

MetaVIn: Meteorological and Visual Integration for Atmospheric Turbulence Strength Estimation

Ripon Kumar Saha
Arizona State University
rsaha8@asu.edu

Scott McCloskey
Kitware, Inc.
scott.mccloskey@kitware.com

Suren Jayasuriya
Arizona State University
sjayasur@asu.edu

Abstract

Long-range image understanding is a challenging task for computer vision due to the presence of atmospheric turbulence. Turbulence can degrade image quality (blur and geometric distortion) due to the medium’s spatio-temporal varying index of refraction bending light rays. The strength of atmospheric turbulence is quantified by the refractive index structure parameter C_n^2 , and estimating it is important both as an indicator of image degradation and is useful for downstream tasks including video restoration and estimating true shape and range/depth. However, traditional methods for estimating C_n^2 involve expensive and complex optical equipment, limiting their practicality. In this paper, we propose MetaVIn: a **M**eteorological and **V**isual **I**ntegration system to predict atmospheric turbulence strength. Our method leverages image quality metrics to capture sharpness and blur, combined with meteorological information within a Kolmogorov Arnold Network (KAN). We demonstrate that this approach provides a more accurate and generalizable estimation of C_n^2 , outperforming previous state-of-the-art methods in both blind image quality assessment and passive video-based turbulence strength estimation on a large dataset of 35,364 image samples with accompanying ground truth scintillometer measurements for C_n^2 . Our method enables better prediction and mitigation of atmospheric image degradation while being useful in applications such as shape and range estimation, enhancing the practical utility of our approach.

1. Introduction

Atmospheric image degradation, a prevalent issue in long-range imaging, significantly impacts the quality of captured images in applications such as surveillance, astrophotography, and remote sensing. This degradation, often referred to as atmospheric turbulence, manifests as image distortions such as increased blur, reduced contrast, image motion, and geometric deformations, adding tilt and

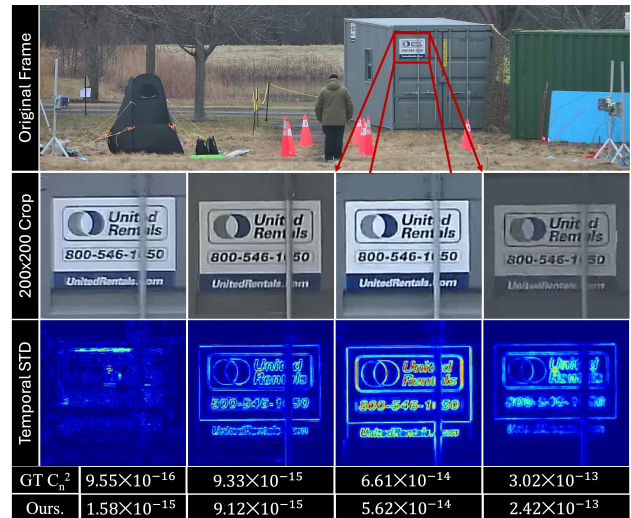


Figure 1. Our MetaVIn algorithm uses image features extracted from a video frame along with measurements from a weather station to estimate the C_n^2 measure of atmospheric turbulence. This parameter predicts the degree of turbulence degradation, as illustrated by the crops and the corresponding temporal standard deviations. Experiments demonstrate that our predictions are more strongly correlated than prior methods with ground truth from an expensive, purpose-built scintillometer.

blur to each pixel in a non-uniform manner across both space and time [9]. The source of these distortions is random fluctuations in the refractive index of the atmosphere, driven by variations in environmental factors such as temperature, humidity, barometric pressure, and wind, leading to phenomena like “image dancing” and the twinkling of stars.

Thus understanding the physical properties of turbulent environments can be very useful for computer vision systems. Of particular interest is the turbulence strength, quantified by the refractive index structure parameter C_n^2 , which has been shown to directly correlate with the degree of image distortion [22]. Lower C_n^2 values indicate very slight atmospheric degradation, whereas higher values indicate se-

vere degradation, causing most conventional and modern image recognition systems to fail or perform poorly [38].

Further, C_n^2 is directly proportional to both the length/depth (L) and pixel variance (α). This means that accurate depth estimation, especially in long-range imaging, is simplified when C_n^2 is known [32]. In astrophotography and in black hole imaging, adaptive optics systems estimate turbulence strength to correct distortions in real-time, significantly enhancing special image clarity to a stunning factor of 10 compared to other ground-based telescopes that don't mitigate turbulence [19]. Accurate and reliable C_n^2 estimation can be potentially useful for computer vision tasks such as image/video restoration [17, 18, 27], segmentation [24], and detection/tracking/recognition [6].

Previous methods for estimating C_n^2 have typically relied on expensive optical setups [39], meteorological sensors [8], or passive video analysis [42] alone. While these approaches have proven effective, they often lack either the practicality or accessibility required for widespread application. In this paper, we propose MetaVIn, a multimodal approach that integrates visual and meteorological data to predict atmospheric turbulence strength. By encoding image quality metrics such as sharpness, Tenenbaum Gradient, and Variance of Gradient, and combining them with meteorological information via a KAN [14], we achieve improved accuracy and generalizability in estimating C_n^2 . This integrated approach leverages the strengths of both data types, offering a comprehensive solution for predicting and mitigating atmospheric image degradation. Our contributions are as follows:

- We introduce MetaVIn, a novel method to infer C_n^2 values from images and weather station data. Our C_n^2 predictions correlate strongly with data from more expensive scintillometers and can be measured remotely without emplacing hardware in the scene.
- We demonstrate that leveraging a KAN instead of a fully Connected Neural Network (FCNN) significantly enhances prediction accuracy and model performance for the multimodal data.
- We provide a detailed analysis of feature importance within the model, allowing practitioners to select a weather station that captures data most relevant to C_n^2 prediction.
- We conduct a thorough evaluation of existing blind image quality assessment (BIQA) methods, recent deep learning models as well as existing image gradient-based C_n^2 estimation, showcasing their limitations in the context of atmospheric turbulence and highlighting the necessity for our proposed method.

To validate our method, we leverage a series of datasets from the Biometric Recognition & Identification at Altitude & Range (BRIAR) program provided by Cornett et al. [4] which features 35,364 videos with accompanying meteo-

rological data as well as ground truth scintillometer measurements for C_n^2 . This large dataset features video captured from several different geographic locations over several years with varying weather/visibility and human subjects moving in the scene as shown in Table 1. To the best of our knowledge, this paper is the largest scale study for C_n^2 estimation conducted in the literature compared to more limited studies carried out over shorter periods and/or fewer locations [28, 41, 42]. We hope that this paper inspires more multimodal fusion between vision and meteorological data for environmental and remote sensing problems in the future.

2. Related Work

2.1. Optical Instrumentation for C_n^2 estimation

The most common approach to estimating C_n^2 utilizes large aperture scintillometers that transmit infrared light over long distances, analyzing the scintillation of the received signal to determine turbulence strength [39]. While highly accurate, these systems are expensive, with costs often running into hundreds of thousands of dollars. Additionally, they require precise calibration and alignment, which can be logistically challenging and time-consuming.

Other optical methods include Hartmann masks for measuring the displacement of light spots caused by turbulence [2] and wavefront sensor-based methods that measure the fluctuation of light waves to estimate turbulence strength [5, 29]. These methods can provide precise turbulence strength measurements essential for real-time image correction but require specialized equipment and complex setups, and are rarely deployed in the field.

Our proposed method, MetaVIn, eliminates the need for costly and complex setups by integrating visual and meteorological data. By using image quality metrics and weather station data within a KAN, we achieve accurate C_n^2 estimation, making the process more affordable and straightforward.

2.2. Estimating C_n^2 from passive video

One-sided systems using image analysis have gained attention for their practicality. These methods estimate turbulence strength by analyzing the gradients and other features of long-range images [28, 41, 42]. Robust, physics-based deep learning models have been introduced to improve estimation accuracy from video imagery. Experimental trials with neuromorphic cameras over a 7 km path demonstrate their effectiveness in reducing data processing needs and offering high dynamic range, enabling real-time turbulence measurement and characterization [23]. However, most of these approaches are tested in controlled environments with minimal vibrations, focusing on a fixed target with visible image gradients while avoiding other object movements in

the region of interest. These methods may not perform well in uncontrolled environments, particularly in the presence of wind.

2.3. Meteorological Sensing for C_n^2 estimation

Meteorological data has been utilized to estimate C_n^2 through various proposed models that leverage parameters such as temperature, humidity, and wind speed [1, 8, 25, 26, 34, 36]. However, this research focuses on path-averaged C_n^2 , which represents the average turbulence strength along a specific propagation path. Since the effect on imaging systems is a path-averaged quantity, corresponding to the integral over the path traveled by light rays [32], and meteorological data only provides nodal measurements (turbulence strength at specific spatial locations), relying solely on meteorological data often fails to provide accurate estimates of the path-averaged quantity.

2.4. Blind Image Quality Assessment

As turbulence degrades image quality proportionally to turbulence strength, C_n^2 can be formulated as a function of image quality assessment. While reference-based techniques like PSNR, SSIM, LPIPS, and Visual Information Fidelity (VIF) are unsuitable metrics for atmospheric image degradation due to the impossibility of obtaining clear reference images, no-reference/blind image quality assessment (BIQA) methods can evaluate overall image quality including sharpness, contrast, distortion, color accuracy, and artifacts. Classical objective methods include BRISQUE [20], NIQE [21], and ILNIQE [43], while deep learning-based approaches for objective image analysis include DBCNN [44], TRES [7], and TOPIQ-NR [3]. Subjective assessment methods like MUSIQ [12], NIMA [30], CLIP-IQ [33], and PAQ2PIQ [40] are also available. While these BIQA methods are valuable for evaluating image quality without references, they are limited in accurately detecting atmospheric degradation strength.

Our proposed method, MetaVIn, builds on these approaches by integrating visual and meteorological data within a KAN [14]. This multimodal approach combines image quality metrics with environmental data, providing a more accurate and robust estimation of atmospheric turbulence strength that can help mitigate its impact on image degradation.

3. Methods

Overview: Our approach integrates visual and meteorological data to predict atmospheric turbulence strength, C_n^2 , and is visualized in Figure 2. In this section, we describe key components of our pipeline including visual feature extraction from single frames, meteorological parameters used for multimodal integration, and the architecture of our KAN. In Section 4, we discuss the training and testing of the

method on the BRIAR datasets [4], which includes long-range video sequences and corresponding weather data as well as C_n^2 values collected using a large aperture scintillometer.

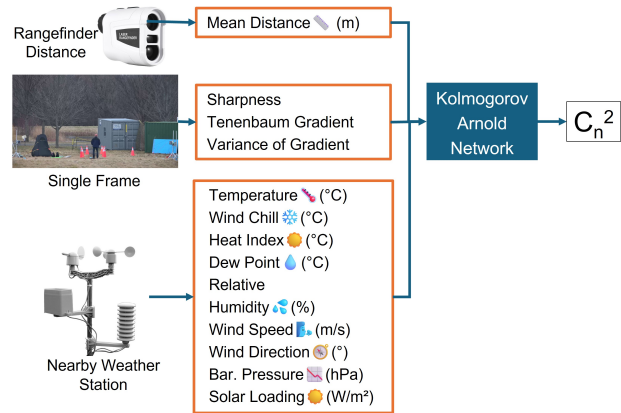


Figure 2. Overview of the proposed model architecture integrating visual and meteorological data for turbulence strength prediction.

Sharpness Features: Since turbulence directly blurs captured images, image sharpness serves as a crucial indicator for estimating C_n^2 , as demonstrated by Zamek et al. [42] and Saha et al. [28] who utilize spatial gradients in their computational approaches. In our pipeline, we extract sharpness metrics from a single video frame (the first frame). We employ three key metrics: Sharpness, Tenenbaum gradient & Variance of Gradient where “Sharpness”, computed as the sum of absolute Laplacian values; “Tenenbaum gradient (Tenengrad)”, which measures edge sharpness [13] using the formula:

$$\frac{\sum |\nabla I_x \cdot \cos(\theta) + \nabla I_y \cdot \sin(\theta)|}{N}$$

where ∇I_x and ∇I_y represent image gradients, θ denotes gradient direction, and N is the pixel count; and “Variance of Gradient”, which quantifies gradient magnitude variability to capture image detail and texture.

While more sophisticated image sharpness and image quality assessment (IQA) measures exist in the literature, our experimental results demonstrate that these methods are not optimal for integration within our KAN. Additionally, using learned visual features from CNNs would significantly expand the feature space dimensionality beyond our low-dimensional classical visual features, necessitating a larger KAN for training. We avoided this approach due to potential overfitting issues and concerns about memory requirements and latency. In Section 5, we compare our method against a baseline CNN trained directly on visual data, demonstrating how classical sharpness features com-

bined with meteorological data in the KAN achieve superior performance over the learned approach.

Meteorological and Distance Data: The meteorological data is collected from nearby weather stations included in the BRIAR datasets [4], including parameters such as Temperature ($^{\circ}\text{C}$), Wind Chill ($^{\circ}\text{C}$), Heat Index ($^{\circ}\text{C}$), Dew Point ($^{\circ}\text{C}$), Relative Humidity (%), Wind Speed (m/s), Wind Direction ($^{\circ}$), Barometric Pressure (hPa), and Solar Loading (W/m^2). Meteorological data can be obtained relatively inexpensively with mobile weather stations that can be co-located with cameras (e.g., the VEVOR YT60234 WiFi Weather Station costs under \$100). Additionally, Sadot et al. demonstrated that C_n^2 could be estimated using macro meteorological measurements taken from nearby standard weather stations [26]. The meteorological data collected by the BRIAR weather stations is normalized using Z-scoring to ensure zero mean and unit variance before being input to the KAN.

In addition to meteorological data, an important optical parameter is the path length or distance from the camera to the target scene, which factors directly into the calculation of C_n^2 [32]. At larger distances, the observed image distortion effects of C_n^2 are magnified, and thus it is important to account for distance when performing the inverse estimation problem. In our method, we leverage a single distance measurement obtained by a laser rangefinder from the camera to the center of the scene in terms of ground truth meters. We found that this single distance value is more practical than obtaining sparse depth via scanning LIDAR systems, which can be complex and expensive, and more accurate than derived depth values from monocular depth estimators (such as DepthAnythingv2 [37]), which do not yield correct metric depths for our turbulent environments.

Kolmogorov Arnold Network (KAN): The computed image sharpness metrics, meteorological data, and distance value are fed into a KAN [14]. KANs represent a novel neural network architecture based on the Kolmogorov-Arnold representation theorem, which demonstrates that any continuous multivariate function can be expressed through a composition of continuous univariate functions. Unlike traditional Multi-Layer Perceptrons (MLPs) that employ fixed activation functions at nodes, KANs utilize learnable activation functions on edges between nodes. This architectural design enables KANs to efficiently approximate complex, high-dimensional functions through their decomposition into simpler univariate functions, resulting in superior accuracy and interpretability compared to traditional MLPs.

The KAN outputs the predicted turbulence strength, C_n^2 , which characterizes the atmospheric turbulence level affecting image quality. Through the integration of both visual and meteorological features, our approach delivers a ro-

Dataset	Location	Date	Samples
BRS1.1	ORNL, TN	Nov 2021	1663
BRS2	Perry, GA	Mar-Apr 2022	9083
BRS3	ORNL, TN	Aug-Sep 2022	8498
BRS4	Glen Ellyn, IL	Jan 2023	7751
BTS3	ORNL, TN	Aug-Sep 2022	4305
BTS4	Glen Ellyn, IL	Jan 2023	4064

Table 1. Summary of the BRIAR datasets used in the experiments, including location, date range, and number of samples.

bust prediction model that effectively handles diverse atmospheric conditions and imaging configurations.

4. Dataset and Implementation

4.1. Datasets

As mentioned earlier, we conducted a large-scale study for C_n^2 estimation spanning multiple geographic locations at different times. We primarily utilized datasets provided by the BRIAR program [4], with training splits (BRIAR Research Set (BRS)) and testing splits (BRIAR Test Set (BTS)). The dataset exhibits C_n^2 values spanning more than four orders of magnitude, ranging from 2.8×10^{-16} to 4.4×10^{-12} , indicating a wide range of atmospheric degradation. The combined dataset consists of more than 20TB of video and weather data, from which we extracted 35,364 frames from distinct videos (we only extract the first frame from each video) with corresponding weather data and C_n^2 values. The details of the datasets are provided in Table 1. For this experiment, we used BRS1.1, BRS2, BRS3, and BRS4 for training, while BTS3 and BTS4 were used for testing. BTS1.1 and BTS2 were excluded from the testing set as they lacked ground truth C_n^2 values.

Data Processing: To prepare the data for analysis, we addressed missing weather data values that occurred due to collection errors. Specifically, BRS1.1 lacked Wind Chill and Heat Index values, while BRS2 had missing wind direction measurements. We employed a random forest regressor model for data imputation, training it exclusively on the available weather data from BRS samples. The imputation process was conducted separately for each dataset, ensuring that neither C_n^2 values nor BTS data were included in the training process. This preprocessing step enabled us to focus on field videos with complete meteorological and image data, yielding a refined dataset of 35,364 samples. Descriptive statistics of the values are presented in Table 2.

4.2. Implementation details

Network training: For our experiments, we implemented a 3-layer KAN architecture. The input layer size was determined by the number of input features: 13 for MetaViN, comprising 9 meteorological values, 3 sharpness values,

Sensor	Min	Max	Mean	Std
Temperature (°C)	-5.3	32.8	14.7	11.1
Wind Chill (°C)	-11.4	32.8	14.0	12.1
Heat Index (°C)	-5.6	38.9	15.3	11.9
Dew Point (°C)	-7.7	24.6	9.2	10.0
Rel. Humidity (%)	29.0	98.0	72.8	18.1
Wind Speed (m/s)	0.0	19.1	2.6	2.2
Wind Dir. (°)	0.0	359	177	96.3
Bar. Pressure (mbar)	990	1026	1011	10.1
Solar Loading (W/m ²)	0.0	1223	381	329
C_n^2 Value (m ^{-2/3})	2e-16	5e-12	7e-14	1e-13

Table 2. Descriptive Statistics of the BRIAR Datasets.

and 1 distance value. The hidden layer contained 8 units, and the output layer consisted of a single unit representing the final C_n^2 value in log scale.

Unlike training fully connected neural networks (FCNNs), we discovered that extremely high learning rates yielded better performance than lower rates, even with extended training durations. For MetaVIN, we employed a learning rate of 0.9. To maintain simplicity, we utilized an L1 loss function with the LBFGS optimizer.

When experimenting with different types of input separately, we explored learning rates ranging from 0.01 to 0.9, grid values from 1 to 9, and hidden layer sizes from 3 to 12. Training the MetaVIN network typically required approximately 3 minutes and ran for 375 epochs, though this duration varied across different input configurations.

Baselines and comparisons: For baseline comparison, we included several blind IQA methods: BRISQUE [20], NIQE [21], IL-NIQE [43], PaQ-2-PiQ [40], DBCNN [44], MUSIQ [12], CNNIQA [11], NIMA [30], TReS [7], CLIP-IQA [33], and TOPIQ-NR [3]. These blind IQA metrics are directly correlated with the logarithmic distribution of C_n^2 values after matching the mean and standard deviation between the two distributions.

Additionally, we compared with the image gradient-based C_n^2 estimation proposed by Zamek et al. [42]. For this passive video-based method, we first performed camera shake stabilization using the Turb-Seg-Res stabilizer [27], which is specifically designed for atmospheric turbulence videos.

For comprehensive evaluation, we implemented two deep learning (DL) approaches: EfficientNetV2 [31] and ViT-Base Patch 16 [35], training them from scratch on the entire BRS dataset. While all images were captured at different resolutions, we first resized them to $1920 \times 1080p$ and then created center-cropped versions of $512 \times 512p$ for EfficientNetV2 and $224 \times 224p$ for ViT to preserve image texture, which is crucial for measuring image degradation.

These models were trained with a batch size of 4 on an A100 GPU for 10 epochs using 26,495 BRS images.

We further expanded our comparison by evaluating several learning-based architectures. This included predictions based on meteorological information alone, image-based values (extracted using sharpness measures), and a simple 4-layer Fully Connected Neural Network (FCNN) with 2753 parameters to investigate the feasibility of C_n^2 estimation using different input modalities.

Metrics: We evaluated our model using three primary metrics: Mean Absolute Error (MAE) on a log scale, Relative Error (RE) on a log scale, and Spearman correlation. For ground truth C_n^2 values, which range from 2.8×10^{-16} to 4.4×10^{-12} , we applied a log10-based normalization to bring the data to a linear scale. Since Blind IQA methods are not inherently on the same scale as the log-transformed C_n^2 , we adjusted their output to follow the same distribution as the log-transformed C_n^2 for fair comparison. Similarly, for the gradient-based C_n^2 estimation, direct computation was not feasible. Instead, we computed a pseudo C_n^2 by excluding distance and camera-specific parameters and then scaled it to match the distribution of the ground truth C_n^2 . The relative error was calculated using the formula:

$$\text{relative_error} = \frac{\sum |\log_{10}(y_{\text{test}}) - \log_{10}(y_{\text{pred}})|}{\sum |\log_{10}(y_{\text{test}}) + \log_{10}(y_{\text{pred}})|}$$

This approach ensured consistent and comparable metrics across all methods and models.

5. Experimental Results

5.1. Main Results

We evaluated all methods on the BRIAR Test Set (BTS3, BTS4), and summarized their results in Table 3.

Blind IQA Results: Among the classic IQA methods, BRISQUE showed the highest Spearman correlation (0.035) but still had high MAE (0.828) and relative error (0.060), indicating limited effectiveness. NIQE and IL-NIQE performed poorly, with negative Spearman correlations and high errors as shown in Table 3. For deep learning-based IQA methods, PaQ-2-PiQ exhibited the best performance with a Spearman correlation of 0.139, an MAE of 0.785, and a relative error of 0.057. MUSIQ showed moderate improvements but were not sufficient for precise C_n^2 estimation, while others showed a negative correlation. These results lend evidence to the observation that blind IQA methods do not correlate well with turbulence strength.

Gradient-Based C_n^2 Estimation: The gradient-based C_n^2 estimation method [42] achieved a Spearman correlation of

0.079 and an MAE of 0.631 as shown in Table 3, demonstrating its limitations in accurately capturing atmospheric turbulence effects. In our investigation, we found that various factors such as rain, snow, camera vibration, and the presence of different pixel movements caused by people, animals, water flow, and rolling shutter effects often misled the gradient-based approach, resulting in inaccurate C_n^2 estimations.

Deep Learning Results: Our results in Table 3 show that EfficientNetV2 [31] achieves a high correlation of 0.762 and a mean absolute error (MAE) of 0.164, outperforming other image quality assessment (IQA) methods. In contrast, the performance of ViT-B16 [35] was relatively poor. This aligns with findings that ViT models require large datasets to outperform CNNs. CNNs like EfficientNetV2 are inherently better suited for this task due to their ability to efficiently process pixel-level details, while Vision Transformers (ViTs) focus on global features through self-attention mechanisms, which may not be as effective for IQA tasks.

KAN/MetaVIn Results: Our proposed MetaVIn method (incorporating meteorological data, image sharpness values, and distance in KAN) significantly outperformed all baselines, achieving a Spearman correlation of 0.943, an MAE of 0.117, and a relative error of 0.0064 as shown in Table 3. For image-based C_n^2 prediction alone, EfficientNetV2 showed the second-best performance with a Spearman correlation of 0.762 (compared to MetaVIn’s 0.943) and MAE of 0.354 (compared to MetaVIn’s 0.177). All other IQA methods demonstrated poor performance in C_n^2 prediction from images. These results validate the superior accuracy and reliability of MetaVIn in estimating atmospheric turbulence strength.

We evaluated the performance of the KAN model using three different combinations: just sharpness, just meteorological data, and a combination of sharpness, meteorological data, and distance (MetaVIn). The evaluation metrics used include Mean Absolute Error (MAE) and Relative Error (RE) on a logarithmic scale. The results are summarized in Table 4. The scatter plots in Figure 3 further illustrate the performance of each model. The MetaVIn model shows the highest correlation between the ground truth and predicted C_n^2 values, indicating its superior predictive capability.

5.2. Ablation Studies

In this subsection, we discuss ablative studies regarding the importance of KAN over FCNN, importance of multimodal features as well as feature selection analysis using Shapley additive explanations (SHAP) [15]. In addition, we have included a study about KAN architecture selection in the supplemental material.

Classic IQA Methods			
Method	Spearman \uparrow	MAE \downarrow	Rel. Error \downarrow
BRISQUE	0.035	0.828	0.060
NIQE	0.015	0.823	0.06
IL-NIQE	-0.152	0.881	0.064
Deep Learning-based IQA Methods			
Method	Spearman \uparrow	MAE \downarrow	Rel. Error \downarrow
PaQ-2-PiQ	0.139	0.785	0.057
DBCNN	-0.130	0.901	0.065
MUSIQ	0.094	0.791	0.057
CNNIQA	-0.04	0.871	0.063
NIMA	-0.051	0.883	0.064
TReS	-0.04	0.868	0.063
CLIP-IQA	-0.04	0.871	0.063
TOPIQ-NR	-0.020	0.856	0.062
C_n^2 Estimation Methods (Trained/Finetuned)			
Method	Spearman \uparrow	MAE \downarrow	Rel. Error \downarrow
Grad C_n^2	0.079	0.631	0.079
ViT B 16	0.320	0.629	0.045
EfficientNetV2	0.762	0.354	0.025
MetaVIn (ours)	0.943	0.177	0.006

Table 3. Benchmark results of various image quality assessment methods for atmospheric image degradation. Classic methods include BRISQUE [20], NIQE [21], and IL-NIQE [43]. Deep learning-based methods include PaQ-2-PiQ [40], DBCNN [44], MUSIQ [12], CNNIQA [11], NIMA [30], TReS [7], CLIP-IQA [33], and TOPIQ-NR [3]. The MAE and relative error were calculated on the log-scaled C_n^2 values after normalizing the methods to match the mean and standard deviation of the log-scaled values.

Features in KAN			Performance Metrics		
Meteor.	Sharp.	Dist.	MAE \downarrow	R.E. \downarrow	Sp. C. \uparrow
	✓		0.606	0.0220	0.341
	✓	✓	0.479	0.0173	0.582
✓			0.245	0.0089	0.892
✓	✓		0.242	0.0087	0.907
✓	✓	✓	0.177	0.0064	0.943
Features in FCNN			Performance Metrics		
			MAE \downarrow	R.E. \downarrow	Sp. C. \uparrow
✓	✓	✓	0.244	0.0087	0.895

Table 4. Evaluation of Feature Combinations on KAN and FCNN Performance for Atmospheric Image Degradation Estimation.

Importance of multimodal features: The KAN model’s performance in predicting atmospheric turbulence strength significantly improves with the combination of diverse features. As shown in Table 4, using only sharpness features




Figure 3. Scatter plots of Ground Truth vs Predicted C_n^2 values for different models. The X-axis represents ground truth C_n^2 whereas the Y-axis shows prediction. (1st from left) Just 3 Sharpness values and 1 distance value, (2nd from left) KAN-based prediction based on Meteorological information, (3rd from left) All Meteorological information with 3 Sharpness values and 1 distance based C_n^2 prediction using FCNN and (Right most) All Meteorological information with 3 Sharpness values and 1 distance based C_n^2 prediction using KAN.

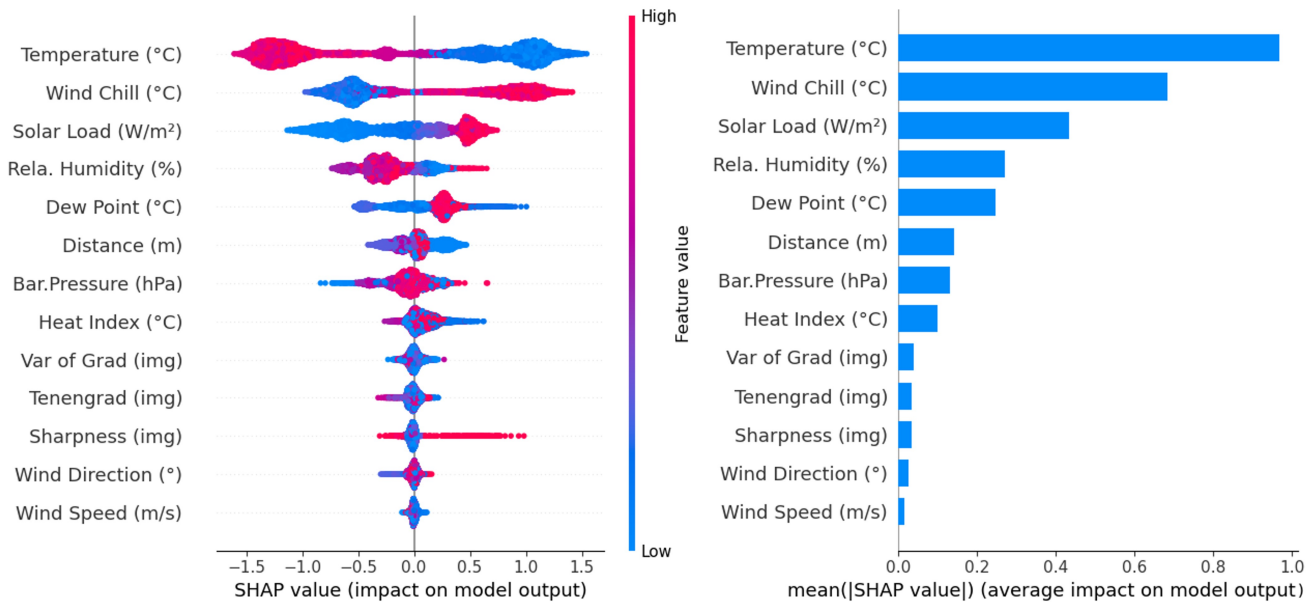


Figure 4. SHAP values indicating the importance of different features in the KAN model. The left plot shows the distribution of SHAP values for each feature, with red indicating high feature values and blue indicating low feature values. The right plot ranks features by their average absolute SHAP values, highlighting the most influential features in predicting C_n^2 . Temperature (°C), Wind Chill (°C), Solar Loading (W/m²), Relative Humidity (%), Dew Point (°C), and Barometric Pressure (hPa) are the most significant meteorological features, while Tenengrad, Sharpness, and Variance of Gradient are the most significant image-based features. Distance seems to be more important than image-based features.

yielded a Spearman correlation (ρ) of 0.341 and a Mean Absolute Error (MAE) of 0.606. Adding distance features enhanced performance ($\rho = 0.582$, MAE = 0.479), while meteorological data alone led to a substantial improvement ($\rho = 0.892$, MAE = 0.245), underscoring the importance of environmental context in turbulence prediction.

Combining meteorological data with sharpness features further improved results ($\rho = 0.907$, MAE = 0.242). However, the best performance was achieved by MetaVIN, which

integrates all features (sharpness, distance, and meteorological data), resulting in ($\rho = 0.943$, MAE = 0.177) as shown in Figure 3 and Table 4. This synergistic effect demonstrates that leveraging multiple feature types leads to more accurate and robust predictions of atmospheric turbulence strength. The KAN model's ability to integrate these diverse features provides a comprehensive understanding of atmospheric turbulence dynamics.

Feature selection: The SHAP (Shapley Additive explanations [16]) feature importance analysis for the KAN model highlights the significance of various features in predicting atmospheric turbulence strength (C_n^2). As shown in Figure 4, features are ordered by their importance, with Temperature being the most influential. The plot’s color scheme represents the feature values, where red indicates high values and blue indicates low values. Overall, higher values of important features tend to increase turbulence strength, while lower values have the opposite effect.

For meteorological features, the KAN model indicates that higher values of Solar Loading, Dew Point, and Barometric Pressure correlate with increased turbulence strength and more image degradation. This is likely due to solar radiation causing thermal convection, higher dew points indicating more moisture, which affects the refractive index, and higher pressure leading to more stable atmospheric layers that can refract light differently. Conversely, higher Temperature values show an opposite trend, potentially due to the stabilizing effect of higher temperatures on the atmospheric layers. Relative Humidity also shows a negative trend. Wind Direction, Wind Chill, and Wind Speed exhibit low feature importance, indicating a lesser impact on turbulence prediction.

Regarding image-based features, the SHAP analysis shows that they contribute less significantly to C_n^2 prediction. For instance, Tenengrad, which captures edge sharpness, demonstrates a negative correlation, meaning that more severe turbulence tends to blur edges and reduce sharpness. The Variance of Gradient (Var of Grad) shows a slight positive correlation with turbulence strength, likely indicating that stronger turbulence leads to greater variations in image gradients. Interestingly, Sharpness shows an overall positive trend, implying that extreme sharpness could represent image artifacts caused by severe turbulence. This lower importance of image features compared to meteorological data suggests that while visual data is relevant, atmospheric factors dominate turbulence predictions.

6. Conclusion

In this paper, we introduced MetaVIn, a multimodal approach for estimating atmospheric turbulence strength by integrating visual and meteorological data. Leveraging image quality metrics and comprehensive meteorological information within a KAN, MetaVIn significantly improves the accuracy and generalizability of C_n^2 estimation compared to traditional and deep learning-based BIQA methods. Validated using the large-scale BRIAR dataset [4], our method demonstrated robust predictions under varied environmental conditions. Overall, MetaVIn offers a comprehensive solution for mitigating atmospheric image degradation, benefiting applications in astrophotography, surveillance, remote sensing, autonomous driving, and environ-

mental monitoring by providing more accurate C_n^2 estimations for better image restoration and analysis.

Limitations: While our proposed MetaVIn model demonstrates significant improvements in estimating atmospheric turbulence strength by integrating visual and meteorological data, there are several limitations to consider. Our method does not leverage temporal information inherently present in video data which could potentially improve the turbulence strength estimation. We did conduct an initial study adding temporal variance of the video pixels as an additional feature value to the KAN, but this did not improve results significantly. Further research is needed to identify what temporal features or network architectures would fully exploit this information and benefit this application. In addition, our method’s latency on a 1080p frame is 15 FPS which is not yet real-time. This is largely due to our unoptimized image sharpness computation in OpenCV and could be accelerated on a GPU via CUDA kernels. However, it remains a systems challenge to configure the weather stations and cameras to stream real-time data into the compute pipeline for end-to-end real-time sensing.

Future Work: In this study, we focused on the problem of turbulence strength estimation. For future work, we wish to utilize C_n^2 estimation as part of downstream tasks such as algorithmically selecting the best lucky frame, and minimizing motion blur and vibration effects. Additionally, we plan to extend our method to automatically identify the optimal region of interest with no object movement, thereby enhancing the robustness of our approach. Furthermore, we intend to develop our model to predict time-series estimations of C_n^2 in real-time, leveraging the entire video duration for continuous and dynamic atmospheric turbulence strength assessment.

Acknowledgements

We wish to thank ASU Research Computing for providing GPU resources to support this research [10]. This research was supported in part by NSF IIS-2232299. In addition, this research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2022-21102100003. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Anthonysamy Arockia Bazil Raj, James Arputha Vijaya Selvi, and Singaravelu Raghavan. Real-time measurement of meteorological parameters for estimating low-altitude atmospheric turbulence strength (cn2). *IET Science, Measurement & Technology*, 8(6):459–469, 2014. [3](#)
- [2] Alexander Boeckenedt, Jack McCrae, Santasri Bose-Pillai, Benjamin Wilson, and Steven Fiorino. Improving on atmospheric turbulence profiles derived from dual beacon hartmann turbulence sensor measurements. *Applied Sciences*, 12(12):5822, 2022. [2](#)
- [3] Chaofeng Chen, Jiadi Mo, Jingwen Hou, Haoning Wu, Liang Liao, Wenxiu Sun, Qiong Yan, and Weisi Lin. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing*, 2024. [3](#), [5](#), [6](#)
- [4] David Cornett, Joel Brogan, Nell Barber, Deniz Aykac, Seth Baird, Nicholas Burchfield, Carl Dukes, Andrew Duncan, Regina Ferrell, Jim Goddard, et al. Expanding accurate person recognition to new altitudes and ranges: The briar dataset. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 593–602, 2023. [2](#), [3](#), [4](#), [8](#)
- [5] Angela Cortés, Benoit Neichel, Andrés Guesalaga, James Osborn, Francois Rigaut, and Dani Guzman. Atmospheric turbulence profiling using multiple laser star wavefront sensors. *Monthly Notices of the Royal Astronomical Society*, 427(3):2089–2099, 2012. [2](#)
- [6] Dawei Du, Cole Hill, Gabriel Bertocco, Mauricio Pamplona Segundo, Wes Robbins, Brandon RichardWebster, Roderic Collins, Sudeep Sarkar, Terrance Boulton, and Scott McCloskey. Doers: Distant observation enhancement and recognition system. In *2023 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–11. IEEE, 2023. [2](#)
- [7] S Alireza Golestaneh, Saba Dadsetan, and Kris M Kitani. No-reference image quality assessment via transformers, relative ranking, and self-consistency. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1220–1230, 2022. [3](#), [5](#), [6](#)
- [8] Xu Hou, Yi Hu, Fujia Du, Michael CB Ashley, Chong Pei, Zhaohui Shang, Bin Ma, Erpeng Wang, and Kang Huang. Machine learning-based seeing estimation and prediction using multi-layer meteorological data at dome a, antarctica. *Astronomy and Computing*, 43:100710, 2023. [2](#), [3](#)
- [9] Claudia S Huebner. An appraisal of suitable evaluation methods for turbulence mitigation algorithms. In *Environmental Effects on Light Propagation and Adaptive Systems V*, volume 12266, pages 83–93. SPIE, 2022. [1](#)
- [10] Douglas M Jennewein, Johnathan Lee, Chris Kurtz, William Dizon, Ian Shaeffer, Alan Chapman, Alejandro Chiquete, Josh Burks, Amber Carlson, Natalie Mason, et al. The sol supercomputer at arizona state university. In *Practice and Experience in Advanced Research Computing*, pages 296–301. 2023. [8](#)
- [11] Le Kang, Peng Ye, Yi Li, and David Doermann. Convolutional neural networks for no-reference image quality assessment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1733–1740, 2014. [5](#), [6](#)
- [12] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5148–5157, 2021. [3](#), [5](#), [6](#)
- [13] Eric Krotkov. Focusing. *International Journal of Computer Vision*, 1(3):223–237, 1988. [3](#)
- [14] Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756*, 2024. [2](#), [3](#), [4](#)
- [15] Scott Lundberg. A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*, 2017. [6](#)
- [16] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4765–4774. Curran Associates, Inc., 2017. [8](#)
- [17] Zhiyuan Mao, Nicholas Chimitt, and Stanley H Chan. Image reconstruction of static and dynamic scenes through anisoplanatic turbulence. *IEEE Transactions on Computational Imaging*, 6:1415–1428, 2020. [2](#)
- [18] Zhiyuan Mao, Ajay Jaiswal, Zhangyang Wang, and Stanley H Chan. Single frame atmospheric turbulence mitigation: A benchmark study and a new physics-inspired transformer model. In *European Conference on Computer Vision*, pages 430–446. Springer, 2022. [2](#)
- [19] Claire E Max, Gabriela Canalizo, and Willem H de Vries. Locating the two black holes in ngc 6240. *Science*, 316(5833):1877–1880, 2007. [2](#)
- [20] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012. [3](#), [5](#), [6](#)
- [21] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2012. [3](#), [5](#), [6](#)
- [22] Mary O’Neill. Strength of turbulence from video imagery. Technical report, Nevada National Security Site/Mission Support and Test Services LLC, 2017. [1](#)
- [23] Ernst Polnau and Mikhail A Vorontsov. Atmospheric turbulence characterization using a neuromorphic camera-based imaging sensor. *Journal of Optics*, 23(12):125608, 2021. [2](#)
- [24] Dehao Qin, Ripon Kumar Saha, Woojeh Chung, Suren Jayasuriya, Jinwei Ye, and Nianyi Li. Unsupervised moving object segmentation with atmospheric turbulence. In *European Conference on Computer Vision*, pages 18–37. Springer, 2024. [2](#)
- [25] A Arockia Bazil Raj, J Arputha Vijaya Selvi, and S Durairaj. Comparison of different models for ground-level atmospheric turbulence strength (cn 2) prediction with a new model according to local weather data for fso applications. *Applied Optics*, 54(4):802–815, 2015. [3](#)
- [26] Dan Sadot and Norman S Kopeika. Forecasting optical turbulence strength on the basis of macroscale meteorology

- and aerosols: models and validation. *Optical Engineering*, 31(2):200–212, 1992. 3, 4
- [27] Ripon Kumar Saha, Dehao Qin, Nianyi Li, Jinwei Ye, and Suren Jayasuriya. Turb-seg-res: A segment-then-restore pipeline for dynamic videos with atmospheric turbulence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25286–25296, 2024. 2, 5
- [28] Ripon Kumar Saha, Esen Salcin, Jihoo Kim, Joseph Smith, and Suren Jayasuriya. Turbulence strength c_n^2 estimation from video using physics-based deep learning. *Optics Express*, 30(22):40854–40870, 2022. 2, 3
- [29] Aleksandr V Sergeev and Michael C Roggemann. Monitoring the statistics of turbulence: Fried parameter estimation from the wavefront sensor measurements. *Applied Optics*, 50(20):3519–3528, 2011. 2
- [30] Hossein Talebi and Peyman Milanfar. Nima: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):3998–4011, 2018. 3, 5, 6
- [31] Mingxing Tan and Quoc Le. Efficientnetv2: Smaller models and faster training. In *International Conference on Machine Learning*, pages 10096–10106. PMLR, 2021. 5, 6
- [32] Yuandong Tian, Srinivasa G Narasimhan, and Alan J Vannevel. Depth from optical turbulence. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 246–253. IEEE, 2012. 2, 3, 4
- [33] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 2555–2563, 2023. 3, 5, 6
- [34] Yao Wang and Sukanta Basu. Using an artificial neural network approach to estimate surface-layer optical turbulence at mauna loa, hawaii. *Optics Letters*, 41(10):2334–2337, 2016. 3
- [35] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph Gonzalez, Kurt Keutzer, and Peter Vajda. Visual transformers: Token-based image representation and processing for computer vision. *arXiv preprint arXiv:2006.03677*, 2020. 5, 6
- [36] Manman Xu, Shiyong Shao, Ningquan Weng, and Qing Liu. Analysis of the optical turbulence model using meteorological data. *Remote Sensing*, 14(13):3085, 2022. 3
- [37] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv:2406.09414*, 2024. 4
- [38] Rajeev Yasarla and Vishal M Patel. Learning to restore images degraded by atmospheric turbulence using uncertainty. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 1694–1698. IEEE, 2021. 2
- [39] Lydia Yatcheva, Rui Barros, Max Segel, Detlev Sprung, Erik Sucher, Christian Eisele, and Szymon Gladysz. Ultimate turbulence experiment: simultaneous measurements of c_n^2 near the ground using six devices and eight methods. In *Optics in Atmospheric Propagation and Adaptive Systems XVIII*, volume 9641, pages 28–35. SPIE, 2015. 2
- [40] Zhenqiang Ying, Haoran Niu, Praful Gupta, Dhruv Mahajan, Deepti Ghadiyaram, and Alan Bovik. From patches to pictures (paq-2-piq): Mapping the perceptual space of picture quality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3575–3585, 2020. 3, 5, 6
- [41] Steve Zamek and Yitzhak Yitzhaky. Turbulence strength estimation and super-resolution from an arbitrary set of atmospherically degraded images. In *Atmospheric Optical Modeling, Measurement, and Simulation II*, volume 6303, pages 27–37. SPIE, 2006. 2
- [42] Steve Zamek and Yitzhak Yitzhaky. Turbulence strength estimation from an arbitrary set of atmospherically degraded images. *JOSA A*, 23(12):3106–3113, 2006. 2, 3, 5
- [43] Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*, 24(8):2579–2591, 2015. 3, 5, 6
- [44] Weixia Zhang, Kede Ma, Jia Yan, Dexiang Deng, and Zhou Wang. Blind image quality assessment using a deep bilinear convolutional neural network. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(1):36–47, 2018. 3, 5, 6