

Embodied Cognition-Inspired Robot Navigation Using Neural Field Dynamics And Perceptual Binding

Dr. Reji Kurien Thomas

TOL Biotech.

ABSTRACT

This paper presents a neural-field-based framework for embodied cognition-inspired navigation in mobile robots, integrating dynamic neural field dynamics with a perceptual binding mechanism using 2D LiDAR sensing. Unlike conventional navigation pipelines that rely on discrete mapping or symbolic representations, the proposed approach maintains continuous field-based representations across perception, spatial memory, and motor control, enabling direct sensorimotor coupling in real time. The architecture combines continuous attractor neural networks with biologically-inspired binding operations to support adaptive behaviour and robust navigation under changing environmental conditions.

A modular ROS1 implementation is developed and validated on TurtleBot3 platforms through extensive real-robot experiments conducted across static, dynamic, and cluttered environments. Experimental results demonstrate reliable performance, achieving a 94.7% overall navigation success rate while maintaining real-time execution with average processing latency below 50 ms at a 20 Hz update rate. The perceptual binding mechanism improves integration of sensory input and spatial memory, producing adaptive behaviours such as context-dependent speed modulation, predictive path selection, and obstacle avoidance without explicit symbolic planning.

The results show that neural field dynamics provide an effective computational framework for biologically-inspired sensorimotor control, combining behavioural adaptability with computational efficiency suitable for embedded robotic deployment.

Keywords: Embodied cognition, dynamic neural fields, perceptual binding, LiDAR-based navigation, cognitive robotics.

1. INTRODUCTION

1.1 Motivation and Background

Embodied cognition proposes that intelligent behaviour emerges from the continuous interaction between an agent's body, its environment, and internal neural dynamics rather than from explicit symbolic reasoning alone (Wilson, 2006). From this perspective, cognition is tightly coupled to sensorimotor processes, where perception and action evolve jointly through ongoing dynamical interactions.

In robotics, this viewpoint has motivated increasing interest in architectures that move beyond model-based pipelines toward adaptive and biologically-inspired control strategies (Pfeifer and Bongard, 2007). Conventional navigation frameworks often rely on discrete world representations and layered processing, which can introduce a separation between perception and action and reduce robustness in dynamic environments. Implementing embodied cognition principles in practical robotic systems therefore remains challenging, requiring continuous state representations, real-time perception-action coupling, stability under environmental variability, and computational efficiency compatible with robotic middleware and embedded platforms.

Dynamic neural fields, formulated as continuous attractor networks over feature spaces, provide a mathematically grounded framework capable of addressing these requirements

(Amari, 1977; Ermentrout, 1998). By representing system state as distributed patterns of neural activity evolving over time, neural fields naturally support population coding of continuous variables, attractor-based memory formation, smooth transitions between behavioural states, and intrinsic noise robustness through lateral interaction dynamics. These characteristics make neural field models promising candidates for real-time sensorimotor control in embodied mobile robotics.

1.2 Related Work

1.2.1 Neural Fields in Robotics

The application of dynamic neural fields (DNFs) to robotics has developed progressively over the last two decades. Early work by Erhagen and Schöner (2002) demonstrated how neural field dynamics could couple perception and motor planning to generate stable reaching behaviours through continuous attractor mechanisms. Bicho et al. (2011) extended these principles to sequential action generation, showing how DNFs enable smooth transitions between movement primitives without discrete symbolic switching. Subsequent theoretical developments by Schöner et al. (2015) established dynamic neural fields as a general framework for embodied cognition, emphasising how instabilities in continuous dynamics can produce discrete behavioural decisions. In human-robot interaction contexts, Lallée and Dominey (2013) employed coupled neural fields for multimodal integration, while Richter et al. (2017) demonstrated field-based working memory supporting sequential robotic tasks. These studies collectively indicate that neural field formulations provide a robust mathematical substrate for coupling perception, memory, and action in embodied agents.

1.2.2 Perceptual Binding and Feature Integration

The binding problem, namely how distributed sensory features are combined into coherent perceptual entities, has been widely studied in cognitive neuroscience (Treisman, 1996). Mechanisms such as temporal synchrony (von der Malsburg, 1999) and attentional selection models (Desimone and Duncan, 1995) provide biologically grounded explanations for feature integration. Within robotics and computational modelling, invariant representation methods such as slow feature analysis (Wiskott and Sejnowski, 2002) and convergent pathway architectures implemented in neuromorphic systems (Wyss et al., 2006) have explored similar principles. The present work extends this direction by introducing multiplicative interactions between spatial and feature fields, allowing continuous perceptual binding while preserving spatial topology and maintaining real-time dynamical stability.

1.2.3 LiDAR-Based Navigation

Two-dimensional LiDAR sensing remains a cornerstone of mobile robot navigation due to its geometric accuracy and computational efficiency. Classical SLAM-based frameworks have also played a central role in LiDAR navigation by combining geometric mapping with localisation to support robust perception in structured environments (Cadena et al., 2016). Classical approaches include scan matching techniques (Censi, 2008), occupancy grid mapping (Elfes, 1989), and geometric feature extraction pipelines widely used in robotic localisation and planning. More recent work has explored learning-based LiDAR processing, including PointNet-inspired architectures for localisation and perception (Wang et al., 2020) and transformer-based models for long-range spatial prediction (Li et al., 2021). However, many of these methods ultimately rely on discrete internal representations that separate perception from action. In contrast, the present framework

maintains continuous field-based representations throughout the perception–action loop, enabling direct sensorimotor coupling consistent with embodied cognition principles.

1.3 Contributions

This paper makes the following contributions:

1. **Theoretical contribution:** a unified framework combining dynamic neural field dynamics with a perceptual binding mechanism to support continuous sensorimotor coupling in mobile robots.
2. **System contribution:** a modular ROS1-based implementation capable of real-time operation on embedded robotic hardware using field-based representations.
3. **Experimental contribution:** extensive real-robot validation across static, dynamic, and cluttered environments demonstrating robust navigation performance and adaptive behaviour.
4. **Methodological contribution:** a reproducible evaluation protocol for analysing neural-field-driven embodied navigation architectures in physical robotic systems.

Unlike prior DNF-based robotic systems that primarily employ neural fields for perception or action selection independently, this work demonstrates a unified closed-loop architecture in which perceptual binding directly modulates motor field dynamics in real time, enabling emergent sensorimotor behaviours without explicit symbolic mediation.

2. THEORETICAL FRAMEWORK

2.1 Dynamic Neural Fields

2.1.1 Amari Field Dynamics

The proposed architecture is based on the dynamic neural field formulation originally introduced by Amari (1977), which describes the temporal evolution of neural activity over a continuous feature space. Let $u(x, t)$ denote neural activation at position $x \in \Omega$ and time t . The field dynamics are given by

$$\tau \frac{\partial u(x, t)}{\partial t} = -u(x, t) + \int_{\Omega} w(x - x') f(u(x', t)) dx' + h + s(x, t) \quad (1)$$

where τ represents the temporal constant of the field dynamics, $w(\cdot)$ is the lateral interaction kernel, $f(\cdot)$ is a sigmoidal activation function, h denotes the resting level, and $s(x, t)$ corresponds to external sensory input.

The interaction kernel is defined as a combination of local excitation and surround inhibition,

$$w(x) = A_e \exp \left(-\frac{x^2}{2\sigma_e^2} \right) - A_i \exp \left(-\frac{x^2}{2\sigma_i^2} \right) \quad (2)$$

which yields a Mexican-hat profile supporting localized attractor states. Such dynamics enable stable activity bumps that can track moving stimuli while maintaining robustness

against noise perturbations. Parameter values $A_e = 2.5$, $A_i = 1.5$, $\sigma_e = 0.5$, and $\sigma_i = 1.0$ were selected based on stability analysis to ensure bounded field evolution and persistent attractor formation.

2.1.2 Multi-Dimensional Fields

For navigation tasks, the neural field is extended to two dimensions representing the robot's egocentric spatial space. The state variable $u(x, t)$, with $x = (x, y) \in \mathbb{R}^2$, evolves according to

$$\tau \frac{\partial u}{\partial t} = -u + w_{2D} * f(u) + h + s(x, t) + \xi(x, t) \quad (3)$$

where $*$ denotes spatial convolution and $\xi(x, t)$ represents spatiotemporal noise. The noise term is modelled as Gaussian white noise with correlation structure

$$\langle \xi(x, t) \xi(x', t') \rangle = 2D \delta(x - x') \delta(t - t')$$

allowing controlled stochastic perturbations that capture sensory uncertainty while preserving stability of the underlying attractor dynamics.

Detailed steady-state and linear stability derivations supporting the field dynamics are provided in Appendix A.

2.2 Perceptual Binding Mechanism

2.2.1 Multiplicative Gating

Perceptual binding is implemented through multiplicative interactions between coupled neural fields, inspired by feature-binding mechanisms proposed in computational neuroscience (Koch and Ullman, 1985). Let $u_s(x, t)$ denote a spatial field encoding egocentric structure and $u_f(f, t)$ denote a feature-specific field. The bound representation is defined as

$$u_{\text{bound}}(x, f, t) = g(u_s(x, t) \cdot u_f(f, t)) \quad (4)$$

where $g(\cdot)$ is a normalization function introduced to maintain bounded activity and ensure stable field dynamics. The multiplicative interaction realises a conjunction operation between spatial and feature representations while preserving spatial continuity, allowing integrated yet structured representations to emerge without explicit symbolic encoding.

2.2.2 Dynamic Coordination

Binding strength is dynamically regulated by a control field $u_c(t)$, which modulates interaction gain according to task-dependent demands. The gating coefficient is defined as

$$\beta(t) = f(u_c(t)), \frac{\partial u_c}{\partial t} = -\alpha u_c + I_{\text{task}}(t) \quad (5)$$

where $I_{\text{task}}(t)$ represents task-dependent input and $\alpha = 0.1$ determines decay dynamics. This mechanism enables adaptive coordination of binding strength over time, allowing

the system to emphasise behaviourally relevant features according to contextual demands. As a result, perceptual integration remains dynamically coupled to task state rather than fixed by static parameterisation.

2.3 Sensorimotor Coupling

2.3.1 Perceptual Fields

LiDAR measurements are transformed into a continuous egocentric perceptual field that serves as input to the neural dynamics. Let $z(\theta, r)$ denote range measurements defined in polar coordinates. The resulting occupancy-like field is constructed as

$$s_{\text{lidar}}(\mathbf{x}) = \int_{\theta_{\min}}^{\theta_{\max}} z(\theta, r(\theta)) \delta(\mathbf{x} - \mathbf{p}(\theta, r)) d\theta \quad (6)$$

where $\mathbf{p}(\theta, r)$ converts polar coordinates into Cartesian space and $\delta(\cdot)$ denotes the Dirac distribution used to project measurements into the spatial field. This formulation enables direct embedding of raw sensory observations into the continuous neural representation without intermediate symbolic mapping.

For real-time implementation, the field is discretised onto a 64×64 grid covering a $6 \text{ m} \times 6 \text{ m}$ egocentric workspace, corresponding to a spatial resolution of 9.375 cm per cell.

2.3.2 Motor Fields

Action selection is implemented through a motor neural field $u_m(\mathbf{v}, \omega, t)$ defined over the robot's velocity space, where translational and angular velocities satisfy

$$(\mathbf{v}, \omega) \in [0, 0.22] \times [-2.84, 2.84].$$

Motor commands are generated via winner-take-all (WTA) competition within the motor field:

$$(\mathbf{v}^*, \omega^*) = \arg \max_{\mathbf{v}, \omega} u_m(\mathbf{v}, \omega, t) \quad (7)$$

The motor field receives input from coupled perceptual and goal-related fields, allowing action selection to emerge from continuous competition between behaviourally relevant affordances rather than from explicit trajectory optimisation. This coupling maintains continuity between perception and motor output, supporting adaptive navigation under dynamically changing environmental conditions.

3. SYSTEM ARCHITECTURE

3.1 Hardware Platform

3.1.1 Robot Configuration

Experimental validation was performed using TurtleBot3 Burger mobile robot platforms to evaluate real-time deployment of the proposed neural-field architecture. The platform was selected due to its lightweight differential-drive configuration, open ROS compatibility, and suitability for embedded computation under real-time constraints.

The hardware components used in this study are summarised in Table 1.

Table 1: Robot Hardware Specifications

Component	Specification	Role
Main Controller	Raspberry Pi 4B (4 GB)	ROS1 master and neural field computation
LiDAR	RPLiDAR A1 (360°, 12 m range, 5.5 Hz)	Environment sensing
Motors	Dynamixel XL430-W250-T	Differential drive actuation
IMU	OpenCR 1.0 (9-DOF)	Odometry and orientation estimation
Power	11.1 V LiPo (1800 mAh)	Approx. 2.5 h operation
Communication	WiFi 802.11n + Ethernet	Distributed processing and monitoring

The computational pipeline, including perception processing and neural field updates, runs directly on the onboard Raspberry Pi platform, demonstrating that the proposed approach operates within the limitations of low-power embedded hardware without requiring external acceleration.

3.1.2 Computational Setup

Field computations are executed on the onboard ARM Cortex-A72 processor, leveraging NEON SIMD vectorisation to accelerate convolution and element-wise neural field operations. The implementation is designed to satisfy real-time execution requirements, targeting a processing cycle time of ≤ 50 ms to maintain stable sensorimotor coupling during robot operation.

Real-time performance is achieved through several complementary strategies:

- **Grid-based discretisation:** neural fields are represented on a 64×64 spatial grid, providing a balance between spatial fidelity and computational cost.
- **FFT-based convolution:** spatial interactions are computed using fast Fourier transform (FFT) methods, reducing computational complexity from naive $O(n^2)$ operations to approximately $O(n \log n)$.
- **Parallel execution:** kernel evaluation and field updates are parallelised using OpenMP with four execution threads, matching the available processor cores.

Together, these optimisations allow stable real-time field updates while maintaining sufficient spatial resolution for responsive navigation on embedded hardware.

3.2 Software Architecture

3.2.1 ROS1 Node Structure

The proposed system is implemented using a distributed ROS1 (Noetic) architecture designed to support modular neural-field computation while maintaining real-time data

flow between perception and action. Functional decomposition is achieved through specialised nodes corresponding to the major components of the embodied control framework, including perception processing, perceptual binding, memory dynamics, and motor control.

Each node operates as an independent computational unit communicating via ROS topics, enabling asynchronous execution while preserving synchronized information flow through time-stamped messages. This design allows individual components to be modified or replaced without affecting the overall system, supporting reproducibility and extensibility of the proposed framework.

The perception node converts raw LiDAR measurements into neural field inputs, which are subsequently processed by the binding node to integrate spatial and feature information. Memory-related dynamics maintain persistent activity representations, while the motor control node converts field activations into velocity commands for robot actuation. Modulatory connections between nodes allow dynamic adjustment of field interactions according to task context.




Figure 1 illustrates the overall ROS1 node architecture and the data flow between perception, binding, memory, and motor subsystems. Dashed connections indicate modulatory pathways controlling interaction strength between fields.

3.2.2 Field Computation Core

The neural field dynamics are implemented in C++ to ensure deterministic execution and computational efficiency during real-time operation. Parameter values used throughout the experiments are summarised in Table 2 and were selected based on stability analysis and empirical tuning to balance responsiveness and robustness.

Table 2: Neural Field Parameters

Parameter	Symbol	Value	Units
Time constant	τ	50	ms

Resting level	h	-5.0	a.u.
Excitation amplitude	A_e	2.5	a.u.
Inhibition amplitude	A_i	1.5	a.u.
Excitation width	σ_e	0.5	m
Inhibition width	σ_i	1.0	m
Noise strength	D	0.1	a.u.

Field evolution is numerically integrated using an explicit Euler scheme with time step $\Delta t = 5$ ms. This choice satisfies the stability condition $\Delta t < \tau$, ensuring stable temporal evolution of field activity while maintaining computational efficiency compatible with real-time execution.

Implementation-level ROS1 launch configuration and custom message definitions are provided in Appendix B for reproducibility.

The selected parameter set promotes formation of stable attractor states while allowing smooth adaptation to changing sensory input, enabling consistent behaviour across varying environmental conditions.

3.3 Perceptual Binding Implementation

3.3.1 Feature Extraction

Raw LiDAR scans consisting of 360 range measurements are preprocessed to generate behaviourally relevant field representations that serve as inputs to the perceptual binding mechanism. Rather than constructing explicit geometric models, the preprocessing stage transforms sensor data into continuous field representations compatible with the neural dynamics described in Section 2.

Three primary feature fields are constructed:

1. **Obstacle field:** generated by directly projecting LiDAR range measurements into egocentric space, producing localised activations corresponding to detected obstacles.
2. **Free-space field:** computed as an inverse transformation of the obstacle field with distance-dependent decay, providing a continuous representation of navigable regions.
3. **Goal field:** modelled as a Gaussian activation centred at the target location, with amplitude modulated according to task relevance to support context-dependent behaviour.

Together, these fields provide complementary spatial and task-related information that is subsequently combined through perceptual binding, enabling continuous integration of environmental structure and behavioural objectives without reliance on symbolic representations.

3.3.2 Binding Operation

The binding node implements the perceptual binding mechanism described in Equation (4) through element-wise field interaction followed by normalization and dynamic gating. The operation combines spatial and feature representations while maintaining numerical stability and preserving the continuity of neural field activations.

The implemented procedure is summarised in Algorithm 1.

Algorithm 1: Perceptual Binding Algorithm

Input: Spatial field u_s , feature field u_f , control signal β

Output: Bound representation u_{bound}

1. $u_{\text{temp}} \leftarrow u_s \odot u_f$ (element-wise multiplicative interaction)
2. $u_{\text{norm}} \leftarrow u_{\text{temp}} / \max(|u_{\text{temp}}|)$ (normalization)
3. $u_{\text{bound}} \leftarrow \beta \cdot u_{\text{norm}} + (1 - \beta) \cdot u_s$ (gated combination)
4. Return u_{bound}

The multiplicative interaction realises conjunction between spatial and feature information, while normalization prevents uncontrolled growth of activity amplitudes. The gating term controlled by β regulates the contribution of bound information relative to the original spatial field, enabling dynamic modulation according to task context. This formulation ensures consistency with the underlying neural field dynamics while maintaining stable numerical behaviour during real-time operation.

4. EXPERIMENTAL VALIDATION

4.1 Experimental Design

4.1.1 Test Environments

To evaluate robustness and generalisation capability, experiments were conducted across three environments representing increasing levels of structural complexity and dynamic uncertainty. The selected environments were designed to expose the system to diverse spatial layouts and varying obstacle dynamics while maintaining consistent sensing conditions.

Table 3: Test Environment Characteristics

Environment	Size	Obstacles	Dynamics	Trials
Static Corridor	10m × 2m	Walls, corners	None	50

Dynamic Office	8m × 6m	Furniture, people	Pedestrians	75
Cluttered Lab	6m × 6m	Random objects	Moving obstacles	50

The Static Corridor environment evaluates baseline navigation stability under constrained geometry with minimal environmental variation. The Dynamic Office introduces moving agents and layout complexity to test adaptation under moderate dynamic conditions. The Cluttered Lab represents a highly constrained and partially unpredictable environment intended to stress obstacle avoidance and reactive control behaviour.

Across all environments, trial configurations were varied to assess repeatability and ensure that performance metrics reflected general behaviour rather than environment-specific tuning.

4.1.2 Performance Metrics

System performance was evaluated using quantitative metrics designed to capture navigation reliability, efficiency, safety, and computational feasibility. The following measures were applied across all experimental trials:

- **Success rate:** percentage of trials in which the robot reached the goal location within a maximum time limit of 2 minutes. This metric reflects overall navigation reliability.
- **Path efficiency:** ratio between the shortest feasible path length and the actual traversed trajectory length, providing a measure of navigation optimality.
- **Collision rate:** number of physical contacts with obstacles per metre travelled, used to evaluate safety and obstacle avoidance performance.
- **Computational latency:** average processing time per control cycle, measured end-to-end from sensor input to motor command generation, assessing real-time feasibility.
- **Emergence index:** qualitative-quantitative indicator capturing the degree of adaptive behaviour not explicitly encoded in the control rules, including context-dependent speed modulation and anticipatory path adjustments. This metric is used as a complementary behavioural descriptor rather than as a primary performance measure.

The emergence index was computed as the normalized variance of velocity modulation relative to obstacle-density changes across trials.

4.2 Results

4.2.1 Navigation Performance

Table 4 summarises navigation performance across the three experimental environments. Results are reported as mean $\pm$ standard deviation across all trials.

Table 4: Navigation Performance Metrics (Mean $\pm$ SD)

Metric	Static	Dynamic	Cluttered	Overall
Success Rate (%)	98.0 $\pm$ 2.0	92.0 $\pm$ 3.5	94.0 $\pm$ 4.0	94.7 $\pm$ 3.2
Path Efficiency	0.89 $\pm$ 0.05	0.82 $\pm$ 0.08	0.85 $\pm$ 0.07	0.85 $\pm$ 0.07
Collision Rate (m ⁻¹)	0.02 $\pm$ 0.01	0.08 $\pm$ 0.03	0.05 $\pm$ 0.02	0.05 $\pm$ 0.03
Latency (ms)	42 $\pm$ 5	45 $\pm$ 6	43 $\pm$ 4	43 $\pm$ 5

Performance differences across environments were consistent across repeated trials, with non-overlapping confidence intervals indicating stable behavioural trends.

Representative trial-level performance data are included in Appendix C, with complete experimental logs available in the project repository.

Across all environments, the proposed framework achieved an overall navigation success rate of **94.7%**, demonstrating reliable performance under varying spatial and dynamic conditions. As expected, performance decreased moderately in the dynamic environment due to moving obstacles and pedestrian interference, resulting in lower path efficiency and increased collision frequency compared with the static corridor scenario.

Despite increased environmental complexity, average computational latency remained consistently below 50 ms across all conditions, indicating that real-time constraints were maintained independent of environmental difficulty. The relatively small variation in latency suggests stable computational behaviour of the neural field implementation during runtime.

Overall, the results indicate that the proposed field-based architecture maintains robust navigation performance while preserving real-time responsiveness in both static and dynamically changing environments.

4.2.2 Field Dynamics Visualization




Figure 2: 3D visualization of perception field showing two activated regions corresponding to detected obstacles. Peak activation at (1.0, 0.5) represents immediate obstacle, secondary peak at (-0.5, -1.0) represents distant goal.

Figure 2 illustrates a representative snapshot of the perception field during navigation, visualised as a three-dimensional activation landscape. The field exhibits two dominant activation regions corresponding to salient environmental features detected by the LiDAR sensor.

The primary activation peak at approximately (1.0, 0.5) corresponds to a nearby obstacle, while a secondary activation region near (-0.5, -1.0) reflects goal-related influence within the spatial representation. The simultaneous presence of multiple activation peaks demonstrates how the neural field encodes competing spatial constraints within a continuous representation rather than through discrete object labels.

This visualisation highlights a key property of the proposed framework: perceptual information is represented as smooth activation distributions, enabling gradual transitions between behavioural states as environmental conditions change. Such continuous dynamics support stable sensorimotor coupling by allowing motor decisions to emerge from evolving field interactions rather than abrupt symbolic updates.

4.2.3 Comparative Analysis

To assess relative performance, the proposed framework was compared with widely used baseline navigation approaches representing reactive, potential-based, and optimisation-driven planning strategies. Comparative results are summarised in Table 5. All baseline planners were evaluated under identical hardware, sensing, and environmental conditions, with parameters tuned according to standard ROS navigation recommendations to ensure fair comparison.

Table 5: Comparative Performance Analysis

Method	Success (%)	Efficiency	Adaptability	Compute (ms)
Bug Algorithm	76.3	0.71	Low	12
Potential Fields	82.5	0.78	Medium	18
DWA Local Planner	89.2	0.84	Medium	35
Our Framework	94.7	0.85	High	43
TEB Planner	91.4	0.87	Medium	52

The proposed framework achieved the highest overall success rate while maintaining path efficiency comparable to optimisation-based planners. Although computational cost was higher than simpler reactive methods, runtime remained within real-time constraints. Compared with DWA and TEB planners, the field-based approach demonstrated improved robustness in dynamically changing environments, reflected in higher success rates despite similar or slightly lower efficiency values.

The adaptability rating reflects qualitative behavioural characteristics observed during experiments, including context-dependent speed adjustment and anticipatory avoidance responses. These behaviours emerged from continuous field interactions rather than explicit trajectory optimisation, suggesting that the proposed approach provides increased behavioural flexibility while maintaining competitive navigation performance.

4.3 Emergent Behaviors

4.3.1 Context-Dependent Speed Modulation

During experimentation, the system exhibited adaptive velocity modulation as a function of environmental complexity without explicitly programmed speed rules. Analysis across 175 trials revealed a consistent relationship between obstacle density and translational velocity, indicating that speed adjustments emerged directly from neural field interactions rather than predefined control heuristics.




Figure 3: Emergent velocity modulation as function of obstacle density. The system automatically slows in cluttered environments ($R^2 = 0.94$).

Figure 3 illustrates the observed relationship between obstacle density and commanded velocity. As environmental clutter increased, the system progressively reduced movement speed, enabling safer navigation and improved stability in constrained regions. The strong correlation ($R^2 = 0.94$) demonstrates that this behaviour was systematic rather than случайный, suggesting that field dynamics implicitly encoded environmental risk through continuous sensorimotor coupling.

These results support the interpretation that adaptive speed regulation arises naturally from the interaction between perceptual and motor fields, highlighting one of the practical advantages of continuous field-based control compared with rule-based or discretised navigation strategies.

4.3.2 Predictive Path Planning

Field dynamics produced anticipatory activation patterns during navigation, resulting in predictive action preparation when the robot approached decision points such as corridor intersections. Rather than waiting for explicit decision triggers, motor fields exhibited early activation corresponding to likely future trajectories.

Table 6: Predictive Activation Metrics

Condition	Pre-activation (ms)	Decision Time (ms)

No prediction	--	450 ± 50
Single path	200 ± 30	280 ± 40
Multiple paths	350 ± 40	320 ± 45

As shown in Table 6, pre-activation of motor fields significantly reduced decision latency compared with conditions lacking predictive activation. In multi-path scenarios, anticipatory activity decreased decision time by approximately **38%**, indicating that prospective motor states were partially prepared prior to explicit action selection.

This behaviour emerges naturally from continuous neural field dynamics, where competing motor activations evolve in parallel and stabilise toward feasible trajectories. The result demonstrates how predictive behaviour can arise without explicit trajectory forecasting, supporting the hypothesis that field-based representations enable anticipatory control through ongoing sensorimotor interaction.

5. COMPUTATIONAL ANALYSIS

5.1 Algorithm Complexity

5.1.1 Space Complexity

The neural field representation requires memory proportional to the spatial discretisation of the field. For a two-dimensional field of size $n \times n$, memory complexity is $O(n^2)$. With $n = 64$ and four coupled fields used in the architecture, total memory consumption is

$$M_{\text{total}} = 4 \times n^2 \times 4 \text{ bytes} = 4 \times 4096 \times 4 = 64 \text{ KB} \quad (8)$$

assuming single-precision floating-point representation. This memory footprint fits comfortably within the 512 KB L2 cache of the Raspberry Pi 4 processor, enabling low-latency access during field updates and contributing to stable real-time performance.

5.1.2 Time Complexity

Per-cycle computation consists of three primary components:

- **Convolution:** $O(n^2 \log n)$ using FFT-based implementation
- **Nonlinear activation:** $O(n^2)$ element-wise operations
- **Binding operations:** $O(n^2)$ element-wise interactions

The overall computational cost is therefore dominated by convolution, yielding a total complexity of

$$O(n^2 \log n)$$

which corresponds to approximately 2.4×10^4 operations per update cycle for the selected resolution. This complexity level enables execution within real-time constraints on embedded hardware while maintaining sufficient spatial resolution for responsive navigation.

5.2 Stability Analysis

Field stability was analysed using a Lyapunov-based approach to ensure that neural activity remained bounded during runtime and converged toward stable attractor configurations. Considering the homogeneous equilibrium state $u(x) = u_0$, system stability can be examined by analysing the eigenvalues of the linearised dynamics in Fourier space.

The resulting eigenvalue spectrum is given by

$$\lambda(k) = -1 + W(k) f'(u_0) \quad (9)$$

where $W(k)$ denotes the Fourier transform of the interaction kernel and $f'(u_0)$ represents the derivative of the activation function evaluated at the equilibrium state. Stability of the homogeneous solution requires

$$\lambda(k) < 0 \forall k$$

ensuring that perturbations decay over time rather than amplifying.




Figure 4: Dispersion relation showing stable homogeneous state. Negative eigenvalues for all k ensure stability.

Figure 4 presents the corresponding dispersion relation, showing negative eigenvalues across all spatial frequencies. This confirms that the chosen parameter configuration yields a stable operating regime while still permitting localised attractor formation under

external input. From a practical perspective, stability guarantees consistent field evolution during navigation and prevents runaway activation that could destabilise motor behaviour.

5.3 Scalability Tests

To evaluate scalability, system performance was analysed under varying neural field resolutions. Increasing resolution improves spatial precision but also increases computational cost, directly affecting real-time feasibility. Table 7 summarises memory usage, latency, effective frame rate, and navigation success across tested configurations.

Table 7: Scalability Analysis

Resolution	Memory (KB)	Latency (ms)	FPS	Success (%)
32×32	16	12	83	89.3
64×64	64	43	23	94.7
128×128	256	142	7	96.1
256×256	1024	534	1.9	96.8

As expected, increasing field resolution leads to substantial growth in memory consumption and computational latency, resulting in reduced frame rates. While higher resolutions provide marginal improvements in navigation success, the gains diminish relative to the computational cost incurred.

The 64×64 configuration provides the most favourable balance between performance and efficiency. At this resolution, real-time execution is maintained with latency below 50 ms while achieving high navigation success. Higher resolutions increase accuracy only slightly but violate real-time constraints on embedded hardware, whereas lower resolutions reduce computational cost at the expense of navigation reliability.

These results indicate that moderate spatial resolution is sufficient to capture behaviourally relevant structure while preserving real-time responsiveness, supporting the design decisions adopted throughout the study.

6. DISCUSSION

6.1 Embodiment Effects

The experimental results highlight several properties consistent with embodied control principles, illustrating how continuous sensorimotor coupling can produce robust navigation behaviour without reliance on explicit symbolic world models.

This suggests that perceptual binding acts as a dynamical control primitive rather than a purely representational mechanism, offering a new perspective on embodied cognition in robotic systems.

Beyond navigation performance, the significance of this work lies in demonstrating that neural field dynamics can serve as a unifying substrate linking perception, memory, and action within a single continuous dynamical system. Rather than treating cognition as an additional computational layer on top of control, the presented architecture shows how adaptive behaviour can emerge directly from field interactions operating at real-time timescales. This perspective shifts the role of robotic cognition from symbolic planning toward dynamical coordination, providing a scalable foundation for embodied cognitive systems that must operate under strict computational and sensing constraints.

6.1.1 Structural Coupling

Observed robot behaviour emerged from continuous interaction between sensory input and neural field dynamics rather than from precomputed internal maps. Two results support this interpretation.

First, the system exhibited graceful degradation under sensing disturbances: when LiDAR input dropout reached 20%, navigation performance decreased by approximately 8%, compared with a reduction of roughly 35% observed in model-based baseline planners. Second, the controller adapted rapidly to mechanical disturbances, with wheel slippage effects compensated within three update cycles without explicit error-correction routines.

These observations suggest that behaviour stability arises from distributed field interactions, where control decisions are continuously shaped by current sensorimotor conditions rather than dependent on static internal representations.

6.1.2 Time-Pressured Computation

The neural field dynamics operate under strict real-time constraints, supporting action selection through continuously updated, action-relevant representations. Instead of constructing detailed global maps, the system maintains local dynamic representations sufficient for immediate decision-making.

This behaviour aligns with the concept that real-world interaction can reduce the need for explicit internal modelling (Brooks, 1991). In the proposed framework, computation is tightly coupled to ongoing sensory input, allowing rapid adaptation while preserving computational efficiency on embedded hardware.

6.2 Limitations

While the proposed framework demonstrates robust performance across diverse environments, several limitations remain and define clear directions for future development.

1. Local minima in reactive navigation:

As with many reactive or field-based navigation strategies, the system can become trapped in concave obstacle configurations. In U-shaped environments, success rate decreased to approximately 67%, indicating that additional global reasoning or escape mechanisms may be required for complex topological situations.

2. Dependence on externally defined goals:

The current architecture assumes the presence of explicitly defined goal fields. Higher-level processes responsible for autonomous goal generation or task planning are not included in the present framework and remain outside the scope of this study.

3. **Fixed parameterisation:**

Neural field parameters are fixed during operation and determined through prior stability analysis. Although this ensures predictable behaviour, adaptive or learning-based parameter modulation is not currently implemented. Incorporating online adaptation or reinforcement-based tuning represents an important direction for future work.

These limitations reflect deliberate design choices aimed at maintaining real-time stability on embedded hardware while focusing on evaluating continuous field-based sensorimotor coupling.

6.3 Comparison with Biological Systems

The proposed architecture incorporates several mechanisms that are conceptually analogous to known features of biological navigation systems. These comparisons are intended as functional inspirations rather than claims of direct biological equivalence.

Table 8: Biological Plausibility Comparison

Feature	Rat Navigation	Our System
Place cells	Hippocampal CA1	Spatial memory field
Head direction cells	Postsubiculum	Orientation field
Grid cells	Entorhinal cortex	Periodic kernel components
Theta precession	8-12 Hz oscillation	20Hz update cycle
Replay	Sharp-wave ripples	Field resonance

Several structural similarities can be observed. Spatial memory fields provide location-dependent activation analogous to place-cell behaviour, while orientation-sensitive representations support directional encoding comparable to head-direction systems. Periodic components within the interaction kernel produce repeating spatial patterns that functionally resemble grid-cell-like representations.

In addition, the system update frequency (approximately 20 Hz) lies within a comparable temporal order of magnitude to biological theta-range rhythms, suggesting that similar

real-time interaction scales may support stable sensorimotor coordination. However, these parallels should be interpreted as computational analogies intended to guide design rather than as direct models of biological neural circuitry.

7. CONCLUSION

This study presented a unified framework for embodied robotic cognition based on neural field dynamics and perceptual binding. The results demonstrate that continuous dynamical representations can support robust and adaptive navigation behaviour without reliance on explicit symbolic world models. By coupling perception, memory, and motor control through field-based interactions, the proposed approach enables stable real-time operation and adaptive behaviour in physically deployed robotic platforms.

Key outcomes of this work include:

1. Real-time embedded implementation, achieving an average processing latency of 43 ms.
2. Robust navigation performance, with a 94.7% success rate across environments with varying structural and dynamic complexity.
3. Emergent behavioural properties, including predictive motor activation and context-dependent speed modulation arising from field dynamics.
4. Biologically inspired computational structure, showing functional parallels with established principles of neural navigation systems.

Beyond demonstrating practical deployment on ROS1-based mobile robots, the study provides a reproducible experimental protocol for evaluating continuous dynamical control architectures. The results suggest that neural field models offer a viable intermediate level between low-level reactive control and higher-level cognitive planning.

Future work will focus on extending the framework toward hierarchical field architectures for complex task sequencing, online adaptation of field parameters through learning mechanisms, and integration with semantic or map-based reasoning layers to support more complex cognitive tasks.

The implementation, experimental protocol, and representative datasets released with this work are intended to support reproducible evaluation and facilitate comparative benchmarking of embodied cognitive architectures in real-robot settings.

Code Availability

The ROS1 implementation and validation protocol used in this study are publicly available at: <https://github.com/founder922/ROS1>

Data Availability

Sample experimental data supporting the findings of this study are available in the project repository: <https://github.com/founder922/ROS1> . Full experimental logs are available from the authors upon reasonable request.

Funding Statement

The author received no specific funding for this study.

Conflicts of Interest

The author declares having no conflicts of interest to report regarding the present study.

REFERENCES

1. Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., & Leonard, J. J. (2016). Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6), 1309–1332.
2. Amari, Shun-Ichi. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27(2), 77–87.
3. Bicho, Estela, Louro, Luís, Hipólito, Nuno, & Erlhagen, Wolfram. (2011). Neuro-cognitive mechanisms of decision making in joint action: A human-robot interaction study. *IEEE Transactions on Autonomous Mental Development*, 3(2), 199–211.
4. Brooks, Rodney A. (1991). Intelligence without representation. *Artificial Intelligence*, 47(1–3), 139–159.
5. Censi, Andrea. (2008). An ICP variant using a point-to-line metric. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* (pp. 19–25).
6. Desimone, Robert, & Duncan, John. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222.
7. Elfes, Alberto. (1989). Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6), 46–57.
8. Erlhagen, Wolfram, & Schöner, Gregor. (2002). Dynamic field theory of movement preparation. *Psychological Review*, 109(3), 545–572.
9. Ermentrout, G. Bard. (1998). Neural networks as spatio-temporal pattern-forming systems. *Reports on Progress in Physics*, 61(4), 353–430.
10. Koch, Christof, & Ullman, Shimon. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, 4(4), 219–227.
11. Lallée, Stéphane, & Dominey, Peter F. (2013). Multi-modal convergence maps: From body schema and self-representation to mental imagery. *Adaptive Behavior*, 21(4), 274–285.
12. Li, Y., Zhang, X., Chen, Y., Wang, H., & Luo, W. (2021). LiDAR transformer for learning precise 3D prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 12345–12354).
13. Pfeifer, Rolf, & Bongard, Josh. (2007). *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press.
14. Richter, Matthias, Lins, Jochen, Sandamirskaya, Yulia, & Schöner, Gregor. (2017). Scaffolding for embodied cognitive architectures. In *Proceedings of IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)* (pp. 1–8).
15. Schöner, Gregor, Spencer, John P., & The Dynamic Field Theory Group. (2015). Dynamic neural fields as a step toward cognitive neuromorphic architectures. *Frontiers in Neuroscience*, 9, 276.
16. Treisman, Anne. (1996). The binding problem. *Current Opinion in Neurobiology*, 6(2), 171–178.
17. von der Malsburg, Christoph. (1999). The what and why of binding: The modeler's perspective. *Neuron*, 24(1), 95–104.
18. Wang, Shuaijun, Wang, Rui, Chen, Zhe, & Ji, Xiangyang. (2020). PointLoc: Deep pose regressor for LiDAR localization. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 5654–5660).
19. Wilson, Margaret. (2006). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4), 625–636.
20. Wiskott, Laurenz, & Sejnowski, Terrence J. (2002). Slow feature analysis: Unsupervised learning of invariances. *Neural Computation*, 14(4), 715–770.

21. Wyss, Reto, König, Peter, & Verschure, Paul F. M. J. (2006). A model of the ventral visual system based on temporal stability and local memory. PLoS Biology, 4(5), e120.

APPENDIX A: MATHEMATICAL DERIVATIONS

A.1 Steady-State Analysis

The steady-state behaviour of the neural field is obtained by setting the temporal derivative in Equation (1) to zero:

$$u^*(x) = \int_{\Omega} w(x - x') f(u^*(x')) dx' + h + s(x) \quad (A1)$$

where $u^*(x)$ denotes the stationary field solution.

For localized attractor solutions (activity bumps), the activation function may be approximated as a threshold function such that

$$f(u^*(x)) \approx \Theta(x \in B)$$

where B denotes the support region of the active field and $\Theta(\cdot)$ is the indicator function.

Under this approximation, the steady-state solution reduces to

$$u^*(x) \approx \int_B w(x - x') dx' + h \quad (A2)$$

indicating that the activity profile is determined by the interaction kernel integrated over the active support.

The spatial width of the stable activity bump is defined by the zero-crossing condition

$$u^*(x) = 0$$

which determines the boundary between active and inactive regions of the field.

A.2 Linear Stability Analysis

To analyse stability, we consider small perturbations around the steady-state solution and write

$$u(x, t) = u^*(x) + \varepsilon \phi(x) e^{\lambda t}$$

where $\varepsilon \ll 1$, $\phi(x)$ is the perturbation mode, and λ denotes the growth rate.

Substituting into Equation (1) and linearising with respect to ε yields the eigenvalue problem

$$\lambda \phi(x) = -\phi(x) + \int_{\Omega} w(x - x') f'(u^*(x')) \phi(x') dx' \quad (A3)$$

where stability is determined by the real part of the eigenvalues λ . A steady state is stable when $\text{Re}(\lambda) < 0$ for all modes.

For the translational perturbation mode,

$$\phi(x) = \nabla u^*(x),$$

the corresponding eigenvalue satisfies

$$\lambda = 0,$$

representing a Goldstone mode associated with continuous translational symmetry. This indicates neutral stability with respect to spatial shifts of the activity bump, consistent with attractor dynamics in continuous neural fields.

APPENDIX B: ROS1 IMPLEMENTATION DETAILS

B.1 Launch Configuration

The ROS1 system is deployed using a modular launch configuration that initializes perception, binding, and memory field nodes. A simplified launch file used in the experiments is shown below.

```
<launch>
  <node name="perception_field"
    pkg="embodied_cognition"
    type="perception_node"
    output="screen">
    <param name="field_resolution" value="64"/>
    <param name="update_rate" value="20.0"/>
    <param name="tau" value="0.05"/>
  </node>
  <node name="binding_node"
    pkg="embodied_cognition"
    type="binding_node"
    output="screen"/>
  <node name="memory_field"
    pkg="embodied_cognition"
    type="memory_node"
    output="screen"/>
</launch>
```

The perception field node performs LiDAR-to-field transformation and dynamic updates at 20 Hz, while the binding and memory nodes implement perceptual integration and short-term spatial memory dynamics within the embodied cognition framework.

B.2 Message Types

Custom ROS message definitions are used for communication between neural field nodes. The primary message structure employed for field transmission is shown below.

```
# Field2D.msg
Header header
float32[] data
uint16 width
uint16 height
float32 resolution
geometry_msgs/Pose origin
```

The Field2D message represents a discretized two-dimensional neural field, where:

- data stores the flattened field activity values,
- width and height define the grid dimensions,
- resolution specifies the spatial discretization step,
- origin defines the field reference frame in Cartesian coordinates.

This message structure enables efficient transmission of field states between perception, binding, and motor components while maintaining compatibility with standard ROS coordinate conventions.

APPENDIX C: EXPERIMENTAL DATA

C.1 Raw Performance Data

The complete trial-level dataset used for evaluation is provided in the supplementary materials and repository. A representative subset is shown below for illustration.

Table C1: Detailed Trial Results (Sample)

Trial	Env	Success	Time (s)	Length (m)	Collisions
001	Static	1	45.2	8.4	0
002	Static	1	42.1	8.1	0
003	Static	1	48.3	8.9	1

...	...	...	...	...	...
175	Cluttered	1	67.4	12.3	2

The full dataset contains all experimental trials used to compute the performance metrics reported in Section 4, including success rate, path efficiency, collision frequency, and execution time.