

CROSS ATTENTION TO EMBED ODDLY SHAPED DATA AND APPLICATIONS IN IONOSPHERIC MODELING

Anonymous authors

Paper under double-blind review

ABSTRACT

It is desirable to have models of many physical phenomena, yet often data for these phenomena are oddly structured. These structures, such as ungridded and arbitrary length data prevents the use of many types of machine learning techniques, such as feed-forward neural networks. It is thus quite desirable to be able to move this data into a fixed size and shape for easier data ingest. We propose a method of using cross attention to do this. An example of oddly shaped data is Total Electron Content (TEC), or the vertical integral of electron density in the atmosphere. TEC data is calculated using both the position of a satellite and a position on the surface of Earth, giving a non-fixed location per sample. This leads to a splattering of points on the globe where measurements exist that change in shape and amount each time step. We apply our technique to TEC in an autoregressive approach. This allows us to both obtain an embedding describing the global TEC and create completed TEC maps, filling in where measurements are not taken. The global embedding can then be further used in other models.

1 INTRODUCTION

1.1 ODDLY SHAPED DATA

Oddly shaped data is data that does not present itself in a gridded or fixed-length form. Such data can vary in length from sample to sample and also can vary in shape (i.e., distribution) from sample to sample. An excellent example of this is Total Electron Content (TEC) data, which, because of the nature of computing it, does not have a set shape or structure (see Figure 1 for a visual on how TEC data changes shape at different timesteps). Additionally, during a single time step, TEC data is not ordered (i.e., to not introduce bias, the order data is processed should not affect the output).

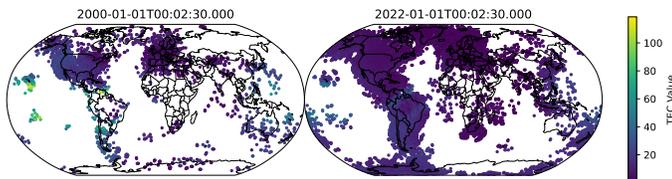


Figure 1: Example of TEC data showing variation in data shape between samples.

Because of the nature of oddly shaped data, it cannot be used in all applications. For example, such data cannot be ingested into feed-forward neural network models. This has led to the lack of using historical TEC data to drive electron density models, despite the advantages that would come with doing so (i.e., giving the model the integral of its output makes computing the output easier).

Despite the lack of using TEC data in electron density models, efforts have been made to “complete” TEC maps, or fill in the missing areas (Pan et al., 2021; Sun et al., 2023). Such a completed TEC map could, in theory, be used as an input to feed-forward neural network models. However, the act of completing data can introduce errors (Schafer & Graham, 2002), as the false data is biased by

054 the techniques used to create it. To avoid this, data can be used as-is, but often this requires some
055 architecture or technique.
056

057 1.2 IONOSPHERE MODELING 058

059 The ionosphere impacts both ground and space assets and is strongly coupled with space weather.
060 Of particular importance is the direct impact on wireless communications (Goodman & Aarons,
061 1990), especially satellite communications. This is because the ionosphere, which is the charged
062 layer of the atmosphere, acts like a plasma and causes electromagnetic waves that pass through it to
063 bend or even reflect. Such a strong impact on communications creates a desire to have an accurate
064 model driven by space weather observations. Such a model would be important to GPS positioning
065 (Dubey et al., 2006; Hoque & Jakowski, 2006), as well as other systems. Additionally, solar storms
066 have can have impacts on power grids and electronics (Eastwood et al., 2017; Kappenman, 2010), as
067 well as human health (Eastwood et al., 2017; Kiznys et al., 2020; Babayev & Allahverdiyeva, 2005).
068 A representative measure of the impact of the ionosphere is the electron density, which directly
069 impacts the aforementioned signal reflection. The electron density of the ionosphere is measurable,
070 but such measurements are very localized, leaving us with a sparse measurement at any one point in
071 time.

072 Electron density is measured in four main ways:

- 073 • GPS Radio Occultation (GPSRO), which analyzes signal paths that go through the iono-
074 sphere between GPSRO satellites and GPS satellites.
- 075 • Ionosondes, which send a variety of frequencies and measure return time.
- 076 • Incoherent Scattering Radar (ISR) (Gordon, 1958), which look at the intensity of the scat-
077 tering of a return radar beam.
- 078 • In-situ measurements, such as Langmuir probes, which directly measure electron density
079 in-place.
080

081 Electron density can also be indirectly measured by measuring the TEC, which is the integral of the
082 electron density along the path between the satellite and receiver used. It is usually measured in the
083 units of $TECU = 1 \times 10^{16} \frac{electrons}{m^2}$.

084 Additionally, there are many correlated or otherwise related measurements we can take of the Earth-
085 Sun system. Many models use indices, which represent magnetic activity, and other common mea-
086 surements.
087

088 1.2.1 EXISTING APPROACHES 089

090 Modeling the ionosphere is done in one of two main ways, either with closed-form physics-based
091 modeling (Anderson et al., 1987; Huba & Krall, 2013; Fuller-Rowell et al., 1996; Scherliess et al.,
092 2009) or empirical-based modeling (Bilitza, 2018; Themens et al., 2017), with both having advan-
093 tages and drawbacks.

094 Physics-based models use known laws of physics, but need functions that are defined everywhere to
095 predict specific outputs (i.e., they can only use directly modeled observations).

096 On the other hand, empirical models can be adjusted to incorporate additional factors, but are not
097 necessarily restrained by physical requirements.
098

099 Recently, Machine Learning (ML) approaches have become popular (Dutta & Cohen, 2022;
100 Sai Gowtam & Tulasi Ram, 2017; Gowtam et al., 2019; Li et al., 2021; Chu et al., 2017; Habarulema
101 et al., 2021). These models perform well compared to existing approaches, but they do not go much
102 further than shallow feed-forward networks.
103

104 2 METHODOLOGY 105

106 2.1 TEC EMBEDDING 107

Our TEC data comes from the Madrigal database (W. Rideout & K. Cariglia).

Our data is split into training, validation, and test sets temporally. Training data is from 15-Oct-1998 to 31-Dec-2020, validation data is from 1-Jan-2021 to 31-Dec-2021, and testing data is from 1-Jan-2022 to 21-May-2023.

We aim to embed TEC data by using an autoencoder. This serves two purposes. First, it allows us to train a model to embed TEC data into a lower, fixed dimension without requiring additional data, and second, it has the side-effect of performing TEC completion, which, while not necessarily desirable for a model input, is useful as a science tool. We will later discuss the TEC completion, but it is not the primary goal of this TEC embedding.

We start our approach to embedding TEC data by following the structure for a transformer outlined by Vaswani et al. (2017). Our model is composed of two parts, the encoder and the decoder or reconstructor. For both we use the same architecture but change which inputs are the keys, values, and queries. For the encoder, we set the keys and values to be the input TEC observation sequences, which have been preprocessed to have their longitude represented by a \sin/\cos pair and latitude scaled to be between -1 and 1. We also normalize TEC. Thus, the observation sequences are scaled latitude, sine of scaled longitude, cosine of scaled longitude, and normalized TEC. This gives the inputs a shape of (batch_size, sequence_length, 4). Importantly, we omit the time (year, day of year, time of day) to force the model to learn the TEC values themselves instead of relying on using time. The query is a parameter of the encoder that acts like a learnable sequence. The multihead attention used is followed by a linear layer that changes from the internal embedding to a target size.

For the reconstructor, we treat the observation sequences as the queries and the embedding as the keys and values. There is a linear layer that goes from the embedded dimension to the attention embedding dimension. For both the embedder and reconstructor we use an internal attention embedding size that is the same.

We also look at chaining these embedding units together like transformer blocks, where the output sequence is used as the input query sequence instead at each step. We notice a lack of improvement by doing this, likely due to the experimental setup, so we do not present such models, but we plan to investigate the possibility of the multi-step embedding in future work.

A graphic of the architecture that we use may be see in Figure 2.

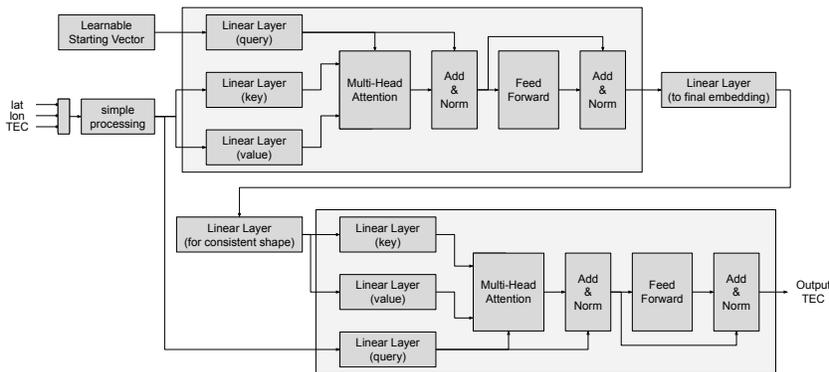


Figure 2: Architecture of autoencoder used to embed and reconstruct TEC data.

We train various versions of this autoencoder with different parameter sizes and compare their performance later in this paper. There we will note the size of various parameters. For all models, an internal attention size of 32 is used.

We also create two versions of a simple TEC predicting model. That is, we have a model that uses location, time, and driving indices (such as those used in electron density modeling that have been shown to be linked to changes in the electron density of the ionosphere) to predict TEC. One is a simple multi-layer feed-forward neural network and the other uses the same structure as the reconstructor used in our autoencoder. These models provide a comparison to the embedding model to demonstrate if more than just cyclical quiet-day information is being encoded. The latter performed much better and thus is the one used in comparisons.

2.2 IONOSPHERE MODELING

We have previously developed ML models for predicting electron density in the ionosphere. We use a simpler version of our model here to compare the benefits of using embedded TEC data to models without such information. These models aim to demonstrate that there is meaningful data captured in the TEC embeddings, and future work aims to improve upon them.

To use the embedded TEC data, we train the TEC model to embed a single timestep’s set of TEC information into a vector of size 16 (we used the smaller sized embedding to save on memory and data storage). For each datapoint used in DINN, the most recent embedded TEC data is concatenated with the standard input vector. Any data where the most recent TEC information is greater than 5 minutes (TEC data is every 5 minutes) is rejected.

Our electron density models are trained on GPSRO data from the CHAMP (UCAR COSMIC Program, 2020a) and GRACE (UCAR COSMIC Program, 2020b) satellites and the COSMIC 1 (UCAR COSMIC Program, 2022) and COSMIC 2 (UCAR COSMIC Program, 2019) satellite constellations. Overall, this data spans approximately 100-600 km altitudes and 20+ years.

We pre-process to reject profiles with peak electron density outside of the 200-450 km altitude range with negative electron densities above 100 km (sub-100km negative densities are instead ignored), an approach introduced by Habarulema et al. (2021) to reject poor-quality samples.

Our data is split into training, validation, and test sets in a similar manner as the TEC data. Training data is from 11-Jun-2001 to 31-Dec-2020, validation data is from 1-Jan-2021 to 31-Dec-2021, and testing data is 1-Jan-2022 to 5-Oct-2022.

For supplemental input data, we use Dst, Kp, ap, and f10.7 from NASA’s OMNI hourly data (Papitashvili & King, 2020). Kp, ap, and Dst help describe the state of Earth’s magnetic field, while f10.7 is the solar flux at a wavelength of 10.7cm.

For the ionosphere models, all models consist of input adjustments (e.g., sin and cos of $\frac{2\pi t}{24}$, where t is the UT hour), similar to the autoencoder of the TEC model. These adjusted inputs are used as the base input tensor. This base input tensor is used in a simple feed-forward neural network we call Deeper Ionosphere Neural Network (DINN), as it has more layers (i.e., it is deeper) than previous feed-forward ionosphere neural networks. This architecture is visualized in Figure 3.

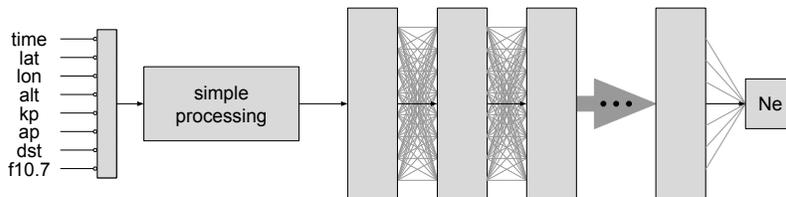


Figure 3: DINN architecture. This model uses feed-forward architecture with fully connected neuron layers, preceded by some simple input adjustment.

Standard model inputs are latitude, longitude, altitude, year, day of year, hour, Kp index, ap index, Dst index, and solar index f10.7. The inputs are adjusted before being used in the feed forward layers as follows: the latitude is scaled between -1 and 1, the longitude is turned into a sine/cosine pair of values. Altitude is scaled to be in megameters. The year is simply used to determine if there is a leap year so that the day of year can be turned into a sine/cosine pair. The hour of the day is also converted into a sine/cosine pair. The Kp index is divided by 10, the Dst and ap indices are divided by 20 and the f10.7 index is divided by 150.

The DINN architecture has an input layer size of 12 (28 when used with embedded TEC, as we have trained it with an embedding size of 16 for the TEC data), with hidden layers of sizes 64, 128, 128, 256, 128, 128, 64, and an output layer of size 1. The feedforward layers use Sigmoid activation functions. The model outputs the log of electron density.

For our models we use Mean Squared Error (MSE) on the log of electron density as our loss function. We train for a total of 100 epochs using a time-grouped batching approach and select the model where validation loss stops decreasing as the trained model. We use a time-grouped batching approach. We use a batch size of 16384 data points.

We also have more complex models, which we use for comparison. We have RDINN, which uses historical K_p , a_p , Dst, and f10.7 values as inputs to RNNs whose subsequent inputs are concatenated like the TEC data is concatenated in DINN. We also have RDINN_XRS, which does the same but also with X-ray flux data. For the latter, GOES X-ray data is used as input data (NOAA), using all available satellites at the time of measurement and averaging satellites together when multiple were available. The historical index values are unscaled. We use hourly data of the past 3 days of observations for all four indices. We also include the past 81 days of daily-averaged f10.7 data. X-ray flux data is used for the past 90 minutes every minute, using both short and long wavelengths.

Each RNN section has a hidden size of 64 with 3 layers followed by a feedforward layer with an input size equal to 64 times the number of types of histories and output size of 16.

RDINN and RDINN_XRS are trained the same way as DINN, although because GOES XRS data is not always present, RDINN_XRS has seen fewer samples (1.28 billion vs 1.36 billion in each epoch).

3 RESULTS/DISCUSSION

3.1 TEC EMBEDDING

We trained a few versions of the autoencoder model. Final models were selected from lowest validation error after it stopped decreasing for more than approximately 5 steps. We have 3 versions of our autoencoder model, each with different embedding sizes. We outline the parameters for each of these in Table 1.

Table 1: Embedded TEC model parameters

NAME	Embedding Sequence Length	Embedding Size
TEC_16	16	1
TEC_64	16	4
TEC_512	64	8

We also trained a decoder/reconstructor that has the same structure as the one used in TEC_512, except instead of encoding we simply use a linear layer to reshape the inputs of the model to what the reconstructor expects. The inputs to this are time, latitude, longitude, K_p , a_p , Dst, and f10.7. Essentially, this model is a TEC predictor. We call this NoEmbedCompare, which we use as a comparison baseline to highlight if the model is capturing anything meaningful in the TEC data besides what can be reconstructed from a set of correlated driving parameters (these are the same as those used to predict electron density in our and other models).

Using these four models, we analyze the reconstruction error on our test data set. This may be seen in Figure 4. Here, we notice that as the size of the embedding increases so too does the accuracy of the reconstructed TEC. Notably, all embedding versions of the model perform noticeably better than the simple prediction model, demonstrating that the embedding operation is capturing relevant information about the specific set of TEC values. From this, we can say that our technique is able to capture information from a sequence of arbitrary length and embed it to a fixed length.

Given that we are able to reconstruct from an embedding to an arbitrary sequence, it stands that we can also use this approach to complete TEC maps. While the accuracy of this completion is out of the scope of this paper (and we plan to do an analysis of this in the future), we present an example of TEC completion in Figure 5, where we can see a match with the expected shape of TEC, particularly the bifurcation along the magnetic field line. Such results are promising for future TEC completion work.

270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323

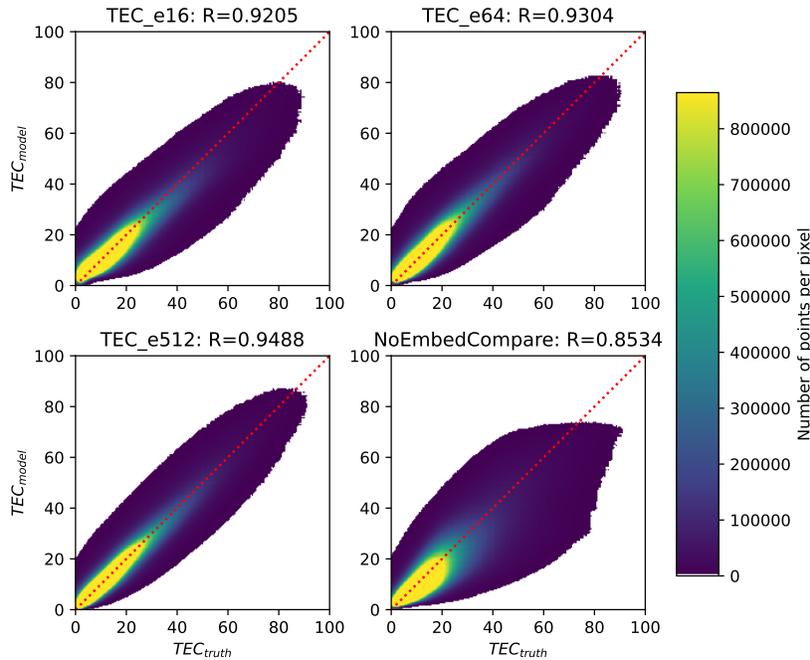


Figure 4: Correlation plots of reconstructed TEC vs testing data, with 2 billion test points. The ideal is represented as a straight line. Yellow data indicates concentrations above the colorbar limit. Included also is the Pearson correlation coefficient, with values closer to 1 representing a distribution closer to falling on the ideal line. Represented here are 3 versions of our autoencoder compared against a simple approach.

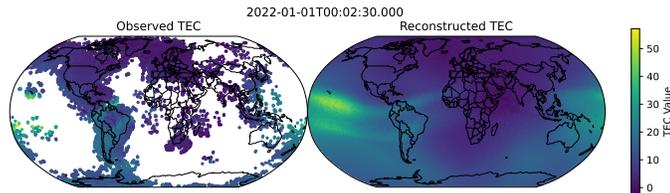


Figure 5: Example of TEC completion using our autoencoder approach. TEC_512 is used in this example.

From these tests, we demonstrate that our embedding technique is successful at embedding an arbitrary length sequence into a meaningful embedding, and that, as expected with larger embedding sizes, improvements can be seen when the embedding space is sufficiently large.

3.2 IONOSPHERE MODELING

We further use the embedded TEC to model the electron density of the ionosphere. In this case, we use TEC_16 for speed and ease of implementation, but acknowledge that model performance would likely increase with larger embeddings. We aim to do this with future work, but the goal of this paper is to present the embedding technique and demonstrate that it provides a functional way to use oddly shaped data, such as TEC. In this testing, we look at our electron density model, DINN, both with and without the size 16 embedding appended after the simple processing step. This gives us DINN_eTEC, which is DINN with the embedded TEC information.

We compare these two approaches using the GPSRO test data, as seen in Figure 6. It is clear that the DINN_TEC model is able to learn a better representation of GPSRO data and it has more accurate predictions.

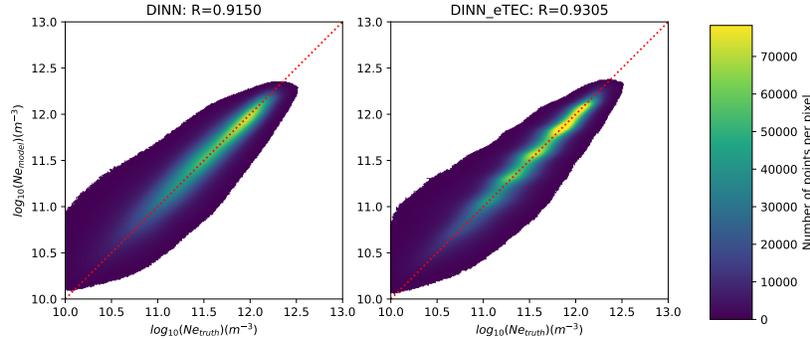


Figure 6: Correlation plots of electron density models vs GPSRO testing data, with 185 million test points. The ideal is represented as a straight line. Yellow data indicates concentrations above the colorbar limit. Included also is the Pearson correlation coefficient, with values closer to 1 representing a distribution closer to falling on the ideal line. Compared here is our DINN model with and without embedded TEC as an additional input.

We also consider ionosonde data, particularly hmf2 and nmf2, or the height and electron density at the peak of the f2 layer of the ionosphere, respectively. The ionosonde data comes from the Global Ionospheric Radio Observatory (GIRO) (Reinisch & Galkin, 2011). We consider high Kp, or storm-time data, with $Kp * 10 \geq 50$. The Pearson correlation coefficient as well as the mean absolute error for this is reported in Figures 7 and 8.

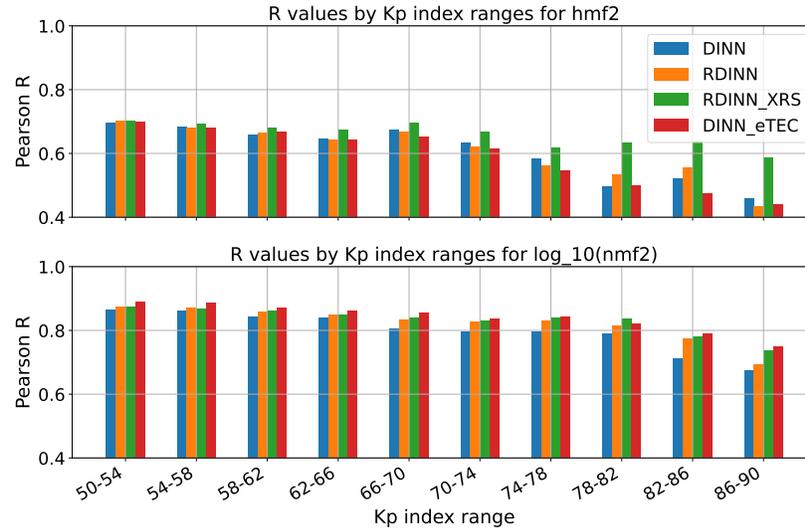


Figure 7: Correlation of the predictions against ionosonde measurements along Kp ranges. Higher is better.

This test demonstrates that the model is able to predict well the peak electron density but not the altitude of the peak electron density. Specifically, we see a decrease at all Kp levels for mean absolute error and an increase in Pearson correlation coefficient when considering nmf2 and comparing DINN

378
 379
 380
 381
 382
 383
 384
 385
 386
 387
 388
 389
 390
 391
 392
 393
 394
 395
 396
 397
 398
 399
 400
 401
 402
 403
 404
 405
 406
 407
 408
 409
 410
 411
 412
 413
 414
 415
 416
 417
 418
 419
 420
 421
 422
 423
 424
 425
 426
 427
 428
 429
 430
 431

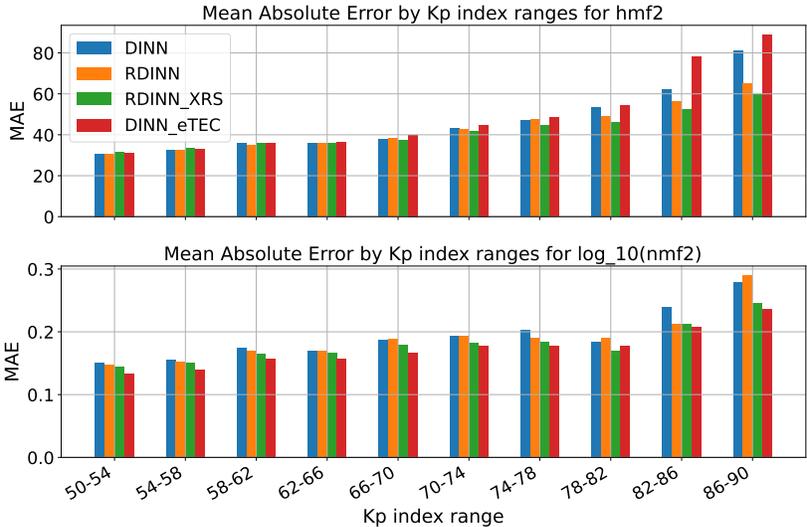


Figure 8: Mean absolute error of the predictions against ionosonde measurements along Kp ranges. Lower is better.

vs DINN_eTEC. We do not see such a trend for hmf2, where often DINN_eTEC performs worse (lower correlation and higher error). We believe that this is because the TEC data has enabled the model to overfit to GPSRO data (we see an odd pattern emerging in Figure 6), as GPSRO data is a product of a slanted TEC measurement and the electron density derived from it is biased by the underlying model used. Additionally, TEC is not as directly correlated with the shape of the 3D electron density profile, but it is strongly correlated with the value of the peak electron density value. Thus, it is reasonable that such a model would perform well in terms of nmf2 predictions but the addition of TEC would not necessarily improve hmf2 predictions.

We include here also results from other versions of our models. We include these because we want to highlight two points. First, that we have lower nmf2 error (and higher correlation) for predictions using the TEC embedding version, even when compared to the more complex models, and second, that hmf2 error can be reduced by including more information about various driving factors of ionization, such as X-ray flux, whose information is not directly captured in historical TEC information. Thus, we have a situation where we demonstrate the limitations of using just embedded TEC information to drive storm predictions while also highlighting its impact in predicting the peaks of electron density in storm predictions. Further work aims to use embedded TEC information alongside X-ray flux to provide even better ionosphere modeling.

4 CONCLUSIONS AND FUTURE WORK

In this paper we proposed a technique to take arbitrarily shaped data and embed it to fixed-length vectors. In particular, we used TEC data, although there are other types of appropriate data that it could be used on, such as weather data. We demonstrated that our technique is able to encode essential information about the sequences, as reconstruction is improved against the lack of an embedding step. Additionally, using larger embeddings increases performance.

Our embedding technique also has the extra benefit of being able to provide completed TEC maps. We aim to further pursue this in future work, but highlight within this paper that the ability to produce reasonable looking TEC maps hints towards the technique’s ability at encoding information.

After demonstrating that information can be successfully encoded using our transformer-like technique, we then demonstrated an example use of embedded arbitrary sequences. By including TEC information in electron density prediction models we see an increase in model performance, particularly in regards to the electron density itself, while we note drawbacks in increasing electron density

432 profile accuracy. We assume that this is due to the loss of vertical information when integrating
 433 (i.e., when measuring TEC the shape of the ionosphere is not captured). We aim to couple the TEC
 434 version of our model with other inputs to make up for the lack vertical information provided by the
 435 TEC in the future.

436 Overall, our paper aims to demonstrate the feasibility of a technique to encode oddly shaped data
 437 and present an example of how such an embedding may be used.
 438

439 REFERENCES

- 440 David N. Anderson, Michael Mendillo, and Bruce Herniter. A semi-empirical low-latitude iono-
 441 spheric model. *Radio Science*, 22(2):292–306, 3 1987. ISSN 1944-799X. doi: 10.1029/
 442 RS022I002P00292. URL [https://agupubs.onlinelibrary.wiley.com/doi/10.](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/RS022i002p00292)
 443 [1029/RS022i002p00292](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/RS022i002p00292).
 444
- 445 Elchin S Babayev and Aysel A Allahverdiyeva. Geomagnetic Storms and their Influence on the
 446 Human Brain Functional State. *Revista CENIC Ciencias Biológicas*, 36, 2005. ISSN 0253-5688.
 447 URL <http://www.redalyc.org/articulo.oa?id=181220525080>.
 448
- 449 Dieter Bilitza. IRI the international standard for the ionosphere. *Advances in Radio Science*, 16:
 450 1–11, 9 2018. ISSN 16849973. doi: 10.5194/ARS-16-1-2018.
 451
- 452 X. Chu, J. Bortnik, W. Li, Q. Ma, R. Denton, C. Yue, V. Angelopoulos, R. M. Thorne, F. Dar-
 453 rouzet, P. Ozhogin, C. A. Kletzing, Y. Wang, and J. Menietti. A neural network model
 454 of three-dimensional dynamic electron density in the inner magnetosphere. *Journal of Geo-*
 455 *physical Research: Space Physics*, 122(9):9183–9197, 9 2017. ISSN 2169-9402. doi:
 456 [10.1002/2017JA024464](https://agupubs.onlinelibrary.wiley.com/doi/10.1002/2017JA024464). URL [https://agupubs.onlinelibrary.wiley.com/doi/](https://agupubs.onlinelibrary.wiley.com/doi/10.1002/2017JA024464)
 457 [10.1002/2017JA024464](https://agupubs.onlinelibrary.wiley.com/doi/10.1002/2017JA024464).
 458
- 459 Smita Dubey, Rashmi Wahi, and A. K. Gwal. Ionospheric effects on GPS positioning. *Advances in*
 460 *Space Research*, 38(11):2478–2484, 1 2006. ISSN 0273-1177. doi: 10.1016/J.ASR.2005.07.030.
 461
- 462 S. Dutta and M. B. Cohen. Improving Electron Density Predictions in the Topside of the Iono-
 463 sphere Using Machine Learning on In Situ Satellite Data. *Space Weather*, 20(9):e2022SW003134,
 464 9 2022. ISSN 1542-7390. doi: 10.1029/2022SW003134. URL [https://agupubs.](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2022SW003134)
 465 [onlinelibrary.wiley.com/doi/10.1029/2022SW003134](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2022SW003134).
 466
- 467 J. P. Eastwood, E. Biffis, M. A. Hapgood, L. Green, M. M. Bisi, R. D. Bentley, R. Wicks, L. A.
 468 McKinnell, M. Gibbs, and C. Burnett. The Economic Impact of Space Weather: Where Do We
 469 Stand? *Risk Analysis*, 37(2):206–218, 2 2017. ISSN 1539-6924. doi: 10.1111/RISA.12765. URL
 470 <https://onlinelibrary.wiley.com/doi/10.1111/risa.12765>.
 471
- 472 T. H. Fuller-Rowell, D. Rees, S. Quegan, R. J. Moffett, M. V. Codrescu, and G. H. Mill-
 473 ward. A coupled thermosphere-ionosphere model (CTIM). In R. W. Schunk (ed.),
 474 *STEP: Handbook of Ionospheric Models*, pp. 217–238. Utah State University, Logan, Utah,
 475 1996. URL [https://www.bc.edu/content/dam/bc1/offices/ISR/SCOSTEP/](https://www.bc.edu/content/dam/bc1/offices/ISR/SCOSTEP/Multimedia/other/ionospheric-models.pdf)
 476 [Multimedia/other/ionospheric-models.pdf](https://www.bc.edu/content/dam/bc1/offices/ISR/SCOSTEP/Multimedia/other/ionospheric-models.pdf).
 477
- 478 John M. Goodman and Jules Aarons. Ionospheric Effects on Modern Electronic Systems. *Proceed-*
 479 *ings of the IEEE*, 78(3):512–528, 1990. ISSN 15582256. doi: 10.1109/5.52228.
 480
- 481 W. E. Gordon. Incoherent Scattering of Radio Waves by Free Electrons with Applications to Space
 482 Exploration by Radar. *Proceedings of the IRE*, 46(11):1824–1829, 1958. ISSN 00968390. doi:
 483 [10.1109/JRPROC.1958.286852](https://doi.org/10.1109/JRPROC.1958.286852).
 484
- 485 V. Sai Gowtam, S. Tulasi Ram, B. Reinisch, and A. Prajapati. A New Artificial Neural Network-
 Based Global Three-Dimensional Ionospheric Model (ANNIM-3D) Using Long-Term Iono-
 spheric Observations: Preliminary Results. *Journal of Geophysical Research: Space Physics*,
 124(6):4639–4657, 6 2019. ISSN 2169-9402. doi: 10.1029/2019JA026540. URL <https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2019JA026540>.

- 486 John Bosco Habarulema, Daniel Okoh, Dalia Burešová, Babatunde Rabi, Mpho Tshisaphungo,
487 Michael Kosch, Ingemar Häggström, Philip J. Erickson, and Marco A. Milla. A global 3-D
488 electron density reconstruction model based on radio occultation data and neural networks. *Journal of Atmospheric and Solar-Terrestrial Physics*, 221:105702, 9 2021. ISSN 1364-6826. doi:
489 10.1016/J.JASTP.2021.105702.
- 491 M. Mainul Hoque and Norbert Jakowski. Higher order ionospheric effects in precise GNSS
492 positioning. *Journal of Geodesy* 2006 81:4, 81(4):259–268, 11 2006. ISSN 1432-1394.
493 doi: 10.1007/S00190-006-0106-0. URL [https://link.springer.com/article/10.](https://link.springer.com/article/10.1007/s00190-006-0106-0)
494 [1007/s00190-006-0106-0](https://link.springer.com/article/10.1007/s00190-006-0106-0).
- 496 J. Huba and J. Krall. Modeling the plasmasphere with SAMI3. *Geophysical Research Letters*, 40
497 (1):6–10, 1 2013. ISSN 1944-8007. doi: 10.1029/2012GL054300. URL [https://agupubs.](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2012GL054300)
498 [onlinelibrary.wiley.com/doi/10.1029/2012GL054300](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2012GL054300).
- 499 John Kappenman. Geomagnetic Storms and Their Impacts on the U.S. Power Grid. Technical report,
500 Metatech Corporation, Goleta, California, 2010.
- 502 Deivydas Kiznys, Jone Vencloviene, and Irena Milvidaitė. The associations of geomagnetic storms,
503 fast solar wind, and stream interaction regions with cardiovascular characteristic in patients with
504 acute coronary syndrome. *Life Sciences in Space Research*, 25:1–8, 5 2020. ISSN 2214-5524.
505 doi: 10.1016/J.LSSR.2020.01.002.
- 506 Wang Li, Dongsheng Zhao, Changyong He, Yi Shen, Andong Hu, and Kefei Zhang. Application
507 of a Multi-Layer Artificial Neural Network in a 3-D Global Electron Density Model Using
508 the Long-Term Observations of COSMIC, Fengyun-3C, and Digisonde. *Space Weather*,
509 19(3):e2020SW002605, 3 2021. ISSN 1542-7390. doi: 10.1029/2020SW002605. URL
510 <https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2020SW002605>.
- 511 NOAA. GOES X-ray Flux — NOAA / NWS Space Weather Prediction Center. URL [https:](https://www.swpc.noaa.gov/products/goes-x-ray-flux)
512 [//www.swpc.noaa.gov/products/goes-x-ray-flux](https://www.swpc.noaa.gov/products/goes-x-ray-flux).
- 514 Yang Pan, Mingwu Jin, Shunrong Zhang, and Yue Deng. TEC Map Completion Through a Deep
515 Learning Model: SNP-GAN. *Space Weather*, 19(11):e2021SW002810, 11 2021. ISSN 1542-
516 7390. doi: 10.1029/2021SW002810. URL [https://onlinelibrary.wiley.com/doi/](https://onlinelibrary.wiley.com/doi/full/10.1029/2021SW002810)
517 [full/10.1029/2021SW002810](https://onlinelibrary.wiley.com/doi/full/10.1029/2021SW002810).
- 518 Natalia E. Papitashvili and Joseph H. King. OMNI Hourly Data, 2020.
- 519 Bodo W. Reinisch and Ivan A. Galkin. Global ionospheric radio observatory (GIRO). *Earth, Planets*
520 *and Space*, 63(4):377–381, 6 2011. ISSN 18805981. doi: 10.5047/EPS.2011.03.001/FIGURES/
521 4. URL [https://earth-planets-space.springeropen.com/articles/10.](https://earth-planets-space.springeropen.com/articles/10.5047/eps.2011.03.001)
522 [5047/eps.2011.03.001](https://earth-planets-space.springeropen.com/articles/10.5047/eps.2011.03.001).
- 524 V. Sai Gowtam and S. Tulasi Ram. An Artificial Neural Network-Based Ionospheric Model to
525 Predict NmF2 and hmF2 Using Long-Term Data Set of FORMOSAT-3/COSMIC Radio Occul-
526 tation Observations: Preliminary Results. *Journal of Geophysical Research: Space Physics*,
527 122(11):743–11, 11 2017. ISSN 2169-9402. doi: 10.1002/2017JA024795. URL [https:](https://agupubs.onlinelibrary.wiley.com/doi/10.1002/2017JA024795)
528 [//agupubs.onlinelibrary.wiley.com/doi/10.1002/2017JA024795](https://agupubs.onlinelibrary.wiley.com/doi/10.1002/2017JA024795).
- 529 Joseph L Schafer and John W Graham. Missing Data: Our View of the State of the Art. 2002. doi:
530 10.1037/1082-989X.7.2.147.
- 532 Ludger Scherliess, Donald C. Thompson, and Robert W. Schunk. Ionospheric dynamics and drivers
533 obtained from a physics-based data assimilation model. *Radio Science*, 44(1):0–32, 2 2009. ISSN
534 1944-799X. doi: 10.1029/2008RS004068. URL [https://agupubs.onlinelibrary.](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2008RS004068)
535 [wiley.com/doi/10.1029/2008RS004068](https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2008RS004068).
- 536 Hu Sun, Yang Chen, Shasha Zou, Jiaen Ren, Yurui Chang, Zihan Wang, and Anthea
537 Coster. Complete Global Total Electron Content Map Dataset based on a Video Impu-
538 tation Algorithm VISTA. *Scientific Data* 2023 10:1, 10(1):1–16, 4 2023. ISSN 2052-
539 4463. doi: 10.1038/s41597-023-02138-7. URL [https://www.nature.com/articles/](https://www.nature.com/articles/s41597-023-02138-7)
[s41597-023-02138-7](https://www.nature.com/articles/s41597-023-02138-7).

540 David R. Themens, P. T. Jayachandran, Ivan Galkin, and Chris Hall. The Empirical Canadian High
541 Arctic Ionospheric Model (E-CHAIM): NmF2 and hmF2. *Journal of Geophysical Research:
542 Space Physics*, 122(8):9015–9031, 8 2017. ISSN 2169-9402. doi: 10.1002/2017JA024398. URL
543 <https://onlinelibrary.wiley.com/doi/full/10.1002/2017JA024398>.
544
545 UCAR COSMIC Program. COSMIC-2 Data Products, 2019.
546 UCAR COSMIC Program. CHAMP Data Products, 2020a.
547 UCAR COSMIC Program. GRACE Data Products, 2020b.
548 UCAR COSMIC Program. COSMIC-1 Data Products, 2022.
549
550 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
551 Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. *Advances in Neural In-
552 formation Processing Systems*, 2017-December:5999–6009, 6 2017. ISSN 10495258. URL
553 <https://arxiv.org/abs/1706.03762v5>.
554
555 W. Rideout and K. Cariglia. CEDAR Madrigal Database. URL [http://cedar.
556 openmadrigal.org](http://cedar.openmadrigal.org).
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593