

Stochastic Parroting in Temporal Attention – Regulating the Diagonal Sink

Victoria Hankemeier¹ and Malte Schilling¹ *

1- University of Münster - Faculty of Mathematics and Computer Science
Autonomous Intelligent Systems Group, Einsteinstr. 62, Münster - Germany

Abstract. Spatio-temporal models analyze spatial structures and temporal dynamics, which makes them prone to information degeneration among space and time. Prior literature has demonstrated that over-squashing in causal attention or temporal convolutions creates a bias on the first tokens. To analyze whether such a bias is present in temporal attention mechanisms, we derive sensitivity bounds on the expected value of the Jacobian of a temporal attention layer. We theoretically show how off-diagonal attention scores depend on the sequence length, and that temporal attention matrices suffer a diagonal attention sink. We suggest regularization methods, and experimentally demonstrate their effectiveness.

1 Introduction

Spatio-temporal deep learning models aim to predict how systems evolve over time and across multiple locations. They leverage historical observations as sequential input as well as using temporal information from multiple locations in order to integrate temporal dependencies as well as spatial ones. A common approach is to use Graph Neural Networks (GNNs) to capture spatial relationships where the node features are usually multivariate time series that are processed using attention or convolution-based layers to learn temporal patterns. A common challenge in spatio-temporal forecasting is to exchange information effectively between both, nearby and distant nodes with two well-known information degeneration issues: over-squashing, which is the insensitivity of a node's features to information contained at distant nodes, and over-smoothing, where node representations become increasingly similar. Most spatio-temporal models either mix temporal and spatial processing at every layer, i.e. Time&Space (T&S), or do all temporal processing first and all spatial processing afterwards, i.e. Time-then-Space(TTS). Both T&S and TTS are affected by spatio-temporal over-squashing because their gradients still multiply across layers in the same way. Over-squashing and over-smoothing are studied extensively for spatial GNNs and Large Language Models (LLMs). However, their role in the temporal component of spatio-temporal models is not yet well understood. Previous work has focused on over-squashing in Temporal Convolutional Networks (TCNs). To address this gap, we provide a theoretical characterization of the Jacobian of Temporal Attention (TA) in which we derive why TA layers can collapse temporal information, which leads to stochastic parroting in TA and showing that this intensifies with increasing sequence length. Guided by these insights, we

*This research has been funded by MAN Truck & Bus SE.

introduce simple architectural adjustments that significantly improve temporal information flow, mitigating over-squashing in a way complementary to existing spatial analysis of GNNs¹.

2 Previous Work on Over-squashing and Over-smoothing

Let $\mathcal{G}_t = (\mathcal{V}, \mathcal{E}, A)$ be the weighted, directed graph at time t , where \mathcal{V} is a set of N nodes that are described by features $X_t \in \mathbb{R}^{N \times d_x}$, the edges are given in \mathcal{E} and the structure is represented by the weighted adjacency matrix $A \in \mathbb{R}^{N \times N}$. The forecasting problem consists of learning a mapping f_θ from past observations $\mathcal{G}_{t-W:t}$ to an expected future sequence $\hat{\mathcal{X}}_{t:t+H} = f_\theta(\mathcal{G}_{t-W:t})$, where W and H denote the input and prediction horizons, respectively.

2.1 Over-squashing and Over-smoothing in GNNs

Di Giovanni et al. [1] laid the groundwork for subsequent research on over-squashing by introducing sensitivity bounds on the Jacobian, which measure how sensitive the final representation of a node v at layer L , $h^{v(L)}$, is to the initial representation of another node u , $h^{u(0)}$. Arroyo et al. [2] further showed that over-smoothing is a representational collapse to a zero fixed point, driven by the same contractive Jacobians responsible for vanishing gradients. This contractivity can lead to signals shrinking before propagating, and over-squashing occurs. For GNNs this has been solved recently ([2], [3], [4], [5]).

2.2 Over-squashing in Temporal Convolution GNNs

Marisca et al. [6] is the only investigation of over-squashing in message-passing Temporal Convolution Network (MPTCNs) known to us. They show that temporal causal convolutions lose focus on recent information, while giving disproportionate weight to inputs from the distant past, creating a “*primacy bias*”.

2.3 Over-squashing in Large Language Transformers

A similar phenomenon can be seen in causal Transformers. In [7], the emergence of an “*attention sink*” in the first token is empirically demonstrated for auto-regressive Large Language Transformers. Barbero et al. [8] examine the Jacobian of auto-regressive decoders, showing that for a number of layers $L \rightarrow \infty$ y_n is only sensitive to the first token. However, they assume attention weights independent of the input, thus focusing only on the value part, and the residual part of the attention layer’s Jacobian, neglecting the key and query.

¹Code is published here: <https://github.com/vicky-hnk/spatio-temp-parroting>,
an extended version here: <https://doi.org/10.48550/arXiv.2602.10956>

3 Over-squashing in Temporal Attention GNNs

The previously shown *primacy bias* does not appear in TA layers with a full attention matrix. As we will derive in the following, the temporal attention has a *diagonal attention sink*, which for very large sequence lengths results in self-copying behavior, i.e. *stochastic parroting*. However, as shown in *SparseBERT* [9] bidirectional Transformers, with a full attention matrix and residual, remain universally expressible even when the diagonal is completely removed, suggesting that the diagonal provides the least important information.

3.1 Deriving Sensitivity Bounds of Temporal Attention

We examine the influence of $\mathbf{x}_j \in \mathbb{R}^{d_x}$ for an input $\mathbf{x}_i \in \mathbb{R}^{d_x}$ on the hidden state $\mathbf{h}_i \in \mathbb{R}^{d_v}$ after one layer by deriving the Jacobian. The index i denotes the source time step, and j the target time step. The hidden state is a weighted sum of the value vectors, where the index m is the summation index that iterates over all source time steps in the sequence, from $m = 1$ to T . With the weight matrices $W^Q \in \mathbb{R}^{d_k \times d_{\text{model}}}$, $W^K \in \mathbb{R}^{d_k \times d_{\text{model}}}$, and $W^V \in \mathbb{R}^{d_v \times d_{\text{model}}}$, we get the query $\mathbf{q}_i = W^Q \mathbf{x}_i$, key $\mathbf{k}_j = W^K \mathbf{x}_j$ and value $\mathbf{v}_j = W^V \mathbf{x}_j$. The scalar attention score of time steps i and j is $e_{ij} = (\mathbf{q}_i^\top \mathbf{k}_j) / \sqrt{d_k}$. The softmax attention is $\alpha_{ij} = \exp(e_{ij}) / \sum_{m=1}^T \exp(e_{im})$, and the hidden state is $\mathbf{h}_i = \sum_{m=1}^T \alpha_{im} \mathbf{v}_m$, with $\sum_{m=1}^T \alpha_{im} = 1$.

$$\frac{\partial \mathbf{h}_i}{\partial \mathbf{x}_j} = \frac{\partial}{\partial \mathbf{x}_j} \left(\sum_{m=1}^T \alpha_{im} \mathbf{v}_m \right) = \sum_{m=1}^T \left[\underbrace{\alpha_{im} \frac{\partial \mathbf{v}_m}{\partial \mathbf{x}_j}}_{\text{Value Path}} + \underbrace{\mathbf{v}_m \frac{\partial \alpha_{im}}{\partial \mathbf{x}_j}}_{\text{Weight Path}} \right] \quad (1)$$

The Jacobian of the **Value Path** is:

$$\sum_{m=1}^T \alpha_{im} \frac{\partial \mathbf{v}_m}{\partial \mathbf{x}_j} = \alpha_{ij} W^V \quad (2)$$

The Jacobian of the **Weight Path** consists of the softmax and the score gradient:

$$\frac{\partial \alpha_{im}}{\partial e_{ik}} = \alpha_{im} (\delta_{mk} - \alpha_{ik}) \quad \frac{\partial e_{ik}}{\partial \mathbf{x}_j} = \frac{1}{\sqrt{d_k}} (\delta_{ij} \mathbf{k}_k^\top W^Q + \delta_{kj} \mathbf{q}_i^\top W^K) \quad (3)$$

The overall Jacobian without residual is thus:

$$\begin{aligned} \frac{\partial \mathbf{h}_i}{\partial \mathbf{x}_j} = & \underbrace{\alpha_{ij} W^V}_{\text{Value Path}} \\ & + \sum_{m=1}^T \left[\mathbf{v}_m \left(\underbrace{\sum_{k=1}^T \alpha_{im} (\delta_{mk} - \alpha_{ik})}_{\text{Softmax Gradient}} \cdot \underbrace{\frac{1}{\sqrt{d_k}} (\delta_{ij} \mathbf{k}_k^\top W^Q + \delta_{kj} \mathbf{q}_i^\top W^K)}_{\text{Score Gradient}} \right) \right] \quad (4) \end{aligned}$$

We derive a uniform expectation of the Jacobian norm over the sequence length $j \in \{1, \dots, T\}$ for the **Value Path**, **Key Path**, and **Query Path** separately:

$$\mathbb{E}[\|J_{ij}^{(\text{value})}\|] = \frac{1}{T} \sum_{j=1}^T \alpha_{ij} \|W^V\| = \frac{\|W^V\|}{T} \quad (5)$$

$$\mathbb{E}[\|J_{ij}^{(\text{key})}\|] = \frac{1}{T\sqrt{d_k}} \sum_{j=1}^T \alpha_{ij} (\|\mathbf{v}_j - \mathbf{h}_i\| \cdot \|(W^K)^\top \mathbf{q}_i\|) \leq \frac{C_K}{T\sqrt{d_k}} \quad (6)$$

$$\mathbb{E}[\|J_{ij}^{(\text{query})}\|] \leq \frac{1}{T\sqrt{d_k}} \sum_{k=1}^T \alpha_{ik} (\|\mathbf{v}_k - \mathbf{h}_i\| \cdot \|(W^Q)^\top \mathbf{k}_k\|) \leq \frac{C_Q}{T\sqrt{d_k}} \quad (7)$$

with $C_K := \max_j \|\mathbf{v}_j - \mathbf{h}_i\| \|(W^K)^\top \mathbf{q}_i\|$, $C_Q := \max_k \|\mathbf{v}_k - \mathbf{h}_i\| \|(W^Q)^\top \mathbf{k}_k\|$. We distinguish the diagonal case ($i = j$), corresponding to local self-updates, and the off-diagonal case ($i \neq j$), corresponding to non-local interactions. The query vector \mathbf{q}_i is defined as $\mathbf{q}_i = W^Q \mathbf{x}_i$. Thus, the derivative of \mathbf{q}_i with respect to \mathbf{x}_j is zero if $i \neq j$, because \mathbf{q}_i is independent of \mathbf{x}_j . We define bounds on the expected value of the Jacobian of the attention layer with a residual connection, using the expected norm for off-diagonal terms ($i \neq j$) and the norm for the diagonal term ($i = j$).

$$\mathbb{E}[\|\delta_{ij}I + J_{ij}^{(\text{total})}\|] \leq \begin{cases} \frac{1}{T} \left(\|W^V\| + \frac{C_K}{\sqrt{d_k}} \right), & i \neq j, \\ \|I\| + \mathbb{E}[\alpha_{ii}] \left(\|W^V\| + \frac{C_K}{\sqrt{d_k}} + \frac{C_Q}{\sqrt{d_k}} \right), & i = j. \end{cases} \quad (8)$$

3.2 Analysis and Regularization of Diagonal Attention Sinks

The $1/T$ -term in the upper part of Eq. 8 shows that the non-local signal strength ($i \neq j$) decays as the sequences are longer ($O(1/T)$). The lower part of Eq. 8 for the diagonal ($i = j$) has a larger bound since it additionally has the query term's bound $C_Q/\sqrt{d_k}$. With a residual path, self-information remains stable regardless of sequence length ($O(1)$). α_{ii} often exceeds α_{ij} as positional encoding (PE) favors information near the diagonal. While emphasizing the diagonal makes sense given the importance of nearby time steps, the extreme case of self-copying does not. To reduce the diagonal attention sink, we evaluated three different regularization methods: a diagonal mask with diagonal entries $-\infty$ and others zero, similar to the approach in *SparseBert* [9], standard dropout on the diagonal of the attention matrix, and a negative scalar penalty added to the raw attention scores.

4 Empirical Evaluation

We evaluate our model on the METR-LA traffic dataset, as described in [10]. The forecasting horizon is 12 steps, using the previous 12 steps as input. Optimization uses AdamW with cosine scheduling, 5 warm-up epochs, 150 epochs in

Table 1: Performance across forecasting horizons (3, 6, 12). MAE, RMSE, and MAPE reported with mean \pm standard deviation.

Model	MAE			RMSE			MAPE		
	3	6	12	3	6	12	3	6	12
(a) No residual	3.165 \pm 0.028	3.671 \pm 0.064	4.627 \pm 0.113	6.190 \pm 0.050	7.414 \pm 0.144	9.315 \pm 0.227	8.705 \pm 0.268	10.286 \pm 0.202	13.406 \pm 0.185
(b) No regularization	2.966 \pm 0.022	3.443 \pm 0.006	3.989 \pm 0.019	5.881 \pm 0.045	7.028 \pm 0.054	8.418 \pm 0.053	7.836 \pm 0.041	9.670 \pm 0.082	11.708 \pm 0.158
(c) Full diagonal mask	2.969 \pm 0.019	3.447 \pm 0.022	3.997 \pm 0.026	5.889 \pm 0.015	7.000 \pm 0.017	8.431 \pm 0.024	7.830\pm0.060	9.666 \pm 0.100	11.734 \pm 0.083
(d) Diagonal dropout	2.959\pm0.001	3.415\pm0.010	3.940\pm0.011	5.845\pm0.029	7.006 \pm 0.051	8.338\pm0.041	7.833 \pm 0.003	9.616\pm0.037	11.570\pm0.047
(e) Diagonal penalty	2.969 \pm 0.027	3.436 \pm 0.037	3.964 \pm 0.041	5.859 \pm 0.034	6.995\pm0.054	8.344 \pm 0.105	7.842 \pm 0.070	9.655 \pm 0.120	11.623 \pm 0.157

total, and an initial learning rate 0.001. We stack a temporal softmax attention block followed by a Graph Convolution Network with learned forward and backward adjacency matrix (T&S). The number of TA heads is 8. We reduce the spatial influence by using a single-layer GCN with about four times fewer parameters than the TA module. The optimal diagonal penalty and dropout were -0.1 and 0.2 , respectively. Since relative PE introduces another bias towards the diagonal, we used absolute PE (tAPE from [11]).

4.1 Results

Table 1 shows the results over four runs with different seeds for a) without residual and without regularization, b) with residual and without regularization, c) with residual and with diagonal mask, d) with residual and with diagonal dropout, and e) with residual and diagonal penalty. The values are the Mean Absolute Error (MAE), Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE) for the forecasted steps 3, 6 and 12. The baseline model without residuals shows substantially higher error than models with residual connections. Using a diagonal mask (c) gives similar results as the variant with residual but without regularization, while dropout and penalty show a significant improvement of around 2.5 percent. Fig. 1 shows the heat maps of final attention matrices after the last epoch. In (a) we see that the attention scores on the diagonal are very large and get smaller with distance if no residual term is used. Fig. 1 (b) shows the attention matrix with residual but without regularization. It is diffuse without clear temporal patterns, as the residual already provides a stable identity mapping. By enforcing regularization on the diagonal elements of the weight matrix, we structurally incentivize the model to allocate learning capacity towards the off-diagonal entries. However, our experiments suggest that the partial control of the diagonal using dropout or penalty is superior to a full diagonal mask. With a full diagonal mask (c) we see similar results as without regularization (b). By masking the diagonal, we suppress the **Query Path**, making the attention less expressive (see Eq. 8). In (d) and (e) we see clear temporal patterns with specific keys attending to specific queries.

5 Conclusion

Recent work on over-squashing in spatio-temporal models has focused on Temporal Convolution [6]. We advance the understanding of information degeneration




Fig. 1: Attention matrices under different regularizations.

in spatio-temporal models by deriving sensitivity bounds on the Jacobian of temporal attention layers. We prove the existence of diagonal attention sinks, induced by the residual and the weight matrix of the query. While full diagonal masks suppress the query part of the softmax attention, resulting in attention emphasizing specific keys for the whole sequence, diagonal penalty and diagonal dropout significantly improve the performance of temporal attention layers.

References

- [1] F. Di Giovanni, L. Giusti, F. Barbero, G. Luise, P. Liò and M. Bronstein, On over-squashing in message passing neural networks: the impact of width, depth, and topology. In *Proc. of the Fortieth Intern. Conf. on Machine Learning (ICML)*, 2023.
- [2] A. Arroyo, A. Gravina, B. Gutteridge, F. Barbero, C. Gallicchio, X. Dong, M. Bronstein and P. Vandergheynst, On Vanishing Gradients, Over-Smoothing, and Over-Squashing in GNNs: Bridging Recurrent and Graph Learning. arXiv:2502.10818, [cs.LG], 2025.
- [3] M. Eliasof, E. Haber and E. Treister, Feature Transportation Improves Graph Neural Networks. *Proc. of the AAAI Conf. on Artificial Intelligence*, 38(11):11874-11882, 2024.
- [4] A. Gravina, M. Eliasof, C. Gallicchio, D. Bacciu and C.-B. Schoenlieb, On Oversquashing in Graph Neural Networks Through the Lens of Dynamical Systems. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(16):16906-16914, 2025.
- [5] S. Heilig, A. Gravina, A. Trenta, C. Gallicchio and D. Bacciu, Port-Hamiltonian Architectural Bias for Long-Range Propagation in Deep Graph Networks. In *Proc. of the Thirteenth International Conference on Learning Representations (ICLR 2025)*, 2025.
- [6] I. Marisca, J. Bamberger, C. Alippi and M. M. Bronstein, Over-squashing in Spatiotemporal Graph Neural Networks. arXiv preprint arXiv:2506.15507, [cs.LG], 2025.
- [7] G. Xiao, Y. Tian, B. Chen, S. Han and M. Lewis, Efficient Streaming Language Models with Attention Sinks. In *Proc. ICLR*, 2024.
- [8] F. Barbero, A. Banino, S. Kapturowski, D. Kumaran, J. G. M. Araújo, A. Vitvitskiy, R. Pascanu and P. Velickovic, Transformers need glasses! Information over-squashing in language tasks. In *Proc. NeurIPS*, 2024.
- [9] H. Shi, J. Gao, X. Ren, H. Xu, X. Liang, Z. Li and J. T. Kwok, SparseBERT: Rethinking the Importance Analysis in Self-attention. In *Proc. ICML*, 2021.
- [10] Y. Li, R. Yu, C. Shahabi, and Y. Liu, Diffusion Convolutional Recurrent Neural Network: Data-driven Traffic Forecasting. In *Proc. Int. Conf. Learn. Represent (ICLR)*, 2018.
- [11] N. M. Foumani, C. W. Tan, G. I. Webb, et al., Improving Position Encoding of Transformers for Multivariate Time Series Classification. *DMKD*, 38:22–48, 2024.