The Cognitive-Dimensional Correspondence Principle: Cognitive Limits and the Dimensional Structure of Reality

Author: Ayodele Arigbabu

KAM Attribution: Ayodele Arigbabu - KAM[CDR-4-F]

The KAM attribution indicates the nature and extent of human AI collaboration.

More information: https://www.kernel-inference.com/guidelines.html

Manuscript

Abstract

This paper critically examines the Cognitive-Dimensional Correspondence Principle (CDCP), which posits that the human mind's typical capacity to process approximately four entities or concepts in working memory, along with related processes like attention and executive function, aligns with the four-dimensional structure of the physical universe (three spatial, one temporal). This reference to “four-ness” is not framed as a numerological claim but as a hypothesis about alignment between cognitive bottlenecks and the effective dimensionality of the most informative environmental manifold; with the 3+1 structure of everyday reality as a salient special case.

The paper clarifies multidimensionality in physics, reviews empirical and neuroscientific evidence on cognitive limits (including their variability), and situates CDCP relative to existing capacity theories (fixed-slot, resource, and population-coding models). The principle is developed as a generative research heuristic, supported by an information-theoretic perspective, and grounded in philosophical and neuroscientific traditions. The paper distinguishes CDCP from alternative framings, clarifies the status of "four-ness," and outlines future empirical and computational tests aimed at building on the speculative but testable nature of CDCP while increasing conceptual rigour and clarity about what would count for or against it.

Introduction

How does the structure of the physical world shape the architecture of the mind? This treatise proceeds from an observation of the apparent correspondence between the number of dimensions used to describe the physical universe—three spatial and one temporal—and the typical upper limit of human working cognitive capacity, often cited as around four distinct entities or concepts. This paper critically examines the Cognitive-Dimensional Correspondence Principle (CDCP), which hypothesises that this convergence reflects an adaptive alignment between cognitive architecture and environmental structure.

Importantly, the CDCP is explicitly not a numerological claim about the metaphysical specialness of the number four, nor does it posit "four-ness" as an invariant or mystical constant. Rather, it advances a principled hypothesis: cognitive bottlenecks for simultaneous integration adaptively align with the effective dimensionality of the organism's high-variance environmental manifold. In human everyday experience, this manifold is dominated by three spatial and one temporal dimension, yielding a modal capacity near four as a natural consequence; not a predetermined magic number. This generalised formulation (developed further in Section 4.4) shifts focus from numerical coincidence to ecological and informational correspondence.

The paper proceeds as follows. Section 2 clarifies the meaning of multidimensionality in physics and introduces convergent metadesign approaches to four-fold structure. Section 3 reviews empirical literature on working memory, attention, and executive function, emphasising capacity limits and their variability. Section 4 formally states the CDCP and refines its scope, including a generalised formulation that de-emphasises any special status of "four-ness" per se. Section 5 discusses neural, computational, and attention-specific mechanisms, addresses alternative explanations, and explicitly situates CDCP with respect to fixed-slot, resource, and population-coding models of capacity. Section 6 offers an information-theoretic framing, including a stylised model that demonstrates, under simple assumptions, how alignment between capacity and environmental dimensionality can arise. It also connects these ideas to dimensionality reduction in machine learning and philosophical theories of closure. Section 7 presents CDCP as a generative heuristic and sharpens discriminative predictions that differentiate it from competing accounts. Section 8 sketches implications for cognitive science, philosophy, education, and artificial intelligence, and outlines directions for future computational work. Section 9 concludes with a balanced assessment of CDCP's promise and limitations.

Throughout, clear delineation between hypothesis, evidence, and speculation is maintained and apophenia or numerological overreach is avoided

Multidimensionality in the Physical World

Physical Dimensions: Definitions and Implications

In physics, a dimension is an independent axis along which position or measurement can be specified. The physical universe is conventionally described as having three spatial dimensions (length, width, height) and one temporal dimension (time), forming the four-dimensional fabric of spacetime (Einstein, 1916). This structure underlies the laws of motion, causality, and the possibility of perception and action.

Locating an object or event in spacetime requires specifying its coordinates along these independent axes. The fact that these dimensions are independent is critical: variation along one cannot, in general, be reduced to or predicted from variation along the others. This independence makes dimensionality a natural candidate for shaping the complexity of the environment that cognitive systems must navigate.

CDCP leans critically on the idea of "effective dimensionality" - the true underlying complexity of an environment's structure, distinct from its surface or "ambient" count of variables.

This may be explicitly explained with the example of tracking the positions of objects in a room, for which ordinarily, there are many variables (positions, colours, speeds, sounds). Most variation however boils down to just a few independent factors: three spatial directions (left-right, up-down, forward-back) and change over time. Features like colour or sound often add little independent information—they may correlate strongly with position (e.g., a red ball's colour doesn't predict much beyond its location in typical scenarios) or be irrelevant to the task, creating redundancies that reduce the effective number of orthogonal (uncorrelated) dimensions needed to capture behaviourally relevant patterns.

Formally, effective dimensionality is the minimal number of orthogonal (uncorrelated) factors required to explain the majority of predictive variance (spread of values) in behaviourally relevant outcomes (e.g., successful navigation or action); accounting for redundancies, correlations, and noise. This can be quantified independently of cognition using tools like dimensionality reduction - e.g., how many principal components capture ≥90% variance (spread of values) in sensorimotor (perception and action) or task data (van der Maaten, Postma, & van den Herik, 2009) - or information-theoretic measures (e.g., independent variables maximising predictive power).

In everyday human sensorimotor (perception and action) reality, three spatial axes plus time dominate such measures, yielding an effective dimensionality near four.

Alternative Dimensional Schemes

Some physical theories such as string theory posit additional spatial dimensions, which are typically compactified or not directly accessible at macroscopic scales. For cognition, it is the directly perceivable, independent variables - three spatial and one temporal - that are most relevant for adaptive behaviour and thus for the shaping of cognitive architecture. The CDCP focuses on the effective dimensionality of the environment as experienced by organisms, not on all dimensions posited by fundamental physics. The question is not how many dimensions exist in principle, but how many independent degrees of freedom are reliably behaviourally relevant.

Convergent Design Principles: Creative Quartets and Tetrahedral Logic

A primary motivation for the cognitive dimensional correspondence hypothesis stems from John Wood's concept of "creative quartets" as framed in his metadesign principles, which gravitate toward four-fold structures as cognitively and relationally tractable units. Wood's concept of "creative quartets" models four interdependent agents or variables as the vertices of a tetrahedron, emphasising the tetrahedron's mnemonic and conceptual optimality for orchestrating complex relations (Wood, 2003). This demonstration of a convergence in design practice on schema based on four independent nodes and the utilisation of their six connecting edges to track relations, whose relational complexity can be collectively apprehended without collapse into unwieldy abstraction, presents a heuristic inspired by cognitive limits near four in relational reasoning (Wood, 2003) and will be explored further in Sections 7–8 as a test suite for validating CDCP.

Perceptual and Cognitive Access to Dimensions

Human perception is fundamentally shaped by the 3+1 structure. Visual, auditory, and somatosensory systems are tuned to extract information about spatial relationships and temporal change. Higher spatial or temporal dimensions, though mathematically conceivable, are not directly accessible to human perception and must be apprehended through abstraction or analogy.

Dimensionality thus enters cognition at two levels: as a structural property of the world, constraining possible trajectories, interactions, and affordances; and as a structural property of representations, shaping how sensory input is encoded, combined, and used for decision-making. The CDCP is concerned with the relationship between these two: how the dimensional structure of the world's high-variance manifold might be mirrored in the dimensional structure of cognitive bottlenecks.

Cognitive Limits: Evidence, Variability, and Context

Working Memory and Simultaneous Processing

Working memory refers to the system responsible for temporary storage and manipulation of information necessary for complex cognitive tasks (Baddeley, 2010). Quantifying its capacity has been a major focus of cognitive psychology. Miller's (1956) classic "seven, plus or minus two" estimate has been revised downward in light of research controlling for chunking and rehearsal, with many studies suggesting a limit closer to four for ungrouped, unrelated items (Cowan, 2001; Luck & Vogel, 1997).

This "magic number ~4" should be understood as a modal central tendency under conditions approximating simple, independent items, and the outcome of a particular class of tasks that often implicitly embody aspects of everyday 3+1 structure (e.g., objects in space over short temporal intervals). From the perspective of CDCP, when the effective dimensionality of the relevant representational space is comparable to 3+1, cognitive bottlenecks will manifest as a capacity in the vicinity of four. In tasks with lower effective dimensionality or with strong redundancy across dimensions, apparent capacity may differ systematically. Conversely, artificially constructed environments with higher effective dimensionality might yield different patterns.

Nuancing the "Magic Number 4"

The "~4" is a modal value, a central tendency, not a hard ceiling. Individual differences, task demands, and context can shift this limit. The CDCP does not claim deterministic or universal "four-ness," but rather that the cognitive system's default architecture aligns with the effective dimensionality of the most salient, independent variables in our typical environment. In our world, that benchmark is naturally three spatial and one temporal dimension.

Attention and Executive Function

Beyond working memory, attention and executive function also exhibit capacity constraints. Multiple object tracking studies suggest an upper limit of approximately four for simultaneous attentional selection (Pylyshyn & Storm, 1988).

Executive function research demonstrates that managing more than four distinct rules or tasks leads to sharp performance declines (Oberauer et al., 2001; see also Halford, Cowan, & Andrews, 2007). These findings suggest that a "four-ness"-like constraint may generalize across core cognitive processes, including relational reasoning where humans reliably process up to quaternary relations.

It is also worth noting that temporal attention is found to show comparable structure, with anticipatory mechanisms organising processing around expected moments in time (Nobre & van Ede, 2018), suggesting that temporal structure, no less than spatial structure, is subject to capacity-limited allocation of attention.

CDCP interprets these convergences as evidence that the same underlying bottlenecks, tuned to a world structured by a small number of independent dimensions, are being probed from different angles: storage, selection, and coordination.

Empirical Evidence, Variability, and Interim Findings

Visual working memory studies by Luck and Vogel (1997) demonstrated that participants could store about four objects with high fidelity. Attention research by Pylyshyn and Storm (1988) found a four-object limit in parallel tracking tasks, with neuroimaging studies implicating the parietal cortex and frontoparietal networks in spatial attention allocation, with activity plateauing around four tracked items (Todd & Marois, 2004). Executive function work by Oberauer et al. (2001) showed that performance drops beyond four concurrent rules or tasks, a pattern that aligns with limits observed in relational reasoning complexity (Halford et al., 2007).

Critical variability is also evident: expertise and chunking can shift the effective limit (Chase & Simon, 1973), and children or individuals with neurodevelopmental differences may exhibit lower or more variable limits (Gathercole & Alloway, 2008). fMRI and EEG studies have shown that the prefrontal cortex and parietal regions exhibit increased activation as more items are held in working memory or attended to, with a plateau or overload response typically observed at around four items (Vogel & Machizawa, 2004; Todd & Marois, 2004).

Theoretical Justification

The convergence on four is often interpreted as reflecting a balance between the need for flexible, combinatorial thought and the constraints of neural architecture (Cowan, 2001). From the CDCP perspective, this balance is not purely internal: it is also shaped by the statistical and dimensional structure of the environment. When the environment's most informative manifold is effectively 3+1-dimensional, an architecture that can stably track a small, comparable number of independent variables is adaptive. This does not rule out more traditional explanations (energy constraints, noisy population codes), but reframes them as proximal mechanisms through which a more distal alignment between mind and world may be realised.

The Cognitive-Dimensional Correspondence Principle (CDCP)

Statement, Scope, and Framing

Scope: The CDCP is focused on working memory capacity, the number of unrelated entities or concepts that can be simultaneously maintained and manipulated in conscious awareness, but also extends to attention and executive function, which exhibit similar capacity constraints.

Statement: The CDCP proposes that cognitive bottlenecks for simultaneous integration tend to align with the effective dimensionality of the high-variance manifold in an organism's environment. In our everyday world, this manifold is dominated by three spatial and one temporal dimension, so a modal capacity in the vicinity of four is expected under standard conditions. The central commitment is thus not to the number four itself, but to an adaptive correspondence between capacity and environmental dimensional structure.

This formulation allows for systematic deviations from four in environments, tasks, or species where the effective dimensionality of behaviourally relevant information differs from 3+1, and provides a richer set of testable predictions about how capacity should shift under controlled changes in environmental structure.

Philosophical and Scientific Foundations

The CDCP draws on traditions that seek correspondence between mind and world. Kant (1781/1998) argued that space and time are the a priori forms of human intuition, structuring all possible experience. Gibson (1979) emphasised the ecological relationship between perception and the affordances of the environment. Roger Shepard's work on the internalisation of environmental regularities (Shepard, 1987) argued that evolutionary pressures shape our minds to mirror the kinematic and geometric properties of the 3D world, though without proposing a specific numerical capacity limit. The CDCP extends these traditions by proposing a testable link between the number of fundamental, independent environmental dimensions that dominate behaviourally relevant variance and the modal cognitive capacity limit for simultaneously integrating distinct entities or variables.

Evolutionary Perspective

From an evolutionary standpoint, organisms that could efficiently track the most salient, independent variables in their environment - namely, the three spatial axes and time - would have a survival advantage. Over evolutionary timescales, selection may have favoured neural architectures that optimise the integration of these four dimensions, resulting in a modal cognitive capacity aligned with the structure of spacetime.

Under the generalised formulation, evolution shapes cognitive systems to allocate limited representational resources across those dimensions that most reliably predict survival and reproduction-relevant outcomes. In a 3+1 world, this naturally points toward a bottleneck tuned to a small number of independent variables approximating that dimensionality, though not necessarily with perfect numerical equality.

Beyond "Four": A Generalized CDCP

Edwin Abbott Abbott’s Flatland: A Romance of Many Dimensions provides a literary thought experiment that imaginatively explores how inhabitants of a two-dimensional world might conceptualise space and struggle to grasp higher dimensions, providing a literary metaphor for reinforcing the idea that cognitive architecture is deeply conditioned by environmental dimensionality (Abbott, 1884).

The generalised CDCP therefore suggests that: in environments where everyday behaviour is effectively two-dimensional (e.g., organisms constrained to surfaces with limited temporal demands), CDCP predicts that modal capacity limits may be closer to the number of effective dimensions governing behaviour in those contexts. In artificially constructed environments or virtual worlds engineered to have higher effective dimensionality (e.g., tasks requiring simultaneous integration of multiple independent feature dimensions beyond space and time), CDCP predicts that the directional pressure on capacity would be toward tracking more than four independent variables.

Clear and replicable deviations from ‘four’ that systematically covary with environmental dimensionality however, would be more informative about CDCP's validity than any single capacity estimate under standard conditions.

Mechanistic Hypotheses, Interim Findings, and Alternative Explanations

Neural, Computational, and Attention-Specific Mechanisms

Several plausible mechanisms could instantiate an alignment between cognitive bottlenecks and environmental dimensionality.

The hippocampus and entorhinal cortex encode spatial information through place and grid cells, which map environmental locations in two and three dimensions (O'Keefe & Nadel, 1978; Hafting et al., 2005). Time cells encode temporal sequences (Eichenbaum, 2014), suggesting a neural substrate for integrating spatial and temporal variables.

The prefrontal cortex supports the maintenance and manipulation of multiple items in working memory (D'Esposito & Postle, 2015). Neural population codes in this region may be optimized for tracking a small number of independent variables, potentially reflecting the environmental dimensionality that most constrains behaviour.

The parietal cortex and dorsal attention network are central to the allocation of spatial attention (Posner & Petersen, 1990). Neuroimaging shows these regions reach a processing plateau at around four items (Todd & Marois, 2004), consistent with an architecture tuned to a limited number of independent spatial-temporal variables.

Neural network models with limited capacity such as recurrent networks with bottleneck layers tend to prioritise the most informative, independent features of input data (Fusi et al., 2016). In environments where three spatial and one temporal variable are most salient, these models may converge on a similar capacity constraint, even without explicit programming.

Current neuroimaging and electrophysiological studies support the idea that both working-memory and attention-related brain regions exhibit capacity limits around four items, with a plateau or overload response at this threshold (Vogel & Machizawa, 2004; Todd & Marois, 2004). CDCP interprets these patterns as potential signatures of an architecture that has adapted to a 3+1 world.

Alternative Explanations and an Experimental Design Example

Alternative accounts suggest that cognitive limits arise from neural bottlenecks (constraints on synaptic connectivity or neural firing rates), metabolic costs (energy limitations on maintaining active representations), bounded rationality (resource-allocation strategies that optimize performance under uncertainty; Simon, 1957), or noise in population codes (capacity reflecting the point at which additional items cause decoding accuracy to drop below functional thresholds).

These explanations do not invoke environmental dimensionality; they locate the origin of capacity entirely within the brain's internal constraints.

A concrete experiment could use a dual-task paradigm in which participants simultaneously track multiple objects (attention) and maintain items in working memory, with task blocks manipulating either the number or salience of spatial/temporal dimensions (e.g., 2D vs. 3D displays, rich vs. impoverished temporal structure) or the metabolic load (e.g., via concurrent physical exertion).

Discriminative Hypothesis A (CDCP): If capacity limits are partially shaped by environmental dimensionality, tasks that increase the effective dimensionality or salience of independent spatial-temporal variables should produce systematic shifts in capacity (e.g., changing plateau points or neural activation profiles), even when metabolic demand is held approximately constant.

Discriminative Hypothesis B (Metabolic/Bottleneck Account): If capacity limits are primarily due to intrinsic neural or metabolic bottlenecks independent of environmental structure, manipulating dimensional salience should produce minimal changes in capacity, whereas changes in metabolic load should reliably shift capacity.

Observing robust dimensionality-driven shifts that cannot be reduced to generic difficulty or energy cost would support the CDCP framework.

Relation to Slot, Resource, and Population-Coding Models

Literature has developed around different formalisations of working-memory capacity: fixed-slot models propose that a small, discrete number of "slots" can each hold one item with high fidelity; continuous resource models treat capacity as a divisible resource that can be spread across items, with precision per item decreasing as more items are maintained; population-coding accounts emphasize that apparent item limits emerge from the properties of noisy population codes and decision criteria.

The CDCP is agnostic about which framework best captures the micro-mechanics of working memory - whether fixed slots, continuous resources, or population codes. For instance, relational binding accounts of complexity limits (Halford et al., 2007) can be viewed as one proximal mechanism through which an ecologically shaped bottleneck might be realised. Instead, it operates at a higher level of description, addressing why the effective bottleneck, whatever its detailed implementation, tends to fall in the observed range and why it might covary with environmental dimensionality.

Fixed-slot accounts can be interpreted as one way of implementing a CDCP-consistent bottleneck: the number of slots may itself be an adaptive response to a world with a small number of independent high-variance dimensions. Continuous resource and population-coding models can similarly be interpreted as mechanisms through which a limited information budget is allocated across variables, with CDCP offering a hypothesis about why the optimal allocation often corresponds to tracking a number of variables comparable to environmental dimensionality.

CDCP adds value to traditional capacity models by emphasizing the match between internal constraints and the structure of the environment, and predicts that when the effective dimensionality of the relevant task manifold is experimentally manipulated, capacity limits should change in systematic, predictable ways, rather than being fixed solely by brain-internal parameters. If future work shows that capacity limits are entirely explained by generic resource constraints and remain invariant under large and carefully controlled changes in environmental effective dimensionality, that would count against the core claim of CDCP.

Thus, CDCP should not be seen as a rival to these capacity models, but as a complementary principle that constrains the types of capacity models that are ecologically and evolutionarily plausible.

Information Theory and the CDCP

Bounded Rationality and Cognitive Bottlenecks

Information theory provides a rigorous framework for understanding cognitive limits. According to bounded rationality, cognitive systems maximise informational efficiency given processing constraints (Simon, 1957). The structure of the environment determines which variables are most informative; the brain's architecture determines how many can be processed at once.

Let us proceed by assuming the cognitive system has an effective information-processing capacity C (in bits) in a given time window. This capacity must be allocated across multiple potential variables in the environment, each with its own entropy and relevance to survival-relevant outcomes.

Environmental Structure and Informational Salience

Suppose the environment presents a set $D = \{d_1, d_2, \ldots, d_n\}$ of candidate variables (e.g., spatial coordinates, temporal intervals, object features), some of which are largely independent, high in variance, and strongly predictive of survival or reward-related outcomes. If the three spatial and one temporal dimension account for a large share of the variance in outcomes relevant to survival, then information theory predicts these will often be the most valuable to track. Over evolutionary time, cognitive systems may have adapted to prioritise these four, because they consistently maximise mutual information with behaviourally relevant outcomes under the constraint C.

This yields a probabilistic tendency, not a deterministic law: other variables can matter, and the importance of any given dimension can vary by context.

Formalisation at a High Level

Let Y denote a survival-relevant outcome (e.g., success of a movement, avoidance of a predator). Let $X = (X_1, X_2, \ldots, X_n)$ denote environmental variables, including spatial and temporal dimensions.

The cognitive system cannot represent all of X with full fidelity. Instead, it must select a subset $S \subseteq \{1, \ldots, n\}$ and allocate its limited capacity C across those variables. A natural objective is to maximize mutual information $I(X_S; Y)$ subject to a capacity constraint:

$\max I(X_S; Y)$ subject to $\mathrm{Cost}(q) \leq C$

where q represents the encoding scheme.

Under plausible assumptions (e.g., that the three spatial coordinates and time dominate variance and predictive power), the optimal subset S* will typically include these four variables. CDCP proposes that the architecture of working memory and attention reflects this long-run optimization, in the sense that it is particularly well-tuned to represent approximately that number of independent variables concurrently.

A Stylized Information-Theoretic Model of Dimensional Alignment

To make these ideas concrete, consider a simplified model.

Assumptions:

Let n denote the effective environmental dimensionality (e.g., $n = 4$ for three spatial plus one temporal dimension). The environment presents n independent, zero-mean Gaussian variables $X_1, \ldots, X_n$, each with variance σ^2.

A survival-relevant outcome Y is a linear function of these variables plus noise: $Y = \Sigma \beta_i X_i + \varepsilon$, where ε is independent noise with variance σ^2_ε.

The cognitive system has an effective capacity C (in bits) that must be allocated across at most k variables simultaneously, where k is the effective "capacity limit."

Representing each variable with sufficient fidelity to be behaviourally useful requires a minimum of c bits per variable (a crude but illustrative threshold).

The system aims to choose k and a subset S of size $|S| = k$ to maximise expected performance, which is monotonically related to $I(X_S; Y)$, subject to $kc \leq C$.

Consequences:

The capacity constraint implies $k \leq \lfloor C/c \rfloor$. If the β_i are all similar in magnitude and the variables are independent, then including more variables increases mutual information, but with diminishing returns once the noise term ε dominates. For a given C, there will be an optimal range of k beyond which either the per-variable precision becomes too low to be useful or additional variables add little information relative to noise.

If we identify D with the effective dimensionality of the environment's high-variance manifold (e.g., $D = 4$ for three spatial plus one temporal dimension) and suppose that the system has evolved such that C is just sufficient to encode these D variables at useful precision, then the model naturally yields $k^* \approx D$ as an adaptive sweet spot: representing roughly one variable per dimension at sufficiently high fidelity.

Importantly, nothing in this model insists that k^* must be exactly equal to D, or that $D = 4$ is fundamentally special. Instead, it shows how, given a bounded capacity C in an environment where a small set of independent variables dominate predictive power, it is rational in an information-theoretic sense to allocate capacity such that the effective bottleneck k^* is of the same order as the environmental dimensionality D. In our everyday world, where behaviourally relevant structure is heavily shaped by three spatial and one temporal dimension, this alignment suggests a modal capacity near four. In environments with different effective dimensionalities, the same logic would point toward different values of k^*, providing a pathway for empirical tests.

This stylised model is not intended as a full proof of CDCP, but as an explicit, assumption-driven demonstration that such an alignment is not merely numerological: it can emerge from coherent informational and ecological considerations.

Dimensionality Reduction in Machine Learning and Connection to CDCP

In high-dimensional settings, both humans and artificial systems face the curse of dimensionality: sparse data, broken distance metrics, computational cost, and difficulty of interpretation. Dimensionality reduction methods such as Principal Component Analysis (PCA) and t-SNE aim to find lower-dimensional representations that preserve important structure—variance, pairwise distances, or local neighborhoods (van der Maaten et al., 2009; van der Maaten & Hinton, 2008). These methods make complex spaces tractable for human understanding and machine learning.

From this perspective, CDCP can be seen as a biological counterpart to dimensionality reduction: rather than representing arbitrarily many features, the cognitive system compresses experience into a small, low-dimensional manifold that preserves those variables most predictive of behaviorally relevant outcomes. The 3+1 structure of everyday reality functions, in effect, as a privileged "principal component space" for human cognition. Both statistical dimensionality reduction and cognitive architecture solve the same fundamental problem: how to represent a high-dimensional world in a compact, usable form.

Closure Theory and CDCP

Hilary Lawson's theory of closure provides a complementary philosophical lens. Lawson (2001) argues that the world is fundamentally open (multiple interpretations or ways of understanding) and this openness is fundamentally navigated cognitively through the imposition of conceptual or structural limits on an 'open' and complex reality to create manageable 'things' or 'facts.'

From this perspective, both statistical dimensionality reduction and everyday cognition can be understood as closure operations that compress a potentially unbounded flux into a manageable, structured manifold.

The CDCP can then be seen as a claim about a particularly pervasive family of closures: those that bind three spatial and one temporal dimension into the stable objects and events of ordinary experience. The tetrahedral framework in metadesign (Section 2.1; extended in 7.2.6, 8) reflects an operative closure for four-way coordination, that can be further understood through the evolutionary correlation the CDCP proposes.

Dimensionality reduction techniques in machine learning offer a formal analogue of such closures, while Lawson's framework reminds us that alternative closure regimes - and thus alternative effective dimensionalities - are possible, with corresponding implications for cognitive architecture. In other words, what we perceive as "dimensional structure" is not independent of the closures we enact; it is co-constituted through ongoing cognitive and cultural practices.

CDCP as a Generative Research Heuristic

The Value of Generative Heuristics

Treating the CDCP as a generative heuristic encourages the formulation of new hypotheses, the design of comparative studies, and the integration of insights from cognitive science, neuroscience, philosophy, and information theory. It is not a claim to be accepted on faith, but a prompt for disciplined speculation and empirical inquiry.

Treating the CDCP as a generative heuristic encourages the formulation of new hypotheses, the design of comparative studies, and the integration of insights from cognitive science, neuroscience, philosophy, and information theory. It is not a claim to be accepted on faith, but a prompt for disciplined speculation and

empirical inquiry.

To ensure rigorous testability and address potential concerns about post-hoc flexibility in the generalized formulation, CDCP is explicitly falsifiable. Observations that would refute the core alignment claim include: (1) absence of systematic capacity shifts (e.g., effect size d 0.8 across conditions).

Contrasting Prediction: Given fixed architectural constraints, ANNs will converge on similar effective capacities (variability d 0.7 gains in 3+1 tasks; weaker in mismatched dimensionality.

Contrasting Prediction (Internal-Bottleneck Only): Scaffolds show generic benefits independent of dimensionality match. Falsifier: No systematic scaling of scaffold efficacy with independently measured task dimensionality.

Comparative Species Tests

CDCP Hypothesis S1: Species in low-D environments (e.g., 2D surface insects, flatfish) show systematically lower modal WM/attention limits (e.g., mean 0.5). Preliminary: Insects manage ~1–4 amid planar tasks; fish spatial WM plateaus low.

Contrasting Prediction: Capacities invariant, driven by brain size/energy alone. Falsifier: No covaration post-controls (e.g., phylogeny).

Mitigation and Incremental Evidence

Given the complexity of fully testing CDCP, incremental studies that provide partial evidence, such as showing a shift in neural activation patterns or behavioural capacity with manipulated dimensional salience, are valuable. Careful control conditions are essential to distinguish dimensional effects from generic difficulty, novelty, or attentional load. The testing regimes would also yield more meaningful outcomes if carried out to test evolutionary outcomes rather than one-off effects. This directly implies iterative simulation based testing protocols and therefore it is pertinent to note that implementing these studies, especially ANN and multidimensional task designs, requires significant computational and methodological resources.

Implications for Research and Real-World Application

For Cognitive Science and Philosophy of Mind

The CDCP invites a re-examination of the relationship between mind and world, perception and cognition, and the evolutionary shaping of cognitive architecture. It encourages interdisciplinary dialogue and the search for unifying principles, but must be pursued with methodological rigour and critical self-awareness.

Cognitive bottlenecks should be studied in relation to environmental structure, not in isolation. Philosophy of mind can benefit from ecological and information-theoretic perspectives that treat capacity limits as partly world-relative.

For Education and Interface Design

Understanding cognitive limits can inform the design of educational materials and user interfaces. Presenting information in groups of around four may, in many contexts, align with natural processing constraints shaped by everyday 3+1 structure, enhancing comprehension and retention.

CDCP can yield operational rubrics for education and interface design by informed adaptation of design practices such as the tetrahedral metadesign system referenced in Section 7.2.5 to dimensionality-matched scaffolds. In regulatory training or control dashboards tracking four coupled variables, explicit pairwise mapping reduces integration load, aligning with CDCP's predicted 3+1 tuning. This approach would enabled streamlined coordination of curricula for 2D-constrained reasoning (e.g., planar navigation, surface biology) via triangular scaffolds (three nodes, three edges), alongside hypergraphs for high-dimensional AI visualization, by scaling input complexity to the environmental manifold dimensionality humans are tuned for.

However, capacity is flexible and context-dependent, especially when chunking or expertise is involved. In domains where the effective dimensionality of the relevant conceptual space is lower or higher, different grouping strategies may be optimal.

Design should be empirically tested rather than assumed, but CDCP provides a useful heuristic: ‘consider how many independent dimensions or concepts are being simultaneously demanded, and whether this aligns with known cognitive bottlenecks’.

For Artificial Intelligence

Insights from CDCP and information theory can guide the design of artificial systems that must operate under information-processing constraints in complex environments. If effective capacity limits should adapt to environmental dimensionality, then architectures and training regimes that allow internal bottlenecks to co-adapt with environmental structure may be more robust. Evaluating ANN capacity solely by internal metrics such as parameter counts may miss important aspects of task and environment-relative capacity.

The mapping between biological and artificial cognition should be empirically validated, not assumed. CDCP suggests a set of design principles and benchmarks for such validation.

Future Computational Work

While the present paper does not include new simulations, it points toward several computational research directions. Train embodied agents in environments with systematically varied effective dimensionalities and constrained representational bottlenecks, measuring how many independent variables can be simultaneously tracked with useful precision. Analyze trained networks to determine the effective dimensionality of internal representations and how these relate to both the environment's state space and behavioural capacity. Combine the stylised information-theoretic model with concrete simulations to test whether alignment between k^* and environmental dimensionality holds under more realistic assumptions.

Such work would provide stronger tests of CDCP and help clarify whether the observed "magic number ~4" in humans is a special case of a broader dimensional alignment principle.

Conclusion

The Cognitive-Dimensional Correspondence Principle offers a lens through which to examine the alignment of cognitive limits and physical dimensionality. In its generalized form, CDCP does not insist on the specialness of "four" as such. Instead, it proposes that cognitive bottlenecks for simultaneous integration tend to align with the effective dimensionality of the environment's high-variance manifold, with the familiar 3+1 structure of everyday sensorimotor reality yielding a natural modal capacity near four.

CDCP is situated more clearly with respect to existing capacity theories as a complementary ecological and information-theoretic principle rather than a competing mechanistic model. A stylized formal model shows how dimensional alignment can arise from bounded rationality and environmental statistics. Sharper, discriminative predictions specify what evidence would support or challenge CDCP.

At the same time, the paper remains clear about the speculative and heuristic nature of the proposal. CDCP is best viewed not as a completed theory, but as a research program: it encourages experiments and models that treat cognitive capacity as implicitly tuned to the structure of the world, and that seek systematic patterns in how capacity changes under controlled manipulations of environmental dimensionality.

The value of CDCP lies less in the claim that "four" is a magic number and more in the intuition that mind and world are co-structured. The apparent limits of thought may bear the imprint of the dimensions along which reality itself varies most. Whether this principle holds up to empirical and computational scrutiny remains an open question, but its pursuit promises to deepen understanding of the interplay between cognitive architecture and the dimensional structure of reality.

References

Abbott, E. A. (1884). Flatland: A romance of many dimensions. Seeley & Co.

Baddeley, A. (2010). Working memory. Current Biology, 20(4), R136–R140. https://doi.org/10.1016/j.cub.2009.12.014

Chase, W. G., & Simon, H. A. (1973). Perception in chess. Cognitive Psychology, 4(1), 55–81. https://doi.org/10.1016/0010-0285(73)90004-2

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. Behavioral and Brain Sciences, 24(1), 87–114. https://doi.org/10.1017/S0140525X01003922

D'Esposito, M., & Postle, B. R. (2015). The cognitive neuroscience of working memory. Annual Review of Psychology, 66, 115–142. https://doi.org/10.1146/annurev-psych-010814-015031

Eichenbaum, H. (2014). Time cells in the hippocampus: A new dimension for mapping memories. Nature Reviews Neuroscience, 15(11), 732–744. https://doi.org/10.1038/nrn3827

Einstein, A. (1916). The foundation of the general theory of relativity. Annalen der Physik, 49, 769–822.

Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. Current Opinion in Neurobiology, 37, 66–74. https://doi.org/10.1016/j.conb.2016.01.010

Gathercole, S. E., & Alloway, T. P. (2008). Working memory and learning: A practical guide for teachers. SAGE.

Gibson, J. J. (1979). The ecological approach to visual perception. Houghton Mifflin.

Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature, 436(7052), 801–806. https://doi.org/10.1038/nature03721

Halford, G. S., Cowan, N., & Andrews, G. (2007). Separating cognitive capacity from knowledge: A new hypothesis. Trends in Cognitive Sciences, 11(6), 236–242. https://doi.org/10.1016/j.tics.2007.04.001

Kant, I. (1998). Critique of pure reason (P. Guyer & A. W. Wood, Trans.). Cambridge University Press. (Original work published 1781)

Lawson, H. (2001). Closure: A story of everything. Routledge.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. Nature, 390(6657), 279–281. https://doi.org/10.1038/36846

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological Review, 63(2), 81–97. https://doi.org/10.1037/h0043158

Nobre, A. C., & van Ede, F. (2018). Anticipated moments: Temporal structure in attention. Nature Reviews Neuroscience, 19(1), 34–48. https://doi.org/10.1038/nrn.2017.141

O'Keefe, J., & Nadel, L. (1978). The hippocampus as a cognitive map. Oxford University Press.

Oberauer, K., Süß, H. M., Wilhelm, O., & Wittmann, W. W. (2001). The multiple faces of working memory: Storage, processing, supervision, and coordination. Intelligence, 29(2), 167–193. https://doi.org/10.1016/S0160-2896(02)00115-0

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. Annual Review of Neuroscience, 13, 25–42. https://doi.org/10.1146/annurev.ne.13.030190.000325

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. Spatial Vision, 3(3), 179–197. https://doi.org/10.1163/156856888X00122

Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. Science, 237(4820), 1317–1323. https://doi.org/10.1126/science.3629243

Simon, H. A. (1957). Models of man: Social and rational. Wiley.

Thomason, M. E., Race, E., Burrows, B., Whitfield-Gabrieli, S., Glover, G. H., & Gabrieli, J. D. (2009). Development of spatial and verbal working memory capacity in the human brain. Journal of Cognitive Neuroscience, 21(2), 316–332. https://doi.org/10.1162/jocn.2008.21028

Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. Nature, 428(6984), 751–754. https://doi.org/10.1038/nature02466

van der Maaten, L., Postma, E., & van den Herik, J. (2009). Dimensionality reduction: A comparative review. Journal of Machine Learning Research, 10, 66–71

van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. Journal of Machine Learning Research, 9, 2579–2605.

Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. Nature, 428(6984), 748–751. https://doi.org/10.1038/nature02447

Wood, J. (2003). Tetrahedral Reasoning. Metadesigners Network. https://metadesigners.org/Tool-1-Tetrahedral-Logic

License

How to cite

Ayodele Arigbabu. "The Cognitive-Dimensional Correspondence Principle: Cognitive Limits and the Dimensional Structure of Reality". Kernel Inference: Journal of Speculative Inference. KI-d3970697
https://kernel-inference.com/p/KI-d3970697

Author bio

AYODELE ARIGBABU is an architect, writer and creative technologist whose work explores the interfaces where design, art and evolving technologies intercept, with speculation as a recurrent strategy.

He is the publisher and editor of LAGOS_2060, an anthology of science fiction from Africa, and was curator of the Lagos edition of a festival on diverse future perspectives of the African continent, produced by Goethe Institut in 2015 tagged African Futures.

He is a PhD candidate in Artistic Research at NTNU’s Trondheim Academy of Fine Art, where his Forward Remembrance project explores latent utopias through the diverse ways in which artificial intelligence mirrors human aspirations.