

Arman Afrasiyabi

Université Laval

Arman.afrasiyabi.1@ulaval.ca

ARTICLE

doi:10.1038/nature20101

Hybrid computing using a neural network with dynamic external memory

Alex Graves^{1*}, Greg Wayne^{1*}, Malcolm Reynolds¹, Tim Harley¹, Ivo Danihelka¹, Agnieszka Grabska-Barwińska¹, Sergio Gómez Colmenarejo¹, Edward Grefenstette¹, Tiago Ramalho¹, John Agapiou¹, Adrià Puigdomènech Badia¹, Karl Moritz Hermann¹, Yori Zwols¹, Georg Ostrovski¹, Adam Cain¹, Helen King², Christopher Summerfield¹, Phil Blunsom¹, Koray Kavukcuoglu¹ & Demis Hassabis¹

Artificial neural networks are remarkably adept at sensory processing, sequence learning and reinforcement learning, but are limited in their ability to represent variables and data structures and to store data over long timescales, owing to the lack of an external memory. Here we introduce a machine learning model called a differentiable neural computer (DNC), which consists of a neural network that can read from and write to an external memory matrix, analogous to the random-access memory in a conventional computer. Like a conventional computer, it can use its memory to represent and manipulate complex data structures, but, like a neural network, it can learn to do so from data. When trained with supervised learning, we demonstrate that a DNC can successfully answer synthetic questions designed to emulate reasoning and inference problems in natural language. We show that it can learn tasks such as finding the shortest path between specified points and inferring the missing links in randomly generated graphs, and then generalize these tasks to specific graphs such as transport networks and family trees. When trained with reinforcement learning, a DNC can complete a moving blocks puzzle in which changing goals are specified by sequences of symbols. Taken together, our results demonstrate that DNCs have the capacity to solve complex, structured tasks that are inaccessible to neural networks without external read–write memory.

Modern computers separate computation and memory. Computation is performed by a processor, which can use an addressable memory to bring operands in and out of play. This confers two important benefits: the use of extensible storage to write new information and the ability to treat the contents of memory as variables. Variables are critical to algorithm generality: to perform the same procedure on one datum or another, an algorithm merely has to change the address it reads from. In contrast to computers, the computational and memory resources of artificial neural networks are mixed together in the network weights and neuron activity. This is a major liability: as the memory demands of a task increase, these networks cannot allocate new storage dynamically, nor easily learn algorithms that act independently of the values realized by the task variables.

Although recent breakthroughs demonstrate that neural networks are remarkably adept at sensory processing¹, sequence learning^{2,3} and reinforcement learning⁴, cognitive scientists and neuroscientists have argued that neural networks are limited in their ability to represent variables and data structures^{5–9}, and to store data over long timescales without interference^{10,11}. We aim to combine the advantages of neural and computational processing by providing a neural network with read–write access to external memory. The access is narrowly focused, minimizing interference among memoranda and enabling long-term storage^{12,13}. The whole system is differentiable, and can therefore be trained end-to-end with gradient descent, allowing the network to learn how to operate and organize the memory in a goal-directed manner.

System overview

A DNC is a neural network coupled to an external memory matrix. (The behaviour of the network is independent of the memory size as long as the memory is not filled to capacity, which is why we view the memory as ‘external.’) If the memory can be thought of as the DNC’s

RAM, then the network, referred to as the ‘controller’, is a differentiable CPU whose operations are learned with gradient descent. The DNC architecture differs from recent neural memory frameworks^{14,15} in that the memory can be selectively written to as well as read, allowing iterative modification of memory content. An earlier form of DNC, the neural Turing machine¹⁶, had a similar structure, but more limited memory access methods (see Methods for further discussion).

Whereas conventional computers use unique addresses to access memory contents, a DNC uses differentiable attention mechanisms^{2,16–18} to define distributions over the N rows, or ‘locations’, in the $N \times W$ memory matrix M . These distributions, which we call weightings, represent the degree to which each location is involved in a read or write operation. The read vector r returned by a read weighting w^r over memory M is a weighted sum over the memory locations: $r = \sum_{j=1}^W M[i, j] w^r[j]$, where the ‘ \cdot ’ denotes all $j = 1, \dots, W$. Similarly, the write operation uses a write weighting w^w to first erase with an erase vector e , then add a write vector v : $M[i, j] \leftarrow M[i, j] (1 - w^w[j] e[j]) + w^w[j] v[j]$. The functional units that determine and apply the weightings are called read and write heads. The operation of the heads is illustrated in Fig. 1 and summarized below; see Methods for a formal description.

Interaction between the heads and the memory

The heads use three distinct forms of differentiable attention. The first is content lookup^{16,17,19–21}, in which a key vector emitted by the controller is compared to the content of each location in memory according to a similarity measure (here, cosine similarity). The similarity scores determine a weighting that can be used by the read heads for associative recall¹⁹ or by the write head to modify an existing vector in memory. Importantly, a key that only partially matches the content of a memory location can still be used to attend strongly to that location.

¹Google DeepMind, 5 New Street Square, London EC4A 3TW, UK.

*These authors contributed equally to this work.

Differentiable Neural Computer (DNC; 2016)

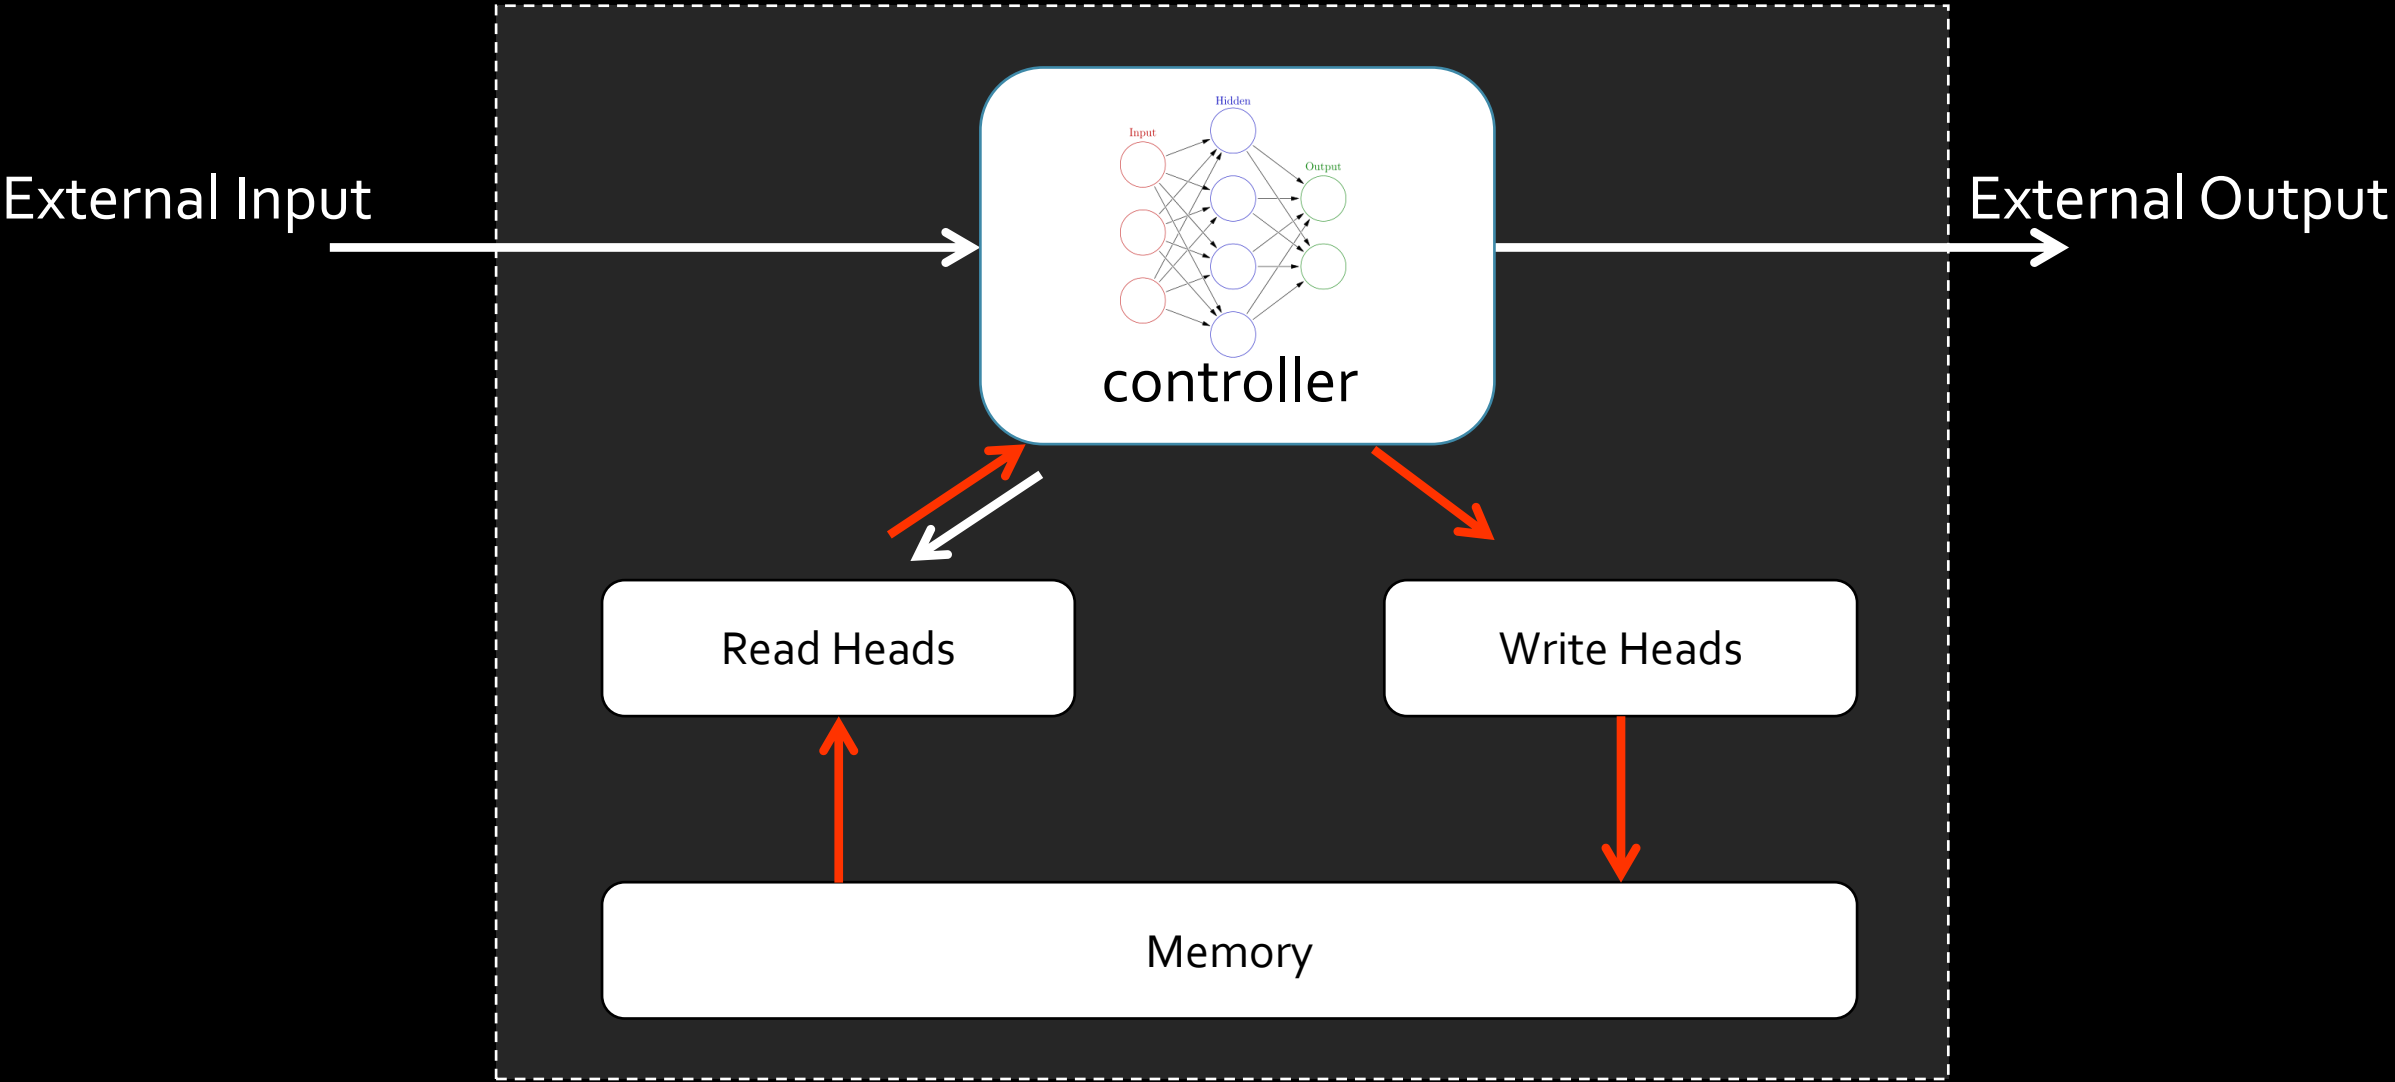
NTM 2.0

Neural Turing Machine (NTM; 2014)

NTM 1.0

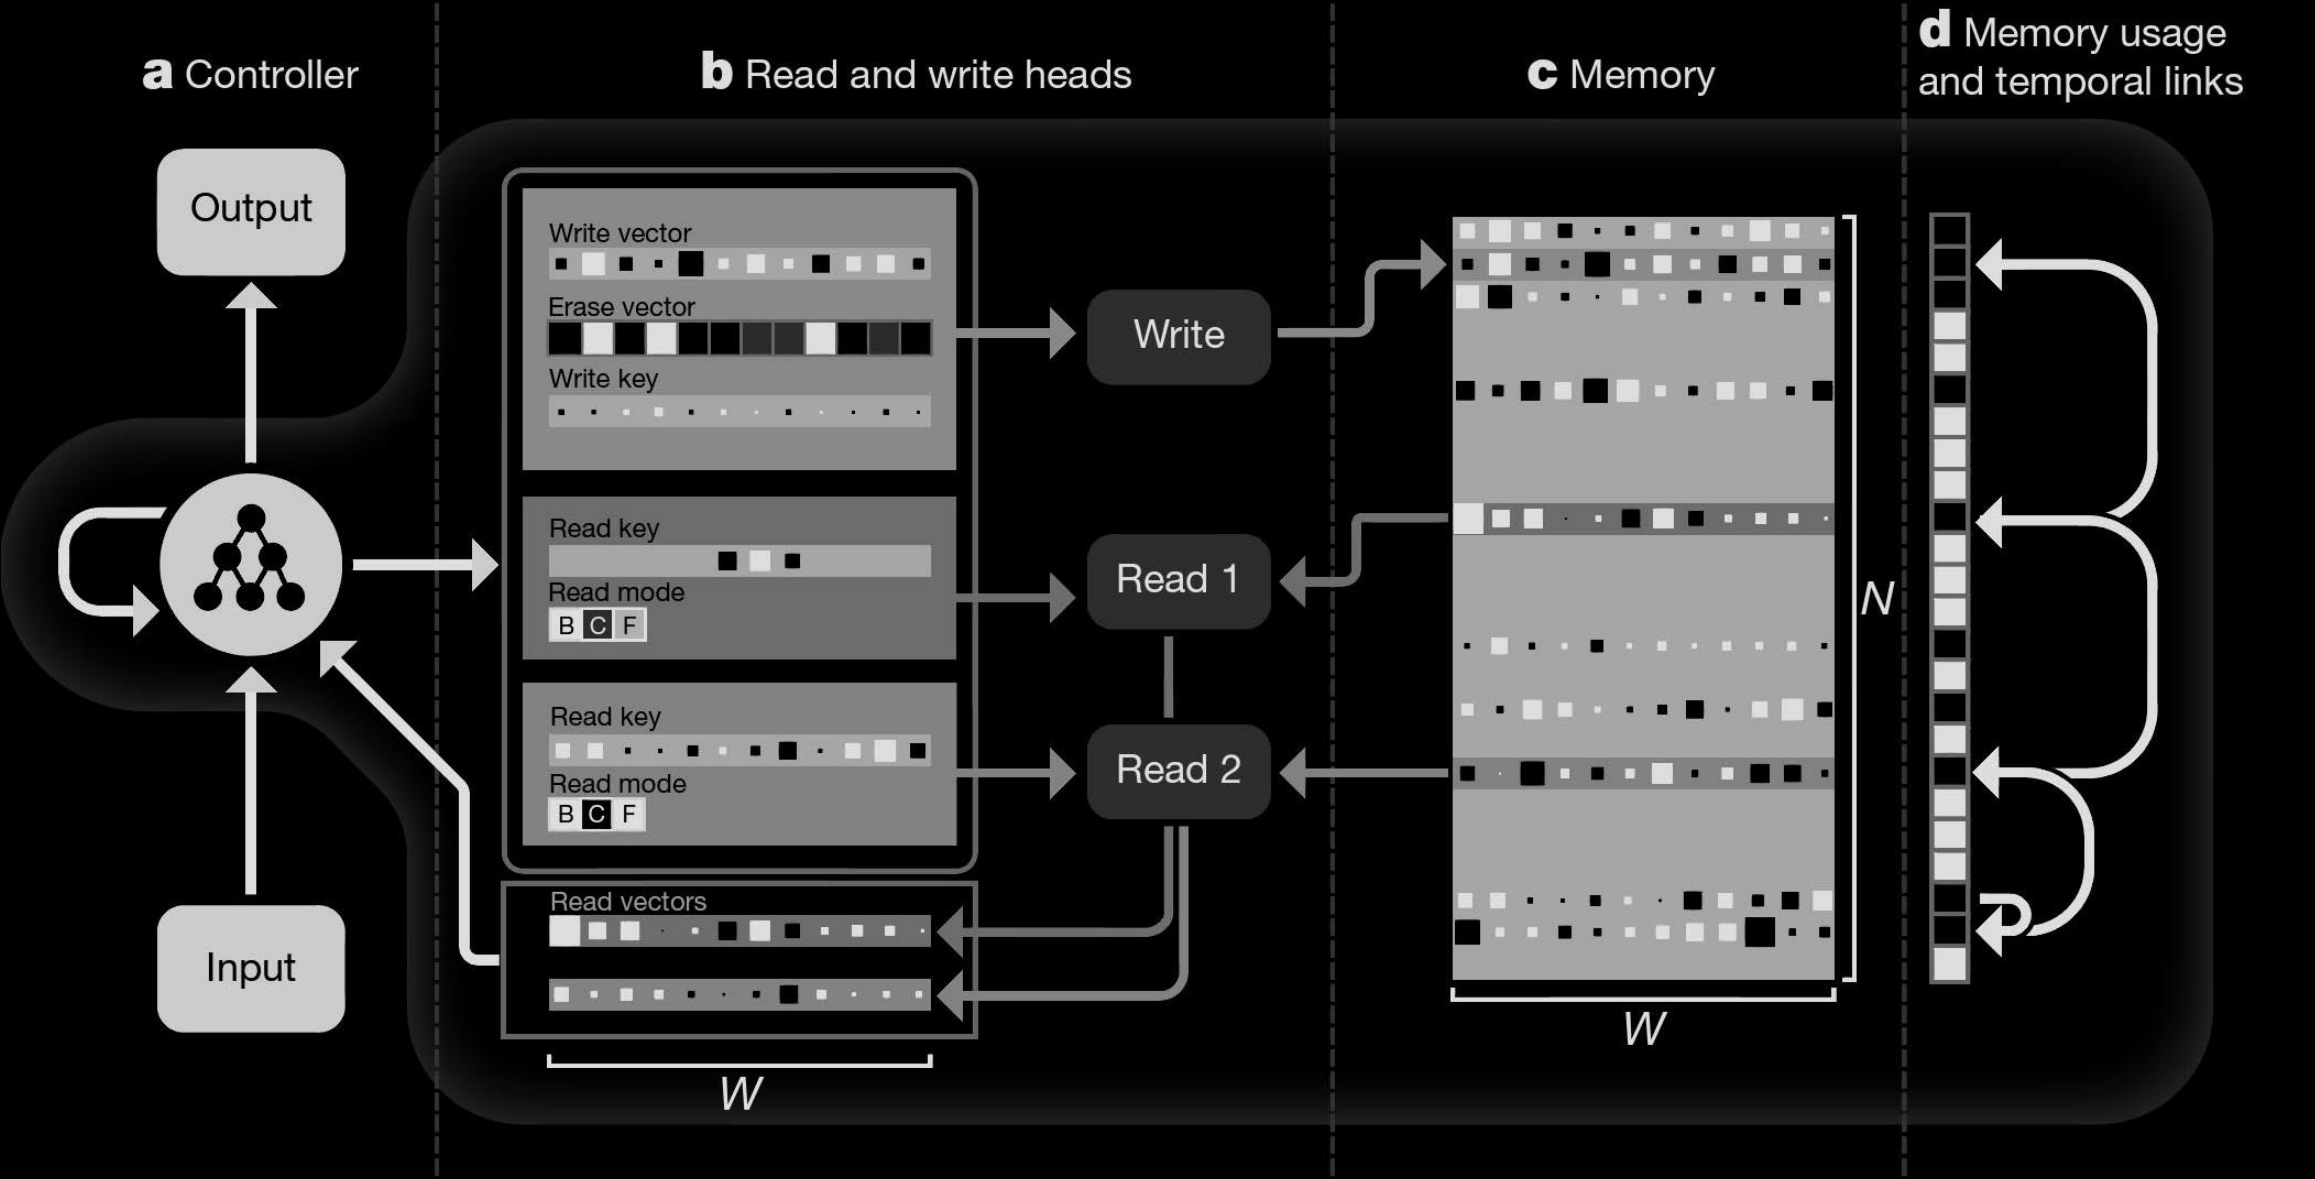
[1] Graves, Alex, et al. "Hybrid computing using a neural network with dynamic external memory." *Nature* 538.7626 (2016): 471-476.

[2] Graves, Alex, Greg Wayne, and Ivo Danihelka. "Neural turing machines." *arXiv preprint arXiv:1410.5401* (2014).

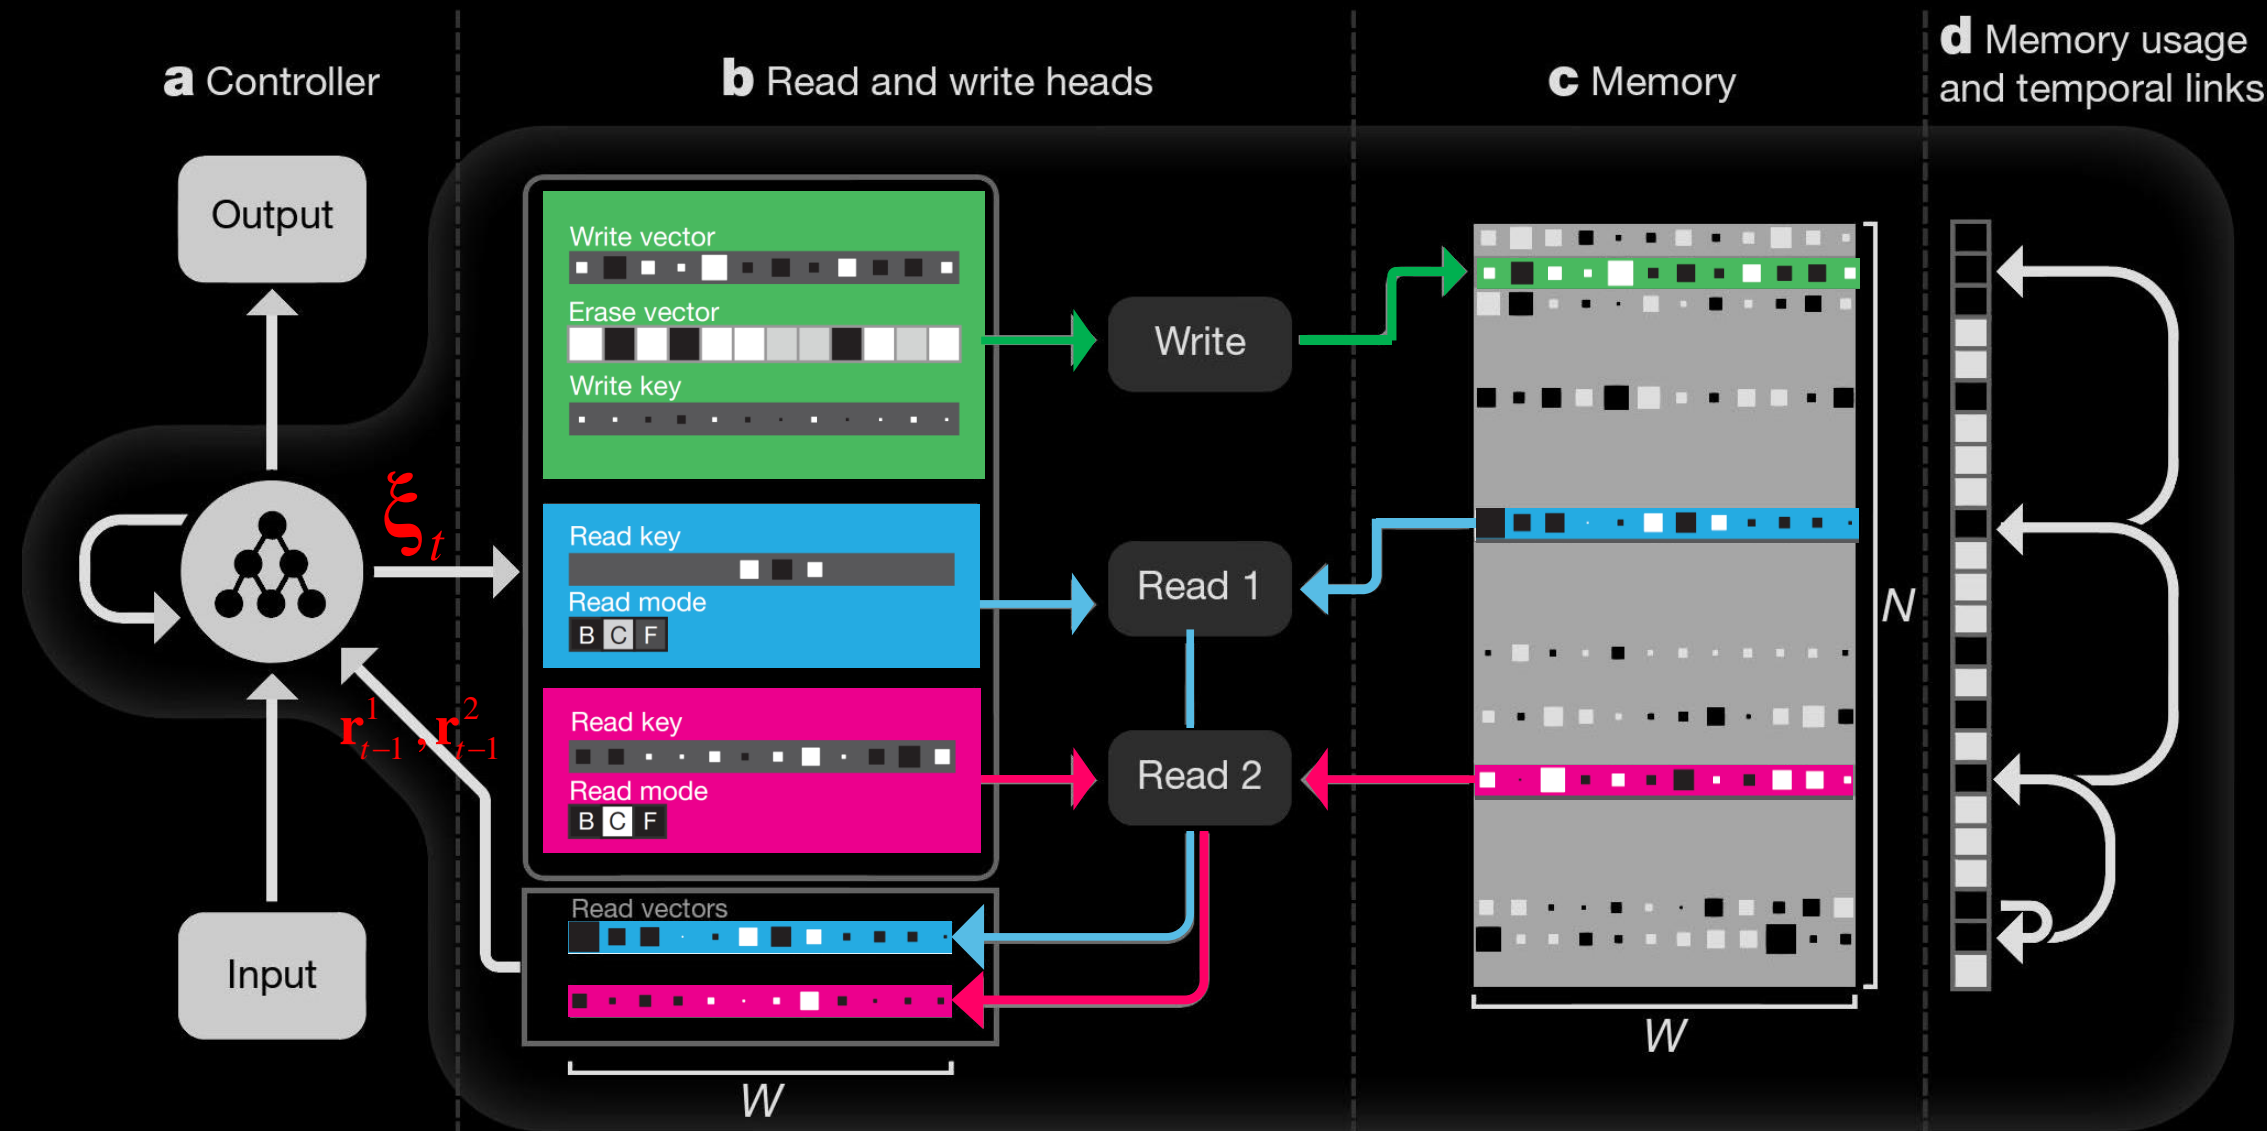


All parts are differentiable!

Architecture- Differentiable Neural Computing (DNC)



Architecture- Differentiable Neural Computing (DNC)



$$\xi_t = \left[\mathbf{k}_t^{\mathbf{r},1}, \dots, \mathbf{k}_t^{\mathbf{r},R}, \hat{\beta}_t^{\mathbf{r},1}, \dots, \hat{\beta}_t^{\mathbf{r},R}, \mathbf{k}_t^{\mathbf{w}}, \hat{\beta}_t^{\mathbf{w}}, \hat{\mathbf{e}}_t, \mathbf{v}_t, \hat{\mathbf{f}}_t^1, \dots, \hat{\mathbf{f}}_t^R, \hat{\mathbf{g}}_t^a, \hat{\mathbf{g}}_t^w, \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$

The controller can be any neural network architecture (in this case a deep LSTM)

Input:

$$\mathbf{X}_t = [x_t; \mathbf{r}_{t-1}^1; \dots; \mathbf{r}_{t-1}^R]$$

Controller:

$$(\mathbf{v}_t, \xi_t) = \mathcal{N}([X_1; \dots, X_t]; \theta)$$

Output:

$$y_t = v_t + \mathbf{W}_r [r_t^1, \dots, r_t^R]$$

In each layer, $l \in \{1, \dots, L\}$ we compute the following functions:

$$i_t^l = \sigma(\mathbf{W}_i^l[\mathbf{X}_t; \mathbf{h}_{t-1}^l; \mathbf{h}_t^{l-1}] + \mathbf{b}_i^l)$$

$$f_t^l = \sigma(\mathbf{W}_f^l[\mathbf{X}_t; \mathbf{h}_{t-1}^l; \mathbf{h}_t^{l-1}] + \mathbf{b}_f^l)$$

$$s_t^l = f_t^l s_{t-1}^l + i_t^l \tanh(\mathbf{W}_s^l[\mathbf{X}_t; \mathbf{h}_{t-1}^l; \mathbf{h}_t^{l-1}] + \mathbf{b}_s^l)$$

$$o_t^l = \sigma(\mathbf{W}_o^l[\mathbf{X}_t; \mathbf{h}_{t-1}^l; \mathbf{h}_t^{l-1}] + \mathbf{b}_o^l)$$

$$\mathbf{h}_t^l = o_t^l \tanh(s_t^l)$$

output vector, inference vector

$$\mathbf{v}_t = \mathbf{W}_y [\mathbf{h}_t^1; \dots; \mathbf{h}_t^L]$$

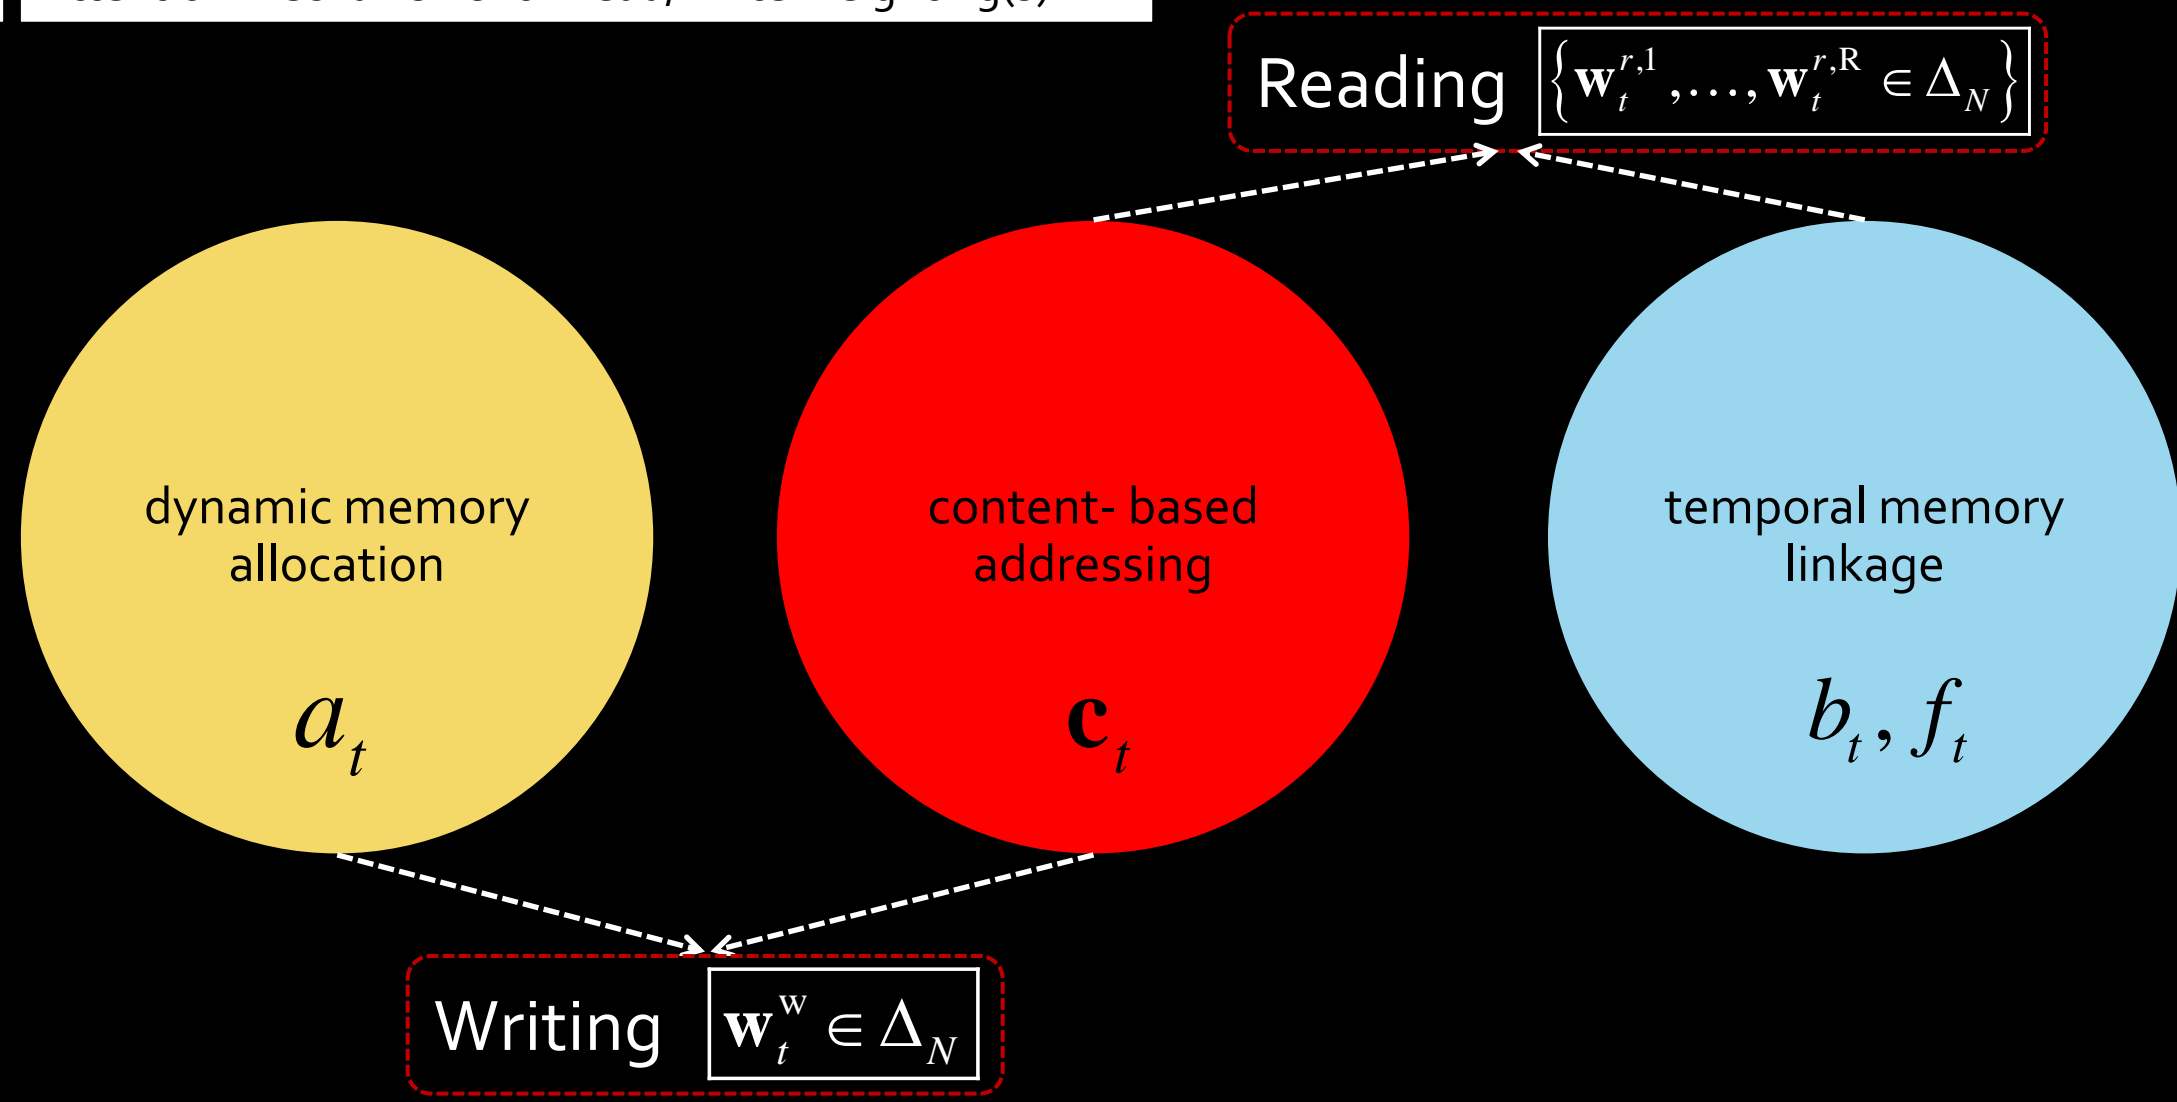
$$\xi_t = \mathbf{W}_\xi [\mathbf{h}_t^1; \dots; \mathbf{h}_t^L]$$

- The **interface vector** contains parameters
- Controller uses interface vector to connect to memory for **Reading/Writing**

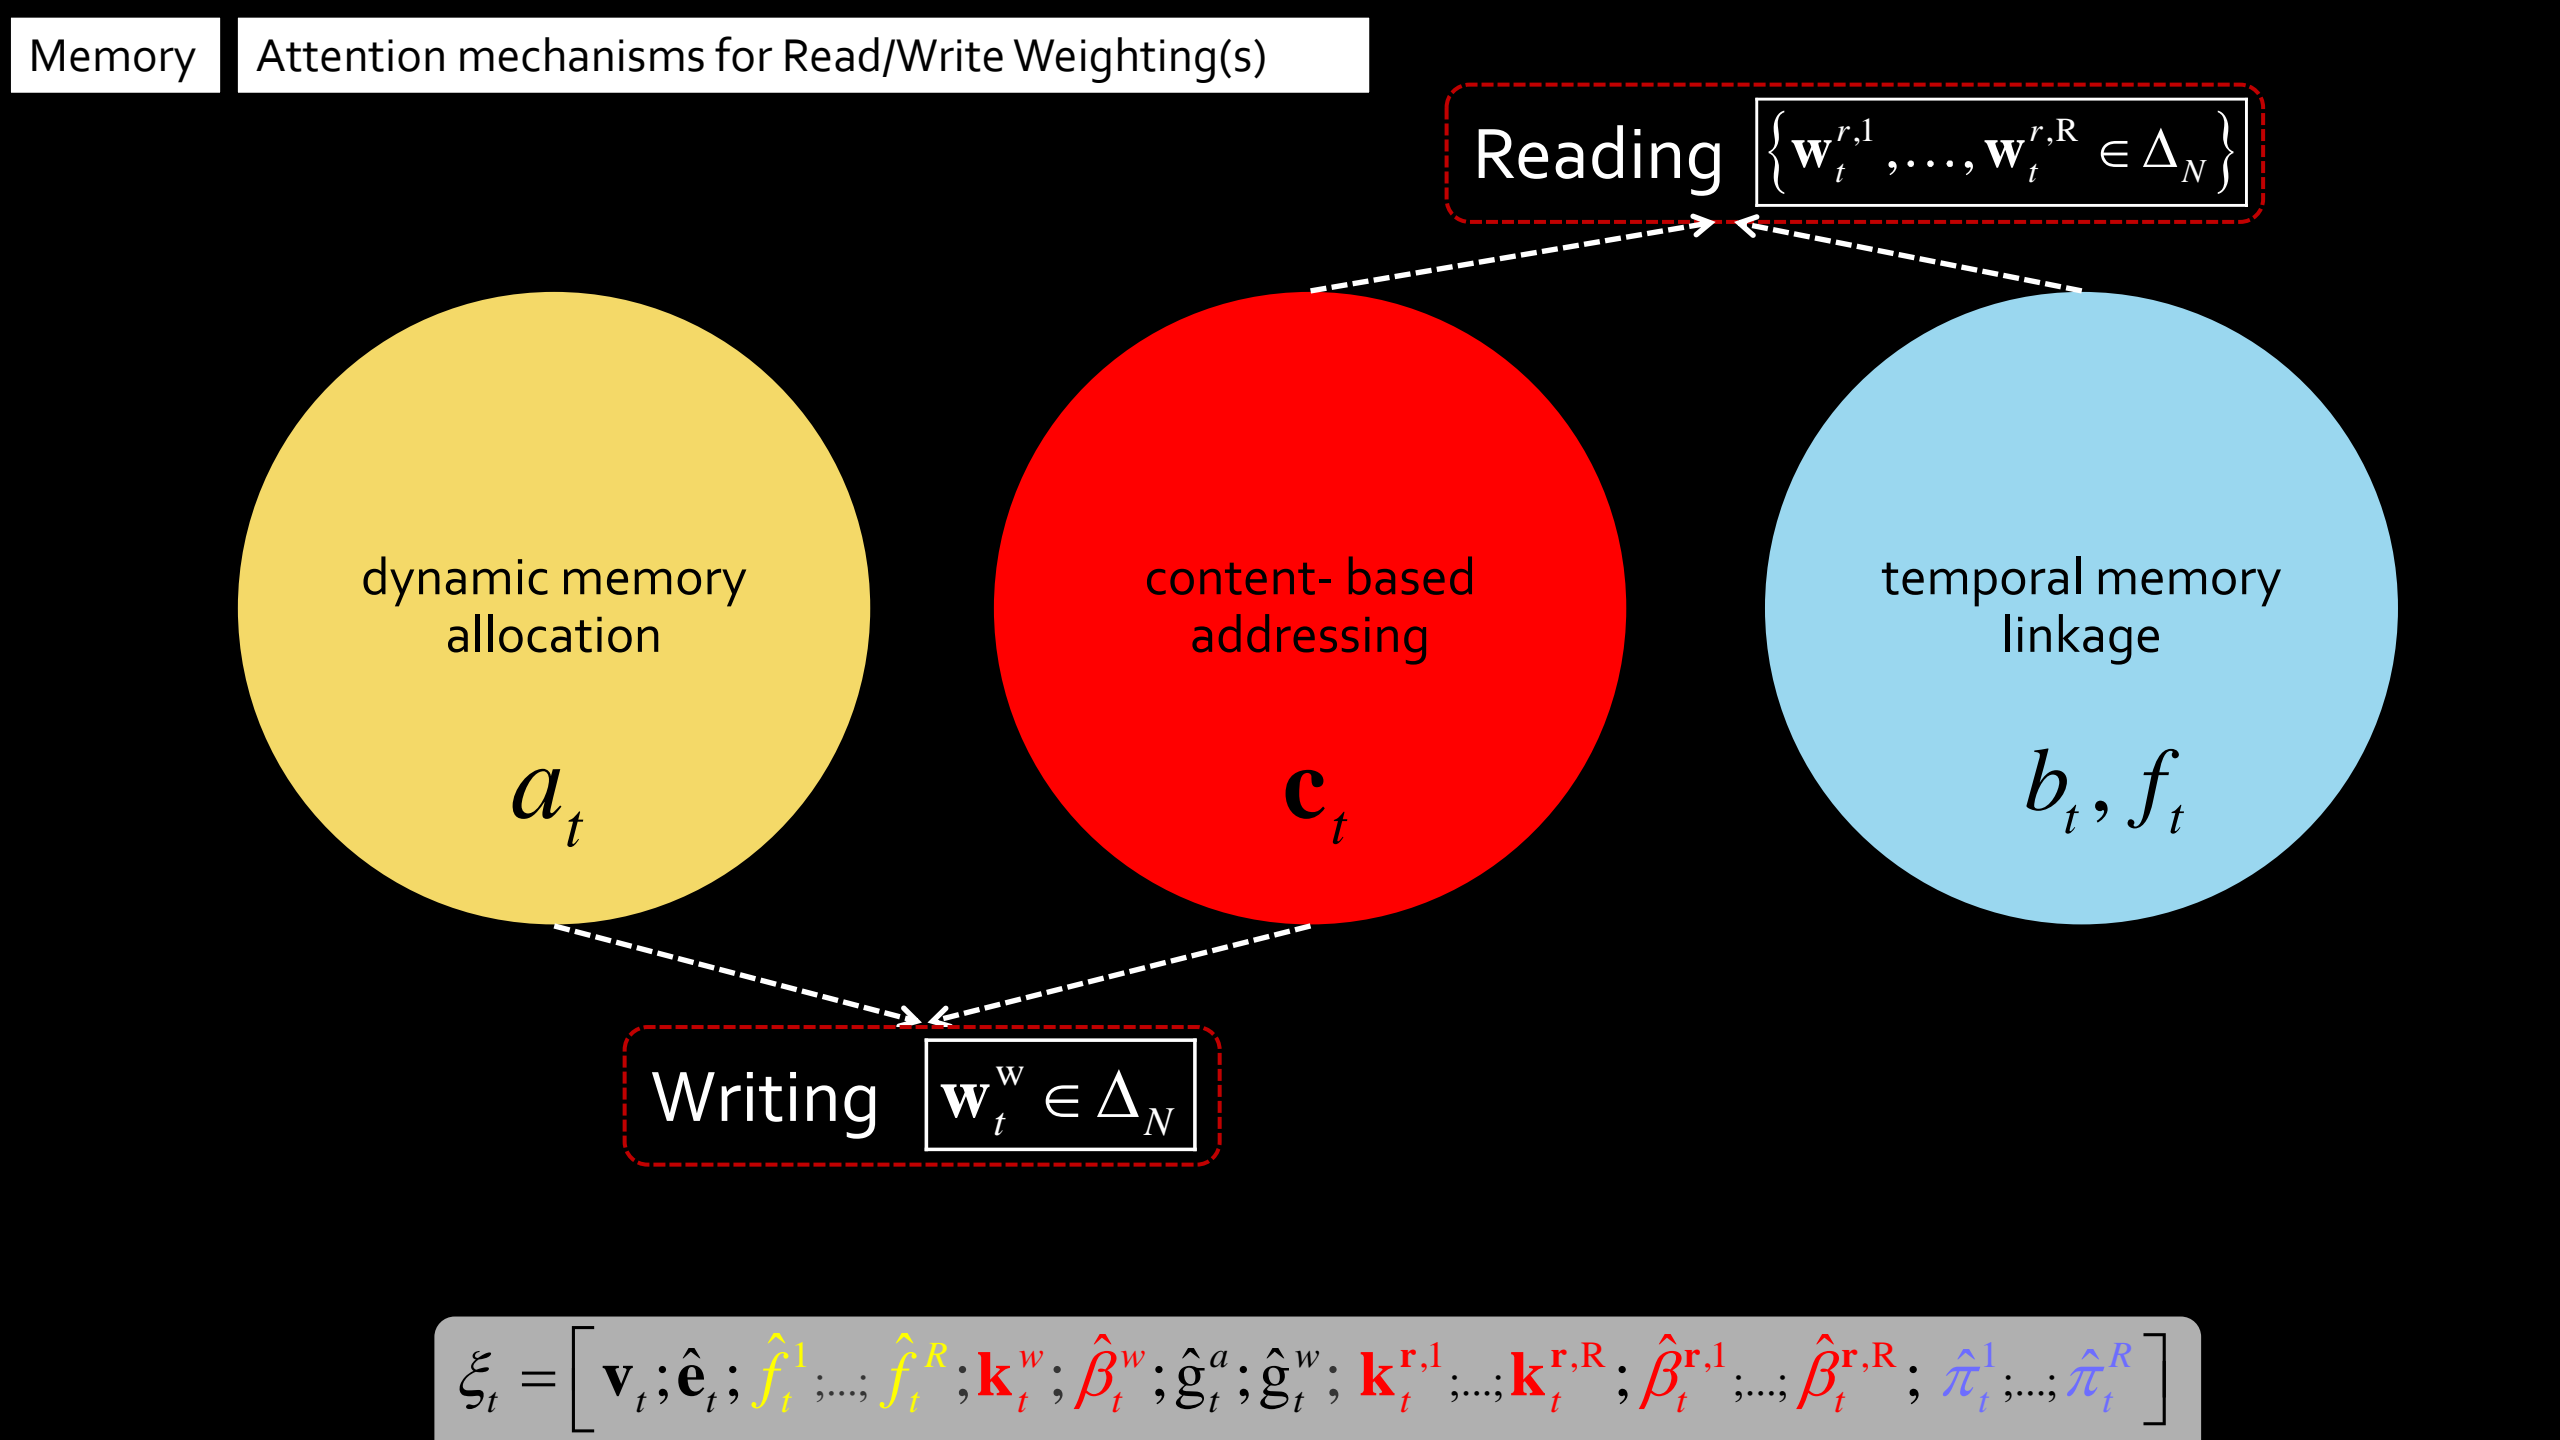
$$\xi_t = \left[\mathbf{k}_t^{\mathbf{r},1}, \dots, \mathbf{k}_t^{\mathbf{r},R}; \hat{\beta}_t^{\mathbf{r},1}, \dots, \hat{\beta}_t^{\mathbf{r},R}; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{e}}_t; \mathbf{v}_t; \hat{\mathbf{f}}_t^1, \dots, \hat{\mathbf{f}}_t^R; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$

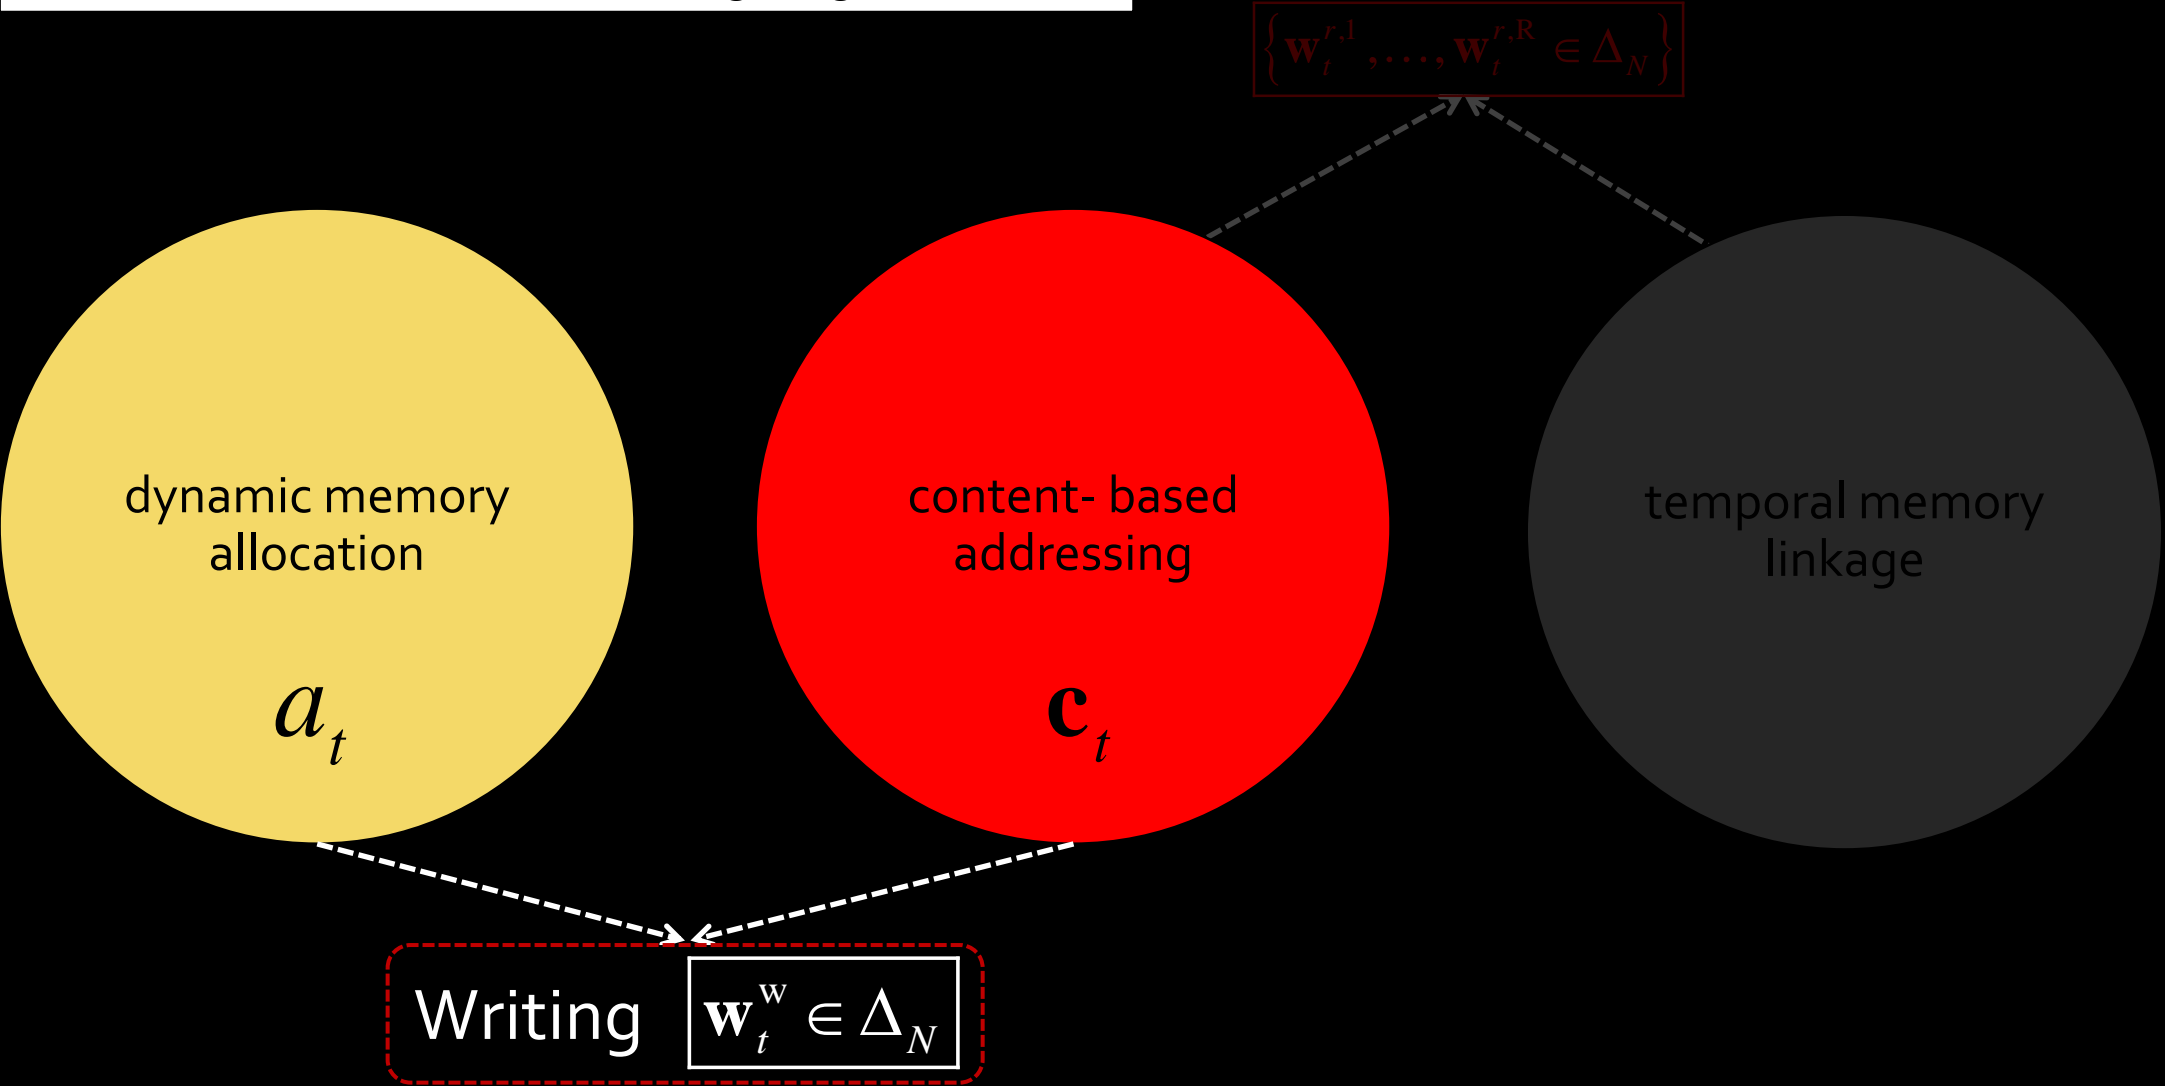
location selection for **reading/writing** depends on **weighting(s)**

$$\Delta_N = \left\{ \boldsymbol{\alpha} \in \mathbb{R}^N : \alpha_i \in [0,1], \sum_{i=1}^N \alpha_i \leq 1 \right\}$$



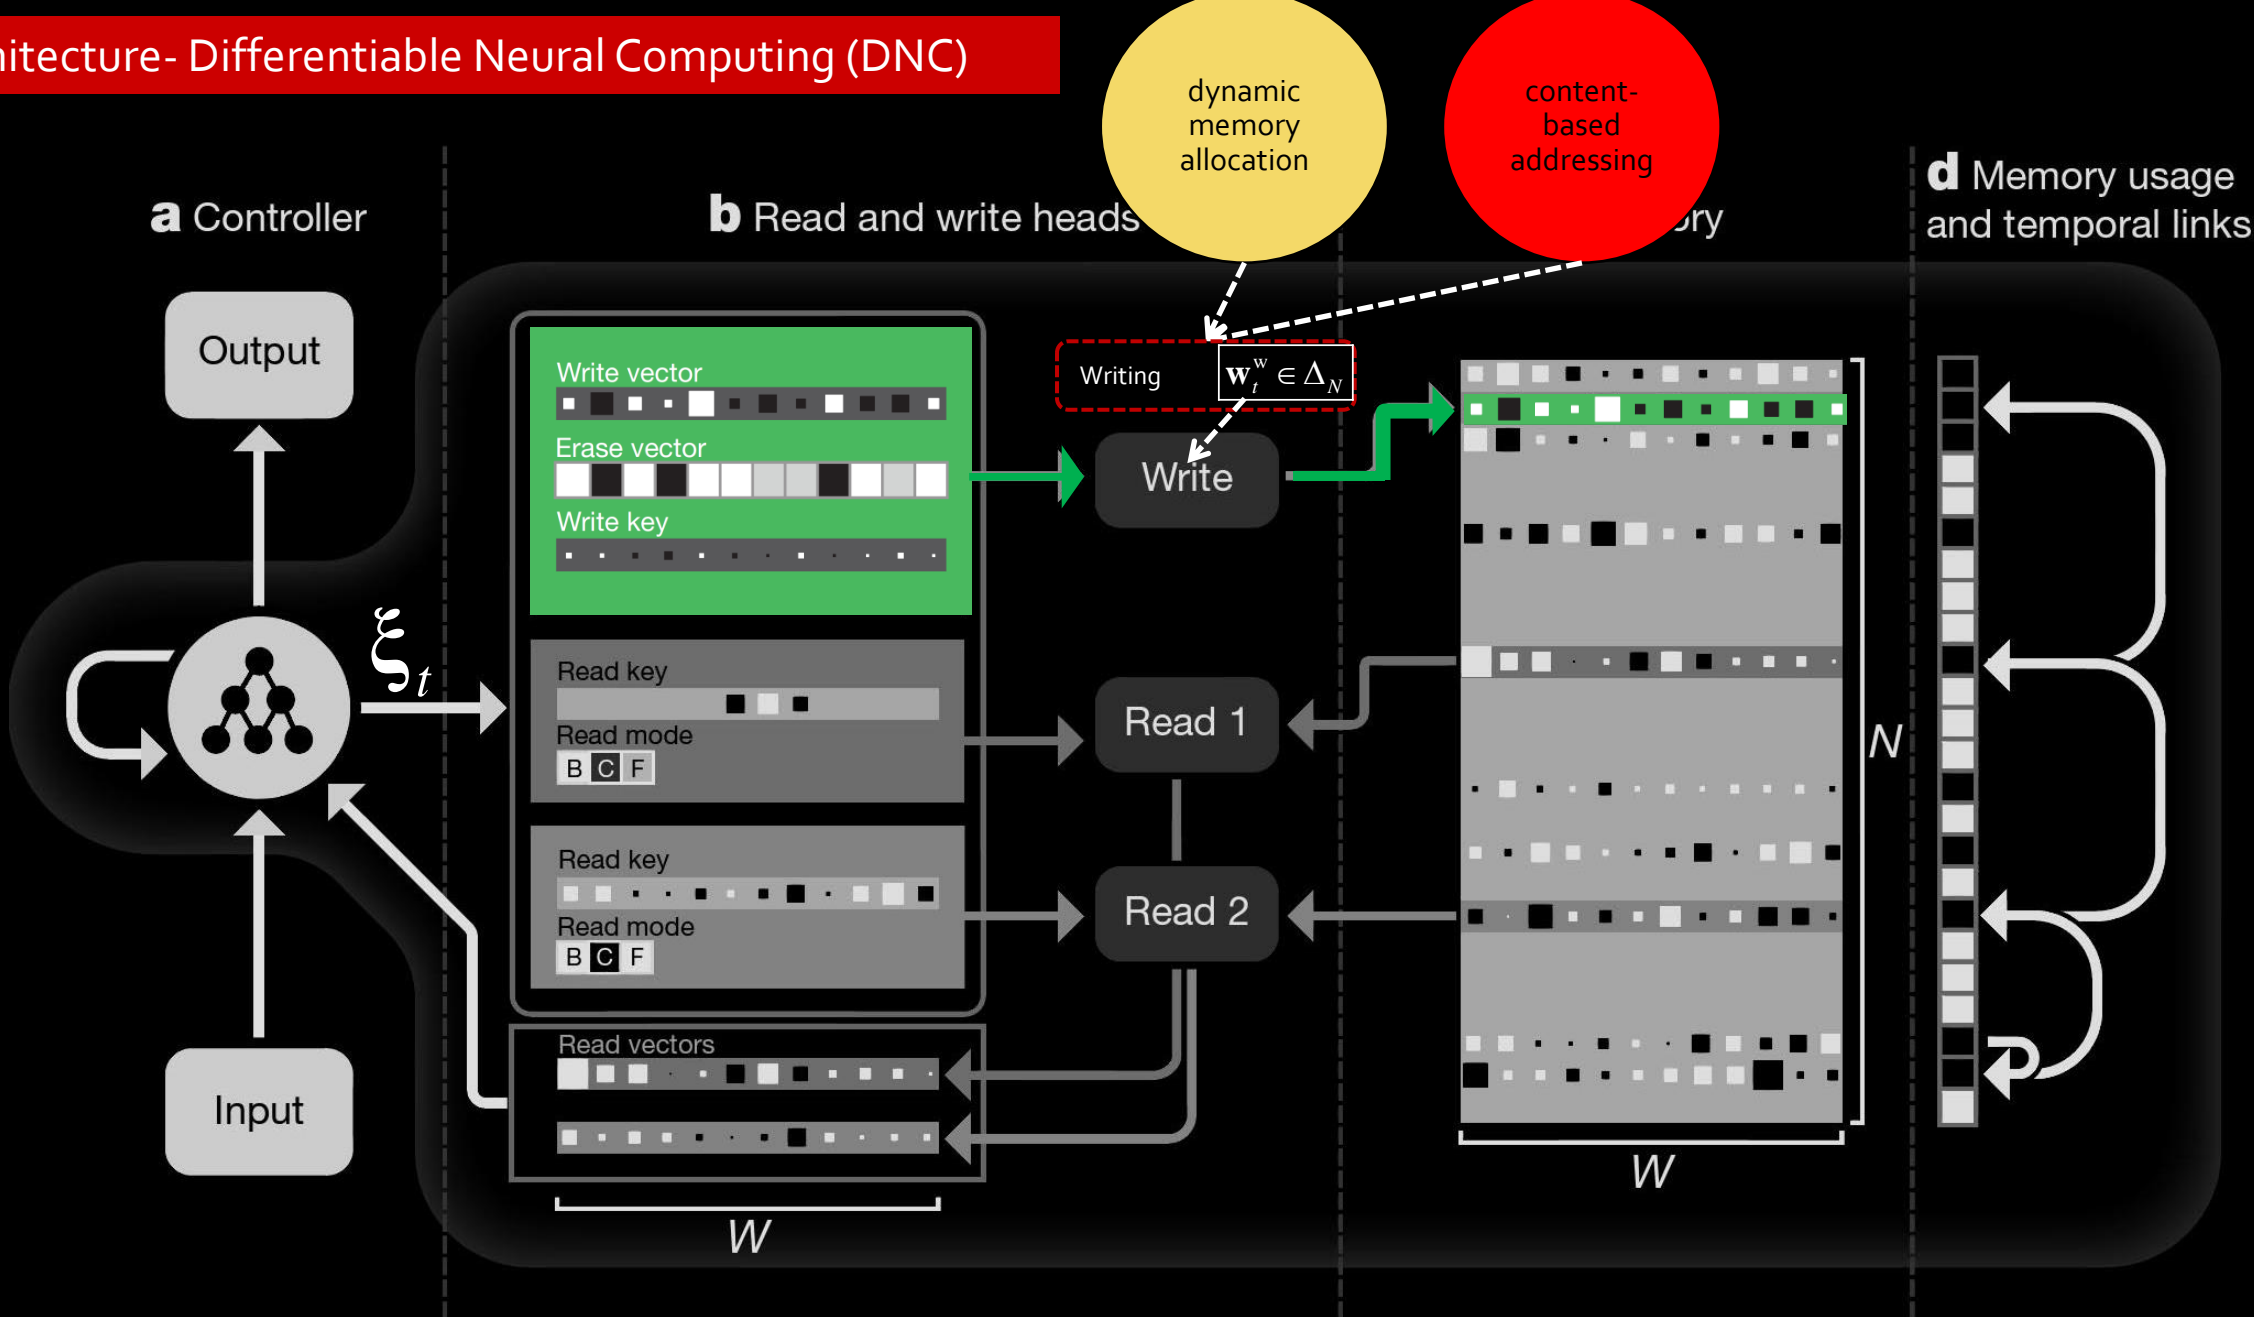
The controller interpolates among these mechanisms using **scalar gates**





$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1; \dots; \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,l}; \dots; \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,l}; \dots; \hat{\beta}_t^{r,R}; \hat{\pi}_t^l; \dots; \hat{\pi}_t^R \right]$$

Architecture- Differentiable Neural Computing (DNC)



$$E = \text{ones}(N \times W)$$

$$\Delta_N = \left\{ \mathbf{a} \in \mathbb{R}^N : \mathbf{a}_i \in [0,1], \sum_{i=1}^N \mathbf{a}_i \leq 1 \right\}$$

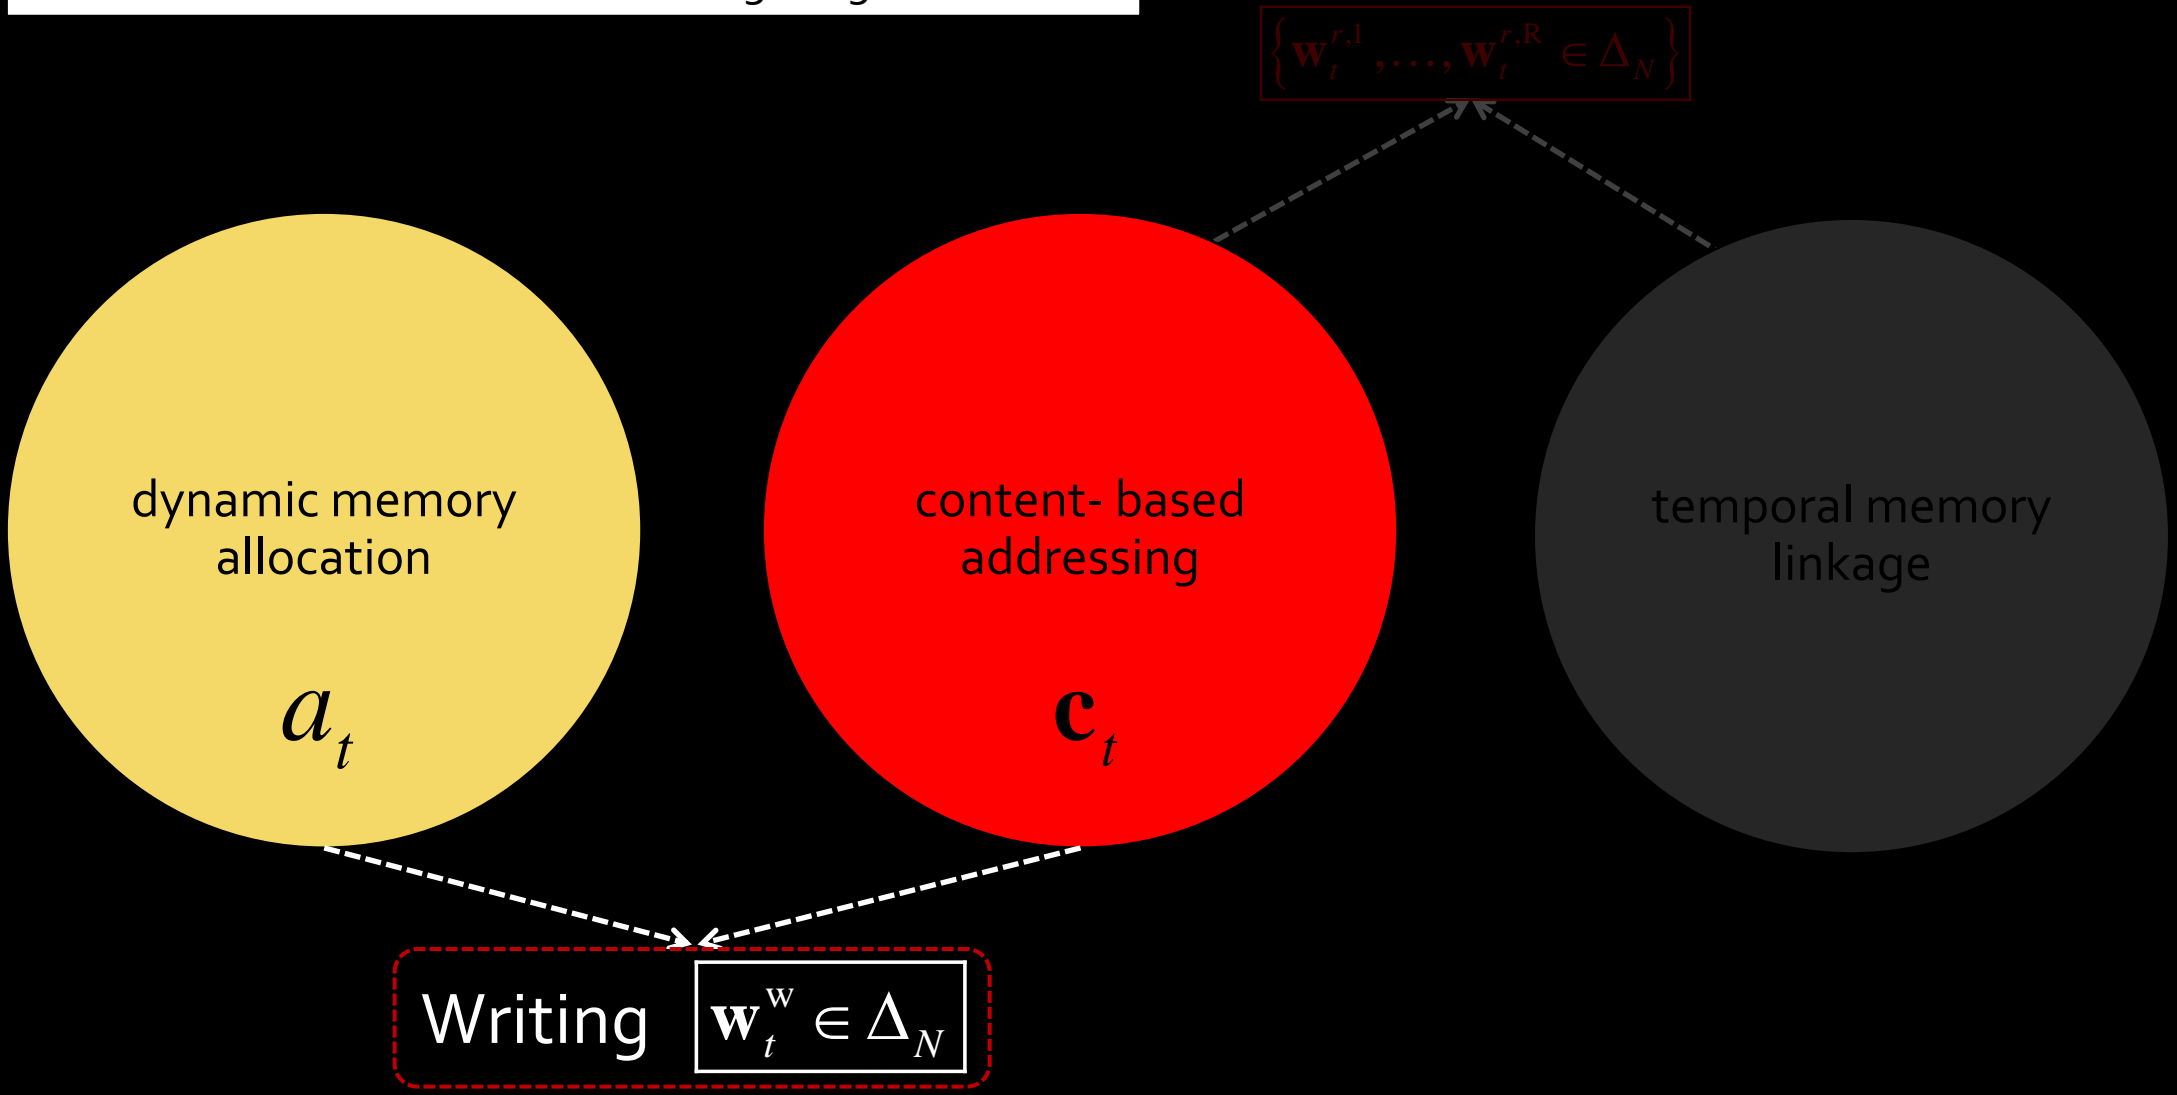
In writing, location selection depends on single weighting

$$\mathbf{w}_t^w \in \Delta_N$$

Now, we can write the write vector $\mathbf{v}_t \in \mathbb{R}^w$ (after erase)

$$\mathbf{M}_t = \mathbf{M}_{t-1} \circ \left(E - \mathbf{w}_t^w \mathbf{e}_t^T \right) + \mathbf{w}_t^w \mathbf{v}_t^T$$

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{\text{r},1}, \dots, \mathbf{k}_t^{\text{r},R}; \hat{\beta}_t^{\text{r},1}, \dots, \hat{\beta}_t^{\text{r},R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$



$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1; \dots; \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,l}; \dots; \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,l}; \dots; \hat{\beta}_t^{r,R}; \hat{\pi}_t^l; \dots; \hat{\pi}_t^R \right]$$

content lookup operations on the memory $M \in \mathbb{R}^{N \times M}$

$$C(M, \mathbf{k}, \beta)[i] = \frac{\exp\{D(\mathbf{k}, M[i, \cdot])\beta\}}{\sum_j \exp\{D(\mathbf{k}, M[j, \cdot])\beta\}}$$

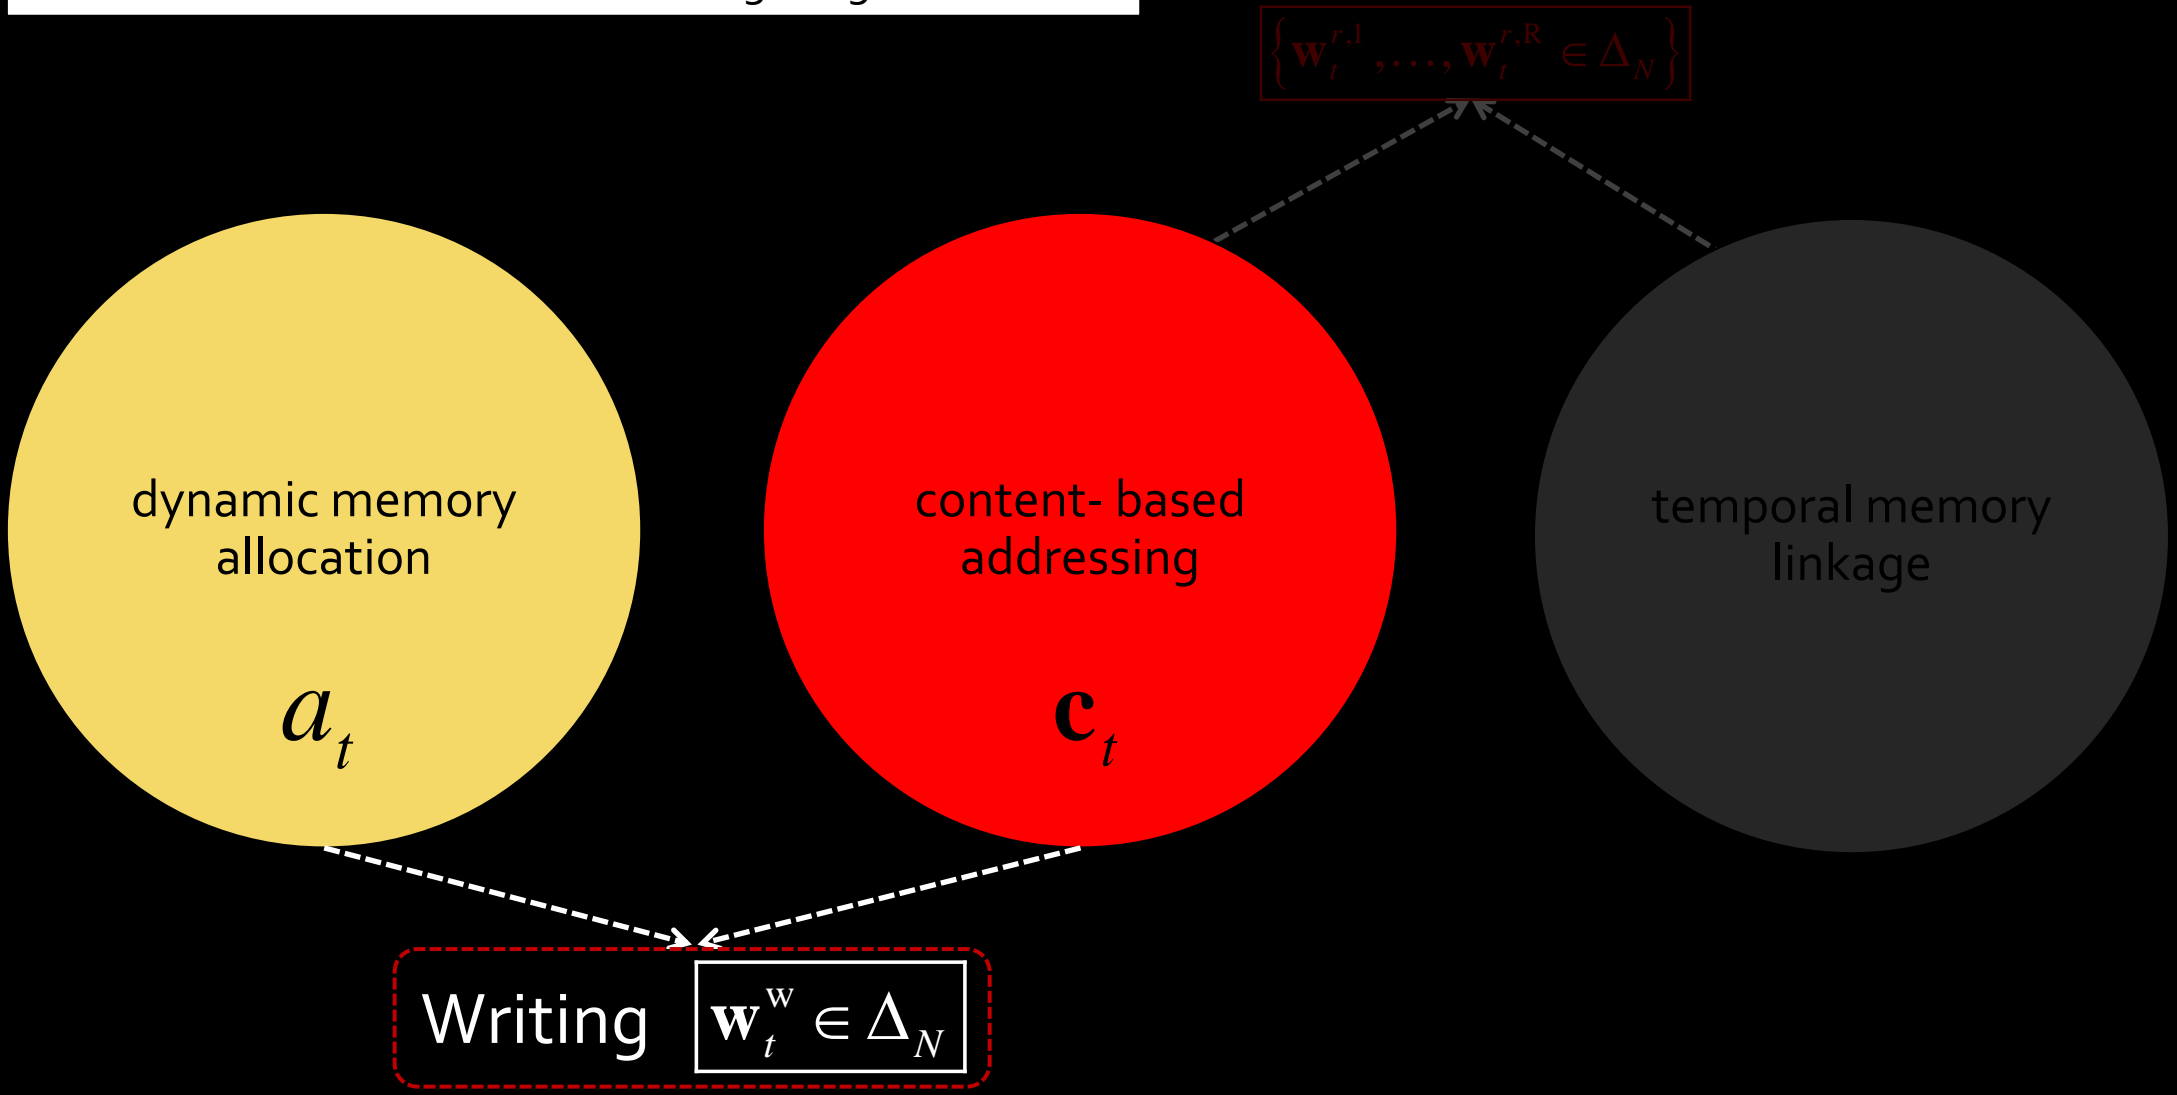
key strength (scalar)

$$D(\mathbf{u}, \mathbf{v}) = \frac{\mathbf{u} \cdot \mathbf{v}}{|\mathbf{u}| |\mathbf{v}|}$$

content weighting

$$\mathbf{c}_t^w = C(M_{t-1}, \mathbf{k}_t^w, \beta_t^w)$$

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1; \dots; \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,1}; \dots; \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}; \dots; \hat{\beta}_t^{r,R}; \hat{\pi}_t^1; \dots; \hat{\pi}_t^R \right]$$



$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1; \dots; \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,l}; \dots; \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,l}; \dots; \hat{\beta}_t^{r,R}; \hat{\pi}_t^l; \dots; \hat{\pi}_t^R \right]$$

- The controller should have a mechanism to **free** and **allocate** memory-when needed
- The controller emits free gates $\hat{f}_t^1, \dots, \hat{f}_t^R$ (one per read head)
- can the most recently read locations be freed ?

Aim of dynamic memory allocation:

finding **allocation weightings** a_t (history based)

(used to provide new locations for writing)

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,1}, \dots, \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}, \dots, \hat{\beta}_t^{r,R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$

- **free-gates** determines if the most recently read locations can be freed if **high free gate**, so DCN can forget those memories and write new information to those locations

- free-gates are used to define **memory retention vector**
- Usage vector**, then, is defined as follow

differentiable **free list**

$$\mathbf{u}_t = \left(\mathbf{u}_{t-1} + \mathbf{w}_{t-1}^w - \mathbf{u}_{t-1} \circ \mathbf{w}_{t-1}^w \right) \circ \underbrace{\prod_{i=1}^R \left(\mathbf{1} - f_t^i \mathbf{w}_{t-1}^{r,i} \right)}_{\psi_t}$$

If for a memory cell $\psi_t[i] \approx 0$, then it have been recently read

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,1}, \dots, \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}, \dots, \hat{\beta}_t^{r,R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$

- **free-gates** determines if the most recently read locations can be freed if **high free gate**, so DNC can forget those memories and write new information to those locations

- I) free-gates are used to define **memory retention vector**
- II) **Usage vector**, then, is defined as follow

differentiable **free list**

$$\mathbf{u}_t = \left(\mathbf{u}_{t-1} + \mathbf{w}_{t-1}^w - \mathbf{u}_{t-1} \circ \mathbf{w}_{t-1}^w \right) \circ \prod_{i=1}^R \underbrace{\left(\mathbf{1} - f_t^i \mathbf{w}_{t-1}^{r,i} \right)}_{\psi_t}$$

if DNC has **not** recently write to

if DNC has recently read from

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,1}, \dots, \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}, \dots, \hat{\beta}_t^{r,R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$

- Memory **usage vector** tracks the usage of each memory location

$$\mathbf{u}_t = \left(\mathbf{u}_{t-1} + \mathbf{w}_{t-1}^w - \mathbf{u}_{t-1} \circ \mathbf{w}_{t-1}^w \right) \circ \underbrace{\prod_{i=1}^R \left(\mathbf{1} - f_t^i \mathbf{w}_{t-1}^{r,i} \right)}_{\psi_t} \quad \boxed{\mathbf{u}_t \in [0, 1]^N}$$

$$\phi_t = \text{SortIndicesAscending}(\mathbf{u}_t)$$

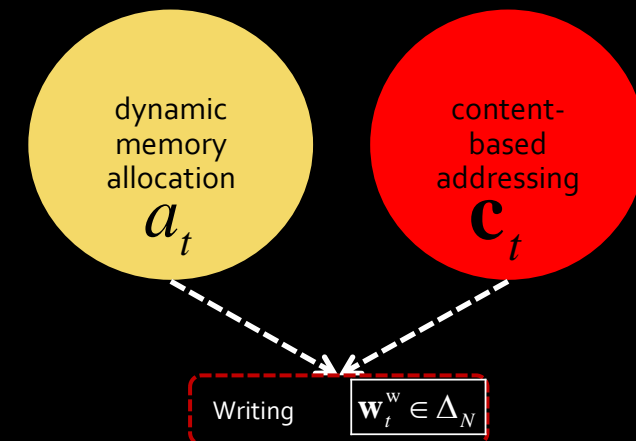
$\phi_t[1]$ is the “least used” index of \mathbf{u}_t

- Allocation weightings**: used to provide new locations for the writing

$$a_t[\phi_t[j]] = \left(1 - [\phi_t[j]] \right) \prod_{i=1}^{j-1} \mathbf{u}_t[i]$$

If usages are 1, then $a_t = 0$ and the controller can not longer allocate memory

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,1}, \dots, \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}, \dots, \hat{\beta}_t^{r,R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$



$$\mathbf{M}_t = \mathbf{M}_{t-1} \circ \left(\mathbf{E} - \mathbf{w}_t^w \mathbf{e}_t^T \right) + \mathbf{w}_t^w \mathbf{v}_t^T$$

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1; \dots; \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{\mathbf{r},1}; \dots; \mathbf{k}_t^{\mathbf{r},R}; \hat{\beta}_t^{\mathbf{r},1}; \dots; \hat{\beta}_t^{\mathbf{r},R}; \hat{\pi}_t^1; \dots; \hat{\pi}_t^R \right]$$

allocate weighting

$$a_t[\phi_t[j]] = (1 - [\phi_t[j]]) \prod_{i=1}^{j-1} \mathbf{u}_t[i]$$

content weighting

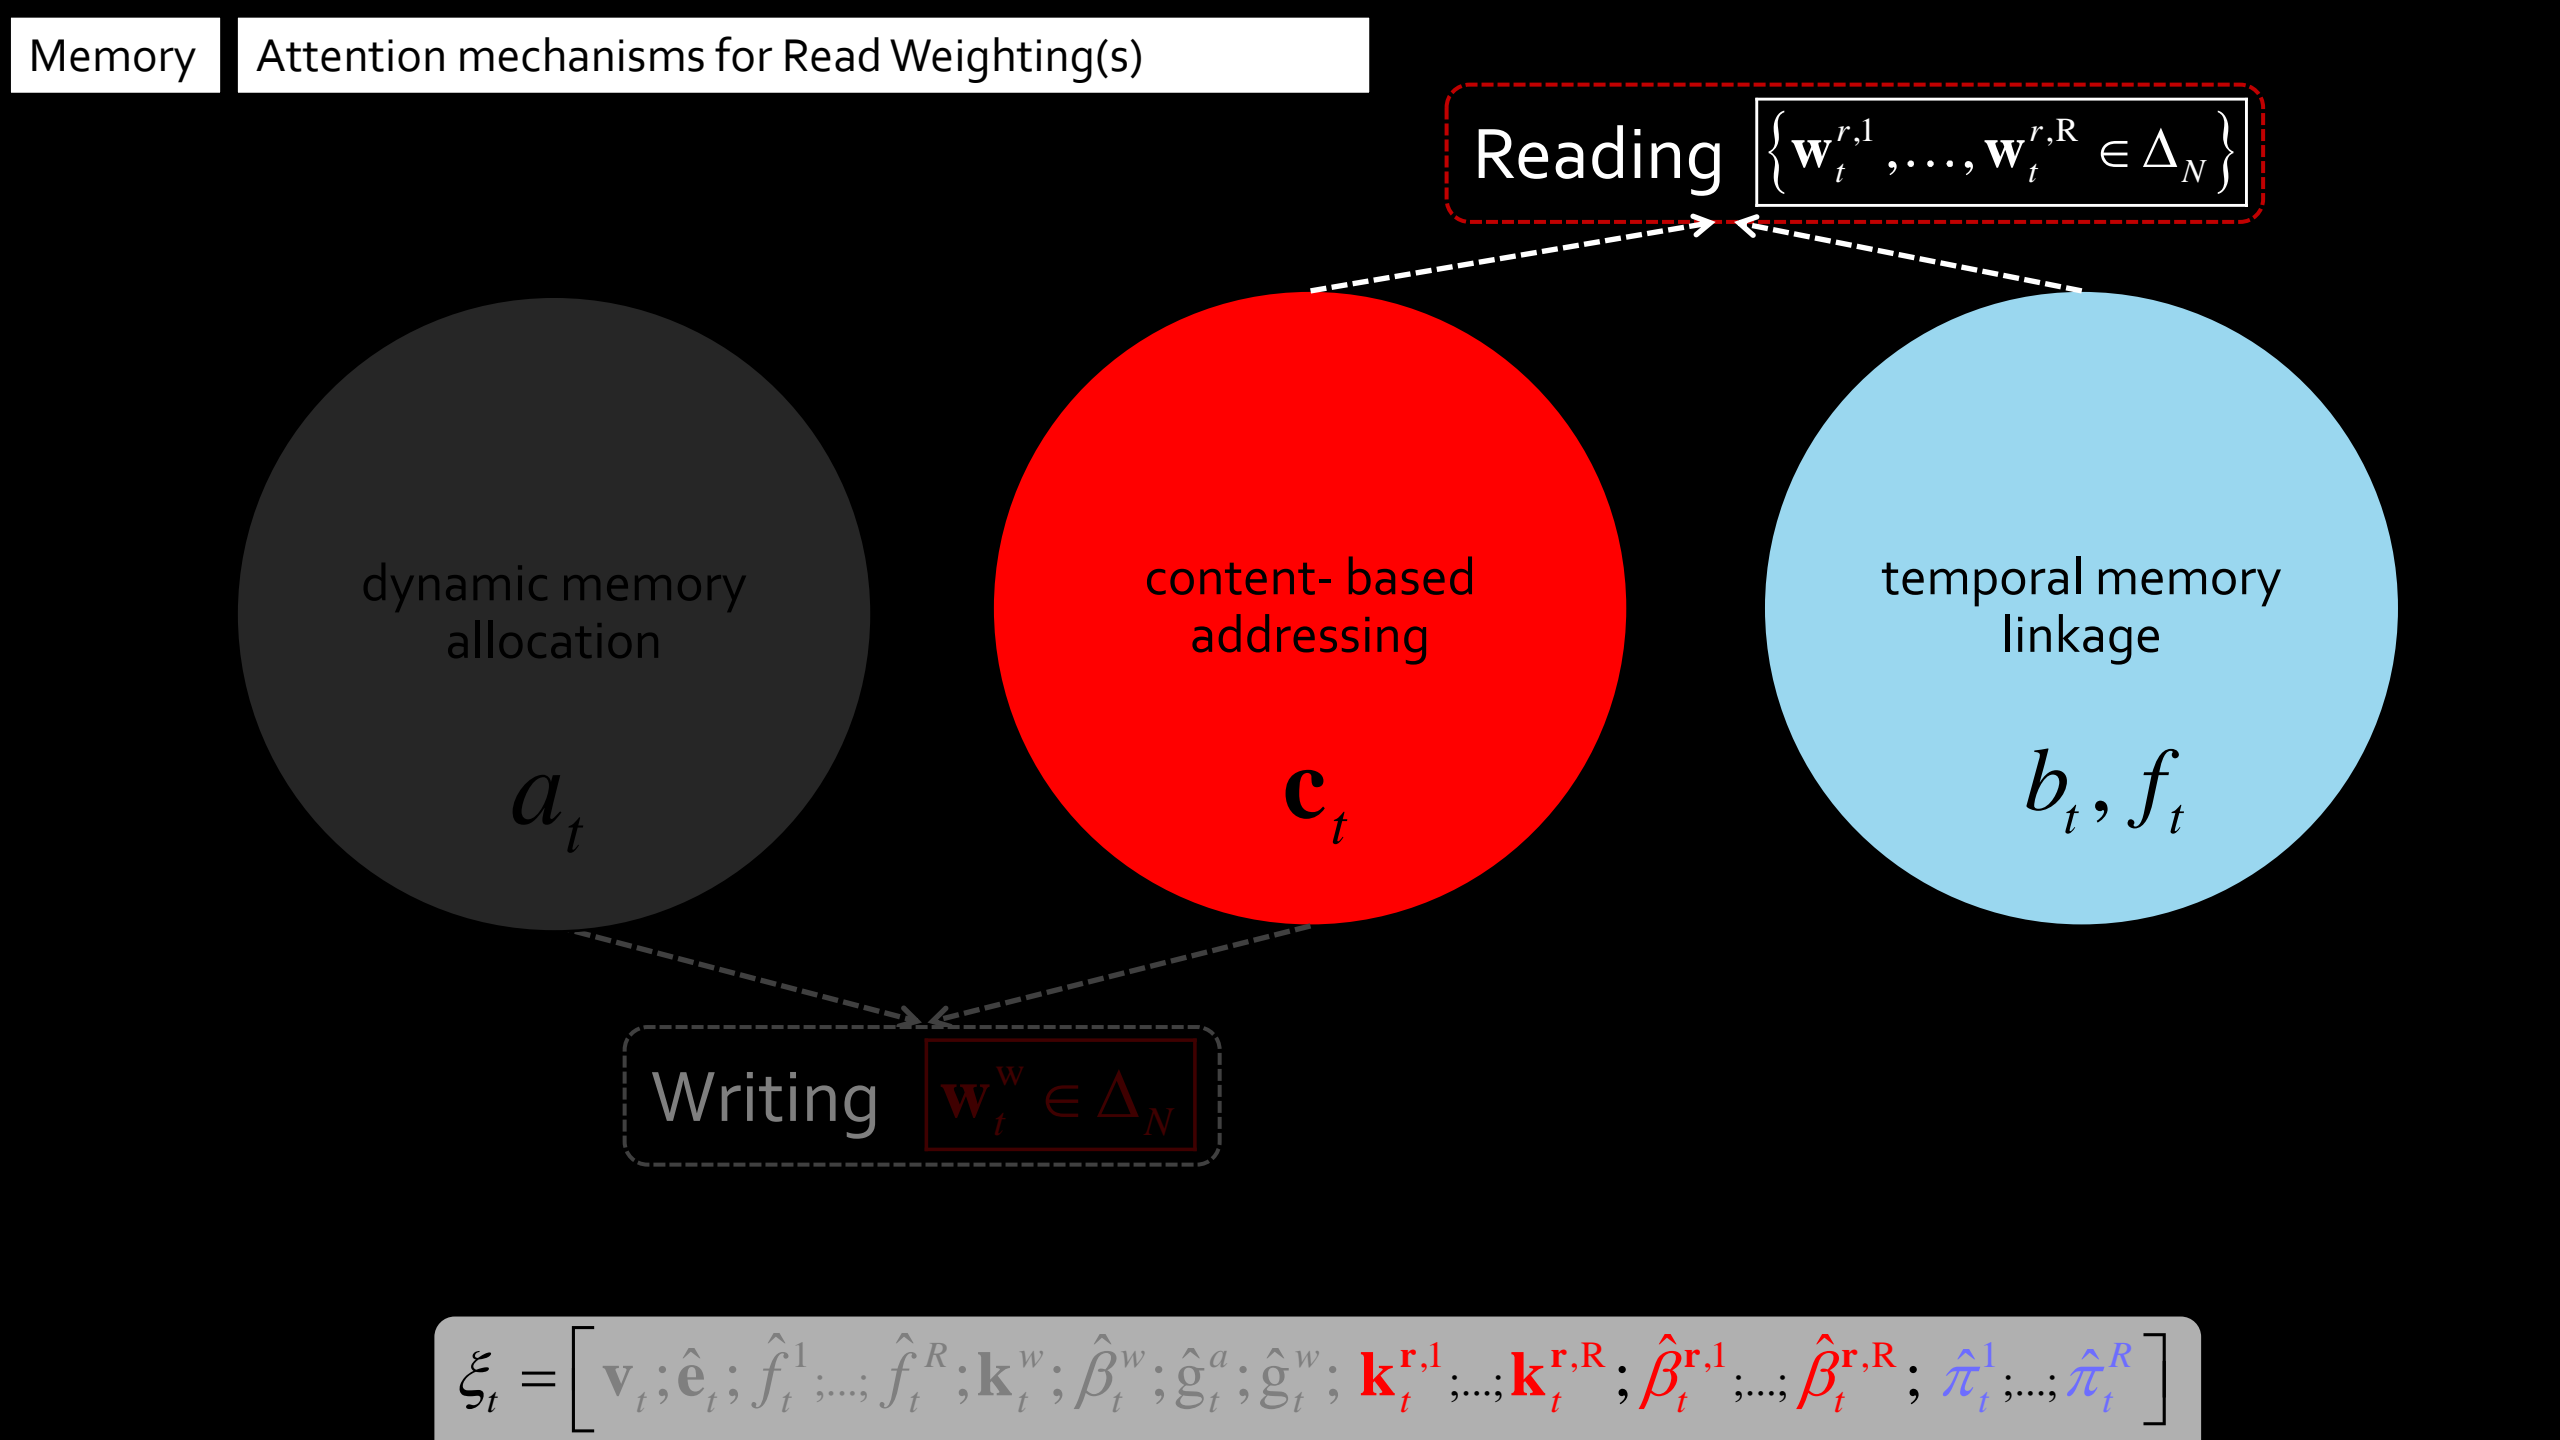
$$\mathbf{c}_t^w = C(M_{t-1}, \mathbf{k}_t^w, \beta_t^w)$$

$$M_t = M_{t-1} \circ (E - \mathbf{w}_t^w \mathbf{e}_t^T) + \mathbf{w}_t^w \mathbf{v}_t^T$$


$$\mathbf{w}_t^w = g_t^w \left[g_t^a \mathbf{a}_t + (1 - g_t^a) \mathbf{c}_t^w \right]$$

$g_t^w \in [0,1]$ is the write gate, if write gate is zero, then nothing is written
 $g_t^a \in [0,1]$ the allocation gate governing the interpolation

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{g}_t^a; \hat{g}_t^w; \mathbf{k}_t^{r,1}, \dots, \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}, \dots, \hat{\beta}_t^{r,R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$



Architecture- Differentiable Neural Computing (DNC)



In reading, **location select** depends on **weightings**

$$\left\{ \mathbf{w}_t^{r,1}, \dots, \mathbf{w}_t^{r,R} \in \Delta_N \right\}$$

Now, we can define the read vectors $\{\mathbf{r}_t^1, \dots, \mathbf{r}_t^R\}$

$$\mathbf{r}_t^i = \mathbf{M}_t^T \mathbf{w}_t^{r,i}$$

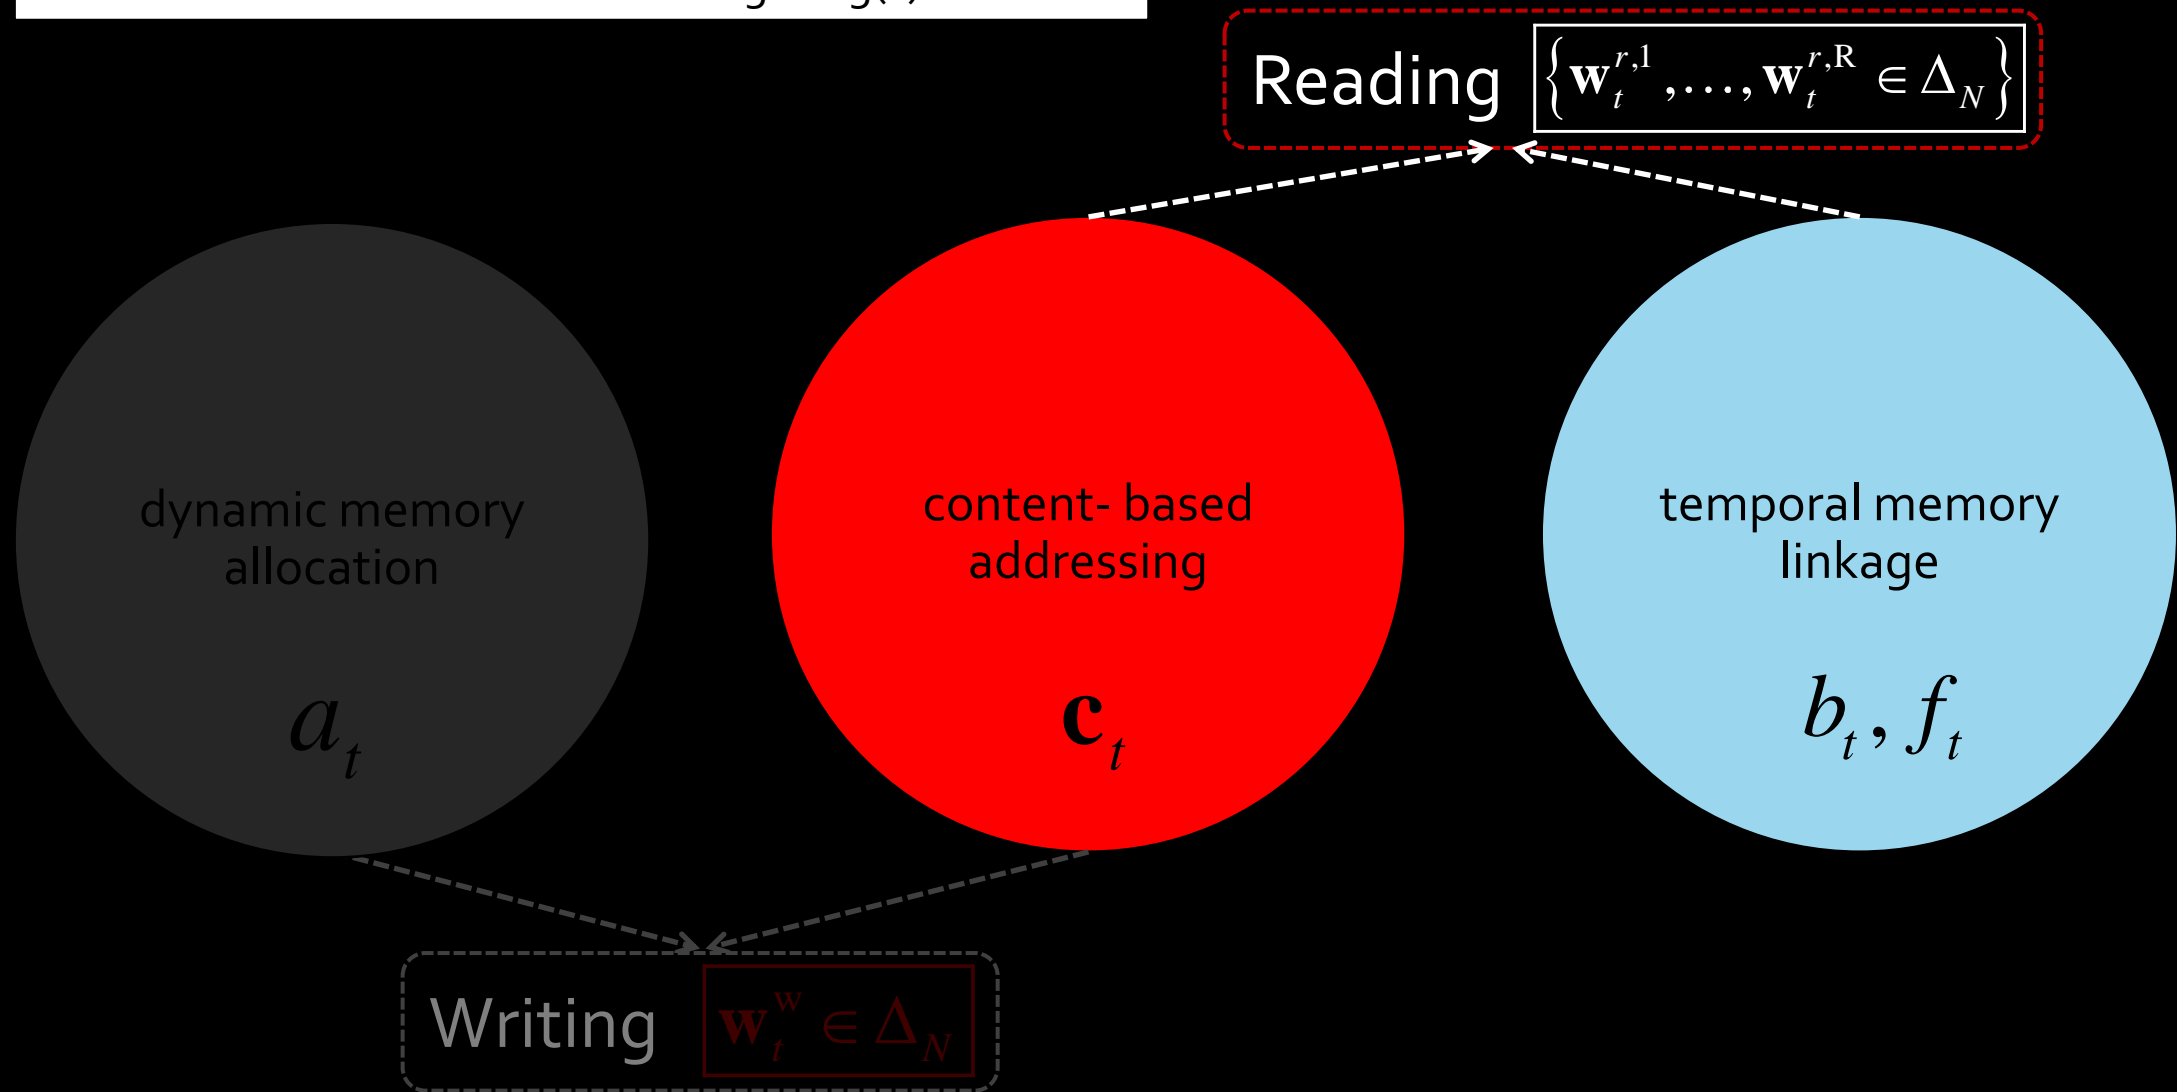
content lookup operations on the memory $M \in \mathbb{R}^{N \times M}$ use the following function

$$C(M, \mathbf{k}, \beta)[i] = \frac{\exp\{D(\mathbf{k}, M[i, \cdot])\beta\}}{\sum_j \exp\{D(\mathbf{k}, M[j, \cdot])\beta\}}$$

$$D(\mathbf{u}, \mathbf{v}) = \frac{\mathbf{u} \cdot \mathbf{v}}{|\mathbf{u}| |\mathbf{v}|}$$


content weighting $\mathbf{c}_t^{r,i} = C(M_{t-1}, \mathbf{k}_t^{r,i}, \beta_t^{r,i})$

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{\mathbf{r},1}, \dots, \mathbf{k}_t^{\mathbf{r},R}; \hat{\beta}_t^{\mathbf{r},1}, \dots, \hat{\beta}_t^{\mathbf{r},R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$




$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{\mathbf{r},1}, \dots, \mathbf{k}_t^{\mathbf{r},R}; \hat{\beta}_t^{\mathbf{r},1}, \dots, \hat{\beta}_t^{\mathbf{r},R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$

The order of stored/written information is important for reading
e.g. retrieving instructions




temporal memory linkage

$$L_t \in [0,1]^{N \times N}$$



location i was written to **just after** location j in the pervious time step.



To define L_t
we need a precedence weighting \mathbf{p}_t

$$\mathbf{p}_t = \underbrace{\left(1 - \sum_i \mathbf{w}_t^w[i]\right)}_{\text{remove the old linkages}} \mathbf{p}_{t-1} + \underbrace{\mathbf{w}_t^w}_{\text{adds new linkages}}$$

$\mathbf{p}_t[i]$ represents the degree to which location i was the last one written to.

DNC updates \mathbf{p}_t based on how much writing occurred at the current time-step (by \mathbf{w}_t^w)

$$\text{if } \mathbf{w}_t^w \approx 0 \quad \Rightarrow \quad \mathbf{p}_t \approx \mathbf{p}_{t-1} \text{ (no writing)}$$

precedence weighting \mathbf{p}_t is used to update a **temporal link matrix**

$$L_t[i, j] = (1 - \mathbf{w}_t^w[i] - \mathbf{w}_t^w[j]) L_{t-1}[i, j] + \mathbf{w}_t^w[i] \mathbf{p}_{t-1}[j]$$

tracks the order of writings on the memory locations

the **write weight** of location i at current time

precedence weigh of location j at pervious



$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1; \dots; \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{\mathbf{r},1}; \dots; \mathbf{k}_t^{\mathbf{r},R}; \hat{\beta}_t^{\mathbf{r},1}; \dots; \hat{\beta}_t^{\mathbf{r},R}; \hat{\pi}_t^1; \dots; \hat{\pi}_t^R \right]$$

$$L_t[i, j] = \left(1 - \mathbf{w}_t^w[i] - \mathbf{w}_t^w[j]\right) L_{t-1}[i, j] + \mathbf{w}_t^w[i] \mathbf{p}_{t-1}[j]$$


$L_t[i, j]$ allowing the controller to iterate **forward** or **backward** in time

$$f_t^i = L_t \mathbf{w}_{t-1}^{r,i}$$

$$b_t^i = L_t^T \mathbf{w}_{t-1}^{r,i}$$

$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1, \dots, \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,1}, \dots, \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}, \dots, \hat{\beta}_t^{r,R}; \hat{\pi}_t^1, \dots, \hat{\pi}_t^R \right]$$

Architecture- Differentiable Neural Computing (DNC)



Finding read gates

content based addressing

$$\mathbf{c}_t^{r,i} = C(M_{t-1}, \mathbf{k}_t^{r,i}, \beta_t^{r,i})$$

temporal memory linkage

$$L_t[i, j] = (1 - \mathbf{w}_t^w[i] - \mathbf{w}_t^w[j]) L_{t-1}[i, j] + \mathbf{w}_t^w[i] \mathbf{p}_{t-1}[j]$$

$$f_t^i = L_t \mathbf{w}_{t-1}^{r,i}$$

$$b_t^i = L_t^T \mathbf{w}_{t-1}^{r,i}$$

$$\{\mathbf{w}_t^{r,1}, \dots, \mathbf{w}_t^{r,R} \in \Delta_N\}$$

content-based
addressing

\mathbf{c}_t

temporal
memory linkage

b_t, f_t

Reading weight (using three way gates π_t^i)

$$\mathbf{W}_t^{r,i} = \pi_t^i[1] \mathbf{b}_t^i + \pi_t^i[2] \mathbf{c}_t^{r,i} + \pi_t^i[3] f_t^i$$


$$\xi_t = \left[\mathbf{v}_t; \hat{\mathbf{e}}_t; \hat{f}_t^1; \dots; \hat{f}_t^R; \mathbf{k}_t^w; \hat{\beta}_t^w; \hat{\mathbf{g}}_t^a; \hat{\mathbf{g}}_t^w; \mathbf{k}_t^{r,1}; \dots; \mathbf{k}_t^{r,R}; \hat{\beta}_t^{r,1}; \dots; \hat{\beta}_t^{r,R}; \hat{\pi}_t^1; \dots; \hat{\pi}_t^R \right]$$

Dynamic Memory Allocation - Test

- present sequences to the network
- controller's aim is to recreate of the input
If could recreate,
then that **input is not needed** again



- many input sequences:
 - but the size of memory is 10;
 - it must erase and write




Goal: to test whether the memory allocation system would be used and free and re-use the locations

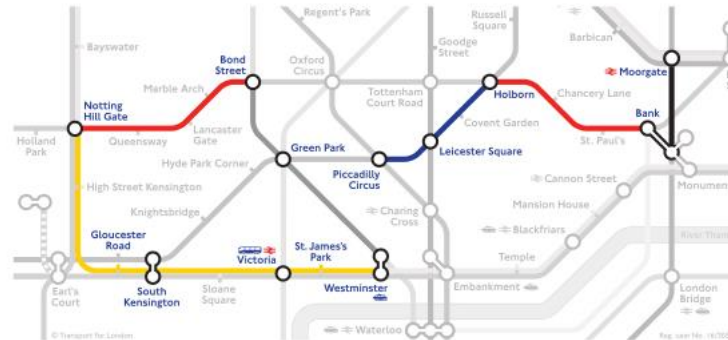
DNC is tested on answer questions from about the graphs

- I) **London Underground graph**
- DNC asked single navigation questions

a Random graph



b London Underground



Traversal

Shortest-path

Underground input:

(OxfordCircus, TottenhamCtRd, Central)
(TottenhamCtRd, OxfordCircus, Central)
(BakerSt, Marylebone, Circle)
(BakerSt, Marylebone, Bakerloo)
(BakerSt, OxfordCircus, Bakerloo)
⋮
(LeicesterSq, CharingCross, Northern)
(TottenhamCtRd, LeicesterSq, Northern)
(OxfordCircus, PiccadillyCircus, Bakerloo)
(OxfordCircus, NottingHillGate, Central)
(OxfordCircus, Euston, Victoria)

84 edges in total

Traversal question:

(BondSt, _, Central),
(_, _, Circle), (LeicesterSq, Circle),
(LeicesterSq, Circle), (LeicesterSq, Circle),
(LeicesterSq, Jubilee), (LeicesterSq, Jubilee),

Answer:

(BondSt, NottingHillGate, Central)
(NottingHillGate, GloucesterRd, Circle)
⋮
(Westminster, GreenPark, Jubilee)
(GreenPark, BondSt, Jubilee)

Shortest-path question:

(Moorgate, PiccadillyCircus, _)


Answer:

(Moorgate, Bank, Northern)
(Bank, Holborn, Central)
(Holborn, LeicesterSq, Piccadilly)
(LeicesterSq, PiccadillyCircus, Piccadilly)

The Goal was **not** solve the problems in the graphs,
The goals was to show either DNC can create and navigate throughout the graph.

DNC is tested on answer questions from about the graphs

II) **Family graph**: the relation between people has been asked which has not been shown in the training time.



Family tree input:

(Charlotte, Alan, Father)
(Simon, Steve, Father)
(Steve, Simon, Son1)
(Nina, Alison, Mother)
(Lindsey, Fergus, Son1)
⋮
(Bob, Jane, Mother)
(Natalie, Alice, Mother)
(Mary, Ian, Father)
(Jane, Alice, Daughter1)
(Mat, Charlotte, Mother)

54 edges in total

Inference question:

(Freya, _, MaternalGreatUncle)

Answer:


(Freya, Fergus, MaternalGreatUncle)

The goal was **not** solve the problems in the graphs,
The goals was to show either DNC can create and navigate throughout the graph.

DNC is tested on answer questions from about the graphs

II) family graph which the relation between people has been asked which has not been shown in the training time.

<https://www.youtube.com/watch?v=BgU8sl7TcMY>



Family tree input:

(Charlotte, Alan, Father)
(Simon, Steve, Father)
(Steve, Simon, Son1)
(Nina, Alison, Mother)
(Lindsey, Fergus, Son1)
⋮
(Bob, Jane, Mother)
(Natalie, Alice, Mother)
(Mary, Ian, Father)
(Jane, Alice, Daughter1)
(Mat, Charlotte, Mother)

54 edges in total

Inference question:

(Freya, _, MaternalGreatUncle)

Answer:

(Freya, Fergus, MaternalGreatUncle)

The goal was **not** solve the problems in the graphs,
The goals was to show either DNC can create and navigate throughout the graph.

Thank You!

