false positive rate of the rover, as demonstrated in Figure 8. We also notice in this case that the window averaging algorithm does not have a substantially lower false positive rate than using the raw data, whereas the SPRT algorithm has reduced the number of false positives by over 50%.

In the second trial, Figure 9, the exemplar virtual bumper algorithm was not well matched to the collected data. However, even when faced with this incredibly noisy data set the SPRT algorithm managed to reduce the number of false positives in the data set. The window averaging algorithm, on the other hand, added false positives.

It is important to highlight that while Tables II and III show the instantaneous true positive rate of the SPRT algorithm is lower than the window averaging algorithm, the SPRT always detects actual obstacles.

We anticipate that with a more robust underlying sensor

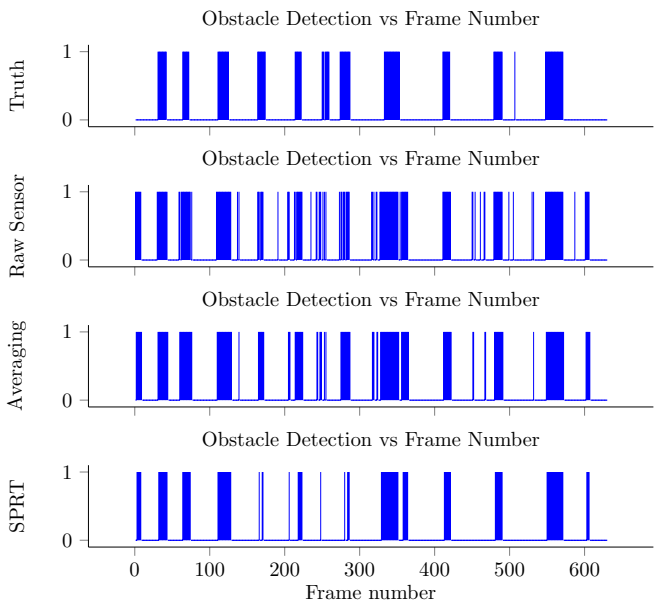

Fig. 8: Trial 1 - The SPRT algorithm reduces the number of false positives, and it does not miss a true obstacle. While individual frames may be falsely registered as clear, the obstacle is still determined to be present. A vertical line represents an image frame that contains an obstacle.

Algorithm	True Positive	False Negative	True Negative	False Positive
Raw	0.870748	0.129252	0.840249	0.159751
SPRT	0.721088	0.278912	**0.935685**	**0.064315**
Window	**0.884354**	**0.115646**	0.854772	0.145228

TABLE II: Trial 1 - The instantaneous classification accuracy of the SPRT and Window averaging algorithms. This data set had 629 frames, 147 of which were obstacles. While there was a cost in terms of instantaneous true positives, there was an over 50% reduction in false positives. Despite the reduction in true positives, the SPRT algorithm did not miss any obstacles.

Algorithm	True Positive	False Negative	True Negative	False Positive
Raw	0.966387	0.033613	0.009756	0.990244
SPRT	0.899160	0.100840	**0.034146**	**0.965854**
Window	**1.000000**	**0.000000**	0.002439	0.997561

TABLE III: Trial 2 - The instantaneous classification accuracy of the SPRT and Window averaging algorithms. The example virtual bumper algorithm failed profoundly on this dataset. Despite the unreliable sensor we still see a reduction in false positives.

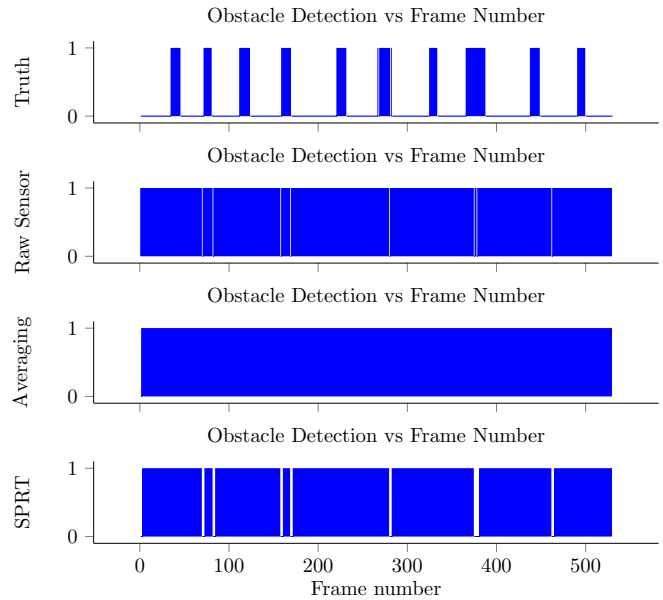

Fig. 9: Trial 2 - The illustrative example of the virtual bumper sensor is not a good fit for this data set. There were 529 frames, 119 of which were obstacles. However even in adversity the SPRT algorithm reduces the number of false positives where the window averaging algorithm does not. A vertical line represents a frame with an obstacle.

the SPRT algorithm will have greater even utility. However the mismatch of the sensor to the environment presents a research avenue adapting the sensor noise model on-line.

C. Modifying Vehicle Speed

To test the ability of the modified SPRT algorithm to slow down the vehicle we first generated terrain from the Shoemaker-Morris rock size distribution. An example of a generated terrain is given in Figure 10. In this experiment we simulated motion of the rover through rocky terrain and if any rock crosses the beam of the virtual bumper then it is reported as an obstacle; if not, it is reported as clear. Before the virtual bumper output is reported to the SPRT module it is corrupted using the sensor characteristics in Table I. Each of the transects are two metres in length.

Control of vehicle speed is given in Algorithm 3. The algorithm always slows the vehicle speed in response to a potential obstacle. The vehicle speed is only increased to the nominal speed once Λ_t crosses below $(a+b)/2$.

There are three possible outcomes covering these transects. The first is that there is a collision between an obstacle $0.1m$ in height or greater. The second is that the rover detects an obstacle and stops before a collision occurs. The third option is that the noise in the sensor causes the rover to stop before an obstacle enters the region between the rover body and the extreme range of the virtual bumper sensor. All transects are guaranteed to have at least one obstacle greater than or equal to $0.1m$ in height. We compare the SPRT algorithm and the window averaging algorithm on the randomly generated transects and compare how many of each category of transect result the algorithms achieve.

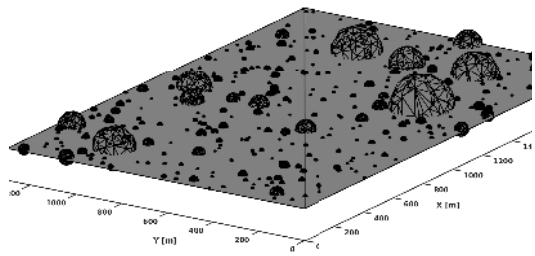

Fig. 10: A sample synthetic terrain generated from the Shoemaker-Morris distribution. [11] The black diagonal line indicates where a 2D linear transect was extracted from the 3D terrain.

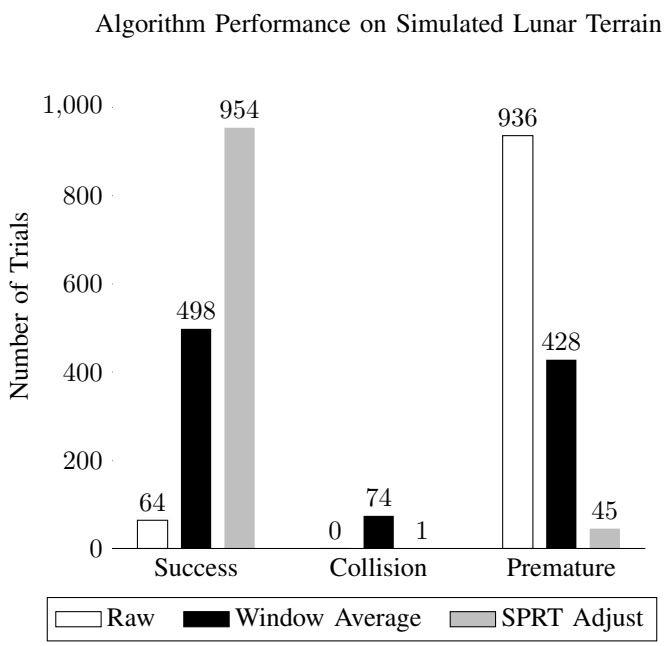

Fig. 11: The SPRT Adjust algorithm reduces the number of collisions from 74 in the case of the window averaging algorithm to 1, a substantial improvement. Similarly it reduces the number of false positive (Premature) incidents considerably.

Using the same α and β parameters as in Section IV-B we demonstrate significant improvement in performance from the basic window averaging to the SPRT Adjust algorithm, shown in Figure 11. There is an 89% reduction in false positives and a 98.6% reduction in collisions. While using the raw data produces no collisions out of 1000 trials, it also renders the rover practically useless due to the overwhelming number of premature halts requested.

What remains to be studied is the possible mission implications of increasing the vehicle speed above nominal $0.1m/s$ when there is confidence the scene is clear. Alternatively, simply stopping the vehicle and accumulating confidence before triggering an intervention may be beneficial to mission performance.

V. CONCLUSIONS

It is well established that more data, at the cost of a time lag, can improve estimation of a noisy random variable. What we have demonstrated in this paper is that the sequential probability ratio test can perform at least as good as the simple window averaging algorithm, in terms of accurate detection of obstacles. More importantly the SPRT algorithm can be tuned to meet mission/time requirements for low false-positive rates, with an achievable reduction in false positives by as much as 98%. Even with fairly low instantaneous detection accuracy we still managed to improve the performance of the sensor while maintaining low false positive rates.

Further, we have provided the criterion for $\mathbb{E}[\Lambda_1]$. If $\mathbb{E}[\Lambda_1] > 0$ then the sensor being used to detect obstacles can be adequately used to confirm or deny the presence or absence of an obstacle. This is valuable in selecting safeguarding sensors. Knowing the distribution of obstacle sizes in the intended environment, the effective frame rate of the sensor, and the quantity $\mathbb{E}[\Lambda_1]$ provides constraints that will help future missions design safe operating speeds and to pick feasible terrain regions to investigate.

Above and beyond protecting the vehicle, the SPRT provides a principled mechanism to alter vehicle speed or sensor frame rate in order to drive safely and with confidence, a novel contribution in safeguarding sensors. When using the speed-adjusting SPRT algorithm, there is a 98.6% reduction in collisions and an 89% reduction in false positives. We demonstrate that the freedom to drive a vehicle in obstacle-strewn terrain increases the robustness and reliability of the system, and thus the trust that may be placed on it by remote operators. This low-cost safeguarding mechanism is an enabling technology in flight missions and other resource constrained operations.

Additionally this technique extends beyond safe guarding systems and could easily be applied to multi-objective science operations, ensuring vehicle speed is satisfying the constraints of all stakeholders during exploration. Anomaly detection approaches such as those used by Thompson et al. [18] and Girdhar et al. [19] could complement a SPRT-based approach.

Based on where the vehicle has travelled and how the sensor previously classified that terrain, the rover could update the sensor characterization from Table I. This will help the rover adapt to regions where the underlying sensor has varying performance, as we observed in the two KRex2 data trials.

Treating the sensor as a Bernoulli random variable can hides important information from the SPRT. Incorporating information estimated range to an obstacle and a sensor error model may yield improved performance.

ACKNOWLEDGEMENT

The authors would like to thank Xavier Bouyssounouse, and Laurence Edwards for their support in developing the virtual bumper mechanism. An especially strong appreciation is

extended to Vinh To for invaluable assistance in experimental setup.

REFERENCES

[1] M. P. Hennessey, C. Shankwitz, and M. Donath, "Sensor-based virtual bumpers for collision avoidance: configuration issues," in *Photonics East'95*. International Society for Optics and Photonics, 1995, pp. 48–59.

[2] R. Simmons, L. Henriksen, L. Chrisman, and G. Whelan, "Obstacle avoidance and safeguarding for a lunar rover," in *AIAA Forum on Advanced Developments in Space Robotics*, 1996.

[3] B. Schiller, V. Morellas, and M. Donath, "Collision avoidance for highway vehicles using the virtual bumper controller," in *Proceedings of the IEEE International Symposium on Intelligent Vehicles*, 1998, pp. 28–30.

[4] A. Gorjestani, M. Donath, L. Alexander, and M. Guidestar, *Radar Based Longitudinal Virtual Bumper Collision Avoidance System Implemented on a Truck*. Minnesota Department of Transportation, Office of Research & Strategic Services, 2000.

[5] C. P. Urmson, M. B. Dias, and R. G. Simmons, "Stereo vision based navigation for sun-synchronous exploration," in *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, vol. 1. IEEE, 2002, pp. 805–810.

[6] T. Carlson, R. Leeb, R. Chavarriaga, and J. del R Millan, "Online modulation of the level of assistance in shared control systems," in *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*. IEEE, 2012, pp. 3339–3344.

[7] M.-C. Hsieh and K.-Y. Young, "Manipulating a multi-dof robot manipulator via virtual tools," *Journal of Information Science and Engineering*, vol. 30, no. 4, pp. 1209–1225, 2014.

[8] J. A. Shah, J. H. Saleh, and J. A. Hoffman, "Analytical basis for evaluating the effect of unplanned interventions on the effectiveness of a human–robot system," *Reliability Engineering & System Safety*, vol. 93, no. 8, pp. 1280–1286, 2008.

[9] M. Maimone, A. Johnson, Y. Cheng, R. Willson, and L. Matthies, *Autonomous navigation results from the Mars Exploration Rover (MER) mission*. Springer, 2006.

[10] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT press, 2005.

[11] Shoemaker and Morris, "Size-frequency distribution of fragmental debris, surveyor project final report: Part ii. science results," Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, Tech. Rep. JPL 32-1265, 1969.

[12] K. S. Pratt and R. R. Murphy, "Protection from human error: Guarded motion methodologies for mobile robots," *Robotics & Automation Magazine, IEEE*, vol. 19, no. 4, pp. 36–47, 2012.

[13] A. Wald, "Sequential tests of statistical hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117–186, 1945.

[14] M. Kulldorff, R. L. Davis, M. Kolczak, E. Lewis, T. Lieu, and R. Platt, "A maximized sequential probability ratio test for drug and vaccine safety surveillance," *Sequential Analysis*, vol. 30, no. 1, pp. 58–78, 2011.

[15] T. Bayer, "Planning for the un-plannable: Redundancy, fault protection, contingency planning and anomaly response for the mars reconnaissance orbiter mission," *AIAA Space 2007*, 2007.

[16] H. W. Stone, "Mars Pathfinder microrover: A low-cost, low-power spacecraft," in *Forum on Advanced Developments in Space Robotics Robotics and Automation*. AIAA, 1996.

[17] J. Carpenter, F. Markley, K. Alfriend, C. Wright, and J. Arcido, "Sequential probability ratio test for collision avoidance maneuver decisions based on a bank of norm-inequality-constrained epoch-state filters," in *AAS/AIAA Astrodynamics Specialist Conference, Girdwood, AK*, 2011.

[18] D. R. Thompson, N. A. Cabrol, M. Furlong, C. Hardgrove, B. K. H. Low, J. Moersch, and D. Wettergreen, "Adaptive sensing of time series with application to remote exploration."

[19] Y. Girdhar, P. Giguere, and G. Dudek, "Autonomous adaptive underwater exploration using online topic modelling," in *International Symposium on Experimental Robotics (ISER)*, 2012.