PREDICTIVE SPIKE TIMING ENABLES DISTRIBUTED SHORTEST PATH COMPUTATION IN SPIKING NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Efficient planning and sequence selection are central to intelligence, yet current approaches remain largely incompatible with biological computation. Classical graph algorithms like Dijkstra's or A* require global state and biologically implausible operations such as backtracing, while reinforcement learning methods rely on slow gradient-based policy updates that appear inconsistent with rapid behavioral adaptation observed in natural systems.

We propose a biologically plausible algorithm for shortest-path computation that operates through local spike-based message-passing with realistic processing delays. The algorithm exploits spike-timing coincidences to identify nodes on optimal paths: Neurons that receive inhibitory-excitatory message pairs earlier than predicted reduce their response delays, creating a temporal compression that propagates backwards from target to source. Through analytical proof and simulations on random spatial networks, we demonstrate that the algorithm converges and discovers all shortest paths using purely timing-based mechanisms. By showing how short-term timing dynamics alone can compute shortest paths, this work provides new insights into how biological networks might solve complex computational problems through purely local computation and relative spike-time prediction. These findings open new directions for understanding distributed computation in biological and artificial systems, with possible implications for computational neuroscience, AI, reinforcement learning, and neuromorphic systems.

1 INTRODUCTION

Retrieving sequences and finding shortest paths is fundamental to both artificial and biological information processing systems. In computer science, this is typically achieved using algorithms like Dijkstra's (Dijkstra, 1959) or Bellman Ford (Bellman, 1958; Ford, 1956; Shimbel, 1955; Moore, 1959) and their hierarchical variants (Hart et al., 1968; Geisberger et al., 2008; Funke & Storandt, 2015; Madkour et al., 2017; Dibbelt et al., 2014; Bast et al., 2016; Delling et al., 2009; Blum, 2019). Sequence selection appears across domains including robotics (Kavraki et al., 1996; Latombe, 1991; Lozano-Pérez & Wesley, 1979; Jafarnia-Jahromi et al., 2023), decision making (Puterman, 1994; Chen et al., 2021), language production (Zganec Gros & Zganec, 2008), memory retrieval Wang et al. (2020b), database queries (Griffiths Selinger et al., 1979; Graefe, 1993), hierarchical data traversal (Tarjan, 1972; Cormen et al., 2009), distributed system analysis (Casavant & Kuhl, 1988; Kwok & Ahmad, 1999), and type systems (Pottier & Rémy, 1998; Agesen, 1994).

Modern AI systems epitomize this search paradigm. Transformers function as sophisticated search engines (Vaswani et al., 2017), identifying token sequences that maximize likelihood functions. This highlights that effective search requires both data storage and efficient retrieval mechanisms—a dual requirement fundamental to both artificial and biological intelligence[1]. However, reinforcement learning typically relies on iterative policy updates through gradient-based optimization (Williams, 1992; Sutton et al., 1999). Despite progress in sample efficiency Fatkhullin et al. (2023); Xu et al. (2020); Wang et al. (2024); Ishfaq et al. (2025), this requires extensive trial-and-error learning (Ilyas et al., 2020), which seems inconsistent with the rapid behavioral adaptation observed in biological

[1]Also proclaimed in Sutton's *The Bitter Lesson*, 2019.

systems, where animals quickly navigate novel environments (Rosenberg et al., 2021; Vale et al., 2017; Zugaro et al., 2003; Widloski & Foster, 2022; Pfeiffer & Foster, 2013).

The brain must solve similar problems for decision making, memory formation, and path planning. Given abundant dynamically changing sequences in cortical networks (Dragoi & Buzsáki, 2006; Dragoi & Tonegawa, 2011; Dragoi, 2020; Buzsáki & Llinás, 2017; Buzsáki & Tingley, 2018; Diba et al., 2014; Bakermans et al., 2025b; Schwartenbeck et al., 2023) and theoretical frameworks for relational information processing (Whittington et al., 2020; Bakermans et al., 2025a; Waniek, 2020), understanding how the brain implements sequence selection for navigation (Pfeiffer & Foster, 2013) or memory retrieval (Oberauer et al., 2018) is appealing. Yet, how the brain solves this efficiently, while allowing quick, contextually modulated re-planning, remains elusive, with the biggest hurdle being the biological *hardware* and its constraints.

Classical algorithms like Dijkstra's – which is considered optimal for data representable as graphs (Haeupler et al., 2024) – propagate from source to target. Then they use *back-tracing* during which they walk backwards along parent nodes that were collected during the forward search. This paradigm of graph traversal with parent tracking extends to many other algorithms, such as the forward-backward algorithm for Hidden Markov Models (Baum et al., 1970). While computationally efficient on classical hardware, implementing these in biological or neuromorphic systems (Furber et al., 2014; Pei et al., 2019; Davies et al., 2018; Richter et al., 2024; Kadway et al., 2023; Akopyan et al., 2015) faces major challenges. Forward propagation is feasible (Ponulak & Hopfield, 2013; Waniek, 2020; Orsher et al., 2024), but back-tracing requires neurons to remember activation sources. This contradicts evidence that neural dynamics are transient (Friston, 1997; Abeles et al., 1995; Amit, 1997; La Camera et al., 2019; Mante et al., 2013; Ponce-Alvarez et al., 2012; Seung, 1996; Wiltschko et al., 2015; Alderson et al., 2020; Buckley & Nowotny, 2012; Durstewitz & Deco, 2007; Rabinovich et al., 2006; Seliger et al., 2003; Brinkman et al., 2022; Mazor & Laurent, 2005; Koch et al., 2024; Tognoli & Kelso, 2014). Moreover, neural signaling, meaning the transduction of information from one neuron to another, is inherently directional and non-reversible (Maass, 1997; Harris & Shepherd, 2015; Douglas & Martin, 2004; Markov & Kennedy, 2013), precluding straightforward backward information propagation required for back-tracing (and, for that matter, error-backprop and automatic differentiation, but see recent progress in Bellec et al. (2020); Lillicrap et al. (2020); Renner et al. (2024); Whittington & Bogacz (2017); Stanojevic et al. (2024)).

We propose a biologically plausible alternative using spike-timing predictions to compute shortest paths without back-tracing. Our contributions include 1) a novel spike-timing protocol for local shortest path inference, 2) analytical convergence proof, 3) simulation results, and 4) discussion of benefits and limitations. This predictive spike-time paradigm could extend to other algorithms, potentially transforming classical approaches into biologically-plausible, distributed variants that advance both neuroscience and adaptive artificial systems.

2 Related Work

Extensive research demonstrates that precise spike timing carries information and drives computation across brain areas (Kayser et al., 2010; Levi et al., 2022; Mainen & Sejnowski, 1995; Montemurro et al., 2007; Shmiel et al., 2006; Rolls et al., 2006; Srivastava et al., 2017), with millisecond-precise patterns encoding stimulus features and behavioral states. Building on this, computational models have proposed temporal coding schemes including rank order codes (Rullen & Thorpe, 2001; Thorpe & Gautrais, 1998) and polychronous networks (Izhikevich & Hoppensteadt, 2009; Szatmáry & Izhikevich, 2010), though these did not demonstrate flexible shortest-path computation. Spike-timing dependent plasticity (STDP) further supports timing-based computation (Masquelier et al., 2008; Bush, 2010; Bi & Poo, 1998; Caporale & Dan, 2008; Gilson, 2010; Markram et al., 2012; Pokorny et al., 2020), operating across timescales for sequence learning and circuit self-organization.

Early theoretical work proposed mechanisms from spike train patterns to synfire chains and gradient fields for sequence generation (Amari, 1972; Abbott & Blum, 1996; Abeles, 1991), with experimental evidence supporting spatiotemporal firing patterns in cortical assemblies (Dabagia et al., 2024; Schrader et al., 2008; Bouhadjar et al., 2022; Abeles et al., 1993). However, these lacked principled shortest-path algorithms with fast goal switching.

Several approaches model sequence learning through reinforcement learning (Samsonovich & Ascoli, 2005) or deep learning techniques for grid and place cells (Banino et al., 2018; Cueva & Wei, 2018; Sorscher et al., 2023), but don't solve shortest paths in biologically relevant timescales. Neural oscillations, especially theta rhythms, organize sequential activity (Wang et al., 2020a; Igata et al., 2021; McNamee et al., 2021; Papale et al., 2016; Parra-Barrero et al., 2021), creating temporal scaffolds for sequence representation and replay phenomena for memory consolidation and planning (Widloski & Foster, 2022; Pfeiffer & Foster, 2013). While demonstrating flexible path computation, the algorithmic mechanisms remain unclear.

Eligibility traces in e-prop (Bellec et al., 2020; Traub et al., 2020) enable biologically plausible backpropagation approximations through local synaptic traces for supervised learning. Our approach instead exploits temporal predictions for unsupervised path-finding without explicit gradients, representing a complementary timing-based paradigm.

Most spiking network path planning approaches use biologically implausible modifications such as weight changes Roth et al. (1997); Ponulak & Hopfield (2013); Schuman et al. (2019), connection removal Davies et al. (2018), or spike-time tables for back-tracing Krichmar et al. (2022). A recent study used spike-threshold adaptation for shortest paths (Dietrich et al., 2025), building on spiking hierarchical temporal memory (Bouhadjar et al., 2022) with separate excitatory/inhibitory populations. While successful in small simulations, their approach differs from ours by modeling distinct populations rather than unified message types, precluding formal convergence analysis and direct biological interpretation.

3 Model Description and Convergence Analysis

Our algorithm operates on a network of spiking neurons where each neuron connects to a local neighborhood, enabling message-passing between adjacent nodes. Moreover, neurons can broadcast inhibitory messages globally and exist in one of two meta-states: *tagged* or *untagged*. The core mechanism relies on *predictive tagging*: neurons become tagged when they receive inhibitory-excitatory (I-E) message pairs earlier than anticipated based on network timing predictions. Initially, only the goal neuron is tagged. When the algorithm begins, excitatory (E) messages propagate from the starting neuron throughout the network. Tagged neurons process messages faster and broadcast both local excitatory and global inhibitory signals, causing their predecessors to receive I-E pairs earlier than expected. This creates a cascading effect where neurons progressively become tagged, propagating the tagging state backward from goal to source until the shortest path emerges. In the following, we describe the neuron model and provide a formal convergence proof.

3.1 Timed State Machine Model of Neural Activity and Neural Interaction

Real neurons typically maintain a resting potential and perform non-linear integration of arriving excitatory inputs until reaching a spiking threshold (Koch, 2004). This threshold depends on factors including neuro-modulation and internal dynamics, while inhibitory inputs can suppress activity and prevent spike generation. We simplify this nonlinear behavior using a timed state machine model where each simulated neuron v exhibits five distinct states: *resting*, *processing*, *spiking*, *refractory*, and *inhibited*. Rather than modeling separate excitatory and inhibitory neural populations following Dale's principle, we represent inhibition and excitation as messages of type I or E, respectively, between model neurons.

A neuron remains *resting* until receiving a message, transitioning to *inhibited* (I messages) or *processing* (E messages). An *inhibited* neuron will remain inhibited for a duration of τ_{inh}, from which it recovers to *resting*. Provided it has not received an I message in the meantime, a *processing* neuron will move to *spiking* after τ^0_{proc}, and emit an E message to all neurons in its local neighborhood, which incurs another short delay τ_{spike}. After moving to *refractory*, it will then recover to *resting* after τ_{ref}. Message transmission incurs axonal delays Δt_I and Δt_E for I and E messages, respectively, upon sending, and a dendritic delay $\Delta t_{\text{dendritic}}$ upon arrival, with $\Delta t_I < \Delta t_E$ reflecting rapid inhibition (Packer & Yuste, 2011) [2].

[2] We further note that biological inhibition often targets the soma directly instead of the dendritic tree of a neuron, which we omit in our simplified model.

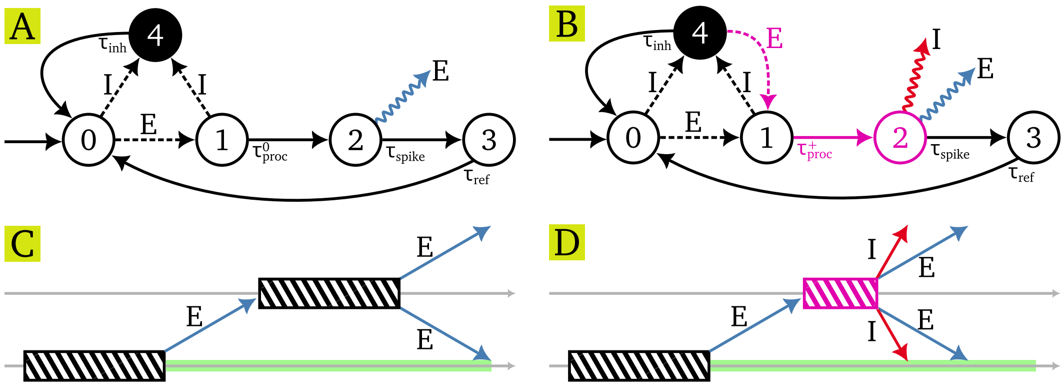

Figure 1: Each model neuron has four states: *resting* (0), *processing* (1), *spiking* (2), *refractory* (3), and *inhibited* (4). A neuron's state machine transits to another state either upon receiving an E or I message (dashed arrows, message type over arrow), or after a temporal delay (solid arrows, delay under arrow). A neuron can emit E and I messages (wavy arrows) when spiking. **(A)** shows the state machine of a *untagged* neuron, which produces only E messages when spiking. **(B)** shows the state machine of a *tagged* neuron, which can move from *inhibited* to *processing*, has a shorter processing delay $\tau^{+}_{\text{proc}} < \tau^{0}_{\text{proc}}$, and also generates I messages when spiking. **(C)** Sequence diagram of two interacting neurons. After the processing delay (hatched box), the neuron in the bottom row sends an E message (blue arrow) to its neighbors (here only top row), with transmission delay Δt_E indicated by the slant of the arrow. It predicts to receive recurrent E messages only after a time window that includes message delivery times and subsequent processing delay (indicated by green bar over the neuron's time axis). **(D)** Sequence diagram where the second neuron (top row) is *tagged*. The tagged neuron has faster processing (magenta hatched box), and emits both I and E messages, with $\Delta t_I < \Delta t_E$ indicated by slant of the arrows. The I-E message pair arrives earlier than predicted at the first neuron (bottom row), i.e. within the spike-delay prediction time window.

Tagged neurons exhibit two key differences. First, they transition more rapidly from processing to spiking using $\tau^{+}_{\text{proc}} < \tau^{0}_{\text{proc}}$, modeling threshold adaptation such as lowering the spiking threshold and local neuro-modulation (Azouz & Gray, 2000; Fontaine et al., 2014; Henze & Buzsáki, 2001). Second, they broadcast global I messages in addition to sending local E messages when spiking, following unspecific inhibition in real networks (Packer & Yuste, 2011). Critically, tagged neurons can transition from *inhibited* to *processing* upon receiving E messages, modeling neuromodulator-gated disinhibition. This enables the backward propagation of tagging states that drives shortest-path discovery.

Figure 1 illustrates the state machines of an *untagged* and *tagged* neuron, and sequence diagrams of neural interactions. While this model substantially simplifies real neuronal dynamics, it captures the essential timing relationships necessary to implement distributed shortest-path computation without requiring global coordination or explicit back-tracing, as neurons self-organize based purely on local timing relationships and predictive coding of spike arrival times.

3.2 Convergence Analysis

To analyze the properties of the distributed, timing-protocol based algorithm, we model the system of interacting neurons as a directed, unweighted graph in which, as introduced above, messages of two distinct types I and E are sent around. Specifically, we have a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of nodes (or neurons) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of directed edges between nodes. A path p from node u to node w written u—w, is the sequence of nodes $(u = u_1, u_2, \dots, u_l = w)$, meaning there exists a directed edge from u_i to u_{i+1}, denoted (u_i, u_{i+1}) or $u_i \to u_{i+1}$. Two nodes $u, v \in \mathcal{V}$ are connected, if there is a path u—w. The distance $k_{u,v} = d(u, v)$ between two nodes u and v is the minimal number of edges that need to be traversed on any path u—w. The neighborhood $\mathcal{N}(v)$ of $v \in \mathcal{V}$ consists of all nodes $u \neq v \in \mathcal{V}$ which are connected by an edge $v \to u$. The k-neighborhood $\mathcal{N}_k(v)$ of $v \in \mathcal{V}$ is the set of all nodes $u \neq v \in \mathcal{V}$ with $d(v, u) \leq k$.

As mentioned above, $v \in \mathcal{V}$ send messages of type E and broadcast messages of type I, if v is tagged, after receiving a message of type E themselves with message travel times Δt_I and Δt_E, respectively, and $\Delta t_I < \Delta t_E$. An I-message that is sent at time t from node v will thus arrive at all nodes $u \in \mathcal{V}$, whereas an E-message will arrive only at nodes $u \in \mathcal{N}(v)$. Without loss of generaltiy, we can omit $\Delta t_{\text{dendritic}}$ in the following analysis.

Recall that at any time t a node $v \in \mathcal{V}$ can be in one of five states and is either tagged or not. If a node is *tagged*, it can emit (forward) messages with short temporal delay τ^+_{proc}, whereas *untagged* nodes require τ^0_{proc} time. Hence, a message of type E that is sent from a tagged node v at time t will arrive at a time $t + \Delta t_E + \tau^+_{\text{proc}}$ at nodes $u \in \mathcal{N}(v)$, while a message of the same type that is sent from an non-tagged node will arrive only at $t + \Delta t_E + \tau^0_{\text{proc}}$ in the neighborhood of the sending node. A message of type I will arrive at all nodes $u \in \mathcal{V}$ at time $t + \Delta t_I + \tau^+_{\text{proc}}$ or $t + \Delta t_I + \tau^0_{\text{proc}}$ after being sent from a tagged or non-tagged node, respectively. A node which is purely *inhibited* and is not tagged itself cannot forward any messages.

A node $v \in \mathcal{V}$ becomes tagged if it produced a message of type E at time t and receives a (recurrent) I-E message pair in a shorter-than-expected time window. More precisely, if v receives

1. an inhibitory message I at time $t_{I,observed} = t_0 + \tau^+_{\text{proc}} + \Delta t_I < t_0 + \tau^0_{\text{proc}} + \Delta t_I = t_{I,\text{expected}}$, and
2. an excitatory message E at time $t_{E,observed} = t_0 + \tau^+_{\text{proc}} + \Delta t_E < t_0 + \tau^0_{\text{proc}} + \Delta t_E = t_{E,\text{expected}}$, where

$t_0 = t + \Delta t_E$ accounts for the travel time of the E-message from v_i to the subsequent node, then v_i itself becomes tagged. In other words, v becomes tagged if it receives an I-E message pair in a time window $\Delta t \leq 2\Delta t_E + \tau^+_{\text{proc}} < 2\Delta t_E + \tau^0_{\text{proc}}$ after sending an E message itself.

In the following, we will show that a system in which a target node t is tagged and the propagation of messages of type E begins at a start node s will eventually converge in finite time to a state where all nodes on a path p from s to t are tagged. We show the convergence and time boundedness using induction, where we show that after k iterations the nodes k-closest, i.e. at a distance at most k to the target, are tagged.

Base Case – Iteration 1 Tagging nodes adjacent to the target w, assuming that a path u—$w = (u = u_1, u_2, \ldots, u_l = w)$ from starting node u to w exists and that no nodes other than the target w are tagged.

1. (Forward Phase) Let $u, w \in \mathcal{V}$ be the start and the target node, respectively, and let propagation of messages of type E start at u. During the first iteration, a message of type E will eventually reach w at time t (after at most $O(N)$ propagations, where N is the number of edges m on the shortest path from u to w), and all neighbors of w that are on the shortest path to w will have sent messages of type E and I at time $t_0 = t - \Delta t_E$.
2. (Inhibitory Control, I-E pair) At time $t_1 = t + \tau^+_{\text{proc}}$, w sends an I-message to all nodes in the network that arrives after Δt_I time, and an E-message to all its neighbors that arrives after Δt_E time.
3. (Tagging Condition) Neighbors $v \in \mathcal{N}(w)$ receive
   (a) an I-message at $t_2 = t_1 + \Delta t_I = t + \tau^+_{\text{proc}} + \Delta t_I$, and
   (b) an E-message at $t_3 = t_1 + \Delta t_E = t + \tau^+_{\text{proc}} + \Delta t_E$.

   Since $t_2 < t_0 + \Delta t_E + \Delta t_I + \tau^0_{\text{proc}}$, $t_3 < t_0 + \Delta t_E + \Delta t_I + \tau^0_{\text{proc}}$, and $t_3 - t \leq \Delta t$, the timing condition for the I-E message pair holds. Those $v \in \mathcal{N}(w)$ which sent an E-message at time t thus will become tagged, while nodes which did not send an E-message will retain their previous state.

Inductive Hypothesis – Iteration *k* Assume that after k iterations, all nodes at a distance k from the target w on the shortest path towards the target are tagged.

Induction Step – Iteration $k+1$ We need to show that nodes at distance $k+1$ from the target will be tagged after iteration $k+1$.

1. (Forward Phase) In iteration $k+1$, the source node u sends messages of type E that propagate through the network. Nodes at distance $k+1$ from the target w will send an E message at time $t_0 = t - \Delta t_E$.
2. (Inhibitory Control, I-E pair) Tagged nodes at a distance k trigger messages of type I and E at time $t_1 = t + \tau^+_{\text{proc}}$, which reach nodes at distance $k+1$ at times $t_2 = t_1 + \Delta t_I$ and $t_3 = t_1 + \Delta t_E$, respectively.
3. (Tagging Condition) Nodes $v \in \mathcal{V}$ at a distance $k+1$ and which are neighbors of a tagged node receive an I-message at t_2 and an E-message at t_3. Since $t_2 < t_0 + \Delta t_E + \Delta t_I + \tau^0_{\text{proc}}$, $t_3 < t_0 + \Delta t_E + \Delta t_I + \tau^0_{\text{proc}}$, and $t_3 - t \leq \Delta t$, the timing condition for the I-E message pair holds, and all nodes v that sent an E-message at t_0 become tagged.

Thereby, nodes at a distance $k+1$ from the target along the shortest path are tagged. By induction, the algorithm tags all nodes along the shortest path. $\square$

4 Results

We demonstrate our algorithm's behavior through simulated navigational shortest-path tasks, which are both intuitive to visualize and directly analogous to spatial navigation paradigms used in neuroscience research on place cells and grid cells – specialized neurons in the hippocampal formation that encode spatial information (Moser et al., 2015). Our simulation environment consists of neurons $v \in V$, distributed randomly across a spatial domain while maintaining a minimum packing distance p_{min} between any pair of neurons to avoid excessive overlap and approximate the distribution of real place cells. Each neuron $v \in V$ is assigned a two-dimensional spatial coordinate $\mathbf{x}(v)$ to determine network connectivity [3]. Specifically, each neuron connects to all neighbors within a spatial annulus, such that $v \in V$ forms connections with all $u \in V$ satisfying $d^2_{\text{min}} < ||\mathbf{x}(v) - \mathbf{x}(u)||^2_2 < d^2_{\text{max}}$. This connectivity pattern replicates a simplified form of the local transition encoding of grid cells previously proposed by Waniek (2018), or more generally an off-center on-surround receptive field.

All simulations contain 1000 neurons and use $p_{\text{min}} = 0.01, d_{\text{min}} = 0.05, d_{\text{max}} = 0.15$ (in meters). Morevover, $\tau^0_{\text{proc}} = 10.0, \tau^+_{\text{proc}} = 5.0, \Delta t_I = 2.0, \Delta t_E = 5.0$ for processing and axonal delays, with additional timing parameters $\tau_{\text{inh}} = 10.0, \tau_{\text{spike}} = 0.1, \tau_{\text{ref}} = 2.0$ and $\Delta t_{\text{dendritic}} = 1.0$ (all times in ms). These parameters were chosen to fall within biologically plausible ranges while providing sufficient network scale and temporal resolution to clearly visualize the algorithm's dynamics and convergence behaviour.

Results for both a square environment and an A-shaped maze are shown in Figure 2. We visualize the algorithm's convergence dynamics as heat maps with overlaid contour lines, where contours encode spike timing relative to algorithm initiation and colors indicate spiking activity and tagging status. This approach enables direct observation of how the *temporal gradient field* evolves as tagged neurons progressively identify the shortest path from source to target in reverse order. To facilitate comparison across iterations, the number of contour levels remains constant throughout each simulation sequence, enabling clear tracking of convergence behavior and emergence of the optimal routing solution.

Initially, neural activity propagates concentrically outward from the starting location, creating circular contours in the temporal gradient field due to uniform processing delays across all neurons. As the algorithm progresses, tagged neurons with accelerated processing times begin to deform this gradient field, biasing the temporal dynamics toward the target location. Global inhibition from tagged neurons suppresses activity along suboptimal paths, effectively pruning exploration in invalid directions. Upon convergence, when the starting neuron becomes tagged, only those neurons that lie precisely on the shortest paths remain active. The algorithm identifies all nodes on shortest paths from start to target, since multiple shortest paths may exist in unweighted graphs, successfully discovering all optimal routes. Additional simulation results for other mazes, different forms

[3] Coordinates are only used for setting up the network connectivity and visualization. They are not used in the algorithm.

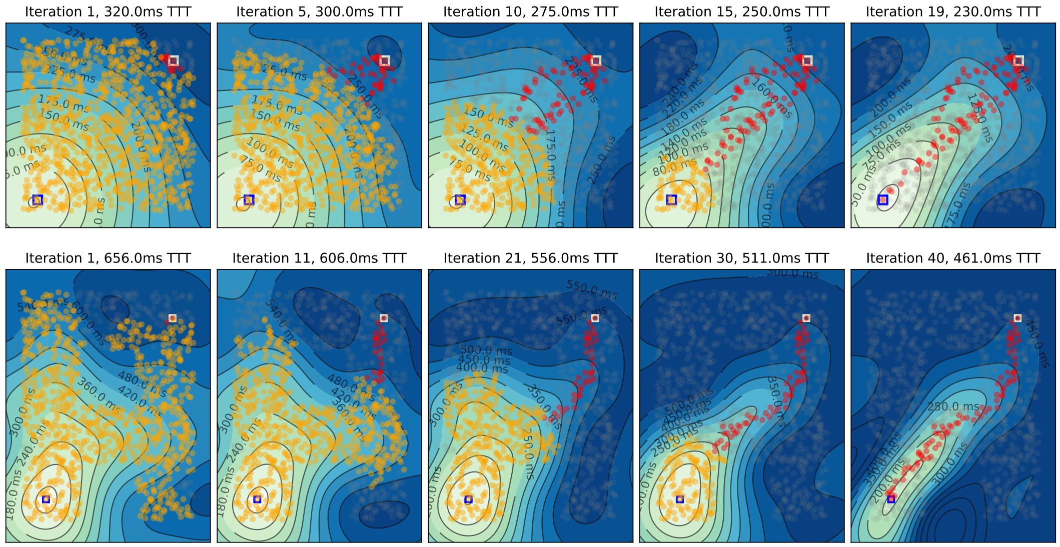

Figure 2: Evolution of the temporal gradient field during iterations of the algorithm. **Top row** shows results for a classical square environment. **Bottom row** shows results for an A-maze, where the spatial domain follows the capital letter A. Starting neuron highlighted by blue box (bottom left corner in each panel), target neuron highlighted by white box (top right). Neurons which spiked during an iteration are colored orange, tagged neurons are red, silent neurons are gray. Contour lines are automatically extracted and indicate time-to-spike for each level. Caption for each panel contains the iteration number, as well as duration during iteration until target neuron spiked. Both rows show the evolution of the temporal gradient field until all nodes on shortest paths are tagged (last column), the decrease in time-to-target (TTT), and the pruning of neural activity along poor directions via global inhibition (silent neurons). The neurons that spike in the final iteration match exactly the shortest-path neurons identified by Dijkstra's algorithm.

of inhibition, and multiple targets can be found in the Appendix. Source code for all simulations is available on github after acceptance.

5 Discussion

Our algorithm demonstrates how local spike-timing predictions can solve shortest path problems without global coordination. The key innovation lies in exploiting temporal coincidences of both excitatory and inhibitory inputs. Specifically, neurons become tagged when receiving inhibitory-excitatory message pairs earlier than expected, creating a backward-propagating wave of transient temporal compression from target to source due to threshold adaptation and without learning in form of synaptic weight changes. This transforms graph traversal from a problem requiring explicit backtracing into one of distributed temporal prediction.

While our algorithm operates on general directed graphs, our simulations utilized locally symmetric neighborhoods in unweighted networks to align with grid cell connectivity patterns observed in spatial navigation. Such symmetric connectivity can emerge naturally during exploration, as demonstrated in the Transition Scale-Space (TSS) model (Waniek, 2018; 2020). This framework suggests that grid cells encode transition relationships across diverse topologies – from Euclidean spaces to general Riemannian manifolds – through biologically plausible spatio-temporal learning kernels. Importantly, our core temporal prediction mechanism generalizes to asymmetric and weighted graphs, indicating broad applicability beyond the spatial domain examined here.

While our algorithm guarantees convergence, its runtime scales linearly with path length (k iterations for length-k sequences), limiting efficiency for long routes. The TSS framework demonstrates how multi-scale hierarchical representations can achieve optimal acceleration of retrieval under biologically plausible constraints that match grid cell representations (Waniek, 2020), but does not address spiking implementations. Combining our spike-based temporal prediction mechanism with

TSS's hierarchical scale-space approach represents a natural next step toward achieving both biological realism and computational efficiency. Preliminary work has already demonstrated the feasibility of this integration, showing significantly accelerated convergence using multiple scales while preserving the distributed, local-computation properties of our algorithm.

A critical aspect of our model is the ability of tagged neurons to overcome inhibition and re-enter processing states. We propose this reflects neuromodulator-mediated disinhibition, in which recently active neurons express transient excitability tags that both lower spike thresholds and weaken GABAergic influence through disinhibitory microcircuits (e.g. VIP $\rightarrow$ SOM pathways) and intrinsic channel modulation (Thiele & Bellgrove, 2018). While conceptually related to synaptic tagging and capture mechanisms (Redondo & Morris, 2011) and threshold plasticity (Pham & Hansel, 2023), our tagging operates on much shorter timescales, corresponding to transient dendritic compartment disinhibition that permits rapid re-engagement with ongoing network dynamics.

Our algorithm's backward propagation of tagged states bears striking resemblance to hippocampal replay phenomena, where spike sequences propagate both forward and backward during sharp-wave ripples (Skaggs & McNaughton, 1996). Recent evidence shows replay occurs not only for experienced trajectories but also for novel, never traversed paths (Gupta et al., 2010; Denovellis et al., 2020), and even during route planning as *preplay* (Pfeiffer & Foster, 2013). These findings suggest the brain may indeed use shortest-path dynamics for path optimization.

Our algorithm demonstrates robust convergence in idealized conditions. Several factors could compromise performance in biological settings, such as network noise, imprecise timing, and heterogeneous neural properties, all of which may degrade the temporal prediction accuracy required for proper tagging. The algorithm's reliance on precise temporal windows means that jitter in synaptic delays or variability in processing times could lead to false tagging or missed detection of early I-E message pairs. However, this jitter could be exploited in future work to further encode path probabilities and other relative information. Additionally, the global inhibition mechanism may become metabolically costly in very large networks, as tagged neurons must broadcast to all other neurons simultaneously. This broadcast requirement could also create informational bottlenecks, particularly in densely connected networks where multiple tagged neurons generate overlapping inhibitory signals. We note that alternatives for inhibition are certainly possible in our algorithm, for instance requiring more than a single node for sending I messages. However, this would require a more elaborate mechanism for inhibitory control, which we excluded for the sake of keeping the timing protocol and its analysis as simple as possible. Furthermore, the algorithm assumes a static network topology during convergence, making it potentially vulnerable to dynamic environments where connections change or obstacles appear during path computation. These limitations suggest that biological implementations may require additional robustness mechanisms, such as temporal averaging of tagging signals or hierarchical organization to reduce global communication overhead, which need to be explored in future work.

Our model makes several testable predictions for biological and behavioral experiments. During spatial navigation tasks, we predict that neurons closer to the goal should show progressively shorter response latencies in successive trials, reflecting the backward propagation of temporal compression from target to source. This latency gradient should emerge even for novel paths, distinguishing our mechanism from simple experience-dependent plasticity. Additionally, neurons on the optimal path should exhibit characteristic timing behavior that differs from neurons on suboptimal routes. In fact, optogenetic experiments could directly test the role of inhibitory signaling by selectively disrupting inhibitory interneurons, which should impair path optimization without preventing basic path finding, while enhancing inhibitory drive should accelerate convergence. Furthermore, our model predicts that replay sequences during rest periods should exhibit the temporal compression signature of tagged neurons, with backward replay showing faster propagation along previously optimized routes. Single-cell recordings could reveal neurons that can overcome inhibition when tagged, displaying reduced spike threshold and altered temporal dynamics. These predictions provide concrete experimental avenues for testing whether biological path-finding systems employ temporal prediction mechanisms similar to our proposed algorithm.

The prevalence of disinhibitory mechanisms in cortical circuits (Reimann et al., 2024; Sridharan & Knudsen, 2015; Williams & Holtmaat, 2019) suggests our approach could be extended to more elaborate timing protocols. Future work should explore how multiple temporal scales and more complex inhibitory dynamics could enable broader classes of graph algorithms to be implemented through local spike-timing predictions, potentially revealing fundamental principles underlying the brain's remarkable computational capabilities.

6 CONCLUSION

We have demonstrated that biologically plausible shortest path computation can emerge from purely local spike-timing dynamics. Our algorithm exploits predictive spike-time coding, where neurons become tagged upon receiving message pairs earlier than anticipated, creating temporal compression that propagates backward from target to source. This transforms distributed graph search into a problem of local temporal prediction.

The biological plausibility stems from reliance on established neural mechanisms such as spike-timing prediction, threshold adaptation, and competitive inhibition. Unlike traditional algorithms that require global state management, our method operates through local message-passing with realistic delays, making it implementable in both biological circuits and neuromorphic hardware. The formal convergence guarantees demonstrate that this is a principled solution, not merely a heuristic.

Our approach to predictive spike time dynamics represents a paradigm that could extend to other algorithms, potentially transforming classical approaches into biologically plausible, locally distributed variants. This could advance both neuroscience and artificial systems for search, planning, and decision making, highlighting that the brain's computational prowess may rely on elegant temporal dynamics and timing codes that emerge from basic neural properties. These principles could also inspire new AI architectures that combine biological realism with computational efficiency.

USE OF LARGE LANGUAGE MODELS

LLMs have been used to polish writing, shorten paragraphs for length constraints, and improve sentence flow while preserving all technical and logical content and references.

REFERENCES

Larry F Abbott and Kenneth I. Blum. Functional Significance of Long-Term Potentiation for Sequence Learning and Prediction. *Cerebral Cortex*, 6(3):406–416, 1996. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/6.3.406.

Moshe Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, Cambridge ; New York, 1991. ISBN 978-0-521-37476-7 978-0-521-37617-4.

Moshe Abeles, Hagai Bergman, Eyal Margalit, and Eilon Vaadia. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, 70(4):1629–1638, 1993.

Moshe Abeles, Hagai Bergman, I. Gat, Isaac Meilijson, E. Seidemann, Naftali Tishby, and Eilon Vaadia. Cortical activity flips among quasi-stationary states. *Proceedings of the National Academy of Sciences of the United States of America*, 92(19):8616–8620, September 1995. ISSN 0027-8424.

Omar Agesen. Type inference in the presence of first-class polymorphism and subtyping. In *SAS / PLDI-era papers on constraint-based analysis (see Agesen 1994 work)*, 1994. Agesen's work is commonly cited for practical constraint-based inference techniques in OO languages.

Filipp Akopyan, Jun Sawada, Andrew Cassidy, Rodrigo Alvarez-Icaza, John Arthur, Paul Merolla, Nabil Imam, Yutaka Nakamura, Pallab Datta, Gi-Joon Nam, Brian Taba, Michael Beakes, Bernard Brezzo, Jente B. Kuang, Rajit Manohar, William P. Risk, Bryan Jackson, and Dharmendra S. Modha. TrueNorth: Design and Tool Flow of a 65 mW 1 Million Neuron Programmable Neurosynaptic Chip. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(10):1537–1557, October 2015. ISSN 0278-0070, 1937-4151. doi: 10.1109/TCAD.2015.2474396.

Thomas H. Alderson, Arun L. W. Bokde, J. A. Scott Kelso, Liam Maguire, and Damien Coyle. Metastable neural dynamics underlies cognitive performance across multiple behavioural paradigms. *Human Brain Mapping*, 41(12):3212–3234, April 2020. ISSN 1065-9471. doi: 10.1002/hbm.25009.

Shun-Ichi Amari. Learning Patterns and Pattern Sequences by Self-Organizing Nets of Threshold Elements. *IEEE Transactions on Computers*, C-21(11):1197–1206, November 1972. ISSN 0018-9340. doi: 10.1109/T-C.1972.223477.

Daniel J Amit. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral Cortex*, 7(3):237–252, April 1997. ISSN 14602199. doi: 10.1093/cercor/7.3.237.

Rony Azouz and Charles M. Gray. Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proceedings of the National Academy of Sciences of the United States of America*, 97(14):8110–8115, July 2000. ISSN 0027-8424. doi: 10.1073/pnas.130200797.

Jacob J. W. Bakermans, Joseph Warren, James C. R. Whittington, and Timothy E. J. Behrens. Constructing future behavior in the hippocampal formation through composition and replay. *Nature Neuroscience*, 28(5):1061–1072, May 2025a. ISSN 1097-6256, 1546-1726. doi: 10.1038/s41593-025-01908-3.

Jacob J. W. Bakermans, Joseph Warren, James C. R. Whittington, and Timothy E. J. Behrens. Constructing future behavior in the hippocampal formation through composition and replay. *Nature Neuroscience*, 28(5):1061–1072, May 2025b. ISSN 1097-6256, 1546-1726. doi: 10.1038/s41593-025-01908-3.

Andrea Banino, Caswell Barry, Benigno Uria, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin J. Chadwick, Thomas Degris, Joseph Modayil, Greg Wayne, Hubert Soyer, Fabio Viola, Brian Zhang, Ross Goroshin, Neil Rabinowitz, Razvan Pascanu, Charlie Beattie, Stig Petersen, Amir Sadik, Stephen Gaffney, Helen King, Koray Kavukcuoglu, Demis Hassabis, Raia Hadsell, and Dharshan Kumaran. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, May 2018. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-018-0102-6.

Hannah Bast, Daniel Delling, Andrew Goldberg, Matthias Müller-Hannemann, Thomas Pajor, Peter Sanders, Dorothea Wagner, and Renato F. Werneck. *Route Planning in Transportation Networks*, pp. 19–80. Springer International Publishing, 2016. ISBN 9783319494876. doi: 10.1007/978-3-319-49487-6_2. URL http://dx.doi.org/10.1007/978-3-319-49487-6_2.

Leonard E. Baum, Ted Petrie, George Soules, and Norman Weiss. A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *The Annals of Mathematical Statistics*, 41(1):164–171, 1970. ISSN 00034851, 21688990. URL http://www.jstor.org/stable/2239727.

Guillaume Bellec, Franz Scherr, Anand Subramoney, Elias Hajek, Darjan Salaj, Robert Legenstein, and Wolfgang Maass. A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications*, 11(1):3625, July 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-17236-y.

Richard Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16(1):87–90, 1958.

Guo-qiang Bi and Mu-ming Poo. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *The Journal of Neuroscience*, 18(24):10464–10472, December 1998. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.18-24-10464.1998.

Johannes Blum. Hierarchy of transportation network parameters and hardness results. In *14th International Symposium on Parameterized and Exact Computation*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi: 10.4230/LIPICS.IPEC.2019.4. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.IPEC.2019.4.

Younes Bouhadjar, Dirk J. Wouters, Markus Diesmann, and Tom Tetzlaff. Sequence learning, prediction, and replay in networks of spiking neurons. *PLOS Computational Biology*, 18(6): e1010233, June 2022. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1010233.

Braden A. W. Brinkman, Han Yan, Arianna Maffei, Il Memming Park, Alfredo Fontanini, Jin Wang, and Giancarlo La Camera. Metastable dynamics of neural circuits and networks. *Applied Physics Reviews*, 9(1):011313, March 2022. ISSN 1931-9401. doi: 10.1063/5.0062603.

Christopher L. Buckley and Thomas Nowotny. Transient dynamics between displaced fixed points: An alternate nonlinear dynamical framework for olfaction. *Brain Research*, 1434:62–72, January 2012. ISSN 00068993. doi: 10.1016/j.brainres.2011.07.032.

Daniel Bush. Spike-timing dependent plasticity and the cognitive map. *Frontiers in Computational Neuroscience*, 4, 2010. ISSN 16625188. doi: 10.3389/fncom.2010.00142.

György Buzsáki and Rodolfo Llinás. Space and time in the brain. *Science*, 358(6362):482–485, October 2017. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aan8869.

György Buzsáki and David Tingley. Space and Time: The Hippocampus as a Sequence Generator. *Trends in Cognitive Sciences*, 22(10):853–869, October 2018. ISSN 13646613. doi: 10.1016/j.tics.2018.07.006.

Natalia Caporale and Yang Dan. Spike Timing–Dependent Plasticity: A Hebbian Learning Rule. *Annual Review of Neuroscience*, 31(1):25–46, July 2008. ISSN 0147-006X, 1545-4126. doi: 10.1146/annurev.neuro.31.060407.125639.

Thomas L. Casavant and Jon G. Kuhl. A taxonomy of scheduling in general-purpose distributed computing systems. *IEEE Transactions on Software Engineering*, 14(2):141–154, 1988. doi: 10.1109/32.4634.

Liyu Chen, Haipeng Luo, and Chen-Yu Wei. Minimax regret for stochastic shortest path with adversarial costs and known transition. In Mikhail Belkin and Samory Kpotufe (eds.), *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 1180–1215. PMLR, 15–19 Aug 2021. URL `https://proceedings.mlr.press/v134/chen21e.html`.

Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, 3rd edition, 2009. ISBN 978-0262033848.

Christopher J Cueva and Xue-Xin Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. In *International Conference for Learning Representations*, pp. 19, 2018.

Max Dabagia, Christos H. Papadimitriou, and Santosh S. Vempala. Computation With Sequences of Assemblies in a Model of the Brain. *Neural Computation*, 37(1):193–233, December 2024. ISSN 0899-7667, 1530-888X. doi: 10.1162/neco_a_01720.

Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, Yuyun Liao, Chit-Kwan Lin, Andrew Lines, Ruokun Liu, Deepak Mathaikutty, Steven McCoy, Arnab Paul, Jonathan Tse, Guruguhanathan Venkataramanan, Yi-Hsin Weng, Andreas Wild, Yoonseok Yang, and Hong Wang. Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro*, 38(1):82–99, January 2018. ISSN 0272-1732, 1937-4143. doi: 10.1109/MM.2018.112130359.

Daniel Delling, Martin Holzer, Kirill Muller, Dorothea Wagner, and Frank Schulz. High-Performance Multi-Level Graphs. In *The Shortest Path Problem*, pp. 73–91. American Mathematical Society, 2009.

Eric L. Denovellis, Anna K. Gillespie, Michael E. Coulter, Marielena Sosa, Jason E. Chung, Uri T. Eden, and Loren M. Frank. Hippocampal replay of experience at real-world speeds. Preprint, Neuroscience, October 2020.

Kamran Diba, Asohan Amarasingham, Kenji Mizuseki, and György Buzsáki. Millisecond Timescale Synchrony among Hippocampal Neurons. *Journal of Neuroscience*, 34(45):14984–14994, November 2014. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1091-14.2014.

Julian Dibbelt, Ben Strasser, and Dorothea Wagner. *Customizable Contraction Hierarchies*, pp. 271–282. Springer International Publishing, 2014. ISBN 9783319079592. doi: 10.1007/978-3-319-07959-2_23. URL http://dx.doi.org/10.1007/978-3-319-07959-2_23.

Robin Dietrich, Tobias Fischer, Nicolai Waniek, Nico Reeb, Michael Milford, Alois Knoll, and Adam D. Hines. Threshold Adaptation in Spiking Networks Enables Shortest Path Finding and Place Disambiguation. In *2025 Neuro Inspired Computational Elements (NICE)*, pp. 1–11, March 2025. doi: 10.1109/NICE65350.2025.11065806.

Edsger W. Dijkstra. A note on two problems in connexion with graphs. 1:269–271, 1959.

Rodney J. Douglas and Kevan A.C. Martin. Neuronal Circuits of the Neocortex. *Annual Review of Neuroscience*, 27(1):419–451, July 2004. ISSN 0147-006X, 1545-4126. doi: 10.1146/annurev.neuro.27.070203.144152.

George Dragoi. Cell assemblies, sequences and temporal coding in the hippocampus. *Current Opinion in Neurobiology*, 64:111–118, October 2020. ISSN 09594388. doi: 10.1016/j.conb.2020.03.003.

George Dragoi and György Buzsáki. Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies. 50(1):145–157, 2006. ISSN 08966273. doi: 10.1016/j.neuron.2006.02.023. URL https://linkinghub.elsevier.com/retrieve/pii/S0896627306001656.

George Dragoi and Susumu Tonegawa. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330):397–401, 2011. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature09633.

Daniel Durstewitz and Gustavo Deco. Computational significance of transient dynamics in cortical networks. *European Journal of Neuroscience*, 27(1):217–227, December 2007. ISSN 0953816X, 14609568. doi: 10.1111/j.1460-9568.2007.05976.x.

Ilyas Fatkhullin, Anas Barakat, Anastasia Kireeva, and Niao He. Stochastic policy gradient methods: Improved sample complexity for Fisher-non-degenerate policies. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 9827–9869. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/fatkhullin23a.html.

Bertrand Fontaine, José Luis Peña, and Romain Brette. Spike-Threshold Adaptation Predicted by Membrane Potential Dynamics In Vivo. *PLoS Computational Biology*, 10(4):e1003560, April 2014. ISSN 1553-734X. doi: 10.1371/journal.pcbi.1003560.

Lester R. Jr. Ford. Network flow theory. Technical Report P-923, RAND Corporation, Santa Monica, California, August 1956.

Karl J. Friston. Transients, Metastability, and Neuronal Dynamics. *NeuroImage*, 5(2):164–171, February 1997. ISSN 10538119. doi: 10.1006/nimg.1997.0259.

Stefan Funke and Sabine Storandt. Provable Efficiency of Contraction Hierarchies with Randomized Preprocessing. In Khaled Elbassioni and Kazuhisa Makino (eds.), *Algorithms and Computation*, volume 9472, pp. 479–490. Springer Berlin Heidelberg, 2015. ISBN 978-3-662-48970-3 978-3-662-48971-0. doi: 10.1007/978-3-662-48971-0_41. URL http://link.springer.com/10.1007/978-3-662-48971-0_41.

Steve B. Furber, Francesco Galluppi, Steve Temple, and Luis A. Plana. The SpiNNaker Project. *Proceedings of the IEEE*, 102(5):652–665, May 2014. ISSN 0018-9219, 1558-2256. doi: 10.1109/JPROC.2014.2304638.

Robert Geisberger, Peter Sanders, Dominik Schultes, and Daniel Delling. Contraction Hierarchies: Faster and Simpler Hierarchical Routing in Road Networks. In Catherine C. McGeoch (ed.), *Experimental Algorithms*, volume 5038, pp. 319–333. Springer Berlin Heidelberg, 2008. ISBN 978-3-540-68548-7 978-3-540-68552-4. doi: 10.1007/978-3-540-68552-4_24. URL http://link.springer.com/10.1007/978-3-540-68552-4_24.

Matthieu Gilson. STDP in recurrent neuronal networks. *Frontiers in Computational Neuroscience*, 4, 2010. ISSN 16625188. doi: 10.3389/fncom.2010.00023.

Goetz Graefe. Query evaluation techniques for large databases. *ACM Computing Surveys / ACM SIGMOD Record (widely circulated as a survey/technical report)*, 1993. URL https://cs.uwaterloo.ca/~david/cs848s13/graefe.pdf. Survey of query execution/optimization techniques; discusses operator order, pipelining, and scheduling.

Patricia Griffiths Selinger, Morton M. Astrahan, Donald Dean Chamberlin, Raymond A. Lorie, and Thomas Gordon Price. Access path selection in a relational database management system. In *Proceedings of the 1979 ACM SIGMOD International Conference on Management of Data*, pp. 23–34. ACM, 1979. doi: 10.1145/582095.582099.

Anoopum S. Gupta, Matthijs A. A. van der Meer, David S. Touretzky, and A. David Redish. Hippocampal Replay is Not a Simple Function of Experience. *Neuron*, 65(5):695–705, March 2010. ISSN 0896-6273. doi: 10.1016/j.neuron.2010.01.034.

Bernhard Haeupler, Richard Hladík, Václav Rozhoň, Robert Tarjan, and Jakub Tětek. Universal Optimality of Dijkstra via Beyond-Worst-Case Heaps, 2024. URL http://arxiv.org/abs/2311.11793.

Kenneth D Harris and Gordon M G Shepherd. The neocortical circuit: Themes and variations. *Nature Neuroscience*, 18(2):170–181, February 2015. ISSN 1097-6256, 1546-1726. doi: 10.1038/nn.3917.

Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. 4(2):100–107, 1968. ISSN 2168-2887. doi: 10.1109/TSSC.1968.300136. URL https://ieeexplore.ieee.org/document/4082128/?arnumber=4082128.

Darrel A. Henze and György Buzsáki. Action potential threshold of hippocampal pyramidal cells in vivo is increased by recent spiking activity. *Neuroscience*, 105(1):121–130, July 2001. ISSN 03064522. doi: 10.1016/S0306-4522(01)00167-1.

Hideyoshi Igata, Yuji Ikegaya, and Takuya Sasaki. Prioritized experience replays on a hippocampal predictive map for learning. *Proceedings of the National Academy of Sciences*, 118(1): e2011266118, January 2021. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.2011266118.

Andrew Ilyas, Logan Engstrom, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. A closer look at deep policy gradients. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=ryxdEkHtPS.

Haque Ishfaq, Guangyuan Wang, Sami Nur Islam, and Doina Precup. Langevin soft actor-critic: Efficient exploration through uncertainty-driven critic learning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=FvQsk3la17.

Eugene M. Izhikevich and Frank C. Hoppensteadt. Polychronous Wavefront Computations. *International Journal of Bifurcation and Chaos*, 19(05):1733–1739, May 2009. ISSN 0218-1274, 1793-6551. doi: 10.1142/S0218127409023809.

Mehdi Jafarnia-Jahromi, Liyu Chen, Rahul Jain, and Haipeng Luo. Posterior sampling-based online learning for the stochastic shortest path model. In Robin J. Evans and Ilya Shpitser (eds.), *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pp. 922–931. PMLR, 31 Jul–04 Aug 2023. URL https://proceedings.mlr.press/v216/jafarnia-jahromi23a.html.

Chetan Kadway, Sounak Dey, Arijit Mukherjee, Arpan Pal, and Gilles Bézard. Low Power & Low Latency Cloud Cover Detection in Small Satellites Using On-board Neuromorphic Processors. In *2023 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, Gold Coast, Australia, June 2023. IEEE. ISBN 978-1-66548-867-9. doi: 10.1109/IJCNN54540.2023.10191569.

Lydia E. Kavraki, P. Svestka, Jean-Claude Latombe, and Mark H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. 12(4):566–580, 1996. ISSN 1042296X. doi: 10.1109/70.508439. URL `http://ieeexplore.ieee.org/document/508439/`.

Christoph Kayser, Nikos K. Logothetis, and Stefano Panzeri. Millisecond encoding precision of auditory cortex neurons. *Proceedings of the National Academy of Sciences*, 107(39):16976–16981, September 2010. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1012656107.

Christof Koch. *Biophysics of Computation: Information Processing in Single Neurons*, volume 74. Oxford University Press, 2004.

Daniel Koch, Akhilesh Nandan, Gayathri Ramesan, and Aneta Koseska. Biological computations: Limitations of attractor-based formalisms and the need for transients, April 2024.

Jeffrey L. Krichmar, Nicholas A. Ketz, Praveen K. Pilly, and Andrea Soltoggio. Flexible Path Planning in a Spiking Model of Replay and Vicarious Trial and Error. In Lola Cañamero, Philippe Gaussier, Myra Wilson, Sofiane Boucenna, and Nicolas Cuperlier (eds.), *From Animals to Animats 16*, Lecture Notes in Computer Science, pp. 177–189. Springer International Publishing, 2022. ISBN 978-3-031-16770-6. doi: 10.1007/978-3-031-16770-6_15.

Yu-Kwong Kwok and Ishfaq Ahmad. Static scheduling algorithms for allocating directed task graphs to multiprocessors. *ACM Computing Surveys*, 31(4):406–471, 1999. doi: 10.1145/344588.344618.

Giancarlo La Camera, Alfredo Fontanini, and Luca Mazzucato. Cortical computations via metastable activity. *Current opinion in neurobiology*, 58:37–45, October 2019. ISSN 0959-4388. doi: 10.1016/j.conb.2019.06.007.

Jean-Claude Latombe. *Robot Motion Planning*. Springer US, Boston, MA, 1991. ISBN 978-0-7923-9206-4 978-1-4615-4022-9. doi: 10.1007/978-1-4615-4022-9.

Amir Levi, Lidor Spivak, Hadas E. Sloin, Shirly Someck, and Eran Stark. Error correction and improved precision of spike timing in converging cortical networks. *Cell Reports*, 40(12):111383, September 2022. ISSN 22111247. doi: 10.1016/j.celrep.2022.111383.

Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman, and Geoffrey Hinton. Backpropagation and the brain. 21(6):335–346, 2020. ISSN 1471-003X, 1471-0048. doi: 10.1038/s41583-020-0277-3. URL `http://www.nature.com/articles/s41583-020-0277-3`.

Tomás Lozano-Pérez and Michael A. Wesley. An algorithm for planning collision-free paths among polyhedral obstacles. *Communications of the ACM*, 22(10):560–570, October 1979. ISSN 0001-0782, 1557-7317. doi: 10.1145/359156.359164.

Wolfgang Maass. Networks of spiking neurons: The third generation of neural network models. 10 (9):1659–1671, 1997. ISSN 08936080. doi: 10.1016/S0893-6080(97)00011-7. URL `https://linkinghub.elsevier.com/retrieve/pii/S0893608097000117`.

Amgad Madkour, Walid G. Aref, Faizan Ur Rehman, Mohamed Abdur Rahman, and Saleh Basalamah. A Survey of Shortest-Path Algorithms. *arXiv:1705.02044 [cs]*, May 2017.

Zachary F. Mainen and Terrence J. Sejnowski. Reliability of Spike Timing in Neocortical Neurons. *Science*, 268(5216):1503–1506, June 1995. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.7770778.

Valerio Mante, David Sussillo, Krishna V. Shenoy, and William T. Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, November 2013. ISSN 0028-0836. doi: 10.1038/nature12742.

Nikola T Markov and Henry Kennedy. The importance of being hierarchical. *Current Opinion in Neurobiology*, 23(2):187–194, April 2013. ISSN 09594388. doi: 10.1016/j.conb.2012.12.008.

Henry Markram, Wulfram Gerstner, and Per Jesper Sjöström. Spike-Timing-Dependent Plasticity: A Comprehensive Overview. *Frontiers in Synaptic Neuroscience*, 4, 2012. ISSN 1663-3563. doi: 10.3389/fnsyn.2012.00002.

Timothée Masquelier, Rudy Guyonneau, and Simon J. Thorpe. Spike Timing Dependent Plasticity Finds the Start of Repeating Patterns in Continuous Spike Trains. *PLoS ONE*, 3(1):e1377, January 2008. ISSN 1932-6203. doi: 10.1371/journal.pone.0001377.

Ofer Mazor and Gilles Laurent. Transient Dynamics versus Fixed Points in Odor Representations by Locust Antennal Lobe Projection Neurons. *Neuron*, 48(4):661–673, November 2005. ISSN 08966273. doi: 10.1016/j.neuron.2005.09.032.

Daniel C. McNamee, Kimberly L. Stachenfeld, Matthew M. Botvinick, and Samuel J. Gershman. Flexible modulation of sequence generation in the entorhinal–hippocampal system. *Nature Neuroscience*, 24(6):851–862, June 2021. ISSN 1097-6256, 1546-1726. doi: 10.1038/s41593-021-00831-7.

Marcelo A. Montemurro, Stefano Panzeri, Miguel Maravall, Andrea Alenda, Michael R. Bale, Marco Brambilla, and Rasmus S. Petersen. Role of Precise Spike Timing in Coding of Dynamic Vibrissa Stimuli in Somatosensory Thalamus. *Journal of Neurophysiology*, 98(4):1871–1882, October 2007. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00593.2007.

Edward F Moore. The shortest path through a maze. In *Proceedings of an International Symposium on the Theory of Switching*, pp. 285–292. Harvard University Press, 1959.

May-Britt Moser, David C. Rowland, and Edvard I. Moser. Place Cells, Grid Cells, and Memory. 7(2):a021808, 2015. ISSN 1943-0264. doi: 10.1101/cshperspect.a021808. URL `http://cshperspectives.cshlp.org/lookup/doi/10.1101/cshperspect.a021808`.

Klaus Oberauer, Stephan Lewandowsky, Edward Awh, Gordon D. A. Brown, Andrew Conway, Nelson Cowan, Christopher Donkin, Simon Farrell, Graham J. Hitch, Mark J. Hurlstone, Wei Ji Ma, Candice C. Morey, Derek Evan Nee, Judith Schweppe, Evie Vergauwe, and Geoff Ward. Benchmarks for models of short-term and working memory. *Psychological Bulletin*, 144(9): 885–958, September 2018. ISSN 0033-2909. doi: 10.1037/bul0000153. URL `http://dx.doi.org/10.1037/bul0000153`.

Yuval Orsher, Ariel Rom, Rotem Perel, Yoav Lahini, Pablo Blinder, and Mark Shein-Idelson. Travelling waves or sequentially activated discrete modules: mapping the granularity of cortical propagation. February 2024. doi: 10.7554/elife.92254.2. URL `http://dx.doi.org/10.7554/eLife.92254.2`.

Adam M. Packer and Rafael Yuste. Dense, Unspecific Connectivity of Neocortical Parvalbumin-Positive Interneurons: A Canonical Microcircuit for Inhibition? *The Journal of Neuroscience*, 31 (37):13260–13271, September 2011. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.3131-11.2011.

Andrew E. Papale, Mark C. Zielinski, Loren M. Frank, Shantanu P. Jadhav, and A. David Redish. Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making. *Neuron*, 92(5):975–982, December 2016. ISSN 08966273. doi: 10.1016/j.neuron.2016.10.028.

Eloy Parra-Barrero, Kamran Diba, and Sen Cheng. Neuronal sequences during theta rely on behavior-dependent spatial maps. *eLife*, 10:e70296, October 2021. ISSN 2050-084X. doi: 10.7554/eLife.70296.

Jing Pei, Lei Deng, Sen Song, Mingguo Zhao, Youhui Zhang, Shuang Wu, Guanrui Wang, Zhe Zou, Zhenzhi Wu, Wei He, Feng Chen, Ning Deng, Si Wu, Yu Wang, Yujie Wu, Zheyu Yang, Cheng Ma, Guoqi Li, Wentao Han, Huanglong Li, Huaqiang Wu, Rong Zhao, Yuan Xie, and Luping Shi. Towards artificial general intelligence with hybrid Tianjic chip architecture. *Nature*, 572(7767): 106–111, August 2019. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-019-1424-8.

Brad E. Pfeiffer and David J. Foster. Hippocampal place-cell sequences depict future paths to remembered goals. 497(7447):74–79, 2013. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature12112. URL `http://www.nature.com/articles/nature12112`.

Tuan Pham and Christian Hansel. Intrinsic threshold plasticity: Cholinergic activation and role in the neuronal recognition of incomplete input patterns. *The Journal of Physiology*, 601(15): 3221–3239, August 2023. ISSN 0022-3751, 1469-7793. doi: 10.1113/JP283473.

Christoph Pokorny, Matias J Ison, Arjun Rao, Robert Legenstein, Christos Papadimitriou, and Wolfgang Maass. STDP Forms Associations between Memory Traces in Networks of Spiking Neurons. *Cerebral Cortex*, 30(3):952–968, March 2020. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhz140.

Adrián Ponce-Alvarez, Verónica Nácher, Rogelio Luna, Alexa Riehle, and Ranulfo Romo. Dynamics of Cortical Neuronal Ensembles Transit from Decision Making to Storage for Later Report. *The Journal of Neuroscience*, 32(35):11956–11969, August 2012. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.6176-11.2012.

Filip Ponulak and John J. Hopfield. Rapid, parallel path planning by propagating wavefronts of spiking neural activity. 7, 2013. ISSN 1662-5188. doi: 10.3389/fncom.2013.00098. URL `http://journal.frontiersin.org/article/10.3389/fncom.2013.00098/abstract`.

François Pottier and Didier Rémy. A framework for type inference with subtyping. In *ICFP (technical report / conference paper)*, 1998. See Pottier's later journal/paper versions on constraint-based inference and HM(X) for detailed algorithms.

M L Puterman. *Markov decision processes*. Wiley Series in Probability & Mathematical Statistics: Applied Probability & Statistics. John Wiley & Sons, Nashville, TN, May 1994.

Mikhail I. Rabinovich, Pablo Varona, Allen I. Selverston, and Henry D. I. Abarbanel. Dynamical principles in neuroscience. *Reviews of Modern Physics*, 78(4):1213–1265, November 2006. ISSN 0034-6861, 1539-0756. doi: 10.1103/RevModPhys.78.1213.

Roger L. Redondo and Richard G. M. Morris. Making memories last: The synaptic tagging and capture hypothesis. *Nature Reviews Neuroscience*, 12(1):17–30, January 2011. ISSN 1471-003X, 1471-0048. doi: 10.1038/nrn2963.

Michael W Reimann, Daniela Egas Santander, András Ecker, and Eilif B Muller. Specific inhibition and disinhibition in the higher-order structure of a cortical connectome. *Cerebral Cortex*, 34(11): bhae433, November 2024. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhae433.

Alpha Renner, Forrest Sheldon, Anatoly Zlotnik, Louis Tao, and Andrew Sornborger. The backpropagation algorithm implemented on spiking neuromorphic hardware. 15(1):9691, 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-53827-9. URL `https://www.nature.com/articles/s41467-024-53827-9`.

Ole Richter, Chenxi Wu, Adrian M Whatley, German Köstinger, Carsten Nielsen, Ning Qiao, and Giacomo Indiveri. DYNAP-SE2: A scalable multi-core dynamic neuromorphic asynchronous spiking neural network processor. *Neuromorphic Computing and Engineering*, 4(1):014003, March 2024. ISSN 2634-4386. doi: 10.1088/2634-4386/ad1cd7.

Edmund T. Rolls, Leonardo Franco, Nikolaos C. Aggelopoulos, and Jose M. Jerez. Information in the first spike, the order of spikes, and the number of spikes provided by neurons in the inferior temporal visual cortex. *Vision Research*, 46(25):4193–4205, November 2006. ISSN 00426989. doi: 10.1016/j.visres.2006.07.026.

Matthew Rosenberg, Tony Zhang, Pietro Perona, and Markus Meister. Mice in a labyrinth show rapid learning, sudden insight, and efficient exploration. *eLife*, 10:e66175, July 2021. ISSN 2050-084X. doi: 10.7554/eLife.66175.

Ulrich Roth, Marc Walker, Arne Hilmann, and Heinrich Klar. Dynamic path planning with spiking neural networks. In José Mira, Roberto Moreno-Díaz, and Joan Cabestany (eds.), *Biological and Artificial Computation: From Neuroscience to Technology*, pp. 1355–1363. Springer, 1997. ISBN 978-3-540-69074-0. doi: 10.1007/BFb0032596.

Rufin Van Rullen and Simon J. Thorpe. Rate Coding Versus Temporal Order Coding: What the Retinal Ganglion Cells Tell the Visual Cortex. *Neural Computation*, 13(6):1255–1283, June 2001. ISSN 0899-7667, 1530-888X. doi: 10.1162/08997660152002852.

Alexei V. Samsonovich and Giorgio A. Ascoli. A simple neural network model of the hippocampus suggesting its pathfinding role in episodic memory retrieval. 12(2):193–208, 2005. ISSN 1072-0502. doi: 10.1101/lm.85205. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1074338/`.

Sven Schrader, Sonja Grün, Markus Diesmann, and George L. Gerstein. Detecting Synfire Chain Activity Using Massively Parallel Spike Train Recording. *Journal of Neurophysiology*, 100(4): 2165–2176, October 2008. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.01245.2007.

Catherine D. Schuman, Kathleen Hamilton, Tiffany Mintz, Md Musabbir Adnan, Bon Woong Ku, Sung-Kyu Lim, and Garrett S. Rose. Shortest Path and Neighborhood Subgraph Extraction on a Spiking Memristive Neuromorphic Implementation. In *Proceedings of the 7th Annual Neuro-inspired Computational Elements Workshop*, NICE '19, pp. 1–6. Association for Computing Machinery, 2019. ISBN 978-1-4503-6123-1. doi: 10.1145/3320288.3320290. URL `https://doi.org/10.1145/3320288.3320290`.

Philipp Schwartenbeck, Alon Baram, Yunzhe Liu, Shirley Mark, Timothy Muller, Raymond Dolan, Matthew Botvinick, Zeb Kurth-Nelson, and Timothy Behrens. Generative replay underlies compositional inference in the hippocampal-prefrontal circuit. *Cell*, 186(22):4885–4897.e14, October 2023. ISSN 00928674. doi: 10.1016/j.cell.2023.09.004.

Philip Seliger, Lev S. Tsimring, and Mikhail I. Rabinovich. Dynamical model of sequential spatial memory: Winnerless competition of patterns. *Physical Review E*, 67(1):011905, January 2003. ISSN 1063-651X, 1095-3787. doi: 10.1103/PhysRevE.67.011905.

H. Sebastian Seung. How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences*, 93(23):13339–13344, November 1996. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.93.23.13339.

Alfonso Shimbel. Structure in communication nets. In *Proceedings of the Symposium on Information Networks (New York, 1954)*, pp. 199–203, Brooklyn, New York, 1955. Polytechnic Press of the Polytechnic Institute of Brooklyn.

Tomer Shmiel, Rotem Drori, Oren Shmiel, Yoram Ben-Shaul, Zoltan Nadasdy, Moshe Shemesh, Mina Teicher, and Moshe Abeles. Temporally Precise Cortical Firing Patterns Are Associated With Distinct Action Segments. *Journal of Neurophysiology*, 96(5):2645–2652, November 2006. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00798.2005.

William E. Skaggs and Bruce L. McNaughton. Replay of Neuronal Firing Sequences in Rat Hippocampus During Sleep Following Spatial Experience. *Science*, 271(5257):1870–1873, March 1996. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.271.5257.1870.

Ben Sorscher, Gabriel C. Mel, Samuel A. Ocko, Lisa M. Giocomo, and Surya Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, 111(1):121–137.e13, January 2023. ISSN 08966273. doi: 10.1016/j.neuron.2022.10.003.

Devarajan Sridharan and Eric I. Knudsen. Selective disinhibition: A unified neural mechanism for predictive and post hoc attentional selection. *Vision Research*, 116:194–209, November 2015. ISSN 00426989. doi: 10.1016/j.visres.2014.12.010.

Kyle H. Srivastava, Caroline M. Holmes, Michiel Vellema, Andrea R. Pack, Coen P. H. Elemans, Ilya Nemenman, and Samuel J. Sober. Motor control by precisely timed spike patterns. *Proceedings of the National Academy of Sciences*, 114(5):1171–1176, January 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1611734114.

Ana Stanojevic, Stanisław Woźniak, Guillaume Bellec, Giovanni Cherubini, Angeliki Pantazi, and Wulfram Gerstner. High-performance deep spiking neural networks with 0.3 spikes per neuron. *Nature Communications*, 15(1):6793, August 2024. ISSN 2041-1723. doi: 10.1038/s41467-024-51110-5.

Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.

Botond Szatmáry and Eugene M. Izhikevich. Spike-Timing Theory of Working Memory. *PLoS Computational Biology*, 6(8):e1000879, August 2010. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1000879.

Robert E. Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1 (2):146–160, 1972. doi: 10.1137/0201010.

Alexander Thiele and Mark A. Bellgrove. Neuromodulation of Attention. *Neuron*, 97(4):769–785, February 2018. ISSN 08966273. doi: 10.1016/j.neuron.2018.01.008.

Simon Thorpe and Jacques Gautrais. Rank Order Coding. In James M. Bower (ed.), *Computational Neuroscience*, pp. 113–118. Springer US, Boston, MA, 1998. ISBN 978-1-4613-7190-8 978-1-4615-4831-7. doi: 10.1007/978-1-4615-4831-7_19.

Emmanuelle Tognoli and J. A. Scott Kelso. The Metastable Brain. *Neuron*, 81(1):35–48, January 2014. ISSN 08966273. doi: 10.1016/j.neuron.2013.12.022.

Manuel Traub, Martin V. Butz, R. Harald Baayen, and Sebastian Otte. Learning precise spike timings with eligibility traces. In Igor Farkaš, Paolo Masulli, and Stefan Wermter (eds.), *Artificial Neural Networks and Machine Learning – ICANN 2020*, pp. 659–669, Cham, 2020. Springer International Publishing. ISBN 978-3-030-61616-8.

Ruben Vale, Dominic A. Evans, and Tiago Branco. Rapid Spatial Learning Controls Instinctive Defensive Behavior in Mice. *Current Biology*, 27(9):1342–1349, May 2017. ISSN 09609822. doi: 10.1016/j.cub.2017.03.031.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. *arXiv:1706.03762 [cs]*, December 2017.

Mengni Wang, David J. Foster, and Brad E. Pfeiffer. Alternating sequences of future and past behavior encoded within hippocampal theta oscillations. *Science*, 370(6513):247–250, October 2020a. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.abb4151.

Shengjie Wang, Shaohuai Liu, Weirui Ye, Jiacheng You, and Yang Gao. Efficientzero v2: mastering discrete and continuous control with limited data. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Yongyue Wang, Beitong Yao, Tianbo Wang, Chunhe Xia, and Xianghui Zhao. A cognitive method for automatically retrieving complex information on a large scale. *Sensors (Basel)*, 20(11):3057, May 2020b.

Nicolai Waniek. Hexagonal Grid Fields Optimally Encode Transitions in Spatiotemporal Sequences. 30(10):2691–2725, 2018. ISSN 0899-7667, 1530-888X. doi: 10.1162/neco_a_01122. URL `https://www.mitpressjournals.org/doi/abs/10.1162/neco_a_01122`.

Nicolai Waniek. Transition Scale-Spaces: A Computational Theory for the Discretized Entorhinal Cortex. 32(2):330–394, 2020. ISSN 0899-7667, 1530-888X. doi: 10.1162/neco_a_01255. URL `https://www.mitpressjournals.org/doi/abs/10.1162/neco_a_01255`.

James C. R. Whittington and Rafal Bogacz. An Approximation of the Error Backpropagation Algorithm in a Predictive Coding Network with Local Hebbian Synaptic Plasticity. *Neural Computation*, 29(5):1229–1262, May 2017. ISSN 0899-7667, 1530-888X. doi: 10.1162/NECO_a_00949.

James C.R. Whittington, Timothy H. Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy E.J. Behrens. The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell*, 183(5):1249–1263.e23, November 2020. ISSN 00928674. doi: 10.1016/j.cell.2020.10.024.

John Widloski and David J. Foster. Flexible rerouting of hippocampal replay sequences around changing barriers in the absence of global place field remapping. *Neuron*, pp. S089662732200109X, February 2022. ISSN 08966273. doi: 10.1016/j.neuron.2022.02.002.

Leena E. Williams and Anthony Holtmaat. Higher-Order Thalamocortical Inputs Gate Synaptic Long-Term Potentiation via Disinhibition. *Neuron*, 101(1):91–102.e4, January 2019. ISSN 08966273. doi: 10.1016/j.neuron.2018.10.049.

Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256, May 1992. ISSN 1573-0565. doi: 10.1007/BF00992696.

Alexander B. Wiltschko, Matthew J. Johnson, Giuliano Iurilli, Ralph E. Peterson, Jesse M. Katon, Stan L. Pashkovski, Victoria E. Abraira, Ryan P. Adams, and Sandeep Robert Datta. Mapping Sub-Second Structure in Mouse Behavior. *Neuron*, 88(6):1121–1135, December 2015. ISSN 0896-6273. doi: 10.1016/j.neuron.2015.11.031.

Pan Xu, Felicia Gao, and Quanquan Gu. Sample efficient policy gradient methods with recursive variance reduction. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=HJlxIJBFDr`.

Jerneja Zganec Gros and Mario Zganec. An efficient unit-selection method for concatenative text-to-speech synthesis systems. *J. Comput. Inf. Technol.*, 16(1):69, 2008.

Michaël B. Zugaro, Angelo Arleo, Alain Berthoz, and Sidney I. Wiener. Rapid Spatial Reorientation and Head Direction Cells. *The Journal of Neuroscience*, 23(8):3478–3482, April 2003. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.23-08-03478.2003.

A Appendix

B Additional Simulation Results

The following contains additional simulation results, in which we explored the algorithm behavior for other environments that are classically used in neuroscience and behavioral studies (Figures 3 and 4, what happens when inhibition is not global but local (Figure 5, when disabling all inhibition which will lead to algorithm failure (Figure 6), and finally the behavior of the algorithm for multiple targets (Figures 7, 8, and 9).

The information presented in the figures follows Figure 2. That is, the starting neuron, meaning the neuron from which activity starts to spread out, is indicated as a blue box and is in all simulations in the bottom left corner for consistency. Each target neuron is indicated by a square white box underlying the neuron itself. The topographical map indicates the temporal landscape of neural activity, and the number of contour levels is kept fixed for each simulation to allow better comparison across iterations. The alpha value of a neuron – and therefore its color intensity – indicates whether a neuron spiked during an iteration or not.

The results show that the algorithm works also with only local inhibition, as long as there is only one target neuron, but shows significantly more remaining activity in the network (Figure 5). Results with multiple simultaneous targets indicate interesting characteristics of the algorithm, in particular in Figure 9, which shows that two of the three selected target neurons establish tagged neurons along a shortest path, but only one remains in the end. Future work will have to explore if additional properties of shortest paths can be encoded in specific relative timing events, and if multiple shortest paths can be successfully super-imposed within one neural population, maybe into different cycles of an oscillation.

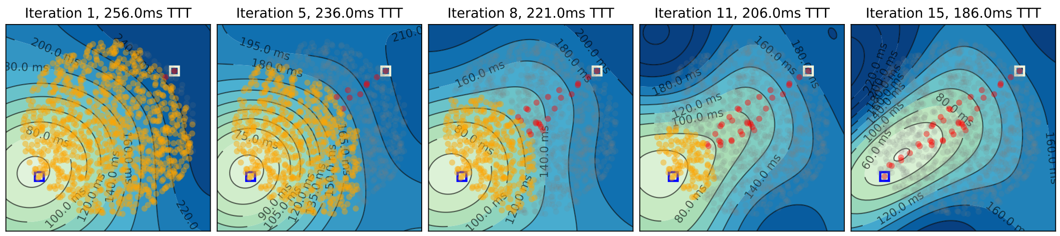

Figure 3: Evolution of the temporal gradient field during iterations of the algorithm in a circular maze. Clearly visible is the concentric propagation of activity during the first iteration (left panel), which over time turns into an elongated ellipse from start towards the target (right panel). This deformation of the contour lines shows the deformation of the temporal gradient field. That is, while the first iteration has neurons that operate alike and therefore propagate activity uniformly in time, the faster *tagged* neurons during the final iteration stretch the temporal gradient from start towards target and block activity on poor paths via inhibition.

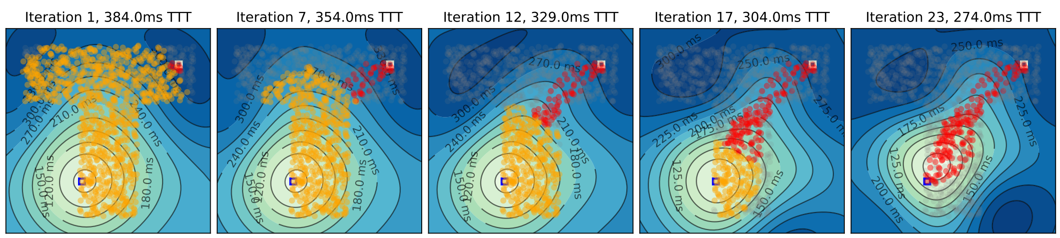

Figure 4: Evolution of the temporal gradient field during iterations of the algorithm in a T-Maze. Due to the high density of neurons in the maze as well as the shape of the maze, there is not a single (or approximately two, as in Figure 3), shortest path but several neurons that are on a broader shortest *band* towards the target. This is expected, given that distance between neurons is considered to be *unweighted* in this implementation of the algorithm.

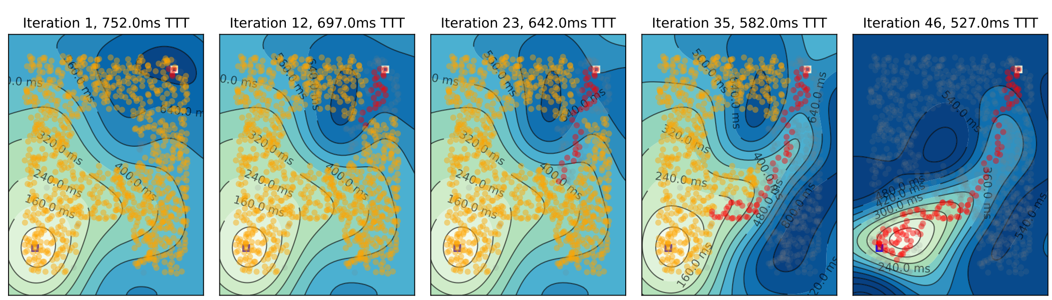

Figure 5: Evolution of the temporal gradient field during iterations of the algorithm in a A-Maze without global inhibition, but with local inhibition. In this simulation, global inhibition was turned off in favor of local inhibition by tagged neurons. Specifically, the local inhibition followed the same connectivity profile as the local excitation, i.e. an annulus around a neuron. The simulation shows that significantly more activity remains active within the network compared to Figure 2, and sub-optimal routes only get extinguished once the starting neuron is tagged. Afterwards, only tagged neurons remain active, and the shortest path appears during the final iteration of the algorithm.

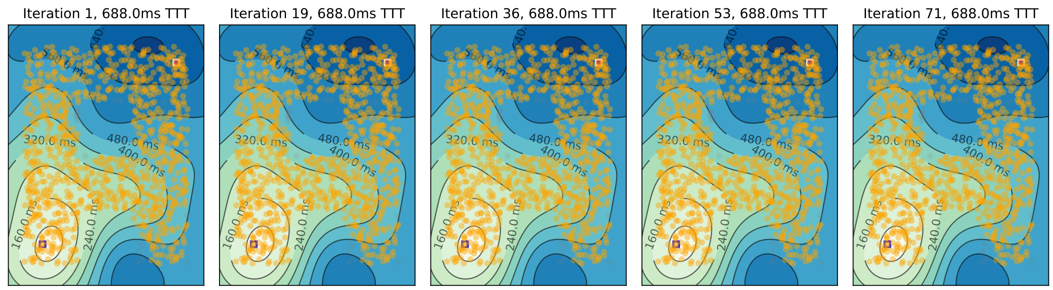

Figure 6: Algorithm failure due to missing inhibition. Without any form of inhibition, the algorithm fails to tag neurons, and the shortest path will not emerge.

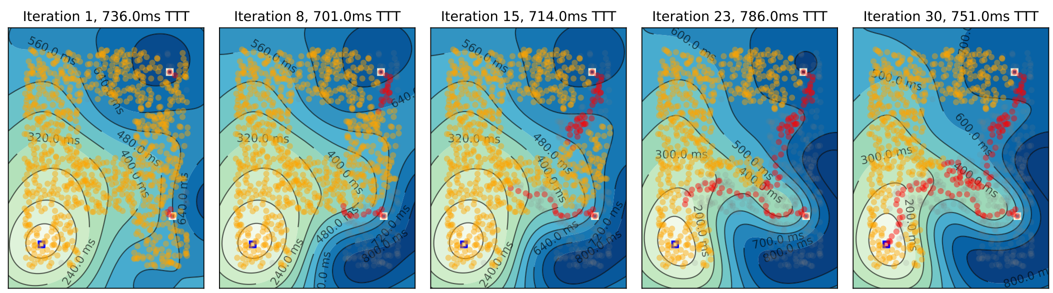

Figure 7: Simulation results for 2 targets and only local inhibition. During this run of the simulation, we disabled global inhibition in favor of local inhibition (see caption of Figure 5 for more details), but added an additional target neuron in the bottom right area. During the final iteration of the algorithm, both trajectories are active. The reason is that neurons on the shortest paths to either target are close enough in the middle section of the paths, so that activity from one shortest path reaches the other. Due to missing global inhibition, which could prevent the secondary shortest path, the activity continues until all paths are found.

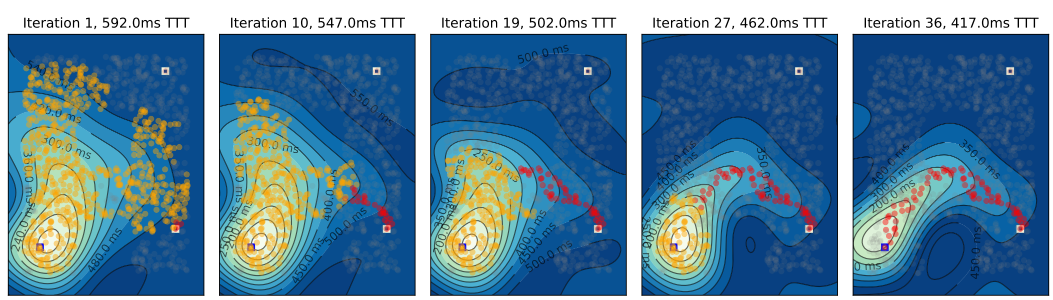

Figure 8: Simulation results for 2 targets and global inhibition. In contrast to Figure 7, we enabled global inhibition. In this case, the global inhibition prevents further propagation of activity towards the top-right target early on. This stops neurons in that direction from getting tagged, and therefore not second shortest-trajectory can emerge in the first place.

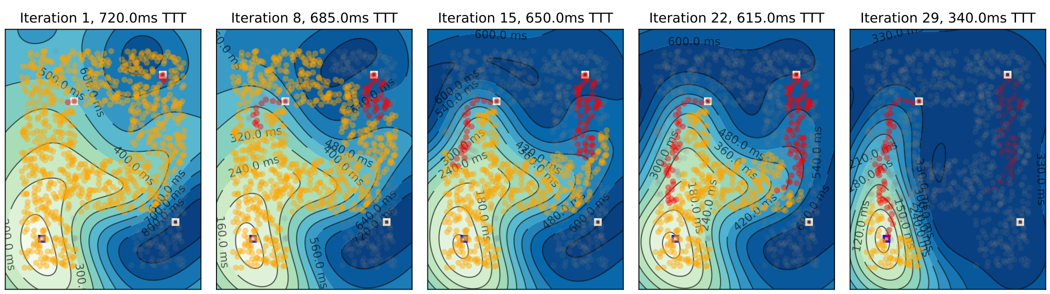

Figure 9: Simulation results for 3 targets, and only local inhibition. During the first iterations, shortest paths towards two of the three targets start to form. After some further iterations, the local inhibition along the trajectory to the target in the top left area extinguishes further excitatory propagation to the secondary remaining shortest path.