Learning distant cause and effect using only local and immediate credit assignment

David Rawlinson
Incubator 491
dave@agi.io

Abdelrahman Ahmed
Incubator 491
abdel@agi.io

Gideon Kowadlo
Incubator 491
gideon@agi.io

December 13, 2019

Abstract

We present a recurrent neural network memory that uses sparse coding to create a combinatoric encoding of sequential inputs. Using several examples, we show that the network can associate distant causes and effects in a discrete stochastic process, predict partially-observable higher-order sequences, and learn to generate many time-steps of video simulations. The network uses only biologically-plausible, local and immediate credit assignment. Typical memory consumption is 10-30x less than conventional RNNs such as LSTM trained by BackProp Through Time (BPTT). The most significant limitation of the memory is generalization to unseen input sequences. We explore this limitation by measuring next-word prediction perplexity on the Penn Treebank dataset.

1. Introduction

Researchers have been unable to find a biological equivalent to the deep error backpropagation used widely in artificial neural networks (O'Reilly, 1996; Luo et al., 2017). This presents a credit assignment problem: How do biological neurons determine the influence of a synapse on an error that may occur many layers distant, or many steps in the future? The answer to this question is of practical significance because it may yield alternative, superior learning rules. Existing biologically-plausible approaches to the distant credit assignment problem have noticeable performance limitations when compared to *implausible* learning rules (Balduzzi et al., 2015; Bengio et al., 2015). In this paper, we hope to overcome some of these limitations while preserving biological plausibility, using the approach of combinatoric sparse coding.

1.1. Sequence learning

To date, the most successful architectures for sequence learning are Recurrent Neural Networks (RNNs) and auto-regressive, feed-forward networks (Oord et al., 2016). Truncated Back-Propagation Through Time (BPTT) (Sutskever, 2013) unrolls the state of an RNN over a fixed number of time steps t, allowing errors to be assigned to the relevant synaptic weights. The network cannot be trained to exploit causes earlier than t time-steps.

Autoregressive, feed-forward networks utilize a fixed moving window of t most recent inputs, that are processed simultaneously through several layers. Error gradients are back-propagated through the layers. Like truncated BPTT, context more distant than t steps cannot be utilized. Techniques such as dilated convolutions allow t to be quite large (100 steps or more).

Credit assignment in the above methods is widely believed to be biologically implausible due to backpropagation through time and layers.

Credit assignment over time can also be achieved by forward-propagating derivatives of hidden states and outputs with respect to each synaptic weight. Real-Time Recurrent Learning (RTRL) (Williams & Zipser, 1995) is an instance of this concept, but is rarely used in practice due to the difficulty of maintaining viable derivatives.

1.2. Gated memory cells

In most RNNs, such as Elman networks, iterative partial derivative updates are likely to produce vanishing or exploding gradients. Gated memory layers such as Long-Short-Term-Memory (LSTM) (Hochreiter & Schmidhuber, 1997) and Gated Recurrent Units (GRUs) (Chung et al., 2014) mitigate this problem by creating a more robustly differentiable relationship between memory state, gate operations and outputs.

Using gated memory cells, causes and effects may be successfully associated despite separation by hundreds of steps. However, the memory requirements of gated memory layers are much larger than BPTT with ordinary RNNs, due to the requirement to train many gate-control parameters.

1.3. Practical requirements

How much context do we actually need for human-like sequence learning performance? Let's take natural language modelling as an example. Chelba et al (Chelba et al., 2013) showed that on the Billion Word benchmark, an LSTM n-gram model where $n = 13$ was as good as LSTMs with access to longer history. In fact, even with simple smoothed n-gram language models such as Kneser-Ney (Ney et al., 1994), longer histories do not yield practical advantage, in part because the additional higher-order context is unlikely to generalize well, especially in smaller datasets.

Recent work by Bai et al (Bai et al., 2018) showed that in general, feed-forward autoregressive models with dilated convolutions tend to outperform RNNs. They also showed that even where very long term context was definitely necessary, RNNs could not exploit it. This implies that the encoding of context is at least as important as the amount of context.

So in summary, for competitive performance on natural language benchmarks it is likely only necessary to perform credit assignment over context of, perhaps, 100 steps. Since natural languages are tailored to human capabilities, this may be indicative of general human limits. Psychological studies of human problem-solving suggest that we work with relatively simple mental simulations of short sequences of highly abstracted features - longer sequences are made shorter by abstraction (Khemlani et al., 2013).

1.4. Biologically plausible criteria

We adopted the following criteria for biological plausibility:

- Only local credit assignment. No back-propagation of errors between cell-layers

- No synaptic memory beyond the current and/or next step

- No time-travel, making use of past or future inputs or hidden states

Note that we only aim to address the distant credit assignment problem, and do not claim that our approach is otherwise an accurate reproduction of biological network learning.

The computational capabilities of single-layer networks are very limited, especially in comparison to two-layer networks. Biological neurons perform "dendrite computation", involving integration and nonlinearities within dendrite subtrees (Guerguiev et al., 2017). This is computationally equivalent to 2 or 3 typical artificial neural network layers. For this reason we allow ourselves to use error backpropagation across two ANN layers, under the assumption that this could approximate dendrite computation within a single biological cell layer, and training signals inside cells.

1.5. Memorization approach

Gated memory layers learn to selectively remember important features. If we can only back-propagate errors across 2 or 3 time steps, how can we learn to retain data that won't be useful for tens or hundreds of steps?

One option is to simply remember everything, which means we could use any past observation to predict a future effect. At first glance this seems impractical, even if we were to limit observations to one hundred steps; but Hawkins & Ahmad (Hawkins & Ahmad, 2016) showed that a form of sparse coding can allow networks to have vast representational capacity, by virtue of all the *combinations* in which cells can be active.

Sparse coding simply means that most cells have zero value - they are inactive (Olshausen & Field, 1997). In Hawkins & Ahmad's HTM model, the few active cells jointly represent not just the current input, but previous inputs as well. This is achieved by applying local inhibition to groups of cells - within each group, different cells fire in response to different sequential contexts. Since a long context is encoded in the current state without filtering for salience, we will refer to this strategy as the memorization approach.

2. Method

We will attempt to implement the memorization approach to sparse sequence encoding using a more conventional RNN, trained with stochastic gradient descent. For convenience we will refer to our approach as recurrent sparse memory (RSM). RSM is a predictive autoencoder derived from the sparse autoencoders developed by Makhzani and Frey (Makhzani & Frey, 2013; 2015). In their algorithm, a fixed top-k sparsity is applied; sparseness is the only non-linearity.

An autoencoder learns to reconstruct the current external input $\mathbf{x}^F(t)$ through a bottleneck. To make an autoencoder learn transitions, we instead train it to generate a prediction $\hat{\mathbf{x}}^F(t)$ of the next external input $\mathbf{x}^F(t + 1)$ using a mean square error loss.

To allow the network to use previous inputs as context for the prediction, we add a recurrent input $\mathbf{x}^R$, containing a copy of layer cell activity from the previous step - similar to an Elman RNN architecture.

In some experiments we use more than one RSM layer. In this case we require an additional input from the adjacent higher layer $\mathbf{x}^B$. The superscript identifiers F, R, B stand for feed-Forward, Recurrent, and feed-Back respectively.
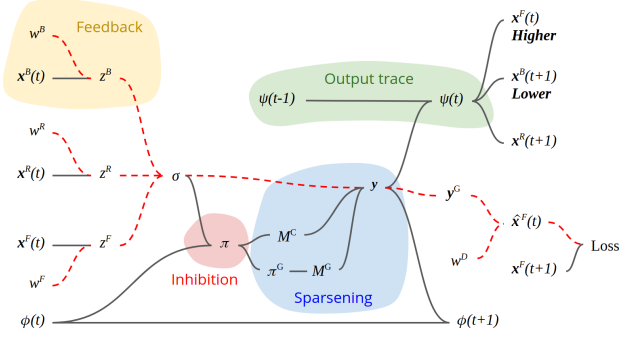
Figure 1: RSM diagram. RSM is trained to generate predictions $\hat{\mathbf{x}}^F(t)$ of its next external input $\mathbf{x}^F(t+1)$. Dashed red lines show how this loss propagates to the learned parameters w^F, w^R, w^B and w^D. The single-layer model is a standard recurrent network with the addition of inhibition (red patch), rank-based sparse masking (blue patch) and integration & normalization of recurrent input (green patch). The feedback input and parameters are only required if stacking more than one RSM layer with bi-directional connections.

Cells within the layer are organized into groups. All cells in a group share the same learned weights for external input $\mathbf{x}^F$. Each cell has its own set of fully-connected learned weights for recurrent input $\mathbf{x}^R$. The parameter g specifies the number of groups. Parameter c is the number of cells in each group.

2.0.1. THEORY OF OPERATION

The network described above is similar to a conventional recurrent layer. However, we add an inhibition term ϕ (defined below) that causes cells to become selectively active in particular sequential contexts. The same observation in different contexts is then represented by a different set of active cells.

When a group of cells responds strongly to an input $\mathbf{x}^F$, the least inhibited cell in the group becomes active. Inhibition due to the refractory period ensures that all cells in a group learn unique contexts in which to fire. Fixed sparsity ensures that cells and groups learn to be active in a huge variety of unique combinations. Each combination of cells encodes not just the current input, but a long history of prior inputs as well. As in (Hawkins & Ahmad, 2016), the memory can represent many long histories using the current set of active cells, in effect learning a higher-order representation of these sequences.

2.0.2. MODEL DETAILS

See figure 1 for an overview of the model. We can calculate the matrix product of inputs and weights as normal for a single ANN layer. Note that each cell has a unique set of

recurrent weights w^R, whereas all cells in a group share the same external input weights w^F:

$$\mathbf{z}^F = w^F \mathbf{x}^F(t) \tag{1}$$

$$z^R = w^R \mathbf{x}^R(t) \tag{2}$$

$$z^B = w^B \mathbf{x}^B(t) \tag{3}$$

σ is a matrix of dimension (groups $\times$ cells) i.e. $g \times c$. The weighted sum σ_{ij} of a cell j in group i is given by:

$$\sigma_{ij} = \mathbf{z}_i^F + z_{ij}^R + z_{ij}^B \tag{4}$$

We shift the weighted sums to positive nonzero values and calculate cell activity π after applying an inhibition term ϕ:

$$\pi_{ij} = (1 - \phi_{ij}(t)) \cdot (\sigma_{ij} - \min(\sigma) + 1) \tag{5}$$

We reduce cell activity π to group activity π^G by taking the max value of the cells in each group:

$$\pi_i^G = \max(\pi_{i1}, \dots, \pi_{in}) \tag{6}$$

The next step is to calculate two sparse binary masks M. M^C indicates the most active cell in each group. M^G indicates the most active groups in the layer. We use a function $\text{topk}(a, b)$ that returns a '1' for the top b elements in the last dimension of argument a, and '0' otherwise.

$$M^C = \text{topk}(\pi, 1) \tag{7}$$

$$M^G = \text{topk}(\pi^G, k) \tag{8}$$

We take the elementwise product of the weighted sum and the two masks to yield the weighted sum of one cell from each of the top-k groups. This is the sparsening step. A nonlinearity is applied.

$$\mathbf{y}_{ij} = \tanh(\sigma_{ij} \cdot M_i^G \cdot M_{ij}^C) \tag{9}$$

Cells that are selected are inhibited in future, mimicking the refractory period observed in biological neurons. The refractory period also ensures good utilization of all cells during training. Inhibition decays exponentially, but since we use ranking to select groups, even tiny inhibitions can have a significant effect over long periods. The period of effective inhibition depends on the number of competing

groups and cells available rather than the rate of decay. The hyperparameter $0 \leq \gamma \leq 1$ determines the inhibition decay rate.

$$\phi_{ij}(t+1) = \max(\phi_{ij}(t) \cdot \gamma), \mathbf{y}_{ij}) \qquad (10)$$

Finally, we allow the option to integrate the recurrent input over time. ϵ determines the decay rate of the integrated encoding ψ. Since the current value of ψ theoretically represents all previous states, when external input statistics are identical between training and test time, ϵ can be zero. However, we found a generalization advantage to nonzero ϵ, discussed later.

$$\psi(t) = \max(\psi(t-1) \cdot \epsilon), \mathbf{y}) \qquad (11)$$

$\mathbf{x}^R$ is both the recurrent input and input to any task-specific classifier (see experiments, below). If more than one RSM layer is used, $\mathbf{x}^B_{lower} = \mathbf{x}^R_{higher}$ and $\mathbf{x}^F_{higher} = \mathbf{x}^R_{lower}$. α is a normalizing scalar such that the sum of $\mathbf{x}^B$ is 1 (since we have many zero values, and few or no negative values, we have not found a more appropriate norm). $\mathbf{x}^R$ is initialized as zeros and updated:

$$\mathbf{x}^R(t+1) = \alpha \cdot \psi(t) \qquad (12)$$

The RSM layer is trained to predict the next input $\mathbf{x}^F(t+1)$. Prediction $\hat{\mathbf{x}}^F$ is generated by "decoding" through a bottleneck of masked hidden layer group activities $\mathbf{y}^G$ by taking the max activity of the cells in each group i:

$$\mathbf{y}^G_i = \max(\mathbf{y}_{i1}, \dots, \mathbf{y}_{in}) \qquad (13)$$

w^D is a set of decoding weights of dimension equal to the transpose of w^F. We observed that tied weights $w^F = (w^D)^T$ were less effective in our experiments.

$$\hat{\mathbf{x}}^F(t) = w^D \mathbf{y}^G \qquad (14)$$

The mean square error between prediction $\hat{\mathbf{x}}^F(t)$ and observation $\mathbf{x}^F(t+1)$ is minimized by stochastic gradient descent. Gradients propagate from the prediction through w^D to the sparse bottleneck $\mathbf{y}^G$, and subsequently to the encoding weights w^F, w^R and w^B. The latter two learn to modulate hidden activity given the sequential context. The backpropagation depth is fixed by this architecture to 2. We conceptualize an RSM layer as a single layer of neurons with two or three independently integrated dendrite trees. Max and ranking functions respectively represent local and regional competition between cells.

The architecture is extremely memory efficient, requiring only a single copy of prior memory state $\mathbf{x}^R$ and inhibition ϕ
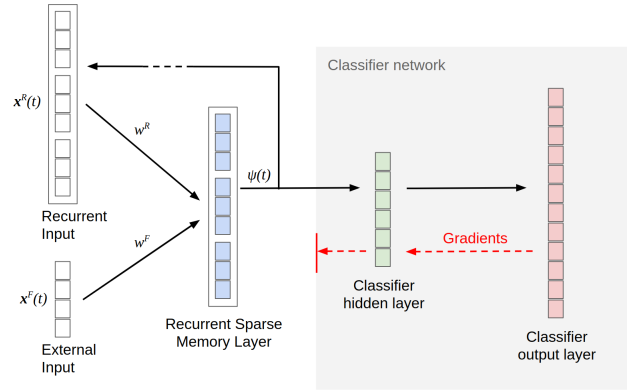

Figure 2: The architecture used in all experiments, except the multi-step prediction task where the classifier network is replaced with a GAN. The classifier has 2 trainable layers, hidden and output. Gradients only propagate within the RSM layer, and within the two classifier layers. The maximum backpropagation depth is 2.

in addition to current external input $\mathbf{x}^F$ for both training and encoding. If c is the number of cells, the asymptotic measure of RSM memory use is $O(c)$. Both truncated BPTT and feed-forward autoregressive approaches such as WaveNet (Oord et al., 2016) require $O(ct)$ where t is the time horizon. With the exception of top-k ranking, all the functions described above are fast, simple arithmetic operations.

2.1. Experimental Architecture

For most of our experiments we use a single RSM layer. To exploit the memory for a discriminatory purpose, we add a "classifier" network of two fully-connected layers (see Figure 2). We use leaky-ReLU nonlinearities in these networks to reduce the accumulation of dead cells due to nonstationary RSM encoding. As per our rules, gradients do not propagate from the classifier into the RSM layer, or across the recurrent or feedback inputs. All layers are trained simultaneously & continuously using only the current input, satisfying the objective of learning using local and immediate credit assignment.

3. Experiments

All experiments use the same architecture described above, with the exception that the classifier is replaced with a Generative Adversarial Network (GAN) in the generative multi-step prediction task.

Published results concerning a sparse memorization approach such as HTM (Cui et al., 2016) focus mostly on network robustness and do not enable easy comparison to more popular ANNs. We will attempt to validate the approach on a range of more easily contextualized tasks & benchmarks.
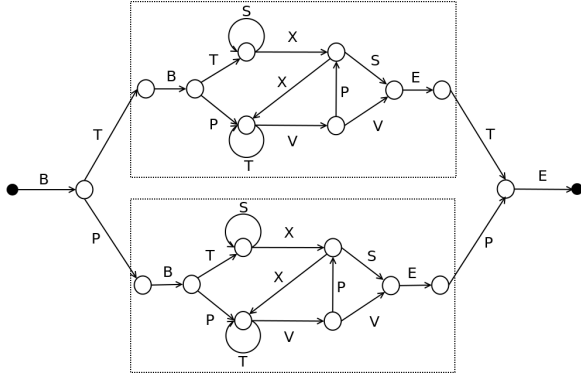
Figure 3: The Embedded Reber Grammar problem. The graph begins with a fork: B-T or B-P. The two fork paths have identical embedded 'Reber' grammars (dashed boxes). After a Reber grammar, T or P must be predicted correctly. The Reber grammar generates long, random, distraction sequences. Successful prediction requires the original fork to be remembered. The Recurrent Sparse Memory enables the classifier to achieve high accuracy.

3.1. Associating distant cause and effect

LSTM was created to solve the problem of vanishing or exploding gradients in RNNs, using gated memory cells. Hochreiter and Schmidhuber demonstrated LSTM using several example problems, including the Embedded Reber Grammar (ERG) (Hochreiter & Schmidhuber, 1997). Since ordinary RNNs were unable to solve the task, it is a good baseline for demonstrating distant cause & effect learning.

The ERG has an initial fork (T,P), followed by the same distraction subsequence on both forks (see figure 3). The final step (T,P) can only be predicted by remembering the pre-fork symbol. For added difficulty, symbols T and P also occur in the distraction sequence.

RSM enables its attached classifier to achieve over 99% accuracy on a random sample from the grammar (see table 1).

The ERG is not deterministic. It generates sequences of minimum length 9. The maximum length is unbounded, but to achieve 99% accuracy, all sequences of length ≤ 30 must be predicted correctly. Without gated memory cells, simple RNNs trained with BPTT and RTRL fail to solve the ERG task (Hochreiter & Schmidhuber, 1997). Cui et al report that HTM achieved 98% accuracy (Cui et al., 2016).

We speculate that RSM achieves high accuracy by memorizing all the sequences frequently generated by the grammar. In a sample of 5000 sequences drawn from the grammar, we observed 601 unique sequences. We used $g = 200$, $c = 6$ and $k = 25$ for this experiment. Since there are only 48 disjoint sets of 25 cells, memorization of 600 or more sequences implies that RSM has learned a combinatoric

Model	Accuracy (%)
Elman RNN (Hochreiter & Schmidhuber, 1997)	0
LSTM (Hochreiter & Schmidhuber, 1997)	100
HTM (Cui et al., 2016)	98.4
RSM	**99.2**

Table 1: Distant cause & effect prediction accuracy on the Embedded Reber Grammar task for various neural models. Hochreiter and Schmidhuber report that a conventional (Elman) RNN fails to solve this task.

model.

RSM prediction is not confounded by stochastic sequences, but note that the statistics of training and test sequences are identical. Both the RSM layer and classifier components of the network achieve the same accuracy, but the classifier allows a specific label to be predicted, rather than an image of a label.

3.2. Modelling partially observable, higher-order sequences

We have demonstrated that RSM is able to model non-deterministic sequences, but what about sequences that are only partially observable? Since an RSM layer is easily derived from a convolutional network layer, and can handle non-deterministic sequences, it is likely RSM is able to simultaneously learn sequential and spatial structure in its input. To test this hypothesis, we presented repeating sequences of MNIST images and measured next-digit appearance and label prediction accuracy. Each step in a sequence has a specific label, but a random image of that label is selected to represent it as input. Due to variation in digit images, the actual sequence is only partially observable. True labels are never observed by the RSM. There are no sequence boundaries and the network is never reset to a baseline state, so all structure must be inferred from the stream of input.

Table 2 shows some of the sequences tested and the minimum history necessary for perfect prediction. Note that some sequences have Markov order > 1, meaning that correct recognition of the current image is insufficient to predict the next label.

As before, both RSM and classifier components acheive accuracy of 99% or better, but the classifier allows labels to be obtained for quantitative measurement. Examination of the next-image predictions is interesting - the RSM generates generic digit-image predictions (see figure 4).

When trained to classify the label of the current image, the

Sequence	Markov order	Accuracy (%)
'0,1,2,3,4,5,6,7,8,9'	1	99.9
'0,1,2,3,4, 0,4,3,2,1'	2	99.9
'0,1,2,3, 0,1,2,3, 0,3,2,1'	5	99.9
None (image classification)	0	**98.2**

Table 2: Partially observable sequence learning results. While the sequences are deterministic, each observation is a random MNIST image of the given class.
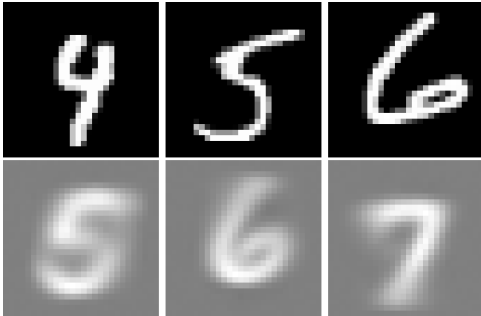

Figure 4: MNIST image sequence prediction experiment. Top row shows current images. Bottom row shows predicted images $\hat{x}^F(t)$ given by RSM, for a sequence 0,1,...,9 repeating. Random exemplar images are selected for each instance. Note that despite the uncertain appearance of each digit, RSM is able to generate a correct "generic" digit prediction. The classifier gives the correct label prediction with $> 99\%$ accuracy.

same fully-connected classifier network achieves a lower accuracy of 98%. We conclude that the RSM is able to combine spatial and higher-order sequential information in a single layer to boost image classification accuracy to over 99% for all sequences tested. In many network architectures incorporating gated memory cells, a convolutional stack is used for spatial dimensionality reduction before consideration of prior state held in memory cells (for example, (Srivastava et al., 2015) uses a pretrained stack to generate "percepts" as input to an LSTM layer). Integrated modelling of spatial and temporal data in a single layer is likely advantageous, and has been investigated (Byeon et al., 2018), but is hampered by the complexity of gated memory layers.

A recent publication by Gordon et al (Gordon et al., 2019) applied a modified version of RSM to a similar task with the additional difficulty of non-deterministic sequences. In this case the modified bRSM was observed to learn the sequence structure better than an LSTM with the same number of parameters.

3.3. Multi-step video prediction

We wished to explore whether RSM can learn temporal dynamics well enough to generate novel sequences, rather than simply next-step predictions. Models trained to perform single-step prediction tend to perform poorly when used for multi-step, self-looped prediction because they lack a persistent representation of the whole production, such as a sentence(Bowman et al., 2015). In addition, we wanted to explore a task with more demanding spatial characteristics rather than a single character centered in the image. Multi-step video prediction is an obvious candidate task and current models are unsatisfactory (Cenzato et al., 2019).

The "Bouncing balls" video prediction task was introduced by Sutskever et al (Sutskever et al., 2009). The objective is to predict the motion of three billiard balls bouncing around in a box. Interactions between balls and the walls of the box are not precisely predictable from the raster image observations provided as input. Researchers evaluate models trained on this dataset in two ways - first, by measuring next-frame prediction mean-square-error (MSE). Second, qualitatively by generating videos of the learned model in self-looped mode. In this condition, a few "priming" frames are produced by the simulation, and passed as input to the learned model. The learned model continuously generates predictions of the next input frame. After priming, the model's own prediction is provided as input. Since tiny angular errors compound rapidly over time, it is expected that sequences generated in this way diverge from the simulation, making an MSE comparison inappropriate (figure 6).

3.3.1. STACKED RSM NETWORKS

The recurrent input $\mathbf{x}^R$ can also be utilized as $\mathbf{x}^F$ for a deeper RSM layer, allowing layers to be stacked. In the stack, credit assignment is still local within each RSM layer (no gradients propagate recurrently or between layers). Stacking up to 3 RSM layers and using the deepest layer as classifier input did not degrade performance in our image sequence experiments, suggesting that RSM layers can be assembled into a hierarchy. In the bouncing balls task, adding a 2nd RSM layer improved prediction accuracy significantly (see table 3). Since a single RSM layer can learn higher-order models, there is little advantage to deeper networks unless it enables more efficient dimensionality reduction, or association of multi-modal inputs.

3.3.2. CONVOLUTIONAL RSM

Given the higher spatial complexity of the bouncing balls task, a convolutional approach is desirable. A convolutional version of RSM can be easily derived from the fully-connected version given above. Learned parameters are shared between all convolutional filter positions, but inhibition ϕ and outputs ψ are unique to each position. The

number of filter positions (and hence, output dimension $[h^1, w^1]$ of a lower layer) is determined by the feed-forward convolution operation. In a convolutional stack architecture, the number of filter positions in the higher layer $[h^2, w^2]$ may be dissimilar to the lower layer $[h^1, w^1]$. A mini-batch of convolution produces a matrix of dimension $[b, h, w, g, c]$. These dimensions are respectively, batch-size, conv. height, conv. width, groups and cells. For the feedback input $\mathbf{x}^B$ from higher layers, we set the receptive field width and height to 1×1 and interpolate from $[h^2, w^2]$ to $[h^1, w^1]$. This ensures that $[h^1, w^1]$ is equal to the number of filter positions required by the feed-forward convolution operation. A 1×1 receptive field is also used for the recurrent input, meaning that each filter position only receives a recurrent input from itself.

3.3.3. GAN RECTIFIER

We observe that RSM generates a prediction covering all possible next input states (e.g. figure 4). Each pixel prediction approximates the expected value of that pixel. When self-looped, the uncertainty in these predictions compounds over time, resulting in degraded input dissimilar to the training input. We noticed that using rudimentary image processing filters to rectify the predictions significantly improved RSM output in self-looped mode. A Generative Adversarial Network (GAN) is well suited to this role, because given a conditioning input it can produce clear, sharp samples that are indistinguishable from genuine input. We trained a GAN to take the lower-layer RSM prediction $\hat{\mathbf{x}}^F$ as conditioning input and output a next-frame sample. Implicitly, the role of the RSM is to predict the distribution of next frame samples, and the role of the GAN is to render one sample from this distribution. The GAN does not have access to any context to allow it to perform a predictive role. Gradients from the GAN do not propagate into the RSM layers.

We used a stacked convolutional autoencoder as the GAN generator, and a fully-connected network as GAN discriminator (see appendix for details). The GAN was trained with both MSE and adversarial losses, copying the Predictive Generative Network (PGN) of Lotter et al (Lotter et al., 2015). However, the PGN architecture combines CNN and LSTM architectures to add the predictive aspect. In our case we rely entirely on the RSM layers for prediction.

3.3.4. VIDEO PREDICTION RESULTS

Table 3 shows comparative performance of RSM and works by other authors. For reference, we include results for RSM without GAN and a single layer RSM + GAN. A two layer RSM + GAN system outperforms all prior works, including a variety of LSTM architectures, in terms of next-frame prediction accuracy. To better understand the generated spatio-temporal dynamics in self-looped mode, we recom-
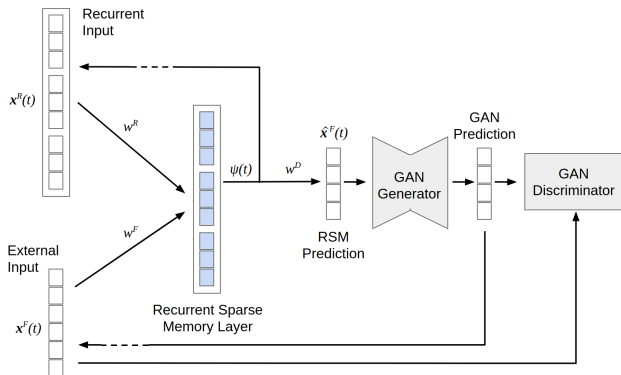

Figure 5: RSM combined with a Generative Adversarial Network (GAN). Although two RSM layers were used, only the lower layer is connected to the GAN and shown here. The upper RSM layer influences the lower RSM layer via the feedback dendrite. When used in self-looping (generative) mode, RSM predictions approximate the expected value of each pixel, resulting in a slightly blurred appearance that covers all possible ball trajectories (there is some uncertainty in the simulation due to the precise timing of ball and wall interactions). The GAN has no predictive ability, but acts as a rectifier, sharpening RSM predictions into a single specific possible future outcome. Passing the GAN-rectified prediction back into RSM produces good simulations many steps into the future.
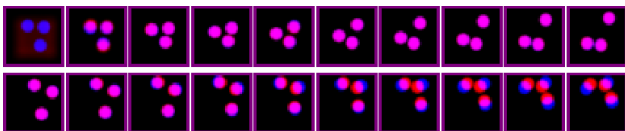

Figure 6: Frame by frame RSM predictions of bouncing ball dynamics (red channel) compared to ground-truth simulation (blue channel). The top row shows RSM in predictive mode from the start of a sequence. Prediction errors are imperceptible after the 5th step. Bottom row: RSM in generative mode - although the dynamics are believable, the trajectories do diverge from the simulation.

mend viewing a video of the output [1]. Although videos of self-looped output are not available for all works, the results are markedly improved over Sutskever et al (RTRBM / SRTRBM) (Sutskever et al., 2009). Whereas the balls wander, change direction or acceleration and stick to each other in the latter, the RSM videos show relatively good conservation of momentum and plausible interactions between objects.

---

[1] For a qualitative understanding of the generated ball dynamics, see video at: `https://www.youtube.com/watch?v=-lBFW1gbokg`

Model	Error
Previous frame (Lotter et al., 2015)	11.82 ± 0.27
RTRBM (Sutskever et al., 2009; Mittelman et al., 2014)	3.88 ± 0.33
SRTRBM (Mittelman et al., 2014)	3.31 ± 0.33
DTSBN (Gan et al., 2015)	2.79 ± 0.39
PGN (MSE) (Lotter et al., 2015)	0.65 ± 0.11
LSTM (Cenzato et al., 2019)	111.09 ± 0.68
ConvLSTM (Cenzato et al., 2019)	0.58 ± 0.16
Seq2seq ConvLSTM (Cenzato et al., 2019)	1.34 ± 0.19
Seq2seq ConvLSTM multi-decoder (Cenzato et al., 2019)	4.55 ± 0.40
Conv RSM 2L	3.72 ± 0.29
Conv RSM 1L + GAN	0.8 ± 0.07
Conv RSM 2L + GAN	$\mathbf{0.41 \pm 0.05}$

Table 3: Next-frame prediction error for the bouncing-balls task. RSM is compared to previous results from the [Structured] Recurrent Temporal Restricted Boltzmann Machine ([S]RTRBM) (Sutskever et al., 2009; Mittelman et al., 2014), Deep Temporal Sigmoid Belief Network (DTSBN) (Gan et al., 2015), Predictive Generative Network (PGN) (Lotter et al., 2015) and a variety of LSTM architectures (Cenzato et al., 2019).

3.4. Language modelling: Next word prediction perplexity

3.4.1. GENERALIZATION

RSM takes the approach of remembering and encoding a long history of prior states in the current state. To achieve this, the same input in different sequential contexts generates highly orthogonal encodings. This is in contrast to gated memory layers, where a few features are *selectively* remembered, and different sequential contexts can generate *similar* encodings. Consequently, we anticipate that changes in sequential context between training and test datasets will severely disrupt RSM. In other words, we don't expect it to generalize well to unseen sequences.

To explore this aspect of generalization, we selected a language modelling task. The Penn TreeBank (PTB) dataset is a corpus of approximately 1 million training words and 80,000 test words. We used Mikolov's preprocessing (Mikolov et al., 2011), which results in a dictionary of 10,000 unique words. PTB is known to present a difficult generalization challenge due to the small size of the corpus. Next-word prediction quality is typically measured using perplexity (PPL). Model output is a distribution over the

words in the dictionary. In our case, model input is a binary vector representing the (random) indices of each word in the dictionary.

3.4.2. GENERALIZATION FEATURES

For RSM, even small changes between training and test sequences can be disruptive. Imagine that the training corpus contains 'The cat sat on the mat' and the test corpus 'The cat sat on *a* mat'. Although the RSM encoding would be identical until 'on', subsequent states would be highly orthogonal, so the classifier cannot predict 'mat' from combinations of active cells never previously observed. To mitigate these types of generalization error, we enabled some minor features.

First, we integrate feedback to the RSM using $\epsilon > 0$. This is not needed to help RSM learn sequences, but it does help RSM *generalize* learned sequences. We apply dropout at rate 0.5 to encourage the RSM to make use of earlier encodings, for example the active cells representing 'The cat sat' are all predictive of a future 'mat'. Second, we provide the integrated RSM state to the classifier. These changes reduce test perplexity by 10-15.

The third change is to "forget" the recurrent state of the network with a fixed probability μ during training. If forgetting is randomly selected, we set integrated activity ψ and inhibition ϕ to zeros. The intention is to expose the classifier to different subsequences, in effect augmenting the training set. This feature improves test perplexity by another 5-10 points.

3.4.3. REGULARIZATION

We attempted to regularize the RSM and classifier networks via conventional techniques - in particular adding an L2 loss term, and dropout. We found neither helpful when added to the RSM, but a small L2 value improved classifier performance.

To understand why these techniques are ineffective with this architecture, consider Cui et al (Cui et al., 2016), which was one of the inspirations for our work. They show that sparse distributed representations are extremely robust to network damage, because the meaning of individual elements is correlated. Therefore, to successfully "ignore" an observation, many cells' activity must be simultaneously excluded. This is difficult for reasonable levels of both dropout and weight penalties. So it seems ineffective regularization may be a consequence of distributed representations that are robust to damage.

3.4.4. MODEL INTERPOLATION

Mikolov et al (Mikolov et al., 2011) showed that linear interpolation of model distributions to produce ensemble

predictions is very often effective at reducing test perplexity. In particular, earlier RNN and LSTM language models (which we are using as a baseline) were often presented as ensembles. 5-gram language models with Kneser-Ney smoothing (KN5) (Chen & Goodman, 1999) are optimal for the PTB corpus; higher n-gram models do not deliver better test perplexity.

When a word occurs, it is likely to occur again in the near future. This general principle leads to cache language models. Mikolov et al(Mikolov et al., 2011) also showed that caches invariably improve ensemble perplexity. We felt that adding a fixed, exponentially-decaying probability mass to words after observation - a primitive cache - meets our biological plausibility criteria. We did not allow 'adaptive' modelling (i.e. allowing the RSM to learn during the test). We report RSM results interpolated with the simplistic cache, and optionally with KN5 for easier comparison and insights on n-gram ensemble complementarity.

3.4.5. RESULTS

With $g = 600$ groups, $k = 20$, $c = 6$ cells per group and 4 epochs of training with batch size 300, we obtained 50% next word prediction accuracy and a perplexity of 9 on the training corpus. Training perplexity and accuracy had not plateaued. This suggests that moderately sized RSM layers can memorize extremely long sequences.

We found test corpus perplexity to be better with $n = 8$ cells per group and only 0.25 epochs of training. We observed next word prediction accuracy of 20.6% on the test corpus. RSM per-word perplexity is bi-modal, showing rapid oscillations between very low and high values. When not combined with KN5, a small uniform mass significantly reduces average RSM test perplexity (reported result without KN5 includes 7% uniform mass). The complementarity of RSM and KN5 - evidenced by RSM+KN5 interpolated model perplexity - suggests that RSM sometimes captures longer range context, but also fails to generalize many shorter range contexts.

Table 4 shows PTB test perplexity for RSM and selected popular language models. RSM is comparable to some early methods such as 5-gram Good-Turing, Random Clusterings, Random Forest and 5-gram Maximum Entropy, both individually and interpolated with KN5. The natural point of comparison is conventional RNNs trained with short-context BPTT. RNNs trained with BPTT are significantly better than RSM, with PPL 124 (RNN) and 105 (RNN+KN5).

Carefully regularized, gated-memory ANN language models such as LSTM reduce PTB test perplexity from 124 to 100 or less (Zaremba et al., 2014), with state of the art scores < 60 (Merity et al., 2017). Models with attentional filtering are even better (Radford et al., 2019). However, these

Model	Perplexity
Random clusterings LM (Mikolov et al., 2011)	170
RSM	166
5-gram Good-Turing (Mikolov et al., 2011)	162
Structured LM (Mikolov et al., 2011)	146
KN5 (our impl.)	143
Max. entropy 5-gram (Mikolov et al., 2011)	142
Feed-forward ANN LM (deep BP) (Mikolov et al., 2011)	140
Random forest LM (Mikolov et al., 2011)	131
RNN LM w. BPTT (Mikolov et al., 2011)	124
RSM+KN5	124
bRSM (Gordon et al., 2019)	103
LSTM variants (Merity et al., 2017)	$58 \leq 86$
Transformer (GPT-2), varying size (Radford et al., 2019)	$35 \leq 65$

Table 4: Test-set perplexity of RSM and bRSM models, KN5, and both combined. Selected other algorithms are shown for context.

improvements required several years' effort by the research community, and exploit selective remembering features that could also be added to RSM over time.

Gordon et al (Gordon et al., 2019) recently published a modified version of RSM (bRSM) that replaces inhibition with a boosting scheme. The resultant PTB perplexity of 103 is significantly better than RNNs trained with BPTT (PPL 124) (Mikolov et al., 2011). This suggests that bRSM overcomes most or all of the generalization issues demonstrated by RSM.

4. Summary

Our primary objective was to develop a strictly local, biologically-plausible credit assignment strategy with comparable practical performance to early deep-backpropagation results. We believe we have been successful in this regard. On several tasks RSM provides results comparable to neural methods trained by deep BP or BPTT, which should not be possible for a simple RNN. For example, we showed that RSM can successfully learn to exploit distant causes when predicting later effects in the ERG task, which an RNN without BPTT cannot do, and in fact requires gated memory such as LSTM (Hochreiter & Schmidhuber, 1997). RSM can learn higher-order sequences, likely with

combinatoric representational efficiency, despite some degree of sequential uncertainty and partial observability. On the "bouncing balls" task, RSM produced new best results beating a range of algorithms trained by deep BP, including LSTM. The learned model is also convincing in self-looped mode, generating arbitrarily long sequences of quite realistic billiard-ball dynamics.

RSM has a number of advantages - an order of magnitude reduction in memory requirements, very large memory capacity, batch-online training, and it is not necessary to specify the maximum time horizon in advance. For complex sequence learning tasks with similar training and test statistics, the RSM approach is appealing due to its simplicity and speed. However, although stochastic processes are not confounding, generalization to *unseen* test sequences with different statistics is clearly limited in its current form, as seen in the PTB task. A later development of RSM by Gordon et al (Gordon et al., 2019) was able to overcome this limitation, giving better results on PTB than RNNs trained by BPTT.

In the image sequence prediction task, RSM was continually exposed to unseen sequences of observations due to the large number of images of each digit. Why was this form of generalization not pathological? First, the memorization approach is unlikely to confound spatial generalization; it is only unseen sequence generalization that is limited. Second, it is not possible to memorize the infinite observed training sequences, and the training regime therefore forces RSM to learn a generalized sequence model. Since both training and test set are drawn from the same distributions, this form of generalization is likely only interpolation rather than extrapolation.

4.1. Selective remembering

One of the most significant differences between RSM and gated memory layers is the strategy of "remember everything" versus "selectively remember useful things". The latter approach is able to generalize better, because it is not distracted by irrelevant features in unseen test data.

Recent work has shown that self-attention is a very effective mechanism for selecting useful features from sequences (Vaswani et al., 2017), with groundbreaking results in neural language models. Radford et al (Radford et al., 2019) uses self-attention in an unsupervised next-word-prediction training regime, like RSM, but autoregressive feed-forward rather than recurrent. In future work we will attempt to use self-attention to selectively remember (sub) sequences using RSM, without violating our biological plausibility constraints.

References

Bai, Shaojie, Kolter, J Zico, and Koltun, Vladlen. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018.

Balduzzi, David, Vanchinathan, Hastagiri, and Buhmann, Joachim. Kickback cuts backprop's red-tape: biologically plausible credit assignment in neural networks. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

Bengio, Yoshua, Lee, Dong-Hyun, Bornschein, Jorg, Mesnard, Thomas, and Lin, Zhouhan. Towards biologically plausible deep learning. *arXiv preprint arXiv:1502.04156*, 2015.

Bowman, Samuel R, Vilnis, Luke, Vinyals, Oriol, Dai, Andrew M, Jozefowicz, Rafal, and Bengio, Samy. Generating sentences from a continuous space. *arXiv preprint arXiv:1511.06349*, 2015.

Byeon, Wonmin, Wang, Qin, Kumar Srivastava, Rupesh, and Koumoutsakos, Petros. Contextvp: Fully context-aware video prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.

Cenzato, Alberto, Testolin, Alberto, and Zorzi, Marco. On the difficulty of learning and predicting the long-term dynamics of bouncing objects. *arXiv preprint arXiv:1907.13494*, 2019.

Chelba, Ciprian, Mikolov, Tomas, Schuster, Mike, Ge, Qi, Brants, Thorsten, Koehn, Phillipp, and Robinson, Tony. One billion word benchmark for measuring progress in statistical language modeling. *arXiv preprint arXiv:1312.3005*, 2013.

Chen, Stanley F and Goodman, Joshua. An empirical study of smoothing techniques for language modeling. *Computer Speech & Language*, 13(4):359–394, 1999.

Chung, Junyoung, Gulcehre, Caglar, Cho, KyungHyun, and Bengio, Yoshua. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

Cui, Yuwei, Ahmad, Subutai, and Hawkins, Jeff. Continuous Online Sequence Learning with an Unsupervised Neural Network Model. *Neural computation*, 28:2474–2504, 2016.

Gan, Zhe, Li, Chunyuan, Henao, Ricardo, Carlson, David E, and Carin, Lawrence. Deep temporal sigmoid belief networks for sequence modeling. In *Advances in Neural Information Processing Systems*, 2015.

Gordon, Jeremy, Rawlinson, David, and Ahmad, Subutai. Long distance relationships without time travel: Boosting the performance of a sparse predictive autoencoder in sequence modeling. *arXiv preprint arXiv:1912.01116*, 2019.

Guerguiev, Jordan, Lillicrap, Timothy P, and Richards, Blake A. Towards deep learning with segregated dendrites. *ELife*, 6:e22901, 2017.

Hawkins, Jeff and Ahmad, Subutai. Why neurons have thousands of synapses, a theory of sequence memory in neocortex. *Frontiers in neural circuits*, 10:23, 2016.

Hochreiter, Sepp and Schmidhuber, Jürgen. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

Khemlani, Sangeet Suresh, Mackiewicz, Robert, Bucciarelli, Monica, and Johnson-Laird, Philip N. Kinematic mental simulations in abduction and deduction. *Proceedings of the National Academy of Sciences*, 110(42):16766–16771, 2013.

Lotter, William, Kreiman, Gabriel, and Cox, David. Unsupervised learning of visual structure using predictive generative networks. *arXiv preprint arXiv:1511.06380*, 2015.

Luo, Hongyin, Fu, Jie, and Glass, James. Adaptive bidirectional backpropagation: Towards biologically plausible error signal transmission in neural networks. *arXiv preprint arXiv:1702.07097*, 2017.

Makhzani, Alireza and Frey, Brendan. K-sparse autoencoders. *arXiv preprint arXiv:1312.5663*, 2013.

Makhzani, Alireza and Frey, Brendan J. Winner-Take-All Autoencoders. In *Advances in Neural Information Processing Systems*, 2015.

Merity, Stephen, Keskar, Nitish Shirish, and Socher, Richard. Regularizing and optimizing lstm language models. *arXiv preprint arXiv:1708.02182*, 2017.

Mikolov, Tomáš, Deoras, Anoop, Kombrink, Stefan, Burget, Lukáš, and Černockỳ, Jan. Empirical evaluation and combination of advanced language modeling techniques. In *Twelfth Annual Conference of the International Speech Communication Association*, 2011.

Mittelman, Roni, Kuipers, Benjamin, Savarese, Silvio, and Lee, Honglak. Structured recurrent temporal restricted boltzmann machines. In *International Conference on Machine Learning*, 2014.

Ney, Hermann, Essen, Ute, and Kneser, Reinhard. On structuring probabilistic dependences in stochastic language modelling. *Computer Speech & Language*, 8(1):1–38, 1994.

Olshausen, Bruno A and Field, David J. Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.

Oord, Aaron van den, Dieleman, Sander, Zen, Heiga, Simonyan, Karen, Vinyals, Oriol, Graves, Alex, Kalchbrenner, Nal, Senior, Andrew, and Kavukcuoglu, Koray. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.

O'Reilly, Randall C. Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural computation*, 8(5): 895–938, 1996.

Radford, Alec, Wu, Jeffrey, Child, Rewon, Luan, David, Amodei, Dario, and Sutskever, Ilya. Language models are unsupervised multitask learners. *OpenAI Blog*, 1:8, 2019.

Srivastava, Nitish, Mansimov, Elman, and Salakhutdinov, Ruslan. Unsupervised learning of video representations using lstms. *arXiv preprint arXiv:1502.04681*, 2015.

Sutskever, Ilya. *Training recurrent neural networks*. University of Toronto Toronto, Ontario, Canada, 2013.

Sutskever, Ilya, Hinton, Geoffrey E, and Taylor, Graham W. The recurrent temporal restricted boltzmann machine. In *Advances in neural information processing systems*, 2009.

Vaswani, Ashish, Shazeer, Noam, Parmar, Niki, Uszkoreit, Jakob, Jones, Llion, Gomez, Aidan N., Kaiser, Lukasz, and Polosukhin, Illia. Attention Is All You Need. (Nips), 2017.

Williams, Ronald J and Zipser, David. Gradient-based learning algorithms for recurrent. *Backpropagation: Theory, architectures, and applications*, 433, 1995.

Zaremba, Wojciech, Sutskever, Ilya, and Vinyals, Oriol. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*, 2014.

A. Appendices

A.1. Bouncing Balls video

We have uploaded a video of the generated dynamics from the RSM 2-layer + GAN algorithm compared to (Sutskever et al., 2009). Unfortunately, there is no available video of the generative dynamics from the PGN method of (Lotter et al., 2015). In the video, our algorithm is first

primed with 50 frames and then switched to self-looping mode for 150 frames, feeding the GAN-rectified prediction back into the RSM. The video is located at: `https://www.youtube.com/watch?v=-lBFW1gbokg`

A.2. Hyperparameters

The following values were used in our experiments. A learning rate of 0.0005 was used with the Adam optimizer in all experiments.

Hyperparameter	Value
Batch size	400
g (groups)	200
c (cells per group)	6
k (sparsity)	25
γ (inhibition decay rate)	0.98
ϵ (recurrent input decay rate)	0
Classifier hidden layer size	500

Table 5: Distant cause & effect (ERG)

Hyperparameter	Value
Batch size	300
g (groups)	200
c (cells per group)	6
k (sparsity)	25
γ (inhibition decay rate)	0.5
ϵ (recurrent input decay rate)	0
Classifier hidden layer size	1200

Table 6: Partially observable sequences (MNIST images)

Hyperparameter	Value
Batch size	300
g (groups)	600
c (cells per group)	8
k (sparsity)	20
γ (inhibition decay rate)	0.8
ϵ (recurrent input decay rate)	0.85
Classifier hidden layer size	1200

Table 7: Language modelling (PTB)

Hyperparameter	Value
Batch size	256
ConvRSM layer 1	
g (groups)	64
c (cells per group)	8
k (sparsity)	3
γ (inhibition decay rate)	0
ϵ (recurrent input decay rate)	0
receptive field	5x5
pool size	2
strides	2
ConvRSM layer 2	
g (groups)	128
c (cells per group)	8
k (sparsity)	5
γ (inhibition decay rate)	0
ϵ (recurrent input decay rate)	0
receptive field	3x3
strides	2
GAN Generator [convolutional AE]	
Enc. receptive field	5x5
Enc. filters	64, 128, 256
Enc. strides	1,2,1
Enc. nonlinearity	Leaky ReLU
Dec. receptive field	5x5
Dec. filters	128, 64, 1
Dec. strides	2, 1, 1
Dec. hidden nonlinearity	Leaky ReLU
Dec. output nonlinearity	Sigmoid
GAN Discriminator [fully connected]	
Hidden layer size	128
Hidden layer nonlinearity	Leaky ReLU
Output layer nonlinearity	Sigmoid

Table 8: Generative modelling (Bouncing Balls)