

# Pseudo-Trilateral Adversarial Training for Domain Adaptive Traversability Prediction

Zheng Chen

Indiana University Bloomington

Durgakant Pushp

Indiana University Bloomington

Lantao Liu (✉ [lantao@iu.edu](mailto:lantao@iu.edu))

Indiana University Bloomington

---

## Research Article

**Keywords:** Domain adaptation, Semantic segmentation, Adversarial training, Class alignment

**Posted Date:** January 30th, 2023

**DOI:** <https://doi.org/10.21203/rs.3.rs-2505079/v1>

**License:** © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

**Additional Declarations:** No competing interests reported.

---

**Version of Record:** A version of this preprint was published at Autonomous Robots on August 2nd, 2023.  
See the published version at <https://doi.org/10.1007/s10514-023-10123-7>.

# Pseudo-Trilateral Adversarial Training for Domain Adaptive Traversability Prediction

Zheng Chen · Durgakant Pushp · Lantao Liu

Received: date / Accepted: date

**Abstract Keywords** Domain adaptation · Semantic segmentation · Adversarial training · Class alignment

Traversability prediction is a fundamental perception capability for autonomous navigation. Deep neural networks (DNNs) have been widely used to predict traversability during the last decade. The performance of DNNs is significantly boosted by exploiting a large amount of data. However, the diversity of data in different domains imposes significant gaps in the prediction performance. In this work, we make efforts to reduce the gaps by proposing a novel pseudo-trilateral adversarial model named CALI to perform *unsupervised domain adaptation* (UDA). Our aim is to transfer the perception model with high data efficiency, eliminate the prohibitively expensive data labeling, and improve the generalization capability during the adaptation from easy-to-access *source domains* to various challenging *target domains*. Existing UDA methods usually adopt a bilateral zero-sum game structure. We prove that our CALI model — a pseudo-trilateral game structure is advantageous over existing bilateral game structures. This proposed work bridges theoretical analyses and algorithm designs, leading to an efficient UDA model with easy and stable training. We show the superiorities of our proposed model over multiple baselines in several challenging domain adaptation setups. To further validate the effectiveness of our model, we then combine our perception model with a visual planner to build a navigation system and show the high reliability of our model in complex natural environments.



Fig. 1: Transferring models from the available domain to the target domain. The existing available data might be from either a simulator or collecting data in certain environments, at a certain time, and with certain sensors. In contrast, the target deployment might have significantly varying environments, time, and sensors.

## 1 Introduction

We consider the deployment of autonomous robots in real-world unstructured field environments, where the environments can be extremely complex involving random obstacles (e.g., big rocks, tree stumps, man-made objects), cross-domain terrains (e.g., combinations of gravel, sand, wet, uneven surfaces), as well as dense vegetation (tall and low grasses, shrubs, trees). Whenever a robot is deployed in such an environment, it needs to understand which area of the captured scene is navigable. A typical solution to this problem is the visual traversability prediction that can be achieved by learning the *scene semantic segmentation*.

Visual traversability prediction has been tackled using deep neural networks where the models are typically trained offline with well-labeled datasets in limited scenarios. However, there might exist a gap between the data used to train the model and the data when testing.

Zheng Chen, Durgakant Pushp, Lantao Liu are with Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN 47408, USA  
E-mail: {zc11, dpushp, lantao}@iu.edu

There are several existing datasets for semantic segmentation, including GTA5 Richter et al. [2016], SYNTHIA Ros et al. [2016], Cityscapes Cordts et al. [2016], ACDC Sakaridis et al. [2021], Dark Zurich Sakaridis et al. [2019], RUGD Wigness et al. [2019], RELIS Jiang et al. [2020], ORFD Min et al. [2022]. Nevertheless, it is usually challenging for existing datasets to well approximate the true distributions of unseen *target* environments where the robot is deployed. Even incrementally collecting and adding new training data on the fly cannot guarantee the target environments to be well *in-distribution* included. In addition, annotating labels for dense predictions, e.g., semantic segmentation, is prohibitively expensive. Therefore, developing a generalization-aware deep model is crucial for the **robustness, trustworthiness, and safety** of robotic systems considering the demands of the practical deployment of deep perception models and the costs/limits of collecting new data in many robotic applications, e.g., autonomous driving, search and rescue, and environmental monitoring.

To tackle this challenge, a broadly studied framework is *transfer learning* Pan and Yang [2009] which aims to transfer models between two domains – *source domain* and *target domain* – that have related but different data distributions. The prediction on target domain can be considered as a strong generalization since testing data (in target domain) might fall out of the independently and identically distributed (i.i.d.) assumption and follow a very different distribution than the training data (in source domain). The “transfer” process has significant meaning to our model development since we can view the available public datasets Richter et al. [2016], Cordts et al. [2016], Wigness et al. [2019], Jiang et al. [2020] as the source domain and treat the data in the to-be-deployed environments as the target domain. In this case, we have access to images and corresponding labels in source domain and images in target domain, but no access to labels in target domain. Transferring models, in this set-up, is called *Unsupervised Domain Adaptation* (UDA) Wilson and Cook [2020], Zhang [2021].

Domain Alignment (DA) Ganin et al. [2016], Hoffman et al. [2016, 2018], Tsai et al. [2018], Vu et al. [2019] and Class Alignment (CA) Saito et al. [2018] are two conventional ways to tackle the UDA problem. DA treats the deep features as a whole. It works well for image-level tasks such as image classification, but has issues with pixel-level tasks such as semantic segmentation, as the alignment of whole distributions ignores the class features and might misalign class distributions, even the whole features from the source domain and target domain are already well-aligned. CA is proposed

to solve this issue for dense predictions with multiple classes.

It is natural and necessary to use CA to tackle the UDA of semantic segmentation as we need to consider aligning class features. However, CA can be problematic and might fail to outperform the DA for segmentation, and in a worse case, might have unacceptable *negative transfer*, which means the performance with adaptation is even degraded than that without adaptation. We empirically found that the training of CA can be unstable and easy to diverge ♣, leading to low performance or even failure of training. Typically, CA adopts a network consisting of a feature extractor and two classification heads. The net is first trained on the source domain such that the decision boundaries of the two heads are both able to well classify source features of different classes. Then the net is trained on the target domain in an adversarial manner, where the discrepancy of the two heads is used as the adversarial objective. During adversarial training, the goal of the feature extractor is to generate features such that the discrepancy between the two heads is minimized; while the goal of the classification heads is to adjust the decision boundaries such that the discrepancy is maximized. By doing so the target features are enforced to be aligned with the trained source features. However, adversarial training can be highly unstable especially when the adversarial objective is not properly valued. We conjecture the reason for the problem of CA ♣ is the value of the adversarial objective — the discrepancy between the two heads is too large due to the domain shift, making the two classification heads in a near-optimal state at the beginning of training. In this case, the training of the two classification heads can quickly converge, breaking the equilibrium of the zero-sum game, and leading to a failure of the whole training.

To solve the issue of CA, our intuition is to apply a DA as a constraint to reduce the general distance between two domains before applying CA. This will also bring a reduction of the discrepancy value between the two classification heads, thus the adversarial training over the discrepancy can be stabilized. An example to explain our intuition is shown in Fig. 2. To achieve this, we investigate the relationship of the upper bounds of the prediction error on the target domain between DA and CA and provide a theoretical analysis of the upper bounds of target prediction error for the two alignments in the UDA setup. Our theoretical analysis justifies the use of DA as a constraint of CA. Guided by the theoretical analysis, we develop a novel UDA model named CALI.

In summary, our contributions include

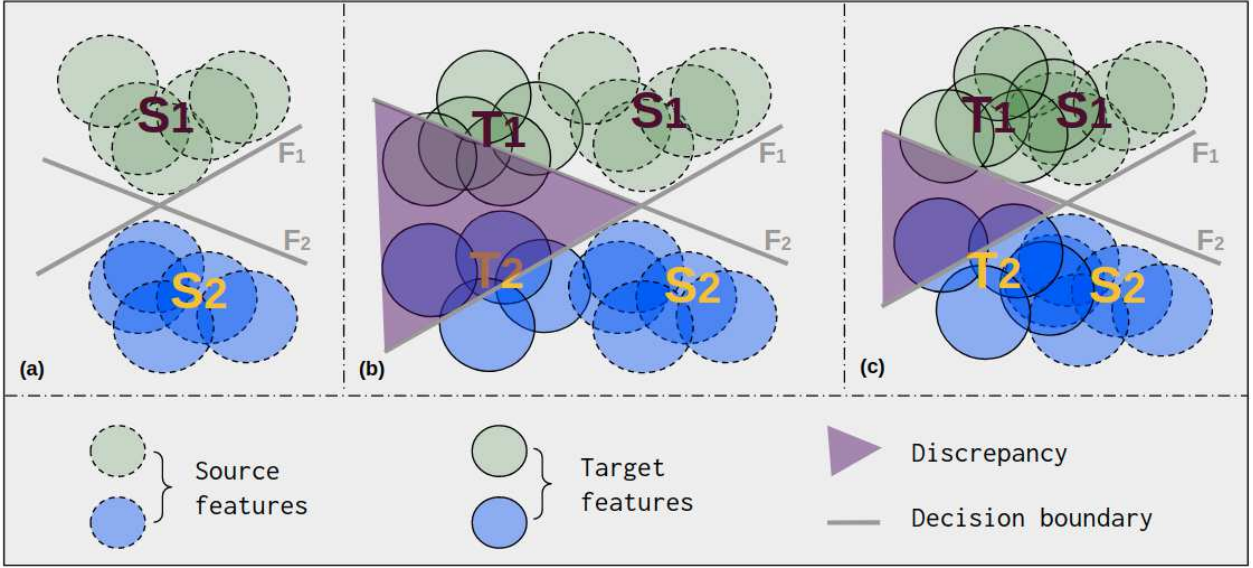


Fig. 2: An example to explain the intuition of our CALI model.  $S_1/T_1$ : Features of class #1 from source/target data;  $S_2/T_2$ : Features of class #2 from source/target data. (a) We first have two well-trained classifiers  $F_1$  and  $F_2$  on the source domain; (b) The discrepancy can be large for the two trained classifiers on the target domain due to the large domain shift; (c) Our CALI model aims to apply domain alignment to reduce the general domain distance between the source domain and the target domain before performing the class alignment. In this case, the adversarial objective for class alignment — the discrepancy between the two classifiers, can be reduced to a proper level (see the shadow area change from (b) to (c)). With a well-regularized adversarial objective, the training of CA can be stabilized and the performance can be improved.

- We prove that with proper assumptions, the upper bound of CA is upper bounded by the upper bound of DA. This indicates that constraining the training of CA using DA can be beneficial. We then propose a novel concept of *pseudo-trilateral game structure* (PTGS) for integrating DA and CA.
- We propose an efficient coarse-to-fine alignment based UDA model, named CALI, for traversability prediction. The new proposal includes a trilateral network structure, novel training losses, and an alternative training process. Our model design is well supported by theoretical analysis. It is also easy and stable to train and converge.
- We show significant advantages of our proposed model compared to several baselines in multiple challenging public datasets and one self-collected dataset. We combine the proposed segmentation model and a visual planner to build a visual navigation system. The results show high safety and effectiveness of our model.

## 2 Related Work

**Semantic Segmentation:** Semantic segmentation aims to predict a unique human-defined semantic class for each pixel in the given images. With the prosperity of deep neural networks, the performance of seman-

tic segmentation has been boosted significantly, especially by the advent of FCN Long et al. [2015] that first proposes to use deep convolutional neural nets to predict segmentation. The following works try to improve the FCN performance by multiple proposals, e.g., using different sizes of kernels or dilation rates to aggregate multi-scale features Chen et al. [2017a,b], Yu and Koltun [2015]; building image pyramids to create multi-resolution inputs Zhao et al. [2017]; applying probabilistic graph to smooth the prediction Liu et al. [2017]; compensating features in deeper level by an encoder-decoder structure Ronneberger et al. [2015]. Recently, Transformers Vaswani et al. [2017], Dosovitskiy et al. [2020] have gained huge popularity for various vision tasks including semantic segmentation. Transformer-based models employ an attention mechanism to capture the long-range dependencies among pixels. Different Transformer-based segmentation models have been developed, e.g., finer-grained and more globally coherent predictions are achieved for dense predictions by assembling tokens from various stages of the vision transformer into image-like representations at various resolutions Ranftl et al. [2021]; Semantic segmentation is treated as a sequence-to-sequence prediction task and a pure transformer without convolution and resolution reduction are used to encode an image as a sequence of patches Zheng et al. [2021]; A novel hierarchically structured Transformer encoder is combined with a lightweight



MLP decoder to build a simple, efficient yet powerful semantic segmentation framework Xie et al. [2021]. Recent SOTA works Ranftl et al. [2021], Zheng et al. [2021], Xie et al. [2021] for semantic segmentation heavily rely on Transformer structure. However, all of those methods belong to fully-supervised learning and the performance might be degraded catastrophically when a domain shift exists between the training data and the data when deploying.

**Unsupervised Domain Adaptation:** The main approaches to tackle UDA include adversarial training (a.k.a., *distribution alignment*) Ganin et al. [2016], Hoffman et al. [2016, 2018], Tsai et al. [2018], Saito et al. [2018], Vu et al. [2019], Luo et al. [2019], Wang et al. [2020] and self-training Zou et al. [2018], Zhang et al. [2017], Mei et al. [2020], Hoyer et al. [2021]. Self-training maintains a teacher-student framework to conduct the knowledge transfer. The teacher model is trained on the source domain and used to predict segmentation for target images. The predictions from the teacher model are then used as pseudo labels to train the student model. During the training of the student model, existing methods use different ways to identify the erroneous regions, including confidence score Zou et al. [2019, 2018] and entropy Xie et al. [2022], Chen et al. [2019], Pan et al. [2020]. Although self-training is becoming a popular method for segmentation UDA in terms of empirical results, it still lacks a sound theoretical foundation. In this paper, we only focus on the alignment-based methods that not only keep close to the UDA state-of-the-art (SOTA) performance but are also well supported by sound theoretical analyses Ben-David et al. [2007], Blitzer et al. [2008], Ben-David et al. [2010].

The alignment-based methods adapt models via aligning the distributions from the source domain and target domain in an adversarial training process, i.e., making the deep features of source images and target images indistinguishable to a discriminator net. Typical alignment-based approaches to UDA include Domain Alignment Ganin et al. [2016], Hoffman et al. [2016, 2018], Tsai et al. [2018], Vu et al. [2019], which aligns the two domains using global features (aligning the feature tensor from source or target *as a whole*) and Class Alignment Saito et al. [2018], Luo et al. [2019], Wang et al. [2020], which only considers aligning features of each class from source and target, no matter whether the domain distributions are aligned or not. In Saito et al. [2018], the authors are inspired by the theoretical analysis of Ben-David et al. [2010] and propose a discrepancy-based model for aligning class features. There is a clear relation between the theory guidance Ben-David et al. [2010] and the design of network, loss,

and training methods. There are some recent works Luo et al. [2019], Wang et al. [2020] similar to our proposed work in spirit and show improved results compared to Saito et al. [2018], but it is still unclear to relate the proposed algorithms with theory and to understand why the structure/loss/training is designed as the presented way.

**Visual Navigation:** To achieve visual navigation autonomously, learning-based methods have been widely studied recently Shen et al. [2019], Bansal et al. [2020]. For example, imitation learning based approaches have been largely explored to train a navigation policy that enables a robot to mimic human behaviors or navigate close to certain waypoints without a prior map Manderson et al. [2020], Hirose et al. [2020]. To fully utilize the known dynamics model of the robot, a semi-learning-based scheme is also proposed Bansal et al. [2020] to combine optimal control and deep neural network to navigate through unknown environments. A large amount of work on visual navigation can also be found in the computer vision community, such as Shen et al. [2019], Gupta et al. [2017a], Chaplot et al. [2020], Gupta et al. [2017b], Wu et al. [2019], all of which use full-learning-based methods to train navigation policies, which work remarkably well when training data is sufficient but can fail frequently if no or very limited data is available.

### 3 Background and Preliminary Materials

#### 3.1 Expected Errors

We consider segmentation tasks where the input space is  $\mathcal{X} \subset \mathbb{R}^{H \times W \times 3}$ , representing the input RGB images, and the label space is  $\mathcal{Y} \subset \{0, 1\}^{H \times W \times K}$ , representing the ground-truth  $K$ -class segmentation images, where the label for a single pixel at  $(h, w)$  is denoted by a one-hot vector  $y^{(h,w)} \in \mathbb{R}^K$  whose elements are by-default 0-valued except at location  $(h, w)$  which is labeled as 1. Domain adaptation has two domain distributions over  $\mathcal{X} \times \mathcal{Y}$ , named source domain  $\mathcal{D}_S$  and target domain  $\mathcal{D}_T$ . In the setting of UDA for segmentation, we have access to  $m_s$  i.i.d. samples with labels  $\mathcal{U}_S = \{\mathbf{x}_{si}, \mathbf{y}_{si}\}_{i=1}^{m_s}$  from  $\mathcal{D}_S$  and  $m_t$  i.i.d. samples without labels  $\mathcal{U}_T = \{\mathbf{x}_{tj}\}_{j=1}^{m_t}$  from  $\mathcal{D}_T$ .

In the UDA problem, we need to reduce the prediction error on the target domain. With a slight abuse of notation, we also use  $h$  to denote a *hypothesis*, which is a function:  $\mathcal{X} \rightarrow \mathcal{Y}$ . We denote the space of  $h$  as  $\mathcal{H}$ . With the loss function  $l(\cdot, \cdot)$ , the expected error of  $h$  on  $\mathcal{D}_S$  is defined as

$$\epsilon_S(h) := \mathbb{E}_{(x,y) \sim \mathcal{D}_S} l(h(x), y). \quad (1)$$

Similarly, we can define the expected error of  $h$  on  $\mathcal{D}_T$  as

$$\epsilon_T(h) := \mathbb{E}_{(x,y) \sim \mathcal{D}_T} l(h(x), y). \quad (2)$$

### 3.2 Upper Bounds for Expected Errors

Two important upper bounds related to the source and target error are given in Ben-David et al. [2010].

① The first upper bound is given in the following theorem. Basically,

**Theorem 1** For a hypothesis  $h$ ,

$$\epsilon_T(h) \leq \epsilon_S(h) + d_1(\mathcal{D}_S, \mathcal{D}_T) + \lambda, \quad (3)$$

where  $d_1(\cdot, \cdot)$  is the  $L^1$  divergence for two distributions, and the constant term  $\lambda$  does not depend on any  $h$ . However, it is claimed in Ben-David et al. [2010] that the bound with  $L^1$  divergence cannot be accurately estimated from finite samples, and using  $L^1$  divergence can unnecessarily inflate the bound. Another divergence measure is thus introduced to replace the  $L^1$  divergence with a new bound derived. The new measure is defined as follows,

**Definition 1** Given two domain distributions  $\mathcal{D}_S$  and  $\mathcal{D}_T$  over  $\mathcal{X}$ , and a hypothesis space  $\mathcal{H}$  that has finite VC dimension, the  $\mathcal{H}$ -divergence between  $\mathcal{D}_S$  and  $\mathcal{D}_T$  is defined as

$$d_{\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) = 2 \sup_{h \in \mathcal{H}} |\mathbb{P}_{x \sim \mathcal{D}_S} [h(x) = 1] - \mathbb{P}_{x \sim \mathcal{D}_T} [h(x) = 1]|, \quad (4)$$

where  $\mathbb{P}_{x \sim \mathcal{D}_S} [h(x) = 1]$  represents the probability of  $x$  belonging to  $\mathcal{D}_S$ . Same to  $\mathbb{P}_{x \sim \mathcal{D}_T} [h(x) = 1]$ .

The  $\mathcal{H}$ -divergence resolves the issues in the  $L^1$  divergence. If we replace  $d_1(\mathcal{D}_S, \mathcal{D}_T)$  in Eq. (3) with  $d_{\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T)$ , then a new upper bound for  $\epsilon_T(h)$ , named as  $\mathbb{UB}_1$ , can be written as

$$\epsilon_T(h) \leq \mathbb{UB}_1, \quad (5)$$

$$\mathbb{UB}_1 = \epsilon_S(h) + d_{\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) + \lambda.$$

An approach to compute the empirical  $\mathcal{H}$ -divergence is also proposed in Ben-David et al. [2010], see the below Lemma 1.

**Lemma 1** For a symmetric hypothesis class  $\mathcal{H}$  (one where for every  $h \in \mathcal{H}$ , the inverse hypothesis  $1 - h$  is also in  $\mathcal{H}$ ) and two sample sets  $\mathcal{U}_S = \{x_i, i = 1, \dots, m_s, x_i \sim \mathcal{D}_S\}$  and  $\mathcal{U}_T = \{x_j, j = 1, \dots, m_t, x_j \sim \mathcal{D}_T\}$ .

$$\hat{d}_{\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) = 2 \left( 1 - \min_{\eta \in \mathcal{H}} \left[ \frac{1}{m_s} \sum_{i=1}^{m_s} \mathbb{I}[\eta(x_i) = 0] + \frac{1}{m_t} \sum_{j=1}^{m_t} \mathbb{I}[\eta(x_j) = 1] \right] \right),$$

where  $\mathbb{I}[a]$  is an indicator function which is 1 if  $a$  is true, and 0 otherwise.

② The second upper bound is based on a new hypothesis called the symmetric difference hypothesis, see the following definition.

**Definition 2** For a hypothesis space  $\mathcal{H}$ , the symmetric difference hypothesis space  $\mathcal{H}\Delta\mathcal{H}$  is the set of hypotheses

$$g \in \mathcal{H}\Delta\mathcal{H} \Leftrightarrow g(x) = h(x) \oplus h'(x) \quad \text{for some } h, h' \in \mathcal{H}, \quad (7)$$

where  $\oplus$  denotes an XOR operation. Then we can define the  $\mathcal{H}\Delta\mathcal{H}$ -distance as

$$d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) = 2 \sup_{h, h' \in \mathcal{H}} |\mathbb{P}_{x \sim \mathcal{D}_S} [h(x) \neq h'(x)] - \mathbb{P}_{x \sim \mathcal{D}_T} [h(x) \neq h'(x)]|. \quad (8)$$

Similar to Eq. (5), if we replace  $d_1(\mathcal{D}_S, \mathcal{D}_T)$  with  $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T)$ , the second upper bound for  $\epsilon_T(h)$ , named as  $\mathbb{UB}_2$ , can be expressed as

$$\epsilon_T(h) \leq \mathbb{UB}_2, \quad (9)$$

$$\mathbb{UB}_2 = \epsilon_S(h) + d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) + \lambda,$$

where  $\lambda$  is the same term as in Eq. (3).

The two bounds (Eq. (5) and Eq. (9)) for the target domain error are separately given in Ben-David et al. [2010]. It has been independently demonstrated that DA corresponds to optimizing over  $\mathbb{UB}_1$  Ganin et al. [2016], where optimization over the upper bound  $\mathbb{UB}_1$  (Eq. (5) with the divergence Eq. (6)) is proved as equivalent to an adversarial learning with Eq. (10) and a supervised learning with the source data, and that CA corresponds to optimizing over  $\mathbb{UB}_2$  Saito et al. [2018], where the  $d_{\mathcal{H}\Delta\mathcal{H}}$  is approximated by the discrepancy between two different classifiers.

Training DA is straightforward since we can easily define binary labels for each domain, e.g., we can use 1 as the source domain label and 0 as the target domain label. Adversarial training over the domain labels can achieve domain alignment. For CA, however, it is difficult to implement as we do not have target labels, hence the target class features are completely unknown to us, thus leading naively using adversarial training over each class impossible. The existing way well supported by theory to perform CA Saito et al. [2018] is to indirectly align class features by devising two different classifier hypotheses. The two classifiers have to be well trained on the source domain and are able to classify different classes in the source domain with different decision boundaries. Then considering

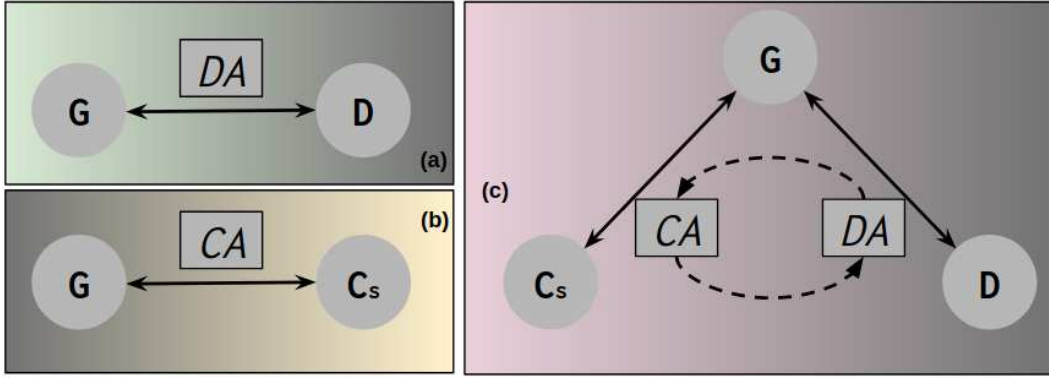


Fig. 3: Different game structures. (a) Bilateral game for DA; (b) Bilateral game for CA; (c) Our proposed pseudo-trilateral game structure (PTGS).

the shift between source and target domain, the trained two classifiers might have disagreements on target domain classes. Note since the two classifiers are already well trained on the source domain, the agreements of the two classifiers represent those features in the target domain that are close to the source domain, while in contrast, the features where disagreements happen indicate that there is a large shift between the source and the target. We use the disagreements to approximate the distance between the source and the target. If we are able to minimize the disagreements of the two classifiers, then features of each class between source and target will be enforced to be well aligned.

### 3.3 Adversarial Training

A standard way to achieve the alignment for deep models is to use the adversarial training method, which is also used in Generative Adversarial Networks (GANs) Goodfellow et al. [2014]. Therefore we explain the key concepts of adversarial training using the example of GANs.

GAN is proposed to learn the distribution  $p_r$  of a set of given data  $\{\mathbf{x}\}$  in an adversarial manner. The architecture consists of two networks - a generator  $G$ , and a discriminator  $D$ . The  $G$  is responsible for generating fake data (with distribution  $p_g$ ) from random noises  $\mathbf{z} \sim p_z$  to fool the discriminator  $D$  that is instead to accurately distinguish between the fake data and the given data. Optimization of a GAN involves a mini-maximization over a joint loss for  $G$  and  $D$ .

$$\min_G \max_D V(G, D)$$

$$V(G, D) = \mathbb{E}_{\mathbf{x} \sim p_r} \log [D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z} \log [1 - D(G(\mathbf{z}))], \quad (10)$$

where we use 1 as the real label and 0 as the fake label. Training with Eq. (10) is a bilateral game where the distribution  $p_g$  is aligned with the distribution  $p_r$ .

## 4 Methodology

In this work we investigate the relationship between the  $\mathcal{UB}_1$  (Eq. (5)) and  $\mathcal{UB}_2$  (Eq. (9)) and prove that  $\mathcal{UB}_1$  turns out to be an upper bound of  $\mathcal{UB}_2$ , meaning DA can be a necessary constraint to CA. This is also consistent with our intuition: DA aligns features globally in a coarse way while CA aligns features locally in a finer way. Constraining CA with DA is actually a coarse-to-fine process. We use the coarse alignment to reduce the general domain distance and regularize the adversarial objective for the fine alignment into a proper level. DA is shown to be a zero-sum game between a feature extractor and a domain discriminator Ganin et al. [2016], while CA is proved to be a zero-sum game between a feature extractor and two classifiers Saito et al. [2018]. Both DA and CA are bilateral games, see (a) and (b) in Fig. 3. In this work, we propose a novel concept, *pseudo-trilateral game structure* (PTGS), for efficiently integrating game structures of DA and CA, see (c) in Fig. 3. Three players are involved in the proposed PTGS, a feature extractor  $G$ , a domain discriminator  $D$ , and a family of classifiers  $C_s$ . The game between  $G$  and  $C_s$  is the CA while the game between  $G$  and  $D$  is the DA. According to the identified relation in Eq. (11), the two upper bounds  $\mathcal{UB}_1$  and  $\mathcal{UB}_2$  need to use the same feature, hence we connect the domain alignment and class alignment using a shared feature extractor. Both  $D$  and  $C_s$  are trying to adjust the  $G$  such that the features between source and target generated from  $G$  could be well aligned globally and locally. There is no game between the  $C_s$  and the  $D$ . The DA and CA in the PTGS are performed in an alternative way during training.

Notations used in this paper are explained as follows. we denote the segmentation model  $h$  as  $h^{\theta, \phi}(x) = C^{\theta}(G^{\phi}(x))$  which consists of a feature extractor  $G^{\phi}$  parameterized by  $\phi$  and a classifier  $C^{\theta}$  parameterized by  $\theta$ , and  $x$  is a sample from  $\mathcal{U}_S$  or  $\mathcal{U}_T$ . If multiple classifiers are used, we denote the  $j^{th}$  classifier as  $C_j$ . We denote the discriminator as  $D^{\psi}$  parameterized by  $\psi$ .

#### 4.1 Bounds Relation

We start by examining the relationship between the DA and the CA from the perspective of target error bound. We propose to use this relation to improve the segmentation performance of class alignment, which is desired for dense prediction tasks. We provide the following theorem:

**Theorem 2** If we assume there is a hypothesis space  $\mathcal{H}$  for segmentation model  $h^{\theta, \phi}$  and a hypothesis space  $\mathcal{H}_D$  for domain classifiers  $D^{\psi}$ , and  $\mathcal{H}\Delta\mathcal{H} \subset \mathcal{H}_D$ , then we have

$$\begin{aligned} \epsilon_T(h) &\leq \hat{\mathcal{U}}\mathbb{B}_2 \leq \hat{\mathcal{U}}\mathbb{B}_1, \\ \hat{\mathcal{U}}\mathbb{B}_1 &= \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}_D}(\mathcal{D}_S, \mathcal{D}_T) + \lambda, \\ \hat{\mathcal{U}}\mathbb{B}_2 &= \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) + \lambda. \end{aligned} \quad (11)$$

The proof of this theorem is provided in Section. 7.1.

Essentially, we limit the hypothesis space  $\mathcal{H}$  and  $\mathcal{H}_D$  in Eq. (11) into the space of deep neural networks. Directly optimizing over  $\hat{\mathcal{U}}\mathbb{B}_2$  might be hard to converge since  $\hat{\mathcal{U}}\mathbb{B}_2$  is a tighter upper bound for the prediction error on the target domain. The bounds relation in Eq. (11) shows that the  $\hat{\mathcal{U}}\mathbb{B}_1$  is an upper bound of  $\hat{\mathcal{U}}\mathbb{B}_2$ . This provides us a clue to improve the training process of class alignment, i.e., *the domain alignment can be a global constraint and narrow down the searching space for the class alignment*. This also implies that integrating the domain alignment and class alignment might boost the training efficiency as well as the prediction performance of UDA. This inspires us to design a new model, and we are explaining next in details about our model structures, losses and training process.

#### 4.2 CALI Structure

Following our proposed PTGS and the identified relation in Eq. (11), we design the structure of our CALI model as shown in Fig. 4. Four networks are involved, a shared feature extractor  $G^{\phi}$ , a domain discriminator  $D^{\psi}$  and two classifiers  $C_1^{\theta}$  and  $C_2^{\theta}$ .  $f$  represents the shared features;  $p_1/o_1$  and  $p_2/o_2$  are the probability/class predictions for  $C_1^{\theta}$  and  $C_2^{\theta}$ , respectively;  $S/T$

represent the source domain label (1) and target domain label (0); and  $L_1$  represents the  $L_1$  distance measure between two probability distributions. The one-way solid arrows indicate the forward propagation of the data flow while the two-way dashed arrows indicate losses are generated. The red arrows represent the source-related data while the blue ones represent the target-related data. Furthermore, as defined in Eq. (8),  $h$  and  $h'$  are two different hypotheses, thus we have to ensure the classifiers  $C_1^{\theta}$  and  $C_2^{\theta}$  are different. Note that the supervision signal from the source domain ( $\mathbf{y}_S$  in Fig. 4) is used to train  $C_1^{\theta}$  and  $C_2^{\theta}$  because both classifiers are expected to generate correct decision boundaries on the source domain. No label from the target domain is used during training.

#### 4.3 CALI Losses

We denote raw images from the source or target domain as  $\mathbf{x}$ , and the label from the source domain as  $\mathbf{y}$ . We use semantic labels in the source domain to train all of the nets, but the domain discriminator, in a supervised way (see the solid red one-way arrow in Fig. 4). We need to minimize the supervised segmentation loss since Eq. (11) and other related Eqs suggest that the source prediction error is also part of the upper bound of target error. In this section, we omit superscripts for all model notations. The supervised segmentation loss for training CALI is defined as

$$\begin{aligned} \mathcal{L}_{seg}(G, C_1, C_2) &= \frac{1}{2} \left( \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_S} [-\mathbf{y} \log(C_1(G(\mathbf{x}))) + \right. \\ &\quad \left. \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_S} [-\mathbf{y} \log(C_2(G(\mathbf{x}))) \right] \\ &= -\frac{1}{2} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}_S} [\mathbf{y}_S \log((C_1(G(\mathbf{x})) \odot (C_2(G(\mathbf{x}))))], \end{aligned} \quad (12)$$

where  $\odot$  represents the element-wise multiplication between two tensors.

To perform domain alignment, we need to define the joint loss function for  $G$  and  $D$

$$\mathcal{V}_1(G, D) = -(\mathcal{CE}_S(\mathbf{x}) + \mathcal{CE}_T(\mathbf{x})), \quad (13)$$

where no segmentation labels but domain labels are used, and we use the standard cross-entropy to compute the domain classification loss for both source ( $\mathcal{CE}_S(\mathbf{x})$ ) and target data ( $\mathcal{CE}_T(\mathbf{x})$ ). We have

$$\begin{aligned} \mathcal{CE}_S(\mathbf{x}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_S} [\mathcal{CE}([1, 0]^T, [D(G(\mathbf{x})), 1 - D(G(\mathbf{x}))]^T)] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_S} [-\log(D(G(\mathbf{x})))] \end{aligned}$$



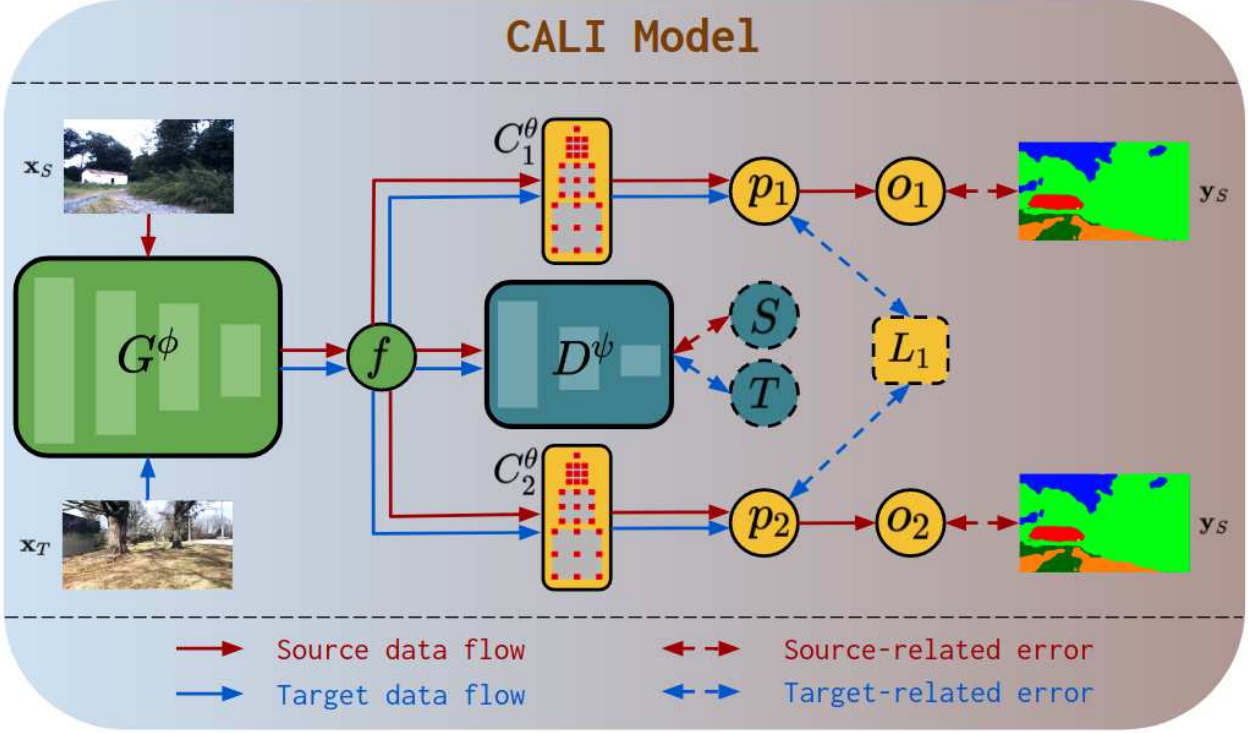


Fig. 4: CALI network structure.  $x_S$  and  $x_T$  represent the source and target images, respectively.  $y_S$  is the label of the source image. See Section 4.2 for more details.

(14)

where  $\mathbf{w}_1$  and  $\mathbf{w}_2$  are the weight vectors of  $C_1$  and  $C_2$ , respectively.

and

$$\begin{aligned} \mathcal{CE}_T(\mathbf{x}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_T} [\mathcal{CE}([0, 1]^T, [D(G(\mathbf{x})), 1 - D(G(\mathbf{x}))]^T)] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_T} [-\log(1 - D(G(\mathbf{x})))] \end{aligned}$$

(15)

Note we include  $G$  in Eq. (14) since both the source data and target data are passed through the feature extractor. This is different than standard GAN, where the real data is directly fed to  $D$ , without passing through the generator.

To perform class alignment, we need to define the joint loss function for  $G$ ,  $C_1$ , and  $C_2$

$$\mathcal{V}_2(G, C_1, C_2) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_T} [d(C_1(G(\mathbf{x})), C_2(G(\mathbf{x})))] \quad (16)$$

where  $d(\cdot, \cdot)$  is the distance measure between two distributions from the two classifiers. In this paper, we use the same  $L_1$  distance in Saito et al. [2018] as the measure, thus  $d(p, q) = \frac{1}{K} |p - q|_1$ , where  $p$  and  $q$  are two distributions and  $K$  is the number of label classes.

To prevent  $C_1$  and  $C_2$  from converging to the same network throughout the training, we use the cosine similarity as a weight regularization to maximize the difference of the weights from  $C_1$  and  $C_2$ , i.e.,

$$\mathcal{WR}(C_1, C_2) = \frac{\mathbf{w}_1 \cdot \mathbf{w}_2}{\|\mathbf{w}_1\| \|\mathbf{w}_2\|} \quad (17)$$

#### 4.4 CALI Training

We integrate the training processes of domain alignment and class alignment to systematically train our CALI model. To be consistent with Eq. (11), we adopt an iterative mechanism that alternates between domain alignment and class alignment. Our training process is pseudo-coded in Algorithm 4.1.

Note the adversarial training order of  $\mathcal{V}_1$  in Algorithm 4.1 is  $\max_{\psi_D} \min_{\phi_G}$ , instead of the  $\min_{\phi_G} \max_{\psi_D}$ , meaning in each training iteration we first train the feature extractor and then the discriminator. The reason for this order is because we empirically find that the feature from  $G$  is relatively easy for  $D$  to discriminate, hence if we train  $D$  first, then the  $D$  might become an accurate discriminator in the early stage of training and there will be no adversarial signals for training  $G$ , thus making the whole training fail. The same order applies to training of the pair of  $G$  and  $C$ s with  $\mathcal{V}_2$ .

#### 4.5 Visual Planner

We design a visual receding horizon planner to achieve feasible visual navigation by combining the learned im-

**Algorithm 4.1** CALI Training Process

---

```

1: Input: Source dataset  $\mathcal{U}_s$ ; Target dataset  $\mathcal{U}_t$ ; Initial model
    $G, C_1, C_2$  and  $D$ ; Maximum iterations  $M$ ; Iteration interval  $I$ .
2: Output: Updated model parameters  $\phi_G, \theta_{C_1}, \theta_{C_2}$  and  $\psi_D$ .
3: Initialization:  $is\_domain=True$ ;  $is\_class=False$ ;
4: for  $m \leftarrow 1$  to  $M$  do
5:   if  $m\%I == 0$  and  $m \neq 0$  then
6:      $is\_domain = \text{not } is\_domain$ ;
7:      $is\_class = \text{not } is\_class$ ;
8:     /* Eq. (12) */
9:      $\min_{\phi_G, \theta_{C_1}, \theta_{C_2}} \mathcal{L}_{seg}(G, C_1, C_2)$ ;
10:    /* Eq. (17) */
11:     $\min_{\theta_{C_1}, \theta_{C_2}} \mathcal{WR}(C_1, C_2)$ ;
12:    if  $is\_domain$  then
13:      /* Eq. (13) */
14:       $\max_{\psi_D} \min_{\theta_G} \mathcal{V}_1(G, D)$ ;
15:    if  $is\_class$  then
16:      /* Eq. (16) */
17:       $\max_{\theta_{C_1}, \theta_{C_2}} \min_{\phi_G} \mathcal{V}_2(G, C_1, C_2)$ ;
18: Return  $\phi_G, \theta_{C_1}, \theta_{C_2}$  and  $\psi_D$ ;

```

---




Fig. 5: Images from the camera on a ground robot in a real indoor environment. *Left column:* RGB images; *Middle column:* Surface normal images; *Right column:* Predicted segmentation images (yellow: navigable; green: non-navigable) from our proposed PSV-Nets. Blue dots are sampled points from the predicted boundary. The orange dash line represents an approximated convex function shape while the red dash line is an approximated concave function shape.

age segmentation. Specifically, first we compute a library of motion primitives Howard and Kelly [2007], Howard et al. [2008]  $\mathcal{M} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n\}$  where each  $\mathbf{p}_* = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m\}$  is a single primitive. We use  $\mathbf{x}_* = [x \ y \ \psi]^T$  to denote a robot pose. Then we project the motion primitives to the image plane and compute the navigation cost function for each primitive based on the evaluation of collision risk in image space and target progress. Finally, we select the primitive with minimal cost to execute. The trajectory selection problem can be defined as:

$$\mathbf{p}_{optimal} = \underset{\mathbf{p}}{\operatorname{argmin}} w_1 \cdot C_c(\mathbf{p}) + w_2 \cdot C_t(\mathbf{p}), \quad (18)$$

where  $C_c(\mathbf{p}) = \sum_j^m c_c^j$  and  $C_t(\mathbf{p}) = \sum_j^m c_t^j$  are the collision cost and target cost of one primitive  $\mathbf{p}$ , and  $w_1, w_2$  are corresponding weights, respectively.




Fig. 6: Illustrations in Gazebo simulation. (a) Surface normal image; (b) Binary navigable space segmentation (the same color meaning as Fig. ??) and different EDFs with varying scale factors of (c)  $\alpha = 0.05$ , (d)  $\alpha = 0.25$ , (e)  $\alpha = 0.55$  and (f)  $\alpha = 1.00$ .

#### 4.5.1 Collision Avoidance

In this work, we propose a Scaled Euclidean Distance Field (SEDF) for obstacle avoidance. Conventional collision avoidance is usually conducted in the map space Gao et al. [2017], Han et al. [2019], where an occupancy map and the corresponding Euclidean Signed Distance Field (ESDF) have to be provided in advance or constructed incrementally in real time. In this work instead we eliminate this expensive map construction process and evaluate the collision risk directly in the image space. Specifically, we first compute a SEDF image  $E'(S)$  based on an edge map  $Edge(S)$  detected in the learned binary segmentation  $S(I)$ , where  $I$  is the input image to our PSV-Nets. We then project the motion primitives from the map space to the image space and evaluate all primitives' projections in  $E'(S)$ .

To perform obstacle avoidance in image space, we have to detect the obstacle boundary in  $Edge(S)$ . To achieve this, we propose to categorize the edges in  $Edge(S)$  into two classes, *Strong Obstacle Boundaries (SOBs)* and *Weak Obstacle Boundaries (WOBs)*. We treat the boundary from the binary segmentation as a function of a single variable in image space, and we use the twin notions of convexity and concavity of functions to define the SOBs and WOBs, respectively. SOBs mean obstacles are near to the robot (e.g., the random furniture closely surrounding the robot) and they cause the boundaries to exhibit an approximated *ConVex Function Shape (CVFS)* (see Fig. 5(c)). WOBs indicate obstacles are far (compared with SOBs) from the robot (e.g., the wall boundaries that the robot makes large clearance from) and they typically make the boundaries reveal an approximated *ConCave Function Shape (CCFS)* (see Fig. 5(f)). In this work, we only consider the obstacles with SOBs and adopt a straightforward way (as in Eq. (19)) to detect boundary segments in  $Edge(S)$  with CVFS. We use a points set  $\Omega$  to represent the




Fig. 7: Qualitative results on adaptation GTA5→Cityscapes. Results of our proposed model is listed in the last second column. GT represents the ground-truth labels.

boundary segments:

$$\Omega = \{(u, v), (u, v) \in \text{Edge}(S) \text{ and } v > v_{thres}\}, \quad (19)$$

where  $(u, v)$  are the coordinates in the image frame, as shown in Fig. 6(a), and  $v_{thres}$  is a pre-defined value for evaluating the boundary convexity. If we use  $\partial\Omega$  to denote the boundary of a set, then in our case we have  $\partial\Omega = \Omega$ . Then the definition of an EDF is:

$$\begin{aligned} E[u, v] &= d((u, v), \partial\Omega) \\ d(x, \partial\Omega) &:= \inf_{y \in \partial\Omega} d(x, y), \end{aligned} \quad (20)$$

where  $d(x, y) = \|x - y\|$  is the Euclidean distance between vectors  $x$  and  $y$ .

However, directly computing an EDF using Eq. (20) in the image space will propagate the obstacles' gradients to the whole image space, which might cause the planning evaluation space to be too limited. To address this, we introduce a scale factor  $\alpha$  to compute a corrected version of EDF:

$$E'[u, v] = \begin{cases} E[u, v], & \text{if } E[u, v] \leq \alpha \cdot d_{max} \\ \alpha \cdot d_{max}, & \text{otherwise} \end{cases} \quad (21)$$

where  $\alpha \in [0, 1]$ , and  $d_{max} = \max_{u \in \mathcal{U}, v \in \mathcal{V}} E[u, v]$ , where  $\mathcal{U}$  and  $\mathcal{V}$  are the rows and columns index sets, respectively.

Some examples of  $E'$  with different  $\alpha$  values can be seen in Fig. 6.

Assuming  $\mathbf{x}^j$  is the  $j^{th}$  pose in one primitive and its image coordinates are  $(u^j, v^j)$ , then the collision risk for  $\mathbf{x}^j$  is

$$c_c^j = E'[u^j, v^j]. \quad (22)$$

#### 4.5.2 Target Progress

To evaluate target progress during the navigation progress, we propose to use the distance on  $SE(3)$  as the metric. We define three types of frames: world frame  $F_w$ , primitive pose frame  $F_{pj}$ , and goal frame  $F_g$ . The transformation of  $F_{pj}$  in  $F_w$  is denoted as  $\mathbf{T}_{wpj}$  while that of  $F_g$  in  $F_w$  is  $\mathbf{T}_{wg}$ . A typical approach to represent the distance is to split a pose into a position and an orientation and define two distances on  $\mathbb{R}^3$  and  $SO(3)$ . Then the two distances can be fused in a weighted manner with two strictly positive scaling factors  $a$  and  $b$  and with an exponent parameter  $p \in [1, \infty]$  Brégier et al.



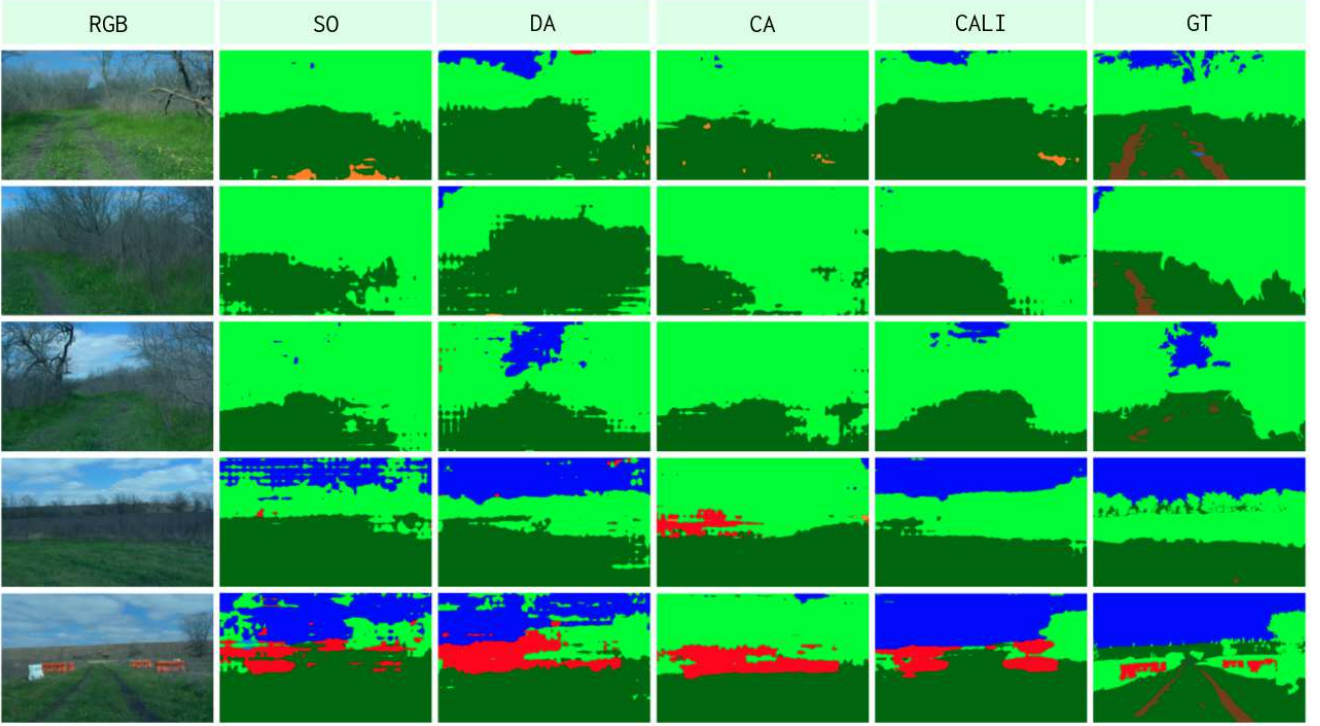


Fig. 8: Qualitative results on adaptation RUGD→RELLIS. Results of our proposed model is listed in the last second column. GT represents the ground-truth labels.

[2018]:

$$d(\mathbf{T}_{wpj}, \mathbf{T}_{wg}) = \left[ a \cdot d_{rot}(\mathbf{R}_{wpj}, \mathbf{R}_{wg})^p + b \cdot d_{trans}(\mathbf{t}_{wpj}, \mathbf{t}_{wg})^p \right]^{1/p}. \quad (23)$$

We use the Euclidean distance as  $d_{trans}(\mathbf{t}_{wpj}, \mathbf{t}_{wg})$ , the Riemannian distance over  $SO(3)$  as  $d_{rot}(\mathbf{R}_{wpj}, \mathbf{R}_{wg})$  and set  $p$  as 2. Then the distance (target cost) between two transformation matrices can be defined Park [1995] as:

$$c_t^j = d(\mathbf{T}_{wpj}, \mathbf{T}_{wg}) = \left[ a \cdot \|\log(\mathbf{R}_{wpj}^{-1} \mathbf{R}_{wg})\|^2 + b \cdot \|\mathbf{t}_{wpj} - \mathbf{t}_{wg}\|^2 \right]^{1/2}. \quad (24)$$

## 5 Experiments

### 5.1 Datasets

We evaluate CALI together with several baseline methods on a few challenging domain adaptation scenarios, where several public datasets, e.g., GTA5 Richter et al.

[2016], Cityscapes Cordts et al. [2016], RUGD Wigness et al. [2019], REllIS Jiang et al. [2020], as well as a small self-collected dataset, named MESH (see the first column of Fig. 9), are investigated. The GTA5 dataset contains 24966 synthesized high-resolution images in the urban environments from a video game and pixel-wise semantic annotations of 33 classes. The Cityscapes dataset consists of 5000 finely annotated images whose label is given for 19 commonly seen categories in urban environments, e.g., road, sidewalk, tree, person, car, etc. The RUGD and REllIS are two datasets that aim to evaluate segmentation performance in off-road environments. The RUGD and the REllIS contain 24 and 20 classes with 8000 and 6000 images, respectively. RUGD and REllIS cover various scenes like trails, creeks, parks, villages, and puddle terrains. Our dataset, MESH, includes features like grass, trees (particularly challenging in winter due to foliage loss and monochromatic colors), mulch, etc. It helps us to further validate the performance of our proposed model for traversability prediction in challenging scenes, particularly the off-road environments.



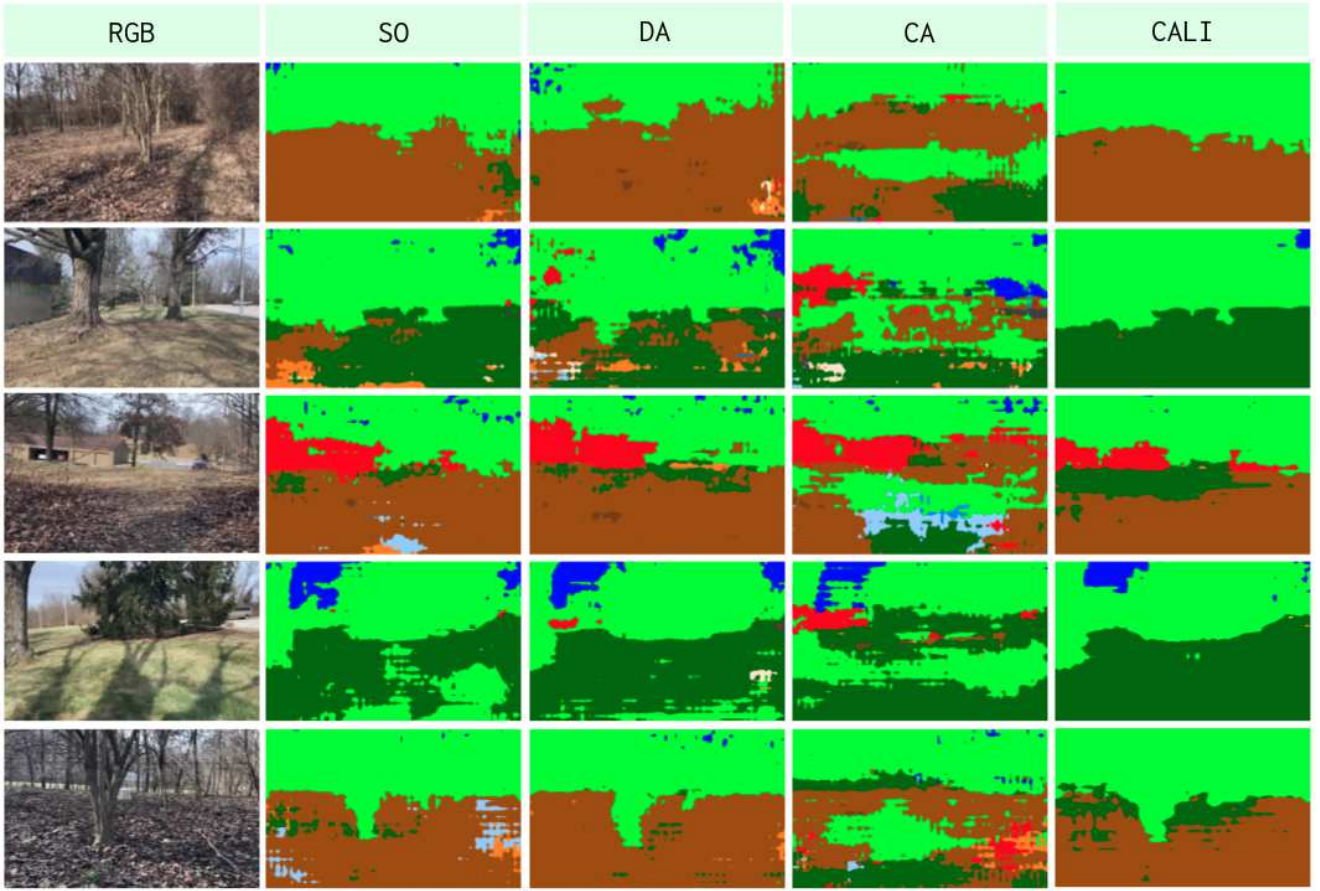


Fig. 9: Qualitative results on adaptation RUGD→MESH. Results of our proposed model is listed in the last column.

Table 1: Quantitative comparison of different methods in UDA of GTA5→Cityscapes. mIoU\* represents the average mIoU over all of classes.

Class	SO	DA	CA	CALI
Road	38.86	52.80	<b>78.56</b>	75.36
Sidewalk	17.47	18.95	2.79	<b>27.12</b>
Building	63.60	61.73	43.51	<b>67.00</b>
Sky	58.08	54.35	46.59	<b>60.49</b>
Vegetation	67.21	64.69	41.48	<b>67.50</b>
Terrain	7.63	7.04	8.37	<b>9.56</b>
Person	<b>16.89</b>	15.45	13.48	15.03
Car	30.32	43.41	31.64	<b>52.25</b>
Pole	11.61	<b>12.38</b>	9.68	11.91
mIoU*	34.63	36.76	30.68	<b>42.91</b>

## 5.2 Implementation Details

To be consistent with our theoretical analysis, the implementation of CALI only adopts the necessary indications by Eq. (11). First, Eq. (11) requires that the input of the two upper bounds (one for DA and the

other one for CA) should be the same. Second, nothing else but only domain classification and hypotheses discrepancy are involved in Eq. (11) and other related analyses (Eq. (3) - Eq. (9)). Accordingly, we strictly follow the guidance of our theoretical analyses. First, CALI performs DA in the intermediate-feature level ( $f$  in Fig. 4), instead of the output-feature level used in Vu et al. [2019]. Second, we exclude the multiple additional tricks, e.g., entropy-based and multi-level features based alignment, and class-ratio priors in Vu et al. [2019] and multi-steps training for feature extractor in Saito et al. [2018]. We also implement baseline methods without those techniques for a fair comparison. To avoid possible degraded performance brought by a class imbalance in the used datasets, we regroup those rare classes into classes with a higher pixel ratio. For example, we treat the building, wall, and fence as the same class; the person and rider as the same class in the adaptation of GTA5→Cityscapes. In the adaptation of RUGD→RELLIS, we treat the tree, bush, and log as the same class, and the rock and rockbed as the same

Table 2: Quantitative comparison of different methods in UDA of RUGD→RELLIS. mIoU\* is the average mIoU over all of classes.

Class	SO	DA	CA	CALI
Dirt	0.00	0.53	<b>3.23</b>	0.01
Grass	64.78	61.63	65.35	<b>67.08</b>
Tree	40.79	45.93	41.51	<b>55.80</b>
Sky	45.07	67.00	2.31	<b>72.99</b>
Building	10.90	<b>12.29</b>	10.91	10.28
mIoU*	32.31	37.48	24.66	<b>41.23</b>

class. Details about remapping can be seen in Fig. 15 and Fig. 16 in Section. 7.2.

We use the PyTorch [Paszke et al. \[2019\]](#) framework for implementation. Training images from source and target domains are cropped to be half of their original image dimensions. The batch size is set to 1 and the weights of all batch normalization layers are fixed. We use the ResNet-101 [He et al. \[2016\]](#) pretrained on ImageNet [Deng et al. \[2009\]](#) as the model  $G$  for extracting features. We use the ASPP module in DeepLabV2 [Chen et al. \[2017a\]](#) as the structure for  $C_1$  and  $C_2$ . We use the similar structure in [Radford et al. \[2015\]](#) as the discriminator  $D$ , which consists of 5 convolution layers with kernel  $4 \times 4$  and with channel size  $\{64, 128, 256, 512, 1\}$  and stride of 2. Each convolution layer is followed by a Leaky-ReLU [Maas et al. \[2013\]](#) parameterized by 0.2, but only the last convolution layer is followed by a Sigmoid function. During the training, we use SGD [Bottou \[2010\]](#) as the optimizer for  $G, C_1$  and  $C_2$  with a momentum of 0.9, and use Adam [Kingma and Ba \[2014\]](#) to optimize  $D$  with  $\beta_1 = 0.9, \beta_2 = 0.99$ . We set all SGD optimizers a weight decay of  $5e-4$ . The initial learning rates of all SGDs for performing domain alignment are set to  $2.5e-4$  and the one of Adam is set as  $1e-4$ . For class alignment, the initial learning rate of SGDs is set to  $1e-3$ . All of the learning rates are decayed by a poly learning rate policy, where the initial learning rate is multiplied by  $(1 - \frac{iter}{max\_iters})^{power}$  with  $power = 0.9$ . All experiments are conducted on a single Nvidia Geforce RTX 2080 Super GPU.

### 5.3 Comparative Studies

We present comparative experimental results of our proposed model, CALI, compared to different baseline methods – Source-Only (SO) method, Domain-Alignment (DA) [Vu et al. \[2019\]](#) method, and Class-Alignment [Saito et al. \[2018\]](#) method. Specifically, we first perform evaluations on a sim2real UDA in city-like environments, where the source domain is represented by GTA5 while

the target domain is the Cityscapes. Then we consider a transfer of real2real in forest environments, where the source domain and target domain are set as RUGD and RELLIS, respectively. All models are trained with full access to the images and labels in the source domain and with only access to the images in the target domain. The labels in target datasets are only used for evaluation purposes. Finally, we further validate our model performance for adapting from RUGD to our self-collected dataset MESH.

To ensure a fair comparison, all the methods use the same feature extractor  $G$ ; both DA and CALI have the same domain discriminator  $D$ ; both CA and CALI have the same two classifiers  $C_1$  and  $C_2$ . We also use the same optimizers and optimization-related hyperparameters if any is used for models under comparison.

We use the mean of Intersection over Union (mIoU) as the metric to evaluate each class and overall segmentation performance on testing images. IoU is computed as  $\frac{n_{tp}}{n_{tp} + n_{fp} + n_{fn}}$ , where  $n_{tp}, n_{tn}, n_{fp}$  and  $n_{fn}$  are true positive, true negative, false positive and false negative, respectively.

#### 5.3.1 GTA5→Cityscapes

Quantitative comparison results of GTA5→Cityscapes are shown in Table. 1, where segmentations are evaluated on 9 classes (as regrouped in Fig. 15). Our proposed method has significant advantages over multiple baseline methods for most categories and overall performance (mIoU\*).

In our testing case, SO achieves the highest score for the class person even without any domain adaptation. One possible reason for this is the deep features of the source person and the target person from the model solely trained on source domain, are already well-aligned. If we try to interfere this well-aligned relation using *unnecessary* additional efforts, the target prediction error might be increased (see the mIoU values of the person from the other three methods). We call this phenomenon as *negative transfer*, which also happens to other classes if we compare SO and DA/CA, e.g., sidewalk, building, sky, vegetation, and so on. In contrast, CALI maintains an improved performance compared to either SO or DA/CA. We validate our analytical method for DA and CA (Section 4.1) by a comparison between CALI and baselines. This indicates either single DA or CA is problematic for semantic segmentation, particularly when we strictly follow what the theory supports and do not include any other training tricks (that might increase the training complexity and make the training unstable). This implies that integration of DA and CA is beneficial to each other with




Fig. 10: Target discrepancy changes during training process of (a) GTA5→Cityscapes; (b) RUGD→RELLIS; and (c) RUGD→MESH.




Fig. 11: Using minmax can cause the collapse of training.




Fig. 12: An example of collapsed trained model using minmax.

significant improvements, and more importantly, CALI is well theoretically supported, and the training process is easy and stable.

Fig. 7 shows the examples of qualitative comparison for UDA of GTA5→Cityscapes. We find that CALI prediction is less noisy compared to the baselines methods as shown in the second and third columns (sidewalk or car on-road), and shows better completeness (part of the car is missing, see the fourth column).

### 5.3.2 RUGD→RELLIS

We show quantitative results of RUGD→RELLIS in Table. 2, where only 5 classes<sup>‡</sup> are evaluated. It shows the same trend as Table. 1. Both tables show that CA has the negative transfer (compared with SO) issue for either sim2real or real2real UDA. However, if we con-

strain the training of CA with DA, as in our proposed model, then the performance will be remarkably improved. Some qualitative results are shown in Fig. 8.

### 5.3.3 RUGD→MESH

Our MESH dataset contains only unlabeled images that restrict us to show only a qualitative comparison for the UDA of RUGD→MESH, as shown in Fig. 9. We have collected data in winter forest environments, which are significantly different than the images in the source domain (RUGD) - collected in a different season, e.g., summer or spring. These cross-season scenarios make the prediction more challenging. However, it is more practical to evaluate the UDA performance of cross-season scenarios, as we might have to deploy our robot at any time, even with extreme weather conditions, but our available datasets might be far from covering every season and every weather condition. From Fig. 9, we can still see the obvious advantages of our proposed CALI model over other baselines.

## 5.4 Discussions

In this section, we aim to discuss our model behaviors in more details. Specifically, first we will explain the advantages of CALI over CA from the perspective of training process. Second, we will show the vital influence of mistakenly using wrong order of adversarial training.

The most important part in CA is the discrepancy between the two classifiers, which is the only training force for the functionality of CA. It has been empirically studied in Saito et al. [2018] that the target prediction accuracy will increase as the target discrepancy

<sup>‡</sup> This is because other classes (in Fig. 16) frequently appearing in source domain (RUGD) are extremely rare in target domain (RELLIS), hence no prediction for those classes occurs especially considering the domain shift.






Fig. 13: Navigation behaviors in MESH#1 environment. The left-most column: top-down view of the environment; Purple triangle: the starting point; Blue star: the target point; We also show the segmentation (top row) and planning results (bottom row) at four different moments during the navigation, as shown from the second column to the last one.




Fig. 14: Navigation behaviors in MESH#2 environment. Same legends with Fig. 13.

is decreasing, hence the discrepancy is also an indicator showing if the training is on the right track. We compare the target discrepancy changes of CALI and our baseline CA in Fig. 10, where the curves for the three UDA scenarios are presented from (a) to (c) and we only show the data before iteration 30k. It can be seen that before around iteration 2k, the target discrepancy of both CALI and CA are drastically decreasing, but thereafter, the discrepancy of CA starts to increase. On the other hand, if we impose a DA constraint over the same CA (iteratively), leading to our proposed CALI, then the target discrepancy will be decreasing as expected. This validates that integrating DA and CA will make the training process of CA more stable, thus improving the target prediction accuracy.

As mentioned in Algorithm 1, we have to use adversarial training order of  $\max_{\psi_D} \min_{\phi_G}$ , instead of using  $\min_{\phi_G} \max_{\psi_D}$ . The reason for this is related to our designed net structure. Following the guidance of Eq. (11), we use the same input to the two classifiers and the domain discriminator, hence the discriminator has to receive the intermediate-level feature as the input. If we use the order of  $\min_{\phi_G} \max_{\psi_D}$  in CALI, then the outputs of the discriminator will be like Fig. 11(a), where the domain discriminator of CALI will quickly converge

to the optimal state and it can accurately discriminate if the feature is from source or target domain. In this case, the adversarial loss for updating the feature extractor will be near 0, hence the whole training fails, which is validated by changes of the target discrepancy curve, as shown in Fig. 11(b), where the discrepancy value is decreasing in a small amount in the first few iterations and then quickly increase to a high level that shows the training is divergent and the model is collapsed. This is also verified by the prediction results at (and after) around iteration 1k, as shown in Fig. 12, where the first row is the source images while the second row is the target images.

### 5.5 Navigation Missions

To further show the effectiveness of our proposed model for real deployments, we build a navigation system by combining the proposed CALI (trained with RUGD  $\rightarrow$  MESH set-up) segmentation model with our visual planner. We test behaviors of our navigation system in two different forest environments (named MESH#1 in Fig. 13 and MESH#2 in Fig. 14), where our navigation system shows high reliability. In navigation tasks, the



image resolution is [400, 300], and the inference time for pure segmentation inference is around 33 frame per second (FPS). However, since a complete perception system requires several post-processing steps, such as navigability definition, noise filtering, Scaled Euclidean Distance Field computation, motion primitive evaluation and so on, the response time for the whole perception pipeline (in python) is around 8 FPS without any engineering optimization. The inference of segmentation for navigation is performed on an Nvidia Tesla T4 GPU. We set the linear velocity as 0.3m/s and control the angular velocity to track the selected motion primitive. The path length is 32.26m in Fig. 13 and 28.63m in Fig. 14. Although the motion speed is slow in navigation tasks, as a proof of concept and with a very basic motion planner, the system behavior is as expected, and we have validated that the proposed CALI model is able to well accomplish the navigation tasks in unstructured environments.

## 6 Conclusion and Future Work

We present CALI, a novel unsupervised domain adaptation model specifically designed for semantic segmentation, which requires fine-grained alignments in the level of class features. We carefully investigate the relationship between a coarse alignment and a fine alignment in theory. The theoretical analysis guides the design of the model structure, losses, and training process. We have validated that the coarse alignment can serve as a constraint to the fine alignment and integrating the two alignments can boost the UDA performance for segmentation. The resultant model shows significant advantages over baselines in various challenging UDA scenarios, e.g., sim2real and real2real. We also demonstrate the proposed segmentation model can be well integrated with our proposed visual planner to enable highly efficient navigation in off-road environments.

To further improve this framework in the future, we observe one problem with our model during deployment — segmentation boundaries can jump (vary fast) when the model makes highly-frequent predictions for images from the video captured by the onboard camera. Our model predicts segmentation for each frame independently and ignores the inter-frame coherence. This sometimes leads to a non-negligible disturbance to smooth navigation, as our visual planner highly relies on the segmentation boundaries to generate the SEDF for obstacle avoidance. To further increase the practicality of our model in real deployments, incorporating temporal correlation during training or inference is necessary to help maintain stable segmentation boundaries. We leave this as our future investigations.

## 7 Appendix

### 7.1 Proof of Theorem 2

For a hypothesis  $h$ ,

$$\begin{aligned}
\epsilon_T(h) &\leq \epsilon_T(h^*) + \epsilon_T(h, h^*) \\
&= \epsilon_T(h^*) + \epsilon_S(h, h^*) - \epsilon_S(h, h^*) + \epsilon_T(h, h^*) \\
&\leq \epsilon_T(h^*) + \epsilon_S(h, h^*) + |\epsilon_T(h, h^*) - \epsilon_S(h, h^*)| \\
&\leq \epsilon_T(h^*) + \epsilon_S(h, h^*) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) \\
&\leq \epsilon_T(h^*) + \epsilon_S(h) + \epsilon_S(h^*) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) \\
&= \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) + \epsilon_S(h^*) + \epsilon_T(h^*) \\
&= \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) + \lambda \\
&= \epsilon_S(h) + \sup_{h, h' \in \mathcal{H}} |\mathbb{P}_{\mathbf{x} \sim \mathcal{D}_S} [h(\mathbf{x}) \neq h'(\mathbf{x})] - \\
&\quad \mathbb{P}_{\mathbf{x} \sim \mathcal{D}_T} [h(\mathbf{x}) \neq h'(\mathbf{x})]| + \lambda \\
&= \epsilon_S(h) + \sup_{g \in \mathcal{H}\Delta\mathcal{H}} |\mathbb{P}_{\mathbf{x} \sim \mathcal{D}_S} [g(\mathbf{x}) = 1] - \\
&\quad \mathbb{P}_{\mathbf{x} \sim \mathcal{D}_T} [g(\mathbf{x}) = 1]| + \lambda \\
&= \epsilon_S(h) + \sup_{g \in \mathcal{H}\Delta\mathcal{H}} |\mathbb{P}_{\mathbf{x} \sim \mathcal{D}_S} [g(\mathbf{x}) = 1] + \\
&\quad \mathbb{P}_{\mathbf{x} \sim \mathcal{D}_T} [g(\mathbf{x}) = 0] - 1| + \lambda \\
&\leq \epsilon_S(h) + \sup_{g \in \mathcal{H}\Delta\mathcal{H}} |\mathbb{P}_{\mathbf{x} \sim \mathcal{D}_S} [g(\mathbf{x}) = 1] + \\
&\quad \mathbb{P}_{\mathbf{x} \sim \mathcal{D}_T} [g(\mathbf{x}) = 0]| - \inf_{g \in \mathcal{H}\Delta\mathcal{H}} 1 + \lambda \\
&= \epsilon_S(h) + \sup_{g \in \mathcal{H}\Delta\mathcal{H}} |\mathbb{P}_{\mathbf{x} \sim \mathcal{D}_S} [g(\mathbf{x}) = 1] + \\
&\quad \mathbb{P}_{\mathbf{x} \sim \mathcal{D}_T} [g(\mathbf{x}) = 0]| + \lambda - 1 \\
&= \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) + 1 + \lambda - 1 \\
&= \epsilon_S(h) + \frac{1}{2}d_{\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) + \lambda,
\end{aligned} \tag{25}$$

where  $\lambda = \epsilon_S(h^*) + \epsilon_T(h^*)$  and  $h^*$  is the ideal joint hypothesis (see the **Definition 2** in Section 4.2 of Ben-David et al. [2010]).

We have the 4<sup>th</sup>, and the 8<sup>th</sup> line because of the **Lemma 3** Ben-David et al. [2010]; the 5<sup>th</sup> line because of the **Theorem 2** Ben-David et al. [2010]; the last second line because of the **Lemma 2** Ben-David et al. [2010]. We have the 11<sup>th</sup> line because  $\sup |f_1 - f_2| = \sup f_1 - \inf f_2 \leq \sup |f_1| - \inf f_2$ . ■

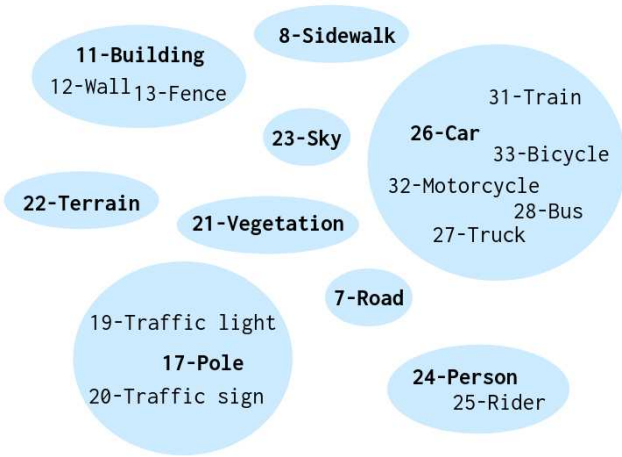


Fig. 15: Lable remapping for GTA5→Cityscapes. Name of each new group is marked as bold.

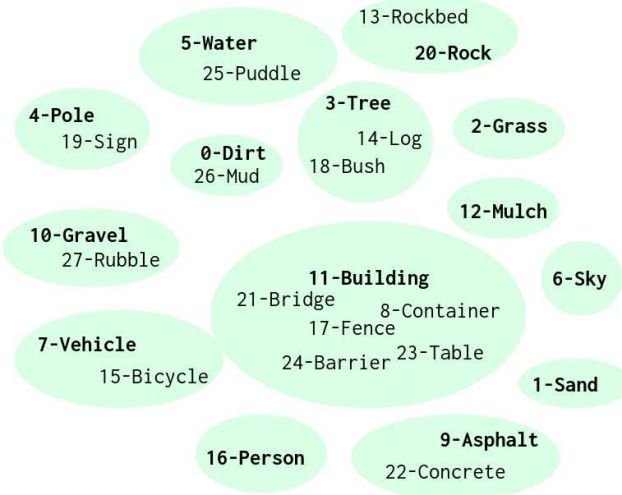


Fig. 16: Lable remapping for RUGD→RELLIS and RUGD→MESH. Name of each new group is marked as bold.

## 7.2 Remapping of Label Space

We regroup the original label classes according to the semantic similarities among classes. In GTA5 and Cityscapes, we cluster the building, wall and fence as the same category; traffic light, traffic sign and pole as the same group; car, train, bicycle, motorcycle, bus and truck as the same class; and treat the person and rider as the same one. See Fig. 15. Similarly, we also have regroupings for classes in RUGD and RELLIS, as can be seen in Fig. 16.

## 8 Statements and Declarations

### Acknowledgements and Funding

We acknowledge the supports of NSF with grant number 2006886 as well as ARL with grant number W911NF-20-2-0099. We are also grateful for the computational resources provided by the Amazon AWS Machine Learning Research Award.

### Authors Contributions

Zheng Chen contributed to the algorithm design and the theoretical analysis. The experiment validation was performed by Zheng Chen and Durgakant Pushp. The first draft of the manuscript was written by Zheng Chen and all authors commented on previous versions of the manuscript. Lantao Liu contributed to the top design and general guidance of this work. All authors read and approved the final manuscript.

### Competing Interests

The authors have no relevant financial or non-financial interests to disclose.

### Ethics Approval

This paper does not involve human or animal subjects. No ethical approval is required.

### Consent For Participation

This paper does not involve human subjects to participate in the study.

### Consent For Publication

The paper does not contain any data/figures collected from human participants. No consent is required to publish the manuscript

## References

- Somil Bansal, Varun Tolani, Saurabh Gupta, Jitendra Malik, and Claire Tomlin. Combining optimal control and learning for visual navigation in novel environments. In *Conference on Robot Learning*, pages 420–429. PMLR, 2020.
- Shai Ben-David, John Blitzer, Koby Crammer, Fernando Pereira, et al. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19:137, 2007.

- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.
- John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman. Learning bounds for domain adaptation. 2008.
- Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.
- Romain Bréquier, Frédéric Devernay, Laetitia Leyrit, and James L Crowley. Defining the pose of any 3d rigid object and an associated distance. *International Journal of Computer Vision*, 126(6):571–596, 2018.
- Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and Saurabh Gupta. Neural topological slam for visual navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12875–12884, 2020.
- Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017a.
- Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017b.
- Minghao Chen, Hongyang Xue, and Deng Cai. Domain adaptation for semantic segmentation with maximum squares loss. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2090–2099, 2019.
- Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- Fei Gao, Yi Lin, and Shaojie Shen. Gradient-based online safe trajectory generation for quadrotor flight in complex environments. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 3681–3688. IEEE, 2017.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2625, 2017a.
- Saurabh Gupta, David Fouhey, Sergey Levine, and Jitendra Malik. Unifying map and landmark based representations for visual navigation. *arXiv preprint arXiv:1712.08125*, 2017b.
- Luxin Han, Fei Gao, Boyu Zhou, and Shaojie Shen. Fiesta: Fast incremental euclidean distance fields for online motion planning of aerial robots. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4423–4430. IEEE, 2019.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Noriaki Hirose, Shun Taguchi, Fei Xia, Roberto Martin-Martin, Yasumasa Tahara, Masanori Ishigaki, and Silvio Savarese. Probabilistic visual navigation with bidirectional image prediction. *arXiv preprint arXiv:2003.09224*, 2020.
- Judy Hoffman, Dequan Wang, Fisher Yu, and Trevor Darrell. Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649*, 2016.
- Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998. PMLR, 2018.
- Thomas M Howard and Alonzo Kelly. Optimal rough terrain trajectory generation for wheeled mobile robots. *The International Journal of Robotics Research*, 26(2):141–166, 2007.

- Thomas M Howard, Colin J Green, Alonzo Kelly, and Dave Ferguson. State space sampling of feasible motions for high-performance mobile robot navigation in complex environments. *Journal of Field Robotics*, 25(6-7):325–345, 2008.
- Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. *arXiv preprint arXiv:2111.14887*, 2021.
- Peng Jiang, Philip Osteen, Maggie Wigness, and Srikanth Saripalli. Rellis-3d dataset: Data, benchmarks and analysis, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Ziwei Liu, Xiaoxiao Li, Ping Luo, Chen Change Loy, and Xiaoou Tang. Deep learning markov random field for semantic segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 40(8):1814–1828, 2017.
- Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- Yawei Luo, Liang Zheng, Tao Guan, Junqing Yu, and Yi Yang. Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2507–2516, 2019.
- Andrew L Maas, Awni Y Hannun, Andrew Y Ng, et al. Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3. Citeseer, 2013.
- Travis Manderson, Juan Camilo Gamboa Higuera, Stefan Wapnick, Jean-François Tremblay, Florian Shkurti, David Meger, and Gregory Dudek. Vision-based goal-conditioned policies for underwater navigation in the presence of obstacles. *arXiv preprint arXiv:2006.16235*, 2020.
- Ke Mei, Chuang Zhu, Jiaqi Zou, and Shanghang Zhang. Instance adaptive self-training for unsupervised domain adaptation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16*, pages 415–430. Springer, 2020.
- Chen Min, Weizhong Jiang, Dawei Zhao, Jiaolong Xu, Liang Xiao, Yiming Nie, and Bin Dai. Orfd: A dataset and benchmark for off-road freespace detection. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2532–2538. IEEE, 2022.
- Fei Pan, Inkyu Shin, Francois Rameau, Seokju Lee, and In So Kweon. Unsupervised intra-domain adaptation for semantic segmentation through self-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3764–3773, 2020.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- Frank C Park. Distance metrics on the rigid-body motions with applications to mechanism design. 1995.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12179–12188, 2021.
- Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M. Lopez. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3723–3732, 2018.
- Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Guided curriculum model adaptation and uncertainty-aware evaluation for semantic nighttime image segmentation. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019.
- Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The adverse conditions dataset with corre-



- spondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021.
- William B Shen, Danfei Xu, Yuke Zhu, Leonidas J Guibas, Li Fei-Fei, and Silvio Savarese. Situational fusion of visual representation for visual navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2881–2890, 2019.
- Yi-Hsuan Tsai, Wei-Chih Hung, Samuel Schuster, Kihyuk Sohn, Ming-Hsuan Yang, and Manmohan Chandraker. Learning to adapt structured output space for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7472–7481, 2018.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019.
- Haoran Wang, Tong Shen, Wei Zhang, Ling-Yu Duan, and Tao Mei. Classes matter: A fine-grained adversarial approach to cross-domain semantic segmentation. In *European Conference on Computer Vision*, pages 642–659. Springer, 2020.
- Maggie Wigness, Sungmin Eum, John G Rogers, David Han, and Heesung Kwon. A rugged dataset for autonomous navigation and visual perception in unstructured outdoor environments. In *International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020.
- Yi Wu, Yuxin Wu, Aviv Tamar, Stuart Russell, Georgia Gkioxari, and Yuandong Tian. Bayesian relational memory for semantic visual navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2769–2779, 2019.
- Binhui Xie, Longhui Yuan, Shuang Li, Chi Harold Liu, and Xinjing Cheng. Towards fewer annotations: Active learning via region impurity and prediction uncertainty for domain adaptive semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8068–8078, 2022.
- Enze Xie, Wenhui Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *arXiv preprint arXiv:2105.15203*, 2021.
- Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*, 2015.
- Yang Zhang, Philip David, and Boqing Gong. Curriculum domain adaptation for semantic segmentation of urban scenes. In *Proceedings of the IEEE international conference on computer vision*, pages 2020–2030, 2017.
- Youshan Zhang. A survey of unsupervised domain adaptation for visual recognition. *arXiv preprint arXiv:2112.06745*, 2021.
- Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6881–6890, 2021.
- Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305, 2018.
- Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5982–5991, 2019.