Coordination Without Command:
Active Inference and the Route to Embodied Intelligence

Dr Alexander D. Shaw

Computational Psychiatry & Neuropharmacological Systems (CPNS) Lab
Department of Psychology, Faculty of Health & Life Sciences
University of Exeter, UK
https://cpnslab.com
a.d.shaw@exeter.ac.uk

Abstract

Embodied intelligence is often discussed in robotics in terms of control, planning, or optimisation; yet real agents must act through bodies that are dynamically constrained, only partially informed, and continuously reshaped by their own movements. In this perspective, we argue that embodied intelligence is better understood not as the operation of a single central controller, but as the coordination of distributed, hierarchical, and plural inferential processes. We develop this argument by first examining why embodiment places pressure on monolithic control architectures, then showing why active inference provides a particularly natural framework for agents that must act under uncertainty while sampling the world through movement. We ground the discussion in two complementary examples: the octopus as a biological instance of intelligence distributed through the body, and a temporally predictive scene-based drone controller as a computational case study in embodied active inference under partial observability. Taken together, these examples suggest that robust embodied behaviour may depend less on exhaustive central resolution than on coordination without command across multiple timescales, interfaces, and inferential demands.

1. Introduction

As artificial agents are deployed into increasingly complex environments, the problem of intelligence begins to look less like one of abstract optimisation and more like one of embodied coordination; because real agents do not act from nowhere, but through bodies that are situated in the world, constrained by morphology, limited by partial observability, and continuously reshaped by the consequences of their own movements. Under these conditions, the computational problem is no longer just to select an optimal action from a well-defined state, but to maintain adaptive behaviour under uncertainty while perception, action, memory, and environmental structure remain tightly coupled [28, 26, 8].

Although modern robotics has made major progress in perception, planning, and control, many of its most successful architectures still inherit a largely centralised view of intelligence - whether this is expressed through model-based control, reinforcement learning, or policy optimisation. In each case, the underlying assumption is often that behaviour should ultimately resolve to a single controller, a single policy, or a single objective (or utility) function whose optimisation governs action selection [34, 19, 36]. That assumption has obvious practical advantages since it supports tractable implementations, clean benchmarks, and strong engineering performance in relatively well-structured settings. Yet, it sits rather uneasily with the realities of embodied agents operating in open-ended and dynamically changing worlds where multiple timescales, competing affordances, uncertain sensory data, and nonlinear body-environment interactions must be managed at once rather than collapsed in advance.

The tension becomes clearer still when the body itself is flexible, high-dimensional, or only partly amenable to detailed central control, because in such cases the problem is not simply one of issuing better commands but of coordinating many local processes whose consequences unfold at different spatial and temporal scales. Biological systems provide many examples in which adaptive behaviour does not appear to arise from a single unified controller directing passive effectors in detail, but instead from the interaction of multiple loops, with local processes handling fast sensorimotor contingencies and more global processes imposing broader contextual or goal-dependent constraints [17, 1, 13]. What emerges in these systems is not (usually) disorder, but a different kind of order; one achieved not through exhaustive central oversight but through structured coordination among semi-autonomous components.

Active inference provides a particularly promising framework in which to rethink these issues, precisely because it treats perception and action as coupled processes of inference under a generative model and therefore offers a principled account of behaviour in which agents reduce uncertainty and fulfil predictions through embodied interaction with the world [8, 3, 24]. For that reason, it is well suited, at least in principle, to the study of embodied intelligence. At the same time, many active inference implementations still retain a hidden centralism insofar as inference is often cast in terms of a single dominant posterior, a single policy distribution, or a single controller operating over an internal model of the world. That may be entirely appropriate in simple settings, but it also risks reproducing within active inference the same *monolithic* assumptions that have constrained other approaches to embodied control.

In this perspective, we argue that embodied intelligence requires distributed, hierarchical, and plural inference. By distributed inference, we mean that the processes responsible for adaptive behaviour need not be concentrated in a single locus of control, but may instead be spread across multiple interacting subsystems at the interface between body and world. By hierarchical inference, we mean that these processes operate across distinct spatial and temporal scales, such that local sensorimotor loops are shaped (but not micromanaged) by broader contextual priors. By plural inference, we mean that adaptive behaviour often depends on maintaining multiple partially compatible hypotheses, control tendencies, or inferential voices in parallel, rather than collapsing prematurely onto a single dominant explanation or policy. In this view, intelligence is better understood as coordination without command than as optimisation by a unitary controller.

We develop this argument by bringing together three strands; first, we situate the problem in relation to embodiment and the limitations of centralised control in robotics. Second, we argue that active inference becomes substantially more compelling as a framework for embodied intelligence when extended in a distributed, hierarchical, and pluralistic direction. Third, we ground these ideas in both biology and computation, using the octopus as a natural example of control pushed into the body itself; and a temporally predictive scene-based drone controller as a computational case study in embodied active inference. The broader claim is not that all intelligent systems must abandon convergence or centralisation entirely, but that when the aim is robust embodied intelligence in rich and uncertain environments, architectures built around distributed, hierarchical, and pluralistic inference may offer a more natural route forward.

2. Why Embodiment Changes the Computational Problem

2.1 The body is part of the inference problem

Once intelligence is treated as something that unfolds in a body rather than above it, the body can no longer be understood as a passive output device that merely executes commands generated elsewhere. Its morphology, material properties, sensor placement, and action capabilities all shape what can be sensed, what can be predicted, and what kinds of control problems the agent is even able to pose in the first place. In that sense, embodiment is not an implementation detail added after the fact, but a crucial part of the computational architecture of the agent itself [29, 30].

This matters because the sensory stream available to an embodied agent is inseparable from the movements through which that stream is sampled. Looking, reaching, turning, probing, and repositioning do not simply follow inference, but participate in it, since action changes what is observed, which in turn changes the evidence available for belief updating. The body therefore does not sit downstream of cognition as a simple actuator. Rather, it closes the loop between prediction and sensation, constraining both the hypotheses the agent can entertain and the means by which those hypotheses can be tested against the world [8, 3, 31].

Seen in these terms, many of the classical separations that have structured robotics and AI begin to look increasingly *artificial*. Perception cannot be cleanly isolated from action when the point of moving is often to resolve uncertainty, and control cannot be cleanly isolated from world-modelling when the body itself determines which parts of the world become available to model. This was one of the central insights of embodied AI and behaviour-based robotics, where intelligence was argued to arise not from abstract internal problem solving alone, but from the structured interaction between agent, body, and environment [2, 29]. What those traditions recognised, and what remains highly relevant for contemporary active inference, is that many apparently difficult control problems can be transformed, simplified, or even dissolved once morphology and sensorimotor coupling are treated as part of the solution rather than as complications to be compensated away.

For embodied agents operating in real environments, this has a further consequence; namely that the body shapes not only how actions are executed, but how hidden states must be inferred over time. Occlusion, self-motion, limited fields of view, delayed consequences, and changing bodily configuration all mean that inference is necessarily perspectival and action-dependent. An agent does not recover a detached, complete description of the world and then decide what to do. Instead, it maintains *beliefs* from a particular embodied vantage point, updating them through movements that are themselves constrained by the body and its current relation to the environment. The inferential problem is therefore always already sensorimotor [8, 3, 27].

From the perspective developed in this paper, that point is more than a philosophical scene-setting because once the body is recognised as part of the inferential machinery, the appeal of distributed, hierarchical, and plural architectures becomes much easier to see. If the world is sampled through multiple sensorimotor interfaces, if different parts of the body face different local contingencies, and if behaviour must be organised across multiple timescales at once, then it becomes increasingly implausible that robust embodied intelligence should always reduce to a single inferential unit, or a single policy selected from nowhere in particular. What embodiment introduces is not simply noise or implementation difficulty, but a structural reason to think that intelligence may need to be spread across coupled processes operating at the interface between body and world.

2.2 Why centralised control scales poorly

Once the body is recognised as part of the inferential problem, it becomes much harder to assume that intelligence can be organised around a single central process that maintains a sufficiently rich model of the world, computes the consequences of all relevant actions, and then issues the correct commands in real time. That picture may remain workable in relatively simple settings, particularly where the state space is tightly constrained and the body behaves in predictable ways; but it becomes increasingly strained as agents move into environments that are only partially observable, dynamically changing, and strongly coupled to the mechanics of their own embodiment [16, 20, 7].

Part of the difficulty lies in the familiar curse of dimensionality, though in embodied systems this problem is not just one of large state spaces in the abstract. What makes the challenge more severe is that the relevant state of the world is often entangled with the agent's own position, posture, sensing configuration, and immediate action possibilities, all of which may change the observations available at the next moment. Under such conditions, centralised control is required not merely to select an action, but to maintain an online account of a world that is continuously being reshaped by the agent's own sensorimotor engagement with it. This complexity therefore grows not only with the environment, but with the body and its relation to the environment [29, 27].

These problems become especially acute in flexible or high-dimensional bodies, where even the local consequences of movement may be difficult to predict from a single global vantage point. Soft robotic systems, continuum manipulators, and biological organisms such as the octopus all illustrate the same underlying point; namely that when morphology introduces large numbers of coupled degrees of freedom, nonlinear deformations, and rich local sensory contingencies, detailed central micromanagement quickly becomes computationally implausible [5, 21, 14]. In those settings, the question is no longer how to design a better central controller for every variable, but how to decompose the problem in a way that allows local structure, bodily dynamics, and context-sensitive constraints to do some of the work.

Partial observability adds a further layer of difficulty because centralised architectures often presume a level of informational access that embodied agents do not possess. Real agents must act despite occlusion, noise, delayed feedback, changing viewpoints, and uncertainty about the hidden causes of their sensations, and this means that control cannot be reduced to the execution of commands over a fully specified external state. Instead, behaviour must be organised around beliefs that are incomplete, revisable, and often action-dependent, with different parts of the agent effectively confronting different local uncertainties at the same time [16, 3]. The result is that a single central controller is not merely burdened with a large computation, but with the impossible task of collapsing heterogeneous, perspective-bound uncertainties into one coherent action space quickly enough to remain adaptive.

Even where centralised solutions can in principle be engineered, the resulting systems often achieve tractability by narrowing the problem in advance; for instance by simplifying the body, constraining the environment, discretising the action space, or outsourcing uncertainty handling to assumptions that do not hold outside controlled benchmarks. None of that is illegitimate, and much of robotics depends on exactly these simplifications, but it does mean that success under centralised control is often purchased by reducing the extent to which the system is genuinely embodied, open-ended, or behaviourally flexible [20, 37, 35]. As soon as those constraints are relaxed, and especially as agents are expected to cope with richer worlds through richer bodies, the limitations of monolithic control begin to re-emerge.

The point here is not that centralisation is always wrong, but that embodiment introduces structural pressures against treating a single inferential unit as the default architecture for intelligence. Once adaptive behaviour depends on multiple sensorimotor interfaces, multiple timescales, and multiple local contingencies unfolding in parallel, it becomes increasingly natural to think in terms of distributed inferential processes, hierarchical organisation, and the preservation of partially distinct control tendencies rather than their immediate collapse. The limitations of centralised control are therefore not simply engineering inconveniences. They are part of the reason why embodied intelligence may require a different computational grammar altogether.

2.3 Why robotics still often defaults to centralisation

If the limitations of centralised control are so familiar, it is worth asking why robotics continues to return to it. Part of the answer is simply that centralisation is often the most straightforward way to make a difficult problem technically manageable, because it allows perception, planning, and action to be organised around a single optimisation loop, a single state estimate, or a single control policy whose behaviour can be analysed, tuned, and benchmarked with relative clarity. In engineering terms, that kind of simplification is not a weakness but a practical achievement, and it has underpinned a great deal of impressive progress across robotics, control, and machine learning.

This is especially true in domains where bodies are relatively rigid, tasks are well specified, and environments can be constrained sufficiently for centralised models to perform reliably. Under those conditions, a monolithic controller often appears not only sufficient but preferable, since it promises unified decision-making, clean interfaces between modules, and an obvious unit at which uncertainty, reward, or task success can be formalised. Much of the success of modern robotics has therefore been built on a productive narrowing of the problem space, in which morphology is simplified, action spaces regularised, sensing curated, and the world is rendered tractable enough for centralised control to work well [35, 20, 37].

There is likely also a deeper conceptual reason why centralisation keeps reappearing; namely that many of the dominant formalisms in AI and control theory are themselves written in a way that encourages it. Reinforcement learning is commonly framed around a single agent maximising a single return, classical control around a policy minimising a single cost functional, and even probabilistic frameworks often culminate in a single posterior summary or a single optimal plan. None of this prevents richer architectures in principle, but it does mean that the default mathematical language of intelligent behaviour often already presumes a unitary unit of integration [37, 38, 16]. Once that assumption is built into the formalism, it becomes much easier for implementation choices to follow the same path.

Benchmark culture reinforces this tendency further: systems are often evaluated in terms of whether they solve a task, maximise a score, or outperform a baseline under controlled conditions, and this naturally favours architectures that are easy to compare, easy to train, and easy to attribute performance to. Distributed, hierarchical, or plural architectures, by contrast, tend to be messier both conceptually and experimentally, because they raise questions about internal coordination, division of labour, and evaluation across multiple interacting processes rather than a single performance metric. As a result, even when richer embodied architectures may be the more natural fit for real-world intelligence, the pressures of experimental design (and publication culture) can still push development toward cleaner, more centralised solutions.

For much the same reason, centralisation also remains attractive because it offers a reassuring picture of where intelligence is supposed to reside. A single controller, policy, or world model provides an identifiable *centre of competence*, a place in the system where one expects the real intelligence to be found. Distributed and plural architectures unsettle that expectation, because they imply that adaptive behaviour may emerge from the interaction of multiple partial processes, none of which is fully in charge on its own. While that is often closer to what biological systems appear to do, it is less comfortable from the standpoint of design, interpretation, and explanation, since it replaces the logic of command with the logic of coordination.

None of this means that robotics has simply made a conceptual error. In many cases, centralisation has been exactly the right trade-off, and there are plenty of tasks for which it will remain entirely appropriate. The point, rather, is that the persistence of centralised control reflects not only its strengths but also the methodological and conceptual habits of the field, including the kinds of bodies that are easiest to build, the kinds of tasks that are easiest to benchmark, and the kinds of formalisms that are easiest to optimise. Once those habits are recognised, it becomes easier to see why embodied intelligence may require a partial break from them. If agents are to operate through richer bodies, in richer worlds, and under richer forms of uncertainty, then the continued default to centralisation may increasingly reflect the limits of our preferred tools rather than the natural form of intelligence itself.

3. Active Inference as a Framework for Embodied Control

3.1 Active inference in brief

Active inference provides a useful framework because it begins from the premise that perception, action, and learning should not be treated as fundamentally separate problems, but as different aspects of the same underlying imperative; namely, to maintain adaptive behaviour by reducing the discrepancy between predicted and encountered sensory states under a generative model of the world [8, 3, 25]. Rather than casting perception as passive estimation followed by an independent control stage, active inference places the agent inside a closed sensorimotor loop in which beliefs are updated through sensory evidence, while actions are selected in ways that change the sensory data the agent is likely to receive.

At the core of the framework is a generative model, which specifies how hidden states of the world and body are expected to give rise to observations. Because the true posterior over those hidden states is generally intractable, the agent instead maintains an approximate posterior, and updates it so as to minimise variational free energy, a quantity that bounds surprise under the model and provides a tractable objective for inference [8, 3]. In standard form this can be written as

$$F[q(s)] = \mathbb{E}_{q(s)}[\ln q(s)] - \mathbb{E}_{q(s)}[\ln p(o, s)], \tag{1}$$

where $q(s)$ is an approximate (*Variational*) posterior over hidden states s, o denotes observations, and $p(o, s)$ is the generative model linking hidden states and observations. In simple terms, perceptual inference corresponds to adjusting beliefs so that predicted observations better match incoming sensory evidence, while taking prior expectations into account. What makes the framework distinctive, however, is that action is brought under the same principle: agents do not merely revise beliefs to fit sensations, but also act so that sensations come to better fit their predictions.

This changes the usual relation between perception and control in an important way because action is no longer treated as the execution of a command computed elsewhere, but as part of the process by which uncertainty is resolved and predictions are fulfilled. In an embodied setting, that matters immediately. Movements of the eyes, head, hand, body, or sensory platform alter what can be sampled, what remains hidden, and what can be inferred next, which means that action is not an output appended to cognition after the fact, but one of the means by which the agent tests and refines its model of the world. Perception and action are therefore coupled not just practically but mechanistically, since both are in the service of reducing prediction error, either by updating beliefs or by changing sensory input [8, 23, 6].

When future actions must be chosen under uncertainty, active inference extends this logic through expected free energy (EFE), which scores prospective policies not only in terms of their likely pragmatic consequences, but also in terms of the uncertainty they are expected to resolve. In schematic form one can write

$$G(\pi) \approx \underbrace{\mathbb{E}_{q(o|\pi)}[-\ln p(o)]}_{\text{pragmatic value}} + \underbrace{\mathbb{E}_{q(o|\pi)}[H(q(s \mid o, \pi))]}_{\text{expected ambiguity / uncertainty}}, \tag{2}$$

such that candidate policies π are valued not only in terms of preferred outcomes, but also in terms of the uncertainty they are expected to reduce. This allows epistemic and goal-directed behaviour to be handled within a common formal scheme rather than being treated as separate drives stitched together after the fact [9, 23, 25]. Agents can therefore be understood as selecting actions that both realise preferred outcomes and place them in states from which the world becomes more predictable, more informative, or less ambiguous.

The appeal of active inference lies in the kind of architecture it invites; it treats behaviour as belief-guided engagement with the world, rather than as detached optimisation over a fixed external state. It is naturally suited to settings in which uncertainty, embodiment, and action-dependent sampling are central. At the same time, the framework does not by itself dictate how inference must be organised within an agent; Active inference can be implemented in highly centralised ways, but it can also be extended toward distributed, hierarchical, and plural forms of organisation. The claim here is that, for embodied intelligence, those richer forms may often be more natural.

3.2 Why active inference is naturally suited to embodiment

What makes active inference especially attractive in the present context is that embodiment is not something that has to be bolted onto the framework from outside, but is already implicit in the way the problem is posed. Because the agent is understood as inferring the hidden causes of its sensations while acting so as to change those sensations, perception and action are coupled from the outset, and the body is therefore placed directly inside that loop rather than treated as a peripheral mechanism for executing decisions made elsewhere [8, 3, 25]. In that sense, active inference begins from a picture of intelligence that is already much closer to real embodied agents than many frameworks that separate world-modelling, decision-making, and control into distinct sequential stages.

This matters because embodied agents do not simply receive information about the world, but actively sample it through movement, orientation, and engagement. What can be inferred at any moment depends on where the agent is, what it is able to sense from that location, what it has chosen to attend to, and what kinds of policies/actions remain available from its current bodily configuration. Active inference captures this naturally, since action is not merely in the service of reward acquisition or command execution, but can also be epistemic, directed toward reducing uncertainty and bringing hidden structure into view [9, 23, 6]. In embodied settings, where uncertainty is often perspective-bound and action-dependent, that is not a minor advantage but a central requirement.

The framework is also well suited to embodiment because it treats behaviour as emerging under a generative model that can, in principle, include the body itself, its sensory surfaces, its expected dynamics, and its relation to environmental structure. This allows the agent's own morphology to enter directly into the inferential problem, rather than appearing only as an external constraint on a planner that is otherwise assumed to operate over a disembodied description of the world. Under active inference, the body is part of what must be modelled, predicted, and regulated, which makes it possible to think about adaptive behaviour in a way that is sensitive to viewpoint, self-motion, occlusion, and the changing sensorimotor contingencies through which the world is encountered [27, 31].

Further merit lies in the fact that active inference offers a common language for phenomena that are often fragmented across other approaches. Perceptual updating, uncertainty reduction, homeostatic regulation, action selection, and preference realisation can all be expressed within the same broad framework, which is particularly useful when dealing with embodied agents whose behaviour cannot easily be decomposed into neat modules without losing something important [8, 25]. This does not mean that all such processes should literally be collapsed into one homogeneous mechanism, but rather they can be related within a single theoretical framework, instead of patched together from incompatible assumptions about sensing, planning, and acting.

For these reasons, active inference is arguably one of the most natural frameworks for thinking about embodied intelligence, especially in settings where the agent must maintain adaptive behaviour under uncertainty while continuously shaping its own evidence through action. Yet, that promise should not be confused with an automatic solution to the architectural questions raised earlier. A framework may be inherently compatible with embodiment while still being implemented in ways that are more centralised, more monolithic, or more collapse-prone than embodiment itself would appear to demand.

3.3 The hidden centralist bias in many active inference implementations

Although active inference is often presented as a framework that naturally accommodates embodiment, uncertainty, and closed-loop interaction with the world, many of its concrete implementations still retain a more centralised architecture than the broader picture might suggest. This is not usually stated as an explicit theoretical commitment, and in many cases it is simply a pragmatic consequence of tractability - but the effect is that inference often ends up being organised around a single dominant posterior, a single policy distribution, or a single unit at which uncertainty is gathered, resolved, and translated into action. In that respect, the framework can remain formally embodied while still being architecturally centralised in practice.

Part of the reason for this is that active inference inherits many of the same pressures that shape other computational approaches. When generative models are simplified to support tractable inference, when policies are evaluated within a single shared action space, when belief updating is expressed through a unified optimisation scheme; it becomes natural to treat the agent as if it were integrating all relevant uncertainty into one coherent internal estimate from which behaviour then follows. That picture can be appropriate in relatively simple tasks, and indeed much of the clarity and elegance of active inference comes from precisely this kind of unification. At the same time, it can also obscure the possibility that, in richer embodied systems, intelligent behaviour may depend less on full internal collapse than on the ongoing coordination of multiple partially distinct inferential processes.

This tendency becomes visible in several ways. In many formulations, for example, competing hypotheses are ultimately resolved through a single approximate posterior, even when the underlying problem is multimodal, only partially decomposable, or distributed across different sensorimotor interfaces. Likewise, policy selection is often expressed in terms of a single posterior over candidate policies or a single expected free energy landscape, such that plurality is treated as something to be integrated-away rather than maintained as a productive feature of the system's organisation [9, 23, 6]. None of this is mathematically unreasonable, but it does mean that the default image of the agent remains one in which coherence is achieved primarily by central resolution rather than by coordination across semi-autonomous components.

The same issue appears at the level of control architecture. Even when a model includes hierarchies, precisions, or multiple latent factors, these are often embedded within a single inferential machine that effectively speaks with one voice at the moment of action. A genuinely embodied agent, particularly one operating through a complex body in a partially observed world, may need to preserve different local uncertainties, different control tendencies, and different inferential timescales without requiring that all of them be collapsed into a single momentary summary before anything useful can happen. What matters in such systems may not be perfect central integration, but viable coordination.

There is a broader conceptual point here as well. Because active inference is frequently introduced as a unified framework, there is a temptation to assume that explanatory unity at the formal level should imply architectural unity at the implementation level. Yet those two claims need not coincide. A single mathematical principle may still be realised through distributed, hierarchical, and plural organisations, just as biological systems often exhibit shared constraints and common imperatives without relying on a single central controller to enforce them in detail. The key point here is not that active inference must abandon integration, nor that all successful agents should be fragmented into loosely coupled modules. Rather, it is that the strongest implications of embodiment may push the framework toward architectures in which inference is distributed across interfaces with the world, structured across multiple levels of timescale and abstraction, and allowed to remain plural where premature collapse would reduce flexibility or obscure uncertainty. Once that possibility is taken seriously, active inference begins to look less like a theory of centralised self-evidencing and more like a general language for coordinated embodied intelligence.

4. Beyond Monolithic Control: Distributed, Hierarchical, and Plural Inference

4.1 Distributed inference

If embodiment places the agent inside a world that must be sampled through multiple sensorimotor interfaces, and if the limitations of monolithic control become more acute as bodies and environments grow more complex, then the next step is not simply to ask how a single controller might be improved, but whether inference itself should be more widely distributed throughout the system. By distributed inference, we do not mean a mere engineering decomposition in which a central intelligence delegates subroutines to peripheral modules while retaining the real explanatory burden for itself. What we mean instead is that different parts of the agent may carry out partially local inferential work in their own right, updating beliefs, responding to local contingencies, and shaping behaviour without requiring that every uncertainty be gathered into one central unit before anything meaningful can happen.

This matters because, in an embodied agent, different interfaces with the world often confront different problems at different timescales. A mobile platform may need to stabilise its posture while simultaneously tracking a target, avoiding obstacles, maintaining orientation, and preserving some estimate of what lies beyond its current field of view. A hand may need to adjust grip, contact force, and trajectory on the basis of fast local feedback that would be difficult to route through a single bottleneck without sacrificing responsiveness. In such settings, distributing inference is not simply a matter of parallelisation for efficiency, although it may of course bring that benefit. More fundamentally, it reflects the fact that the world is encountered locally, unevenly, and under multiple overlapping constraints, such that adaptive behaviour may be better served by several coupled inferential processes than by a single unified estimate that attempts to absorb everything at once.

Biological systems make this especially clear: sensorimotor control is organised through layered and partly decentralised loops, with peripheral circuits, local reflex-like mechanisms, and intermediate control structures handling fast contingencies while broader contextual influences are imposed over longer timescales and larger spatial domains [14, 17, 39]. Even where there is clear large-scale coordination, that does not imply that every aspect of behaviour is computed centrally in detail. Distributed inference, in this sense, is not the absence of organisation, but a particular way of organising it.

Within active inference, this idea can be expressed quite naturally. If behaviour is guided by generative models that relate hidden states, observations, and actions, there is no principled reason why all such modelling must occur at one place or in one homogeneous inferential stream. Different subsystems may maintain different local beliefs about the states most relevant to their own coupling with the world, while exchanging enough information to remain mutually constraining rather than independent. The important point is that coordination need not depend on complete central collapse. A system may remain coherent even when its inferential labour is distributed across multiple sites, provided those sites are coupled through shared priors, mutual constraints, or ongoing message passing.

Seen in this way, distributed inference offers a different answer to the question of where intelligence resides. Rather than locating intelligence in a single central model that represents the world on behalf of the whole agent, it allows intelligence to emerge from the interaction of multiple inferential processes situated at different interfaces between body and environment. Some of those processes may be highly local, others more integrative, but what matters is that adaptive behaviour arises from their coordination rather than from the detailed command of one over the others.

4.2 Hierarchical inference

If distributed inference addresses the question of where inferential work is carried out, hierarchical inference addresses the question of how that work is organised across the system. The importance of hierarchy in embodied agents lies not simply in the fact that some processes are more abstract than others, but in the fact that different forms of uncertainty, control, and prediction unfold over very different spatial and temporal scales. Some contingencies must be handled almost immediately at the level of posture, contact, or local orientation, while others concern broader behavioural context, longer-horizon expectations, or more slowly changing beliefs about the structure of the environment. A viable embodied architecture therefore cannot rely on a single inferential timescale without either becoming too slow for local control or too *myopic* for coherent behaviour over time.

What hierarchy makes possible is a division of inferential labour in which faster, more local processes remain responsive to immediate sensorimotor contingencies, while slower and more global processes shape the broader conditions under which those local adjustments occur. In that sense, hierarchy should not be considered the reintroduction of monolithic command - the role of higher levels is not to specify every detail that lower levels must implement, but to provide contextual priors, constraints, or expectations that modulate local inference without replacing it. Lower levels, in turn, do not merely execute instructions. They resolve fine-grained uncertainties that would be inaccessible, irrelevant, or computationally wasteful for higher levels to represent directly.

This is especially important in embodied settings, because bodies do not encounter the world all at once or at a single scale. A moving agent may need, within the same moment, to stabilise itself locally, track salient objects, preserve a sense of orientation, and maintain some longer-term expectation about where it is heading or what task it is currently pursuing. Those are not just multiple tasks, but multiple inferential horizons. If all are forced into one level of description, then either fine-grained control becomes overloaded with long-range structure, or long-range reasoning becomes swamped by transient local detail. Hierarchical inference offers a way of preventing that collapse by allowing different levels of the system to carry different burdens while remaining coupled to one another.

Biological systems again provide a useful reference here. Across nervous systems, one repeatedly finds patterns in which local loops handle rapid sensorimotor variation while broader contextual organisation is maintained over longer timescales and larger functional domains [17, 39, 14]. The significance of this is not just that biology happens to be hierarchical, but that hierarchy appears to be one of the ways embodied systems remain both responsive and coherent without requiring exhaustive central oversight. In the octopus, for example, peripheral control at the level of the arms is not independent of the rest of the organism, yet neither is it micromanaged in full detail by a central controller. What emerges instead is a pattern of bounded autonomy in which local processes remain locally adaptive under broader behavioural constraints.

Active inference is particularly well suited to this kind of organisation, because generative models can in principle be arranged hierarchically, with higher levels encoding slower, more abstract, or more contextual expectations, and lower levels handling faster, more immediate prediction errors at the interface with the world [8, 10, 25]. Under that arrangement, higher levels do not need to represent every sensory fluctuation directly, and lower levels do not need to carry a full account of the agent's broader goals or environmental structure. Instead, information can flow in both directions, with top-down predictions constraining local inference and bottom-up prediction errors updating more global beliefs. What matters is not that the hierarchy converges on a single perfect model at every moment, but that it allows behaviour to remain coordinated across scales that would otherwise interfere with one another.

Hierarchical inference therefore does more than add another layer of organisation to a distributed system; it provides one of the main reasons why distribution need not degenerate into fragmentation. Once inferential processes are structured across levels of timescale and abstraction, local responsiveness and global coherence no longer have to compete so directly.

4.3 Plural inference

Even once inference is distributed across the body and organised hierarchically across levels of timescale and abstraction, there remains a further tendency in many models of intelligence to assume that adaptive behaviour must ultimately be achieved by collapsing uncertainty onto a single dominant estimate, a single prevailing hypothesis, or a single policy selected as best overall. In some settings that may be entirely appropriate, and there are many cases in which convergence on one workable interpretation or action tendency is both efficient and sufficient. Yet as environments become more uncertain, more open-ended, and more tightly coupled to the agent's own embodied perspective, it is fair to assume that intelligence may sometimes depend not on eliminating plurality as quickly as possible, but on maintaining it in a structured and productive form.

Plural inference, then, means that multiple partially compatible hypotheses, control tendencies, or inferential voices may remain active in parallel, each capturing different aspects of the agent's relationship to the world and each continuing to exert some influence on behaviour without being reduced immediately to a single winner. This does not imply that all alternatives are weighted equally, nor that the system remains indefinitely undecided. Rather, adaptive organisation may depend on preserving some degree of heterogeneity within the inferential process itself, such that alternative interpretations, priorities, or action tendencies remain available while the agent continues to engage with the world. In that sense, plurality is not the opposite of intelligence, but may under some conditions be one of its enabling conditions.

This becomes especially important in embodied systems, because different aspects of behaviour often place partially competing demands on the agent at the same time. A robot may need to remain goal-directed while preserving safety margins, maintaining orientation, resolving uncertainty about an occluded target, and conserving enough flexibility to respond if the scene changes unexpectedly. A manipulator may need to commit to a trajectory while still remaining sensitive to contact, slippage, or new local evidence that alters what the task requires. In such cases, forcing all relevant structure into a single scalar objective or a single dominant policy may achieve decisiveness at the cost of brittleness. What is lost is not just alternative action, but the capacity to remain poised between multiple partially valid modes of engagement with the world.

Biological systems again suggest that this kind of plurality is not merely a theoretical possibility. Neural and behavioural organisation appears to involve overlapping representations, partially coexisting action tendencies, and transiently competing interpretations that are not instantly suppressed the moment a better alternative becomes available [39, 32, 11]. More generally, cognition in uncertain environments frequently seems to rely on the retention of partially unresolved structure, whether in perception, action preparation, or the coordination of multiple behavioural imperatives. From this perspective, intelligence is not always best understood as the rapid elimination of alternatives, but as the controlled maintenance of alternatives under shared constraints.

Within active inference, the importance of plurality moves beyond the assumption that all inferential labour must culminate in a single momentary resolution. If policies are evaluated under uncertainty, if different parts of the agent encounter different local contingencies, and if behaviour must remain flexible enough to absorb new evidence without becoming unstable, then there is no reason in principle why all meaningful inferential structure should be collapsed into a single dominant posterior or policy at every step. A system may instead preserve multiple partially distinct trajectories through belief- or policy-space, allowing them to constrain one another, coexist, and contribute jointly to behaviour without any one of them being granted absolute control [33].

Taken together, distributed, hierarchical, and plural inference describe not three separate embellishments to an otherwise standard controller, but three aspects of a different architectural framework. Inference is distributed because the world is encountered through multiple sensorimotor interfaces, hierarchical because those interfaces and uncertainties unfold across different scales, and plural because adaptive behaviour may depend on preserving multiple partially distinct inferential tendencies rather than resolving them too early.

4.4 Coordination without command

What kind of organisation is adequate to embodied agents acting in uncertain worlds? The argument developed thus far points toward a picture in which coherence does not depend on the existence of a single unit that gathers all uncertainty, resolves all competition, and issues the final command, but instead emerges from the interaction of multiple inferential processes that remain locally responsive, globally constrained, and only partially collapsed into one another. In that sense, the contrast is not between order and disorder, or integration vs. fragmentation, but between two different models of organisation: one based on central resolution, and another based on coordination among semi-autonomous processes.

What makes this second option viable is that distributed, hierarchical, and plural inference need not imply arbitrariness. A system can remain coherent without being monolithic, provided its constituent processes are coupled strongly enough to constrain one another, weakly enough to preserve local responsiveness, and organised in a way that allows behaviour to remain mutually compatible across scales. In compact form, one can express this kind of organisation through a coordinated free-energy functional of the form

$$\mathcal{F}_{\text{coord}} = \sum_{k=1}^{K} \pi_k F_k + \sum_{i<j} \lambda_{ij}\, C(q_i, q_j), \tag{3}$$

where F_k denotes the local inferential objective of process k, π_k its bounded influence on the whole, and $C(q_i, q_j)$ a compatibility (or covariance) term coupling pairs of local posteriors or belief states. The point is not that all embodied systems literally implement this exact functional, but that global coherence can be understood as emerging from multiple coupled inferential processes rather than from one controller collapsing uncertainty. Coordination, in this sense, is neither the absence of control nor a diluted form of command. It is a mode of organisation in which stability is achieved through ongoing relational adjustment, with local processes shaping one another under shared priors, bodies, environments, and behavioural demands.

This is important because many of the phenomena that matter most for embodied intelligence are not well captured by the image of a single controller solving a well-posed problem from above. Real agents must maintain posture while pursuing goals, remain sensitive to unexpected local evidence while preserving longer-term direction, and continue acting under conditions in which the world is only partly known and only partly knowable from any one vantage point. Under such conditions, coherence is less likely to emerge from the full elimination of uncertainty than from the ability to remain adaptively organised in its presence. Coordination without command is precisely this: a system in which multiple inferential tendencies remain active, mutually constraining, and behaviourally aligned enough to support action without requiring complete internal unanimity.

From this perspective, *viability* begins to displace *optimality* as the more relevant criterion. The question is no longer simply whether the agent has identified the single best hypothesis, policy, or control state, but whether it can remain poised in a configuration that is sufficiently coherent, flexible, and responsive to continue engaging the world adaptively as conditions evolve. That does not exclude local optimisation, nor does it deny that some forms of convergence may occur within parts of the system, but what it resists is the assumption that intelligence as such should always be understood as the collapse of multiplicity into one final solution. In embodied systems, the ability to preserve and coordinate multiple partially distinct processes may itself be part of what makes intelligent behaviour possible.

For active inference, this reframing has significant consequences. If the framework is taken seriously as a theory of embodied engagement with the world, then its implementation need not be tied to a single central inferential machine that speaks with one voice at every moment of action. It can instead be understood as a more general language for organising coupled belief-updating processes across bodies, timescales, and behavioural demands. Active inference therefore becomes less a theory of unitary internal resolution and more a theory of how adaptive systems maintain coherence under uncertainty through structured patterns of mutual constraint. That is much closer to the kind of intelligence suggested by flexible biological systems, and much closer, we argue, to the kind of architecture that embodied robotics will increasingly require.

5. Biological Inspiration: The Octopus as a Model of Embodied Inference

5.1 A nervous system distributed through the body

Among biological systems, the octopus is especially striking because it makes visible, in unusually clear form, an organisational principle that is often only implicit elsewhere: that intelligence need not be concentrated in a single central unit in order to remain coherent, adaptive, and behaviourally sophisticated. What makes the octopus so compelling in this regard is not simply that it is flexible, dexterous, or behaviourally rich, but that its nervous system is distributed through the body to an extent that sharply departs from the vertebrate template. A substantial proportion of its neurons are located not in the central brain, but in the arms themselves, embedded within peripheral circuits that support local sensing, control, and action selection [14, 15, 12].

This architecture matters because it changes what it means for the animal to perceive and act. The arms are not merely passive effectors awaiting detailed commands from a centralised controller, but semi-autonomous sensorimotor systems that can explore, grasp, probe, and adjust in close relation to local contingencies. Tactile, proprioceptive information is gathered and acted upon at the edge of the body, where the relevant uncertainty is encountered, rather than being routed in full to a single processing unit for exhaustive central interpretation [14, 22, 4]. In computational terms, this makes the octopus difficult to understand through the image of a unitary controller solving a complete inverse problem for the whole body at once. What one sees instead is a nervous system in which inferential labour appears to be distributed across multiple bodily interfaces with the world.

This is particularly important given the morphology of the octopus itself. Each arm is a highly flexible muscular hydrostat with a vast number of degrees of freedom capable of bending, elongating, twisting, and stiffening in ways that resist simple central parametrisation. If control were organised in a strictly centralised fashion, the computational demands would be enormous, since the brain would need to continuously specify and update a detailed mapping from goals to local motor states across a body whose shape and coupling to the environment are changing moment by moment. The distributed organisation of the octopus nervous system offers a way around that problem, because much of the required sensorimotor adaptation can occur locally within the limbs themselves, without every detail having to be represented and resolved centrally [18, 7, 14].

What follows from this is not that the octopus lacks central organisation, nor that the arms function as fully independent agents. The central brain clearly remains crucial for broader behavioural context, task selection, integration across modalities, and the coordination of whole-animal behaviour. The point, rather, is that the body itself participates directly in the production of intelligence, with peripheral circuits carrying out forms of local sensing and control that would be difficult to explain if one assumed that all meaningful computation had to occur centrally. In that respect, the octopus provides a vivid biological instance of the claims developed so far: adaptive embodied behaviour may depend on inference being distributed across the organism, rather than concentrated entirely within one commanding centre.

Seen in this way, the octopus is not merely a curiosity for comparative neurobiology, but a useful model for thinking about embodied inference more generally. Its nervous system shows that coherent behaviour can emerge from a system in which control is pushed far into the periphery, local processes remain active at the interface with the world, and intelligence is realised through the coordination of distributed bodily subsystems rather than through detailed central micromanagement. That makes it an unusually instructive example for robotics, where many of the same questions arise once agents are expected to act through flexible morphologies, rich sensorimotor couplings, and environments that cannot be mastered from a single vantage point.

5.2 Hierarchical bounded autonomy

What makes the octopus especially informative for the present argument is not just that its nervous system is distributed, but that this distribution is organised in a way that preserves both local autonomy and broader behavioural coherence. The arms are not independent agents pursuing wholly separate goals, nor are they passive appendages awaiting detailed central instruction. Instead, they appear to operate under *hierarchical bounded autonomy*, in which local sensorimotor processes retain the capacity to respond adaptively to immediate contingencies, while broader behavioural context is shaped by more central systems concerned with the state of the organism as a whole.

This is important because it illustrates a form of organisation that is neither fully centralised nor fragmented. At the level of the arms, the octopus can rely on peripheral circuits to handle rapid, fine-grained adjustments in posture, contact, reaching, and exploration, drawing on local sensory signals that would be cumbersome or unnecessary to route through a single central bottleneck. At the same time, these local processes do not unfold in an unconstrained vacuum. They are modulated by higher-level conditions relating to what the animal is doing, whether it is hunting, escaping, orienting, or interacting with some feature of the environment, and by the need for behaviour across the whole organism to remain sufficiently coherent to support those broader aims [14, 15, 12].

From an active inference perspective, the appeal of this arrangement is fairly obvious: lower levels can be understood as resolving fast local uncertainties at the interface between body and world, while higher levels encode slower, broader priors concerning context, goals, and expected behavioural regimes. That does not mean that the central brain computes a complete solution and passes it down for execution. Rather, it shapes the inferential landscape within which local processes operate, biasing peripheral control toward some classes of action and away from others without specifying every detail in advance. In that sense, higher levels constrain rather than command, while lower levels adapt rather than merely obey.

This pattern helps explain how the octopus can remain both flexible and coherent despite the enormous control problem posed by its body. A fully centralised solution would be computationally burdensome and likely too slow or too brittle to handle the moment-to-moment contingencies encountered by highly flexible arms interacting with an uncertain environment. A fully decentralised solution, by contrast, would risk behavioural fragmentation, with local processes competing or drifting in ways that undermine whole-animal goals. Hierarchical bounded autonomy occupies the middle ground. It allows local systems to carry out the inferential work they are best placed to do, while ensuring that their activity remains embedded within a larger behavioural organisation [7, 22].

What is especially striking here is that the relation between levels is not well described by the language of top-down command. The central brain does not appear to micromanage each degree of freedom in the arms, and the arms do not merely transmit raw information upward for all meaningful interpretation to occur elsewhere. Instead, intelligence is distributed across levels, with both local and global processes participating in reciprocal organisation and constraint.

For robotics, this matters because many embodied agents face a closely analogous problem: rich sensorimotor control often depends on fast local adjustments that are best handled near the interface with the world, yet those adjustments must still remain compatible with broader task demands, safety constraints, and longer-horizon behavioural organisation. The octopus shows that this is not a biological curiosity but a viable architectural principle. Local autonomy need not undermine coherence, provided it is bounded by higher-level context, and higher-level organisation need not collapse into micromanagement in order to remain effective.

6. Computational Case Study: A Temporally Predictive Scene-Based Drone Controller

6.1 Why a drone?

The drone[1] is a particularly useful platform for the present case study because it brings several of the central challenges of embodied intelligence into sharp focus at once. It must act through a body that is dynamically unstable, move through a world that is only ever partially observed, and maintain coherent behaviour under conditions in which sensing, prediction, and control are tightly entangled from moment to moment.

[1] A video of the controller in operation is available at https://cpnslab.com/drone.html.

Unlike more static platforms, an aerial agent cannot easily afford to stop, deliberate, and then resume action from a stable resting configuration. Instead it must remain continuously engaged in the regulation of its own movement, orientation, and relation to the environment, even while pursuing higher-level behavioural aims such as tracking, navigation, target maintenance, or scene interpretation. In that respect, the drone makes visible something that is sometimes easier to ignore in slower or more constrained systems - namely, that embodied intelligence often depends on the simultaneous coordination of multiple inferential demands unfolding at different timescales. Local stabilisation, short-horizon adjustment, and broader task-directed behaviour are not separate phases, but overlapping aspects of one ongoing control problem.

The drone is also a particularly useful platform because it encounters the world *perspectivally*. What it can infer depends directly on viewpoint, motion, occlusion, and the geometry of its own sensing, and those factors are themselves altered by action. A target may disappear behind an obstacle, reappear from a different angle, or remain only partially specified while the agent must nevertheless continue to behave coherently. This makes the platform a natural setting in which to study active inference as embodied evidence-gathering rather than as abstract state estimation alone, because action does not merely implement a decision, but helps determine what can be known next and how uncertain parts of the scene can be brought back into view.

For these reasons, the drone provides a useful bridge between simple reactive control and richer scene-based forms of embodied inference. It offers a compact but demanding setting in which memory, prediction, action-dependent uncertainty, and multi-timescale coordination are not optional extras, but practical necessities.

6.2 From reactive control to scene-based active inference

If we consider the drone as an embodied agent moving through a partially observed world, the limitations of purely reactive control become immediately apparent: a controller organised only as a mapping from current observation to current action can respond effectively only to what is presently visible and salient. That may be adequate in simple settings where all task-relevant structure remains continuously available and the consequences of action are short-lived and transparent, but it becomes much less satisfactory once objects pass in and out of view, scenes have hidden depth, and behaviour depends on preserving some grasp of what is no longer directly observed.

This matters because embodied agents do not act in instantaneous worlds. They act in scenes that unfold over time, within which visible evidence is only one part of a larger and often partially latent structure. A target may be briefly occluded while still remaining behaviourally relevant, and a movement may matter not only because of what it reveals now, but because of how it positions the agent with respect to what is likely to become visible next. Under such conditions, the controller cannot remain tied to the sensory information alone. Instead, it must maintain some inferential relation to a scene that persists beyond immediate observation.

What follows is that active inference, in this setting, has to become more than a mechanism for online sensorimotor correction. It must support the maintenance and updating of latent scene structure over time, so that action can be guided not only by present evidence, but by beliefs about what continues to exist, where uncertainty is accumulating, and how the world is likely to

Controller at a glance

1. Observe self-state, target cue, and local ray geometry
2. Update fast beliefs over self, target, and local map
3. Update slow scene beliefs over visibility, progress, affordance, and memory
4. Roll out short-horizon candidate actions
5. Predict how the scene will evolve under each candidate
6. Score actions by expected free energy plus scene value
7. Execute best action and repeat

Figure 1: Compact overview of the temporally predictive scene-based drone controller. Fast local beliefs and slower scene-level beliefs are updated online, candidate actions are rolled out over a short horizon, and policy evaluation depends on both expected free energy and predicted scene improvement. Right: representative frame from the PyBullet drone simulation.

evolve as the agent moves through it. The importance of temporal prediction follows directly from this; a controller that represents only what is visible now is always in danger of becoming myopic, especially when behaviour depends on continuity across moments in which relevant structure is intermittent or partly concealed.

The shift from reactive control to scene-based active inference therefore changes what the controller is taken to be doing. Rather than merely reducing instantaneous mismatch between observation and action, the agent is engaged in maintaining a structured relationship to a world that extends beyond what is currently seen. Action becomes a way not only of correcting current error, but of preserving contact with latent scene variables, managing uncertainty over time, and moving into positions from which hidden structure can be recovered or stabilised. For the present case study, that shift is not simply a matter of making the controller more elaborate. It is what allows the architecture to address the actual demands of aerial behaviour in a partially observed environment.

6.3 A temporally predictive scene-based controller

The drone controller was designed to move beyond reactive control while remaining lightweight enough for continuous embodied operation. Rather than mapping current observation directly to immediate action, it maintains a structured belief state in which fast sensorimotor estimates are coupled to a slower scene-level description of what kind of situation the agent is in and how that situation is likely to evolve. In that sense, it is scene-based not simply because it processes perceptual input, but because it preserves an inferential relationship to a partially latent environment whose relevant structure extends beyond what is visible at any one moment.

At the faster level, the controller maintains beliefs about the drone, the target, and local spatial structure derived from ray-based sensing, and immediate observation. These support the moment-to-moment regulation required for aerial behaviour including position, orientation, target localisation, and obstacle-sensitive adjustment. The key architectural addition is a slower layer of scene beliefs, which tracks latent variables such as visibility, progress, affordance, broader behavioural context, and target-memory reliability. These slower variables function as part of the controller's working generative 'picture' of the scene, allowing it to distinguish between, for example, clear progress, poor vantage, or partial occlusion with only degraded target memory.

What makes the controller temporally predictive is that its scene variables evolve. The controller estimates not only the current scene, but how that scene is likely to change under candidate actions. Visibility can improve or degrade, progress can accumulate or stall, and memory reliability can strengthen or decay. Policy evaluation is therefore shaped not only by immediate error reduction, but by whether an action is expected to move the agent toward a better future scene.

In practical terms, candidate actions are evaluated through an expected-free-energy-like objective that combines pragmatic and stability considerations with epistemic value over the evolving scene:

$$J(\pi) = -\,G(\pi) + \alpha\,\mathrm{Vis}(\pi) + \beta\,\mathrm{Prog}(\pi) + \delta\,\mathrm{MemRel}(\pi), \tag{4}$$

where policies π are valued according to expected free energy together with predicted visibility, progress, and memory reliability over a short horizon. Candidate actions are therefore favoured when they are expected to recover line of sight, improve vantage, reduce scene uncertainty, preserve target memory under occlusion, or move the drone into states from which further progress becomes more likely. This is the main reason the controller behaves less like a reflexive tracker and more like an embodied agent maintaining contact with a partially hidden world.

The resulting organisation is already some distance from monolithic reactive control; inferential labour is distributed across fast and slow components rather than collapsed into a single flat state estimate. It is hierarchical in the sense that local sensorimotor variables and slower scene-level beliefs operate on different timescales while remaining mutually constraining. And it is at least incipiently plural, because behaviour is shaped not only by immediate target pursuit, but by the interaction of visibility recovery, progress maintenance, uncertainty reduction, and memory preservation.

To keep the main text focused on the architectural logic, a fuller mathematical specification of the controller and a compact algorithmic summary are provided in the Supplementary Materials.

6.4 Why this controller instantiates distributed, hierarchical, and plural inference

This controller departs from monolithic reactive control in three recognisable ways. It does not collapse all uncertainty into a single flat state estimate, nor does it reduce action selection to one immediate control demand. Instead, it distributes inferential labour across fast and slow components, organises those components across different timescales, and allows several partially distinct behavioural pressures to remain active together during policy evaluation.

It is distributed because fast local estimates relating to the drone, the target, and nearby spatial structure continue to operate at the interface with the environment, while slower scene-level variables track broader aspects of visibility, progress, affordance, context, and memory reliability. These correspond to different aspects of the agent's coupling with the world, each carrying its own uncertainties and behavioural relevance, rather than being folded immediately into one homogeneous controller state.

It is hierarchical because those components do not operate on the same timescale or at the same level of abstraction. The fast layer is concerned with immediate regulation and local geometry, while the slower scene layer tracks more persistent and behaviourally consequential features of the situation, including visibility, progress, navigability, and the reliability of target memory under occlusion. Scene-level beliefs therefore shape what kinds of actions become valuable, while fast states continue to absorb local contingencies at the sensory frontier.

It is plural because behaviour is not governed by a single immediate imperative, but by the joint influence of several partially distinct control tendencies that remain active during policy evaluation; target pursuit, visibility recovery, progress maintenance, obstacle sensitivity, uncertainty reduction, and memory preservation do not always point in exactly the same direction.

Taken together, these properties reinforce one another. Distribution without hierarchy would risk little more than loosely connected local estimates. Hierarchy without plurality could still collapse into rigid multi-level command. Plurality without either of the other two could become unstable or incoherent. Here, by contrast, fast and slow scene-sensitive processes are coupled in a way that allows multiple behavioural pressures to remain active without losing enough organisation for the drone to act coherently. The controller is therefore best understood as a computational stepping stone: a proof of principle that once scene persistence, temporal prediction, and multiple partially distinct behavioural demands are allowed into the loop, even a relatively compact embodied agent begins to depart from the logic of monolithic reactive control.

6.5 Behavioural consequences

The most immediate consequence of this architecture is that behaviour becomes less tightly bound to the sensory present. Because the controller does not rely solely on what is currently visible and instead maintains beliefs about a scene that persists through occlusion, motion, and changing viewpoint, the drone is less brittle when relevant structure temporarily disappears from view. Rather than collapsing into aimless search or reverting instantly to a purely local reflex, it can continue to act under a weaker but still meaningful estimate of where the target is likely to be, how reliable that estimate remains, and what kinds of movement are most likely to recover a better vantage point.

This makes a noticeable difference when target visibility is intermittent or when direct pursuit would otherwise lead the agent into poor geometrical positions. A purely reactive controller will often chase whatever is currently salient, even when doing so reduces future visibility or leads the agent into states from which recovery becomes harder. By contrast, a controller that values scene-level variables such as visibility, progress, affordance, and memory reliability can behave more strategically without requiring a heavy planner in the classical sense. It may sacrifice a small immediate gain in target alignment in order to preserve line of sight, maintain navigability, or move into a position from which uncertainty can be reduced more effectively over subsequent steps.

A second consequence is that behaviour becomes more continuous across changes in the scene. Because the controller maintains slower beliefs about the evolving situation, sudden shifts in visibility or local geometry do not have to produce equally sudden shifts in the entire control policy. Instead, the agent can absorb such changes into an ongoing inferential process in which local correction and broader scene interpretation remain coupled. This supports smoother recovery under partial occlusion, more stable reorientation when the target reappears from a different angle, and less oscillatory behaviour when immediate local signals are temporarily ambiguous.

The plural structure of the controller also changes behaviour in a useful way. Because policy evaluation is shaped by multiple partially distinct pressures rather than a single instantaneous imperative, the drone is less prone to pathological commitment when one objective dominates too strongly. Target pursuit remains important, but it is tempered by other concerns, including the preservation of visibility, the avoidance of scene states in which progress stalls, and the maintenance of enough inferential flexibility to recover when current evidence becomes unreliable. Behaviour therefore looks less like the relentless execution of one fixed drive and more like the ongoing balancing of partially competing demands under uncertainty.

More broadly, the architecture shifts the qualitative character of the agent from tracking to scene maintenance. The drone is not merely trying to minimise the distance between itself and whatever it currently sees. It is trying to remain in an epistemically and behaviourally workable relationship to a world that is partly hidden, partly changing, and only intermittently available at the sensory surface.

7. Design Implications for Embodied Robotics

The argument developed here suggests a few practical implications for embodied robotics. First, some inferential labour will often need to remain close to the point at which the world is actually encountered, rather than being routed immediately into a single central process. In systems operating under partial observability, rich local sensorimotor coupling, or fast-changing geometry, local inference can preserve responsiveness without sacrificing larger-scale coherence, provided it remains embedded within broader contextual organisation.

Second, embodied control is likely to benefit from architectures organised across multiple timescales and levels of abstraction. Fast local processes and slower scene-level or task-level processes do not solve the same problem, and forcing them into one inferential layer can produce controllers that are either too myopic or too overburdened to remain adaptive. Hierarchical organisation matters here not as a doctrinal commitment, but as a practical response to the fact that embodied intelligence unfolds simultaneously in the fast present and the slower continuity of an evolving scene.

Third, robust behaviour may depend on preserving some bounded plurality within the controller rather than collapsing all uncertainty and all behavioural pressures into a single dominant objective too early. In embodied settings, visibility, stability, progress, safety, memory, and uncertainty reduction may all remain behaviourally relevant at once, and useful control may depend less on rigid arbitration than on ‘soft coordination’ among partially distinct inferential pressures. Seen in this light, variables such as occlusion, memory reliability, and structured uncertainty are not secondary implementation details, but part of the architecture through which embodied agents remain connected to a world that is only intermittently available.

Taken together, these considerations suggest that embodied robotics may often be better served by aiming for viable behaviour than for narrow optimality alone. The issue is not whether optimisation still matters, but whether intelligent embodied action is always best understood as maximising one quantity at a time. In many of the settings that matter most, the more relevant question may be whether the agent can remain coherently, flexibly, and adaptively organised as the scene changes and uncertainty unfolds.

8. Open Questions and Outlook

Although the argument developed here points toward distributed, hierarchical, and plural inference as a promising route for embodied intelligence, a number of important questions remain open. One concerns scale: it is one thing to show, in conceptual terms and in relatively compact computational examples, that inference can be distributed across interfaces with the world, structured across timescales, and coordinated without immediate collapse. But it is another to determine how such architectures behave as bodies become more complex, scenes become richer, and the number of interacting inferential processes grows. Questions of computational cost, stability, and tractable coordination do not disappear once monolithic control is relaxed.

A second open issue concerns learning. Much of the discussion in this paper has focused on architectural principles and proof-of-concept organisation rather than fully learned systems. Yet for embodied robotics, one of the central challenges will be how local inferential processes, scene-level variables, coupling schemes, and policy-evaluation structures are acquired, adapted, or refined through experience rather than specified entirely in advance. Relatedly, it remains unclear how best to determine which uncertainties should remain local, which should be elevated into slower contextual beliefs, and which forms of plurality are genuinely useful rather than merely adding complexity.

There are also questions of transfer and evaluation; for example, architectures that look promising in simulation may behave quite differently once exposed to the friction, latency, sensing imperfections, and morphological unpredictability of real robotic platforms. More importantly, some of the capacities emphasised here, including persistence under occlusion, recovery of visibility, and the maintenance of coherent behaviour under uncertainty, are not especially well captured by metrics centred on reward, efficiency, or immediate task completion. If embodied intelligence is to be studied in richer terms, it may also need to be evaluated in richer ones.

What seems clear, however, is that the broader architectural question is worth taking seriously. As robotics moves toward richer bodies, denser local sensing, and more open-ended environments, the assumption that intelligence should default to a single central controller or a single dominant inferential trajectory may become increasingly hard to sustain. Whether or not future systems are explicitly framed in terms of active inference, the pressures introduced by embodiment are likely to favour architectures that distribute inferential labour, organise it across multiple timescales, and preserve enough plurality to remain flexible when the world does not yield a single stable answer all at once.

9. Conclusion

Embodied intelligence is not well captured by the image of a single controller operating over a passive body and a fully specified world. Once agents are understood as acting through bodies that shape what can be sensed, predicted, and controlled, the computational problem looks different. Perception, action, memory, uncertainty, and scene persistence become tightly coupled, and the assumptions that support monolithic control begin to strain.

Active inference is useful here because it places the agent inside a closed sensorimotor loop and treats action as part of the process by which uncertainty is managed. But the strongest implications of embodiment push beyond active inference in the abstract toward architectures in which inference is distributed across interfaces with the world, organised hierarchically across timescales and levels of abstraction, and allowed to remain plural where premature collapse would reduce flexibility. In this view, intelligent behaviour is better understood as coordination without command than as optimisation by a unitary controller.

The octopus provides a biological example of this logic, showing how coherent behaviour can emerge when control is pushed into the body and local autonomy is bounded by broader behavioural context. The temporally predictive scene-based drone controller provides a corresponding computational case study, illustrating how even a relatively compact embodied agent begins to behave differently once scene persistence, uncertainty over occlusion, and multiple inferential demands are allowed to shape action. If that picture is right, then the future of embodied intelligence may depend less on ever more detailed central control and more on how well artificial agents can distribute, organise, and coordinate inferential labour across body and world.

References

[1] Michael Breakspear. Dynamic models of large-scale brain activity. *Nature Neuroscience*, 20(3):340–352, 2017.

[2] Rodney A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47(1–3):139–159, 1991.

[3] Christopher L. Buckley, Chang Sub Kim, Simon McGregor, and Anil K. Seth. The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81:55–79, 2017.

[4] Sabrina Carls-Diamante. Embodied cognition and the octopus: Integrating distributed nervous systems, morphological computation, and soft-bodied intelligence. *BioSystems*, 220:104749, 2022.

[5] Matteo Cianchetti, Cecilia Laschi, Arianna Menciassi, and Paolo Dario. Biomedical applications of soft robotics. *Nature Reviews Materials*, 3:143–153, 2018.

[6] Lancelot Da Costa, Thomas Parr, Noor Sajid, Sebastijan Veselic, Victor Neacsu, and Karl J. Friston. Active inference on discrete state-spaces: A synthesis. *Journal of Mathematical Psychology*, 99:102447, 2020.

[7] Tamar Flash and Luigi Zullo. Motor control in the octopus arm. *Current Biology*, 33(10):R429–R441, 2023.

[8] Karl Friston. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

[9] Karl J. Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo. Active inference: A process theory. *Neural Computation*, 29(1):1–49, 2017.

[10] Karl J. Friston, Thomas FitzGerald, Francesco Rigoli, Philipp Schwartenbeck, and Giovanni Pezzulo. Active inference: A process theory. *Neural Computation*, 29(1):1–49, 2017.

[11] Stefano Fusi, Earl K. Miller, and Mattia Rigotti. Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, 37:66–74, 2016.

[12] Melina E. Hale. Octopus as a comparative model for understanding the neural control of limb movement and limb-based behaviors. *Current Opinion in Neurobiology*, 91:102982, 2025.

[13] Binyamin Hochner. An embodied view of octopus neurobiology. *Current Biology*, 22(20):R887–R892, 2012.

[14] Binyamin Hochner. An embodied view of octopus neurobiology. *Current Biology*, 22(20):R887–R892, 2012.

[15] Binyamin Hochner, Letizia Zullo, Tal Shomrat, Guy Levy, and Nir Nesher. Embodied mechanisms of motor control in the octopus. *Current Biology*, 33(20):R1119–R1125, 2023.

[16] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998.

[17] J. A. Scott Kelso. *Dynamic Patterns: The Self-Organization of Brain and Behavior*. MIT Press, Cambridge, MA, 1995.

[18] William M. Kier and A. M. Smith. Tongues, tentacles and trunks: The biomechanics of movement in muscular-hydrostats. *Zoological Journal of the Linnean Society*, 83(4):307–324, 1985.

[19] Jens Kober, J. Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.

[20] Jens Kober, J. Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.

[21] Kohei Nakajima, Helmut Hauser, Rongjie Kang, Emanuele Guglielmino, Darwin G. Caldwell, and Rolf Pfeifer. A soft body as a reservoir: Case studies in a dynamic model of octopus-inspired soft robotic arm. *Frontiers in Computational Neuroscience*, 7:91, 2013.

[22] Lauren Olson and Clifton W. Ragsdale. The arm nervous system of cephalopods. *Current Opinion in Neurobiology*, 80:102741, 2023.

[23] Thomas Parr and Karl J. Friston. Generalised free energy and active inference. *Biological Cybernetics*, 113(5–6):495–513, 2019.

[24] Thomas Parr, Giovanni Pezzulo, and Karl J. Friston. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, Cambridge, MA, 2022.

[25] Thomas Parr, Giovanni Pezzulo, and Karl J. Friston. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, Cambridge, MA, 2022.

[26] Giovanni Pezzulo, Thomas Parr, Paul Cisek, Andy Clark, and Karl Friston. Generating meaning: Active inference and the scope and limits of passive AI. *Trends in Cognitive Sciences*, 28(2):97–112, 2024.

[27] Giovanni Pezzulo, Thomas Parr, Paul Cisek, Andy Clark, and Karl Friston. Generating meaning: active inference and the scope and limits of passive AI. *Trends in Cognitive Sciences*, 28(2):97–112, 2024.

[28] Rolf Pfeifer and Josh Bongard. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, Cambridge, MA, 2006.

[29] Rolf Pfeifer and Josh Bongard. *How the Body Shapes the Way We Think: A New View of Intelligence*. MIT Press, Cambridge, MA, 2007.

[30] Rolf Pfeifer and Fumiya Iida. Embodied artificial intelligence: Trends and challenges. In Fumiya Iida, Rolf Pfeifer, Luc Steels, and Yasuo Kuniyoshi, editors, *Embodied Artificial Intelligence*, volume 3139 of *Lecture Notes in Computer Science*, pages 1–26. Springer, Berlin, Heidelberg, 2004.

[31] Matteo Priorelli, Ivilin P. Stoianov, and Giovanni Pezzulo. Embodied decisions as active inference. *PLOS Computational Biology*, 21(6):e1013180, 2025.

[32] Mattia Rigotti, Omri Barak, Melissa R. Warden, Xiao-Jing Wang, Nathaniel D. Daw, Earl K. Miller, and Stefano Fusi. The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451):585–590, 2013.

[33] Alexander D. Shaw. Polyphonic intelligence: Constraint-based emergence, pluralistic inference, and non-dominating integration, 2025. Preprint.

[34] Bruno Siciliano and Oussama Khatib, editors. *Springer Handbook of Robotics*. Springer, Cham, 2 edition, 2016.

[35] Bruno Siciliano and Oussama Khatib, editors. *Springer Handbook of Robotics*. Springer, Cham, 2 edition, 2016.

[36] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 2 edition, 2018.

[37] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 2 edition, 2018.

[38] Emanuel Todorov and Michael I. Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235, 2002.

[39] Emmanuelle Tognoli and J. A. Scott Kelso. The metastable brain. *Neuron*, 81(1):35–48, 2014.