

SIDE: Socially Informed Drought Estimation Toward Understanding Societal Impact Dynamics of Environmental Crisis

Lanyu Shang^{1,3}, Bozhang Chen¹, Shiwei Liu¹, Yang Zhang¹, Ruohan Zong¹, Anav Vora²
Ximing Cai², Na Wei², Dong Wang¹

¹School of Information Sciences, University of Illinois Urbana-Champaign, Champaign, IL, USA

²Civil and Environmental Engineering, University of Illinois Urbana-Champaign, Urbana, IL, USA

³Computer Science, Loyola Marymount University, Los Angeles, CA, USA

lanyu.shang@lmu.edu, {bozhang4, shiwei2, yzhangnd, rzong2, amvora3, xmcai, nawei2, dwang24}@illinois.edu

Abstract

Drought has become a critical global threat with significant societal impact. Existing drought monitoring solutions primarily focus on assessing drought severity using quantitative measurements, overlooking the diverse societal impact of drought from human-centric perspectives. Motivated by the collective intelligence on social media and the computational power of AI, this paper studies a novel problem of socially informed AI-driven drought estimation that aims to leverage social and news media information to jointly estimate drought severity and its societal impact. Two technical challenges exist: 1) How to model the implicit temporal dynamics of drought societal impact. 2) How to capture the social-physical interdependence between the physical drought condition and its societal impact. To address these challenges, we develop SIDE, a socially informed AI-driven drought estimation framework that explicitly quantifies the societal impact of drought and effectively models the social-physical interdependency for joint severity-impact estimation. Experiments on real-world datasets from California and Texas demonstrate SIDE's superior performance compared to state-of-the-art baselines in accurately estimating drought severity and its societal impact. SIDE offers valuable insights for developing human-centric drought mitigation strategies to foster sustainable and resilient communities.

Introduction

With the increasing frequency and severity, drought has become a critical global threat and poses significant challenges to our societies, leading to detrimental impacts on human well-being, environmental sustainability, and socio-economic development (Sugg et al. 2020). For example, drought has cost the agricultural sector of California an estimated \$1.7 billion in direct revenue losses and over 19,000 job losses in 2022 (Medellin-Azuara et al. 2022). Existing solutions have mainly focused on assessing drought severity with quantitative measurements, such as meteorological and hydrological records, and agricultural yield data (Dantas, da Silva, and Santos 2020; Becker and Sparks 2020). However, the underlying *societal impact* of drought (e.g., public health, ecosystem, recreation) has not been systematically explored, especially from human-centric perspectives. Motivated by the collective intelligence on social media in cap-

turing human observations and experiences and the computational power of AI, this paper focuses on a novel problem of socially informed AI-driven drought estimation. Our goal is to jointly estimate drought severity and identify salient determinants of the societal impact precipitated by drought.

A few recent efforts have been made to investigate the societal impact posed by drought (Sugg et al. 2020). These solutions have mainly focused on scientific measurements, such as meteorological and hydrological data, or economic losses in specific sectors like agriculture (Naumann et al. 2021). However, these approaches often overlook the diverse nature of drought's impact on human society, which often extends beyond quantifiable metrics and encompasses various aspects of the human community, such as public health, ecosystem, recreation, and tourism (Savelli et al. 2022). Moreover, traditional drought monitoring and assessment methods heavily rely on physical sensors and remote sensing technologies, which may not capture the diverse and localized perceptions of communities affected by drought (Savelli et al. 2022). The lack of human-centric perspectives in drought severity assessment can lead to an incomplete understanding of the societal impact and hinder the development of effective mitigation strategies that address the diverse needs and concerns of communities affected by drought.

To bridge this gap, in this paper, we study the problem of socially informed AI-driven drought estimation by leveraging the collective intelligence on social and news media to *jointly estimate the severity and societal impact of drought*. Social media platforms and online news outlets have become vital sources of timely, localized, and human-centric information during various natural disasters and crises (Bathaiy et al. 2021). Figure 1 shows an overview of the socially informed drought severity estimation problem. Our proposed framework integrates physical drought severity indicators, social media data, and news media content to provide a comprehensive and human-centric estimation of drought severity and its societal impact. Specifically, while social media inputs reflect the observations and experiences of individual users, news media data offers broader coverage and more neutral perspectives of drought conditions and their consequences analyzed by professional journalists and domain experts. However, two technical challenges exist in developing such a framework.

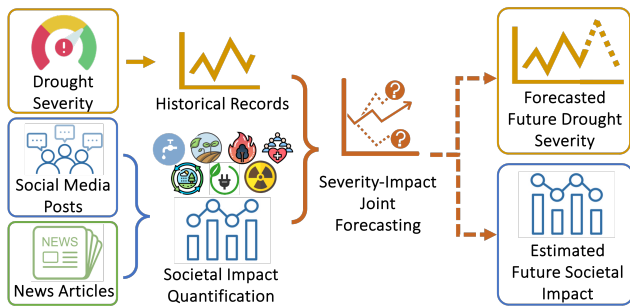


Figure 1: Socially Informed Drought Estimation Overview

Implicit Temporal Dynamics. The first challenge lies in the implicit temporal dynamic of drought’s societal impact. The physical conditions of drought dynamically change over time and can be measured with a set of observable indicators (e.g., precipitation, temperature, soil moisture). However, the societal impact of drought often encompasses a spectrum of factors (e.g., the income of farms and small businesses, drought resilience based on the mitigation strategies, physical and mental health of individuals) that are often dynamically changed and cannot be directly assessed. For example, farmers are more concerned about the drought’s societal impact on their income during growing seasons as compared to the off-season, when their attention is directed toward strategies for potential drought mitigation. More importantly, the dynamic changes of societal impact often cannot be directly observed or measured in the real world. Therefore, it remains a challenge to accurately capture the dynamic variation of the implicit societal impact to accurately estimate and explain drought’s impact on the societal dimension.

Social-Physical Interdependence. The second challenge arises from the complex interdependence between the physical condition and societal impact of drought. On one hand, the physical manifestation of drought (e.g., a period of below-normal rainfall or precipitation) can lead to reduced water availability, affecting both surface water and groundwater resources. As a consequence, such water scarcity directly affects agricultural productivity and water quality, resulting in significant societal impact, including reduced income, increased poverty, and food insecurity among these drought-impacted communities. On the other hand, the societal impact subsequently shapes community responses to drought. For example, communities lacking public awareness about drought risks and the importance of water conservation may overconsume water resources in an attempt to mitigate the impact of drought on agriculture, which may further intensify the severity of drought. Therefore, it is crucial to consider the interdependence between the physical condition and societal impact when developing a comprehensive drought estimation framework.

To address the above challenges, this paper develops SIDE, a socially informed AI-driven drought estimation framework that aims to jointly estimate the drought severity and the societal impact precipitated by drought. To address the first challenge, SIDE explicitly models the dynamics of

various determinants of drought’s societal impact by exploring drought-related social media discourse and news articles. To address the second challenge, SIDE designs a social-physical cross-attention mechanism to effectively capture the interdependence between physical drought conditions and societal responses to accurately estimate the severity and societal impact of drought. Evaluation results demonstrate that SIDE outperforms state-of-the-art baselines in accurately estimating the drought severity and the societal impact of drought crises. The proposed SIDE framework could be adopted for real-world deployment by integrating with the Dialogue and Information Platform for a Drought Event (DIP-Drought Platform) (DIP-Drought 2024) developed by our team to provide actionable insights for stakeholders and decision-makers. The platform will leverage SIDE’s capabilities to offer real-time drought severity and societal impact estimations, which will allow drought stakeholders (e.g., farmers, business owners, and policymakers) to access timely and comprehensive information about ongoing and potential drought crises.

Related Work

Societal Impact of Drought

Traditional approaches for estimating drought’s societal impact often rely on professionally collected data, such as hydrological/meteorological measurements (Dantas, da Silva, and Santos 2020), economic loss (Becker and Sparks 2020), surveys/interviews (Edwards, Gray, and Hunter 2019), and literature reviews (Lester, Flatau, and Kyron 2022). For example, Dantas, da Silva, and Santos (2020) utilized meteorological data (e.g., precipitation and temperature), to assess the severity and spatial extent of drought and link these physical indicators to potential societal impact, such as reduced agricultural productivity and increased food insecurity. Edwards, Gray, and Hunter (2019) conducted interviews with farmers to understand the social and emotional impacts (e.g., stress, anxiety, and mental health) of drought during prolonged dry periods. However, these traditional approaches often fall short of capturing the complex and dynamic nature of drought’s societal impact due to their narrow focus (e.g., only using physical indicators), limited scalability (e.g., small-scale interviews), and lack of timely information (e.g., previously published studies). To address these limitations, this paper designs a determinant-driven societal impact quantification approach that explicitly explores the timely social and news media data to quantify the societal impact of drought.

Time Series Forecasting

Time series forecasting has been widely studied in various domains, including weather forecasting, stock market prediction, and disease outbreak detection (Lim and Zohren 2021). Traditional time series forecasting methods, such as autoregressive integrated moving average (ARIMA) (Shumway et al. 2017) and exponential smoothing (Hyndman et al. 2008), have been extensively used for predicting future values based on historical patterns. However, these methods often struggle to capture complex non-

linear relationships and long-term dependencies. Recent advancements in deep learning techniques have shown promising results in capturing intricate patterns and dependencies in time series data. For example, Wu et al. (2022) proposed a time-aware graph neural network that incorporates temporal dynamics and graph structure for time series forecasting. Liu et al. (2023) introduced a hierarchical attention mechanism to capture both inter-series and intra-series dependencies for multivariate time series forecasting. However, most existing approaches focus on univariate or multivariate time series data from a single domain which are insufficient to address the socially informed drought estimation problem that integrates multiple heterogeneous data sources (i.e., drought severity, social and news media data). To overcome such a limitation, we propose a social-physical cross-attention mechanism that aims to capture the complex interactions and dependencies between the physical conditions and the societal impact determinants to jointly estimate the drought severity and its societal impact.

Problem Statement

Definition 1 Time Step t : the collection period (e.g., weekly) during which the meteorological, social media, and news media data are collected. In particular, we define $T \in \mathbb{Z}^+$ as the total number of time steps in our study and t is the t^{th} time step.

Definition 2 Drought Severity D_t : the intensity of drought conditions ($D_t \in \mathbb{R}$) in the studied area at time step t . In particular, we adopt the Drought Severity and Coverage Index (DSCI) (Smith et al. 2020) as our primary meteorological metric for evaluating the extent and intensity of drought conditions in the studied areas. DSCI is a comprehensive numerical measurement that measures both drought intensity and its impacted area, and has been widely adopted by leading drought monitor agencies (e.g., drought.gov and U.S. Drought Monitor).

Definition 3 Social Input S_t : a set of drought-related social media posts published during the time step t .

Definition 4 News Input \mathcal{N}_t : a set of drought-related news articles collected from leading news publishers during time step t .

Definition 5 Societal Impact M_t : the societal impact M_t of drought at time step t is the distribution of a set of δ determinants specified by the National Integrated Drought Information System (National Integrated Drought Information System), including *Agriculture, Ecosystems, Energy, Hazard Planning & Preparedness, Manufacturing, Navigation and Transportation, Public Health, Recreation and Tourism, Water Utilities, Wildfire Management, and Other*. Formally, let $M_t = [m_{t,1}, m_{t,2}, \dots, m_{t,\delta}]$ where $m_{t,i} \forall i \in \{1, 2, \dots, \delta\}$ represents the impact score for the i -th determinant at time step t . Each $m_{t,i}$ can be a real number between 0 and 1, where 0 indicates no impact and 1 indicates the most severe impact. We elaborate on the details of quantifying the societal impact from the social and news input in the Solution section.

Definition 6 Lookback Window T_L : the historical time period consisting of T_L , $1 \leq T_L \leq T - 1$, time steps within which the input data are used to predict future drought severity and societal impact. T_L is the lookback window size.

Definition 7 Prediction Window T_P : the future time period consisting of one or more time steps ($T_P \geq 1$) for which the model predicts drought severity and the societal impact. T_P is the prediction window size.

Using the above definition, we formulate the socially informed drought severity estimation problem as a time series regression problem. In particular, for a time step t , given the historical drought severity $D_{t-T_L:t} = [D_{t-T_L}, D_{t-T_L+1}, \dots, D_t]$, social input $S_{t-T_L:t} = [S_{t-T_L}, S_{t-T_L+1}, \dots, S_t]$, and news input $\mathcal{N}_{t-T_L:t} = [\mathcal{N}_{t-T_L}, \mathcal{N}_{t-T_L+1}, \dots, \mathcal{N}_t]$ during the past T_L time steps, our goal is to predict the future drought severity $D_{t+1:t+T_P} = [D_{t+1}, \dots, D_{t+T_P}]$ and societal impact $M_{t+1:t+T_P} = [M_{t+1}, \dots, M_{t+T_P}]$ for the next T_P time steps. Formally, our problem is formulated as

$$\theta^* = \arg \min_{\theta, \varphi} \left(\frac{1}{T_P} \sum_{i=1}^{T_P} f_{\theta}(\hat{D}_{t+i}, D_{t+i}) + g_{\varphi}(\hat{M}_{t+i}, M_{t+i}) \right. \\ \left. \mid D_{t-T_L:t}, S_{t-T_L:t}, \mathcal{N}_{t-T_L:t} \right) \quad (1)$$

where f_{θ} and g_{φ} are the loss function (e.g., mean squared error, mean absolute error) that measures the difference between the predicted drought severity \hat{D}_{t+i} and the actual drought severity D_{t+i} , and the predicted societal impact \hat{M}_{t+i} and the actual societal impact M_{t+i} , respectively. θ and φ represent the parameters of the estimation model that need to be optimized.

Solution

In this section, we present SIDE, a socially informed drought estimation framework that jointly estimates drought severity and the underlying societal impact. Figure 2 shows an overview of the SIDE framework. In particular, SIDE contains three key modules: 1) a *Determinant-Driven Societal Impact Quantification (DSIQ)* module that explicitly quantifies the societal impact of drought by accurately transforming the social and news media data into a comprehensive set of determinants; 2) a *Social-Physical Interdependency Extraction (SPIE)* module that effectively captures the complex interdependency between the distribution of the extracted social determinants and physical drought severity indicator; 3) a *Severity-Impact Joint Estimation (SIJE)* module that leverages the cross-attended features from SPIE to jointly predict future drought severity and its societal impact.

Determinant-Driven Societal Impact Quantification

The determinant-driven societal impact quantification (DSIQ) module is designed to effectively quantify the societal impact of drought by explicitly analyzing social input (Definition 3) and news input (Definition 4). We observe

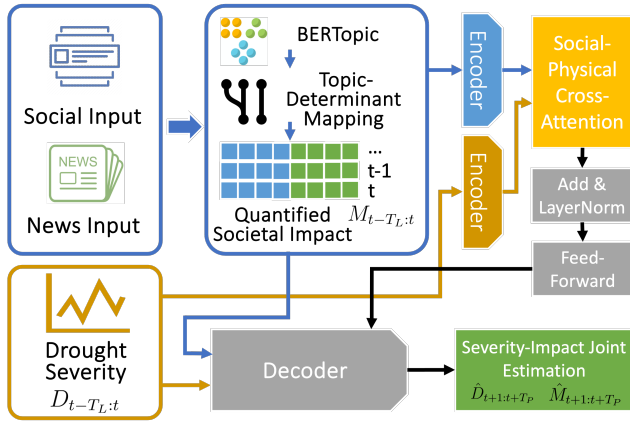


Figure 2: Overview of SIDE

that the social input from social media discourse can provide valuable insights into public sentiment, concerns, and experiences related to drought. Similarly, news input, such as reports from local and national news outlets, can offer a comprehensive overview of the broader societal consequences of drought. However, the social input and news input are primarily unstructured text data, presenting challenges in extracting and quantifying the relevant information for assessing the societal impact of drought. For example, social media posts contributed by individual users often contain a non-trivial amount of noise, such as irrelevant information, opinions, and emotions that may not directly contribute to understanding the societal impact of drought.

To address such a challenge, we design a hierarchical determinant-driven information distillation strategy to effectively extract drought-related posts/articles from the social and news inputs and quantify their distribution across different societal impact determinants. In particular, the hierarchical determinant-driven information distillation strategy consists of two main steps. First, we extract the fine-grained topics from each social and news input to capture the diverse aspects of drought’s societal impact on human communities. We employ the advanced unsupervised neural topic model, BERTopic (Grootendorst 2022), to discover the latent topics in the noisy text data from social and news inputs. Specifically, considering the varied linguistic characteristics of social media posts and news articles (e.g., language formality, writing style), we train two separate topic models using the social input and news input, denoted as Φ_s and Φ_n , respectively. For each topic model, we cluster the corresponding posts/articles into K topic clusters, denoted as \mathcal{T}_s and \mathcal{T}_n .

Second, we map the extracted topics to a pre-defined set of high-level societal impact determinants (Definition 5) and learn the distribution of the posts/articles across the determinants to explicitly quantify the societal impact of drought. A straightforward solution is to manually annotate the extracted topics with their corresponding determinants. However, such an approach is not only unscalable but also prone to the dynamic nature of the topics in the social media discourse and news articles as drought condition evolves. To address this challenge, we propose a large language model

(LLM) assisted topic-determinant mapping approach that aims to incorporate the commonsense knowledge and natural language understanding capability of LLM to adaptively map the dynamic topics to the relevant societal impact determinants. In particular, we provide the extracted topics and the determinants as prompts and obtain the likelihood score for each topic-determinant pair. Finally, we assign each topic to the determinant with the highest likelihood score and aggregate the posts/articles associated with the topics to compute the distribution of the social media posts and news articles across the determinants. Formally, for a given time step t , we define the quantified societal impact as $M_t = [M_{s,t} || M_{n,t}]$ where $M_{s,t}, M_{n,t} \in \mathbb{R}^\delta$ are the normalized distribution of posts/articles across the determinants from the social input and news input, respectively. $||$ is the concatenation operation. The quantified societal impact within the lookback window (Definition 6) will be utilized to estimate the drought severity and predict future societal impact.

Social-Physical Interdependency Extraction

Our next objective is to effectively capture the interdependency between the quantified societal impact and drought severity (Definition 2). In particular, we observe that the quantified societal impact and drought severity exhibit complex interactions and dependencies. On one hand, the severity of drought directly influences the societal impact, as more severe drought conditions tend to lead to more significant consequences on various aspects of human society, such as agriculture, economy, and public health. On the other hand, the societal impact of drought can also provide valuable insights and early signals for estimating and predicting drought severity. For example, an increase in social media posts and news articles related to agricultural losses and water scarcity may indicate a worsening drought condition.

To capture such complex interdependency between the quantified societal impact and drought severity, we design a Social-Physical Interdependency Extraction (SPIE) module that designs a social-physical cross-attention mechanism to explicitly learn the hidden relation between the quantified societal impact and drought severity. In particular, we first encode the time series of quantified societal impact $M_{t-T_L:t}$ and drought severity $D_{t-T_L:t}$ within a lookback window of T_L time steps using two separate encoders, f_{enc}^M and f_{enc}^D :

$$H_M = f_{enc}^M(M_{t-T_L:t}), H_D = f_{enc}^D(D_{t-T_L:t}) \quad (2)$$

where $H_M, H_D \in \mathbb{R}^{T_L \times d}$ are the encoded representations of $M_{t-T_L:t}$ and $D_{t-T_L:t}$, and d is the hidden dimension. f_{enc}^M and f_{enc}^D are the transformer-based sequence encoders for the time series of the societal impact and drought severity, respectively.

Next, we employ a social-physical cross-attention mechanism to capture the hidden relations between H_M and H_D . Specifically, we compute the query, key, and value matrices from H_M and H_D as:

$$Q_M = H_M W_Q^M; K_D = H_D W_K^D; V_D = H_D W_V^D \quad (3)$$

$$Q_D = H_D W_Q^D; K_M = H_M W_K^M; V_M = H_M W_V^M \quad (4)$$

where $W_Q^M, W_K^D, W_V^D, W_Q^D, W_K^M, W_V^M \in \mathbb{R}^{d \times d}$ are learnable weights. Then, we compute the cross-attention weights A_{MD} and A_{DM} as:

$$A_{MD} = \text{softmax}\left(\frac{Q_M K_D^T}{\sqrt{d}}\right); A_{DM} = \text{softmax}\left(\frac{Q_D K_M^T}{\sqrt{d}}\right) \quad (5)$$

The cross-attention weights A_{MD} and A_{DM} capture the hidden relations between the encoded representations of quantified societal impact and drought severity. Specifically, each element in A_{MD} represents the attention weight from a time step in H_M to a time step in H_D , and vice versa for A_{DM} . Finally, we compute the cross-attended representation of the quantified societal impact H_{MD} and the drought severity H_{DM} as:

$$H_{MD} = A_{MD} V_D; H_{DM} = A_{DM} V_M \quad (6)$$

The cross-attended representation H_{MD} and H_{DM} incorporate the hidden relations between the quantified societal impact and drought severity, which are then fed into the subsequent Severity-Impact Joint Estimation module to enhance the joint prediction of drought severity and societal impact.

Severity-Impact Joint Estimation

The severity-impact joint estimation module is designed to jointly predict future drought severity and societal impact by leveraging the cross-attended representation learned from the SPIE module. Traditional time series forecasting models often focus on predicting single or multiple independent target variables based on their historical observations. However, in our problem of socially informed drought estimation, it is crucial to simultaneously estimate both the drought severity and the associated societal impact, as they are inherently interconnected and influence each other dynamically.

To address this challenge, we propose an encoder-decoder architecture that incorporates the cross-attended features from the SPIE module to capture the complex interdependency between drought severity and societal impact. Specifically, let $H_{MD} \in \mathbb{R}^{T_L \times d}$ and $H_{DM} \in \mathbb{R}^{T_L \times d}$ denote the cross-attended representation learned from the SPIE module, where T_L is the lookback window size and d is the hidden dimension. We first concatenate H_{MD} and H_{DM} to obtain a comprehensive representation $H = [H_{MD} || H_{DM}]$ that captures the bidirectional interactions between drought severity and societal impact.

We then utilize a decoder network f_{dec} to jointly predict the future drought severity and societal impact based on the cross-attended representation H . f_{dec} takes H as input and generates the predictions for the next T_P time steps:

$$f_{dec}(H) = [\hat{D}_{t+1:t+T_P}, \hat{M}_{t+1:t+T_P}] \quad (7)$$

where $\hat{D}_{t+i} \in \mathbb{R}$ and $\hat{M}_{t+i} \in \mathbb{R}^\delta$ are the predicted drought severity and quantified societal impact at time step $t+i$, respectively, and T_P is the prediction window size.

We optimize the SIDE framework by jointly considering the estimation error for both drought severity and societal

impact. Specifically, we define the loss function as:

$$\mathcal{L} = \sum_{i=1}^{T_P} (\lambda_D \mathcal{L}_D(\hat{D}_{t+i}, D_{t+i}) + \lambda_M \mathcal{L}_M(\hat{M}_{t+i}, M_{t+i})) \quad (8)$$

where \mathcal{L}_D and \mathcal{L}_M are the loss functions for drought severity and societal impact, respectively (e.g., mean squared error), and λ_D and λ_M are hyperparameters that control the relative importance of each task.

By jointly predicting the future drought severity and societal impact, the SIJE module provides more comprehensive estimation for drought severity and societal impact, aiming to support well-informed decision-making and resource allocation for drought mitigation and management.

Evaluation

In this section, we evaluate the performance of SIDE on estimating drought severity and its societal impact in different geographical areas. Evaluation results demonstrate that SIDE achieves substantial performance gains in accurately forecasting the drought severity.

Data

We adopt the publicly available SocialDrought dataset (Shang et al. 2024) for evaluating the performance of SIDE. In particular, the SocialDrought dataset is a comprehensive drought dataset that contains the meteorological records of drought severity in the U.S. along with the relevant social media discourse from Twitter/X and credible news articles from mainstream news media. In our experiments, we focus on two states in the U.S., namely California (CA) and Texas (TX), as the primary geographical areas in our study due to their significant agricultural outputs, large and diverse populations, and vulnerability to drought conditions (Becker and Sparks 2020; Farahmand, Reager, and Madani 2021).

We retrieve the weekly DSCI (Definition 2), social media posts, and news articles from the SocialDrought dataset for each studied state (i.e., CA and TX) to be used as the ground-truth drought severity, social input, and news input, respectively. The social and news inputs are further utilized to compute the ground-truth societal impact using the quantification method introduced in the DSIQ module. In addition, we notice that only a limited amount of the social media posts and news articles in the SocialDrought dataset contain explicit geolocation information that could be used to retrieve relevant posts and articles in each state. To address this issue, we further filter the posts and articles with a list of state-specific location entities to retrieve posts and articles that are relevant to each state. This list of location entities includes major cities, counties, and landmarks specific to California and Texas, such as Los Angeles, Fresno, and Dallas. This approach allows us to obtain a comprehensive and state-specific dataset for our analysis, ensuring that the social media posts and news articles accurately reflect the drought conditions and public discourse within each state.

In our study, we primarily focus on the period from January 2017 to April 2023 where both California and Texas




Figure 3: Temporal Variation

experienced several significant drought events (Kam, Stowers, and Kim 2019; Zhang et al. 2022). For example, California faced severe drought conditions, particularly in 2021, which was recorded as one of the driest years in the state’s history (Seager et al. 2022). Similarly, Texas encountered extreme drought conditions in 2022, with many parts of the state experiencing exceptional drought, the highest level on the U.S. Drought Monitor scale (Weaver et al. 2023).

We plot the temporal variation of the DSCI and the amount of weekly social media posts and news articles of California and Texas in Figure 3. We observe that the temporal patterns of drought severity, social media activity, and news coverage exhibit notable correlations and fluctuations over time. Periods of increased drought severity often coincide with spikes in the volume of drought-related social media posts and news articles in both California (e.g., summer of 2021) and Texas (e.g. summer of 2022). This suggests that social and news media show great potential for providing valuable insights into the societal impact of drought.

Baselines and Experimental Settings

Baselines We compare SIDE with a set of state-of-the-art multivariate time series forecasting solutions.

- **iTransformer** (Liu et al. 2023): a transformer-based time series forecasting method that applies the attention and feed-forward network on inverted dimensions.
- **Koopa** (Liu et al. 2024): a deep learning approach that utilizes Koopman theory to separate time-variant and time-invariant components in time series forecasting.
- **TimesNet** (Wu et al. 2022): a neural network architecture that leverages tensor transformation to capture intra- and inter-period variations for time series forecasting.
- **One-Fits-All** (Zhou et al. 2024): a pre-trained language and vision model with fine-tuning for downstream time series forecasting tasks.
- **RWKV-TS** (Hou and Yu 2024): an RNN-based sequential forecasting approach that is designed to capture the long-term temporal dependencies in time series data.

Method	MAE (↓)	MSE (↓)	RMSE (↓)	MFA (↑)
SIDE	30.07	1823.20	42.69	0.94
iTransformer	49.25	5913.26	76.89	0.93
Koopa	55.87	6389.19	79.93	0.90
TimesNet	40.78	2026.89	45.02	0.88
One-Fits-All	73.67	26981.50	164.26	0.92
RWKV-TS	74.98	27727.57	166.51	0.91

Table 1: Drought Severity Estimation Performance – CA

Experimental Settings To ensure a fair comparison, we use the same input in SIDE and all compared baseline models, including the DSCI records, transformed social and news inputs, and their corresponding time steps. We follow the original implementation of the compared baselines. We set the lookback window as 52 weeks and the prediction window as 5 weeks. We split each dataset into training, validation, and testing sets using a ratio of 7:1:2. We set $\delta = 11$ based on the latest number of societal impact determinants as specified by the National Integrated Drought Information System. We set the number of topics in the DSIQ module as 50. We train all models using the Adam optimizer with a decay of 0.5. We run each model with a maximum of 20 epochs and apply early stopping with a patience of 10 epochs to prevent overfitting. We run the experiments on Ubuntu 20.04 with four NVIDIA A16.

Drought Severity Estimation Performance

To evaluate the drought severity estimation performance of SIDE, we adopt a series of evaluation metrics that are commonly used for evaluating the performance of time series forecasting, including *Mean Absolute Error (MAE)*, *Mean Squared Error (MSE)*, *Root Mean Squared Error (RMSE)*, and *Median Forecast Accuracy (MFA)* (Chatfield 2000). A higher value (↑) of MFA and lower values (↓) of MAE, MSE, and RMSE indicate better estimation performance. We compare the predicted drought severity with the ground-truth drought severity in the prediction window.

We summarize the evaluation results for California and Texas in Table 1 and Table 2, respectively. We observe that SIDE achieves significant performance gains compared to the baseline methods across all evaluation metrics on both the California and Texas datasets, demonstrating its ability to accurately estimate drought severity. For example, SIDE outperforms the best-performing baseline (i.e., TimesNet) on the California dataset by 26.25% and 10.05% in terms of reducing the MAE and MSE, respectively. We observe similar performance improvements on the Texas dataset. The improvement can be attributed to the social-physical cross-attention mechanism that effectively captures the complex interdependence between dynamic physical conditions and societal impact influencing drought severity.

Method	MAE (↓)	MSE (↓)	RMSE (↓)	MFA (↑)
SIDE	48.08	4369.98	66.10	0.85
iTransformer	50.50	4693.32	68.50	0.84
Koopa	55.87	6389.19	79.93	0.76
TimesNet	116.06	17306.74	131.55	0.60
One-Fits-All	89.30	93827.28	306.31	0.75
RWKV-TS	99.63	101685.62	318.88	0.72

Table 2: Drought Severity Estimation Performance – TX

Societal Impact Estimation Performance

We evaluate the societal impact estimation performance of the SIJE module. We hold out the societal impact quantified by DSIQ on the prediction window (Definition 7), and only use the quantified societal impact from the lookback window (Definition 6) as input to SIJE to predict the future societal impact in the prediction window, which will be compared with the ground-truth societal impact computed by DSIQ. We show the average predicted and ground-truth societal impact of California and Texas in Figure 4. We observe that SIJE accurately captures the distribution of the societal impact determinants for both states. In addition, we observe that, while “water utilities” is the most concerning determinant for both states, California and Texas exhibit distinct patterns in the relative importance of other determinants. For example, “agriculture” is shown to be more concerned in California while “ecosystem” and “public health” are more prominent in Texas. More importantly, we notice that the distribution of the societal impact determinants extracted from the social input and news input also varied (e.g., “wildfire management” in California), demonstrating the importance of considering both social and news inputs to comprehensively estimate the societal impact of drought.

Ablation Study

We further conduct an ablation study to investigate the contribution of the key components in the SIDE framework. The detailed results and analysis are provided in the Appendix of the full paper at: <https://arxiv.org/pdf/2412.12575>.

Discussion

Deployment

The proposed SIDE framework can be deployed as a real-time drought monitoring and forecasting system. For example, by integrating SIDE with the U.S. Drought Monitor (USDM 2024) or DIP-Drought Platform (DIP-Drought 2024) that are developed by leading institutions, we can provide community members with actionable insights for proactive drought management. Moreover, SIDE can support various stakeholders, such as government agencies, water resource managers, and agricultural organizations. These stakeholders can utilize the timely drought severity estimation and societal impact assessments provided by SIDE to

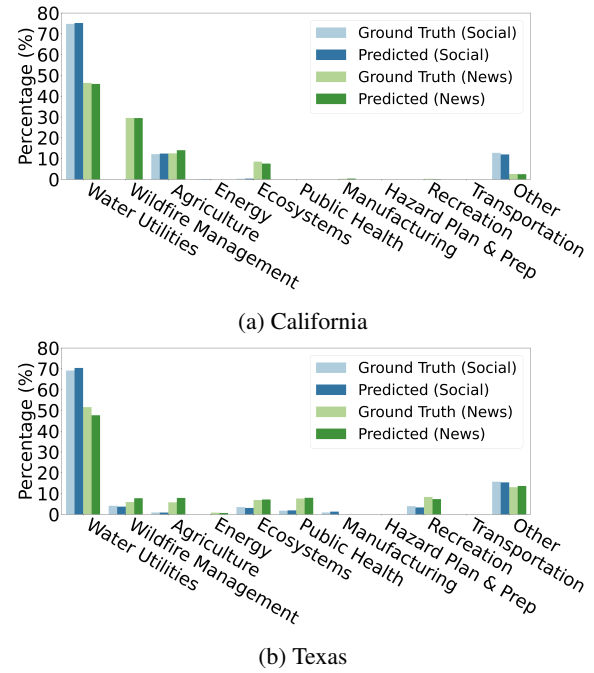


Figure 4: Societal Impact Estimation Performance

make informed decisions regarding water allocation, crop management, and drought mitigation solutions. For example, water resource managers can use the insights from SIDE to optimize water distribution strategies and prioritize water conservation in areas severely affected by drought.

Broader Social Impact

The development of SIDE has broader implications for understanding and addressing the societal impact of environmental crises. By incorporating human-centric perspectives from social and news media, SIDE offers a more comprehensive approach to assessing the diverse impact of drought on human communities. This framework can be adapted to other environmental crises, such as floods, wildfires, and extreme weather events, to gain insights into their societal consequences. Furthermore, the insights derived from SIDE can inform the development of more effective and targeted mitigation strategies and policies that address the specific needs and concerns of affected populations.

Conclusion

This paper introduces SIDE, a socially informed drought estimation framework that leverages social and news media information to jointly estimate drought severity and its societal impact. SIDE addresses the challenges of modeling the implicit temporal dynamics of drought societal impact and capturing the social-physical interdependency between drought conditions and societal impact. Extensive experiments on two real-world datasets from California and Texas demonstrate that SIDE substantially outperforms state-of-the-art baselines in accurately estimating drought severity and its societal impact.

Acknowledgments

This research is supported in part by the National Science Foundation under Grant No. CNS-2427070, IIS-2331069, CHE-2105032, IIS-2130263, CNS-2131622, CNS-2140999. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

- Bathaiy, S. S.; Chizari, M.; Sadighi, H.; and Alambeigi, A. 2021. Social media and farmer's resilience to drought as an environmental disaster: A moderation effect. *International Journal of Disaster Risk Reduction*, 59: 102209.
- Becker, S.; and Sparks, P. 2020. "It never rains in California": Constructions of drought as a natural and social phenomenon. *Weather and Climate Extremes*, 29: 100257.
- Chatfield, C. 2000. *Time-series forecasting*. Chapman and Hall/CRC.
- Dantas, J. C.; da Silva, R. M.; and Santos, C. A. G. 2020. Drought impacts, social organization, and public policies in northeastern Brazil: a case study of the upper Paraíba River basin. *Environmental monitoring and assessment*, 192: 1–21.
- DIP-Drought. 2024. Dialogue and Information Platform for a Drought Event (DIP-Drought). <https://droughtweb.web.illinois.edu/>. Accessed: August 3, 2024.
- Edwards, B.; Gray, M.; and Hunter, B. 2019. The social and economic impacts of drought. *Australian Journal of Social Issues*, 54(1): 22–31.
- Farahmand, A.; Reager, J.; and Madani, N. 2021. Drought cascade in the terrestrial water cycle: Evidence from remote sensing. *Geophysical Research Letters*, 48(14): e2021GL093482.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- Hou, H.; and Yu, F. R. 2024. RWKV-TS: Beyond Traditional Recurrent Neural Network for Time Series Tasks. *arXiv preprint arXiv:2401.09093*.
- Hyndman, R.; Koehler, A. B.; Ord, J. K.; and Snyder, R. D. 2008. *Forecasting with exponential smoothing: the state space approach*. Springer Science & Business Media.
- Kam, J.; Stowers, K.; and Kim, S. 2019. Monitoring of drought awareness from google trends: a case study of the 2011–17 California drought. *Weather, Climate, and Society*, 11(2): 419–429.
- Lester, L.; Flatau, P.; and Kyron, M. 2022. Understanding the Social Impacts of Drought. *Perth: The University of Western Australia*.
- Lim, B.; and Zohren, S. 2021. Time-series forecasting with deep learning: a survey. *Philosophical Transactions of the Royal Society A*, 379(2194): 20200209.
- Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; and Long, M. 2023. itransformer: Inverted transformers are effective for time series forecasting. *arXiv preprint arXiv:2310.06625*.
- Liu, Y.; Li, C.; Wang, J.; and Long, M. 2024. Koopa: Learning non-stationary time series dynamics with koopman predictors. *Advances in Neural Information Processing Systems*, 36.
- Medellin-Azuara, J.; Escriva-Bou, A.; Rodriguez-Flores, J. M.; Cole, S. A.; Abatzoglou, J.; Viers, J. H.; Santos, N.; Summer, D. A.; Medina, C.; Arevalo, R.; et al. 2022. Economic Impacts of the 2020–22 Drought on California Agriculture.
- National Integrated Drought Information System. 2024. <https://www.drought.gov/>. Accessed: August 3, 2024.
- Naumann, G.; Cammalleri, C.; Mentaschi, L.; and Feyen, L. 2021. Increased economic drought impacts in Europe with anthropogenic warming. *Nature Climate Change*, 11(6): 485–491.
- Savelli, E.; Rusca, M.; Cloke, H.; and Di Baldassarre, G. 2022. Drought and society: Scientific progress, blind spots, and future prospects. *Wiley Interdisciplinary Reviews: Climate Change*, 13(3): e761.
- Seager, R.; Ting, M.; Alexander, P.; Nakamura, J.; Liu, H.; Li, C.; and Simpson, I. R. 2022. Mechanisms of a meteorological drought onset: Summer 2020 to spring 2021 in southwestern North America. *Journal of Climate*, 35(22): 7367–7385.
- Shang, L.; Chen, B.; Vora, A.; Zhang, Y.; Cai, X.; and Wang, D. 2024. SocialDrought: A Social and News Media Driven Dataset and Analytical Platform towards Understanding Societal Impact of Drought. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, 2051–2062.
- Shumway, R. H.; Stoffer, D. S.; Shumway, R. H.; and Stoffer, D. S. 2017. ARIMA models. *Time series analysis and its applications: with R examples*, 75–163.
- Smith, K. H.; Tyre, A. J.; Tang, Z.; Hayes, M. J.; and Akyuz, F. A. 2020. Calibrating human attention as indicator monitoring# drought in the Twittersphere. *Bulletin of the American Meteorological Society*, 101(10): E1801–E1819.
- Sugg, M.; Runkle, J.; Leeper, R.; Bagli, H.; Golden, A.; Handwerger, L. H.; Magee, T.; Moreno, C.; Reed-Kelly, R.; Taylor, M.; et al. 2020. A scoping review of drought impacts on health and society in North America. *Climatic Change*, 162: 1177–1195.
- USDM. 2024. United States Drought Monitor. <https://droughtmonitor.unl.edu/>. Accessed: August 3, 2024.
- Weaver, S. M.; Guinan, P. E.; Semenova, I. G.; Aloysius, N.; Lupo, A. R.; and Hunt, S. 2023. A case study of drought during summer 2022: A large-scale analyzed comparison of dry and moist summers in the Midwest USA. *Atmosphere*, 14(9): 1448.

Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; and Long, M. 2022. Timesnet: Temporal 2d-variation modeling for general time series analysis. In *The eleventh international conference on learning representations*.

Zhang, B.; Salem, F. K. A.; Hayes, M. J.; and Tadesse, T. 2022. Quantitative Assessment of Drought Impacts Using XGBoost based on the Drought Impact Reporter. *arXiv preprint arXiv:2211.02768*.

Zhou, T.; Niu, P.; Sun, L.; Jin, R.; et al. 2024. One fits all: Power general time series analysis by pretrained lm. *Advances in neural information processing systems*, 36.