POPULOUS: A MULTIMODAL GEOSPATIAL AI MODEL FOR UNDERSTANDING THE CLIMATE-DRIVEN INSURANCE CRISIS IN THE U.S.

Tristan Ballard, Gopal Erinjippurath, Michael Sierks, & Peter Sousounis
Sust Global
San Francisco, California

ABSTRACT

The U.S. home insurance market is experiencing growing instability due to rising climate risks, surging premiums, and widespread non-renewals. Despite the urgent need for data-driven insights, publicly available datasets remain limited, making it difficult to assess and mitigate climate-driven insurance risk. This study introduces an AI-driven modeling framework that predicts home insurance premiums and non-renewals at high spatial resolution. By integrating multimodal data—including Google's Population Dynamics Foundation Model embeddings (leveraging transfer learning), socioeconomic indicators, historical loss records, and physical climate risk—the model first predicts county-level trends before applying embedding-based super-resolution to refine predictions at the zipcode level. Results show strong predictive performance, achieving 93% top-2 accuracy in premium classification and explaining 61% of the variability in non-renewals at the county level. Feature importance analysis highlights climate hazards, such as hurricane and wildfire risk, as key drivers of recent insurance market shifts. By offering fine-scale, data-driven insights, this study not only enhances understanding of how climate change is reshaping the insurance industry but also provides actionable guidance to incentivize resilient building, discourage high-risk development, and inform policies aimed at mitigating climate-related insurance volatility.

1 INTRODUCTION

The U.S. home insurance market is facing unprecedented instability, driven by rising climate risks, soaring premiums, and widespread non-renewals (1). From 2020 to 2023, home insurance premiums have surged by 33% nationwide, with states like Florida and California seeing increases exceeding 50% in some areas (2)(3). Insurers are pulling back from high-risk regions, leaving homeowners with fewer coverage options or pushing them into costly last-resort plans (3). These shifts worsen housing affordability and contribute to broader economic instability (4).

This crisis is exacerbated by a surge in climate-driven disasters (5). As extreme weather events become more frequent and severe, the stability of the home insurance market faces growing uncertainty, raising urgent questions about how climate risk is priced and managed (6)(1). Moreover, climate risk not only drives insurance instability but also exacerbates disparities in coverage, with underinsurance disproportionately affecting high-risk and lower-income communities (7)(5).

Despite the urgency of these challenges, publicly available data on insurance premiums and non-renewals is still limited, hindering efforts to fully assess the impact of climate change on home insurance markets (7)(2). In December 2024, the U.S. Senate released the first nationwide dataset on county-level non-renewals from 2018 to 2023, calling for greater consistency and transparency from insurers moving forward (3). However, datasets on insurance premiums were not included, pre-2018 data is unavailable, and it remains unclear whether such data will be released in the future.

This research introduces a novel AI-driven approach to predicting premiums and non-renewals at high spatial resolution, leveraging multimodal data and foundation model embeddings. The model first predicts county-level trends using a machine learning model with multimodal inputs, then applies embedding-based super-resolution (SR) to refine predictions at the zipcode level for greater precision.

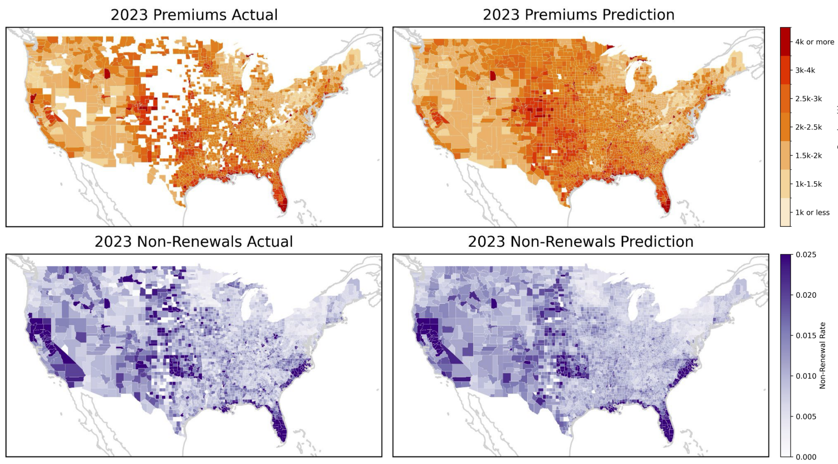

Figure 1: Our multimodal model shows strong performance in reproducing observed 2023 insurance premiums (top) and non-renewals (bottom) at the county level. Our zipcode-level model then refines these predictions to a higher spatial resolution (Fig. 2). Missing data shown in white.

2 DATA

This study utilizes multiple novel datasets, including recently released (December 2024) county-level non-renewal data (3), high-resolution climate risk data, and embeddings from the Population Dynamics Foundation Model (PDFM), a foundation model released in November 2024 (8), which captures spatial and temporal patterns in population and environmental dynamics. Leveraging transfer learning, the model integrates these PDFM embeddings with multimodal data such as housing, demographic data, resilience scores, loss estimates, and climate hazard data (fire, flood, cyclone) (9). We analyze 2023 county-level premium data (2) and non-renewal rate data from 2018-2023 (3), covering roughly 65% of the U.S. homeowners insurance market, and zipcode-level data from California's 10 largest insurers, representing 56% of policies (10)(11). Future climate risk (2024–2050) was also included (9). Further details on the datasets are available in the Appendix.

3 METHODOLOGY

Our insurance super-resolution modeling framework consists of two stages: (1) a county-level model to predict either home insurance premiums or non-renewal rates and (2) a zipcode-level model that refines these predictions to a higher spatial resolution (Fig. A1).

3.1 COUNTY-LEVEL MODELING

We first train a predictive model at the county level. To ensure robust evaluation, we randomly allocate 90% of counties (premiums: $n = 2,160$; non-renewals: $n = 2,763$) for training and 10% (premiums: $n = 240$; non-renewals: $n = 308$) for testing. We explore multiple candidate models—including regularized regression, random forest, and gradient-boosted machines like XGBoost—and select the best-performing model using 10-fold cross-validation with hyperparameter optimization (12). The final model is then evaluated on the test set.

3.2 ZIPCODE-LEVEL SUPER-RESOLUTION

We train a second model to predict at the zipcode level in California, using fine-grained inputs and county-level outputs from the first stage (Fig. A1). While we focus on California due to limited zipcode data in other states, the methodology is applicable to any region. This model incorporates county-level predicted premiums or non-renewal rates as an additional feature, along with input variables from the first stage at the zipcode scale when available. We apply the same model optimization and evaluation process, with the same 90% (premiums: $n = 1,365$; non-renewals: $n = 1,359$) and 10% (premiums: $n = 152$; non-renewals: $n = 151$) train-test split strategy.

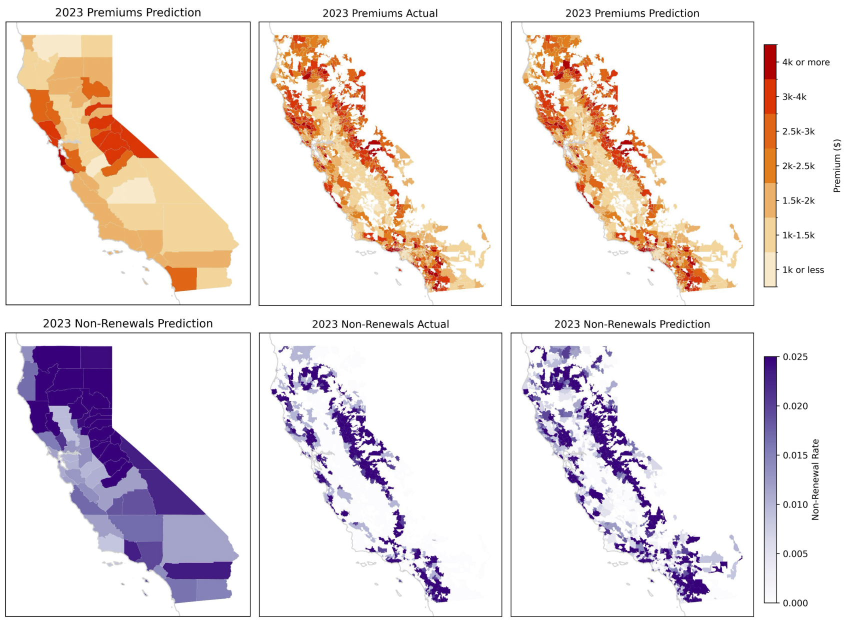

Figure 2: The model generates county-level predictions (left, Fig. 1) that are then super-resolved to the zipcode level (right), matching well the observations (middle). The model shows strong performance in reproducing observed 2023 insurance premiums (top) and non-renewals (bottom) across California at the zipcode level. Test set-only values shown in Figure A3.

4 RESULTS

4.1 COUNTY-LEVEL INSURANCE PREMIUMS AND NON-RENEWALS

The first stage of our model predicts premiums and non-renewals at the county level, capturing dominant spatial patterns and variability both quantitatively and qualitatively (Fig. 1). Since premium data is provided in binned categorical form, we treat this as a classification task and evaluate performance using top-2 accuracy and standard accuracy. Top-2 accuracy measures whether the true premium category falls within the model's two highest-confidence predictions, helping account for cases where the true value lies near a bin edge and reducing undue penalization for predicting an adjacent category. For non-renewal rates, which are available as continuous values, we assess performance using R^2 and RMSE.

For premiums, the model has a very high 93.3% top-2 accuracy score on the test set (Table 1, Fig. A2). Likewise, when applied to predict all U.S. counties, the model effectively captures both the magnitude and spatial variability of the actual premiums (Fig. 1). The model also predicts plausible premiums in counties where there is no groundtruth data (Fig. 1). For non-renewals, the model likewise has strong performance, able to explain the majority (61.4%) of the variability in non-renewals in the test set (Table 1, Fig. A2). When applied to predict all U.S. counties, it too captures well the magnitude and spatial variability in the true non-renewal rates (Fig. 1).

4.2 ZIPCODE-LEVEL SUPER-RESOLUTION (SR) OF INSURANCE PREMIUMS AND NON-RENEWALS

The second stage of our model predicts premiums and non-renewals at the zipcode level, incorporating county-level predictions from the first stage (Fig. 2). We focus on California zipcodes due to data availability. Overall we find that the model captures spatial variability well, with slightly lower performance than the first stage models at the county level (Table 1). For premiums, the model has a very high 89.5% top-2 accuracy score on the test set (Table 1, Fig. A3). For non-renewals, the model has decent performance, able to explain 49.8% of the variability in California non-renewals in the test set (Table 1). Applied across all of California, the zipcode level predictions quantitatively

Premiums Test Set Performance			Non-Renewals Test Set Performance		
County-Level ($n = 240$)			County-Level ($n = 308$)		
	Top-2 Accuracy	Accuracy		R^2	RMSE
Full Model	93.3%	50.4%	Full Model	61.4%	0.0072
Climate Features Removed	92.5%	50.4%	Climate Features Removed	51.2%	0.0081
Zipcode-Level ($n = 152$)			Zipcode-Level ($n = 151$)		
Full Model	89.5%	54.0%	Full Model	49.8%	0.047
Climate Features Removed	89.5%	54.6%	Climate Features Removed	40.1%	0.066

Table 1: Insurance premiums (left) and non-renewals (right) test set performance of the U.S. county-level and California zipcode-level models. Non-renewals performance drops considerably when omitting climate hazard features, suggesting these have had a strong influence.

and qualitatively align with the true values (Fig. 2). We discuss additional input features, such as the price of reinsurance, that could further improve model skill in Section 5.

4.3 Assessing the Role of Climate Change

We conduct two experiments to examine the impact of climate hazards on insurance dynamics. The first analyzes feature importance in the final XGBoost model to identify the most influential variables, with a focus on whether climate features significantly affect premiums and non-renewals. The second experiment removes climate features to assess model performance degradation, though this method may be less robust due to potential compensation by correlated non-climate features.

Feature importance analysis shows that both past and future climate risks, particularly tropical cyclones (hurricanes), significantly influence premiums and non-renewals (Fig. A4). For premiums, past observed cyclones and future cyclone risk are key factors, while non-renewals are heavily impacted by a combination of cyclone, flooding, and wildfire risk (Fig. A4). At the zipcode level in California, wildfire risk is especially influential (Fig. A5). Removing climate features reduces model performance by 17% at the county level and 19% at the zipcode level for non-renewals, providing additional evidence of climate influence, but has little effect on premiums, likely due to the categorical nature of the data (Table 1).

5 Discussion

This study introduces the first US-wide AI-driven framework for predicting home insurance premiums and non-renewal rates, using a multimodal approach that integrates PDFM embeddings, housing and demographic data, historical loss records, and climate risk metrics. The model refines county-level predictions to the zipcode level using embedding-based super-resolution, offering unprecedented spatial precision crucial for policymakers and analysts.

The model performs strongly, with 93.3% top-2 accuracy for premiums at the county level and 61.4% explained variability for non-renewals. At the zipcode level in California, it achieves 89.5% accuracy for premiums and 49.8% explained variability for non-renewals, though these finer-scale predictions still face challenges. Feature importance analysis reveals that climate hazards—particularly hurricanes, wildfires, and flooding—drive recent trends in insurance markets, underscoring the growing impact of climate change. The model could be enhanced by incorporating reinsurance costs, regulatory factors, and additional zipcode-level data beyond California. Its flexibility and super-resolution capability make it adaptable for future improvements, providing insights into how climate and socioeconomic factors shape insurance markets.

This AI-driven framework fills critical data gaps by generating historical insurance estimates and laying the groundwork for understanding future changes in response to extreme events and long-term climate shifts. It offers a robust, data-driven view of local climate-driven insurance risks, with implications for more equitable risk pricing in an era of escalating climate uncertainty. By enabling more accurate pricing and informed decision-making, this model could help mitigate the destabilizing effects of climate shocks on insurance and property markets. Moreover, it could incentivize better building practices and discourage development in high-risk areas, ultimately contributing to climate change mitigation by fostering more resilient communities and reducing vulnerability to climate impacts.

REFERENCES

[1] Aparna Gupta, Abena Owusu, and Jue Wang. Assessing US insurance firms' climate change impact and response. *The Geneva Papers on Risk and Insurance-Issues and Practice*, 49(3):571–604, 2024.

[2] Benjamin J Keys and Philip Mulder. Property insurance and disaster risk: New evidence from mortgage escrow data. Technical report, National Bureau of Economic Research, 2024.

[3] Senate Budget Committee. Next to fall: The climate-driven insurance crisis is here–and getting worse. Technical report, U.S. Senate Budget Committee, December 2024. Staff Report.

[4] Devika Hazra and Patricia Gallagher. Role of insurance in wildfire risk mitigation. *Economic Modelling*, 108:105768, 2022.

[5] Will Jeziorski and Shanthi Ramnath. Homeowners insurance and climate change. *Chicago Fed Letter*, (460):1–6, 2021.

[6] Paulo Issler, Richard Stanton, Carles Vergara-Alert, and Nancy Wallace. Mortgage markets with climate-change risk: Evidence from wildfires in California. *Available at SSRN 3511843*, 2020.

[7] Parinitha Sastry, Ishita Sen, Ana-Maria Tenekedjieva, and Therese C Scharlemann. Climate risk and the us insurance gap: Measurement, drivers and implications. *Drivers and Implications*(May 31, 2024)*, 2024.

[8] Mohit Agarwal, Mimi Sun, Chaitanya Kamath, Arbaaz Muslim, Prithul Sarker, Joydeep Paul, Hector Yee, Marcin Sieniek, Kim Jablonski, Yael Mayer, et al. General geospatial inference with a population dynamics foundation model. *arXiv preprint arXiv:2411.07207*, 2024.

[9] Gopal Erinjippurath and Tristan Ballard. Climate scenario analysis and risk exposure assessments at high resolution, May 18 2023. US Patent App. 17/529,670.

[10] Megan Fan Munce and Sriharsha Devulapalli. California insurance crisis: Map shows which areas are hardest hit by nonrenewals, November 2024.

[11] Megan Fan Munce and Sriharsha Devulapalli. California insurance crisis: Map shows what home insurance costs in every California zip code, September 2024.

[12] Trevor Hastie. The elements of statistical learning: data mining, inference, and prediction, 2009.

[13] U.S. Census Bureau. American community survey. Technical report, U.S. Census Bureau, 2024.

[14] Susan L Cutter, Christopher G Burton, and Christopher T Emrich. Disaster resilience indicators for benchmarking baseline conditions. *Journal of homeland security and emergency management*, 7(1), 2010.

[15] Casey Zuzak, Matthew Mowrer, Emily Goodenough, Jordan Burns, Nicholas Ranalli, and Jesse Rozelle. The national risk index: establishing a nationwide baseline for natural hazard risk in the US. *Natural Hazards*, 114(2):2331–2355, 2022.

[16] Louis Giglio, Tatiana Loboda, David P Roy, Brad Quayle, and Christopher O Justice. An active-fire based burned area mapping algorithm for the MODIS sensor. *Remote sensing of environment*, 113(2):408–420, 2009.

[17] Fritz Policelli, Dan Slayback, Bob Brakenridge, Joe Nigro, Alfred Hubbard, Ben Zaitchik, Mark Carroll, and Hahn Jung. The NASA global flood mapping system. *Remote sensing of hydrological extremes*, pages 47–63, 2017.

[18] Kenneth R Knapp, Michael C Kruk, David H Levinson, Howard J Diamond, and Charles J Neumann. The international best track archive for climate stewardship (IBTrACS) unifying tropical cyclone data. *Bulletin of the American Meteorological Society*, 91(3):363–376, 2010.

[19] Daniel Melser, Trinh Le, and Ummul Ruthbah. Climate change and its impact on home insurance uptake in Australia. *Ecological Economics*, 222:108195, 2024.

[20] Tristan Ballard, Matthew Cooper, Chris Lowrie, and Gopal Erinjippurath. Widespread increases in future wildfire risk to global forest carbon offset projects revealed by explainable AI. *arXiv preprint arXiv:2305.02397*, 2023.

[21] Tristan Ballard and Gopal Erinjippurath. Contrastive learning for climate model bias correction and super-resolution. *arXiv preprint arXiv:2211.07555*, 2022.

A APPENDIX

A.1 EXTENDED DATA DESCRIPTION

A.1.1 INSURANCE PREMIUMS AND NON-RENEWALS

At the county level, we collected county-level insurance premium data for 2023, which is based on millions of individual policy premiums (2). Counties with fewer than 20 observations were removed from the analysis. This data is publicly available only in categorical bins, such as $1,500-2,000$ (Fig. 1), rather than exact values. We collected county-level insurance non-renewal rate (0 to 1) data released in December 2024 by the U.S. Senate Budget Committee (3). The annual non-renewal rate dataset spans 2018 to 2023 and encompasses approximately 65% of the homeowners insurance market. We use the 2023 data to align with the insurance premiums modeling and categorize the zipcode-level data into the same bins as the county-level data for consistency.

At the zipcode level, we analyzed premium and non-renewal data for California, collected from the 10 largest home insurance companies in the state by the California Department of Insurance and aggregated from 2019 to 2024 (10)(11). The data represents 56% of policies in the state, and zipcodes with less than 10% of policies reported were removed from the analysis.

A.1.2 POPULATION DYNAMICS FOUNDATION MODEL (PDFM) EMBEDDINGS

The Population Dynamics Foundation Model (PDFM) model, released by Google Research in November 2024, was trained on a wide variety of data modalities, such as maps, Google search trends, weather patterns, and many others, across zipcodes and counties in the U.S. (8). By using a graph neural network to model the spatial and contextual relationships between these datasets, PDFM generates 330 embeddings that represent the dynamic nature of populations and environments. These embeddings are then used as input to our model, enabling it to leverage the learned connections between human behavior and local factors in predicting insurance pricing. Due to the high dimensionality of the embeddings features, we perform a principal component analysis (PCA) to reduce the number of features while maintaining 67% of the variability in the original embeddings.

A.1.3 HOUSING AND POPULATION FEATURES

We included multiple input features related to housing characteristics and demographics thought to influence insurance policies. These include population, median household income, percent of population with higher education, available at both the county and zipcode level (13). Other features included available only at the county level were the median home age and total homes from the U.S. Census (13), two housing and infrastructure resilience scores from the Baseline Resilience Indicators for Communities index (14), as well as climate-driven historical annual loss estimates from the FEMA national risk index (15).

A.1.4 OBSERVED CLIMATE HAZARD EVENTS

We use observed climate events data provided by Sust Global (9). This includes over 30TB of daily fire and flood data provided by the National Aeronautics and Space Administration (NASA) based on imagery from NASA's MODIS satellites (16)(17). NASA's global fire data, available from 2001 to present at 500m resolution, indicates whether fire was observed at a given pixel. The global flood data, available from 2012 to present at 250m resolution, likewise indicates whether inland or coastal flooding was observed at a given pixel. This satellite data is then aggregated to zipcode and county level. For observed tropical cyclones (hurricanes), the historical track and wind speed data from the International Best Track Archive for Climate Stewardship archive, recording the number of Category 1 or stronger storms hitting annually from 2010 to present (18).

A.1.5 HISTORICAL AND FUTURE CLIMATE RISK EXPOSURE

We use historical and projected climate risk exposure data provided by Sust Global to represent baseline historical climate risk exposure and future exposure(9). Sust Global climate risk data has been used in prior research to understand the current and future impacts of climate hazards on Australia's home insurance market (19). The data includes high-resolution, bias-corrected wildfire, inland and coastal flood, and cyclone probabilities annually over the historical period as well as projected to 2100 under different climate scenarios (9)(20)(21). We use the annual projections from 2024 to 2050 under the Shared Socioeconomic Pathway (SSP5-8.5) high emissions scenario to represent future climate risk. Rather than use the raw hazard probabilities, we use expected annual loss estimates for each hazard, which combine the hazard probabilities with damage curves to assess the likely structural damage from each event class.

B ASSESSING THE ROLE OF CLIMATE CHANGE ADDITIONAL NOTES

We conduct two experiments to explore the relationship between past and future climate hazards, driven by climate change, and insurance dynamics. The first examines feature importance from each model to identify the most influential variables, with the hypothesis that if climate features rank highly, they have had a substantial impact on insurance premiums or non-renewals. In the XGBoost model used here, feature importance is measured by gain, which quantifies the average improvement in model performance when a feature is used for splitting, highlighting its predictive contribution.

The second experiment evaluates how model performance degrades when climate features are removed, based on the hypothesis that a significant drop in performance would indicate their influence. However, this approach is less robust than the first, as correlated non-climate features may partially compensate for the missing climate variables, and the remaining embeddings may still encode some historical weather information.

Through feature importance analysis, we find that both historical climate events and future climate risks have significantly influenced insurance premiums and non-renewals nationwide, particularly tropical cyclones (hurricanes) (Fig. A4). A history of observed tropical cyclones or wildfires accounts for three of the top five most influential features for premiums, alongside future cyclone risk (Fig. A4, left). Similarly, for non-renewals, future tropical cyclone risk and past observed cyclones rank among the top three most important factors, with flooding and wildfire risk also playing key roles (Fig. A4, right). At the zipcode level in California, past and future wildfire risk emerge as some of the most influential factors (Fig. A5).

Consistently, we find that both observed and modeled climate features are critical for predicting non-renewals, with model performance declining by 17% at the county level and 19% at the zipcode level when these features are removed (Table 1). Somewhat unexpectedly, the performance of the premiums model remains largely unchanged when climate features are excluded (Table 1). We suspect this is due to the categorical nature of the premium data, which lacks numerical granularity. If numerical premium data were available—something we are actively working to acquire—we could assess this relationship with greater precision.

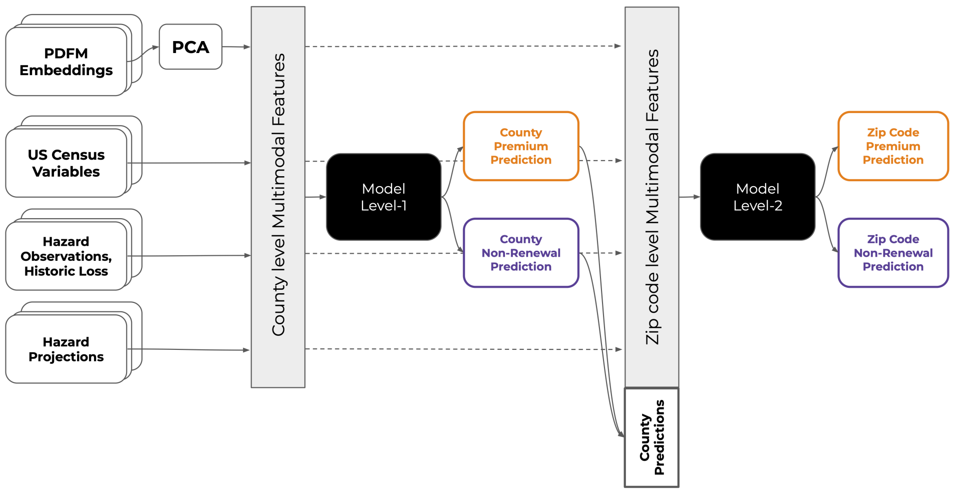

Figure A1: The model is composed of a two-stage architecture used for predicting home insurance premiums and non-renewal rates. The first stage (level 1) trains a county-level predictive model using multimodal inputs, followed by a second stage (level 2) applying zipcode-level super-resolution to refine predictions with county-level outputs and additional fine-grained features.

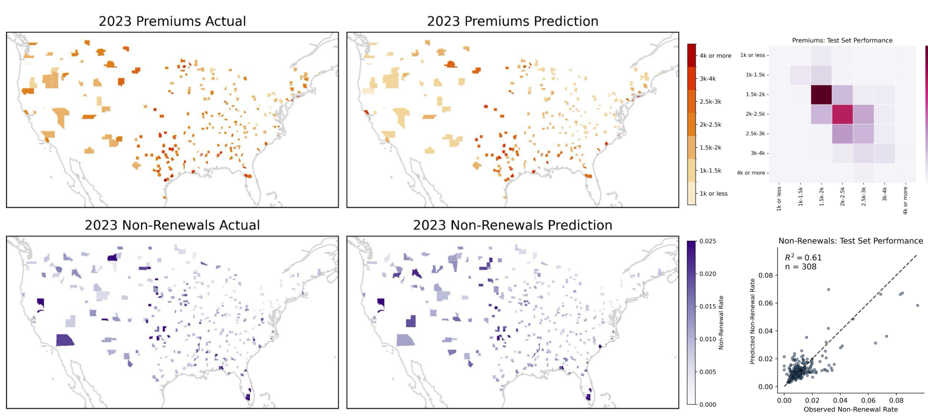

Figure A2: As in Figure 1 but test-set counties only.

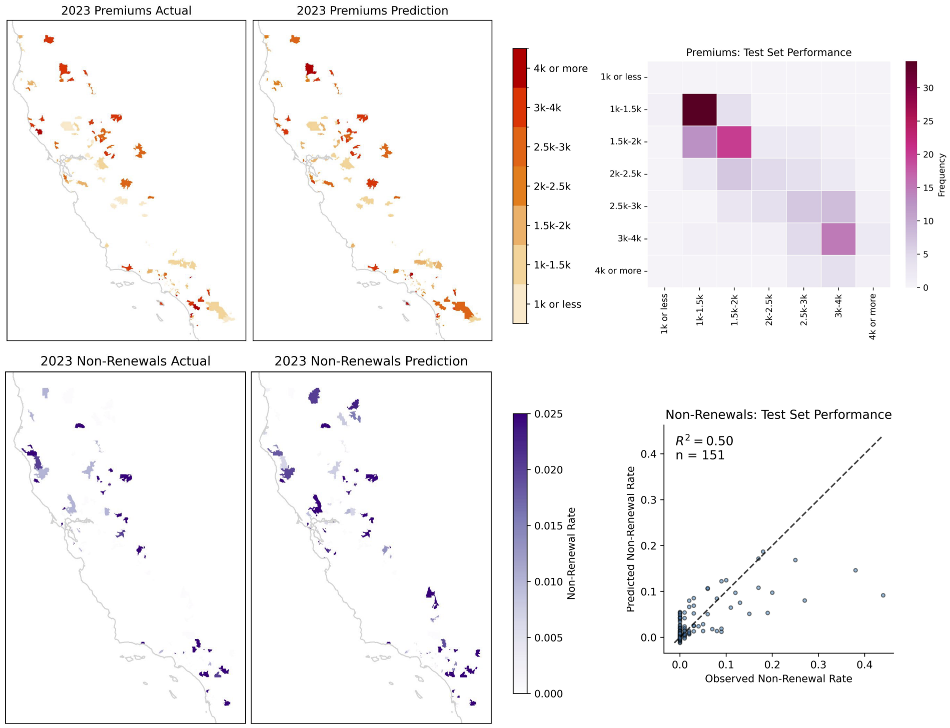

Figure A3: As in Figure 2 but test-set zipcodes only.

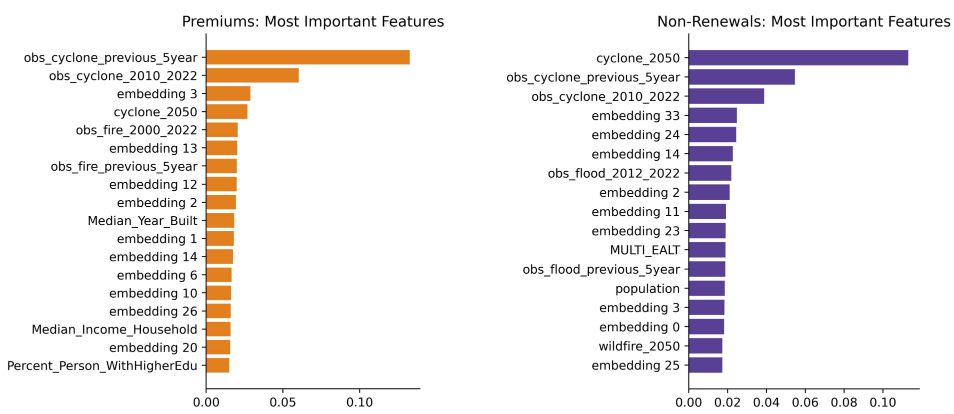

Figure A4: Past climate hazards and future climate risk are among the most influential features for nationwide county-level premiums (left) and non-renewals (right), particularly tropical cyclones (hurricanes). A similar analysis for the zipcode-level model is available in Figure A5.

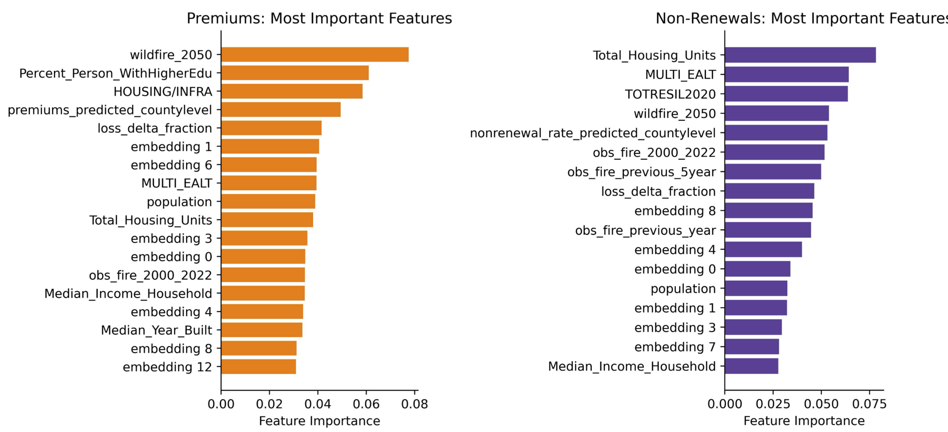

Figure A5: Future wildfire risk (*wildfire_2050*) and past wildfires events are among the most influential features for California zipcode-level premiums (left) and non-renewals (right).