Leveraging Physics Anomaly Knowledge and Conditional Contrastive Learning for Region-Agnostic Landslide Prediction

Ren Ozeki[1,2], **Hirozumi Yamaguchi**[1],
[1]The University of Osaka
[2]Japan Meteorological Agency
r-ozeki@ist.osaka-u.ac.jp, h-yamagu@ist.osaka-u.ac.jp

Abstract

Global climate change has increased the frequency of rainfall-induced disasters, such as landslides, even in previously safe regions, highlighting the need for accurate and region-generalizable prediction systems. Landslide datasets are imbalanced and regionally heterogeneous, as events are rare and influenced by local rainfall, soil, and vegetation. Existing methods either use physically grounded anomaly detection, which performs reasonably in data-scarce regions but may miss complex spatiotemporal patterns, or deep learning models, which capture complex patterns but often misinterpret rare safe conditions. To combine their strengths, we propose a hybrid framework that distills knowledge from physical models and applies region-conditioned supervised contrastive learning to learn latent representations generalizable across regions. Using a three-layer tank model to simulate the Soil Water Index, the framework jointly trains on physical outputs and disaster labels, integrating interpretable cues with deep representations. Evaluated on rainfall and terrain data from 22 Japanese regions using a Leave-K-Region-Out scheme, the proposed method achieved 50% higher PR-AUC in unseen regions compared to state-of-the-art baselines, demonstrating its potential for region-agnostic, physically informed landslide prediction.

Introduction

Global climate change has made the development and enhancement of disaster management systems a pressing worldwide issue (Malik et al. 2022; Newman and Noy 2023; Association et al. 2022). In recent years, large-scale disasters have frequently occurred even in areas previously considered safe, and the impact of these disasters is expanding (Association et al. 2022). Although hard countermeasures for disasters, such as levees and other physical infrastructures, have advanced significantly due to technological progress, it has become clear that these measures alone are insufficient (Kazama and Noda 2012). Therefore, effective disaster information is crucial to reduce human casualties and minimize social impacts (Suppasri et al. 2013).

The effectiveness of disaster information depends on accurate disaster prediction, which allows people to secure time for evacuation and move to safe locations. Among natural hazards, rainfall-induced disasters (e.g., landslides, floods) are particularly sensitive to climate change, and shifts in rainfall patterns may cause disasters to occur in regions that were previously unaffected (Association et al. 2022). Especially, landslides (including slope failures and debris flows) are responsible for approximately 40% of fatalities from natural disasters, representing a direct threat to human life(Association et al. 2022). This study focuses on predicting landslides, which are among the most life-threatening disasters.

Landslide disasters are rare events, which leads to datasets being dominated by non-disaster observations and class imbalance between positive and negative samples. For instance, in a weather-related dataset in one year, only a few landslide events may be recorded, while normal observations can number in the thousands. In some regions, historical disaster records may be absent due to the lack of past landslide occurrences (Fig. 1). In this work, we refer to such regions as ”non-disaster regions.”

In addition, the occurrence of hydrological landslide depends heavily on region-specific factors, such as rainfall intensity and duration, soil properties, vegetation conditions, and infrastructure development (Dhakal et al. 2025; Korup, Seidemann, and Mohr 2019). As a result, the mechanisms, frequency, and effective predictive indicators for landslides vary by region. Even within the same region, regional characteristics may change over time due to seismic activity, climate change, and urbanization. Thus, the data distributions of landslides are non-independent and non-identically distributed (non-IID) across regions, which makes it challenging to directly transfer a prediction model trained in one region to another or to fine-tune a model in a new area with scarce occurrence data. While landslide prediction by binary classification may work in a single region where sufficient landslide occurrence data is accumulated, it is a significant challenge to predict landslides across multiple regions, which include non-disaster regions.

To address this issue, existing studies (Teja, Dikshit, and Satyam 2019; KURAMOTO et al. 2001) have widely adopted region-specific anomaly detection models that are grounded in physical knowledge of landslide mechanisms. These models focus on a limited set of explanatory variables, such as soil water index and cumulative rainfall, that are physically known to exhibit a positive correlation with landslide occurrence(KURAMOTO et al. 2001; Yue et al. 2025). By restricting the analysis to physically meaningful variables, these models ensure that the computed anomaly scores are aligned with actual landslide risk, enabling reasonable predictions even when occurrence data are scarce. Recently, deep learning models for anomaly detection(Collini et al. 2022; Xie, Zhou, and Chai 2019; Wang et al. 2017; Khalili et al. 2024; Collini et al. 2022) have been explored to handle more complex spatio-temporal structures and diverse observational data. However, when raw spatio-temporal inputs are provided directly, the resulting anomaly scores often reflect statistical rarity rather than true hazard potential. For example, rare but safe patterns, such as unusual terrain formations or atypical rainfall distributions, may be incorrectly assigned high anomaly scores. Consequently, the predictive performance of such purely data-driven approaches is often limited.

Thus, the two approaches exhibit complementary strengths and limitations: physically grounded anomaly detection models provide wide applicability, whereas deep learning models offer strong representational capacity for capturing complex spatiotemporal patterns. To bridge these advantages, we propose a hybrid framework that distills knowledge from physically grounded anomaly detection models while incorporating region-conditioned supervised contrastive learning to obtain latent representations that generalize across regions. For the physical component, we employ the three-layer tank model proposed by (KURAMOTO et al. 2001), which has been widely used by meteorological agencies worldwide(Yue et al. 2025). This model simulates the Soil Water Index (SWI) and quantifies the statistical deviation of current rainfall events from historical observations as an indicator of hydrological abnormality. Our prediction model is trained using both the outputs of this physical model and the actual disaster occurrence labels. This joint training enables the model to learn effectively even in non-disaster regions where no past disaster events are recorded. To further improve regional generalization, we proposed region-conditioned contrastive learning as a pre-training method, which prevents the latent representations from separating by region. By combining the broad applicability and physical interpretability of the anomaly detection model with the strong representational capacity of deep learning models for capturing complex spatiotemporal patterns, our framework effectively addresses the scarcity and regional heterogeneity of landslide occurrence data, which leads to region-agnostic and physically grounded landslide prediction.

To evaluate the predictive performance of the proposed method, rainfall and terrain data spanning two years across Japan were collected, and data from 22 regions where landslides occurred were used for evaluation. The proposed method was assessed using a Leave-K-Region-Out manner, in which K regions are excluded from training and the model is evaluated on these held-out regions. To appropriately measure performance on class-imbalanced disaster data, PR-AUC and Precision were employed as the main metrics. The proposed method achieved a 50% higher PR-AUC in unseen regions compared to other state-of-the-art methods.

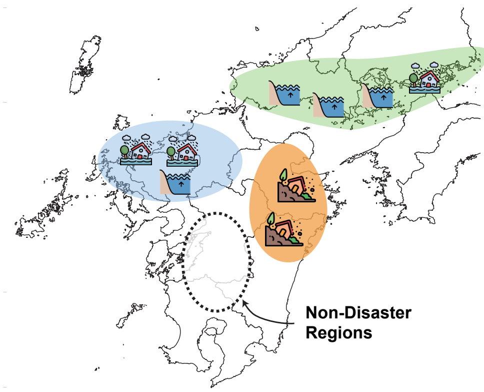

Figure 1: Landslide Data Distribution

Related Work

Landslide Prediction

With the recent global climate change, the frequency of landslides has increased even in regions that were previously considered at low risk, highlighting the growing importance of accurate and reliable landslide prediction. Traditional landslide prediction methods have primarily been algorithm-based approaches (Teja, Dikshit, and Satyam 2019; Ho and Lee 2017; Nguyen et al. 2019; Liu et al. 2020). For example, (Teja, Dikshit, and Satyam 2019) proposed a statistical early warning method using an Indian dataset to evaluate the collapse volume causing landslides in specific regions. However, these approaches do not account for the spatiotemporal dependencies of landslides, such as soil distribution, elevation, and water accumulation due to rainfall, and thus are limited in terms of regional generalization and prediction accuracy.

To address this limitation, machine learning methods that capture the spatiotemporal dependencies of landslides have been proposed (Collini et al. 2022; Xie, Zhou, and Chai 2019; Zhang et al. 2021). For instance, (Collini et al. 2022) employed CNNs to model spatial dependencies for landslide prediction, while (Xie, Zhou, and Chai 2019) used LSTMs to predict ground displacement as an indicator of impending disasters. Although these methods show promise in terms of accuracy, real-world deployment of landslide prediction systems requires ensuring high recall to prevent situations where residents cannot evacuate. At the same time, disaster information must maintain high precision to minimize false alarms, as low-accuracy predictions may undermine public trust. Since landslides are rare and extreme events, achieving both high recall and high precision is challenging from the perspective of class imbalance, yet many existing studies do not sufficiently consider the class imbalance problem or the precision requirements specific to disaster prevention systems. Moreover, most prior studies are limited to case studies, providing insufficient evaluation of regional generalization.

In contrast, the proposed system addresses class imbalance in landslide prediction and the precision requirements specific to disaster prevention systems. The system's perfor-

mance is evaluated using PR-AUC and precision at 80% recall on a large-scale dataset rather than on isolated case studies. Specifically, rainfall, topography, and landslide data spanning two years across Japan were collected, and the system was evaluated on 22 regions where disasters occurred.

Learning from Scarce and Imbalanced Data

Disaster prediction requires handling limited and imbalanced data due to the uneven frequency of events. Extensive research has addressed this class imbalance problem (Johnson and Khoshgoftaar 2019; Valverde-Albacete and Peláez-Moreno 2014; Tasci et al. 2022; Alsaui et al. 2022; Singh, Ranjan, and Tiwari 2022). Representative approaches include data-level and algorithm-level methods. Data-level methods balance the training data by adjusting class ratios through undersampling or oversampling. SMOTE (Chawla et al. 2002) and its variants are widely used and have been reported to be effective for simple machine learning models. Algorithm-level methods include cost-sensitive learning, which introduces class weights into the loss function, and Focal Loss (Lin et al. 2017), which focuses on difficult samples. Ensemble learning methods (Liu et al. 2017; Díez-Pastor et al. 2015) that combine multiple classifiers have also been shown to handle imbalanced data effectively. However, these approaches are mainly effective for shallow networks or single-modality data, and their utility is limited for large-scale deep learning models that need to capture complex spatiotemporal features such as subsurface water accumulation. This limitation arises because it is difficult to stably learn high-level latent representations from minority data, which are strongly affected by sampling bias.

To overcome this challenge, this study proposes a new learning framework combining a hybrid architecture that integrates physical model knowledge with supervised contrastive pretraining. The proposed architecture uses the physical model to reproduce fundamental hydrological behaviors, while the deep learning model learns the deviation from the physical model output. This design appropriately constrains the search space of the machine learning model and reduces the relative impact of class imbalance by learning deviations from the physical model rather than directly predicting the occurrence of events. Furthermore, supervised contrastive learning with spatiotemporal data augmentation enables the model to acquire robust and generalizable latent representations even from limited disaster occurrence data.

Domain Adaptation Techniques

When applying disaster prediction models to different regions, distribution shifts caused by differences in topography, climate, and land use are known to significantly degrade performance, even if the model achieves high accuracy in the training region (Zhou et al. 2022; Ozeki et al. 2024). To address this, regional adaptation methods using transfer learning or meta-learning frameworks have been proposed (Jin, Chen, and Yang 2022; Dong et al. 2024; Chen et al. 2022). These methods adapt to new regions via fine-tuning on a small amount of labeled data and are effective when sufficient disaster occurrence data are available in the target region. However, because disasters are rare events, new regions often lack such data, limiting the applicability of these methods.

To overcome this limitation, zero-shot domain adaptation, known as domain generalization, has been studied (Zhou et al. 2022; Wang et al. 2025). Domain generalization typically aims to learn domain-invariant latent representations by introducing adversarial learning against a region classifier (Ozeki et al. 2023). However, adversarial learning-based methods face challenges as the number of regions increases: the region classifier can become unstable, and training may fail to converge.

In this study, we introduce a constraint during the pairing stage of contrastive learning that prevents data from the same region from being pulled together. This approach suppresses excessive separation of latent representations by region, allowing the proposed method to learn features that are robust to inter-regional distribution shifts and to generalize effectively to unseen regions.

Proposed Method

The proposed landslide prediction system is designed to address two major challenges in landeslide prediction: (1) the scarcity of landslide occurrences data, and (2) data heterogeneity among regions. To tackle these challenges, our framework combines region-conditioned pre-training and fine-tuning with knowledge distillation from a physically grounded anomaly detection model. Figure 2 provides an overview of the proposed system. The prediction pipeline consists of two main stages: *pre-training* and *fine-tuning*. In the pre-training stage, a machine learning model is trained using *region-conditioned supervised contrastive learning*, which encourages the model to acquire latent representations primarily driven by class (i.e., disaster occurrence) while preventing the representations from being dominated by region-specific features. This enables the model to generalize across regions, including those without prior landslide occurrences.

Following pre-training, the model undergoes a *fine-tuning* stage in which knowledge is distilled from a physics-based anomaly detection model. Specifically, the outputs of the three-layer tank model, which simulates hydrological conditions and computes anomaly scores based on the Soil Water Index and rainfall deviations, are used as soft targets. This knowledge distillation allows our prediction model to learn meaningful patterns even in non-disaster regions, ensuring stable and physically interpretable predictions where historical landslide data are absent.

The key strengths of the proposed system are twofold: first, the pre-training method provides *region-agnostic latent representations* that capture the essential characteristics of landslide events; second, fine-tuning with knowledge distillation from a physically grounded model ensures that the predictive model remains accurate and interpretable, even in regions without prior disaster occurrences. In the following sections, we describe the proposed system in detail.

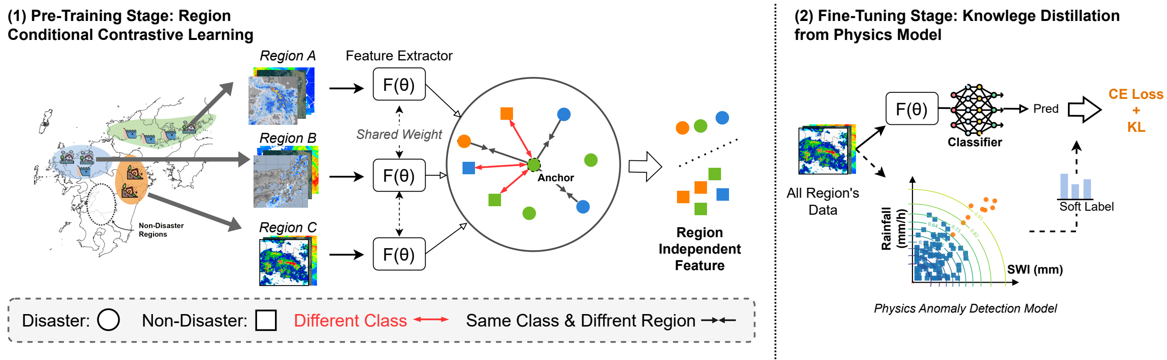

Figure 2: System Overview

Model Architecture

The machine learning model of the proposed system is designed to integrate both physically grounded knowledge and raw observational data to effectively predict landslide occurrence. The model receives three types of inputs: (1) outputs from the physical anomaly detection model, (2) spatial data such as terrain and soil properties, and (3) spatiotemporal rainfall intensity data. The output of the physical model represents the degree of abnormality for a given rainfall event, serving as a meaningful indicator for landslide prediction. By incorporating this information as part of the ML model input, the system leverages physically interpretable signals to guide learning, particularly in regions with scarce historical landslide events.

Spatial features are extracted from terrain-related data using a two-layer Convolutional Neural Network (CNN), which captures local spatial patterns and correlations. Spatio-Temporal features from rainfall observations are processed using CNN-Transformer which consists of two-layer CNN and three-layer Transformer. The CNN-Transformer encodes temporal dependencies and complex interactions in the rainfall sequence. The Spatio-temporal features are summarized by applying a mean pooling operation over the sequence length to obtain a fixed-size representation. Finally, the spatial features from the CNN, the summarized spatio-temporal features from the CNN-Transformer, and the physical model output are concatenated into a single feature vector, which serves as the input to downstream prediction heads (e.g., for binary landslide occurrence classification or metrics learning). This design ensures that the model effectively integrates heterogeneous information from physical simulations and observational data, providing a robust basis for the subsequent pre-training and fine-tuning stages.

Pre-training Stage: Region-Conditional Contrastive Learning (RCCL)

Landslide prediction is influenced by region-specific factors such as topography, soil properties, and vegetation, causing data distributions to vary across regions. Hence, it is crucial for a trained model to generalize to unseen regions, i.e., to have strong regional generalization performance. To address this, we propose Region-Conditional Contrastive

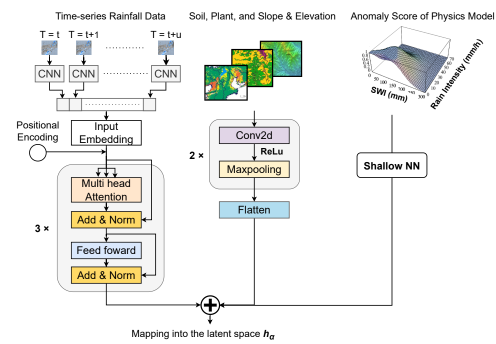

Figure 3: Model Architecture

Learning (RCCL) as a pretraining approach. RCCL aims to learn environment-invariant class representations across multiple source domains by extending positive pairs in supervised contrastive learning (e.g., SupCon(Khosla et al. 2020)) based on regional conditions.

In conventional SupCon, the latent representations are learned to separate samples by class. However, in practice, samples from the same class but different trained models must generalize from different classes in the same region (Fig. 2). This occurs because regional features dominate, causing latent representations to reflect regional differences more than class-specific features, which decreases regional generalization. To resolve this, pretraining enforces class invariance in the feature extractor f_θ output, minimizing the conditional contrastive loss L_{MI} to bring same-class samples from different regions closer. For an anchor sample x_u in batch B, pairs are defined as:

- **Positive pairs** ($pos(u)$)**:**All samples x_v that belong to the same class but originate from different regions $[(y_u = y_v) \wedge (r_u \neq r_v)]$ are defined as positive pairs, where y denotes the class label and r denotes the region. This design encourages class-level feature alignment across regions, thereby mitigating regional bias.
- **Negative pairs** ($neg(u)$)**:** All samples x_k of a different class ($y_u \neq y_k$) are negative pairs. This maintains discriminability between classes independent of region.

This pairing allows the model to learn discriminative and robust features independent of regions. The conditional contrastive loss L_{MI} is defined as:

$$\mathcal{L}^u_{\text{MI}} = \frac{-1}{|pos(i)|} \sum_{p \in pos(i)} \log \frac{\exp\left(\text{sim}(\mathbf{z}_i, \mathbf{z}_p)/\tau\right)}{\sum_{a \in neg(i)} \exp\left(\text{sim}(\mathbf{z}_i, \mathbf{z}_a)/\tau\right)} \tag{1}$$

Fine-Tuning Stage: Knowledge Distillation from Physics Anomaly Detection Model

After the region-conditioned pre-training stage, the model is fine-tuned to perform landslide prediction in a supervised manner. In this stage, the feature extractor obtained from pre-training is used as a backbone, while a classification head is trained, and the feature extractor is further adjusted. The fine-tuning process is guided by three main loss components. At first, to address the severe class imbalance between landslide and non-landslide samples, we employ the Focal Loss (Lin et al. 2017), which emphasizes the contribution of hard-to-classify, minority class samples. Compared to feature extractors, classifiers have fewer parameters and are prone to overfitting. Thus, we introduce Stochastic Feature Augmentation (SFA) in the latent space for classifier fine-tuning. SFA perturbs the latent feature vector with random noise to improve model robustness, providing a simpler and more efficient alternative to data-space augmentation (Li et al. 2021). SFA enables learning robust decision boundaries that are stable under distribution shifts across regions. Specifically, for a latent feature vector z_i obtained by the feature extractor F_Θ, an augmented feature $\hat{z}_i$ is generated as:

$$\hat{z}_i = A(z_i) = \alpha \odot z_i + \beta$$

where the noise $\alpha \sim \mathcal{N}(1, \sigma_1 I)$ and $\beta \sim \mathcal{N}(0, \sigma_2 I)$ are sampled from data-independent multivariate Gaussian distributions. I is the identity matrix, and σ_1 and σ_2 are scalar hyper-parameters controlling noise intensity. During training, the augmented features $\hat{z}_i$ are used to train the classifier.

To incorporate physically grounded information, we distill knowledge from the anomaly detection model. Let q denote the output probability from the physical model and p the output from the ML model. The distillation loss is given by the Kullback–Leibler divergence:

$$\mathcal{L}_{\text{KD}} = \text{KL}(q||p) = \sum_i q_i \log \frac{q_i}{p_i} \tag{2}$$

This loss encourages the ML model to mimic the anomaly scores of the physical model, ensuring physically interpretable predictions even in non-disaster regions. The total fine-tuning loss is a weighted combination of the three components:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{Focal}} + \lambda_1 L_{\text{Focal-SFA}} + \lambda_2 \mathcal{L}_{\text{KD}} \tag{3}$$

where $\mathcal{L}_{\text{Focal}}$ is ordinal Focal loss for labeled data, $L_{\text{Focal-SFA}}$ is the Focal Loss computed on the augmented latent vectors, $\mathcal{L}_{\text{KD}}$ is distilliation loss from anomaly detection model, and λ_1, λ_2 are hyperparameters that control the relative importance of SFA and knowledge distillation, respectively. In addition, the Focal Loss itself has hyperparameters α_t and γ. In this study, $\lambda_1 = 3$ is used for fine-tuning the landslide prediction model. During inference, the SFA module is disabled, and F_Θ operates deterministically.

Evaluation

Data Collection and Configuration

To evaluate the proposed method, we collected data on landslide occurrences in Japan over two years (2021 and 2022). Since landslides depend on multiple factors, we focused on standard variables widely used in previous studies to construct the dataset. Unlike some prior systems, our proposed system does not rely on additional observation devices and can operate using generally accessible data. We selected 22 regions where more than three landslide events were confirmed during the observation period. Each region covered an area of 1 km $\times$ 1 km, and the regional division was based on Japan's MESH3 boundaries.

Landslide Events: We collected 253 landslide events that occurred in Japan between 2021 and 2022. Each event is annotated with the occurrence date and location (latitude and longitude up to six decimal places). The temporal resolution of the data is daily, and the proposed system predicts landslide occurrences at this daily granularity, which is sufficient for evacuation planning.

Rainfall: As shown in previous studies (Xie, Zhou, and Chai 2019), time-series rainfall data are crucial for landslide prediction. Even small cumulative rainfall over long periods can trigger landslides due to short-term intense rainfall, and vice versa. The dataset includes rainfall data from January 1, 2021, to December 31, 2022.

Soil Properties: Soil moisture is an important factor for landslide occurrences (Xie, Zhou, and Chai 2019; KURAMOTO et al. 2001). Soil slipperiness and drainage capacity significantly influence disaster risk. The dataset represents soil in 10 categories, including artificial pavements.

Vegetation: Roots of plants and trees stabilize groundwater and soil, so surrounding vegetation also affects landslide risk. In this study, 11 vegetation categories are defined.

Elevation and Slope Angle: Elevation and slope angle closely relate to water accumulation and gravitational effects, making them important for landslide prediction. Steeper slopes are generally associated with higher disaster risk. Elevation is represented as a continuous value, while slope angle is classified into 8 directions.

An example of the dataset is shown in Fig. 4. The data were collected at a 1 km mesh resolution based on Japan's MESH3 boundaries. For training and evaluation, days without landslide occurrences were treated as negative cases. To ensure the quality of negative samples, we used all data from rainy days in each region, resulting in a total of 3,567 samples in our dataset. The experimental settings are summarized in Table 1.

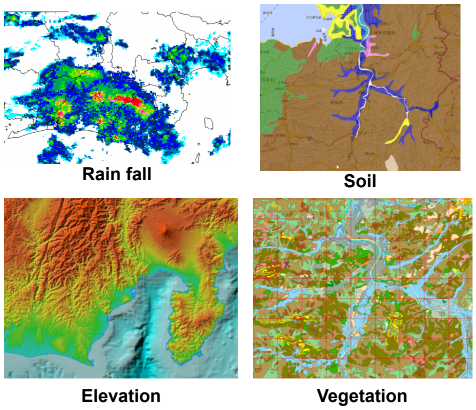

Figure 4: Positive sample from the dataset on July 3, 2021, in Shizuoka Prefecture. Top-left shows rainfall, with color intensity indicating rainfall amount. Bottom-left shows elevation, with red indicating higher areas. Top-right shows soil properties, and the bottom-right shows vegetation categories.

Table 1: Experimental settings

Parameter	Search range (**bold**: default value)
Batch size	{8, 16, 32, 64, **128**}
Temperature τ	{0.3, **0.5**, 0.8}
Learning rate	{1-e3, **1-e4**, 1-e5}
Optimizer	AdamW
Pretraining epochs	100
Fine-tuning epochs	50
CTLGNet (Transformer) layers	{**3**, 6}

Evaluation Metrics

Our proposed system aims to predict landslides accurately. In disaster prediction, recall is especially important as it directly affects residents' evacuation actions. At the same time, maintaining precision is necessary to avoid excessive false alarms, which can undermine trust. Thus, we adopt *precision at 80% recall* as the main evaluation metric, and we refer to the precision at 80% recall simply as precision hereafter. To further assess whether the method correctly ranks the severity of each event, we also use PR-AUC on the test data. PR-AUC reflects the model's ability to rank actual events higher than non-events and evaluates the reliability of predictions beyond simple correctness. Correct ranking is critical in disaster prevention, as events with a higher likelihood should be prioritized for warning residents.

Comparison with State-of-the-Art Methods

We evaluated the proposed method by comparing prediction accuracy with state-of-the-art approaches. Prediction performance was measured both within trained regions and on K unseen regions (Leave-K-Region-Out). We compared the proposed method with SWI-EWS (KURAMOTO et al. 2001), CTLGNet (Zhao et al. 2024), PGNN-0 (Daw et al. 2022), and MMD-AAE (Li et al. 2018).

SWI-EWS (KURAMOTO et al. 2001): Uses a fixed-parameter three-layer tank model to calculate soil water index (SWI) from rainfall, then statistically evaluates the abnormality of rainfall events in the target region using past rainfall data to predict landslides.

CTLGNet (Zhao et al. 2024): A deep learning model combining CNN and Transformer to capture spatio-temporal features of landslides. If no model is specified for other baselines, CTLGNet is used.

PGNN-0 (Daw et al. 2022): Incorporates hydrological model predictions as inputs to ML models, leveraging physical knowledge for the target task. In this study, SWI-EWS outputs were used as inputs.

MMD-AAE (Li et al. 2018): A region generalization framework that simultaneously learns the main task and adversarially enforces region-invariant feature extraction.

As mentioned in the Introduction, anomaly detection models can be trained without disaster occurrence data, making them more applicable to non-disaster regions than binary classification-based landslide prediction models. However, deep spatio-temporal anomaly detection models tend to produce false detections in patterns that are unusual but not hazardous. To examine the characteristics of both binary classification and anomaly detection approaches using deep learning, the anomaly detection models with the aforementioned comparison methods are also evaluated. In Table 2, the anomaly detection versions of the comparison methods are denoted as "AD". In the specific implementation, each target model was first trained in an end-to-end manner, after which only the feature extractor was reused to perform training using non-occurrence data as normal samples.

As described in the Evaluation Metrics section, we used PR-AUC and precision at 80% recall to evaluate model performance. Table 2 presents these metrics for each method in the training regions. To prevent time-series leakage, the data were aggregated daily and stratified according to disaster occurrence labels, with 20% of the data reserved as the test set. The results demonstrate that the proposed method outperforms all baseline approaches, achieving 52% higher precision than the best alternative. Compared with SWI-EWS and CTLGNet, the proposed method effectively leverages limited and imbalanced data by integrating hydrological knowledge with deep learning. Compared with PGNN-0, it not only incorporates hydrological outputs as inputs but also employs skip connections to constrain the output space with hydrological predictions, enabling more effective feature extraction and deeper integration of physical and deep learning models. As a result, precision improved by 56% relative to PGNN-0. Furthermore, Table 2 indicates that landslide prediction using anomaly detection among the compared methods exhibits lower accuracy than the other approaches. This is likely because the inputs consist of complex spatio-temporal data, where the degree of abnormality does not necessarily correspond to the level of hazard, making it difficult to accurately predict landslide occurrences.

For regional generalization, the proposed method achieved 66% precision on target regions, compared with 17% for MMD-AAE, representing approximately 50% im-

Table 2: Performance comparison of the proposed and state-of-the-art methods (including anomaly detection variants)

Method	All-in		Leave-K-Out	
	PR-AUC	Precision	PR-AUC	Precision
SWI-EWS (KURAMOTO et al. 2001)	0.21	0.08	0.21	0.08
CTLGNet (Zhao et al. 2024)	0.62	0.29	0.53	0.16
CTLGNet + AD	0.01	0.01	0.02	0.02
PGNN-0 (Daw et al. 2022)	0.61	0.27	0.60	0.37
PGNN-0 + AD	0.02	0.02	0.02	0.02
MMD-AAE (Li et al. 2018)	0.58	0.19	0.42	0.17
MMD-AAE + AD	0.02	0.02	0.01	0.01
Proposed	**0.86**	**0.83**	**0.72**	**0.60**

provement. MMD-AAE aims for domain-independent prediction using adversarial learning, but increasing the number of domains can lead to unstable training and lower performance in target regions. In contrast, the proposed method learns meaningful representation through RCCL, which doesn't rely on adversarial training, and stochastic feature augmentation contributes to improved regional generalization. Compared with PGNN-0, the proposed method achieves 29% higher precision, demonstrating that while incorporating hydrological knowledge contributes to the model, pretraining with RCCL and fine-tuning with SFA are crucial for obtaining region-independent latent representations and achieving region-generalized landslide prediction.

Ablation Study

The proposed method integrates three key components: knowledge distillation from physics-based anomaly detection models, region-generalized contrastive learning (RCCL), and latent-space data augmentation. All of these modules enable our system to predict landslides accurately, even in non-disaster regions. To assess the contribution of each component, we conducted an ablation study in which individual modules were selectively disabled or replaced. Specifically, RCCL was replaced with supervised contrastive learning (SupCon) (Khosla et al. 2020) during the contrastive pretraining stage, and SFA was replaced with standard fine-tuning based on focal loss. Knowledge distillation from physics models was simply replaced with fine-tuning using normal focal loss. Figure 5 presents the results of these experiments. The full method consistently outperformed all ablated variants, demonstrating that each module contributes to overall performance. These findings indicate that the combination of knowledge distillation, region-generalized contrastive learning, and latent-space augmentation effectively addresses the challenges inherent in landslide prediction.

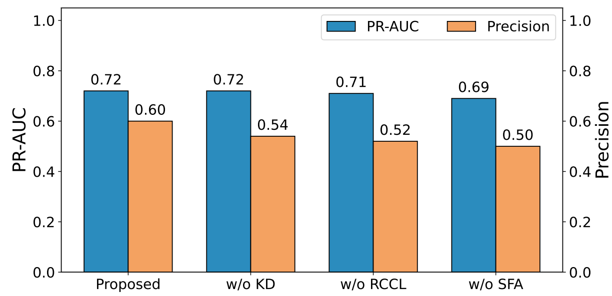

Figure 5: Ablation study on module contributions. ”w/o” indicates a variant without a specific module. For example, ”w/o RCCL” uses SupCon instead of RCCL for pretraining.

Conclusion and Future Work

In this study, we proposed a hybrid method combining physical models and deep learning to achieve region-independent landslide prediction. By incorporating hydrological knowledge based on a three-layer tank model, the proposed method stabilizes learning while maintaining physical consistency, even under limited and imbalanced disaster data conditions. Moreover, supervised contrastive learning that considers regional information suppresses the separation of latent representations caused by non-IID characteristics across regions, thereby improving generalization to unseen areas. Experiments using data from 22 regions across Japan demonstrate that the proposed method achieves up to 54% improvement in precision compared to conventional deep learning and statistical models for unseen regions. These results indicate that the proposed method can detect early signs of disaster occurrences without strongly depending on regional characteristics. In future work, we will extend the model to incorporate subsurface structural data from sources such as borehole data and apply the approach to other types of disasters, such as floods and inundations.

Acknowledgment

This work was partially funded by JST CREST Grant JPMJCR21M5, JST ACT-X JPMJAX25M6, JSPS KAKENHI Grant 22K12011, and NVIDIA award. The authors appreciate JMA, MLIT, JMBSC, the Ministry of Environment, and the Technical Report of the Geospatial Information Authority of Japan, which provides datasets to the authors. The JMA data used in this study was provided by way of the ”Meteorological Research Consortium”, a framework for research cooperation of JMA and MSJ.

References

Alsaui, A. A.; Alghofaili, Y. A.; Alghadeer, M.; and Alharbi, F. H. 2022. Resampling techniques for materials informatics: limitations in crystal point groups classification. *Journal of Chemical Information and Modeling*, 62(15): 3514–3523.

Association, W. M.; et al. 2022. State of the global climate 2021.

Chawla, N. V.; Bowyer, K. W.; Hall, L. O.; and Kegelmeyer, W. P. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16: 321–357.

Chen, Y.; Li, L.; Whiting, M.; Chen, F.; Sun, Z.; Song, K.; and Wang, Q. 2022. Observations and Geoinformation. *International Journal of Applied Earth Observation and Geoinformation*, 112: 102907.

Collini, E.; Palesi, L. I.; Nesi, P.; Pantaleo, G.; Nocentini, N.; and Rosi, A. 2022. Predicting and understanding landslide events with explainable AI. *IEEE Access*, 10: 31175–31189.

Daw, A.; Karpatne, A.; Watkins, W. D.; Read, J. S.; and Kumar, V. 2022. Physics-guided neural networks (pgnn): An application in lake temperature modeling. In *Knowledge guided machine learning*, 353–372. Chapman and Hall/CRC.

Dhakal, D.; Singh, K.; Onyelowe, K. C.; Cazco, S. A. S.; Sharma, A.; Alarifi, N.; Islam, F.; Randeep; Arunachalam, K. P.; and Youssef, Y. M. 2025. Enhancing landslide disaster prediction by evaluating non landslide area sampling in machine learning models for Spiti Valley India. *Scientific Reports*, 15(1): 12242.

Dong, A.; Dou, J.; Li, C.; Chen, Z.; Ji, J.; Xing, K.; Zhang, J.; and Daud, H. 2024. Accelerating cross-scene co-seismic landslide detection through progressive transfer learning and lightweight deep learning strategies. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–13.

Díez-Pastor, J. F.; Rodríguez, J. J.; García-Osorio, C.; and Kuncheva, L. I. 2015. Random Balance: Ensembles of variable priors classifiers for imbalanced data. *Knowledge-Based Systems*, 85: 96–111.

Ho, J.-Y.; and Lee, K. T. 2017. Performance evaluation of a physically based model for shallow landslide prediction. *Landslides*, 14(3).

Jin, Y.; Chen, K.; and Yang, Q. 2022. Selective cross-city transfer learning for traffic prediction via source city region re-weighting. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, 731–741.

Johnson, J. M.; and Khoshgoftaar, T. M. 2019. Survey on deep learning with class imbalance. *Journal of Big Data*, 6(1): 1–54.

Kazama, M.; and Noda, T. 2012. Damage statistics (Summary of the 2011 off the Pacific Coast of Tohoku Earthquake damage). *Soils and Foundations*, 52(5): 780–792.

Khalili, M. A.; Voosoghi, B.; Calcaterra, D.; Kouchakkapourchali, A.; Di Muro, C.; Madadi, S.; Tufano, R.; and Di Martire, D. 2024. Advanced Machine Learning Strategies for Landslide Detection. In *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*, 1755–1759. IEEE.

Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33: 18661–18673.

Korup, O.; Seidemann, J.; and Mohr, C. H. 2019. Increased landslide activity on forested hillslopes following two recent volcanic eruptions in Chile. *Nature Geoscience*, 12(4): 284–289.

KURAMOTO, K.; TETSUGA, H.; HIGASHI, N.; ARAKAWA, M.; NAKAYAMA, H.; and FURUKAWA, K. 2001. A study on a method for determining non-linear critical line of slope failures during heavy rainfall based on RBF network. *Doboku Gakkai Ronbunshu*, 2001(672): 117–132.

Li, H.; Pan, S. J.; Wang, S.; and Kot, A. C. 2018. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5400–5409.

Li, P.; Li, D.; Li, W.; Gong, S.; Fu, Y.; and Hospedales, T. M. 2021. A simple feature augmentation for domain generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8886–8895.

Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollar, P. 2017. Focal Loss for Dense Object Detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Liu, S.; Wang, Y.; Zhang, J.; Chen, C.; and Xiang, Y. 2017. Addressing the class imbalance problem in Twitter spam detection using ensemble learning. *Computers&Security*, 69: 35–49.

Liu, Z.; Guo, D.; Lacasse, S.; Li, J.-h.; Yang, B.; and Choi, J. C. 2020. Algorithms for intelligent prediction of landslide displacements.

Malik, A.; Li, M.; Lenzen, M.; Fry, J.; Liyanapathirana, N.; Beyer, K.; Boylan, S.; Lee, A.; Raubenheimer, D.; Geschke, A.; et al. 2022. Impacts of climate change and extreme weather on food supply chains cascade across sectors and regions in Australia. *Nature Food*, 3(8): 631–643.

Newman, R.; and Noy, I. 2023. The global costs of extreme weather that are attributable to climate change. *Nature Communications*, 14(1): 6103.

Nguyen, P. T.; Tuyen, T. T.; Shirzadi, A.; Pham, B. T.; Shahabi, H.; Omidvar, E.; Amini, A.; Entezami, H.; Prakash, I.; Phong, T. V.; et al. 2019. Development of a novel hybrid intelligence approach for landslide spatial prediction. *Applied Sciences*, 9(14): 2824.

Ozeki, R.; Yonekura, H.; Baimbetova, A.; Rizk, H.; and Yamaguchi, H. 2023. One Model Fits All: Cross-Region Taxi-Demand Forecasting. *In The 31st ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '23)*.

Ozeki, R.; Yonekura, H.; Rizk, H.; and Yamaguchi, H. 2024. Decentralized Landslide Disaster Prediction for Imbalanced and Distributed Data. In *2024 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 143–150. IEEE.

Singh, A.; Ranjan, R. K.; and Tiwari, A. 2022. Credit card fraud detection under extreme imbalanced data: a comparative study of data-level algorithms. *Journal of Experimental & Theoretical Artificial Intelligence*, 34(4): 571–598.

Suppasri, A.; Shuto, N.; Imamura, F.; Koshimura, S.; Mas, E.; and Yalciner, A. C. 2013. Lessons learned from the 2011 Great East Japan tsunami: performance of tsunami countermeasures, coastal buildings, and tsunami evacuation in Japan. *Pure and Applied Geophysics*, 170.

Tasci, E.; Zhuge, Y.; Camphausen, K.; and Krauze, A. V. 2022. Bias and Class Imbalance in Oncologic Data—Towards Inclusive and Transferrable AI in Large Scale Oncology Data Sets. *Cancers*, 14(12): 2897.

Teja, T. S.; Dikshit, A.; and Satyam, N. 2019. Determination of Rainfall Thresholds for Landslide Prediction Using an Algorithm-Based Approach: Case Study in the Darjeeling Himalayas, India. *Geosciences*, 9(7).

Valverde-Albacete, F. J.; and Peláez-Moreno, C. 2014. 100% classification accuracy considered harmful: The normalized information transfer factor explains the accuracy paradox. *PloS one*, 9(1): e84217.

Wang, H.; Chen, J.; Pan, T.; Dong, Z.; Jiang, R.; and Song, X. 2025. Evaluating the generalization ability of spatiotemporal model in urban scenario. *IEEE Transactions on Mobile Computing*.

Wang, Y.; Liu, Z.; Wang, D.; Li, Y.; and Yan, J. 2017. Anomaly detection and visual perception for landslide monitoring based on a heterogeneous sensor network. *IEEE Sensors Journal*, 17(13): 4248–4257.

Xie, P.; Zhou, A.; and Chai, B. 2019. The application of long short-term memory (LSTM) method on displacement prediction of multifactor-induced landslides. *IEEE Access*, 7: 54305–54311.

Yue, X.; Dai, F.; Ke, Z.; Zhu, J.; and Cheng, W. 2025. An improved numerical model for landslide-induced waves and its application to the Huangtian landslide in the XW reservoir, China. *Scientific Reports*, 15(1): 27366.

Zhang, Y.-g.; Tang, J.; He, Z.-y.; Tan, J.; and Li, C. 2021. A novel displacement prediction method using gated recurrent unit model with time series analysis in the Erdaohe landslide. *Natural Hazards*, 105.

Zhao, Z.; Chen, T.; Dou, J.; Liu, G.; and Plaza, A. 2024. Landslide susceptibility mapping considering landslide local-global features based on CNN and transformer. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17: 7475–7489.

Zhou, K.; Liu, Z.; Qiao, Y.; Xiang, T.; and Loy, C. C. 2022. Domain generalization: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(4): 4396–4415.