LAGUNA: LAnguage Guided UNsupervised Adaptation with structured spaces

Anxhelo Diko[1*]   Antonino Furnari[2*]   Luigi Cinque[1]   Giovanni Maria Farinella[2]

[1]La Sapienza University of Roma [2]University of Catania

{diko,cinque}@di.uniroma1.it {antonino.furnari,giovanni.farinella}@unict.it

Abstract

Unsupervised domain adaptation remains a critical challenge in enabling the knowledge transfer of models across unseen domains. Existing methods struggle to balance the need for domain-invariant representations with preserving domain-specific features, which is often due to alignment approaches that impose the projection of samples with similar semantics close in the latent space despite their drastic domain differences. We introduce LAGUNA - LAnguage Guided UNsupervised Adaptation with structured spaces, a novel approach that shifts the focus from aligning representations in absolute coordinates to aligning the relative positioning of equivalent concepts in latent spaces. LAGUNA defines a domain-agnostic structure upon the semantic/geometric relationships between class labels in language space and guides adaptation, ensuring that the organization of samples in visual space reflects reference inter-class relationships while preserving domain-specific characteristics. Remarkably, LAGUNA surpasses previous works in 18 different adaptation scenarios across four diverse image and video datasets with average accuracy improvements of +3.32% on DomainNet, +5.75% in Geo-Places, +4.77% on GeoImnet, and +1.94% mean class accuracy improvement on EgoExo4D.

1. Introduction

Domain shift challenges trained models to generalize across scenarios with differing distributions and presents a significant hurdle for supervised learning in computer vision [42]. While fine-tuning with labeled data from the target domain can mitigate this problem, obtaining such labels often proves prohibitive. Unsupervised domain adaptation (UDA) [4] offers a compelling alternative, enabling knowledge transfer to novel domains without relying on expensive labeled data, which has garnered significant attention [70], promising cost-effective solutions for real-world applications prone to domain shift. The typical UDA setting
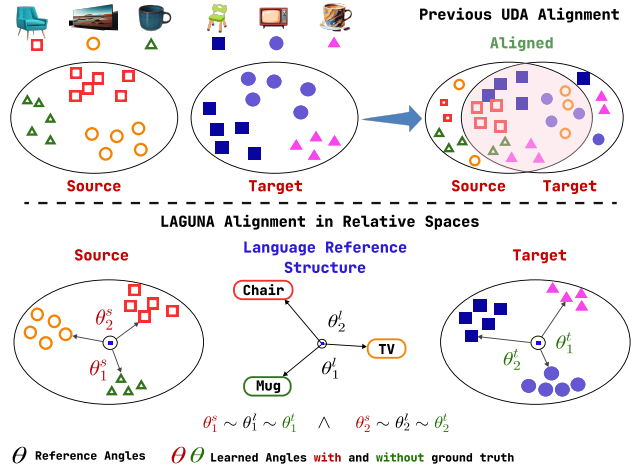
---

*Equal Contribution

Figure 1. **Top**: Existing Unsupervised Domain Adaptation approaches align source and target spaces in absolute coordinates, which may overlook the peculiarities of the distinct domains resulting in partial alignment. **Bottom**: LAGUNA aligns source (left) and target (right) spaces in relative terms, allowing data points to be distinct in absolute coordinates (e.g., circles in source and target), while θ_i^s, θ_i^t, their angles with other data points in the original spaces, match θ_i^l, those of the reference domain-agnostic structure derived from language ($\theta_i^s \sim \theta_i^l \sim \theta_i^t$), encouraging the development of similar geometric-semantic relations.

considers the availability of a labeled source domain and an unlabeled target domain. In general, source and target domains are semantically equivalent but drawn from distinct data distributions, e.g., real images versus clipart of chairs, TVs, and mugs. Thus, the main challenge of UDA is to effectively mitigate the distribution shift between domains, which is often addressed by reducing the distribution discrepancy of source and target representation spaces either by minimizing some discrepancy measure [21, 58, 63], using adversarial learning [31, 33, 40, 49, 67, 68, 77], aligning data around centroids [33], or leveraging pseudo-labels [27, 33]. These methods aim to align source and target representation spaces in a shared coordinate system, pushing feature vectors of equivalent semantic concepts close to each other in the embedding space, which may happen at the expense of the representation power of in-

dividual domain-specific spaces. For instance, bright colors and rounded shapes might be important to encode for representing a clipart, while nuanced shadows, reflections, and textures might be important to represent a real image. As a result, the aligned space may become overly generic, correctly encoding only a subset of the data (see Figure 1(top)). Recent work [46] showed that semantically equivariant representation spaces of similarly trained neural networks exhibit distinct representation spaces with matching geometrical structures. For instance, two data points (x_1, x_2) may be mapped to distinct vector pairs in the two representation spaces (v_1^1, v_2^1) and (v_1^2, v_2^2) (e.g., $||v_1^1 - v_1^2||_2 >> 0$ and $||v_2^1 - v_2^2||_2 >> 0$), while angles between the two pairs will be similar in their own representation spaces (e.g., $\angle(v_1^1, v_2^1) \sim \angle(v_1^2, v_2^2))$[1]. This suggests that pushing representation spaces to overlap in absolute coordinates, as done in current domain adaptation approaches, is not necessary to obtain equivariant representation spaces.

Based on this observation, we introduce LAGUNA - LAnguage Guided UNsupervised Adaptation with structured spaces, a novel approach to domain adaptation which guides source and target spaces to develop semantic-geometric inter-relationships reflecting the structure of a reference space (Figure 1(bottom)). As shown in recent works, language can provide a semantic space agnostic to the nuances of visual observations enabling robust zero-shot generalization [10, 54] and supporting domain robustness [14, 23, 27, 45, 66], hence we choose language to build our reference structured space in LAGUNA. This assumes the availability for both source and target samples of natural language descriptions, which can be generated from available captioning models [32] or collected from the web at a fraction of the cost of typical human labeling procedures [27]. Specifically, LAGUNA employs a 3-stage approach to structurally align the source and target representation spaces while allowing each domain to preserve its typical patterns. In Stage 1, textual class labels are mapped to a domain-agnostic reference space representing their relationships at the conceptual level. In Stage 2, a language model is trained to map textual captions to the reference latent space in order to provide pseudo-labels for target samples Finally, in Stage 3, a cross-domain classifier is trained to encourage domain-specific representations to follow the reference structure.

Experiments demonstrate the superiority of LAGUNA over existing state-of-the-art methods across 18 different domain adaptation scenarios sourced from four diverse image and video datasets, with gains of +3.32% on DomanNet [50], +5.75% on GeoPlaces [26], +4.77% on GeoImnet [26] and +1.94% mean per class accuracy on EgoExo4D [20]. We further report ablations to analyze the specific contributions of our design choices.

In sum, our main contributions are as follows: 1) we investigate the use of relative representations for UDA, showing its advantages with respect to the current absolute alignment paradigm; 2) we propose LAGUNA, a three-stage method which learns a cross-domain classifier where individual source and target spaces are distinct yet aligned to a reference structure derived from language; 3) through extensive ablations and comparisons with state-of-the-art models, we show the superiority of the proposed approach. Code available at: https://github.com/ADiko1997/LAGUNA.git

2. Related Work

2.1. UDA in Computer Vision

UDA seeks to transfer knowledge from a labeled source domain to an unlabeled target domain [2, 16, 40]. This is tackled through various approaches, prominently, discrepancy-based methods that minimize distribution differences using techniques like MMD [28, 31, 39, 58, 61, 63], and adversarial learning [4, 6, 16, 38, 40, 57, 64, 64, 67, 68]. Other approaches also investigated class-conditional distributions alignment [29, 43, 49, 68, 72], clustering similar instances across domains [9, 24, 36, 48], instance-specific adaptation[25, 60, 65], and self-training leveraging pseudo-labeling to improve target domain performance [15, 27, 36, 52, 62]. More recently, transformer-based approaches utilize patch composition for domain-invariant representations through intermediate domains [77]. Video domain adaptation tackles the unique challenges of temporal dynamics and consistency in videos [8, 56, 69] with a particular focus on adapting models across the exocentric and egocentric domains [22, 34, 53, 74]. Existing approaches sought to align source and target representation spaces in absolute coordinates, often falling short in bridging complicated forms of domain shift [26, 52]. Contrarily, LAGUNA aligns representations in relative terms, encouraging source and target visual spaces to share a similar structure as a reference space derived from language while allowing them to encode domain-specific peculiarities.

2.2. Language Guidance in Vision UDA

Vision-language models like CLIP [54] have shown promising zero-shot transfer capabilities [10]. However, when fine-tuned to a specific scenario [1, 51], they often lose the ability to generalize to new domains [30, 71]. To address this issue, recent works leveraged textual information to bridge the gap between domains [14, 18, 19, 23, 45, 66]. In particular, the seminal work of [27] proposed to use rich textual captions to provide pseudo-labels for the target domain and addressed UDA by training a joint classifier on source and target domains. Similarly to [27], we generate pseudo-labels from textual captions, but, rather than seek-
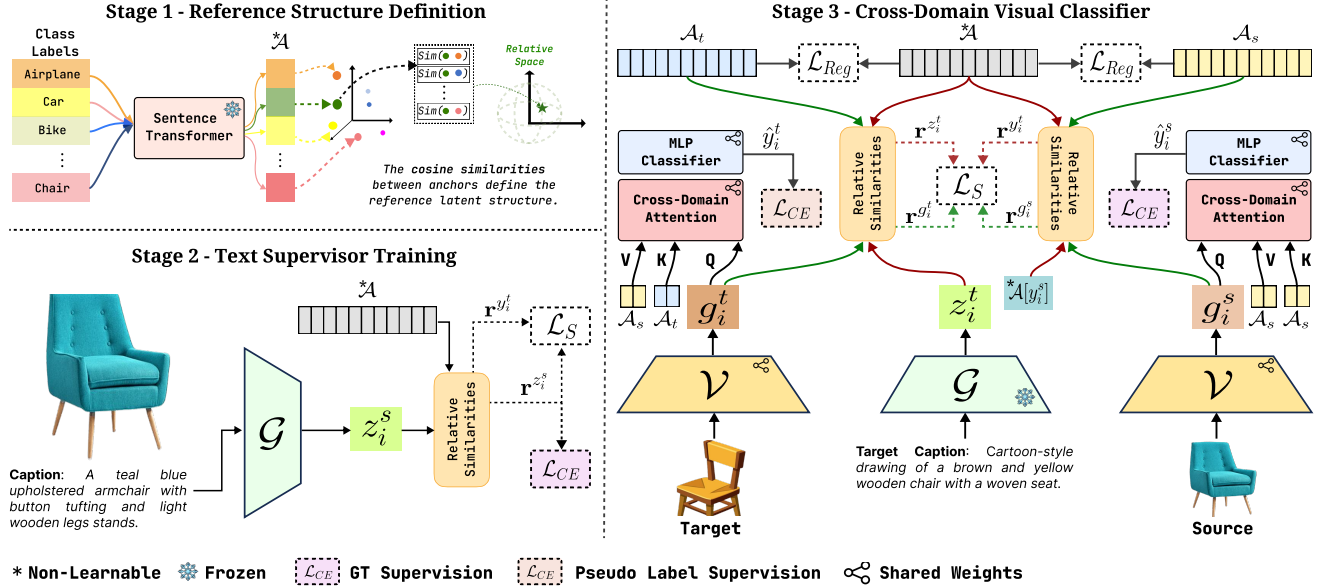
---

[1]The reader is referred to [46] for empirical verifications of such claim.

Figure 2. The 3 stages of LAGUNA. In Stage 1, a domain-agnostic semantic structure in the form of a set of anchors $\mathcal{A}$ is defined. In Stage 2, a language model $\mathcal{G}$ is trained on source captions-label pairs to provide pseudo-labels for the target domain. The model is trained to be structurally aligned to anchors $^*\mathcal{A}$ with a structural loss $\mathcal{L}_S$ and to predict class labels with a cross-entropy loss $\mathcal{L}_{CE}$. In Stage 3, a cross-domain visual classifier is trained on labeled target and unlabeled source images. A visual encoder $\mathcal{V}$ is used to represent target and source samples as g_i^t and g_i^s. The cosine similarities $\mathbf{r}^{g_i^t}$ between the target visual representation g_i^t and all learnable target anchors $\mathcal{A}_t$ are bound to be similar to the cosine similarities $\mathbf{r}^{z_i^t}$ between the textual representation of the target caption z_i^t predicted through the language model $\mathcal{G}$ and the reference domain-agnostic anchors $^*\mathcal{A}$ by means of the structural loss $\mathcal{L}_S$. A similar processing is applied to the visual representation of source sample g_i^s, the textual representation of the ground truth class $^*A[y_i^S]$ and the learnable source anchors $\mathcal{A}_s$. A regularization loss $\mathcal{L}_{Reg}$ is used to avoid collapse during anchor learning. Representations g_i^t and g_i^s are hence grounded into the learnable source anchors $\mathcal{A}_s$ through a novel Cross-Domain Attention layer taking visual representations as queries, source/target anchors as keys and source anchors as values. Classification is finally achieved through an MLP classifier supervised through a cross-entropy loss $\mathcal{L}_{CE}$ using ground truth labels in the source domain and pseudo-labels in the target domain. The visual encoder, Cross-Domain Attention, and MLP layers share weights across domains.

ing to align the visual representation spaces of source and target domains to language in an absolute reference frame, LAGUNA uses language to guide a *relative* alignment between source and target visual spaces.

2.3. Relative Encodings

Recent work [46] has shown that equivalent latent spaces of similarly trained networks tend to be misaligned in absolute terms but share similar internal geometric relationships. As a result, representing data points with relative encodings, obtained as vectors of cosine similarity values of a data point with a predefined set of anchors, allows to perform zero-shot model stitching [46], translate representation spaces across models [44] or modalities [47]. As shown in [5], predefined invariances can be incorporated into the learned representation to enable specific forms of relative representations. In [11] relative encodings were used to tackle the downstream task of action anticipation. LAGUNA builds on relative encodings to represent the semantic inter-relations between classes in a domain-agnostic

language-based reference space and support the development of aligned yet specialized source and target domains.

3. Method

In Stage 1, a language-based reference structure space is defined. We represent this structure as a set of anchors, which we will use as a reference to guide representation learning. In Stage 2, a text-supervisor language model is trained to map image/video captions to class categories in order to provide pseudo-labels for the unlabeled target domain. The language model is also trained to keep the internal text embeddings aligned to the domain-agnostic structure defined in Stage 1. Finally, in Stage 3, a cross-domain visual classifier is trained to categorize source and target samples by learning domain-specific action anchors bound to be structurally similar to the domain-agnostic structure defined in Stage 1.

3.1. Problem Setup

We follow the formulation of [27] and assume a labeled source domain $\mathcal{X}_s : \{(x_i^s, y_i^s, l_i^s)\}_{i=1}^{N_s}$, where samples x_i^s are paired both with labels y_i^s and language descriptions l_i^s, whereas the target domain $\mathcal{X}_t : \{x_i^t, l_i^t\}_{i=1}^{N_t}$ contains unlabeled samples x_i^t paired with language descriptions l_i^t. These textual descriptions can be readily obtained from associated metadata or generated using image-to-text models [32]. N_s and N_t represent the number of source and target samples, respectively, and the two domains share the same high-level semantics and number of classes N_c.

3.2. Stage 1 - Reference Structure Definition

Assuming shared classes across domains, in Stage 1 (Fig. 2), we construct a language-based, domain-agnostic reference structure in the form of a set of vectors $\mathcal{A} \in \mathbb{R}^{N_c \times D_l}$. Each vector is obtained by encoding the N_c class label names through a pre-trained *SentenceTransformer* model with output dimensionality D_l trained for semantic similarity [55]. Following literature on relative representations [46], we treat the vectors in $\mathcal{A}$ as a reference set of *anchors*. Given a vector $v \in \mathbb{R}^{D_l}$, we define its relative encoding with respect to anchors $\mathcal{A}$ as $\mathbf{r}^v = rel(v, \mathcal{A}) = [\cos(v, \mathcal{A}[1]), \ldots, \cos(v, \mathcal{A}[N_c])]$, where $\cos(\cdot, \cdot)$ is the cosine similarity (i.e., the cosine of the angle θ between two vectors), and $\mathcal{A}[i]$ is the anchor in $\mathcal{A}$ associated to class i. Intuitively, vector $\mathbf{r}^v$ encodes the geometrical/semantic relationships between vector v and the reference language anchors $\mathcal{A}$. Following this logic, the set of all reference affinities is defined as $\mathbf{r}^{\mathcal{A}} = [rel(\mathcal{A}[i], \mathcal{A}), \ldots, rel(\mathcal{A}[N_c], \mathcal{A})]$ where each class anchors $\mathcal{A}[i]$ is represented according to its relationship with respect to the other anchors in $\mathcal{A}$. During the training of the cross-domain visual classifier (Stage 3), these encodings are used to enforce the learned latent space to follow a similar structure as the one induced by $\mathcal{A}$.

3.3. Stage 2: Training of the Language Supervisor

Similar to [27], we train a language model to map image captions to semantic labels, with the aim to provide pseudo-labels for target samples, which are associated with captions but not class labels. Differently from [27], we use the language supervisor both to provide pseudo-labels for target samples, useful to train the classifier, and textual representations semantically structured as the domain-agnostic anchors $\mathcal{A}$, useful to encourage alignment to the reference structure. Hence, rather than training a regular classifier, we directly supervise a language model $\mathcal{G}$ to provide representations that are 1) geometrically aligned to $\mathcal{A}$ and 2) suitable for predicting class labels. Specifically, given a pair of source caption and corresponding label (l_i^s, y_i^s), $\mathcal{G}$ processes l_i^s to produce a vector representation z_i^s:

$$z_i^s = \mathcal{G}(l_i^s). \tag{1}$$

Next, we encourage the vector representation z_i^s to be geometrically aligned to the anchor $\mathcal{A}[y_i^s]$ corresponding to ground truth action y_i^s. To do so, we compute $\mathbf{r}^{z_i^s} = rel(z_i^s, \mathcal{A})$, the relative encoding of z_i^s, and supervise it to be similar to $\mathbf{r}^{y_i^s} = rel(\mathcal{A}[y_i^s], \mathcal{A})$, the relative encoding of the anchor $\mathcal{A}[y_i^s]$, with the following structure-preserving loss:

$$\mathcal{L}_S = |\mathbf{r}^{z_i^s} - \mathbf{r}^{y_i^s}|. \tag{2}$$

This loss encourages the encodings z_i^s of text descriptions l_i^s associated with labels y_i^s to preserve the same geometrical associations as their corresponding anchor, encouraging the latent space learned by $\mathcal{G}$ to mirror the structure defined by $\mathcal{A}$. To further favor alignment to $\mathcal{A}$, rather than employing a classification head, we predict class probabilities directly by Softmax-normalizing relative encodings:

$$p_j^{z_i^s} = \frac{e^{\mathbf{r}_j^{z_i^s}}}{\sum_k e^{\mathbf{r}_k^{z_i^s}}}. \tag{3}$$

Finally, we train the model using a combined loss $\mathcal{L}$:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{CE}(p^{z_i^s}, y_i^s) + \lambda_2 \mathcal{L}_S(\mathbf{r}^{z_i^s}, \mathbf{r}^{y_i^s}), \tag{4}$$

where $\mathcal{L}_{CE}$ is the cross-entropy loss, while λ_1 and λ_2 are hyperparameter weights to calibrate the magnitude of each loss. We refer to the label predicted by $\mathcal{G}$ from l_i^t as $\overline{y}_i^t$.

3.4. Stage 3: Cross-Domain Visual Classifier

Stage 3 aims to train a cross-domain visual classifier aligning representations extracted through a visual encoder $\mathcal{V}$ to the structure imposed by the domain-agnostic language anchors $\mathcal{A}$. We allow each domain to develop its own latent space, but we encourage both spaces to be aligned to two sets of learnable anchors $\mathcal{A}_t \in \mathbb{R}^{N_c \times D_v}$ and $\mathcal{A}_s \in \mathbb{R}^{N_c \times D_v}$ specific to the target and source domain respectively. We further ensure that the structures of $\mathcal{A}_t$ and $\mathcal{A}_s$ are aligned to the domain-agnostic, fixed anchors $\mathcal{A}$ through the supervision provided by ground truth labels y_i^s in the source domain and the text supervision provided by $\mathcal{G}$ in the target domain in the form of the predicted vector z_i^t and the predicted pseudo-label $\overline{y}_i^t$. The source and target datasets are merged ($\hat{\mathcal{X}} = \mathcal{X}_t + \mathcal{X}_s$) and used for training with a total of $M = N_s + N_t$ samples. Source samples comprise an image, caption, and label triplet (x_i^s, l_i^s, y_i^s), whereas target samples $(x_i^t, l_i^t, \overline{y}_i^t)$ replace label y_i^s with the pseudo-label $\overline{y}_i^t$. Target and source images are processed by $\mathcal{V}$ to obtain latent representations g_i^t and g_i^s, while target captions l_i^t are processed by $\mathcal{G}$ to obtain latent representations z_i^t as follows:

$$g_i^t = \mathcal{V}(x_i^t), z_i^t = \mathcal{G}(l_i^t), g_i^s = \mathcal{V}(x_i^s). \tag{5}$$

These representations and the anchors are used to achieve two main training objectives: 1) structure learning and 2) classification training.

3.4.1. Structure Learning

To ensure the learned visual representations adhere to the structure defined by $\mathcal{A}$, LAGUNA employs a supervised learning approach based on relative encodings. For target domain samples (without ground truth labels), we first compute the relative encoding $\mathbf{r}^{z_i^t} = rel(z_i^t, \mathcal{A})$. For source domain samples, we compute the relative encoding of the anchor associated to the ground truth class $\mathcal{A}[y_i^s]$: $\mathbf{r}^{y_i^s} = rel(\mathcal{A}[y_i^s], \mathcal{A})$. These encodings represent the relative positions of textual representations with respect to the reference space $\mathcal{A}$ and are not trainable as $\mathcal{G}$ is fixed. On the visual side, we compute the relative encodings of g_i^t and g_i^s with respect to the learnable anchors $\mathcal{A}_t$ and $\mathcal{A}_s$, i.e., $\mathbf{r}^{g_i^t} = rel(g_i^t, \mathcal{A}_t)$ and $\mathbf{r}^{g_i^s} = rel(g_i^s, \mathcal{A}_s)$. We encourage the visual relative encodings $(\mathbf{r}^{g_i^t}, \mathbf{r}^{g_i^s})$ to match the text relative encodings $(\mathbf{r}^{z_i^t}, \mathbf{r}^{y_i^s})$ with an L1 loss:

$$\mathcal{L}_S = \begin{cases} |\mathbf{r}^{g_i^t} - \mathbf{r}^{z_i^t}|, & \text{if Target Domain} \\ |\mathbf{r}^{g_i^s} - \mathbf{r}^{y_i^s}|, & \text{otherwise} \end{cases}. \quad (6)$$

This loss encourages the visual representations to maintain relationships with their respective domain-specific visual anchors that mirror the relationships expressed by the domain-invariant anchors $\mathcal{A}$. This process supervises both the learnable anchors $\mathcal{A}_s, \mathcal{A}_t$ and the visual encoder $\mathcal{V}$. However, since this loss alone might lead to the collapse of anchor representations into a small region, hindering effective classification, we introduce a volume (or spread) regularization loss. This loss treats each set of anchors as a multidimensional parallelotope whose spread or volume occupied in the latent space can be measured by the determinant of its Gram matrix. Given the Gram matrices for the three sets of anchors:

$$\gamma = \mathcal{A}\mathcal{A}^T, \gamma_s = \mathcal{A}_s\mathcal{A}_s^T, \gamma_t = \mathcal{A}_t\mathcal{A}_t^T, \quad (7)$$

we encourage the volume of each domain-specific parallelotope to be approximately equal to the volume of the reference anchors $\mathcal{A}$ with the following regularization loss: e

$$\begin{aligned} \mathcal{L}_{Reg} = & |logDet(\gamma_t) - logDet(\gamma)| + \\ & |logDet(\gamma_s) - logDet(\gamma)|, \end{aligned} \quad (8)$$

where we consider log-determinants instead of determinants for numerical stability. This regularization loss ensures the visual anchors occupy a volume similar to the one of $\mathcal{A}$, preventing collapse and promoting better separation of classes in the representation space.

3.4.2. Classification Training

The network undergoes classification training concurrently with structure learning. To ensure that source and target visual representations are grounded in the structure imposed by $\mathcal{A}$, while still being free to capture domain-specific nuances, we propose a novel cross-domain attention layer

that aims to ground domain-specific visual representations g_i^t and g_i^s to the structure of source anchors $\mathcal{A}_s$. We include two Cross-Domain attention layers sharing the same weights for the target and source branches of the architecture (See Fig. 2). In the target branch, we use g_i^t as queries, target anchors $\mathcal{A}_t$ as keys, and source anchors $\mathcal{A}_s$ as values. The output is summed to g_i^t with a residual connection to include both domain-specific and cross-domain information in the final representation:

$$f_i^t = \text{Attention}(Q = g_i^t, K = \mathcal{A}_t, V = \mathcal{A}_s) + g_i^t \quad (9)$$

A similar processing is applied to the source encoding g_i^s:

$$f_i^s = \text{Attention}(Q = g_i^s, K = \mathcal{A}_s, V = \mathcal{A}_s) + g_i^s \quad (10)$$

Crucially, the cross-domain attention layer constructs an attention map by leveraging the domain-specific relations between the encoded representations and the anchors (i.e., QK^T). These relations are expected to be similar across domains due to the influence of the structural loss $\mathcal{L}_S$. This process results in a unified representation since the output of the attention layer always depends on values coming from the source anchors $\mathcal{A}_s$. The residual connection included in the attention layer is introduced to account for domain-dependent class characteristics, which results in final feature vectors f_i^s and f_i^t. Finally, an MLP classification head processes f_i^t and f_i^s to output predictions $\hat{y}_i^t$ and $\hat{y}_i^s$. The MLP has shared weights across domains and is optimized with a cross-entropy loss using the ground truth label y_i^s for source examples and the pseudo-label $\overline{y}_i^t$ predicted by $\mathcal{G}$ in the case of target samples. The overall training objective of the classifier is the following:

$$\mathcal{L} = \lambda_1\mathcal{L}_{CE}(\hat{y}_i, y_i^*) + \lambda_2\mathcal{L}_S(\mathbf{r}^*, \mathbf{r}^{g_i}) + \lambda_3\mathcal{L}_{Reg}(\gamma, \gamma_*), \quad (11)$$

where λ_1, λ_2, and λ_3 are hyperparameter weights to calibrate the magnitude of each loss value, and '*' signifies that the argument is domain-dependent[2].

4. Experiments

4.1. Experimental Setup

Datasets. We evaluate LAGUNA on four comprehensive UDA benchmarks: DomainNet [50], GeoPlaces [26] GeoImnet [26], and Ego2Exo [20, 27]. DomainNet consists of $400K$ images across 345 classes and is used to evaluate performance in 12 domain transfer settings [27, 68]. GeoImnet and GeoPlaces are subsets of the larger GeoNet dataset [26], contain over 750K images and focus on geographic disparities in image classification (GeoImNet with 600 classes) and scene recognition (GeoPlaces with 205 classes). Finally, Ego2Exo is a subset of EgoExo-4D curated for domain adaptation. It involves a total of 9086

---

[2]For example, for $\mathcal{L}_{CE}$ the ground truth can either be y_i^s or $\overline{y}_i^t$.

| Source→ | Real | | | Clipart | | | Sketch | | | Painting | | | Avg. |
Target→	C	S	P	R	S	P	R	C	P	R	C	S	
Source Only	63.02	49.47	60.48	70.52	56.09	52.53	70.42	65.91	54.47	73.34	60.09	48.25	60.38
MCD [58]	39.40	25.20	41.20	44.60	31.20	25.50	34.50	37.30	27.20	48.10	31.10	22.80	34.01
MDD [75]	52.80	41.20	47.80	52.50	42.10	40.70	54.20	54.30	43.10	51.20	43.70	41.70	47.11
CGDM [13]	49.40	38.20	47.20	53.50	36.90	35.30	55.60	50.10	43.70	59.40	37.70	33.50	45.04
SCDA [33]	54.00	42.50	51.90	55.00	44.10	39.30	53.20	55.60	44.70	56.20	44.10	42.00	48.55
SSRT-B [62]	69.90	58.90	66.00	75.80	59.80	60.80	73.20	70.60	62.20	71.40	61.70	55.20	65.41
MemSAC [25]	63.49	42.14	60.32	72.33	54.92	46.14	73.46	68.04	52.75	74.42	57.79	43.57	59.11
CDTrans [73]	66.20	52.90	61.50	72.60	58.10	57.10	72.50	69.00	59.00	72.10	62.90	53.90	63.16
PMTrans [77]	74.10	61.10	_70.00_	79.30	63.70	62.70	77.50	73.80	62.60	79.80	69.70	61.20	69.63
UDA with Language Guidance													
TextMatch [27]	71.36	64.30	65.32	_81.25_	65.65	64.85	81.09	72.65	63.94	_81.08_	70.84	_64.17_	70.64
nGramMatch [27]	68.92	59.82	63.15	76.35	61.72	62.87	76.35	69.28	62.51	76.04	68.52	60.52	67.17
LaGTran [27]	_77.30_	_68.25_	67.35	81.31	_67.03_	_66.81_	80.78	_75.62_	_68.08_	79.23	_73.80_	63.44	_72.41_
LAGUNA	**80.34**	**70.68**	**71.92**	**83.07**	**69.51**	**70.59**	**83.34**	**79.71**	**70.51**	**83.32**	**77.47**	**68.32**	**75.73**
Improvement	+3.04	+2.43	+1.92	+1.82	+2.48	+3.78	+2.25	+4.09	+2.43	+2.24	+3.67	+4.15	+3.32

Table 1. Results on DomainNet [50], comprising 4 domains: Real, Clipart, Sketch, and Painting, for a total of 12 domain adaptation scenarios. All models use the Swin-B [37] backbone. Best results are reported in **bold**, whereas second best results are underlined.

videos and 24 action categories. DomainNet and GeoNet include per-image captions derived from metadata and BLIP-2 [32], respectively, while Ego2Exo provides video descriptions from the original EgoExo-4D dataset.

Model and training details. Our method comprises three base components: a language model for structure definition (Stage 1), a text supervisor model (Stage 2), and a visual backbone (Stage 3). We use a pre-trained *SentenceTransformer* [55] for structure definition and a BERT base model [59] as text supervisor. The latter is trained on the source domain of each scenario for 5 epochs with a batch size of 64, a learning rate of 1e-4, and the AdamW [41] optimizer. The visual encoders are tailored to each dataset. For DomainNet, we adapt the Swin-B backbone [37], following prior work [25, 73, 77]. For GeoImnet and GeoPLaces, we utilize the ViT-B backbone [12] following [68, 77]. For Ego2Exo, we employ pre-extracted Omnivore-L features [20] as in [27]. All vision models undergo joint training with the other network components for 10 epochs with a batch size of 32, an initial learning rate of $1e-4$, and cosine scheduling. We set the loss weights to λ_1=1.0 λ_2=0.1 and λ_3=0.001 to normalize the relative magnitude of the related loss functions.

4.2. Results

We benchmark our model against various domain adaptation methods that have reported results on the considered datasets [7, 25, 68, 73, 75, 77]. Besides classic UDA methods, we also compare with recent approaches using language guidance [27]. Following the evaluation protocols of each benchmark, we report accuracy for DomainNet, GeoImNet, and GeoPlaces, while for Ego2Exo, we use per-class mean accuracy as reported in [27].

DomainNet. Table 1 presents the results on the Domain-

| Model | GeoImnet | | GeoPlaces | | Avg. |
	U→A	A→U	U→A	A→U	
Source Only	52.46	51.91	44.90	36.85	46.53
CDAN [40]	54.48	53.87	42.88	36.21	46.86
MemSAC [25]	53.02	54.37	42.05	38.33	46.94
ToAlign [68]	55.67	55.92	42.32	38.40	48.08
MDD [75]	51.57	50.73	42.54	39.23	46.02
DALN [7]	55.36	55.77	41.06	40.41	48.15
PMTrans [77]	56.76	57.60	46.18	40.33	50.22
UDA with Language Guidance					
TextMatch [27]	49.68	54.82	53.06	50.11	51.92
nGramMatch [27]	49.53	51.02	51.70	49.87	50.93
LaGTran [27]	_63.67_	_64.16_	_56.14_	_57.02_	_60.24_
LAGUNA	**67.39**	**69.97**	**61.15**	**63.51**	**65.40**
Improvement	+3.72	+5.81	+5.01	+6.49	+5.26

Table 2. Results on GeoImnet [26] and GeoPlaces [26] with 4 adaptation scenarios. Best results are in **bold**, whereas second best results are underlined. All methods use the ViT-B [12] backbone.

Net dataset, encompassing 12 domain adaptation scenarios across four distinct domains: Real (R), Clipart (C), Sketch (S), and Painting (P). LAGUNA is compared against previous state-of-the-art methods utilizing the Swin-B backbone to ensure a fair comparison. Notably, LAGUNA surpasses all prior work in all 12 scenarios, achieving an average accuracy improvement of +3.32%.

GeoImnet & GeoPlaces. Table 2 presents the results on GeoImnet and GeoPlaces covering adaptation between USA (U) and Asia (A) geographical domains. LAGUNA consistently outperforms previous methods, achieving improvements of +3.72% (U→A) and +5.81% (A→U) on GeoImnet, and +5.01% (U→A) and +6.49% (A→U) on GeoPlaces. This translates to an average improvement of +5.26% over prior art across all scenarios.

Ego2Exo. LAGUNA demonstrates consistent performance

Model	Ego→Exo	Exo→Ego	Avg.
Source Only	8.39	15.66	12.03
TA3N [8]	6.92	27.95	17.44
TransVAE [69]	12.06	23.34	17.70
Zero-shot Video Recognition			
EgoVLP [35]	5.89	19.35	12.62
LaVILA [76]	5.86	23.16	14.51
UDA with Language Guidance			
TextMatch [27]	10.36	13.57	11.97
nGramMatch [27]	11.50	15.46	13.98
LaGTran [27]	12.34	30.76	21.55
LAGUNA	**13.52**	**33.45**	**33.45**
Improvement	**+1.18**	**+2.69**	**+1.94**

Table 3. Results on the Ego2Exo benchmark from EgoExo-4D [20]. Best results are in **bold**, whereas second best results are underlined. All methods use official pre-extracted Omnivore-L features [17, 20] except EgoVLP [35] and LaVILA [76], which use Timesformer-B [3] backbone.

Setting	$\mathcal{L}_S$	$\mathcal{A}$	$\mathcal{A}_{t/s}$	CD Attn.	$\mathcal{L}_{Reg}$	Avg. Acc.	Rel. Imp.	Abs. Imp.
(1)	-	-	-	-	-	63.21	-	-
(2)	-	✓	-	-	-	63.99	+0.78	+0.78
(3)	✓	✓	✓*	-	-	65.22	+1.32	+2.10
(4)	✓	✓	✓	-	-	67.08	+1.77	+3.87
(5)	✓	✓	✓	✓	-	67.52	+0.44	+4.31
LAGUNA	✓	✓	✓	✓	✓	68.68	+1.16	+5.47

Table 4. Ablation on LAGUNA's methodological components. '*' means that the learnable anchors $\mathcal{A}_{t/s}$ are domain-independent (i.e., the same set of learnable anchors for target and source).

gains also on the challenging recent Ego2Exo action recognition benchmark based on EgoExo4D [20], as introduced in [27] (Table 3). Following the definition of this benchmark [27], we report per-class average accuracy due to the dataset's imbalanced class distribution. LAGUNA surpasses all prior work, achieving improvements of +1.18% (Ego→Exo) and +2.69% (Exo→Ego), resulting in an overall average gain of +1.94%.

These consistent improvements across all tested datasets, with a total of 18 scenarios, underscore the effectiveness of our approach both when considering classic UDA algorithms and recent methods exploring the use of language guidance to assist domain adaptation.

4.3. Ablation Study

In this section, we analyze the effect of the individual components of LAGUNA to provide an in-depth assessment of model performance and justify our design choices. Specifically, we investigate (a) the impact of our methodological elements, (b) the effects of different reference structures, (c) the behavior of LAGUNA with different ratios of pseudo-labels, and (d) the qualitative illustration of the alignment of relative and absolute spaces. More ablations are reported in the *Supplementary Material*.

Methodological Elements. In Table 4, we ablate the main methodological elements of LAGUNA on GeoImnet. We set up this ablation by removing the core elements of LAGUNA, namely the structure loss ($\mathcal{L}_S$ - Eq. (6)), the use of domain-agnostic reference anchors $\mathcal{A}$, the learnable anchors $\mathcal{A}_{t/s}$, the cross-domain attention (CD Attn. - Eq. (9), (10)), and the regularization loss ($\mathcal{L}_{Reg}$ - Eq. (12)). We then progressively add the considered elements until we reach LAGUNA (settings (1)-(5)). For all settings, we report the

relative improvement (w.r.t. previous row) and the absolute one (w.r.t. first row). In setting (1), the visual classifier is trained on source and target samples using labels and pseudo-labels predicted by the text supervisor, which leads to a baseline model akin to LaGTran [27]. In setting (2), we add the reference anchors $\mathcal{A}$ and train the visual classifier as in setting (1), but we also add a loss to align visual representations with the related class anchors in absolute coordinates through cosine similarity. This yields an improvement of +0.87 with respect to (1), suggesting that imposing a reference structure beyond pseudo-labels is beneficial to performance but that absolute alignment is too restrictive. In setting (3), we extend (2) with a set of domain-independent learnable anchors (✓* in Table 4, indicating that the weights of $\mathcal{A}_t$ and $\mathcal{A}_s$ are shared) and add the structure-preserving loss $\mathcal{L}_S$ (Eq. 6). The introduction of learnable anchors enables the relative alignment of the visual domain with the reference language domain, leading to an improvement of +1.32 with respect to (2) and +2.10 over (1), which highlights the benefits of the proposed relative alignment over an absolute one. Considering domain-specific anchors in setting (4) allows source and target domains to focus on the individual characteristics of the respective domains, leading to a further improvement of +1.77 with respect to (3) and a robust +3.87 with respect to the baseline (1). Adding cross-domain attention in setting (5) allows to bridge the gap between target and source representations, leading to an additional improvement of +0.44 with respect to (4) and an overall improvement of +4.31 with respect to (1). We achieve the final configuration in LAGUNA, where we add the regularization loss, which prevents collapse in learning domain-specific anchors and leads to a further improvement of +1.16 with respect to (5) and an overall gain of +5.47 with respect to the baseline (1). Together, these improvements show the advantages of the careful design of LAGUNA, highlighting the benefits of learning domain-specific visual representations still aligned to a reference language structure.

Ablation on reference structures. Fig. 3 delves into the importance of defining a semantic structure for our method (Stage 1). We ablate three models: *Sentence-Transformer*[55], CLIP [54], and BERT [59], each with

Figure 3. Ablation on three language models, *SentenceTransformer*, CLIP, and BERT, used to define the reference structure. We visualize the similarity maps of 100 randomly selected classes from GeoImnet (green to yellow areas indicate higher similarity) and the average accuracies on all classes achieved by using each model to define the reference space. The best result is in **bold**, and the second best is underlined.

Dataset	10%	20%	30%	50%	75%	100%
GeoImnet	61.40	63.43	64.21	66.42	66.71	68.68
GeoPlaces	55.35	57.85	58.95	60.02	60.75	62.33
Avg.	58.38	60.64	61.58	63.22	63.72	65.40

Table 5. Ablation using different target data ratios. We report average accuracy on GeoImnet and GeoPlaces. The result in Green indicates the ratio needed for LAGUNA to obtain SOTA accuracy.

varying capabilities of extracting semantic relationships. The figure shows semantic similarity maps for 100 randomly selected classes from GeoImnet and the average accuracy obtained by LAGUNA on the entire dataset when using the specified model. As can be seen, *SentenceTransformer* keeps low similarities for distinct classes (low values off-diagonal), BERT maintains high similarities for most classes but also presents some spurious correlations between distinct classes, while the behavior of CLIP is in-between. *SentenceTransformer*, which can better differentiate between semantic relations, leads to the best accuracy.

Ablation on the ratio of target samples. LAGUNA demonstrates strong generalization capabilities even with limited target domain data, as shown in Table 5. We progressively increase the amount of pseudo-labeled target data used for training (10%, 20%, 30%, 50%, 75%) and observe the impact on performance. Notably, LAGUNA achieves high accuracy even with only 10% of the target samples and attains state-of-the-art results with just 20%. This efficiency underscores the benefits of our structure-driven approach, which organizes sample projections within domain-specific representation spaces while leveraging a reference geometrical structure to guide their placement.

Qualitative illustration of embedding spaces. LAGUNA uniquely aligns representation structures without enforcing alignment in absolute coordinates but rather leverag-


Figure 4. Visualization of *Source* and *Target* (Cross-Domain) visual features of 1000 randomly selected samples from GeoImnet with t-SNE. We illustrate their position in latent space from two perspectives: left) absolute coordinates and right) relative coordinates. Notably, in absolute coordinates, the two representation spaces are quite distinct, adhering to our training strategy without direct alignment. Relative encodings w.r.t $\mathcal{A}_s$ and $\mathcal{A}_t$ are aligned.

ing a reference structure defined by a language model and the class labels to guide the alignment of inter-class affinities. Figure 4 provides a qualitative validation of this approach. Using t-SNE, we visualize the projection of 1000 randomly selected samples from the GeoImnet validation set, where the source and the target domains belong to the USA and Asia, respectively. Notably, the features are well separated in absolute coordinates, reflecting their unique domain-related visual properties. However, when using relative encodings with domain-specific learnable anchors, the features blend together due to the structure-preserving loss $\mathcal{L}_S$. This demonstrates that semantic alignment can be achieved without forcing representations into a shared space, allowing each domain to preserve its unique characteristics while adhering to a common structure.

5. Conclusions

This work introduces LAGUNA, a novel domain adaptation approach leveraging geometrical structures of semantically equivariant spaces to guide adaptation. LAGUNA defines a reference representation space structure based on domain-agnostic class semantic similarities encoded by a language model, ensuring the organization of sample projections reflects this structure while preserving domain-specific characteristics. By conditioning the classifier to adhere to this structure, LAGUNA encourages structural similarity across domain-specific latent spaces, retaining unique features for improved classification. This is achieved through pseudo-labeling and learnable domain-specific anchors, guided by a loss function that prioritizes mimicking geometrical associations over direct representation alignment. Extensive experiments demonstrate LAGUNA's superior performance compared to existing methods. This research highlights the importance of structural alignment and language-guided learning in domain adaptation. Future work could explore extending LAGUNA to more complex scenarios and multi-modal data.

References

[1] Anders Andreassen, Yasaman Bahri, Behnam Neyshabur, and Rebecca Roelofs. The evolution of out-of-distribution robustness throughout fine-tuning. *TMLR*, 2022, 2021. 2

[2] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *NeurIPS*, 19, 2006. 2

[3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? *ICML*, abs/2102.05095, 2021. 7

[4] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. *NeurIPS*, 29, 2016. 1, 2

[5] Irene Cannistraci, Luca Moschella, Marco Fumero, Valentino Maiorca, and Emanuele Rodolà. From bricks to bridges: Product of invariances to enhance latent space communication. In *ICLR*, 2024. 3

[6] Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *CVPR*, pages 627–636, 2019. 2

[7] Lin Chen, H. Chen, Zhixiang Wei, Xin Jin, Xiao Tan, Yi Jin, and Enhong Chen. Reusing the task-specific classifier as a discriminator: Discriminator-free adversarial domain adaptation. *CVPR*, pages 7171–7180, 2022. 6

[8] Min-Hung Chen, Zsolt Kira, Ghassan AlRegib, Jaekwon Yoo, Ruxin Chen, and Jian Zheng. Temporal attentive alignment for large-scale video domain adaptation. In *ICCV*, pages 6321–6330, 2019. 2, 7

[9] Zhijie Deng, Yucen Luo, and Jun Zhu. Cluster alignment with a teacher for unsupervised domain adaptation. In *ICCV*, pages 9944–9953, 2019. 2

[10] Benjamin Devillers, Romain Bielawski, Bhavin Choksi, and Rufin van Rullen. Does language help generalization in vision models? In *Conference on Computational Natural Language Learning*, 2021. 2

[11] Anxhelo Diko, Danilo Avola, Bardh Prenkaj, Federico Fontana, and Luigi Cinque. Semantically guided representation learning for action anticipation. In *ECCV*, 2024. 3

[12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929, 2020. 6

[13] Zhekai Du, Jingjing Li, Hongzu Su, Lei Zhu, and Ke Lu. Cross-domain gradient discrepancy minimization for unsupervised domain adaptation. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3936–3945, 2021. 6

[14] Lisa Dunlap, Clara Mohri, Devin Guillory, Han Zhang, Trevor Darrell, Joseph E. Gonzalez, Aditi Raghunanthan, and Anja Rohrbach. Using language to extend to unseen domains. *ICLR*, abs/2210.09520, 2022. 2

[15] Geoff French, Michal Mackiewicz, and Mark Fisher. Self-ensembling for visual domain adaptation. In *ICLR*, 2018. 2

[16] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015. 2

[17] Rohit Girdhar, Mannat Singh, Nikhil Ravi, Laurens van der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A single model for many visual modalities. *CVPR*, pages 16081–16091, 2022. 7

[18] Tejas Gokhale, Rushil Anirudh, Bhavya Kailkhura, Jayaraman J. Thiagarajan, Chitta Baral, and Yezhou Yang. Attribute-guided adversarial training for robustness to natural perturbations. In *AAAI*, 2020. 2

[19] Sachin Goyal, Ananya Kumar, Sankalp Garg, Zico Kolter, and Aditi Raghunathan. Finetune like you pretrain: Improved finetuning of zero-shot vision models. *CVPR*, pages 19338–19347, 2022. 2

[20] Kristen Grauman, Andrew Westbury, and et al. Lorenzo Torresani. Ego-exo4d: Understanding skilled human activity from first- and third-person perspectives. *CVPR*, pages 19383–19400, 2023. 2, 5, 6, 7

[21] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *JMLR*, 13(25):723–773, 2012. 1

[22] Yifei Huang, Guo Chen, Jilan Xu, Mingfang Zhang, Lijin Yang, Baoqi Pei, Hongjie Zhang, Lu Dong, Yali Wang, Limin Wang, and Yu Qiao. Egoexolearn: A dataset for bridging asynchronous ego- and exo-centric view of procedural activities in real world. *CVPR*, pages 22072–22086, 2024. 2

[23] Zeyi Huang, Andy Zhou, Zijian Lin, Mu Cai, Haohan Wang, and Yong Jae Lee. A sentence speaks a thousand images: Domain generalization through distilling clip with language guidance. *ICCV*, pages 11651–11661, 2023. 2

[24] Tarun Kalluri and Manmohan Chandraker. Cluster-to-adapt: Few shot domain adaptation for semantic segmentation across disjoint labels. In *CVPR*, pages 4121–4131, 2022. 2

[25] Tarun Kalluri, Astuti Sharma, and Manmohan Chandraker. Memsac: Memory augmented sample consistency for large scale domain adaptation. In *ECCV*, pages 550–568. Springer, 2022. 2, 6

[26] Tarun Kalluri, Wangdong Xu, and Manmohan Chandraker. Geonet: Benchmarking unsupervised adaptation across geographies. *CVPR*, pages 15368–15379, 2023. 2, 5, 6

[27] Tarun Kalluri, Bodhisattwa Prasad Majumder, and Manmohan Chandraker. Tell, don't show!: Language guidance eases transfer across domains in images and videos. *arXiv preprint arXiv:2403.05535*, 2024. 1, 2, 4, 5, 6, 7

[28] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *ICCV*, pages 4893–4902, 2019. 2

[29] Abhishek Kumar, Prasanna Sattigeri, Kahini Wadhawan, Leonid Karlinsky, Rogerio Feris, Bill Freeman, and Gregory Wornell. Co-regularized alignment for unsupervised domain adaptation. *NeurIPS*, 31, 2018. 2

[30] Ananya Kumar, Aditi Raghunathan, Robbie Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained

features and underperform out-of-distribution. *ArXiv*, abs/2202.10054, 2022. 2

[31] Jingjing Li, Erpeng Chen, Zhengming Ding, Lei Zhu, Ke Lu, and Heng Tao Shen. Maximum density divergence for domain adaptation. *IEEE TPAMI*, 43:3918–3930, 2020. 1, 2

[32] Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. 2, 4, 6

[33] Shuang Li, Mixue Xie, Fangrui Lv, Chi Harold Liu, Jian Liang, Chen Qin, and Wei Li. Semantic concentration for domain adaptation. In *ICCV*, pages 9102–9111, 2021. 1, 6

[34] Yanghao Li, Tushar Nagarajan, Bo Xiong, and Kristen Grauman. Ego-exo: Transferring visual representations from third-person to first-person videos. *CVPR*, pages 6939–6949, 2021. 2

[35] Kevin Lin, Alex Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Z. Xu, Difei Gao, Rong-Cheng Tu, Wenzhe Zhao, Weijie Kong, Chengfei Cai, Hongfa Wang, Dima Damen, Bernard Ghanem, Wei Liu, and Mike Zheng Shou. Egocentric video-language pretraining. *NeurIPS*, abs/2206.01670, 2022. 7

[36] Hong Liu, Jianmin Wang, and Mingsheng Long. Cycle self-training for domain adaptation. *NeurIPS*, 34:22968–22981, 2021. 2

[37] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *ICCV*, pages 9992–10002, 2021. 6

[38] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, 2017. 2

[39] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, pages 2208–2217, 2017. 2

[40] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. *NeurIPS*, 31, 2018. 1, 2, 6

[41] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6

[42] 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. *CVPR*, pages 2502–2511, 2018. 1

[43] 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 *CVPR*, pages 2507–2516, 2019. 2

[44] Valentino Maiorca, Luca Moschella, Antonio Norelli, Marco Fumero, Francesco Locatello, and Emanuele Rodolà. Latent space translation via semantic alignment. In *NeurIPS*, 2023. 3

[45] Seonwoo Min, Nokyung Park, Siwon Kim, Seunghyun Park, and Jinkyu Kim. Grounding visual representations with texts for domain generalization. In *ECCV*, 2022. 2

[46] L Moschella et al. Relative representations enable zero-shot latent space communication (2023). *ICML*, pages 1–10, 2023. 2, 3, 4

[47] Antonio Norelli, Marco Fumero, Valentino Maiorca, Luca Moschella, Emanuele Rodola, and Francesco Locatello. Asif: Coupled data turns unimodal models to multimodal without training. *Advances in Neural Information Processing Systems*, 36:15303–15319, 2023. 3

[48] Changhwa Park, Jonghyun Lee, Jaeyoon Yoo, Minhoe Hur, and Sungroh Yoon. Joint contrastive learning for unsupervised domain adaptation. *arXiv preprint arXiv:2006.10297*, 2020. 2

[49] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. In *AAAI*, 2018. 1, 2

[50] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. *ICCV*, pages 1406–1415, 2018. 2, 5, 6

[51] Hieu Pham, Zihang Dai, Golnaz Ghiasi, Hanxiao Liu, Adams Wei Yu, Minh-Thang Luong, Mingxing Tan, and Quoc V. Le. Combined scaling for zero-shot transfer learning. *Neurocomputing*, 555:126658, 2021. 2

[52] Viraj Prabhu, Ramprasaath R Selvaraju, Judy Hoffman, and Nikhil Naik. Can domain adaptation make object recognition work for everyone? In *CVPR*, pages 3981–3988, 2022. 2

[53] Camillo Quattrocchi, Antonino Furnari, Daniele Di Mauro, Mario Valerio Giuffrida, and Giovanni Maria Farinella. Synchronization is all you need: Exocentric-to-egocentric transfer for temporal action segmentation with unlabeled synchronized video pairs. *ECCV*, 2023. 2

[54] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICLR*, 2021. 2, 7, 1

[55] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP*, 2019. 4, 6, 7, 1

[56] Aadarsh Sahoo, Rutav Shah, Rameswar Panda, Kate Saenko, and Abir Das. Contrast and mix: Temporal contrastive video domain adaptation with background mixing. *NeurIPS*, 34: 23386–23400, 2021. 2

[57] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Adversarial dropout regularization. In *ICLR*, 2018. 2

[58] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, pages 3723–3732, 2018. 1, 2, 6

[59] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108, 2019. 6, 7, 1

[60] Astuti Sharma, Tarun Kalluri, and Manmohan Chandraker. Instance level affinity-based transfer for unsupervised domain adaptation. In *CVPR*, pages 5361–5371, 2021. 2

[61] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, pages 443–450, 2016. 2

[62] Tao Sun, Cheng Lu, Tianshuo Zhang, and Haibin Ling. Safe self-refinement for transformer-based domain adaptation. In *CVPR*, pages 7191–7200, 2022. 2, 6

[63] Shuhan Tan, Xingchao Peng, and Kate Saenko. Class-imbalanced domain adaptation: An empirical odyssey. In *ECCV*, pages 585–602, 2020. 1, 2

[64] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, pages 7167–7176, 2017. 2

[65] Rui Wang, Zuxuan Wu, Zejia Weng, Jingjing Chen, Guo-Jun Qi, and Yu-Gang Jiang. Cross-domain contrastive learning for unsupervised domain adaptation. *IEEE Transactions on Multimedia*, 25:1665–1673, 2022. 2

[66] Zhenbin Wang, Lei Zhang, Lituan Wang, and Minjuan Zhu. Landa: Language-guided multi-source domain adaptation. *ArXiv*, abs/2401.14148, 2024. 2

[67] Guoqiang Wei, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Metaalign: Coordinating domain alignment and classification for unsupervised domain adaptation. *CVPR*, pages 16638–16648, 2021. 1, 2

[68] Guoqiang Wei, Cuiling Lan, Wenjun Zeng, Zhizheng Zhang, and Zhibo Chen. Toalign: Task-oriented alignment for unsupervised domain adaptation. *NeurIPS*, 34:13834–13846, 2021. 1, 2, 5, 6

[69] Pengfei Wei, Lingdong Kong, Xinghua Qu, Yi Ren, Zhiqiang Xu, Jing Jiang, and Xiang Yin. Unsupervised video domain adaptation for action recognition: A disentanglement perspective. *NeurIPS*, 36:17623–17642, 2023. 2, 7

[70] Garrett Wilson and Diane J. Cook. A survey of unsupervised deep domain adaptation. *ACM Trans. Intell. Syst. Technol.*, 11(5):1–47, 2020. 1

[71] Mitchell Wortsman, Gabriel Ilharco, Mike Li, Jong Wook Kim, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, and Ludwig Schmidt. Robust fine-tuning of zero-shot models. *CVPR*, pages 7949–7961, 2021. 2

[72] Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *ICML*, pages 5423–5432, 2018. 2

[73] Tongkun Xu, Weihua Chen, Pichao Wang, Fan Wang, Hao Li, and Rong Jin. Cdtrans: Cross-domain transformer for unsupervised domain adaptation. *arXiv preprint arXiv:2109.06165*, 2021. 6

[74] Zihui Xue and Kristen Grauman. Learning fine-grained view-invariant representations from unpaired ego-exo videos via temporal alignment. *NeurIPS*, pages 1–8, 2023. 2

[75] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael I. Jordan. Bridging theory and algorithm for domain adaptation. In *ICML*, pages 1–10, 2019. 6

[76] Yue Zhao, Ishan Misra, Philipp Krahenbuhl, and Rohit Girdhar. Learning video representations from large language models. *CVPR*, pages 6586–6597, 2022. 7

[77] Jinjing Zhu, Haotian Bai, and Lin Wang. Patch-mix transformer for unsupervised domain adaptation: A game perspective. In *CVPR*, pages 3561–3571, 2023. 1, 2, 6

LAGUNA: LAnguage Guided UNsupervised Adaptation with structured spaces

Supplementary Material

A. Introduction

In this document, we report additional ablation studies regarding the choice of language models and the structural guidance of the target domain. Moreover, we provide a more detailed explanation of the motivation behind our regularization loss $\mathcal{L}_{Reg}$ in Eq. (8). Finally, we attach a Python file, LAGUNA.py, demonstrating the implementation of LAGUNA in PyTorch. We will release the full code and configurations upon acceptance.

B. Ablation on Language Models

In addition to the language model used for defining the reference structure through $\mathcal{A}$ in Stage 1, LAGUNA employs language models in Stage 2 for encoding image/video captions trained for text classification and capturing semantic structure. This model is then used to generate pseudo-labels and serve as structure guidance for the target domain in Stage 3. We ablate the choice of language model in these Stages, comparing LAGUNA's performance when employing *SentenceTransformer* [55], CLIP [54], or BERT [59] on GeoImnet dataset in two settings: (1) using *SentenceTransformer*-defined $\mathcal{A}$ (Table 6) as it can better model semantic relationships (recall Fig. 3 from the main manuscript) and (2) defining $\mathcal{A}$ with the same language model as the one chosen for training (Table 7).

In setting (1), BERT outperforms CLIP and *SentenceTransformer*, motivating our choice for the language model. In setting (2), CLIP achieves the highest accuracy, but its overall performance remains lower than BERT's in setting (1). These results not only emphasize the best model and justify our choice but also demonstrate that using different models for structure definition and training can boost performance since the choices are tailored for their specific characteristics. Specifically, the Stage 1 model is selected for its ability to model meaningful relationships between classes, while Stage 2 is intended for building better representations of captions that learn the defined structure and can also produce better pseudo-labels.

C. Ablation on structure guidance for target domain in Stage-3

During Stage-3 training of LAGUNA, in addition to classification training, we train our model also for structure-preserving as defined from $\mathcal{A}$ using the structural loss $\mathcal{L}_S$. In this particular training loss, for the source domain, we supervise the visual relative representation $\mathbf{r}^{g_i^s} = rel(g_i^s, A_s)$ with the language relative encoding of the class correspond-

Scenario	*SentenceTransformer*	CLIP	BERT
U→A	64.40	66.81	67.39
A→U	67.35	69.40	69.97
Avg.	65.87	<u>68.11</u>	**68.68**

Table 6. Ablation on three language models used in stage 2 and 3, *SentenceTransformer*, CLIP, and BERT with *Sentence-Transformer*-defined $\mathcal{A}$ on GeoImnet with two domains: Asia (A) and Usa (U). The best result is in **bold**, and the second best is <u>underlined</u>.

Scenario	*SentenceTransformer*	CLIP	BERT
U→A	64.40	66.48	66.32
A→U	67.35	69.28	68.72
Avg.	65.87	**67.88**	<u>67.52</u>

Table 7. Ablation on three language models used in stage 2 and 3, *SentenceTransformer*, CLIP, and BERT, used for Stage-2 training on GeoImnet with two domains: Asia (A) and Usa (U). In this scenario, the same pre-trained model encodes $\mathcal{A}$ (stage 1). The best result is in **bold**, and the second best is <u>underlined</u>.

ing language anchor $\mathbf{r}^{y_i^s} = rel(\mathcal{A}[y_i^s], \mathcal{A})$. On the other hand, in the target domain, since we do not have a ground truth label but only the estimated pseudo label $\overline{y}_i^t$, we supervise the visual relative representation $\mathbf{r}^{g_i^t} = rel(g_i^t, A_t)$ with the language relative encoding of the sample captioning $\mathbf{r}^{z_i^t} = rel(z_i^t, \mathcal{A})$ to benefit from the richer information expressed in the language encoding compared to the estimated pseudo-label. In this ablation, we investigate the benefits of using the $\mathbf{r}^{z_i^t}$ compared to relying on the pseudo-label and supervise the visual encoding with $\mathbf{r}^{\overline{y}_i^t} = rel(\mathcal{A}[y_i^t], \mathcal{A})$. Specifically, we ablate on the GeoImnet and Ego2Exo datasets and report the results in Table 8. Notably, $\mathbf{r}^{z_i^t}$ constantly gives better performances than $\mathbf{r}^{\overline{y}_i^t}$ in all scenarios motivating our choice to use the textual encoding rather than the pseudo-label corresponding anchor to create the relative representations for structure supervision in the target domain.

D. Regularization Loss Extended

In Section 3.4.1 of the main manuscript, we explain how the structure loss $\mathcal{L}_S$ is supported by a regularization loss $\mathcal{L}_{Reg}$ to avoid feature collapsing. The logic behind $\mathcal{L}_{Reg}$ is built upon two base concepts of linear algebra: Gram matrix

GeoImnet				Ego2Exo			
U→A		A→U		Ego→Exo		Exo→Ego	
$\mathbf{r}^{z_i^t}$	$\mathbf{r}^{\overline{y}_i^t}$	$\mathbf{r}^{z_i^t}$	$\mathbf{r}^{\overline{y}_i^t}$	$\mathbf{r}^{z_i^t}$	$\mathbf{r}^{\overline{y}_i^t}$	$\mathbf{r}^{z_i^t}$	$\mathbf{r}^{\overline{y}_i^t}$
67.39	66.77	**69.97**	68.14	**13.52**	13.49	**33.45**	32.88

Table 8. Ablation on structure guidance representations on GeoImnet and Ego2Exo. We consider the comparison between relative representations obtained from using the language encoding $\mathbf{r}^{z_i^t}$ and from using $\mathbf{r}^{\overline{y}_i^t}$. The best result for each scenario is in **bold**. Note that for Ego2Exo the reported metric is mean per class accuracy.

and determinant. To better explain this logic and its relation with volume and feature collapsing, in this section, we give more information about what is a Gram matrix, what its determinant represents, how it is connected to volume, and why it serves as a collapsing indicator to motivate our approach.

Gram matrix. Given a set of vectors $v_1, v_2, \ldots, v_n$ in $\mathbb{R}^d$, the Gram matrix γ is defined as:

$$\gamma = V^T V$$

where V is the matrix whose columns are the vectors v_i, and γ is an $n \times n$ symmetric matrix whose elements are the pairwise inner products of the vectors:

$$\gamma_{ij} = \langle v_i, v_j \rangle$$

Additionally, its rank corresponds to the number of linearly independent vectors, and its determinant (if full-rank) relates to the volume of the parallelepiped spanned by the vectors. If γ is not full rank, the vectors used to construct it are linearly dependent, so they do not span the full n-dimensional space. Geometrically, this implies that the volume of the parallelepiped they define is zero, as they lie in a lower-dimensional subspace. In other words, a lower rank or a volume of zero of the vectors composing γ indicates collapsing in a lower subspace. In LAGUNA, we use these properties of the Gram matrix to control and prevent the collapsing of the anchor vectors composing $\mathcal{A}_s$ and $\mathcal{A}_t$. Particularly, from Eq. (7) in the manuscript, we calculate $\gamma_s = A_s^T A_s$ and $\gamma_s = A_t^T A_t$.

Determinant of Gram matrix. To determine if a set of vectors has collapsed, one can either compute the rank of its Gram matrix or evaluate its volume in space. In LAGUNA, we choose the latter approach as we can compare it with a reference volume that helps prevent collapse during training. Specifically, to calculate the volume of the parallelepiped formed by the set of vectors (anchors) in $\mathcal{A}_{s/t}$, we utilize the determinant. The determinant of a Gram matrix represents the square of the volume of the parallelepiped defined by the vectors used to compute the matrix (e.g.,

$Det(\gamma_s)$ gives the square of the volume occupied by all anchors in $\mathcal{A}_s$). Thus, the volume for $\mathcal{A}_{s/t}$ can be expressed as:

$$\text{Volume}_{s/t} = \sqrt{Det(\gamma_{s/t})}.$$

Importantly, if the anchors in $\mathcal{A}_{s/t}$ are linearly dependent (i.e., not full rank), $Det(\gamma_{s/t}) = 0$. To prevent collapse, we can regularize $\mathcal{L}_S$ by requiring the volume to remain greater than zero. However, in higher-dimensional spaces, volumes can have very large values. For numerical stability, we replace $Det(\gamma_{s/t})$ with $logDet(\gamma_{s/t})$ in Eq. (8) and enforce that the space occupied by anchors in $\mathcal{A}_{s/t}$ remains log comparable to that of the reference anchors in $\mathcal{A}$ preventing their collapse into a zero volume. Specifically, Eq. (8) in the manuscript is defined as follows:

$$\begin{aligned} \mathcal{L}_{Reg} = & \ |logDet(\gamma_t) - logDet(\gamma)| + \\ & \ |logDet(\gamma_s) - logDet(\gamma)|, \end{aligned} \quad (12)$$

where $\gamma = \mathcal{A}^T \mathcal{A}$.