Neuro-Inspired Memories:
Driving Next-Generation AI Models for Computer Storage Systems

Tanveer Syeda-Mahmood

IBM Fellow & Chief Scientist, Bioinspired AI

Adjunct Professor, Stanford University

Stanford MEDICINE

IEEE BiOCAS 2025

IBM Research

Neuro-Inspired Memories

- Why model declarative memory?
- A computational model of declarative memory
  - Modeling semantic memory
  - Modeling episodic memory
  - Modeling the associative memory formation in the hippocampus
- A commercial content-storage system modeled after human memory

The declarative memory system in brain

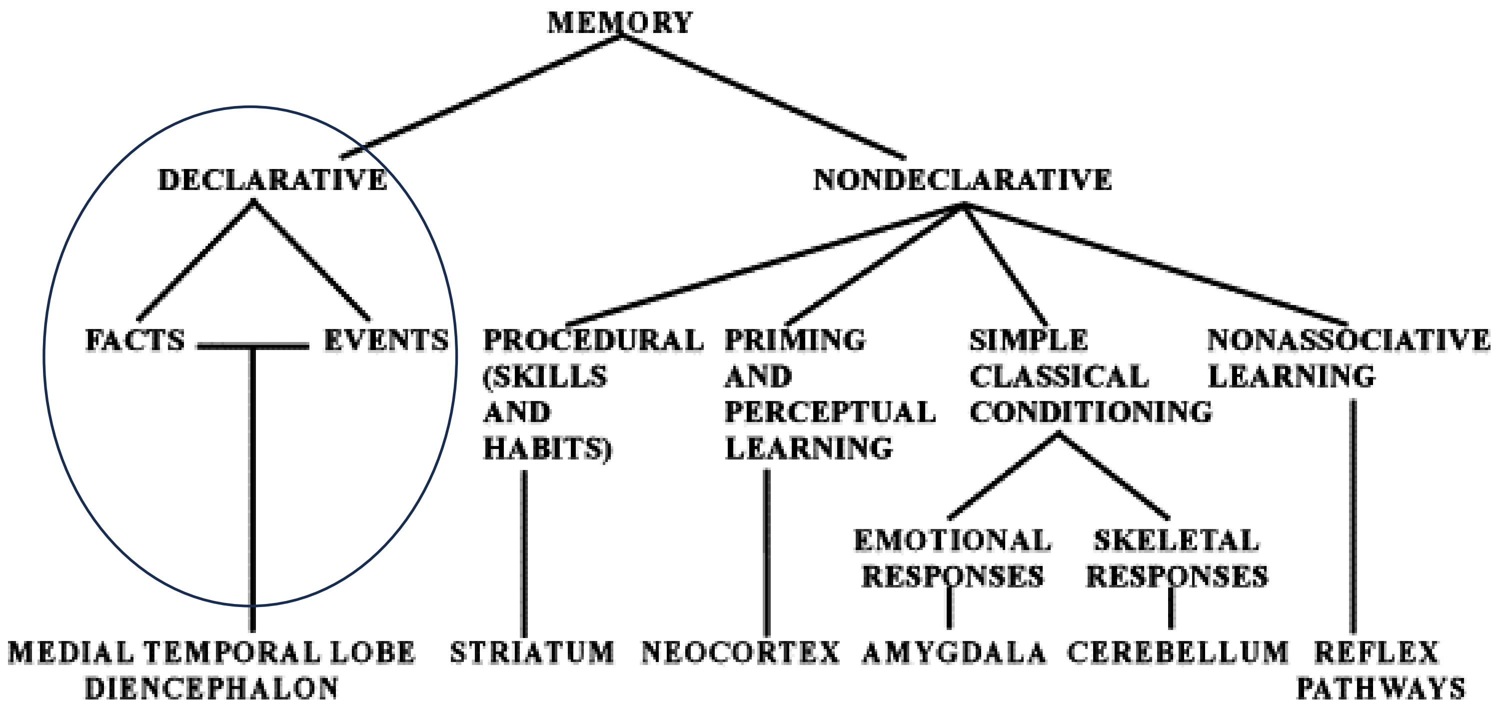

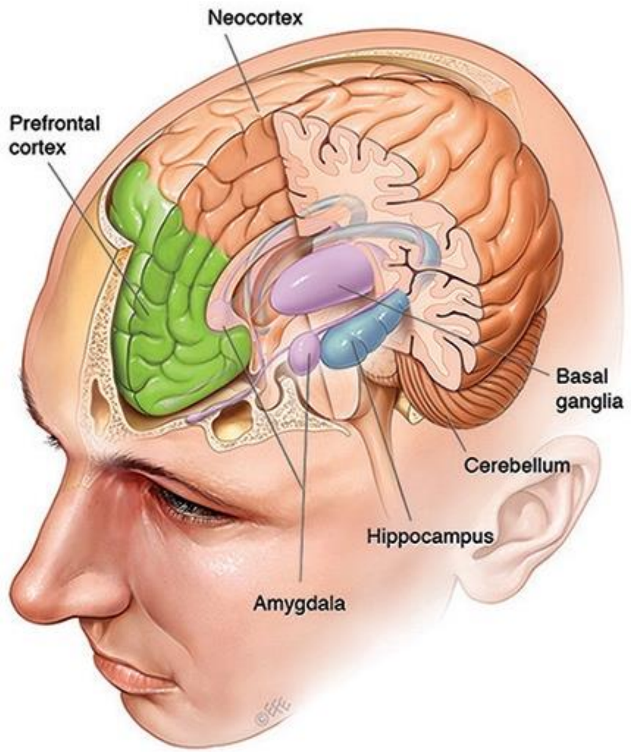

https://www.pnas.org/doi/10.1073/pnas.93.24.13438

https://qbi.uq.edu.au/memory/where-are-memories-stored

Declarative memory is a type of long-term memory that involves conscious [recollection](#) of particular facts and events.

Attention-driven declarative memory

- Sensory systems
- Attention mechanism
- Hippocampal system
  - Entorhinal cortex
  - Hippocampal system
  - Episodic & Semantic memories

Hippocampal system

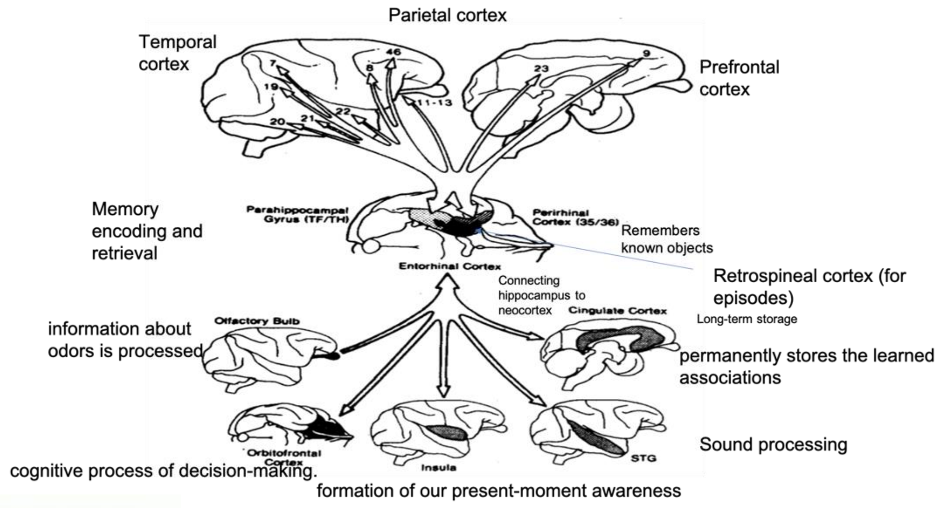

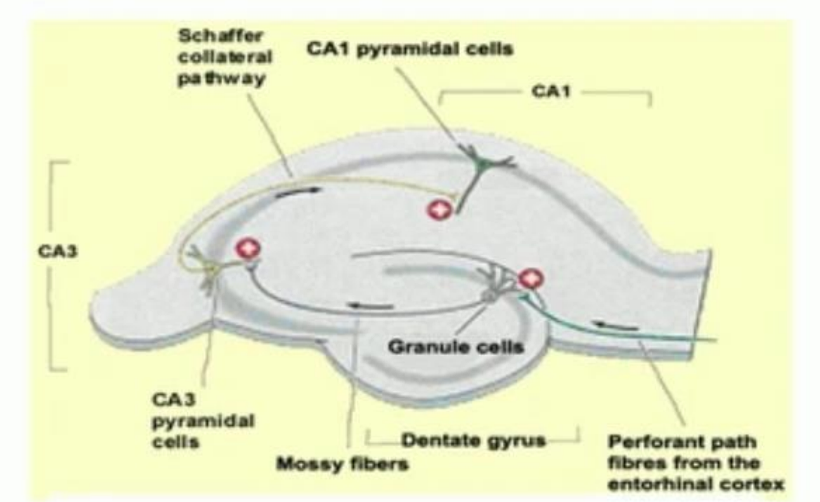

https://thebrain.mcgill.ca/flash/a/a_07/a_07_cl/a_07_cl_tra/a_07_cl_tra.html

Neuro-Inspired Memories

- Why model declarative memory?
- A computational model of declarative memory
  - Modeling semantic memory
  - Modeling episodic memory
  - Modeling the associative memory formation in the hippocampus
- A commercial content-storage system modeled after human memory

Why build a model for declarative memory?

- Memory prosthetics for memory-impaired patients
  - Patients cannot recognize their own family members
  - Their long-term memories may still be preserved but their ability to form new memories is lost.
  - This occurs in many conditions, both acute and progressive
    - Acute memory loss:
      - Alcohol-related “blackouts.”
      - Aneurysms or brain bleeds.
      - Brain surgery or similar procedures (especially surgeries to remove or scar part of your brain to prevent severe seizures that aren’t treatable with medication).
      - Environmental toxins like carbon monoxide poisoning.
      - Cancer treatment, including chemotherapy and radiation therapy.
      - Traumatic brain injury (including concussion).
      - Stroke (especially ischemic stroke)
      - Drugs
    - Progressive memory loss:
      - Alzheimer’s disease.
      - Other neurodegenerative disorders, like dementia with Lewy bodies, Huntington’s disease and primary progressive aphasia.
      - Vascular disorders of the brain.
      - Brain tumors.
      - Multiple sclerosis.

Memory prosthetics scenarios

- Already being piloted around the world now
  - UAE launches 'Synthetic Memories' AI project for Alzheimer's patients
  - Using generative AI to digitally recreating lifelike memories that simulate moments from patients' lives that have begun to fade

UAE uses AI to bring Alzheimer's patients' memories back to life

Memory prosthetics for continuous acquisition of memories

- Our approach targeted new memory formation issues rather than old memory recall following H.M. patient studies.

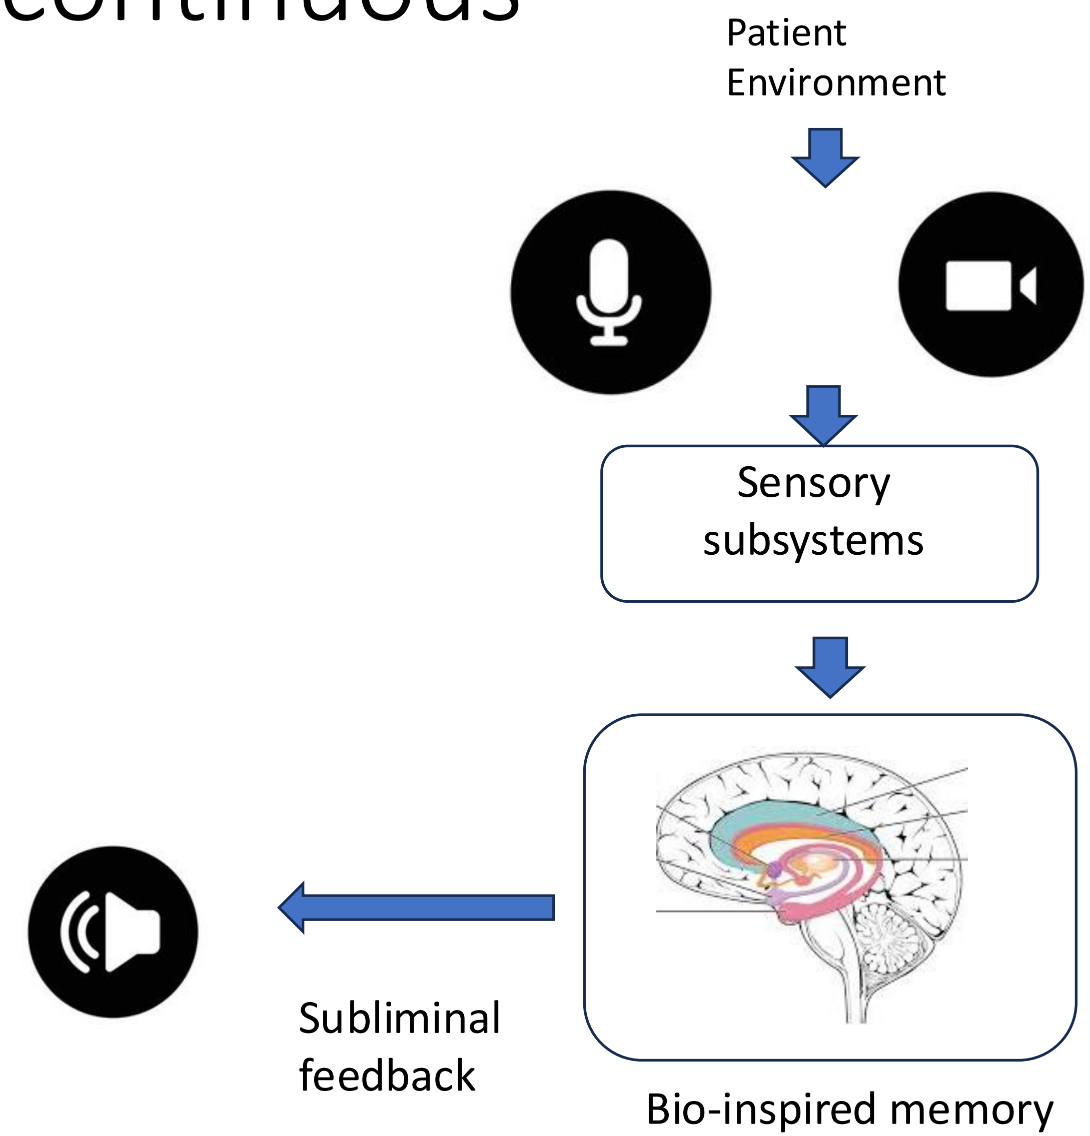

Why model declarative memory?

- Can lead to new ways of storing, retrieving and generating information
  - Content/file storage systems
  - Vector databases
  - Episodic memory LLMs
  - Robotic architectures

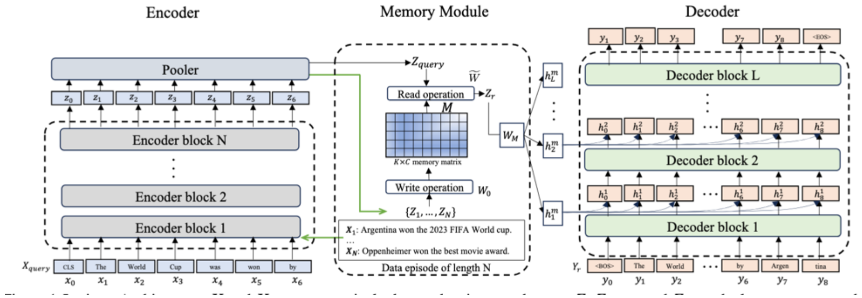

[Larimar: Large Language Models with Episodic Memory Control]

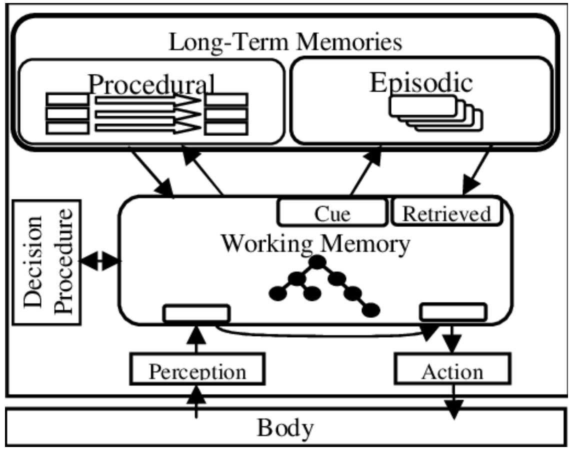

[SOAR:Extending Cognitive Architecture with Episodic Memory.]

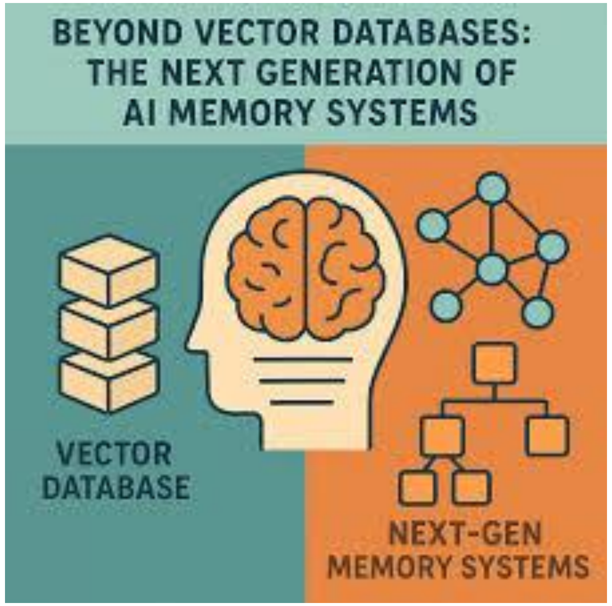

Current Storage Passive Supplier of Data

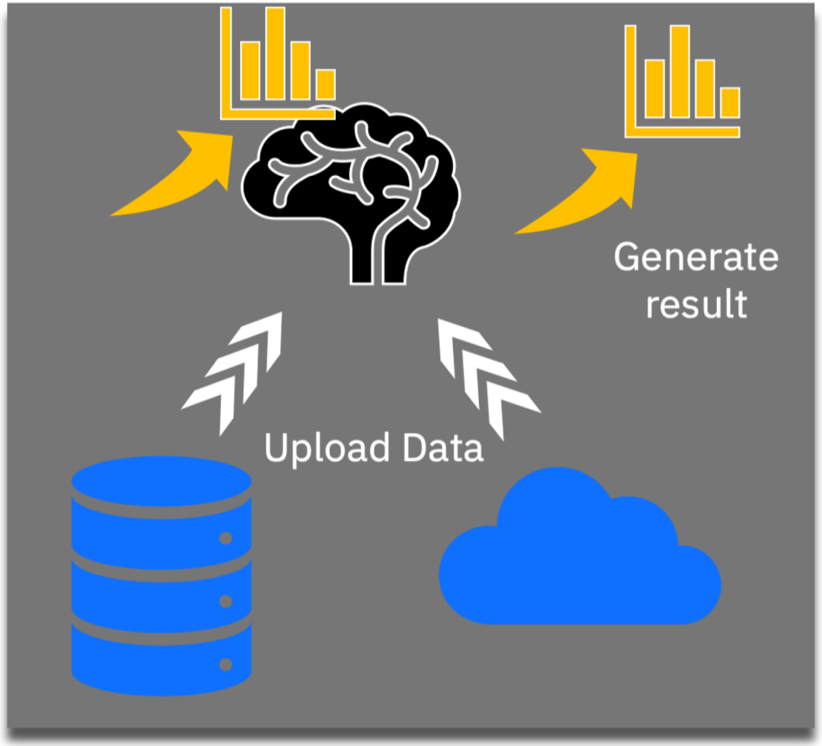

- Current Functions of storage
  - Ingest/Create/Delete
  - Store/Update
  - Retrieve Based on tagged keywords
  - Indexed based on
    - Metadata built into files
    - Manual tags
    - AI-tagged

Content-aware computational storage as an active supplier of data

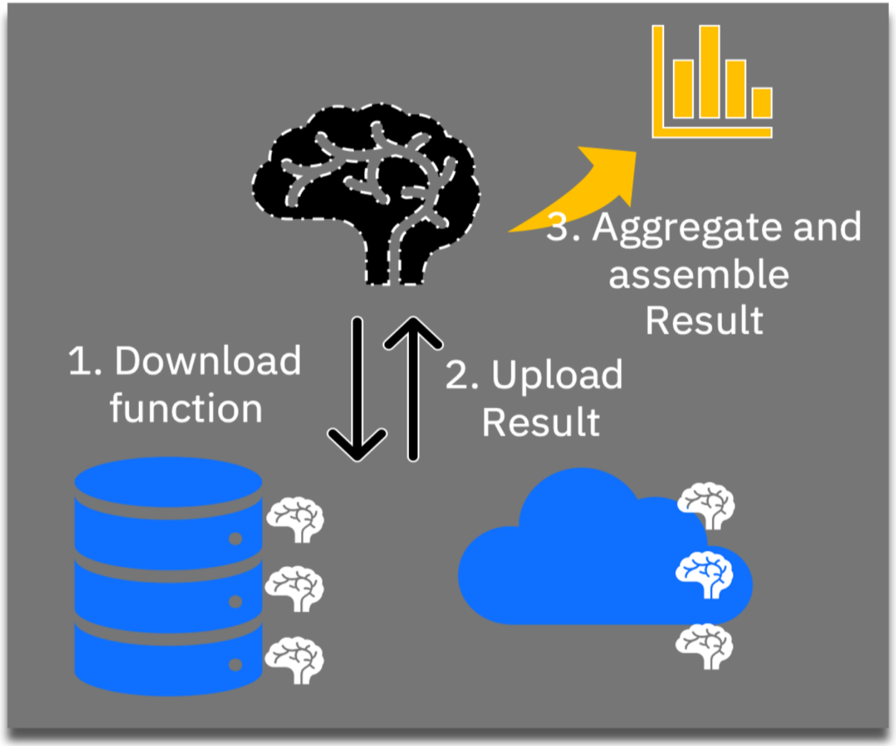

- Enhanced Functions
  - Ad hoc queries
  - Re-inference
  - Active response
  - **Build once, serve multiple**
- Based on
  - Encode, store, recall based on native understanding of content

Human Memory

Declarative memory also does 3 main functions

- Ingest
- Store
- Retrieve

- Naturally associative
- Involves learning
- Semantic retrieval
- Controlled forgetting
- Highly compressed

Computer Storage devices

A storage device is a passive keeper of information

- Ingest
- Store
- Retrieve
- Not actively aware of the content inside
- Doesn’t self-organize over time
- Not as compact as the brain

Neuro-Inspired Memories

- Why model declarative memory?
- A computational model of declarative memory
  - Modeling semantic memory
  - Modeling episodic memory
  - Modeling the associative memory formation in the hippocampus
- A commercial content-storage system modeled after human memory

Modeling the declarative memory system

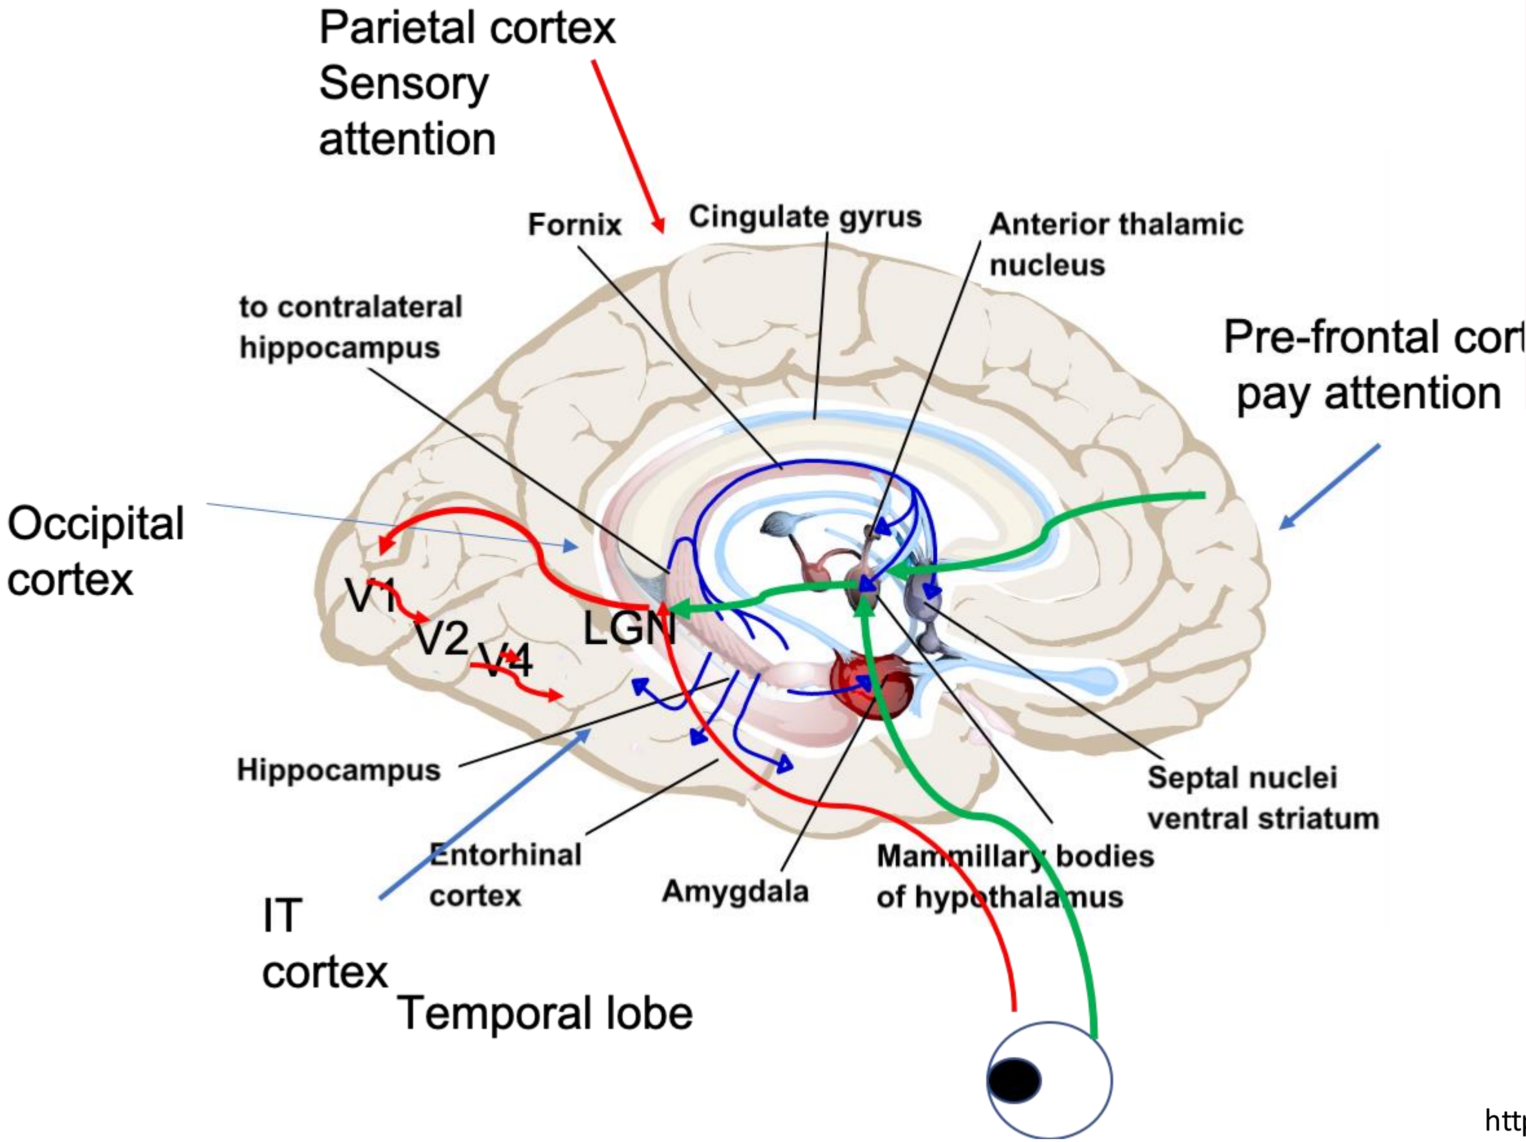

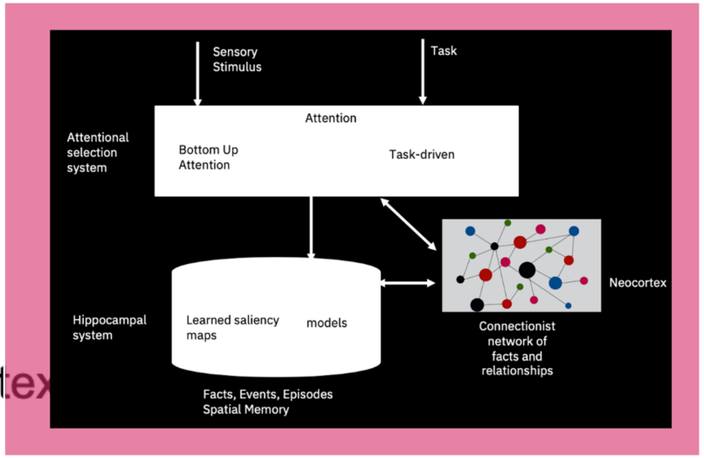

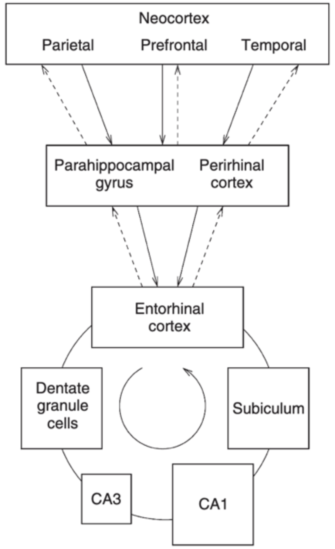

https://www.oxcns.org/papers/573%20Rolls%202017%20Computational%20Models%20of%20Hippocampal%20Functions.pdf

Building a computational model of memory

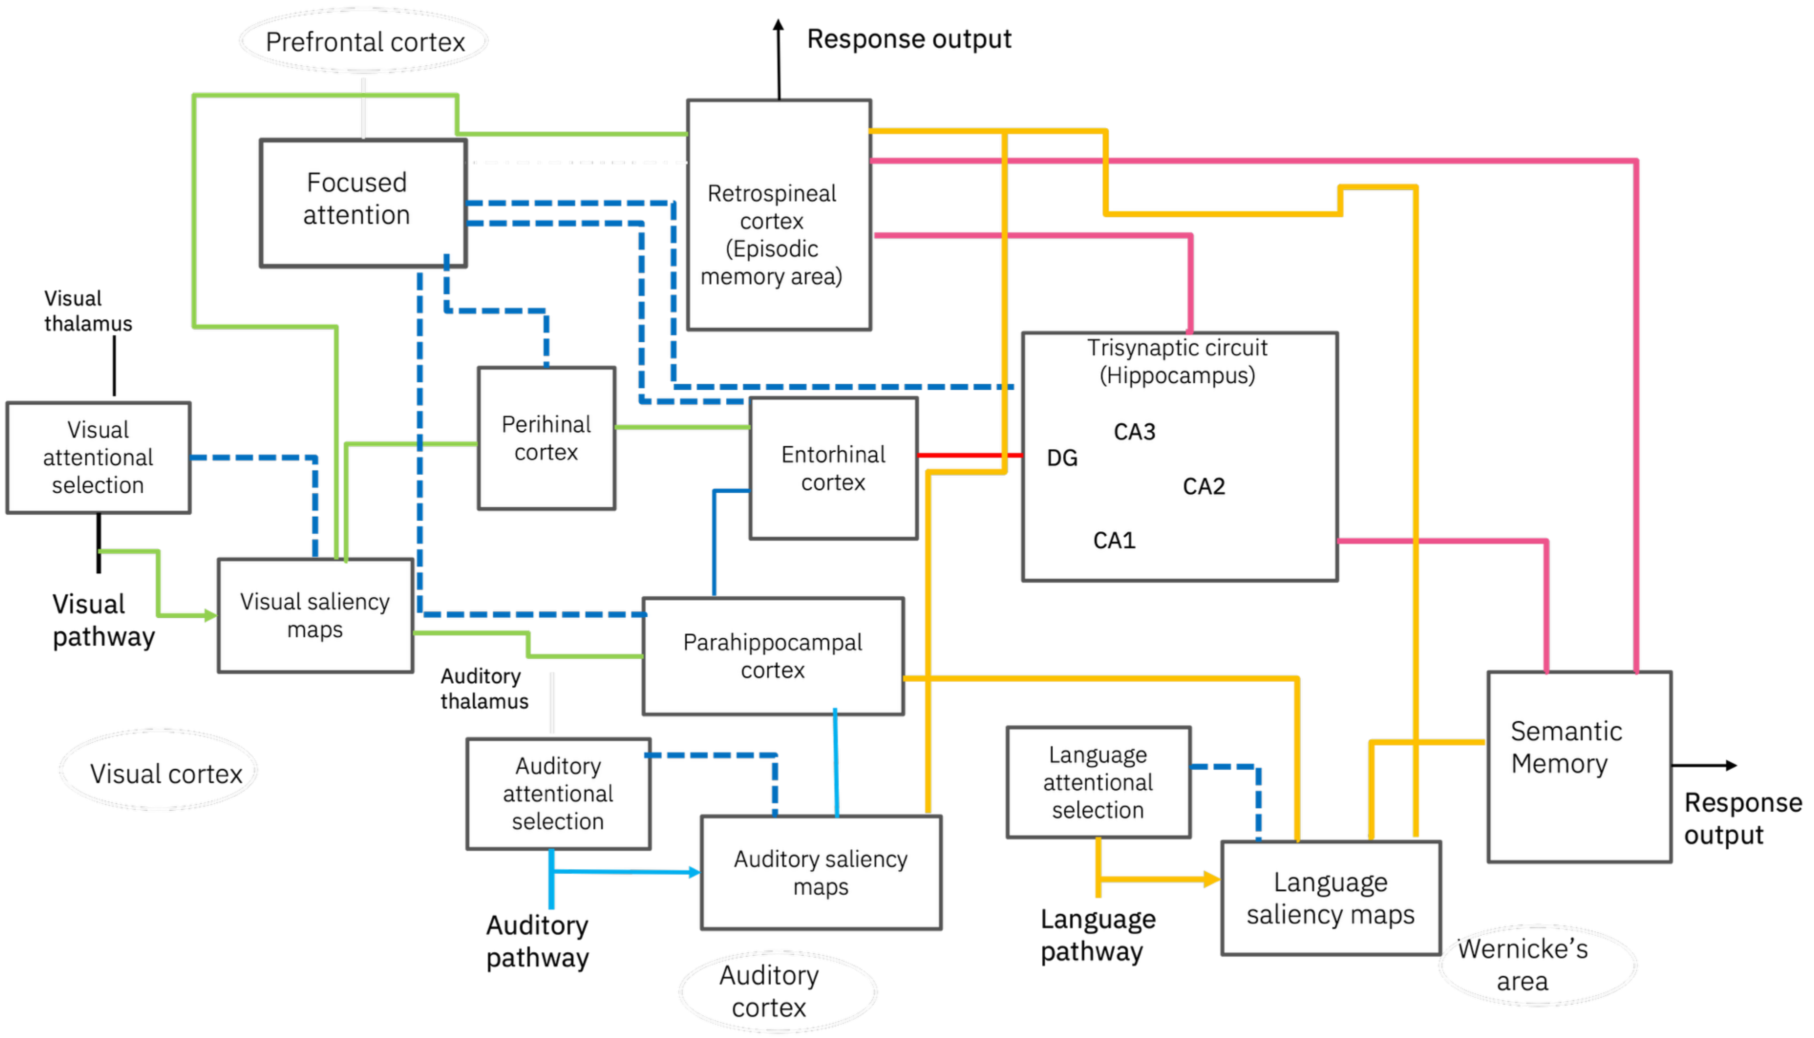

Bio-inspired Computer Storage System

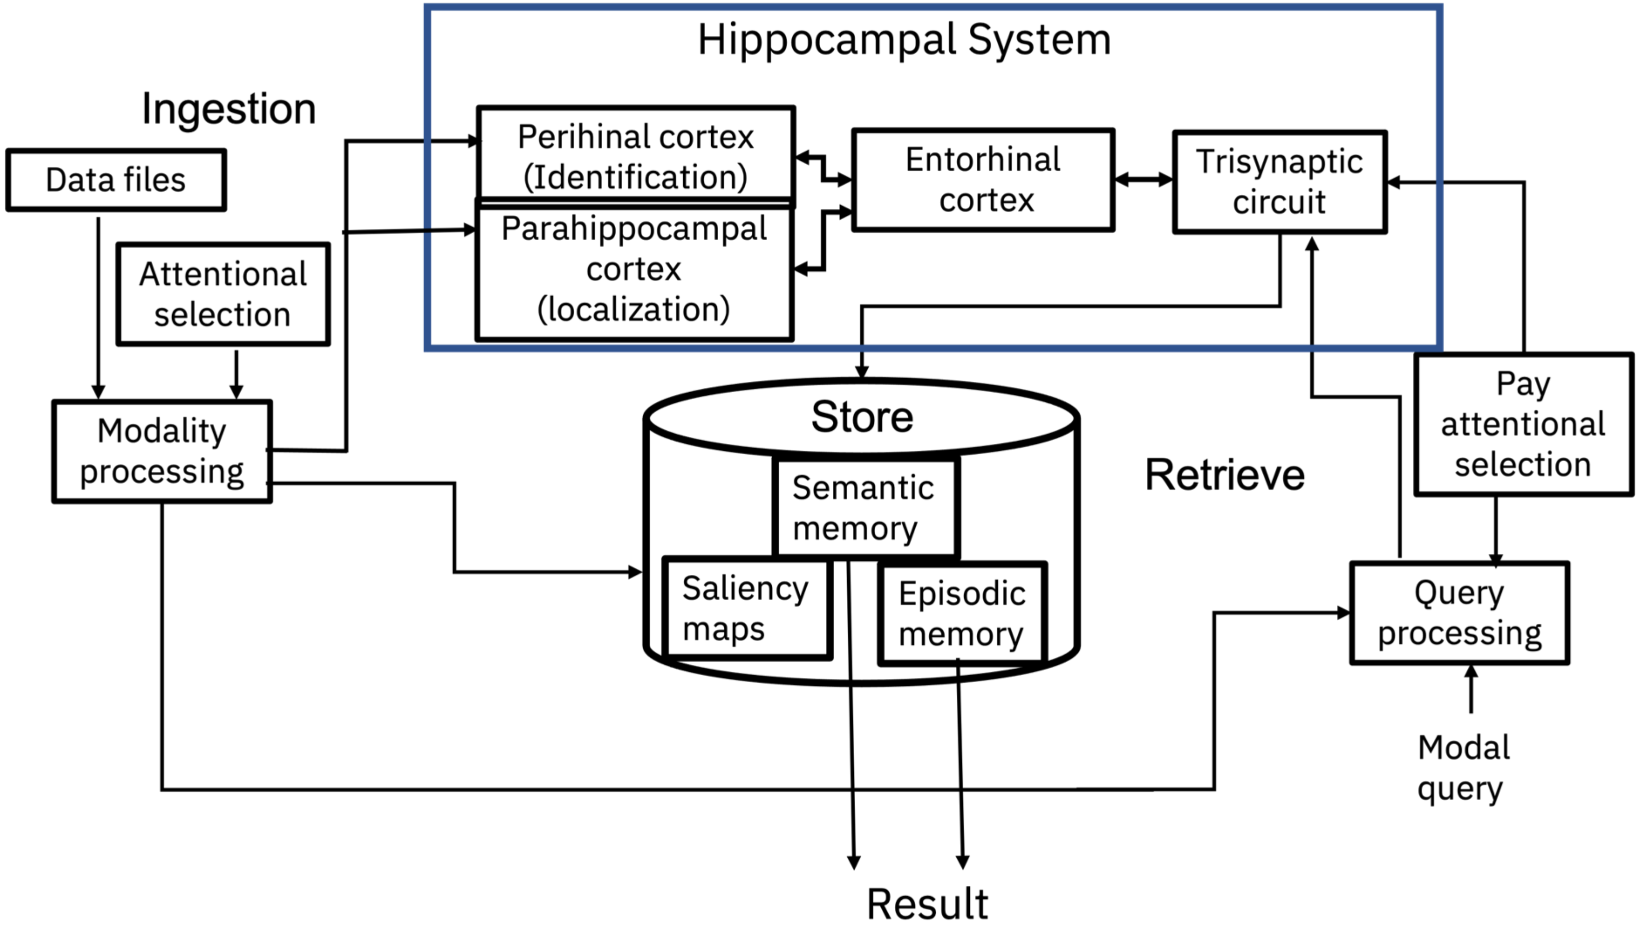

Neuro-Inspired Memories

- Why model declarative memory?
- A computational model of declarative memory
  - Modeling semantic memory
  - Modeling episodic memory
  - Modeling the associative memory formation in the hippocampus
- A commercial content-storage system modeled after human memory

Modeling semantic memory

- Computational models have so far considered :
  - Knowledge graphs
  - Semantic networks
  - Graph embeddings
- Are there any neural models of semantic memory that capture semantic similarity?
  - Only concepts similar in meaning are close to each other

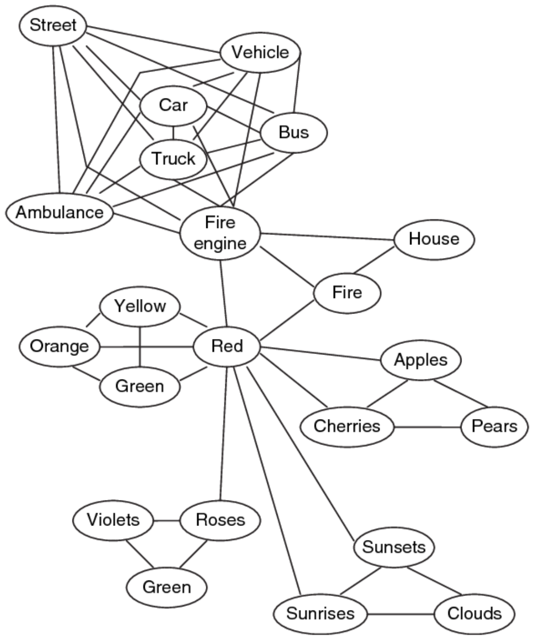

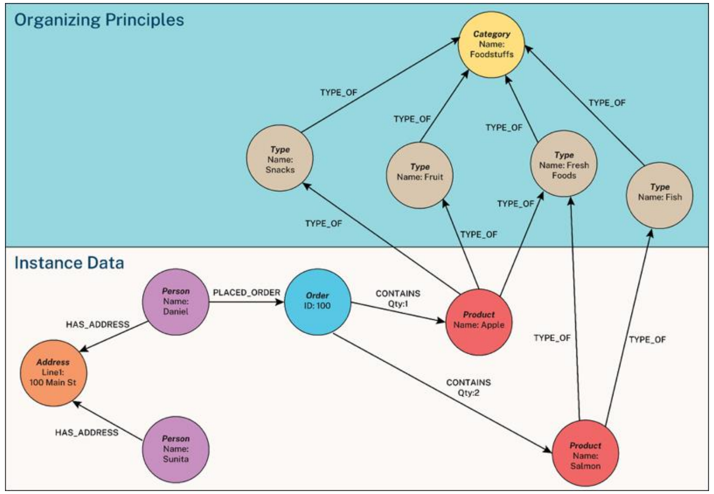

https://neo4j.com/blog/knowledge-graph/what-is-knowledge-graph/

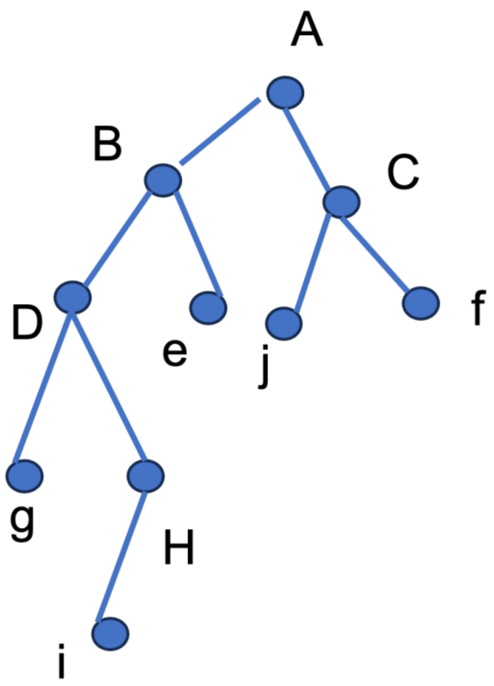

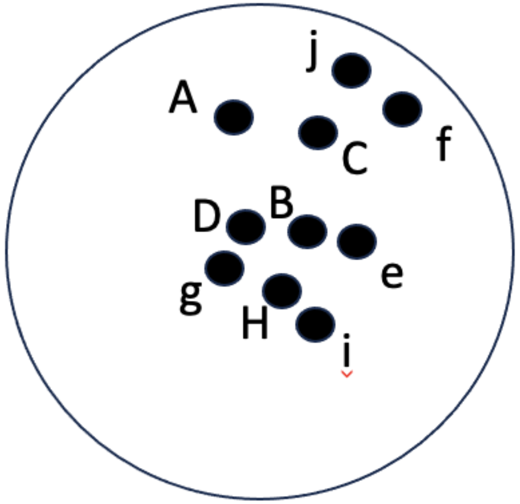

Neural modeling of domain-specific ontologies

- Existing approaches leverage textual information associated with nodes
- Can the structure alone be used to capture the semantic similarity?

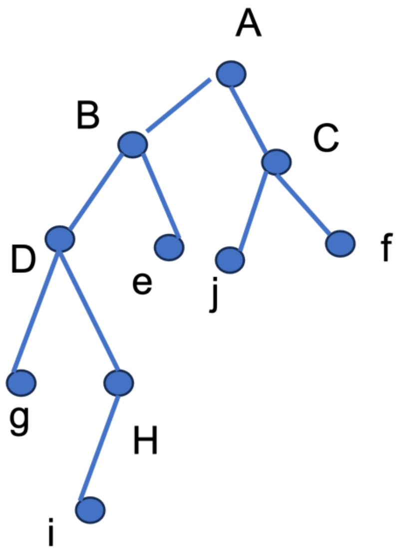

K	K Nearest Neighbors	Shortest Paths of lengh<K	Shortest Paths of lengh=K
1	D	NA	g->D
2	D,H,B	g->D	g->D->H,g->D->B
3	$D,H,B,A,$ i,e	g->D,g->D->H, g->D->B	g->D->B->A, g->D->H->i, g->D->B->e
4	$D,H,B,A,$ i,e,C	g->D,g->D->H, g->D->B,g->D->B->A, g->D->H->i,g->D->B->e	g->D->B->A->C
5	$D,H,B,A,$ i,e,C,j,f	g->D,g->D->H, g->D->B,g->D->B->A, g->D->H->i,g->D->B->e, g->D->B->A->C	g->D->B->A->C->i, g->D->B->A->C->f

ChatGPT

Hamper
Bin
Crate
Box
Carrier
Tote
Wicker basket
Chest
Pannier
Caddy

Words that are similar in sense

Nearest words	Score	Nearest words	Score
basket	0.99	team	1.0
handbasket	0.95	squad	0.98
creel	0.84	defending_team	0.91
frail	0.84	first_team	0.88
breadbasket	0.83	offence	0.87
skep	0.79	basketball_team	0.86
shopping_basket	0.77	A-team	0.85
hamper	0.77	powerhouse	0.85
bushel_basket	0.77	junior varsity	0.85
wicker_basket	0.75	minor-league_team	0.83

ChatGPT

Group
Squad
Crew
Tribe
Gang
Party
Band
Unit
Company
Force

Learning semantic memory embedding

- using supervised contrastive learning

$$\ell_{\text{contrast}}(S_i) = -\sum_{W_j \in S_i} \log \frac{\exp(z_i \cdot z_j/\tau)}{\sum_{a \in V} \exp(z_i \cdot z_a/\tau)}$$

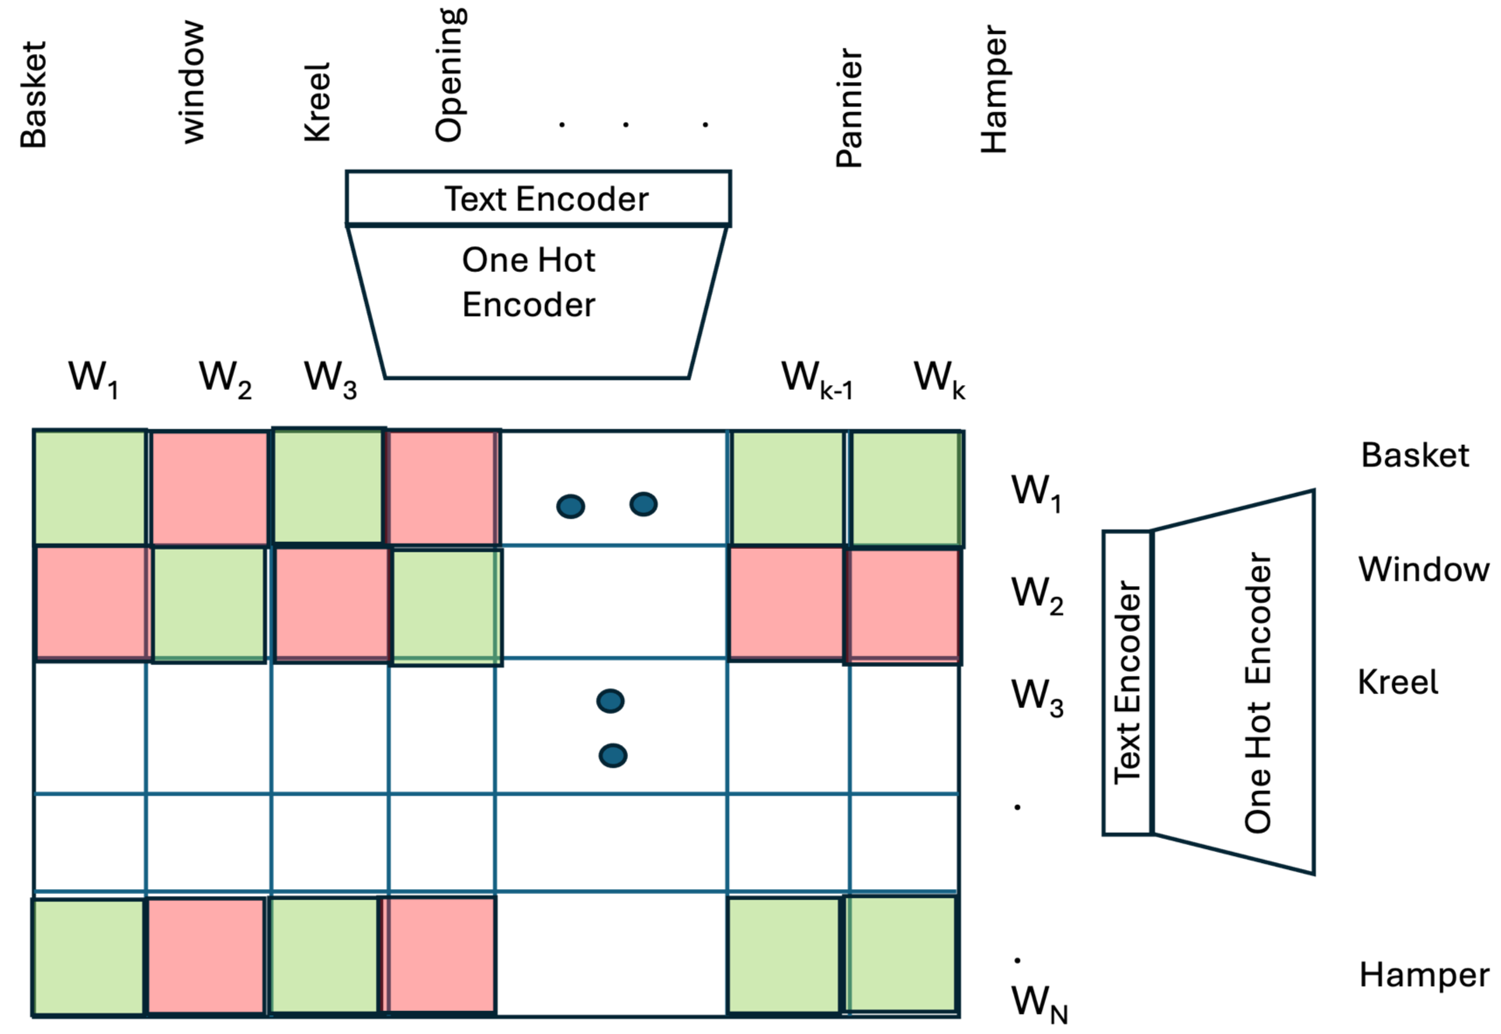

Learning associations in semantic memory

Query	**Search in Wordnet**	**Search in Semantic Embedding**
’hood	{’hood’, vicinity}	{’hood’, ’proximity’, ’gold_coast’, ’locality’, ’neighbourhood’, ’neck_of_the_woods’, ’neighborhood’, ’place’, ’section’, ’vicinity’}
abdominal cavity	{’abdominal_cavity’, ’cavity’, ’abdomen’}	{’pit_of_the_stomach’, ’orbital_cavity’, ’glenoid_cavity’, ’cavity’, ’axillary_fossa’, ’abdomen’, ’orbit’, ’cavum’, ’abdominal_cavity’, ’bodily_cavity’}
erosion	{’erosion’, ’ablation’}	{’erosion’, ’deflation’, ’wearing_away’, ’ablation’, ’detrition’, ’eroding’, ’abrasion’, ’attrition’, ’eating_away’, ’wearing’}
dorm room	{’dorm_room’, ’dormitory_room’, ’dormitory’, ’bedroom’}	{’dormitory_room’, ’chamber’, ’sleeping_accommodation’, ’master_bedroom’, ’dormitory’, ’dorm_room’, ’sleeping_room’, ’bedroom’, ’bedchamber’, ’guestroom’}

Domain-specific semantic memory models

- 13 benchmark datasets
- Semantic similarity is respected for datasets that contain antonyms

Spearman correlation coefficient

Datasets	Original Word #	WordNet Filtered #	Word2Vec	Glove	SBERT	Path2Vec	SME	SME (similar subset)
EM_SIMLEX_SYNS	297	297	0.285	0.240	0.145	0.301	0.265	**0.570**
EN-MC-30	30	30	**0.789**	0.702	0.410	**0.782**	0.650	0.650
EN-MEN-TR-3k	3000	2657	**0.776**	0.743	0.310	0.366	0.257	**0.780**
EN-MTurk-287	287	243	0.767	0.705	0.435	0.317	0.300	**0.810**
EN-MTurk-771	771	771	0.671	0.649	0.335	0.404	0.466	**0.760**
EN-RG-65	65	64	0.761	0.770	0.446	0.723	0.640	**0.820**
EN-RW-STANFORD	2034	910	0.492	0.341	0.226	0.194	0.217	**0.590**
EN-SIMLEX	666	666	0.452	0.397	0.233	0.505	0.398	**0.670**
EN-WS-353-REL	252	248	**0.626**	0.578	0.159	0.136	0	0
EN-WS-353-SIM	203	201	0.774	0.659	0.388	0.599	**0.820**	**0.820**
EN-YP-130	130	43	0.542	0.545	0.326	0.029	0.426	**0.660**
EW-WS-353-Syns	99	98	0.507	0.507	0.366	0.616	**0.655**	**0.655**
EN-WS-353-ALL	352	348	0.694	0.607	0.256	0.406	0.303	**0.720**

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)}$$

ρ = Spearman's rank correlation coefficient

d_i = difference between the two ranks of each observation

n = number of observations

Neuro-Inspired Memories

- Why model declarative memory?
- A computational model of declarative memory
    - Modeling semantic memory
    - Modeling episodic memory
    - Modeling the associative memory formation in the hippocampus
- A commercial content-storage system modeled after human memory

Modeling Episodic Memory

- Episodic memory is an autobiographical memory of facts and events
  - Remembers objects, shapes, colors, actions, the names, etc.
- Interacts with semantic memory during formation and recall
  - Relates instances to broader concepts
- Demonstrated through
  - Behavioral, lesion, and functional imaging studies
- Semantic memory located closer to the hippocampal regions
- Episodic memories distributed through the cortex
- What is known:
  - Semantic memory facilitates the acquisition of new episodic memories
  - Episodic memory facilitates the addition of new information to the semantic store.
  - episodic memories facilitate the retrieval of information from semantic memory
  - Semantic memories are the basic material from which complex and detailed episodic memories are constructed

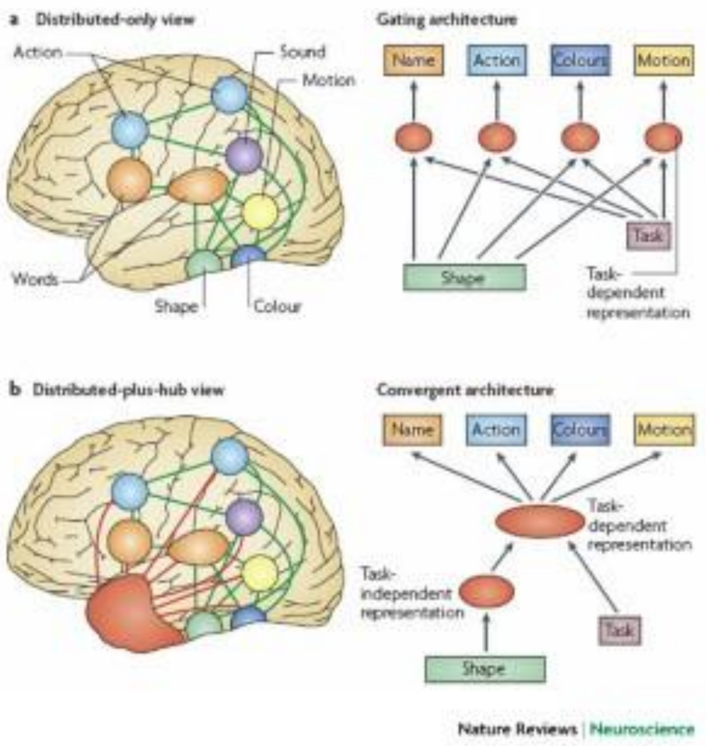

A vision-language model can be seen as an episodic memory

- Associates images with language

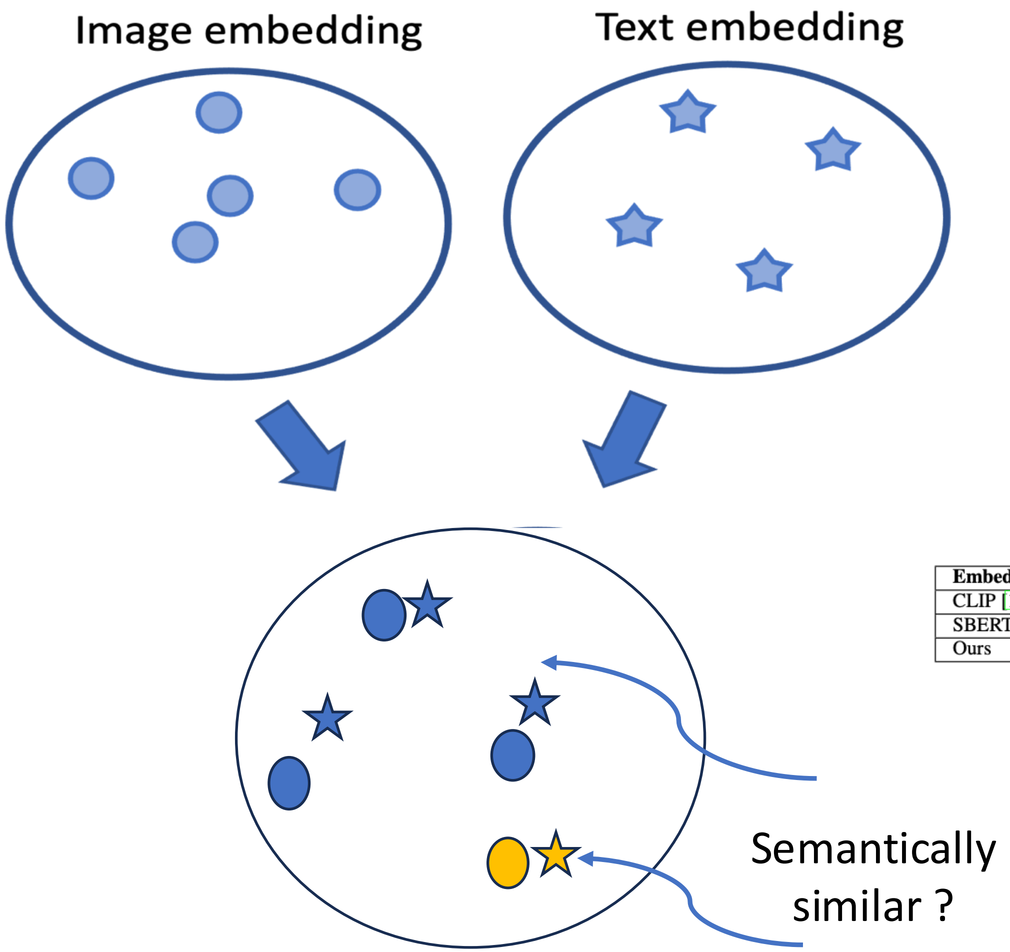

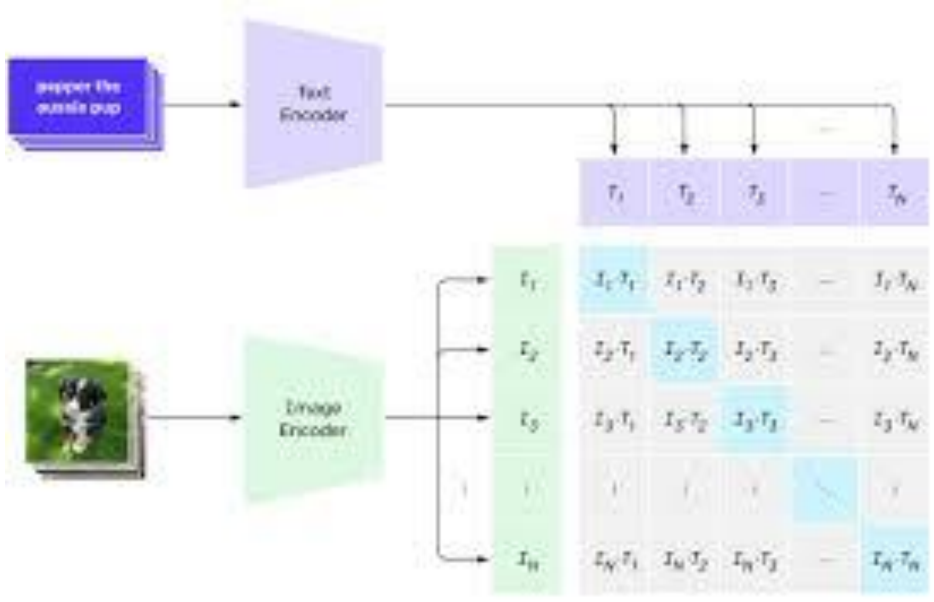

Table 1: Illustration of synonym recognition across text embeddings.

Embedding	# Queries	Synonyms in Top10	%age synonyms covered
CLIP [15]	71895	28070	49.27%
SBERT [17]	71895	37888	52.7%
Ours	71895	67309	87.7%

Need for semantic memory associations

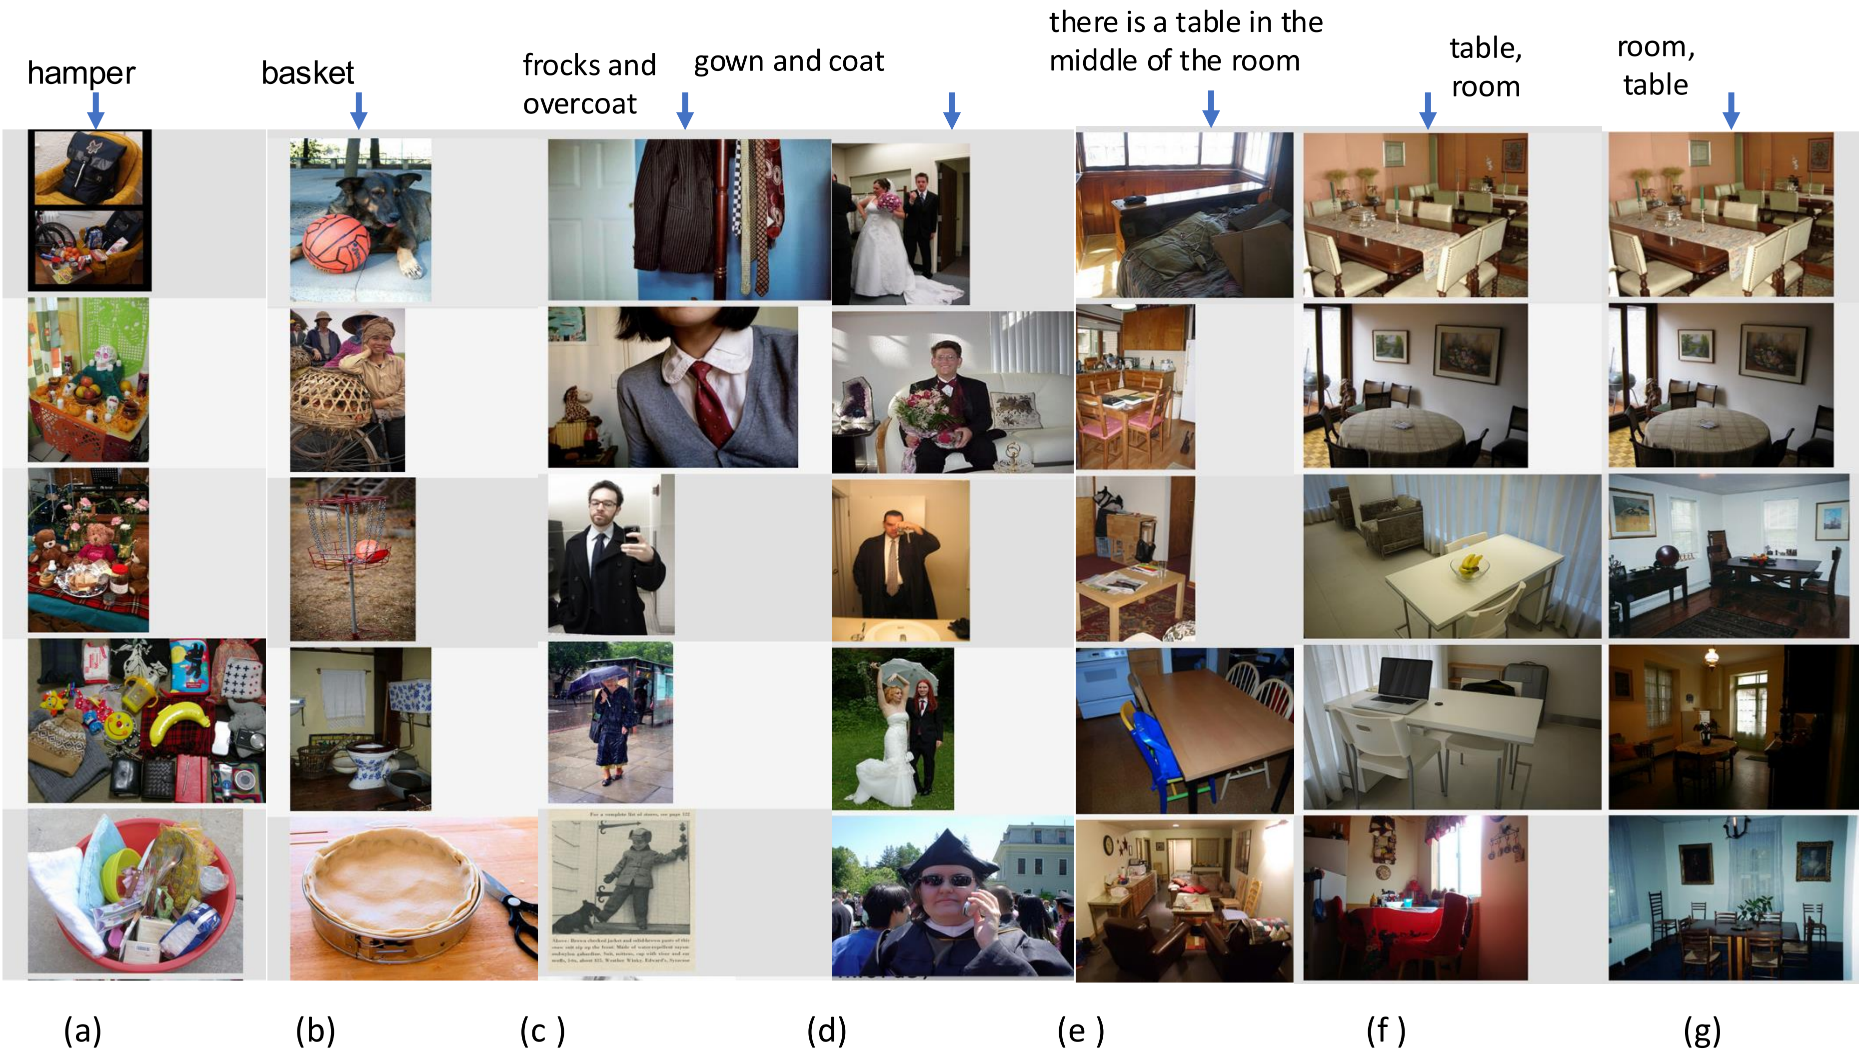

- Recall sensitive to exact query
- Semantics not well understood
- Order seems to matter

Aligning VLM with Semantic Memory

Project VLM onto the semantic memory embedding space

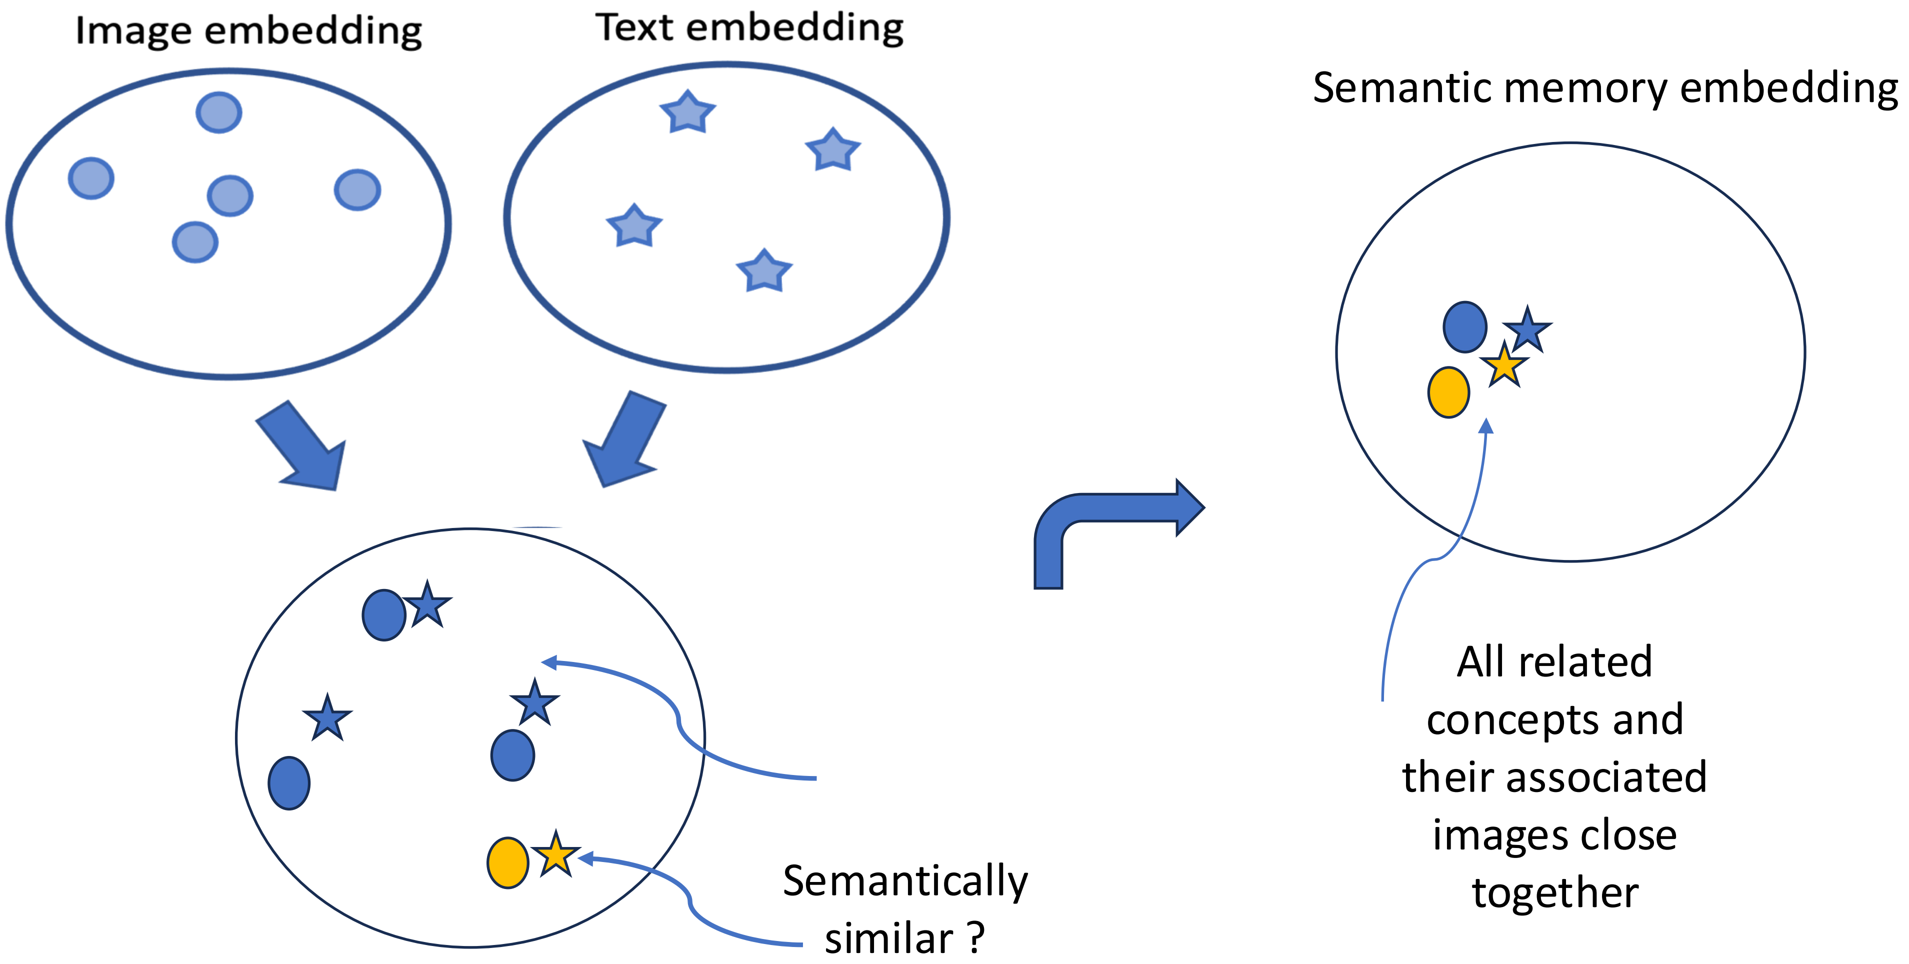

SemCLIP – A new episodic memory model built on semantic memory substrate

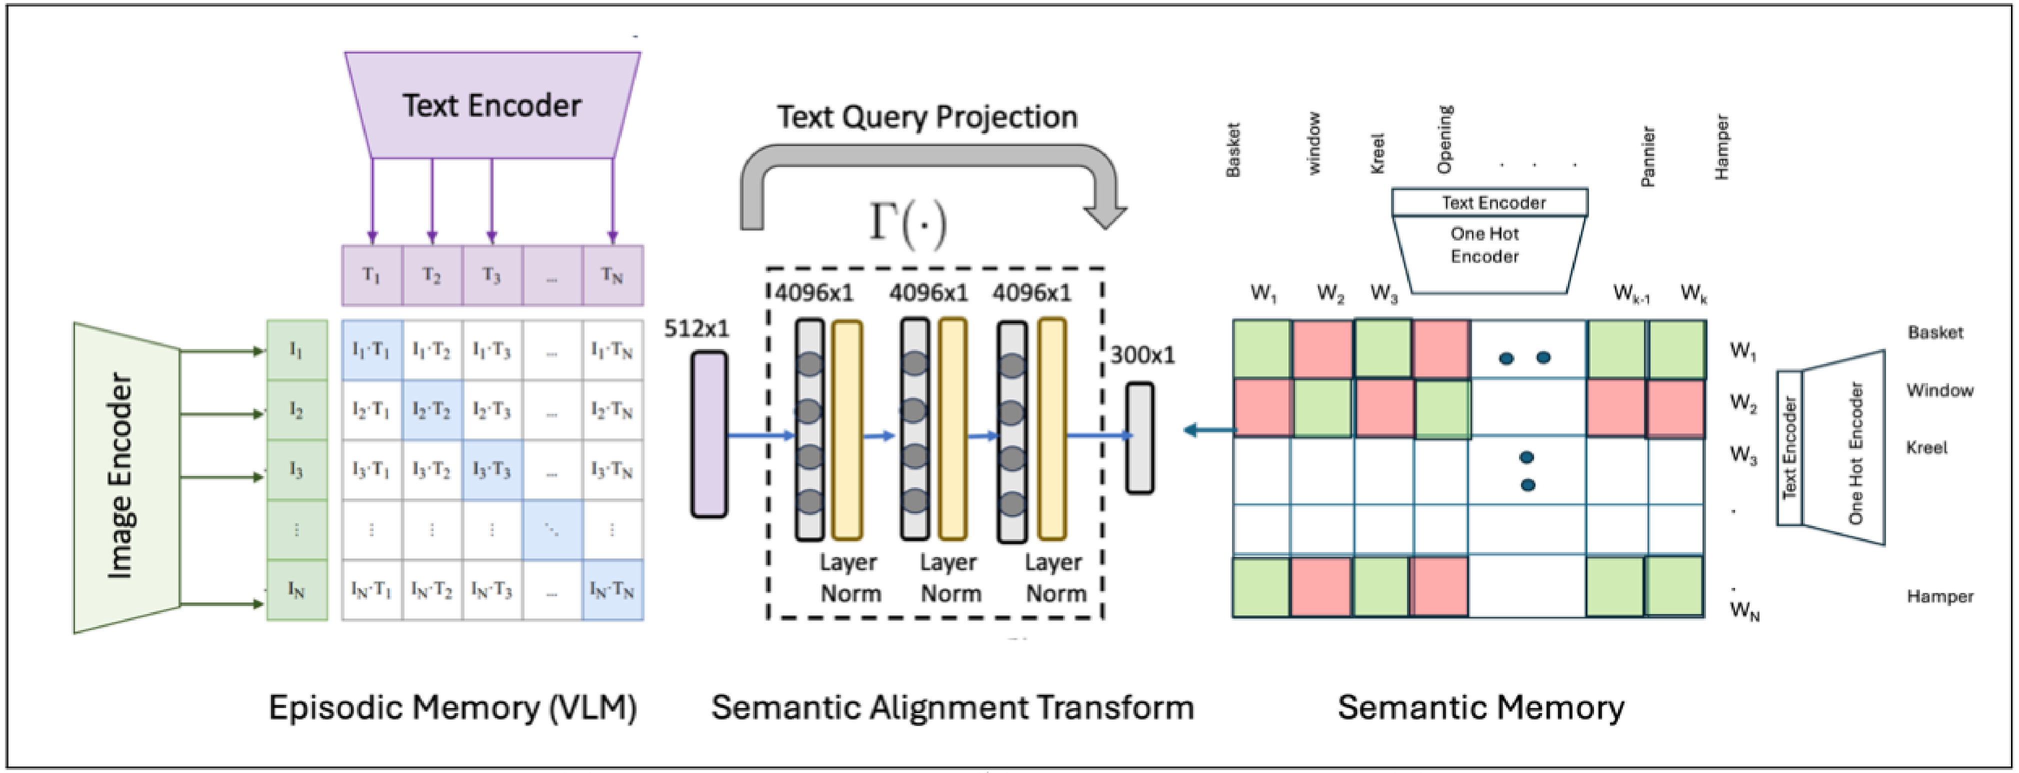

(a) (b) (c)

$$\ell(i,j) = -\log \frac{\exp(s_{i,j}/\tau)}{\sum_{k=1}^{2N} \mathbb{1}_{[k\neq i]} \exp(s_{i,k}/\tau)}$$

Transform: $\Gamma_t(\cdot) = \mathbf{FC}_3(\mathbf{\Phi}_{\text{relu}}(\mathbf{FC}_2(\mathbf{\Phi}_{\text{relu}}(\mathbf{FC}_1((\cdot))))))$ Loss: $\mathcal{L} = ||\Gamma_t(C_t) - \mathbf{C}_t^{'}||_2^2$

$$\ell_{\text{contrast}}(S_i) = -\sum_{W_j \in S_i} \log \frac{\exp(z_i \cdot z_j/\tau)}{\sum_{a \in V} \exp(z_i \cdot z_a/\tau)}$$

Comparing episodic memory model to VLMs

Dataset	Images/Queries	Method	Overlap@1	Overlap@5	Overlap@10	Overlap@50
Visual Genome	7794 / 14513	**SemCLIP**	**0.551**	**0.532**	**0.517**	**0.523**
		CLIP	0.119	0.056	0.038	0.026
		OpenCLIP	0.129	0.055	0.040	0.025
		BLIP	0.119	0.056	0.037	0.024
		NegCLIP	0.109	0.063	0.038	0.024
CUB	11788 / 200	**SemCLIP**	**0.812**	**0.783**	**0.732**	**0.715**
		CLIP	0.118	0.072	0.062	0.053
		OpenCLIP	0.149	0.079	0.062	0.053
		BLIP	0.181	0.084	0.066	0.053
		NegCLIP	0.119	0.078	0.066	0.055
SUN	16657 / 567	**SemCLIP**	**0.554**	**0.531**	**0.523**	**0.511**
		CLIP	0.092	0.048	0.034	0.025
		OpenCLIP	0.070	0.039	0.028	0.021
		BLIP	0.091	0.05	0.035	0.025
		NegCLIP	0.095	0.045	0.032	0.023
AWA2	6985 / 10	**SemCLIP**	**0.751**	**0.723**	**0.702**	**0.715**
		CLIP	0.086	0.051	0.038	0.028
		OpenCLIP	0.139	0.067	0.046	0.032
		BLIP	0.139	0.057	0.039	0.028
		NegCLIP	0.101	0.046	0.034	0.025

An episodic memory model is more semantically stable

An episodic memory model has higher retrieval quality

Dataset	Images / Labels	Model	t2i: NDCG / mAP / Recall (@10)	i2t: NDCG / mAP / Recall (@10)
Visual Genome	7794 / 14513	**SemCLIP**	**0.192 / 0.172**	**0.254 / 0.185**
		CLIP	0.053 / 0.060 / 0.065	0.050 / 0.129 / 0.028
		OpenCLIP	0.066 / 0.075 / 0.081	0.061 / 0.159 / 0.035
		BLIP	0.072 / 0.081 / 0.088	0.069 / 0.167 / 0.039
		NegCLIP	0.063 / 0.072 / 0.077	0.060 / 0.150 / 0.035
		PosCLIP	0.074 / 0.084 / 0.091	0.060 / 0.146 / 0.034
CUB	11788 / 200	**SemCLIP**	**0.721 / 0.845**	**0.891 / 0.812**
		CLIP	0.513 / 0.621 / 0.084	0.619 / 0.554 / 0.826
		OpenCLIP	0.669 / 0.744 / 0.111	0.777 / 0.726 / 0.935
		BLIP	0.204 / 0.341 / 0.033	0.326 / 0.260 / 0.543
		NegCLIP	0.406 / 0.535 / 0.065	0.472 / 0.403 / 0.694
		PosCLIP	0.488 / 0.580 / 0.080	0.553 / 0.479 / 0.788
SUN	16657 / 567	**SemCLIP**	**0.686 / 0.712**	**0.810 / 0.671**
		CLIP	0.414 / 0.562 / 0.191	0.458 / 0.415 / 0.595
		OpenCLIP	0.549 / 0.664 / 0.260	0.514 / 0.476 / 0.634
		BLIP	0.413 / 0.535 / 0.194	0.426 / 0.384 / 0.557
		NegCLIP	0.429 / 0.553 / 0.201	0.425 / 0.380 / 0.567
		PosCLIP	0.463 / 0.602 / 0.214	0.445 / 0.399 / 0.589
AWA2	6985 / 10	**SemCLIP**	0.967 / 0.987	0.995 / 0.989
		CLIP	0.993 / 0.999 / 0.016	0.992 / 0.989 / 1.000
		OpenCLIP	**1.000 / 1.000** / 0.016	**0.994 / 0.991** / 1.000
		BLIP	**1.000 / 1.000** / 0.016	0.991 / 0.988 / 1.000
		NegCLIP	**1.000 / 1.000** / 0.016	0.987 / 0.982 / 1.000
		PosCLIP	**1.000 / 1.000** / 0.016	0.983 / 0.977 / 1.000
COCO	5000 / 80	**SemCLIP**	**0.940** / 0.91	**0.895 / 0.923**
		CLIP	0.810 / 0.869 / 0.081	0.706 / 0.771 / 0.745
		OpenCLIP	0.866 / 0.926 / 0.087	0.756 / 0.788 / 0.799
		BLIP	0.897 / 0.942 / 0.089	0.774 / 0.852 / 0.781
		NegCLIP	0.834 / 0.892 / 0.083	0.766 / 0.797 / 0.808
		PosCLIP	0.919 / 0.946 / 0.088	0.807 / 0.834 / 0.843
Flickr30k	31014 / 158391	**SemCLIP**	**0.571 / 0.523**	**0.580 / 0.572**
		CLIP	0.354 / 0.306 / 0.510	0.314 / 0.430 / 0.320
		OpenCLIP	0.427 / 0.377 / 0.588	0.376 / 0.489 / 0.382
		BLIP	0.562 / 0.510 / 0.725	0.475 / 0.587 / 0.480
		NegCLIP	0.425 / 0.373 / 0.591	0.354 / 0.465 / 0.360
		PosCLIP	0.323 / 0.279 / 0.468	0.273 / 0.367 / 0.283

Neuro-Inspired Memories

- Why model declarative memory?
- A computational model of declarative memory
  - Modeling semantic memory
  - Modeling episodic memory
  - Modeling the associative memory formation in the hippocampus
- A commercial content-storage system modeled after human memory

Modeling the associative learning in the hippocampus

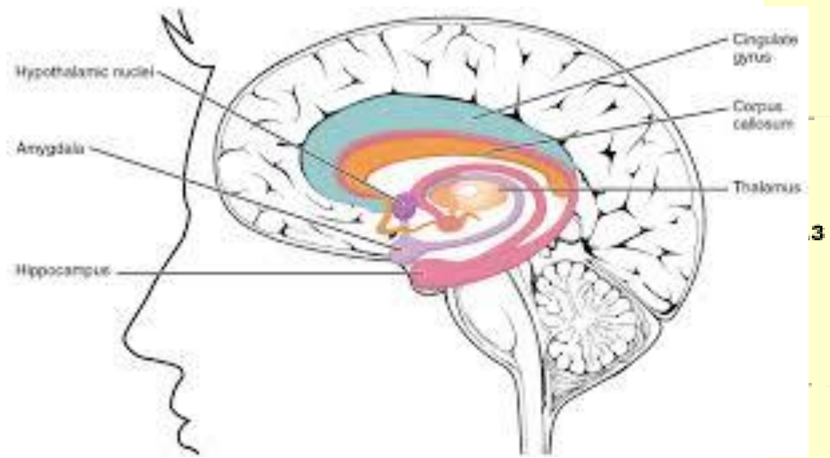

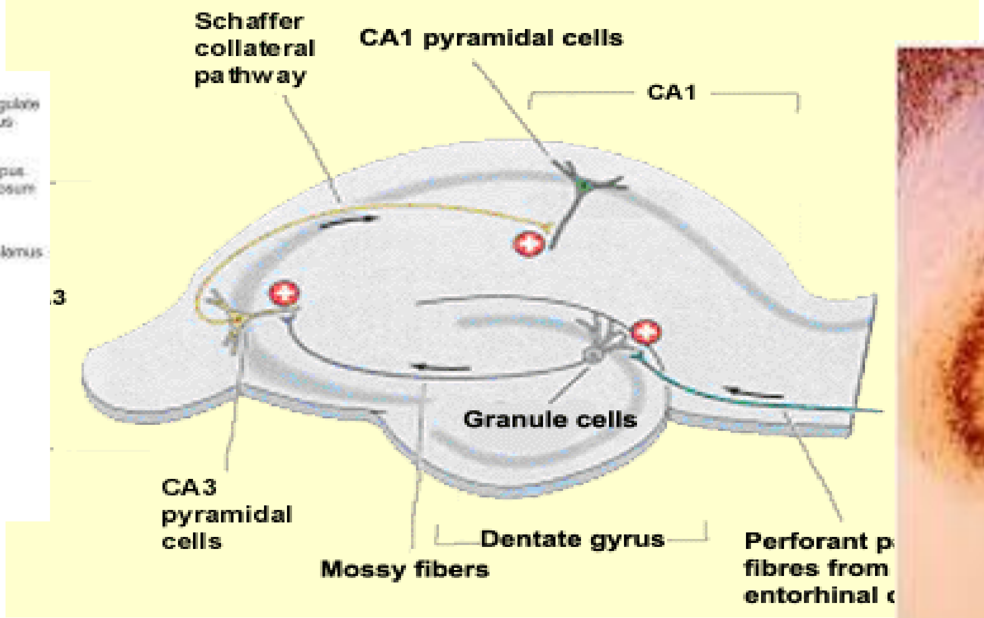

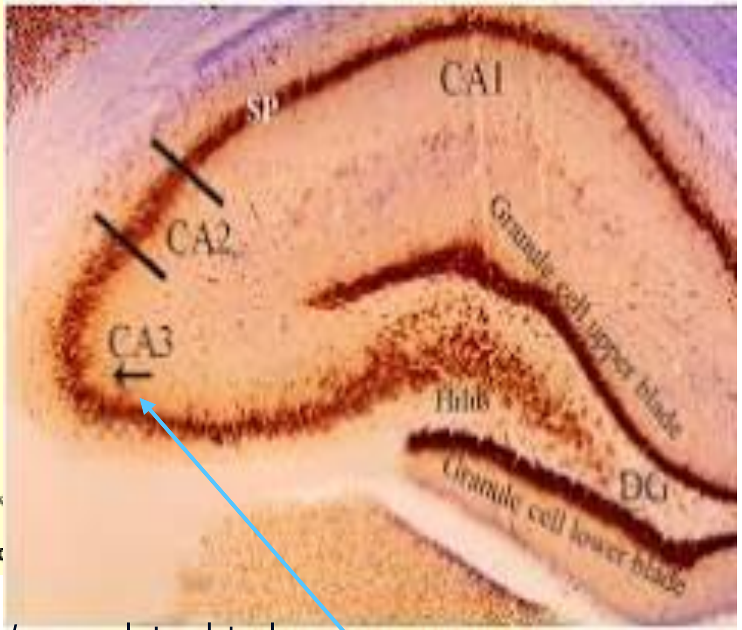

https://thebrain.mcgill.ca/flash/a/a_07/a_07_cl/a_07_cl_tra/a_07_cl_tra.html

Dentate Gyrus (DG):
•Distinguishes between similar experiences, a process known as "pattern separation".
Cornu Ammonis (CA):
•**CA3:** Contains "attractor networks" that help retrieve complete memories from partial cues.
CA1: Receives input from CA3 and the entorhinal cortex projects to the subiculum to send information back out of the hippocampus.

Subiculum:
•The final stage of the hippocampal circuit, it combines information from CA1 and the entorhinal cortex to project to other brain areas, helping to consolidate and retrieve memories.

Hopfield Networks

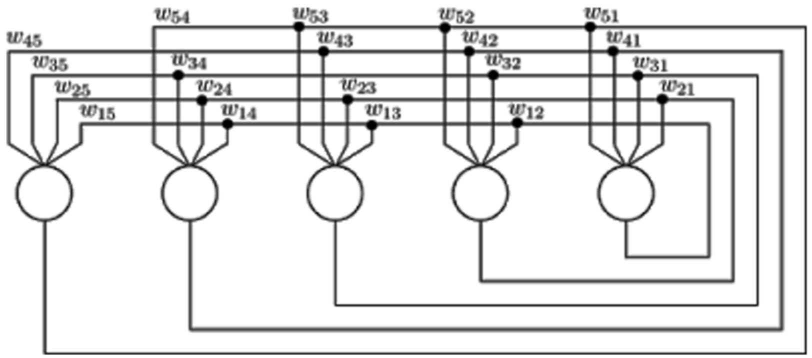

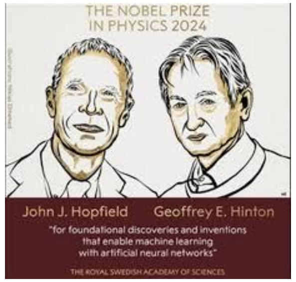

- CA3 cells in the Hippocampus have been known to perform Hebbian learning.
- Content-addressable memory systems store data as attractor states in an energy landscape.
- Can hold only a limited number of memories

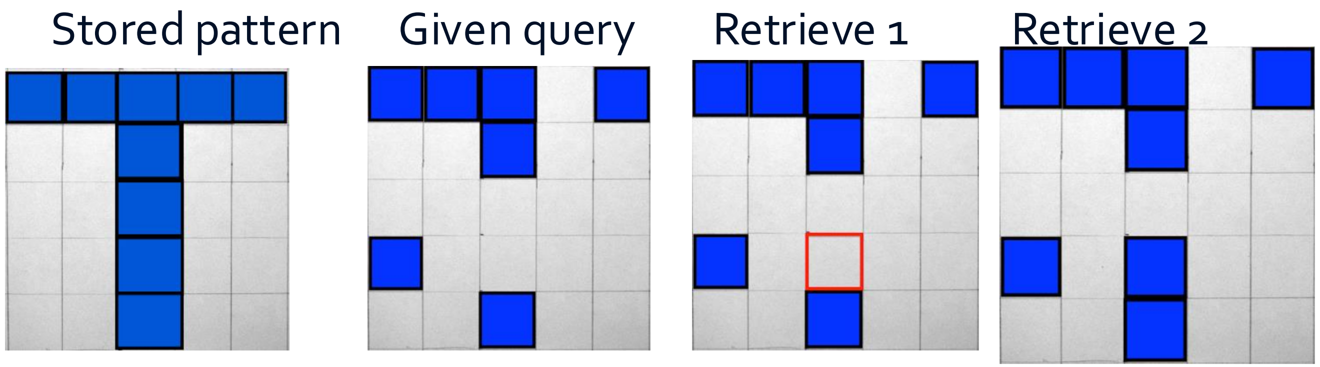

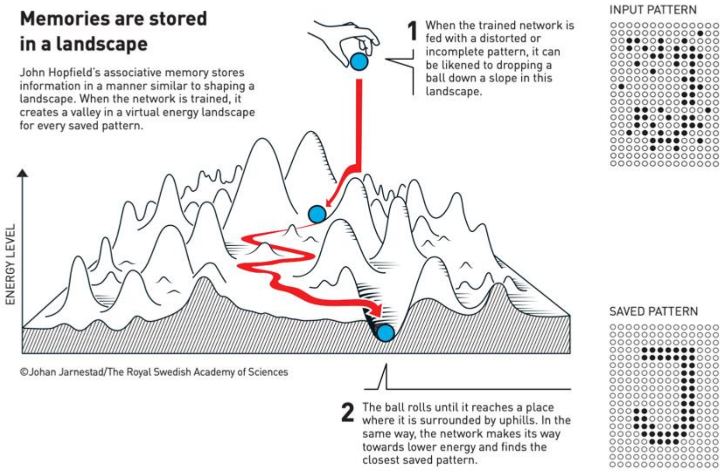

$$E = -\frac{1}{2}\sum_i a_i\left(\sum_{j\neq i} w_{ij}a_j\right)$$

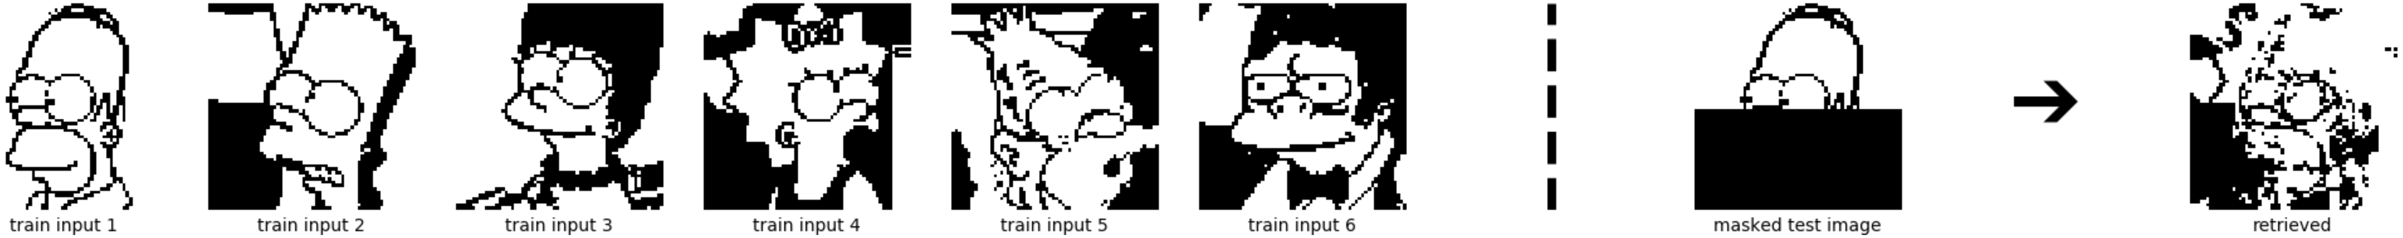

Modern Hopfield Networks

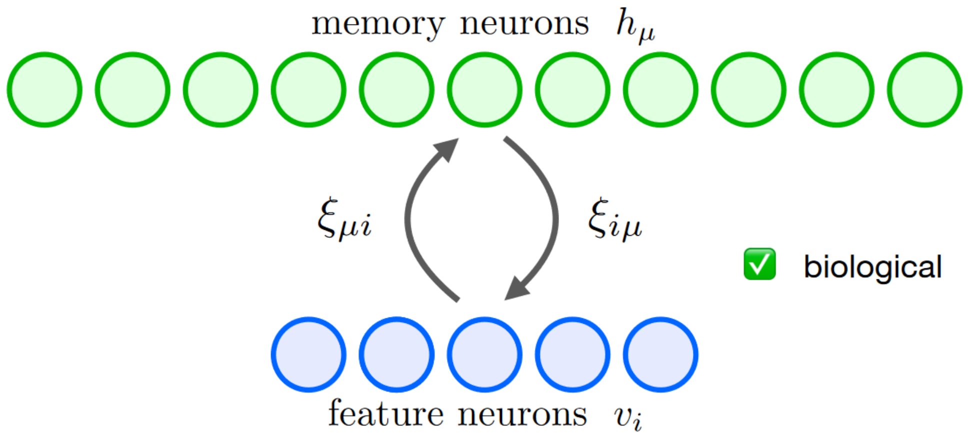

✅ biological

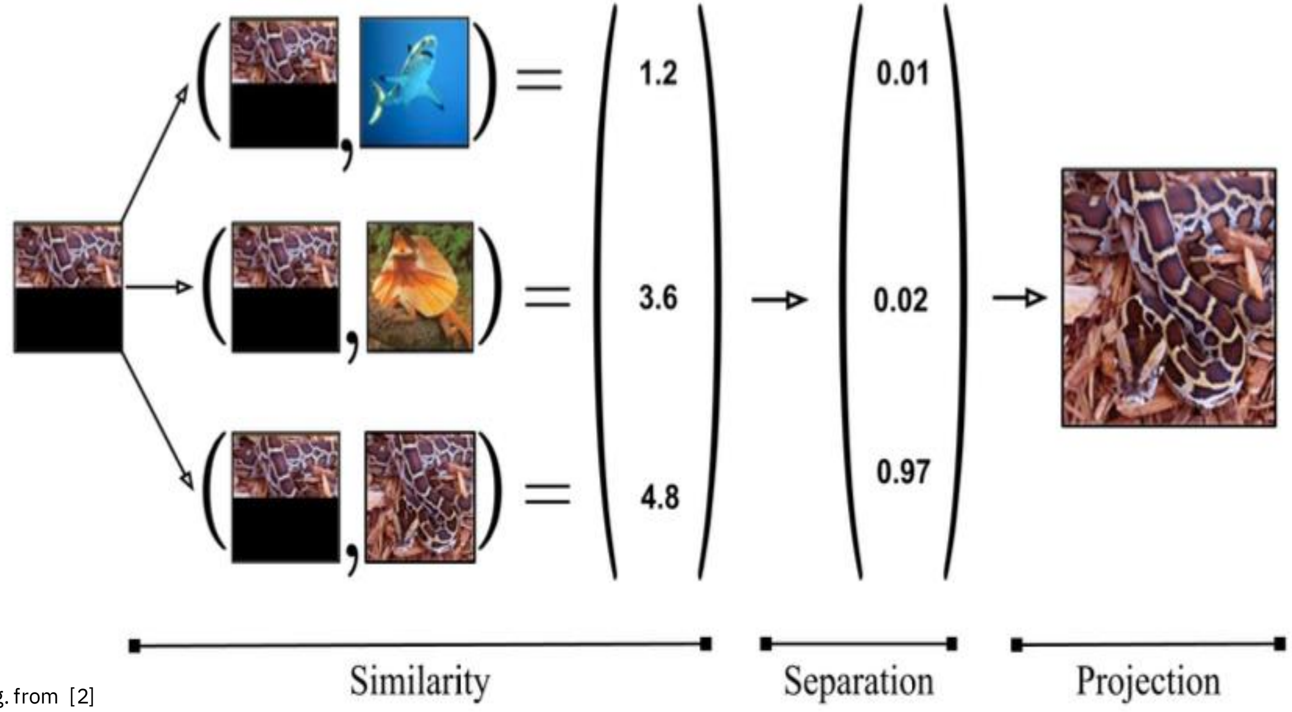

Fig. from [2]

- Modern Hopfield Networks address the limited storage capacity by using am expontential energy function

$$\mathrm{E} = -\mathrm{lse}\left(\beta, \boldsymbol{X}^T\boldsymbol{\xi}\right) + \frac{1}{2}\boldsymbol{\xi}^T\boldsymbol{\xi} + \beta^{-1}\log N + \frac{1}{2}M^2 \,,$$

$$\boxed{s = \underbrace{P}_{\text{Projection}} \cdot \underbrace{\text{sep}}_{\text{Separation}} \underbrace{(\text{sim}(M, q))}_{\text{Similarity}}}$$

$$\boxed{\mathbf{s}^{(t+1)} = \boldsymbol{\Xi}\,\text{softmax}(\beta\boldsymbol{\Xi}^T\mathbf{s}^{(t)})}$$

- Reduces to pattern matching
- As many neurons as the patterns!
- Suffers from metastable states
- Requires partial patterns to be specified

[1] Krotov D, Hopfield JJ. Large Associative Memory Problem in Neurobiology and Machine Learning.
[2] Millidge, Beren, et al. "Universal hopfield networks: A general framework for single-shot associative memory models."

Our key ideas

- Enhanced separability through latent representations can reduce metastable states.
  - Evidence from prior encoding in DG cells before CA3 in the tri-synaptic circuit suggests this as a plausible mechanism.
- Partial content specification can be avoided using a language association.
  - Integration of language cue with content in the ET cortex prior to entering the hippocampus suggests hetero-association.

Encoding Hopfield Networks

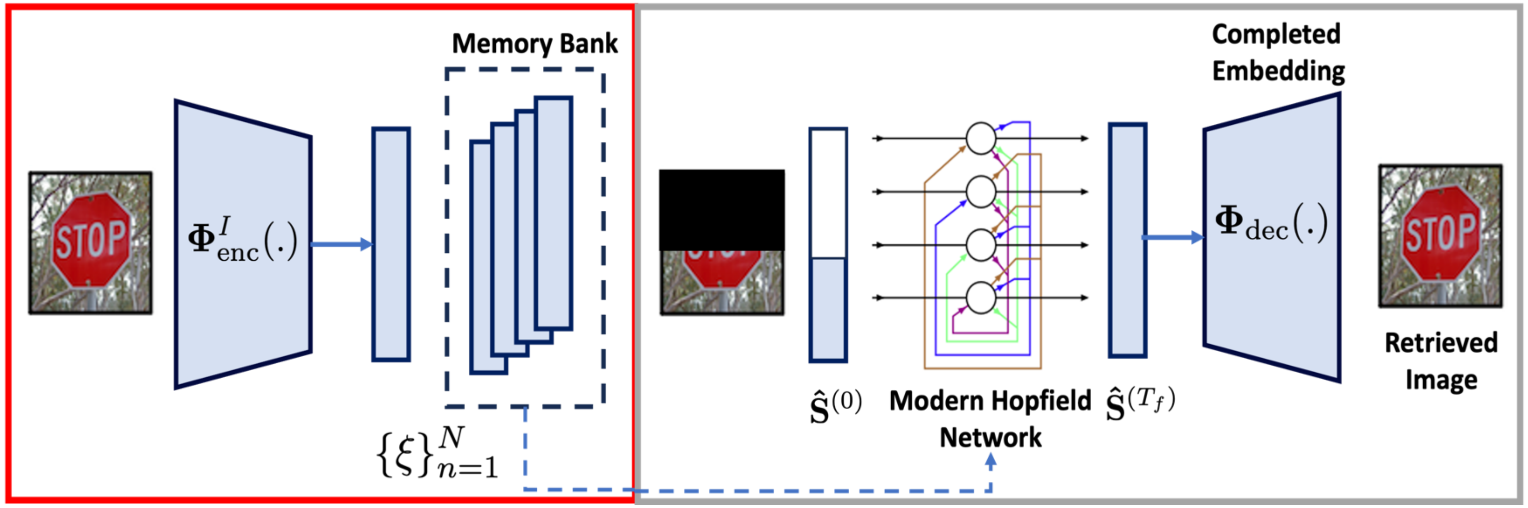

MHN State update:

$s^{(t+1)} = \mathrm{softmax}(\beta \Xi^{T} s^{(t)})$

HEN latent transformation:

$s^{(t+1)} = \mathrm{softmax}(\beta \Phi_{enc}(\Xi)^{T} \Phi_{enc}(s^{(t)}))$.

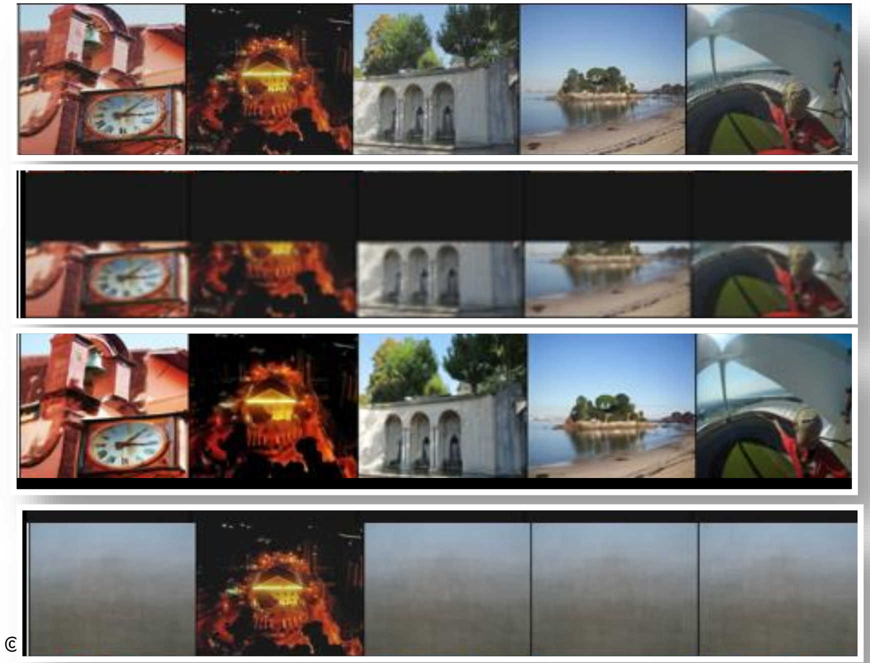

Memorized images

Retrieval cues

HEN - Recovered from memory

MHN- Recovered from memory

– Would any encoding work?

  – Separability in the embedding space is important

Separability of various embeddings in Hopfield Encoding Network

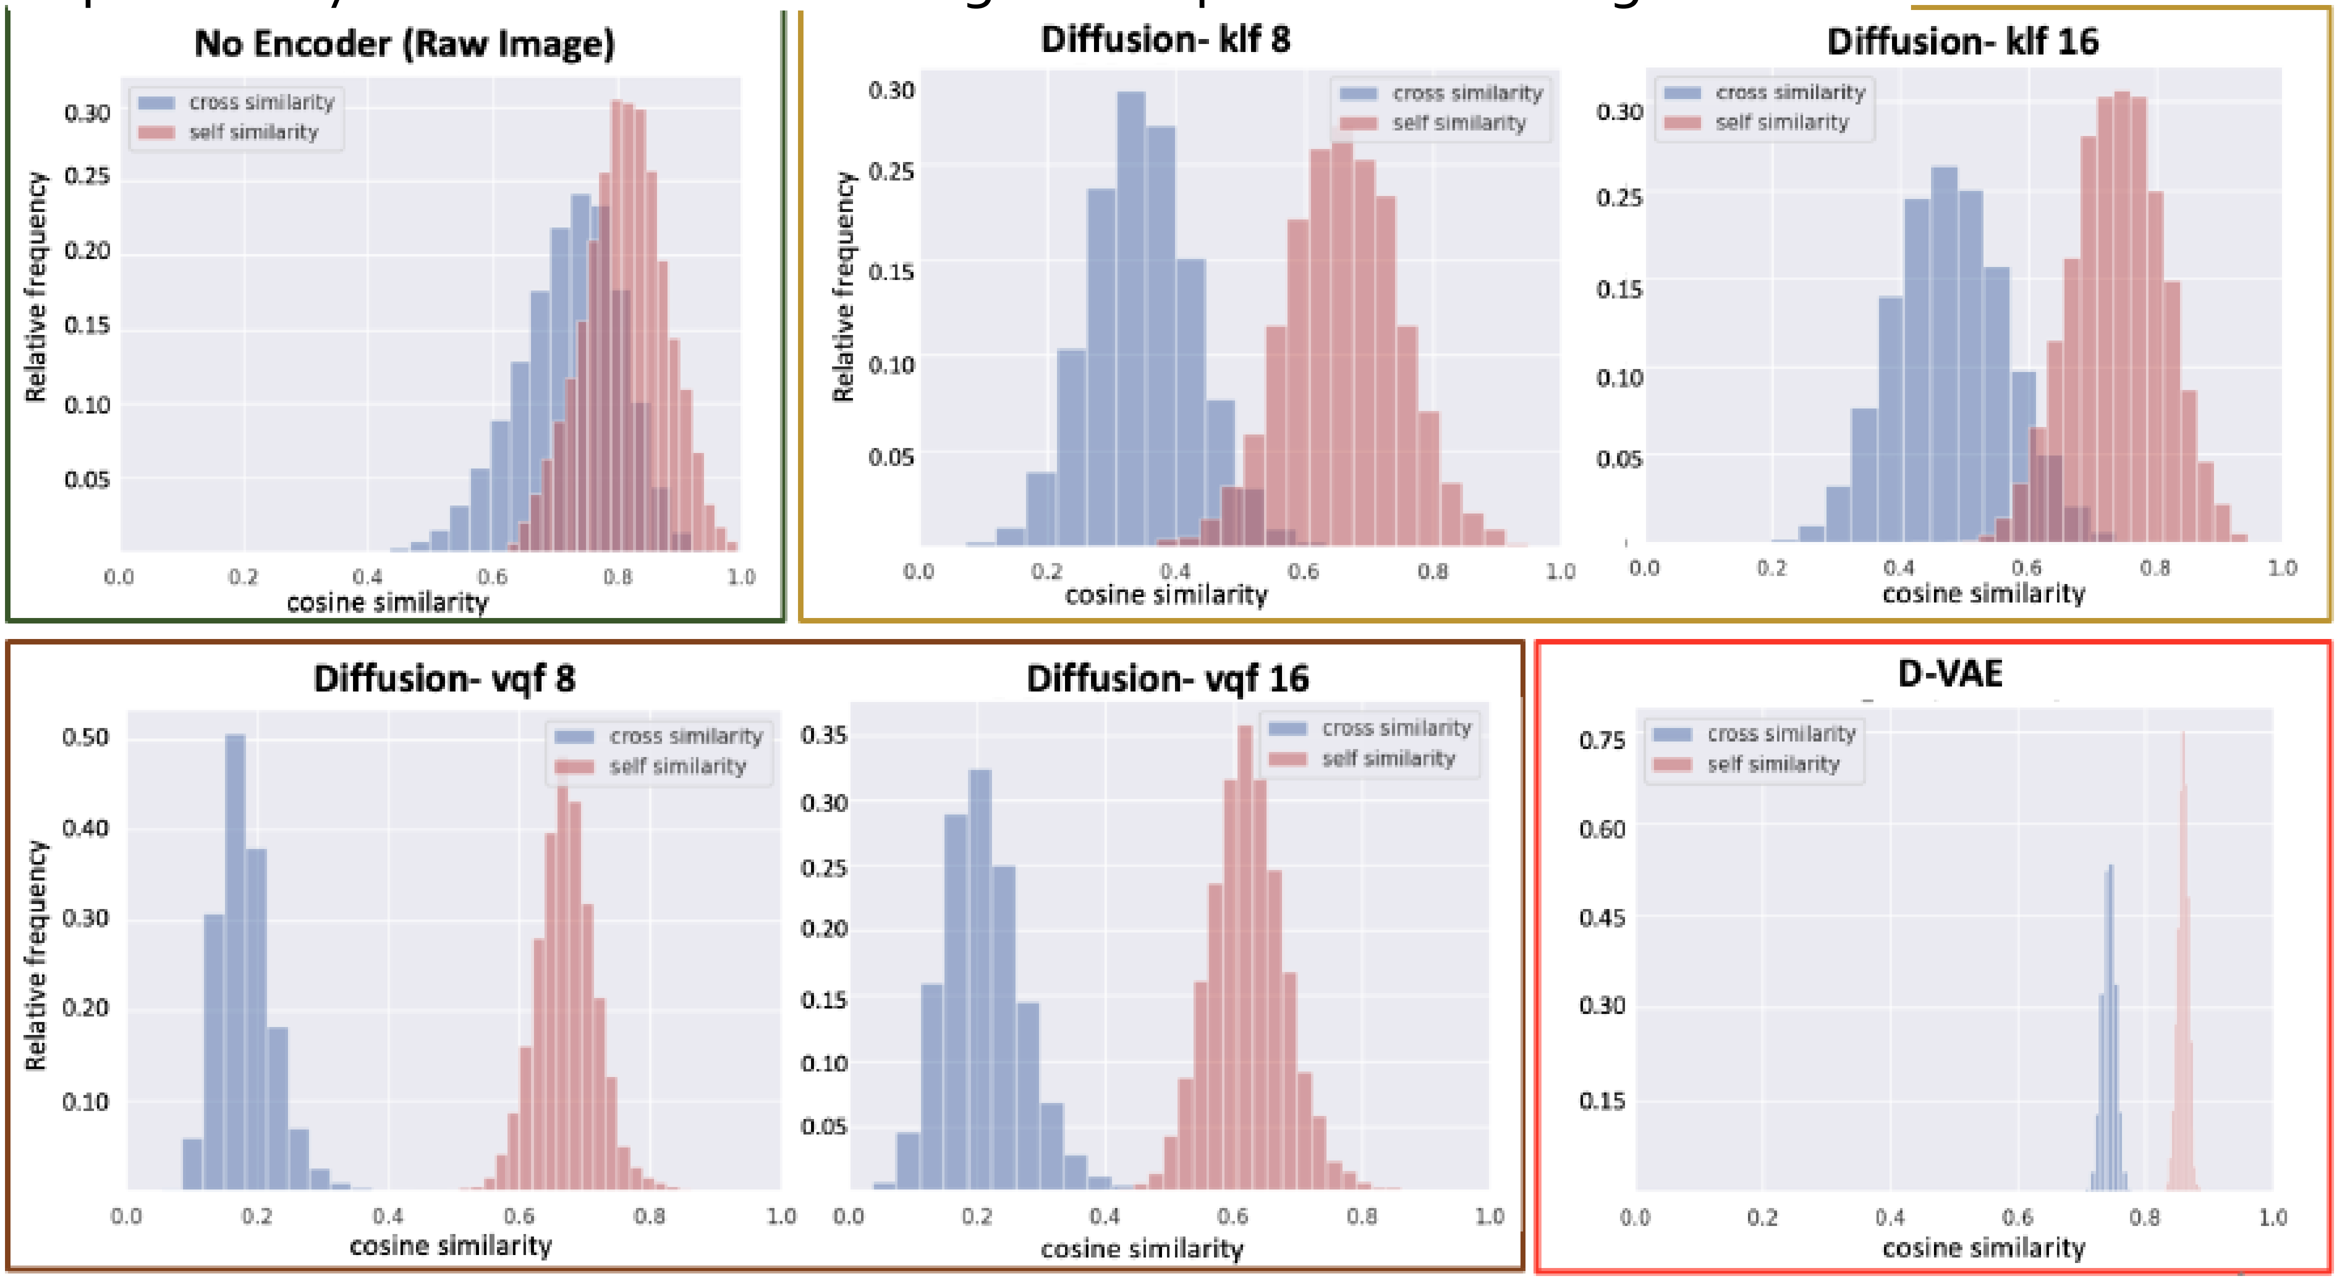

Reduce spurious attractors

Similarity Metrics: dot product and L2 losses
Pretrained Encoders: Diffusion models[1] – vqf8, vqf16, discrete VAE[2], Image
Eval Metrics: MSE, 1-SSIM

$$s = \underbrace{P}_{\text{Projection}} \cdot \underbrace{\text{sep}}_{\text{Separation}} \underbrace{(\text{sim}(M, q))}_{\text{Similarity}}$$

β Ablation MS-COCO[3] 1-SSIM, N=6000

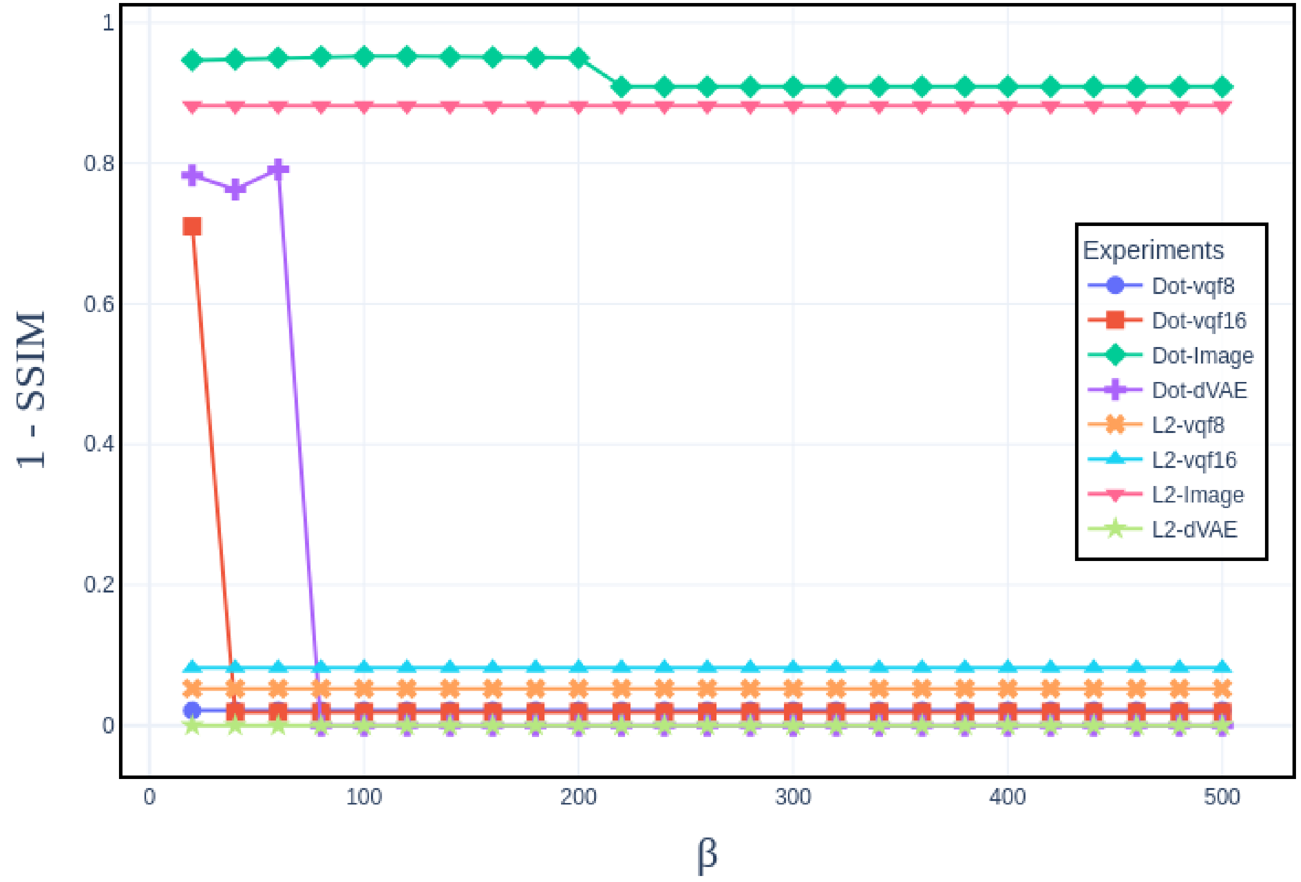

In NeurIPS UniReps Workshop 2024

MS-COCO[3] with β = 150, Scale Up

Table 1. All encoder-based methods can recover image representations without quality loss, even as more images are stored in $\{\xi_n\}$.

NUM IMAGES 1-SSIM, MSE	vqf8	vqf16	Image	D-VAE
6000	0.021, 0.000	0.019, 0.004	0.836, 0.064	**0.000, 0.000**
8000	0.019, 0.000	0.046, 0.004	0.835, 0.067	**0.000, 0.000**
10000	0.019, 0.000	0.047, 0.004	0.835, 0.064	**0.000, 0.000**
15000	0.019, 0.000	0.048, 0.004	0.836, 0.066	**0.000, 0.000**

[1] Taming Transformers for High-Resolution Image Synthesis
[2] Dall-E: https://github.com/openai/DALL-E
[3] Lin, Tsung-Yi, et al. "Microsoft coco: Common objects in context." Computer Vision–ECCV 2014

Hopfield encoding network results

Dense associative memory bank

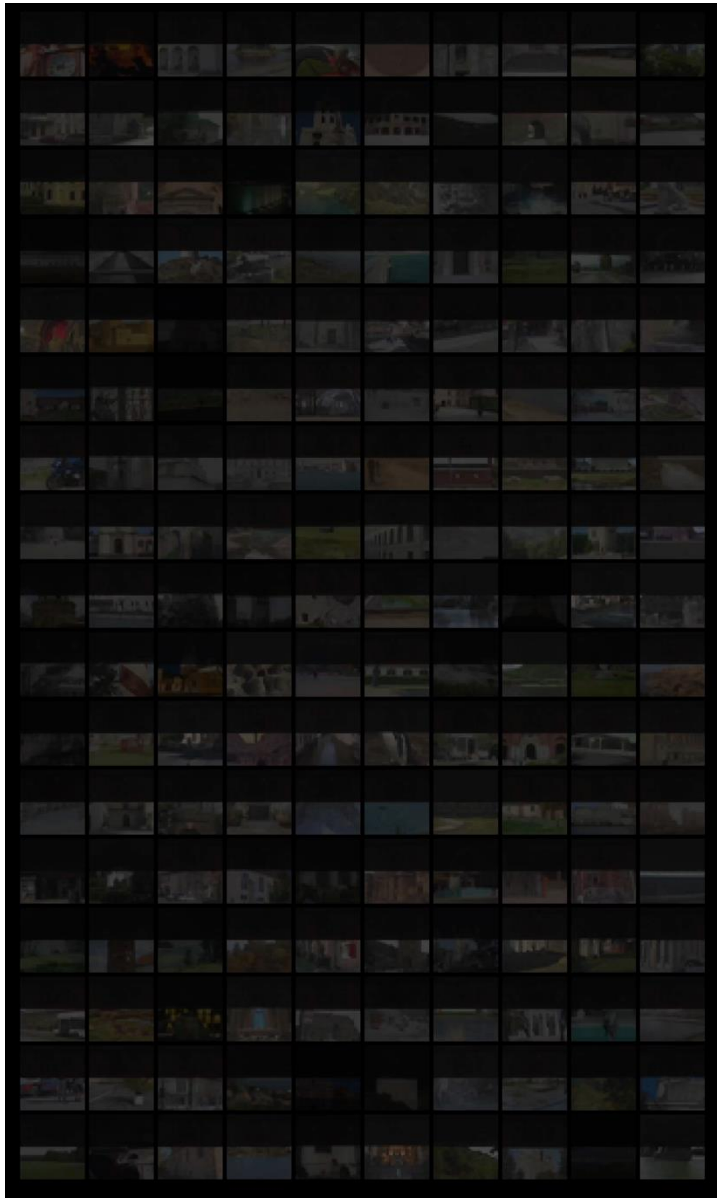

Modern Hopfield Network Retrieval

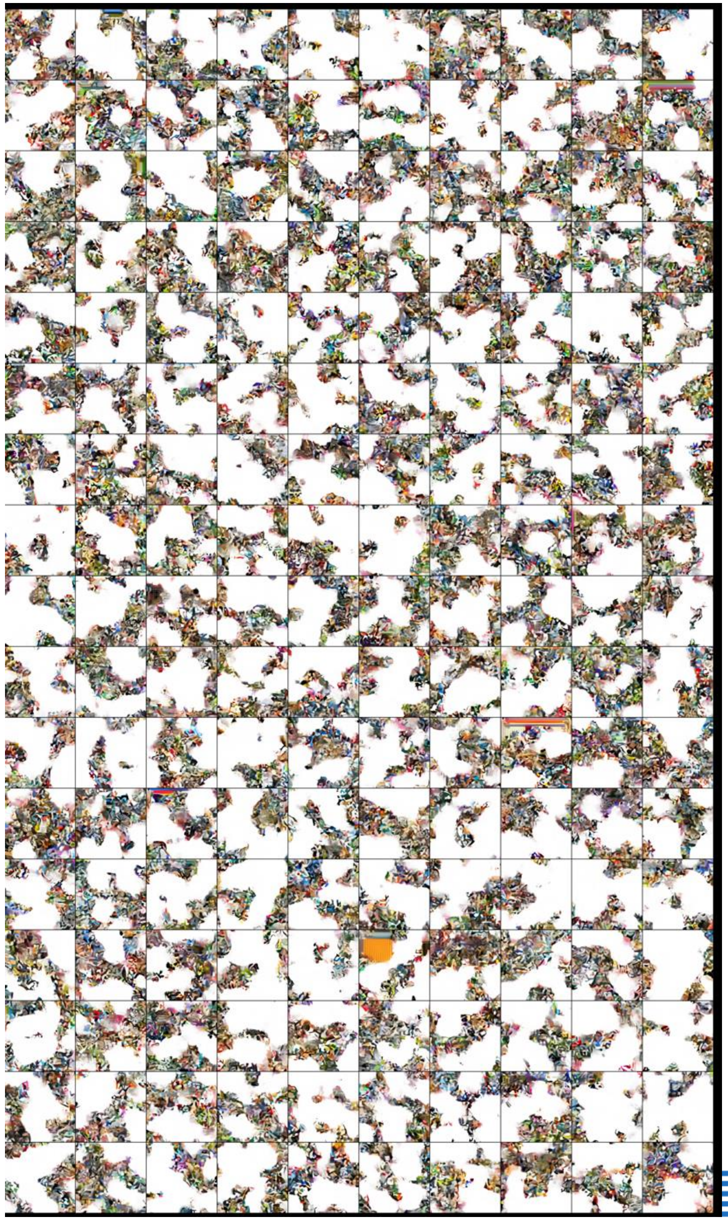

Hopfield Encoding Network Retrieval

Cross-modal Hopfield Networks

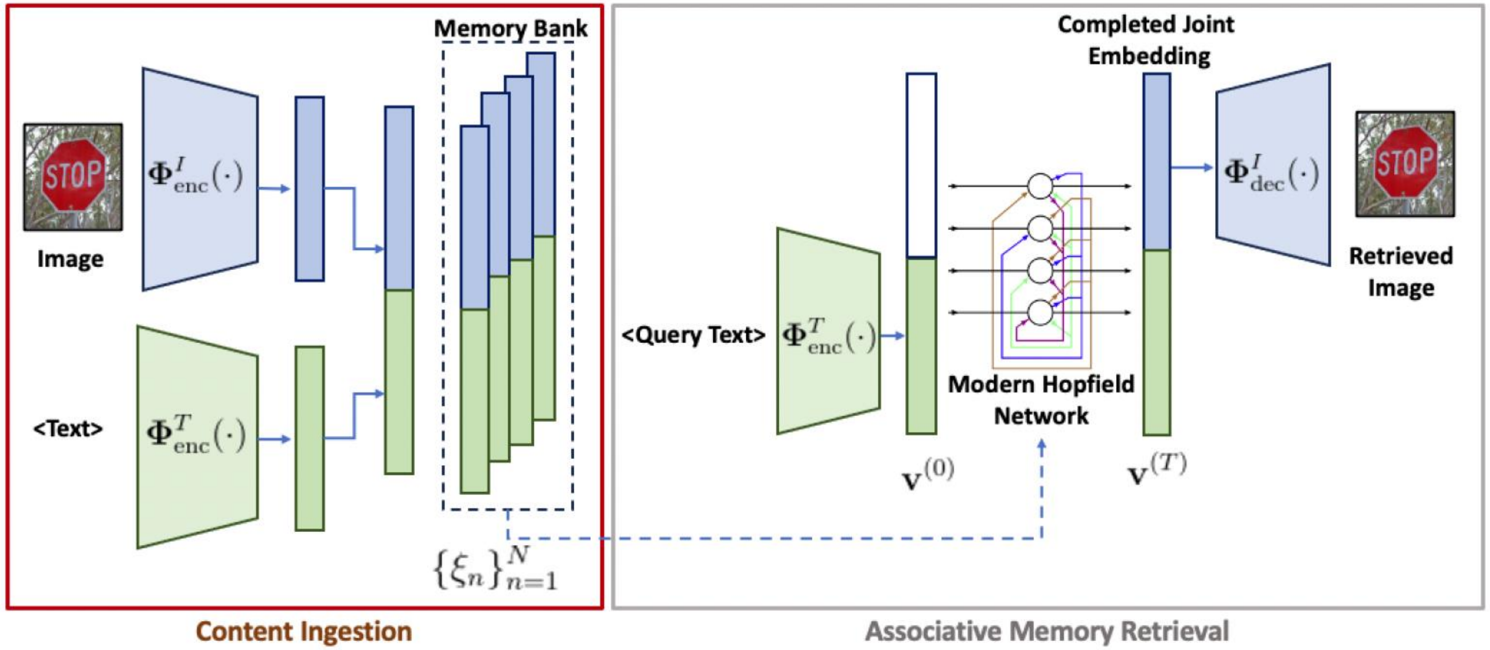

- Can the association be through language?
  - Yes
- Should the associative cue be in the same modality form?
  - No
- Can the associative cue be separately encoded?
  - Yes
- Should the associative cue be unique?
  - Yes

- Memory vectors are created by concatenating embeddings from both modalities.
- Queries in one modality retrieve corresponding content in another modality.
- Encoded latent representations ensure separability even across modalities

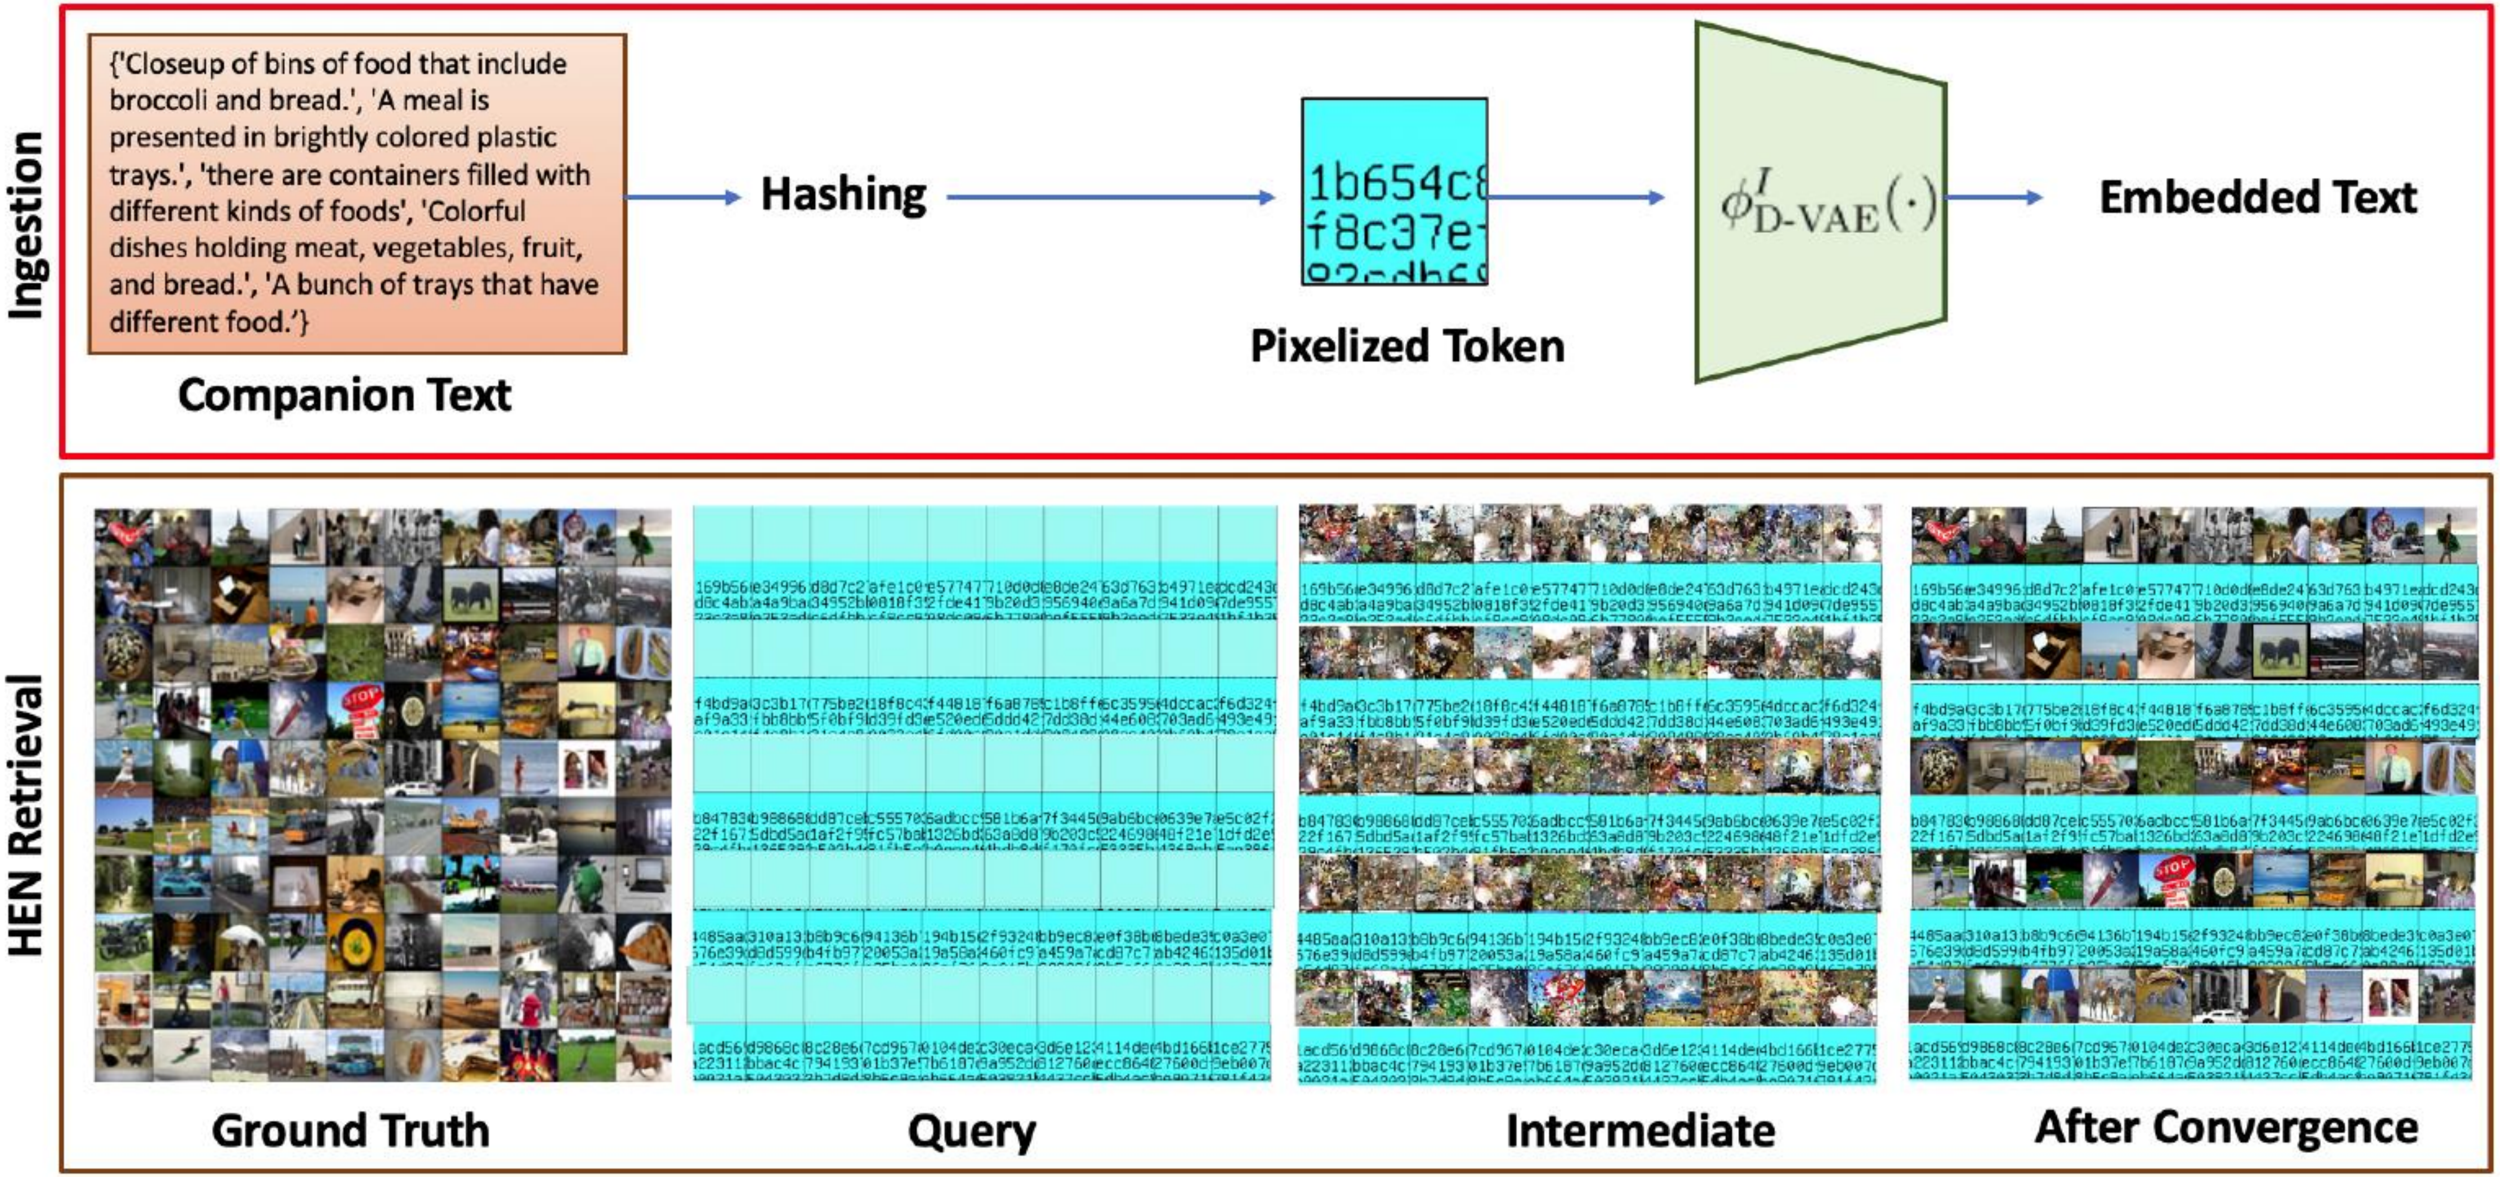
Ingestion
{'Closeup of bins of food that include broccoli and bread.', 'A meal is presented in brightly colored plastic trays.', 'there are containers filled with different kinds of foods', 'Colorful dishes holding meat, vegetables, fruit, and bread.', 'A bunch of trays that have different food.'}
Companion Text
Hashing
Pixelized Token
$\phi^{I}_{\text{D-VAE}}(\cdot)$
Embedded Text
HEN Retrieval
Ground Truth
Query
Intermediate
After Convergence

Ingestion

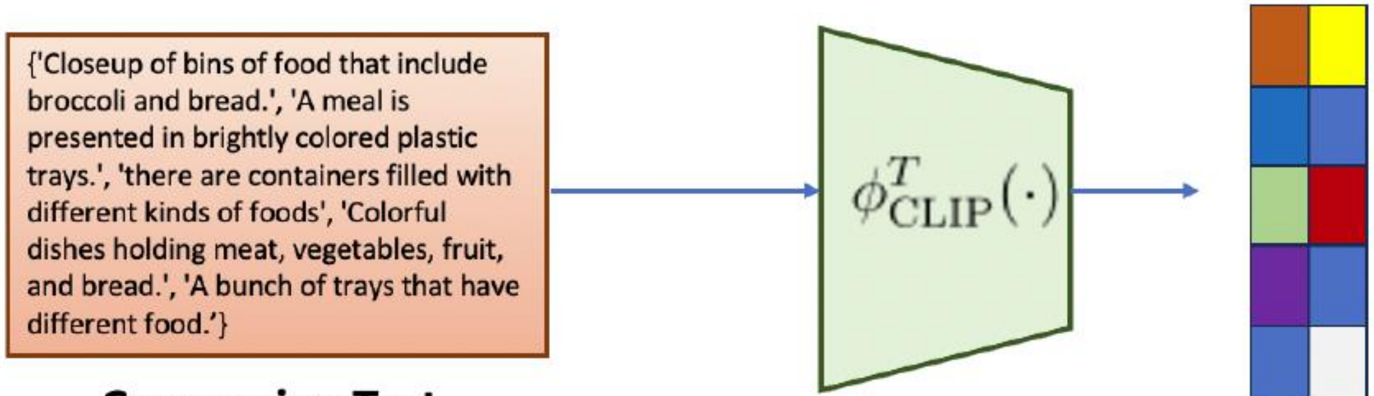

Companion Text

Embedded Text Tokens

NUM IMAGES 1-SSIM, MSE	vqf8	vqf16	Image	D-VAE
6000-CLIP	0.016, 0.000	0.024, 0.000	0.681, 0.118	**0.000, 0.000**
6000	0.016, 0.000	0.024, 0.000	0.952, 0.214	**0.000, 0.000**
8000	0.016, 0.000	0.023, 0.000	0.952, 0.215	**0.000, 0.000**
10000	0.015, 0.000	0.024, 0.000	0.952, 0.215	**0.000, 0.000**
15000	0.015, 0.000	0.024, 0.000	0.952, 0.215	**0.000, 0.000**

HEN Retrieval

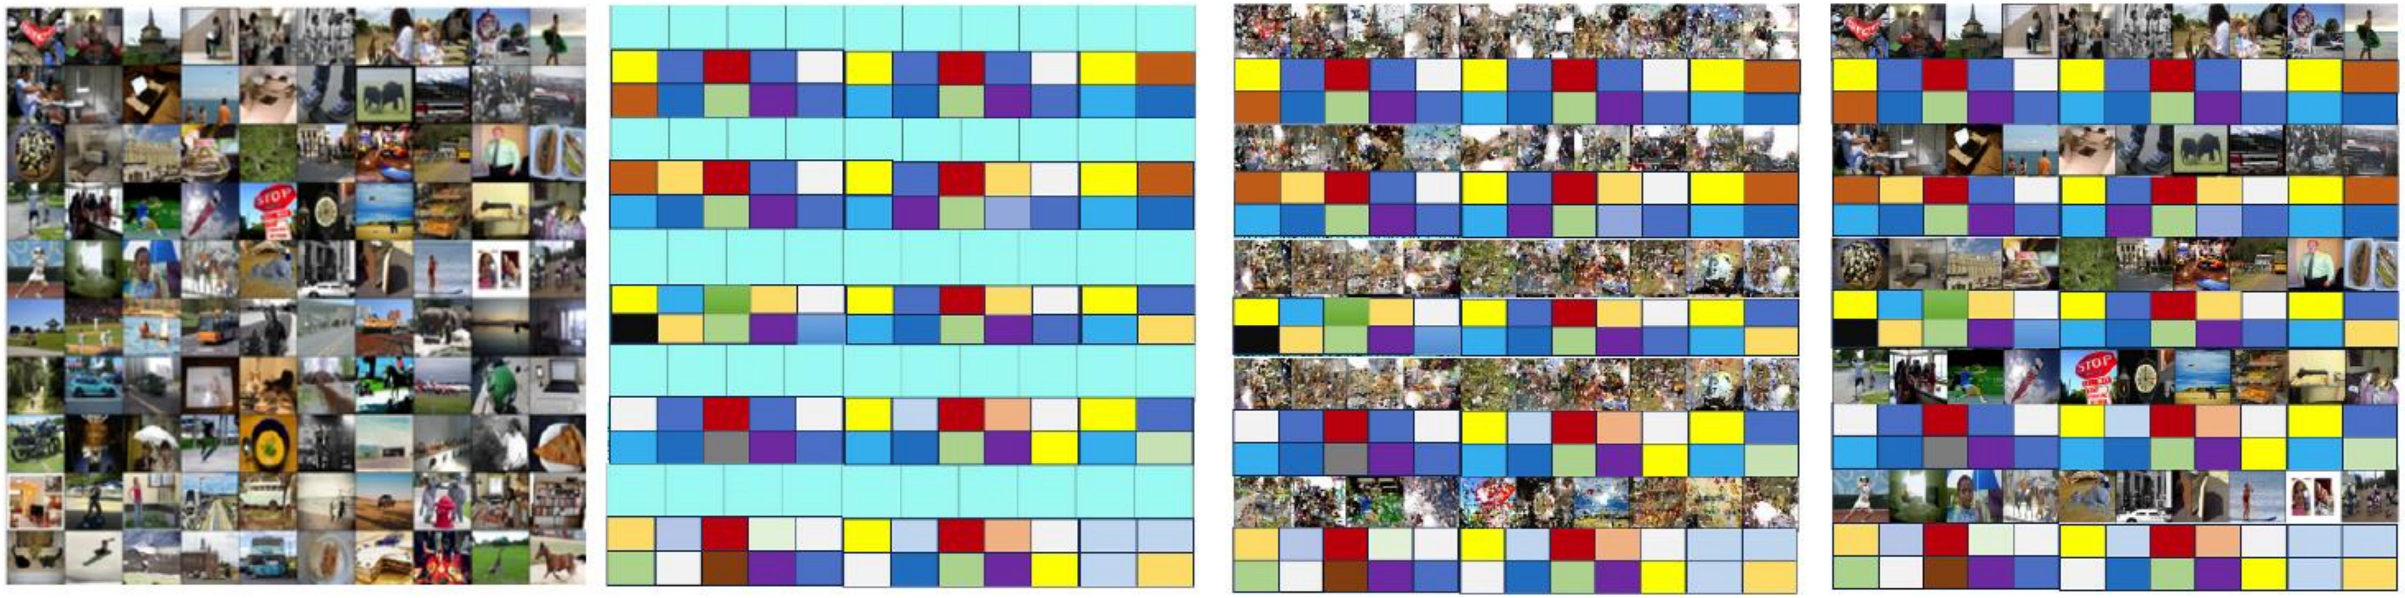

Ground Truth **Query** **Intermediate** **After Convergence**

Need for unique associations

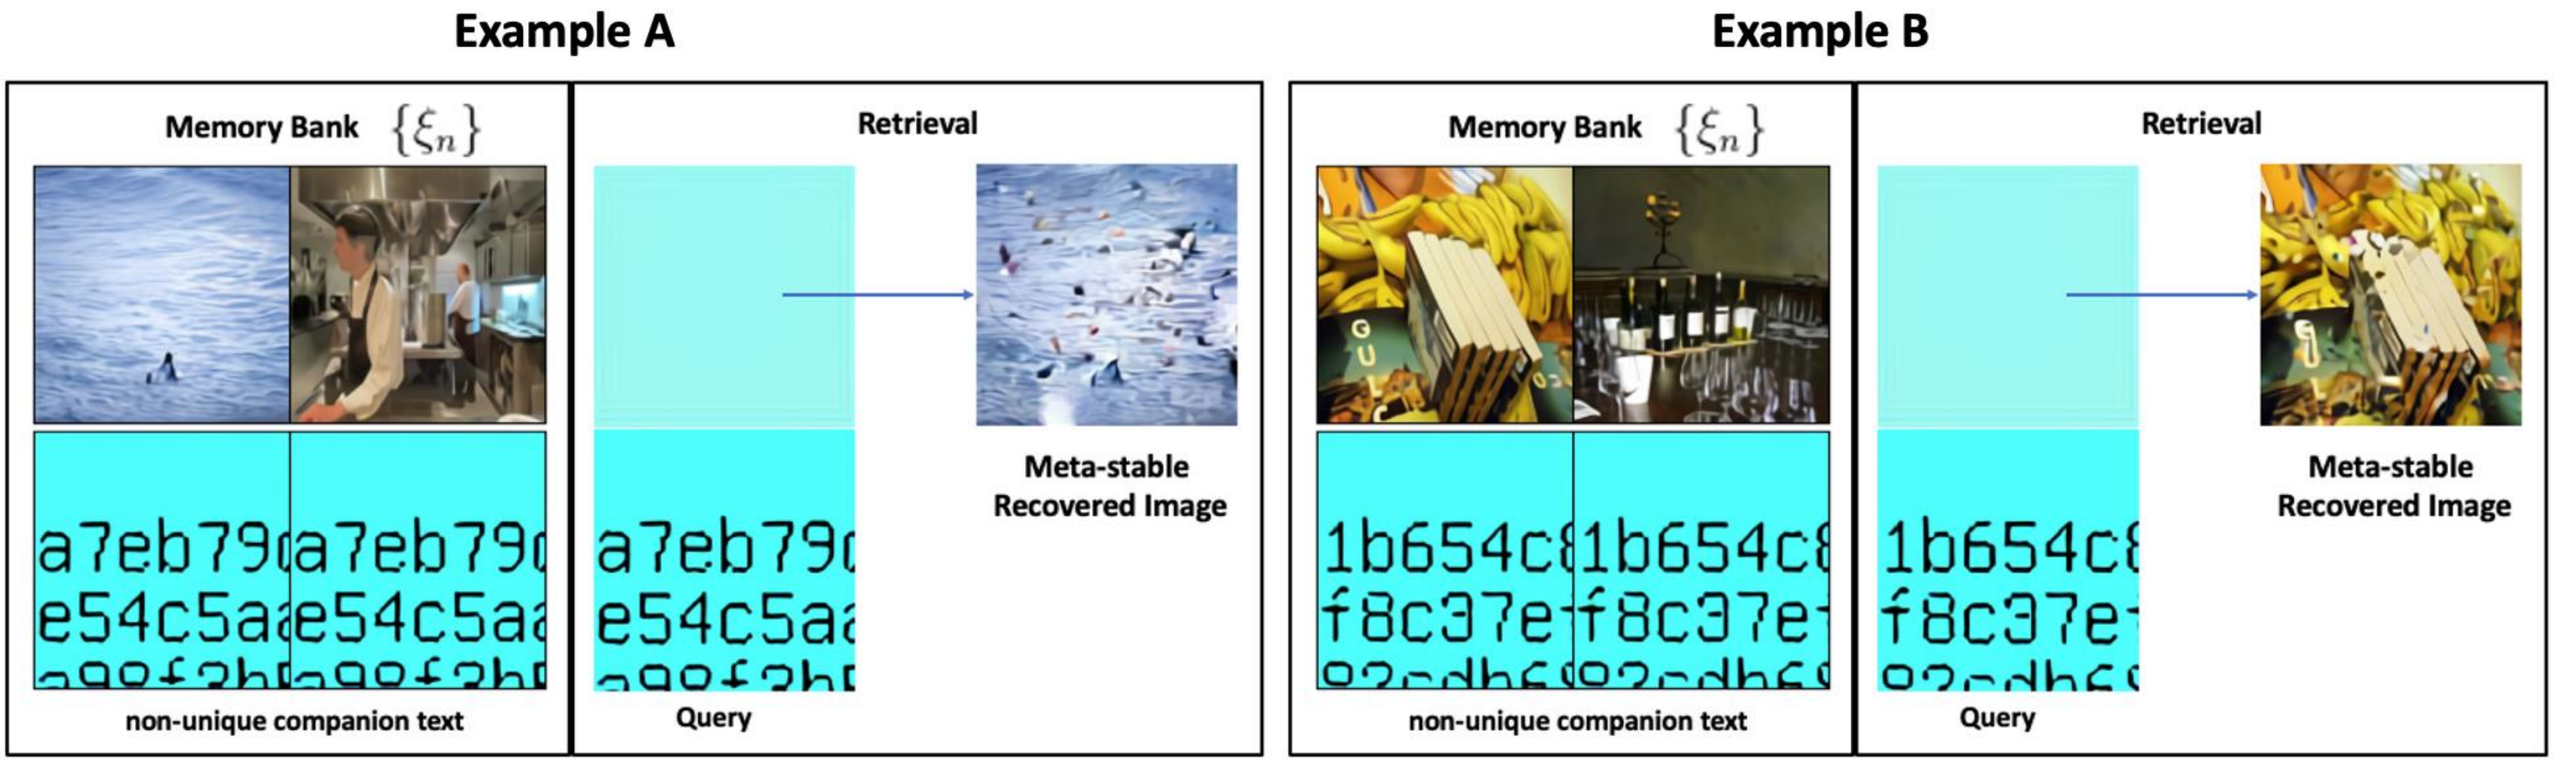

Neuro-Inspired Memories

- Why model declarative memory?
- A computational model of declarative memory
  - Modeling semantic memory
  - Modeling episodic memory
  - Modeling the associative memory formation in the hippocampus
- A commercial content-storage system modeled after human memory

Ingestion Processing

File D1

Saliency Map

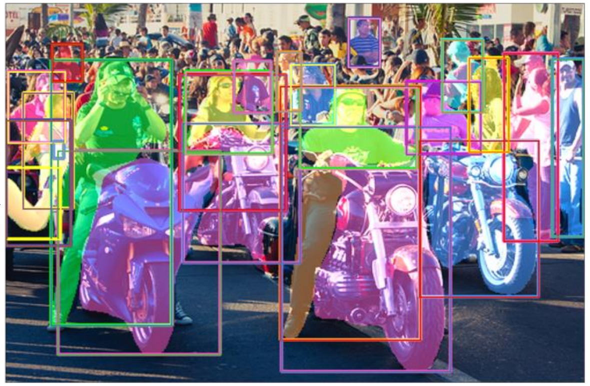

Attended Regions

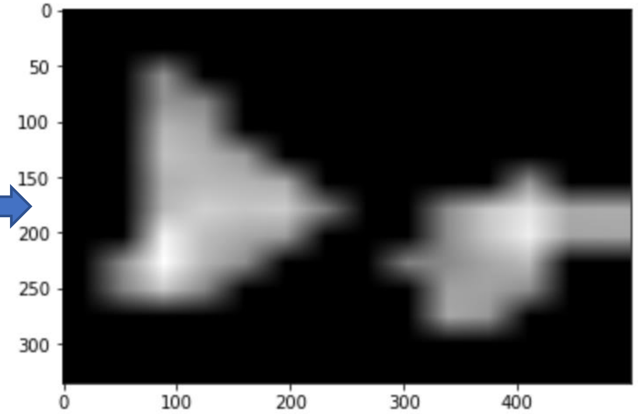

Selected Region S1

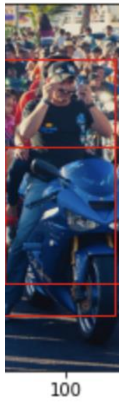

Nearby Object Labels L1,L2..

'motorcycles': '0.19647014',
'biker': '0.19616504',
'bike rider': '0.14070328',
'red hat': '0.13334015',
'motor bike': '0.13019492',
'motorcycle': '0.11932715',
'black motorcycle': '0.11693948',
'motorbike': '0.09740791',
'handlebars': '0.061977874',
'red cap': '0.055556886',
'handlebar': '0.052708704',
'blue helmet': '0.04897222',
'yellow cap': '0.04508579',
'two wheels': '0.044758428',
'motorcycle helmet': '0.04200281',
'bikes': '0.034488477',
'red helmet': '0.029659964',
'bike': '0.028195346',
'motorcycle seat': '0.0274861',

Semantic Memory

helmet->{'red helmet', 'baseball helmet', 'green helmet', 'safety covering', 'white helmet', 'blue helmet', 'yellow helmet', 'bike helmet', 'motorcycle helmet'} hair->{brown hair, 'curly hair', 'woman's hair', 'man's hair', 'gray hair', 'dark hair', 'girl's hair', 'long hair', 'blond hair', 'black hair', 'hair', 'white hair', 'red hair', 'blonde hair'}

Episodic memory

Episodic Memory indexing

Motorcycles->S1->D1, ->S3->D1

Blue helmet->S1->D1

....

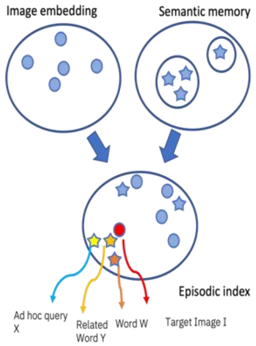

Retrieval Processing

File D1

Selected Region S1

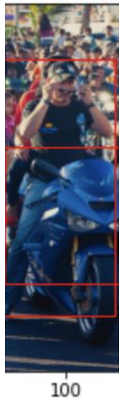

Retrieve images files depicting safety covering

Query attentional processing

"Safety Covering"

Semantic Memory

helmet->{'red helmet', 'baseball helmet', 'green helmet', 'safety covering', 'white helmet', 'blue helmet', 'yellow helmet', 'bike helmet', 'motorcycle helmet'}

hair->{brown hair, 'curly hair', 'woman's hair', 'man's hair', 'gray hair', 'dark hair', 'girl's hair', 'long hair', 'blond hair', 'black hair', 'hair', 'white hair', 'red hair', 'blonde hair'}

S1, D1

Bioinspired Retrieval

Episodic Memory indexing

Motorcycles->S1->D1, ->S3->D1

Blue helmet->S1->D1

…….

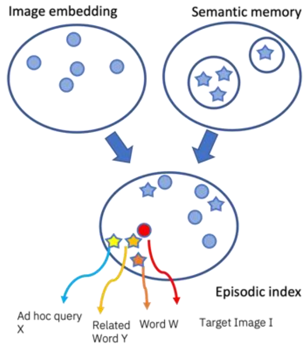

IBM CAS - Content Aware Storage

- GA in March 2025
- At NVIDIA GTC conference

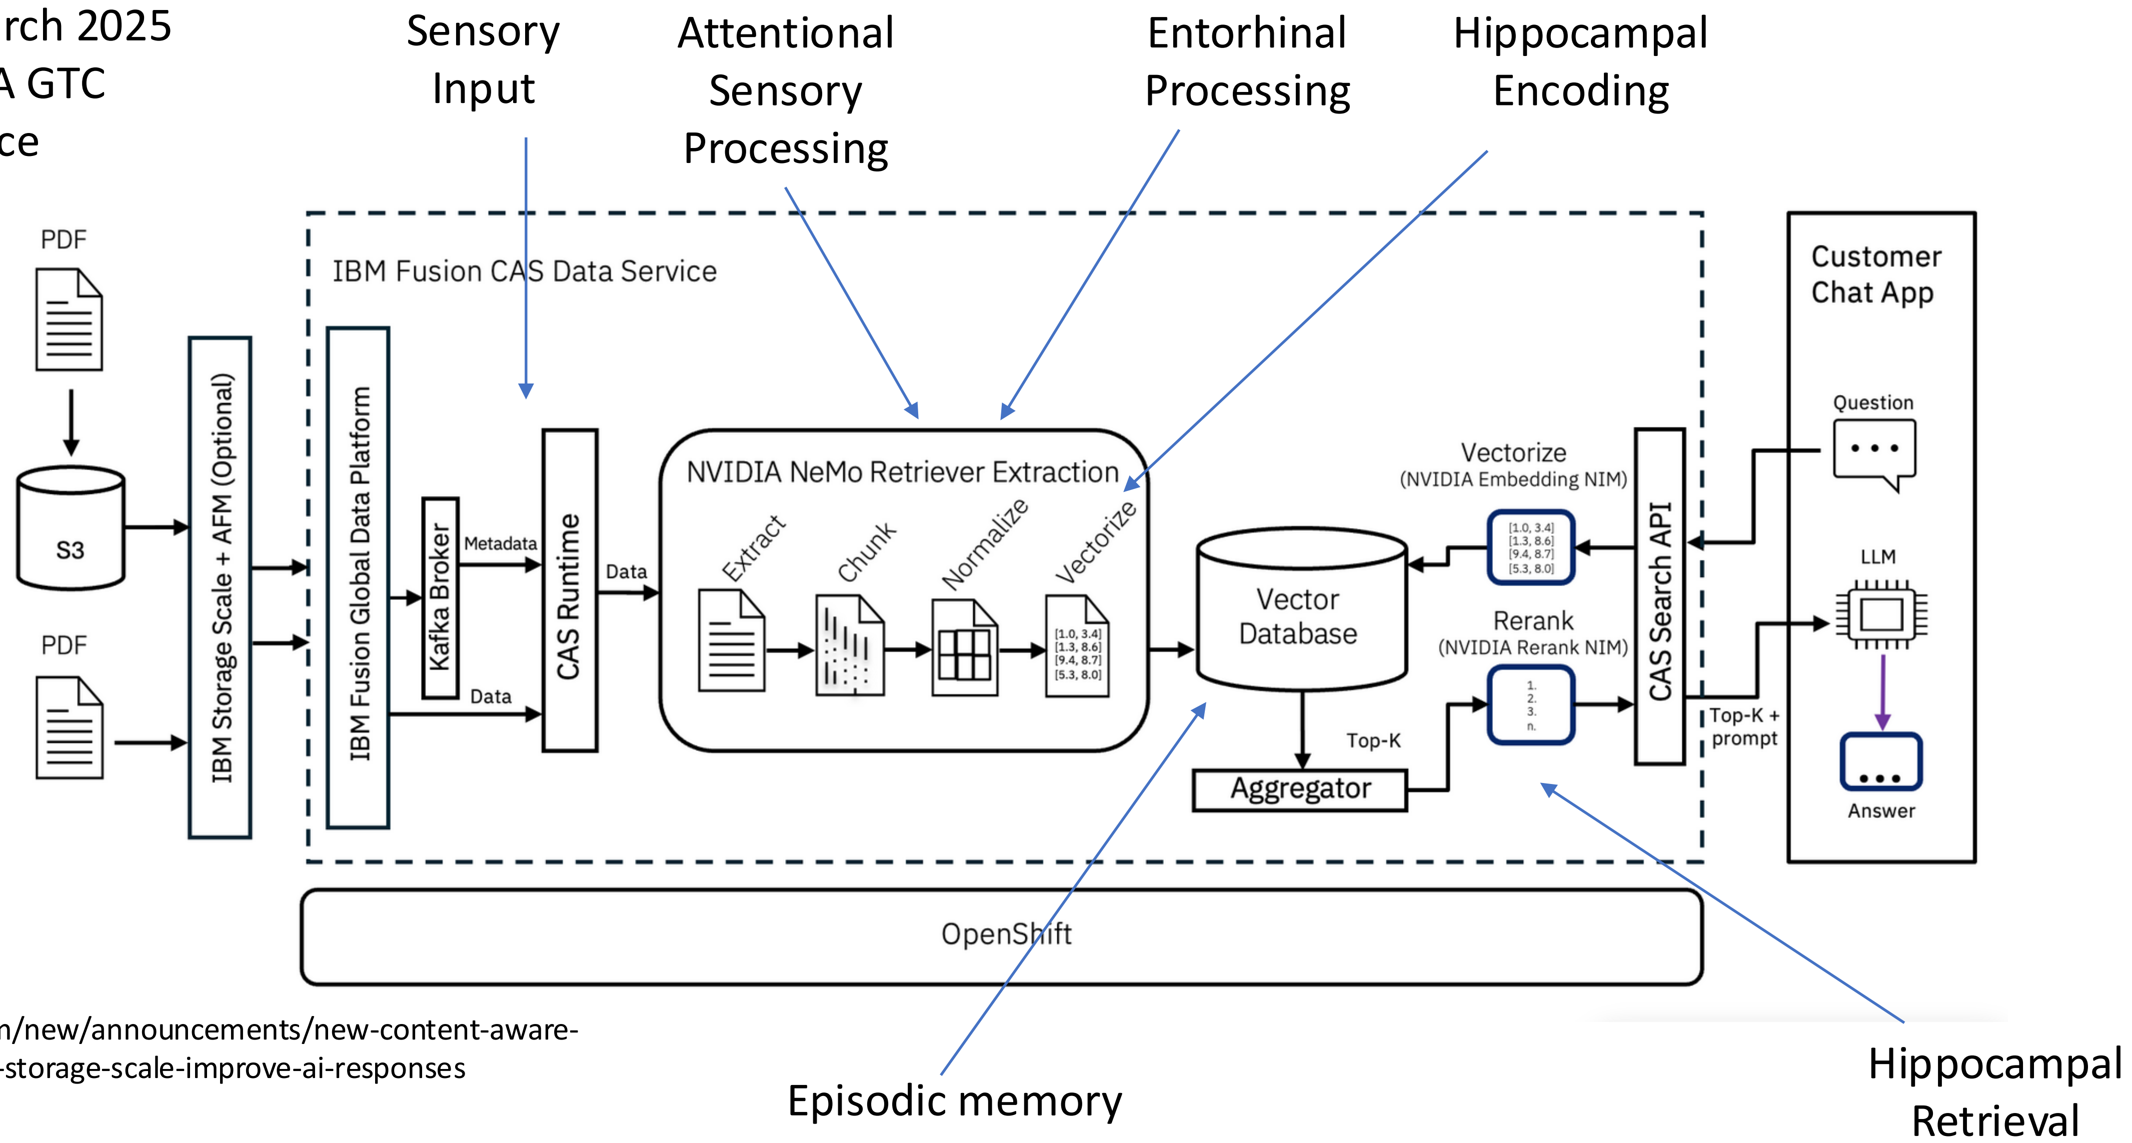

https://www.ibm.com/new/announcements/new-content-aware-capabilities-help-ibm-storage-scale-improve-ai-responses

Features of CAS

- Natural language ad hoc semantic retrieval
- High accuracy multimodal search
- Incremental update to memory
- Naturally associative between text and images

Evaluation results on enterprise datasets

Dataset	Chunks	Queries	GT Method	Recall@5	Recall@5 difference	Recall@1	Recall@1 difference
AAPOT	53	182	Auto	0.93/0.89	4.4%	0.72/0.55	30%
GT20	2247	1000	Auto	0.83/0.73	13.6%	0.63/0.47	34%
IndustrySample	3902	1000	Auto	0.88/0.84	4.7%	0.72/0.60	19.4%
Patent107	88398	1000	Auto	0.76/0.67	11.9%	0.63/0.45	39%

Summary

- A computational model of declarative memory can be a viable way to implement a content-aware storage platform
- Important aspects of human memory simulated through novel neural networks:
  - Semantic memory developed using supervised contrastive learning from ontological structure
  - Episodic memory developed using semantic memory substrate
  - Cross-modal Hopfield Networks improve capacity, offer higher compression, and increased accuracy of retrieval requiring uniqueness of association
- Future work:
  - The model implementation still under development with the all key aspects need to be tied back together
  - The transition from memory formation to long-term memory store needs to be explored
  - Engram theory could be a possible tie-in