128 ideas from “The molecular memory code and synaptic plasticity: A synthesis” By Sam Gershman

The molecular memory code and synaptic plasticity: A synthesis - ScienceDirect

1. Synapses are conceived as storage sites for the parameters of an approximate posterior over latent causes.
2. Intracellular molecules are conceived as storage sites for the parameters of a generative model.
3. The model stipulates how these two components work together as part of an integrated algorithm for learning and inference.
4. To understand the biological basis of memory, we must answer: What is the content of memory? What is the structure of memory?
5. Content refers to the information encoded in memory.
6. Structure refers to the code that maps information into a physical trace, or “engram.”
7. Gallistel has argued that the content of memory consists of facts, i.e. explicit representations of variables in the environment, rather than associations.
8. These facts are read-out by the spiking activity of neurons, …
9. … and thereby made computationally accessible to downstream neurons.
10. Synaptic plasticity, in this view, plays no role in memory storage.
11. I hypothesize that: (i) this inference model is implemented by spiking activity in a network of neurons
12. (ii) its parameters are stored at the synapse
13. (iii) synaptic plasticity updates these parameters to optimize the inference model
14. Animals appear to use representations of time to guide their behavior, even in simple Pavlonian conditioning protocols.
15. The brain cannot compute with information that it does not represent in memory.
16. Models that make time and space computationally accessible but that do not solve the problem of long-term storage: …
17. …How can information be retrieved long after persistent activity has subsided?
18. Spike timing-dependent plasticity can be used to learn a “navigation map” …
19. … in which firing fields shift in the animal’s traversal direction.
20. This solution presupposes that the enduring memories of goal location are stored outside of the hippocampus; …
21. … the hippocampal neurons are, in effect, a readout of these memories.
22. the problem of storing the relevant content in a computationally accessible format
23. In the hippocampus, the lifetime of spines is about 1-2 weeks.
24. How can behaviorally expressed memories be more persistent than their putative synaptic storage sites?
25. Systems consolidation: the process by which memories are gradually transferred from a temporary store in the hippocampus to a more durable representation in the neocortex.
26. Changes in intrinsic excitability are often coupled with changes in synaptic strength.
27. On the other hand, changes in intrinsic excitability and changes in synaptic strength can sometimes be dissociated, suggesting that they may play distinct functional roles.

28. Purkinje cells appear to store information about the timing of input signals in a cell-intrinsic format (Johansson et al., 2014).
29. McConnell & collaborators in 1959 showed that planarians can retain conditioned fear of a light stimulus across decapitation.
30. In 1962, McConnell took advantage of naturally occurring cannibalism between planarians, feeding trained subjects to untrained subjects, and finding that the untrained subjects exhibited conditioned fear of the light stimulus on the first day of training.
31. The untrained subjects had apparently ingested a memory!
32. (Jacobson et al. 1966) – directly injecting RNA extracted from trained planarians into untrained planarians can induce a transfer effect
33. Jennings demonstrated that when repeatedly stimulated, Stentor exhibited a regular sequence of distinct avoidance behavior
34. This suggests that the response to stimulation was memory-based: …
35. … the same external stimulus produced a different behavioral output as a function of stimulation history.
36. Amazingly, a single cell can maintain multiple memory traces, …
37. … as evidenced by the fact that habituation to cholinergic stimulation has no effect on habituation to depolarization.
38. Landauer (1964) conceived of neurons as electronic filters tuned to particular input frequencies.
39. If the frequency tuning of a neuron is determined by some biochemical signature, …
40. … then changes in RNA composition would generate (via protein synthesis) changes in the neuron's signature and hence its frequency tuning.
41. During periods of spiking activity, a neuron's membrane becomes permeable to RNA molecules originating in glial cells.
42. If the spiking activity of the neurons could somehow select RNA that produces particular signatures, …
43. … then the neuron could acquire long-term frequency tuning.
44. Landauer's ideas, which were ahead of his time and remain neglected, appear to presage more recent theories of bioelectric memory codes (Fields & Levin, 2018).
45. RNA is responsible for the transfer of pathogen avoidance between nematodes (Moore et al., 2021).
46. This development rehabilitated the plausibility of memory transfer.
47. This development indicated a mechanism by which transfer could occur under natural conditions, …
48. … namely, by packaging RNA into virus-like particles, i.e. capsids.
49. Peripheral receptors in the nervous system collect sensory observations, …
50. … and it is the job of the brain to make sense of these observations.
51. Bayesian theories of the brain hold that this sense-making process…
52. … corresponds to probabilistic inference over the latent causes of observations.
53. The phenomenon of rapid reacquisition: animals switch to the high reward option more quickly after repeated reversals.
54. This suggests that the animals gradually learn the parameters governing the reversal task, …

55. … leaving only uncertainty about the latent variable, …
56. … which can be resolved quickly after a few choices.
57. Exact Bayesian inference is equivalent to finding a conditional distribution that minimizes free energy.
58. The separation of updates for ϕ and λ implies a factorized, or “mean-field,” posterior approximation: …
59. $q(z, \theta \mid x) = q_{\phi}(z \mid x) q_{\lambda}(\theta | x)$
60. Using these gradients, approximate inference and learning can be carried out by stochastic gradient descent.
61. The simplest form of stochastic gradient descent is given by: $\Delta\Phi = -\alpha_{\phi} \nabla_{\phi} F$ and $\Delta\lambda = -\alpha_{\lambda} \nabla_{\lambda} F$ where $\alpha_{\phi}, \alpha_{\lambda}$ are learning rates.
62. The algorithmic solution provides high-level answers to the content and structure questions about memory.
63. Content consists of the variational parameters ϕ and λ.
64. From the perspective of free energy optimization, there is no substantive difference between these parameters.
65. However, from the perspective of biology, there is a substantive difference: …
66. … these parameters may be encoded by different substrates: one synaptic (ϕ) and one molecular (λ).
67. The structure of memory consists of writing and reading operations for these two types of content.
68. Memory is written via the variational update equations.
69. Memory is read via sampling q_{ϕ} and q_{λ}.
70. A complete algorithmic model of learning and memory should explain why learning proceeds more quickly under some conditions than others.
71. “Associability” – the speed at which two stimuli enter into association.
72. I hypothesize a direct coding scheme in which the spiking of individual neurons…
73. … reports a random sample of a variable conditional on the neuron’s input.
74. Two alternatives to direct coding: (1) predictive coding (2) probabilistic population coding
75. Predictive coding: a neuron reports a prediction error
76. Probabilistic population coding: a variable is represented by the spiking activity of a neural population
77. Post-synaptically silent synapse: lacks AMPA receptors but possesses NMDA receptors
78. If synaptic plasticity is mediated by the trafficking of AMPA receptors to the postsynaptic membrane, …
79. … then an inference model that depends on AMPA receptors can be separated from a generative model that depends on NMDA receptors.
80. Specifically, a cell could receive inputs about its probabilistic dependencies through NMDA receptors without eliciting AMPA currents that primarily control spiking activity (i.e. sampling from the inference model).
81. Departing from Rezende & Gerstner (2014), I propose that the L_i term is a signal generated from querying the generative model.

82. “Querying the generative model” consists of evaluating the probability of h under the joint probability distribution parameterized by theta.
83. If RNA encodes theta, then one possibility is that signals encoding x and z bind to RNA, …
84. … and then a specialized molecule reads the bound complex to report the scalar log joint probability L_i, …
85. … and making it accessible to the plasticity rules described above.
86. Learning “facts” about the environment does not require synaptic plasticity according to my account (facts are stored inside the cell), …
87. … but synapses are necessary for the expression of memory for these facts, …
88. … and synaptic plasticity exists to optimize this expression.
89. This is why synaptic plasticity accompanies learning.
90. This is why the disruption of synaptic plasticity interferes with the expression of memory, …
91. … even though synapses may not be the sites of memory storage.
92. When synapses are destroyed, the generative parameters are preserved in the RNA codes, …
93. … which can circulate between cells in virus-like capsids.
94. The RNA (or protein) codes need not be durable if they are being continuously sampled via transcription or post-translational modification of newly synthesized proteins.
95. The terminal storage sites are hypothesized to be nuclear marks, such as histone acetylation or DNA methylation, …
96. … which are sufficiently durable that they can be transmitted intergenerationally (although not always faithfully).
97. If the brain represents a generative model, …
98. … then this representation must be implemented by populations of cells, …
99. … with the parameters of the generative model stored in the form of synaptic strengths.
100. This means that if individual cells, including unicellular organisms, represent generative models, …
101. … then they must be using a fundamentally different mechanism.
102. In other words, the conventional view implies a radical evolutionary discontinuity, …
103. … yet the evidence seems to point towards continuity rather than discontinuity.
104. If individual cells are able to form local generative models of their microenvironment, …
105. … then these local models can be linked up to form more complex generative models.
106. This would necessitate continuous exchange of information about the generative model in order to maintain representational alignment -- …
107. … a function that may be carried out by the transfer of RNA between cells.
108. Belief bias is a natural consequence of amortized variational inference (Learning to Infer, 2021).
109. If the function, e.g. a neural network, used to approximate Bayesian inference has limited capacity, …

110. … then it inevitably has to concentrate its capacity on the distributions it encounters frequently.
111. In many domains, memory is used not only to store information about the world, but also to store information about how to think.
112. A person may know the rules of chess or the axioms of mathematics but may not be a particularly good chess player or mathematician.
113. These are skills that are acquired from experience, but the nature of this experience is not about fact learning, …
114. … in the traditional sense of observing the world, …
115. … since the relevant knowledge is acquired by thinking more.
116. Learning to think is conceptually distinct from, and complementary to, fact learning.
117. Four main conclusions: (1) The available evidence makes it extremely unlikely that synapses are the site of long-term memory storage for representational content, …
118. … i.e. memory for “facts” about quantities like space, time, and number.
119. (2) Fact memories, or more generally probabilistic beliefs about facts, …
120. … are plausibly stored in an intracellular molecular format.
121. (3) Synapses may be the site of long-term memory storage for computational parameters…
122. … that facilitate fast context-dependent belief updating and communication.
123. (4) These two forms of memory (representational and computational) work synergistically to optimize a common objective function (free energy).
124. This synergy is realized by interactions between synaptic and molecular processes within a cell.
125. It was not the case that a non-associative hypothesis was considered and then rejected, …
126. … but rather that it was never considered at all; presumably, because no one could imagine what that would look like.
127. This story testifies to the power of theory, even when implicit, …
128. … to determine how we interpret experimental data and ultimately what experiments we do.