Quantifying Temporal Entropy in Neuromorphic Memory Forgetting: Exploring Advanced Forgetting Models for Robust Long-term Information Storage

Link to publication record in Ulster University Research Portal

Quantifying Temporal Entropy in Neuromorphic Memory Forgetting: Exploring Advanced Forgetting Models for Robust Long-term Information Storage

S. Harrigan, S. Coleman, D. Kerr, J. Quinn, L. Lindsay, K. Madden, S. Rahman, B. Henderson, S. Liu

Ulster University
CRL-SCEIS
United Kingdom

Abstract

This paper presents a progression of a popular neuromorphic memory structure by exploring advanced forgetting models for robust long-term information storage. Inspired by biological neuronal systems, neuromorphic sensors efficiently capture and transmit sensory information using event-based communication. Managing the decay of information over time is a critical aspect, and forgetting models play a vital role in this process. Building upon the foundation of an existing popular neuromorphic memory structure, this study introduces and evaluates four advanced forgetting models: ROT, adaptive, emotional memory enhancement, and context-dependent memory forgetting models. Each model incorporates different factors to modulate the rate of decay or forgetting. Through rigorous experimentation and analysis, these models are compared with the original ROT forgetting model to assess their effectiveness in retaining relevant information while discarding irrelevant or outdated data. The results provide insights into the strengths, limitations, and potential applications of these advanced forgetting models in the context of neuromorphic memory systems, thereby contributing to the progression of this popular neuromorphic memory structure.

Keywords: Neuromorphic, Forgetting Model, Data Structure, Imaging, Machine Vision

1 Introduction

Neuromorphic sensors aim to replicate the spike-based communication paradigm found in biological systems [Lichtsteiner et al., 2006, Brandli et al., 2014a]. By emulating the spiking behaviour of neurons, these sensors capture and transmit sensory information more efficiently and selectively. Instead of continuously sampling and transmitting data, they generate spikes only when significant changes or events occur in the input signals. This event-driven approach reduces data redundancy and enables real-time processing [Gallego et al., 2020, Delbruck, 2008, Lagorce et al., 2016, Alzugaray and Chli, 2018]. Mathematics plays a crucial role in understanding and modelling the behaviour of neuromorphic sensors. Spiking neural network models, based on mathematical equations, simulate the dynamics of artificial neurons and their interactions. These models represent the generation, propagation, and integration of spikes, allowing analysis and prediction of sensor responses to different stimuli. Neuromorphic sensors have a wide range of applications. In computer vision, they excel at detecting and tracking moving objects with high precision and low latency. In robotics, these sensors enable biologically-inspired and efficient perception and interaction with the environment. They also contribute to advancements in artificial intelligence, bio-informatics, and neuromorphic engineering, driving innovation and discovery.

A core component of a neuromorphic sensor is the event datatype, represented as $e = \langle t, c \rangle$ [Gallego et al., 2020,

Wang et al., 2019]. The timestamp t indicates the event's detection or generation time, providing temporal context for precise timing-based computations. The event content c varies based on the sensor type and captured information. It can be a scalar value like intensity or a multidimensional vector representing features or attributes of the input. In event-based vision data, such as dynamic vision neuromorphic sensors (DVS), the content c is represented as $c = \langle x, y, p \rangle$ [Lichtsteiner et al., 2006, Brandli et al., 2014a, Shrestha et al., 2022]. Here, x and y denote the spatial coordinates of the event, and p represents its polarity. Consider all events collected to be $\mathscr{E}$, we consider our working collection of events $\mathbb{E} = \{e \in \mathscr{E} | 0 \leq t \leq \text{end}\}$. By efficiently encoding and transmitting visual information in a sparse and asynchronous manner, event-based neuromorphic vision sensors reduce data redundancy and processing overhead. This approach is particularly suitable for real-time object recognition and visual navigation. Neuromorphic data processing has also led to the development of specialised data structures, such as the neuromorphic ring buffer, time-surface [Lagorce et al., 2016, Sironi et al., 2018, Manderscheid et al., 2019], and the Reduction-Over-Time (ROT) tree [Harrigan et al., 2021a]. The ROT tree, inspired by the brain's structure and functioning, efficiently processes large-scale spatial-temporal data. It combines a hierarchical organisation with forgetting models to determine which information to retain or discard over time. By focusing on relevant data and capturing temporal and spatial patterns, the ROT tree optimises processing capabilities. The ROT tree, and other data structures, can be considered as a novel class of data structures called neuromorphic memory structures, which are designed to retain and manage decay of information over time in neuromorphic event state. In summary, neuromorphic sensors replicate the spike-based communication paradigm of biological systems, enabling efficient and selective sensory information processing. Mathematical models, such as spiking neural networks, play a crucial role in understanding and simulating their behaviour. These sensors find applications in computer vision, robotics, artificial intelligence, and other fields. The event datatype, consisting of both a timestamp and content, captures essential information, and the ROT tree provides an effective computational structure for processing large-scale spatial-temporal data.

2   Forgetting Models

The forgetting models in the following subsections share some similarities in their underlying principles:
Initial Strength: Each model includes an initial strength parameter, representing the strength or clarity of the memory at its inception. This initial strength determines the starting point of the memory's decay or forgetting process. Initial strength is always denoted with χ.
Exponential Decay: All these models are based on exponential decay functions, where the memory strength decreases over time in an exponential manner. The exponential decay reflects the general pattern observed in memory forgetting, where memories tend to fade more rapidly initially and then stabilise at a slower rate over time. Exponential is denoted as exp.
Time Factor: Time plays a crucial role in all these models. The decay or forgetting of memory strength is influenced by the passage of time, represented by the variable t in the equations. As time increases from encoding or retrieval, the more the memory strength diminishes. In this paper we quantify temporal entropy energy as follows:

$$\psi(\hat{x}, \hat{y}, T) = \int_0^T \|\mathbb{E}\| \mid \forall \mathbb{E} \text{ where } \hat{x} = x, \hat{y} = y, t \in [0, T] \tag{1}$$

where ψ is the energy function of a pixel x, y located at time T such that the energy is equal to the cumulative number of events in time $[0, T]$. Using the ROT tree this equals, with decay, the number of nodes currently balanced. Furthermore, to compute the difference in time we treat the difference in time as:

$$\Delta\psi = \psi(\hat{x}, \hat{y}, T) - \psi(\hat{x}, \hat{y}, T - 1) \tag{2}$$

Factors Modulating Decay: While the core decay function is exponential, each model incorporates additional factors that modulate the rate of decay or forgetting. These factors differ between models and include interference factors, difficulty of information, emotional valence, context influence, relevance, significance, and spatial

cues. These additional factors introduce variations in the decay rates, reflecting different aspects of memory forgetting influenced by specific conditions or contexts.

2.1 ROT Forgetting Model

The theory behind the ROT forgetting model is the emulation decay properties exhibited in thermodynamics and reinforced by the forgetting curve [Murre and Dros, 2015]. Imagine a hot cup of coffee cooling down over time; the temperature decreases exponentially, starting from its initial high temperature at a rate determined by the rate of cooling. Similarly, in this model, the memory strength decreases exponentially over time, with the initial strength and forgetting rate influencing the decay pattern. The model is described by:

$$S_i = \chi \times \exp^{(-R^f \times t_i)} \tag{3}$$

The model has three parameters: t which denotes timestamp values, χ which is an initial strength value, and R^f which denotes the forgetting rate within the model. The original rot-Harris [Harrigan et al., 2021a], on which this work is based, made use of the ROT forgetting model to drive forgetting within the ROT tree structure; we provide this definition here for context and this method is used to evaluate the methods that follow by comparing metrics produced from the new methods and this method.

2.2 Adaptive Forgetting Forgetting Model

An analogy for adaptive forgetting [Zaidi et al., 2020] can be found in the process of learning new skills or acquiring knowledge. Imagine you are learning to play a musical instrument, such as a guitar. Initially, as a beginner, you may find it challenging to remember and execute the correct finger placements and chord progressions. However, as you practice and gain experience, your ability to retain and recall this information improves. The model is described by:

$$S_i = \chi \times \exp^{-(0.1+0.5 \times D) \times t_i} \tag{4}$$

The model has three parameters: t which denotes thetimestamp values, χ which is an initial strength value, and D which denotes the difficulty factor. The core difference of this model compared to the ROT forgetting model is the introduction of additional constants as well as the difficulty scoring.

2.3 Emotional Memory Enhancement Forgetting Model

Imagine attending a thrilling roller coaster ride. The intense emotions felt during the ride, such as excitement or fear, can have a profound impact on memory forgetting. In this analogy, emotional memory enhancement [Levens and Phelps, 2008] suggests that the emotional valence associated with an event, like the roller coaster ride, can enhance the strength of memory formation. The initial strength of the memory is multiplied by an emotion factor, which amplifies the memory's intensity. Consequently, these emotionally charged memories may be more vivid and have a lasting impact compared to neutral or less emotionally significant memories. The model is described by:

$$S_i = \chi \times \mathbb{E} \times \exp^{-0.1 \times t_i} \tag{5}$$

The model has three parameters: t denotes the timestamp values, χ is an initial strength value, and $\mathbb{E}$ -is the emotional valence model based on random probability. The emotional memory enhancement forgetting model introduces an emotional valence factor which will cause large variations in the memory forgetting over time.

2.4 Context-Dependent Memory Forgetting Model

As an analogy for Context-Dependent Memory [Raaijmakers and Shiffrin, 1981], it can be compared to a key and lock system. Imagine you have a set of keys, each representing a specific memory. The lock represents the

context or environment in which the memory was initially encoded. When you try to recall a particular memory, the effectiveness of your memory retrieval is influenced by whether the context or environment matches the one in which the memory was encoded. The model is described by:

$$S_i = \chi \times (1 - \mathbb{P}) \times \exp^{-0.1 \times t_i} \tag{6}$$

The model has three parameters: t denotes thetimestamp values, χ is an initial strength value, and $\mathbb{P}$ is a value from a random normalised distribution representing context over time based on area activity. By mixing probabilistic entropy into the equation as **P**, we introduce a measure of variability into the models forgetting over time supplementing harsher pruning activities within the ROT tree structure.

2.5 Multi-dimensional Memory Forgetting Model

The multidimensional memory [Carpenter and Grossberg, 1987] forgetting can be compared to a complex web of interconnected memories, where multiple factors contribute to the strength of forgetting. Imagine exploring a vibrant city for the first time. As you encounter various landmarks, events, and experiences, each memory is influenced by different factors. In this analogy, multidimensional memory forgetting suggests that the relevance, significance, and spatial cues associated with each memory contribute to its overall forgetting strength. The initial strength of the memory is multiplied by a forgetting factor that incorporates these multidimensional factors. This implies that memories with high relevance, significance, and spatial cues are more likely to be retained strongly over time. The exponential decay factor further accounts for the gradual forgetting of these multidimensional memories as time progresses. The model is described by:

$$S_i = \chi \times (r_i \times s_i \times c_i) \times \exp^{-0.1 \times t_i} \tag{7}$$

The model has five parameters: t denotes the timestamp values, χ is an initial strength value, r_i is the relevance factor for the memory at time i (based on recent activity), s_i denotes the significance factor for the memory at time i (based on recent activity), and c_i denotes the spatial cue factor for the memory at time i (based on recent activity).

3 Experiments

In this section, we describe the experiments conducted to evaluate the forgetting models introduced in Section 2. Our evaluation aims to assess the effectiveness of these models by comparing them to the original forgetting curve models. To ensure a fair and consistent assessment, we followed the experimental setting used in the original ROT-Harris paper [Harrigan et al., 2021a]. The ROT-Harris paper introduces a variant of the original Harris corner detection [Harris and Stephens, 1988] which is designed to operate over the neuromorphic ROT tree data structure; the ROT forgetting model is used to drive forgetting of data over time by maintaining information whose difference in time has not yet approached a zero value in the forgetting model response. The ROT-Harris is built on the original ROT [Harrigan et al., 2021b] which compared corner detection methods using rich neuromorphic datasets with corner detection being the key metric of concern, and it is determined that the ROT tree achieves a good balance between accuracy and time when processing the neuromorphic image data.

By adopting the same experimental setting, we aim to establish a direct comparison between the new forgetting models and the original forgetting model in terms of accuracy, F1 score, and the time required to make decisions compared to the original forgetting model. Through rigorous experimentation, we gathered empirical evidence on the performance of each forgetting model. The evaluation process involved running the models on a carefully selected dataset, designed to encompass a wide range of scenarios and challenges.

By conducting experiments under the same conditions as with the original ROT-Harris, we maintain consistency and comparability between the new models and the established baseline. This methodology ensures that our evaluation provides valuable insights into the advancements offered by the new forgetting models in terms

Figure 1: An example output from each of the algorithms with corners formed over frame 107 of the dataset. ROT (green), adaptive (blue), emotional (red), context (cyan), and multi-dimensional (yellow)

of accuracy, F1 score, and the efficiency of decision-making. The comparison of these metrics served as a reliable basis for assessing the performance of the new proposed forgetting models.

To evaluate the forgetting models, we utilised a widely used and publicly available neuromorphic image database [Mueggler et al., 2017], which was also employed with the original ROT-Harris. Specifically, we focus on the shapes datasets within this database. The datasets were captured using a DAVIS240-C camera [Brandli et al., 2014b], which is a hybrid camera capable of both event-based and frame-based imaging. The captured data have a resolution of 240 × 180 pixels. The database reports the ground truth per frame for each of the datasets for analysis. By using this established database and camera, we ensured consistency with the original ROT-Harris and allow for a meaningful comparison of the forgetting models' performance.

From the original ORT-Harris work we inherit an experiment which involves loading neuromorphic vision data into a spatial ROT tree which, using a selected forgetting model, will automatically remove nodes (the neuromorphic data) from the forgetting model as they approach a zero response value; the model is provided with the difference in time between the last time the node was active (last insertion) and the current running time. This difference value is provided to the forgetting model as the tree is searched in time on-demand. Table 3 provides a comprehensive overview of the individual metric scores for each of the forgetting models evaluated. The reported metrics include F1 score, normalised accuracy and time-to-decision for each forgetting model. Figure 1 shows the corner outputs mapped to a 2D image (frame 107 of the dataset) with corners computer from ROT tree whose pruning behaviour is dictated by the models set out in this paper.

The time-to-decision metric measures the time taken by each forgetting model to make a decision or prediction. A lower time-to-decision indicates a faster and more efficient decision-making process. The accuracy metric reflects the overall correctness of the model's predictions, considering both true positive and true negative instances. A higher accuracy score indicates a more accurate and reliable forgetting model.

In Table 3, the data show the performance metrics for the different memory forgetting models. The Multi-dimensional Memory Forgetting model stands out as the top performer, achieving the highest F1 score of 94.8% and the highest accuracy of 90.1%. It also exhibits a relatively low false negative rate of 9.9%, indicating its proficiency in correctly identifying positive cases. This model strikes a balance between high performance and a reasonable time-to-decision of 102 nanoseconds. The Emotional Memory Enhancement model performs reasonably well with a balanced true positive rate and false negative rate, resulting in an F1 score of 66.6% and an accuracy of 50%. It shows potential in enhancing emotional memory retention. On the other hand, the Adaptive Forgetting model demonstrates a relatively low true positive rate and a high false negative rate, resulting in a lower F1 score of 45.8% and an accuracy of 29.8%. Although it has the shortest time-to-decision of 35 nanoseconds, its performance metrics suffer as a trade-off. The Context-Dependent Memory model performs the least effectively among the models. It exhibits a low true positive rate, high false negative rate, and the lowest F1 score of 33.1% and accuracy of 19.9%. Additionally, it has the longest time-to-decision of 139 nanoseconds, making it less desirable in terms of both accuracy and processing speed.

The adaptive forgetting model is the most similar to the original ROT forgetting model so it is surprising to note the contrast in the performance between the two approaches. While adaptive forgetting aims to remove irrelevant or less important information, there is a possibility of discarding valuable information as well. The

fine balance between forgetting irrelevant details and preserving essential knowledge can be challenging to achieve, potentially leading to the loss of important memories. By introducing the constants alongside a difficulty factor it is possible that a softer or harder forgetting model is produced causing data retention/forgetting behaviour. Additionally the poor results observed in the context-dependent memory model can be attributed to the challenges associated with accurately identifying and representing contextual information. In experimental settings, it can be difficult to precisely define and capture the relevant contextual cues for memory retrieval. This ambiguity in context representation can lead to inconsistencies and inaccuracies in the association between memory items and their respective contexts, resulting in reduced performance.

Furthermore, the context-dependent memory model requires significant computational resources due to the complexities of the underlying algorithm. The process of storing and retrieving contextual information for each memory item adds to the computational burden, impacting the overall efficiency and scalability of the system. The higher computation time required by the context-dependent memory model may have contributed to its poorer performance compared to other models. In summary, based on the updated data, the Multi-dimensional Memory Forgetting model remains the best overall performer when compared against the original ROT-Harris forgetting model, providing a balance between high F1 score, accuracy, and a reasonable time-to-decision. The Emotional Memory Enhancement model demonstrates good performance, while the Adaptive Forgetting and Context-Dependent Memory model exhibits comparatively lower accuracy in detecting relevant information.

In Table 2 we report the optimised values discovered during the experiment for each of the forgetting models, it must also be noted that the initial strength factor for each of the models was always set to 1. The experimental hardware used in the study consisted of a system equipped with a 12th Gen Intel(R) Core(TM) i7-1265U processor running at a base frequency of 1.80 GHz. The system was configured with 16.0 GB of installed RAM. In the experiment, all algorithms were developed using the Java programming language. The implementation of the memory forgetting models, as well as the data processing and analysis, were completed in Java. Additionally, the statistical computations for evaluating the performance metrics were conducted offline using Python.

Model	F1 Score	Accuracy	Time-to-Decision [nS]
ROT Forgetting Model	0.82	0.719	68
Adaptive Forgetting	0.458	0.298	**35**
Emotional Memory Enhancement	0.666	0.500	87
Context-Dependent Memory	0.331	0.199	139
Multi-dimensional Memory Forgetting	**0.948**	**0.901**	102

Table 1: Overall Metric Averages for Memory Forgetting Models

Function	Variable	Value
Adaptive Forgetting	D	0.326
Emotional Memory Enhancement	$\mathbb{E}$	1.2
Context Dependent Memory	$\mathbb{P}$	0.379
Multi-dimensional Memory Forgetting	r	0.830
	s	0.120
	c	0.285

Table 2: Optimised values of the variables within each of the forgetting models.

4 Conclusion

The forgetting models presented, including the ROT forgetting model, adaptive forgetting model, emotional memory enhancement model, context-dependent memory model, and multi-dimensional memory forgetting model, aim to capture different aspects of memory forgetting influenced by specific conditions or contexts. Through an evaluation comparing these models to the original forgetting curve model within the ROT tree framework, their effectiveness in driving forgetting was assessed. The results of the experiment demonstrated that the multi-dimensional memory forgetting model outperformed the other forgetting models and the original model in terms of accuracy and time efficiency. By incorporating factors such as relevance, significance, and spatial cues, the multi-dimensional model exhibited a more nuanced decay rate, capturing the complexities of memory forgetting in a comprehensive manner. This finding highlights the importance of considering multiple dimensions when designing forgetting models for neuromorphic memory data structures. The implications of this study are significant for the field of neuromorphic memory data structures and their applications. This can lead to more efficient processing of large-scale spatial-temporal data and improve the performance of tasks such as object recognition, visual navigation, and robotics. Future research in this area will further explore the combination of multiple forgetting models or the development of hybrid models that integrate different factors and principles. Investigating the impact of varying parameters within the forgetting models and assessing their adaptability to different application domains would deepen our understanding of memory forgetting in neuromorphic systems and contribute to further advancements in this field. In conclusion, the advanced multi-dimensional forgetting model presented offers valuable insights into robust long-term information storage in neuromorphic memory. By capturing various aspects of memory forgetting influenced by specific conditions or contexts, this model provides a promising avenue for optimizing memory systems and advancing the field of neuromorphic engineering.

References

[Alzugaray and Chli, 2018] Alzugaray, I. and Chli, M. (2018). Asynchronous corner detection and tracking for event cameras in real time. *IEEE Robotics and Automation Letters*, 3(4):3177–3184.

[Brandli et al., 2014a] Brandli, C., Berner, R., Yang, M., Liu, S.-C., and Delbruck, T. (2014a). A 240× 180 130 db 3 μs latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341.

[Brandli et al., 2014b] Brandli, C., Berner, R., Yang, M., Liu, S. C., and Delbruck, T. (2014b). A 240 x 180 130 dB 3 μs latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341.

[Carpenter and Grossberg, 1987] Carpenter, G. A. and Grossberg, S. (1987). A massively parallel architecture for a self-organizing neural pattern recognition machine. *Computer vision, graphics, and image processing*, 37(1):54–115.

[Delbruck, 2008] Delbruck, T. (2008). Frame-free dynamic digital vision. In *Proceedings of Intl. Symp. on Secure-Life Electronics, Advanced Electronics for Quality Life and Society*, volume 1, pages 21–26. Citeseer.

[Gallego et al., 2020] Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A. J., Conradt, J., Daniilidis, K., et al. (2020). Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180.

[Harrigan et al., 2021a] Harrigan, S., Coleman, S., Ker, D., Yogarajah, P., Fang, Z., and Wu, C. (2021a). Rotharris: A dynamic approach to asynchronous interest point detection. In *2021 17th International Conference on Machine Vision and Applications (MVA)*, pages 1–6. IEEE.

[Harrigan et al., 2021b] Harrigan, S., Coleman, S., Kerr, D., Yogarajah, P., Fang, Z., and Wu, C. (2021b). Reducing-over-time tree for event-based data. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 1765–1772. IEEE.

[Harris and Stephens, 1988] Harris, C. and Stephens, M. (1988). A Combined Corner and Edge Detector. *Procedings of the Alvey Vision Conference 1988*, pages 23.1–23.6.

[Lagorce et al., 2016] Lagorce, X., Orchard, G., Galluppi, F., Shi, B. E., and Benosman, R. B. (2016). Hots: a hierarchy of event-based time-surfaces for pattern recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(7):1346–1359.

[Levens and Phelps, 2008] Levens, S. M. and Phelps, E. A. (2008). Emotion processing effects on interference resolution in working memory. *Emotion*, 8(2):267.

[Lichtsteiner et al., 2006] Lichtsteiner, P., Posch, C., and Delbruck, T. (2006). A 128 x 128 120db 30mw asynchronous vision sensor that responds to relative intensity change. In *2006 IEEE International Solid State Circuits Conference-Digest of Technical Papers*, pages 2060–2069. IEEE.

[Manderscheid et al., 2019] Manderscheid, J., Sironi, A., Bourdis, N., Migliore, D., and Lepetit, V. (2019). Speed invariant time surface for learning to detect corner points with event-based cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10245–10254.

[Mueggler et al., 2017] Mueggler, E., Rebecq, H., Gallego, G., Delbruck, T., and Scaramuzza, D. (2017). The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and SLAM. *The International Journal of Robotics Research*, 36(2):142–149.

[Murre and Dros, 2015] Murre, J. M. and Dros, J. (2015). Replication and analysis of ebbinghaus' forgetting curve. *PloS one*, 10(7):e0120644.

[Raaijmakers and Shiffrin, 1981] Raaijmakers, J. G. and Shiffrin, R. M. (1981). Search of associative memory. *Psychological review*, 88(2):93.

[Shrestha et al., 2022] Shrestha, A., Fang, H., Mei, Z., Rider, D. P., Wu, Q., and Qiu, Q. (2022). A survey on neuromorphic computing: Models and hardware. *IEEE Circuits and Systems Magazine*, 22(2):6–35.

[Sironi et al., 2018] Sironi, A., Brambilla, M., Bourdis, N., Lagorce, X., and Benosman, R. (2018). HATS: Histograms of Averaged Time Surfaces for Robust Event-based Object Classification. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1731–1740.

[Wang et al., 2019] Wang, L., Ho, Y.-S., Yoon, K.-J., et al. (2019). Event-based high dynamic range image and very high frame rate video generation using conditional generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10081–10090.

[Zaidi et al., 2020] Zaidi, A., Caines, A., Moore, R., Buttery, P., and Rice, A. (2020). Adaptive forgetting curves for spaced repetition language learning. In *Artificial Intelligence in Education: 21st International Conference, AIED 2020, Ifrane, Morocco, July 6–10, 2020, Proceedings, Part II 21*, pages 358–363. Springer.