

# W-RIZZ: A Weakly-Supervised Framework for Relative Traversability Estimation in Mobile Robotics

Andre Schreiber<sup>1</sup>, Arun N. Sivakumar<sup>1</sup>, Peter Du<sup>1</sup>, Mateus V. Gasparino<sup>1</sup>,  
Girish Chowdhary<sup>1</sup>, *Senior Member, IEEE*, and Katherine Driggs-Campbell<sup>1</sup>, *Member, IEEE*

**Abstract**—Successful deployment of mobile robots in unstructured domains requires an understanding of the environment and terrain to avoid hazardous areas, getting stuck, and colliding with obstacles. Traversability estimation—which predicts where in the environment a robot can travel—is one prominent approach that tackles this problem. Existing geometric methods may ignore important semantic considerations, while semantic segmentation approaches involve a tedious labeling process. Recent self-supervised methods reduce labeling tedium, but require additional data or models and tend to struggle to explicitly label untraversable areas. To address these limitations, we introduce a weakly-supervised method for relative traversability estimation. Our method involves manually annotating the relative traversability of a small number of point pairs, which significantly reduces labeling effort compared to traditional segmentation-based methods and avoids the limitations of self-supervised methods. We further improve the performance of our method through a novel cross-image labeling strategy and loss function. We demonstrate the viability and performance of our method through deployment on a mobile robot in outdoor environments.

**Index Terms**—Field robots, deep learning for visual perception, vision-based navigation.

## I. INTRODUCTION

WITH the rapidly expanding use of mobile robots in areas like agriculture [1] and delivery [2], it is critical that such robots are able to effectively understand their environments in order to navigate and successfully complete their tasks. Traversability estimation, which involves determining where in the environment a robot can travel, is one common method for solving this problem in unstructured environments [3]. Due to

Manuscript received 23 January 2024; accepted 14 April 2024. Date of publication 1 May 2024; date of current version 8 May 2024. This letter was recommended for publication by Associate Editor R. Boutheau and Editor P. Vasseur upon evaluation of the reviewers' comments. This work was supported in part by National Robotics Initiative 2.0 under Grant NIFA#2021-67021-33449, in part by AIFARMS through the Agriculture and Food Research Initiative (AFRI) under Grant 2020-67021-32799 Project Accession Number 1024178 from the USDA/NIFA, and in part by the UIUC's Center for Autonomous Construction in Manufacturing at Scale. (*Corresponding author: Andre Schreiber.*)

The authors are with the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Champaign, IL 61820 USA (e-mail: andrems2@illinois.edu; av7@illinois.edu; peterdu2@illinois.edu; mvalve2@illinois.edu; girishc@illinois.edu; krde@illinois.edu).

Code is available at: <https://github.com/andreschreiber/W-RIZZ>.

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2024.3396095>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2024.3396095

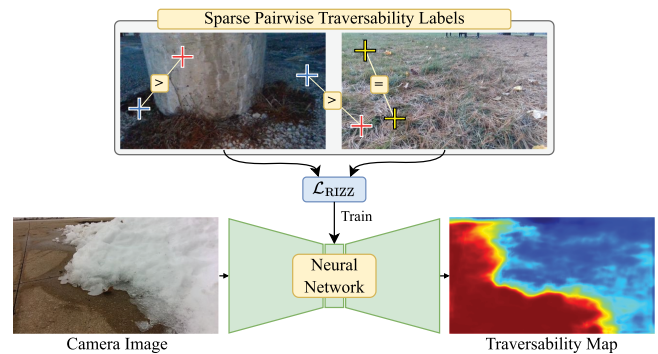


Fig. 1. Our proposed framework trains a model for traversability estimation using sparse pairwise labels of relative traversability. The camera images are from the dataset introduced by Gasparino et al. [4]. In the traversability map, warmer colors indicate higher anticipated ease of traversal.

advances in computer vision and visual perception systems in recent years, vision-based methods for traversability estimation have shown significant promise [4], [5], [6], [7], [8].

Vision-based methods for environmental understanding and traversability estimation can typically be classified into strongly-supervised methods based on semantic segmentation and self-supervised methods. Strongly-supervised semantic segmentation approaches [8], [9], [10], [11], [12], [13] involve using a semantic segmentation model to segment the environment into classes relevant to navigation and traversability. However, these methods suffer from a tedious and expensive manual labeling process as training images are densely annotated by a human labeler. Self-supervised approaches avoid this labeling burden by annotating data automatically [4], [5], [6], [7], but typically involve additional assumptions, data, or models. In addition, these self-supervised approaches generally require that the robot experience conditions of interest to label them, which creates difficulty in explicitly labeling untraversable conditions in a safe manner [6], [7].

Motivated by the limitations of these existing approaches, we propose a weakly-supervised framework for visual traversability estimation, termed Weak-Relative Inference of haZard Zones (W-RIZZ). We draw inspiration from work in relative depth estimation [14], [15], [16] and intrinsic decomposition [15], [17], and propose annotating sparse pairs of points with relative traversability labels (as shown in Fig. 1). The sparse labeling greatly reduces the tedium seen in strongly-supervised methods, while enabling easy explicit labeling of untraversable regions

and not requiring robot experience during training. Our method can also provide continuous predictions of traversability (rather than discrete classes [8]), and does not require the additional assumptions, data, or models needed in self-supervised methods [4], [5], [6], [7].

We summarize our contributions as follows:

- 1) We introduce a weakly-supervised framework for training traversability estimation neural networks that employs sparse pairwise labels of relative traversability.
- 2) We propose a new loss for training traversability estimation models and a novel cross-image pairwise labeling scheme that improves consistency of traversability predictions across different images and environments.
- 3) We deploy our approach on a robot and demonstrate that our model outperforms a state-of-the-art self-supervised method for visual traversability prediction, showing greater generalization capability and improved navigation success rate while only using a small number of annotated pixels per image.

## II. RELATED WORKS

Numerous potential approaches to mobile robot navigation and traversability estimation have been proposed and are relevant to the work presented in this letter.

### A. Geometric and Semantic Approaches to Navigation

A classical approach to mobile robot navigation involves the creation of an environmental map using sensors like LiDAR, localizing the robot in the map using SLAM, and planning an obstacle-free path through the map [18], [19]. However, only considering geometry in navigation can be problematic since it fails to consider non-rigid and traversable obstacles like tall grass and may ignore characteristics like the bumpiness of a traversed surface.

These limitations can be partially resolved by integrating semantic information about the environment. For example, Valada et al. [9] introduce a robust segmentation architecture that uses a convoluted mixture of deep experts model. Maturana et al. [11] utilize geometric and semantic information, where environment semantics are predicted using a neural network and semantic classes are assigned navigation costs.

Several datasets exist for semantic segmentation in outdoor unstructured environments. For example, the Robot Unstructured Ground Driving (RUGD) dataset [12] and Freiburg Forest dataset [10] provide images that are densely labeled with semantic classes relevant for outdoor navigation. The CaT: CAVS [8] dataset features images collected from an off-road car, and directly labels regions as traversable by a specific vehicle platform. However, these datasets only contain a few hundred [10] to a few thousand [8], [12] images, which is significantly smaller than traditional datasets for segmentation in computer vision [20].

Methods based on semantic segmentation can be appealing as they can provide more fine-grained understanding than purely geometric approaches. Moreover, numerous architectures for semantic segmentation exist (e.g., PSPNet [21]) and can be implemented for the task of semantic segmentation in unstructured environments. However, densely labeling datasets for segmentation is tedious, making it difficult to apply this method in a new domain. In addition, choosing an appropriate set of semantic classes may be challenging and can vary for different applications. It may also be difficult to map these semantic

classes into costs that can be used for control and motion planning [22]. Segmenting images directly for traversability can also be problematic, as traversability segmentation labels typically represent only binary indications of traversability and may only apply to a specific robotic platform [8]. Finally, each of these forms of semantic segmentation fails to consider the possibility of traversability characteristics varying within semantic classes.

### B. Self-Supervised Navigation and Traversability

To overcome the limitations of geometric and segmentation approaches to mobile robot navigation, much recent work has focused on developing self-supervised methods for navigation and traversability estimation. BADGR [23] collects and labels off-policy data using a simple navigation policy, and trains an action-conditioned neural network to predict navigation events. However, BADGR requires a large dataset to perform well, as it learns both the navigation characteristics of the environment and dynamics of the robot.

Other methods [4], [24] overcome this limitation by directly incorporating a dynamics model. For example, WayFAST [4] uses the kinodynamic model of a skid-steer robot to measure traction coefficients. These traction coefficients are then projected into camera images and used to train a traversability prediction neural network. For quadrupeds, Wellhausen et al. [5] propose projecting footholds into images and retroactively assign labels using proprioceptive measurements. Similarly, Castro et al. [22] combine proprioceptive sensing of terrain interactions with exteroceptive sensor data to learn costmaps for off-road driving. Other proposed methods involve labeling traversed paths as traversable and training a classifier on such paths [6], or training an autoencoder only on traversed regions and using reconstruction error as a proxy for traversability [7].

These self-supervised methods are appealing by eliminating the burden of manual labeling through self-supervision. However, such methods are typically only weakly-supervised (they do not label all training image pixels) and can be difficult to apply on existing image-only datasets as they often require additional data sources. The need for the robot to experience conditions to label them also creates difficulty in labeling non-traversable areas, since subjecting the robot to such conditions may not be possible or can lead to damage. As a result, these methods may assume unfamiliar regions as non-traversable [7] and are often biased towards a positive and unlabeled (PU) [6], [25] data regime.

### C. Weakly-Supervised Learning

Learning with sparse labels (weakly-supervised learning) is also heavily related to our proposed method. Various works have proposed methods for reducing the labeling burden in semantic segmentation with approaches using image-level labels [26] or using annotations such as points [27] or squiggles [28]. For mobile robot navigation in particular, Gao et al. [29] introduce a contrastive learning approach for traversability estimation, and Schreiber [30] presents a method for weakly-supervised traversability segmentation using point annotation. Methods have also been introduced for boosting the accuracy of weakly-supervised models [31], [32] and have successfully been applied in field robotics [5], [30].

Beyond semantic segmentation, weakly-supervised learning has been utilized with pairwise labels for tasks like intrinsic image decomposition [15], [17] and depth estimation [14], [15],

[16]. In these methods, relative annotations for pairs of points are provided and are used with ranking-based losses to produce continuous outputs (in contrast to the discrete class outputs seen in segmentation). We draw inspiration from such pairwise labeling schemes and adapt this idea to weakly-supervised traversability prediction.

### III. METHOD

We propose a framework (W-RIZZ) for weakly-supervised relative traversability estimation. The proposed framework involves manually labeling a small number of point pairs in each training image, where one point in a pair is labeled as more, less, or equally traversable as compared with the other point in the pair. By using sparse annotations, the labeling burden compared to a strongly-supervised approach is greatly reduced, with the time to annotate an image in W-RIZZ taking seconds rather than minutes. The use of sparse manual annotations (rather than self-supervision) means that our method does not require any additional data, models, or assumptions and can be used on existing image-only datasets. Finally, the lack of self-supervision in W-RIZZ means that untraversable regions can be labeled explicitly without risking damage to the robot or its environment.

#### A. Data Annotation

Inspired by relative annotation strategies for tasks like depth estimation [14], [15], [16] and intrinsic decomposition [15], [17], we label images using a small number of pairwise annotations of sparse points. Specifically, given an image to be annotated, a labeler is prompted with a pair of points and asked to indicate which point in the pair is more traversable (or if the points are equally traversable). The locations of the points in a pair are selected randomly in order to capture the natural traversability statistics of images in the dataset and since random sampling has been shown to be particularly effective in weakly-supervised semantic segmentation [27].

If relative traversability labels are only provided for pairs where both pixels belong to the same image, the scale of the traversability predictions across different images may vary. This can lead to inconsistent predictions and may require tweaks to the controller that uses the traversability predictions. Thus, unlike the relative pairwise labeling strategies described in other works (which only label relationships within the same image) [14], [15], [16], [17], we propose to also label pairs where each point belongs to a different image.

For each image in a dataset, we label the relative traversability of one pair of pixels within the image (an ‘‘intra-image’’ label), and we also provide a ‘‘cross-image’’ label (where one pixel belongs to the image and the other pixel in the pair belongs to a different image). The pairings for cross-image labels are chosen randomly, and one cross-image label is provided per image. The annotation strategy leads to 3 pairwise labels for every 2 images (1 intra-image pair label for each image and 1 cross-image pair label). For intra-image labels, we limit selected pixels for an annotation pair to be at least 5% of  $\min(W, H)$  apart (where  $W$  and  $H$  are the width and height of the image) to avoid labeling pairs where the pixels represent nearly the same location. Fig. 2 illustrates this labeling strategy on example images from the dataset introduced by Gasparino et al. [4]. Labeling the



Fig. 2. Illustration of our annotation strategy on the WayFAST dataset [4], showing intra-image labeling (top) and cross-image labeling (bottom). Yellow crosshairs indicate both points in a pair are equally traversable, while the blue point is more traversable in pairs with blue and red crosshairs.

16558 images from this dataset with our annotation method took approximately 3 standard working days and produced 24837 pairwise annotations.

#### B. Loss Function

We train a model in our proposed framework using ordinal annotations of relative traversability. While various works describe methods for learning vision tasks using relative relationships [14], [15], [16], [17], the method of Chen et al. [14] is most similar to our desired application, presenting a ranking loss specifically for end-to-end training of depth estimation neural networks using sparse pairwise annotations. The loss introduced by Chen et al. [14] can be described as follows:

$$\mathcal{L}_{\text{DIW}}(p_a, p_b, t) = \begin{cases} \ln(1 + \exp(-t(p_b - p_a))) & t \neq 0 \\ (p_b - p_a)^2 & t = 0 \end{cases} \quad (1)$$

where  $p_a$  is the prediction value of the first point in the pair (point  $a$ ),  $p_b$  is the prediction value of the second point in the pair (point  $b$ ), and  $t \in \{-1, 0, 1\}$  is the ordinal annotation of the point pair (with  $t = 0$  being an equality relation,  $t = 1$  meaning  $p_b$  should be greater than  $p_a$ , and  $t = -1$  meaning  $p_b$  should be less than  $p_a$ ).

In the case of inequality labels ( $t \neq 0$ ), the  $\mathcal{L}_{\text{DIW}}$  loss features a term reminiscent of the log-loss, which pushes prediction values in inequality pairs to be infinitely far apart. For our task of relative traversability prediction, we found that this leads to traversability predictions mostly at the edges of the prediction range, which is unattractive as it degrades the predictions to binary (go/no-go) traversability prediction.

Later work in relative depth and surface normal estimation also describes this problem of the loss encouraging predictions in inequality pairs to be infinitely far apart [16], and proposes a modified loss ( $\mathcal{L}_{\text{SNOW}}$ ) that clamps the difference  $p_b - p_a$  in the original  $\mathcal{L}_{\text{DIW}}$  loss to remedy the issue. We instead propose to introduce a new loss ( $\mathcal{L}_{\text{RIZZ}}$ ) based on a squared hinge loss. Our use of the squared hinge loss allows for squared penalty for both inequality and equality relations, while also resolving the issue of producing mostly extreme predictions as the hinge loss is zero once predictions are sufficiently correct. We have found that this loss provides improved performance over both  $\mathcal{L}_{\text{DIW}}$  and  $\mathcal{L}_{\text{SNOW}}$ . Our  $\mathcal{L}_{\text{RIZZ}}$  loss is calculated as follows:

$$\mathcal{L}_{\text{RIZZ}}(p_a, p_b, t) = \begin{cases} \max\{0, L - t(p_b - p_a)\}^2 & t \neq 0 \\ (p_b - p_a)^2 & t = 0 \end{cases} \quad (2)$$

where  $p_a$ ,  $p_b$ , and  $t$  are defined as before, and  $L$  is a margin hyperparameter.

### C. Training and Network Architecture

We train a model in our framework using our proposed  $\mathcal{L}_{\text{RIZZ}}$  loss. However, like other work in weakly-supervised learning for field robotics [5], [30], we use mean teacher [32] to improve the performance of our model when trained with sparse annotations. We maintain a teacher network and a student network that share the same architecture, where the teacher network weights are an exponential moving average of the student network weights. The final loss used for training a model with W-RIZZ is a weighted sum of an accuracy-based loss ( $\mathcal{L}_{\text{RIZZ}}$ ) and a mean teacher consistency loss (mean squared error). The accuracy-based loss is computed with the student network predictions and is used to learn the relative traversability score, while the consistency loss penalizes differences between the teacher and student network predictions.

Our method is agnostic to the network architecture that is used. Any architecture for per-pixel regression tasks can be used within the W-RIZZ framework, and different models can be selected to balance accuracy and inference speed. For the experiments in this letter, we adopt a modified version of the TravNet architecture introduced by Gasparino et al. [4]. We only use RGB camera images as input and we keep the same encoder (based on ResNet-18 [33]). We modify the TravNet decoder by replacing the transpose convolutional layers with standard  $3 \times 3$  convolution layers that are followed by nearest-neighbor interpolation, and as our final layer we use a  $1 \times 1$  convolutional layer with a sigmoid activation to yield a traversability score between 0 and 1 for each pixel. The network can be viewed as a function  $f_\theta : \mathbb{R}^{H \times W \times 3} \mapsto [0, 1]^{H \times W}$ . We modified the decoder to reduce checkerboarding artifacts seen in the predictions of the original TravNet.

In training, we apply data augmentations in the form of color jitter, random horizontal flipping, and random cropping. As described by Wellhausen et al. [5], geometric augmentations (cropping and flipping) are applied in the same way to the inputs for the teacher and student network, while color jitter is sampled independently for the teacher and student.

## IV. EXPERIMENTAL RESULTS

We validate our method using the WayFAST dataset [4]. The dataset contains 16558 images across a variety of conditions, such as snow, tall grass, forest-like terrain, and semi-urban areas. We label the dataset with 24837 annotations of relative traversability (3 pairwise annotations for every 2 images) for a total of 49674 labeled pixels for the dataset.

For our ablation studies, we split the dataset into a training and validation split, with 13248 training images and 3310 validation images. In all of our experiments, we use a resolution of  $H \times W = 240 \times 424$  for the inputs and outputs of our network. Our dataset contains more equality ordinal labels than inequality labels, and we oversample inequality labels to achieve a more balanced number of labels. For all offline results, we use a computer with an RTX 3080 GPU and neural network inference takes 3 ms per image.

TABLE I  
RESULTS FOR DIFFERENT VARIANTS OF W-RIZZ

Model	Loss	HDR <sub>0.25</sub>	HDR <sub>0.25</sub> <sup>-</sup>	HDR <sub>0.25</sub> <sup>≠</sup>
INIT	$\mathcal{L}_{\text{RIZZ}}$	0.141	0.074	<b>0.266</b>
MT	$\mathcal{L}_{\text{RIZZ}}$	0.129	0.047	0.281
MT+PT	$\mathcal{L}_{\text{RIZZ}}$	<b>0.125</b>	<b>0.044</b>	0.275
MT+PT	$\mathcal{L}_{\text{DIW}}$	0.153	0.067	0.311
MT+PT	$\mathcal{L}_{\text{RIZZ-L1}}$	0.176	0.054	0.401
MT+PT	$\mathcal{L}_{\text{SNOW}}$	0.141	0.044	0.320
WayFAST		0.311	0.124	0.658

### A. Quantitative Results

We measure the quantitative performance of our model using Human Disagreement Rate (HDR), which is an adaptation of the Weighted Human Disagreement Rate (WHDR) metric given by Zoran et al. [15]:

$$\text{HDR}_\tau(p, t) = \frac{\sum_i^N \mathbf{1}(L_\tau(p_{i,a}, p_{i,b}) \neq t_i)}{N} \quad (3)$$

where  $p \in [0, 1]^{N \times 2}$  is the set of neural network predictions at the labeled points,  $t \in \{-1, 0, 1\}^N$  is the set of ground-truth ordinal labels, and  $N$  is the total number of labels.  $L_\tau$  maps prediction pairs to ordinal labels according to threshold  $\tau$  and is defined as follows:

$$L_\tau(p_a, p_b) = \begin{cases} 1, & p_b - p_a > \tau \\ 0, & |p_b - p_a| \leq \tau \\ -1, & p_b - p_a < -\tau \end{cases} \quad (4)$$

We analyze the following model variants: a model with no pretraining on RUGD [12] and no mean teacher [32] (*INIT*); a model using mean teacher but with no pretraining on RUGD (*MT*); and a model using mean teacher and that is pretrained on RUGD (*MT+PT*). For pretraining, we train our network for multi-class segmentation on RUGD and use the weights of all but the last layer to initialize our W-RIZZ model. We test the *MT+PT* model with our  $\mathcal{L}_{\text{RIZZ}}$  loss, as well as  $\mathcal{L}_{\text{DIW}}$  [14],  $\mathcal{L}_{\text{SNOW}}$  [16], and  $\mathcal{L}_{\text{RIZZ-L1}}$  (a variant of our loss that does not square the hinge loss and that uses absolute error for equality pairs). Any hyperparameters for losses are chosen via grid search (all experiments use a value of  $L = 0.5$  for  $\mathcal{L}_{\text{RIZZ}}$ ). In addition to the W-RIZZ variants, we analyze the performance of the self-supervised WayFAST [4], which is trained on the publicly available labels in the WayFAST dataset with the same network architecture used by W-RIZZ (which uses only RGB camera data and is pretrained on RUGD).

We show quantitative results on the validation set for HDR <sub>$\tau$</sub>  (overall human disagreement rate), HDR <sub>$\tau$</sub> <sup>-</sup> (human disagreement rate on pairs labeled as equality), and HDR <sub>$\tau$</sub> <sup>≠</sup> (human disagreement rate on pairs labeled as inequality) for our selected variants at a threshold of  $\tau = 0.25$  in Table I. In Table II, we provide results for WayFAST and the *MT+PT* variant of W-RIZZ (with several different losses) at different values of the disagreement threshold  $\tau$ .

In Table I, we see that the *MT+PT* model with  $\mathcal{L}_{\text{RIZZ}}$  performs best on HDR<sub>0.25</sub> and HDR<sub>0.25</sub><sup>-</sup> (and performs second best on HDR<sub>0.25</sub><sup>≠</sup>, showing only slightly higher error than the *INIT* model). Such a result demonstrates that using our  $\mathcal{L}_{\text{RIZZ}}$  loss along with mean teacher [32] and pretraining on RUGD [12]

TABLE II  
RESULTS FOR MULTIPLE DISAGREEMENT THRESHOLDS

Method	Loss	HDR <sub>0.1</sub>	HDR <sub>0.25</sub>	HDR <sub>0.5</sub>
W-RIZZ	$\mathcal{L}_{DIW}$	0.193	0.153	<b>0.164</b>
W-RIZZ	$\mathcal{L}_{SNOW}$	0.195	0.141	0.266
W-RIZZ	$\mathcal{L}_{RIZZ-L1}$	0.182	0.176	0.293
W-RIZZ	$\mathcal{L}_{RIZZ}$	<b>0.165</b>	<b>0.125</b>	0.225
WayFAST		0.370	0.311	0.310

TABLE III  
INTRA- VS. CROSS-IMAGE LABELING

Labeling Method	HDR <sub>0.25</sub>	HDR <sub>0.25</sub> <sup>=</sup>	HDR <sub>0.25</sub> <sup>≠</sup>
Intra	0.165	<b>0.025</b>	0.425
Intra+Cross	<b>0.133</b>	0.042	<b>0.301</b>

yields the best results. However, the *MT* model (which is not pretrained on RUGD) performs only slightly worse than the *MT+PT* model, indicating that pretraining using existing segmentation datasets is beneficial if possible but is not required for our approach to be successful.

The results in Tables I and II show that our  $\mathcal{L}_{RIZZ}$  loss outperforms the other losses across nearly all metrics. Only for HDR<sub>0.5</sub> (which involves a high threshold for equality) does  $\mathcal{L}_{RIZZ}$  not perform best of all losses, and only  $\mathcal{L}_{DIW}$  performs better on this metric. Moreover,  $\mathcal{L}_{DIW}$  shows relatively similar error across equality thresholds, which suggests that it mostly produces predictions at the extremes of the prediction range (i.e., highly traversable or untraversable). In Tables I and II, our method also shows significantly better HDR than the self-supervised WayFAST [4] (which is trained with over 3000× more labeled pixels), likely due to the added label noise introduced via self-supervised labeling. However, we do acknowledge that WayFAST is not specifically trained for optimizing human disagreement, so lower performance is somewhat expected.

We analyze the effect of the cross-image labeling in Table III. For this experiment, we train models using mean teacher, pre-training on RUGD, and  $\mathcal{L}_{RIZZ}$  loss. We provide results when training using only intra-image labels (*Intra*) and when training using cross-image and intra-image labels (*Intra+Cross*). To ensure fair comparison, we train the two variants using the same number of annotations (so each method uses 2 labels for every 2 images). The results in Table III demonstrate that using cross-image and intra-image labels outperforms using solely intra-image labels both on overall HDR<sub>0.25</sub> and on HDR<sub>0.25</sub><sup>≠</sup>. On the other hand, we see that using only intra-image labels leads to slightly improved error on HDR<sub>0.25</sub><sup>=</sup> (although both labeling schemes show relatively low error on this metric). Nonetheless, the better overall HDR<sub>0.25</sub> indicates that using cross-image and intra-image labels is beneficial by significantly boosting the accuracy on pairs where the ground-truth ordinal relationship is inequality. Moreover, the use of cross-image and intra-image labels is beneficial as it produces more consistently calibrated traversability predictions (since predictions are optimized across images, not just within a given image).

As labeling efficiency is a critical motivation behind our work, in Fig. 3 we analyze the effect of the number of annotated training images on the validation set HDR. For this analysis, we use the *MT+PT* model with  $\mathcal{L}_{RIZZ}$  loss (where each training image




Fig. 3. Validation set HDR<sub>τ</sub> at various thresholds as a function of the number of annotated training images.

has 3 annotated pixels). The results in Fig. 3 demonstrate that increasing the number of annotated images greatly improves HDR when the number of annotated images is small. As the size of the training dataset grows, the effect of adding more annotated images decreases. For example, training on 5000 or 10000 images shows similar HDR to training on the entire 13248 training images, indicating that labeling the entire dataset is not crucial for high performance. However, this behavior will likely vary for different datasets (e.g., datasets with low complexity or low image variation may require fewer annotations for high performance compared to more complex datasets).

## B. Qualitative Results

We visualize several validation set images predicted using W-RIZZ (using the *MT+PT* variant with  $\mathcal{L}_{RIZZ}$  loss) in Fig. 4. The results shown in Fig. 4 demonstrate that W-RIZZ provides accurate relative traversability predictions. Our method also captures nuanced notions of traversability due to its relative labeling, like the smooth concrete being more traversable than bumpy grass in Fig. 4(a) and Fig. 4(f). Despite using only three annotated pixels per training image, W-RIZZ can also capture small obstacles like the metal post in Fig. 4(d) and the legs of the table in Fig. 4(e).

In Fig. 5, we show predictions when using only intra-image labeling and when using cross-image and intra-image labeling. For these predictions, we followed the same training procedure as used for the results in Table III. The results in Fig. 5 demonstrate that using cross-image labeling significantly improves prediction consistency. For example, if only intra-image labeling is used, the snow mound in Fig. 5(c) has approximately the same traversability score as the smooth concrete in Fig. 5(a), and the untraversable tree in Fig. 5(b) has approximately the same traversability score as the traversable grass in Fig. 5(a). However, when cross-image labeling is incorporated, the resulting traversability predictions are significantly more consistent.

In Fig. 6 we provide predictions on sample images from the WayFAST dataset [4] when trained using WayFAST and W-RIZZ, where we train WayFAST as described in Section IV-A. While our method uses a strategy of weakly-supervised manual labeling (with only 3 pixels per image labeled), WayFAST labels traction coefficients in a self-supervised manner and produces many more annotated pixels per image. Although W-RIZZ requires human annotation, our annotation strategy is efficient and takes only seconds per image. Our method also does not require any data beyond images, whereas WayFAST assumes a kinodynamic model and requires additional data/models to predict traction coefficients. Finally, WayFAST requires a robot




Fig. 4. Example inference visualizations from our validation split of the WayFAST dataset [4]. The input color images are shown in the top row, and the corresponding traversability predictions are shown in the bottom row. The traversability score ranges from 0 to 1 for each image, with a higher traversability score indicating that a region is more easily traversed.

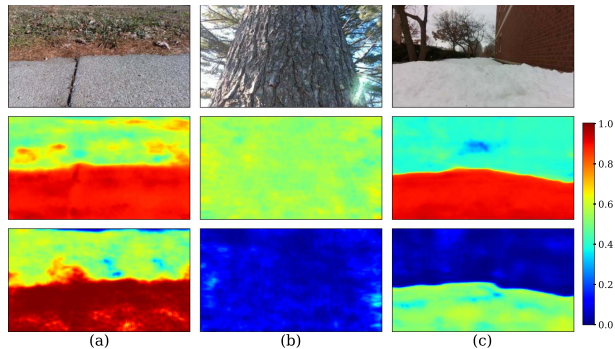


Fig. 5. Predictions on images from the WayFAST dataset [4], showing color images (top), as well as predictions from a model trained without cross-image labeling (middle) and with cross-image labeling (bottom).




Fig. 6. Example predictions on images from the WayFAST dataset [4], showing color images (top), predictions from WayFAST (middle), and predictions from W-RIZZ (bottom).

to experience traction conditions to label them, which makes it difficult to safely label untraversable zones.

The results provided in Fig. 6 highlight important distinctions between the predictions from WayFAST and W-RIZZ. For example, W-RIZZ predicts higher traversability for concrete than grass (which is beneficial as smooth concrete is likely a more appealing surface to traverse than bumpy grass if all else is equal), while WayFAST fails to make a distinction between the two surfaces. In addition, the difficulty of labeling untraversable areas with the self-supervised strategy leads to erroneous traversability predictions in Fig. 6(b)–(e). For example, WayFAST fails to capture the distant obstacle in Fig. 6(b), the right leg of the metal table in Fig. 6(c), and the concrete pillar in Fig. 6(e). WayFAST also produces odd, overly-cautious predictions regarding the tree in Fig. 6(d). By comparison, W-RIZZ produces significantly more accurate predictions for obstacles.

TABLE IV  
W-RIZZ VS. STRONGLY-SUPERVISED SEGMENTATION ON RUGD

Method	mIOU	fw-mIOU	mAcc	fw-mAcc
W-RIZZ	0.668	0.764	0.767	0.856
Segmentation	0.692	0.783	0.785	0.879

### C. Comparison to Strongly-Supervised Segmentation

We additionally compare our sparse annotation method to strongly-supervised segmentation on the RUGD dataset [12]. This dataset consists of images which are labeled for semantic segmentation using 25 classes (if the void class is included) relevant to outdoor navigation in unstructured environments, and we divide its semantic annotations into 4 traversability tiers: tier 3 (concrete and asphalt), tier 2 (dirt, sand, grass, gravel, mulch, and rock-bed), tier 1 (water and bush), and tier 0 (all other classes), where low tiers correspond to lower traversability. We use these tiers to create automatic annotations for W-RIZZ (with 3 annotated pixels per image) which we use to train our model. We use automated annotations to ensure consistent comparison to the strongly-supervised labeling. As over 90% of the pixels in the top half of training set images correspond to sky or tree, we bias our sampling such that 90% of samples must be in the bottom half of the image (with the remaining 10% being unrestricted) to provide more informative labels. We train the *MT*W-RIZZ model (i.e., we do not pretrain on RUGD) using the sparse automatically-generated annotations, and compare our method to a strongly-supervised model using the same network architecture (where the last layer is modified to produce 25-class segmentation outputs).

We provide results for our method and the strongly-supervised approach on the RUGD test set in Table IV. We compute our metrics assuming 4-class segmentation of the traversability tiers described earlier, and report mean intersection-over-union (*mIOU*) and mean pixel-wise accuracy (*mAcc*), where the means are computed over the 4 classes. We also provide frequency-weighted metrics where the mean is weighted by the relative frequency of each class (*fw-mIOU* and *fw-mAcc*). To convert the traversability scores from W-RIZZ to discrete traversability tiers, we compute the cutoff that differentiates tier  $N$  from tier  $N - 1$  ( $N \in \{3, 2, 1\}$ ) as  $\text{cutoff}_N = \frac{(\mu_N - \sigma_N) + (\mu_{N-1} + \sigma_{N-1})}{2}$ , where  $\mu_N$  and  $\sigma_N$  are the mean and standard deviation of the W-RIZZ traversability scores for tier  $N$  on the training set. The segmentation model produces 25-class outputs, and we remap these to the 4 traversability tiers for our metric calculations.

The results in Table IV highlight that our sparse annotation strategy shows reasonably high performance despite using very few labeled pixels in training. Even though only three annotated pixels are provided for each training image, W-RIZZ shows results within 4% of the strongly-supervised segmentation approach for each metric, despite the strongly-supervised model being trained with significantly more annotations (as the strongly-supervised segmentation model is trained using images where every pixel is labeled).

#### D. Real-World Experiments

We further analyze our method by conducting real-world experiments on a TerraSentia robot (a skid-steer robot developed by EarthSense Inc.). For these experiments, we utilize the traversability predictions from W-RIZZ with a non-linear model predictive controller (NMPC) described by Gasparino et al. [4]. The neural network inference and controller are run on the robot’s onboard computer (Jetson AGX Orin) at 15 Hz. While we use the controller introduced by Gasparino et al. [4] for these experiments, W-RIZZ is agnostic to the control strategy and can be used with other controllers that consume traversability prediction images.

We compare our approach against the following baselines:

- 1) *LiDAR*: an obstacle-avoiding approach based on purely geometric information acquired by a 2D LiDAR unit.
- 2) *WayFAST* [4]: an RGB-only variant of a state-of-the-art self-supervised visual traversability prediction method.
- 3) *WayFAST\** [4]: an enhanced version of WayFAST trained on a different dataset and using the RGB-D WayFAST network architecture.

The W-RIZZ and WayFAST models are trained with the same neural network architecture (which uses only RGB images) on the full dataset from Gasparino et al. [4]. We train W-RIZZ using our relative traversability labels, and we train WayFAST using the labels from the WayFAST dataset.

The robot used for our experiments had a camera setup that differed slightly from that of the original WayFAST dataset (the WayFAST dataset has a different camera that is angled more towards the ground). While this did not cause issues with W-RIZZ, it led to poor performance of the original WayFAST model. Thus, we provide results with WayFAST\*, which is trained on a new dataset that more closely matches the camera setup of the robot used for the experiments and that uses the RGB-D network architecture from WayFAST (which utilizes RGB and depth images). The W-RIZZ model used in these experiments is only trained on the original WayFAST dataset and does not use a depth camera.

We perform our navigation experiments in a grove of trees where a robot must navigate approximately 20 m between obstacles to reach a goal; the setup of the experiment is shown in Fig. 7. We perform 5 navigation trials for each method, and the results are shown in Table V. All methods took approximately the same time to reach the goal (if the goal was reached). Videos from the experiments and for a longer navigation task are provided as supplemental material.

The results from Table V demonstrate that both the LiDAR baseline and W-RIZZ are able to successfully navigate the course in all attempts. The WayFAST model (trained on the same data as W-RIZZ) crashed into a tree in every navigation attempt. These results suggest greater generalization capability




Fig. 7. Navigation experiment setup, showing the TerraSentia robot (left), as well as a sample camera image from the robot (top right) and the associated traversability prediction from W-RIZZ (bottom right).

TABLE V  
ROBOT NAVIGATION EXPERIMENTAL RESULTS

Method	Successful Trials	Total Trials	Success Rate
LiDAR	5	5	100%
WayFAST	0	5	0%
WayFAST*	4	5	80%
W-RIZZ	5	5	100%

with W-RIZZ when compared to WayFAST. The results for WayFAST\* show improved performance compared to WayFAST, demonstrating that incorporating depth and retraining on a new dataset help to overcome the domain shift. However, even when given such training advantages over W-RIZZ, WayFAST\* still showed lower navigation success rate than W-RIZZ (in one of its navigation attempts WayFAST\* failed to avoid an obstacle, crashing into a tree).

While both W-RIZZ and LiDAR show 100% success rate in our experiments, we note the difference in the perception method of RGB camera versus 2D LiDAR. The 2D LiDAR has a greater field of view and produces simpler measurement data that only provide a purely geometric view of the environment, which can fail for traversable obstacles like tall grass and cannot capture features like smoothness of terrain. The richer vision-based sensing modality used by W-RIZZ can capture such nuances as shown in Section IV-B. The 2D LiDAR can also display measurement limitations, as obstacles above or below the measurement plane of the LiDAR unit cannot be detected.

## V. CONCLUSION

We have introduced a novel method of relative traversability estimation that uses a small number of sparsely annotated point pairs which are labeled with annotations of relative traversability. We enhance the performance of our method by introducing a new loss for relative traversability estimation and a novel cross-image labeling scheme. We have validated our approach through offline evaluation in various settings (including snow, tall grass, forests, and semi-urban areas) and through deployment on an actual mobile robot platform.

Although our method does not depend on the robotic platform that is used, it shows the limitation of requiring relabeling of data if a robot’s capabilities change significantly. In addition, while our method still requires manual effort for labeling, the sparse relative annotation scheme is simple, does not need extensive training, and demands significantly less effort than labeling for strongly-supervised segmentation

(while maintaining high accuracy). Our use of manual relative traversability labels also circumvents the need to define semantic classes and does not assume constant traversability characteristics within semantic classes (in contrast to semantic segmentation approaches), allowing for direct prediction of a continuous traversability score while not requiring robot experience or the data, models, or assumptions needed in self-supervised approaches. Our framework can be used with any neural network for dense prediction tasks and is compatible with techniques like pretraining to further improve results.

Our online experiments demonstrate that our method can outperform a state-of-the-art self-supervised approach and shows greater robustness to domain shifts like varying camera angles. Possible future work includes incorporating additional methods for weakly-supervised learning (e.g., pseudo-labeling) and combining our method with foundation models to improve accuracy and generalization performance.

#### ACKNOWLEDGMENT

The robot platforms and data were provided by the Illinois Autonomous Farm and the Illinois Center for Digital Agriculture.

#### REFERENCES

- [1] E. Kayacan, Z. Zhang, and G. Chowdhary, "Embedded high precision control and corn stand counting algorithms for an ultra-compact 3D printed field robot," in *Proc. Robot.: Sci. Syst.*, 2018. [Online]. Available: <https://www.roboticsproceedings.org/rss14/p36.html>
- [2] D. Lee, G. Kang, B. Kim, and D. H. Shim, "Assistive delivery robot application for real-world postal services," *IEEE Access*, vol. 9, pp. 141981–141998, 2021.
- [3] P. Papadakis, "Terrain traversability analysis methods for unmanned ground vehicles: A survey," *Eng. Appl. Artif. Intell.*, vol. 26, no. 4, pp. 1373–1385, 2013.
- [4] M. V. Gasparino et al., "WayFAST: Navigation with predictive traversability in the field," *IEEE Robot. Automat. Lett.*, vol. 7, no. 4, pp. 10651–10658, Oct. 2022.
- [5] L. Wellhausen, A. Dosovitskiy, R. Ranftl, K. Walas, C. Cadena, and M. Hutter, "Where should I walk? Predicting terrain properties from images via self-supervised learning," *IEEE Robot. Automat. Lett.*, vol. 4, no. 2, pp. 1509–1516, Apr. 2019.
- [6] J. Seo, S. Sim, and I. Shim, "Learning off-road terrain traversability with self-supervisions only," *IEEE Robot. Automat. Lett.*, vol. 8, no. 8, pp. 4617–4624, Aug. 2023.
- [7] R. Schmid et al., "Self-supervised traversability prediction by learning to reconstruct safe terrain," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 12419–12425.
- [8] S. Sharma et al., "CaT: CAVS traversability dataset for off-road autonomous driving," *IEEE Access*, vol. 10, pp. 24759–24768, 2022.
- [9] A. Valada, J. Vertens, A. Dhall, and W. Burgard, "AdapNet: Adaptive semantic segmentation in adverse environmental conditions," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 4644–4651.
- [10] A. Valada, G. Oliveira, T. Brox, and W. Burgard, "Deep multispectral semantic scene understanding of forested environments using multimodal fusion," in *Int. Symp. Exp. Robot.*, 2016, pp. 465–477.
- [11] D. Maturana, P.-W. Chou, M. Uenoyama, and S. Scherer, "Real-time semantic mapping for autonomous off-road navigation," in *Proc. Int. Conf. Field Service Robot.*, 2017, pp. 335–350.
- [12] M. Wigness, S. Eum, J. G. Rogers, D. Han, and H. Kwon, "A RUGD dataset for autonomous navigation and visual perception in unstructured outdoor environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 5000–5007.
- [13] P. Jiang, P. Osteen, M. Wigness, and S. Saripalli, "RELLIS-3D dataset: Data, benchmarks and analysis," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 1110–1116.
- [14] W. Chen, Z. Fu, D. Yang, and J. Deng, "Single-image depth perception in the wild," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 730–738.
- [15] D. Zoran, P. Isola, D. Krishnan, and W. T. Freeman, "Learning ordinal relationships for mid-level vision," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 388–396.
- [16] W. Chen, D. Xiang, and J. Deng, "Surface normals in the wild," in *Proc. IEEE Int. Conf. Computer Vis.*, 2017, pp. 1557–1566.
- [17] S. Bell, K. Bala, and N. Snavely, "Intrinsic images in the wild," *ACM Trans. Graph.*, vol. 33, no. 4, pp. 1–12, 2014.
- [18] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA, USA: MIT Press, Aug. 2005.
- [19] R. Siegwart, I. R. Nourbakhsh, and D. Scaramuzza, *Introduction to Autonomous Mobile Robots*, 2nd ed. Cambridge, MA, USA: MIT Press, Feb. 2011.
- [20] T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in *Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
- [21] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 6230–6239.
- [22] M. G. Castro et al., "How does it feel? Self-supervised costmap learning for off-road vehicle traversability," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 931–938.
- [23] G. Kahn, P. Abbeel, and S. Levine, "BADGR: An autonomous self-supervised learning-based navigation system," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 1312–1319, Apr. 2021.
- [24] X. Cai, M. Everett, L. Sharma, P. R. Osteen, and J. P. How, "Probabilistic traversability model for risk-aware motion planning in off-road environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 11297–11304.
- [25] J. Seo, T. Kim, K. Kwak, J. Min, and I. Shim, "ScaTE: A scalable framework for self-supervised traversability estimation in unstructured environments," *IEEE Robot. Automat. Lett.*, vol. 8, no. 2, pp. 888–895, Feb. 2023.
- [26] A. Vezhnevets, V. Ferrari, and J. M. Buhmann, "Weakly supervised semantic segmentation with a multi-image model," in *Proc. Int. Conf. Comput. Vis.*, 2011, pp. 643–650.
- [27] A. Bearman, O. Russakovsky, V. Ferrari, and L. Fei-Fei, "What's the point: Semantic segmentation with point supervision," in *Eur. Conf. Comput. Vis.*, 2016, pp. 549–565.
- [28] D. Lin, J. Dai, J. Jia, K. He, and J. Sun, "ScribbleSup: Scribble-supervised convolutional networks for semantic segmentation," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 3159–3167.
- [29] B. Gao, S. Hu, X. Zhao, and H. Zhao, "Fine-grained off-road semantic segmentation and mapping via contrastive learning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 5950–5957.
- [30] A. M. Schreiber, "Weakly-supervised traversability estimation for mobile robots using sparse point annotation," Master's thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA, 2023.
- [31] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," in *Proc. Int. Conf. Learn. Representations*, 2017.
- [32] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 1195–1204.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.