Fisher Discriminant Analysis for Extracting Interpretable Phenological Information from Multivariate Time Series Data

Authors: Conor T. Doherty[a], Meagan S. Mauter[a]

[a] Department of Civil and Environmental Engineering, Stanford University, Jerry Yang & Akiko Yamazaki Environment & Energy Building, 473 Via Ortega, Stanford, CA 94305

Corresponding Author: Conor T. Doherty, dohertyc@stanford.edu

NOTE ON THE STATUS OF THIS MANUSCRIPT: This manuscript has been submitted to REMOTE SENSING OF ENVIRONMENT. It has not yet undergone peer review and has not yet been accepted for publication. Subsequent versions of the manuscript may differ from this version. If accepted, the final version of the manuscript will available via a publication DOI link. The authors welcome feedback on the current version of the manuscript.

Abstract

For many applications in environmental remote sensing, the interpretation of a given measurement depends strongly on what time of year the measurement was taken. This is particularly the case for phenology studies concerned with identifying when plant developmental transitions occur, but it is also true for a wide range of applications including vegetation species classification, crop yield estimation, and more. This study explores the use of Fisher Discriminant Analysis (FDA) as a method for extracting time-resolved information from multivariate environmental time series data. FDA is useful because it can be applied to multivariate input data and, for phenological estimation problems, produces a transformation that is physically interpretable. This work contains both theoretical and applied components. First, we use FDA to demonstrate the time-resolved nature of phenological information. Where curve-fitting and other commonly used data transformations that are sensitive to variation throughout a full time series, we show how FDA identifies application-relevant variation in specific variables at specific points in time. Next, we apply FDA to estimate county-average corn planting dates in the United States corn belt. We find that using multivariate data inputs can reduce prediction RMSE (in days) by 20% relative to models using only univariate inputs. We also compare FDA (which is linear) to nonlinear planting date estimation models based on curve-fitting and random forest estimators. We find that multivariate FDA models significantly improve on univariate curve-fitting and have comparable performance when using the same univariate inputs (despite the linearity of FDA). We also find that FDA-based approaches have lower RMSE than random forest in all configurations. Finally, we interpret FDA coefficients for individual measurements sensitive to vegetation density, land surface temperature, and soil moisture by relating them to physical mechanisms indicative of earlier or later planting.

1. Introduction

Many applications of environmental remote sensing rely on methods for extracting information from a time series of measurements of one or more variables. Analyses of phenology, which Lieth (1974) defines as the study of recurring plant life cycle stages, estimate the timing of a phenological event or predict the value of another variable known to be associated with phenology. Past work has comprehensively reviewed phenological information extraction (PIE) methods, finding that subtle differences in assumptions and methodologies can significantly alter estimates of the timing of phenological events (Zeng et al. 2020). There are many combinations of applications and PIE methods, making it difficult to evaluate them all solely on the basis of empirical comparisons without an overarching theoretical framework that can explain when and why particular methods are likely to be effective.

This paper proposes Fisher Discriminant Analysis (FDA) as both a practical method for PIE and as a mathematical framework for analyzing a broad class of PIE methods. Given a set of observations with associated class labels, FDA produces a linear transformation that maximizes the ratio of between-class variance to within-class variance. Intuitively, the aim is to concentrate observations with similar class labels close to one another and to separate the centroids of dissimilar classes in the transformed space. Class labels can be realizations of any variable of interest including categorical variables, like crop type or irrigation status, or variables defined over ordered sets, like planting date or yield. For a problem with k distinct classes, FDA gives a transformation into spaces with dimension up to $k - 1$. In this paper, we will focus exclusively on projections into one dimension because this type of transformation has a clear interpretation for PIE problems, where the one-dimensional space corresponds to time.

In addition to a theoretical analysis of linear transformations for PIE, we demonstrate the utility of FDA by applying it to estimate county average planting dates for corn in the United States Corn Belt. Crop planting date estimation is a useful test of PIE methods for several reasons. First, it exemplifies a class of application where structure in time is a defining feature of the problem. This is both because the response variable itself is time-resolved, and also because the interpretation of measurements of covariates depends strongly on when (e.g. what day of year) the measurements were taken. Second, the application is well-suited to multivariate analysis because the relationship between planting date and vegetation development is influenced by multiple different physical variables. For example, variables like soil temperature, air temperature, and soil moisture affect the viability and rate of development of corn after planting (Abendroth et al. 2011, Watts 1972). Both temperature and soil moisture are physical variables not captured by a single vegetation index (VI) time series.

Beyond its value as a case study for PIE methodology, planting date is also a variable of environmental and economic interest. Planting date influences other important variables including water demand and expected yield, given that these variables depend on the timing of vegetation development relative to the growing season. Relatedly, planting date is also a fundamental variable through which climate and weather events influence agriculture. For example, an unusually wet spring in 2019 delayed planting in much of the US Corn Belt (Rippey, 2019).

Many prior efforts to estimate crop planting dates, and extract phenological information more generally, apply variants of the following method. First, observations from the visible and near-infrared spectra are combined into a vegetation index (VI) and a function of predetermined form is fit to the resulting univariate time series. Parameter estimates from the fitted curve are then used to estimate the timing of phenological milestones, such as "peak vegetation density" or the point in time when a crop "greens up." There are several curve-fitting functional forms that have been used (see Zeng et al. (2020) for examples), but a common approach is to use one or more logistic curves to model vegetation development. This approach has been used to identify phenological transition points such as the onset of vegetation green-up (Guan et al., 2014; Wardlow et al., 2006; Zhang et al. 2003). The estimated timings of these transition point are also used as a proxy for unobserved events like crop planting, where the date the crop was planted is assumed to be correlated with some point on the fitted curve (Lobell et al., 2013; Urban et al., 2018). Other approaches use the timing of the phenological event as one of multiple sources of information. For example, Dong et al. (2019) used temperature measurements from ground weather stations to calculate growing degree days, which they used to complement the information extracted by VI curve-fitting in order to estimate canola planting dates.

In this paper, we present FDA both as a method that can be applied to remote sensing problems in practice, and also as a tool for analyzing the structure of environmental time series. To demonstrate its practical utility for phenological information extraction, we apply FDA to the problem of corn planting date estimation using multivariate time series of measurements sensitive to vegetation density, land surface temperature, and soil moisture. For context, we compare FDA with two other methods that are widely used in environmental remote sensing: VI

curve fitting, which is easily interpretable but univariate, and Random Forest, which is multivariate but less interpretable. Before getting into the application, we describe what FDA is, how it works, and what it can tell us about the structure of environmental time series.

2. Theory

2.1 Background

Remotely sensed measeurements are noisy and often are not direct observations of the target variable of interest. Zeng et al. (2020) describe approaches that deal with these challenges as "data smoothing methods" and "phenology extraction methods," respectively. Curve-fitting is a form of data smoothing but is also used for information extraction, where the timing of phenological transition points are estimated using some function of curve-fitting parameter estimates.

For smoothing and information extraction methods that can be expressed as linear transformations, there is an implied mathematical framework through which we can reason about many type of methods jointly; that is, transformations that can be expressed as matrix-vector products. Many methods commonly used to process remotely sensed time series data fall under this category, including approaches based on the Discrete Fourier Transform (Adams et al. 2020; Filippelli et al. 2020; Jakubauskas et al. 2002; Mingwei et al. 2008; Wang et al. 2019), Discrete Cosine Transform (Garcia 2010; Guan et al. 2014; Urban et al. 2018), Discrete Wavelet Transforms (Sakamoto et al. 2005), Savitsky-Golay filter (Chen et al. 2004; Kandasamy et al. 2013), Whittaker filter (Atzberger and Eilers 2010; Kandasamy et al. 2013), and others. It should be noted that nonlinear methods are also commonly used in phenology analysis; logistic curve-

fitting is a nonlinear transformation, for example. However, linear transformations encompass a large set of commonly used methods and focusing on them is a mathematically tractable step toward formulating a more general analytical framework. In addition, linear methods offer other benefits such as interpretability, reduced computational complexity, and often having well-understood analytical solutions.

Many data transformations that are commonly used in remote sensing rely on the assumption that environmental time series have underlying structure that can be approximated using a low dimension representation. For example, suppose we have n observational time series $x_1, \dots, x_n$ where $x_i \in \mathbb{R}^D$, each composed of measurements of D unique combinations of measurement types (e.g. spectral bands or VIs) and timesteps. Then we can define a transformation P: $\mathbb{R}^D \rightarrow \mathbb{R}^d$ that projects the data into a lower dimensional space spanned by $d < D$ orthonormal vectors (such that the space is isomorphic to $\mathbb{R}^d$). We can also define a second transformation R: $\mathbb{R}^d \rightarrow \mathbb{R}^D$ that reconstructs the time series in the original dimension. Many data smoothing transformations, for example, can be expressed as $R \circ P(x_i)$, where $\circ$ denotes composition. For linear transformations, this is simply RPx_i, where R and P are matrices and x_i is the time series written as a column vector.

While it is a common practice to apply an initial smoothing step that reconstructs the time series in the original dimension, the same information is contained in the lower dimensional representation in $\mathbb{R}^d$ as in the reconstruction. This is the logic underlying the method used by Jakubauskas et al. (2002) and Wang et al. (2019), for example, where they project remotely sensed time series data onto a select set of Fourier modes (frequencies) and then use the

coordinates in the lower dimensional space as features for machine learning models. This procedure is equivalent to applying a DFT-based smoothing filter, but without applying the inverse transform to return the time series to its full dimension in the time domain.

The low-dimensional representation of a time series, and corresponding information extracted, is influenced by the choice of basis vectors onto which the time series are projected. Many prior DFT-based approaches apply a form of low pass filter that preserves the low frequency components and removes the high frequency "noise." However, there is no prescribed canonical set of frequencies associated with this task. Even if the transformation projects the data onto predefined basis vectors from a generic dictionary (DFT, wavelet, etc.), this still leaves an infinite "design space" for the transformation. The selection of the particular basis vectors, from a generic dictionary or otherwise, must be informed by some combination of prior knowledge, assumptions, and labeled data. However, there is no reason to assume *a priori* that a given generic basis will be optimal for a particular task. An alternative to selecting basis vectors from a generic dictionary is to select a basis computed from the measurement data themselves.

Data-derived transformations work by defining a criterion that quantifies some property of the data, and then computing the vectors that maximize (or minimize) the criterion. The resulting transformation preserves or amplifies certain features of the data corresponding to the criterion. For example, one of the most widely used methods in data-intensive sciences is Principal Component Analysis (PCA), which finds the set of orthogonal directions over which the data have maximal variance. By construction, PCA minimizes squared reconstruction error among all possible d-dimensional orthogonal linear encodings for a given dataset because it preserves as

much variance as is possible using only d orthogonal vectors. This property makes PCA well-suited to general purpose dimensionality reduction and data compression. However, it does not necessarily mean that it will be useful for PIE.

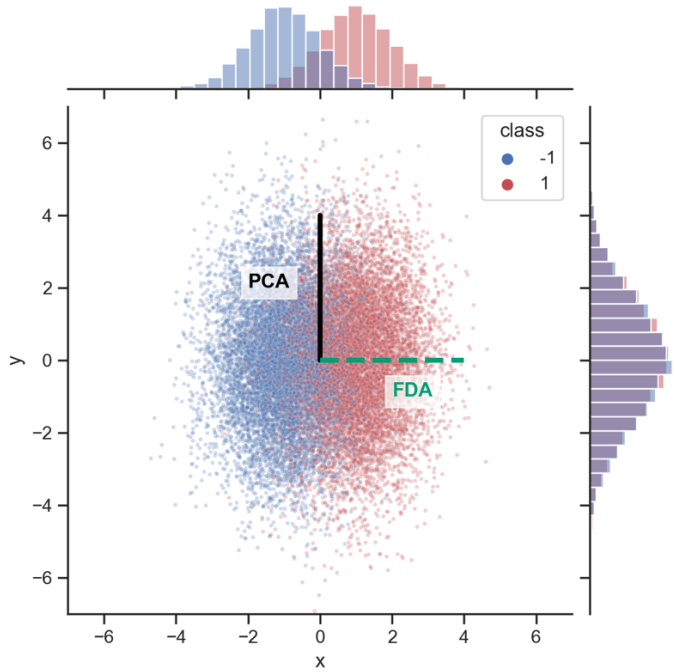

Figure 1*: Example of PCA and FDA for 2-dimensional simulated data. Data are generated from two Gaussian distributions that differ only in the x-coordinate of the mean (y-coordinate of mean, covariance matrices are identical for both classes). Distributions are specified such that the PCA vector is parallel to the y-axis and the FDA vector is parallel to the x-axis. The histogram on the left shows the marginal distribution of the y-coordinates, which is equivalent to the distribution of the data projected onto the PCA vector. The histogram on the top shows the marginal distribution of the x-coordinates, which is equivalent to the distribution of the data projected onto the FDA vector. Projection onto the PCA vector preserves more total variance, but projection onto the FDA vector preserves more (all) of the variance that differentiates the two classes.*

The utility of a data-derived transformation depends on the application for which it is used. Figure 1 provides a simple illustration of PCA and FDA using simulated data. The data are drawn from two Gaussian distributions that differ only in the x coordinate of their mean. The two

classes have identical covariances and are specified such that the PCA vector is parallel to the y-axis and the FDA vector is parallel to the x-axis. The histogram on the righthand side is the marginal distribution of the y-coordinates of the data, which is equivalent to the distribution of the data projected onto the PCA vector. This is the direction with the most variance and, as such, will result in the smallest reconstruction error. However, looking at the histogram on the right, we can see that projection onto the PCA vector makes it impossible to distinguish between the two classes as the marginal distributions are identical. The histogram on top of the scatter plot is the marginal distribution of the x-coordinates, which is also the distribution of the data project onto the FDA vector. Both classes have less variance in this direction, and so using the FDA vector will result in a larger reconstruction error than the PCA vector. However, this is the direction that maximizes the "separation" of the two classes, as evidenced by the distinctly multi-modal marginal distribution. If the goal is to predict which class an unlabeled data point belongs to, then projecting onto the FDA vector is clearly more useful than projecting onto the PCA vector.

For PIE tasks, we must define a transformation that facilitates distinguishing between time series that differ in some phenological variable. In real environmental remote sensing applications, the data will have much higher dimension and determining what constitutes "useful information" will be much less obvious than the toy example in Figure 1. FDA provides a principled approach to representing data in a low dimensional space such that similar observations are close one another and dissimilar observations are far apart.

2.2 Mathematical formulation of FDA

The following section summarizes the mathematical formulation of FDA. Suppose the data consist of n observations each of length D, each with an associated class label from a set of k classes where class j contains n_j elements. Then we define the following:

$$\text{Observation matrix: } X \in \mathbb{R}^{n \times D}$$

$$\text{Overall mean: } \mathbb{R}^D \ni \mu = \frac{1}{n} \sum_{i=1}^{n} x_i$$

$$\text{Class means: } \mathbb{R}^D \ni \mu_j = \frac{1}{n_j} \sum_{i=1}^{n_j} x_i^{(j)}$$

$$\text{Between-class scatter: } \mathbb{R}^{D \times D} \ni S_b = \sum_{j=1}^{k} n_j \left(\mu_j - \mu\right) (\mu_j - \mu)^T$$

$$\text{Within-class scatter: } \mathbb{R}^{D \times D} \ni S_w = \sum_{j=1}^{k} \sum_{i=0}^{n_j} \left(x_i^{(j)} - \mu_j\right) (x_i^{(j)} - \mu_j)^T$$

FDA is defined as the linear transformation that maximizes the ratio of between-class scatter (BCS) to within-class scatter (WCS). Intuitively, it seeks to transform the data such that observations within a class are tightly clustered around a centroid, and that the class centroids are far away the overall mean. The transformation given by FDA is "optimal" in the sense of maximizing the following criterion:

Find $U \in \mathbb{R}^{D \times k}$ that satisfies: $\max_U \frac{\mathrm{tr}(U^T S_b U)}{\mathrm{tr}(U^T S_w U)}$, where *tr* is the matrix trace operator

We will refer to this objective function as the "Fisher Criterion." This optimization problem has an analytical solution and is equivalent to a solving the generalized eigenvalue-eigenvector problem (Ghojogh et al. 2019) of the form:

$$S_b U = S_W U \Lambda$$

Where Λ is a diagonal matrix with the generalized eigenvalues along the diagonal and the columns of $U \in \mathbb{R}^{D \times k}$ define an orthonormal basis. This system is solved using standard functions from a numerical programming library.

We extend the basic formulation of FDA slightly by applying a simple form of regularization. There are several ways of implementing regularized FDA, but we select one originally described by Friedman (1989) in which we replace S_b with $S_b + \alpha I$, where α is a small scalar and I is the identity matrix. Zhang et al. (2010) derive formal equivalences between regularized FDA and ridge regression, which further justifies our application of FDA to an ordered response variable. However, there is also a simpler and more intuitive interpretation of how FDA works when applied to the planting date prediction problem.

2.3 Interpretation of FDA

FDA can produce interpretable transformations for a range of problems, particularly for those where the response space is ordered and one-dimensional. While the general notion of an "optimal" transformation is fairly abstract, considering specific examples will make it more concrete. It will also illustrate connections with transformations based on generic dictionaries like the DFT. Figure 2 shows five examples of pairs of classes of time series and their corresponding WCS matrix, BCS matrix, and first FDA vector. For each example, we specify two classes of random vectors with known jointly Gaussian distributions. Properties of the

distributions, and which properties differentiate the classes, vary from example to example. In each example, the covariance structure is specified in a way that emulates autocorrelation in time. While these examples are much simpler than real remotely sensed measurement time series, they are carefully chosen to illustrate both theoretical and practical considerations for phenological information extraction.

The first three rows are examples where a given DFT frequency mode is optimal with respect to the Fisher Criterion. That is, the first FDA vector corresponds with the DFT mode with the highest spectral density. (The correspondence in not exact equality in the third example because of the simulated autocorrelation, although it would be exact if the covariance was spherical.) In each of these three examples, the class means μ_k are constant or periodic in time (constant can also be thought of as periodic with frequency zero). The variance is also constant in these first three examples.

In Figure 2A, the class means do not vary in time. The first FDA vector is also constant, which can also be thought of as the constant column of the DFT matrix. Multiplication by this vector is equivalent to summing (or averaging) the all measurements in the time series. Time series like these, where neither the mean nor variance depend on time, are known as “covariance stationary” (Gray 2006). The BCS matrix is approximately constant, up to numerical precision, prior to regularization because there is no meaningful cross-class covariance.

In Figure 2B, the two class means vary with the same frequency and have the same amplitude but are out of phase with one another. While these classes cannot be separated with simple

averaging like in Figure 2A, they can be easily distinguished when projected onto the DFT mode corresponding to their frequency (five oscillations per sample period); one class will give positive values and the other will give negative values. The BCS matrix has a checkerboard pattern because the class means always have opposite signs and alternate at each timestep.

In Figure 2C, the class means vary with the same frequency and are in phase with one another but have different amplitudes. Similar to Figure 2B, the classes can be separated by projecting onto a cosine with the proper frequency (one oscillation per period). However, in the case of Figure 2C, that the first FDA vector is not exactly a cosine due to the nonzero off-diagonal covariance. It would be exact if all off-diagonal covariances were zero, but this scenario is unlikely with time series data where autocorrelation in the time is common. The BCS matrix for Figure 2C looks like the BCS matrix for Figure 2B but “zoomed in”, due to the time series in Figure 2C having lower frequency than the ones in Figure 2B. Looking at Figures 1A, 1B, and 1C, we can see that DFT modes can separate classes that differ in amplitude, phase, or mean.

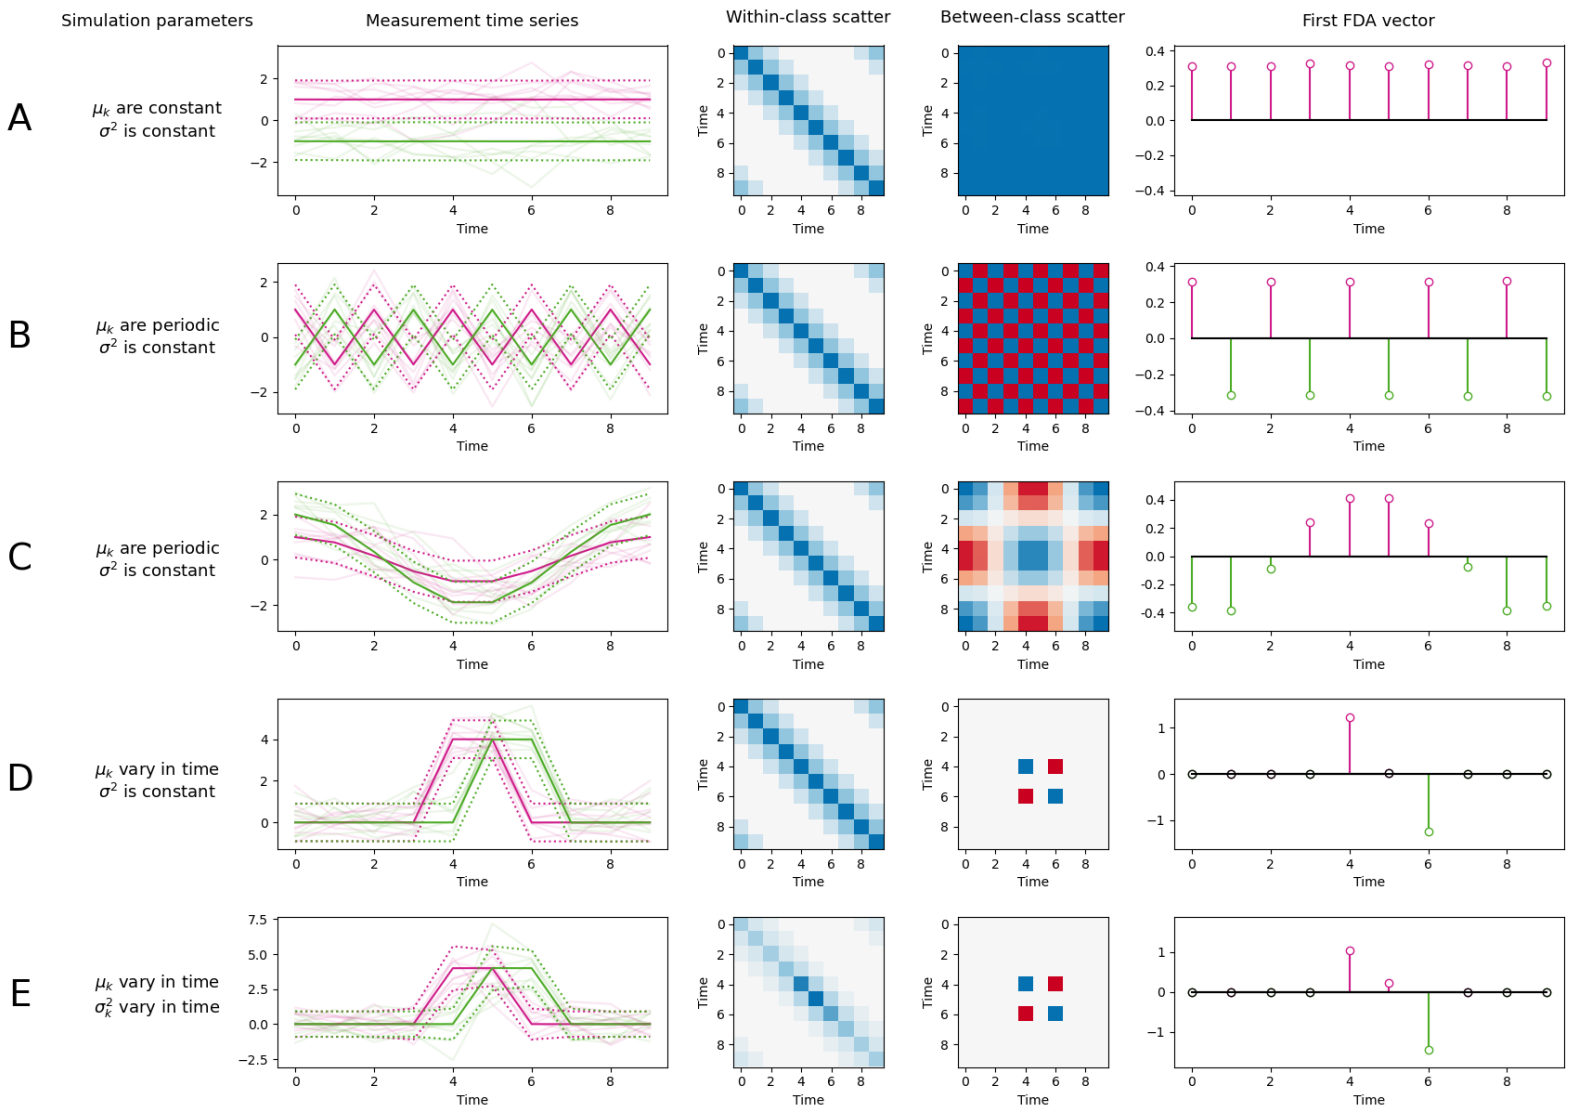

***Figure 2**: Example time series with corresponding scatter matrices and FDA vectors. Each row shows two classes of time series that differ in some respect, the corresponding WCS and BCS matrices, and the FDA vector that separates the classes in 1-dimensions space. The "measurement time series" column shows the mean (solid line) and 0.25/0.75 quantiles (dotted line) of each class. Row A shows classes that differ in (constant) mean. Row B shows classes with periodic means that differ in initial phase. Row C shows classes with periodic means that differ in amplitude. The FDA vectors for rows A-C correspond to DFT (frequency) modes. Rows D-E are more similar to real intra-annual time series in that the means and (for Row E) variance depend idiosyncratically on time. The FDA vectors for Rows D-E are only nonzero at two (Row D) or three (Row E) points in time because these are the only observations that are "informative" in terms of distinguishing between classes.*

Each of the simulated time series in Figures 1A-1C are specifically designed to produce FDA vectors that are (exactly or approximately) equal to DFT frequency modes. Similar reverse engineering could be performed to construct time series corresponding to, for example, a wavelet function or any other element of a generic basis. However, when working with real data, it is not known ahead of time which vectors from a given generic dictionary are best suited to distinguishing between classes of time series. In practice, one must either select basis vectors using prior knowledge or select from among a set of candidate vectors using the measurement data themselves. For example, Jakubauskas et al. (2002) consider a set of DFT modes where the lowest frequency (excluding the constant term) completes one oscillation per period. Wang et al. (2018) expand the set of candidate basis vectors to include lower frequencies that complete less than one oscillation per period. They then perform supervised feature selection on the candidate frequencies using labelled data and statistical tests of feature relevance. These examples show how, even when using generic transformations, there is often a degree of "supervision" involved in selecting the transformation best suited to the particular data and task.

The examples in Figures 1D and 1E show time series where, like real remotely sensed environmental time series, properties of the data depend strongly on time. In both examples, the simulated time series have expected value of zero for most of the duration but have nonzero expected value for particular timesteps. These simulated time series are inspired by, though simpler than, typical vegetation index time series for agricultural fields. For example, the expected NDVI in April is different from the expected NDVI in July for corn fields in the US Midwest. In both Figures 1D and 1E, the two classes differ in the timing of observations with

nonzero expected value. This is similar to how one might expect, for example, NDVI in early July to be different for a field planted in early April versus a field planted in mid-May.

Figure 2D shows an example of time series with time-varying mean, but that still have time-invariant covariance structure. The two classes differ only in mean; they have identical covariance matrices. In fact, the covariance matrices of the classes in Figure 2D are identical to those used to simulate the time series in Figures 1A-1C. This is why the WCS matrices are the same in Figures 1A-1D. In the BCS matrix, we can see that the only meaningful cross-class variance occurs at timesteps four and six. These are the only two timesteps where the expected values of the two classes are not equal to one another. Accordingly, the FDA vector is zero everywhere except for at timesteps four and six. The logic in this example translates to applications with real remote sensing data. We expect the variation in measurements stemming from planting date to be confined to a relatively narrow portion of the year, specifically certain parts of the growing season. Measurements from outside this time window will not have significant weights, positive or negative, because they contain little to no information about when the crop was planted.

The example in Figure 2E is the most similar to real agricultural time series because both the mean and variance depend on time. Time-dependent variance is a fundamental feature of agricultural time series. For example, one would expect the NDVI variance to be relatively high during the early part of the growing season, when some fields are “ahead” of others in terms of vegetation development, or at the end of the season, when some fields have been harvested and some have not. In contrast, one would expect the variance to be lower during the middle of the

season, when most fields have a full canopy. Changes in variances are reflected in Figure 2E in the width of the quantile bands, which are narrower at timesteps where the mean is zero and wider at the timesteps with nonzero mean. Because Figure 2E is the only example with time-varying variance, its WCS matrix is different from the WCS matrices in Figures 1A-1D. However, because Figure 2E has the same sequence of means as Figure 2D, these two figures have the same BCS matrices. Comparing Figure 2D to Figure 2E shows how the FDA vector is a function of both the WCS and BCS. Like Figure 2D, the FDA vector in Figure 2E is zero at most timesteps where the observations are not informative for distinguishing between the classes. Unlike Figure 2D, the FDA vector value at timestep five is nonzero, despite the classes having the same mean at this timestep, due to the effects of the time dependence of the variance. The FDA vector values at timesteps four and six do not have the same magnitude, unlike in Figure 2D, again due to differing covariance structure of the data generating process.

An important takeaway from Figure 2 is the fundamental, and visually apparent, differences between Figures 1A-1C and Figures 1D-1E. The BCS matrices in Figures 1A-1C have translational symmetry, which corresponds to time invariance. In contrast, the BCS matrices in Figures 1D-1E have no nontrivial translational symmetries; the scatter matrices and corresponding FDA vectors are idiosyncratic with respect to time. Strong dependence on time (or time-of-year) is a fundamental characteristic of many types of environmental data, and particularly of agricultural time series. This is not to say that generic transformations based on, for example, the DFT cannot be beneficially utilized for environmental time series analysis. However, to be useful, one must determine which frequency mode or modes (columns of DFT matrix) are well-suited to the task using the measurement data themselves. This limits the utility

of generic transformations since the transformation still must be informed by the measurement data. Moreover, there is no reason to assume *a priori* that any given DFT mode, wavelet mode, or any other vector from a generic dictionary will optimally discriminate between classes of time series. In fact, for linear transformations, the "best-case" scenario is that a generic linear transformation will be equivalent to the FDA transformation for a given application.

3. Methods

In this empirical study, we test combinations of remotely sensed time series measurements, data transformations, and estimation methods on the problem of county-year average planting date estimation. We intend for the implementations of the various methods to approximate "average" model performance, while acknowledging that predictive models can often be more finely tuned. Model tuning can take many forms including predictor variable selection (spectral bands, days-of-year), data transformations (VI construction, smoothing, dimensionality reduction, "feature engineering"), and hyperparameter selection (regularization, algorithm-specific parameters). That said, our implementations are clearly defined, reproducible, and apply established methods using publicly available data and software. The primary purpose of the empirical study is to test whether different types of remotely sensed measurement time series can be combined in a multivariate model to improve prediction accuracy. The accuracy metrics from the various approaches provide context for one another but are not claimed to demonstrate that any one method is superior for all real-world applications.

3.1 Remote Sensing Data

We use two MODIS data products as inputs to the models: Terra 8-day composite surface reflectance product with 500 m resolution (MOD09A1) (Vermote, Eric, 2015) and Aqua 8-day composite Land Surface Temperature and Emissivity product with 1 km resolution (MYD21A1) (Hulley et al. 2022). We use the land surface temperature product from Aqua due to the dead detector for Terra band 29 starting in 2006. From MOD09A1, we use bands 1, 2, 6, and 7 representing observations in red (620-670 nm), near-infrared (841-876 nm), and two shortwave infrared ranges (1628-1652 nm, 2105-2155 nm), respectively. We use bands 1 and 2 to construct the Normalized Difference Vegetation Index (NDVI) as (B2-B1)/(B2+B1). We use bands 6 and 7 to construct an index from Yue et al. (2019) as (B6-B7)/B6, which is sensitive to variation in soil moisture. Following Yue et al., we refer to this index as Normalized Shortwave-infrared Difference Soil moisture Index (NSDSI). From MYD21A1, we use the bands corresponding to daytime and nighttime land surface temperature. The LST values, expressed in degrees Kelvin, are divided by 100 to be on the same order of magnitude and NDVI and NSDSI. This avoids numerical issues and does not change the affine relationships between variables. For both MOD09A1 and MYD21A1, we filter the data using the provided quality assurance bands. For MOD09A1, we mask pixels that are marked as containing cloud, cloud shadow, or cirrus, or that are not marked as “highest quality” for a given spectral band. For MYD21A1, we kept pixels marked as having errors of less than or equal to 2 K (“marginal” performance or better).

3.2 County-Year Aggregation

We extract MODIS pixels believed to be majority corn before averaging pixel values within county boundaries. To do this, we first use the USDA Cropland Data Layer (CDL) (USDA-

NASS, 2022) to create a mask of a mask of corn versus "not corn" pixels in the native resolution of the CDL (30 m). Then we resample this raster to the MODIS grid (500 m for MOD09A1, 1 km for MOD11A1) using 'mode' resampling. This procedure returns a large (i.e. MODIS-scale) corn pixel only if the pixel is majority corn. The resampling and reprojection were performed using the `gdalwarp` function from Geospatial Data Abstraction Library (GDAL/OGR contributors, 2022). The remaining pixels are averaged within county boundaries. We work with this filtered data in an attempt to reduce the effects of unobserved confounding factors (like the ratio of corn to soy in a county, as observed by Urban et al. 2018) that do not apply at field scale. Finally, we exclude county-year combinations with any observations with fewer than 10 pixels after quality assurance filtering and crop masking.

3.3 Ground Truth Data

We use data published by Lobell et al. (2014) that provides planting date identified at the county level. The dataset contains planting dates for 100 anonymous, randomly sampled fields per county per year in Indiana, Illinois, and Iowa (Lobell et al., 2014). To train and test our models, we used 10 years of observations from 2003 to 2012 (2003 is the first full season of MYD21 data, and the Lobell et al. dataset runs through 2012). Planting date varies at the field level, not the county level, so this training and validation strategy assumes that the relationship between planting date and remote sensing observations is at least partially preserved when averaging both at the county scale. Urban et al. (2018) present empirical evidence supporting the validity of this assumption.

3.4 Data Partitioning and Validation

This planting date dataset exhibits significant variation in both space and time relative to its size, so we employ time- and space-aware validation schemes in order to estimate "average" model performance. We test two different cross-validation schemes, both of which have the property that the training and test sets have no overlap in space or time. That is for a given record corresponding to county C in year Y in the test set, the training set will not contain county C in any year nor will it contain any observations from year Y. In both schemes, we partition the data by year, where the test samples are drawn from a single year and the training samples are drawn from the remaining 9 years. We spatially partition the data in two different ways: by state (Iowa, Illinois, Indiana) and randomly. We present results for all models using both partitioning schemes. Partitioning by state tests the models' ability to transfer to a nearby but distinct region (we call this "partitioned-by-state" prediction). Partitioning counties randomly is, intuitively, an "easier" prediction problem because the training sample should be more similar to the test sample with respect to geographic variation. However, this "randomly-partitioned" prediction problem is also closer to many real-world applications where the task is to estimate planting dates for a specific agricultural region for which prior measurement data are available.

Every modeling approach we test estimates some parameter values from the training data. The curve-fitting approach estimates the timing of an observable proxy for planting (e.g. vegetation emergence or "green-up") rather than the actual planting event itself. It is still necessary to estimate planting date as a function of this proxy variable (even if "the function" is simply adding a fixed number of days). For both FDA and curve fitting, we model the projected time

series data as being linearly related to planting date. To fit the final step of the predictive model, we estimate parameters corresponding to a scalar coefficient and constant (intercept) using Ordinary Least Squares regression. One primary difference between curve-fitting and FDA is that the curve-fitting generally applies a predetermined curve functional form whereas FDA computes the transformation from training data.

For each model and validation scheme, we compute a number of summary statistics including root mean squared error (RMSE), bias, and unbiased root mean squared error (ubRMSE) as a function of the predicted planting DOY p_i and the true planting DOY t_i for $i = 1, \dots, N$.

$$RMSE = \sqrt{\frac{1}{N}\sum_{i}^{N}(p_i - t_i)^2}$$

$$Bias = \frac{1}{N}\sum_{i}^{N} p_i - t_i$$

$$ubRMSE = \sqrt{\frac{1}{N}\sum_{i}^{N}\left((p_i - \overline{p_i}) - (t_i - \overline{t_i})\right)^2}$$

RMSE can be decomposed into bias and ubRMSE using the relation $RMSE^2 = Bias^2 + ubRMSE^2$ (Entekhabi et al. 2009). For an unbiased predictor, ubRMSE = RMSE.

3.5 Method Implementation Details

3.5.1 Fisher Discriminant Analysis (FDA)

We implement regularized FDA as described in the Theory section, using `eig` from `scipy.linalg` to compute the generalized eigenvectors (SciPy 1.0 Contributors, 2020). We define a regularization parameter α where the FDA directions are computed as the generalized eigenvectors of the pair of matrices $(S_w, S_b + \alpha I)$, where α is a scalar and I is identity matrix. The model is not overly sensitive to α, which is similar to the regularization parameter in ridge regression; see Zhang et al. (2010) for a more in-depth discussion. We use the value $\alpha = 1$ for all FDA models, having considered the values 0.1, 1, and 10.

For input data, we select contiguous (in time) sequences of measurements for four variables: NDVI, NSDSI, and daytime and nighttime land surface temperate (dLST and nLST, respectively). We select temporal ranges for each variable for which we can interpret NDVI, NSDSI, and LST as being sensitive to land surface characteristics including vegetation density, moisture, and temperature, respectively. We use measurements ranging over the following composite periods (identified by composite start day-of year): 121-225 (early May to mid-August) for NDVI and 97-129 (early April to early May) for NSDSI, dLST and nLST. One could take more sophisticated approaches to variable selection, like applying a stepwise procedure or penalized estimation, and possibly achieve higher test accuracy. However, we find our simple approach to be sufficiently robust and also physically defensible. The NDVI measurements are the same data that would be input into a curve-fitting model. The LST measurements cover a time period where temperature may affect the rate of vegetation development. NSDSI

measurements are selected from a window of time where unusually wet soil may lead to a delay in planting.

3.5.2 Univariate Curve Fitting

We implement a curve-fitting planting date estimation model using methods that have been previously validated and are widely used in the field. In particular, we follow Zhang et al. (2003) and Wardlow et al. (2006) and fit a logistic curve to a portion of the NDVI times series that starts at the early part of the growing season; we found that starting the time series at the composite period beginning DOY 113 produced the best results. We truncate each county-year time series at the first point where the slope of a five-observation moving average turns negative. Then we fit a logistic curve, as a function of time, to the remaining observations. To compute the parameter estimates for the logistic model, we use the `minimize` function from the SciPy optimization library `scipy.optimize` with the 'Nelder-Mead' solver. To facilitate convergence of the of the nonlinear program, we select reasonable starting points.

After obtaining parameter estimates for the logistic curve, we compute "metrics" for phenology extraction. We consider several possible metrics from previously published studies. First, we test points where the curve reaches p percent of the range of NDVI during the green-up period (for $p \in \{10, 20, \ldots, 100\}$) (Lobell et al. 2013; Urban et al. 2018). We also test the inflection point in the logistic curve (Urban et al., 2018; Zhang et al., 2003). For each metric, we then estimate the parameters a and b in the linear model $y = ax + b$ using Ordinary Least Squares, where y is the county-year average planting DOY and x is the metric computed by curve-fitting. The parameters estimated on the training set are used to predict planting DOY as a function of metric

values computed on test set. In practice, we find that point of 50% of full green-up has the highest correlation with planting date, which is what we use for the model comparison.

3.5.3 Random Forest

For an additional point of comparison, we tested Random Forest (RF) predictors using the same combinations of univariate and multivariate input data as we tested with FDA. RF predictors (Breiman 2001) are popular, easy to use, and in practice have been shown to perform well across a wide range of tasks (Hastie et al. 2009). Individual decision trees are nonlinear predictors. Predicting labels for new observation consists of aggregating predictions made by an ensemble of trees. While there are methods for quantifying individual variable importance and RFs can be visualized using proximity maps (Hastie et al., 2009), the internal logic of the trained predictor itself is not readily apparent.

We use the `RandomForestRegressor` class from the SciKit-Learn (Pedregosa et al. 2011) ensemble learning library `sklearn.ensemble`, keeping all of the default argument values (100 trees, mean squared error split criterion, no maximum depth). We test the same combinations of input variables as we tested with FDA.

4. Results

We evaluate the performance of the models for each combination of input data, validation scheme, and prediction method. A complete set of summary statistics is given in Table 1.

		Counties partitioned by state			Counties partitioned randomly		
Model type	**Data input**	**RMSE**	**Bias**	**ubRMSE**	**RMSE**	**Bias**	**ubRMSE**
Curve-fitting	NDVI	6.48	0.29	6.47	6.24	0.05	6.24
FDA	NDVI	6.61	0.66	6.57	6.22	0.05	6.22
	NDVI, NSDSI	5.98	0.52	5.96	5.68	0.06	5.68
	NDVI, LST	5.39	0.57	5.36	5.03	- 0.01	5.03
	NDVI, NSDSI, LST	5.28	0.77	5.22	4.81	0.10	4.81
Random Forest	NDVI	6.79	0.27	6.78	6.31	- 0.27	6.30
	NDVI, NSDSI	6.33	0.36	6.31	6.12	- 0.08	6.12
	NDVI, LST	6.27	0.15	6.27	5.91	- 0.43	5.89
	NDVI, NSDSI, LST	5.96	0.18	5.96	5.70	- 0.31	5.69

Table 1*: Error summary statistics for all combinations of model type, input data, and validation scheme. The unit for all errors is "days."*

Figure 3 shows the RMSE for each combination of model type and input data. Figure 3 also includes a "simple mean" bar, which represents the RMSE for a naive model that takes no remotely sensed inputs; for each cross-validation fold, it predicts the label of every record in the test set to be the mean of the training set.

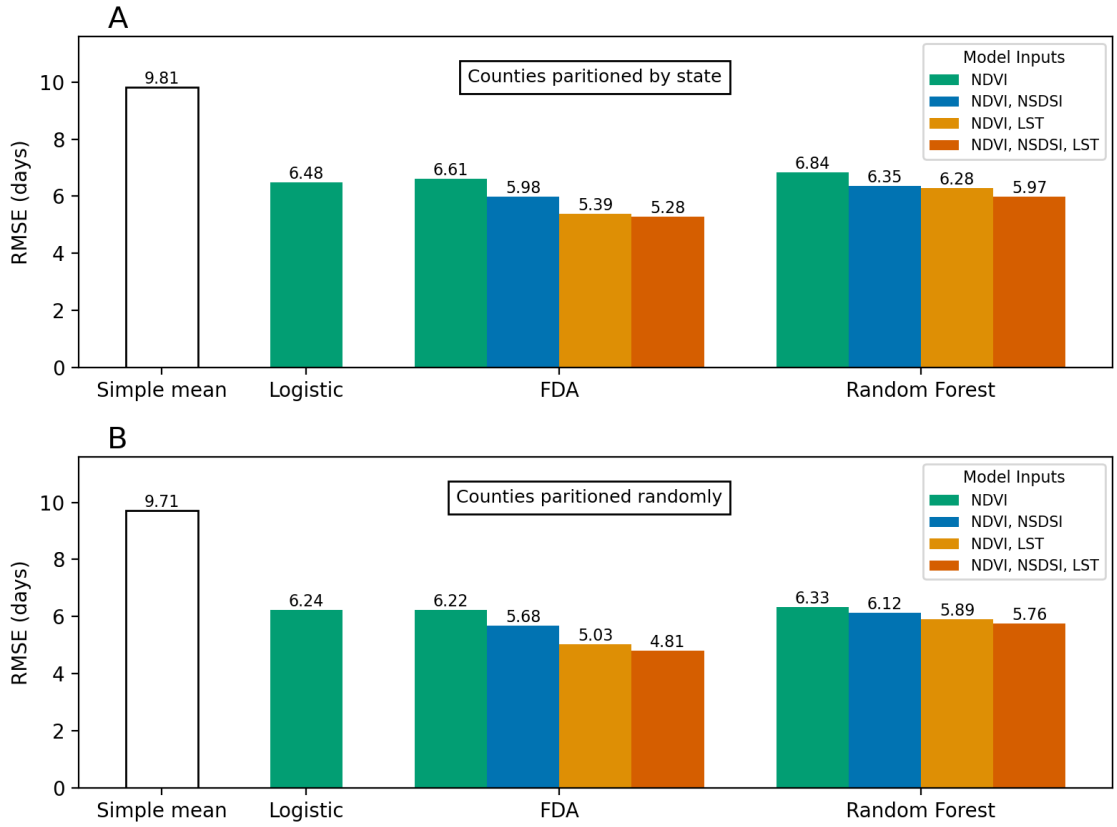

***Figure 3**: RMSE for all model specifications and validation schemes. Results are given for each model type, set of model inputs, and validation scheme. Panel A shows results for the partitioned-by-state scheme. Panel B shows results for the partitioned-randomly scheme. The results indicate that incorporating model inputs beyond NDVI can meaningfully reduce the magnitude of prediction error. This effect is observed with both multivariate predictors (FDA and random forest), but the effect is larger using FDA.*

Figure 3 shows that using additional physically meaningful data inputs, beyond NDVI, can reduce average prediction error. For example, the FDA predictions using all data inputs reduced

RMSE by 1.3 days (20%) and 1.4 days (23%) for the partitioned-by-state and randomly partitioned schemes, respectively, when compared to NDVI-only models. A similar, though smaller, effect can be seen in the random forest results. Using all data inputs reduced RMSE by 0.8 days (12%) and 0.6 days (10%) for the partitioned-by-state and randomly partitioned schemes, respectively, compared to NDVI-only.

When considering the data sample as a whole and without controlling for omitted variables, the magnitude of "random" error dominates the bias error for all models. Even the largest bias observed in the study of 0.77 days was small compared to the corresponding ubRMSE of 5.22 days for that model specification (in the partitioned-by-state scheme). In fact, this model specification (FDA, all data inputs) had the lowest overall RMSE in both validation schemes. This result is indicative of how the introduction of small bias can cause a proportionally larger reduction in random error resulting in a reduction in total error.

The most significant reductions in prediction RMSE from additional data inputs can be attributed to LST reducing the error variance. This phenomenon is apparent in the FDA results and also present, to a lesser extent, in the random forest results. Figure 4 shows the distribution of errors for FDA using different model inputs (similar plots for curve-fitting and random forest are provided in Supplementary Information Section A). The distributions are estimated using kernel density estimation with a Gaussian kernel. In Figure 4, note how the results from models that include LST have distributions with smaller tails and more of the probability mass concentrated in the center. The inclusion of NSDSI also reduces the error variance relative to NDVI-only models, but to a smaller degree than the inclusion of LST inputs does.

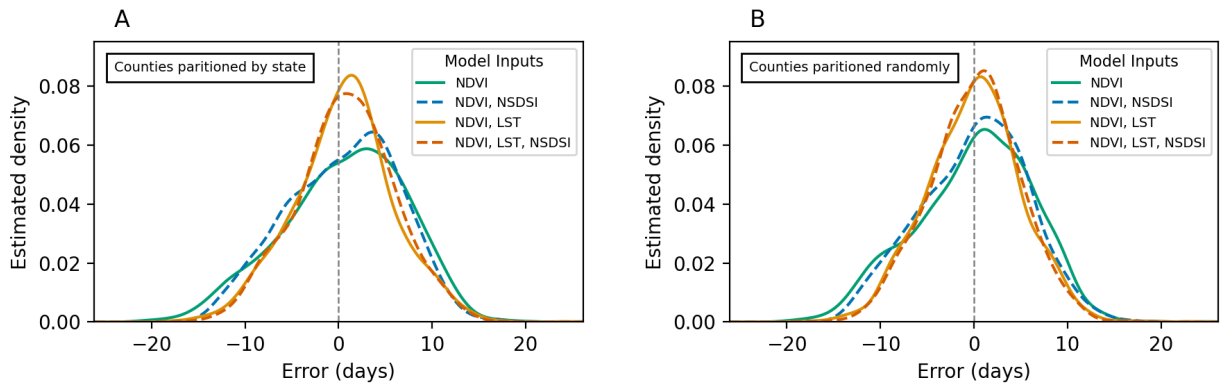

Figure 4*: Error distributions for FDA models. Panel A shows the error distribution for the partitioned-by-state validation scheme. Panel B shows the error distribution for the randomly partitioned scheme. The distributions are generated by kernel density estimation using a Gaussian kernel. Note that the inclusion of LST as a model input reduces the error variance, as indicated by more of the probability mass being concentrated around zero.*

Much of what appears as "random" error in aggregate statistics can be attributed to multiplicative biases associated with omitted variables or with the response variable itself. For example, with all model specifications we observe a negative correlation between planting day and signed error. That is, we observe that all models on average overestimate early planting dates and underestimate late planting dates, with predictions tending toward the mean of a given training distribution. The magnitudes of this multiplicative bias are given in Table 2 in the "Planting DOY" columns.

		Counties partitioned by state			Counties partitioned randomly		
Model type	**Data input**	**Planting DOY**	**Air Temperature**	**Intercept**	**Planting DOY**	**Air Temperature**	**Intercept**
Curve-fitting	NDVI	- 0.51	- 0.46	0.29	- 0.47	- 0.44	0.05
FDA	NDVI	- 0.52	- 0.54	0.66	- 0.45	- 0.53	0.05
	NDVI, NSDSI	- 0.39	- 0.06	0.52	- 0.35	0.00	0.06
	NDVI, LST	- 0.39	- 0.51	0.57	- 0.33	- 0.50	- 0.01
	NDVI, NSDSI, LST	- 0.34	- 0.21	0.77	- 0.29	- 0.17	0.10
Random Forest	NDVI	- 0.48	- 0.39	0.27	- 0.43	- 0.42	- 0.27
	NDVI, NSDSI	- 0.41	0.05	0.36	- 0.36	0.02	- 0.08
	NDVI, LST	- 0.48	- 0.10	0.15	- 0.40	- 0.18	- 0.43
	NDVI, NSDSI, LST	- 0.45	0.03	0.18	- 0.38	- 0.03	- 0.31

Table 2*: Multiplicative biases for each model type, set of data inputs, and validation scheme. Note that Planting DOY and Air Temperature values have been centered (subtracted the mean) so that the intercept is exactly the bias from the previous table.*

We also observe a negative correlation between air temperature and signed error, but that that including NSDSI as a model input reduces or, in one configuration, eliminates this bias. The relationship between VI-based planting date prediction error and air temperature has been observed in previous studies (Urban et al. 2018). Uncovering this relationship in our experimental results requires controlling for the mean-reverting multiplicative bias described previously; failing to account for this relationship makes the correlation appear to be positive between prediction error and air temperature. This is due to the negative correlation between planting date and air temperature resulting in omitted variable bias when not controlling for the former.

Figure 5 shows the relationship between prediction error and air temperature for the randomly partitioned validation scheme. Specifically, the plots show the residuals from regressing prediction error on planting DOY versus the residuals from regressing air temperature on planting DOY. This is equivalent to regressing prediction error on air temperature while controlling for planting DOY. The slope of the linear trend is most negative in the NDVI-only model (-0.53 days/degree C) and only slightly less negative in the NDVI+LST model (-0.50 days/degree C). For the NDVI+NSDSI model, there is no significant linear relationship between error and air temperature. The NDVI+NSDSI+LST model has a nonzero slope (-0.17 days/degree C), but one that is still much smaller in magnitude than the model configurations that do not include NSDSI.

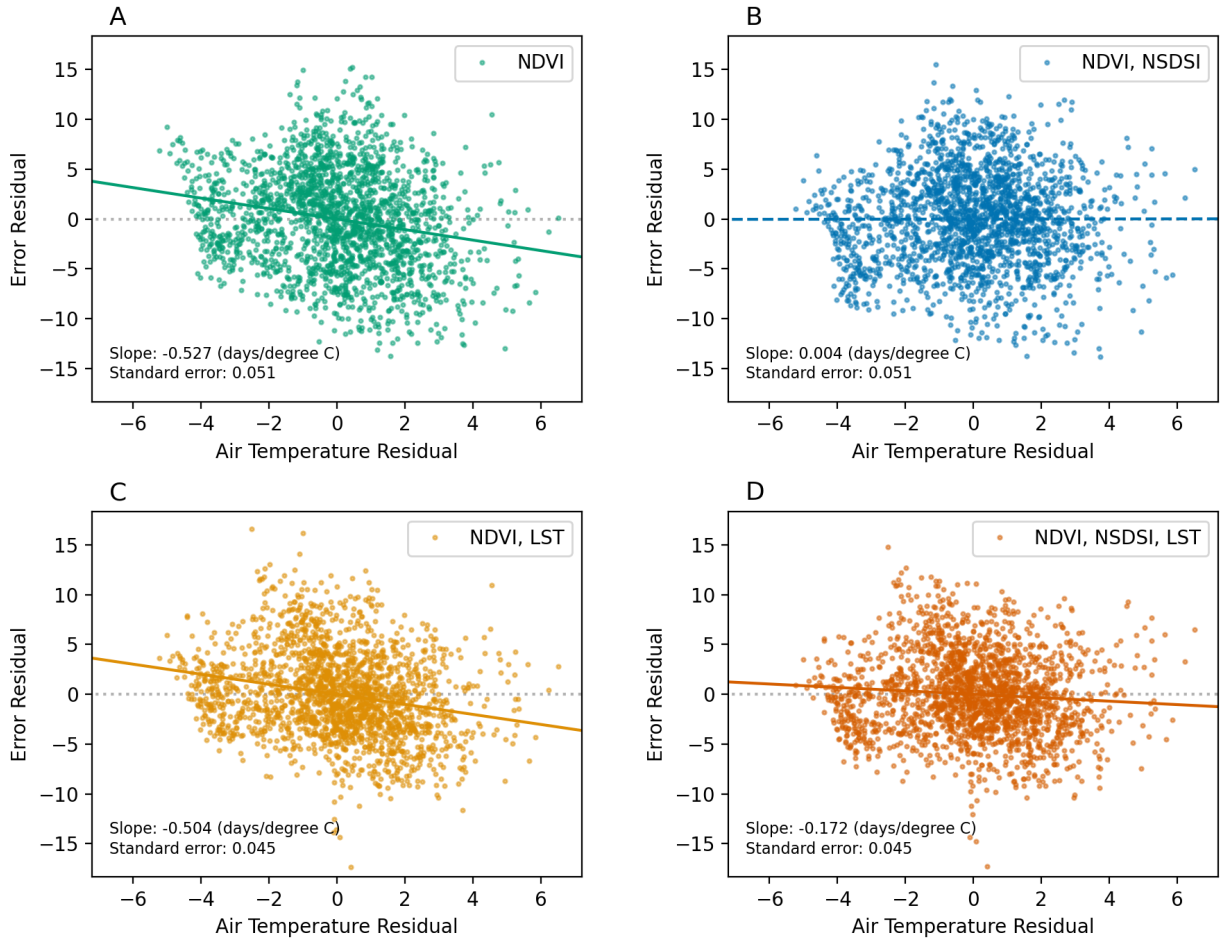

Figure 5*: Prediction error versus air temperature while controlling for planting DOY. These plots are for the randomly partitioned validation scheme. The slope of the regression line and the standard error of the estimate are given. Standard errors are robust to heteroskedasticity. Panel A shows results using only NDVI inputs. Panel B shows results using NDVI and NSDSI inputs (the dashed line indicates that the linear relationship between air temperature and error is not statistically significant). Panel C shows results using NDVI and LST inputs. Panel D shows results using NDVI, NSDSI, and LST inputs.*

We can also examine and interpret the FDA transformations computed from the data. Because FDA is linear and planting date labels are ordered, we can straightforwardly interpret how variation in a given model input influence predictions. Figure 6 shows the FDA transformations for different sets of model inputs computed using randomly partitioned counties and scaled by

the standard deviation for each measurement type and composite period. The transformations are similar for the divided-by-state validation scheme, which are shown in Supplementary Information Section B. Positive coefficients push the prediction of planting date later and negative coefficients push it earlier. Scaling by the standard deviation gives a representation of the average "contribution" of a given observation to the overall prediction.

Many of the FDA coefficients have clear physical interpretations. For example, the NDVI coefficients for composite periods starting DOY 161 have the largest magnitude in each model configuration. They are negative, meaning that higher NDVI in this period pushed the estimated planting date earlier. This period corresponds to the middle of June, when virtually all corn fields have been planted but have not yet reached peak vegetation density. Therefore, the NDVI during this observation period is likely to be determined by the crop itself, rather than weeds or other non-crop vegetation. Furthermore, the amount of crop biomass will be positively correlated with the time since planting. Similar logic can be applied to the consistently negative NDVI coefficients for periods beginning DOY 153, 169, and 177, although on average these observations have less influence on the final prediction. Conversely, the NDVI coefficients for DOY 225, and to a lesser extent 209 and 217, are consistently positive. If the crop is still at or near peak vegetation during these observation periods (early to mid-August), that means the leaves are not yet senescing and the crop was likely planted late. Despite fundamental differences in the approaches, the idea of extracting information from early and late season NDVI measurements using FDA is similar to the underlying logic of curve-fitting approaches. However, the interpretation of FDA coefficients is not limited to VI time series.

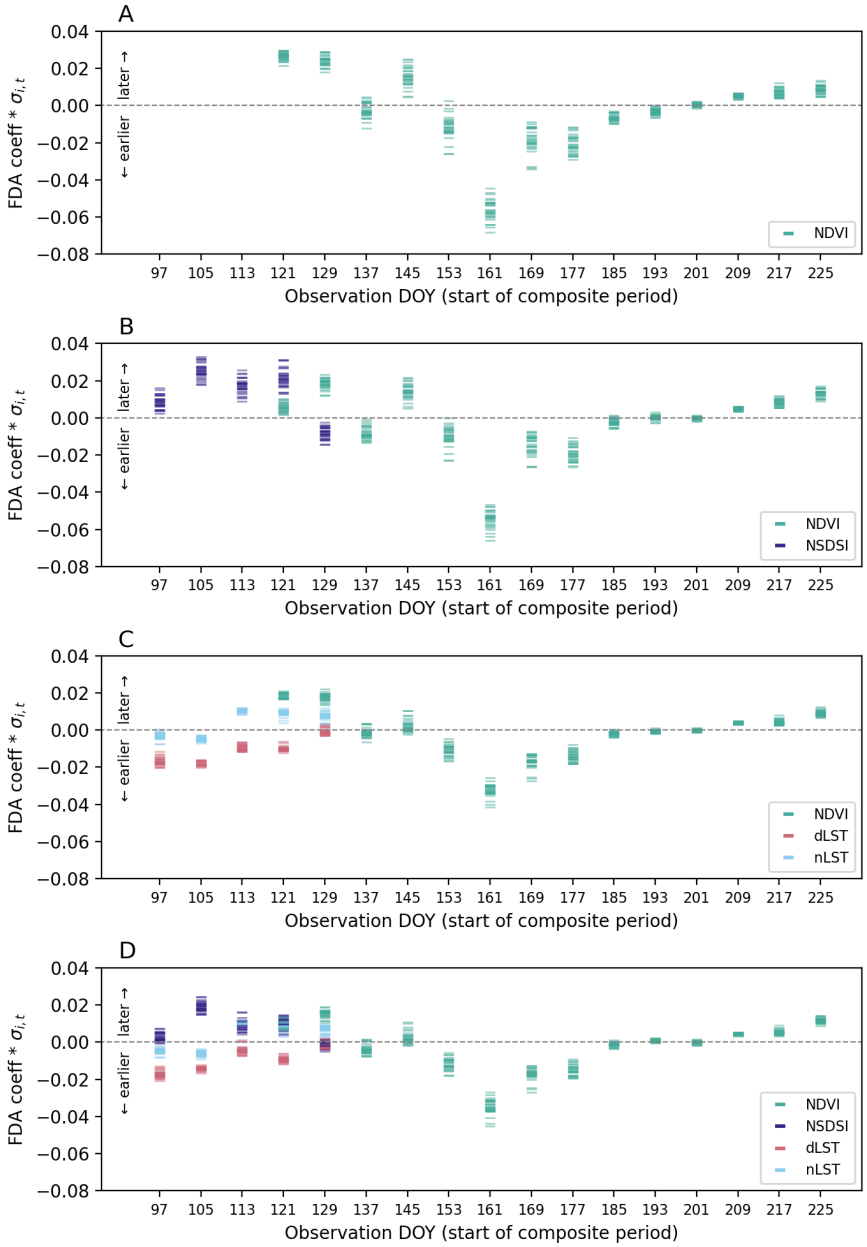

Figure 6*: FDA coefficients for different combinations of input data. These coefficients are for the randomly partitioned validation scheme. Coefficients are scaled by the standard deviation of each input-DOY combination such that the magnitude represents the "average contribution" of an observation to the final prediction. Panel A shows the coefficients for NDVI-only input. Panel B shows the coefficients for NDVI and NSDSI inputs. Panel C shows the coefficients for NDVI and LST inputs, where dLST and nLST refer to daytime and nighttime LST, respectively. Panel D shows the coefficients for NDVI, NSDSI, and both LST inputs.*

FDA is a more general approach and is able to incorporate information not captured by univariate curve-fitting. One clear example is the interpretation of NSDSI coefficients for observation periods beginning DOY 105, 113, and 121. All of these coefficients are consistently positive, meaning that higher NSDSI pushes planting date predictions later. NSDSI is positively correlated with soil moisture (Yue et al. 2019). Therefore, this association tracks the physical mechanism where overly wet fields can delay planting. While we may not expect soil moisture to linearly influence planting date, we would expect it to have an effect and we find that, overall the inclusion of NSDSI reduced prediction error in both validation schemes. Similarly, the NDVI coefficients for the early part of the growing season are also positive. These values are likely driven by weeds, cover crops, or other vegetation indicating that fields have not yet been planted. The effects of early season non-crop vegetation can confound curve-fitting approaches (Urban et al. 2018, Wardlow et al. 2006) due to the NDVI profile not matching the assumed functional form (e.g. logistic). An advantage of FDA is that it does not make prior assumptions about the shape of the time series, so this information from these early season observations can be incorporated without, for example, biasing a curve-fitting procedure.

We can also observe how LST measurement data influence predictions. For observations periods beginning DOY 97 and 105, both dLST and nLST coefficients are negative meaning that higher temperatures push predictions earlier. Starting at DOY 113, the sign of nighttime LST changes to positive but dLST coefficients remains negative (or approximately zero at DOY 129). The relationships between soil temperature, air temperature, vegetation development, and farmer behavior are complex, and we do not propose a definitive mechanism explaining these LST coefficients. However, the following information may be useful in interpreting them. First, nLST

was more highly correlated (0.58) than dLST (0.35) with average air temperature over the period DOY 90-120. Air temperature has been shown to be positively correlated with the rate of leaf extension in corn (Watts 1972). Therefore, warmer air temperature can lead to faster canopy development and, as a result, an NDVI profile of a crop that "appears" to have been planted earlier. In addition, we also observe that nLST to be positively correlated with NDVI and dLST negatively correlated with NDVI. For the periods beginning 113, 121, and 129, the correlations with NDVI of dLST and nLST were -0.31 and 0.25, -0.23 and 0.40, and -0.20 and 0.14. If vegetation leads to cooler daytime and warmer nighttime LST, then LST coefficients could be indicative of a similar phenomenon as with very early season NDVI. Cooler daytime and warmer nighttime LST could indicate that the soil is covered by non-crop vegetation, and planting has not yet occurred.

Finally, while this work is primarily concerned with the benefits of multivariate model inputs, it is notable that the errors for the NDVI-only inputs are similar across model types. One may expect a priori that the nonlinear approaches (logistic curve-fitting and random forest) to perform better than linear FDA given the same input data the given greater modeling flexibility. However, though that is not what we observe. Instead, the RMSEs for each model are approximately equal to one another within each validation scheme. The result that the nonlinear predictors are not significantly outperforming the linear predictor, given the same input, underscores the need for a more general mathematical framework in which to analyze phenology analysis tools.

5. Discussion

The analysis and results presented in this study have practical and theoretical implications for a range of applications in environmental remote sensing, particularly applications relating to phenological information extraction. In terms of practical takeaways, this study demonstrates how using multivariate time series data can lead to more precise extraction of phenological information such as crop planting date. It also establishes FDA as a method capable of extracting information from multivariate times series, the use of which can improve on widely used univariate prediction methods. Finally, this analysis demonstrates that application-relevent variation in the data (useful for the purpose of PIE) is often highly localized in time. Our analysis using FDA highlights this often-overlooked fact and our application of FDA presents a relatively simple method of PIE that accounts for it.

One concrete takeaway from this study is the demonstration that using non-VI data inputs can improve the precision of phenological information extraction. We demonstrate this empirically with the example application of planting date estimation. This result is not surprising, given that ground-based measurements of physical variables like air temperature and soil temperature have been used to improve estimates of planting date (Dong et al. 2019) and the onset of the growing season (Leeper et al. 2021). However, our study shows that this information can also be extracted from remotely sensed measurements. This has implications for the applications previously mentioned, but also to any phenological or agricultural management event that cannot be directly observed in a VI time series.

In addition, the potential applications of FDA extend beyond characterization of land surface phenology. The application in this study, estimation of county-year average planting date, has a response variable that takes values on an ordered continuum and so required putting data points into discrete bins. A similar procedure could be used to estimate quantities like crop yield. There are also many applications with categorical response variables, where the application of FDA is even more straightforward. For example, other potential applications of FDA include crop or land cover classification, identification of management practices like the use of irrigation or cover crops, and more. For each of these applications, the time when a measurement is taken is important to the interpretation of that measurement, which indicates that FDA may be of use.

For phenological information extraction and other potential applications, there are also many possible extensions to the basic use of FDA presented here. In this study, we only consider FDA projections into one dimensional space. However, FDA computes optimal (with respect to the Fisher Criterion) transformations into spaces with dimension up to $k - 1$ for a training data set containing k distinct classes (Ghojogh et al. 2019). Making beneficial use of a $d > 1$ dimensional representation requires a nonlinear predictor mapping the d-dimensional representation to labels (by construction, if there existed a superior linear predictor $M: \mathbb{R}^d \to \mathbb{R}$, then the first FDA vector would not be optimal). Strategies like this have been used with other dimensionality reduction techniques, like harmonic regression (Adams et al. 2020, Wang et al. 2019), where the lower-dimensional representation of the data is passed to a nonlinear predictor such as a random forest. This type of dimensionality reduction seems to be a natural fit for FDA given the relationship between FDA and the DFT explained in section 2.3. Similarly, it may be beneficial to employ ensembles of FDA transformations fit to different sets of predictors. This

kind of approach could also be useful in the presence of missing data. One can simply fit FDA transformations to sets of input data with particular observations removed, given that FDA does not make prior assumptions about the covariance structure of the input data (i.e. it does not require samples to be evenly spaced in time).

Beyond the applications of FDA, perhaps the most important point raised by this study is the need for representations of environmental time series that are localized in time. Figure 2 and the surrounding discussion show how to think about generic time series transformations, like the DFT, with respect to the Fisher Criterion. Representations that give equal weight the whole time series will necessarily be influenced by observations that contain little information about a given phenological event or other property of the time series. Despite the significant limitation of linearity, FDA is able to effectively extract phenological information because it allows for observations of different variables at different times to be weighted different when it comes to the final prediction. Nonlinear curve-fitting with logistic curves allows for some localization in time, but still relies on irrelevant data points to stabilize the estimation procedure. It also does not lend itself to multivariate data inputs, like soil moisture and LST, that do not lend themselves to fixed functional forms.

Author Credit

CD conceived the paper and processed the data. CD and MM performed the analysis and wrote the paper.

Declaration of Interest

The authors declare that they have no known competing interests that could have appeared to influence the work reported in this paper.

Data availability

All data used in this study is publicly available at the sources described by the associated references.

Funding

This work was supported by the Stanford Data Science Institute and the National Science Foundation (NSF) CBET 1604853.

References

Abendroth, L., Elmore, R.W., Boyer, M.J., Marlay, S.K., 2011. Corn Growth and Development. Iowa State University Extension, Ames.

Adams, B., Iverson, L., Matthews, S., Peters, M., Prasad, A., Hix, D.M., 2020. Mapping Forest Composition with Landsat Time Series: An Evaluation of Seasonal Composites and Harmonic Regression. Remote Sensing 12, 610. https://doi.org/10.3390/rs12040610

Breiman, L., 2001. Random Forests. Machine Learning 45, 5–32.

Chen, J., Jönsson, Per., Tamura, M., Gu, Z., Matsushita, B., Eklundh, L., 2004. A simple method for reconstructing a high-quality NDVI time-series data set based on the Savitzky–Golay filter. Remote Sensing of Environment 91, 332–344. https://doi.org/10.1016/j.rse.2004.03.014

Dong, T., Shang, J., Qian, B., Liu, J., Chen, J.M., Jing, Q., McConkey, B., Huffman, T., Daneshfar, B., Champagne, C., Davidson, A., MacDonald, D., 2019. Field-Scale Crop Seeding Date Estimation from MODIS Data and Growing Degree Days in Manitoba, Canada. Remote Sensing 11, 1760. https://doi.org/10.3390/rs11151760

Entekhabi, D., Reichle, R.H., Koster, R.D., Crow, W.T., 2010. Performance Metrics for Soil Moisture Retrievals and Application Requirements. Journal of Hydrometeorology 11, 832–840. https://doi.org/10.1175/2010JHM1223.1

Filippelli, S.K., Vogeler, J.C., Falkowski, M.J., Meneguzzo, D.M., 2020. Monitoring conifer cover: Leaf-off lidar and image-based tracking of eastern redcedar encroachment in central

Nebraska. Remote Sensing of Environment 248, 111961. https://doi.org/10.1016/j.rse.2020.111961

Friedman, J.H., 1989. Regularized Discriminant Analysis. Journal of the American Statistical Association 84, 165–175. https://doi.org/10.1080/01621459.1989.10478752

Garcia, D., 2010. Robust smoothing of gridded data in one and higher dimensions with missing values. Computational Statistics & Data Analysis 54, 1167–1178. https://doi.org/10.1016/j.csda.2009.09.020

GDAL/OGR contributors, 2022. Geospatial Data Abstraction software Library.

Ghojogh, B., Karray, F., Crowley, M., 2019. Fisher and Kernel Fisher Discriminant Analysis: Tutorial. arXiv:1906.09436 [cs, stat].

Gray, R.M., 2006. Toeplitz and Circulant Matrices: A Review. Foundations and Trends in Communications and Information Theory 2, 155–239.

Guan, K., Medvigy, D., Wood, E.F., Caylor, K.K., Li, S., Jeong, S.-J., 2014. Deriving Vegetation Phenological Time and Trajectory Information Over Africa Using SEVIRI Daily LAI. IEEE Transactions on Geoscience and Remote Sensing 52, 1113–1130. https://doi.org/10.1109/TGRS.2013.2247611

Hastie, T., Tibshirani, R., Friedman, J., 2008. The Elements of Statistical Learning, 2nd ed. Springer.

Hulley, G., Freepartner, R., Malakar, N., Sarkar, S., 2022. Moderate Resolution Imaging Spectroradiometer (MODIS) Land Surface Temperature and Emissivity Product (MxD21) User Guide Collection 6.1 37.

Jakubauskas, M.E., Legates, D.R., Kastens, J.H., 2002. Crop identification using harmonic analysis of time-series AVHRR NDVI data. Computers and Electronics in Agriculture 37, 127–139. https://doi.org/10.1016/S0168-1699(02)00116-3

Kandasamy, S., Baret, F., Verger, A., Neveux, P., Weiss, M., 2013. A comparison of methods for smoothing and gap filling time series of remote sensing observations – application to MODIS LAI products. Biogeosciences 10, 4055–4071. https://doi.org/10.5194/bg-10-4055-2013

Leeper, R.D., Matthews, J.L., Cesarini, M.S., Bell, J.E., 2021. Evaluation of Air and Soil Temperatures for Determining the Onset of Growing Season. JGR Biogeosciences 126. https://doi.org/10.1029/2020JG006171

Leith, H., 1974. Purposes of a Phenology Book, in: Phenology and Seasonality Modeling. Springer-Verlag, New York, pp. 3–19.

Lobell, D.B., Ortiz-Monasterio, J.I., Sibley, A.M., Sohu, V.S., 2013. Satellite detection of earlier wheat sowing in India and implications for yield trends. Agricultural Systems 115, 137–143. https://doi.org/10.1016/j.agsy.2012.09.003

Lobell, D.B., Roberts, M.J., Schlenker, W., Braun, N., Little, B.B., Rejesus, R.M., Hammer, G.L., 2014. Greater Sensitivity to Drought Accompanies Maize Yield Increase in the U.S. Midwest. Science 344, 516–519. https://doi.org/10.1126/science.1251423

Mingwei, Z., Qingbo, Z., Zhongxin, C., Jia, L., Yong, Z., Chongfa, C., 2008. Crop discrimination in Northern China with double cropping systems using Fourier analysis of time-series MODIS data. International Journal of Applied Earth Observation and Geoinformation 10, 476–485. https://doi.org/10.1016/j.jag.2007.11.002

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., 2011. Scikit-learn: Machine Learning in Python. MACHINE LEARNING IN PYTHON 6.

Rippey, B., 2019. Nation's Wettest 12-Month Period on Record Slows Down 2019 Planting Season [WWW Document]. US Department of Agriculture Blog. URL https://www.usda.gov/media/blog/2019/06/14/nations-wettest-12-month-period-record-slows-down-2019-planting-season (accessed 3.22.20).

Roger, P.J.C., Ray, J.P., Vermote, E.F., 2020. MODIS Collection 6.1 (C61) LSR Product User Guide 36.

Sakamoto, T., Yokozawa, M., Toritani, H., Shibayama, M., Ishitsuka, N., Ohno, H., 2005. A crop phenology detection method using time-series MODIS data. Remote Sensing of Environment 96, 366–374. https://doi.org/10.1016/j.rse.2005.03.008

SciPy 1.0 Contributors, Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S.J., Brett, M., Wilson, J., Millman, K.J., Mayorov, N., Nelson, A.R.J., Jones, E., Kern, R., Larson, E., Carey, C.J., Polat, İ., Feng, Y., Moore, E.W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E.A., Harris, C.R., Archibald, A.M., Ribeiro, A.H., Pedregosa, F., van Mulbregt, P., 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat Methods 17, 261–272. https://doi.org/10.1038/s41592-019-0686-2

Urban, D., Guan, K., Jain, M., 2018. Estimating sowing dates from satellite data over the U.S. Midwest: A comparison of multiple sensors and metrics. Remote Sensing of Environment 211, 400–412. https://doi.org/10.1016/j.rse.2018.03.039

USDA-NASS, 2022. USDA National Agricultural Statistics Service Cropland Data Layer.

Wang, S., Azzari, G., Lobell, D.B., 2019. Crop type mapping without field-level labels: Random forest transfer and unsupervised clustering techniques. Remote Sensing of Environment 222, 303–317. https://doi.org/10.1016/j.rse.2018.12.026

Wardlow, B.D., Kastens, J.H., Egbert, S.L., 2006. Using USDA Crop Progress Data for the Evaluation of Greenup Onset Date Calculated from MODIS 250-Meter Data. photogramm eng remote sensing 72, 1225–1234. https://doi.org/10.14358/PERS.72.11.1225

Watts, W.R., 1972. Leaf Extension in Zea mays. J Exp Bot 23, 713–721. https://doi.org/10.1093/jxb/23.3.713

Yue, J., Tian, J., Tian, Q., Xu, K., Xu, N., 2019. Development of soil moisture indices from differences in water absorption between shortwave-infrared bands. ISPRS Journal of Photogrammetry and Remote Sensing 154, 216–230. https://doi.org/10.1016/j.isprsjprs.2019.06.012

Zeng, L., Wardlow, B.D., Xiang, D., Hu, S., Li, D., 2020. A review of vegetation phenological metrics extraction using time-series, multispectral satellite data. Remote Sensing of Environment 237, 111511. https://doi.org/10.1016/j.rse.2019.111511

Zhang, X., Friedl, M.A., Schaaf, C.B., Strahler, A.H., Hodges, J.C.F., Gao, F., Reed, B.C., Huete, A., 2003. Monitoring vegetation phenology using MODIS. Remote Sensing of Environment 84, 471–475. https://doi.org/10.1016/S0034-4257(02)00135-9

Zhang, Z., Dai, G., Xu, C., Jordan, M.I., n.d. Regularized Discriminant Analysis, Ridge Regression and Beyond 30.

Supplementary Information

A. Error Density Plots for Curve-Fitting and Random Forest Predictors

This section contains error density plots for the curve-fitting and random forest predictors. The plots are analogous to the plots in Figure 4, which shows the error density plots for the FDA predictors. Figure 7 shows the error densities for the curve-fitting predictors. Note that the only input is NDVI because the curve-fitting approach takes only a univariate input. The shape of the estimated density functions is similar to the NDVI-only error density functions shown in Figure 4.

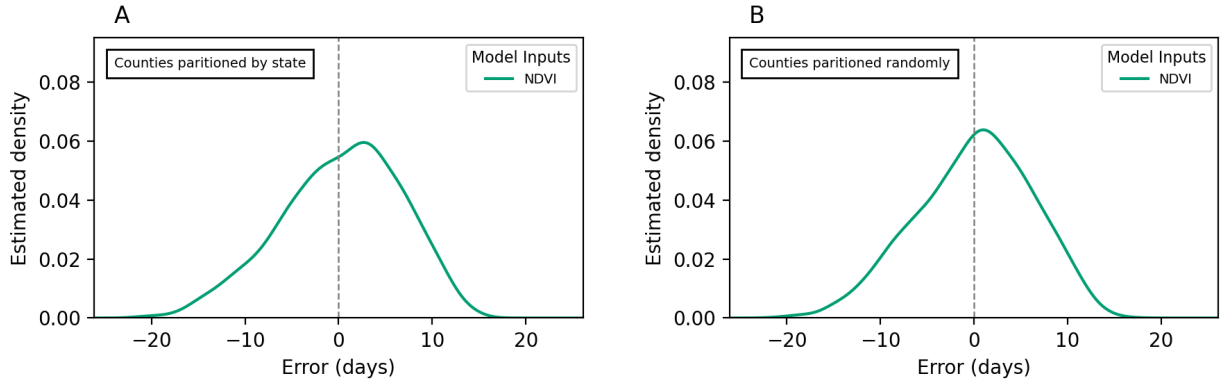

***Figure 7**: Error distributions for curve-fitting models. Panel A shows the error distribution for the partitioned-by-state validation scheme. Panel B shows the error distribution for the randomly partitioned scheme. The distributions are generated by kernel density estimation using a Gaussian kernel. The shape of the error density functions is similar to that of the NDVI-only FDA predictors shown in Figure 4.*

Figure 8 shows the error densities for the random forest predictors. The estimated density functions for NDVI-only inputs are fairly similar the NDVI densities for FDA and curve-fitting predictors. Technically the error variance (estimated as ubRMSE) is larger for random forest, but the difference is not large enough to be visually apparent. In contrast, it is apparent that the reduction in error variance when using LST is smaller than the reduction for FDA-based

predictors. The "narrowing" of the estimated error density functions for models taking LST inputs is more significant for FDA-based predictors, as seen in Figure 4, than it is for random forest predictors.

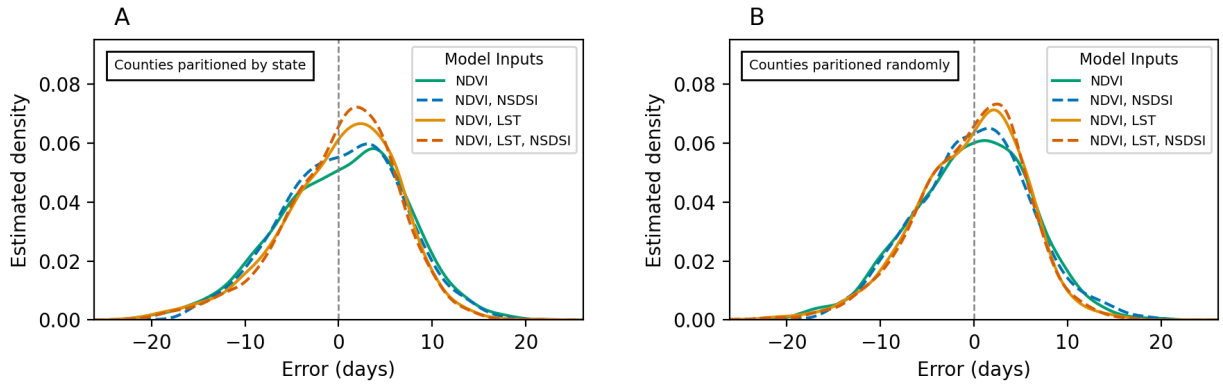

Figure 8*: Error distributions for random forest models. Panel A shows the error distribution for the partitioned-by-state validation scheme. Panel B shows the error distribution for the randomly partitioned scheme. The distributions are generated by kernel density estimation using a Gaussian kernel. Note that the use of LST data inputs results in a less significant reduction in error variance than for comparable FDA-based predictors. This can be observed visually in how, in Figure 4, more of the probably mass is concentrated around 0 for the orange and red curves. In Figure 8, there is some reduction in error variance, but the reduction is visibly smaller than in Figure 4.*

B. FDA Coefficients for Partitioned-by-State Validation Scheme

This section contains Figure 9, which is the plot of the FDA coefficients computed in the portioned-by-state validation scheme. Figure 9 is analogous to Figure 6, which shows the FDA coefficients but for the randomly-partitioned validation scheme. The plots in Figure 9 are very similar to the Figure 6, but with slightly more variation between the training data sets across cross-validation partitions. This is because the measurement data and average planting dates both vary spatially, and the resulting FDA coefficient reflect this variation. However, the overall "shape" of the transformations is very similar to the transformations in the randomly-partitioned

validation scheme. This indicates that FDA is identifying roughly the same "signals" in both validation schemes.

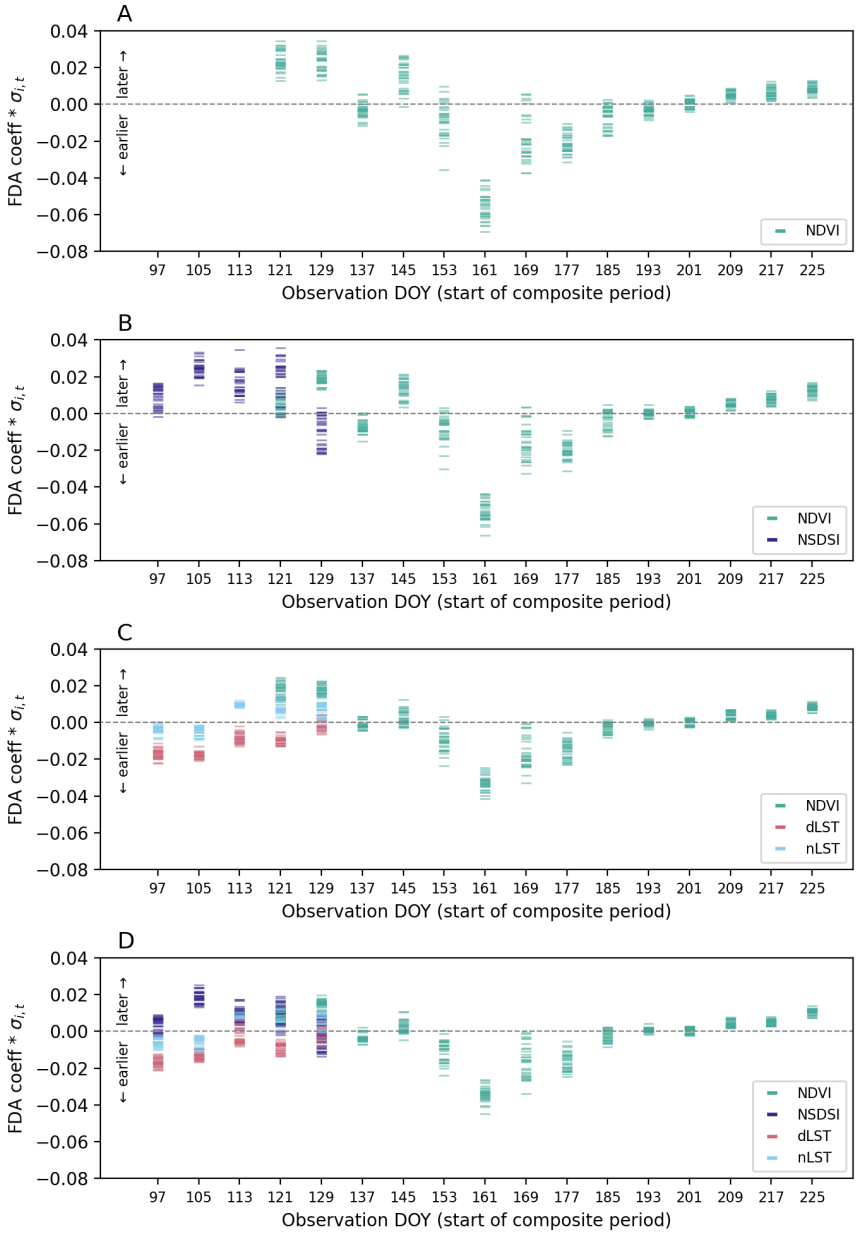

***Figure 9**: FDA coefficients for different combinations of input data for the partitioned-by-state validation scheme. Coefficients are scaled by the standard deviation of each input-DOY combination such that the magnitude represents the "average contribution" of an observation to the final prediction. Panel A shows the coefficients for NDVI-only input. Panel B shows the coefficients for NDVI and NSDSI inputs. Panel C shows the coefficients for NDVI and LST inputs, where dLST and nLST refer to daytime and nighttime LST, respectively. Panel D shows the*

coefficients for NDVI, NSDSI, and both LST inputs. Overall, the coefficients computed in the portioned-by-state validation scheme are very similar to those computed in the randomly-partitioned validation scheme, as shown in Figure 6.