Perching Failure Detection and Recovery with Onboard Sensing

Hao Jiang, Morgan T. Pope, Matthew A. Estrada, Bobby Edwards, Mark Cuson,
Elliot W. Hawkes and Mark R. Cutkosky

***Abstract*— Perching on a vertical surface carries the risk of severe damage to the vehicle if the maneuver fails, especially if failure goes undetected. We present a detection method using an onboard 3-axis accelerometer to discriminate between perching success and failure. An analytical model was developed to calculate acceleration differences for success and failure and set decision times. Two distinct decision times were shown to be effective, corresponding to properly engaging the gripper and overloading the gripper's capabilities. According to a machine learning feature selection algorithm, the maximum Z axis acceleration of the quadrotor and the presence of near-zero readings are the most relevant features within these two time frames. Using these features, the detection algorithm discriminated between success and failure with a 91% accuracy at 40 ms, and 94% at 80 ms. Real-time detection and failure recovery experiments with a 20 g quadrotor verify the detection method. An improved approach that combines various decision times correctly identified success/failure for all 20 trials with an average total falling distance of 0.8 m during recovery. We discuss the feasibility of extending our method to other quadrotor platforms.**

I. Introduction

Endowing Micro Air Vehicles (MAVs) with the ability to perch greatly extends mission life, enables close inspection of surfaces, and allows them to wait out undesirable flying conditions. Additionally, perching in a dynamic maneuver provides robustness against wind disturbances. Recently various perching mechanisms have been developed for different MAV platforms to adhere to smooth surfaces [1]–[4], rough surfaces [5]–[7] and even tubes [8,9]. While indoor perching with adhesives and velcro has been accomplished, the flight controls of most dynamic maneuvers rely largely on Vicon positioning systems [10,11]. Off-Vicon position control is still not accurate enough for closed-loop perching applications [12].

If perching fails, the MAV must recognize it as soon as possible; any time spent falling brings it closer to collision with the ground. Furthermore, if an MAV is able to recover from one or more failed trials, it might finally achieve a perching success. Robust perching and recovery on flat surfaces have been demonstrated with Vicon [7]. The falling distances before failure recognition and full recovery in this work are sizable, relying on registering the platform dropping a set distance. A multi-accelerometer based, off-Vicon sensing strategy has been developed to detect the incipient tilt of an avian-foot-style perched quadrotor on tubes and tree branches [13]. Such a method would not allow detection of failure on a flat wall, since incipient tilt does not necessarily indicate failure. Here we demonstrate recovery from perching failure on flat surfaces utilizing only onboard sensing with less detection time and falling distance than previous work.

All authors are with the Dept. of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA jianghao@stanford.edu

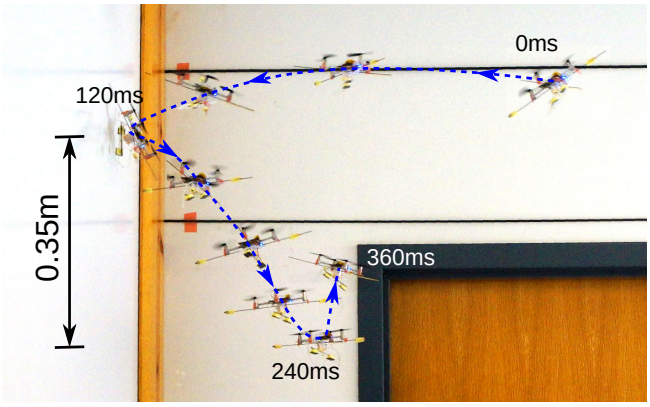

Fig. 1: Time sequence of a perching failure recovery. The platform shown is the 20 gram Crazyflie quadrotor used in real-time experiments.

Accelerometer-based sensing and tactile sensing have been widely used in robotic end-effectors to measure contact forces and vibrations. For gripping applications, such sensing methods can achieve quick detection before relative movement can be measured [14,15]. For object characterization applications accelerometers were explored to collect vibration data "blindly" [16]–[19]. More thorough reviews of such sensing methods can be found in [20].

Thus, onboard accelerometers are a likely candidate for fast detection of perching failures and perching surface characterization. The approach does not rely on the measurement of free-fall distance or prior knowledge of impact initial conditions captured by vision systems, and takes advantage of the fact that an onboard accelerometer is standard equipment on most quadrotors.

This paper focuses on using an onboard 3-axis accelerometer on a Crazyflie [21] quadrotor with a perching mechanism introduced in [1] to detect perching success/failure quickly and accurately. For dynamic perching on surfaces, a quadrotor usually has more control authority to perch on inverted surfaces than vertical or nearly vertical surfaces, which makes the latter case more challenging to recover from failure. Hence, this paper focuses on the harder case. To the best of the authors' knowledge, this is the first work to report sensing of and recovery from failed adhesive perching using onboard sensing and simple machine learning. In this paper, we first present a simplified model to describe

the perching post-impact behavior. Next, we use captured accelerometer data from perching successes and failures to train a Support Vector Machine (SVM) to select an optimal set of features. Then, we implement and verify our perching failure detection method with real-time flight tests. Finally, we discuss extensions to other quadrotor platforms.

II. Modeling

The attachment mechanism consists of a pair of controllable adhesives for adhering to a smooth surface, a rebound spring to absorb rebound kinetic energy, and several pieces of damping foam to mitigate the impact, which are shown in Fig. 2. Our model assumes that the vertical surface perching maneuver can be described in a 2-dimensional plane normal and tangential to the vertical wall; lateral motion along the surface is ignored. Detailed modeling is described in [22].

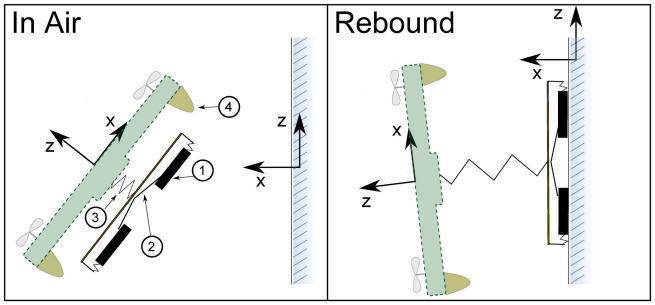

Fig. 2: The perching mechanism used in this paper. (1) Directional and controllable adhesive pads; (2) Tendons; (3) Rebound spring; (4) Damping foam. This is a lightweight (2 g) version of the mechanism used in [1,22].

We use a further simplified model to describe the characteristics of perching success and failure. For a perching success, the adhesive pads attach to the surface during impact and the quadrotor moves in a 2D plane as a rigid body constrained by the rebound spring and perching surface. Failure comes in two flavors: either the adhesive pads failed to engage on a surface (engagement failure) or the rebound spring bottoms out and rips the pads off the surface (adhesive failure). Very occasionally ($\approx$ 1% of experiments), the adhesive pads successfully attach but gradually lose adhesion and fail (delayed failure). Assuming the quadrotor does not contact the surface during rebound, the equations of motion are as follows:

$$m\ddot{x} = -kx \tag{1}$$

$$m\ddot{z} = kz - mg \tag{2}$$

where m is the mass of the quadrotor and k is the stiffness of the rebound spring. The initial conditions are $x = 0, \dot{x} = v_{reb\_x}, z = 0, \dot{z} = v_{reb\_z}$, and the constraints are $x \geq 0$ and $k\sqrt{(x^2 + z^2)} \leq F_{limit}$, where F_{limit} is the adhesion limit. Thus the quadrotor's linear motion is as follows:

$$x(t) = v_{reb\_x}\frac{1}{\omega}sin\omega t \tag{3}$$

$$z(t) = -v_{reb\_z}\frac{1}{\omega}sin\omega t + \frac{g}{\omega^2}cos\omega t - \frac{g}{\omega^2} \tag{4}$$

where $\omega = \sqrt{\frac{k}{m}}$. The quadrotor experiences maximum acceleration when the rebound spring has fully extended. The time of maximum acceleration in the surface X axis direction is:

$$t_{max\_acc} = \frac{\pi}{2\omega} \tag{5}$$

The time that the quadrotor returns to the surface in a perching success is $2t_{max\_acc}$. Assuming that the quadrotor impacts the wall without gross misalignment, the surface X axis corresponds to the quadrotor body Z axis. In actual experiments, the possible error introduced by this assumption was on the order of 5%. In the rest of the paper, surface X and body Z axes will be assumed interchangeable.

The rebound acceleration and velocity of a successful perch can be calculated given an initial rebound velocity. For an engagement failure, the quadrotor will experience negligible rebound acceleration. The difference between an engagement failure and a success will then be greatest at t_{max_acc}. An adhesive failure will still resemble a success at this time, but at time $2t_{max\_acc}$ should be clearly differentiable from success as it enters free-fall. Fig. 3 illustrates perching successes, different types of perching failures, and the corresponding accelerometer characteristics.

III. Prediction Results

Acceleration data was collected for the various perching results. The impact and rebound segments of all the data were extracted, and a list of features were constructed with analysis. An SVM algorithm selects the most relevant features to achieve a high prediction accuracy. Learning errors are subsequently discussed. More experiments were conducted to verify the detection method on another quadrotor platform.

A. Data Acquisition

150 experiments were conducted to acquire accelerometer data for perching successes and failures (74 and 76 trials respectively). Each perching attempt consisted of four phases: approaching, impact, rebound, and stay/free-fall, illustrated in Fig. 3. The perching platform used in the experiments is a 20 g Crazyflie quadrotor [21] equipped with a 2 g adhesive mechanism, providing a maximum normal adhesion of 2 N. The stiffness of the rebound spring is approximately 60 N/m. The onboard accelerometer samples at 100 Hz. In each trial the quadrotor was launched at a vertical glass plate with various initial conditions. The motors and propellers were disabled to launch with a ballistic trajectory, which is similar to real perching scenarios. Perching successes were achieved by launching the quadrotor at optimal conditions: clean adhesive pads, a velocity of 1-2 m/s, and an angular misalignment less than 20°. Perching failures were achieved

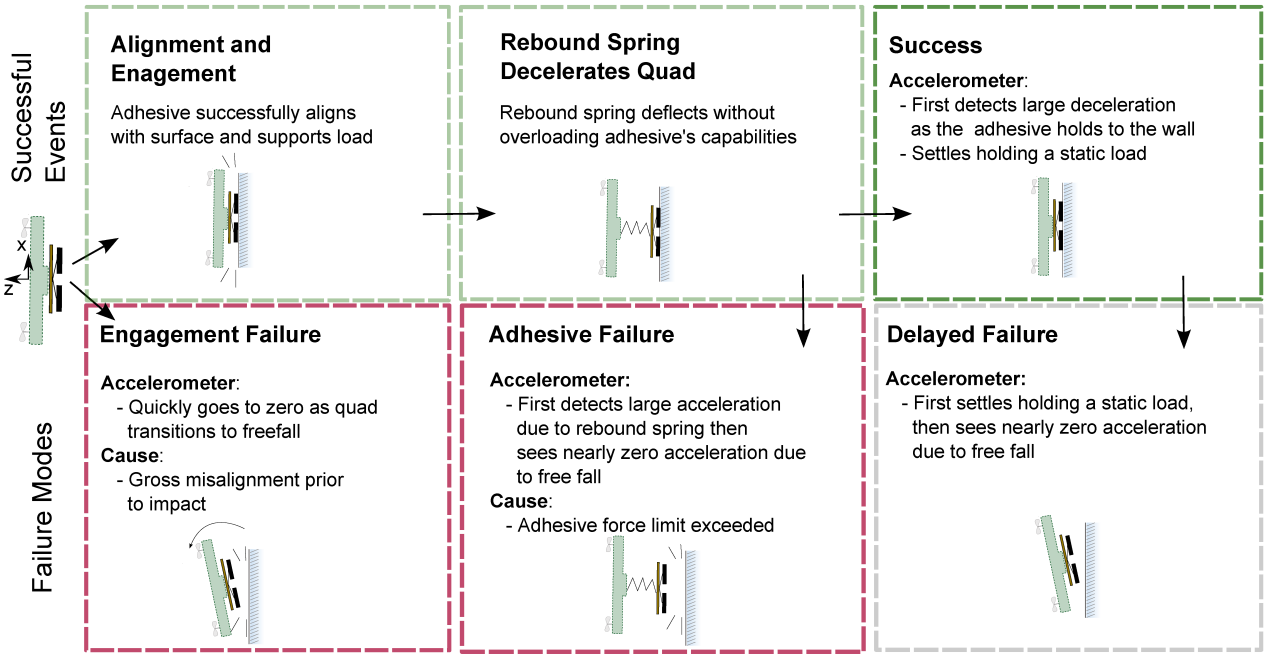

Fig. 3: Illustration of different types of perching failures comparing to perching successes.

by either occluding the adhesive, perching upon dirty surfaces, applying velocities outside the gripper's envelope of acceptable conditions, or applying a significant angular misalignment before impact. Delayed failures were difficult to induce and happened by chance. This method focuses on analyzing accelerations after the impact; the initial conditions serve as a black box and are not presented.

B. Data Processing

Raw data of representative cases are shown in Fig. 4, with the exception of a delayed failure, which typically occurred after data collection ceased. The accelerometer data were not filtered; a typical perching maneuver included less than 30 data points and did not show much noise. Using impact as the starting data point, there are two decision time choices: 40 ms (5 data points) and 80 ms (9 data points) after the impact. According to the model, t_{max_acc} is at about 40 ms after the impact, and is the earliest decision time that accelerometer data show a noticeable difference between perching success and engagement failure. 80 ms after the impact corresponds to $2t_{max\_acc}$, and is the earliest time for reliably differentiating between success and adhesive failure.

C. Feature Construction

Potential features are "constructed," or aggregated, to have their relevance assessed by a machine learning algorithm. According to the analysis in Section II, accelerometer data includes enough information for detecting perching success/failure. Thus all the 3-axis accelerometer data points between the impact point and the decision point are selected as relevant features for detection. Since the duration of perches is inconsistent, a window spanning impact to decision cutoff was set to check for relevant signal features. Such relevant features include but are not limited to (1) large acceleration

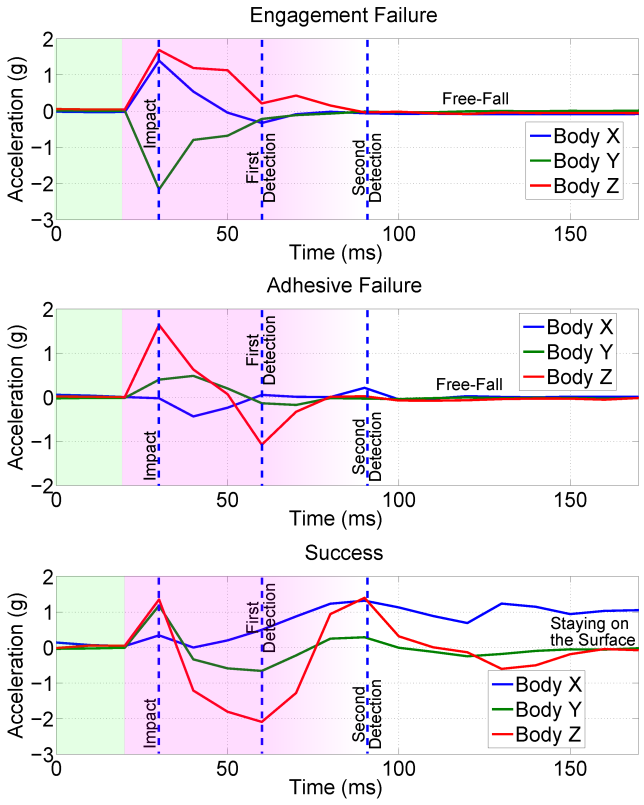

Fig. 4: Raw accelerometer data from a typical engagement failure (top), a typical adhesive failure (middle), and a typical perching success (bottom). The green shaded region is the ballistic flying stage, and the red shaded region is the impact and rebound stage. Two detection time choices illustrate significant difference between success and failure. The first decision time helps to differentiate between successes and engagement failures, and the second decision time helps to differentiate between successes and adhesive failures.

Feature Index, 40ms	Feature	Feature Index, 80ms
1 - 5	Acceleration in body X	1 - 9
6 - 10	Acceleration in body Y	10 - 18
11- 15	Acceleration in body Z	19 - 27
16 - 20	Acceleration magnitude	28 - 36
21	Max acceleration in body X	37
22	Max acceleration in body Z	38
23	Nearly zero acceleration	39
24 - 27	Jerk in body X	40 - 47
28 - 31	Jerk in body Y	48 - 55
32 - 35	Jerk in body Z	56 - 63
36 - 38	Velocity change	64 - 66

Fig. 5: A list of all relevant features input to the SVM. The algorithm scored the most relevant features as seen in Fig. 6.

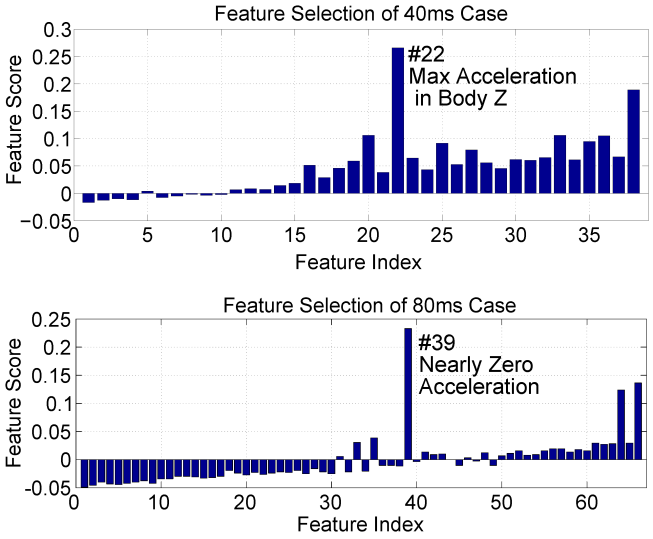

Fig. 6: Scores of all possible features. Upper: Detection at 40 ms after the impact. Lower: Detection at 80 ms after the impact. Higher score corresponds to more direct correlation to success and failure detection, and a negative score suggests overfitting problem by adding that feature.

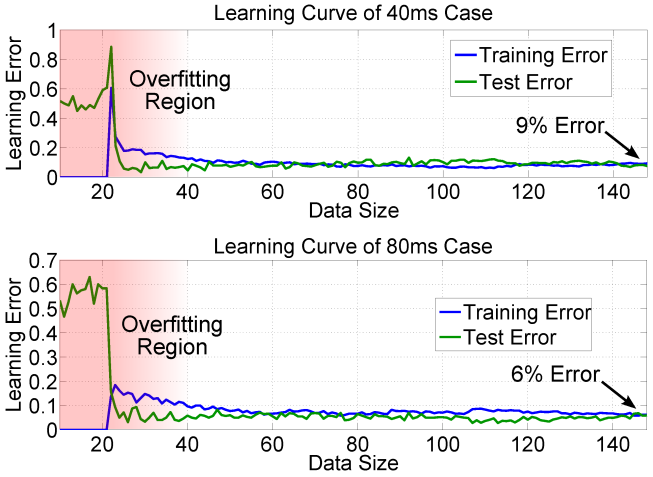

Fig. 7: The learning curves for both decision time choices. Upper: Detection at 40 ms after the impact. Lower: Detection at 80 ms after the impact. The training error and test error converge and become stable as the total data size increases, suggesting an adequate data size.

in the body Z axis due to force from the rebound spring and (2) at least two data points of nearly zero acceleration caused by free-fall, to avoid mistakenly regarding successes as failures. The difference in neighbored accelerations (jerk) is also selected as relevant to capture the stretching speed of the rebound spring. Velocity changes in perching successes are also larger than in perching failures due to the effect of the rebound spring, and thus the integration of acceleration in the second segment could be relevant for detection. Delayed failures were not considered in feature construction since their time to failure is typically tens of seconds after impact. Instead, a simple, reoccurring check for accelerometer readings close to zero after perching screens for delayed failures. Figure 5 shows a list of all relevant features based on analysis.

D. Feature Selection

With the pre-selected features listed above, we used a Support Vector Machine (SVM) developed by [23] to evaluate each feature and select a boundary to predict success or failure based just on that feature. We then compared the prediction with the real result to calculate prediction accuracy. While all the features listed in Table 5 can contribute to the detection, using every feature causes an over-fitting problem. Training error becomes abnormally low while the test error becomes high. With a backward searching algorithm, each feature was assigned a score indicating the relevance to correct detection. Two histograms listing the scores of all possible features for both decision time choices are shown in Fig. 6. A higher score corresponds to higher relevance. Only a few of the highest scoring features were ultimately selected for prediction.

A cross validation algorithm within the 150 data sets was implemented to verify the performance of the selected features. 70% of the data were used as training data with the remaining 30% used as test data. The training set and test set were randomly selected for every cross validation, and the validation script ran 50 times. Training and test errors were averaged afterwards. A learning curve showing the trend of learning errors was plotted for both decision times in Fig. 7 as the size of the combined training and test sets varied from 10 to 150. The curves of training and test errors decrease and converge as the total data size increases, suggesting that the SVM algorithm can predict the result of a new trial with similar accuracy, using all 150 trials as the training set.

Once we have identified the most important feature at each decision time, we can calculate a threshold for discrimination based on the simple model developed in Section II and see how our analysis matches the thresholds predicted by the SVM algorithm. For the 40 ms case, the most relevant feature is large acceleration in the body Z direction. The initial

rebound velocity of typical successful perchings is between 0.2 m/s and 0.6 m/s, corresponding to a calculated maximum acceleration in body Z direction of 1.1 g to 3.3 g. Engagement failures would show negligible acceleration in the body Z direction, so a threshold set around the lower calculated acceleration should discriminate effectively between success and engagement failure. The SVM computes a decision boundary of 1 g, supporting this analysis. For the 80 ms case, the most relevant feature is near-zero acceleration magnitude. This indicates free-fall, and the decision boundary can be selected based on the noise of accelerometer data. The SVM algorithm computes a boundary of 0.3 g.

At 40 ms after impact the SVM algorithm predicts success/failure with an accuracy of 91% by detecting the presence of a greater than 1 g acceleration in the body Z axis. With decision time at 80 ms after the impact, the most relevant feature is the presence of at least two data points with acceleration magnitude smaller than 0.3 g. The corresponding prediction accuracy is 94%. For both cases adding other important features boosts the accuracy by about 1% but needs proper weighting among features, which is difficult to predict with the model. Acceleration magnitude near zero is not very relevant for the 40 ms case since it is not always enough time for a free-fall to occur, but 80 ms almost always includes this information. There are tradeoffs: a 40 ms decision time is less accurate than an 80 ms decision, but results in a shorter free-fall distance (within 100 mm) before detection. For an 80 ms decision, the free-fall distance can be several hundred millimeters before detection. Other features are also relevant but have limited contribution to the prediction accuracy, and are thus neglected.

E. Causes and Considerations for Error

Even if the prediction accuracies for both decision time choices are high, potential prediction errors can still cause an MAV to crash. Within the 14 error cases for 40 ms, there are 8 data sets where the algorithm failed to pick the right impact point due to some abnormal impact (e.g. almost no ballistic flying stage or multiple impacts). The remaining error cases were usually caused by the fact that the data were not completely separable within the dimension of the selected feature. For 80 ms, the selected feature that detects free-fall separates the data better. Note that there are several other features that also have relatively high scores according to Fig. 6. Adding these features can further boost the prediction accuracy to 95%, but requires proper weightings among features. Such weightings can be computed by the SVM algorithm but are difficult to calculate by analysis.

Learning errors can be further divided into successes being regarded as failures (false negative) and failures being regarded as successes (false positive). If a false negative detection occurs, the quadrotor will be fully powered on and try to pry the pads off the surface. With the current experimental setup the quadrotor does not have enough thrust to overcome the adhesive limit of the mechanism, so this type of error does not carry high risk.[1] If a false positive detection occurs, the quadrotor will power down and fall to the ground. This may damage the quadrotor and thus should be avoided. By intentionally tuning the thresholds in feature construction, we can bias the results towards false negatives instead of false positives, possibly incurring a higher total error. With some adjustment, false positive frequency was reduced to 20% lower than false negative frequency, with negligible decrease in total accuracy.

F. Extension to A Larger Quadrotor

The machine learning algorithm verifies the simplified model for a Crazyflie platform in terms of acceleration and decision time choices, but a larger vehicle might behave differently. To begin to evaluate the applicability of such detection methods to other platforms, more experiments were conducted.

The new platform weighs 150 g and is connected to a pair of opposed adhesive pads (8 N normal adhesion limit) with a rebound spring (stiffness 50 N/m). The onboard accelerometer is identical to that of the Crazyflie platform. Based on the new mass and stiffness parameters, the model developed in Section II predicts the time that the vehicle experiences maximum acceleration in the body Z direction is about 90 ms after impact, which correspond to 10 data points. The time that the quadrotor returns to the surface in a perching success is then 180 ms, corresponding to 19 data points. These two decision times were verified by visual inspection of the acceleration data plots, and 55 data sets were recorded for machine learning and feature selection.

For 90 ms after impact, the maximum acceleration is still one of the five most relevant features. The most relevant feature shifts from maximum acceleration to the specific acceleration at 90 ms, but the difference in prediction error from using the original feature is small (2%). The range of maximum acceleration in a perching success is calculated to be between 1.4 g and 2.8 g, with some uncertainty due to imprecision in velocity measurement. The decision boundary for using maximum acceleration in body Z axis computed by the machine is 1.5 g, which is indeed close to the lower acceleration boundary predicted by the model. The prediction accuracy is 89%, which is similar to the Crazyflie 40 ms decision time. For 180 ms after the impact, the most relevant feature is still acceleration magnitude near zero, and the prediction accuracy is 96% - similar to Crazyflie 80 ms case. Thus, it seems feasible to adopt the model and the detection method to other quadrotor platforms.

IV. Real-time Experimental Results

Real-time experiments were conducted using the Crazyflie quadrotor to verify the feasibility of the selected features to detect perching success and failure. For each experiment, the quadrotor was commanded to first fly at the wall at 2 m above the ground and then pitch back to 90° to present the opposed adhesive mechanism to a smooth surface. The maneuver used

[1]When take-off is desired, the gripper relaxes so that little force is required for detachment.

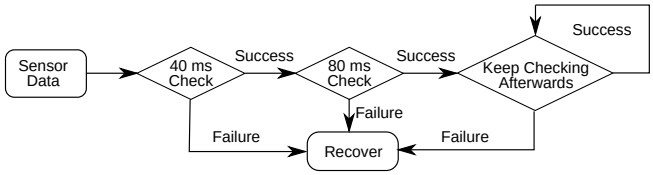

Fig. 8: An improved strategy that combines the detection at 40 ms and 80 ms to enable high detection accuracy.

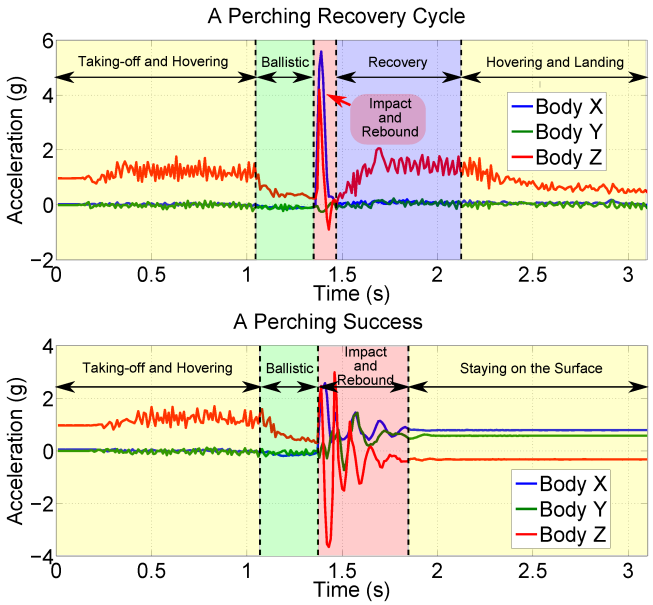

Fig. 9: The raw accelerometer data from a complete perching recovery cycle and a perching success. The perching recovery cycle includes taking-off, hovering, ballistic approaching, impact and rebound, recovering, hovering, and landing. For the perching success, the quadrotor stayed on the surface after rebound.

open loop control without the aid of a Vicon system. The accelerometer data of a complete perching recovery cycle and a perching success are illustrated in Fig. 9.

The time of impact is determined by looking for a period of low-acceleration ballistic motion followed by an acceleration spike on impact. In the real-time experiments, the ballistic phase had slightly higher accelerations because the rotors were still spinning down, and the acceleration spike tended to be larger because of higher average perching velocities. No part of the failure detection algorithm was changed. Once an impact is detected, the quadrotor begins to spin up its rotors for the recovery maneuver. If the algorithm determines that the perch is successful at 40 ms or 80 ms, the rotors power down. The algorithm continues to monitor the accelerometer data in case of a delayed failure. This strategy is illustrated in Fig. 8. Computation took place on a nearby laptop communicating wirelessly with the quadrotor. The recovery manuever simply utilized the quadrotor's built-in stabilization routine to return to level flight.

Figure. 10 shows the statistical performance of twenty real-time experiments. Among the twenty experiments, there are fifteen failures and five successes. The maneuver was biased towards producing initial conditions leading to failure

Specs	Average	STD
Post-impact Attitude (Degrees)	86	19
Maximum Trust-to-Weight Ratio	1.4	0.2
Post-Impact Tangential Velocity (m/s)	0.6	0.5
Time before Detection (ms)	45	14
Falling Distance before Detection (m)	0.04	0.03
Time before Zero Pitch (ms)	310	60
Falling Distance before Zero Pitch (m)	0.4	0.2
Time before Full Recovery (ms)	660	170
Falling Distance before Full Recovery (m)	0.8	0.4

Fig. 10: A list of statistical performance of perching success/failure detection. The high standard deviation of falling distance is due to the high standard deviation of post-impact attitude, thrust and tangential velocity, which were difficult to keep consistent due to open loop control.

to facilitate exploration of the post-failure recovery behavior of the quadrotor. In all twenty experiments, the failure detection algorithm correctly identified the perching result. For the failure cases, engagement failures had an average falling distance before detection of 3 cm, while adhesive failures had an average falling distance of 6 cm. The vehicle accelerates downward until the attitude has been corrected to level, which takes an average of 0.3 s, then decelerates until it stops falling and begins to move upward. The Crazyflie used in the experiments has an average thrust-to-weight ratio of 1.4, which corresponds to a maximum deceleration of 0.4 g. On average, the vehicle took 0.8 m to reverse velocity and begin moving upward. Total falling distance varies depending on the post-impact attitude, thrust level, and tangential velocity of the quadrotor.

V. Discussion

A. Perching on Vertical Surfaces

Results show that the detection method of setting a decision time, selecting the most relevant features (maximum acceleration and nearly zero acceleration), and setting a boundary is applicable to other quadrotor platforms. The experiments using a larger platform detailed in Section III show consistency of relevant features across the two different vehicles analyzed. They also show that knowledge of system mass and rebound spring stiffness enable calculation of appropriate detection times. Furthermore, if post-impact rebound velocity can be determined from expected pre-impact conditions and the coefficient of restitution for collision, the boundary value for the acceleration in body Z can also be predicted. Thus, given a new quadrotor, it should not be necessary to conduct many experiments to take data and do machine learning to set decision boundaries. A few launching tests with the propellers off should be sufficient to compare with the model, settle the exact boundary, and detect the perching result with high accuracy.

B. Strategies for Extension to Other Types of Perching

While the proposed sensing strategy has been verified on dynamic perching on vertical surface, it can also be extended to perching on various other surface orientations and curvatures with some modifications. The simple model can be adapted by changing the initial rebound velocity vector to capture the dynamics of various perching scenarios. Combined with launching tests and machine learning, the final crucial parameters should be similar to vertical surface perching with slightly different values. However, for floor perching where there is little difference between perching success and failure, additional sensors need to be included to detect the adhesion quality to detect perching success/failures.

VI. CONCLUSIONS

We present a success/failure discrimination method for quadrotor vertical surface perching by using onboard accelerometer data. With several most relevant features selected by a support vector machine, the method is able to determine a perching outcome with 91% accuracy after 40 ms and with 94% accuracy after 80 ms. Real-time experiments verify the detection method. Perching failure is detected after an average falling distance of 4 cm, and recovery from failure is accomplished with an average total falling distance of 0.8 m.

The simple analytical model shows that critical decision times rely upon the mass of the system and the stiffness of the rebound spring. With these parameters and the typical rebound velocity known, appropriate boundaries for failure detection can be calculated for other platforms. Additionally, the rapid discrimination of success and failure allows for repeated trials and the identification of unsuitable perching surfaces.

In the future a full flying-perching-recovery-reperching cycle will be demonstrated. Accelerometer arrays will be implemented on the quadrotor to provide more information about misalignment with the surface, and thus inform perching success/failure earlier. New force sensors integrated directly on perching mechanisms will be made to further shorten the decision time and identify surface suitability for perching. Especially for micro-spine based out-door perching, a new strategy will be developed to differentiate between an improper perching initial condition and a bad choice of perching surface, and therefore, inform whether it is worthwhile to reperch on the same surface.

ACKNOWLEDGMENTS

Support for this work was provided by ARL MAST MCE 15-4.4 and NSF 1161679. Matt A. Estrada is supported by NSF Graduate Fellowship. We thank Shiquan Wang, Li Tao and all the members from the Biomimetics and Dextrous Manipulation Laboratory for the help of implementing the SVM algorithm and video recording.

REFERENCES

[1] E. W. Hawkes, D. L. Christensen, E. V. Eason, M. A. Estrada, M. Heverly, E. Hilgemann, H. Jiang, M. T. Pope, A. Parness, and M. R. Cutkosky, "Dynamic surface grasping with directional adhesion," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pp. 5487–5493, IEEE, 2013.

[2] L. Daler, A. Klaptocz, A. Briod, M. Sitti, and D. Floreano, "A perching mechanism for flying robots using a fibre-based adhesive," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pp. 4433–4438, IEEE, 2013.

[3] A. Peyvandi, P. Soroushian, and J. Lu, "A versatile perching mechanism employing shape memory wires and bio-inspired adhesives," *Journal of Micro-Bio Robotics*, pp. 1–13, 2014.

[4] A. Kalantari, K. Mahajan, D. Ruffatto III, and M. Spenko, "Autonomous perching and take-off on vertical walls for a quadrotor micro air vehicle," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pp. 4669–4674, IEEE, 2015.

[5] A. L. Desbiens, A. T. Asbeck, and M. R. Cutkosky, "Landing, perching and taking off from vertical surfaces," *The International Journal of Robotics Research*, 2011.

[6] M. Kovač, J. Germann, C. Hürzeler, R. Y. Siegwart, and D. Floreano, "A perching mechanism for micro aerial vehicles," *Journal of Micro-Nano Mechatronics*, vol. 5, no. 3-4, pp. 77–91, 2009.

[7] D. Mellinger, M. Shomin, and V. Kumar, "Control of quadrotors for robust perching and landing," in *Proceedings of the International Powered Lift Conference*, pp. 205–225, 2010.

[8] C. E. Doyle, J. J. Bird, T. A. Isom, J. C. Kallman, D. F. Bareiss, D. J. Dunlop, R. J. King, J. J. Abbott, and M. A. Minor, "An avian-inspired passive mechanism for quadrotor perching," *Mechatronics, IEEE/ASME Transactions on*, vol. 18, no. 2, pp. 506–517, 2013.

[9] P. Xie, O. Ma, L. Zhang, and Z. Zhao, "A bio-inspired uav leg-foot mechanism for landing, grasping and perching tasks," 2015.

[10] D. Mellinger, N. Michael, and V. Kumar, "Trajectory generation and control for precise aggressive maneuvers with quadrotors," *The International Journal of Robotics Research*, 2012.

[11] S. Lupashin, A. Schollig, M. Sherback, and R. D'Andrea, "A simple learning strategy for high-speed quadrocopter multi-flips," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 1642–1648, IEEE, 2010.

[12] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadrocopter state estimation," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pp. 1730–1736, IEEE, 2015.

[13] K. L. Crandall and M. A. Minor, "Uav fall detection from a dynamic perch using instantaneous centers of rotation and inertial sensing," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pp. 4675–4679, IEEE, 2015.

[14] M. R. Tremblay and M. R. Cutkosky, "Estimating friction using incipient slip sensing during a manipulation task," in *Robotics and Automation, 1993. Proceedings., 1993 IEEE International Conference on*, pp. 429–434, IEEE, 1993.

[15] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. J. Kuchenbecker, "Human-inspired robotic grasp control with tactile sensing," *Robotics, IEEE Transactions on*, vol. 27, no. 6, pp. 1067–1079, 2011.

[16] J. M. Romano and K. J. Kuchenbecker, "Creating realistic virtual textures from contact acceleration data," *Haptics, IEEE Transactions on*, vol. 5, no. 2, pp. 109–119, 2012.

[17] K. J. Kuchenbecker, J. Fiene, and G. Niemeyer, "Event-based haptics and acceleration matching: Portraying and assessing the realism of contact," in *Eurohaptics Conference, 2005 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2005. World Haptics 2005. First Joint*, pp. 381–387, IEEE, 2005.

[18] J. Windau and W.-M. Shen, "An inertia-based surface identification system," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pp. 2330–2335, IEEE, 2010.

[19] P. Wisanuvej, J. Liu, C.-M. Chen, and G.-Z. Yang, "Blind collision detection and obstacle characterisation using a compliant robotic arm," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pp. 2249–2254, IEEE, 2014.

[20] M. Cutkosky, R. Howe, and W. Provancher, *Springer Handbook of Robotics*, ch. Force and Tactile Sensors, pp. 455–476. Springer.

[21] "http://www.bitcraze.se/crazyflie/,"

[22] H. Jiang, M. T. Pope, E. W. Hawkes, D. L. Christensen, M. A. Estrada, A. Parlier, R. Tran, and M. R. Cutkosky, "Modeling the dynamics of perching with opposed-grip mechanisms," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pp. 3102–3108, IEEE, 2014.

[23] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "Liblinear: A library for large linear classification," *The Journal of Machine Learning Research*, vol. 9, pp. 1871–1874, 2008.