Forecasting Human West Nile Virus Cases Using Eco-Climate and Demographic Surveillance Data

Harsita Vigneshwaran

Abstract

Currently, West Nile Virus is the leading mosquito-borne arbovirus in the United States, causing seasonal outbreaks that overwhelm public health surveillance systems. Although better forecasting could support earlier vector-control interventions, the only publicly available source of county-level human West Nile Virus cases—CDC's ArboNET—reports data annually rather than at finer temporal scales, limiting short-term prediction. In this study, we evaluated whether county-level annual West Nile Virus case counts could be predicted one year in advance using prior-year case history, demographics, land use, and weather variables. We hypothesized that higher prior-year growing degree days would be associated with increased case counts in the following year. Gradient Boosting with Poisson loss performed best, attaining a mean absolute error (MAE) of 49.4 cases per year; that's a 13% improvement over the strongest baseline (per-county climatology, MAE = 56.6). Despite this improvement, overall predictive power was limited, with a high root-mean-squared error (RMSE = 175.2) and near-zero R^2 across models, largely driven by a small number of extreme outbreak years. Performance varied substantially by outbreak status, with strong accuracy in non-outbreak years (MAE = 26) but poor performance during outbreak years (MAE = 960). Prior-year growing degree days emerged as the most important predictor based on permutation importance, consistent with our hypothesis. However, adding additional weather variables did not substantially improve model effectiveness, suggesting that growing degree days primarily capture stable geographic differences between counties rather than interannual variation in risk.

Introduction

West Nile Virus (WNV) is a mosquito-borne arbovirus first discovered in Uganda in 1937, that then made its way to the United States in 1999, and rapidly became the leading cause of mosquito-borne illnesses in the country. WNV transmission is influenced by local temperature, rainfall, and land use, since these factors affect how mosquitoes develop which leads to how the virus multiplies. As climates change and urban growth continues to shift these environmental conditions, WNV patterns are also changing. However, most public health responses still tend to happen only after people start showing serious complications, rather than preventing outbreaks before they spread.

Trying to predict where and when WNV will show up isn't easy, mostly because the data are limited. For example, the CDC's ArboNET system only publishes case counts for each county once a year, so it's hard to spot trends as they happen. Recent attempts to use machine learning for prediction have had mixed success. One group compared several models and found that even the most complex approaches didn't do much better than simply looking at historical case numbers (5).

Another study found that newer techniques, like graph neural networks that map out spatial relationships between mosquito trap sites, can outperform more basic methods—but only when detailed weekly data from traps are available (6). In short, better predictions require both higher-quality data and models that can handle complex patterns in time and space.

We hypothesize that higher prior-year annual growing degree days are associated with higher county-level human WNV case counts in the following year. This study tests that hypothesis using annual ArboNET totals as the prediction target and prior-year case history, demographics, land use, and weather features as predictors, benchmarked against four baselines under leave-one-year-out cross-validation.

Results

Overall Model Performance

Across all 120 county-year observations under leave-one-year-out cross-validation, Gradient Boosting with a Poisson loss objective achieved the lowest mean absolute error of any predictor (MAE = 49.4 cases per year), compared to the strongest baseline of per-county climatology (MAE = 56.6). This represents a 13% reduction in mean absolute error, indicating modest but real predictive value from the machine learning model. However, the other evaluation metrics tell a subtler story. The root mean squared error was 3.5 times larger than the mean absolute error for Gradient Boosting (RMSE = 175.2), and the coefficient of fit was approximately zero or negative across all predictors, including baselines. Pearson correlations between predicted and observed values were small (between −0.40 and 0.43) for every predictor evaluated. Linear regression underperformed relative to all baselines (MAE = 99.6), reflecting multicollinearity among the weather features and the model's inability to capture nonlinear associations (Table 1).

Table 1. Overall predictive performance under leave-one-year-out cross-validation for the six-county WNV dataset (2005–2024, n = 120 county-years). Each of the four baseline predictors (global mean, per-county climatology, persistence, three-year rolling mean) and three machine learning models (linear regression, random forest, gradient boosting with a Poisson loss objective) was evaluated by training on all years except one and generating predictions for the held-out year, repeated for each of the 20 years in the analysis window. Metrics reported are mean absolute error (MAE), root mean squared error (RMSE), the ratio RMSE/MAE (a diagnostic for heavy-tailed error distributions), Pearson correlation among predicted and observed values (r), and coefficient of explanation (R^2). Lower MAE and RMSE indicate better performance. The values shown are computed across all 120 county-year predictions and pooled.

Predictor	MAE	RMSE	RMSE/MAE	Pearson r	R^2	Type
Global mean	57.1	168.9	2.96	−0.40	−0.02	Baseline
Per-county climatology	56.6	166.1	2.93	0.19	0.02	Baseline
Persistence (last year)	58.1	180.2	3.10	0.43	−0.16	Baseline
3-year rolling mean	61.5	185.3	3.01	0.20	−0.22	Baseline
Linear regression	99.6	221.1	2.22	−0.05	−0.74	ML

Predictor	MAE	RMSE	RMSE/MAE	Pearson r	R^2	Type
Random Forest	61.1	180.4	2.95	0.03	−0.16	ML
Gradient Boosting (Poisson)	49.4	175.2	3.55	−0.04	−0.10	ML (best)

The ratio of RMSE to MAE was consistently high—usually between 2.5 and 3.5—across all models. This points to the fact that just a few unusually large outbreaks made the squared error much higher than the average error. The outbreaks in Maricopa County in 2021 and 2022 (with 1,487 and 1,126 cases) are responsible for most of the RMSE seen in every model. Because of these rare but massive outbreaks, none of the models could explain much of the variation in the data, which is why R^2 values hovered around zero. So, while the best machine learning model did beat the simplest baseline, the improvement was fairly small.

Per-County Performance

Disaggregating mean absolute error by county revealed substantial heterogeneity in model performance, largely driven by the volatility of each county's case time series (Table 2).

Table 2. Mean absolute error (MAE), in cases per year, disaggregated by county for one baseline (climatology), one simple time-series baseline (persistence), and three machine learning models under leave-one-year-out cross-validation (n = 20 county-years per county). Values are pooled across the 2005–2024 analysis window and reflect the average absolute difference between predicted and observed annual case counts. Lower values indicate better predictive accuracy. Maricopa, AZ, shows an order-of-magnitude larger MAE across all predictors due to the 2021 (1,487 cases) and 2022 (1,126 cases) outbreak years.

County	Climatology	Persistence	Linear reg.	Random Forest	Grad. Boost.
Boulder, CO	11.3	14.6	33.8	22.9	12.2
Cook, IL	27.6	31.6	74.4	30.1	31.6
Dallas, TX	40.2	51.4	87.0	46.7	36.7
Larimer, CO	13.1	17.0	30.5	15.6	14.8
Los Angeles, CA	14.2	28.8	75.5	23.0	14.2
Maricopa, AZ	233.3	205.4	296.6	228.0	186.7

In low-volatility counties such as Boulder and Larimer, where year-to-year case counts stayed within a relatively narrow range, per-county climatology was either similar to or better than the best machine learning model. Dallas County was where Gradient Boosting most clearly outperformed climatology, while Cook County showed the opposite pattern: climatology slightly outperformed the machine learning model after weather features were included. This may be because Cook County's case history is driven more by a 2002 outbreak and a 2012 spike than by consistent weather-related patterns. Maricopa County stands apart from the other counties. All models produced MAE values that were much larger than those seen in other counties, mainly due to the very large 2021–2022 outbreak. Gradient Boosting achieved the lowest MAE in Maricopa among the machine learning models (186.7), but it still did not match the climatology baseline during the outbreak years.

Outbreak Versus Non-Outbreak Performance

The prominent result is that during outbreak years, Gradient Boosting actually performed worse than simply repeating the previous year's case count (MAE 959.8 vs. 725.0). That's because big outbreaks often happen back-to-back, which led to last year's number ends up being a strong predictor. By comparison, the machine learning models tend to predict case numbers near the long-term average, which means they regularly underestimate when there's a sudden surge. In non-outbreak years, though, Gradient Boosting did well, posting the lowest mean absolute error (26.0 cases per year) and handling routine, year-to-year variation. To get these results, we split the 120 county-year observations into outbreak years (at least 200 cases, n = 3) and non-outbreak years (fewer than 200 cases, n = 117), then calculated mean absolute error and root mean squared error for each group (see Table 3).

Table 3. Predictive performance by outbreak status. Outbreak county-years are shown as those with at least 200 reported cases (n = 3: Dallas, TX 2012 [398 cases], Maricopa, AZ 2021 [1,487 cases], and Maricopa, AZ 2022 [1,126 cases]); non-outbreak county-years have fewer than 200 cases (n = 117). Mean absolute error (MAE) and root mean squared error (RMSE) were calculated under leave-one-year-out cross-validation.

MAE and RMSE values are presented separately for outbreak and non-outbreak years. The machine learning model shows higher error in outbreak years relative to non-outbreak years, reflecting reduced accuracy when case counts are large.

Predictor	MAE outbreak (n=3)	RMSE outbreak	MAE normal (n=117)	RMSE normal
Climatology	891.8	985.1	35.2	58.5
Persistence	725.0	902.1	41.0	111.5
Random Forest	932.4	1034.1	38.7	77.1
Gradient Boosting (Poisson)	959.8	1067.6	26.0	47.6

Feature Importance

Permutation feature importance was computed for the Gradient Boosting model, measuring the mean increase in cross-validated MAE when each feature was randomly shuffled (Table 4).

Table 4. Top ten features ranked by permutation importance for the gradient boosting (Poisson loss) model, computed on the full six-county 2005–2024 dataset (n = 120 county-years). Permutation importance was calculated by randomly shuffling each feature's values, recomputing predictions, and measuring the resulting increase in mean absolute error (MAE), averaged over 20 repetitions. Higher values indicate features whose shuffling degrades predictive performance more, and therefore contribute more predictive information. All features shown are lagged by at least one year relative to the prediction target (denoted by the "prev_yr" prefix). Feature abbreviations: wx = weather (NOAA), cdl = Cropland Data Layer (USDA), saipe = Small Area Income and Poverty Estimates (Census Bureau).

Feature	Permutation Importance (MAE increase)
prev_yr_wx_gdd_annual (annual growing degree days)	19.76
prev_yr_cdl_pasture_hay_pct	19.64
prev_yr_wx_tmax_annual (annual max temperature)	15.66
prev_yr_saipe_poverty_rate	14.24
prev_yr_cdl_developed_pct	14.19
prev_yr_pct_65_plus	10.72
cases_3yr_ago	9.69
prev_yr_cases	8.58
prev_yr_wx_tavg_annual (annual mean temperature)	7.84
prev_yr_saipe_median_income	6.04

Three weather features were in the top ten by permutation importance: annual growing degree days (rank 1, importance 19.76), annual maximum temperature (rank 3, importance 15.66), and annual mean temperature (rank 9, importance 7.84). Land-use variables were also important, with prior-year percent pasture and hay land ranked second and prior-year percent developed land ranked fifth. Demographic variables ranked third and fourth, with prior-year poverty rate and prior-year percent of the population aged 65 and older both in the top six features. Prior-year case count, which was the fifth most important feature in a model without weather data, dropped to eighth after weather features were added. This suggests that some of the predictive power from past case counts may actually come from underlying climate conditions that affect both past and current transmission. These results differ from some previous WNV machine learning studies, where mosquito surveillance features are often the most important predictors. In those studies, mosquito data is often closely related to the human case data, which can make its importance appear higher than it really is.

Marginal Value of Weather Features

Although weather variables were among the most important predictors according to permutation importance, adding them to the feature set did not clearly improve prediction accuracy on new data. A sensitivity analysis was done by running the same modeling pipeline with and without the eleven weather features. The Gradient Boosting model without weather achieved MAE = 49.2 cases per year and RMSE = 176.4, while the same model with weather features achieved MAE = 49.4 and RMSE = 175.2. This change in MAE is about 0.3 percent, which is very small and likely within the normal variation observed in cross-validation. Random Forest results were also very similar (MAE 60.9 versus 61.1). Linear regression performed worse when weather features were added (MAE increased from 85.6 to 99.6). This may be because many of the weather variables are strongly related to one another, and linear regression cannot readily capture these relationships or nonlinear effects.

The combination of high feature significance with little improvement regarding prediction accuracy is not a contradiction. Weather variables remain key drivers of West Nile virus transmission: for example, temperature affects mosquito reproduction and the time it takes for the virus to develop inside mosquitoes (3). The model correctly identifies these variables as important. However, at the yearly county level, weather does not vary much within a single county from year to year compared to the differences between counties. Those long-term differences are already captured by the county fixed effects. In other words, the model already "knows" that places like Maricopa County are generally hotter than places like Boulder County, even without explicit weather features. Because of this, most of the predictive value of weather is already absorbed by the county indicators, leaving only a small additional effect from year-to-year weather changes. This suggests that weather data may not add much value for forecasting WNV at an annual scale. More detailed time scales, such as monthly or weekly data, may be needed to more effectively capture weather effects in predictions.

Discussion

This study evaluated whether annual county-level human WNV case counts could be predicted one year in advance using prior-year case history, demographic structure, land-use composition, and weather conditions. Four main findings emerged. First, machine learning provides a modest improvement over simple baselines. Gradient Boosting with a Poisson loss reduced MAE by 13% compared to the best baseline per county climatology. Second, performance depends on the situation: the model performs well in non-outbreak years (MAE = 26 cases) but performs poorly in outbreak years (MAE = 960 cases), where a simple persistence-based baseline performs better. Third, weather variables are among the most important features according to permutation importance, but they do not improve prediction effectiveness when added to a model that already includes county fixed effects. This suggests that much of the weather signal overlaps with long-term differences between counties. Fourth, all evaluation measures show patterns expected for heavy-tailed epidemiological data, with RMSE values 2.5 to 3.5 times larger than MAE values and R^2 values near zero, consistent with earlier work on WNV forecasting (5).

The differences between MAE, RMSE, and R^2 in this study reflect the structure of WNV case data rather than problems with the models. MAE treats all errors equally, so it stays relatively low when most county-years have typical case counts. RMSE, however, gives extra weight to large errors because it squares them before averaging. As a result, a small number of outbreak years dominate RMSE. For example, the 2021 and 2022 Maricopa County outbreaks (1,487 and 1,126 cases) produce very large squared errors when underpredicted. These few observations account for most of the total RMSE across all 120 county-years. R^2 behaves differently because it compares model error to overall variance in the data. In this dataset, much of the variance also comes from rare outbreak years. Since no model precisely predicts these extreme values, both the error and the total variance are dominated by the same observations, causing R^2 to stay near zero or become negative. This pattern is expected in real epidemiological count data and is consistent with results reported in previous WNV forecasting studies.

Each data source contributed to model effectiveness, but in ways that differed from prior multi-source WNV machine learning studies. Weather features ranked highly in permutation importance, with annual growing degree days as the most important predictor. This aligns with established biology, as temperature affects mosquito development and accelerates viral replication in mosquitoes (3). However, the added predictive value of weather was small once county fixed effects were included. Land-use features were the second most important group of predictors, especially percent pasture and hay land and percent developed land. These likely reflect long-term environmental conditions that influence mosquito habitats and change slowly over time. Demographic features were also important, especially the poverty rate and the percentage of the population aged 65 and older. Age structure may be more related to how cases are detected and reported, since severe WNV disease is more common in older adults. Poverty rate may reflect differences in housing, exposure to outdoor environments, and access to medical care. Notably, mosquito surveillance data have been identified in prior work as a useful early-warning signal that often precedes human cases by one to several weeks (7), but were not directly usable here because the available weekly file was derived from annual ArboNET totals rather than from real trap data. Real weekly mosquito pool positivity, published by some county vector-control programs, is a promising input for future higher-resolution WNV forecasting.

Maricopa County experienced unusually large WNV outbreaks in 2021 and 2022, with reported case counts of 1,487 and 1,126, respectively. These two years alone account for more cases than the entire period before 2021 in Maricopa County combined. No model in this study, including all baselines and machine learning methods, could predict the 2021 outbreak using only prior-year features. This is not a limitation specific to these models; it reflects the difficulty of predicting rare and extreme outbreak events. The 2021 outbreak has been linked in subsequent studies to a combination of unusually wet conditions, very high mosquito populations, and local exposure factors not evident in prior-year data. This case study shows a key limitation of relying solely on past-year data for WNV prediction. Large outbreaks likely depend on current conditions, especially current mosquito populations, mosquito infection rates, and within-season weather patterns. These factors are not captured by annual lagged features. Because of this, early warning systems would likely need real-time mosquito surveillance data in addition to historical data. However, such surveillance data is not available in all counties across the United States.

Several limitations should be considered when interpreting the results. First, the analysis is limited to six counties, chosen for data availability rather than representativeness. As a result, it is unclear how well the findings apply to other U.S. counties, and local recalibration would likely be needed. Second, real-time mosquito surveillance data, which have been shown in prior studies to be useful for early warning, were excluded because they were not consistently available across all study counties. Subsequent research should incorporate standardized mosquito surveillance data where possible. Third, although weather variables were included, they were aggregated to the annual level to match the yearly case data. This likely removes important within-year variation that is more directly related to transmission. Fourth, the modeling approach assumes that the WNV system is stable over time. However, changes in mosquito populations, climate conditions, or human exposure patterns shall affect model effectiveness and would not be captured by cross-validation alone. Finally, the dataset is relatively small, with 120 county-year observations and only three outbreak years. This limits the ability to evaluate outbreak prediction. Even with good models, it is difficult to learn reliable patterns from so few extreme events, so larger datasets would be needed for stronger conclusions about outbreak forecasting.

Within these limitations, the data show multiple practical conclusions. For counties with stable seasonal WNV patterns and limited year-to-year variation, machine learning models provide a small but consistent improvement over simple climatology forecasts. This may be useful for routine planning such as resource assignment, vector control budgeting, and seasonal staffing, where errors of tens of cases per year are acceptable. For outbreak prediction, however, prior-year data alone is not sufficient. Effective early-warning systems would require real-time information, such as mosquito, environmental, and bird surveillance data, which were not included in this study. Overall, reporting performance separately for outbreak and non-outbreak years is more informative for public health use than relying only on overall summary metrics. Several further studies would be needed before these outcomes could be applied operationally: weekly forecasting models for counties with real weekly case data, integration of actual mosquito surveillance from vector-control programs, expansion to 30–50 counties to test generalizability, evaluation of methods designed for rare-event prediction such as ensemble approaches used successfully for other vector-borne diseases (8) or time-series methods such as long short-term memory networks (9) , and operational validation in which forecasts are deployed and their impact on public health decisions is measured.

Materials and Methods

Study Area and County Selection

This study analyzed six U.S. counties — Boulder, CO; Cook, IL; Dallas, TX; Larimer, CO; Los Angeles, CA; and Maricopa, AZ— chosen to span very different climates and outbreak patterns. Maricopa and LA both see huge episodic outbreaks, while Boulder and Larimer have more stable, low-level transmission most years. All six have had continuous CDC ArboNET reporting since 1999.

Data Sources

This analysis used four main data sources, all based on actual measurements rather than simulated data. The four main data sources each provided a different piece of the picture. For human WNV cases, we relied on annual county-level totals from CDC's ArboNET system, which distinguishes between neuroinvasive, non-neuroinvasive, and fatal cases and spans 1999 to

2024 (1). Demographic details—median household income, poverty, population size and density, and the proportion of residents over 65—were drawn from the Census Bureau's Small Area Income and Poverty Estimates and the decennial Census (10). To capture landscape differences across counties, we included the USDA Cropland Data Layer (2008–2021) for the proportions in developed land, wetlands, cropland, and pasture/hay (11). Weather data, covering daily temperature, precipitation, and where available, wind, dewpoint, and humidity, came from NOAA records for 2000–2024 (12). These daily weather measurements were summarized into yearly statistics at the county level, as detailed in the Feature Engineering section.

A fifth data file, which included weekly mosquito surveillance data, was not used in this study. According to the documentation, the weekly mosquito pool positivity and infection rates were not measured directly from mosquito traps. Instead, these values were constructed by spreading the annual ArboNET case totals across the season using a mathematical weighting method. Using this file would have caused data leakage because these predictor values would just reflect the annual case numbers that were already being predicted. Real weekly mosquito trap data are reported by some county vector-control agencies, but at the time of this analysis, a combined dataset covering all six study counties was not available. Including real mosquito surveillance data is a priority for future research.

Feature Engineering

For each county and year, we set out to predict the total number of reported human WNV cases. All predictor variables were based strictly on information from the previous year or earlier, to avoid any risk of looking ahead. Case history features included the number of cases, neuroinvasive cases, and deaths in the prior year, as well as values from two and three years earlier, three- and five-year rolling averages, and year-over-year changes. Demographic variables—like median income, poverty rate, total population, population density, and percent of residents over 65—came from SAIPE and the Census. Land use was represented by the proportion of each county in developed land, wetlands, cropland, or pasture/hay, based on the USDA Cropland Data Layer. For years beyond its 2008–2021 coverage, missing land-use values were filled in with the county median. Weather features were built from daily NOAA data, summarized each year for each county: mean, maximum, and minimum temperatures, total precipitation, temperature and precipitation for June–September, winter minimum temperature, annual growing degree days (base 10°C), and counts of hot (30°C or above) and wet (at least 5 mm) days. We also included county fixed effects (one-hot encoded) to account for persistent differences between counties. The final dataset included 120 county-year observations and 32 features.

Models and Baselines

Three machine learning models were evaluated. Linear regression with no regularization was included as a simple model baseline to assess the contribution of nonlinearity to predictive performance. Random forest regression used 400 trees and a maximum depth of 8, selected as a representative bagging-based ensemble method. Gradient boosting regression used a Poisson loss objective, 300 boosting iterations, and a maximum depth of 4. The Poisson loss is appropriate for non-negative count data and is widely used in epidemiological modeling. All machine learning models were implemented in Python using the scikit-learn library, version 1.8.0.

Four baseline predictors were also evaluated, all computed under the same cross-validation protocol as the machine learning models. Global mean predicted the overall mean of training-set case counts across all training county-years. Per-county climatology predicted the mean of training-set case counts within each county. Persistence predicted the previous year's case count for each county. Three-year rolling mean predicted the mean of the previous three years' case counts for each county. Inclusion of these baselines, particularly per-county climatology and persistence, was essential for evaluating whether machine learning offers meaningful improvement over straightforward heuristics.

Cross-Validation and Evaluation Metrics

Models were evaluated using leave-one-year-out cross-validation, where each year from 2005–2024 was held out in turn while the remaining years were used for training. Performance was assessed using MAE, RMSE, Pearson correlation (r), and R^2, with the RMSE/MAE ratio reported to assess error distribution. Results were also analyzed by county and outbreak status (≥200 cases). Permutation feature importance was calculated as the average increase in MAE after randomly shuffling each feature over 20 repetitions.

References

[1] Centers for Disease Control and Prevention. "West Nile Virus Statistics and Maps." CDC, 2025, www.cdc.gov/west-nile-virus/data-maps/.

[2] Petersen, Lyle R., et al. "West Nile Virus: Review of the Literature." JAMA, vol. 310, no. 3, 2013, pp. 308–315.

[3] Kilpatrick, A. Marm, et al. "Temperature, Viral Genetics, and the Transmission of West Nile Virus by Culex pipiens Mosquitoes." PLoS Pathogens, vol. 4, no. 6, 2008, e1000092.

[4] Ruiz, Marilyn O., et al. "Environmental and Social Determinants of Human Risk during a West Nile Virus Outbreak in the Greater Chicago Area, 2002." International Journal of Health Geographics, vol. 3, no. 1, 2004, article 8.

[5] Holcomb, Karen M., et al. "Multi-Model Prediction of West Nile Virus Neuroinvasive Disease with Machine Learning for Identification of Important Regional Climatic Drivers." GeoHealth, vol. 7, no. 11, 2023, e2023GH000906.

[6] Tonks, Adam, et al. "Forecasting West Nile Virus with Graph Neural Networks: Harnessing Spatial Dependence in Irregularly Sampled Geospatial Data." GeoHealth, vol. 8, no. 7, 2024, e2023GH000784.

[7] Eisen, Lars, and Rebecca J. Eisen. "Using Geographic Information Systems and Decision Support Systems for the Prediction, Prevention, and Control of Vector-Borne Diseases." Annual Review of Entomology, vol. 56, 2011, pp. 41–61.

[8] Chen, Yirong, et al. "Neighbourhood Level Real-Time Forecasting of Dengue Cases in Tropical Urban Singapore." BMC Medicine, vol. 16, 2018, article 129.

[9] Hochreiter, Sepp, and Jürgen Schmidhuber. "Long Short-Term Memory." Neural Computation, vol. 9, no. 8, 1997, pp. 1735–1780.

[10] U.S. Census Bureau. "Small Area Income and Poverty Estimates (SAIPE) Program." U.S. Census Bureau, 2024, www.census.gov/programs-surveys/saipe.html.

[11] U.S. Department of Agriculture, National Agricultural Statistics Service. "Cropland Data Layer." USDA NASS, 2021, www.nass.usda.gov/Research_and_Science/Cropland/.

[12] National Oceanic and Atmospheric Administration. "Climate Data Online: National Climatic Data Center Daily Summaries." NOAA, 2025, www.ncei.noaa.gov/cdo-web/.