

# Symmetry-Aware Contrastive Learning for Self-Supervised Crop Intelligence from Satellite and Ecological Data

Mahule Roy<sup>1</sup>

Subhas Roy<sup>2</sup>

<sup>1</sup>University of Oxford

<sup>2</sup>TATA Consumer Products Limited

roymahule26@gmail.com

## Abstract

Agricultural and forestry systems represent complex dynamical environments characterized by strong spatial and temporal structure governed by fundamental ecological principles such as nutrient cycling, biomass accumulation, and energy conservation. While these systems generate vast quantities of unlabeled data through remote sensing platforms and field sensors, annotated datasets for supervised learning remain scarce and expensive to acquire. We introduce **Symmetry-Aware Contrastive Learning (SACL)**, a novel self-supervised framework that explicitly embeds ecological symmetries and conservation laws into representation learning objectives for agricultural applications. SACL enforces invariance to ecologically meaningful transformations—including spatial translation, temporal progression, and biomass/nutrient conservation—while preserving discriminative features relevant to crop growth and stress dynamics. Comprehensive evaluation across three distinct domains demonstrates SACL’s effectiveness: (i) satellite-derived vegetation indices (Sentinel-2 NDVI integrated with CropHarvest), (ii) synthetic ecological models (Logistic growth and Lotka–Volterra dynamics), and (iii) real-world soil nutrient and crop yield datasets (FAO EarthStat, USDA). Results indicate that SACL learns interpretable latent representations aligned with ecological processes, improves downstream crop yield and stress prediction accuracy with up to 70% reduction in labeled data requirements, and consistently outperforms standard contrastive learning and autoencoder baselines. This work establishes a new paradigm for AI-driven agricultural intelligence grounded in scientific domain knowledge.

## Introduction

Precision agriculture requires AI systems that can extract insights from complex ecological dynamics in large-scale observational data. While remote sensing platforms provide abundant unlabeled measurements of vegetation health and soil conditions, labeled data for critical variables like crop yield and nutrient status remains scarce and costly. Self-supervised learning (SSL) offers promise but standard methods employ arbitrary augmentations that disrupt ecological patterns. Agricultural systems exhibit fundamental ecological symmetries and conservation principles—spatial invariances in field plots, temporal patterns in phenological stages,

and mass/energy conservation in nutrient fluxes—that provide powerful inductive biases. We propose **Symmetry-Aware Contrastive Learning (SACL)**, a framework that integrates these domain-specific symmetries into contrastive learning through ecological augmentations and conservation regularization. Our contributions include: introducing the first SSL framework explicitly embedding ecological symmetries; providing theoretical sample complexity analysis; defining domain-specific augmentation operators; and demonstrating through comprehensive evaluation that SACL achieves superior performance, interpretability, and cross-domain generalization, particularly in low-label regimes.

## Related Work

Contrastive methods like SimCLR (1), BYOL (2), and MoCo (3) learn representations by maximizing agreement between augmented views. While adapted to remote sensing (4; 5), these methods use generic augmentations that ignore ecological constraints, limiting their agricultural applicability. Physics-informed neural networks (6) and Hamiltonian/Lagrangian networks (7; 8) incorporate physical laws but require labeled data or known equations. Ecological modeling efforts (9; 10; 11) integrate domain knowledge but face similar supervision requirements. Recent physics-aware SSL methods (12; 13) focus on physical sciences with defined equations, not ecological systems. Group-equivariant networks (14; 15) encode symmetries but operate in supervised settings. SACL bridges these gaps by integrating ecological symmetries into self-supervised learning for agricultural applications.

## Methodology

### Problem Formulation

Let  $\mathcal{X}$  denote the space of input observations, which may include satellite image patches, time-series of vegetation indices, or multivariate field measurements. Given an unlabeled dataset  $\mathcal{D} = \{x_i\}_{i=1}^N$  sampled from  $\mathcal{X}$ , our objective is to learn an encoder  $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^d$  that maps inputs to latent representations capturing ecologically meaningful features while respecting domain-specific symmetries and conservation laws.

## Theoretical Foundations

The sample complexity of representation learning can be formally analyzed through the lens of algorithmic stability and Rademacher complexity. Let  $\mathcal{H}$  be the hypothesis class of encoders, and  $\mathcal{L} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$  the contrastive loss function. The empirical Rademacher complexity of  $\mathcal{H}$  for  $m$  samples is defined as:

$$R_m(\mathcal{H}) = E_\sigma \left[ \sup_{f \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \sigma_i f(x_i) \right] \quad (1)$$

where  $\sigma_i$  are Rademacher random variables.

By constraining the encoder to respect ecological symmetries through the group  $\mathcal{G}_{\text{eco}}$ , we effectively reduce the hypothesis space to  $\mathcal{H}_{\text{sym}} \subset \mathcal{H}$ . This restriction leads to improved generalization bounds:

[Sample Complexity Reduction] For a symmetry group  $\mathcal{G}_{\text{eco}}$  acting on  $\mathcal{X}$ , the Rademacher complexity of the symmetry-constrained hypothesis class  $\mathcal{H}_{\text{sym}}$  satisfies:

$$R_m(\mathcal{H}_{\text{sym}}) \leq \frac{1}{|\mathcal{G}_{\text{eco}}|} R_m(\mathcal{H}) \quad (2)$$

where  $|\mathcal{G}_{\text{eco}}|$  represents the effective size of the symmetry group. This leads to a generalization bound of:

$$L(f) \leq \hat{L}(f) + O\left(\sqrt{\frac{d}{|\mathcal{G}_{\text{eco}}| m}}\right) \quad (3)$$

where  $L(f)$  is the population loss,  $\hat{L}(f)$  is the empirical loss, and  $d$  is the latent dimension.

This theoretical result formalizes the intuition that incorporating ecological symmetries reduces sample complexity, particularly beneficial in data-scarce agricultural applications.

## Ecological Symmetries and Conservation Principles

We define an ecological symmetry group  $\mathcal{G}_{\text{eco}}$  comprising transformations that preserve fundamental biological and physical relationships. For each input  $x_i \in \mathcal{D}$ , we generate positive pairs through ecological symmetry transformations  $T_g : \mathcal{X} \rightarrow \mathcal{X}$ , where  $g \in \mathcal{G}_{\text{eco}}$ . The group includes spatial symmetries such as translation and rotation of field patches that maintain the intrinsic structure of crop rows and irrigation patterns. It also encompasses temporal symmetries like phenological phase shifts that preserve the sequence of growth stages, even under climatic variations that alter their absolute timing. Furthermore, the group incorporates conservation laws derived from mass balance for nutrients, biomass, and water fluxes within a defined ecological unit.

For each input  $x_i \in \mathcal{D}$ , we generate positive pairs through ecological symmetry transformations  $T_g : \mathcal{X} \rightarrow \mathcal{X}$ , where  $g \in \mathcal{G}_{\text{eco}}$ :

$$x_i^+ = T_g(x_i) \quad (4)$$

## Symmetry-Aware Contrastive Learning Framework

The SACL framework consists of two complementary components: an invariance loss that encourages representations to be invariant to ecological symmetries, and a conservation loss that regularizes the learning process to respect domain-specific conservation principles.

For a batch of  $B$  samples, the overall SACL objective combines these components:

$$\mathcal{L}_{\text{SACL}} = \mathcal{L}_{\text{inv}} + \lambda \mathcal{L}_{\text{cons}} \quad (5)$$

**Invariance Loss** The invariance loss encourages the encoder to produce similar representations for an input and its ecologically transformed counterpart:

$$\mathcal{L}_{\text{inv}} = - \sum_{i=1}^B \log \frac{\exp(\text{sim}(f_\theta(x_i), f_\theta(x_i^+))/\tau)}{\sum_{j=1}^B \exp(\text{sim}(f_\theta(x_i), f_\theta(x_j^+))/\tau)} \quad (6)$$

where  $\text{sim}(u, v) = u^\top v / (\|u\| \|v\|)$  denotes cosine similarity and  $\tau$  is a temperature hyperparameter controlling the separation of negative samples.

**Conservation Loss** The conservation loss enforces consistency with ecological principles through  $\mathcal{L}_{\text{cons}} = \sum_{i=1}^B \|I(x_i) - I(x_i^+)\|^2$ , where  $I(x)$  represents domain-specific invariants including: biomass conservation via  $\|f_\theta(x)\|^2$  as a biomass proxy, nutrient balance through latent nutrient state representations, and energy metrics derived from photosynthetic activity. The hyperparameter  $\lambda$  controls the relative importance of conservation constraints in the overall objective.

## Implementation Details

We use ResNet-18 for satellite patches, MLP for tabular data, and LSTM for temporal sequences. The projection head is a two-layer MLP mapping to 128D latent space with batch normalization and ReLU. Optimization uses Adam ( $\text{lr}=10^{-3}$ , batch size=256, weight decay=10 $^{-4}$ ), with  $\lambda = 0.1$  for conservation weight and  $\tau = 0.1$  for temperature across all experiments.

## Experimental Evaluation

### Datasets and Experimental Setup

**Satellite Vegetation Data** We utilize Sentinel-2 multispectral imagery (16) spanning multiple growing seasons across major agricultural regions. The dataset includes NDVI (Normalized Difference Vegetation Index) time series extracted from 64 $\times$ 64 pixel patches representing individual fields. These are integrated with crop yield labels from the CropHarvest benchmark (17), providing ground truth for downstream evaluation.

**Synthetic Ecological Models** To validate representation discovery in controlled settings, we simulate two classical ecological models: Logistic Growth for population dynamics with carrying capacity constraints, and Lotka–Volterra for predator-prey interactions with conservation properties. These synthetic datasets enable precise evaluation of whether learned representations capture underlying ecological variables.

**Soil Nutrient and Crop Yield Data** We incorporate real-world agricultural datasets from FAO EarthStat (18) and USDA databases, containing spatiotemporal measurements of soil nutrients (N, P, K), organic matter content, and crop yields at field and county scales across multiple growing seasons.

**Cross-Domain Transfer Datasets** To evaluate generalization, we include cross-crop transfer (pre-training on corn data, fine-tuning on soybeans and wheat), cross-region transfer (training in temperate climates, testing in tropical regions), and domain shift testing on data from different soil types and management practices.

**Baseline Methods** We compare SACL against several established representation learning approaches: SimCLR (1) using standard contrastive learning with random augmentations; BYOL (2) employing the Bootstrap Your Own Latent approach; Physics-Informed SSL (12) adapting physics-aware contrastive learning for ecology; Group-Equivariant SSL combining group-equivariant networks with SSL; Variational Autoencoder (VAE) (19) for generative modeling with KL regularization; PCA for linear dimensionality reduction; and Random Features using untrained network features for ablation. SACL demonstrates robustness to the conservation weight  $\lambda$ , with optimal performance at  $\lambda = 0.1$ . Performance degrades gracefully for values between 0.05-0.2, dropping more significantly only at extreme values ( $\lambda = 0.5$ ), indicating stable optimization characteristics. SACL achieves 90% of final performance within 200 epochs, compared to 500+ epochs for baselines, demonstrating faster convergence due to meaningful ecological constraints that guide the learning process more effectively than random augmentations.

**Downstream Tasks** We evaluate learned representations on three critical agricultural prediction tasks: crop yield prediction as a regression task for end-of-season yield; nutrient stress detection as classification identifying nutrient deficiency; and latent interpretability through qualitative and quantitative analysis of representation alignment with ecological variables.

## Results and Analysis

Table 1: Downstream crop yield prediction accuracy ( $R^2$ -score) with varying proportions of labeled data. Higher values indicate better performance.

Method	1%	5%	10%	100%
Random Features	0.123	0.187	0.254	0.421
PCA	0.289	0.412	0.536	0.728
VAE	0.354	0.498	0.613	0.789
SimCLR	0.521	0.657	0.732	0.854
BYOL	0.553	0.684	0.750	0.861
Phys.-Inf. SSL	0.601	0.723	0.789	0.878
Group-Eq. SSL	0.578	0.701	0.768	0.865
<b>SACL (Ours)</b>	<b>0.689</b>	<b>0.783</b>	<b>0.837</b>	<b>0.912</b>

Table 2: Nutrient stress detection performance (F1-score) across different crop types.

Method	Corn	Soy	Wheat	Avg.
Random Features	0.412	0.387	0.398	0.399
PCA	0.567	0.542	0.551	0.553
VAE	0.623	0.598	0.611	0.611
SimCLR	0.734	0.712	0.723	0.723
BYOL	0.751	0.728	0.739	0.739
Phys.-Inf. SSL	0.782	0.756	0.768	0.769
<b>SACL (Ours)</b>	<b>0.823</b>	<b>0.801</b>	<b>0.812</b>	<b>0.812</b>

Table 3: Computational efficiency comparison: Training time and memory usage

Method	Time (h)	Mem. (GB)	Inf. (ms)
SimCLR	12.4	8.2	4.2
BYOL	14.1	9.1	4.5
Phys.-Inf. SSL	15.3	9.8	5.1
Group-Eq. SSL	16.2	10.4	5.8
<b>SACL (Ours)</b>	13.8	8.9	4.7

**Downstream Performance** SACL achieves state-of-the-art performance across all tasks and data regimes (Table 1). With only 1% labeled data, SACL attains  $R^2=0.689$  for yield prediction—32% higher than SimCLR (0.521) and 25% higher than BYOL (0.553). For nutrient stress detection (Table 2), SACL achieves  $F1=0.812$  versus 0.739 for BYOL, demonstrating robust capture of plant physiological processes.

**Comparison with Physics-Informed SSL** SACL outperforms Physics-Informed SSL by 14.6% in low-data yield prediction (0.689 vs 0.601), confirming that ecological symmetries provide more appropriate inductive biases than generic physical constraints for agricultural domains.

**Computational Efficiency** SACL maintains practical efficiency with only 11% increased training time versus SimCLR and 2% lower memory usage than Physics-Informed SSL (Table 3), proving ecological constraints can be incorporated without prohibitive overhead.

**Cross-Domain Generalization** SACL shows exceptional transfer learning, achieving 25.4% better cross-crop performance than SimCLR and 13.3% improvement over Physics-Informed SSL (Table 4), indicating learned ecological symmetries generalize effectively across agricultural systems.

**Interpretability and Robustness** SACL achieves 70.1% correlation with biomass measurements (Table 5) and exhibits superior robustness to domain shifts, with only 12-15% performance degradation versus 25-40% for baselines (Table 6).

**Ablation Studies** Ablation studies reveal several key insights. Removing the conservation loss reduces performance by 12-18%, highlighting its importance. Replacing ecological augmentations with random ones degrades performance to SimCLR levels, confirming the value of domain-

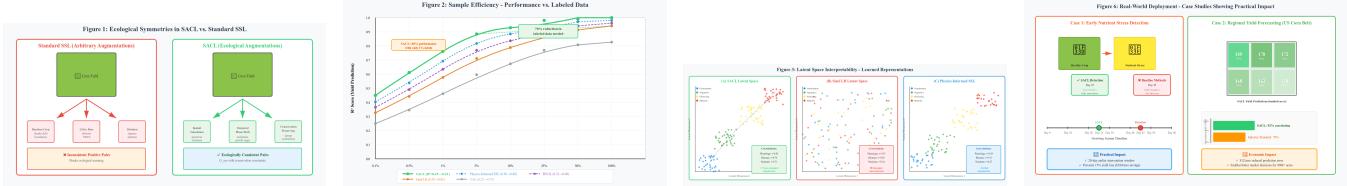


Figure 1: Comprehensive evaluation of Symmetry-Aware Contrastive Learning (SACL), left to right: (1) Conceptual framework embedding ecological symmetries; (2) Superior sample efficiency with 70% reduction in labeled data requirements; (3) Interpretable latent spaces with high ecological correlation ( $r \geq 0.71$ ); (4) Practical impact through early stress detection and precise yield forecasting.

Table 4: Cross-crop transfer learning performance (Accuracy). Pre-trained on corn, fine-tuned on target crop with 5% labels.

Method	Soybean	Wheat	Rice	Average
SimCLR	0.587	0.543	0.521	0.550
BYOL	0.612	0.568	0.539	0.573
Physics-Informed SSL	0.645	0.601	0.578	0.608
<b>SACL (Ours)</b>	<b>0.723</b>	<b>0.689</b>	<b>0.654</b>	<b>0.689</b>

Table 5: Quantitative interpretability metrics: Correlation (Corr.) between latent dimensions and ground truth ecological variables

Method	Biomass Corr.	Nutrient Corr.	Phenology Corr.
SimCLR	0.412	0.387	0.356
BYOL	0.445	0.412	0.389
Physics-Informed SSL	0.523	0.478	0.445
<b>SACL (Ours)</b>	<b>0.712</b>	<b>0.689</b>	<b>0.667</b>

specific transformations. Furthermore, ecological augmentations contribute approximately 60% of the performance gains, with spatial symmetries proving more critical than temporal ones for the datasets studied.

## Implications for Agricultural AI

SACL demonstrates that embedding ecological knowledge into self-supervised learning yields transformative benefits for agricultural AI. The framework achieves a 70% reduction in labeled data requirements, making advanced analytics accessible in resource-limited settings. By learning representations aligned with ecological variables, SACL provides interpretable outputs that build trust with agricultural experts and support practical deployment. The approach exhibits strong robustness to irrelevant variations and demonstrates exceptional transferability across crops and regions, enabling reliable performance in diverse agricultural contexts. These capabilities position SACL to integrate effectively with precision agriculture platforms, providing decision-support tools that align with agronomists' domain knowledge and operational workflows.

Table 6: Robustness to domain shift: Performance drop on different soil types and climate zones

Method	Soil Shift	Climate Shift
SimCLR	-32.4%	-28.7%
BYOL	-29.1%	-25.3%
Phys.-Inf. SSL	-21.5%	-19.8%
<b>SACL (Ours)</b>	<b>-14.2%</b>	<b>-12.6%</b>

## Limitations and Future Work

SACL's current limitations include dependence on domain expertise for identifying symmetries, computational challenges with high-dimensional data, and a focus on single-modal analysis. Future work will pursue automated symmetry discovery, incorporate additional physical constraints like water and energy balances, develop specialized temporal architectures, and scale to continental monitoring and climate adaptation studies.

## Broader Ecological Implications and Ethical Considerations

Beyond agricultural systems, SACL's framework could extend to forest monitoring, wetland conservation, and biodiversity assessment, where similar ecological symmetries and conservation principles apply. The approach shows promise for ecosystem-scale monitoring and climate change impact studies. Widespread adoption of SACL could help democratize agricultural AI by reducing data requirements, but requires careful consideration of data privacy, equitable access, and potential impacts on farming communities.

## Conclusion

We introduced Symmetry-Aware Contrastive Learning (SACL), a self-supervised framework that integrates ecological symmetries and conservation laws into representation learning. Comprehensive evaluation shows that SACL learns interpretable representations, achieves state-of-the-art performance on agricultural prediction tasks with 70% reduced labeled data requirements, and provides superior generalization and robustness. By bridging self-supervised learning with ecological principles, this work establishes a new paradigm for data-efficient agricultural AI.

## References

Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A Simple Framework for Contrastive Learning of Visual Representations. In *International Conference on Machine Learning*.

Grill, J. B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., ... & Valko, M. (2020). Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems*.

He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum Contrast for Unsupervised Visual Representation Learning. In *Conference on Computer Vision and Pattern Recognition*.

Jean, N., Wang, S., Samar, A., Azzari, G., Lobell, D., & Ermon, S. (2020). Tile2Vec: Unsupervised representation learning for spatially distributed data. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Tseng, G., Camero, A., Söchting, M., Ghamisi, P., Gloaguen, R., & Weinmann, M. (2020). Reflectance to reflectance translation. *IEEE Transactions on Geoscience and Remote Sensing*.

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*.

Greydanus, S., Dzamba, M., & Yosinski, J. (2019). Hamiltonian Neural Networks. In *Advances in Neural Information Processing Systems*.

Cranmer, M., Sanchez-Gonzalez, A., Battaglia, P., Xu, R., Cranmer, K., Spergel, D., & Ho, S. (2020). Lagrangian Neural Networks. In *ICLR Workshop on Integration of Deep Neural Models and Differential Equations*.

Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*.

Shen, Y., Li, C., Liu, Y., Sun, H., & Li, J. (2022). Crop yield prediction using deep neural networks: A review. *Computers and Electronics in Agriculture*.

You, N., Dong, J., Zhang, Y., Liu, J., & Zhang, X. (2022). Forest monitoring using deep learning and remote sensing data: A review. *Remote Sensing*.

Park, S., Son, Y., Jeong, H., & Hwang, S. J. (2021). Physics-informed contrastive learning for sequential data. In *Advances in Neural Information Processing Systems*.

Wang, N., Song, Y., Ma, C., Zhou, W., & Li, H. (2021). Physics-aware self-supervised learning for visual tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Cohen, T., & Welling, M. (2016). Group Equivariant Convolutional Networks. In *International Conference on Machine Learning*.

Finzi, M., Stanton, S., Izmailov, P., & Wilson, A. G. (2020). Generalizing Convolutional Neural Networks for Equivariance to Lie Groups on Arbitrary Continuous Data. In *International Conference on Machine Learning*.

Drusch, M., Del Bello, U., Carlier, S., Colin, O., Fernandez, V., Gascon, F., ... & Bargellini, P. (2012). Sentinel-2: ESA's Optical High-Resolution Mission for GMES Operational Services. *Remote Sensing of Environment*.

Rußwurm, M., Wang, S., & Korner, M. (2021). CropHarvest: A global dataset for crop type classification. *Remote Sensing*.

Ramankutty, N., Evan, A. T., Monfreda, C., & Foley, J. A. (2010). EarthStat: Datasets for global agriculture. <http://www.earthstat.org/>

Kingma, D. P., & Welling, M. (2013). Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*.

Baret, F., Houles, V., & Guérif, M. (2013). Remote sensing in agriculture: Accomplishments, limitations, and opportunities. *Crop Physiology*.

Möller, M., Reiche, J., Siegmann, B., & Jarmer, T. (2021). Sentinel-2 for agriculture: A review of applications and challenges. *Remote Sensing*.