

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CARL: CAMERA-AGNOSTIC REPRESENTATION LEARN- ING FOR SPECTRAL IMAGE ANALYSIS

Anonymous authors

Paper under double-blind review

ABSTRACT

Spectral imaging offers promising applications across diverse domains, including medicine and urban scene understanding, and is already established as a critical modality in remote sensing. However, variability in channel dimensionality and captured wavelengths among spectral cameras impede the development of AI-driven methodologies, leading to camera-specific models with limited generalizability and inadequate cross-camera applicability. To address this bottleneck, we introduce **CARL**, a model for Camera-Agnostic Representation Learning across RGB, multispectral, and hyperspectral imaging modalities. To enable the conversion of a spectral image with any channel dimensionality to a camera-agnostic representation, we introduce a novel spectral encoder, featuring a self-attention-cross-attention mechanism, to distill salient spectral information into learned spectral representations. Spatio-spectral pre-training is achieved with a novel feature-based self-supervision strategy tailored to CARL. Large-scale experiments across the domains of medical imaging, autonomous driving, and satellite imaging demonstrate our model’s unique robustness to spectral heterogeneity, outperforming on datasets with simulated and real-world cross-camera spectral variations. The scalability and versatility of the proposed approach position our model as a backbone for future spectral foundation models.

1 INTRODUCTION

Spectral imaging, including RGB, multispectral, and hyperspectral imaging, capture channel-wise reflectance information for camera-specific wavelengths. The enriched spectral information, contained in a few to hundreds of channels, enables applications in a variety of fields, including segmentation and classification tasks in medicine (Seidlitz et al., 2022; Ayala et al., 2023), urban scene perception (Theisen et al., 2024; Shen et al.), and remote sensing (Lu et al., 2020; Thenkabail et al., 2018). To develop robust solutions for these tasks, data-driven models have emerged as the prevailing standard, maximizing performance through the utilization of all available images, regardless of camera characteristics. However, the evolution of spectral imaging technology has resulted in significant variability in camera devices (Qian, 2021), leading to the formation of camera-specific data silos. These silos share valuable domain-specific geometric information but differ in spectral characteristics such as channel dimensionality and covered wavelengths. Conventional imaging models such as Convolutional Neural Networks (CNNs) (He et al., 2016) cannot accommodate these variations, resulting in camera-specific models and absent knowledge transfer between these data silos. Therefore, such models ignore large amounts of data, limiting their robustness and cross-applicability.

Table 1: **Comparison of state-of-the-art spectral image encoding approaches.** The proposed model is the only one that incorporates all four desirable characteristics simultaneously.

Model	Wavelength-awareness	Channel-invariance	Spatio-spectral encoding	Spatio-spectral SSL pre-training
SpectralGPT ⁺	✗	✗	✓	✓
Spectral Earth	✗	✓	✓	✗
DOFA	✓	✓	✗	✗
Copernicus-FM	✓	✓	✗	✗
SMARTIES	✓	✓	✗	✗
CARL (Ours)	✓	✓	✓	✓

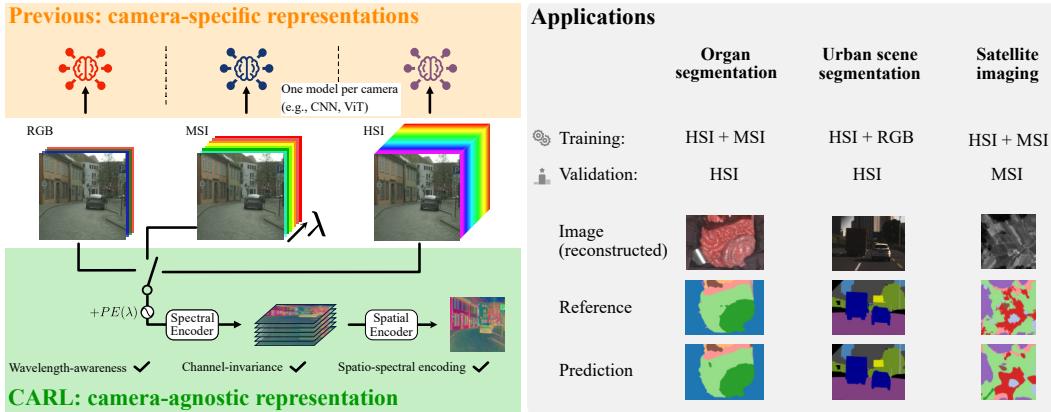


Figure 1: **CARL addresses spectral camera variations by learning camera-agnostic representations.** Unlike existing methods that require retraining for each channel configuration, CARL generalizes across cameras and outperforms both camera-specific and channel-invariant approaches across domains. The model processes one image at a time, ensuring flexibility without dependence on fusion strategies.

Furthermore, supervised downstream models are inherently limited by the availability of application-specific annotations. Given that manual labeling is time-intensive and often infeasible for large-scale datasets, self-supervised pre-training has emerged as a powerful alternative (He et al., 2022; Devlin et al., 2019; Caron et al., 2021; He et al., 2020). Empirical findings in Natural Language Processing have demonstrated that the effectiveness of self-supervised-learning (SSL) scales with the amount of training samples (Kaplan et al., 2020). This motivates the use of extensive cross-silo datasets to enhance pre-training. However, existing strategies are not camera-agnostic, restricting pre-training to camera-specific data silos and limiting their effectiveness. To overcome these obstacles, we propose a novel camera-agnostic model with a tailored SSL strategy that is capable of unlocking the data treasures of different cameras that are not yet accessible (Fig. 1). Our contribution is threefold:

1. **First approach to spatio-spectral camera-agnostic representation learning:** We propose the first method that enables spatio-spectral encoding in a camera-agnostic manner. To this end, we introduce wavelength positional encoding for establishing cross-camera channel correspondences, and learnable spectral representations for efficient representation learning.
2. **First camera-agnostic spatio-spectral self-supervision framework:** We propose a novel spectral feature-based SSL strategy tailored to CARL, which can be seamlessly combined with I-JEPA spatial pre-training (Assran et al., 2023) to form an end-to-end framework for camera-agnostic spatio-spectral self-supervised pre-training.
3. **Large-scale cross-domain validation:** We validated the proposed model in three application areas, specifically medical imaging, automotive vision, and satellite imaging. Across all experiments, our approach outperformed both camera-specific and channel-invariant baselines, demonstrating superior cross-modality knowledge transfer and unique robustness to spectral heterogeneity arising from simulated and real-world camera variations.

2 RELATED WORK

Feature extraction strategies for spectral imaging Generating rich image representations remains a fundamental challenge in computer vision, with significant implications for downstream tasks such as image segmentation. For 2D spectral images, the encoding process inherently spans three dimensions: two spatial dimensions and one spectral dimension. In datasets with uniform spectral properties, conventional models such as CNNs and Vision Transformers (ViTs) are commonly employed (Dosovitskiy et al., 2021; Theisen et al., 2024). However, these models focus solely on spatial encoding (2D projections, ViT blocks) and assume a fixed channel dimension. Recent approaches in remote sensing have introduced models that jointly encode spatial and spectral information, for example, by forming spatio-spectral patches (Cong et al., 2022; Hong et al., 2024). Yet, these methods lack

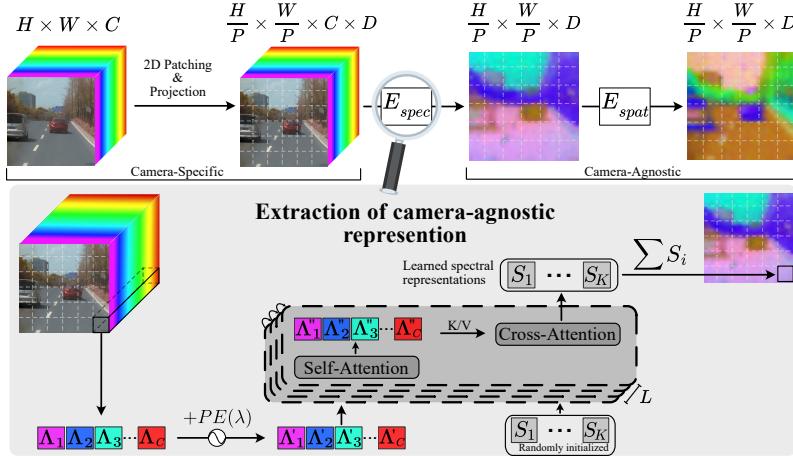


Figure 2: **Conversion of a camera-specific spectral image into a camera-agnostic representation.** To address the heterogeneity in camera-dependent spectral properties, a dedicated spectral encoder extracts a camera-agnostic representation by leveraging spectral tokens encoding wavelength information. A spectral image of dimension $H \times W \times C$ is divided into patches of size P and projected band-wise into a D -dimensional feature space. The spectral encoder E_{spec} processes each patch individually, and hereby resolves the spectral dimension. In particular, E_{spec} encodes the wavelength λ_i of channel i as positional encoding $PE(\lambda_i)$ and adds it to the embedded patch Λ_i . Self-Attention across spectral tokens $(\Lambda_i)_{i \leq C}$ and Cross-Attention with K learned spectral representations yield enriched representations $(S_j)_{j \leq K}$. After aggregation into a camera-agnostic representation, a standard image encoder, E_{spat} , captures spatial relationships.

invariance to the channel dimension, preventing their applicability to cross-camera datasets. Braham et al. (2024) addresses this by introducing a Spectral Adapter that resolves the channel dimension through 1D convolutions and pooling operations. However, it overlooks channel relationships derived from camera-specific wavelength information. To overcome this limitation, recent work has proposed channel-adaptive 2D projection layers that learn wavelength-conditioned projection matrices (Li et al., 2025a; Xiong et al., 2024; Varga et al., 2023; Wang et al., 2025). While effective, these methods rely on spatial operations and do not explicitly encode salient spectral information, which may reduce robustness on spectrally heterogeneous datasets. Alternative strategies encode pixel- or patch-wise information only along the channel dimension (Hong et al., 2021; Hang et al., 2019; Seidlitz et al., 2022). However, these designs cannot model geometric structures, which significantly impairs downstream performance. Fusion-based methods offer another solution by aligning multi-modal data with varying channel dimensions through early, mid, or late fusion (Lin et al., 2023; Audebert et al., 2018). Typically, these architectures incorporate modality-specific projection layers or entire encoders before feeding into a multi-modal encoder (Astruc et al., 2025; Tseng et al., 2025; Jakubik et al., 2025; Fuller et al., 2023). Crucially, such methods assume access to all modalities during both training and inference. While this assumption may hold in remote sensing with standardized sensors, it is impractical for industrial or medical spectral imaging applications, where sensor diversity is broader, often unknown, and sensor-specific data silos contain relatively few samples.

Self-supervised learning strategies for spectral imaging With growing compute resources and data availability, SSL strategies have gained importance in recent years. In RGB imaging, masked image modeling has emerged as a central paradigm, where input patches are randomly masked and reconstructed at the pixel level using a strong encoder paired with a lightweight decoder (He et al., 2022). This paradigm has been extended to spectral imaging, encompassing camera-specific spatio-spectral encoding (Hong et al., 2024; Cong et al., 2022) and camera-agnostic spatial modeling (Xiong et al., 2024; Wang et al., 2025; Sumbul et al., 2025). Building on advances in RGB vision, feature-based SSL has proven more efficient than pixel-based approaches (Caron et al., 2021; Assran et al., 2023), a benefit that is particularly important in spectral imaging, where pixel values are highly sensitive to factors such as atmospheric conditions in satellite data or illumination calibration in laboratory settings (Baumann et al., 2024). Accordingly, feature-based SSL has been adapted to

162 spectral imaging models, though existing methods remain restricted to spatial encoding and spatial
 163 self-supervision (Tseng et al., 2025; Astruc et al., 2025; Waldmann et al., 2025). Notably, no SSL
 164 framework—whether camera-agnostic, feature-based, or both—has yet been designed to capture
 165 spatio-spectral encoding. An overview of existing approaches is provided in Tab. 1.

168 3 FRAMEWORK FOR CAMERA-AGNOSTIC SPECTRAL IMAGE ANALYSIS

170 In this paper, we present CARL, a novel model for spectral image processing, designed to unlock the
 171 potential of camera-specific data silos. As illustrated in Fig. 2, the proposed framework transforms
 172 camera-dependent spectral information into a camera-agnostic representation through a novel spectral
 173 encoder E_{spec} , followed by the extraction of geometric information through a standard spatial
 174 encoder E_{spat} . To establish cross-camera channel correspondences, we translate the concept of
 175 positional encoding, traditionally used for discrete token positions within transformers (Vaswani
 176 et al., 2017), to channel-specific wavelengths. To facilitate efficient representation learning along
 177 the spectral dimension, we propose a novel encoder that distills channel information into a sparse
 178 set of spectral representations. Ultimately, the encoder produces a camera-agnostic feature map
 179 enriched with spectral attributes, which can be seamlessly forwarded to established transformer-based
 180 spatial encoders (e.g., ViT). The spatial encoder E_{spat} operates subsequent to E_{spec} and enhances
 181 feature representation by capturing spatial relationships. Two learning paradigms are proposed for
 182 optimizing the spectral representations derived in E_{spec} . Specifically, they are either learned implicitly
 183 by minimizing a downstream task-specific loss or explicitly by minimizing a self-supervised loss
 (CARL-SSL, as described in Sec. 3.2 and Fig. 3).

185 3.1 ARCHITECTURE OF CARL

187 Given a spectral image $I \in \mathbb{R}^{H \times W \times C}$ with arbitrary channel dimension, C , the objective is to derive
 188 a camera-agnostic representation that contains salient spectral information. To project I to feature
 189 space, each channel is processed by a shared 2D convolution with kernel size and stride equal to
 190 the patch size P and output channels D , yielding a tensor of dimensionality $\frac{H}{P} \times \frac{W}{P} \times C \times D$. A
 191 patch, denoted as $\Lambda = (\Lambda_1, \dots, \Lambda_C) \in \mathbb{R}^{C \times D}$, is then independently processed by the spectral
 192 encoder, E_{spec} , to generate a camera-agnostic representation. We first construct a positional encoding
 193 $PE(\lambda) \in \mathbb{R}^{C \times D}$ where $\lambda = (\lambda_1, \dots, \lambda_C)$ and λ_i corresponds to the wavelength of channel i such
 194 that the model is capable of establishing channel correspondences across cameras with different
 195 wavelength specifications. In this work, we employ the sinusoidal Fourier Features (Tancik et al.,
 196 2020) to encode positional information within the spectral dimension, defined as:

$$198 PE(\lambda_i) = [\cos(2\pi\alpha\lambda_i B), \sin(2\pi\alpha\lambda_i B)]^T \in \mathbb{R}^D \quad (1)$$

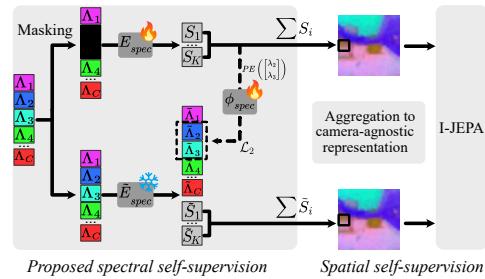
200 where $\alpha \in \mathbb{R}$ is a scaling factor, $B \sim \mathcal{N}(0, \sigma^2 \mathbf{I}) \in \mathbb{R}^{D/2}$, and $\sigma \in \mathbb{R}$. Here, both α and σ are
 201 hyperparameters. Subsequently, $PE(\lambda)$ is added to the patch Λ , thereby encoding the position
 202 of each Λ_i along the wavelength axis. As illustrated in Fig. 2, a self-attention-cross-attention
 203 module is introduced to process the spectral tokens, $(\Lambda_i)_{i \leq C}$, and derive spectral representations.
 204 Specifically, K learnable D -dimensional spectral representations, denoted as $(S_j)_{j \leq K}$, are initialized
 205 from a truncated normal distribution. Following a self-attention block applied to the spectral tokens,
 206 the spectral representations, $(S_j)_{j \leq K}$, attend to the spectral tokens, $(\Lambda_i)_{i \leq C}$, via cross-attention,
 207 effectively distilling the most salient information. This self-attention-cross-attention module is iterated
 208 L times to learn enriched spectral representations, $(S_j)_{j \leq K}$. Subsequently, a readout function, in
 209 this instance summation, is applied to $(S_j)_{j \leq K}$ to aggregate the information into a camera-agnostic
 210 representation for the patch Λ . As the spectral encoder generates such a representation for each patch
 211 independently, the incorporation of spatial relationships necessitates the utilization of a subsequent
 212 spatial encoder. It is noteworthy that since E_{spec} has encoded device-dependent spectral properties
 213 within the feature space, most common transformer-based spatial encoders, such as ViT, may be
 214 employed for spatial encoding. To ensure dimensional compatibility between the spectral and spatial
 215 encoder, layer normalization and a linear transformation are applied prior to spatial encoding. After
 capturing inter-patch relationships, a task-specific head can be added for the intended downstream
 application.

216 3.2 SELF-SUPERVISED TRAINING STRATEGY
217

218 Tailored to CARL, we propose a self-supervised pre-training strategy, CARL-SSL, to leverage
219 large-scale unlabeled datasets. As illustrated in Fig. 3, the procedure is disentangled into spectral
220 and spatial self-supervised pre-training within an end-to-end framework. While I-JEPA (Assran
221 et al., 2023) is adapted for spatial self-supervision, we introduce a novel feature-based spectral SSL
222 strategy. Given student encoders E_{spec} and E_{spat} from Sec. 3.1 along with their teacher counterparts,
223 \tilde{E}_{spec} and \tilde{E}_{spat} , which are updated via exponential moving average, we apply a masking strategy
224 to specific regions of the students’ input. The remaining tokens are then encoded by the student
225 networks. The SSL objective is to predict the masked features generated by the teacher encoders,
226 using dedicated predictors ϕ_{spec} and ϕ_{spat} . Specifically, for an image $I \in \mathbb{R}^{H \times W \times C}$, the initial
227 convolution is applied as described in Sec. 3.1. A spectral mask, denoted by $M \subseteq \{1, \dots, C\}$,
228 containing the masked channel indices, is sampled for a patch Λ , and the unmasked tokens, de-
229 noted by $(\Lambda_i)_{i \notin M}$, are forwarded to the student spectral encoder, E_{spec} , to generate spectral rep-
230 resentations, $(S_j)_{j \leq K}$ (see Fig. 3). Conversely, the teacher spectral encoder receives all spec-
231 tral tokens as input, producing learned spectral tokens, $(\tilde{\Lambda}_i)_{i \leq C}$, via self-attention, and spectral
232 representations, $(\tilde{S}_j)_{j \leq K}$. The objective of spectral pre-training is then to predict the masked
233 spectral tokens, $(\tilde{\Lambda}_i)_{i \in M}$, based on the student spectral representations, $(S_j)_{j \leq K}$, and the pos-
234 tional encoding of the masked wavelengths. To this end, a transformer-based predictor, denoted by
235 ϕ_{spec} , is employed, receiving as input a sequence with the spectral representations and dedicated
236 mask tokens (Devlin et al., 2019). The mask tokens are $|M|$ copies of a shared, learnable embed-
237 ding and are summed with wavelength positional encoding $(PE(\lambda_i))_{i \in M}$, corresponding to the
238 masked wavelengths. Subsequently, ϕ_{spec} processes this sequence through multiple self-attention
239 blocks, resulting in learned mask tokens as predictions for $(\tilde{\Lambda}_i)_{i \in M}$. Network optimization uses the
240 VICReg loss (Bardes et al., 2021) on the spectral
241 predictions, denoted as $\mathcal{L}_{\text{spec}}$, which comprises invari-
242 ance, variance, and covariance terms. The invariance
243 term minimizes the mean-squared error between pre-
244 dicted and target spectral tokens, while the variance
245 and covariance terms contribute to training stability
246 and the prevention of feature collapse. For joint spatial
247 training, the spectral representations from both
248 the student and teacher encoders are aggregated into
249 2D camera-agnostic representations. Subsequently,
250 a 2D region of the student’s feature representation
251 is masked, and the remaining spatial tokens are pro-
252 cessed by E_{spat} . Analogous to spectral pre-training,
253 the spatial predictor, ϕ_{spat} , receives the student fea-
254 tures and the positional encoding of the masked to-
255 kens as input, and predicts the masked features gen-
256 erated by the teacher spatial encoder, \tilde{E}_{spat} . The spatial
257 loss function, $\mathcal{L}_{\text{spat}}$, is likewise defined using the VI-
258 CReg loss on the spatial predictions. Finally, the over-
259 all training objective is given by $\mathcal{L} = \mathcal{L}_{\text{spat}} + \mathcal{L}_{\text{spec}}$ to
260 jointly optimize the student encoders and predictors.
261 Further details are provided in appendix A.2.
262

263 3.3 IMPLEMENTATION DETAILS
264

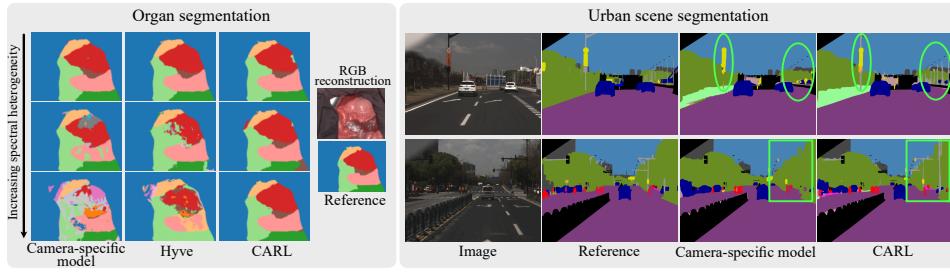
265 Following ablation studies (see Sec. 5), we set $\sigma = 3$, as defined in Eq. (1) for the wavelength
266 positional encoding, and the number of spectral representations within the spectral encoder to
267 $K = 8$. Importantly, these hyperparameters are held constant across all three application domains,
268 underscoring their generality. Unless otherwise stated in the experiments, we employed EVA-02
269 (Fang et al., 2024) as spatial encoder, which is a modern version of ViT (Dosovitskiy et al., 2021).
270 Further implementation details of the model are outlined in appendix A.



271 **Figure 3: CARL-SSL enables joint learning of camera-agnostic representations and spatial
272 relations.** Spectral self-supervision involves recon-
273 struction of masked spectral channels in feature
274 space. The student E_{spec} extracts spectral rep-
275 resentations $(S_j)_{j \leq K}$ from a spectrally masked input
276 patch, while the predictor ϕ_{spec} predicts the masked
277 spectral tokens using masked wavelengths (here:
278 λ_2, λ_3), and $(S_j)_{j \leq K}$. Target tokens are gen-
279 erated by the teacher \tilde{E}_{spec} from the complete input.
280 The aggregated camera-agnostic representations
281 are subsequently processed by I-JEPA.

270
 271 **Table 2: CARL-SSL demonstrates superior performance compared with both camera-specific**
 272 **and camera-agnostic models.** The mIoU scores with the 95 % confidence intervals on the HSICity
 273 test set. While the camera-specific model was pre-trained on Cityscapes and fine-tuned exclusively
 274 on HSICity, the other models are channel-invariant adaptations which were concurrently trained on
 both datasets.

	Camera-specific model	Spectral Adapter	HyperFree	Hyve	DOFA	CARL	CARL-SSL
mIoU	44.6 [40.9; 47.3]	43.4 [41.0; 45.2]	44.6 [42.2; 46.5]	48.0 [45.4; 50.0]	49.6 [46.8; 51.6]	48.6 [45.6; 51.0]	50.1 [47.2; 52.4]



280
 281 **Figure 4: Our model enables cross-modality knowledge transfer.** **(Left)** With increasing spectral
 282 heterogeneity in the training set, both the camera-specific model and Hyve exhibit a notable rise
 283 in prediction noise. In contrast, CARL consistently provides accurate predictions. **(Right)** In two
 284 HSICity test set examples, the HSI-specific model fails to segment "poles" (gray labels) due to their
 285 absence in the HSICity training set. CARL, however, jointly trained on RGB and HSI data, effectively
 286 leverages "pole" annotations from Cityscapes to inform its predictions on HSICity.

295 4 EXPERIMENTS AND RESULTS

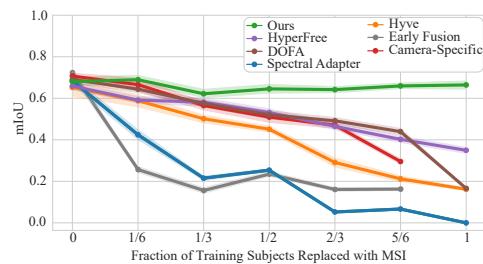
296 The experiments aimed to address the following research questions pertaining to the proposed model:

300 (RQ1) In silico proof of concept: Does the spectral representation learning approach enable effective
 301 knowledge transfer across cameras?
 302 (RQ2) Real-world cross-domain generalization: To what extent do the proposed model and the
 303 spectral self-supervision approach help in handling real-world cross-camera variations?

305 Classification performance is reported as overall accuracy (OA), whereas segmentation performance
 306 is measured by intersection-over-union (IoU).

308 4.1 IN SILICO PROOF OF CONCEPT: MEDICAL IMAGING

310 **Datasets.** Synthetic multispectral images were generated from a real hyperspectral dataset, enabling isolated control over spectral variations while preserving spatial context. We employed a private collection of 311 porcine organ images, semantically annotated into 19 classes (Seidlitz et al., 2022). Acquisition was 312 performed with a Tivita® Tissue HSI camera (Disapctive Vision GmbH, Am Salzhaff, Germany), capturing 313 100 spectral channels spanning 500 nm to 1,000 nm. The training set comprises 254 images 314 from 12 subjects. To emulate realistic multispectral acquisition with optical filters, we modeled each filter 315 response as a Gaussian function. Specifically, the 316 number of channels C in a virtual multispectral camera 317 was first sampled uniformly from $\{10, \dots, 25\}$. The 318 corresponding C center wavelengths, serving as 319



320 **Figure 5: Our model shows unique robustness**
 321 **to spectral heterogeneity in the organ experiments.** As spectral heterogeneity increases with the
 322 multispectral replacements in the training set, CARL uniquely maintains a high mIoU score on
 323 the hyperspectral test set. Shaded area: 95 % confidence interval.

324 Table 3: **CARL excels at cross-sensor learning across different application domains.** (a)
325 Hyperspectral images per subject (S) transformed to multispectral images using real-world filters (F).
326 CARL achieves stable performance in contrast to existing methods. (b) CARL exhibits the small
327 drop in IoU when traffic light and sign classes were removed from HSI training set. (c) Ablation
328 study on the loss components of CARL-SSL. Overall accuracy (OA) is reported on the m-forestnet
329 validation set using a feature-based k -NN classifier.

(a) Medical Data				(b) Automotive Data				(c) Ablation on CARL-SSL	
Training Data	Hyve	DOFA	CARL	Class	Hyve	DOFA	CARL	SSL Strategy	OA
Synthetic MSI (4S, 1F) Real HSI (8S)	52.1	58.1	64.6	Traffic Light (Δ)	15.5 -37.5	8.4 -50.4	29.2 -24.5	Spatial SSL $\mathcal{L}_{\text{spat}}$ + Spectral SSL $\mathcal{L}_{\text{spec}}$	22.1 32.6
Synthetic MSI (4S, 2F) Real HSI (8S)	47.7	49.2	60.3	Traffic Sign (Δ)	10.9 -43.8	14.9 -45.0	31.7 -26.7	= CARL-SSL	
Synthetic MSI (6S, 1F) Real HSI (6S)	35.4	52.0	62.1	mIoU (Δ)	42.7 -5.3	42.7 -6.9	46.2 -2.4		

340 the Gaussian means, were selected within 550 nm to 950 nm using farthest point sampling. To obtain
341 a realistic range of variances, we fitted Gaussian functions to the filter responses of a commercial
342 multispectral camera and sampled C variance values from this range. Given the sampled means
343 and variances of C channels, the filter functions were constructed and applied to an hyperspectral
344 image to generate the corresponding multispectral image via matrix multiplication. In this way,
345 we simulated six distinct camera configurations and progressively replaced hyperspectral images
346 in the training set with their multispectral counterparts on a per-subject basis, while keeping the
347 hyperspectral validation and test sets unchanged. This protocol isolates spectral variability and allows
348 for a rigorous assessment of model robustness to spectral heterogeneity. Additional details on the
349 data generation can be found in appendix F.

350 **Baseline methods.** The model’s performance was benchmarked against a camera-specific baseline
351 and six channel-invariant methods, which can be grouped into three categories: spatio-spectral en-
352 coding (Spectral Adapter (Braham et al., 2024)); channel-adaptive embedding layers (DOFA (Xiong
353 et al., 2024), Hyve (Varga et al., 2023), and HyperFree (Li et al., 2025a)); and camera-specific
354 embedding layers (Early Fusion (Astruc et al., 2025)). The camera-specific model employs a standard
355 U-Net, representing the state of the art on the original hyperspectral dataset (Seidlitz et al., 2022), and
356 was trained exclusively on the hyperspectral subset of each training set variant. The other methods
357 were integrated with either a U-Net or a ViT-based architecture, depending on which proved most
358 compatible. For ViT-based methods, including CARL, we adopted the ViT-Adapter (Chen et al.,
359 2023) for hierarchical features and UperNet (Xiao et al., 2018) for segmentation. All models followed
360 the same training protocol.

361 **Results.** The mIoU scores as a function of the fraction of multispectral subjects within the training
362 set is presented in Fig. 5. The proposed method uniquely maintained a high mIoU across all training
363 set variants. This is qualitatively confirmed in Fig. 4, where prediction noise of the baseline methods
364 increases with spectral heterogeneity, while our model remains stable and accurate. While simulated
365 filter functions enable scalable evaluation under spectral heterogeneity, realistic modeling remains
366 crucial. We therefore conducted an additional experiment in which multispectral images were syn-
367 thesized using real-world filters instead of Gaussian approximations. As shown in Tab. 3a, CARL
368 consistently outperforms baseline methods across all training data configurations.

369 4.2 REAL-WORLD EVALUATION: AUTOMOTIVE

370 The purpose of this experiment was to test the capabilities of our approach under real-world conditions
371 in the context of autonomous driving. Specifically, we investigated whether our model can effectively
372 leverage RGB and hyperspectral images for urban scene segmentation.

373 **Datasets.** The Cityscapes RGB dataset (Cordts et al., 2016), comprising semantic annotations for
374 19 classes, was employed alongside its hyperspectral counterpart, HSICity (Shen et al.). HSICity
375 contains images with 128 channels (450 nm to 950 nm) and shares the same labels. However, its
376 training set suffers from coarse annotations and class imbalance; for instance, the "pole" class is
377 present in the test set but absent from the training set. We therefore added finely annotated
378 Cityscapes images containing "pole" labels, resulting in 4,029 training images (1,054 from HSICity).

378 Table 4: **CARL learns strong in-distribution features.** Linear-probing results on four Sentinel-2
 379 benchmarks. CARL attains the highest accuracy on three of the four datasets and the best average
 380 rank across all eleven benchmark datasets.

	m-bigearthnet	m-eurosat	m-cashew	m-SA-crop-type	Rank over 11 datasets
SpectralGPT ⁺	45.0	69.9	14.5	13.7	5.5
Galileo	49.8	84.2	10.5	19.3	5.5
Croma	59.5	86.6	12.1	25.2	5.0
DOFA	61.0	89.9	18.2	21.7	3.2
Copernicus-FM	62.1	87.2	14.5	26.5	2.6
SMARTIES	62.0	92.6	12.7	24.3	2.6
CARL	69.0	84.4	18.9	26.5	1.6

390 Table 5: **CARL produces robust features on unseen sensors.** Linear-probing results on four out-of-
 391 distribution sensors. Despite large spectral heterogeneity, CARL achieves the best performance on
 392 three of the four datasets, demonstrating strong cross-sensor generalization.

	LoveDA Urban	m-forestnet	WHU-OHS	Wuhan
Sensor	RGB (3 bands)	LandSat-8 (6 bands)	Orbita (32 bands)	Gaofen-5 (116 bands)
DOFA	12.6	43.8	1.5	20.3
Copernicus-FM	15.4	44.8	1.5	18.1
SMARTIES	13.5	49.8	1.5	18.8
CARL	29.0	47.0	21.7	21.5

402 To assess cross-modality learning, models were trained on this combined dataset and evaluated on
 403 the HSICity test set. For self-supervised pre-training, we leveraged a collection of heterogeneous
 404 urban datasets, including Cityscapes, HSICity, and the multispectral datasets HyKo-VIS (Winkens
 405 et al., 2017) and HSIDrive (Basterretxea et al., 2021).

406 **Baseline methods.** A Swin Transformer (Liu et al., 2021) with a Mask2Former head (Cheng et al.,
 407 2022), referred to as SwinMask2Former, pre-trained on the Cityscapes dataset, was adapted through
 408 the integration of a channel-invariant module. Our spectral encoder, serving as this module, was
 409 compared with channel-adaptive layers of HyperFree, DOFA, and Hyve, as well as with the Spectral
 410 Adapter. Additionally, as a camera-specific baseline, the RGB-pre-trained SwinMask2Former was
 411 trained exclusively on HSICity. All models adhered to an identical training protocol.

412 **Results.** The mIoU scores on the HSICity test set are presented in Tab. 2. CARL-SSL demonstrated
 413 superior performance compared to the baseline methods. Due to the absence of the "pole" class in the
 414 HSICity training set, the camera-specific model failed to segment any poles in the test set, despite
 415 RGB-pre-training (see Fig. 4). In contrast, our model effectively transferred knowledge from the
 416 "pole" labels in Cityscapes to improve its predictions on HSICity, achieving the highest IoU for this
 417 class. To further assess cross-modality learning, we removed the "traffic light" and "traffic sign"
 418 classes from the HSICity training set, and re-trained the models. As shown in Tab. 3b, CARL most
 419 effectively leveraged RGB supervision to achieve superior HSI predictions for the excluded classes
 420 on the test set compared to the baselines.

421 4.3 REAL-WORLD EVALUATION: SATELLITE IMAGING

422 The third experiment aimed to evaluate the capabilities of CARL-SSL in satellite imaging.

423 **Dataset & Baseline methods.** To facilitate benchmarking against strong pre-trained baselines such
 424 as SpectralGPT⁺ and DOFA, we scaled up pre-training to a corpus of approximately 800,000 images,
 425 comprising Sentinel-2 multispectral data (Sumbul et al., 2019) and EnMAP hyperspectral data (Fuchs
 426 and Demir, 2023; Braham et al., 2024). **Feature quality was assessed via linear probing on eleven**
 427 **datasets (Lacoste et al., 2023; Li et al., 2022; Wang et al., 2021; Hong et al., 2023).** These include five
 428 in-distribution datasets (Sentinel-2) and six datasets captured by sensors not present in the pre-training
 429 set. Further dataset details are given in Tabs. 7, 8. We compared against six remote-sensing foundation
 430 models, three of which are camera-agnostic and therefore can be evaluated on every dataset.
 431 **Results.** Results are summarized in Tabs. 4, 5, 9. CARL delivers consistently strong performance
 432 across the benchmarks and achieves the best average rank across all eleven datasets. In particular,

432 **Table 6: Ablation studies on CARL.** **(a)** Wavelength positional encoding (PE) is essential, with
 433 $\sigma = 3$ yielding the best performance. **(b)** Summation proves to be the most effective strategy for
 434 aggregating spectral representations, and **(c)** using $K = 8$ spectral representations suffices to distill
 435 the channels. **(d)** Moreover, CARL benefits from larger embedding dimensions.

(a) Positional Encoding		(b) Aggregation		(c) # Spectral Rep.		(d) Feature Dim.	
Method	mIoU	Method	mIoU	# Spectral Rep.	mIoU	Size	mIoU
No PE	18.3	Summation	62.7	$K = 1$	57.8	$D = 384$	64.4
PE ($\sigma = 1$)	55.1	Concatenation	61.8	$K = 4$	58.2	$D = 768$	66.2
PE ($\sigma = 3$)	61.5	Maximum	61.8	$K = 8$	63.9		
PE ($\sigma = 10$)	57.2	Attention Pooling	60.0	$K = 16$	62.2		

444 CARL exhibits robust generalization to OOD sensors, outperforming the second-best method by
 445 a substantial margin on several datasets. We attribute this advantage to CARL’s camera-agnostic
 446 spatio-spectral encoding scheme, which is unique in comparison to the baseline methods.

448 5 ABLATION STUDY

450 The ablation study in Tab. 6 examines the contribution of CARL’s key architectural components.
 451 Training was conducted on a dataset variant from Sec. 4.1, while evaluation was performed on an
 452 HSI validation split. Removing the wavelength positional encoding severely impairs the model’s
 453 ability to align channels across different cameras. For aggregating the spectral representations, simple
 454 summation achieves the highest accuracy while remaining computationally efficient. With respect
 455 to the number of spectral representations, we find that $K = 8$ tokens are sufficient to capture the
 456 essential spectral information. Notably, performance gains beyond this point are primarily achieved
 457 by increasing the embedding dimension rather than K . We also conducted an analysis of CARL-
 458 SSL’s two self-supervision tasks—spectral and spatial—as summarized in Tab. 3c. To isolate their
 459 individual contributions, we performed two pre-training variants under a reduced training budget:
 460 one employing only spatial self-supervision, and the other utilizing our proposed spatio-spectral
 461 strategy. The resulting image embeddings were evaluated using a k -NN classifier on the m-forestnet
 462 validation set (Lacoste et al., 2023). The model trained with spatial self-supervision alone exhibited
 463 a collapse in the spectral representations, leading to significantly reduced accuracy. In contrast,
 464 incorporating spectral self-supervision effectively mitigated this collapse and yielded substantially
 465 stronger representations, resulting in a +10.5 OA improvement.

466 6 DISCUSSION

468 We introduced CARL, to our knowledge the first camera-agnostic framework that unifies spatio-
 469 spectral encoding with spatio-spectral SSL pre-training. Our approach tackles a critical gap in spectral
 470 image processing: the lack of a representation learning framework that generalizes across spectrally
 471 heterogeneous datasets. We demonstrated its effectiveness in both traditional satellite imaging and
 472 in domains such as medical imaging, where sensor variability is particularly pronounced due to
 473 the diversity of commercial camera manufacturers. Adaptive embedding approaches, including
 474 HyperFree, Hyve, and DOFA, rely solely on spatial operations and neglect crucial inter-channel
 475 relationships, resulting in limited performance. Alternative models, such as SpectralGPT⁺, intro-
 476 duce spatio-spectral encoding, but depend on fixed channel dimensions, preventing generalization
 477 across different spectral cameras. In contrast, CARL employs wavelength-aware spatio-spectral
 478 encoding that is independent of channel dimensionality, enabling robust generalization under spectral
 479 heterogeneity and scalability across modalities, as demonstrated in both pre-training (Tabs. 2, 4, 5,
 480 9) and downstream tasks (Figs. 4, 5, Tab. 3). Limitations of our work include higher computational
 481 cost compared to channel-adaptive embedding approaches (see appendix D), as well as challenges
 482 from sensor heterogeneity beyond spectral properties, such as differences in spatial resolution. The
 483 latter may be mitigated by incorporating additional sensor metadata. Despite these limitations, our
 484 approach successfully integrates real-world cross-camera datasets and outperforms existing methods.
 485 By unlocking the untapped potential of cross-sensor spectral datasets, CARL paves the way toward a
 486 more universal and accessible future in spectral imaging.

486 7 REPRODUCIBILITY STATEMENT

488 To ensure reproducibility, we have provided a detailed description of our method in Sec. 3, with
 489 additional implementation details in appendix A. Pseudo-code illustrating the forward passes of
 490 CARL and CARL-SSL can be found in algorithms 1, 2. Furthermore, all datasets used in Secs. 4.2,
 491 4.3 are publicly available. Finally, full per-class IoU scores are reported in Tabs. 10, 12.

493 REFERENCES

495 Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun,
 496 and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In
 497 *CVPR*, pages 15619–15629, 2023.

498 Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. Omnisat: Self-supervised modality
 499 fusion for earth observation. In *European Conference on Computer Vision*, pages 409–427. Springer, 2024.

500 Guillaume Astruc, Nicolas Gonthier, Clement Mallet, and Loic Landrieu. Anysat: One earth observation model
 501 for many resolutions, scales, and modalities. In *Proceedings of the Computer Vision and Pattern Recognition
 502 Conference*, pages 19530–19540, 2025.

503 Nicolas Audebert, Bertrand Le Saux, and Sébastien Lefèvre. Beyond rgb: Very high resolution urban remote
 504 sensing with multimodal deep networks. *ISPRS journal of photogrammetry and remote sensing*, 140:20–32,
 505 2018.

506 Leonardo Ayala, Tim J Adler, Silvia Seidlitz, Sebastian Wirkert, Christina Engels, Alexander Seitel, Jan Sellner,
 507 Alexey Aksenov, Matthias Bodenbach, Pia Bader, et al. Spectral imaging enables contrast agent-free real-time
 508 ischemia monitoring in laparoscopic surgery. *Science advances*, 9(10):eadd6778, 2023.

509 Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization for
 510 self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.

511 Koldo Basterretxea, Victoria Martínez, Javier Echanobe, Jon Gutiérrez-Zaballa, and Inés Del Campo. Hsi-drive:
 512 A dataset for the research of hyperspectral image processing applied to autonomous driving systems. In *2021
 513 IEEE Intelligent Vehicles Symposium (IV)*, pages 866–873. IEEE, 2021.

514 Alexander Baumann, Leonardo Ayala, Alexander Studier-Fischer, Jan Sellner, Berkin Özdemir, Karl-Friedrich
 515 Kowalewski, Slobodan Ilic, Silvia Seidlitz, and Lena Maier-Hein. Deep intra-operative illumination calibration
 516 of hyperspectral cameras. In *International Conference on Medical Image Computing and Computer-Assisted
 517 Intervention*, pages 120–131. Springer, 2024.

518 Nassim Ait Ali Braham, Conrad M Albrecht, Julien Mairal, Jocelyn Chanussot, Yi Wang, and Xiao Xiang Zhu.
 519 Spectralearth: Training hyperspectral foundation models at scale. *arXiv preprint arXiv:2408.08447*, 2024.

520 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin.
 521 Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9650–9660, 2021.

522 Zhe Chen, Yuchen Duan, Wenhui Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer
 523 adapter for dense predictions. In *The Eleventh International Conference on Learning Representations*, 2023.
 524 URL <https://openreview.net/forum?id=plKu2GByCNW>.

525 Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention
 526 mask transformer for universal image segmentation. In *CVPR*, pages 1290–1299, 2022.

527 Yezhen Cong, Samar Khanna, Chenlin Meng, Patrick Liu, Erik Rozi, Yutong He, Marshall Burke, David Lobell,
 528 and Stefano Ermon. Satmae: Pre-training transformers for temporal and multi-spectral satellite imagery. In
 529 S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *NeurIPS*, volume 35, pages 197–
 530 211. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/01c561df365429f33fc7a7faa44c985-Paper-Conference.pdf.

531 Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe
 532 Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In
 533 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.

534 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical
 535 image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee,
 536 2009.

540 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional
 541 transformers for language understanding. In *Proceedings of the 2019 conference of the North American*
 542 *chapter of the association for computational linguistics: human language technologies, volume 1 (long and*
 543 *short papers)*, pages 4171–4186, 2019.

544 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner,
 545 Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby.
 546 An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. URL
 547 <https://openreview.net/forum?id=YicbFdNTTy>.

548 Yuxin Fang, Quan Sun, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva-02: A visual
 549 representation for neon genesis. *Image and Vision Computing*, 149:105171, 2024.

550 Martin Hermann Paul Fuchs and Begüm Demir. Hyspecnet-11k: A large-scale hyperspectral dataset for
 551 benchmarking learning-based hyperspectral image compression methods. In *IGARSS 2023-2023 IEEE*
 552 *International Geoscience and Remote Sensing Symposium*, pages 1779–1782. IEEE, 2023.

553 Anthony Fuller, Koreen Millard, and James Green. Croma: Remote sensing representations with contrastive
 554 radar-optical masked autoencoders. *Advances in Neural Information Processing Systems*, 36:5506–5538,
 555 2023.

556 Yuval Garini, Ian T Young, and George McNamara. Spectral imaging: principles and applications. *Cytometry*
 557 *part a: the journal of the international society for analytical cytology*, 69(8):735–747, 2006.

558 Renlong Hang, Qingshan Liu, Danfeng Hong, and Pedram Ghamisi. Cascaded recurrent neural networks for
 559 hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 57(8):5384–5394,
 560 2019.

561 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In
 562 *CVPR*, pages 770–778, 2016.

563 Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised
 564 visual representation learning. In *CVPR*, pages 9729–9738, 2020.

565 Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are
 566 scalable vision learners. In *CVPR*, pages 16000–16009, 2022.

567 Danfeng Hong, Zhu Han, Jing Yao, Lianru Gao, Bing Zhang, Antonio Plaza, and Jocelyn Chanussot. Spectral-
 568 former: Rethinking hyperspectral image classification with transformers. *IEEE Transactions on Geoscience*
 569 *and Remote Sensing*, 60:1–15, 2021.

570 Danfeng Hong, Bing Zhang, Hao Li, Yuxuan Li, Jing Yao, Chenyu Li, Martin Werner, Jocelyn Chanussot,
 571 Alexander Zipf, and Xiao Xiang Zhu. Cross-city matters: A multimodal remote sensing benchmark dataset
 572 for cross-city semantic segmentation using high-resolution domain adaptation networks. *Remote Sensing of*
 573 *Environment*, 299:113856, 2023.

574 Danfeng Hong, Bing Zhang, Xuyang Li, Yuxuan Li, Chenyu Li, Jing Yao, Naoto Yokoya, Hao Li, Pedram
 575 Ghamisi, Xiuping Jia, et al. Spectralgpt: Spectral remote sensing foundation model. *IEEE TPAMI*, 2024.

576 Johannes Jakubik, Felix Yang, Benedikt Blumenstiel, Erik Scheurer, Rocco Sedona, Stefano Maurogiovanni,
 577 Jente Bosmans, Nikolaos Dionelis, Valerio Marsocci, Niklas Kopp, et al. Terramind: Large-scale generative
 578 multimodality for earth observation. *arXiv preprint arXiv:2504.11171*, 2025.

579 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray,
 580 Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint*
 581 *arXiv:2001.08361*, 2020.

582 Alexandre Lacoste, Nils Lehmann, Pau Rodriguez, Evan Sherwin, Hannah Kerner, Björn Lütjens, Jeremy Irvin,
 583 David Dao, Hamed Alemohammad, Alexandre Drouin, et al. Geo-bench: Toward foundation models for earth
 584 monitoring. *Advances in Neural Information Processing Systems*, 36:51080–51093, 2023.

585 Jiayi Li, Xin Huang, and Lilin Tu. Whu-ohs: A benchmark dataset for large-scale herseptral image classification.
 586 *International Journal of Applied Earth Observation and Geoinformation*, 113:103022, 2022.

587 Jingtao Li, Yingyi Liu, Xinyu Wang, Yunning Peng, Chen Sun, Shaoyu Wang, Zhendong Sun, Tian Ke, Xiao
 588 Jiang, Tangwei Lu, et al. Hyperfree: A channel-adaptive and tuning-free foundation model for hyperspectral
 589 remote sensing imagery. *arXiv preprint arXiv:2503.21841*, 2025a.

594 Xuyang Li, Chenyu Li, Pedram Ghamisi, and Danfeng Hong. Fleximo: A flexible remote sensing foundation
 595 model. *arXiv preprint arXiv:2503.23844*, 2025b.

596

597 Junyan Lin, Feng Gao, Xiaochen Shi, Junyu Dong, and Qian Du. Ss-mae: Spatial–spectral masked autoencoder
 598 for multisource remote sensing image classification. *IEEE Transactions on Geoscience and Remote Sensing*,
 61:1–14, 2023.

599

600 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin
 601 transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021.

602 Bing Lu, Phuong D Dao, Jiangui Liu, Yuhong He, and Jiali Shang. Recent advances of hyperspectral imaging
 603 technology and applications in agriculture. *Remote Sensing*, 12(16):2659, 2020.

604 Stefan Niewiadomski. *Filter handbook: a practical design guide*. Newnes, 2013.

605

606 Maxime Oquab, Timothée Darzet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez,
 607 Daniel Haziza, Francisco Massa, Alaeldin El-Nouby, et al. Dinov2: Learning robust visual features without
 608 supervision. *arXiv preprint arXiv:2304.07193*, 2023.

609 Shen-En Qian. Hyperspectral satellites, evolution, and development history. *IEEE Journal of Selected Topics in
 610 Applied Earth Observations and Remote Sensing*, 14:7032–7056, 2021.

611 Silvia Seidlitz, Jan Sellner, Jan Odenthal, Berkin Özdemir, Alexander Studier-Fischer, Samuel Knödler, Leonardo
 612 Ayala, Tim J Adler, Hannes G Kenngott, Minu Tizabi, et al. Robust deep learning-based semantic organ
 613 segmentation in hyperspectral images. *Medical Image Analysis*, 80:102488, 2022.

614 Qiu Shen, Yuxing Huang, Tianqi Ren, Ying Fu, and Shaodi You. Urban scene understanding via hyperspectral
 615 images: Dataset and benchmark. Available at SSRN 4560035.

616

617 Gencer Sumbul, Marcela Charfuelan, Begüm Demir, and Volker Markl. Bigearthnet: A large-scale benchmark
 618 archive for remote sensing image understanding. In *IGARSS 2019-2019 IEEE international geoscience and
 619 remote sensing symposium*, pages 5901–5904. IEEE, 2019.

620 Gencer Sumbul, Chang Xu, Emanuele Dalsasso, and Devis Tuia. Smarties: Spectrum-aware multi-sensor
 621 auto-encoder for remote sensing images. *arXiv preprint arXiv:2506.19585*, 2025.

622 Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal,
 623 Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency
 624 functions in low dimensional domains. *NeurIPS*, 33:7537–7547, 2020.

625 Nick Theisen, Robin Bartsch, Dietrich Paulus, and Peer Neubert. Hs3-bench: A benchmark and strong baseline
 626 for hyperspectral semantic segmentation in driving scenarios. In *2024 IEEE/RSJ International Conference on
 627 Intelligent Robots and Systems (IROS)*, pages 5895–5901. IEEE, 2024.

628 Prasad S Thenkabail, John G Lyon, and Alfredo Huete. Advances in hyperspectral remote sensing of vegetation
 629 and agricultural crops. In *Fundamentals, sensor systems, spectral libraries, and data mining for vegetation*,
 630 pages 3–37. CRC press, 2018.

631 Gabriel Tseng, Anthony Fuller, Marlena Reil, Henry Herzog, Patrick Beukema, Favyen Bastani, James R Green,
 632 Evan Shelhamer, Hannah Kerner, and David Rolnick. Galileo: Learning global & local features of many
 633 remote sensing modalities. *arXiv preprint arXiv:2502.09356*, 2025.

634 Leon Amadeus Varga, Martin Messmer, Nuri Benbarka, and Andreas Zell. Wavelength-aware 2d convolutions
 635 for hyperspectral imaging. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer
 636 Vision*, pages 3788–3797, 2023.

637 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser,
 638 and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach,
 639 R. Fergus, S. Vishwanathan, and R. Garnett, editors, *NeurIPS*, volume 30. Curran Associates, Inc., 2017.

640 Leonard Waldmann, Ando Shah, Yi Wang, Nils Lehmann, Adam Stewart, Zhitong Xiong, Xiao Xiang Zhu,
 641 Stefan Bauer, and John Chuang. Panopticon: Advancing any-sensor foundation models for earth observation.
 642 In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2204–2214, 2025.

643 Junjue Wang, Zhuo Zheng, Ailong Ma, Xiaoyan Lu, and Yanfei Zhong. Loveda: A remote sensing land-cover
 644 dataset for domain adaptive semantic segmentation. *arXiv preprint arXiv:2110.08733*, 2021.

645

646 Yi Wang, Zhitong Xiong, Chenying Liu, Adam J Stewart, Thomas Dujardin, Nikolaos Ioannis Bountos, Angelos
 647 Zavras, Franziska Gerken, Ioannis Papoutsis, Laura Leal-Taixé, et al. Towards a unified copernicus foundation
 model for earth vision. *arXiv preprint arXiv:2503.11849*, 2025.

648 Christian Winkens, Florian Sattler, Veronika Adams, and Dietrich Paulus. Hyko: A spectral dataset for scene
649 understanding. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages
650 254–261, 2017.

651 Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene
652 understanding. In *Proceedings of the European conference on computer vision (ECCV)*, pages 418–434,
653 2018.

654 Zhitong Xiong, Yi Wang, Fahong Zhang, Adam J Stewart, Joëlle Hanna, Damian Borth, Ioannis Papout-
655 sis, Bertrand Le Saux, Gustau Camps-Valls, and Xiao Xiang Zhu. Neural plasticity-inspired multimodal
656 foundation model for earth observation. *arXiv preprint arXiv:2403.15356*, 2024.

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 703 704 705 706 707 Appendix

708 A IMPLEMENTATION DETAILS

709 We begin by outlining general implementation details of CARL. In the experiments from Secs. 4.1, 4.2,
 710 we employed a small version of CARL, comprising overall 12 attention blocks with an embedding
 711 dimension of 384. For the satellite experiments in Sec. 4.3, we switched to the base version of
 712 CARL, keeping the same 12-block depth but increasing the embedding dimension of the spatial
 713 encoder to 768 to align with SpectralGPT⁺ and DOFA. The blocks were structured into $L = 4$ self-
 714 attention-cross-attention modules within the spectral encoder and 8 self-attention blocks within the
 715 spatial encoder. To incorporate wavelength information in the positional encoding, the wavelengths
 716 were given in nanometers, and scaled by a factor of $\alpha = 10^{-3}$, to obtain position coordinates
 717 approximately in the range of one. In accordance with the ablation studies presented in the main
 718 manuscript, the hyperparameter σ in the wavelength positional encoding was set to 3 and $K = 8$
 719 spectral representations were employed. The initial convolution in our model utilized a kernel size
 720 and stride of 8, resulting in patch dimensions of 8×8 . The spectral representations within the spectral
 721 encoder were implemented as learnable embeddings, initialized using a truncated normal distribution
 722 with a mean of zero and a standard deviation of 0.5. Furthermore, a 1D sinusoidal positional encoding
 723 scheme, based on discrete token positions, was implemented to represent the positions of $(S_j)_{j \leq K}$.
 724 Pseudo-code of a forward step of CARL can be found in algorithm 1.

725 A.1 IMPLEMENTATION DETAILS OF DOWNSTREAM TASKS

726 In the first and third experiment, CARL was integrated with the ViTAdapter (Chen et al., 2023) to
 727 generate hierarchical features, which were subsequently forwarded to the UperNet segmentation
 728 head (Xiao et al., 2018). The ViTAdapter applies lightweight convolutions to the input and facilitates
 729 information exchange with the spatial transformer via injector and extractor modules. To maintain
 730 invariance with respect to the channel dimensionality, a single channel of the given spectral image was
 731 used as input to these convolutions. The rationale behind this approach is to leverage the enhanced
 732 spatial resolution provided by the convolutions, as the encoded spectral information is contained
 733 within the camera-agnostic representation of the proposed model. As segmentation loss, an equally
 734 weighted sum of the cross entropy loss and the dice loss was employed. To enhance training stability,
 735 the attention blocks within the spectral encoder were initialized using Dinov2 weights (Oquab
 736 et al., 2023), which provide a robust checkpoint derived from self-supervised training on a large-
 737 scale dataset. The spatial encoder, EVA-02 (Fang et al., 2024), was initialized with self-supervised
 738 weights obtained through masked image modeling (He et al., 2022) on ImageNet (Deng et al., 2009).
 739 Furthermore, the model was optimized using the AdamW optimizer with an initial learning rate of
 740 10^{-4} . An exponential learning rate scheduler was employed to reduce the learning rate throughout
 741 the training process. In the second experiment, the training procedure followed that of the original
 742 Mask2Former (Cheng et al., 2022) training on Cityscapes (Cordts et al., 2016). In all experiments, a
 743 channel sampling strategy was employed to reduce the GPU memory consumption during training.
 744 Specifically, in instances where the channel dimension of the spectral image surpassed 32, a random
 745 subsampling of 32 channels was performed. Crucially, this did not affect performance, as the spectral
 746 tokens are ordered via wavelength positional encoding. As a result, all experiments were successfully
 747 executed on a single GPU endowed with a memory capacity of up to 40GB.

748 A.2 IMPLEMENTATION DETAILS OF CARL-SSL

749 The model subject to pre-training consisted of the proposed spectral encoder, in conjunction with
 750 the EVA-02 spatial encoder. As outlined in the main manuscript, the pre-training strategy employs
 751 spectral and spatial masking with reconstruction objectives in the feature space through predictor
 752 networks. Pseudo-code of CARL-SSL is outlined in algorithm 2. To manage compute resources, we
 753 downsampled each hyperspectral image to 64 channels. For spectral self-supervision on hyperspectral
 754 EnMAP data, we applied a single mask covering 15 % to 30 % of the channels; for Sentinel-2
 755 (multispectral) we masked two to three channels. Our spatial SSL strategy used two masks, each
 spanning 30 % to 50 % of the image area. The predictors are transformer architectures with a depth

756
757**Algorithm 1:** Pseudo-code of a forward pass of CARL.

```

758 # Input:
759 # x: batch of images of shape (B, C, H, W)
760 # w: wavelengths corresponding of shape (B, C)
761
762 def forward(self, x, w):
763     # ----- Projection -----
764     x = rearrange(x, "B C H W -> (B C) 1 H W")
765     x = self.projection(x) # shape: (BxC, D, h, w)
766     x = rearrange(x, "(B C) D h w -> (B h w) C D")
767     # ----- Spectral Positional Encoding -----
768     spec_pe = self.spec_pe(w) # shape: (B, C, D)
769     spec_pe = repeat(spec_pe, "B C D -> (B h w) C D")
770     x = x + spec_pe
771     spec_reps = repeat(self.spec_reps, "1 K D -> (B h w) K D")
772     spec_rep_pe = repeat(self.spec_rep_pe, "1 K D -> (B h w) K D")
773     spec_reps = spec_reps + spec_rep_pe
774     # ----- Spectral Attention -----
775     for self_attn, cross_attn in self.spectral_blk:
776         x = self_attn(x)
777         spec_reps = cross_attn(spec_reps, x, x)
778     # ----- Spectral-to-Spatial Transition -----
779     spec_rep = spec_reps.sum(dim=1) # shape: (B*h*w, D)
780     spec_rep = self.linear(self.norm(spec_rep))
781     spec_rep = rearrange(spec_rep, "(B h w) D -> B (h w) D")
782     # ----- Spatial Positional Encoding -----
783     spat_pe = self.spat_pe(h, w) # shape: (1, h*w, D)
784     spat_pe = repeat(spat_pe, "1 N D -> B N D")
785     spec_rep = spec_rep + spat_pe
786     # ----- Spatial Attention -----
787     for self_attn in self.spatial_blk:
788         spec_rep = self_attn(spec_rep)
789     spec_rep = self.norm(spec_rep) # shape: (B, h*w, D)
790     # ----- Output Head -----
791     out = self.head(spec_rep)
792     return out

```

776

777

778

of 3 and an embedding dimension of 384. Both, spectral and spatial self-supervision leverage the VICReg loss (Bardes et al., 2021), which is decomposed into invariance, variance, and covariance terms. Following the original recommendations, weights of 1 were assigned to the invariance and variance terms, and a weight of 0.05 was assigned to the covariance term. Let \hat{y} and y be the predicted and target tokens, both of shape (B, N, D) —with B as batch size, N sequence length, and D embedding dimension. We then compute:

784
$$\mathcal{L} = \mathcal{L}_{\text{spec}} + \mathcal{L}_{\text{spat}} \quad (2)$$

785
$$= \mathcal{L}_{\text{vicreg}}(\hat{y}_{\text{spec}}, y_{\text{spec}}) + \mathcal{L}_{\text{vicreg}}(\hat{y}_{\text{spat}}, y_{\text{spat}}) \quad (3)$$

786
$$\mathcal{L}_{\text{vicreg}}(\hat{y}, y) = \mathcal{L}_{\text{inv}}(\hat{y}, y) + \mathcal{L}_{\text{var}}(\hat{y}) + 0.05 \cdot \mathcal{L}_{\text{cov}}(\hat{y}) \quad (4)$$

787
$$\mathcal{L}_{\text{inv}}(\hat{y}, y) = \frac{1}{B \cdot N \cdot D} \sum_{b=1}^B \sum_{n=1}^N \sum_{d=1}^D (\hat{y}_{b,n,d} - y_{b,n,d})^2 \quad (5)$$

788
$$\mathcal{L}_{\text{var}}(\hat{y}) = \frac{1}{N \cdot D} \sum_{n=1}^N \sum_{d=1}^D \max(0, 1 - \sqrt{\text{Var}(\hat{y}_{n,d})}) \quad (6)$$

789
$$\mathcal{L}_{\text{cov}}(\hat{y}) = \frac{1}{N \cdot D} \sum_{n=1}^N \sum_{\substack{i,j \leq D \\ i \neq j}} \text{Cov}(\hat{y}_n)_{i,j}^2 \quad (7)$$

790
$$\text{Cov}(\hat{y}_n) = \frac{1}{B-1} \sum_{b=1}^B (\hat{y}_{b,n} - \bar{\hat{y}_n})(\hat{y}_{b,n} - \bar{\hat{y}_n})^T \text{ where } \bar{\hat{y}_n} = \frac{1}{B} \sum_{b=1}^B \hat{y}_{b,n} \quad (8)$$

801

While \mathcal{L}_{inv} encourages \hat{y} to resemble y , \mathcal{L}_{var} promotes diversity among generated tokens by increasing their variance across the samples, preventing feature collapse. Finally, the covariance term minimizes off-diagonal absolute values in the feature covariance matrix, encouraging independence between feature components. This reduces redundancy and increases information across the feature dimension.

802

803

A.3 CARL-SSL ON SATELLITE IMAGES

804

805

806

807

808

809

The pre-training dataset consisted of three distinct datasets, specifically HySpecNet-11k (Fuchs and Demir, 2023), SpectralEarth (Braham et al., 2024), and BigEarthNet (Sumbul et al., 2019), whose detailed composition is summarized in Tab. 7. The student and predictor networks were optimized

810
 811 Table 7: **Composition of our remote sensing SSL-data.** Three datasets were used for self-supervised
 812 pre-training on satellite images. The table reports the number of images, sensor, number of spectral
 813 channels, and covered wavelength range. ¹Only a subset of the full SpectralEarth dataset was used in
 814 our pre-training.

Dataset	# Images	Sensor	# Channels	Wavelength Range
HySpecNet-11k	11,483	EnMap	202	418 nm to 2,445 nm
SpectralEarth	247,030 ¹	EnMap	202	418 nm to 2,445 nm
BigEarthNet	549,488	Sentinel-2	12	443 nm to 2,202 nm

815
 816 **Algorithm 2:** Pseudo-code of a training step of CARL-SSL.

```

817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2178
2179
2180
2181
2182
2183
2184
2185
2186
2186
2187
2188
2189
2189
2190
2191
2192
2193
2194
2195
2195
2196
2197
2198
2199
2199
2200
2201
2202
2203
2204
2204
2205
2206
2207
2208
2209
2209
2210
2211
2212
2213
2214
2215
2215
2216
2217
2218
2219
2219
2220
2221
2222
2223
2224
2225
2225
2226
2227
2228
2229
2229
2230
2231
2232
2233
2234
2234
2235
2236
2237
2238
2238
2239
2240
2241
2242
2243
2244
2244
2245
2246
2247
2248
2248
2249
2250
2251
2252
2252
2253
2254
2255
2256
2256
2257
2258
2259
2259
2260
2261
2262
2263
2263
2264
2265
2266
2266
2267
2268
2269
2269
2270
2271
2272
2273
2273
2274
2275
2276
2276
2277
2278
2279
2279
2280
2281
2282
2283
2283
2284
2285
2286
2286
2287
2288
2289
2289
2290
2291
2292
2293
2293
2294
2295
2296
2296
2297
2298
2299
2299
2300
2301
2302
2303
2303
2304
2305
2306
2306
2307
2308
2309
2309
2310
2311
2312
2313
2313
2314
2315
2316
2316
2317
2318
2319
2319
2320
2321
2322
2323
2323
2324
2325
2326
2326
2327
2328
2329
2329
2330
2331
2332
2333
2333
2334
2335
2336
2336
2337
2338
2339
2339
2340
2341
2342
2343
2343
2344
2345
2346
2346
2347
2348
2349
2349
2350
2351
2352
2353
2353
2354
2355
2356
2356
2357
2358
2359
2359
2360
2361
2362
2363
2363
2364
2365
2366
2366
2367
2368
2369
2369
2370
2371
2372
2373
2373
2374
2375
2376
2376
2377
2378
2379
2379
2380
2381
2382
2383
2383
2384
2385
2386
2386
2387
2388
2389
2389
2390
2391
2392
2393
2393
2394
2395
2396
2396
2397
2398
2399
2399
2400
2401
2402
2403
2403
2404
2405
2406
2406
2407
2408
2409
2409
2410
2411
2412
2413
2413
2414
2415
2416
2416
2417
2418
2419
2419
2420
2421
2422
2423
2423
2424
2425
2426
2426
2427
2428
2429
2429
2430
2431
2432
2433
2433
2434
2435
2436
2436
2437
2438
2439
2439
2440
2441
2442
2443
2443
2444
2445
2446
2446
2447
2448
2449
2449
2450
2451
2452
2453
2453
2454
2455
2456
2456
2457
2458
2459
2459
2460
2461
2462
2463
2463
2464
2465
2466
2466
2467
2468
2469
2469
2470
2471
2472
2473
2473
2474
2475
2476
2476
2477
2478
2479
2479
2480
2481
2482
2483
2483
2484
2485
2486
2486
2487
2488
2489
2489
2490
2491
2492
2493
2493
2494
2495
2496
2496
2497
2498
2499
2499
2500
2501
2502
2503
2503
2504
2505
2506
2506
2507
2508
2509
2509
2510
2511
2512
2513
2513
2514
2515
2516
2516
2517
2518
2519
2519
2520
2521
2522
2523
2523
2524
2525
2526
2526
2527
2528
2529
2529
2530
2531
2532
2533
2533
2534
2535
2536
2536
2537
2538
2539
2539
2540
2541
2542
2543
2543
2544
2545
2546
2546
2547
2548
2549
2549
2550
2551
2552
2553
2553
2554
2555
2556
2556
2557
2558
2559
2559
2560
2561
2562
2563
2563
2564
2565
2566
2566
2567
2568
2569
2569
2570
2571
2572
2573
2573
2574
2575
2576
2576
2577
2578
2579
2579
2580
2581
2582
2583
2583
2584
2585
2586
2586
2587
2588
2589
2589
2590
2591
2592
2593
2593
2594
2595
2596
2596
2597
2598
2599
2599
2600
2601
2602
2603
2603
2604
2605
2606
2606
2607
2608
2609
2609
2610
2611
2612
2613
2613
2614
2615
2616
2616
2617
2618
2619
2619
2620
2621
2622
2623
2623
2624
2625
2626
2626
2627
2628
2629
2629
2630
2631
2632
2633
2633
2634
2635
2636
2636
2637
2638
2639
2639
2640
2641
2642
2643
2643
2644
2645
2646
2646
2647
2648
2649
2649
2650
2651
2652
2653
2653
2654
2655
2656
2656
2657
2658
2659
2659
2660
2661
2662
2663
2663
2664
2665
2666
2666
2667
2668
2669
2669
2670
2671
2672
2673
2673
2674
2675
2676
2676
2677
2678
2679
2679
2680
2681
2682
2683
2683
2684
2685
2686
2686
2687
2688
2689
2689
2690
2691
2692
2693
2693
2694
2695
2696
2696
2697
2698
2699
2699
2700
2701
2702
2703
2703
2704
2705
2706
2706
2707
2708
2709
2709
2710
2711
2712
2713
2713
2714
2715
2716
2716
2717
2718
2719
2719
2720
2721
2722
2723
2723
2724
2725
2726
2726
2727
2728
2729
2729
2730
2731
2732
2733
2733
2734
2735
2736
2736
2737
2738
2739
2739
2740
2741
2742
2743
2743
2744
2745
2746
2746
2747
2748
2749
2749
2750
2751
2752
2753
2753
2754
2755
2756
2756
2757
2758
2759
2759
2760
2761
2762
2763
2763
2764
2765
2766
2766
2767
2768
2769
2769
2770
2771
2772
2773
2773
2774
2775
2776
2776
2777
2778
2779
2779
2780
2781
2782
2783
2783
2784
2785
2786
2786
2787
2788
2789
2789
2790
2791
2792
2793
2793
2794
2795
2796
2796
2797
2798
2799
279
```

864
 865 Table 8: **Details of benchmark datasets for remote sensing experiments.** The first block lists five
 866 in-distribution Sentinel-2 datasets that are present in CARL’s pretraining set. The second block lists
 867 datasets captured by sensors that were unseen during pretraining, suitable for evaluating cross-sensor
 868 generalization.

869 Dataset	869 Sensor / Platform	869 Channels	869 Wavelength (nm)	869 Classes	869 Task	869 Zero-shot
<i>870 In-distribution sensors (present in pretraining)</i>						
871 SegMunich	871 Sentinel-2	871 10	871 442–2202	871 –	871 Seg	871 ✗
872 m-bigearthnet	872 Sentinel-2	872 12	872 442–2202	872 –	872 Cls	872 ✗
873 m-eurosat	873 Sentinel-2	873 13	873 442–2202	873 –	873 Cls	873 ✗
874 m-cashew	874 Sentinel-2	874 13	874 442–2202	874 –	874 Seg	874 ✗
875 m-SA-crop-type	875 Sentinel-2	875 13	875 442–2202	875 –	875 Seg	875 ✗
<i>876 Out-of-distribution sensors (unseen during pretraining)</i>						
877 LoveDA Urban	877 Google Earth	877 3	877 RGB	877 –	877 Seg	877 ✓
878 LoveDA Rural	878 Google Earth	878 3	878 RGB	878 –	878 Seg	878 ✓
879 m-forestnet	879 LandSat-8	879 6	879 482–2200	879 –	879 Cls	879 ✓
880 WHU-OHS	880 Orbita hyperspectral sat.	880 32	880 466–940	880 –	880 Seg	880 ✓
881 Wuhan	881 Gaofen-5	881 116	881 420–2400	881 –	881 Seg	881 ✓
882 Beijing	882 Gaofen-5	882 116	882 420–2400	882 –	882 Seg	882 ✓

883
 884 or their variants. These architectures are designed for RGB inputs and therefore do not natively
 885 handle sensors with arbitrary numbers of spectral channels. Early adaptations for spectral imaging
 886 add multiple projection layers to accommodate different channel counts (Tseng et al., 2025; Astruc
 887 et al., 2025), or replace the single projection layer with complete modality-specific encoders (Jakubik
 888 et al., 2025; Fuller et al., 2023; Astruc et al., 2024). Such designs—often implemented as early-
 889 or mid-fusion models—work well for the sensors seen during training but cannot generalize to
 890 unseen sensors with different spectral dimensions. To address this limitation, **channel-invariant**
 891 approaches have been proposed. These include projection-weight interpolation (Sumbul et al., 2025;
 892 Varga et al., 2023), and channel-adaptive projection layers (Xiong et al., 2024; Wang et al., 2025;
 893 Li et al., 2025b). Several methods also exploit known wavelength information (for example via
 894 wavelength positional encodings) to establish cross-sensor channel relationships; we refer to such
 895 methods as **wavelength-aware** (Xiong et al., 2024; Wang et al., 2025; Li et al., 2025b; Sumbul et al.,
 896 2025; Varga et al., 2023). When a model is both channel-invariant and wavelength-aware, we call
 897 it **camera-agnostic**, since it can in principle generalize to any spectral sensor, whether seen during
 898 training or not.

899 Most camera-agnostic designs handle spectral inputs only at the projection stage and then perform
 900 purely spatial operations in feature space to learn spatial structure Xiong et al. (2024); Wang et al.
 901 (2025); Sumbul et al. (2025); Waldmann et al. (2025); Jakubik et al. (2025); Fuller et al. (2023);
 902 Tseng et al. (2025); critically, they do not learn spectral relationships within that feature space. We
 903 categorize them into **spatial encoding schemes** By contrast, **spatio-spectral encoding schemes**
 904 learn joint spatial *and* spectral relations in the feature space. This capability can be essential in
 905 spectrally heterogeneous settings where the model must align spectral signatures across different
 906 sensors. Examples include spatio-spectral patching with self-attention (Hong et al., 2024) or dedicated
 907 spectral encoders (Braham et al., 2024). However, existing spatio-spectral encoding schemes are not
 908 camera-agnostic.

909 To address this gap, we propose CARL—to the best of our knowledge, the first approach to unify
 910 explicit spatio-spectral feature encoding with a camera-agnostic design. CARL achieves this by
 911 introducing a channel-invariant, wavelength-aware spectral encoder that compresses variable-length
 912 channel inputs into fixed-length spectral representations, enabling the model to learn rich joint
 913 spatio-spectral features and to generalize strongly out-of-distribution without per-camera retraining.

914 **Self-supervised learning strategies for spectral imaging** Self-supervised learning (SSL) is in-
 915 creasingly important as unlabeled data and compute scale up. Most current spectral-imaging pipelines
 916 inherit spatial encoding schemes from RGB models and therefore apply SSL that only learns spatial
 917 relations. We call that **spatial self-supervision**. These methods typically adapt RGB SSL recipes
 (e.g., MAE, DINOV2, iJEP) to remote-sensing foundation models (He et al., 2022; Oquab et al.,

918
 919
 920
 921
 922
 923 Table 9: **CARL learns robust representations during pre-training.** Linear-probing mIoU results
 924 on three evaluation datasets. Two of the datasets are acquired by sensors unseen during pre-training
 925 (Beijing, LoveDA). CAR shows competitive in-distribution performance and superior generalization
 926 to out-of-distribution sensors, achieving the best average rank across all 11 benchmark datasets.
 927
 928
 929
 930

Sensor	SegMunich (mIoU)	Beijing (mIoU)	LoveDA Rural (mIoU)	Avg. Rank over all 11 datasets
	Sentinel-2 (10 bands)	Gaofen-5 (116 bands)	RGB (3 bands)	
SpectralGPT	27.9	-	-	5.5
Galileo	35.3	-	-	5.5
Croma	-	-	-	5.0
DOFA	38.2	17.4	14.9	3.2
Copernicus-FM	38.4	14.5	17.6	2.6
SMARTIES	39.1	17.1	17.3	2.6
CARL	38.9	19.1	29.8	1.6

931
 932 2023; Assran et al., 2023; Fuller et al., 2023; Jakubik et al., 2025; Xiong et al., 2024; Waldmann et al.,
 933 2025; Tseng et al., 2025). In contrast, **spatio-spectral self-supervision** learns spatio-spectral relations
 934 during pretraining. For example, SpectralGPT performs reconstruction of 3D spatio-spectral patches
 935 (Hong et al., 2024). Scaling pretraining requires camera-agnostic pretraining, but this is not always
 936 feasible: some approaches require decoder heads tied to specific channel counts, which prevents
 937 straightforward cross-sensor pretraining (Sumbul et al., 2025). Moreover, SpectralGPT as backbone
 938 is not camera-agnostic. Additionally, pixel-reconstruction objectives are particularly sensitive in
 939 spectral imagery because pixel noise from atmospheric effects, illumination variation, and sensor
 940 distortions are stronger than in RGB data. For this reason, feature-based SSL methods (e.g., I-JEPA
 941 (Assran et al., 2023), DINOv2 (Caron et al., 2021)) — which learn robust latent representations rather
 942 than raw pixel reconstructions — may be better suited to spectral data

943 To the best of our knowledge, there is currently no SSL strategy that is both camera-agnostic
 944 *and* spatio-spectral, nor a strategy that is both feature-based *and* spatio-spectral. In this work
 945 we introduce CARL-SSL: a feature-based, camera-agnostic spatio-spectral self-supervision frame-
 946 work. CARL-SSL enables scalable pretraining across heterogeneous sensors and yields robust joint
 947 spatio-spectral representations.

949 C ADDITIONAL REMOTE SENSING EXPERIMENTS

950
 951 In addition to Tabs. 4, 5, we provide further results on our linear probing evaluation in Tab. 9.
 952 In particular, we evaluated segmentation performance on the three land cover datasets. While
 953 SegMunich is acquired by the Sentinel-2 sensor, and is therefore in-distribution, Beijing and LoveDA
 954 were acquired by unseen sensors (hyperspectral and RGB), and therefore emphasizes on cross-
 955 generalizability. Further details to the datasets can be seen in Tab. 8. While every model except
 956 of Croma can be applied on the SegMunich images that have 10 channels, only three baseline
 957 methods can process hyperspectral images. Particularly in those scenarios, CARL’s outperformance is
 958 pronounced. Furthermore, we report the average rank over all eleven datasets of all baseline methods.
 959 CARL ranks best with an average rank of 1.6.

960 As the SegMunich is sufficiently large for full fine-tuning, we performed supervised fine-tuning. Due
 961 to the required compute, we limit the baseline methods to DOFA, and SpectralGPT⁺. As can be seen
 962 in Tab. 10, CARL keeps its outperformance and achieves a score of 50.9.

964 D COMPUTATIONAL COMPLEXITY

965
 966 Encoding high-dimensional spatio-spectral data imposes computational and memory demands;
 967 therefore, controlling excessive floating-point operations (FLOPs) is essential. For the models
 968 considered in Sec. 4.3, we report parameter counts and FLOP estimates for the input sizes used in our
 969 remote-sensing evaluation. All experiments use the base parameter size of each model to provide a
 970 fair comparison across architectures. Two architectures, Galileo and Croma, were excluded from this
 971 complexity study because they were specifically designed for Sentinel-2 and cannot be reasonably
 972 adapted to other sensors without substantial redesign. By contrast, SpectralGPT depends on the

972 number of channels only through a single linear layer, making it straightforward to adapt to different
 973 sensor channel counts; accordingly, we include it in our study. The remaining models are sensor-
 974 agnostic and can be applied for the different input size out of the box. To make comparisons consistent,
 975 we fixed a patch grid of 16×16 and evaluated four channel configurations corresponding to RGB,
 976 Sentinel-2, OHS, and Gaofen-5 sensors (3, 12, 32, and 116 channels, respectively, see Tab. 8). We
 977 report FLOPs for multiple input sizes because the models scale differently with spectral dimensionality.
 978 Finally, we analyze how computational complexity relates to model design choices—such as encoding
 979 schemes—and to average model performance, to highlight trade-offs between efficiency and accuracy

980 In Tab. 11 we summarize the measured costs. Overall, purely spatial encoding schemes are
 981 computationally cheaper than spatio-spectral approaches; the differences in scaling behavior explain
 982 most of this gap. SpectralGPT uses 3D patching and full self-attention across both spatial and
 983 spectral dimensions for twelve blocks. Consequently, the complexity of each attention block grows
 984 as $\mathcal{O}((H \cdot W \cdot C)^2)$.

985 CARL takes a different design that yields substantially more favorable scaling by exploiting two
 986 principles: (1) disentangled spectral and spatial encoding, and (2) cross-attention with K learned
 987 spectral tokens. Concretely, CARL first extracts rich features along the channel axis, aggregates
 988 them, and then applies spatial-only transformer blocks. This produces spectral encoding stages whose
 989 attention cost scales as $\mathcal{O}(C^2)$ (four blocks), and spatial-only stages whose complexity scales as
 990 $\mathcal{O}((H \cdot W)^2)$ (eight blocks). Beyond lower FLOPs, disentanglement also simplifies spectral- and
 991 spatial-specific design choices (e.g., wavelength positional encoding) while preserving rich spatio-
 992 spectral representations. To further reduce costs, CARL interleaves a small number of full spectral
 993 self-attention operations with cheaper cross-attention blocks inside the spectral encoder. While
 994 spectral self-attention scales as $\mathcal{O}(C^2)$, a cross-attention over K learned spectral representations
 995 scales as $\mathcal{O}(C \cdot K)$. Since K is fixed and small in our experiments (we use $K = 8$), these cross-
 996 attention blocks are much less expensive than full spectral self-attention when C is large (as in
 997 hyperspectral imagery). Although CARL remains more expensive than purely spatial encoders,
 998 the additional cost is offset by substantially richer features and markedly better out-of-distribution
 999 generalization. This trade-off is reflected in CARL’s strong empirical performance (average rank 1.6
 in our evaluation; see Tab. 11).

1000 Because CARL is camera-agnostic, we can exploit this property to reduce training cost using a
 1001 simple channel-subsampling strategy. For the hyperspectral organ dataset (originally 100 channels),
 1002 we randomly sampled 16 channels at each training step and optimized the model on this reduced
 1003 input. This approach reduced training FLOPs by roughly $\sim 75\%$ in our experiments while preserving
 1004 validation performance: we obtained a mIoU of 68.8, compared to the original 69.1.

1005 Two points explain why subsampling works well here. First, CARL’s spectral encoder explicitly learns
 1006 relationships across channels and compresses spectral information into a fixed set of representations.
 1007 Therefore, the model can integrate information across different sampled subsets and still recover rich
 1008 spectral features. Second, random channel sampling acts as a form of stochastic regularization: by
 1009 seeing many different channel subsets during training, the model becomes more robust to missing
 1010 or shifted spectral bands and generalizes better to out-of-distribution sensors. Practically, channel
 1011 subsampling is an easy-to-implement, low-overhead augmentation.

1017 E URBAN SCENE SEGMENTATION

1021 The detailed class-wise IoU scores on the HSICity test set are depicted in Tab. 12. CARL-SSL
 1022 exhibited superior performance benchmarked against camera-specific and channel-invariant spectral
 1023 imaging models. As the HSICity training set does not contain any “pole” annotation, the camera-
 1024 specific model exhibits a “pole” IoU score of 0. Notably, CARL and CARL-SSL achieved the
 1025 best IoU scores for the “pole” class, indicating superior capability of translating RGB labels from
 Cityscapes to hyperspectral imagery.

1026 Table 10: **Fully fine-tuned CARL yields superior downstream performance.** Class-wise IoU
 1027 scores with 95% confidence intervals on the SegMunich multispectral land cover dataset were
 1028 benchmarked against strong SSL-pretrained models.

Method	Arable land	Perm. Crops	Pastures	Forests	Water	Shrub	Open spaces	Wetlands	Mine, dump	Agr. vsg.	Urban fabric	Buildings	mIoU
SatMAE	72.3	15.8	49.7	81.9	74.0	13.6	27.5	41.2	37.9	18.3	65.3	50.6	45.7
SMARTIES	71.6	19.6	49.4	85.2	72.9	17.3	44.8	39.3	37.3	22.2	65.6	53.2	48.20 [47.6; 48.7]
SpectralGPT ⁺	72.2	21.9	51.0	86.3	76.5	17.2	44.0	39.0	38.2	22.9	66.8	53.3	49.1 [48.6; 49.6]
DOFA	72.0	21.5	50.7	86.1	75.3	18.2	45.4	40.1	39.1	23.1	67.3	54.9	49.5 [48.9; 49.9]
CARL	72.9	24.9	51.9	86.6	76.5	21.0	43.9	42.0	42.0	25.6	68.6	54.8	50.9 [50.4; 51.4]

1040 Table 11: **Computational complexity of compared models.** For each sensor in our remote sensing
 1041 evaluation, we report the number of parameters and estimated GFLOPs for a subset of models.
 1042 “Avg. Rank” denotes the mean performance rank (lower is better). CARL offers a more favorable
 1043 compute–performance balance than the spatio-spectral baseline SpectralGPT, while spatial-only
 1044 encodings are cheaper but generalize poorly to unseen sensors.

Model	Params (M)	GFLOPs (per sensor)				Spatio-spectral	Avg. Rank
		RGB	Sentinel-2	OHS	Gaofen-5		
SpectralGPT	85.4	23	107	339	2573	✓	5.5
DOFA	111.1	23	24	26	32	✗	3.2
Copernicus-FM	139.3	23	24	26	32	✗	2.6
SMARTIES	88.4	23	24	26	55	✗	2.6
CARL	71.5	34	53	96	286	✓	1.6

F SYNTHETIC MULTISPECTRAL DATA GENERATION

1058 As outlined in Section 4.1 of the main manuscript, we synthesized multispectral images from given
 1059 hyperspectral images to simulate spectral heterogeneity within the training set. This section provides
 1060 a more in-depth explanation of the data generation procedure.

1061 In spectral imaging, optical filters play a crucial role in isolating specific wavelength bands, allowing
 1062 the system to capture reflectance information with spectral specificity, as discussed in (Garini et al.,
 1063 2006). These filters are designed to selectively transmit light within defined wavelength ranges,
 1064 determined by their material properties and design specifications. The transmission characteristics of
 1065 an optical filter are typically described by its filter function, which quantifies the transmitted intensity
 1066 as a function of wavelength. In practice, these filter functions often exhibit smooth, bell-shaped
 1067 curves centered around a target wavelength (Niewiadomski, 2013).

1068 To simulate this behavior in a proof-of-concept setting, we modeled the filter functions as normalized
 1069 Gaussian distributions, where the mean corresponds to the center wavelength and the variance controls
 1070 the bandwidth of the filter. This approach enables the generation of virtual multispectral cameras
 1071 with tunable spectral profiles. Using these Gaussian-modeled filters, we synthesized multispectral
 1072 images from hyperspectral data, while preserving the spatial context of the scene.

1073 Specifically, we simulated a multispectral channel by first sampling the corresponding filter’s center
 1074 wavelength μ through farthest point sampling within [550 nm, 950 nm]. To define a realistic range
 1075 for the filter bandwidth, we analyzed a real near-infrared multispectral camera (Ximea® MQ022HG-
 1076 IM-SM5X5 NIR), which features 25 spectral channels. Particularly, we fitted Gaussian curves to the
 1077 camera’s filter functions to estimate plausible variance values. Based on this analysis, the variance of
 1078 each simulated Gaussian filter was then uniformly sampled from the interval [5, 25]. As the given
 1079 hyperspectral images exhibited wavelengths from 500 nm to 1,000 nm with 5 nm steps, we discretized
 the wavelength axis accordingly and set $\lambda = (500, 505, \dots, 995)^T \in \mathbb{R}^{100}$. Then, we defined the

1080
 1081 **Table 12: The proposed spectral encoder demonstrates superior performance as a camera-**
 1082 **agnostic model.** The class-wise IoU scores with the 95 % confidence intervals of the mIoU scores on
 1083 the HSICity test set. While the camera-specific model was pre-trained on Cityscapes and fine-tuned
 1084 exclusively on HSICity, the other models are channel-invariant adaptations which were concurrently
 1085 trained on both datasets. Notably, our spectral encoder performs best among the presented adaptation
 1086 methods, and significantly benefits from self-supervised pre-training.

	Camera-specific model	Spectral Adapter	HyperFree	Hyve	DOFA	CARL	CARL-SSL
Road	93.4	93.6	93.7	94.0	94.1	94.7	95.0
Sidewalk	32.8	33.5	38.3	44.3	47.6	43.5	47.8
Building	69.8	54.9	65.0	69.4	71.9	71.1	71.1
Wall	55.1	43.9	54.7	54.4	52.4	55.4	55.2
Fence	14.1	5.3	11.6	11.5	10.4	11.3	13.9
Pole	0.0	29.6	15.1	20.8	30.9	31.0	31.8
Traffic light	51.0	50.4	47.8	53.0	58.8	55.5	57.2
Traffic sign	53.4	49.4	49.1	54.6	59.9	59.1	61.5
Vegetation	80.9	72.5	79.5	80.9	82.3	82.0	81.9
Terrain	3.2	9.0	5.7	3.6	3.9	6.6	5.1
Sky	85.8	79.9	83.7	87.2	88.5	88.4	88.7
Person	30.9	28.0	25.9	36.0	31.1	29.6	31.9
Rider	34.0	34.3	35.8	44.5	37.6	32.2	37.8
Car	86.0	88.1	86.2	88.0	89.5	89.3	90.3
Truck	53.7	50.8	57.0	60.6	73.6	57.5	63.0
Bus	67.6	80.0	79.3	80.3	83.0	87.6	87.4
Train	-	-	-	-	-	-	-
Motorcycle	0.0	0.4	0.0	0.0	0.1	0.0	0.0
Bicycle	36.0	11.3	19.2	29.1	26.5	29.1	33.1
mIoU	44.6 [40.9; 47.3]	43.4 [41.0; 45.2]	44.6 [42.2; 46.5]	48.0 [45.4; 50.0]	49.6 [46.8; 51.6]	48.6 [45.6; 51.0]	50.1 [47.2; 52.4]

1101
 1102 **Table 13: Semantic content drives CARL’s feature variance.** Proportion of variance in the learned
 1103 feature embeddings explained by semantic content (organ class) versus imaging sensor, as measured
 1104 by mean R^2 across all embedding dimensions.

Feature Variance Source	Explained Variance
Confounding Variable (Sensor)	0.6 %
Semantic Content (Organ)	61.6 %

1109
 1110 L1-normalized filter function $\tilde{F}_{\mu, \sigma}$ as following:

$$F_{\mu, \sigma}(\lambda_i) = e^{-\frac{(\lambda_i - \mu)^2}{2\sigma^2}} \quad (9)$$

$$\tilde{F}_{\mu, \sigma}(\lambda_i) = \frac{F_{\mu, \sigma}(\lambda_i)}{\sum_{j=1}^{100} |F_{\mu, \sigma}(\lambda_j)|} \quad (10)$$

1117 By uniformly sampling the number of channels, C , within [10, 25], we obtained C channel-specific
 1118 filter functions, $(\tilde{F}_{i, \mu_i, \sigma_i})_{i \leq C}$, as described above. These functions can be collectively represented in
 1119 matrix form as follows:

$$\tilde{F} = \begin{bmatrix} \tilde{F}_{1, \mu_1, \sigma_1}(\lambda)^T \\ \vdots \\ \tilde{F}_{C, \mu_C, \sigma_C}(\lambda)^T \end{bmatrix} \in \mathbb{R}^{C \times 100} \quad (11)$$

1124 Finally, for a hyperspectral pixel $P_{HSI} \in \mathbb{R}^{100}$, we simulated the corresponding multispectral pixel
 1125 by:

$$P_{MSI} = \tilde{F} \cdot P_{HSI} \in \mathbb{R}^C \quad (12)$$

1128 This matrix-vector multiplication can be performed for each pixel, leading to an multispectral image
 1129 with C channels. Notably, we only altered spectral properties of the images, while preserving
 1130 geometric information.

1131 In this way, six camera configurations were simulated, resulting in six sets of synthetic multispectral
 1132 images. To systematically introduce spectral heterogeneity into the original hyperspectral training
 1133 data, a progressive substitution strategy with the synthetic multispectral images was employed.
 In each iteration, hyperspectral data from two additional porcine subjects was substituted with

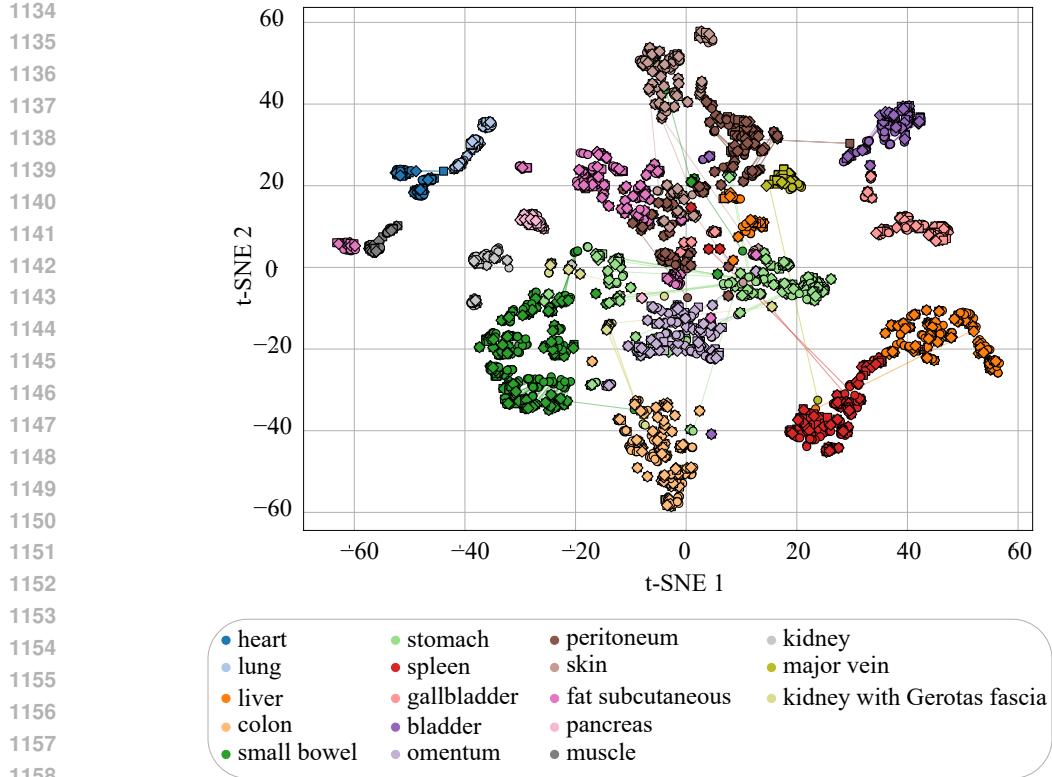


Figure 6: **CARL effectively disentangles organ semantics from camera variability.** t-SNE projection of mean feature embeddings for each organ region across the original hyperspectral image and two simulated multispectral variants. Points are color-coded by organ class and shaped by camera type (one hyperspectral, two unseen multispectral). Embeddings from the same region are connected into triangles. The dominant clustering by color and the scarcity of modality-connecting triangles demonstrate that the learned features are strongly organ-specific while remaining largely invariant to camera variations.

the corresponding synthetic multispectral images from a different simulated camera configuration. This process produced six augmented datasets that exhibit increasing spectral heterogeneity while preserving the surgical scene content. Model generalization to hyperspectral imagery was evaluated on the original hyperspectral test set, encompassing 166 images from five different porcine subjects.

G ANALYSIS ON FEATURE REPRESENTATIONS

In Sec. 4.1, we generated spectrally heterogeneous training datasets while keeping the hyperspectral evaluation set unchanged to avoid any modifications during testing. We now extend the evaluation by adding two sets of synthetic multispectral images, generated via simulated filter responses as described in Sec. 4.1. For each hyperspectral sample in the evaluation set, we produced two corresponding multispectral versions. After training our model on a dataset variant from Sec. 4.1, we performed inference on this enlarged, spectrally diverse evaluation set. Importantly, the multispectral cameras from the evaluation set are unseen during training. For each labeled organ region, we computed the mean feature vector—obtained from the spatial encoder—over its ground-truth mask. This yielded three embeddings per region: one from the original hyperspectral image and two from the simulated multispectral variants. To visualize these embeddings, we applied t-SNE to project them into a two-dimensional space (Fig. 6). Embeddings corresponding to the same region but different modalities were connected to form a triangle. Marker shapes indicate camera type, while colors denote organ class. The strong color-based grouping, rather than marker-based grouping, demonstrates that our feature representations are organ-specific and largely camera-agnostic. This is further supported by the scarcity of visible triangles.

1188 To quantify this observation, we performed a variance decomposition analysis using linear regression.
1189 Specifically, the feature embeddings were independently regressed against one-hot encoded predictors
1190 of either organ class or imaging sensor, and the coefficient of determination R^2 was computed for
1191 each regression. The average R^2 across all feature dimensions was then taken as the proportion
1192 of variance explained by the corresponding factor. As reported in Table 13, organ class explains
1193 61.6 % of the feature variance, while the imaging sensor accounts for only 0.6 %. This indicates
1194 that the representations are strongly driven by semantic content while remaining largely invariant to
1195 the confounding sensor domain, suggesting robust disentanglement of task-relevant features from
1196 acquisition-specific artifacts.
1197

H LLM USAGE STATEMENT

1200 Large Language Models (LLMs) were used exclusively to assist with paper writing. Specifically,
1201 they were employed to correct grammar errors and enhance the clarity and style of existing text.
1202 Importantly, LLMs were not used to generate section drafts or even write entire sections. Their role
1203 was limited to refining and improving written material.
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241