High-Cadence Thermospheric Density Estimation enabled by Machine Learning on Solar Imagery

Shreshth A. Malik* University of Oxford James Walsh* University of Cambridge Giacomo Acciarini University of Surrey

Thomas E. Berger University of Colorado at Boulder Atılım Güneş Baydin University of Oxford

Abstract

Accurate estimation of thermospheric density is critical for precise modeling of satellite drag forces in low Earth orbit (LEO). Improving this estimation is crucial to tasks such as state estimation, collision avoidance, and re-entry calculations. The largest source of uncertainty in determining thermospheric density is modeling the effects of space weather driven by solar and geomagnetic activity. Current operational models rely on ground-based proxy indices which imperfectly correlate with the complexity of solar outputs and geomagnetic responses. In this work, we directly incorporate NASA's Solar Dynamics Observatory (SDO) extreme ultraviolet (EUV) spectral images into a neural thermospheric density model to determine whether the predictive performance of the model is increased by using space-based EUV imagery data instead of, or in addition to, the ground-based proxy indices. We demonstrate that EUV imagery can enable predictions with much higher temporal resolution and replace ground-based proxies while significantly increasing performance relative to current operational models. Our method paves the way for assimilating EUV image data into operational thermospheric density forecasting models for use in LEO satellite navigation processes.

1 Introduction

As the number of space objects in the low-Earth orbit (LEO) grows [1], there is a need for a more accurate model of the thermospheric density variability in order reduce uncertainty in trajectory calculations and improve the accuracy of collision probabilities. The largest source of uncertainty in LEO satellite trajectory calculations is the drag force, which in turn is driven largely by uncertainties in thermospheric density estimation [2]. Currently, the large uncertainties characteristic of operational thermospheric density models translate directly into more frequent conjunction warnings as resident space object trajectory covariances frequently exceed warning thresholds. A major goal of modern space traffic coordination (STC) is to reduce the number of low-probability conjunction warnings issued to satellite operators and thus reduce their workload in adjudicating and mitigating possible collisions. Hence, it is crucial to reduce the uncertainties in thermospheric density models as well as to improve the tracking frequency and characterization of satellites and debris in LEO.

Basal thermospheric density levels are primarily set by heating from the absorption of extreme ultraviolet (EUV) photons from the Sun. In addition, geomagnetic activity due to the interaction of solar wind transients and solar Coronal Mass Ejections (CMEs) with the Earth's magnetosphere can cause impulsive heating and large deviations from basal levels, particularly during geomagnetic storms. Current operational empirical models rely on simple daily proxy indices to represent the basal

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

^{*}Equal contribution.

and transient solar impacts to the thermosphere, e.g. the F10.7 proxy index [3] which imperfectly correlates with true EUV output, and the Kp or Ap indices [4] to represent geomagnetic activity. Both indices are fairly crude representations of physical interactions, lacking in both spatial and temporal resolution that results in underestimating the dynamics of the solar-terrestrial system. The High Accuracy Satellite Drag Model (HASDM) is the operational thermospheric density model used by the US Space Force (USSF) for their STC mission and is based on the empirical JB08 density model [5]. HASDM compensates the shortcomings due to use of proxy indices in the JB08 model by deriving temperature corrections driven by assimilation of calibration satellite trajectory data. However the model does not forecast well during high solar activity periods due to the limitations of using proxy data inputs. For some time now, real-time EUV irradiance data has been available from NOAA weather satellites and from NASA research satellites. It is therefore interesting to investigate whether the use of this high-cadence data in a predictive model would improve density estimation during high solar activity, e.g., during solar flare events that can increase the EUV and X-ray irradiance of the Sun by an order of magnitude over time periods as short as an hour.

In this work, we directly incorporate data from NASA's Solar Dynamics Observatory (SDO) EUV images [6, 7] into a neural thermospheric density model to investigate and quantify how additional information of the EUV spectrum and the solar activity can benefit the prediction of high-cadence thermospheric density variations [8]. We demonstrate that compressed latent space embeddings of SDO imagery generated by a variational autoencoder can replace ground-based proxies while improving predictive accuracy and temporal resolution. This provides a pathway to near real-time, high fidelity density estimation.

2 Background and Related Work

Drivers of thermospheric density change The dynamic changes in the thermospheric density are caused by various space weather processes including EUV irradiance from the Sun and energy deposition from the Earth's magnetosphere during geomagnetic storm periods [9]. EUV photons are deposited in the Earth's upper atmosphere at thermospheric heights (200–1000 km above sea level) where they provide the majority of the baseline energy input and heating. During geomagnetic storms, thermospheric density can rapidly change by tens of percent on top of the baseline set by the solar irradiance with concomitant changes in the drag force experienced by LEO resident space objects. The position of the Earth relative to the Sun, as well as the Earth's rotation cause diurnal and seasonal variations of the EUV irradiance, and short term solar variation from solar flares (the electromagnetic radiation following magnetic eruptions in the solar outer atmosphere) can last 1–4 hours and cause density perturbations of up to 20% [10].

Existing thermospheric density models Full physics-based models, while highly accurate, are computationally intractable for real-time operations [11]. Instead, operational (empirical) models such as JB08 [5], MSIS [12] and HASDM [13] are simplified physics-constrained functions that have been fit to observational data. However, these rely on proxy indices to represent solar EUV irradiance. While computationally efficient, proxies imperfectly capture complex spatiotemporal solar dynamics, limiting density prediction accuracy and cadence to daily predictions. Recent machine-learning-based approaches [14, 15] such as Karman² [16, 17] demonstrate improved performance over operational models. However, while there has been an addition of FISM2 EUV irradiance Stan bands in Karman, these models still do not account for the full complexity of the solar activity.

3 Data

Here we outline the input features and density targets used in this work. For more information, a detailed table is provided in the Appendix A.

SDO data, solar proxies and geomagnetic indices The Solar Dynamics Observatory provides detailed solar imagery from the Atmospheric Imaging Assembly (AIA) [6], EUV Variability Experiment (EVE), and Helioseismic and Magnetic Imager (HMI) [7] instruments. The AIA and HMI data are incorporated through a machine-learning-ready derivative dataset, SDOML [18], containing

²https://github.com/spaceml-org/karman, date of access: September 2023.



Figure 1: F10.7 predictions on test set against ground truth. Parity plot between averaged daily F10.7 predictions from the model against the ground truth (left). Time series of F10.7 values and corresponding high-cadence model predictions (right).

512x512 imagery spanning two ultraviolet and seven extreme ultraviolet wavelengths, as well as magnetic vector field components B_x , B_y , B_z , resulting in a 12 channel image input (example in Appendix A). Data is currently available at a 12 minute time resolution. Empirical models use solar proxy indices: F10.7, M10.7, S10.7, & Y10.7 [8], and geomagnetic indices: Kp, Ap & Dst [4] to account for geomagnetic storm activity, which are all provided daily.

Thermospheric density ground truth The primary source of ground truth for training tasks is TUDelft's thermosphere data [19]. These data contains precise orbit determination-derived thermospheric drag values for CHAMP, GRACE, GRACE-FO, GOCE, and Swarm missions [19, 20, 21]. It contains satellite time, altitude, and location information as well which is used by models for prediction. The data was sourced from multiple providers, with observations occurring at different time intervals. GOCE & Swarm at 10 seconds, CHAMP & GRACE recorded once per 30 seconds. We removed orbits with outlier density measurements caused by missing data (Appendix A).

Dataset preprocessing The dataset was aligned based on a forward-nearest fill policy. All experiments were carried out at a 12-minute cadence, beginning on 4 May 2010, and iterated forward until 29 November 2018. Following Karman [17], non-SDO inputs to models were 0/1 normalized and cyclical features were produced for spatiotemporal inputs. The validation and test subsets were selected by choosing one month of the calendar year, changing per year to the following month.

4 Methods

We follow a three step method for investigating whether EUV imagery can be used for accurate thermospheric density modeling. First we generate general-purpose compressed embeddings of the imagery. Then we add these features as input into a physics-informed deep learning model. Finally, we systematically evaluate the performance of models with and without solar proxies and SDO data in an ablation study. For more details on models, please refer to Appendix B.

Solar embeddings We first extract valuable features from SDOML data. Building on prior work [22], we employ a variational autoencoder (VAE) trained to reconstruct SDOML data from a lowerdimensional encoding, offering a highly compressed representation of the Sun's state for downstream tasks, such as thermospheric density prediction. Our VAE architecture is also enhanced with residual blocks [23], enabling variable latent dimension testing [24].

SDO thermospheric density model In Karman [17], a simple feed-forward network (FFN) was used to process the input empirical model features. Here we additionally embed SDO imagery with the VAE as described and mean-aggregate embeddings from a 4 hour history at a 12 minute resolution



Figure 3: Ablation study model performance. Distribution of absolute percentage error (APE) for each model across altitude ranges (left), and solar activity levels (right).

for prediction. We concatenate this SDO time series embedding with the empirical model feature embedding and process this through a further FFN for density prediction (Figure 2).

Physics-informed residual learning Karman [17] directly predicted the (log) density from input features. Here, we also investigated predicting residuals to a physics-informed model. We use an exponential atmosphere (EXP) here $[25]^3$, which captures the simplest physics, but in principle any baseline can be used, such as another empirical model. This allows the model to focus on correcting the model for the additional information provided by the SDOML images while staying close to the underlying physics. Density predictions vary by orders of magnitude with altitude, so we use residuals in the form: $\mathbf{R}(x, y) = \log(y) - \log(\mathbf{EXP}(x))$, where x are the input features and y is the true density, as targets to train our network. We plan to include more physical inductive biases in the future, such as smoothness in space and time.



Figure 2: Model pipeline diagram. SDOML images are encoded using the VAE (Section 4), aggregated, then concatenated with the other input features in a feed forward neural network (FFN).

5 Results

We show that the inclusion of direct EUV imagery not only allows for higher-cadence reconstruction of solar indices to use with existing models, but also accurate end-to-end density prediction.

Solar embeddings can reconstruct solar indices at high temporal resolution. To validate the embeddings provide relevant information on solar activity, we use the VAE latent means as inputs into a simple feed-forward neural network (Appendix B) to predict the F10.7 solar index. We find that we are able to predict F10.7 at a high accuracy using just SDO embeddings, with a root mean squared error (RMSE) of 8.32. This is significantly better than random embeddings (RMSE = 32.7). Note this is not an auto-regressive model. If we were to include previous F10.7 values, we would expect performance to increase even further. A key byproduct of this investigation is that we now have a higher cadence ML-informed replacement for F10.7 for use in existing operational models.

³https://github.com/lcx366/ATMOS/blob/master/pyatmos/standardatmos/expo.py date of access: September 2023

Model	Solar Imagery	Solar Proxies	Physics informed	MAPE	MAE / 10 ⁻¹²	RMSE / 10 ⁻¹²	MSA
JB08 [5]	X	1	1	21.61	3.99	7.22	44.8
MSIS [12]	X	\checkmark	1	21.54	3.04	5.82	44.7
Karman [17]	X	\checkmark	×	15.61	1.79	3.60	41.0
Karman - Proxies	X	X	×	23.54	2.69	5.27	43.4
Karman - Proxies + EXP. Residual	X	X	√	23.45	2.30	4.85	42.8
Karman + EXP. Residual	×	\checkmark	1	14.18	1.72	4.32	40.2
SDO	1	×	X	16.67	1.93	3.63	41.6
SDO + Proxies	\checkmark	\checkmark	×	15.69	1.80	3.56	41.1
SDO + EXP. residual	\checkmark	X	1	14.50	1.90	4.36	40.8
SDO + Proxies + EXP. residual	1	√	1	14.00	1.72	4.23	40.3

Table 1: Ablation study results on test fold. The mean absolute percentage error (MAPE), mean absolute error (MAE), root mean squared error (RMSE), and median symmetric accuracy (MSA [26]) is shown. Lower is better for all.

F10.7 is only given daily but SDOML imagery is available every 12 minutes, so we can also predict inter-day values (Figure 1).

Solar embeddings can replace proxies in thermospheric density models. Having now validated the efficacy of the embeddings, we utilize them for thermospheric density prediction. We show results for the best performing models on a validation fold across each model type in Table 1. We find that an ML model using SDO has comparative performance to one using solar proxies (Karman [17]), and greatly outperforms the standard empirical models. Including the proxies as well as SDO data also improves performance. The physics-informed method of predicting residuals to a simple physics model results in the best performance, significantly reducing the error in the baseline Karman model as well. Figure 3 shows error rates for each model across altitudes and solar activities. We find that SDO models work better in periods of high solar activity.

6 Conclusions

Contributions In this work, we have shown that we can replace daily ground-based proxies for EUV irradiance such as F10.7 with solar imagery from the SDO mission for thermospheric density predictions. By doing so, we enable accurate, higher cadence, near real-time density estimation for LEO operations. From a science perspective, this shows that SDO imagery contains the relevant information for accurate thermospheric density prediction for the first time.

Further work To ascertain whether our model can effectively capture short-term solar variability such as flares, we should investigate time periods where flares have had a significant impact on density [27]. An initial investigation of the ground truth density data found that flare effects were insignificant compared to noise (Appendix C). Thus we must increase the fidelity of the measurements to get the most benefit from the high temporal resolution EUV data. Research into extracting densities from LEO satellite constellations and engineering data will be a promising area to investigate.

Acknowledgements

This work is the research product of FDL-X (fdlxhelio.org) an initiative of Frontier Development Lab (FDL.ai). FDL is a public/private partnership between NASA, Trillium Technologies and commercial AI partners Google Cloud and NVIDIA Corporation. FDL-X and its outputs have been designed, managed and delivered by Trillium Technologies Inc (trillium.tech). Trillium is a research and development company with a focus on intelligent systems and collaborative communities for planetary stewardship, space exploration and human health. We express our gratitude to Google Cloud for providing extensive computational resources, as well as to NVIDIA Corporation for

use of DGX Cloud. The material is based upon work supported by NASA under award No(s) 80ARC018D0010-80GSFC18F0114. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Aeronautics and Space Administration. FDL aspires to ensure that the latest tools and techniques in Artificial Intelligence (AI) and Machine Learning (ML) are applied to developing open science for all Humankind.

References

- [1] Jonathan C McDowell. "The low earth orbit satellite population and impacts of the SpaceX Starlink constellation". In: *The Astrophysical Journal Letters* 892.2 (2020), p. L36.
- [2] Thomas E Berger et al. "Flying through uncertainty". In: *Space Weather* 18.1 (2020), e2019SW002373.
- [3] KF Tapping. "The 10.7 cm solar radio flux (F10. 7)". In: *Space weather* 11.7 (2013), pp. 394–406.
- [4] Michel Menvielle et al. "Geomagnetic indices". In: *Geomagnetic observations and models*. Springer, 2010, pp. 183–228.
- [5] Bruce Bowman et al. "A new empirical thermospheric density model JB2008 using new solar and geomagnetic indices". In: *AIAA/AAS astrodynamics specialist conference and exhibit*. 2008, p. 6438.
- [6] James R Lemen et al. "The atmospheric imaging assembly (AIA) on the solar dynamics observatory (SDO)". In: *Solar Physics* 275 (2012), pp. 17–40.
- [7] Philip Hanby Scherrer et al. "The helioseismic and magnetic imager (HMI) investigation for the solar dynamics observatory (SDO)". In: *Solar Physics* 275 (2012), pp. 207–227.
- [8] A Vourlidas and S Bruinsma. "EUV irradiance inputs to thermospheric density models: Open issues and path forward". In: *Space Weather* 16.1 (2018), pp. 5–15.
- [9] DJ Knipp, W Kent Tobiska, and BA Emery. "Direct and indirect thermospheric heating sources for solar cycles 21–23". In: *Solar Physics* 224 (2004), pp. 495–505.
- [10] Huijun Le et al. "Global thermospheric disturbances induced by a solar flare: a modeling study". In: *Earth, Planets and Space* 67 (2015), pp. 1–14.
- [11] Liying Qian et al. "The NCAR TIE-GCM: A community model of the coupled thermosphere/ionosphere system". In: *Modeling the ionosphere-thermosphere system* (2014), pp. 73– 83.
- [12] JM Picone et al. "NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues". In: *Journal of Geophysical Research: Space Physics* 107.A12 (2002), SIA–15.
- [13] Mark F Storz et al. "High accuracy satellite drag model (HASDM)". In: Advances in Space Research 36.12 (2005), pp. 2497–2505.
- [14] Richard J Licata and Piyush M Mehta. "Uncertainty quantification techniques for data-driven space weather modeling: thermospheric density application". In: *Scientific Reports* 12.1 (2022), p. 7256.
- [15] Herbert Turner et al. "Machine learning algorithms for improved thermospheric density modeling". In: Dynamic Data Driven Applications Systems: Third International Conference, DDDAS 2020, Boston, MA, USA, October 2-4, 2020, Proceedings 3. Springer. 2020, pp. 143– 151.
- [16] Stefano Bonasera et al. "Dropout and ensemble networks for thermospheric density uncertainty estimation". In: *Bayesian Deep Learning Workshop, NeurIPS*. 2021.
- [17] Giacomo Acciarini et al. "Karman a Machine Learning Software Package for Benchmarking Thermospheric Density Models". In: Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, 19–22 September 2023. 2023.
- [18] Richard Galvez et al. "A machine-learning data set prepared from the NASA solar dynamics observatory mission". In: *The Astrophysical Journal Supplement Series* 242.1 (2019), p. 7.
- [19] Eelco Doornbos. *Thermospheric density and wind determination from satellite dynamics*. Springer Science & Business Media, 2012.
- [20] Christian Siemes et al. "Swarm accelerometer data processing from raw accelerations to thermospheric neutral densities". In: *Earth, Planets and Space* 68 (2016), pp. 1–16.

- [21] Jose van den IJssel et al. "Thermosphere densities derived from Swarm GPS observations". In: *Advances in Space Research* 65.7 (2020), pp. 1758–1771.
- [22] Edward Brown et al. "Learning the solar latent space: sigma-variational autoencoders for multiple channel solar imaging". In: Dec. 2021.
- [23] Kaiming He et al. "Deep residual learning for image recognition". In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.
- [24] Jun-Yan Zhu et al. "Unpaired image-to-image translation using cycle-consistent adversarial networks". In: *Proceedings of the IEEE international conference on computer vision*. 2017, pp. 2223–2232.
- [25] David A Vallado. *Fundamentals of astrodynamics and applications*. Vol. 12. Springer Science & Business Media, 2001.
- [26] Steven Karl Morley, Thiago Vasconcelos Brito, and Daniel T Welling. "Measures of model performance based on the log accuracy ratio". In: *Space Weather* 16.1 (2018), pp. 69–88.
- [27] C Briand, K Doerksen, and F Deleflie. "Solar EUV-Enhancement and Thermospheric Disturbances". In: *Space Weather* 19.12 (2021), e2021SW002840.
- [28] Diederik P Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: *arXiv* preprint arXiv:1412.6980 (2014).

A Data

A detailed breakdown of the data sources used in this work is presented in Table 2. An example of the 12 channel SDO imagery used in this work is shown in Figure 4.

Data Cleaning We originally detected the need for a data cleaning process due to high prediction error in limited circumstances. The cleaning process occurred by first detecting periods of time greater than 30 seconds of continuously missing data. These brief instances in 1.05% of the dataset contributed to examples of sub 10^{-14} kgm⁻³ recordings of density, which are outliers compared to the usual behavior of the thermospheric density across the satellites' orbits. While we have not processed the precise orbit determination data ourselves, and therefore do not possess all elements to assert if these values were actually the thermospheric density experienced by the satellite during its motion, the lack of data during these periods as well as the fact that these values are significant outliers (i.e., by two orders of magnitude) were enough to invalidate these orbits. An example is provided in Figure 5, whereby the full orbit would then be rejected from the dataset time series.



Figure 4: SDO imagery from the SDOML dataset [18], ordered as three magnetic vector components, seven extreme ultraviolet and two ultraviolet. The final image was unavailable due to sensor protection during geomagnetic storming.

B Models

Here we describe the embedding model, proxy prediction model, and thermospheric density model proposed here.

VAE The model architecture and training pipeline is shown in Figure 6. We trained this architecture for ten epochs using the Adam optimizer [28] and a learning rate of 0.001. We used a latent embedding size of 4,096 in this work and a validation size of 0.1%, roughly 2 days.

Name & Type	Description	Granularity & Source	
SDOML (v2) Training Data	Atmospheric Imaging Assembly (AIA; Lemen et al. 2012) 2 ultraviolet, 1600 & 1700 Å 7 extreme ultraviolet, 94, 131, 171, 193, 211, 304,	2010-2020 12 second cadence for AIA	
Imagery 512x512/0.6, 512x512/0.5 arcsec	and 335 A Helioseismic and Magnetic Imager (HMI; Schou et al. 2012) - visible filtergrams processed into: photospheric Dopplergrams, line-of-sight magnetograms and vector magnetograms	NASA Solar Dynamics Observatory (SDO)	
CHAMP Ground Truth	CHAMP, a LEO satellite launched on July 15, 2000, orbits Earth in 90 minutes with an initial altitude of 456 km and an orbit inclination of 87.3°. CHAMP needs 131 days to cover all local times (longitudinal cycle). CHAMP re-entered in	May 2000-Sep 2010 30 second cadence GFZ TUDelft	
GOCE Ground Truth	September 2010. Includes satellite location data. GOCE is a LEO satellite launched on 17 March 2009 at 270km and 96.5° orbit inclination. The mass density estimates were determined from the along-track acceleration of the GOCE satellite's ion propulsion array with changing thrust. The retrieved GOCE mass densities are made available by the ESA. GOCE re-entered in October 2013. Includes satellite location data.	Nov 2009-Oct 2013 10 second cadence ESA Earth Online TUDelft	
GRACE/ GRACE-FO Ground Truth	GRACE-A and GRACE-B were twin satellites launched in March 2002. GRACE-B kept 220 km from GRACE-A. Both spacecraft derived comparable mass densities. GRACE started at 500 km. GRACE had a 95-minute orbit and 89.5° declination. GRACE longitudinal coverage is about 160 days. GRACE reentered in March 2018. GRACE-FO mission maintains GRACE legacy. GRACE-FO is a NASA-GFZ collaboration. Includes satellite location data.	Apr 2002- Nov 2009 30 second cadence NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC) TUDelft	
SWARM-A&B Ground Truth	Alpha, Bravo, and Charlie (A, B, and C) are identical satellites launched into near-polar low Earth orbit simultaneously for the swarm project.SWARM-A,-B, and -C precess at different local times. For example, SWARM-A and -B precess through 12 hours of local time in 133 days, whereas SWARM-C precesses through 12 hours in 144 days, resulting in a local time gap of 10 hours after 5 years of lunch.	European Space Agency (ESA) Swarm Mission FTP from TUDelft	
10.7 Solar Proxies	F, M, S, & Y proxies F10.7 - Flux of solar radiation at 10.7cm wavelength. A long running index. M10.7 - Derived from the Mg II core-to-wing ratio, recording middle-UV near 280nm S10.7 - integrated 26-34nm irradiance measured by the Solar Extreme-UV Monitor above Solar and Heliospheric Observatory Y10.7 - The normalised 81-day F10.7 defines a weighted sum of a solar X-ray region without flare and Lyman-alpha.	Space Environment Technologies	
Dst & Ap Geomagnetic Proxies	Ap index - Enumerated Kp index, a 3 hour quasi-logarithmic index of magnetic activity relative to a calm day curve. Disturbance Storm Time (Dst) index - measure of geomagnetic activity from observatories near the equator, intensity of the globally symmetrical equatorial electrojet.	Celetrack	

Table 2: Overview of data sources used in this work.



Figure 5: Instance of two orbits whereby a greater than 30 second gap in data occurred, leading to an erroneously low density measurement.



Figure 6: ResNetVAE architecture and training.

Proxy indices prediction Model hyperparameters are given in Table 3. We train/validation/test split by month as discussed in Section 4.

SDO thermospheric density model We experimented with a range of time resolutions and histories. Including more SDOML history and resolution did not significantly improve results, so we opt to include 4 hours of history at a 12 minute resolution such that short term solar activity such as flares are able to be captured by the model [10]. We use mean pooling to aggregate the embeddings. We additionally tried other aggregation methods such as max pooling, a learnable linear combination, and attention-weighted pooling, but these did not significantly improve performance. Model hyperparameters are given in Table 4.

C Effects of short-term EUV irradiance are difficult to observe in the ground truth density data

Figure 7 shows ground truth densities during a flare event and geomagnetic storm. We find that in many cases we cannot see the flare in the underlying noise, and hence it is difficult to ascertain whether our model captures this short-term activity. Large (X-flare) events without the presence of CMEs or a geomagnetic storm are rare, and these are the only type of short-term variation that will significantly change thermospheric density and thus be relevant for satellite operations [27].

Hyperparameter	Value	
Optimization		
Loss Function Target Optimizer Learning Rate Batch Size Early Stopping Patience (epochs) Max Epochs	Mean Squared Error (MSE) F10.7 Adam [28] 0.0001 512 10 100	
Architecture		
Hidden Layer Non-linearities Encoder Hidden Dimensions SDO pooling	Leaky-ReLU (slope=0.01) [512,256,256] Mean	
Data		
SDO Embeddings SDO Time Resolution SDO Sequence History	ddingsVAE latent mean (4,096 dimensional)Resolution1 hourence History1 day	

Table 3: Hyperparameter configuration for F10.7 prediction model.

Table 4: Hyperparameter configuration for SDO thermospheric density model.

Hyperparameter	Value		
Optimization			
Loss Function	Mean Squared Error (MSE)		
Target	[Log density] or [log(density) - log(baseline model prediction)]		
Optimizer	Adam [28]		
Learning Rate	0.0001		
Batch Size	64		
Early Stopping Patience (epochs)	2		
Max Epochs	4		
Architecture			
Hidden Layer Non-linearities	Leaky-ReLU (slope=0.01)		
Empirical Features Encoder Hidden	[256,256,256]		
Dimensions			
Combined Encoder Hidden Dimen-	[256,256,256]		
sions			
SDO pooling	Mean		
Data			
SDO Embeddings	VAE latent mean (4,096 dimensional)		
SDO Time Resolution	12 minutes		
SDO Sequence History	4 hours		

Nevertheless, by including high-cadence SDO images, our model is theoretically capable of capturing this. Up-sampling flare events may be necessary for the model to adequately learn from these outliers though. An additional by-product of this observation is that changing the history length and time resolution of SDO images used for prediction did not change accuracies significantly.



Figure 7: Time series ground truth density during flares and geomagnetic storm periods. Solar flares (orange), periods where the Ap geomagnetic index was 8 or greater (red).