

Knowledge-Informed Deep Learning for Irrigation Type Mapping from Remote Sensing

Oishee Bintey Hoque^{1,*}, Nibir Chandra Mandal^{1,*}, Abhijin Adiga², Samarth Swarup², Sayjro Kossi Nouwakpo³, Amanda Wilson², Madhav Marathe^{1,2}

¹Dept. of Computer Science, University of Virginia

²Biocomplexity Institute, University of Virginia

³US Department of Agriculture

{gza5dr, wyr6fx, abhijin, swarup, alw4ey, marathe}@virginia.edu,
Kossi.Nouwakpo@usda.gov

Abstract

Accurate mapping of irrigation methods is crucial for sustainable agricultural practices and food systems. However, existing models that rely solely on spectral features from satellite imagery are ineffective due to the complexity of agricultural landscapes and limited training data, making this a challenging problem. We present Knowledge-Informed Irrigation Mapping (KIIM), a novel Swin-Transformer based approach that uses (i) a specialized projection matrix to encode crop to irrigation probability, (ii) a spatial attention map to identify agricultural lands from non-agricultural lands, (iii) bi-directional cross-attention to focus complementary information from different modalities, and (iv) a weighted ensemble for combining predictions from images and crop information. Our experimentation on five states in the US shows up to 22.9% (IoU) improvement over baseline with a 71.4% (IoU) improvement for hard-to-classify drip irrigation. In addition, we propose a two-phase transfer learning approach to enhance cross-state irrigation mapping, achieving a 51% IoU boost in a state with limited labeled data. The ability to achieve baseline performance with only 40% of the training data highlights its efficiency, reducing the dependency on extensive manual labeling efforts and making large-scale, automated irrigation mapping more feasible and cost-effective. Code: <https://github.com/Nibir088/KIIM>

1 Introduction

Mapping Irrigation Assets: A Social Good Problem.

Irrigation is a crucial component of agricultural management, supporting approximately 40% of global food production [WWAP, 2019]. As a dominant freshwater-use practice, irrigation accounts for nearly 90% of global consumptive freshwater use [Döll *et al.*, 2009; Meier *et al.*, 2018; Dieter *et al.*, 2018; Zhou *et al.*, 2020], significantly shaping regional and global hydrological cycles [de Vrese *et al.*, 2016; Leng *et al.*, 2014]. In regions like northwestern China and




Figure 1: Irrigation mapping from satellite imagery: (left) RGB image showing citrus, alfalfa, and jojoba fields (red borders); (right) irrigation classification mask with Sprinkler, Flood, and Drip irrigation methods.

the US High Plains, excessive irrigation has caused substantial declines in river discharge and groundwater levels, highlighting the impact of inefficiency of irrigation methods [Hao *et al.*, 2015; Pérez-Blanco *et al.*, 2020]. The irrigation type (e.g., drip, sprinkler, or flood irrigation) determines how extracted water is distributed across irrigated areas, affecting water quantity and quality [Ippolito *et al.*, 2019]. While accurate mapping of irrigation methods can facilitate identification of current practices and sustainable upgrades, traditional approaches only distinguish irrigated from non-irrigated fields or focus only on small areas or a single type of irrigation; large-scale cross-region generalization remains underexplored [Tang *et al.*, 2021; Nouwakpo *et al.*, 2024; Hoque *et al.*, 2024]. Thus, effective irrigation mapping aligns with the United Nations Sustainable Development Goals 2 and 8 [Nations, 2015], which aim to promote sustainable agricultural practices and food systems, while also supporting the goal of “Leave No One Behind” [WWAP, 2019].

Team. This work is an interdisciplinary collaboration between computer scientists (at universities) and an agricultural scientist (at the US Department of Agriculture).

Challenges. Due to differences in climate, crop types, and water availability, irrigation practices can vary significantly from one region to another [Nie *et al.*, 2021]. As a result, traditional methods that perform well locally may struggle to accurately classify irrigation across multiple regions, making large-scale, consistent irrigation mapping an ongoing challenge. Many regions lack sufficient labeled data for

training robust irrigation classification models, as collecting ground truth across vast farmland areas is costly and time-consuming. Remote sensing offers a scalable solution using multispectral satellite imagery, but challenges persist due to spectral ambiguity, resolution limits, class imbalance, and regional differences. Furthermore, available data are often highly imbalanced, with drip-irrigated fields constituting only a small fraction of samples, making it difficult for the model to learn minority classes effectively. For instance, drip irrigation accounts for just **0.08%** of Utah’s irrigated land¹.

Our Contributions. To address these challenges, we propose the **Knowledge-Informed Irrigation Mapping (KIIM)** model (Fig. 2), which leverages the Normalized Difference Vegetation Index (NDVI), the Normalized Difference Water Index (NDWI), and the Normalized Difference Tillage Index (NDTI) derived from additional remote sensing bands (details in Appendix) to capture plant health, water content, and soil conditions. KIIM also uses land-use data and crop-type information, from the USDA Cropland Data Layer, to refine predictions by focusing on cultivated areas and incorporating crop-irrigation relationships, improving the identification of underrepresented irrigation methods like drip irrigation. Our main contributions are:

- We propose a vision transformer-based multi-stream learning framework that integrates RGB and vegetation indices using a *Bidirectional Cross-Attention* module (Fig. 2). Vegetation indices are closely related to irrigation methods; for example, NDWI identifies flood-irrigated fields better, while NDVI distinguishes healthy crop growth associated with sprinkler or drip irrigation [Dempsey, 2024; Allen *et al.*, 2021]. While straightforward single-stream channel stacking may lead to the loss of modality-specific information due to early fusion, our framework instead guides the model to capture complementary relationships allowing each stream to query relevant information from the other stream’s perspective.
- We encode crop-irrigation relationships in a state-specific projection matrix, capturing historical irrigation preferences, and use weighted ensemble method with the prediction from the multi-stream module. Crop-type information further refines predictions, as certain crops are historically associated with specific irrigation methods (e.g., vineyards with drip irrigation, alfalfa with flood irrigation; see Fig 1).
- We incorporate a spatial attention map to enhance agricultural land segmentation by generating pixel-level attention map. This assigns higher weights to agricultural regions and field boundaries while suppressing non-agricultural areas, effectively highlighting irrigation-relevant features.
- Our extensive evaluations across five states (Arizona (AZ), Colorado (CO), Utah (UT), Washington (WA), and Florida (FL)) show that each module in KIIM improves performance individually, with the best results achieved when all modules are combined. KIIM consistently outperforms the baseline, achieving an average IoU improvement of 18.1% across states and a 71.4% improvement in challenging cases (i.e., drip). KIIM demonstrates strong generaliza-

tion to unseen state data, achieving impressive performance in both zero-shot and few-shot settings. Impressively, KIIM achieves this performance using Landsat’s 30m resolution, demonstrating its ability to learn irrigation patterns despite coarse spatial granularity.

2 Related Work

Remote Sensing for Mapping Agricultural Infrastructure. Deep learning models in agricultural remote sensing have been applied to crop classification, field boundary detection, and irrigation mapping [Jin *et al.*, 2017; Weiss *et al.*, 2020; Paolini *et al.*, 2022]. CNNs have been used for classifying invasive species [Hung *et al.*, 2014], segmenting mixed crops [Mortensen *et al.*, 2016], and detecting weeds [Milioto *et al.*, 2017; Di Cicco *et al.*, 2017]. Attention-based models [Wang *et al.*, 2020a; Zheng *et al.*, 2021] and channel-wise feature selection techniques [Cheng *et al.*, 2021; Tao *et al.*, 2020] improve segmentation accuracy in complex landscapes. However, CNNs and ViTs struggle to differentiate irrigation methods due to spectral similarity. Multi-stream fusion of spectral indices (NDVI, NDWI, NDTI) enhances segmentation robustness but still demands large labeled datasets for reliable generalization [He *et al.*, 2017; Touvron *et al.*, 2021]. While progress has been made in other agricultural infrastructure, advancements in irrigation mapping remain limited.

Multi-Channel Representation Learning. Recent advances in deep learning have significantly improved semantic segmentation, evolving from FCNs to encoder-decoder models like U-Net, LinkNet, and DeepLabv3+ [Long *et al.*, 2015; Ronneberger *et al.*, 2015; Chen *et al.*, 2017; Chaurasia and Culurciello, 2017]. Multi-scale feature extraction is enhanced by FPN and PAN [Li *et al.*, 2018], while transformer-based models (ViTs, Swin-Transformers) improve global context learning [Dosovitskiy *et al.*, 2021; He *et al.*, 2022]. Attention mechanisms, such as CBAM [Woo *et al.*, 2018], refine feature representation, and multi-stream architectures leverage attention for modality fusion in remote sensing [Wang *et al.*, 2020b; Bastidas and Tang, 2019]. Unlike prior works that stack agricultural indices with RGB [Hoque *et al.*, 2024; Nouwakpo *et al.*, 2024], our approach employs bidirectional cross-attention within a multi-stream architecture to enable dynamic feature interaction between streams.

Domain-Aware Segmentation Models. Remote sensing in agriculture faces domain shifts due to variations in soil, cropping patterns, irrigation, and climate [Raei *et al.*, 2022; WANG, 2024; Chen *et al.*, 2024]. Transfer learning techniques, including feature adaptation and domain-aware fine-tuning, help align feature representations across regions [Zhuang *et al.*, 2020; Bosilj *et al.*, 2020; Coulibaly *et al.*, 2019]. Recent studies show that incorporating domain knowledge enhances generalization and reduces dependence on large labeled datasets [Shi *et al.*, 2021].

3 Knowledge-Informed Irrigation Mapping

3.1 Problem Formulation

We formulate irrigation mapping from satellite imagery as a semantic segmentation problem. Given a satellite image

¹<https://dwre-utahdnr.opendata.arcgis.com/pages/wrlu-data>

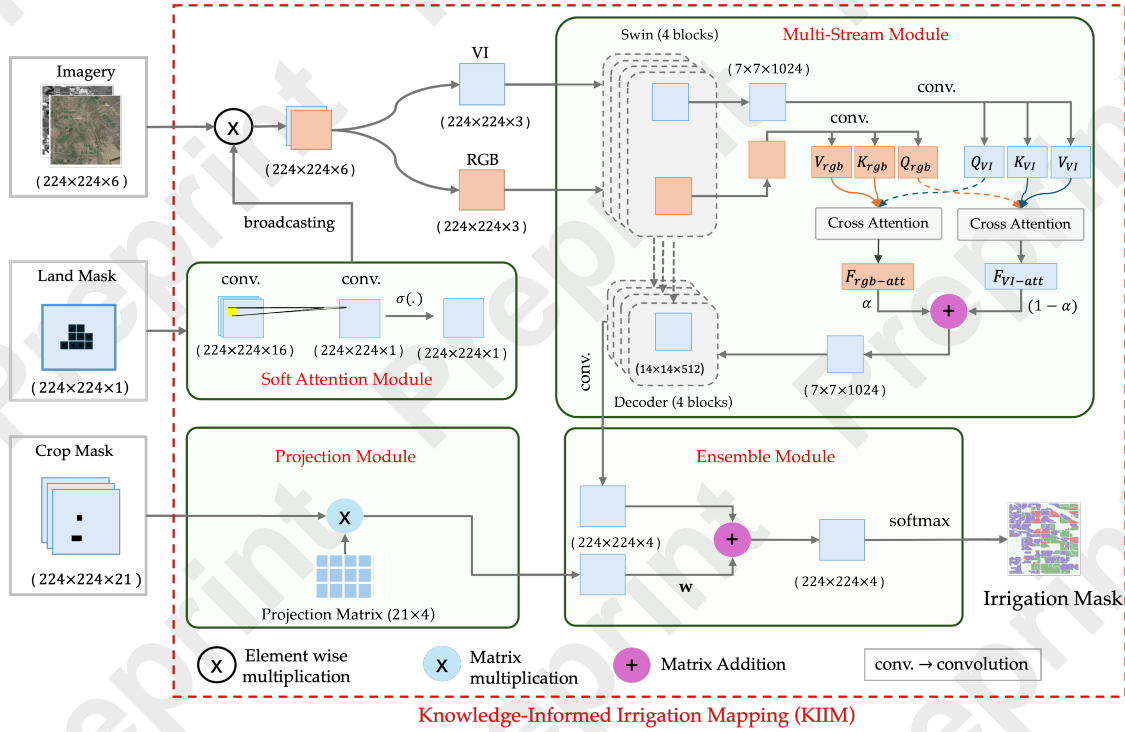


Figure 2: Overview of our Knowledge-Informed Irrigation Mapping model. (i) A soft attention module refines spatial focus by highlighting irrigated areas. (ii) The multi-stream module processes RGB and vegetation indices through a shared Swin Transformer and fuses features using bidirectional cross-attention, allowing adaptive feature interaction. (iii) A projection module incorporates domain knowledge by mapping crop types to irrigation probabilities using a predefined projection matrix. (iv) Finally, an ensemble module balances satellite-derived predictions with knowledge-informed irrigation likelihoods, optimizing weights through end-to-end training.

$X \in \mathbb{R}^{H \times W \times C}$, crop mask $M \in \{0, 1\}^{H \times W \times G}$, and land mask $L \in \{0, 1\}^{H \times W}$, where H, W, C , and G denote height, width, spectral bands, and crop groups, respectively, our goal is to classify each pixel into one of K irrigation methods.

Let $Y \in \mathcal{Y}^{H \times W}$ be the ground truth irrigation labels, where $Y_{i,j}$ represents an irrigation type for pixel (i, j) . We aim to learn a mapping function:

$$f_{\theta} : \mathbb{R}^{H \times W \times C} \times \{0, 1\}^{H \times W \times G} \times \{0, 1\}^{H \times W} \rightarrow [0, 1]^{H \times W \times K}$$

parameterized by θ , that outputs a probability distribution over irrigation methods. The optimal parameters θ^* are obtained by minimizing a loss function \mathcal{L} :

$$\theta^* = \arg \min_{\theta} \mathcal{L}(f_{\theta}(X, M, L), Y) \quad (1)$$

3.2 Methodology

Our knowledge-informed irrigation mapping (KIIM) model addresses the key challenges in irrigation mapping through a specialized architecture that combines satellite imagery with crop knowledge. The model takes three inputs to capture different aspects of agricultural landscapes: (i) Land mask ($224 \times 224 \times 1$), a binary mask defining agricultural boundaries; (ii) Satellite bands ($224 \times 224 \times B$), multi-spectral imagery providing spectral reflectance across B wavelength bands; and (iii) Crop mask ($224 \times 224 \times 21$), a one-hot encoded spatial representation where each channel corresponds

to a specific crop type. These inputs are processed through specialized modules designed to maximize irrigation mapping accuracy. Our architecture comprises four main modules: (i) soft attention module, (ii) multi-stream module, (iii) projection module, and (iv) ensemble module. Figure 2 illustrates the complete architecture.

- **Soft Attention Module** applies sequential 3×3 convolutions followed by sigmoid activation on the land mask to generate a spatial attention map that highlights irrigated areas. In contrast to hard attention, it preserves spatial continuity, refining boundaries and preventing abrupt transitions between irrigated and non-irrigated areas. By adaptively scaling pixel-wise attention weights, it filters out background noise while increasing the contrast between irrigated and non-irrigated regions.
- **Multi-Stream Module** processes two complementary input streams through a shared Swin Transformer ϕ : RGB data ($X_{\text{RGB}} \in \mathbb{R}^{H \times W \times 3}$) and derived vegetation indices (NDVI, NDTI, and NDWI) ($X_{\text{VI}} \in \mathbb{R}^{H \times W \times 3}$). The extracted feature maps $F_{\text{RGB}}, F_{\text{VI}}$ preserve modality-specific features while enhancing fusion and reducing overfitting compared to a single-stream model. Instead of naïve concatenation or averaging for merging these feature maps before it goes to the decoder, we introduce **bidirectional cross-attention**, allowing RGB and vegetation indices to exchange information and dynamically prioritize relevant

features across streams, leading to more context-aware feature fusion. Unlike spatial and self-attention, bidirectional cross-attention enables explicit feature interaction between RGB and spectral streams, allowing each stream to dynamically query and integrate complementary information from the other stream [Tan and Bansal, 2019], rather than just focusing on spatial relationships within individual feature maps or single feature spaces. To achieve this, we first transform RGB and VI features into queries (Q), keys (K), and values (V) using 1×1 convolutions for both streams. The RGB stream is projected into $Q_{\text{rgb}}, K_{\text{rgb}}, V_{\text{rgb}}$, while the VI stream is mapped to $Q_{\text{vi}}, K_{\text{vi}}, V_{\text{vi}}$ with analogous definitions for $Q_{\text{aux}}, K_{\text{rgb}}, V_{\text{rgb}}$. Next, we calculate the standard scaled dot-product attention operation as done by [Vaswani *et al.*, 2017]. For brevity, consider a single query–key–value triplet $(Q, K, V) \in \mathbb{R}^{B \times C \times H \times W}$ to compute $\text{Attention}(Q, K, V) = \text{softmax}(\frac{QK^T}{\sqrt{C}})V$. In our two-stream setting, RGB attends to vegetation indices as $F_{\text{VI-att}} = \text{Attention}(Q_{\text{rgb}}, K_{\text{vi}}, V_{\text{vi}})$, while vegetation indices attend to RGB as $F_{\text{rgb-att}} = \text{Attention}(Q_{\text{vi}}, K_{\text{rgb}}, V_{\text{rgb}})$. We then combine these attended feature maps with a learnable fusion parameter $\alpha \in \mathbb{R}$ (initialized to 0.8), resulting in the final fused representation $F_{\text{fused}} = \alpha F_{\text{rgb-att}} + (1 - \alpha) F_{\text{VI-att}}$. Lastly, we employ a U-Net-style decoder [Ronneberger *et al.*, 2015] with skip connections to reconstruct pixel-level predictions. A final 1×1 convolution generates logits for segmentation into four irrigation classes.

• **Projection Module** maps domain knowledge (i.e., crop information) to irrigation probabilities using a predefined projection matrix $P \in [0, 1]^{(G \times K)}$. The projection matrix transforms crop information into irrigation method probabilities based on regional farming practices. Let $G = \{g_1, \dots, g_n\}$ be the set of n crop groups and $\mathcal{Y} = \{1, \dots, K\}$ the set of K irrigation methods. The total area of crop group $g \in G$ under irrigation type $i \in \mathcal{Y}$ is denoted by $A_{g,i}$. The probability $P_{g,i}$ of irrigation type i for crop group g is given by:

$$P_{g,i} = \begin{cases} \frac{A_{g,i}}{\sum_{j \in \mathcal{Y}} A_{g,j}}, & \text{if } \sum_{j \in \mathcal{Y}} A_{g,j} > 0 \\ \frac{1}{K}, & \text{otherwise} \end{cases} \quad (2)$$

where $\sum_{i \in \mathcal{Y}} P_{g,i} = 1$ for each $g \in G$. The resulting $P \in [0, 1]^{n \times K}$ matrix is structured as:

$$P = \begin{bmatrix} P_{1,1} & \dots & P_{1,K} \\ \vdots & \vdots & \vdots \\ P_{n,1} & \dots & P_{n,K} \end{bmatrix} \quad (3)$$

where each row represents a crop group’s probability distribution across irrigation methods. This predefined matrix offers three key advantages: (i) it incorporates comprehensive historical knowledge about crop-irrigation relationships that cannot be learned from limited satellite imagery; (ii) it provides reliable predictions even when irrigation infrastructure is not visible in the imagery; and (iii) it enables easy updates of irrigation statistics without retraining the model.

• **Ensemble Module** combines data-driven predictions from Multi-Stream Module and knowledge-informed irrigation likelihoods using learnable weights w . Given two inputs ($224 \times 224 \times K$), the module applies weighted summation to aggregate predictions, followed by softmax activation to normalize the final probability distribution. The weights are optimized jointly with the model parameters through end-to-end training. This enables the model to adaptively integrate spectral features with domain knowledge for more informed predictions.

3.3 Loss Function

We use a composite loss \mathcal{L} integrating cross-entropy loss (\mathcal{L}_c) [Jadon, 2020] for per-pixel classification and constrained Dice loss (\mathcal{L}_d) [Milletari *et al.*, 2016] to enforce spatial consistency:

$$\mathcal{L} = \alpha \mathcal{L}_c + (1 - \alpha) \mathcal{L}_d, \quad (4)$$

where $\alpha \in [0, 1]$ balances the two components.

For a given irrigation label Y and prediction \bar{Y} , **cross-entropy loss** (\mathcal{L}_c) ensures pixel-wise class separation:

$$\mathcal{L}_c(Y, \bar{Y}) = -\frac{1}{H \times W} \sum_{k=1}^K \sum_{i,j} Y_{i,j,k} \log(\bar{Y}_{i,j,k}). \quad (5)$$

Land-masked dice loss \mathcal{L}_d enhances segmentation quality by constraining predictions to agricultural land.

$$\mathcal{L}_d(Y, \bar{Y}) = 1 - \frac{1}{K} \sum_{k=1}^K \frac{2 \sum_{i,j} \bar{Y}_{i,j,k} \cdot Y_{i,j,k} \cdot L_{i,j}}{\sum_{i,j} L_{i,j} (Y_{i,j,k} + \bar{Y}_{i,j,k})}, \quad (6)$$

where L denotes the landmask. Land-masked dice loss enhances spatial coherence by considering neighboring pixels. In addition, it adjusts for class imbalance for underrepresented irrigation methods, such as drip irrigation, by normalizing over the sum of predictions and ground truth. Moreover, it constrains optimization to agricultural regions to prioritize learning irrigation patterns within farmland.

4 Experimentation and Results

4.1 Dataset

Our study integrates multi-source geospatial datasets to construct an irrigation and crop mapping framework across five U.S. states, using the Utah Water-Related Land Use (WRLU) Dataset (2023) for Utah², the USGS Verified Irrigated Agricultural Lands Dataset (2002–2017) for Arizona and Florida³, the Washington State Department of Agriculture Agricultural Land Use dataset for Washington⁴, and the Colorado Division of Water Resources GIS dataset for Colorado⁵. In addition,

²<https://dwre-utahdnr.opendata.arcgis.com/pages/wrlu-data>

³<https://catalog.data.gov/dataset/verified-irrigated-agricultural-lands-for-the-united-states-200217>

⁴<https://agr.wa.gov/departments/land-and-water/natural-resources/agricultural-land-use>

⁵<https://dwr.colorado.gov/services/data-information/gis>

Model	AZ				UT				WA				CO			
	MIoU	Drip			MIoU	Drip			MIoU	Drip			MIoU	Drip		
		P	R	IoU		P	R	IoU		P	R	IoU		P	R	IoU
ResNet50	0.880	0.916	0.922	0.850	0.542	0.625	0.117	0.110	0.505	0.616	0.485	0.372	0.740	0.866	0.661	0.600
LinkNet	0.878	0.914	0.925	0.851	0.456	0.755	0.035	0.035	0.539	0.612	0.583	0.425	0.733	0.748	0.771	0.612
PAN	0.856	0.907	0.897	0.821	0.561	0.590	0.394	0.309	0.550	0.650	0.550	0.425	0.699	0.743	0.617	0.508
FPN	0.836	0.897	0.871	0.792	0.556	0.625	0.311	0.262	0.566	0.718	0.535	0.442	0.722	0.782	0.684	0.575
DeepLabV3+	0.873	0.913	0.915	0.841	0.564	0.448	0.316	0.228	0.606	0.700	0.646	0.506	0.747	0.796	0.712	0.602
ViT	0.751	0.873	0.831	0.742	0.467	0.427	0.178	0.144	0.478	0.600	0.398	0.314	0.626	0.782	0.491	0.432
FarSeg	0.878	0.930	0.903	0.846	0.575	0.575	0.374	0.293	0.554	0.682	0.556	0.441	0.769	0.850	0.742	0.656
SegFormer	0.867	0.918	0.901	0.833	0.577	0.580	0.425	0.325	0.558	0.720	0.528	0.438	0.727	0.833	0.651	0.576
Swin	0.896	0.942	0.913	0.865	0.640	0.472	0.584	0.353	0.645	0.752	0.663	0.544	0.786	0.810	0.797	0.671
KIIM	0.988	0.988	0.993	0.982	0.791	0.664	0.873	0.605	0.770	0.796	0.776	0.647	0.931	0.820	0.967	0.798

Table 1: Macro-IoU (averaged over Flood, Sprinkler, Drip, and non-irrigated land) and Drip-only performance (without Dice) for each model in AZ, UT, WA, and CO.

we map various irrigation practices to three primary methods (drip, sprinkler, and flood), as the original datasets contained multiple irrigation subtypes that were unified for consistency in this study. Crop data are derived from these sources, consolidating 143 distinct crop types into 20 standardized categories (details in Appendix). In addition, we collect Landsat-8 satellite imagery to generate irrigation masks with a spatial resolution of 30 meters.

Following standard remote sensing procedures, we segment satellite images into non-overlapping patches of size 224×224 pixels. The crop mask is created by assigning each pixel to a crop group based on the mapped crop type, while the land mask is derived by categorizing pixels into agricultural and nonagricultural lands. The final dataset comprises 36738 image patches, including Arizona (7154 patches), Utah (6062 patches), Washington (3557 patches), Florida (1230 patches) and Colorado (18735 patches). Among the state datasets, Utah and Florida have the lowest percentage (1.8% and 0.08%) of patches with drip irrigation. We discuss further dataset collections, dataset details, preprocessing steps, and projection matrix formulation in Appendix.

4.2 Evaluation Metrics

We evaluate the irrigation mapping task using four standard segmentation metrics: Intersection over Union (IoU), Precision (P), Recall (R), and Dice Score (D). Let Y and \bar{Y} be the ground truth and predicted masks, respectively, for an image of size $H \times W$, where each pixel (i, j) is assigned a class $k \in \mathcal{Y}$. The ground truth and predicted pixel sets for class k are defined as:

$$T_k = \{(i, j) \mid Y_{i,j} = k\}, \quad M_k = \{(i, j) \mid \bar{Y}_{i,j} = k\}. \quad (7)$$

The evaluation metrics precision, recall, Dice, and IoU are computed as:

$$P_k = \frac{|M_k \cap T_k|}{|M_k|}, \quad R_k = \frac{|M_k \cap T_k|}{|T_k|}, \quad (8)$$

$$D_k = \frac{2 \times P_k \times R_k}{P_k + R_k}, \quad \text{IoU}_k = \frac{|M_k \cap T_k|}{|M_k \cup T_k|}, \quad (9)$$

Precision measures the proportion of correctly predicted irrigated pixels among all predicted as class k , while recall

quantifies the fraction of correctly identified irrigated pixels out of all actual class k pixels. On the contrary, Dice Score computes the harmonic mean of precision and recall. IoU is defined as the ratio of intersection to union, provides a more balanced spatial evaluation by penalizing both over-segmentation (false positives) and under-segmentation (false negatives).

4.3 Experimentation Setting

We split each state’s dataset into 85% training and 15% testing, except for Florida, where we used a 50%-50% split due to limited drip irrigation samples. Model implementation was conducted using PyTorch and executed on NVIDIA A40 GPU. We performed 5-fold cross-validation and optimized hyperparameters through grid search over learning rates $\{1e-4, 2e-4, 5e-4\}$, batch sizes $\{16, 32, 64\}$, and loss weight α values $\{0, 0.4, 0.5, 0.6, 1\}$. The optimal configuration was selected based on IoU performance on the validation set.

4.4 Effectiveness of KIIM Model

We first evaluate KIIM’s ability to classify irrigation methods, particularly drip irrigation, and compare it against nine state-of-the-art segmentation models, including transformer-based architectures (Swin, ViT, SegFormer) and the remote sensing-specific model FarSeg. We train each model on state-specific irrigation datasets and evaluate on the corresponding test dataset across Arizona (AZ), Utah (UT), Washington (WA) and Colorado (CO). In Table 1, we show the performance of our KIIM model and state-of-the-art models. Across all states, state-of-the-art models struggle to identify drip-irrigated lands due to extreme class imbalance. The best baseline (Swin) achieves only 0.353 and 0.544 IoU in Utah and Washington which highlights the challenges of identifying drip irrigation. Due to its sparse presence and high spectral similarity with other irrigated lands, state-of-the-art models fail to identify drip irrigated lands. In contrast, the higher prevalence of drip irrigation in Arizona makes it easier for models to learn its spatial patterns and distinguish it from other irrigation methods.

We notice that the KIIM model effectively captures spatial dependencies and distinguishes underrepresented irrigation methods, which leads to significantly improved segmentation




Figure 3: Visual comparison of our model (KIIM) against top-performing baselines (Swin and DeepLabV3+). KIIM accurately segments farmland and correctly classifies irrigation methods, whereas baseline models struggle to delineate agricultural farmland and misclassify irrigation methods. The higher MIoU scores underscores KIIM model’s effectiveness in detecting and identifying different irrigation methods.

accuracy. Specifically, KIIM outperforms the best baseline (Swin) in macro-IoU by 10.3% in AZ (0.988 vs. 0.896), 19.6% in CO (0.931 vs. 0.778), 22.9% in UT (0.791 vs. 0.644), and 19.4% in WA (0.770 vs. 0.645). Moreover, KIIM achieves a 13.5% improvement in AZ (0.982 vs. 0.865), 18.9% in CO (0.798 vs. 0.671), 71.4% in UT (0.605 vs. 0.353), and 19.0% in WA (0.647 vs. 0.544). The most significant gain is observed in Utah, where baseline models struggle due to severe class imbalance, but KIIM improves drip IoU from 0.353 (Swin) to 0.605 (a 71.4% improvement). In Figure 3, we present an example of KIIM model’s predictions, demonstrating the model’s ability to accurately classify irrigation methods. These results highlight KIIM’s robustness in handling extreme class imbalances and its effectiveness in different geographical regions.

4.5 Transfer Learning for Cross-State Irrigation Mapping

Mapping irrigation methods across different states is challenging due to limited labeled data, extreme class imbalance, and substantial regional variations in irrigation practices. Therefore, training separate state-specific models is often impractical, as some irrigation methods (e.g., drip irrigation) are severely underrepresented in many states. Moreover, in some cases, sufficient training data is entirely unavailable. For example, Florida has only 1230 training samples, with only 11 samples containing drip irrigation, which makes it impossible to train and test a reliable state-specific model. To overcome this, we leverage transfer learning that enables a model to learn common irrigation features from a diverse multi-state dataset and adapt them to specific states with minimal labeled data.

In this work, we create a multi-state training dataset which maintains enough drip irrigated samples in the training data. Following [Buda *et al.*, 2018], we construct the multi-state training dataset (8880 samples where 4440 samples have drip irrigation) where the total sample size is twice the number of drip-irrigated samples, maintaining an imbalance ratio of (2 : 1). The training samples were selected from the Arizona, Utah, Colorado, and Washington training data.

To train our model, we follow a two-phase learning approach (training and state-adaptive fine-tuning). In the **training step**, we train KIIM on the multi-state dataset to learn

State Model		Flood		Sprinkler		Drip	
		Dice	IoU	Dice	IoU	Dice	IoU
AZ	KIIM	0.991	0.983	0.995	0.989	0.991	0.982
	w/o FT	0.982	0.964	0.993	0.986	0.985	0.971
	w FT	0.986	0.973	0.994	0.988	0.988	0.976
UT	KIIM	0.876	0.780	0.893	0.806	0.754	0.605
	w/o FT	0.722	0.565	0.766	0.620	0.613	0.442
	w FT	0.884	0.792	0.900	0.819	0.835	0.717
WA	KIIM	0.759	0.611	0.926	0.863	0.786	0.647
	w/o FT	0.702	0.541	0.928	0.866	0.780	0.640
	w FT	0.820	0.695	0.945	0.896	0.838	0.722
CO	KIIM	0.983	0.966	0.981	0.963	0.887	0.798
	w/o FT	0.950	0.904	0.932	0.873	0.852	0.743
	w FT	0.984	0.969	0.980	0.961	0.937	0.882

Table 2: Performance of KIIM model for state-wise training (denoted as KIIM), cross-state training without state-adaptive fine-tuning (denoted as w/o FT), and cross-state training with state-adaptive fine-tuning (denoted as w FT) for Flood, Sprinkler, and Drip (Dice and IoU).

universal irrigation patterns (e.g., circular sprinkler layouts) that are consistent across states. In the **state-adaptive fine-tuning** step, we initialize state-specific models with pre-trained weights (from the training step) and fine-tune them using state-specific data to adapt to regional farm sizes, irrigation preferences, and class distributions while keeping the architecture unchanged. This hierarchical learning strategy enables KIIM to generalize across states while capturing local irrigation nuances and improving segmentation performance with minimal state-level labeled data.

Effectiveness of cross-state transfer learning: To assess the effectiveness of cross-state knowledge transfer, we compare two training strategies for KIIM: (i) a two-step transfer learning approach, where the model is first trained on a balanced multi-state dataset and then fine-tuned on state-specific data, and (ii) state-specific training only (KIIM). Additionally, we evaluate the model without the fine-tuning setting (zero-shot), where the model is pretrained on the multi-state dataset but not fine-tuned on the target state. The results in Table 2 demonstrate that transfer learning effectively enhances both majority and minority class segmentation. Performance on the majority class (sprinkler) remains consistent across approaches (AZ: 0.989, UT: 0.819, WA: 0.863, and CO: 0.961 IoU). However, minority class performance improves significantly (particularly for drip irrigation), where IoU increases from 0.605 to 0.717 in Utah (18.5%), and from 0.647 to 0.722 in Washington (11.5%). Similarly, in Washington, flood IoU improves from 0.611 to 0.695 (13.7%). This indicates that transfer learning effectively captures underrepresented irrigation patterns. It is noteworthy that the model without state-specific adaptations achieves similar performance as compared to the baseline (state-wise training only) for the majority class (e.g., sprinkler). These results validate that cross-state transfer learning retains general irrigation knowledge while adapting to state-specific variations.

Generalization irrigation mapping across states: The




Figure 4: KIIM model performance for the state of Florida for different training approaches. ‘FL-Only Training’ indicates that KIIM is trained solely on the Florida training dataset, whereas ‘Base + X% data’ indicates that KIIM is trained on a combined dataset (without Florida) and fine-tuned on X% of the Florida training dataset. Notably, for drip irrigation, KIIM attains a 0.678 IoU score with cross-state transfer learning which indicates 51% improvement over state-wise training.

Florida data set is very sparse and has severe class imbalance which makes state-only training ineffective for irrigation mapping. To address this, we implement cross-state transfer learning, where KIIM is pretrained on a multi-state irrigation dataset and fine-tuned with incremental portions of Florida training data (30%, 40%, 60%, 80%, and 100%). This approach allows the model to leverage universal irrigation features learned from diverse states (e.g., AZ, WA, CO, and UT) and improve the mapping task even with small amounts of labeled data. Figure 4 shows that our model achieves IoU of 0.56 and 0.86 for flood and sprinkler irrigation without seeing any data from Florida. This demonstrates that irrigation structures exhibit transferable patterns across regions. Moreover, tuning the model with Florida data, the performance improves significantly. For drip irrigation (minority class), the model achieves IoU 0.678 when tuned on all the training data. This underscores the necessity of state-specific fine-tuning to capture local irrigation practices. In comparison to the baseline (state-wise training), the IoU improves from 0.447 to 0.678 (51%) for drip irrigation. Notably, even with 40% fine-tuning, performance is comparable to full-state training. This suggests our model can be effectively used for irrigation mapping when very limited data are available. These findings show that cross-state learning enables the model to generalize irrigation patterns across regions with minimal labeled data.

5 Ablation Study

To evaluate the individual contribution of each architectural component and validate our design choices, we conduct an ablation study by systematically removing different modules from our KIIM model. Our results demonstrate that each module plays an important role in irrigation type mapping.

	AM	PM	LDL	MSM	Dice	IoU
✓	✓	✓	✓	cross	0.937	0.883
✓	✓	✓	✓	self	0.931	0.873
✗	✓	✓	✓	cross	0.850	0.747
✗	✗	✓	✓	cross	0.842	0.736
✗	✓	✓	✗	cross	0.853	0.753
✗	✗	✗	✗	cross	0.844	0.739
✗	✗	✗	✗	✗	0.826	0.712

Table 3: Performance of KIIM on validation data while varying different modules. AM indicates attention module, PM denotes projection module, LDL indicates land-masked dice loss, and MSM indicates multi-stream module technique. Checkmarks (✓) indicate inclusion of the respective module, while crosses (✗) indicate exclusion. Note that we report macro IoU, and macro Dice scores in the Table.

From Table 3, we show that the model, without any specialized modules, achieves a macro IoU score of 0.712, while our complete architecture incorporating all components reaches 0.883. Similarly, the exclusion of any module reduces the macro IoU score. This highlights the importance of land mask information, crop information, land-masked Dice loss, and multistream module for identifying irrigation patterns.

6 Discussion

While accurate irrigation mapping is crucial for identifying current practices and guiding sustainable upgrades, existing approaches rely on manual surveys (e.g., USGS data) or lack generalization across different regions and irrigation methods. To address this, we propose the Knowledge-Informed Irrigation Mapping (KIIM) model, a multi-stream framework that integrates RGB and agriculture-specific indices through a bidirectional attention module for enhanced feature fusion across different modalities of input streams. Additionally, KIIM incorporates land-use and crop data, enabling the model to focus on agriculture-specific pixels and leverage historical crop-irrigation relationships for improved classification. Our findings highlight the effectiveness of the proposed approach in improving irrigation mapping across multiple states, particularly in challenging cases like drip irrigation. The substantial performance gains over the baseline demonstrate the model’s ability to capture complex spatial patterns and stream-specific relationships. Therefore, KIIM enables timely identification of irrigation, which directly contributes to SDG 2 by supporting sustainable food production systems and resilient agricultural practices in water-stressed regions.

However, part of our model’s performance relies on historical crop-irrigation data, which may change over time. Also, crop-type (and land-use) masks may be erroneous or outdated in certain regions. Future work could explore more representative region-specific datasets and extending this framework to more diverse agricultural landscapes and refining it for even greater adaptability across varying irrigation practices.

Contribution Statement

Oishee Bintey Hoque and Nibir Chandra Mandal contributed equally to this research and share first authorship.

Acknowledgments

This material is based upon work supported by the AI Research Institutes program supported by NSF and USDA-NIFA under the AI Institute: Agricultural AI for Transforming Workforce and Decision Support (AgAID) award No. 2021-67021-35344. This work was partially supported by University of Virginia Strategic Investment Fund award number SIF160.

References

- [Allen *et al.*, 2021] L. Niel Allen, Alfonso Torres-Rua, Anastasia Thayer Hassett, Ryan Larsen, and Matt Yost. Utah agricultural water optimization - water savings from drip irrigation. Technical report, Utah State University, December 2021.
- [Bastidas and Tang, 2019] Alexei A. Bastidas and Hanlin Tang. Channel attention networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- [Bosilj *et al.*, 2020] Petra Bosilj, Erchan Aptoula, Tom Duckett, and Grzegorz Cielniak. Transfer learning between crop types for semantic segmentation of crops versus weeds in precision agriculture. *Journal of Field Robotics*, 37(1):7–19, 2020.
- [Buda *et al.*, 2018] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural networks*, 106:249–259, 2018.
- [Chaurasia and Culurciello, 2017] Abhishek Chaurasia and Eugenio Culurciello. Linknet: Exploiting encoder representations for efficient semantic segmentation. In *2017 IEEE Visual Communications and Image Processing (VCIP)*, pages 1–4. IEEE, 2017.
- [Chen *et al.*, 2017] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv:1706.05587*, 2017.
- [Chen *et al.*, 2024] Jie Chen, Dongyang Hou, Changxian He, Yaoting Liu, Ya Guo, and Bin Yang. Change detection with cross-domain remote sensing images: A systematic review. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024.
- [Cheng *et al.*, 2021] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *Advances in Neural Information Processing Systems*, volume 34, pages 17864–17875, 2021.
- [Coulibaly *et al.*, 2019] Solemane Coulibaly, Bernard Kamsu-Foguem, Dantouma Kamissoko, and Daouda Traore. Deep neural networks with transfer learning in millet crop images. *Computers in industry*, 108:115–120, 2019.
- [de Vrese *et al.*, 2016] Philipp de Vrese, Stefan Hagemann, and Martin Claussen. Asian irrigation, african rain: Remote impacts of irrigation. *Geophysical Research Letters*, 43(8):3737–3745, 2016.
- [Dempsey, 2024] Caitlin Dempsey. Normalized difference water index (ndwi) and its use in flood detection. *Geography Realm*, 2024. Accessed: February 7, 2024.
- [Di Cicco *et al.*, 2017] Maurilio Di Cicco, Ciro Potena, Giorgio Grisetti, and Alberto Pretto. Automatic model based dataset generation for fast and accurate crop and weeds detection. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5188–5195. IEEE, 2017.
- [Dieter *et al.*, 2018] C. A. Dieter, M. Maupin, R. Caldwell, M. Harris, T. Ivahnenko, J. Lovelace, et al. Estimated use of water in the united states in 2015. Report Circular 1441, U.S. Geological Survey, 2018.
- [Dosovitskiy *et al.*, 2021] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2021.
- [Döll *et al.*, 2009] P. Döll, K. Fiedler, and J. Zhang. Global-scale analysis of river flow alterations due to water withdrawals and reservoirs. *Hydrology and Earth System Sciences*, 13(12):2413–2432, 2009.
- [Gao, 1996] Bo-Cai Gao. Ndw— a normalized difference water index for remote sensing of vegetation liquid water from space. *Remote sensing of environment*, 58(3):257–266, 1996.
- [Gu *et al.*, 2007] Yingxin Gu, Jesslyn F Brown, James P Verdin, and Brian Wardlaw. A five-year analysis of modis ndvi and ndwi for grassland drought assessment over the central great plains of the united states. *Geophysical Research Letters*, 34(6), 2007.
- [Hao *et al.*, 2015] Zhenchun Hao, Sichun Chen, Zehua Li, Zhongbo Yu, Quanxi Shao, Fei Yuan, and Fangxin Shi. Quantitative assessment of the impacts of irrigation on surface water fluxes in the tarim river, china. *Hydrology Research*, 46(6):996–1007, 2015.
- [He *et al.*, 2017] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2961–2969, 2017.
- [He *et al.*, 2022] Xin He, Yong Zhou, Jiaqi Zhao, Di Zhang, Rui Yao, and Yong Xue. Swin transformer embedding unet for remote sensing image semantic segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–15, 2022.
- [Hoque *et al.*, 2024] Oishee Bintey Hoque, Samarth Swarup, Abhijin Adiga, Sayjro Kossi Nouwakpo, and Madhav Marathe. Irrnet: Advancing irrigation mapping with incremental patch size training on remote sensing imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5460–5469, 2024.
- [Hung *et al.*, 2014] Calvin Hung, Zhe Xu, and Salah Sukkarieh. Feature learning based approach for weed classification using high resolution aerial images from a digital

- camera mounted on a uav. *Remote Sensing*, 6(12):12037–12054, 2014.
- [Ippolito et al., 2019] James A. Ippolito, Dave L. Bjorneberg, Steve W. Blecker, and Michael S. Massey. Mechanisms responsible for soil phosphorus availability differences between sprinkler and furrow irrigation. *Journal of Environmental Quality*, 48(5):1370–1379, 2019.
- [Jadon, 2020] Shruti Jadon. A survey of loss functions for semantic segmentation. In *2020 IEEE conference on computational intelligence in bioinformatics and computational biology (CIBCB)*, pages 1–7. IEEE, 2020.
- [Jin et al., 2017] Zhenong Jin, George Azzari, Marshall Burke, Stephen Aston, and David B Lobell. Mapping smallholder yield heterogeneity at multiple scales in eastern africa. *Remote Sensing*, 9(9):931, 2017.
- [Kogan, 1995] Felix N Kogan. Application of vegetation index and brightness temperature for drought detection. *Advances in space research*, 15(11):91–100, 1995.
- [Leff et al., 2004] Billie Leff, Navin Ramankutty, and Jonathan A Foley. Geographic distribution of major crops across the world. *Global biogeochemical cycles*, 18(1), 2004.
- [Leng et al., 2014] Guoyong Leng, Maoyi Huang, Qihong Tang, Huilin Gao, and L Ruby Leung. Modeling the effects of groundwater-fed irrigation on terrestrial hydrology over the conterminous united states. *Journal of Hydrometeorology*, 15(3):957–972, 2014.
- [Li et al., 2018] Hanchao Li, Pengfei Xiong, Jie An, and Lingxue Wang. Pyramid attention network for semantic segmentation. In *British Machine Vision Conference (BMVC)*, 2018.
- [Long et al., 2015] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [McFeeters, 1996] Stuart K McFeeters. The use of the normalized difference water index (ndwi) in the delineation of open water features. *International journal of remote sensing*, 17(7):1425–1432, 1996.
- [Meier et al., 2018] J. Meier, F. Zabel, and W. Mauser. A global approach to estimate irrigated areas—a comparison between different data and statistics. *Hydrology and Earth System Sciences*, 22(2):1119–1133, 2018.
- [Milioto et al., 2017] Andres Milioto, Philipp Lottes, and Cyrill Stachniss. Real-time blob-wise sugar beets vs weeds classification for monitoring fields using convolutional neural networks. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4:41–48, 2017.
- [Milletari et al., 2016] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016.
- [Mortensen et al., 2016] Anders Krogh Mortensen, Mads Dyrmann, Henrik Karstoft, Rasmus Nyholm Jørgensen, and René Gislum. Semantic segmentation of mixed crops using deep convolutional neural network. 2016.
- [Nations, 2015] United Nations. Transforming our world: The 2030 agenda for sustainable development, 2015. Accessed: 2014-01-15.
- [Nie et al., 2021] Wanshu Nie, Benjamin F. Zaitchik, Matthew Rodell, Sujay V. Kumar, Kristi R. Arsenault, and Hamada S. Badr. Irrigation water demand sensitivity to climate variability across the contiguous united states. *Water Resources Research*, 57(3), 2021.
- [Nouwakpo et al., 2024] S. K. Nouwakpo, D. Bjorneberg, K. McGwire, and O. Hoque. Mapping irrigation methods in the northwestern us using deep learning classification. *Water Resources Research*, 60(8):e2023WR036155, 2024.
- [Paolini et al., 2022] Giovanni Paolini, Maria Jose Escorihuela, Olivier Merlin, Magí Pamies Sans, and Joaquim Bellvert. Classification of different irrigation systems at field scale using time-series of remote sensing data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:10055–10072, 2022.
- [Pérez-Blanco et al., 2020] C Dionisio Pérez-Blanco, Arthur Hraest-Essenfelder, and Chris Perry. Irrigation technology and water conservation: A review of the theory and evidence. *Review of Environmental Economics and Policy*, 2020.
- [Pettorelli et al., 2005] Nathalie Pettorelli, Jon Olav Vik, Atle Mysterud, Jean-Michel Gaillard, Compton J Tucker, and Nils Chr Stenseth. Using the satellite-derived ndvi to assess ecological responses to environmental change. *Trends in ecology & evolution*, 20(9):503–510, 2005.
- [Raei et al., 2022] Ehsan Raei, Ata Akbari Asanjan, Mohammad Reza Nikoo, Mojtaba Sadegh, Shokoufeh Pourshahabi, and Jan Franklin Adamowski. A deep learning image segmentation model for agricultural irrigation system classification. *Computers and Electronics in Agriculture*, 198:106977, 2022.
- [Ronneberger et al., 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [Shi et al., 2021] Chunlei Shi, Xianwei Xin, and Jiakai Zhang. Domain adaptation using a three-way decision improves the identification of autism patients from multisite fmri data. *Brain Sciences*, 11(5):603, 2021.
- [Tan and Bansal, 2019] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5100–5111. Association for Computational Linguistics, 2019.