# Self-Supervised Online Learning for Safety-Critical Control using Stereo Vision

Ryan K. Cosner<sup>\*</sup>, Ivan D. Jimenez Rodriguez<sup>\*</sup>, Tamas G. Molnar, Wyatt Ubellacker, Yisong Yue, Aaron D. Ames, and Katherine L. Bouman

Abstract—With the increasing prevalence of complex vision sensing methods for use in obstacle identification and state estimation, characterizing environment-dependent measurement errors has become a difficult and essential part of modern robotics. This paper presents a self-supervised learning approach to safety-critical control. In particular, the uncertainty associated with stereo vision is estimated, and adapted online to new visual environments, wherein this estimate is leveraged in a safety-critical controller in a robust fashion. To this end, we propose an algorithm that exploits the structure of stereo-vision to learn an uncertainty estimate without the need for ground-truth data. We then robustify existing Control Barrier Function-based controllers to provide safety in the presence of this uncertainty estimate. We demonstrate the efficacy of our method on a quadrupedal robot in a variety of environments. When not using our method safety is violated. With offline training alone we observe the robot is safe, but overly-conservative. With our online method the quadruped remains safe and conservatism is reduced.

#### I. INTRODUCTION

Accounting for vision-based uncertainty is particularly important for modern safety-critical robotic applications such as autonomous vehicles, health care, and manufacturing [1]. The importance of safety for these systems demands that their controllers are designed to provide robust safety in the presence uncertainty. *Control Barrier Functions (CBFs)* [2], [3] have become a popular tool for achieving provable safety. However, standard CBF theory requires accurate state estimation. This motivates the need for a method that provides safety when using stereoscopic sensor measurements that are inherently noisy.

Computer vision has become an important tool in robotics for sensing environments and identifying obstacles. They are often an integral component of robotics applications such as simultaneous localization and mapping (SLAM [4]). Despite their utility and ubiquity, using vision sensors to achieve robust safety is difficult due to the complex environmentdependent error that they generate. For example, error patterns are highly correlated with the textures and appearance of a scene. Supervised methods can identify and model error as it affects the CBF [5], [6], [7], [8], [9]; however, supervised approaches require ground-truth training data that



**Fig. 1.** These space-time images display our quadrupedal robot throughout the course of an experiment. The robot is considered safe if it remains left of the yellow line. The standard control barrier function (CBF) condition fails to keep the robot safe due to errors in stereo vision; the robust CBF condition keeps the robot safe, but is conservative; and our proposed method, Robust CBFs with Online Uncertainty Estimation, keeps the robot safe without remaining overly conservative.

may be difficult or impossible to obtain. Additionally, such frameworks often inaccurately estimate errors in regions of the state space that were under-sampled or generally not present in the training set; training on indoor environments or synthetic data with easy-to-access ground truth often does not translate well to outdoor environments.

In this paper, we focus on two main challenges related to automatic vision-based safety critical control: (i) the perception of the safe region is uncertain, and (ii) the uncertainty is often unknown and may depend heavily on the environment. We tackle these issues in the context of stereo vision-based obstacle avoidance on a quadrupedal robot. Namely, our goal is to avoid obstacles seen in stereo cameras mounted on the robot, while taking into account that vision-based measurements are uncertain. We use a multibaseline stereo vision system on the robot to record stereo images. Then, we determine the position of the objects associated with each image pixel, and subsequently infer the uncertainty of these positions through a self-supervised error estimation algorithm that frames the problem as online learning [10]. Online learning frameworks have shown success in a variety of robotic applications [11], [12], [13] Finally, we use the position and uncertainty estimates for safety critical control,

<sup>\*</sup> These authors contributed equally.

This research is supported in part by the National Science Foundation, CPS Award #1932091, BP, Aerovironment and Dow (#227027AT).

Authors are with the Department of Control and Dynamical Systems, the Department of Mechanical and Civil Engineering, and with the Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA, {rkcosner, ivan.jimenez, tmolnar, wubellac, yyue, ames, klbouman}@caltech.edu.



Fig. 2. The overarching structure of our approach. It begins on the bottom right by capturing three time-synchronized images that are then fed into an uncertainty estimation pipeline and also used to generate a 3D point cloud. There are three possible CBF filters that result in the three possible robot realizations shown. From top to bottom, the standard filter only takes into account the noisy point-cloud in avoiding obstacles. The robust CBF safety filters use the estimate of the uncertainty  $\mathcal{P}$  to compensate for noise in the point cloud. Finally, the "Robust with Retraining" filter refines the model of uncertainty to the current environment in real-time.

wherein we achieve robust safety via CBFs. A visualization of this method can be found in Fig. 2.

The contributions of this work are three-fold. First, we present and evaluate an online, self-supervised method for characterizing the uncertainty of disparity errors generated by stereo vision algorithms in novel environments (Section II). Second, we develop a robustified CBF-based control method for utilizing this error estimate for obstacle avoidance in robotic systems (Section III). And third, we demonstrate the proposed methods of error estimation and obstacle avoidance on a quadrupedal robot operating in real time (Section IV).

# II. STEREO VISION UNCERTAINTY QUANTIFICATION

We begin by revisiting stereo vision-based depth estimation. We then propose an approach for learning the uncertainty of a black box stereo-matching algorithm. The proposed self-supervised learning approach can be trained online and takes advantage of geometric structure in stereo disparity maps so as not to require ground truth data.

#### A. Background in Stereo Vision

Stereo vision is a popular tool for determining depth from images. These methods compute a *disparity*: the shift observed in an object's projection onto two camera planes. Using a geometric understanding of the camera setup, pixelbased disparity maps can be converted to depth maps. Errors in the final depth-map result from a combination of pixel-mismatch in disparity estimation and error in the camera parameters used to convert from disparity to depth. The errors in the intrinsic and extrinsic parameters of the camera are usually small and their effect on the resulting depth distribution is easy to compute. On the other hand, pixel matching errors are much larger and the result of a much more complicated stereo matching procedure whose effect on the resulting disparity is difficult to quantify and environment-dependent.

For standard stereo vision we adopt the model of [14] for two cameras (left and right) and assume that they are perfectly rectified, vertically aligned and evenly spaced with known distance  $b \in \mathbb{R}_{>0}$  between each camera. Pixel coordinates within an image are given by the tuple  $p \triangleq (u, v) \in K$ , where  $K \triangleq \{0, \ldots, W\} \times \{0, \ldots, H\}$  for image width  $W \in \mathbb{N}_{>0}$  and image height  $H \in \mathbb{N}_{>0}$ .

Stereo algorithms such as Block Matching, Semi-Global Block Matching, and Efficient Large-Scale Stereo [15] compute disparities by determining the discrete pixel distance between matching regions of two images. Since the disparity represents a shift between pixels of two images, the measured disparity  $\hat{d}$  must be a finite integer value. Assuming that the true disparity d is a finite integer implies that the error  $e \triangleq \hat{d} - d$  must also be a finite integer. Prior work has been done to interpolate disparities for non-integer subpixel accuracy [16]; however, we restrict our attention to integer disparity values to highlight the error in pixel-matching.

#### B. Self-supervised Error Estimation

To learn the error in disparity, we introduce a three-camera multibaseline stereo system which produces multiple disparity maps that are related through simple functions; deviations from the ideal relationship indicate error in the estimated disparities. By analyzing the correlation of image appearance with these errors, a function that estimates disparity error from appearance is learned and used to specify state errorbounds in real-time for use in a robustified CBF.

We introduce a three-element camera system, whose central camera is assumed to be perfectly rectified and vertically aligned with the other two cameras as shown in Fig. 2. This third camera is placed between the left and right cameras such that it has a baseline of b/2 with both. The three cameras produce a time-synchronized grayscale image triple  $(I_1, I_2, I_3)$  where  $I_i \in \mathbb{N}^{W \times H}$  for  $i \in 1, 2, 3$  and 1, 2, 3 correspond with left, center, and right, respectively. The disparity between any image pair  $(I_i, I_j)$  for i < j is obtained using the stereo-vision algorithm  $\mathcal{D} : \mathbb{N}^{W \times H} \times \mathbb{N}^{W \times H} \to \Gamma^{W \times H}$ , so that  $\hat{d}_{i,j} = \mathcal{D}(I_i, I_j)$ . Here,  $\Gamma \subset \mathbb{N}_{\geq 0}$  is the set of possible disparity values.

Given the measurement  $d_{i,j}$ , the error appears as follows  $\hat{d}_{i,j} = d_{i,j} + e_{i,j}$  with error distribution  $e_{i,j} \sim \mathcal{P}(I_i, I_j)$ . Here,  $e_{i,j} \in \Gamma^{W \times H}$  represents the error between the true disparity  $d_{i,j} \in \Gamma^{W \times H}$  and the measured disparity  $\hat{d}_{i,j} \in \Gamma^{W \times H}$ . We model this error as a discrete random variable with probability  $\mathcal{P}(I_i, I_j)$  on  $\Gamma$ . This model of disparity errors contrasts sharply with other common error models, such as punctual observation, uniform observation, and Gaussian observation [14], in that it accounts for the discrete nature of stereo-pixel matching algorithms. If a ground-truth measurement of  $d_{i,j}$  exists, then supervised learning methods can be implemented to directly estimate this error term. However, it is often the case that such ground-truth estimates are unavailable; particularly when a domain transfer must occur during operation. Thus we seek a general method to estimate e for any black-box disparity algorithm without the need for ground-truth disparity data.

We leverage the known geometric relationships between the three cameras to learn a mapping between image appearance and disparity error uncertainty distribution that can adapt during operation in new environments. Given a multibaseline stereo system, if one ignores occlusions, it is possible to completely reconstruct each disparity map from the other two maps. The relationship to reconstruct  $\hat{d}_{1,3}$ from  $\hat{d}_{1,2}$  and  $\hat{d}_{2,3}$  is shown in Algorithm 1; we denote this reconstruction as  $\bar{d}_{1,3} \triangleq \hat{d}_{1,2} \oplus \hat{d}_{2,3}$ .

 Algorithm 1 Disparity Reconstruction:  $\overline{d}_{1,3} = \widehat{d}_{1,2} \oplus \widehat{d}_{2,3}$  

 1:  $\overline{d}_{1,3} \leftarrow \mathbf{0}_{H \times W}$  

 2: for  $v \in [1, ..., H]$  do

 3: for  $u \in [1, ..., W]$  do

 4:  $\widehat{u} \leftarrow n + \widehat{d}_{1,2}(u, v)$  

 5:  $\overline{d}_{1,3}(u, v) \leftarrow \widehat{d}_{1,2}(u, v) + \widehat{d}_{2,3}(u, \widehat{v})$  

 6: end for

 7: end for

We use the reconstructed disparity  $\overline{d}_{1,3}$  to learn the parameters  $\theta$  of a function  $\mathcal{P}^{\theta}$  that approximates error distribution  $\mathcal{P}$  (refer to Algorithm 2). Since this method does not require ground truth information, Algorithm 2 can be run online during operation to adapt  $\mathcal{P}^{\theta}$  to new visual environments. Recall that the disparity error,  $e_{1,3}$  is discrete in nature. Therefore, the pixel-wise reconstruction error  $re(p) \triangleq ||\widehat{d}_{1,3}^p - \overline{d}_{1,3}^p||_1$ will also be discrete. For this reason, optimizing the loss L reduces to a pixel-wise classification problem similar to image segmentation. Thus, as is done in image segmentation, we use pixel-wise cross entropy as the loss function L. This method is shown in Algorithm 2. In Line 8, for each pixel p

Algorithm 2 Self-Supervised Stereo Error Estimation Adaptation

1:  $L \leftarrow 0$ while robot is running do 2:  $(I_1, I_2, I_3) \leftarrow \text{Capture Current Frame}$ 3:  $d_{1,2} \leftarrow \mathcal{D}(I_1, I_2)$ 4: 5: 6: 7:  $\begin{aligned} & re(p) \leftarrow \left| \hat{d}_{1,3}^p - \overline{d}_{1,3}^p \right| \\ & L \leftarrow -\frac{1}{H \times W} \sum_p \mathbb{E}_{\mathbb{1}(re(p))}[\log \mathcal{P}^{\theta}(I_i, I_k)] \\ & \theta_{t+1} \leftarrow \theta_t - \eta \frac{\partial L}{\partial \theta} \end{aligned}$ 8: 9: 10: 11: end while

of the disparity  $\hat{d}_{1,3}$  the corresponding reconstruction error is computed. The loss function in Line 9, then is equivalent to the expected negative log likelihood of each pixel under the proposed model  $\mathcal{P}^{\theta}$ . An example visualization of lines 3-8 can be found in Fig. 3. Although this algorithm focuses on the reconstructed disparity  $\overline{d}_{1,3}$ , it can be easily extended



Fig. 3. Lines 3-8 of Algorithm 2 illustrated from left to right. Starting from three time-synchronized images three pairwise disparities are computed as shown in the middle column. Two of these disparities are used to build a reconstruction of the third disparity shown in the top right which can then be used to estimate the pixel-wise error of the stereo algorithm shown in the bottom right image. These steps of the algorithm correctly identify that the back of the closest chair is a high-error region without using ground truth information. This information is used to learn a correspondence between visual features and error distributions.

to similar reconstructions of  $d_{1,2}$  and  $d_{2,3}$ .

Supervised methods have been used in the past to estimate uncertainty in robotic applications by computing the covariance of state estimates [17]. Our approach differs in that we do not require ground truth and we take advantage of the discrete structure of images to learn a discrete, rather than a Gaussian, distribution, which is better suited to stereo measurements.

#### III. SAFE VISION-BASED CONTROL

In this section we review Control Barrier Functions (CBFs) [18] as a tool for guaranteeing the safety of dynamical systems. We then propose CBFs that rely on the position of pixels provided by stereoscopic sensing. Finally, we incorporate the proposed self-supervised error estimates of Section II to enforce robust safety.

#### A. Control Barrier Functions

First we give a brief introduction to CBFs which follows our description in [19], where additional technical details can be found. In this work we consider the safety of robotic systems with control affine dynamics

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u}, \quad \mathbf{x} \in \mathbb{R}^n, \ \mathbf{u} \in \mathbb{R}^m,$$
 (1)

where **x** is the state of the system, **u** is the input,  $\mathbf{f}: \mathbb{R}^n \to \mathbb{R}^n$  is the drift dynamics, and  $\mathbf{g}: \mathbb{R}^n \to \mathbb{R}^{n \times m}$  is the input matrix. We assume that **f** and **g** are locally Lipschitz continuous. Given a locally Lipschitz continuous statefeedback controller  $\mathbf{k}: \mathbb{R}^n \to \mathbb{R}^m$ , the closed-loop dynamics are governed by:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}). \tag{2}$$

For any initial condition  $\mathbf{x}(0) = \mathbf{x}_0 \in \mathbb{R}^n$  there exists a unique solution  $\mathbf{x}(t)$  to (2), which we assume to exist  $\forall t \in [0, \infty)$ .

The notion of safety is formalized by defining a *safe set*  $C \subset \mathbb{R}^n$  in the state space that the system must remain within. In particular, consider the set C as the 0-superlevel set of a continuously differentiable function  $h : \mathbb{R}^n \to \mathbb{R}$ :

$$\mathcal{C} \triangleq \{ \mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \ge 0 \},\tag{3}$$

where  $h(\mathbf{x}) = 0 \implies \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}) \neq 0$  and C is non-empty and has no isolated points. Safety is defined as the *forward invariance* of C, i.e., if  $\mathbf{x}_0 \in C$ , then  $\mathbf{x}(t) \in C$  for all  $t \geq 0$ .

To synthesize controllers that ensure safety, we use Control Barrier Functions (CBFs) [2] defined as follows:

**Definition 1** (*Control Barrier Function (CBF)*). Let  $C \subset \mathbb{R}^n$  be a safe set given by (3). The function *h* is a *Control Barrier Function* (CBF) for (1) on C if there exists  $\gamma \in \mathcal{K}_{\infty,e}^{-1}$  such that for all  $\mathbf{x} \in C$ :

$$\sup_{\mathbf{u}\in\mathbb{R}^{m}}\dot{h}(\mathbf{x},\mathbf{u})\triangleq\underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{f}(\mathbf{x})}_{L_{\mathbf{f}}h(\mathbf{x})}+\underbrace{\frac{\partial h}{\partial \mathbf{x}}(\mathbf{x})\mathbf{g}(\mathbf{x})}_{L_{\mathbf{g}}h(\mathbf{x})}\mathbf{u}\geq-\gamma(h(\mathbf{x})),$$
(4)

where  $L_{\mathbf{f}}h : \mathbb{R}^n \to \mathbb{R}$  and  $L_{\mathbf{g}}h : \mathbb{R}^n \to \mathbb{R}^m$  are the Lie derivatives of h with respect to  $\mathbf{f}$  and  $\mathbf{g}$  respectively.

A main result in [18], [20] relates CBFs to the safety of the closed-loop system (2) with respect to C:

**Theorem 1.** Given a safe set  $C \subset \mathbb{R}^n$ , if h is a CBF for (1) on C, then any locally Lipschitz continuous controller  $\mathbf{k} : \mathbb{R}^n \to \mathbb{R}^m$  satisfying

$$L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{k}(\mathbf{x}) \ge -\gamma(h(\mathbf{x}))$$
(5)

for all  $\mathbf{x} \in C$ , renders the system (2) safe w.r.t. C.

Given a nominal (but not necessarily safe) locally Lipschitz continuous controller  $\mathbf{k}_d : \mathbb{R}^n \to \mathbb{R}^m$  and a CBF h, the CBF-Quadratic Program (CBF-QP) [2] is a controller that guarantees the system's safety:

$$\mathbf{k}(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m}{\operatorname{argmin}} \quad \frac{1}{2} \|\mathbf{u} - \mathbf{k}_d(\mathbf{x})\|_2^2 \quad (\text{CBF-QP})$$
  
s.t.  $L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} \ge -\gamma(h(\mathbf{x})).$ 

#### B. Control Barrier Functions for Safe Vision-Based Control

Next we apply CBFs to achieve safe obstacle avoidance for robotic systems based on stereo vision. First we construct CBFs for safe vision-based control. Let  $\rho_p \in \mathbb{R}^3$  represent the true three-dimensional position of the portion of the scene which generated pixel p. Using this, we can define a CBF  $h: \mathbb{R}^n \times \mathbb{R}^3 \to \mathbb{R}$  that relies on both the state  $\mathbf{x}$  and three dimensional pixel position  $\rho_p$ . The pixel position is a geometric function of the true disparity,  $\rho_p \triangleq T(\mathbf{x}, r(p, d_{1,3}^p))$ where  $r: \mathbb{N}^2 \times \mathbb{N}$  is the stereo reprojection function and  $T: \mathbb{R}^n \times \mathbb{R}^3 \to \mathbb{R}^3$  is the transformation mapping from the robot's state and relative pixel position to pixel position.

In order to relate the output of the stereoscopic sensor with safety, we make the following assumptions:

**Assumption 1.** The environment is static, so the time derivative of the pixelized environment is zero:  $\frac{d\rho_p}{dt} = 0$  for all  $p \in K$ .

**Assumption 2.** Ensuring safety with respect to the true three dimensional pixel locations is sufficient to ensure safety with respect to the environment. That is, the safe set for the system is given by:

$$\mathcal{C}_{K} = \{ \mathbf{x} \in \mathbb{R}^{n} \mid h(\mathbf{x}, \rho_{p}) \ge 0, \ \forall p \in K \}$$
(6)

 ${}^{l}\mathcal{K}_{\infty,e}$  denotes the set of extended class- $\mathcal{K}$  infinity functions, wherein  $\gamma \in \mathcal{K}_{\infty,e}$  satisfies  $\gamma : \mathbb{R} \to \mathbb{R}$  is strictly increasing,  $\gamma(0) = 0$ , and  $\lim_{r \to \infty} \gamma(r) = \infty$ ,  $\lim_{r \to -\infty} \gamma(r) = -\infty$ .

where  $h : \mathbb{R}^n \times \mathbb{R}^3 \to \mathbb{R}$  is the CBF for pixel p.

Although it is not outlined in this work, Assumption 1 can be relaxed to include moving environments by calculating  $\dot{h}$  accordingly and estimating the motion of obstacles in the environment. Assumption 2 simplifies the surrounding environment from infinite- to finite-dimensional by assuming that the environment is smooth between a sufficiently dense coverage of pixels. It also implies that the system only has to stay safe with respect to objects that can be seen in the cameras' field of view.<sup>2</sup>

Based on Assumptions 1 and 2 and Theorem 1, synthesizing the control input **u** such that

$$L_{\mathbf{f}}h(\mathbf{x},\rho_p) + L_{\mathbf{g}}h(\mathbf{x},\rho_p)\mathbf{u} \ge -\gamma(h(\mathbf{x},\rho_p)), \qquad (7)$$

 $\forall p \in K$ , is sufficient to guarantee safety. Considering each pixel  $p \in K$ , however, may be computationally intractible, therefore we seek a condition with fewer required constraints.

To combine the constraints, we apply Boolean composition to each CBF h to produce a single nonsmooth CBF  $h_{ns}$ ,

$$h_{\rm ns}(\mathbf{x}) \triangleq \min_{p \in K} h(\mathbf{x}, \rho_p),\tag{8}$$

and simply enforce the CBF constraint associated with the pixels whose CBFs have the smallest value [21]. In particular, to achieve safety it is sufficient to enforce only the constraints whose indeces appear in the locally-encapsulating index set:

$$\Lambda = \{ p \in K : h(\mathbf{x}, \rho_p) \le h_{\rm ns}(\mathbf{x}) + \delta \},\tag{9}$$

for some  $\delta > 0$ , as stated formally below.

**Theorem 2** ([21], Prop III.6). Let  $h : \mathbb{R}^n \times \mathbb{R}^3 \to \mathbb{R}^3$  be a locally Lipschitz function and  $h_{ns}$  be as in (8). If there exists a locally Lipschitz extended class  $\mathcal{K}$  function  $\gamma$  and a measurable and locally bounded controller  $\mathbf{k} : \mathbb{R}^n \to \mathbb{R}^m$ that satisfies:

$$\min_{p \in \Lambda} \left\{ L_{\mathbf{f}} h(\mathbf{x}, \rho_p) + L_{\mathbf{g}} h(\mathbf{x}, \rho_p) \mathbf{k}(\mathbf{x}) \right\} \ge -\gamma(h_{\text{ns}}(\mathbf{x})).$$
(10)

Then  $h_{ns}$  is a valid nonsmooth CBF and the closed loop dynamics (2) with controller **k** are safe with respect to  $C_K$ .

This theorem indicates that enforcing the CBF condition only for the "least safe" pixel is sufficient to guarantee the safety of the system.

#### C. Robustness to Uncertainty

Error in the disparity propagates to the controller in the form of the measured 3D pixel position  $\hat{\rho}_p$ . The measured value  $\hat{\rho}_p$  lies in a neighborhood  $\mathcal{E}_p$  of the true value  $\rho_p$ , which is characterized by the error distribution  $\mathcal{P}(I_i, I_j)$ . We assume that the distribution  $\mathcal{P}(I_i, I_j)$  is symmetric about the measured value and define the pixelwise uncertainty set:

$$\mathcal{E}_p \triangleq \left\{ \rho \in \mathbb{R}^3 \mid \begin{array}{c} \rho = T(\mathbf{x}, r(p, \xi)), \quad \xi \in \Gamma \\ \mathcal{P}^{\theta}(e_{1,3}(p) < |\xi - \hat{d}(p)|; I_1, I_3) \ge \sigma \end{array} \right\}$$
(11)

<sup>2</sup>The field of view aspect of Assumption 2 can be overcome by tracking features that leave the frame as done in Simultaneous Localization and Mapping (SLAM) algorithms [4].

where  $\sigma > 0$  is a parameter defining the desired uncertainty robustness.

To achieve safety, one must determine which pixels are safety-critical given  $\mathcal{E}_p$  and then enforce robust safety with respect to those pixels. The safety-critical pixels can be determined by expanding the index set  $\Lambda$  using the uncertainty:

$$\Lambda \subseteq \left\{ p \in K \left| h(\mathbf{x}, \rho_p) \le \max_{\rho_p \in \mathcal{E}_p} \min_{p \in K} h(\mathbf{x}, \rho_p) + \delta \right\}.$$
(12)

This can further be expanded to an easily calculable index set  $\widehat{\Lambda} \supseteq \Lambda$  by minimizing the left-hand-side of the inequality condition and using the max-min inequality [22]:

$$\widehat{\Lambda} = \left\{ p \in K \middle| \min_{\rho_p \in \mathcal{E}_p} h(\mathbf{x}, \rho_p) \le \min_{p \in K} \max_{\rho_p \in \mathcal{E}_p} h(\mathbf{x}, \rho_p) + \delta \right\}.$$
(13)

This expanded uncertain index set  $\widehat{\Lambda}$  indicates which pixels are safety critical and whose constraints must be enforced to achieve safety given the pixelwise uncertainty sets  $\mathcal{E}_p$ .

Measurement-Robust Control Barrier Functions (MR-CBFs) as outlined in [23] are a general method for accounting for state uncertainty in CBFs. We can use this for each pixel  $p \in \widehat{\Lambda}$  to ensure that the safety constraint is satisfied despite the uncertainty. The resulting constraint is:

$$L_{\mathbf{f}}h(\mathbf{x},\widehat{\rho}_{p}) + L_{\mathbf{g}}h(\mathbf{x},\widehat{\rho}_{p})\mathbf{u} - \left(\mathfrak{L}_{L_{f}h} + \mathfrak{L}_{\gamma\circ h_{ns}} + \mathfrak{L}_{L_{g}h} \|\mathbf{u}\|_{2}\right)\epsilon_{p} \geq -\gamma(h_{ns}(\mathbf{x})), \quad \forall p\in\widehat{\Lambda} \quad (14)$$

where  $\mathfrak{L}$  is the Lipschitz constant of the subscript and

$$\epsilon_p \ge \max_{\rho_p \in \mathcal{E}_p} \|\rho_p - \widehat{\rho}_p\|_2 \tag{15}$$

is a bound on the uncertainty. Since  $\Lambda \subseteq \widehat{\Lambda}$  and the MR-CBF condition implies the CBF condition (5), satisfying (14) also satisfies (10) providing safety of the system if  $\sigma = 1$  and  $\mathcal{P}^{\theta} = \mathcal{P}$ .

## IV. APPLICATION: OBSTACLE AVOIDANCE ON A QUADRUPEDAL ROBOT

We evaluate our approach on a quadrupedal robotic platform. With these experiments we aim to demonstrate: 1) Our method is capable of keeping the system safe in a simple do-not-collide task, and 2) Our method can adapt online to measurement uncertainty in different environments without ground-truth data.

#### A. Hardware System

For the hardware experiments we designed a custom camera array with three equally spaced inexpensive CMOS, global shutter, time-synchronized Arducam cameras. An Nvidia Jetson Nano is used to capture, downsize, and greyscale the stereo images. The images are then sent to an external computer that receives the images and outputs the filtered control input at a frequency of at least 10 Hz. The robot used in this experiment is a Unitree A1 quadrupedal robot that receives inputs of velocity and angle rate,  $\mathbf{u} = \begin{bmatrix} v & \omega \end{bmatrix}^{\mathsf{T}}$ . A 1 kHz Inverse Dynamics Quadratic Program (ID-QP) walking controller designed using the concepts in [24], is used to track these inputs. Stereo pixelmatching calculations were performed using Efficient LArgescale Stereo (ELAS) [15].

#### B. Learning Method and Model

The architecture of the model used to estimate  $\mathcal{P}^{\theta}$  is a modified version of the Hierarchical Multi-Scale Attention for Semantic Segmentation introduced in [25]; this model is relatively lightweight, consisting of only 196 thousand parameters (e.g., network weights). The robustness threshold used was  $\sigma = 0.99$  and the online learning rate was 0.001. We pretrain the model until convergence on a dataset of 6000 stereo image triples collected by manually walking the robot in a variety of environments.

## C. Dynamics Model and Control

In the context of safety we consider a reduced order model of the system dynamics given by the standard unicycle model. The specific form of (2) for this system is:

$$\underbrace{\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix}}_{\dot{\mathbf{x}}} = \underbrace{\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} \cos \theta & 0 \\ \sin \theta & 0 \\ 0 & 1 \end{bmatrix}}_{\mathbf{g}(\mathbf{x})} \underbrace{\begin{bmatrix} v \\ \omega \end{bmatrix}}_{\mathbf{k}(\mathbf{x})}$$
(16)

A formal analysis of CBFs which utilize reduced-order velocity input models is described in [26].

For this system we consider the pixelwise CBFs,

$$h(\mathbf{x}, \rho_p) = \frac{1}{2} \left( \left\| \begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} \rho_{p,x} \\ \rho_{p,y} \end{bmatrix} \right\|_2^2 - c^2 \right)$$
(17)

where  $\rho_{p,x}$  and  $\rho_{p,y}$  indicate the global real-world x and y positions of pixel p. This function characterizes safety as remaining a planar distance c > 0 from  $\rho_p$ . This can be thought of as buffering surfaces in the environment by a radius c.

#### D. Robustness to Uncertainty

To illustrate the efficacy of our method we use two controllers in our experiments. A standard, unrobustified controller:

$$\mathbf{k}_{cbf} = \underset{\mathbf{u} \in \mathbb{R}^{2}}{\operatorname{argmin}} \quad \frac{1}{2} \| \mathbf{k}_{des}(\mathbf{x}) - \mathbf{u} \|_{2}^{2}$$
(18)  
s.t. 
$$\underbrace{-\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^{\mathsf{T}} r(p, \widehat{d}_{p}) v}_{\hat{h}} \ge -\gamma(\underset{p \in K}{\min} h(\mathbf{x}, \widehat{\rho_{p}})), \forall p \in \Lambda$$

and a robustified controller:

$$\mathbf{k}_{\mathsf{cbf}}^* = \underset{\mathbf{u} \in \mathbb{R}^2}{\operatorname{argmin}} \quad \frac{1}{2} \| \mathbf{k}_{\mathsf{des}}(\mathbf{x}) - \mathbf{u} \|_2^2 \tag{19}$$
$$\text{s.t.} \quad -v \ge \frac{-\gamma(\min_{p \in K} h(\mathbf{x}, \rho_p^*))}{\left[1 \quad 0 \quad 0\right]^{\mathsf{T}} r(p, d_p^*)}, \quad \forall p \in \widehat{\Lambda}.$$

where  $\mathbf{k}_{\text{des}} : \mathbb{R}^m \to \mathbb{R}^n$  is a desired controller,  $d^*$  is the maximum disparity for any  $\rho_p \in \mathcal{E}_p$ , and  $\rho_p^*$  is pixel location associated with  $d_p^*$ .



Fig. 4. Demonstration of our method in a variety of environments. From left to right the goal is to maintain a safe distance from (A) a tree, (B) a backpack, (C) a blue box, and (D) a glass window. The distance to the barrier is measured and marked on the floor with a yellow tape for visualization purposes – we emphasize this tape is not used for depth estimation. Notice that the barrier is assumed to be a sphere around an obstacle but in the case on the glass, this sphere degenerates into a plane. The quadrupedal robot is given a desired control input of 0.2 m/s. In all cases, a naive barrier implementation that simply takes the noisy measurements from a stereo vision system fails to keep the system safe. The robustified controller (19) with a pretrained model consistently shows overly conservative behavior. Finally, with online learning, the robot converges to the barrier due to the perceived uncertainty. The (A-D) corresponding plots below show the control input filtered by the barrier in each of the three robustification cases.

Although the MR-CBF-based condition (14) is a general method for providing robustness, we opt for a simpler condition that can be found using the system dynamics (16), the chosen CBF (17), and Assumption 2. In particular these imply that the denominator of (19) is positive and the right side of the inequality is decreasing for  $d_p^* \in \Gamma$  so remaining safe with respect to the worst case closest pixel location would ensure safety with respect to the true pixel location.

## E. Experimental Results

The system was run in 4 different environments (see Fig. 4). The CBF (17) was used with a safe radius of c = 0.33 m. The intended obstacle in the 4 different environments were (**A**) a tree, (**B**) a backpack, (**C**) a blue box, (**D**) and a glass window. A desired constant forward velocity v = 0.2 m/s was used in each experiment and the robot was started approximately 1.3 m away from the obstacle. Since ground-truth measurements were unavailable, we use a yellow line on the ground to indicate the true location of the barrier.

For each environment three different tests were performed. First, controller (18) was used. Since this did not consider measurement uncertainty it failed to achieve safety in every environment; in all experiments the stereo vision overestimated the distance to objects at some point during the run and the quadruped ran directly into the obstacles. Second, the controller (19) was used with an error estimate computed through a pretrained function  $\mathcal{P}^{\theta}$ ; this succeeded in providing safety, but was found to be overly conservative and did not allow the quadruped to approach the obstacle as desired. Third, the controller (19) was used with a  $\mathcal{P}^{\theta}$  that adapted to the environment according to Algorithm 2. In this case, safety of the system was generally maintained and over time the system was able to approach the boundary of the safe set. Even when small safety violations occurred, the system eventually corrected and came to rest at a safe steadystate. These results can be seen in Figure 4. A video of the experiments can be found at [27].

## V. CONCLUSION AND FUTURE WORK

We presented a framework for achieving safety of a stereo vision-based system using self-supervised online uncertainty estimation and robustified CBFs. Refining the uncertainty estimate model online was shown to achieve significantly better performance. We validated our online learning approach across several environments and successfully achieved robust safety with minimal violations and conservatism.

Future work involves providing mathematical guarantees for our method and extending this theory and application to dynamic environments. We also note that our online uncertainty estimation method, as outlined in Algorithm 2, is a general method that can be coupled with other control or state estimation techniques. Finally, implementing a highperformance on-board version of our hardware system would remove the need for a remote computer.

#### REFERENCES

- [1] J. C. Knight, "Safety Critical Systems: Challenges and Directions,"
- 9. C. Kingin, Surely Chicki Systems Linear C. A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control Barrier Function Based Quadratic Programs for Safety Critical Systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, 2027 (Automatic Control, vol. 62, no. 8, 2027). [2] pp. 3861–3876, Aug. 2017. [Online]. Available: http://ieeexplore.ieee. org/document/7782377/
- [3] Y. Emam, P. Glotfelter, and M. Egerstedt, "Robust barrier functions for a fully autonomous, remotely accessible swarm-robotics testbed," in 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 3984–3990.
- [4] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age," *IEEE*
- *Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, Dec. 2016. S. Dean, N. Matni, B. Recht, and V. Ye, "Robust Guarantees for Perception-Based Control," *arXiv:1907.03680 [cs, math, stat]*, Dec. 2016. [5] 1907.03680
- 1907.03680
  [6] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, "Learning for Safety-Critical Control with Control Barrier Functions," p. 10.
  [7] N. Csomay-Shanklin, R. K. Cosner, M. Dai, A. J. Taylor, and A. D. Ames, "Episodic Learning for Safe Bipedal Locomotion with Control Barrier Functions and Projection-to-State Safety," arXiv:2105.01697 [cs], May 2021, arXiv: 2105.01697. [Online]. Available: http://arxiv.org/abs/2105.01697
  [8] M. Ohnishi, L. Wang, G. Notomista, and M. Egerstedt, "Barrier-certified adaptive reinforcement learning with applications to brushbot navigation," *IEEE Transactions on Robotics*, vol. 35, no. 5, pp. 1186–1205, 2019.
  [9] J. Choi, E. Castaneda, C. J. Tomlin, and K. Sreenath, "Beinforcement
- J. Choi, F. Castaneda, C. J. Tomlin, and K. Sreenath, "Reinforcement [9] learning for safety-critical control under model uncertainty, using
- [10] E. Hazan, "Introduction to online control barrier functions," arXiv preprint arXiv:2004.07584, 2020.
  [10] E. Hazan, "Introduction to online convex optimization," 2019.
  [11] J. H. Gillula and C. J. Tomlin, "Guaranteed safe online learning via reachability: tracking a ground target using a quadrotor," in 2012 IEEE International Conference on Robotics and Automation, 2012, pp. 2723-2730.
- [12] B. Sofman, E. Lin, J. A. Bagnell, J. Cole, N. Vandapel, and A. Stentz, "Improving robot navigation through self-supervised online learning," *Journal of Field Robotics*, vol. 23, no. 11-12, pp. 1059–1075, 2006.
  [13] A. Koppel, J. Fink, G. Warnell, E. Stump, and A. Ribeiro, "Online learning for characterizing unknown environments in ground robotic vehicle models," in 2016 IEEE/RSJ International Conference on University Packar and Contemp (IPOS) 2016, pp. 626–633. Intelligent Robots and Systems (IROS), 2016, pp. 626-633.

- [14] M. Perrollaz, A. Spalanzani, and D. Aubert, "Probabilistic representation of the uncertainty of stereo-vision and application to obstacle detection," in 2010 IEEE Intelligent Vehicles Symposium. IEEE, 2010, pp. 313–318.
- [15] A. Geiger, M. Roser, and R. Urtasun, "Efficient large-scale stereo matching," in Asian conference on computer vision. Springer, 2010, pp. 25–38.
- [16] R. Szeliski and D. Scharstein, "Symmetric sub-pixel stereo matching," in European Conference on Computer Vision. Springer, 2002, pp. 525-540
- K. Liu, K. Ok, W. Vega-Brown, and N. Roy, "Deep Inference for Covariance Estimation: Learning Gaussian Noise Models for State Estimation," in 2018 IEEE International Conference on Robotics [17] and Automation (ICRA). Brisbane, QLD: IEEE, May 2018, pp. 1436–1443. [Online]. Available: https://ieeexplore.ieee.org/document/ 8461047/
- [18] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control,'
- based quadratic programs with application to adaptive cruise control, in 53rd IEEE Conference on Decision and Control. Los Angeles, CA, USA: IEEE, Dec. 2014, pp. 6271–6278. [Online]. Available: http://ieeexplore.ieee.org/document/7040372/
  R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, and A. D. Ames, "Measurement-Robust Control Barrier Functions: Certainty in Safety with Uncertainty in State," arXiv:2104.14030 [cs, eess], Apr. 2021, arXiv: 2104.14030. [Online]. Available: http://urxiv.gev/aba/10020 [19]
- arXiv:2104.14050 [cs, eess], Apr. 2021, arXiv: 2104.14030. [Online]. Available: http://arXiv.org/abs/2104.14030
  [20] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Robustness of Control Barrier Functions for Safety Critical Control," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015, arXiv: 1612.01554.
  [20] P. Glotfelter, J. Cortes, and M. Egerstedt, "A Nonsmooth Approach to Control Restrong Serversion Sciences Serversed Test Control to Control Restrong Serversions." *IEEE Transactions*
- Controller Synthesis for Boolean Specifications," IEEE Transactions
- S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004. [22]
- S. Dean, A. J. Taylor, R. K. Cosner, B. Recht, and A. D. Ames, "Guaranteeing Safety of Learned Perception Modules via Measurement-Robust Control Barrier Functions," *arXiv:2010.16001* [23] [cs, eess, math, stat], Oct. 2020, arXiv: 2010.16001 [value]. Available: http://arxiv.org/abs/2010.16001 J. Buchli, M. Kalakrishnan, M. Mistry, P. Pastor, and S. Schaal, "Compliant quadruped locomotion over rough terrain," in *IEEE/RSJ* Interesting Conference on Local Construct Path
- [24] International Conference on Intelligent Robots and Systems, 2009, pp. 814 - 820
- [25] A. Tao, K. Sapra, and B. Catanzaro, "Hierarchical multi-scale attention for semantic segmentation," 2020.
- T. G. Molnar, R. K. Cosner, A. W. Singletary, W. Ubellacker, and A. D. Ames, "Model-free safety-critical control for robotic systems," *IEEE Robotics and Automation Letters*, 2021, submitted. [Online]. Available: http://ames.caltech.edu/2021.ReducedOrderModels.pdf Supplementary\_uidea\_https://wines.came/forSuppl272 [26]
- [27] Supplementary video, https://vimeo.com/605281037.